
Flow Forecast

Release 0.0.1

Aug 11, 2020

General:

1 Utilities	1
2 Model Evaluator	3
3 Long Train	7
4 Model Dictionaries	9
5 Pre Dictionaries	11
6 PyTorch Training	13
7 Time Model	15
8 Trainer	17
9 Interpolate Preprocessing	19
10 Build Dataset	21
11 Closest Station	23
12 Data Converter	25
13 Preprocess DA RNN	27
14 Preprocess Metadata	29
15 Process USGS	31
16 PyTorch Loaders	33
17 Temporal Features	35
18 Custom Optimizations	37
19 Dummy Torch Model	69
20 Lower Upper Configuration	77

21 Simple Multi Attention Head Model	101
22 Basic Transformer	111
23 Transformer XL	137
24 Basic GCP Utils	185
25 Train da	187
26 Utils	189
27 Custom Types	191
28 Model	193
29 Modules	203
30 Indices and tables	219
Python Module Index	221
Index	223

CHAPTER 1

Utilities

```
flood_forecast.utils.flatten_list_function(input_list)

class flood_forecast.utils.EarlyStopper(patience: int, min_delta: float = 0.0, cumulative_delta: bool = False)
Bases: object

EarlyStopping handler can be used to stop the training if no improvement after a given number of events. Args:
    patience (int): Number of events to wait if no improvement and then stop the training.
    score_function (callable): It should be a function taking a single argument, an Engine object, and return a score float. An improvement is considered if the score is higher.
    trainer (Engine): trainer engine to stop the run if no improvement.
    min_delta (float, optional): A minimum increase in the score to qualify as an improvement, i.e. an increase of less than or equal to min_delta, will count as no improvement.
    cumulative_delta (bool, optional): If True, min_delta defines an increase since the last patience reset, otherwise, it defines an increase after the last event. Default value is False.
```

Examples: .. code-block:: python

```
from ignite.engine import Engine, Events from ignite.handlers import EarlyStopping def score_function(engine):
    val_loss = engine.state.metrics['nll'] return -val_loss

handler = EarlyStopping(patience=10, score_function=score_function, trainer=trainer) # Note: the handler is attached to an Evaluator (runs one epoch on validation dataset). evaluator.add_event_handler(Events.COMPLETED, handler)

check_loss (model, validation_loss) → bool
save_model_checkpoint (model)
```


CHAPTER 2

Model Evaluator

```
flood_forecast.evaluator.stream_baseline(river_flow_df: pandas.core.frame.DataFrame,  
                                         forecast_column: str, hours_forecast=336)  
                                         -> (<class 'pandas.core.frame.DataFrame'>,  
                                         <class 'float'>)
```

Function to compute the baseline MSE by using the mean value from the train data.

```
flood_forecast.evaluator.plot_r2(river_flow_preds: pandas.core.frame.DataFrame) → float  
We assume at this point river_flow_preds already has a predicted_baseline and a predicted_model column
```

```
flood_forecast.evaluator.get_model_r2_score(river_flow_df: pandas.core.frame.DataFrame,  
                                              model_evaluate_function: Callable, forecast_column: str, hours_forecast=336)
```

model_evaluate_function should call any necessary preprocessing

```
flood_forecast.evaluator.get_r2_value(model_mse, baseline_mse)
```

```
flood_forecast.evaluator.get_value(the_path: str) → None
```

```
flood_forecast.evaluator.metric_dict(metric: str) → Callable
```

```
flood_forecast.evaluator.evaluate_model(model: Type[flood_forecast.time_model.TimeSeriesModel],  
                                         model_type: str, target_col: List[str], evaluation_metrics: List[T], inference_params: Dict[KT, VT], eval_log: Dict[KT, VT]) → Tuple[Dict[KT, VT], pandas.core.frame.DataFrame, int, pandas.core.frame.DataFrame]
```

A function to evaluate a model. Requires a model of type TimeSeriesModel

```
flood_forecast.evaluator.infer_on_torch_model(model, test_csv_path: str = None,
                                              datetime_start: datetime.datetime
                                              = datetime.datetime(2018, 9, 22,
                                              0, 0), hours_to_forecast: int
                                              = 336, decoder_params=None,
                                              dataset_params: Dict[KT, VT]
                                              = {}, num_prediction_samples:
                                              int = None) -> (<class 'pan-
                                              das.core.frame.DataFrame'>,
                                              <class 'torch.Tensor'>, <class
                                              'int'>, <class 'int'>, <class
                                              'flood_forecast.preprocessing.pytorch_loaders.CSVTestLoader'>,
                                              <class 'pan-
                                              das.core.frame.DataFrame'>)
```

Function to handle both test evaluation and inference on a test dataframe. :returns

df: df including training and test data end_tensor: the final tensor after the model has finished predictions history_length: num rows to use in training forecast_start_idx: row index to start forecasting test_data: CSVTestLoader instance df_prediction_samples: has same index as df, and num cols equal to num_prediction_samples

or no columns if num_prediction_samples is None

```
flood_forecast.evaluator.generate_predictions(model: Type[flood_forecast.time_model.TimeSeriesModel],
                                              df: pan-
                                              das.core.frame.DataFrame, test_data:
                                              flood_forecast.preprocessing.pytorch_loaders.CSVTestLoader,
                                              history: torch.Tensor, device:
                                              torch.device, forecast_start_idx: int,
                                              forecast_length: int, hours_to_forecast:
                                              int, decoder_params: Dict[KT, VT]) ->
                                              torch.Tensor
```

```
flood_forecast.evaluator.generate_predictions_non_decoded(model:
                                              Type[flood_forecast.time_model.TimeSeriesModel],
                                              df: pan-
                                              das.core.frame.DataFrame,
                                              test_data:
                                              flood_forecast.preprocessing.pytorch_loaders.CS-
                                              history_dim:
                                              torch.Tensor, fore-
                                              cast_length: int,
                                              hours_to_forecast: int)
                                              -> torch.Tensor
```

```
flood_forecast.evaluator.generate_decoded_predictions(model:
                                              Type[flood_forecast.time_model.TimeSeriesModel],
                                              test_data:
                                              flood_forecast.preprocessing.pytorch_loaders.CSVTest-
                                              forecast_start_idx: int,
                                              device: torch.device, his-
                                              tory_dim: torch.Tensor,
                                              hours_to_forecast: int, de-
                                              coder_params: Dict[KT,
                                              VT]) -> torch.Tensor
```

```
flood_forecast.evaluator.generate_prediction_samples(model:  
    Type[flood_forecast.time_model.TimeSeriesModel],  
    df:                                         pandas.core.frame.DataFrame,  
    test_data:  
    flood_forecast.preprocessing.pytorch_loaders.CSVTestL  
    history:          torch.Tensor,  
    device:           torch.device,  
    forecast_start_idx:      int,  
    forecast_length:       int,  
    hours_to_forecast:     int, de-  
    coder_params: Dict[KT, VT],  
    num_prediction_samples: int)  
→ numpy.ndarray
```


CHAPTER 3

Long Train

```
flood_forecast.long_train.split_on_letter(s: str) → List[T]
flood_forecast.long_train.loop_through(data_dir: str, intermittent_gcs: bool = False,
                                         use_transfer: bool = True, start_index: int = 0,
                                         end_index: int = 25) → None
    Function that makes and executes a set of config files This is since we have over 9k files.
flood_forecast.long_train.make_config_file(flow_file_path: str, gage_id: str, station_id:
                                         str, weight_path=None)
flood_forecast.long_train.main()
```


CHAPTER 4

Model Dictionaries

```
flood_forecast.model_dict_function.generate_square_subsequent_mask (sz: int) →  
                                torch.Tensor  
Generate a square mask for the sequence. The masked positions are filled with float('inf'). Unmasked positions  
are filled with float(0.0).
```


CHAPTER 5

Pre Dictionaries

CHAPTER 6

PyTorch Training

```
flood_forecast.pytorch_training.train_transformer_style(model:  
    flood_forecast.time_model.PyTorchForecast,  
    training_params: Dict[KT,  
    VT], takes_target=False,  
    forward_params: Dict[KT,  
    VT] = {}, model_filepath:  
    str = 'model_save') →  
    None
```

Function to train any PyTorchForecast model :model The initialized PyTorchForecastModel :training_params_dict A dictionary of the parameters needed to train model :takes_target boolean: Determines whether to pass target during training :forward_params: A dictionary for additional forward parameters (for instance target)

```
flood_forecast.pytorch_training.torch_single_train(model:  
    flood_forecast.time_model.PyTorchForecast,  
    opt:  
    torch.optim.optimizer.Optimizer,  
    criterion:  
    Type[torch.nn.modules.loss._Loss],  
    data_loader:  
    torch.utils.data.dataloader.DataLoader,  
    takes_target: bool, for-  
    ward_params: Dict[KT, VT]  
    = {}) → float
```

```
flood_forecast.pytorch_training.compute_validation(validation_loader:  
    torch.utils.data.dataloader.DataLoader,  
    model, epoch: int, sequence_size: int, criterion:  
    Type[torch.nn.modules.loss._Loss],  
    device: torch.device, decoder_structure=False,  
    use_wandb: bool = False,  
    val_or_test='validation_loss') → float
```

Function to compute the validation or test loss

CHAPTER 7

Time Model

```
class flood_forecast.time_model.TimeSeriesModel(model_base: str, training_data: str,  
                                                validation_data: str, test_data: str,  
                                                params: Dict[KT, VT])
```

Bases: abc.ABC

An abstract class used to handle different configurations of models + hyperparams for training, test, and predict functions. This class assumes that data is already split into test train and validation at this point.

load_model (model_base: str, model_params: Dict[KT, VT], weight_path=None) → object

This function should load and return the model this will vary based on the underlying framework used

make_data_load (data_path, params: Dict[KT, VT], loader_type: str) → object

Initializes a data loader based on the provided data_path. This may be as simple as a pandas dataframe or as complex as a custom PyTorch data loader.

save_model (output_path: str)

Saves a model to a specific path along with a configuration report of the parameters and data info.

upload_gcs (save_path: str, name: str, file_type: str, epoch=0, bucket_name=None)

Function to upload model checkpoints to GCS

wandb_init()

```
class flood_forecast.time_model.PyTorchForecast(model_base: str, training_data, val-  
                                                idation_data, test_data, params_dict:  
                                                Dict[KT, VT])
```

Bases: *flood_forecast.time_model.TimeSeriesModel*

load_model (model_base: str, model_params: Dict[KT, VT], weight_path: str = None, strict=True)

This function should load and return the model this will vary based on the underlying framework used

save_model (final_path: str, epoch: int) → None

Function to save a model to a given file path

upload_gcs (save_path: str, name: str, file_type: str, epoch=0, bucket_name=None)

Function to upload model checkpoints to GCS

wandb_init()

```
make_data_load(data_path: str, dataset_params: Dict[KT, VT], loader_type: str,  
the_class='default')
```

Initializes a data loader based on the provided data_path. This may be as simple as a pandas dataframe or as complex as a custom PyTorch data loader.

CHAPTER 8

Trainer

```
flood_forecast.trainer.train_function(model_type: str, params: Dict[KT, VT])
```

Function to train a Model(TimeSeriesModel) or da_rnn. Will return the trained model
model_type str: Type of the model (for now) must be da_rnn or :params dict: Dictionary containing all the parameters needed to run the model

```
flood_forecast.trainer.main()
```

Main function which is called from the command line. Entrypoint for all ML models.

CHAPTER 9

Interpolate Preprocessing

```
flood_forecast.preprocessing.interpolate_preprocess.fix_timezones(csv_path:
                                                               str) → pan-
                                                               das.core.frame.DataFrame
Basic function to fix intial data bug related to NaN values in non-eastern-time zones due to UTC conversion.

flood_forecast.preprocessing.interpolate_preprocess.split_on_na_chunks(df:
                                                               pan-
                                                               das.core.frame.DataFrame)
                                                               →
                                                               None

flood_forecast.preprocessing.interpolate_preprocess.interpolate_missing_values(df:
                                                               pan-
                                                               das.core.frame.D
                                                               →
                                                               pan-
                                                               das.core.frame.D)
```

Function to fill missing values with nearest value. Should be run only after splitting on the NaN chunks.

CHAPTER 10

Build Dataset

```
flood_forecast.preprocessing.buil_dataset.build_weather_csv(json_full_path,  
asos_base_url,  
base_url_2,  
econet_data,    vis-  
ited_gages_path,  
start=0,  
end_index=100)  
flood_forecast.preprocessing.buil_dataset.join_data(weather_csv,    meta_json_file,  
flow_csv)  
flood_forecast.preprocessing.buil_dataset.create_visited()  
flood_forecast.preprocessing.buil_dataset.get_eco_netset(directory_path: str) →  
set  
Econet data was supplied to us by the NC State climate office. They gave us a directory of CSV files in following  
format LastName_First_station_id_Hourly.txt This code simply constructs a set of stations based on what is in  
the folder.  
flood_forecast.preprocessing.buil_dataset.combine_data(flow_df:           pandas.core.frame.DataFrame,  
                                         precip_df:          pandas.core.frame.DataFrame)  
flood_forecast.preprocessing.buil_dataset.create_usgs(meta_data_dir: str,  pre-  
cip_path: str, start: int, end:  
int)
```


CHAPTER 11

Closest Station

```
flood_forecast.preprocessing.closest_station.get_closest_gage(gage_df:      pan-
                                                               das.core.frame.DataFrame,
                                                               station_df:  pan-
                                                               das.core.frame.DataFrame,
                                                               path_dir:    str,
                                                               start_row:   int,
                                                               end_row:    int)

flood_forecast.preprocessing.closest_station.haversine(lon1, lat1, lon2, lat2)
Calculate the great circle distance between two points on the earth (specified in decimal degrees)

flood_forecast.preprocessing.closest_station.get_weather_data(file_path:      str,
                                                               econet_gages:
                                                               Set[T], base_url:
                                                               str)
Function that retrieves if station has weather data for a specific gage either from ASOS or ECONet

flood_forecast.preprocessing.closest_station.format_dt(date_time_str: str) → date-
                                                               time.datetime

flood_forecast.preprocessing.closest_station.convert_temp(temparature:  str) →
                                                               float
Note here temp could be a number or 'M' which stands for missing. We use 50 at the moment to fill missing
values.

flood_forecast.preprocessing.closest_station.process_asos_data(file_path:      str,
                                                               base_url:    str)
                                                               → Dict[KT, VT]

Function that saves the ASOS data to CSV uses output of get weather data.

flood_forecast.preprocessing.closest_station.process_asos_csv(path: str)
```


CHAPTER 12

Data Converter

A set of function aimed at making it easy to convert other time series datasets to our format for transfer learning purposes

```
flood_forecast.preprocessing.data_converter.make_column_names(df)
```


CHAPTER 13

Preprocess DA RNN

```
flood_forecast.preprocessing.preprocess_da_rnn.format_data(dat,      targ_column:  
                                         List[str])           →  
                                         flood_forecast.da_rnn.custom_types.TrainData  
  
flood_forecast.preprocessing.preprocess_da_rnn.make_data(csv_path: str, target_col:  
                                         List[str],      test_length:  
                                         int, relevant_cols=['cfs',  
                                         'temp',   'precip']) →  
                                         flood_forecast.da_rnn.custom_types.TrainData
```

Returns full preprocessed data. Does not split train/test that must be done later.

CHAPTER 14

Preprocess Metadata

```
flood_forecast.preprocessing.preprocess_metadata.make_gage_data_csv(file_path:  
str)
```


CHAPTER 15

Process USGS

```
flood_forecast.preprocessing.process_usgs.make_usgs_data(start_date: date-time.datetime, end_date: date-time.datetime, site_number: str) → pandas.core.frame.DataFrame
flood_forecast.preprocessing.process_usgs.process_response_text(file_name: str) → Tuple[str, Dict[KT, VT]]
flood_forecast.preprocessing.process_usgs.df_label(usgs_text: str) → str
flood_forecast.preprocessing.process_usgs.create_csv(file_path: str, params_names: dict, site_number: str)
    Function that creates the final version of the CSV file
flood_forecast.preprocessing.process_usgs.get_timezone_map()
flood_forecast.preprocessing.process_usgs.process_intermediate_csv(df: pandas.core.frame.DataFrame) -> (<class 'pandas.core.frame.DataFrame'>, <class 'int'>, <class 'int'>, <class 'int'>)
```


CHAPTER 16

PyTorch Loaders

```
class flood_forecast.preprocessing.pytorch_loaders.CSVDataLoader(file_path:  
    str,      fore-  
    cast_history:  
    int,      fore-  
    cast_length:  
    int,      tar-  
    get_col:  
    List[T],  rel-  
    evant_cols:  
    List[T],  scal-  
    ing=None,  
    start_stamp:  
    int = 0,  
    end_stamp:  
    int = None,  
    interp-  
    late_param=True)  
Bases: torch.utils.data.dataset.Dataset
```

A data loader that takes a CSV file and properly batches for use in training/eval a PyTorch model :param file_path: The path to the CSV file you wish to use. :param forecast_history: This is the length of the historical time series data you wish to

utilize for forecasting

Parameters

- **forecast_length** – The number of time steps to forecast ahead (for transformer this must equal history_length)
- **relevant_cols** – Supply column names you wish to predict in the forecast (others will not be used)
- **target_col** – The target column or columns you to predict. If you only have one still use a list ['cfs']

- **scaling** – (highly recommended) If provided should be a subclass of sklearn.base.BaseEstimator

and sklearn.base.TransformerMixin) i.e StandardScaler, MaxAbsScaler, MinMaxScaler, etc) Note without a scaler the loss is likely to explode and cause infinite loss which will corrupt weights :param start_stamp int: Optional if you want to only use part of a CSV for training, validation

or testing supply these

Parameters int (end_stamp) – Optional if you want to only use part of a CSV for training, validation, or testing supply these

inverse_scale (result_data: Union[torch.Tensor, pandas.core.series.Series, numpy.ndarray]) → torch.Tensor

```
class flood_forecast.preprocessing.pytorch_loaders.CSVTestLoader(df_path:
    str, forecast_total:
    int, use_real_precip=True,
    use_real_temp=True, tar-
    get_supplied=True,
    interpole-
    late=False,
    **kwargs)
```

Bases: *flood_forecast.preprocessing.pytorch_loaders.CSVDataLoader*

Parameters df_path (str) –

A data loader for the test data.

inverse_scale (result_data: Union[torch.Tensor, pandas.core.series.Series, numpy.ndarray]) → torch.Tensor

get_from_start_date (forecast_start: int)

convert_real_batches (the_col: str, rows_to_convert)

A helper function to return properly divided precip and temp values to be stacked with forecasted cfs.

convert_history_batches (the_col: Union[str, List[str]], rows_to_convert: pandas.core.frame.DataFrame)

A helper function to return dataframe in batches of size (history_len, num_features)

Args: the_col (str): column names rows_to_convert (pd.DataFrame): rows in a dataframe to be converted into batches

CHAPTER 17

Temporal Features

```
flood_forecast.preprocessing.temporal_feats.make_temporal_features(features_list:  
    Dict[KT,  
        VT],  
    dt_column:  
        str,      df:  
        pan-  
        das.core.frame.DataFrame)  
    →  pan-  
    das.core.frame.DataFrame
```

Function to create features

```
flood_forecast.preprocessing.temporal_feats.get_day(x: datetime.datetime) → int  
flood_forecast.preprocessing.temporal_feats.get_month(x: datetime.datetime)  
flood_forecast.preprocessing.temporal_feats.get_hour(x: datetime.datetime)  
flood_forecast.preprocessing.temporal_feats.get_weekday(x: datetime.datetime)
```


CHAPTER 18

Custom Optimizations

```
flood_forecast.custom.custom_opt.warmup_cosine(x, warmup=0.002)
flood_forecast.custom.custom_opt.warmup_constant(x, warmup=0.002)
    Linearly increases learning rate over  $warmup^*t_{total}$  (as provided to BertAdam) training steps. Learning rate
    is 1. afterwards.
flood_forecast.custom.custom_opt.warmup_linear(x, warmup=0.002)
    Specifies a triangular learning rate schedule where peak is reached at  $warmup^*t_{total}$ -th (as provided to
    BertAdam) training step. After  $t_{total}$ -th training step, learning rate is zero.

class flood_forecast.custom.custom_opt.RMSELoss
    Bases: torch.nn.modules.module.Module
    Returns RMSE using: target -> True y output -> Prediction by model source: https://discuss.pytorch.org/t/rmse-loss-function/16540/3
    forward(target: torch.Tensor, output: torch.Tensor)
        T_destination = ~T_destination
        add_module(name: str, module: torch.nn.modules.module.Module) → None
            Adds a child module to the current module.
            The module can be accessed as an attribute using the given name.
        Args:
            name (string): name of the child module. The child module can be accessed from this module
            using the given name
            module (Module): child module to be added to the module.
        apply(fn: Callable[[Module], None]) → T
            Applies fn recursively to every submodule (as returned by .children()) as well as self. Typical use
            includes initializing the parameters of a model (see also nn-init-doc).
        Args: fn (Module -> None): function to be applied to each submodule
        Returns: Module: self
```

Example:

```
>>> @torch.no_grad()
>>> def init_weights(m):
>>>     print(m)
>>>     if type(m) == nn.Linear:
>>>         m.weight.fill_(1.0)
>>>         print(m.weight)
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))
>>> net.apply(init_weights)
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
        [ 1.,  1.]])
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
        [ 1.,  1.]])
Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
```

bfloat16() → T

Casts all floating point parameters and buffers to bfloat16 datatype.

Returns: Module: self

buffers(reuse: bool = True) → Iterator[torch.Tensor]

Returns an iterator over module buffers.

Args:

reuse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer

Example:

```
>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

children() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module

cpu() → T

Moves all model parameters and buffers to the CPU.

Returns: Module: self

cuda(device: Union[int, torch.device, None] = None) → T

Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self

double () → T

Casts all floating point parameters and buffers to double datatype.

Returns: Module: self

dump_patches = False

eval () → T

Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with `self.train(False)`.

Returns: Module: self

extra_repr () → str

Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float () → T

Casts all floating point parameters and buffers to float datatype.

Returns: Module: self

half () → T

Casts all floating point parameters and buffers to half datatype.

Returns: Module: self

load_state_dict (state_dict: Dict[str, torch.Tensor], strict: bool = True)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict` is True, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Arguments:

state_dict (dict): a dict containing parameters and persistent buffers.

strict (bool, optional): whether to strictly enforce that the keys in `state_dict` match the keys returned by this module's `state_dict()` function. Default: True

Returns:

NamedTuple with missing_keys and unexpected_keys fields:

- `missing_keys` is a list of str containing the missing keys

- `unexpected_keys` is a list of str containing the unexpected keys

modules () → Iterator[torch.nn.modules.module.Module]

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, 1 will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

named_buffers (*prefix: str = ”, recurse: bool = True*) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: *prefix* (str): prefix to prepend to all buffer names. *recurse* (bool): if True, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (string, torch.Tensor): Tuple containing the name and buffer

Example:

```
>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())
```

named_children () → Iterator[Tuple[str, torch.nn.modules.module.Module]]

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple containing a name and child module

Example:

```
>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)
```

named_modules (*memo: Optional[Set[Module]] = None, prefix: str = ”*)

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, 1 will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.named_modules()):
    print(idx, '->', m)

0 -> ('', Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
))
```

(continues on next page)

(continued from previous page)

```
) )
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))
```

named_parameters (prefix: str = "", recurse: bool = True) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: prefix (str): prefix to prepend to all parameter names. recurse (bool): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

parameters (recurse: bool = True) → Iterator[torch.nn.parameter.Parameter]

Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

register_backward_hook (hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]) → torch.utils.hooks.RemovableHandle

Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex Module that perform many operations. In some failure cases, grad_input and grad_output will only contain the gradients for a subset of the inputs and outputs. For such Module, you should use `torch.Tensor.register_hook()` directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook(module, grad_input, grad_output) -> Tensor or None
```

The grad_input and grad_output may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that

will be used in place of `grad_input` in subsequent computations. `grad_input` will only correspond to the inputs given as positional arguments.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

`register_buffer(name: str, tensor: torch.Tensor, persistent: bool = True) → None`

Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's `running_mean` is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting `persistent` to `False`. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

`name (string): name of the buffer. The buffer can be accessed` from this module using the given name

`tensor (Tensor): buffer to be registered. persistent (bool): whether the buffer is part of this module's` `state_dict`.

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

`register_forward_hook(hook: Callable[[], None]) → torch.utils.hooks.RemovableHandle`

Registers a forward hook on the module.

The hook will be called every time after `forward()` has computed an output. It should have the following signature:

```
hook(module, input, output) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the output. It can modify the input inplace but it will not have effect on forward since this is called after `forward()` is called.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

`register_forward_pre_hook(hook: Callable[[], None]) → torch.utils.hooks.RemovableHandle`

Registers a forward pre-hook on the module.

The hook will be called every time before `forward()` is invoked. It should have the following signature:

```
hook(module, input) -> None or modified input
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_parameter (*name: str, param: torch.nn.parameter.Parameter*) → None
Adds a parameter to the module.

The parameter can be accessed as an attribute using given name.

Args:

name (string): name of the parameter. The parameter can be accessed from this module using the given name

param (Parameter): parameter to be added to the module.

requires_grad_ (*requires_grad: bool = True*) → T

Change if autograd should record operations on parameters in this module.

This method sets the parameters' `requires_grad` attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

requires_grad (bool): whether autograd should record operations on parameters in this module.

Default: `True`.

Returns: Module: `self`

share_memory () → T

state_dict (*destination=None, prefix=”, keep_vars=False*)

Returns a dictionary containing a whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

dict: a dictionary containing a whole state of the module

Example:

```
>>> module.state_dict().keys()
['bias', 'weight']
```

to (**args, **kwargs*)

Moves and/or casts the parameters and buffers.

This can be called as

to (*device=None, dtype=None, non_blocking=False*)

to (*dtype, non_blocking=False*)

to (*tensor, non_blocking=False*)

to (*memory_format=torch.channels_last*)

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with `dtypes` unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

device (torch.device): the desired device of the parameters and buffers in this module
dtype (torch.dtype): the desired floating point type of the floating point parameters and buffers in this module
tensor (torch.Tensor): Tensor whose dtype and device are the desired dtype and device for all parameters and buffers in this module
memory_format (torch.memory_format): the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]])
>>> linear.to(torch.double)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]], dtype=torch.float64)
>>> gpu1 = torch.device("cuda:1")
>>> linear.to(gpu1, dtype=torch.half, non_blocking=True)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16)
```

train (mode: bool = True) → T

Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (**True**) or evaluation mode (**False**). Default: **True**.

Returns: Module: self

type (*dst_type*: *Union[torch.dtype, str]*) → T
 Casts all parameters and buffers to *dst_type*.

Arguments: *dst_type* (type or string): the desired type

Returns: Module: self

zero_grad() → None
 Sets gradients of all model parameters to zero.

class *flood_forecast.custom.custom_opt.MAPELoss*
 Bases: *torch.nn.modules.module.Module*

Returns MAPE using: target -> True y output -> Prediction by model

forward(*target*: *torch.Tensor*, *output*: *torch.Tensor*)

T_destination = ~**T_destination**

add_module(*name*: str, *module*: *torch.nn.modules.module.Module*) → None
 Adds a child module to the current module.

The module can be accessed as an attribute using the given name.

Args:

- **name (string): name of the child module. The child module can be** accessed from this module using the given name
- **module (Module): child module to be added to the module.**

apply(*fn*: *Callable[[Module], None]*) → T
 Applies *fn* recursively to every submodule (as returned by *.children()*) as well as self. Typical use includes initializing the parameters of a model (see also nn-init-doc).

Args: *fn* (Module -> None): function to be applied to each submodule

Returns: Module: self

Example:

```
>>> @torch.no_grad()
>>> def init_weights(m):
>>>     print(m)
>>>     if type(m) == nn.Linear:
>>>         m.weight.fill_(1.0)
>>>         print(m.weight)
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))
>>> net.apply(init_weights)
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
```

(continues on next page)

(continued from previous page)

```
(1): Linear(in_features=2, out_features=2, bias=True)
)
```

bfloat16() → T

Casts all floating point parameters and buffers to bfloat16 datatype.

Returns: Module: self

buffers(reuse: bool = True) → Iterator[torch.Tensor]

Returns an iterator over module buffers.

Args:

reuse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer

Example:

```
>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

children() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module

cpu() → T

Moves all model parameters and buffers to the CPU.

Returns: Module: self

cuda(device: Union[int, torch.device, None] = None) → T

Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self

double() → T

Casts all floating point parameters and buffers to double datatype.

Returns: Module: self

dump_patches = False**eval()** → T

Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with `self.train(False)`.

Returns: Module: self

extra_repr () → str

Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float () → T

Casts all floating point parameters and buffers to float datatype.

Returns: Module: self

half () → T

Casts all floating point parameters and buffers to half datatype.

Returns: Module: self

load_state_dict (state_dict: Dict[str, torch.Tensor], strict: bool = True)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict` is True, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Arguments:

state_dict (dict): a dict containing parameters and persistent buffers.

strict (bool, optional): whether to strictly enforce that the keys in `state_dict` match the keys returned by this module's `state_dict()` function. Default: True

Returns:

NamedTuple with missing_keys and unexpected_keys fields:

- **missing_keys** is a list of str containing the missing keys
- **unexpected_keys** is a list of str containing the unexpected keys

modules () → Iterator[torch.nn.modules.module.Module]

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

named_buffers (prefix: str = "", recurse: bool = True) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: prefix (str): prefix to prepend to all buffer names. recurse (bool): if True, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (string, torch.Tensor): Tuple containing the name and buffer

Example:

```
>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())
```

named_children() → Iterator[Tuple[str, torch.nn.modules.module.Module]]

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple containing a name and child module

Example:

```
>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)
```

named_modules(memo: Optional[Set[Module]] = None, prefix: str = '')

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.named_modules()):
>>>     print(idx, '->', m)

0 -> ('', Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
))
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))
```

named_parameters(prefix: str = '', recurse: bool = True) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: prefix (str): prefix to prepend to all parameter names. recurse (bool): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

parameters(recurse: bool = True) → Iterator[torch.nn.parameter.Parameter]

Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

register_backward_hook (hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]) → torch.utils.hooks.RemovableHandle

Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex Module that perform many operations. In some failure cases, `grad_input` and `grad_output` will only contain the gradients for a subset of the inputs and outputs. For such Module, you should use `torch.Tensor.register_hook()` directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook(module, grad_input, grad_output) -> Tensor or None
```

The `grad_input` and `grad_output` may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that will be used in place of `grad_input` in subsequent computations. `grad_input` will only correspond to the inputs given as positional arguments.

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_buffer (name: str, tensor: torch.Tensor, persistent: bool = True) → None

Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's `running_mean` is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting `persistent` to `False`. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

name (string): name of the buffer. The buffer can be accessed from this module using the given name

tensor (Tensor): buffer to be registered. persistent (bool): whether the buffer is part of this module's `state_dict`.

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

register_forward_hook (hook: Callable[..., None]) → torch.utils.hooks.RemovableHandle
Registers a forward hook on the module.

The hook will be called every time after `forward()` has computed an output. It should have the following signature:

```
hook(module, input, output) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the output. It can modify the input inplace but it will not have effect on forward since this is called after `forward()` is called.

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_forward_pre_hook (hook: Callable[..., None]) → torch.utils.hooks.RemovableHandle
Registers a forward pre-hook on the module.

The hook will be called every time before `forward()` is invoked. It should have the following signature:

```
hook(module, input) -> None or modified input
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_parameter (name: str, param: torch.nn.parameter.Parameter) → None
Adds a parameter to the module.

The parameter can be accessed as an attribute using given name.

Args:

name (string): name of the parameter. The parameter can be accessed from this module using the given name

param (Parameter): parameter to be added to the module.

requires_grad_ (requires_grad: bool = True) → T
Change if autograd should record operations on parameters in this module.

This method sets the parameters' `requires_grad` attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

requires_grad (bool): whether autograd should record operations on parameters in this module.
Default: True.

Returns: Module: self

share_memory() → T

state_dict (*destination=None*, *prefix=*"", *keep_vars=False*)

Returns a dictionary containing a whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

dict: a dictionary containing a whole state of the module

Example:

```
>>> module.state_dict().keys()
['bias', 'weight']
```

to (*args, **kwargs)

Moves and/or casts the parameters and buffers.

This can be called as

to (*device=None*, *dtype=None*, *non_blocking=False*)

to (*dtype*, *non_blocking=False*)

to (*tensor*, *non_blocking=False*)

to (*memory_format=torch.channels_last*)

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with dtypes unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

device (torch.device): the desired device of the parameters and buffers in this module

dtype (torch.dtype): the desired floating point type of the floating point parameters and buffers in this module

tensor (torch.Tensor): Tensor whose `dtype` and `device` are the desired `dtype` and `device` for all parameters and buffers in this module

memory_format (torch.memory_format): the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
```

(continues on next page)

(continued from previous page)

```

        [-0.5113, -0.2325]])
>>> linear.to(torch.double)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]], dtype=torch.float64)
>>> gpul = torch.device("cuda:1")
>>> linear.to(gpul, dtype=torch.half, non_blocking=True)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16)

```

train(*mode: bool = True*) → T

Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (True) or evaluation mode (False). Default: True.

Returns: Module: self**type**(*dst_type: Union[torch.dtype, str]*) → TCasts all parameters and buffers to *dst_type*.**Arguments:** *dst_type* (type or string): the desired type**Returns:** Module: self**zero_grad()** → None

Sets gradients of all model parameters to zero.

class flood_forecast.custom.custom_opt.**GaussianLoss**(*mu, sigma*)

Bases: torch.nn.modules.module.Module

Compute the negative log likelihood of Gaussian Distribution From <https://arxiv.org/abs/1907.00235>**forward**(*x*)**T_destination** = ~**T_destination****add_module**(*name: str, module: torch.nn.modules.module.Module*) → None

Adds a child module to the current module.

The module can be accessed as an attribute using the given name.

Args:

name (string): name of the child module. The child module can be accessed from this module using the given name

module (Module): child module to be added to the module.

apply (fn: Callable[[Module], None]) → T

Applies fn recursively to every submodule (as returned by .children()) as well as self. Typical use includes initializing the parameters of a model (see also nn-init-doc).

Args: fn (Module → None): function to be applied to each submodule

Returns: Module: self

Example:

```
>>> @torch.no_grad()
>>> def init_weights(m):
>>>     print(m)
>>>     if type(m) == nn.Linear:
>>>         m.weight.fill_(1.0)
>>>         print(m.weight)
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))
>>> net.apply(init_weights)
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
```

bfloat16 () → T

Casts all floating point parameters and buffers to bfloat16 datatype.

Returns: Module: self

buffers (reuse: bool = True) → Iterator[torch.Tensor]

Returns an iterator over module buffers.

Args:

reuse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer

Example:

```
>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

children() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module

cpu() → T

Moves all model parameters and buffers to the CPU.

Returns: Module: self

cuda(device: Union[int, torch.device, None] = None) → T

Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self

double() → T

Casts all floating point parameters and buffers to double datatype.

Returns: Module: self

dump_patches = False

eval() → T

Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with `self.train(False)`.

Returns: Module: self

extra_repr() → str

Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float() → T

Casts all floating point parameters and buffers to float datatype.

Returns: Module: self

half() → T

Casts all floating point parameters and buffers to half datatype.

Returns: Module: self

load_state_dict(state_dict: Dict[str, torch.Tensor], strict: bool = True)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict` is True, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Arguments:

state_dict (dict): a dict containing parameters and persistent buffers.

strict (bool, optional): whether to strictly enforce that the keys in `state_dict` match the keys returned by this module's `state_dict()` function. Default: True

Returns:

NamedTuple with `missing_keys` and `unexpected_keys` fields:

- `missing_keys` is a list of str containing the missing keys
- `unexpected_keys` is a list of str containing the unexpected keys

modules() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

named_buffers (`prefix: str = ''`, `recurse: bool = True`) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: `prefix` (str): prefix to prepend to all buffer names. `recurse` (bool): if True, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (string, torch.Tensor): Tuple containing the name and buffer

Example:

```
>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())
```

named_children() → Iterator[Tuple[str, torch.nn.modules.module.Module]]

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple containing a name and child module

Example:

```
>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)
```

named_modules (`memo: Optional[Set[Module]] = None`, `prefix: str = ''`)

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.named_modules()):
    print(idx, '->', m)

0 -> ('', Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
))
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))
```

named_parameters (*prefix: str = ”, recurse: bool = True*) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: *prefix* (str): prefix to prepend to all parameter names. *recurse* (bool): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

parameters (*recurse: bool = True*) → Iterator[torch.nn.parameter.Parameter]

Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

register_backward_hook (*hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]*) → torch.utils.hooks.RemovableHandle

Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex Module that perform many operations. In some failure cases, `grad_input` and `grad_output` will only contain the gradients for a subset of the inputs and outputs. For such Module, you should use `torch.Tensor.register_hook()` directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook(module, grad_input, grad_output) -> Tensor or None
```

The `grad_input` and `grad_output` may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that will be used in place of `grad_input` in subsequent computations. `grad_input` will only correspond to the inputs given as positional arguments.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

`register_buffer(name: str, tensor: torch.Tensor, persistent: bool = True) → None`

Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's `running_mean` is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting `persistent` to `False`. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

`name (string)`: name of the buffer. The buffer can be accessed from this module using the given name

`tensor (Tensor)`: buffer to be registered. `persistent (bool)`: whether the buffer is part of this module's `state_dict`.

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

`register_forward_hook(hook: Callable[..., None]) → torch.utils.hooks.RemovableHandle`

Registers a forward hook on the module.

The hook will be called every time after `forward()` has computed an output. It should have the following signature:

```
hook(module, input, output) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the output. It can modify the input inplace but it will not have effect on `forward` since this is called after `forward()` is called.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

`register_forward_pre_hook(hook: Callable[..., None]) → torch.utils.hooks.RemovableHandle`

Registers a forward pre-hook on the module.

The hook will be called every time before `forward()` is invoked. It should have the following signature:

```
hook(module, input) -> None or modified input
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

`register_parameter(name: str, param: torch.nn.parameter.Parameter)` → None

Adds a parameter to the module.

The parameter can be accessed as an attribute using given name.

Args:

`name (string): name of the parameter. The parameter can be accessed` from this module using the given name

`param (Parameter): parameter to be added to the module.`

`requires_grad_(requires_grad: bool = True)` → T

Change if autograd should record operations on parameters in this module.

This method sets the parameters' `requires_grad` attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

`requires_grad (bool): whether autograd should record operations on` parameters in this module.

Default: True.

Returns: Module: self

`share_memory()` → T

`state_dict(destination=None, prefix='', keep_vars=False)`

Returns a dictionary containing a whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

`dict: a dictionary containing a whole state of the module`

Example:

```
>>> module.state_dict().keys()  
['bias', 'weight']
```

`to(*args, **kwargs)`

Moves and/or casts the parameters and buffers.

This can be called as

`to(device=None, dtype=None, non_blocking=False)`

`to(dtype, non_blocking=False)`

to (*tensor, non_blocking=False*)
to (*memory_format=torch.channels_last*)

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with dtypes unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

device (torch.device): the desired device of the parameters and buffers in this module
dtype (torch.dtype): the desired floating point type of the floating point parameters and buffers in this module
tensor (torch.Tensor): Tensor whose dtype and device are the desired dtype and device for all parameters and buffers in this module
memory_format (torch.memory_format): the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]])
>>> linear.to(torch.double)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]], dtype=torch.float64)
>>> gpu1 = torch.device("cuda:1")
>>> linear.to(gpu1, dtype=torch.half, non_blocking=True)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16)
```

train (*mode: bool = True*) → T
Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (True) or evaluation mode (False). Default: True.

Returns: Module: self

type (dst_type: Union[torch.dtype, str]) → T

Casts all parameters and buffers to dst_type.

Arguments: dst_type (type or string): the desired type

Returns: Module: self

zero_grad() → None

Sets gradients of all model parameters to zero.

class flood_forecast.custom.custom_opt.**QuantileLoss** (quantiles)

Bases: torch.nn.modules.module.Module

From <https://medium.com/the-artificial-impostor/quantile-regression-part-2-6fdb26b2629>

forward (preds, target)

T_destination = ~T_destination

add_module (name: str, module: torch.nn.modules.module.Module) → None

Adds a child module to the current module.

The module can be accessed as an attribute using the given name.

Args:

name (string): name of the child module. The child module can be accessed from this module using the given name

module (Module): child module to be added to the module.

apply (fn: Callable[[Module], None]) → T

Applies fn recursively to every submodule (as returned by .children()) as well as self. Typical use includes initializing the parameters of a model (see also nn-init-doc).

Args: fn (Module -> None): function to be applied to each submodule

Returns: Module: self

Example:

```
>>> @torch.no_grad()
>>> def init_weights(m):
>>>     print(m)
>>>     if type(m) == nn.Linear:
>>>         m.weight.fill_(1.0)
>>>         print(m.weight)
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))
>>> net.apply(init_weights)
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
        [ 1.,  1.]])
Linear(in_features=2, out_features=2, bias=True)
```

(continues on next page)

(continued from previous page)

```

Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)

```

`bfloat16()` → T

Casts all floating point parameters and buffers to `bfloat16` datatype.

Returns: Module: self

`buffers(reuse: bool = True)` → Iterator[torch.Tensor]

Returns an iterator over module buffers.

Args:

reuse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer

Example:

```

>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)

```

`children()` → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module

`cpu()` → T

Moves all model parameters and buffers to the CPU.

Returns: Module: self

`cuda(device: Union[int, torch.device, None] = None)` → T

Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self

`double()` → T

Casts all floating point parameters and buffers to `double` datatype.

Returns: Module: self

`dump_patches = False`

eval() → T

Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with `self.train(False)`.

Returns: Module: self

extra_repr() → str

Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float() → T

Casts all floating point parameters and buffers to float datatype.

Returns: Module: self

half() → T

Casts all floating point parameters and buffers to half datatype.

Returns: Module: self

load_state_dict(state_dict: Dict[str, torch.Tensor], strict: bool = True)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict` is True, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Arguments:

state_dict (dict): a dict containing parameters and persistent buffers.

strict (bool, optional): whether to strictly enforce that the keys in `state_dict` match the keys returned by this module's `state_dict()` function. Default: True

Returns:

NamedTuple with missing_keys and unexpected_keys fields:

- `missing_keys` is a list of str containing the missing keys
- `unexpected_keys` is a list of str containing the unexpected keys

modules() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
...     print(idx, '->', m)

0 -> Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

named_buffers (*prefix: str = ”, recurse: bool = True*) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: *prefix* (str): prefix to prepend to all buffer names. *recurse* (bool): if True, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (string, torch.Tensor): Tuple containing the name and buffer

Example:

```
>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())
```

named_children () → Iterator[Tuple[str, torch.nn.modules.module.Module]]

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple containing a name and child module

Example:

```
>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)
```

named_modules (*memo: Optional[Set[Module]] = None, prefix: str = ”*)

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, *l* will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.named_modules()):
>>>     print(idx, '->', m)

0 -> ('', Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
))
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))
```

named_parameters (*prefix: str = ”, recurse: bool = True*) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: *prefix* (str): prefix to prepend to all parameter names. *recurse* (bool): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

parameters (*recurse: bool = True*) → Iterator[torch.nn.parameter.Parameter]

Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

register_backward_hook (*hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]*) → torch.utils.hooks.RemovableHandle

Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex Module that perform many operations. In some failure cases, `grad_input` and `grad_output` will only contain the gradients for a subset of the inputs and outputs. For such Module, you should use `torch.Tensor.register_hook()` directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook(module, grad_input, grad_output) -> Tensor or None
```

The `grad_input` and `grad_output` may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that will be used in place of `grad_input` in subsequent computations. `grad_input` will only correspond to the inputs given as positional arguments.

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_buffer (*name: str, tensor: torch.Tensor, persistent: bool = True*) → None

Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's `running_mean` is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting `persistent` to `False`. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

name (string): name of the buffer. The buffer can be accessed from this module using the given name
tensor (Tensor): buffer to be registered. **persistent (bool):** whether the buffer is part of this module's *state_dict*.

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

register_forward_hook (hook: Callable[[], None]) → torch.utils.hooks.RemovableHandle
Registers a forward hook on the module.

The hook will be called every time after *forward()* has computed an output. It should have the following signature:

```
hook(module, input, output) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the *forward*. The hook can modify the output. It can modify the input inplace but it will not have effect on forward since this is called after *forward()* is called.

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling *handle.remove()*

register_forward_pre_hook (hook: Callable[[], None]) → torch.utils.hooks.RemovableHandle
Registers a forward pre-hook on the module.

The hook will be called every time before *forward()* is invoked. It should have the following signature:

```
hook(module, input) -> None or modified input
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the *forward*. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling *handle.remove()*

register_parameter (name: str, param: torch.nn.parameter.Parameter) → None
Adds a parameter to the module.

The parameter can be accessed as an attribute using given name.

Args:

name (string): name of the parameter. The parameter can be accessed from this module using the given name

param (Parameter): parameter to be added to the module.

requires_grad_ (requires_grad: bool = True) → T
Change if autograd should record operations on parameters in this module.

This method sets the parameters' *requires_grad* attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

requires_grad (bool): whether autograd should record operations on parameters in this module.
Default: True.

Returns: Module: self

share_memory () → T

state_dict (destination=None, prefix="", keep_vars=False)
Returns a dictionary containing a whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

dict: a dictionary containing a whole state of the module

Example:

```
>>> module.state_dict().keys()  
['bias', 'weight']
```

to (*args, **kwargs)

Moves and/or casts the parameters and buffers.

This can be called as

to (device=None, dtype=None, non_blocking=False)

to (dtype, non_blocking=False)

to (tensor, non_blocking=False)

to (memory_format=torch.channels_last)

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with dtypes unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

device (torch.device): the desired device of the parameters and buffers in this module

dtype (torch.dtype): the desired floating point type of the floating point parameters and buffers in this module

tensor (torch.Tensor): Tensor whose dtype and device are the desired `dtype` and `device` for all parameters and buffers in this module

memory_format (torch.memory_format): the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]])
>>> linear.to(torch.double)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]], dtype=torch.float64)
>>> gpu1 = torch.device("cuda:1")
>>> linear.to(gpu1, dtype=torch.half, non_blocking=True)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16)
```

train(mode: bool = True) → T

Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (True) or evaluation mode (False). Default: True.

Returns: Module: self

type(dst_type: Union[torch.dtype, str]) → T

Casts all parameters and buffers to dst_type.

Arguments: dst_type (type or string): the desired type

Returns: Module: self

zero_grad() → None

Sets gradients of all model parameters to zero.

```
class flood_forecast.custom.custom_opt.BertAdam(params, lr=<required parameter>, warmup=-1, t_total=-1, schedule='warmup_linear', b1=0.9, b2=0.999, e=1e-06, weight_decay=0.01, max_grad_norm=1.0)
```

Bases: torch.optim.optimizer.Optimizer

Implements BERT version of Adam algorithm with weight decay fix. Params:

lr: learning rate warmup: portion of t_total for the warmup, -1 means no warmup. Default: -1 t_total: total number of training steps for the learning

rate schedule, -1 means constant learning rate. Default: -1

schedule: schedule to use for the warmup (see above). Default: ‘warmup_linear’ b1: Adams b1. Default: 0.9 b2: Adams b2. Default: 0.999 e: Adams epsilon. Default: 1e-6 weight_decay: Weight decay. Default: 0.01 max_grad_norm: Maximum norm for the gradients (-1 means no clipping). Default: 1.0

get_lr () → List[T]

step (closure=None)

Performs a single optimization step. Arguments:

closure (callable, optional): A closure that reevaluates the model and returns the loss.

add_param_group (param_group)

Add a param group to the Optimizer's *param_groups*.

This can be useful when fine tuning a pre-trained network as frozen layers can be made trainable and added to the Optimizer as training progresses.

Arguments: param_group (dict): Specifies what Tensors should be optimized along with group specific optimization options.

load_state_dict (state_dict)

Loads the optimizer state.

Arguments:

state_dict (dict): optimizer state. Should be an object returned from a call to *state_dict ()*.

state_dict ()

Returns the state of the optimizer as a *dict*.

It contains two entries:

- state - a dict holding current optimization state. Its content differs between optimizer classes.
- param_groups - a dict containing all parameter groups

zero_grad()

Clears the gradients of all optimized *torch.Tensor*s.

CHAPTER 19

Dummy Torch Model

A dummy model specifically for unit and integration testing purposes

`class flood_forecast.transformer_xl.dummy_torch.DummyTorchModel (forecast_length: int)`

Bases: `torch.nn.modules.module.Module`

`forward (x: torch.Tensor, mask=None)`

`T_destination = ~T_destination`

`add_module (name: str, module: torch.nn.modules.module.Module) → None`

Adds a child module to the current module.

The module can be accessed as an attribute using the given name.

Args:

`name (string): name of the child module. The child module can be accessed from this module using the given name`

`module (Module): child module to be added to the module.`

`apply (fn: Callable[[Module], None]) → T`

Applies `fn` recursively to every submodule (as returned by `.children()`) as well as `self`. Typical use includes initializing the parameters of a model (see also `nn-init-doc`).

Args: `fn (Module -> None): function to be applied to each submodule`

Returns: `Module: self`

Example:

```
>>> @torch.no_grad()
>>> def init_weights(m):
>>>     print(m)
>>>     if type(m) == nn.Linear:
>>>         m.weight.fill_(1.0)
>>>         print(m.weight)
```

(continues on next page)

(continued from previous page)

```
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))
>>> net.apply(init_weights)
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
```

bfloat16() → T

Casts all floating point parameters and buffers to bfloat16 datatype.

Returns: Module: self

buffers(reuse: bool = True) → Iterator[torch.Tensor]

Returns an iterator over module buffers.

Args:

reuse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer

Example:

```
>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

children() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module

cpu() → T

Moves all model parameters and buffers to the CPU.

Returns: Module: self

cuda(device: Union[int, torch.device, None] = None) → T

Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self

double() → T

Casts all floating point parameters and buffers to `double` datatype.

Returns: Module: self

dump_patches = False

eval() → T

Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with `self.train(False)`.

Returns: Module: self

extra_repr() → str

Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float() → T

Casts all floating point parameters and buffers to `float` datatype.

Returns: Module: self

half() → T

Casts all floating point parameters and buffers to `half` datatype.

Returns: Module: self

load_state_dict(`state_dict: Dict[str, torch.Tensor]`, `strict: bool = True`)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict` is `True`, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Arguments:

state_dict (dict): a dict containing parameters and persistent buffers.

strict (bool, optional): whether to strictly enforce that the keys in `state_dict` match the keys returned by this module's `state_dict()` function. Default: True

Returns:

NamedTuple with `missing_keys` and `unexpected_keys` fields:

- `missing_keys` is a list of str containing the missing keys
- `unexpected_keys` is a list of str containing the unexpected keys

modules() → Iterator[`torch.nn.modules.module.Module`]

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)
```

(continues on next page)

(continued from previous page)

```

0 -> Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)

```

`named_buffers (prefix: str = "", recurse: bool = True) → Iterator[Tuple[str, torch.Tensor]]`

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: prefix (str): prefix to prepend to all buffer names. recurse (bool): if True, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (string, torch.Tensor): Tuple containing the name and buffer

Example:

```

>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())

```

`named_children () → Iterator[Tuple[str, torch.nn.modules.module.Module]]`

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple containing a name and child module

Example:

```

>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)

```

`named_modules (memo: Optional[Set[Module]] = None, prefix: str = "")`

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, 1 will be returned only once.

Example:

```

>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.named_modules()):
>>>     print(idx, '->', m)

0 -> ('', Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
))
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))

```

`named_parameters (prefix: str = "", recurse: bool = True) → Iterator[Tuple[str, torch.Tensor]]`

Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: prefix (str): prefix to prepend to all parameter names. recurse (bool): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

parameters (recurse: bool = True) → Iterator[torch.nn.parameter.Parameter]

Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

register_backward_hook (hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]) → torch.utils.hooks.RemovableHandle

Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex Module that perform many operations. In some failure cases, grad_input and grad_output will only contain the gradients for a subset of the inputs and outputs. For such Module, you should use torch.Tensor.register_hook() directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook(module, grad_input, grad_output) -> Tensor or None
```

The grad_input and grad_output may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that will be used in place of grad_input in subsequent computations. grad_input will only correspond to the inputs given as positional arguments.

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling handle.remove()

register_buffer (name: str, tensor: torch.Tensor, persistent: bool = True) → None

Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's `running_mean` is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting `persistent` to `False`. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

- name (string): name of the buffer. The buffer can be accessed** from this module using the given name
- tensor (Tensor): buffer to be registered. persistent (bool): whether the buffer is part of this module's `state_dict`.**

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

register_forward_hook (hook: Callable[..., None]) → torch.utils.hooks.RemovableHandle

Registers a forward hook on the module.

The hook will be called every time after `forward()` has computed an output. It should have the following signature:

```
hook(module, input, output) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the output. It can modify the input inplace but it will not have effect on forward since this is called after `forward()` is called.

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_forward_pre_hook (hook: Callable[..., None]) → torch.utils.hooks.RemovableHandle

Registers a forward pre-hook on the module.

The hook will be called every time before `forward()` is invoked. It should have the following signature:

```
hook(module, input) -> None or modified input
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_parameter (name: str, param: torch.nn.parameter.Parameter) → None

Adds a parameter to the module.

The parameter can be accessed as an attribute using given name.

Args:

name (string): name of the parameter. The parameter can be accessed from this module using the given name

param (Parameter): parameter to be added to the module.

requires_grad_ (*requires_grad: bool = True*) → T

Change if autograd should record operations on parameters in this module.

This method sets the parameters' `requires_grad` attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

requires_grad (bool): whether autograd should record operations on parameters in this module.

Default: True.

Returns: Module: self

share_memory () → T

state_dict (*destination=None, prefix=”, keep_vars=False*)

Returns a dictionary containing a whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

dict: a dictionary containing a whole state of the module

Example:

```
>>> module.state_dict().keys()
['bias', 'weight']
```

to (*args, **kwargs)

Moves and/or casts the parameters and buffers.

This can be called as

to (device=None, dtype=None, non_blocking=False)

to (dtype, non_blocking=False)

to (tensor, non_blocking=False)

to (memory_format=torch.channels_last)

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with `dtypes` unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

device (torch.device): the desired device of the parameters and buffers in this module

dtype (torch.dtype): the desired floating point type of the floating point parameters and buffers in this module

tensor (torch.Tensor): Tensor whose **dtype** and **device** are the desired **dtype** and **device** for all parameters and buffers in this module

memory_format (torch.memory_format): the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]])
>>> linear.to(torch.double)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]], dtype=torch.float64)
>>> gpul = torch.device("cuda:1")
>>> linear.to(gpul, dtype=torch.half, non_blocking=True)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16)
```

train (mode: bool = True) → T

Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (**True**) or evaluation mode (False). Default: True.

Returns: Module: self

type (dst_type: Union[torch.dtype, str]) → T

Casts all parameters and buffers to dst_type.

Arguments: dst_type (type or string): the desired type

Returns: Module: self

zero_grad() → None

Sets gradients of all model parameters to zero.

CHAPTER 20

Lower Upper Configuration

```
flood_forecast.transformer_xl.lower_upper_config.initial_layer(layer_type: str,  
                                layer_params:  
                                Dict[KT, VT],  
                                layer_number:  
                                int = 1)  
  
flood_forecast.transformer_xl.lower_upper_config.variable_forecast_layer(layer_type,  
                                layer_params)  
  
class flood_forecast.transformer_xl.lower_upper_config.PositionwiseFeedForward(d_in,  
                                d_hid,  
                                dropout=0.1)
```

Bases: torch.nn.modules.module.Module

A two-feed-forward-layer module Take from DSANET

forward (x)

T_destination = ~T_destination

add_module (name: str, module: torch.nn.modules.module.Module) → None

Adds a child module to the current module.

The module can be accessed as an attribute using the given name.

Args:

name (string): name of the child module. The child module can be accessed from this module
using the given name

module (Module): child module to be added to the module.

apply (fn: Callable[[Module], None]) → T

Applies fn recursively to every submodule (as returned by .children()) as well as self. Typical use includes initializing the parameters of a model (see also nn-init-doc).

Args: fn (Module -> None): function to be applied to each submodule

Returns: Module: self

Example:

```
>>> @torch.no_grad()
>>> def init_weights(m):
>>>     print(m)
>>>     if type(m) == nn.Linear:
>>>         m.weight.fill_(1.0)
>>>         print(m.weight)
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))
>>> net.apply(init_weights)
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
        [ 1.,  1.]])
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
        [ 1.,  1.]])
Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
```

bfloat16() → T

Casts all floating point parameters and buffers to bfloat16 datatype.

Returns: Module: self

buffers (*recurse: bool = True*) → Iterator[torch.Tensor]

Returns an iterator over module buffers.

Args:

recurse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer

Example:

```
>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

children() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module

cpu() → T

Moves all model parameters and buffers to the CPU.

Returns: Module: self

cuda (*device: Union[int, torch.device, None] = None*) → T

Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self

double () → T

Casts all floating point parameters and buffers to double datatype.

Returns: Module: self

dump_patches = False

eval () → T

Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with `self.train(False)`.

Returns: Module: self

extra_repr () → str

Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float () → T

Casts all floating point parameters and buffers to float datatype.

Returns: Module: self

half () → T

Casts all floating point parameters and buffers to half datatype.

Returns: Module: self

load_state_dict (state_dict: Dict[str, torch.Tensor], strict: bool = True)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict` is True, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Arguments:

state_dict (dict): a dict containing parameters and persistent buffers.

strict (bool, optional): whether to strictly enforce that the keys in `state_dict` match the keys returned by this module's `state_dict()` function. Default: True

Returns:

NamedTuple with missing_keys and unexpected_keys fields:

- `missing_keys` is a list of str containing the missing keys

- `unexpected_keys` is a list of str containing the unexpected keys

modules () → Iterator[torch.nn.modules.module.Module]

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, 1 will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

named_buffers (*prefix: str = ”, recurse: bool = True*) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: *prefix* (str): prefix to prepend to all buffer names. *recurse* (bool): if True, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (string, torch.Tensor): Tuple containing the name and buffer

Example:

```
>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())
```

named_children () → Iterator[Tuple[str, torch.nn.modules.module.Module]]

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple containing a name and child module

Example:

```
>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)
```

named_modules (*memo: Optional[Set[Module]] = None, prefix: str = ”*)

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, 1 will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.named_modules()):
    print(idx, '->', m)

0 -> ('', Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
```

(continues on next page)

(continued from previous page)

```
) )
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))
```

named_parameters (prefix: str = "", recurse: bool = True) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: prefix (str): prefix to prepend to all parameter names. recurse (bool): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

parameters (recurse: bool = True) → Iterator[torch.nn.parameter.Parameter]

Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

register_backward_hook (hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]) → torch.utils.hooks.RemovableHandle

Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex Module that perform many operations. In some failure cases, grad_input and grad_output will only contain the gradients for a subset of the inputs and outputs. For such Module, you should use `torch.Tensor.register_hook()` directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook(module, grad_input, grad_output) -> Tensor or None
```

The grad_input and grad_output may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that

will be used in place of `grad_input` in subsequent computations. `grad_input` will only correspond to the inputs given as positional arguments.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

`register_buffer(name: str, tensor: torch.Tensor, persistent: bool = True) → None`

Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's `running_mean` is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting `persistent` to `False`. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

`name (string): name of the buffer. The buffer can be accessed` from this module using the given name

`tensor (Tensor): buffer to be registered. persistent (bool): whether the buffer is part of this module's` `state_dict`.

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

`register_forward_hook(hook: Callable[[], None]) → torch.utils.hooks.RemovableHandle`

Registers a forward hook on the module.

The hook will be called every time after `forward()` has computed an output. It should have the following signature:

```
hook(module, input, output) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the output. It can modify the input inplace but it will not have effect on forward since this is called after `forward()` is called.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

`register_forward_pre_hook(hook: Callable[[], None]) → torch.utils.hooks.RemovableHandle`

Registers a forward pre-hook on the module.

The hook will be called every time before `forward()` is invoked. It should have the following signature:

```
hook(module, input) -> None or modified input
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_parameter (*name: str, param: torch.nn.parameter.Parameter*) → None
Adds a parameter to the module.

The parameter can be accessed as an attribute using given name.

Args:

name (string): name of the parameter. The parameter can be accessed from this module using the given name

param (Parameter): parameter to be added to the module.

requires_grad_ (*requires_grad: bool = True*) → T

Change if autograd should record operations on parameters in this module.

This method sets the parameters' `requires_grad` attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

requires_grad (bool): whether autograd should record operations on parameters in this module.

Default: `True`.

Returns: Module: `self`

share_memory () → T

state_dict (*destination=None, prefix=”, keep_vars=False*)

Returns a dictionary containing a whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

dict: a dictionary containing a whole state of the module

Example:

```
>>> module.state_dict().keys()
['bias', 'weight']
```

to (**args, **kwargs*)

Moves and/or casts the parameters and buffers.

This can be called as

to (*device=None, dtype=None, non_blocking=False*)

to (*dtype, non_blocking=False*)

to (*tensor, non_blocking=False*)

to (*memory_format=torch.channels_last*)

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with `dtypes` unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

device (torch.device): the desired device of the parameters and buffers in this module
dtype (torch.dtype): the desired floating point type of the floating point parameters and buffers in this module
tensor (torch.Tensor): Tensor whose dtype and device are the desired dtype and device for all parameters and buffers in this module
memory_format (torch.memory_format): the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]])
>>> linear.to(torch.double)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]], dtype=torch.float64)
>>> gpu1 = torch.device("cuda:1")
>>> linear.to(gpu1, dtype=torch.half, non_blocking=True)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16)
```

train (mode: bool = True) → T

Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (**True**) or evaluation mode (**False**). Default: **True**.

Returns: Module: self

type (*dst_type*: *Union[torch.dtype, str]*) → T
 Casts all parameters and buffers to *dst_type*.

Arguments: *dst_type* (type or string): the desired type

Returns: Module: self

zero_grad() → None
 Sets gradients of all model parameters to zero.

class *flood_forecast.transformer_xl.lower_upper_config.AR*(*window*)
 Bases: *torch.nn.modules.module.Module*

forward(*x*)

T_destination = ~**T_destination**

add_module(*name*: str, *module*: *torch.nn.modules.module.Module*) → None
 Adds a child module to the current module.
 The module can be accessed as an attribute using the given name.

Args:

name (string): name of the child module. The child module can be accessed from this module using the given name

module (Module): child module to be added to the module.

apply(*fn*: *Callable[[Module], None]*) → T
 Applies *fn* recursively to every submodule (as returned by *.children()*) as well as self. Typical use includes initializing the parameters of a model (see also nn-init-doc).

Args: *fn* (Module -> None): function to be applied to each submodule

Returns: Module: self

Example:

```
>>> @torch.no_grad()
>>> def init_weights(m):
>>>     print(m)
>>>     if type(m) == nn.Linear:
>>>         m.weight.fill_(1.0)
>>>         print(m.weight)
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))
>>> net.apply(init_weights)
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
```

bfloat16() → T
Casts all floating point parameters and buffers to bfloat16 datatype.

Returns: Module: self

buffers (*reuse: bool = True*) → Iterator[torch.Tensor]
Returns an iterator over module buffers.

Args:

reuse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer

Example:

```
>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

children() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module

cpu() → T
Moves all model parameters and buffers to the CPU.

Returns: Module: self

cuda (*device: Union[int, torch.device, None] = None*) → T
Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self

double() → T
Casts all floating point parameters and buffers to double datatype.

Returns: Module: self

dump_patches = False

eval() → T
Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with `self.train(False)`.

Returns: Module: self

extra_repr() → str
Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float() → T

Casts all floating point parameters and buffers to float datatype.

Returns: Module: self

half() → T

Casts all floating point parameters and buffers to half datatype.

Returns: Module: self

load_state_dict(state_dict: Dict[str, torch.Tensor], strict: bool = True)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict` is True, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Arguments:

state_dict (dict): a dict containing parameters and persistent buffers.

strict (bool, optional): whether to strictly enforce that the keys in `state_dict` match the keys returned by this module's `state_dict()` function. Default: True

Returns:

NamedTuple with missing_keys and unexpected_keys fields:

- **missing_keys** is a list of str containing the missing keys
- **unexpected_keys** is a list of str containing the unexpected keys

modules() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, 1 will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

named_buffers(prefix: str = "", recurse: bool = True) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: prefix (str): prefix to prepend to all buffer names. recurse (bool): if True, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (string, torch.Tensor): Tuple containing the name and buffer

Example:

```
>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())
```

named_children() → Iterator[Tuple[str, torch.nn.modules.module.Module]]

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple containing a name and child module

Example:

```
>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)
```

named_modules(memo: Optional[Set[Module]] = None, prefix: str = '')

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.named_modules()):
>>>     print(idx, '->', m)

0 -> ('', Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
))
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))
```

named_parameters(prefix: str = '', recurse: bool = True) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: prefix (str): prefix to prepend to all parameter names. recurse (bool): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

parameters(recurse: bool = True) → Iterator[torch.nn.parameter.Parameter]

Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

register_backward_hook (hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]) → torch.utils.hooks.RemovableHandle

Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex Module that perform many operations. In some failure cases, grad_input and grad_output will only contain the gradients for a subset of the inputs and outputs. For such Module, you should use torch.Tensor.register_hook () directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook(module, grad_input, grad_output) -> Tensor or None
```

The grad_input and grad_output may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that will be used in place of grad_input in subsequent computations. grad_input will only correspond to the inputs given as positional arguments.

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling handle.remove ()

register_buffer (name: str, tensor: torch.Tensor, persistent: bool = True) → None
Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's running_mean is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting persistent to False. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

name (string): name of the buffer. The buffer can be accessed from this module using the given name

tensor (Tensor): buffer to be registered. **persistent (bool):** whether the buffer is part of this module's `state_dict`.

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

register_forward_hook (hook: Callable[..., None]) → torch.utils.hooks.RemovableHandle
Registers a forward hook on the module.

The hook will be called every time after `forward()` has computed an output. It should have the following signature:

```
hook(module, input, output) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the output. It can modify the input inplace but it will not have effect on forward since this is called after `forward()` is called.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

register_forward_pre_hook (hook: Callable[..., None]) → torch.utils.hooks.RemovableHandle
Registers a forward pre-hook on the module.

The hook will be called every time before `forward()` is invoked. It should have the following signature:

```
hook(module, input) -> None or modified input
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

register_parameter (name: str, param: torch.nn.Parameter) → None
Adds a parameter to the module.

The parameter can be accessed as an attribute using given name.

Args:

name (string): name of the parameter. The parameter can be accessed from this module using the given name

param (Parameter): parameter to be added to the module.

requires_grad_ (requires_grad: bool = True) → T
Change if autograd should record operations on parameters in this module.

This method sets the parameters' `requires_grad` attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

requires_grad (bool): whether autograd should record operations on parameters in this module.
Default: True.

Returns: Module: self

share_memory() → T

state_dict (*destination=None*, *prefix=*"", *keep_vars=False*)

Returns a dictionary containing a whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

dict: a dictionary containing a whole state of the module

Example:

```
>>> module.state_dict().keys()
['bias', 'weight']
```

to (*args, **kwargs)

Moves and/or casts the parameters and buffers.

This can be called as

to (*device=None*, *dtype=None*, *non_blocking=False*)

to (*dtype*, *non_blocking=False*)

to (*tensor*, *non_blocking=False*)

to (*memory_format=torch.channels_last*)

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with dtypes unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

device (`torch.device`): the desired device of the parameters and buffers in this module

dtype (`torch.dtype`): the desired floating point type of the floating point parameters and buffers in this module

tensor (`torch.Tensor`): Tensor whose `dtype` and `device` are the desired `dtype` and `device` for all parameters and buffers in this module

memory_format (`torch.memory_format`): the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
```

(continues on next page)

(continued from previous page)

```

        [-0.5113, -0.2325]])
>>> linear.to(torch.double)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]], dtype=torch.float64)
>>> gpul = torch.device("cuda:1")
>>> linear.to(gpul, dtype=torch.half, non_blocking=True)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16)

```

train(*mode: bool = True*) → T

Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (True) or evaluation mode (False). Default: True.

Returns: Module: self**type**(*dst_type: Union[torch.dtype, str]*) → T

Casts all parameters and buffers to dst_type.

Arguments: dst_type (type or string): the desired type**Returns:** Module: self**zero_grad()** → None

Sets gradients of all model parameters to zero.

```

class flood_forecast.transformer_xl.lower_upper_config.MetaEmbedding(meta_vector_dim,
out-
put_dim,
predic-
tor_number,
predic-
tor_order)

```

Bases: torch.nn.modules.module.Module

T_destination = ~T_destination**add_module**(*name: str, module: torch.nn.modules.module.Module*) → None

Adds a child module to the current module.

The module can be accessed as an attribute using the given name.

Args:

name (string): name of the child module. The child module can be accessed from this module using the given name

module (Module): child module to be added to the module.

apply (fn: Callable[[Module], None]) → T

Applies fn recursively to every submodule (as returned by .children()) as well as self. Typical use includes initializing the parameters of a model (see also nn-init-doc).

Args: fn (Module -> None): function to be applied to each submodule

Returns: Module: self

Example:

```
>>> @torch.no_grad()
>>> def init_weights(m):
>>>     print(m)
>>>     if type(m) == nn.Linear:
>>>         m.weight.fill_(1.0)
>>>         print(m.weight)
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))
>>> net.apply(init_weights)
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
```

bfloat16 () → T

Casts all floating point parameters and buffers to bfloat16 datatype.

Returns: Module: self

buffers (reuse: bool = True) → Iterator[torch.Tensor]

Returns an iterator over module buffers.

Args:

reuse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer

Example:

```
>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
```

(continues on next page)

(continued from previous page)

```
<class 'torch.Tensor'> (20L, )
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

children() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module

cpu() → T

Moves all model parameters and buffers to the CPU.

Returns: Module: self

cuda(device: Union[int, torch.device, None] = None) → T

Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self

double() → T

Casts all floating point parameters and buffers to double datatype.

Returns: Module: self

dump_patches = False

eval() → T

Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with `self.train(False)`.

Returns: Module: self

extra_repr() → str

Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float() → T

Casts all floating point parameters and buffers to float datatype.

Returns: Module: self

forward(*input) → None

half() → T

Casts all floating point parameters and buffers to half datatype.

Returns: Module: self

load_state_dict(state_dict: Dict[str, torch.Tensor], strict: bool = True)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict` is `True`, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Arguments:

state_dict (dict): a dict containing parameters and persistent buffers.
strict (bool, optional): whether to strictly enforce that the keys in `state_dict` match the keys returned by this module's `state_dict()` function. Default: True

Returns:

NamedTuple with missing_keys and unexpected_keys fields:

- **missing_keys** is a list of str containing the missing keys
- **unexpected_keys** is a list of str containing the unexpected keys

modules () → Iterator[torch.nn.modules.module.Module]

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, 1 will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

named_buffers (prefix: str = "", recurse: bool = True) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: prefix (str): prefix to prepend to all buffer names. recurse (bool): if True, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (string, torch.Tensor): Tuple containing the name and buffer

Example:

```
>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())
```

named_children () → Iterator[Tuple[str, torch.nn.modules.module.Module]]

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple containing a name and child module

Example:

```
>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)
```

`named_modules` (*memo: Optional[Set[Module]] = None, prefix: str = ''*)

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, 1 will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.named_modules()):
    print(idx, '->', m)

0 -> ('', Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
))
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))
```

`named_parameters` (*prefix: str = '', recurse: bool = True*) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: prefix (str): prefix to prepend to all parameter names. recurse (bool): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

`parameters` (*recurse: bool = True*) → Iterator[torch.nn.parameter.Parameter]

Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

`register_backward_hook` (*hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]*) → torch.utils.hooks.RemovableHandle

Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex Module that perform many operations. In some failure cases, `grad_input` and `grad_output` will only contain the gradients for a subset of the inputs and outputs. For such Module, you should use `torch.Tensor.register_hook()` directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook(module, grad_input, grad_output) -> Tensor or None
```

The `grad_input` and `grad_output` may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that will be used in place of `grad_input` in subsequent computations. `grad_input` will only correspond to the inputs given as positional arguments.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

register_buffer(`name: str, tensor: torch.Tensor, persistent: bool = True`) → `None`
Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's `running_mean` is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting `persistent` to `False`. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

`name (string)`: name of the buffer. The buffer can be accessed from this module using the given name

`tensor (Tensor)`: buffer to be registered. `persistent (bool)`: whether the buffer is part of this module's `state_dict`.

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

register_forward_hook(`hook: Callable[[], None]`) → `torch.utils.hooks.RemovableHandle`
Registers a forward hook on the module.

The hook will be called every time after `forward()` has computed an output. It should have the following signature:

```
hook(module, input, output) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the output. It can modify the input inplace but it will not have effect on forward since this is called after `forward()` is called.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

register_forward_pre_hook (*hook*: *Callable*[...], *None*) → *torch.utils.hooks.RemovableHandle*
Registers a forward pre-hook on the module.

The hook will be called every time before *forward()* is invoked. It should have the following signature:

```
hook (module, input) -> None or modified input
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the *forward*. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling *handle.remove()*

register_parameter (*name*: str, *param*: *torch.nn.Parameter*) → None
Adds a parameter to the module.

The parameter can be accessed as an attribute using given name.

Args:

name (string): name of the parameter. The parameter can be accessed from this module using the given name

param (Parameter): parameter to be added to the module.

requires_grad_ (*requires_grad*: bool = True) → T
Change if autograd should record operations on parameters in this module.

This method sets the parameters' *requires_grad* attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

requires_grad (bool): whether autograd should record operations on parameters in this module.
Default: True.

Returns: Module: self

share_memory_ () → T

state_dict (*destination*=None, *prefix*=",", *keep_vars*=False)
Returns a dictionary containing a whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

dict: a dictionary containing a whole state of the module

Example:

```
>>> module.state_dict().keys()
['bias', 'weight']
```

to (**args*, ***kwargs*)
Moves and/or casts the parameters and buffers.

This can be called as

```
to(device=None, dtype=None, non_blocking=False)
to(dtype, non_blocking=False)
to(tensor, non_blocking=False)
to(memory_format=torch.channels_last)
```

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with `dtypes` unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

- device (torch.device):** the desired device of the parameters and buffers in this module
- dtype (torch.dtype):** the desired floating point type of the floating point parameters and buffers in this module
- tensor (torch.Tensor):** Tensor whose `dtype` and `device` are the desired `dtype` and `device` for all parameters and buffers in this module
- memory_format (torch.memory_format):** the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]])
>>> linear.to(torch.double)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]], dtype=torch.float64)
>>> gpu1 = torch.device("cuda:1")
>>> linear.to(gpu1, dtype=torch.half, non_blocking=True)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
```

(continues on next page)

(continued from previous page)

```
Parameter containing:  
tensor([[ 0.1914, -0.3420],  
       [-0.5112, -0.2324]], dtype=torch.float16)
```

train (*mode: bool = True*) → T

Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (True) or evaluation mode (False). Default: True.

Returns: Module: self

type (*dst_type: Union[torch.dtype, str]*) → T

Casts all parameters and buffers to *dst_type*.

Arguments: *dst_type* (type or string): the desired type

Returns: Module: self

zero_grad() → None

Sets gradients of all model parameters to zero.

CHAPTER 21

Simple Multi Attention Head Model

```
class flood_forecast.transformer_xl.multi_head_base.MultiAttnHeadSimple(number_time_series:  
    int,  
    seq_len=10,  
    out-  
    put_seq_len=None,  
    d_model=128,  
    num_heads=8,  
    fore-  
    cast_length=None,  
    dropout=0.1,  
    sig-  
    moid=False)
```

Bases: torch.nn.modules.module.Module

A simple multi-head attention model inspired by Vaswani et al.

T_destination = ~T_destination

add_module (name: str, module: torch.nn.modules.module.Module) → None

Adds a child module to the current module.

The module can be accessed as an attribute using the given name.

Args:

name (string): name of the child module. The child module can be accessed from this module
using the given name

module (Module): child module to be added to the module.

apply (fn: Callable[[Module], None]) → T

Applies fn recursively to every submodule (as returned by .children()) as well as self. Typical use includes initializing the parameters of a model (see also nn-init-doc).

Args: fn (Module -> None): function to be applied to each submodule

Returns: Module: self

Example:

```
>>> @torch.no_grad()
>>> def init_weights(m):
>>>     print(m)
>>>     if type(m) == nn.Linear:
>>>         m.weight.fill_(1.0)
>>>         print(m.weight)
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))
>>> net.apply(init_weights)
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
        [ 1.,  1.]])
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
        [ 1.,  1.]])
Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
```

bfloat16() → T

Casts all floating point parameters and buffers to bfloat16 datatype.

Returns: Module: self

buffers(reuse: bool = True) → Iterator[torch.Tensor]

Returns an iterator over module buffers.

Args:

reuse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer

Example:

```
>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

children() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module

cpu() → T

Moves all model parameters and buffers to the CPU.

Returns: Module: self

cuda(device: Union[int, torch.device, None] = None) → T

Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self

double () → T

Casts all floating point parameters and buffers to double datatype.

Returns: Module: self

dump_patches = False

eval () → T

Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with `self.train(False)`.

Returns: Module: self

extra_repr () → str

Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float () → T

Casts all floating point parameters and buffers to float datatype.

Returns: Module: self

forward (x: torch.Tensor, mask=None) → torch.Tensor

Parameters `torch.Tensor (x)` – of shape (B, L, M)

Where B is the batch size, L is the sequence length and M is the number of time :returns a tensor of dimension (B, forecast_length)

half () → T

Casts all floating point parameters and buffers to half datatype.

Returns: Module: self

load_state_dict (state_dict: Dict[str; torch.Tensor], strict: bool = True)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict` is `True`, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Arguments:

state_dict (dict): a dict containing parameters and persistent buffers.

strict (bool, optional): whether to strictly enforce that the keys in `state_dict` match the keys returned by this module's `state_dict()` function. Default: `True`

Returns:

NamedTuple with missing_keys and unexpected_keys fields:

- `missing_keys` is a list of str containing the missing keys

- **unexpected_keys** is a list of str containing the unexpected keys

modules () → Iterator[torch.nn.modules.module.Module]

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, 1 will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

named_buffers (prefix: str = "", recurse: bool = True) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: prefix (str): prefix to prepend to all buffer names. recurse (bool): if True, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (string, torch.Tensor): Tuple containing the name and buffer

Example:

```
>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())
```

named_children () → Iterator[Tuple[str, torch.nn.modules.module.Module]]

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple containing a name and child module

Example:

```
>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)
```

named_modules (memo: Optional[Set[Module]] = None, prefix: str = "")

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, 1 will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.named_modules()):
    print(idx, '->', m)

0 -> ('', Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
))
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))
```

named_parameters (prefix: str = "", recurse: bool = True) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: prefix (str): prefix to prepend to all parameter names. recurse (bool): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

parameters (recurse: bool = True) → Iterator[torch.nn.parameter.Parameter]

Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

register_backward_hook (hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]) → torch.utils.hooks.RemovableHandle

Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex Module that perform many operations. In some failure cases, grad_input and grad_output will only contain the gradients for a subset of the inputs and outputs. For such Module, you should use torch.Tensor.register_hook() directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook(module, grad_input, grad_output) -> Tensor or None
```

The `grad_input` and `grad_output` may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that will be used in place of `grad_input` in subsequent computations. `grad_input` will only correspond to the inputs given as positional arguments.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

`register_buffer(name: str, tensor: torch.Tensor, persistent: bool = True) → None`

Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's `running_mean` is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting `persistent` to `False`. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

`name (string)`: name of the buffer. The buffer can be accessed from this module using the given name

`tensor (Tensor)`: buffer to be registered. `persistent (bool)`: whether the buffer is part of this module's `state_dict`.

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

`register_forward_hook(hook: Callable[..., None]) → torch.utils.hooks.RemovableHandle`

Registers a forward hook on the module.

The hook will be called every time after `forward()` has computed an output. It should have the following signature:

```
hook(module, input, output) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the output. It can modify the input inplace but it will not have effect on `forward` since this is called after `forward()` is called.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

`register_forward_pre_hook(hook: Callable[..., None]) → torch.utils.hooks.RemovableHandle`

Registers a forward pre-hook on the module.

The hook will be called every time before `forward()` is invoked. It should have the following signature:

```
hook(module, input) -> None or modified input
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

`register_parameter(name: str, param: torch.nn.parameter.Parameter)` → None

Adds a parameter to the module.

The parameter can be accessed as an attribute using given name.

Args:

`name (string): name of the parameter. The parameter can be accessed` from this module using the given name

`param (Parameter): parameter to be added to the module.`

`requires_grad_(requires_grad: bool = True)` → T

Change if autograd should record operations on parameters in this module.

This method sets the parameters' `requires_grad` attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

`requires_grad (bool): whether autograd should record operations on` parameters in this module.

Default: True.

Returns: Module: self

`share_memory()` → T

`state_dict(destination=None, prefix= "", keep_vars=False)`

Returns a dictionary containing a whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

`dict: a dictionary containing a whole state of the module`

Example:

```
>>> module.state_dict().keys()
['bias', 'weight']
```

`to(*args, **kwargs)`

Moves and/or casts the parameters and buffers.

This can be called as

`to(device=None, dtype=None, non_blocking=False)`

`to(dtype, non_blocking=False)`

```
to (tensor, non_blocking=False)
to (memory_format=torch.channels_last)
```

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with `dtypes` unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

device (torch.device): the desired device of the parameters and buffers in this module
dtype (torch.dtype): the desired floating point type of the floating point parameters and buffers in this module
tensor (torch.Tensor): Tensor whose `dtype` and `device` are the desired `dtype` and `device` for all parameters and buffers in this module
memory_format (torch.memory_format): the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]])
>>> linear.to(torch.double)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]], dtype=torch.float64)
>>> gpu1 = torch.device("cuda:1")
>>> linear.to(gpu1, dtype=torch.half, non_blocking=True)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16)
```

train (mode: bool = True) → T

Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (True) or evaluation mode (False). Default: True.

Returns: Module: self

type (dst_type: Union[torch.dtype, str]) → T

Casts all parameters and buffers to dst_type.

Arguments: dst_type (type or string): the desired type

Returns: Module: self

zero_grad() → None

Sets gradients of all model parameters to zero.

CHAPTER 22

Basic Transformer

```
class flood_forecast.transformer_xl.transformer_basic.SimpleTransformer(number_time_series:  
    int,  
    seq_length:  
    int  
    =  
    48,  
    out-  
    put_seq_len:  
    int  
    =  
    None,  
    d_model:  
    int  
    =  
    128,  
    n_heads:  
    int  
    =  
    8,  
    dropout=0.1,  
    for-  
    ward_dim=2048,  
    sig-  
    moid=False)
```

Bases: torch.nn.modules.module.Module

Full transformer model

```
forward(x: torch.Tensor, t: torch.Tensor, tgt_mask=None, src_mask=None)  
basic_feature(x: torch.Tensor)  
encode_sequence(x, src_mask=None)  
decode_seq(mem, t, tgt_mask=None, view_number=None) → torch.Tensor
```

T_destination = ~T_destination
add_module(name: str, module: torch.nn.modules.module.Module) → None
 Adds a child module to the current module.

The module can be accessed as an attribute using the given name.

Args:

name (string): name of the child module. The child module can be accessed from this module using the given name

module (Module): child module to be added to the module.

apply(fn: Callable[[Module], None]) → T

Applies fn recursively to every submodule (as returned by .children()) as well as self. Typical use includes initializing the parameters of a model (see also nn-init-doc).

Args: fn (Module -> None): function to be applied to each submodule

Returns: Module: self

Example:

```
>>> @torch.no_grad()
>>> def init_weights(m):
>>>     print(m)
>>>     if type(m) == nn.Linear:
>>>         m.weight.fill_(1.0)
>>>         print(m.weight)
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))
>>> net.apply(init_weights)
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
```

bfloat16() → T

Casts all floating point parameters and buffers to bfloat16 datatype.

Returns: Module: self

buffers(reuse: bool = True) → Iterator[torch.Tensor]

Returns an iterator over module buffers.

Args:

reuse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer

Example:

```
>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

children() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module

cpu() → T

Moves all model parameters and buffers to the CPU.

Returns: Module: self

cuda(device: Union[int, torch.device, None] = None) → T

Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self

double() → T

Casts all floating point parameters and buffers to double datatype.

Returns: Module: self

dump_patches = False

eval() → T

Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with `self.train(False)`.

Returns: Module: self

extra_repr() → str

Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float() → T

Casts all floating point parameters and buffers to float datatype.

Returns: Module: self

half() → T

Casts all floating point parameters and buffers to half datatype.

Returns: Module: self

load_state_dict(state_dict: Dict[str, torch.Tensor], strict: bool = True)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict`

is True, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Arguments:

- **state_dict (dict):** a dict containing parameters and persistent buffers.
- **strict (bool, optional):** whether to strictly enforce that the keys in `state_dict` match the keys returned by this module's `state_dict()` function. Default: True

Returns:

NamedTuple with missing_keys and unexpected_keys fields:

- **missing_keys** is a list of str containing the missing keys
- **unexpected_keys** is a list of str containing the unexpected keys

modules () → Iterator[torch.nn.modules.module.Module]

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

named_buffers (prefix: str = "", recurse: bool = True) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: prefix (str): prefix to prepend to all buffer names. recurse (bool): if True, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (string, torch.Tensor): Tuple containing the name and buffer

Example:

```
>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())
```

named_children () → Iterator[Tuple[str, torch.nn.modules.module.Module]]

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple containing a name and child module

Example:

```
>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)
```

named_modules (memo: *Optional[Set[Module]]* = None, prefix: *str* = '')

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.named_modules()):
>>>     print(idx, '->', m)

0 -> ('', Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
))
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))
```

named_parameters (prefix: *str* = '', recurse: *bool* = True) → Iterator[Tuple[*str*, *torch.Tensor*]]

Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: prefix (*str*): prefix to prepend to all parameter names. recurse (*bool*): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

parameters (recurse: *bool* = True) → Iterator[*torch.nn.parameter.Parameter*]

Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

```
register_backward_hook(hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]) → torch.utils.hooks.RemovableHandle
```

Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex `Module` that perform many operations. In some failure cases, `grad_input` and `grad_output` will only contain the gradients for a subset of the inputs and outputs. For such `Module`, you should use `torch.Tensor.register_hook()` directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook(module, grad_input, grad_output) -> Tensor or None
```

The `grad_input` and `grad_output` may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that will be used in place of `grad_input` in subsequent computations. `grad_input` will only correspond to the inputs given as positional arguments.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

```
register_buffer(name: str, tensor: torch.Tensor, persistent: bool = True) → None
```

Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's `running_mean` is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting `persistent` to `False`. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

`name (string): name of the buffer. The buffer can be accessed` from this module using the given name

`tensor (Tensor): buffer to be registered. persistent (bool): whether the buffer is part of this module's` `state_dict`.

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

```
register_forward_hook(hook: Callable[[...], None]) → torch.utils.hooks.RemovableHandle
```

Registers a forward hook on the module.

The hook will be called every time after `forward()` has computed an output. It should have the following signature:

```
hook(module, input, output) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the output. It can modify the input inplace but it will not have effect on forward since this is called after `forward()` is called.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

<code>register_forward_pre_hook</code> (<code>hook: Callable[[], None]</code>) →	<code>torch.utils.hooks.RemovableHandle</code>
--	--

Registers a forward pre-hook on the module.

The hook will be called every time before `forward()` is invoked. It should have the following signature:

<code>hook(module, input) -> None or modified input</code>

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

<code>register_parameter</code> (<code>name: str, param: torch.nn.parameter.Parameter</code>) → <code>None</code>

Adds a parameter to the module.

The parameter can be accessed as an attribute using given name.

Args:

name (string): name of the parameter. The parameter can be accessed from this module using the given name

`param` (`Parameter`): parameter to be added to the module.

<code>requires_grad_</code> (<code>requires_grad: bool = True</code>) → <code>T</code>
--

Change if autograd should record operations on parameters in this module.

This method sets the parameters' `requires_grad` attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

requires_grad (bool): whether autograd should record operations on parameters in this module.
Default: `True`.

Returns: `Module: self`

`share_memory()` → `T`

<code>state_dict</code> (<code>destination=None, prefix='', keep_vars=False</code>)

Returns a dictionary containing a whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

`dict`: a dictionary containing a whole state of the module

Example:

```
>>> module.state_dict().keys()  
['bias', 'weight']
```

to(*args, **kwargs)

Moves and/or casts the parameters and buffers.

This can be called as

to(device=None, dtype=None, non_blocking=False)

to(dtype, non_blocking=False)

to(tensor, non_blocking=False)

to(memory_format=torch.channels_last)

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with dtypes unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

device (torch.device): the desired device of the parameters and buffers in this module

dtype (torch.dtype): the desired floating point type of the floating point parameters and buffers in this module

tensor (torch.Tensor): Tensor whose `dtype` and `device` are the desired `dtype` and `device` for all parameters and buffers in this module

memory_format (torch.memory_format): the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)  
>>> linear.weight  
Parameter containing:  
tensor([[ 0.1913, -0.3420],  
       [-0.5113, -0.2325]])  
>>> linear.to(torch.double)  
Linear(in_features=2, out_features=2, bias=True)  
>>> linear.weight  
Parameter containing:  
tensor([[ 0.1913, -0.3420],  
       [-0.5113, -0.2325]], dtype=torch.float64)  
>>> gpu1 = torch.device("cuda:1")  
>>> linear.to(gpu1, dtype=torch.half, non_blocking=True)  
Linear(in_features=2, out_features=2, bias=True)
```

(continues on next page)

(continued from previous page)

```
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16)
```

train(mode: bool = True) → T

Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (True) or evaluation mode (False). Default: True.

Returns: Module: self**type(dst_type: Union[torch.dtype, str]) → T**

Casts all parameters and buffers to dst_type.

Arguments: dst_type (type or string): the desired type

Returns: Module: self**zero_grad() → None**

Sets gradients of all model parameters to zero.

```
class flood_forecast.transformer_xl.transformer_basic.CustomTransformerDecoder(seq_length:
                                int,
                                out-
                                put_seq_length:
                                int,
                                n_time_series:
                                int,
                                d_model=128,
                                out-
                                put_dim=1,
                                n_layers_encode
                                for-
                                ward_dim=2048,
                                dropout=0.1,
                                use_mask=False,
                                n_heads=8)
```

Bases: torch.nn.modules.module.Module

Uses a number of encoder layers with simple linear decoder layer

forward(x: torch.Tensor) → torch.Tensor

Performs forward pass on tensor of (batch_size, sequence_length, n_time_series) Return tensor of dim (batch_size, output_seq_length)

T_destination = ~T_destination

add_module (*name: str, module: torch.nn.modules.module.Module*) → None

Adds a child module to the current module.

The module can be accessed as an attribute using the given name.

Args:

name (string): name of the child module. The child module can be accessed from this module using the given name

module (Module): child module to be added to the module.

apply (*fn: Callable[[Module], None]*) → T

Applies fn recursively to every submodule (as returned by .children()) as well as self. Typical use includes initializing the parameters of a model (see also nn-init-doc).

Args: fn (Module -> None): function to be applied to each submodule

Returns: Module: self

Example:

```
>>> @torch.no_grad()
>>> def init_weights(m):
>>>     print(m)
>>>     if type(m) == nn.Linear:
>>>         m.weight.fill_(1.0)
>>>         print(m.weight)
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))
>>> net.apply(init_weights)
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
```

bfloat16() → T

Casts all floating point parameters and buffers to bfloat16 datatype.

Returns: Module: self

buffers (*re recurse: bool = True*) → Iterator[torch.Tensor]

Returns an iterator over module buffers.

Args:

recurse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer

Example:

```
>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

children() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module

cpu() → T

Moves all model parameters and buffers to the CPU.

Returns: Module: self

cuda(device: Union[int, torch.device, None] = None) → T

Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self

double() → T

Casts all floating point parameters and buffers to double datatype.

Returns: Module: self

dump_patches = False

eval() → T

Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with `self.train(False)`.

Returns: Module: self

extra_repr() → str

Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float() → T

Casts all floating point parameters and buffers to float datatype.

Returns: Module: self

half() → T

Casts all floating point parameters and buffers to half datatype.

Returns: Module: self

load_state_dict(state_dict: Dict[str, torch.Tensor], strict: bool = True)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict` is True, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Arguments:

- **state_dict (dict):** a dict containing parameters and persistent buffers.
- **strict (bool, optional):** whether to strictly enforce that the keys in `state_dict` match the keys returned by this module's `state_dict()` function. Default: True

Returns:

NamedTuple with missing_keys and unexpected_keys fields:

- **missing_keys** is a list of str containing the missing keys
- **unexpected_keys** is a list of str containing the unexpected keys

modules () → Iterator[torch.nn.modules.module.Module]

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, 1 will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

named_buffers (prefix: str = "", recurse: bool = True) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: prefix (str): prefix to prepend to all buffer names. recurse (bool): if True, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (string, torch.Tensor): Tuple containing the name and buffer

Example:

```
>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())
```

named_children () → Iterator[Tuple[str, torch.nn.modules.module.Module]]

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple containing a name and child module

Example:

```
>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)
```

named_modules (*memo: Optional[Set[Module]] = None, prefix: str = ''*)

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, 1 will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.named_modules()):
    print(idx, '->', m)

0 -> ('', Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
))
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))
```

named_parameters (*prefix: str = '', recurse: bool = True*) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: prefix (str): prefix to prepend to all parameter names. recurse (bool): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

parameters (*recurse: bool = True*) → Iterator[torch.nn.parameter.Parameter]

Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

register_backward_hook (*hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]*) → torch.utils.hooks.RemovableHandle

Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex Module that perform many operations. In some failure cases, `grad_input` and `grad_output` will only contain the gradients for a subset of the inputs and outputs. For such Module, you should use `torch.Tensor.register_hook()` directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook(module, grad_input, grad_output) -> Tensor or None
```

The `grad_input` and `grad_output` may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that will be used in place of `grad_input` in subsequent computations. `grad_input` will only correspond to the inputs given as positional arguments.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

register_buffer(`name: str, tensor: torch.Tensor, persistent: bool = True`) → `None`
Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's `running_mean` is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting `persistent` to `False`. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

`name (string): name of the buffer. The buffer can be accessed` from this module using the given name

`tensor (Tensor): buffer to be registered. persistent (bool): whether the buffer is part of this module's` `state_dict`.

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

register_forward_hook(`hook: Callable[[], None]`) → `torch.utils.hooks.RemovableHandle`
Registers a forward hook on the module.

The hook will be called every time after `forward()` has computed an output. It should have the following signature:

```
hook(module, input, output) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the output. It can modify the input inplace but it will not have effect on `forward` since this is called after `forward()` is called.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

register_forward_pre_hook (*hook*: *Callable*[...], *None*) → *torch.utils.hooks.RemovableHandle*
Registers a forward pre-hook on the module.

The hook will be called every time before *forward()* is invoked. It should have the following signature:

```
hook (module, input) -> None or modified input
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the *forward*. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling *handle.remove()*

register_parameter (*name*: str, *param*: *torch.nn.Parameter*) → None
Adds a parameter to the module.

The parameter can be accessed as an attribute using given name.

Args:

name (string): name of the parameter. The parameter can be accessed from this module using the given name

param (Parameter): parameter to be added to the module.

requires_grad_ (*requires_grad*: bool = True) → T
Change if autograd should record operations on parameters in this module.

This method sets the parameters' *requires_grad* attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

requires_grad (bool): whether autograd should record operations on parameters in this module.
Default: True.

Returns: Module: self

share_memory_ () → T

state_dict (*destination*=None, *prefix*=",", *keep_vars*=False)
Returns a dictionary containing a whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

dict: a dictionary containing a whole state of the module

Example:

```
>>> module.state_dict().keys()  
['bias', 'weight']
```

to (**args*, ***kwargs*)
Moves and/or casts the parameters and buffers.

This can be called as

```
to(device=None, dtype=None, non_blocking=False)  
to(dtype, non_blocking=False)  
to(tensor, non_blocking=False)  
to(memory_format=torch.channels_last)
```

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with dtypes unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

device (torch.device): the desired device of the parameters and buffers in this module
dtype (torch.dtype): the desired floating point type of the floating point parameters and buffers in this module
tensor (torch.Tensor): Tensor whose `dtype` and `device` are the desired `dtype` and `device` for all parameters and buffers in this module
memory_format (torch.memory_format): the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]])
>>> linear.to(torch.double)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]], dtype=torch.float64)
>>> gpu1 = torch.device("cuda:1")
>>> linear.to(gpu1, dtype=torch.half, non_blocking=True)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
```

(continues on next page)

(continued from previous page)

```
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16)
```

train(*mode: bool = True*) → T

Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (True) or evaluation mode (False). Default: True.

Returns: Module: self**type**(*dst_type: Union[torch.dtype, str]*) → TCasts all parameters and buffers to *dst_type*.**Arguments:** *dst_type* (type or string): the desired type**Returns:** Module: self**zero_grad**() → None

Sets gradients of all model parameters to zero.

```
class flood_forecast.transformer_xl.transformer_basic.SimplePositionalEncoding(d_model,
                                                                           dropout=0.1,
                                                                           max_len=5000)
```

Bases: torch.nn.modules.module.Module

forward(*x: torch.Tensor*) → torch.Tensor

Creates a basic positional encoding

T_destination = ~T_destination**add_module**(*name: str, module: torch.nn.modules.module.Module*) → None

Adds a child module to the current module.

The module can be accessed as an attribute using the given name.

Args:

name (string): name of the child module. The child module can be accessed from this module using the given name

module (Module): child module to be added to the module.**apply**(*fn: Callable[[Module], None]*) → TApplies *fn* recursively to every submodule (as returned by *.children()*) as well as self. Typical use includes initializing the parameters of a model (see also nn-init-doc).**Args:** *fn (Module -> None): function to be applied to each submodule***Returns:** Module: self

Example:

```
>>> @torch.no_grad()
>>> def init_weights(m):
>>>     print(m)
>>>     if type(m) == nn.Linear:
```

(continues on next page)

(continued from previous page)

```
>>>     m.weight.fill_(1.0)
>>>     print(m.weight)
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))
>>> net.apply(init_weights)
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
```

bfloat16() → T

Casts all floating point parameters and buffers to bfloat16 datatype.

Returns: Module: self

buffers(reuse: bool = True) → Iterator[torch.Tensor]

Returns an iterator over module buffers.

Args:

reuse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer

Example:

```
>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

children() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module

cpu() → T

Moves all model parameters and buffers to the CPU.

Returns: Module: self

cuda(device: Union[int, torch.device, None] = None) → T

Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self

double () → T

Casts all floating point parameters and buffers to double datatype.

Returns: Module: self

dump_patches = False

eval () → T

Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with self.train(False).

Returns: Module: self

extra_repr () → str

Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float () → T

Casts all floating point parameters and buffers to float datatype.

Returns: Module: self

half () → T

Casts all floating point parameters and buffers to half datatype.

Returns: Module: self

load_state_dict (state_dict: Dict[str, torch.Tensor], strict: bool = True)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict` is True, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Arguments:

state_dict (dict): a dict containing parameters and persistent buffers.

strict (bool, optional): whether to strictly enforce that the keys in state_dict match the keys returned by this module's state_dict() function. Default: True

Returns:

NamedTuple with missing_keys and unexpected_keys fields:

- **missing_keys** is a list of str containing the missing keys
- **unexpected_keys** is a list of str containing the unexpected keys

modules () → Iterator[torch.nn.modules.module.Module]

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

`named_buffers` (*prefix: str = ”, recurse: bool = True*) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: *prefix* (str): prefix to prepend to all buffer names. *recurse* (bool): if True, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (string, torch.Tensor): Tuple containing the name and buffer

Example:

```
>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())
```

`named_children` () → Iterator[Tuple[str, torch.nn.modules.module.Module]]

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple containing a name and child module

Example:

```
>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)
```

`named_modules` (*memo: Optional[Set[Module]] = None, prefix: str = ”*)

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.named_modules()):
    print(idx, '->', m)

0 -> ('', Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
))
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))
```

named_parameters (*prefix: str = ”, recurse: bool = True*) → Iterator[Tuple[str, torch.Tensor]]
 Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: *prefix* (str): prefix to prepend to all parameter names. *recurse* (bool): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

parameters (*recurse: bool = True*) → Iterator[torch.nn.parameter.Parameter]
 Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

register_backward_hook (*hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]*) → torch.utils.hooks.RemovableHandle
 Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex Module that perform many operations. In some failure cases, `grad_input` and `grad_output` will only contain the gradients for a subset of the inputs and outputs. For such Module, you should use `torch.Tensor.register_hook()` directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook(module, grad_input, grad_output) -> Tensor or None
```

The `grad_input` and `grad_output` may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that will be used in place of `grad_input` in subsequent computations. `grad_input` will only correspond to the inputs given as positional arguments.

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_buffer (`name: str, tensor: torch.Tensor, persistent: bool = True`) → None

Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's `running_mean` is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting `persistent` to `False`. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

name (string): name of the buffer. The buffer can be accessed from this module using the given name

tensor (Tensor): buffer to be registered. persistent (bool): whether the buffer is part of this module's `state_dict`.

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

register_forward_hook (`hook: Callable[..., None]`) → `torch.utils.hooks.RemovableHandle`

Registers a forward hook on the module.

The hook will be called every time after `forward()` has computed an output. It should have the following signature:

```
hook(module, input, output) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the output. It can modify the input inplace but it will not have effect on `forward` since this is called after `forward()` is called.

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_forward_pre_hook (`hook: Callable[..., None]`) → `torch.utils.hooks.RemovableHandle`

Registers a forward pre-hook on the module.

The hook will be called every time before `forward()` is invoked. It should have the following signature:

```
hook(module, input) -> None or modified input
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_parameter (*name: str, param: torch.nn.parameter.Parameter*) → None

Adds a parameter to the module.

The parameter can be accessed as an attribute using given name.

Args:

name (string): name of the parameter. The parameter can be accessed from this module using the given name

param (Parameter): parameter to be added to the module.

requires_grad_ (*requires_grad: bool = True*) → T

Change if autograd should record operations on parameters in this module.

This method sets the parameters' `requires_grad` attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

requires_grad (bool): whether autograd should record operations on parameters in this module.

Default: True.

Returns: Module: self

share_memory_ () → T

state_dict (*destination=None, prefix=”, keep_vars=False*)

Returns a dictionary containing a whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

dict: a dictionary containing a whole state of the module

Example:

```
>>> module.state_dict().keys()
['bias', 'weight']
```

to (*args, **kwargs)

Moves and/or casts the parameters and buffers.

This can be called as

to (*device=None, dtype=None, non_blocking=False*)

to (*dtype, non_blocking=False*)

to (*tensor, non_blocking=False*)

to (*memory_format=torch.channels_last*)

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with `dtypes` unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

device (torch.device): the desired device of the parameters and buffers in this module
dtype (torch.dtype): the desired floating point type of the floating point parameters and buffers in this module
tensor (torch.Tensor): Tensor whose dtype and device are the desired dtype and device for all parameters and buffers in this module
memory_format (torch.memory_format): the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]])
>>> linear.to(torch.double)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]], dtype=torch.float64)
>>> gpu1 = torch.device("cuda:1")
>>> linear.to(gpu1, dtype=torch.half, non_blocking=True)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16)
```

train (mode: bool = True) → T

Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (True) or evaluation mode (False). Default: True.

Returns: Module: self

type (*dst_type*: *Union[torch.dtype, str]*) → T

Casts all parameters and buffers to *dst_type*.

Arguments: *dst_type* (type or string): the desired type

Returns: Module: self

zero_grad() → None

Sets gradients of all model parameters to zero.

```
flood_forecast.transformer_xl.transformer_basic.generate_square_subsequent_mask(sz: int) → torch.Tensor
```

Generate a square mask for the sequence. The masked positions are filled with float('inf'). Unmasked positions are filled with float(0.0).

```
flood_forecast.transformer_xl.transformer_basic.greedy_decode(model, src: torch.Tensor, max_len: int, real_target: torch.Tensor, unsqueeze_dim=1, device='cpu')
```

Mechanism to sequentially decode the model :src Historical time series values :real_target The real values (they should be masked), however if want can include known real values. :returns tensor

CHAPTER 23

Transformer XL

Model from Keita Kurita. Not useable https://github.com/keitakurita/Practical_NLP_in_PyTorch/blob/master/deep_dives/transformer_xl_from_scratch.ipynb

```
class flood_forecast.transformer_xl.transformer_xl.MultiHeadAttention(d_input:  
                           int,  
                           d_inner:  
                           int,  
                           n_heads:  
                           int  
                           = 4,  
                           dropout:  
                           float  
                           = 0.1,  
                           dropouts:  
                           float  
                           =  
                           0.0)  
Bases: torch.nn.modules.module.Module
```

```
forward(input_: torch.FloatTensor, pos_embs: torch.FloatTensor, memory: torch.FloatTensor, u:  
        torch.FloatTensor, v: torch.FloatTensor, mask: Optional[torch.FloatTensor] = None)
```

pos_embs: we pass the positional embeddings in separately because we need to handle relative positions

input shape: (seq, bs, self.d_input) pos_embs shape: (seq + prev_seq, bs, self.d_input) output shape: (seq, bs, self.d_input)

```
T_destination = ~T_destination
```

```
add_module(name: str, module: torch.nn.modules.module.Module) → None  
    Adds a child module to the current module.
```

The module can be accessed as an attribute using the given name.

Args:

name (string): name of the child module. The child module can be accessed from this module using the given name

module (Module): child module to be added to the module.

apply (fn: Callable[[Module], None]) → T

Applies fn recursively to every submodule (as returned by .children()) as well as self. Typical use includes initializing the parameters of a model (see also nn-init-doc).

Args: fn (Module -> None): function to be applied to each submodule

Returns: Module: self

Example:

```
>>> @torch.no_grad()
>>> def init_weights(m):
>>>     print(m)
>>>     if type(m) == nn.Linear:
>>>         m.weight.fill_(1.0)
>>>         print(m.weight)
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))
>>> net.apply(init_weights)
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
```

bfloat16 () → T

Casts all floating point parameters and buffers to bfloat16 datatype.

Returns: Module: self

buffers (reuse: bool = True) → Iterator[torch.Tensor]

Returns an iterator over module buffers.

Args:

reuse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer

Example:

```
>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

children() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module

cpu() → T

Moves all model parameters and buffers to the CPU.

Returns: Module: self

cuda(device: Union[int, torch.device, None] = None) → T

Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self

double() → T

Casts all floating point parameters and buffers to double datatype.

Returns: Module: self

dump_patches = False

eval() → T

Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with `self.train(False)`.

Returns: Module: self

extra_repr() → str

Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float() → T

Casts all floating point parameters and buffers to float datatype.

Returns: Module: self

half() → T

Casts all floating point parameters and buffers to half datatype.

Returns: Module: self

load_state_dict(state_dict: Dict[str, torch.Tensor], strict: bool = True)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict` is True, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Arguments:

state_dict (dict): a dict containing parameters and persistent buffers.

strict (bool, optional): whether to strictly enforce that the keys in `state_dict` match the keys returned by this module's `state_dict()` function. Default: True

Returns:

NamedTuple with **missing_keys** and **unexpected_keys** fields:

- **missing_keys** is a list of str containing the missing keys
- **unexpected_keys** is a list of str containing the unexpected keys

modules() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

named_buffers (*prefix: str = ”, recurse: bool = True*) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: *prefix* (str): prefix to prepend to all buffer names. *recurse* (bool): if True, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (string, torch.Tensor): Tuple containing the name and buffer

Example:

```
>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())
```

named_children() → Iterator[Tuple[str, torch.nn.modules.module.Module]]

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple containing a name and child module

Example:

```
>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)
```

named_modules (*memo: Optional[Set[Module]] = None, prefix: str = ”*)

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.named_modules()):
    print(idx, '->', m)

0 -> ('', Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
))
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))
```

named_parameters (*prefix: str = ”, recurse: bool = True*) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: *prefix* (str): prefix to prepend to all parameter names. *recurse* (bool): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

parameters (*recurse: bool = True*) → Iterator[torch.nn.parameter.Parameter]

Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

register_backward_hook (*hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]*) → torch.utils.hooks.RemovableHandle

Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex Module that perform many operations. In some failure cases, `grad_input` and `grad_output` will only contain the gradients for a subset of the inputs and outputs. For such Module, you should use `torch.Tensor.register_hook()` directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook(module, grad_input, grad_output) -> Tensor or None
```

The `grad_input` and `grad_output` may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that will be used in place of `grad_input` in subsequent computations. `grad_input` will only correspond to the inputs given as positional arguments.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

`register_buffer(name: str, tensor: torch.Tensor, persistent: bool = True) → None`

Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's `running_mean` is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting `persistent` to `False`. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

`name (string)`: name of the buffer. The buffer can be accessed from this module using the given name

`tensor (Tensor)`: buffer to be registered. `persistent (bool)`: whether the buffer is part of this module's `state_dict`.

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

`register_forward_hook(hook: Callable[..., None]) → torch.utils.hooks.RemovableHandle`

Registers a forward hook on the module.

The hook will be called every time after `forward()` has computed an output. It should have the following signature:

```
hook(module, input, output) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the output. It can modify the input inplace but it will not have effect on `forward` since this is called after `forward()` is called.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

`register_forward_pre_hook(hook: Callable[..., None]) → torch.utils.hooks.RemovableHandle`

Registers a forward pre-hook on the module.

The hook will be called every time before `forward()` is invoked. It should have the following signature:

```
hook(module, input) -> None or modified input
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

`register_parameter(name: str, param: torch.nn.parameter.Parameter)` → None

Adds a parameter to the module.

The parameter can be accessed as an attribute using given name.

Args:

`name (string): name of the parameter. The parameter can be accessed` from this module using the given name

`param (Parameter): parameter to be added to the module.`

`requires_grad_(requires_grad: bool = True)` → T

Change if autograd should record operations on parameters in this module.

This method sets the parameters' `requires_grad` attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

`requires_grad (bool): whether autograd should record operations on` parameters in this module.

Default: True.

Returns: Module: self

`share_memory()` → T

`state_dict(destination=None, prefix= "", keep_vars=False)`

Returns a dictionary containing a whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

`dict: a dictionary containing a whole state of the module`

Example:

```
>>> module.state_dict().keys()
['bias', 'weight']
```

`to(*args, **kwargs)`

Moves and/or casts the parameters and buffers.

This can be called as

`to(device=None, dtype=None, non_blocking=False)`

`to(dtype, non_blocking=False)`

```
to (tensor, non_blocking=False)
to (memory_format=torch.channels_last)
```

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with `dtypes` unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

device (torch.device): the desired device of the parameters and buffers in this module
dtype (torch.dtype): the desired floating point type of the floating point parameters and buffers in this module
tensor (torch.Tensor): Tensor whose `dtype` and `device` are the desired `dtype` and `device` for all parameters and buffers in this module
memory_format (torch.memory_format): the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]])
>>> linear.to(torch.double)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]], dtype=torch.float64)
>>> gpu1 = torch.device("cuda:1")
>>> linear.to(gpu1, dtype=torch.half, non_blocking=True)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16)
```

train (mode: bool = True) → T

Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (True) or evaluation mode (False). Default: True.

Returns: Module: self

type (dst_type: Union[torch.dtype, str]) → T

Casts all parameters and buffers to dst_type.

Arguments: dst_type (type or string): the desired type

Returns: Module: self

zero_grad() → None

Sets gradients of all model parameters to zero.

```
class flood_forecast.transformer_xl.transformer_xl.PositionwiseFF(d_input,
                                                               d_inner,
                                                               dropout)
```

Bases: torch.nn.modules.module.Module

forward (input_: torch.FloatTensor) → torch.FloatTensor

T_destination = ~T_destination

add_module (name: str, module: torch.nn.modules.module.Module) → None

Adds a child module to the current module.

The module can be accessed as an attribute using the given name.

Args:

name (string): name of the child module. The child module can be accessed from this module using the given name

module (Module): child module to be added to the module.

apply (fn: Callable[[Module], None]) → T

Applies fn recursively to every submodule (as returned by .children()) as well as self. Typical use includes initializing the parameters of a model (see also nn-init-doc).

Args: fn (Module -> None): function to be applied to each submodule

Returns: Module: self

Example:

```
>>> @torch.no_grad()
>>> def init_weights(m):
>>>     print(m)
>>>     if type(m) == nn.Linear:
>>>         m.weight.fill_(1.0)
>>>         print(m.weight)
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))
>>> net.apply(init_weights)
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
        [ 1.,  1.]])
Linear(in_features=2, out_features=2, bias=True)
```

(continues on next page)

(continued from previous page)

```
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
```

`bfloat16()` → T

Casts all floating point parameters and buffers to `bfloat16` datatype.

Returns: Module: self

`buffers(reuse: bool = True)` → Iterator[torch.Tensor]

Returns an iterator over module buffers.

Args:

reuse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer

Example:

```
>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

`children()` → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module

`cpu()` → T

Moves all model parameters and buffers to the CPU.

Returns: Module: self

`cuda(device: Union[int, torch.device, None] = None)` → T

Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self

`double()` → T

Casts all floating point parameters and buffers to `double` datatype.

Returns: Module: self

`dump_patches = False`

eval () → T

Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with `self.train(False)`.

Returns: Module: self

extra_repr () → str

Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float () → T

Casts all floating point parameters and buffers to float datatype.

Returns: Module: self

half () → T

Casts all floating point parameters and buffers to half datatype.

Returns: Module: self

load_state_dict (state_dict: Dict[str, torch.Tensor], strict: bool = True)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict` is True, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Arguments:

state_dict (dict): a dict containing parameters and persistent buffers.

strict (bool, optional): whether to strictly enforce that the keys in `state_dict` match the keys returned by this module's `state_dict()` function. Default: True

Returns:

NamedTuple with missing_keys and unexpected_keys fields:

- **missing_keys** is a list of str containing the missing keys
- **unexpected_keys** is a list of str containing the unexpected keys

modules () → Iterator[torch.nn.modules.module.Module]

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, 1 will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

named_buffers (*prefix: str = ”, recurse: bool = True*) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: *prefix* (str): prefix to prepend to all buffer names. *recurse* (bool): if True, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (string, torch.Tensor): Tuple containing the name and buffer

Example:

```
>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())
```

named_children () → Iterator[Tuple[str, torch.nn.modules.module.Module]]

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple containing a name and child module

Example:

```
>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)
```

named_modules (*memo: Optional[Set[Module]] = None, prefix: str = ”*)

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.named_modules()):
>>>     print(idx, '->', m)

0 -> ('', Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
))
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))
```

named_parameters (*prefix: str = ”, recurse: bool = True*) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: *prefix* (str): prefix to prepend to all parameter names. *recurse* (bool): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

parameters (*recurse: bool = True*) → Iterator[torch.nn.parameter.Parameter]

Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

register_backward_hook (*hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]*) → torch.utils.hooks.RemovableHandle

Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex Module that perform many operations. In some failure cases, `grad_input` and `grad_output` will only contain the gradients for a subset of the inputs and outputs. For such Module, you should use `torch.Tensor.register_hook()` directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook(module, grad_input, grad_output) -> Tensor or None
```

The `grad_input` and `grad_output` may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that will be used in place of `grad_input` in subsequent computations. `grad_input` will only correspond to the inputs given as positional arguments.

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_buffer (*name: str, tensor: torch.Tensor, persistent: bool = True*) → None

Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's `running_mean` is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting `persistent` to `False`. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

name (string): name of the buffer. The buffer can be accessed from this module using the given name
tensor (Tensor): buffer to be registered. persistent (bool): whether the buffer is part of this module's *state_dict*.

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

register_forward_hook (hook: Callable[[], None]) → torch.utils.hooks.RemovableHandle
Registers a forward hook on the module.

The hook will be called every time after *forward()* has computed an output. It should have the following signature:

```
hook(module, input, output) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the *forward*. The hook can modify the output. It can modify the input inplace but it will not have effect on forward since this is called after *forward()* is called.

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling *handle.remove()*

register_forward_pre_hook (hook: Callable[[], None]) → torch.utils.hooks.RemovableHandle
Registers a forward pre-hook on the module.

The hook will be called every time before *forward()* is invoked. It should have the following signature:

```
hook(module, input) -> None or modified input
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the *forward*. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling *handle.remove()*

register_parameter (name: str, param: torch.nn.parameter.Parameter) → None
Adds a parameter to the module.

The parameter can be accessed as an attribute using given name.

Args:

name (string): name of the parameter. The parameter can be accessed from this module using the given name

param (Parameter): parameter to be added to the module.

requires_grad_ (requires_grad: bool = True) → T
Change if autograd should record operations on parameters in this module.

This method sets the parameters' *requires_grad* attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

requires_grad (bool): whether autograd should record operations on parameters in this module.
Default: True.

Returns: Module: self

share_memory () → T

state_dict (destination=None, prefix='', keep_vars=False)
Returns a dictionary containing a whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

dict: a dictionary containing a whole state of the module

Example:

```
>>> module.state_dict().keys()
['bias', 'weight']
```

to (*args, **kwargs)

Moves and/or casts the parameters and buffers.

This can be called as

to (device=None, dtype=None, non_blocking=False)

to (dtype, non_blocking=False)

to (tensor, non_blocking=False)

to (memory_format=torch.channels_last)

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with dtypes unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

device (torch.device): the desired device of the parameters and buffers in this module

dtype (torch.dtype): the desired floating point type of the floating point parameters and buffers in this module

tensor (torch.Tensor): Tensor whose dtype and device are the desired `dtype` and `device` for all parameters and buffers in this module

memory_format (torch.memory_format): the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]])
>>> linear.to(torch.double)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]], dtype=torch.float64)
>>> gp1 = torch.device("cuda:1")
>>> linear.to(gp1, dtype=torch.half, non_blocking=True)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16)
```

train(mode: bool = True) → T

Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (True) or evaluation mode (False). Default: True.

Returns: Module: self

type(dst_type: Union[torch.dtype, str]) → T

Casts all parameters and buffers to dst_type.

Arguments: dst_type (type or string): the desired type

Returns: Module: self

zero_grad() → None

Sets gradients of all model parameters to zero.

```
class flood_forecast.transformer_xl.transformer_xl.DecoderBlock(n_heads,
                                                               d_input,
                                                               d_head_inner,
                                                               d_ff_inner,
                                                               dropout,
                                                               dropouts=0.0)
```

Bases: torch.nn.modules.module.Module

```
forward(input_: torch.FloatTensor, pos_embs: torch.FloatTensor, u: torch.FloatTensor, v: torch.FloatTensor, mask=None, mems=None)
```

T_destination = ~T_destination

add_module(name: str, module: torch.nn.modules.module.Module) → None

Adds a child module to the current module.

The module can be accessed as an attribute using the given name.

Args:

name (string): name of the child module. The child module can be accessed from this module using the given name

module (Module): child module to be added to the module.

apply(fn: Callable[[Module], None]) → T

Applies fn recursively to every submodule (as returned by .children()) as well as self. Typical use includes initializing the parameters of a model (see also nn-init-doc).

Args: fn (Module -> None): function to be applied to each submodule

Returns: Module: self

Example:

```
>>> @torch.no_grad()
>>> def init_weights(m):
>>>     print(m)
>>>     if type(m) == nn.Linear:
>>>         m.weight.fill_(1.0)
>>>         print(m.weight)
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))
>>> net.apply(init_weights)
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
```

bfloat16() → T

Casts all floating point parameters and buffers to bfloat16 datatype.

Returns: Module: self

buffers(reuse: bool = True) → Iterator[torch.Tensor]

Returns an iterator over module buffers.

Args:

recurse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer

Example:

```
>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

children () → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module

cpu () → T

Moves all model parameters and buffers to the CPU.

Returns: Module: self

cuda (device: Union[int, torch.device, None] = None) → T

Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self

double () → T

Casts all floating point parameters and buffers to double datatype.

Returns: Module: self

dump_patches = False

eval () → T

Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with `self.train(False)`.

Returns: Module: self

extra_repr () → str

Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float () → T

Casts all floating point parameters and buffers to float datatype.

Returns: Module: self

half () → T

Casts all floating point parameters and buffers to half datatype.

Returns: Module: self

load_state_dict (*state_dict*: *Dict[str; torch.Tensor]*, *strict*: *bool* = *True*)

Copies parameters and buffers from *state_dict* into this module and its descendants. If *strict* is *True*, then the keys of *state_dict* must exactly match the keys returned by this module's *state_dict()* function.

Arguments:

state_dict (*dict*): a dict containing parameters and persistent buffers.

strict (*bool*, optional): whether to strictly enforce that the keys in *state_dict* match the keys returned by this module's *state_dict()* function. Default: *True*

Returns:

NamedTuple with missing_keys and unexpected_keys fields:

- **missing_keys** is a list of str containing the missing keys
- **unexpected_keys** is a list of str containing the unexpected keys

modules () → *Iterator[torch.nn.modules.module.Module]*

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, 1 will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

named_buffers (*prefix*: *str* = "", *recurse*: *bool* = *True*) → *Iterator[Tuple[str, torch.Tensor]]*

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: *prefix* (*str*): prefix to prepend to all buffer names. *recurse* (*bool*): if *True*, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (*string, torch.Tensor*): Tuple containing the name and buffer

Example:

```
>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())
```

named_children () → *Iterator[Tuple[str, torch.nn.modules.module.Module]]*

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (*string, Module*): Tuple containing a name and child module

Example:

```
>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)
```

named_modules (*memo: Optional[Set[Module]]* = *None*, *prefix: str* = '')

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.named_modules()):
>>>     print(idx, '->', m)

0 -> ('', Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
))
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))
```

named_parameters (*prefix: str* = '', *recurse: bool* = *True*) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: *prefix* (str): prefix to prepend to all parameter names. *recurse* (bool): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

parameters (*recurse: bool* = *True*) → Iterator[torch.nn.parameter.Parameter]

Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

register_backward_hook (*hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]*) → `torch.utils.hooks.RemovableHandle`
Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex `Module` that perform many operations. In some failure cases, `grad_input` and `grad_output` will only contain the gradients for a subset of the inputs and outputs. For such `Module`, you should use `torch.Tensor.register_hook()` directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook(module, grad_input, grad_output) -> Tensor or None
```

The `grad_input` and `grad_output` may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that will be used in place of `grad_input` in subsequent computations. `grad_input` will only correspond to the inputs given as positional arguments.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

register_buffer (*name: str, tensor: torch.Tensor, persistent: bool = True*) → `None`
Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's `running_mean` is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting `persistent` to `False`. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

name (string): name of the buffer. The buffer can be accessed from this module using the given name

tensor (Tensor): buffer to be registered. **persistent (bool):** whether the buffer is part of this module's `state_dict`.

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

register_forward_hook (*hook: Callable[[...], None]*) → `torch.utils.hooks.RemovableHandle`
Registers a forward hook on the module.

The hook will be called every time after `forward()` has computed an output. It should have the following signature:

```
hook(module, input, output) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the output. It can modify the input inplace but it will not have effect on forward since this is called after `forward()` is called.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

<code>register_forward_pre_hook(hook: Callable[..., torch.utils.hooks.RemovableHandle]</code>	<code>None])</code>	→
Registers a forward pre-hook on the module.		

The hook will be called every time before `forward()` is invoked. It should have the following signature:

<code>hook(module, input) -> None or modified input</code>

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

<code>register_parameter(name: str, param: torch.nn.parameter.Parameter)</code>	→ None
Adds a parameter to the module.	

The parameter can be accessed as an attribute using given name.

Args:

name (string): name of the parameter. The parameter can be accessed from this module using the given name

param (Parameter): parameter to be added to the module.

<code>requires_grad_(requires_grad: bool = True)</code>	→ T
Change if autograd should record operations on parameters in this module.	

This method sets the parameters' `requires_grad` attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

requires_grad (bool): whether autograd should record operations on parameters in this module.
Default: True.

Returns: Module: self

`share_memory()` → T

<code>state_dict(destination=None, prefix= "", keep_vars=False)</code>	→
Returns a dictionary containing a whole state of the module.	

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

dict: a dictionary containing a whole state of the module

Example:

```
>>> module.state_dict().keys()
['bias', 'weight']
```

to(*args, **kwargs)

Moves and/or casts the parameters and buffers.

This can be called as

```
to(device=None, dtype=None, non_blocking=False)
to(dtype, non_blocking=False)
to(tensor, non_blocking=False)
to(memory_format=torch.channels_last)
```

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with dtypes unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

device (torch.device): the desired device of the parameters and buffers in this module

dtype (torch.dtype): the desired floating point type of the floating point parameters and buffers in this module

tensor (torch.Tensor): Tensor whose `dtype` and `device` are the desired `dtype` and `device` for all parameters and buffers in this module

memory_format (torch.memory_format): the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]])
>>> linear.to(torch.double)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]], dtype=torch.float64)
>>> gpu1 = torch.device("cuda:1")
>>> linear.to(gpu1, dtype=torch.half, non_blocking=True)
Linear(in_features=2, out_features=2, bias=True)
```

(continues on next page)

(continued from previous page)

```
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16)
```

train(*mode: bool* = *True*) → *T*

Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (*True*) or evaluation mode (*False*). Default: *True*.

Returns: Module: self**type**(*dst_type: Union[torch.dtype, str]*) → *T*

Casts all parameters and buffers to *dst_type*.

Arguments: *dst_type* (type or string): the desired type

Returns: Module: self**zero_grad()** → None

Sets gradients of all model parameters to zero.

class flood_forecast.transformer_xl.transformer_xl.**PositionalEmbedding**(*d*)

Bases: torch.nn.modules.module.Module

forward(*positions: torch.LongTensor*)**T_destination** = ~**T_destination****add_module**(*name: str, module: torch.nn.modules.module.Module*) → None

Adds a child module to the current module.

The module can be accessed as an attribute using the given name.

Args:

name (string): name of the child module. The child module can be accessed from this module using the given name

module (Module): child module to be added to the module.

apply(*fn: Callable[[Module], None]*) → *T*

Applies *fn* recursively to every submodule (as returned by *.children()*) as well as self. Typical use includes initializing the parameters of a model (see also nn-init-doc).

Args: *fn (Module -> None): function to be applied to each submodule***Returns:** Module: self

Example:

```
>>> @torch.no_grad()
>>> def init_weights(m):
>>>     print(m)
>>>     if type(m) == nn.Linear:
>>>         m.weight.fill_(1.0)
>>>         print(m.weight)
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))
>>> net.apply(init_weights)
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
```

bfloat16() → T

Casts all floating point parameters and buffers to bfloat16 datatype.

Returns: Module: self

buffers(reuse: bool = True) → Iterator[torch.Tensor]

Returns an iterator over module buffers.

Args:

reuse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer

Example:

```
>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

children() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module

cpu() → T

Moves all model parameters and buffers to the CPU.

Returns: Module: self

cuda(device: Union[int, torch.device, None] = None) → T

Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self

double () → T

Casts all floating point parameters and buffers to double datatype.

Returns: Module: self

dump_patches = False

eval () → T

Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with `self.train(False)`.

Returns: Module: self

extra_repr () → str

Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float () → T

Casts all floating point parameters and buffers to float datatype.

Returns: Module: self

half () → T

Casts all floating point parameters and buffers to half datatype.

Returns: Module: self

load_state_dict (state_dict: Dict[str, torch.Tensor], strict: bool = True)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict` is True, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Arguments:

state_dict (dict): a dict containing parameters and persistent buffers.

strict (bool, optional): whether to strictly enforce that the keys in `state_dict` match the keys returned by this module's `state_dict()` function. Default: True

Returns:

NamedTuple with missing_keys and unexpected_keys fields:

- **missing_keys** is a list of str containing the missing keys

- **unexpected_keys** is a list of str containing the unexpected keys

modules () → Iterator[torch.nn.modules.module.Module]

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, 1 will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

named_buffers (*prefix: str = ”, recurse: bool = True*) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: *prefix* (str): prefix to prepend to all buffer names. *recurse* (bool): if True, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (string, torch.Tensor): Tuple containing the name and buffer

Example:

```
>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())
```

named_children () → Iterator[Tuple[str, torch.nn.modules.module.Module]]

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple containing a name and child module

Example:

```
>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)
```

named_modules (*memo: Optional[Set[Module]] = None, prefix: str = ”*)

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, 1 will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.named_modules()):
    print(idx, '->', m)

0 -> ('', Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

(continues on next page)

(continued from previous page)

```
) )
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))
```

named_parameters (*prefix: str = ”, recurse: bool = True*) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: *prefix* (str): prefix to prepend to all parameter names. *recurse* (bool): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

parameters (*recurse: bool = True*) → Iterator[torch.nn.parameter.Parameter]

Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

register_backward_hook (*hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]*) → torch.utils.hooks.RemovableHandle

Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex Module that perform many operations. In some failure cases, `grad_input` and `grad_output` will only contain the gradients for a subset of the inputs and outputs. For such Module, you should use `torch.Tensor.register_hook()` directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook(module, grad_input, grad_output) -> Tensor or None
```

The `grad_input` and `grad_output` may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that

will be used in place of `grad_input` in subsequent computations. `grad_input` will only correspond to the inputs given as positional arguments.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

`register_buffer(name: str, tensor: torch.Tensor, persistent: bool = True) → None`

Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's `running_mean` is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting `persistent` to `False`. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

`name (string): name of the buffer. The buffer can be accessed` from this module using the given name

`tensor (Tensor): buffer to be registered. persistent (bool): whether the buffer is part of this module's` `state_dict`.

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

`register_forward_hook(hook: Callable[[], None]) → torch.utils.hooks.RemovableHandle`

Registers a forward hook on the module.

The hook will be called every time after `forward()` has computed an output. It should have the following signature:

```
hook(module, input, output) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the output. It can modify the input inplace but it will not have effect on forward since this is called after `forward()` is called.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

`register_forward_pre_hook(hook: Callable[[], None]) → torch.utils.hooks.RemovableHandle`

Registers a forward pre-hook on the module.

The hook will be called every time before `forward()` is invoked. It should have the following signature:

```
hook(module, input) -> None or modified input
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_parameter (`name: str, param: torch.nn.parameter.Parameter`) → None

Adds a parameter to the module.

The parameter can be accessed as an attribute using given name.

Args:

name (string): name of the parameter. The parameter can be accessed from this module using the given name

param (Parameter): parameter to be added to the module.

requires_grad_ (`requires_grad: bool = True`) → T

Change if autograd should record operations on parameters in this module.

This method sets the parameters' `requires_grad` attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

requires_grad (bool): whether autograd should record operations on parameters in this module.

Default: True.

Returns: Module: self

share_memory () → T

state_dict (`destination=None, prefix=”, keep_vars=False`)

Returns a dictionary containing a whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

dict: a dictionary containing a whole state of the module

Example:

```
>>> module.state_dict().keys()  
['bias', 'weight']
```

to (*args, **kwargs)

Moves and/or casts the parameters and buffers.

This can be called as

to (`device=None, dtype=None, non_blocking=False`)

to (`dtype, non_blocking=False`)

to (`tensor, non_blocking=False`)

to (`memory_format=torch.channels_last`)

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with dtypes unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

device (torch.device): the desired device of the parameters and buffers in this module
dtype (torch.dtype): the desired floating point type of the floating point parameters and buffers in this module
tensor (torch.Tensor): Tensor whose dtype and device are the desired dtype and device for all parameters and buffers in this module
memory_format (torch.memory_format): the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]])
>>> linear.to(torch.double)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]], dtype=torch.float64)
>>> gpu1 = torch.device("cuda:1")
>>> linear.to(gpu1, dtype=torch.half, non_blocking=True)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16)
```

train (mode: bool = True) → T

Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (**True**) or evaluation mode (**False**). Default: **True**.

Returns: Module: self

type (*dst_type*: *Union[torch.dtype, str]*) → T
Casts all parameters and buffers to *dst_type*.

Arguments: *dst_type* (type or string): the desired type

Returns: Module: self

zero_grad() → None
Sets gradients of all model parameters to zero.

class *flood_forecast.transformer_xl.transformer_xl.StandardWordEmbedding* (*num_embeddings*,
em-
bed-
ding_dim,
div_val=1,
sam-
ple_softmax=False)

Bases: *torch.nn.modules.module.Module*

forward (*input_*: *torch.LongTensor*)

T_destination = ~**T_destination**

add_module (*name*: str, *module*: *torch.nn.modules.module.Module*) → None
Adds a child module to the current module.

The module can be accessed as an attribute using the given name.

Args:

name (string): name of the child module. The child module can be accessed from this module
using the given name

module (Module): child module to be added to the module.

apply (*fn*: *Callable[[Module], None]*) → T

Applies *fn* recursively to every submodule (as returned by *.children()*) as well as self. Typical use includes initializing the parameters of a model (see also nn-init-doc).

Args: *fn* (*Module* → None): function to be applied to each submodule

Returns: Module: self

Example:

```
>>> @torch.no_grad()
>>> def init_weights(m):
>>>     print(m)
>>>     if type(m) == nn.Linear:
>>>         m.weight.fill_(1.0)
>>>         print(m.weight)
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))
>>> net.apply(init_weights)
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Sequential()
```

(continues on next page)

(continued from previous page)

```
(0): Linear(in_features=2, out_features=2, bias=True)
(1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
(0): Linear(in_features=2, out_features=2, bias=True)
(1): Linear(in_features=2, out_features=2, bias=True)
)
```

bfloat16() → T

Casts all floating point parameters and buffers to bfloat16 datatype.

Returns: Module: self**buffers(reuse: bool = True)** → Iterator[torch.Tensor]

Returns an iterator over module buffers.

Args:

reuse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer

Example:

```
>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

children() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module**cpu()** → T

Moves all model parameters and buffers to the CPU.

Returns: Module: self**cuda(device: Union[int, torch.device, None] = None)** → T

Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self**double()** → T

Casts all floating point parameters and buffers to double datatype.

Returns: Module: self**dump_patches = False****eval()** → T

Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with `self.train(False)`.

Returns: Module: self

extra_repr() → str

Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float() → T

Casts all floating point parameters and buffers to float datatype.

Returns: Module: self

half() → T

Casts all floating point parameters and buffers to half datatype.

Returns: Module: self

load_state_dict(state_dict: Dict[str, torch.Tensor], strict: bool = True)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict` is True, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Arguments:

state_dict (dict): a dict containing parameters and persistent buffers.

strict (bool, optional): whether to strictly enforce that the keys in state_dict match the keys returned by this module's `state_dict()` function. Default: True

Returns:

NamedTuple with missing_keys and unexpected_keys fields:

- **missing_keys** is a list of str containing the missing keys
- **unexpected_keys** is a list of str containing the unexpected keys

modules() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

named_buffers(prefix: str = "", recurse: bool = True) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: prefix (str): prefix to prepend to all buffer names. recurse (bool): if True, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (string, torch.Tensor): Tuple containing the name and buffer

Example:

```
>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())
```

named_children() → Iterator[Tuple[str, torch.nn.modules.module.Module]]

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple containing a name and child module

Example:

```
>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)
```

named_modules(memo: Optional[Set[Module]] = None, prefix: str = '')

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, l will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.named_modules()):
>>>     print(idx, '->', m)

0 -> ('', Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
))
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))
```

named_parameters(prefix: str = '', recurse: bool = True) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: prefix (str): prefix to prepend to all parameter names. recurse (bool): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

parameters (*recurse: bool = True*) → Iterator[torch.nn.parameter.Parameter]

Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

register_backward_hook (*hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]*) → torch.utils.hooks.RemovableHandle

Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex Module that perform many operations. In some failure cases, `grad_input` and `grad_output` will only contain the gradients for a subset of the inputs and outputs. For such Module, you should use `torch.Tensor.register_hook()` directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook(module, grad_input, grad_output) -> Tensor or None
```

The `grad_input` and `grad_output` may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that will be used in place of `grad_input` in subsequent computations. `grad_input` will only correspond to the inputs given as positional arguments.

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_buffer (*name: str, tensor: torch.Tensor, persistent: bool = True*) → None

Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's `running_mean` is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting `persistent` to `False`. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

name (string): name of the buffer. The buffer can be accessed from this module using the given name
tensor (Tensor): buffer to be registered. **persistent (bool):** whether the buffer is part of this module's *state_dict*.

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

register_forward_hook (hook: Callable[[], None]) → torch.utils.hooks.RemovableHandle
Registers a forward hook on the module.

The hook will be called every time after *forward()* has computed an output. It should have the following signature:

```
hook(module, input, output) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the *forward*. The hook can modify the output. It can modify the input inplace but it will not have effect on forward since this is called after *forward()* is called.

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling *handle.remove()*

register_forward_pre_hook (hook: Callable[[], None]) → torch.utils.hooks.RemovableHandle
Registers a forward pre-hook on the module.

The hook will be called every time before *forward()* is invoked. It should have the following signature:

```
hook(module, input) -> None or modified input
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the *forward*. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling *handle.remove()*

register_parameter (name: str, param: torch.nn.parameter.Parameter) → None
Adds a parameter to the module.

The parameter can be accessed as an attribute using given name.

Args:

name (string): name of the parameter. The parameter can be accessed from this module using the given name

param (Parameter): parameter to be added to the module.

requires_grad_ (requires_grad: bool = True) → T
Change if autograd should record operations on parameters in this module.

This method sets the parameters' *requires_grad* attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

requires_grad (bool): whether autograd should record operations on parameters in this module.
Default: True.

Returns: Module: self

share_memory () → T

state_dict (destination=None, prefix="", keep_vars=False)
Returns a dictionary containing a whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

dict: a dictionary containing a whole state of the module

Example:

```
>>> module.state_dict().keys()  
['bias', 'weight']
```

to (*args, **kwargs)

Moves and/or casts the parameters and buffers.

This can be called as

to (device=None, dtype=None, non_blocking=False)

to (dtype, non_blocking=False)

to (tensor, non_blocking=False)

to (memory_format=torch.channels_last)

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with dtypes unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

device (torch.device): the desired device of the parameters and buffers in this module

dtype (torch.dtype): the desired floating point type of the floating point parameters and buffers in this module

tensor (torch.Tensor): Tensor whose dtype and device are the desired `dtype` and `device` for all parameters and buffers in this module

memory_format (torch.memory_format): the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]])
>>> linear.to(torch.double)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]], dtype=torch.float64)
>>> gpu1 = torch.device("cuda:1")
>>> linear.to(gpu1, dtype=torch.half, non_blocking=True)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16)
```

train(mode: bool = True) → T

Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (True) or evaluation mode (False). Default: True.

Returns: Module: self

type(dst_type: Union[torch.dtype, str]) → T

Casts all parameters and buffers to dst_type.

Arguments: dst_type (type or string): the desired type

Returns: Module: self

zero_grad() → None

Sets gradients of all model parameters to zero.

```
class flood_forecast.transformer_xl.transformer_xl.TransformerXL(num_embeddings,  

    n_layers,  

    n_heads,  

    d_model,  

    d_head_inner,  

    d_ff_inner,  

    dropout=0.1,  

    dropouth=0.0,  

    seq_len:  

        int = 0,  

    mem_len: int  

        = 0)
```

Bases: torch.nn.modules.module.Module

init_memory(*device*=*device(type='cpu')*) → torch.FloatTensor

update_memory(*previous_memory*: List[torch.FloatTensor], *hidden_states*: List[torch.FloatTensor])

reset_length(*seq_len*, *ext_len*, *mem_len*)

forward(*idxs*: torch.LongTensor, *target*: torch.LongTensor, *memory*: Optional[List[torch.FloatTensor]] = None) → Dict[str, torch.Tensor]

T_destination = ~**T_destination**

add_module(*name*: str, *module*: torch.nn.modules.module.Module) → None

Adds a child module to the current module.

The module can be accessed as an attribute using the given name.

Args:

name (string): **name of the child module. The child module can be** accessed from this module using the given name

module (Module): child module to be added to the module.

apply(*fn*: Callable[[Module], None]) → T

Applies *fn* recursively to every submodule (as returned by *.children()*) as well as self. Typical use includes initializing the parameters of a model (see also nn-init-doc).

Args: *fn* (Module -> None): function to be applied to each submodule

Returns: Module: self

Example:

```
>>> @torch.no_grad()  
>>> def init_weights(m):  
>>>     print(m)  
>>>     if type(m) == nn.Linear:  
>>>         m.weight.fill_(1.0)  
>>>         print(m.weight)  
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))  
>>> net.apply(init_weights)  
Linear(in_features=2, out_features=2, bias=True)  
Parameter containing:  
tensor([[ 1.,  1.],  
       [ 1.,  1.]])  
Linear(in_features=2, out_features=2, bias=True)  
Parameter containing:  
tensor([[ 1.,  1.],
```

(continues on next page)

(continued from previous page)

```
[ 1.,  1.])
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
```

bfloating16() → T

Casts all floating point parameters and buffers to bfloat16 datatype.

Returns: Module: self

buffers(reuse: bool = True) → Iterator[torch.Tensor]

Returns an iterator over module buffers.

Args:

reuse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer

Example:

```
>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

children() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module

cpu() → T

Moves all model parameters and buffers to the CPU.

Returns: Module: self

cuda(device: Union[int, torch.device, None] = None) → T

Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self

double() → T

Casts all floating point parameters and buffers to double datatype.

Returns: Module: self

dump_patches = False**eval()** → T

Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with `self.train(False)`.

Returns: Module: self

extra_repr() → str

Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float() → T

Casts all floating point parameters and buffers to float datatype.

Returns: Module: self

half() → T

Casts all floating point parameters and buffers to half datatype.

Returns: Module: self

load_state_dict(state_dict: Dict[str, torch.Tensor], strict: bool = True)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict` is True, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Arguments:

state_dict (dict): a dict containing parameters and persistent buffers.

strict (bool, optional): whether to strictly enforce that the keys in state_dict match the keys returned by this module's `state_dict()` function. Default: True

Returns:

NamedTuple with missing_keys and unexpected_keys fields:

- **missing_keys** is a list of str containing the missing keys
- **unexpected_keys** is a list of str containing the unexpected keys

modules() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

named_buffers(prefix: str = "", recurse: bool = True) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: prefix (str): prefix to prepend to all buffer names. recurse (bool): if True, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (string, torch.Tensor): Tuple containing the name and buffer

Example:

```
>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())
```

named_children() → Iterator[Tuple[str, torch.nn.modules.module.Module]]

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple containing a name and child module

Example:

```
>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)
```

named_modules(memo: Optional[Set[Module]] = None, prefix: str = '')

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, l will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.named_modules()):
>>>     print(idx, '->', m)

0 -> ('', Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
))
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))
```

named_parameters(prefix: str = '', recurse: bool = True) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: prefix (str): prefix to prepend to all parameter names. recurse (bool): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

parameters (*recurse: bool = True*) → Iterator[torch.nn.parameter.Parameter]

Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

register_backward_hook (*hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]*) → torch.utils.hooks.RemovableHandle

Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex Module that perform many operations. In some failure cases, `grad_input` and `grad_output` will only contain the gradients for a subset of the inputs and outputs. For such Module, you should use `torch.Tensor.register_hook()` directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook(module, grad_input, grad_output) -> Tensor or None
```

The `grad_input` and `grad_output` may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that will be used in place of `grad_input` in subsequent computations. `grad_input` will only correspond to the inputs given as positional arguments.

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_buffer (*name: str, tensor: torch.Tensor, persistent: bool = True*) → None

Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's `running_mean` is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting `persistent` to `False`. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

name (string): name of the buffer. The buffer can be accessed from this module using the given name
tensor (Tensor): buffer to be registered. **persistent (bool):** whether the buffer is part of this module's *state_dict*.

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

register_forward_hook (hook: Callable[[], None]) → torch.utils.hooks.RemovableHandle
Registers a forward hook on the module.

The hook will be called every time after *forward()* has computed an output. It should have the following signature:

```
hook(module, input) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the *forward*. The hook can modify the output. It can modify the input inplace but it will not have effect on forward since this is called after *forward()* is called.

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling *handle.remove()*

register_forward_pre_hook (hook: Callable[[], None]) → torch.utils.hooks.RemovableHandle
Registers a forward pre-hook on the module.

The hook will be called every time before *forward()* is invoked. It should have the following signature:

```
hook(module, input) -> None or modified input
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the *forward*. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling *handle.remove()*

register_parameter (name: str, param: torch.nn.parameter.Parameter) → None
Adds a parameter to the module.

The parameter can be accessed as an attribute using given name.

Args:

name (string): name of the parameter. The parameter can be accessed from this module using the given name

param (Parameter): parameter to be added to the module.

requires_grad_ (requires_grad: bool = True) → T
Change if autograd should record operations on parameters in this module.

This method sets the parameters' *requires_grad* attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

requires_grad (bool): whether autograd should record operations on parameters in this module.
Default: True.

Returns: Module: self

share_memory () → T

state_dict (destination=None, prefix="", keep_vars=False)
Returns a dictionary containing a whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

dict: a dictionary containing a whole state of the module

Example:

```
>>> module.state_dict().keys()  
['bias', 'weight']
```

to (*args, **kwargs)

Moves and/or casts the parameters and buffers.

This can be called as

to (device=None, dtype=None, non_blocking=False)

to (dtype, non_blocking=False)

to (tensor, non_blocking=False)

to (memory_format=torch.channels_last)

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with dtypes unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

device (torch.device): the desired device of the parameters and buffers in this module

dtype (torch.dtype): the desired floating point type of the floating point parameters and buffers in this module

tensor (torch.Tensor): Tensor whose dtype and device are the desired `dtype` and `device` for all parameters and buffers in this module

memory_format (torch.memory_format): the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]])
>>> linear.to(torch.double)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]], dtype=torch.float64)
>>> gp1 = torch.device("cuda:1")
>>> linear.to(gp1, dtype=torch.half, non_blocking=True)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16)
```

train(mode: bool = True) → T

Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (True) or evaluation mode (False). Default: True.

Returns: Module: self

type(dst_type: Union[torch.dtype, str]) → T

Casts all parameters and buffers to dst_type.

Arguments: dst_type (type or string): the desired type

Returns: Module: self

zero_grad() → None

Sets gradients of all model parameters to zero.

CHAPTER 24

Basic GCP Utils

```
flood_forecast.gcp_integration.basic_utils.get_storage_client()           →  
                                                               google.cloud.storage.client.Client  
Utility function to return a properly authenticated GCS storage client whether working in Colab, CircleCI, or  
other environment.  
  
flood_forecast.gcp_integration.basic_utils.upload_file(bucket_name:      str,  
                                                       file_name:       str,    up-  
                                                       load_name:     str,    client:  
                                                               google.cloud.storage.client.Client)  
flood_forecast.gcp_integration.basic_utils.create_file_environ()
```


CHAPTER 25

Train da

```
flood_forecast.da_rnn.train_da.da_rnn(train_data: flood_forecast.da_rnn.custom_types.TrainData,  
    n_targs:     int,   encoder_hidden_size=64,   de-  
    coder_hidden_size=64,   T=10,   learning_rate=0.01,  
    batch_size=128,           param_output_path="",  
    save_path:   str   =   None)   → Tuple[dict,  
    flood_forecast.da_rnn.custom_types.DaRnnNet]  
n_targs: The number of target columns (not steps) T: The number timesteps in the window  
  
flood_forecast.da_rnn.train_da.train(net:   flood_forecast.da_rnn.custom_types.DaRnnNet,  
    train_data: flood_forecast.da_rnn.custom_types.TrainData,  
    t_cfg: flood_forecast.da_rnn.custom_types.TrainConfig,  
    train_config=",",   n_epochs=10,   save_plots=True,  
    wandb=False, tensorboard=False)  
  
flood_forecast.da_rnn.train_da.prep_train_data(batch_idx:   numpy.ndarray,   t_cfg:  
    flood_forecast.da_rnn.custom_types.TrainConfig,  
    train_data:  
    flood_forecast.da_rnn.custom_types.TrainData)  
    → Tuple  
  
flood_forecast.da_rnn.train_da.adjust_learning_rate(net:  
    flood_forecast.da_rnn.custom_types.DaRnnNet,  
    n_iter: int) → None  
  
flood_forecast.da_rnn.train_da.train_iteration(t_net: flood_forecast.da_rnn.custom_types.DaRnnNet,  
    loss_func:   Callable,   X,   y_history,  
    y_target)  
  
flood_forecast.da_rnn.train_da.predict(t_net: flood_forecast.da_rnn.custom_types.DaRnnNet,  
    t_dat: flood_forecast.da_rnn.custom_types.TrainData,  
    train_size:   int,   batch_size:   int,   T:   int,  
    on_train=False)
```


CHAPTER 26

Utils

```
flood_forecast.da_rnn.utils.setup_log(tag='VOC_TOPICS')  
flood_forecast.da_rnn.utils.save_or_show_plot(file_nm: str, save: bool, save_path='')  
flood_forecast.da_rnn.utils.numpy_to_tvar(x)
```


CHAPTER 27

Custom Types

```
class flood_forecast.da_rnn.custom_types.TrainConfig(T, train_size, batch_size,  
loss_func)
```

Bases: tuple

Create new instance of TrainConfig(T, train_size, batch_size, loss_func)

T

Alias for field number 0

train_size

Alias for field number 1

batch_size

Alias for field number 2

loss_func

Alias for field number 3

count()

Return number of occurrences of value.

index()

Return first index of value.

Raises ValueError if the value is not present.

```
class flood_forecast.da_rnn.custom_types.TrainData(feats, targs)
```

Bases: tuple

Create new instance of TrainData(feats, targs)

feats

Alias for field number 0

targs

Alias for field number 1

count()

Return number of occurrences of value.

index()

Return first index of value.

Raises `ValueError` if the value is not present.

Bases: tuple

Create new instance of DaRnnNet(encoder, decoder, enc_opt, dec_opt)

count ()

Return number of occurrences of value.

dec_opt

Alias for field number 3

decoder

Alias for field number 1

enc_opt

Alias for field number 2

encoder

Alias for field number 0

index()

Return

Rais

Raises `ValueError` if the value

CHAPTER 28

Model

```
class flood_forecast.da_rnn.model.DARNN(input_size: int, hidden_size_encoder: int, T: int,
                                         decoder_hidden_size: int, out_feats=1)
Bases: torch.nn.modules.module.Module
input size: number of underlying factors (81) T: number of time steps (10) hidden_size: dimension of the hidden state
forward(x: torch.Tensor, y_history: torch.Tensor)
    will implement
T_destination = ~T_destination
add_module(name: str, module: torch.nn.modules.module.Module) → None
    Adds a child module to the current module.
    The module can be accessed as an attribute using the given name.
Args:
    name (string): name of the child module. The child module can be accessed from this module using the given name
    module (Module): child module to be added to the module.
apply(fn: Callable[[Module], None]) → T
    Applies fn recursively to every submodule (as returned by .children()) as well as self. Typical use includes initializing the parameters of a model (see also nn-init-doc).
Args: fn (Module -> None): function to be applied to each submodule
Returns: Module: self
Example:
```

```
>>> @torch.no_grad()
>>> def init_weights(m):
>>>     print(m)
>>>     if type(m) == nn.Linear:
```

(continues on next page)

(continued from previous page)

```
>>> m.weight.fill_(1.0)
>>> print(m.weight)
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))
>>> net.apply(init_weights)
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
```

bfloat16() → T

Casts all floating point parameters and buffers to bfloat16 datatype.

Returns: Module: self

buffers(reuse: bool = True) → Iterator[torch.Tensor]

Returns an iterator over module buffers.

Args:

reuse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer

Example:

```
>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

children() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module

cpu() → T

Moves all model parameters and buffers to the CPU.

Returns: Module: self

cuda(device: Union[int, torch.device, None] = None) → T

Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self

double () → T

Casts all floating point parameters and buffers to double datatype.

Returns: Module: self

dump_patches = False

eval () → T

Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with self.train(False).

Returns: Module: self

extra_repr () → str

Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float () → T

Casts all floating point parameters and buffers to float datatype.

Returns: Module: self

half () → T

Casts all floating point parameters and buffers to half datatype.

Returns: Module: self

load_state_dict (state_dict: Dict[str, torch.Tensor], strict: bool = True)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict` is True, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Arguments:

state_dict (dict): a dict containing parameters and persistent buffers.

strict (bool, optional): whether to strictly enforce that the keys in state_dict match the keys returned by this module's state_dict() function. Default: True

Returns:

NamedTuple with missing_keys and unexpected_keys fields:

- **missing_keys** is a list of str containing the missing keys
- **unexpected_keys** is a list of str containing the unexpected keys

modules () → Iterator[torch.nn.modules.module.Module]

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

`named_buffers` (*prefix: str = ”, recurse: bool = True*) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: *prefix* (str): prefix to prepend to all buffer names. *recurse* (bool): if True, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (string, torch.Tensor): Tuple containing the name and buffer

Example:

```
>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())
```

`named_children` () → Iterator[Tuple[str, torch.nn.modules.module.Module]]

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple containing a name and child module

Example:

```
>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)
```

`named_modules` (*memo: Optional[Set[Module]] = None, prefix: str = ”*)

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> ('', Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
))
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))
```

named_parameters (*prefix: str = ”, recurse: bool = True*) → Iterator[Tuple[str, torch.Tensor]]
 Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: *prefix* (str): prefix to prepend to all parameter names. *recurse* (bool): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

parameters (*recurse: bool = True*) → Iterator[torch.nn.parameter.Parameter]

Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

register_backward_hook (*hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]*) → torch.utils.hooks.RemovableHandle
 Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex Module that perform many operations. In some failure cases, `grad_input` and `grad_output` will only contain the gradients for a subset of the inputs and outputs. For such Module, you should use `torch.Tensor.register_hook()` directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook(module, grad_input, grad_output) -> Tensor or None
```

The `grad_input` and `grad_output` may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that will be used in place of `grad_input` in subsequent computations. `grad_input` will only correspond to the inputs given as positional arguments.

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_buffer (`name: str, tensor: torch.Tensor, persistent: bool = True`) → `None`
Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's `running_mean` is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting `persistent` to `False`. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

name (string): name of the buffer. The buffer can be accessed from this module using the given name

tensor (Tensor): buffer to be registered. **persistent (bool):** whether the buffer is part of this module's `state_dict`.

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

register_forward_hook (`hook: Callable[..., None]`) → `torch.utils.hooks.RemovableHandle`
Registers a forward hook on the module.

The hook will be called every time after `forward()` has computed an output. It should have the following signature:

```
hook(module, input, output) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the output. It can modify the input inplace but it will not have effect on `forward` since this is called after `forward()` is called.

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_forward_pre_hook (`hook: Callable[..., None]`) → `torch.utils.hooks.RemovableHandle`
Registers a forward pre-hook on the module.

The hook will be called every time before `forward()` is invoked. It should have the following signature:

```
hook(module, input) -> None or modified input
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_parameter (*name: str, param: torch.nn.parameter.Parameter*) → None

Adds a parameter to the module.

The parameter can be accessed as an attribute using given name.

Args:

name (string): name of the parameter. The parameter can be accessed from this module using the given name

param (Parameter): parameter to be added to the module.

requires_grad_ (*requires_grad: bool = True*) → T

Change if autograd should record operations on parameters in this module.

This method sets the parameters' `requires_grad` attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

requires_grad (bool): whether autograd should record operations on parameters in this module.

Default: True.

Returns: Module: self

share_memory_ () → T

state_dict (*destination=None, prefix=”, keep_vars=False*)

Returns a dictionary containing a whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

dict: a dictionary containing a whole state of the module

Example:

```
>>> module.state_dict().keys()
['bias', 'weight']
```

to (*args, **kwargs)

Moves and/or casts the parameters and buffers.

This can be called as

to (*device=None, dtype=None, non_blocking=False*)

to (*dtype, non_blocking=False*)

to (*tensor, non_blocking=False*)

to (*memory_format=torch.channels_last*)

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with `dtypes` unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

device (torch.device): the desired device of the parameters and buffers in this module
dtype (torch.dtype): the desired floating point type of the floating point parameters and buffers in this module
tensor (torch.Tensor): Tensor whose dtype and device are the desired dtype and device for all parameters and buffers in this module
memory_format (torch.memory_format): the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]])
>>> linear.to(torch.double)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]], dtype=torch.float64)
>>> gpu1 = torch.device("cuda:1")
>>> linear.to(gpu1, dtype=torch.half, non_blocking=True)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16)
```

train (mode: bool = True) → T

Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (True) or evaluation mode (False). Default: True.

Returns: Module: self

type (*dst_type*: *Union[torch.dtype, str]*) → T

Casts all parameters and buffers to *dst_type*.

Arguments: *dst_type* (type or string): the desired type

Returns: Module: self

zero_grad() → None

Sets gradients of all model parameters to zero.

CHAPTER 29

Modules

```
flood_forecast.da_rnn.modules.init_hidden(x, hidden_size: int)
    Train the initial value of the hidden state: https://r2rt.com/non-zero-initial-states-for-recurrent-neural-networks.html
```

```
class flood_forecast.da_rnn.modules.Encoder(input_size: int, hidden_size: int, T: int)
    Bases: torch.nn.modules.module.Module
    input size: number of underlying factors (81) T: number of time steps (10) hidden_size: dimension of the hidden state
```

```
        forward(input_data: torch.Tensor)
```

```
        T_destination = ~T_destination
```

```
        add_module(name: str, module: torch.nn.modules.module.Module) → None
```

```
            Adds a child module to the current module.
```

```
            The module can be accessed as an attribute using the given name.
```

Args:

```
    name (string): name of the child module. The child module can be accessed from this module using the given name
```

```
    module (Module): child module to be added to the module.
```

```
        apply(fn: Callable[[Module], None]) → T
```

```
            Applies fn recursively to every submodule (as returned by .children()) as well as self. Typical use includes initializing the parameters of a model (see also nn-init-doc).
```

Args: fn (Module -> None): function to be applied to each submodule

Returns: Module: self

Example:

```
>>> @torch.no_grad()
>>> def init_weights(m):
```

(continues on next page)

(continued from previous page)

```
>>> print(m)
>>> if type(m) == nn.Linear:
>>>     m.weight.fill_(1.0)
>>>     print(m.weight)
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))
>>> net.apply(init_weights)
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
```

bfloat16() → T

Casts all floating point parameters and buffers to bfloat16 datatype.

Returns: Module: self

buffers(recurse: bool = True) → Iterator[torch.Tensor]

Returns an iterator over module buffers.

Args:

recurse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer

Example:

```
>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

children() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module

cpu() → T

Moves all model parameters and buffers to the CPU.

Returns: Module: self

cuda(device: Union[int, torch.device, None] = None) → T

Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self

double () → T

Casts all floating point parameters and buffers to double datatype.

Returns: Module: self

dump_patches = False

eval () → T

Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with `self.train(False)`.

Returns: Module: self

extra_repr () → str

Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float () → T

Casts all floating point parameters and buffers to float datatype.

Returns: Module: self

half () → T

Casts all floating point parameters and buffers to half datatype.

Returns: Module: self

load_state_dict (state_dict: Dict[str, torch.Tensor], strict: bool = True)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict` is True, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Arguments:

state_dict (dict): a dict containing parameters and persistent buffers.

strict (bool, optional): whether to strictly enforce that the keys in `state_dict` match the keys returned by this module's `state_dict()` function. Default: True

Returns:

NamedTuple with missing_keys and unexpected_keys fields:

- **missing_keys** is a list of str containing the missing keys
- **unexpected_keys** is a list of str containing the unexpected keys

modules () → Iterator[torch.nn.modules.module.Module]

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, 1 will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

`named_buffers` (*prefix: str = ”, recurse: bool = True*) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: *prefix* (str): prefix to prepend to all buffer names. *recurse* (bool): if True, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (string, torch.Tensor): Tuple containing the name and buffer

Example:

```
>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())
```

`named_children` () → Iterator[Tuple[str, torch.nn.modules.module.Module]]

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple containing a name and child module

Example:

```
>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)
```

`named_modules` (*memo: Optional[Set[Module]] = None, prefix: str = ”*)

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> ('', Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
))
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))
```

named_parameters (*prefix: str = ”, recurse: bool = True*) → Iterator[Tuple[str, torch.Tensor]]
 Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: *prefix* (str): prefix to prepend to all parameter names. *recurse* (bool): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

parameters (*recurse: bool = True*) → Iterator[torch.nn.parameter.Parameter]
 Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

register_backward_hook (*hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]*) → torch.utils.hooks.RemovableHandle
 Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex Module that perform many operations. In some failure cases, `grad_input` and `grad_output` will only contain the gradients for a subset of the inputs and outputs. For such Module, you should use `torch.Tensor.register_hook()` directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook(module, grad_input, grad_output) -> Tensor or None
```

The `grad_input` and `grad_output` may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that will be used in place of `grad_input` in subsequent computations. `grad_input` will only correspond to the inputs given as positional arguments.

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_buffer (`name: str, tensor: torch.Tensor, persistent: bool = True`) → `None`
Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's `running_mean` is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting `persistent` to `False`. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

name (string): name of the buffer. The buffer can be accessed from this module using the given name

tensor (Tensor): buffer to be registered. persistent (bool): whether the buffer is part of this module's `state_dict`.

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

register_forward_hook (`hook: Callable[..., None]`) → `torch.utils.hooks.RemovableHandle`
Registers a forward hook on the module.

The hook will be called every time after `forward()` has computed an output. It should have the following signature:

```
hook(module, input, output) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the output. It can modify the input inplace but it will not have effect on `forward` since this is called after `forward()` is called.

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_forward_pre_hook (`hook: Callable[..., None]`) → `torch.utils.hooks.RemovableHandle`
Registers a forward pre-hook on the module.

The hook will be called every time before `forward()` is invoked. It should have the following signature:

```
hook(module, input) -> None or modified input
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_parameter (*name: str, param: torch.nn.parameter.Parameter*) → None

Adds a parameter to the module.

The parameter can be accessed as an attribute using given name.

Args:

name (string): name of the parameter. The parameter can be accessed from this module using the given name

param (Parameter): parameter to be added to the module.

requires_grad_ (*requires_grad: bool = True*) → T

Change if autograd should record operations on parameters in this module.

This method sets the parameters' `requires_grad` attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

requires_grad (bool): whether autograd should record operations on parameters in this module.

Default: True.

Returns: Module: self

share_memory_ () → T

state_dict (*destination=None, prefix=”, keep_vars=False*)

Returns a dictionary containing a whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

dict: a dictionary containing a whole state of the module

Example:

```
>>> module.state_dict().keys()
['bias', 'weight']
```

to (*args, **kwargs)

Moves and/or casts the parameters and buffers.

This can be called as

to (*device=None, dtype=None, non_blocking=False*)

to (*dtype, non_blocking=False*)

to (*tensor, non_blocking=False*)

to (*memory_format=torch.channels_last*)

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with `dtypes` unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

device (torch.device): the desired device of the parameters and buffers in this module
dtype (torch.dtype): the desired floating point type of the floating point parameters and buffers in this module
tensor (torch.Tensor): Tensor whose dtype and device are the desired dtype and device for all parameters and buffers in this module
memory_format (torch.memory_format): the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]])
>>> linear.to(torch.double)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]], dtype=torch.float64)
>>> gpu1 = torch.device("cuda:1")
>>> linear.to(gpu1, dtype=torch.half, non_blocking=True)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16)
```

train (mode: bool = True) → T

Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (True) or evaluation mode (False). Default: True.

Returns: Module: self

type (*dst_type*: *Union[torch.dtype, str]*) → T
 Casts all parameters and buffers to *dst_type*.

Arguments: *dst_type* (type or string): the desired type

Returns: Module: self

zero_grad() → None
 Sets gradients of all model parameters to zero.

```
class flood_forecast.da_rnn.modules.Decoder(encoder_hidden_size: int, decoder_hidden_size: int, T: int, out_feats=1)
Bases: torch.nn.modules.module.Module
```

forward (*input_encoded*, *y_history*)

T_destination = ~T_destination

add_module (*name*: str, *module*: *torch.nn.modules.module.Module*) → None
 Adds a child module to the current module.
 The module can be accessed as an attribute using the given name.

Args:

- name (string): name of the child module. The child module can be** accessed from this module using the given name
- module (Module): child module to be added to the module.**

apply (*fn*: *Callable[[Module], None]*) → T
 Applies *fn* recursively to every submodule (as returned by *.children()*) as well as self. Typical use includes initializing the parameters of a model (see also nn-init-doc).

Args: *fn* (Module -> None): function to be applied to each submodule

Returns: Module: self

Example:

```
>>> @torch.no_grad()
>>> def init_weights(m):
>>>     print(m)
>>>     if type(m) == nn.Linear:
>>>         m.weight.fill_(1.0)
>>>         print(m.weight)
>>> net = nn.Sequential(nn.Linear(2, 2), nn.Linear(2, 2))
>>> net.apply(init_weights)
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Linear(in_features=2, out_features=2, bias=True)
Parameter containing:
tensor([[ 1.,  1.],
       [ 1.,  1.]])
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
  (1): Linear(in_features=2, out_features=2, bias=True)
)
Sequential(
  (0): Linear(in_features=2, out_features=2, bias=True)
```

(continues on next page)

(continued from previous page)

```
(1): Linear(in_features=2, out_features=2, bias=True)
)
```

bfloat16() → T

Casts all floating point parameters and buffers to bfloat16 datatype.

Returns: Module: self**buffers(reuse: bool = True)** → Iterator[torch.Tensor]

Returns an iterator over module buffers.

Args:

reuse (bool): if True, then yields buffers of this module and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: torch.Tensor: module buffer**Example:**

```
>>> for buf in model.buffers():
>>>     print(type(buf), buf.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

children() → Iterator[torch.nn.modules.module.Module]

Returns an iterator over immediate children modules.

Yields: Module: a child module**cpu()** → T

Moves all model parameters and buffers to the CPU.

Returns: Module: self**cuda(device: Union[int, torch.device, None] = None)** → T

Moves all model parameters and buffers to the GPU.

This also makes associated parameters and buffers different objects. So it should be called before constructing optimizer if the module will live on GPU while being optimized.

Arguments:

device (int, optional): if specified, all parameters will be copied to that device

Returns: Module: self**double()** → T

Casts all floating point parameters and buffers to double datatype.

Returns: Module: self**dump_patches = False****eval()** → T

Sets the module in evaluation mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

This is equivalent with `self.train(False)`.

Returns: Module: self

extra_repr () → str

Set the extra representation of the module

To print customized extra information, you should reimplement this method in your own modules. Both single-line and multi-line strings are acceptable.

float () → T

Casts all floating point parameters and buffers to float datatype.

Returns: Module: self

half () → T

Casts all floating point parameters and buffers to half datatype.

Returns: Module: self

load_state_dict (state_dict: Dict[str, torch.Tensor], strict: bool = True)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict` is True, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Arguments:

state_dict (dict): a dict containing parameters and persistent buffers.

strict (bool, optional): whether to strictly enforce that the keys in `state_dict` match the keys returned by this module's `state_dict()` function. Default: True

Returns:

NamedTuple with missing_keys and unexpected_keys fields:

- **missing_keys** is a list of str containing the missing keys
- **unexpected_keys** is a list of str containing the unexpected keys

modules () → Iterator[torch.nn.modules.module.Module]

Returns an iterator over all modules in the network.

Yields: Module: a module in the network

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.modules()):
    print(idx, '->', m)

0 -> Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
)
1 -> Linear(in_features=2, out_features=2, bias=True)
```

named_buffers (prefix: str = "", recurse: bool = True) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself.

Args: prefix (str): prefix to prepend to all buffer names. recurse (bool): if True, then yields buffers of this module

and all submodules. Otherwise, yields only buffers that are direct members of this module.

Yields: (string, torch.Tensor): Tuple containing the name and buffer

Example:

```
>>> for name, buf in self.named_buffers():
>>>     if name in ['running_var']:
>>>         print(buf.size())
```

named_children() → Iterator[Tuple[str, torch.nn.modules.module.Module]]

Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple containing a name and child module

Example:

```
>>> for name, module in model.named_children():
>>>     if name in ['conv4', 'conv5']:
>>>         print(module)
```

named_modules(memo: Optional[Set[Module]] = None, prefix: str = '')

Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself.

Yields: (string, Module): Tuple of name and module

Note: Duplicate modules are returned only once. In the following example, `l` will be returned only once.

Example:

```
>>> l = nn.Linear(2, 2)
>>> net = nn.Sequential(l, l)
>>> for idx, m in enumerate(net.named_modules()):
>>>     print(idx, '->', m)

0 -> ('', Sequential(
    (0): Linear(in_features=2, out_features=2, bias=True)
    (1): Linear(in_features=2, out_features=2, bias=True)
))
1 -> ('0', Linear(in_features=2, out_features=2, bias=True))
```

named_parameters(prefix: str = '', recurse: bool = True) → Iterator[Tuple[str, torch.Tensor]]

Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself.

Args: prefix (str): prefix to prepend to all parameter names. recurse (bool): if True, then yields parameters of this module

and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: (string, Parameter): Tuple containing the name and parameter

Example:

```
>>> for name, param in self.named_parameters():
>>>     if name in ['bias']:
>>>         print(param.size())
```

parameters(recurse: bool = True) → Iterator[torch.nn.parameter.Parameter]

Returns an iterator over module parameters.

This is typically passed to an optimizer.

Args:

recurse (bool): if True, then yields parameters of this module and all submodules. Otherwise, yields only parameters that are direct members of this module.

Yields: Parameter: module parameter

Example:

```
>>> for param in model.parameters():
>>>     print(type(param), param.size())
<class 'torch.Tensor'> (20L,)
<class 'torch.Tensor'> (20L, 1L, 5L, 5L)
```

register_backward_hook (hook: Callable[[Module, Union[Tuple[torch.Tensor, ...], torch.Tensor], Union[Tuple[torch.Tensor, ...], torch.Tensor]], Union[None, torch.Tensor]]) → torch.utils.hooks.RemovableHandle

Registers a backward hook on the module.

Warning: The current implementation will not have the presented behavior for complex Module that perform many operations. In some failure cases, `grad_input` and `grad_output` will only contain the gradients for a subset of the inputs and outputs. For such Module, you should use `torch.Tensor.register_hook()` directly on a specific input or output to get the required gradients.

The hook will be called every time the gradients with respect to module inputs are computed. The hook should have the following signature:

```
hook (module, grad_input, grad_output) -> Tensor or None
```

The `grad_input` and `grad_output` may be tuples if the module has multiple inputs or outputs. The hook should not modify its arguments, but it can optionally return a new gradient with respect to input that will be used in place of `grad_input` in subsequent computations. `grad_input` will only correspond to the inputs given as positional arguments.

Returns:

torch.utils.hooks.RemovableHandle: a handle that can be used to remove the added hook by calling `handle.remove()`

register_buffer (name: str, tensor: torch.Tensor, persistent: bool = True) → None

Adds a buffer to the module.

This is typically used to register a buffer that should not to be considered a model parameter. For example, BatchNorm's `running_mean` is not a parameter, but is part of the module's state. Buffers, by default, are persistent and will be saved alongside parameters. This behavior can be changed by setting `persistent` to `False`. The only difference between a persistent buffer and a non-persistent buffer is that the latter will not be a part of this module's `state_dict`.

Buffers can be accessed as attributes using given names.

Args:

name (string): name of the buffer. The buffer can be accessed from this module using the given name

tensor (Tensor): buffer to be registered. persistent (bool): whether the buffer is part of this module's `state_dict`.

Example:

```
>>> self.register_buffer('running_mean', torch.zeros(num_features))
```

register_forward_hook (hook: Callable[..., None]) → torch.utils.hooks.RemovableHandle
Registers a forward hook on the module.

The hook will be called every time after `forward()` has computed an output. It should have the following signature:

```
hook(module, input, output) -> None or modified output
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the output. It can modify the input inplace but it will not have effect on forward since this is called after `forward()` is called.

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

register_forward_pre_hook (hook: Callable[..., None]) → torch.utils.hooks.RemovableHandle
Registers a forward pre-hook on the module.

The hook will be called every time before `forward()` is invoked. It should have the following signature:

```
hook(module, input) -> None or modified input
```

The input contains only the positional arguments given to the module. Keyword arguments won't be passed to the hooks and only to the `forward`. The hook can modify the input. User can either return a tuple or a single modified value in the hook. We will wrap the value into a tuple if a single value is returned(unless that value is already a tuple).

Returns:

`torch.utils.hooks.RemovableHandle`: a handle that can be used to remove the added hook by calling `handle.remove()`

register_parameter (name: str, param: torch.nn.parameter.Parameter) → None
Adds a parameter to the module.

The parameter can be accessed as an attribute using given name.

Args:

name (string): name of the parameter. The parameter can be accessed from this module using the given name

param (Parameter): parameter to be added to the module.

requires_grad_ (requires_grad: bool = True) → T
Change if autograd should record operations on parameters in this module.

This method sets the parameters' `requires_grad` attributes in-place.

This method is helpful for freezing part of the module for finetuning or training parts of a model individually (e.g., GAN training).

Args:

requires_grad (bool): whether autograd should record operations on parameters in this module.
Default: True.

Returns: Module: self

share_memory() → T

state_dict (*destination=None*, *prefix=*"", *keep_vars=False*)

Returns a dictionary containing a whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names.

Returns:

dict: a dictionary containing a whole state of the module

Example:

```
>>> module.state_dict().keys()
['bias', 'weight']
```

to (*args, **kwargs)

Moves and/or casts the parameters and buffers.

This can be called as

to (*device=None*, *dtype=None*, *non_blocking=False*)

to (*dtype*, *non_blocking=False*)

to (*tensor*, *non_blocking=False*)

to (*memory_format=torch.channels_last*)

Its signature is similar to `torch.Tensor.to()`, but only accepts floating point desired `dtype`s. In addition, this method will only cast the floating point parameters and buffers to `dtype` (if given). The integral parameters and buffers will be moved `device`, if that is given, but with dtypes unchanged. When `non_blocking` is set, it tries to convert/move asynchronously with respect to the host if possible, e.g., moving CPU Tensors with pinned memory to CUDA devices.

See below for examples.

Note: This method modifies the module in-place.

Args:

device (`torch.device`): the desired device of the parameters and buffers in this module

dtype (`torch.dtype`): the desired floating point type of the floating point parameters and buffers in this module

tensor (`torch.Tensor`): Tensor whose `dtype` and `device` are the desired `dtype` and `device` for all parameters and buffers in this module

memory_format (`torch.memory_format`): the desired memory format for 4D parameters and buffers in this module (keyword only argument)

Returns: Module: self

Example:

```
>>> linear = nn.Linear(2, 2)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
```

(continues on next page)

(continued from previous page)

```

        [-0.5113, -0.2325]])
>>> linear.to(torch.double)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1913, -0.3420],
       [-0.5113, -0.2325]], dtype=torch.float64)
>>> gpul = torch.device("cuda:1")
>>> linear.to(gpul, dtype=torch.half, non_blocking=True)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16, device='cuda:1')
>>> cpu = torch.device("cpu")
>>> linear.to(cpu)
Linear(in_features=2, out_features=2, bias=True)
>>> linear.weight
Parameter containing:
tensor([[ 0.1914, -0.3420],
       [-0.5112, -0.2324]], dtype=torch.float16)

```

train(*mode: bool = True*) → T

Sets the module in training mode.

This has any effect only on certain modules. See documentations of particular modules for details of their behaviors in training/evaluation mode, if they are affected, e.g. Dropout, BatchNorm, etc.

Args:

mode (bool): whether to set training mode (True) or evaluation mode (False). Default: True.

Returns: Module: self**type**(*dst_type: Union[torch.dtype, str]*) → TCasts all parameters and buffers to *dst_type*.**Arguments:** *dst_type* (type or string): the desired type**Returns:** Module: self**zero_grad()** → None

Sets gradients of all model parameters to zero.

CHAPTER 30

Indices and tables

- genindex
- modindex
- search

Python Module Index

f

flood_forecast, ??
flood_forecast.custom, 35
flood_forecast.custom.custom_opt, 37
flood_forecast.da_rnn, 185
flood_forecast.da_rnn.custom_types, 191
flood_forecast.da_rnn.model, 193
flood_forecast.da_rnn.modules, 203
flood_forecast.da_rnn.train_da, 187
flood_forecast.da_rnn.utils, 189
flood_forecast.evaluator, 3
flood_forecast.gcp_integration, 183
flood_forecast.gcp_integration.basic_utils,
 185
flood_forecast.long_train, 7
flood_forecast.model_dict_function, 9
flood_forecast.pre_dict, 11
flood_forecast.preprocessing, 17
flood_forecast.preprocessing.buil_dataset,
 21
flood_forecast.preprocessing.closest_station,
 23
flood_forecast.preprocessing.data_converter,
 25
flood_forecast.preprocessing.interpolate_preprocess,
 19
flood_forecast.preprocessing.preprocess_da_rnn,
 27
flood_forecast.preprocessing.preprocess_metadata,
 29
flood_forecast.preprocessing.process_usgs,
 31
flood_forecast.preprocessing.pytorch_loaders,
 33
flood_forecast.preprocessing.temporal_feats,
 35
flood_forecast.pytorch_training, 13
flood_forecast.time_model, 15
flood_forecast.trainer, 17

flood_forecast.transformer_xl, 68
flood_forecast.transformer_xl.dummy_torch,
 69
flood_forecast.transformer_xl.lower_upper_config,
 77
flood_forecast.transformer_xl.multi_head_base,
 101
flood_forecast.transformer_xl.transformer_basic,
 111
flood_forecast.transformer_xl.transformer_xl,
 137
flood_forecast.utils, 1

Index

A

A

```
add_module() (flood_forecast.transformer_xl.transformer_xl.StandardV
add_module() (flood_forecast.custom.custom_opt.GaussianLoss method), 168
    method), 52
add_module() (flood_forecast.transformer_xl.transformer_xl.Transform
    method), 176
add_module() (flood_forecast.custom.custom_opt.MAPELoss method), 45
add_param_group()
add_module() (flood_forecast.custom.custom_opt.QuantileLoss method), 60
(flood_forecast.custom.custom_opt.BertAdam
method), 68
add_module() (flood_forecast.custom.custom_opt.RMSELoss adjust_learning_rate() (in module
method), 37
flood_forecast.da_rnn.train_da), 187
add_module() (flood_forecast.da_rnn.model.DARNN apply() (flood_forecast.custom.custom_opt.GaussianLoss
method), 193
method), 53
add_module() (flood_forecast.da_rnn.modules.Decoder apply() (flood_forecast.custom.custom_opt.MAPELoss
method), 211
method), 45
add_module() (flood_forecast.da_rnn.modules.Encoder apply() (flood_forecast.custom.custom_opt.QuantileLoss
method), 203
method), 60
add_module() (flood_forecast.transformer_xl.dummy_tokenize) apply() (flood_forecast.custom.custom_opt.RMSELoss
method), 69
method), 37
add_module() (flood_forecast.transformer_xl.lower_upper_upper_config.AR apply() (flood_forecast.da_rnn.model.DARNN
method), 85
method), 193
add_module() (flood_forecast.transformer_xl.lower_upper_upper_config.MetaEmbedding) apply() (flood_forecast.da_rnn.modules.Decoder
method), 92
method), 211
add_module() (flood_forecast.transformer_xl.lower_upper_upper_config.PositionwiseFeedForward
method), 77
apply() (flood_forecast.da_rnn.modules.Encoder
method), 203
add_module() (flood_forecast.transformer_xl.multi_head_base) apply() (flood_forecast.transformer_xl.dummy_torch.DummyTorchMode
method), 101
method), 69
add_module() (flood_forecast.transformer_xl.transformer) apply() (flood_forecast.transformer_xl.lower_upper_config.AR
method), 119
method), 85
add_module() (flood_forecast.transformer_xl.transformer) apply() (flood_forecast.transformer_xl.lower_upper_config.MetaEmbedding
method), 127
method), 93
add_module() (flood_forecast.transformer_xl.transformer) apply() (flood_forecast.transformer_xl.lower_upper_config.Positionwise
method), 112
method), 77
add_module() (flood_forecast.transformer_xl.transformer) apply() (flood_forecast.transformer_xl.multi_head_base.MultiAttnHead
method), 153
method), 101
add_module() (flood_forecast.transformer_xl.transformer) apply() (flood_forecast.transformer_xl.transformer_basic.CustomTransfo
method), 137
method), 120
add_module() (flood_forecast.transformer_xl.transformer) apply() (flood_forecast.transformer_xl.transformer_basic.SimplePosition
method), 160
method), 127
add_module() (flood_forecast.transformer_xl.transformer) apply() (flood_forecast.transformer_xl.transformer_basic.SimpleTransfo
method), 145
method), 112
apply() (flood_forecast.transformer_xl.transformer_xl.DecoderBlock
```

```

        method), 153
apply() (flood_forecast.transformer_xl.transformer_xl.MultiHeadAttention),
        method), 153
apply() (flood_forecast.transformer_xl.transformer_xl.PositionalEmbedding),
        method), 138
apply() (flood_forecast.transformer_xl.transformer_xl.PositionwiseFFN),
        method), 160
apply() (flood_forecast.transformer_xl.transformer_xl.StandardWordEmbedding),
        method), 168
apply() (flood_forecast.transformer_xl.transformer_xl.TransformerXLMethod),
        method), 169
AR(class in flood_forecast.transformer_xl.lower_upper_config),
        method), 177
AR.to() (in module
        method), 53
batch_size(flood_forecast.da_rnn.custom_types.TrainConfig),
        (flood_forecast.da_rnn.model.DARNN
attribute), 191
BertAdam(class in flood_forecast.custom.custom_opt),
        buffers() (flood_forecast.da_rnn.modules.Decoder
method), 67
bf16(flood_forecast.custom.custom_opt.GaussianLoss),
        buffers() (flood_forecast.da_rnn.modules.Encoder
method), 53
bf16(flood_forecast.custom.custom_opt.MAPELoss),
        buffers() (flood_forecast.transformer_xl.dummy_torch.DummyTorchModule
method), 46
bf16(flood_forecast.custom.custom_opt.QuantileLoss),
        buffers() (flood_forecast.transformer_xl.lower_upper_config.AR
method), 61
bf16(flood_forecast.custom.custom_opt.RMSELoss),
        buffers() (flood_forecast.transformer_xl.lower_upper_config.MetaEmbedding
method), 38
bf16(flood_forecast.da_rnn.model.DARNN),
        buffers() (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFFN
method), 194
bf16(flood_forecast.da_rnn.modules.Decoder),
        buffers() (flood_forecast.transformer_xl.multi_head_base.MultiAttnHead
method), 212
bf16(flood_forecast.da_rnn.modules.Encoder),
        buffers() (flood_forecast.transformer_xl.transformer_basic.CustomTransformer
method), 204
bf16(flood_forecast.transformer_xl.dummy_torch.DummyTorchModule),
        (flood_forecast.transformer_xl.transformer_basic.SimplePositionalEmbedding
method), 70
bf16(flood_forecast.transformer_xl.lower_upper_boundConfig),
        (flood_forecast.transformer_xl.transformer_basic.SimpleTransformer
method), 85
bf16(flood_forecast.transformer_xl.lower_upper_boundMetaEmbedding),
        (flood_forecast.transformer_xl.decoder.DecoderBlock
method), 93
bf16(flood_forecast.transformer_xl.lower_upper_boundPositionalEmbedding),
        (flood_forecast.transformer_xl.multi_head_base.MultiAttnHead
method), 78
bf16(flood_forecast.transformer_xl.multi_head_base.MultiAttnHead),
        (flood_forecast.transformer_xl.transformer_xl.PositionalEmbedding
method), 102
bf16(flood_forecast.transformer_xl.transformer_xl.MultiHeadAttention),
        (flood_forecast.transformer_xl.transformer_xl.PositionwiseFFN
method), 138
bf16(flood_forecast.transformer_xl.transformer_xl.PositionwiseFFN),
        (flood_forecast.transformer_xl.transformer_xl.StandardWordEmbedding
method), 161
bf16(flood_forecast.transformer_xl.transformer_xl.SimplePositionalEmbedding),
        (flood_forecast.transformer_xl.transformer_xl.TransformerXLMethod
method), 169
bf16(flood_forecast.transformer_xl.transformer_xl.SimpleTransformer),
        (flood_forecast.transformer_xl.transformer_xl.TransformerXLMethod
method), 177

```

```

build_weather_csv()           (in      module   convert_history_batches()
    flood_forecast.preprocessing.buil_dataset),
    21                               (flood_forecast.preprocessing.pytorch_loaders.CSVTestLoader
                                         method), 34
                                         convert_real_batches()
                                         (flood_forecast.preprocessing.pytorch_loaders.CSVTestLoader
                                         method), 34
C
check_loss()      (flood_forecast.utils.EarlyStopper
    method), 1
children() (flood_forecast.custom.custom_opt.GaussianLoss
    method), 53
children() (flood_forecast.custom.custom_opt.MAPELoss)
    count () (flood_forecast.da_rnn.custom_types.DaRnnNet
    method), 192
children() (flood_forecast.custom.custom_opt.QuantileLoss)
    count () (flood_forecast.da_rnn.custom_types.TrainConfig
    method), 191
children() (flood_forecast.custom.custom_opt.RMSELoss)
    count () (flood_forecast.da_rnn.custom_types.TrainData
    method), 191
children() (flood_forecast.da_rnn.model.DARNN
    method), 194
children() (flood_forecast.da_rnn.modules.Decoder
    method), 212
children() (flood_forecast.da_rnn.modules.Encoder
    method), 204
children() (flood_forecast.transformer_xl.dummy_torch.DummyTorchModel)
    config (flood_forecast.custom.custom_opt.RMSELoss
    method), 38
children() (flood_forecast.transformer_xl.lower_upper_config.AR
    method), 86
children() (flood_forecast.transformer_xl.lower_upper_config.MetaEmbedding)
    method), 94
children() (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFeedForward)
    method), 78
children() (flood_forecast.transformer_xl.multi_head_base.MultiHeadSim
    method), 102
children() (flood_forecast.transformer_xl.transformer_basic.Ci
    method), 121
children() (flood_forecast.transformer_xl.transformer_basic.Si
    method), 128
children() (flood_forecast.transformer_xl.transformer_basic.Si
    method), 113
children() (flood_forecast.transformer_xl.transformer_xl.DecoderBlock)
    method), 154
children() (flood_forecast.transformer_xl.transformer_xl.MultiHeadAttention)
    method), 138
children() (flood_forecast.transformer_xl.transformer_xl.PositionalEmbedding)
    method), 161
children() (flood_forecast.transformer_xl.transformer_xl.PositionwiseFF)
    method), 146
children() (flood_forecast.transformer_xl.transformer_xl.StandardWordEmbedding)
    method), 169
children() (flood_forecast.transformer_xl.transformer_xl.Transformers)
    method), 177
combine_data()           (in      module   cpu ()
    flood_forecast.preprocessing.buil_dataset),
    21                               (flood_forecast.transformer_xl.transformer_xl.PositionalEmbedding
                                         method), 161
                                         cpu ()
                                         (flood_forecast.transformer_xl.transformer_xl.PositionwiseFF
                                         method), 146
                                         cpu ()
                                         (flood_forecast.transformer_xl.transformer_xl.StandardWordEmbedding
                                         method), 139

```

method), 169
 cpu () (flood_forecast.transformer_xl.transformer_xl.TransformerXLmethod), 161
 method), 177
 create_csv () (in module flood_forecast.preprocessing.process_usgs), 31
 create_file_environ () (in module flood_forecast.gcp_integration.basic_utils), 185
 create_usgs () (in module flood_forecast.preprocessing.buil_dataset), 21
 create_visited () (in module flood_forecast.preprocessing.buil_dataset), 21
 CSVDataLoader (class in flood_forecast.preprocessing.pytorch_loaders), 33
 CSVTTestLoader (class in flood_forecast.preprocessing.pytorch_loaders), 34
 cuda () (flood_forecast.custom.custom_opt.GaussianLoss method), 54
 cuda () (flood_forecast.custom.custom_opt.MAPELoss method), 46
 cuda () (flood_forecast.custom.custom_opt.QuantileLoss method), 61
 cuda () (flood_forecast.custom.custom_opt.RMSELoss method), 38
 cuda () (flood_forecast.da_rnn.model.DARNN method), 194
 cuda () (flood_forecast.da_rnn.modules.Decoder method), 212
 cuda () (flood_forecast.da_rnn.modules.Encoder method), 204
 cuda () (flood_forecast.transformer_xl.dummy_torch.DummyTorchModel method), 70
 cuda () (flood_forecast.transformer_xl.lower_upper_config.AR method), 86
 cuda () (flood_forecast.transformer_xl.lower_upper_config.MetaEmbedding method), 94
 cuda () (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFeedForward method), 78
 cuda () (flood_forecast.transformer_xl.multi_head_base.MultiAttrHeadSimpleForecast.custom.custom_opt.GaussianLoss method), 54
 cuda () (flood_forecast.transformer_xl.multi_head_base.MultiAttrHeadSimpleForecast.custom.custom_opt.MAPELoss method), 46
 cuda () (flood_forecast.transformer_xl.transformer_basic.CustomTransformerDecoder.custom.custom_opt.QuantileLoss method), 61
 cuda () (flood_forecast.transformer_xl.transformer_basic.SimplePositionalEmbeddingSimpleForecast.custom.custom_opt.RMSELoss method), 39
 cuda () (flood_forecast.transformer_xl.transformer_basic.SimpleTransformer(flood_forecast.da_rnn.model.DARNN method), 195
 cuda () (flood_forecast.transformer_xl.transformer_xl.DecoderBlock) (flood_forecast.da_rnn.modules.Decoder method), 212
 cuda () (flood_forecast.transformer_xl.transformer_xl.MultiHeadAttention(flood_forecast.da_rnn.modules.Encoder method), 205

D

da_rnn () (in module flood_forecast.da_rnn.train_da), 187
 DARNN (class in flood_forecast.da_rnn.model), 193
 DARNN.to () (in module flood_forecast.da_rnn.model), 199
 DaRnnNet (class in flood_forecast.da_rnn.custom_types), 192
 dec_opt (flood_forecast.da_rnn.custom_types.DaRnnNet attribute), 192
 decode_seq () (flood_forecast.transformer_xl.transformer_basic.Simple method), 111
 Decoder (class in flood_forecast.da_rnn.modules), 211
 decoder (flood_forecast.da_rnn.custom_types.DaRnnNet attribute), 192
 Decoder.to () (in module flood_forecast.da_rnn.modules), 217
 DecoderBlock (class in flood_forecast.transformer_xl.transformer_xl),
 DecoderBlock.to () (in module flood_forecast.transformer_xl.transformer_xl), 159
 MetaEmbedding (in module flood_forecast.preprocessing.process_usgs), 31
 PositionwiseFeedForward (flood_forecast.custom.custom_opt.GaussianLoss method), 54
 PositionwiseFeedForward (flood_forecast.custom.custom_opt.MAPELoss method), 46
 PositionwiseFeedForward (flood_forecast.custom.custom_opt.QuantileLoss method), 61
 PositionwiseFeedForward (flood_forecast.custom.custom_opt.RMSELoss method), 39
 PositionwiseFeedForward (flood_forecast.da_rnn.model.DARNN method), 195
 PositionwiseFeedForward (flood_forecast.da_rnn.modules.Decoder method), 212
 PositionwiseFeedForward (flood_forecast.da_rnn.modules.Encoder method), 205

double () (flood_forecast.transformer_xl.dummy_torch.DummyTorchModel (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFF method), 70
 double () (flood_forecast.transformer_xl.lower_upper_config.ARpatches (flood_forecast.transformer_xl.multi_head_base.MultiAttentionHead method), 86
 double () (flood_forecast.transformer_xl.lower_upper_config.MetaHeadedit (flood_forecast.transformer_xl.transformer_basic.CustomHead method), 94
 double () (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFF (flood_forecast.transformer_xl.transformer_basic.SimplePositionwiseFF method), 79
 double () (flood_forecast.transformer_xl.multi_head_base.MultiAttentionHeadSimple (flood_forecast.transformer_xl.transformer_basic.SimplePositionwiseFF method), 103
 double () (flood_forecast.transformer_xl.transformer_basic.GhostPatchEncoder (flood_forecast.transformer_xl.transformer_xl.DecoderBlock method), 121
 double () (flood_forecast.transformer_xl.transformer_basic.SimplePositionwiseFF (flood_forecast.transformer_xl.transformer_xl.MultiHead method), 129
 double () (flood_forecast.transformer_xl.transformer_basic.SimpleTransformer (flood_forecast.transformer_xl.transformer_xl.PositionwiseFF method), 113
 double () (flood_forecast.transformer_xl.transformer_xl.PositionwiseFF (flood_forecast.transformer_xl.transformer_xl.PositionwiseFF method), 139
 double () (flood_forecast.transformer_xl.transformer_xl.DecoderBlock (flood_forecast.transformer_xl.transformer_xl.PositionwiseFF method), 154
 double () (flood_forecast.transformer_xl.transformer_xl.MultiHeadAttention (flood_forecast.transformer_xl.transformer_xl.StandardWiseAttention method), 139
 double () (flood_forecast.transformer_xl.transformer_xl.PositionwiseFF (flood_forecast.transformer_xl.transformer_xl.Transformers method), 162
 double () (flood_forecast.transformer_xl.transformer_xl.PositionwiseFF (flood_forecast.transformer_xl.transformer_xl.PositionwiseFF method), 146
 double () (flood_forecast.transformer_xl.transformer_xl.StandardWordEmbedding (flood_forecast.utils), 1
 double () (flood_forecast.transformer_xl.transformer_xl.TransformerXL (flood_forecast.da_rnn.custom_types.DaRnnNet attribute), 192
 DummyTorchModel (class in flood_forecast.transformer_xl.dummy_torch), 69
 DummyTorchModel.to () (in module flood_forecast.transformer_xl.dummy_torch), 75
 dump_patches (flood_forecast.custom.custom_opt.GaussianLoss attribute), 54
 dump_patches (flood_forecast.custom.custom_opt.MAPELoss attribute), 46
 dump_patches (flood_forecast.custom.custom_opt.QuantileLoss attribute), 61
 dump_patches (flood_forecast.custom.custom_opt.RMSELoss attribute), 39
 dump_patches (flood_forecast.da_rnn.model.DARNN attribute), 195
 dump_patches (flood_forecast.da_rnn.modules.Decoder attribute), 212
 dump_patches (flood_forecast.da_rnn.modules.Encoder attribute), 205
 dump_patches (flood_forecast.transformer_xl.dummy_torch.DummyTorchModel attribute), 71
 dump_patches (flood_forecast.transformer_xl.lower_upper_config.ARmethod), 71
 dump_patches (flood_forecast.transformer_xl.lower_upper_config.MetaHeadSimple (flood_forecast.transformer_xl.lower_upper_config.AR attribute), 86
 dump_patches (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFF (flood_forecast.transformer_xl.lower_upper_config.AR attribute), 94

```

eval() (flood_forecast.transformer_xl.lower_upper_config.MetaEmbedding (flood_forecast.transformer_xl.transformer_basic.Simple
method), 94
eval() (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFeedForward (flood_forecast.transformer_xl.transformer_xl.DecoderB
method), 154
eval() (flood_forecast.transformer_xl.multi_head_base.MultiAttnHeadSimple (flood_forecast.transformer_xl.transformer_xl.MultiHead
method), 103
eval() (flood_forecast.transformer_xl.transformer_basic.CustomTransformerDecoder (flood_forecast.transformer_xl.transformer_xl.Positional
method), 121
eval() (flood_forecast.transformer_xl.transformer_basic.SimplePositionalEncoding (flood_forecast.transformer_xl.transformer_xl.Positional
method), 129
eval() (flood_forecast.transformer_xl.transformer_basic.SimpleTransformerDecoder (flood_forecast.transformer_xl.transformer_xl.Positional
method), 113
eval() (flood_forecast.transformer_xl.transformer_xl.PositionwiseFeedForward (flood_forecast.transformer_xl.transformer_xl.DecoderB
method), 139
eval() (flood_forecast.transformer_xl.transformer_xl.StandardWordEmbedding (flood_forecast.transformer_xl.transformer_xl.StandardW
method), 162
eval() (flood_forecast.transformer_xl.transformer_xl.DecoderBlock (flood_forecast.transformer_xl.transformer_xl.Transform
method), 178
eval() (flood_forecast.transformer_xl.transformer_xl.MultiHeadAttention
F
eval() (flood_forecast.transformer_xl.transformer_xl.PositionwiseFeedForward (flood_forecast.da_rnn.custom_types.TrainData
method), 162
eval() (flood_forecast.transformer_xl.transformer_xl.PositionwiseFeedForward (in module
method), 146 flood_forecast.preprocessing.interpolate_preprocess),
eval() (flood_forecast.transformer_xl.transformer_xl.StandardWordEmbedding
method), 169
eval() (flood_forecast.transformer_xl.transformer_xl.TransformerXL
flood_forecast.utils), 1
method), 177
evaluate_model() (in module
flood_forecast.evaluator), 3
extra_repr() (flood_forecast.custom.custom_opt.GaussianLoss
method), 47
extra_repr() (flood_forecast.custom.custom_opt.GaussianLoss
method), 54
extra_repr() (flood_forecast.custom.custom_opt.MAPELoss
method), 62
extra_repr() (flood_forecast.custom.custom_opt.RMSELoss
method), 39
extra_repr() (flood_forecast.custom.custom_opt.QuantileLoss
method), 62
extra_repr() (flood_forecast.custom.custom_opt.RMSELoss
method), 195
extra_repr() (flood_forecast.da_rnn.model.DARN
method), 195
extra_repr() (flood_forecast.da_rnn.modules.Decoder
method), 213
extra_repr() (flood_forecast.da_rnn.modules.Encoder
method), 205
extra_repr() (flood_forecast.da_rnn.modules.Decoder
method), 212
extra_repr() (flood_forecast.transformer_xl.dummy_torch.DummyTorchMode
method), 71
extra_repr() (flood_forecast.transformer_xl.lower_upper_config.AR
method), 86
extra_repr() (flood_forecast.transformer_xl.lower_upper_config.MetaEmbedding
method), 94
extra_repr() (flood_forecast.transformer_xl.lower_upper_config.MetaEmbedding
method), 94
extra_repr() (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFeedForward
method), 79
extra_repr() (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFeedForward
method), 79
extra_repr() (flood_forecast.transformer_xl.multi_head_base.MultiAttnHeadSimple
method), 103
extra_repr() (flood_forecast.transformer_xl.transformer_basic.CustomTransformerDecoder
method), 121
extra_repr() (flood_forecast.transformer_xl.transformer_basic.SimplePositionalEncoding
method), 129

```

```

float() (flood_forecast.transformer_xl.transformer_xl.DecoderBlock
    method), 154                                     flood_forecast.time_model(module), 15
float() (flood_forecast.transformer_xl.transformer_xl.MultiHeadAttention
    method), 139                                     MultiHeadAttention.trainer(module), 17
float() (flood_forecast.transformer_xl.transformer_xl.PositionalEmbedding
    method), 162                                     PositionalEmbedding.transformer_xl.dummy_torch
                                                    (module), 69
float() (flood_forecast.transformer_xl.transformer_xl.PositionwiseFeedForward
    method), 147                                     PositionwiseFeedForward.transformer_xl.lower_upper_config
                                                    (module), 77
float() (flood_forecast.transformer_xl.transformer_xl.StandardWordEmbedding
    method), 170                                     StandardWordEmbedding.transformer_xl.multi_head_base
                                                    (module), 101
float() (flood_forecast.transformer_xl.transformer_xl.TransformerXL
    method), 178                                     TransformerXL.transformer_xl.transformer_basic
                                                    (module), 111
flood_forecast(module), 1                           flood_forecast.transformer_xl.transformer_xl
                                                    (module), 137
flood_forecast.custom(module), 35                  flood_forecast.utils(module), 1
flood_forecast.custom.custom_opt(mod
    ule), 37                                     format_data() (in module
                                                    flood_forecast.preprocessing.preprocess_da_rnn),
                                                    27
flood_forecast.da_rnn(module), 185                format_dt() (in module
                                                    flood_forecast.preprocessing.closest_station),
                                                    23
flood_forecast.da_rnn.custom_types(mod
    ule), 191                                     forward() (flood_forecast.custom.custom_opt.GaussianLoss
                                                    method), 52
flood_forecast.da_rnn.model(module), 193           forward() (flood_forecast.custom.custom_opt.MAPELoss
                                                    method), 45
flood_forecast.da_rnn.modules(module), 203         forward() (flood_forecast.custom.custom_opt.QuantileLoss
                                                    method), 60
flood_forecast.da_rnn.train_da(module), 187        forward() (flood_forecast.custom.custom_opt.RMSELoss
                                                    method), 37
flood_forecast.da_rnn.utils(module), 189           forward() (flood_forecast.da_rnn.model.DARNN
                                                    method), 193
flood_forecast.evaluator(module), 3                 forward() (flood_forecast.da_rnn.modules.Decoder
                                                    method), 211
flood_forecast.gcp_integration(module), 183        forward() (flood_forecast.da_rnn.modules.Encoder
                                                    method), 203
flood_forecast.gcp_integration.basic_utils(mod
    ule), 185                                     forward() (flood_forecast.transformer_xl.dummy_torch.DummyTorchModule
                                                    method), 69
flood_forecast.long_train(module), 7               forward() (flood_forecast.transformer_xl.lower_upper_config.AR
                                                    method), 85
flood_forecast.model_dict_function(mod
    ule), 9                                      forward() (flood_forecast.transformer_xl.lower_upper_config.MetaEmbedding
                                                    method), 94
flood_forecast.pre_dict(module), 11                forward() (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFeedForward
                                                    method), 77
flood_forecast.preprocessing(module), 17           forward() (flood_forecast.transformer_xl.multi_head_base.MultiAttnHead
                                                    method), 103
flood_forecast.preprocessing.buil
    d_dataset(module), 21                                     forward() (flood_forecast.transformer_xl.transformer_basic.CustomTrainer
                                                    method), 119
flood_forecast.preprocessing.closest_stations(mod
    ule), 23                                     forward() (flood_forecast.transformer_xl.transformer_basic.SimplePositional
                                                    method), 127
flood_forecast.preprocessing.data_converter(mod
    ule), 25                                     forward() (flood_forecast.transformer_xl.transformer_basic.SimpleTransformer
                                                    method), 111
flood_forecast.preprocessing.interpolate(pr
    ocess, 19                                     forward() (flood_forecast.transformer_xl.transformer_xl.DecoderBlock
                                                    method), 152
flood_forecast.preprocessing.preprocess_darw
    and(mod
    ule), 27                                     forward() (flood_forecast.transformer_xl.transformer_xl.MultiHeadAttention
                                                    method), 154
flood_forecast.preprocessing.preprocess_fet
    rmat(mod
    ule), 29                                     forward() (flood_forecast.transformer_xl.transformer_xl.MultiHeadAttention
                                                    method), 154
flood_forecast.preprocessing.process_usgs(forward() (flood_forecast.transformer_xl.transformer_xl.MultiHeadAttention
    module), 31                                     method), 154
flood_forecast.preprocessing.pytorch_loader(forward() (flood_forecast.transformer_xl.transformer_xl.MultiHeadAttention
    module), 33                                     method), 154
flood_forecast.preprocessing.temporal_featur
    er(forward() (flood_forecast.transformer_xl.transformer_xl.MultiHeadAttention
    module), 35                                     method), 154
flood_forecast.pytorch_training(module),  forward() (flood_forecast.transformer_xl.transformer_xl.MultiHeadAttention
                                                    method), 154

```

<p><i>method), 137</i></p> <p><i>forward() (flood_forecast.transformer_xl.transformer_xl.PositionwiseFF method), 160</i></p> <p><i>forward() (flood_forecast.transformer_xl.transformer_xl.PositionwiseFF (in module flood_forecast.evaluator), 3 method), 145</i></p> <p><i>get_weather_data() (in module flood_forecast.preprocessing.closest_station), 23</i></p> <p><i>forward() (flood_forecast.transformer_xl.transformer_xl.StandardScaler (in module flood_forecast.preprocessing.closest_station), 23</i></p> <p><i>forward() (flood_forecast.transformer_xl.transformer_xl.TransformerXL (in module flood_forecast.preprocessing.temporal_feats), 35</i></p>	<p><i>185</i></p> <p><i>flood_forecast.preprocessing.process_usgs), 31</i></p> <p><i>(in module flood_forecast.evaluator), 3</i></p> <p><i>(in module flood_forecast.preprocessing.closest_station), 23</i></p> <p><i>(in module flood_forecast.preprocessing.temporal_feats), 35</i></p>
<p>G</p> <p><i>GaussianLoss (class in flood_forecast.custom.custom_opt), 52</i></p> <p><i>GaussianLoss.to() (in module flood_forecast.custom.custom_opt), 58, 59</i></p> <p><i>generate_decoded_predictions() (in module flood_forecast.evaluator), 4</i></p> <p><i>generate_prediction_samples() (in module flood_forecast.evaluator), 4</i></p> <p><i>generate_predictions() (in module flood_forecast.evaluator), 4</i></p> <p><i>generate_predictions_non_decoded() (in module flood_forecast.evaluator), 4</i></p> <p><i>generate_square_subsequent_mask() (in module flood_forecast.model_dict_function), 9</i></p> <p><i>generate_square_subsequent_mask() (in module flood_forecast.transformer_xl.transformer_basic), 135</i></p> <p><i>get_closest_gage() (in module flood_forecast.preprocessing.closest_station), 23</i></p> <p><i>get_day() (in module flood_forecast.preprocessing.temporal_feats), 35</i></p> <p><i>get_eco_netset() (in module flood_forecast.preprocessing.buil_dataset), 21</i></p> <p><i>get_from_start_date() (flood_forecast.preprocessing.pytorch_loaders.CSVTestLoader method), 103</i></p> <p><i>get_hour() (in module flood_forecast.preprocessing.temporal_feats), 35</i></p> <p><i>get_lr() (flood_forecast.custom.custom_opt.BertAdam method), 68</i></p> <p><i>get_model_r2_score() (in module flood_forecast.evaluator), 3</i></p> <p><i>get_month() (in module flood_forecast.preprocessing.temporal_feats), 35</i></p> <p><i>get_r2_value() (in module flood_forecast.evaluator), 3</i></p> <p><i>get_storage_client() (in module flood_forecast.gcp_integration.basic_utils),</i></p>	<p><i>greedy_decode() (in module flood_forecast.transformer_xl.transformer_basic), 135</i></p> <p>H</p> <p><i>half() (flood_forecast.custom.custom_opt.GaussianLoss method), 54</i></p> <p><i>half() (flood_forecast.custom.custom_opt.MAPELoss method), 47</i></p> <p><i>half() (flood_forecast.custom.custom_opt.QuantileLoss method), 62</i></p> <p><i>half() (flood_forecast.custom.custom_opt.RMSELoss method), 39</i></p> <p><i>half() (flood_forecast.da_rnn.model.DARNN method), 195</i></p> <p><i>half() (flood_forecast.da_rnn.modules.Decoder method), 213</i></p> <p><i>half() (flood_forecast.da_rnn.modules.Encoder method), 205</i></p> <p><i>half() (flood_forecast.transformer_xl.dummy_torch.DummyTorchModel method), 71</i></p> <p><i>half() (flood_forecast.transformer_xl.lower_upper_config.AR method), 87</i></p> <p><i>half() (flood_forecast.transformer_xl.lower_upper_config.MetaEmbedding method), 94</i></p> <p><i>half() (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFF method), 79</i></p> <p><i>half() (flood_forecast.transformer_xl.multi_head_base.MultiAttnHeadSinusoidal method), 103</i></p> <p><i>half() (flood_forecast.transformer_xl.transformer_basic.CustomTransformer method), 121</i></p> <p><i>half() (flood_forecast.transformer_xl.transformer_basic.SimplePositionwiseFF method), 129</i></p> <p><i>half() (flood_forecast.transformer_xl.transformer_basic.SimpleTransformer method), 113</i></p> <p><i>half() (flood_forecast.transformer_xl.transformer_xl.DecoderBlock method), 154</i></p> <p><i>half() (flood_forecast.transformer_xl.transformer_xl.MultiHeadAttention method), 139</i></p> <p><i>half() (flood_forecast.transformer_xl.transformer_xl.PositionalEmbedding method), 162</i></p> <p><i>half() (flood_forecast.transformer_xl.transformer_xl.PositionwiseFF method), 147</i></p>

half () (flood_forecast.transformer_xl.transformer_xl.StandardWordEmbedding (),
 method), 170
 half () (flood_forecast.transformer_xl.transformer_xl.TransformerXLmethod), 39
 method), 178
 haversine () (in module flood_forecast.preprocessing.closest_station),
 23
 |
 index () (flood_forecast.da_rnn.custom_types.DaRnnNet load_state_dict (),
 method), 192
 index () (flood_forecast.da_rnn.custom_types.TrainConfig method), 205
 method), 191
 index () (flood_forecast.da_rnn.custom_types.TrainData (flood_forecast.transformer_xl.dummy_torch.DummyTorchModel
 method), 71
 infer_on_torch_model () (in module load_state_dict (),
 flood_forecast.evaluator), 3
 init_hidden () (in module load_state_dict (),
 flood_forecast.da_rnn.modules), 203
 init_memory () (flood_forecast.transformer_xl.transformer_xl.TransformXLmethod), 94
 method), 176
 initial_layer () (in module load_state_dict (),
 flood_forecast.transformer_xl.lower_upper_config),
 77
 interpolate_missing_values () (in module load_state_dict (),
 flood_forecast.preprocessing.interpolate_preprocess),
 19
 inverse_scale () (flood_forecast.preprocessing.pytorch_loaders.CSVDataLoader
 method), 34
 inverse_scale () (flood_forecast.preprocessing.pytorch_loaders.TSNTDataLoader
 method), 34
 J
 join_data () (in module load_state_dict (),
 flood_forecast.preprocessing.buil_dataset),
 21
 L
 load_model () (flood_forecast.time_model.PyTorchForecast load_state_dict (),
 method), 15
 load_model () (flood_forecast.time_model.TimeSeriesModel load_state_dict (),
 method), 15
 load_state_dict () (flood_forecast.transformer_xl.DecoderBlock
 method), 154
 load_state_dict () (flood_forecast.transformer_xl.MultiHeadAttention
 method), 139
 load_state_dict () (flood_forecast.transformer_xl.PositionalEmbedding
 method), 162
 load_state_dict () (flood_forecast.transformer_xl.PositionwiseFF
 method), 147
 load_state_dict () (flood_forecast.transformer_xl.StandardWordEmbedding
 method), 170
 load_state_dict () (flood_forecast.transformer_xl.TransformerXL
 method), 178

```

loop_through()           (in      module   modules() (flood_forecast.transformer_xl.dummy_torch.DummyTorchMethod), 71
flood_forecast.long_train), 7
loss_func(flood_forecast.da_rnn.custom_types.TrainConfig.modules() (flood_forecast.transformer_xl.lower_upper_config.ARmethod), 87
attribute), 191
modules() (flood_forecast.transformer_xl.lower_upper_config.MetaEmbeddingMethod), 95
modules() (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFeedForwardMethod), 79
modules() (flood_forecast.transformer_xl.multi_head_base.MultiAttnHeadSimpleMethod), 104
modules() (flood_forecast.transformer_xl.transformer_basic.CustomTransformerBasicMethod), 122
modules() (flood_forecast.transformer_xl.transformer_basic.SimplePositionwiseFFNMethod), 129
modules() (flood_forecast.transformer_xl.transformer_basic.SimpleTransformerBasicMethod), 114
modules() (flood_forecast.transformer_xl.transformer_xl.DecoderBlockMethod), 155
make_data_load() (flood_forecast.time_model.PyTorchForecastModel.modules() (flood_forecast.transformer_xl.transformer_xl.MultiHeadAttentionMethod), 140
method), 15
make_data_load() (flood_forecast.time_model.TimeSeriesModel.modules() (flood_forecast.transformer_xl.transformer_xl.MultiHeadAttentionMethod), 140
method), 15
make_gage_data_csv() (in      module   modules() (flood_forecast.transformer_xl.transformer_xl.PositionalEmbeddingMethod), 162
flood_forecast.preprocessing.preprocess_metadata), 29
modules() (flood_forecast.transformer_xl.transformer_xl.PositionwiseFeedForwardMethod), 147
make_temporal_features() (in      module   modules() (flood_forecast.transformer_xl.transformer_xl.StandardWordEmbeddingMethod), 170
flood_forecast.preprocessing.temporal_feats), 35
modules() (flood_forecast.transformer_xl.transformer_xl.TransformerXLMethod), 178
make_usgs_data() (in      module   MultiAttnHeadSimple (class      in
flood_forecast.preprocessing.process_usgs), 31
flood_forecast.transformer_xl.multi_head_base), 101
MAPELoss (class in flood_forecast.custom.custom_opt), 45
MAPELoss.to() (in      module   MultiAttnHeadSimple.to() (in      module
flood_forecast.custom.custom_opt), 51
flood_forecast.transformer_xl.multi_head_base), 107, 108
MetaEmbedding (class      in
flood_forecast.transformer_xl.lower_upper_config), 92
MultiHeadAttention (class      in
flood_forecast.transformer_xl.transformer_xl), 137
MetaEmbedding.to() (in      module   MultiHeadAttention.to() (in      module
flood_forecast.transformer_xl.lower_upper_config), 99
flood_forecast.transformer_xl.transformer_xl), 143, 144
metric_dict () (in module flood_forecast.evaluator), 3
modules() (flood_forecast.custom.custom_opt.GaussianLoss
method), 55
modules() (flood_forecast.custom.custom_opt.MAPELoss.named_buffers() (flood_forecast.custom.custom_opt.GaussianLoss
method), 47
method), 55
modules() (flood_forecast.custom.custom_opt.QuantileLoss.named_buffers() (flood_forecast.custom.custom_opt.MAPELoss
method), 62
method), 47
modules() (flood_forecast.custom.custom_opt.RMSELoss.named_buffers() (flood_forecast.custom.custom_opt.QuantileLoss
method), 39
method), 62
modules() (flood_forecast.da_rnn.model.DARNN.named_buffers() (flood_forecast.custom.custom_opt.RMSELoss
method), 195
method), 40
modules() (flood_forecast.da_rnn.modules.Decoder.named_buffers() (flood_forecast.da_rnn.model.DARNN
method), 213
method), 196
modules() (flood_forecast.da_rnn.modules.Encoder.named_buffers() (flood_forecast.da_rnn.modules.Decoder
method), 205
method), 213

```

named_buffers () (flood_forecast.da_rnn.modules.Encoder) (flood_forecast.transformer_xl.transformer_basic.SimplifiedForwardingForecastMethod) (method), 206
 named_children () (flood_forecast.transformer_xl.dummy_torch.DummyTorchMethod) (method), 122
 named_buffers () (flood_forecast.transformer_xl.lower_upper_config.AR) (flood_forecast.transformer_xl.transformer_basic.SimplifiedForwardingForecastMethod) (method), 72
 named_buffers () (flood_forecast.transformer_xl.lower_upper_config.AR) (flood_forecast.transformer_xl.transformer_basic.SimplifiedForwardingForecastMethod) (method), 130
 named_buffers () (flood_forecast.transformer_xl.lower_upper_config.AR) (flood_forecast.transformer_xl.transformer_basic.SimplifiedForwardingForecastMethod) (method), 87
 named_buffers () (flood_forecast.transformer_xl.lower_upper_config.AR) (flood_forecast.transformer_xl.transformer_basic.SimplifiedForwardingForecastMethod) (method), 114
 named_buffers () (flood_forecast.transformer_xl.lower_upper_config.MetaEmbeddingForecastMethod) (flood_forecast.transformer_xl.transformer_xl.Decoder) (method), 95
 named_buffers () (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFeedForwardTransformerXL) (flood_forecast.transformer_xl.MultiplexedDecoder) (method), 155
 named_buffers () (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFeedForwardTransformerXL) (flood_forecast.transformer_xl.MultiplexedDecoder) (method), 80
 named_buffers () (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFeedForwardTransformerXL) (flood_forecast.transformer_xl.PositionwiseFeedForward) (method), 140
 named_buffers () (flood_forecast.transformer_xl.multi_head_base.MultiHeadAttention) (flood_forecast.transformer_xl.transformer_xl.PositionwiseFeedForward) (method), 104
 named_buffers () (flood_forecast.transformer_xl.transformer_basic.CustomTanh) (flood_forecast.transformer_xl.transformer_xl.PositionwiseFeedForward) (method), 122
 named_buffers () (flood_forecast.transformer_xl.transformer_basic.SimplePositionForwarding) (flood_forecast.transformer_xl.transformer_xl.StandardForwarding) (method), 148
 named_buffers () (flood_forecast.transformer_xl.transformer_basic.SimplePositionForwarding) (flood_forecast.transformer_xl.transformer_xl.StandardForwarding) (method), 130
 named_buffers () (flood_forecast.transformer_xl.transformer_basic.SimplePositionForwarding) (flood_forecast.transformer_xl.transformer_xl.Transformers) (method), 171
 named_buffers () (flood_forecast.transformer_xl.transformer_basic.SimplePositionForwarding) (flood_forecast.transformer_xl.transformer_xl.Transformers) (method), 114
 named_buffers () (flood_forecast.transformer_xl.transformer_xl.DecoderBlock) (flood_forecast.custom.custom_opt.GaussianLoss) (method), 155
 named_buffers () (flood_forecast.transformer_xl.transformer_xl.MultiHeadAttention) (flood_forecast.custom.custom_opt.MAPELoss) (method), 140
 named_buffers () (flood_forecast.transformer_xl.transformer_xl.PositionwiseFeedForward) (flood_forecast.custom.custom_opt.QuantileLoss) (method), 163
 named_buffers () (flood_forecast.transformer_xl.transformer_xl.PositionwiseFeedForward) (flood_forecast.custom.custom_opt.RMSELoss) (method), 147
 named_buffers () (flood_forecast.transformer_xl.transformer_xl.StandardWordEmbedding) (flood_forecast.da_rnn.model.DARNN) (method), 170
 named_buffers () (flood_forecast.transformer_xl.transformer_xl.TransformerXL) (flood_forecast.da_rnn.modules.Decoder) (method), 214
 named_children () (flood_forecast.custom.custom_opt.GaussianLoss) (flood_forecast.da_rnn.modules.Encoder) (method), 55
 named_children () (flood_forecast.custom.custom_opt.MAPELoss) (flood_forecast.transformer_xl.dummy_torch.DummyTorchMethod) (method), 72
 named_children () (flood_forecast.custom.custom_opt.QuantileLoss) (flood_forecast.transformer_xl.lower_upper_config.AR) (method), 63
 named_children () (flood_forecast.custom.custom_opt.RMSELoss) (flood_forecast.transformer_xl.lower_upper_config.AR) (method), 95
 named_children () (flood_forecast.da_rnn.model.DARNN) (flood_forecast.transformer_xl.lower_upper_config.AR) (method), 196
 named_children () (flood_forecast.da_rnn.modules.Decoder) (flood_forecast.transformer_xl.multi_head_base.MultiHeadAttention) (method), 80
 named_children () (flood_forecast.da_rnn.modules.Decoder) (flood_forecast.transformer_xl.multi_head_base.MultiHeadAttention) (method), 214
 named_children () (flood_forecast.da_rnn.modules.Encoder) (flood_forecast.transformer_xl.transformer_basic.CustomTanh) (method), 206
 named_children () (flood_forecast.transformer_xl.dummy_torch.DummyTorchMethod) (flood_forecast.transformer_xl.transformer_basic.SimplePositionForwarding) (method), 72
 named_children () (flood_forecast.transformer_xl.lower_upper_config.AR) (flood_forecast.transformer_xl.transformer_xl.Decoder) (method), 130
 named_children () (flood_forecast.transformer_xl.lower_upper_config.AR) (flood_forecast.transformer_xl.transformer_xl.Transformers) (method), 88
 named_children () (flood_forecast.transformer_xl.lower_upper_config.AR) (flood_forecast.transformer_xl.transformer_xl.Transformers) (method), 95
 named_children () (flood_forecast.transformer_xl.lower_upper_config.MetaEmbedding) (flood_forecast.transformer_xl.transformer_xl.Decoder) (method), 140
 named_children () (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFeedForward) (flood_forecast.transformer_xl.transformer_xl.MultiplexedDecoder) (method), 80
 named_children () (flood_forecast.transformer_xl.multi_head_base.MultiHeadAttention) (flood_forecast.transformer_xl.transformer_xl.PositionwiseFeedForward) (method), 104

```

named_modules() (flood_forecast.transformer_xl.transformer_xl.PositionwiseFF)
    method), 148
named_modules() (flood_forecast.transformer_xl.transformer_xl.StandardWordEmbedding
    method), 171
named_modules() (flood_forecast.transformer_xl.transformer_xl.TransformerXL.transformer_xl.PositionalEmbedding
    method), 179
named_parameters()
    (flood_forecast.custom.custom_opt.GaussianLoss
    method), 56
named_parameters()
    (flood_forecast.custom.custom_opt.MAPELoss
    method), 48
named_parameters()
    (flood_forecast.custom.custom_opt.QuantileLoss
    method), 63
named_parameters()
    (flood_forecast.custom.custom_opt.RMSELoss
    method), 41
named_parameters()
    (flood_forecast.da_rnn.model.DARNN
    method), 196
named_parameters()
    (flood_forecast.da_rnn.modules.Decoder
    method), 214
named_parameters()
    (flood_forecast.da_rnn.modules.Encoder
    method), 206
named_parameters()
    (flood_forecast.transformer_xl.dummy_torch.DummyTorchModel
    method), 72
named_parameters()
    (flood_forecast.transformer_xl.lower_upper_config.AR
    method), 88
named_parameters()
    (flood_forecast.transformer_xl.lower_upper_config.MetaEmbedding
    method), 96
named_parameters()
    (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFeedForward
    method), 81
named_parameters()
    (flood_forecast.transformer_xl.multi_head_base.MultiAttnHeadSimple
    method), 105
named_parameters()
    (flood_forecast.transformer_xl.transformer_basic.CustomTransformerDecoder
    method), 123
named_parameters()
    (flood_forecast.transformer_xl.transformer_basic.SimplePositionalEncoding
    method), 130
named_parameters()
    (flood_forecast.transformer_xl.transformer_basic.SimpleTransformerEncoder
    method), 115
named_parameters()
    (flood_forecast.transformer_xl.DecoderBlock
    method), 156
named_parameters()
    (flood_forecast.transformer_xl.MultiHeadAttention
    method), 141

```

P

```

parameters() (flood_forecast.custom.custom_opt.GaussianLoss
    method), 56
parameters() (flood_forecast.custom.custom_opt.MAPELoss
    method), 48
parameters() (flood_forecast.custom.custom_opt.QuantileLoss
    method), 64
parameters() (flood_forecast.custom.custom_opt.RMSELoss
    method), 41
parameters() (flood_forecast.da_rnn.model.DARNN
    method), 197
parameters() (flood_forecast.da_rnn.modules.Decoder
    method), 214
parameters() (flood_forecast.da_rnn.modules.Encoder
    method), 207
parameters() (flood_forecast.transformer_xl.dummy_torch.DummyTorchModel
    method), 73
parameters() (flood_forecast.transformer_xl.lower_upper_config.AR
    method), 88
parameters() (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFeedForward
    method), 96
parameters() (flood_forecast.transformer_xl.lower_upper_config.MetaEmbedding
    method), 105
parameters() (flood_forecast.transformer_xl.multi_head_base.MultiAttnHeadSimple
    method), 115
parameters() (flood_forecast.transformer_xl.transformer_basic.SimplePositionalEncoding
    method), 115
parameters() (flood_forecast.transformer_xl.transformer_basic.SimpleTransformerEncoder
    method), 156
parameters() (flood_forecast.transformer_xl.MultiHeadAttention
    method), 141

```

parameters() (flood_forecast.transformer_xl.transformer_xl.~~PositionalEmbedding~~.hook(),
 method), 164
 parameters() (flood_forecast.transformer_xl.transformer_xl.~~PositionalEmbedding~~.hook(),
 method), 149
 parameters() (flood_forecast.transformer_xl.transformer_xl.StanfordWfdEndeavourgm.custom_opt.QuantileLoss
 method), 172
 parameters() (flood_forecast.transformer_xl.transformer_xl.~~TransformerXLward~~.hook(),
 method), 180
 plot_r2() (in module flood_forecast.evaluator), 3
 PositionalEmbedding (class in register_backward_hook()
 flood_forecast.transformer_xl.transformer_xl), 160
 PositionalEmbedding.to() (in module register_backward_hook()
 flood_forecast.transformer_xl.transformer_xl), 166
 PositionwiseFeedForward (class in register_backward_hook()
 flood_forecast.transformer_xl.lower_upper_config), 77
 PositionwiseFeedForward.to() (in module register_backward_hook()
 flood_forecast.transformer_xl.lower_upper_config), 83
 PositionwiseFF (class in register_backward_hook()
 flood_forecast.transformer_xl.transformer_xl), 145
 PositionwiseFF.to() (in module register_backward_hook()
 flood_forecast.transformer_xl.transformer_xl), 151
 predict() (in module register_backward_hook()
 flood_forecast.da_rnn.train_da), 187
 prep_train_data() (in module register_backward_hook()
 flood_forecast.da_rnn.train_da), 187
 process_asos_csv() (in module register_backward_hook()
 flood_forecast.preprocessing.closest_station), 23
 process_asos_data() (in module register_backward_hook()
 flood_forecast.preprocessing.closest_station), 23
 process_intermediate_csv() (in module register_backward_hook()
 flood_forecast.preprocessing.process_usgs), 31
 process_response_text() (in module register_backward_hook()
 flood_forecast.preprocessing.process_usgs), 31
 PyTorchForecast (class in register_backward_hook()
 flood_forecast.time_model), 15
Q
 QuantileLoss (class in register_backward_hook()
 flood_forecast.custom.custom_opt), 60
 QuantileLoss.to() (in module register_backward_hook()
 flood_forecast.custom.custom_opt), 66
R
 register_backward_hook() (flood_forecast.custom.custom_opt.GaussianLoss
 method), 56
 (flood_forecast.transformer_xl.transformer_xl.PositionwiseFF
 method), 149

```
register_backward_hook ()                                register_buffer ()  
    (flood_forecast.transformer_xl.transformer_xl.StandardWordEmbeddingForecast.transformer_xl.transformer_xl.MultiHeadAttention  
     method), 172  
register_backward_hook ()                                register_buffer ()  
    (flood_forecast.transformer_xl.transformer_xl.TransformerXL.flood_forecast.transformer_xl.transformer_xl.PositionalEmbedding  
     method), 180  
register_buffer ()                                     register_buffer ()  
    (flood_forecast.custom.custom_opt.GaussianLoss  
     method), 57  
register_buffer ()                                     register_buffer ()  
    (flood_forecast.custom.custom_opt.MAPELoss  
     method), 49  
register_buffer ()                                     register_buffer ()  
    (flood_forecast.custom.custom_opt.QuantileLoss  
     method), 64  
register_buffer ()                                     register_forward_hook ()  
    (flood_forecast.custom.custom_opt.RMSELoss  
     method), 42  
register_buffer ()                                     register_forward_hook ()  
    (flood_forecast.da_rnn.model.DARNN  
     method), 198  
register_buffer ()                                     register_forward_hook ()  
    (flood_forecast.da_rnn.modules.Decoder  
     method), 215  
register_buffer ()                                     register_forward_hook ()  
    (flood_forecast.da_rnn.modules.Encoder  
     method), 208  
register_buffer ()                                     register_forward_hook ()  
    (flood_forecast.transformer_xl.dummy_torch.DummyTorchModel.flood_forecast.da_rnn.model.DARNN  
     method), 198  
register_buffer ()                                     register_forward_hook ()  
    (flood_forecast.transformer_xl.lower_upper_config.AR  
     method), 89  
register_buffer ()                                     register_forward_hook ()  
    (flood_forecast.transformer_xl.lower_upper_config.MetaEmbeddingForecast.da_rnn.modules.Encoder  
     method), 97  
register_buffer ()                                     register_forward_hook ()  
    (flood_forecast.transformer_xl.lower_upper_config.PositionEmbeddingTransformer_xl.dummy_torch.DummyTorchModel  
     method), 82  
register_buffer ()                                     register_forward_hook ()  
    (flood_forecast.transformer_xl.multi_head_base.MultiAttnHead.flood_forecast.transformer_xl.lower_upper_config.AR  
     method), 106  
register_buffer ()                                     register_forward_hook ()  
    (flood_forecast.transformer_xl.transformer_basic.CustomTransformer_xl.lower_upper_config.MetaEmbedding  
     method), 124  
register_buffer ()                                     register_forward_hook ()  
    (flood_forecast.transformer_xl.transformer_basic.SimplePositionEmbeddingTransformer_xl.lower_upper_config.PositionwiseFF  
     method), 132  
register_buffer ()                                     register_forward_hook ()  
    (flood_forecast.transformer_xl.transformer_basic.SimpleTransformer_xl.multi_head_base.MultiAttnHeadSimpleTransformer  
     method), 116  
register_buffer ()                                     register_forward_hook ()  
    (flood_forecast.transformer_xl.transformer_xl.DecoderBlock.flood_forecast.transformer_xl.transformer_basic.CustomTransfo  
     method), 157
```

```

register_forward_hook()           register_forward_pre_hook()
(flood_forecast.transformer_xl.transformer_basic.SimplePositionwiseFF
method), 132                     method), 82

register_forward_hook()           register_forward_pre_hook()
(flood_forecast.transformer_xl.transformer_basic.SimpleTrafficForecastTransformerXL
method), 116                      method), 106

register_forward_hook()           register_forward_pre_hook()
(flood_forecast.transformer_xl.transformer_xl.DecoderBlock
method), 157                      method), 124

register_forward_hook()           register_forward_pre_hook()
(flood_forecast.transformer_xl.transformer_xl.MultiHeadAttention
method), 142                      method), 132

register_forward_hook()           register_forward_pre_hook()
(flood_forecast.transformer_xl.transformer_xl.PositionalEncoder
method), 165                      method), 117

register_forward_hook()           register_forward_pre_hook()
(flood_forecast.transformer_xl.transformer_xl.PositionwiseFF
method), 150                      method), 158

register_forward_hook()           register_forward_pre_hook()
(flood_forecast.transformer_xl.transformer_xl.StandardWordEmbedding
method), 173                      method), 142

register_forward_hook()           register_forward_pre_hook()
(flood_forecast.transformer_xl.transformer_xl.TransformerXL
method), 181                      method), 165

register_forward_pre_hook()       register_forward_pre_hook()
(flood_forecast.custom.custom_opt.GaussianLoss
method), 57                         method), 150

register_forward_pre_hook()       register_forward_pre_hook()
(flood_forecast.custom.custom_opt.MAPELoss
method), 50                         method), 173

register_forward_pre_hook()       register_forward_pre_hook()
(flood_forecast.custom.custom_opt.QuantileLoss
method), 65                         method), 181

register_forward_pre_hook()       register_parameter()
(flood_forecast.custom.custom_opt.RMSELoss
method), 42                         method), 58

register_forward_pre_hook()       register_parameter()
(flood_forecast.da_rnn.model.DARNN
method), 198                         method), 50

register_forward_pre_hook()       register_parameter()
(flood_forecast.da_rnn.modules.Decoder
method), 216                         method), 65

register_forward_pre_hook()       register_parameter()
(flood_forecast.da_rnn.modules.Encoder
method), 208                         method), 43

register_forward_pre_hook()       register_parameter()
(flood_forecast.transformer_xl.dummy_torch.DummyTorchModel
method), 74                           method), 198

register_forward_pre_hook()       register_parameter()
(flood_forecast.transformer_xl.lower_upper_config.AR
method), 90                           method), 216

register_forward_pre_hook()       register_parameter()
(flood_forecast.transformer_xl.lower_upper_config.MetaEncoder
method), 97                           method), 208

```

```

register_parameter()                               requires_grad_() (flood_forecast.da_rnn.modules.Encoder
    (flood_forecast.transformer_xl.dummy_torch.DummyTorchModule), 209
    method), 74                                requires_grad_() (flood_forecast.transformer_xl.dummy_torch.Dumm
register_parameter()                               requires_grad_() (flood_forecast.transformer_xl.lower_upper_config
    (flood_forecast.transformer_xl.lower_upper_config.ARquires_grad_() (flood_forecast.transformer_xl.lower_upper_config
    method), 90                                method), 90
register_parameter()                               requires_grad_() (flood_forecast.transformer_xl.lower_upper_config
    (flood_forecast.transformer_xl.lower_upper_config.MetaEmbedding, 98
    method), 98                                requires_grad_() (flood_forecast.transformer_xl.lower_upper_config
register_parameter()                               requires_grad_() (flood_forecast.transformer_xl.lower_upper_config
    (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFeedForward), 107
    method), 83                                method), 83
register_parameter()                               requires_grad_() (flood_forecast.transformer_xl.multi_head_base.M
    (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFeedForwod_forecast.transformer_xl.multi_head_base.M
    method), 107                                method), 107
register_parameter()                               requires_grad_() (flood_forecast.transformer_xl.transformer_basic.C
    (flood_forecast.transformer_xl.multi_head_base.MultiAttnHeadSimple25
    method), 107                                requires_grad_() (flood_forecast.transformer_xl.transformer_basic.S
register_parameter()                               requires_grad_() (flood_forecast.transformer_xl.transformer_basic.C
    (flood_forecast.transformer_xl.transformer_basic.CustomTransformerDecoder), 117
    method), 125                                method), 133
register_parameter()                               requires_grad_() (flood_forecast.transformer_xl.transformer_xl.Dec
    (flood_forecast.transformer_xl.transformer_basic.SimplePositionalEncoding
    method), 132                                requires_grad_() (flood_forecast.transformer_xl.transformer_xl.Mul
register_parameter()                               requires_grad_() (flood_forecast.transformer_xl.transformer_xl.Dec
    (flood_forecast.transformer_xl.transformer_basic.SimpleTransformer)
    method), 117                                method), 143
register_parameter()                               requires_grad_() (flood_forecast.transformer_xl.transformer_xl.Posi
    (flood_forecast.transformer_xl.transformer_xl.PositionalEmbedding)
    method), 150                                method), 150
register_parameter()                               requires_grad_() (flood_forecast.transformer_xl.transformer_xl.Stan
    (flood_forecast.transformer_xl.decoderBlock)
    method), 158                                method), 173
register_parameter()                               requires_grad_() (flood_forecast.transformer_xl.transformer_xl.Trans
    (flood_forecast.transformer_xl.MultiHeadAttention)
    method), 143                                method), 181
register_parameter()                               reset_length() (flood_forecast.transformer_xl.transformer_xl.Trans
    (flood_forecast.transformer_xl.PositionalEmbedding), 176
    method), 166                                RMSELoss (class in flood_forecast.custom.custom_opt),
register_parameter()                               37
    (flood_forecast.transformer_xl.PositionalEmbedding)
    method), 150                                module
register_parameter()                               SaveModelF.to()          (in           module
    (flood_forecast.transformer_xl.StandardWordEmbedding)
    method), 173                                flood_forecast.custom.custom_opt), 43
register_parameter()                               save_model() (flood_forecast.time_model.PyTorchForecast
    (flood_forecast.transformer_xl.TransformerXL)
    method), 181                                method), 15
register_parameter()                               SaveModel.checkpoint()          (flood_forecast.time_model.TimeSeriesModel
    (flood_forecast.custom.custom_opt.GaussianLoss)
    method), 58                                method), 15
register_parameter()                               EarlyStopper.method,
    (flood_forecast.custom.custom_opt.MAPELoss)
    method), 50                                save_or_show_plot()          (in           module
register_parameter()                               flood_forecast.utils.EarlyStopper.method,
    (flood_forecast.custom.custom_opt.QuantileLoss)
    method), 65                                setup_log() (in module flood_forecast.da_rnn.utils),
register_parameter()                               share_memory() (flood_forecast.custom.custom_opt.GaussianLoss
    (flood_forecast.da_rnn.model.DARNN)
    method), 199                                method), 58
register_parameter()                               share_memory() (flood_forecast.custom.custom_opt.MAPELoss
    (flood_forecast.da_rnn.modules.Decoder)
    method), 216                                method), 50

```

```

share_memory() (flood_forecast.custom.custom_opt.QuantileLoss) 19
    method), 66                                     StandardWordEmbedding      (class      in
share_memory() (flood_forecast.custom.custom_opt.RMSELoss)  flood_forecast.transformer_xl.transformer_xl),
    method), 43                                     168
share_memory() (flood_forecast.da_rnn.model.DARNN) StandardWordEmbedding.to()      (in      module
    method), 199                                     flood_forecast.transformer_xl.transformer_xl),
share_memory() (flood_forecast.da_rnn.modules.Decoder) 174
    method), 216                                     state_dict() (flood_forecast.custom.custom_opt.BertAdam
share_memory() (flood_forecast.da_rnn.modules.Encoder)  method), 68
    method), 209                                     state_dict() (flood_forecast.custom.custom_opt.GaussianLoss
share_memory() (flood_forecast.transformer_xl.dummy_torch.DummyTorch) 50
    method), 75                                     state_dict() (flood_forecast.custom.custom_opt.MAPELoss
share_memory() (flood_forecast.transformer_xl.lower_upper_configAR) 51
    method), 90                                     state_dict() (flood_forecast.custom.custom_opt.QuantileLoss
share_memory() (flood_forecast.transformer_xl.lower_upper_configM) 56
    method), 98                                     state_dict() (flood_forecast.custom.custom_opt.RMSELoss
share_memory() (flood_forecast.transformer_xl.lower_upper_configR) 43
    method), 83                                     state_dict() (flood_forecast.da_rnn.model.DARNN
share_memory() (flood_forecast.transformer_xl.multi_head_base.M) 50
    method), 107                                     state_dict() (flood_forecast.da_rnn.modules.Decoder
share_memory() (flood_forecast.transformer_xl.transformer_basinHead) 117
    method), 125                                     state_dict() (flood_forecast.da_rnn.modules.Encoder
share_memory() (flood_forecast.transformer_xl.transformer_basinP) 209
    method), 133                                     state_dict() (flood_forecast.transformer_xl.dummy_torch.DummyTorch
share_memory() (flood_forecast.transformer_xl.transformer_basinT) 75
    method), 117                                     state_dict() (flood_forecast.transformer_xl.lower_upper_config.AR
share_memory() (flood_forecast.transformer_xl.transformer_xl.D) 158
    method), 158                                     state_dict() (flood_forecast.transformer_xl.lower_upper_config.Meta
share_memory() (flood_forecast.transformer_xl.transformer_xl.M) 143
    method), 143                                     state_dict() (flood_forecast.transformer_xl.lower_upper_config.Positional
share_memory() (flood_forecast.transformer_xl.transformer_xl.P) 166
    method), 166                                     state_dict() (flood_forecast.transformer_xl.multi_head_base.MultiAttention
share_memory() (flood_forecast.transformer_xl.transformer_xl.P) 151
    method), 151                                     state_dict() (flood_forecast.transformer_xl.transformer_basic.Custom
share_memory() (flood_forecast.transformer_xl.StandardW) 174
    method), 174                                     state_dict() (flood_forecast.transformer_xl.transformer_basic.Simple
share_memory() (flood_forecast.transformer_xl.TransformerXL) 182
    method), 182                                     state_dict() (flood_forecast.transformer_xl.transformer_basic.Simple
SimplePositionalEncoding (class      in      method), 117
    flood_forecast.transformer_xl.transformer_basic).state_dict() (flood_forecast.transformer_xl.transformer_xl.DecoderB
    127                                     158
SimplePositionalEncoding.to() (in module state_dict() (flood_forecast.transformer_xl.transformer_xl.MultiHead
    flood_forecast.transformer_xl.transformer_basic),      method), 143
    133                                     state_dict() (flood_forecast.transformer_xl.transformer_xl.Positional
SimpleTransformer (class      in      method), 166
    flood_forecast.transformer_xl.transformer_basic).state_dict() (flood_forecast.transformer_xl.transformer_xl.Positionw
    111                                     151
SimpleTransformer.to() (in      module state_dict() (flood_forecast.transformer_xl.StandardW
    flood_forecast.transformer_xl.transformer_basic),      method), 174
    118                                     state_dict() (flood_forecast.transformer_xl.transformer_xl.Transform
split_on_letter() (in      module state_dict() (flood_forecast.long_train), 7
    method), 182                                     step() (flood_forecast.custom.custom_opt.BertAdam
split_on_na_chunks() (in      module state_dict() (flood_forecast.preprocessing.interpolate_preprocess), 68
    method), 68                                     team_baseline() (in      module
    flood_forecast.preprocessing.interpolate_preprocess)

```

`flood_forecast.evaluator), 3`
T
`T (flood_forecast.da_rnn.custom_types.TrainConfig attribute), 191`
`T_destination (flood_forecast.custom.custom_opt.GaussianLoss method), 43`
`attribute), 52`
`T_destination (flood_forecast.custom.custom_opt.MAPELoss attribute), 199`
`attribute), 45`
`T_destination (flood_forecast.custom.custom_opt.QuantileLoss attribute), 217`
`attribute), 60`
`T_destination (flood_forecast.custom.custom_opt.RMSELoss attribute), 209`
`attribute), 37`
`T_destination (flood_forecast.da_rnn.model.DARNN attribute), 193`
`T_destination (flood_forecast.da_rnn.modules.Decoder attribute), 211`
`T_destination (flood_forecast.da_rnn.modules.Encoder attribute), 203`
`T_destination (flood_forecast.transformer_xl.dummy_torch.DummyTorchModel attribute), 69`
`T_destination (flood_forecast.transformer_xl.lower_upper_config.AR attribute), 107`
`attribute), 85`
`T_destination (flood_forecast.transformer_xl.lower_upper_config.MetaEmbedding attribute), 92`
`T_destination (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFeedForward attribute), 77`
`T_destination (flood_forecast.transformer_xl.multi_head_base.MultiAttnHeadSimple attribute), 101`
`T_destination (flood_forecast.transformer_xl.transformer_basic.CustomTransformer attribute), 119`
`T_destination (flood_forecast.transformer_xl.transformer_basic.SimplePositionalEncoding attribute), 127`
`T_destination (flood_forecast.transformer_xl.transformer_basic.SimpleTransformer attribute), 111`
`T_destination (flood_forecast.transformer_xl.decoder_block.DecoderBlock attribute), 153`
`T_destination (flood_forecast.transformer_xl.transformer_decoder.CUSTOMTransformerDecoder attribute), 119`
`T_destination (flood_forecast.transformer_xl.transformer_decoder.MultiHeadAttention attribute), 143`
`T_destination (flood_forecast.transformer_xl.transformer_decoder.PositionalEncoding attribute), 127`
`T_destination (flood_forecast.transformer_xl.transformer_decoder.PositionalEmbedding attribute), 166`
`T_destination (flood_forecast.transformer_xl.transformer_decoder.PositionwiseFF attribute), 151`
`T_destination (flood_forecast.transformer_xl.transformer_decoder.StandardWordEmbedding attribute), 153`
`T_destination (flood_forecast.transformer_xl.transformer_decoder.MultiHeadAttention attribute), 174`
`T_destination (flood_forecast.transformer_xl.transformer_decoder.TransformerXL attribute), 137`
`T_destination (flood_forecast.transformer_xl.PositionEmbedding attribute), 182`
`attribute), 160`
`T_destination (flood_forecast.transformer_xl.PositionwiseFF attribute), 145`
`attribute), 145`
`T_destination (flood_forecast.transformer_xl.StandardWordEmbedding attribute), 168`
`attribute), 168`
`T_destination (flood_forecast.transformer_xl.TransformerXL attribute), 176`
`targs (flood_forecast.da_rnn.custom_types.TrainData attribute), 191`
`TimeSeriesModel (class in flood_forecast.time_model), 15`
`to () (flood_forecast.custom.custom_opt.GaussianLoss method), 58`
`to () (flood_forecast.custom.custom_opt.MAPELoss method), 51`
`to () (flood_forecast.custom.custom_opt.QuantileLoss method), 66`
`to () (flood_forecast.custom.custom_opt.RMSELoss method), 44`
`to () (flood_forecast.da_rnn.model.DARNN method), 200`
`to () (flood_forecast.da_rnn.modules.Decoder method), 218`
`to () (flood_forecast.transformer_xl.dummy_torch.DummyTorchModel method), 75`
`to () (flood_forecast.transformer_xl.lower_upper_config.AR method), 91`
`to () (flood_forecast.transformer_xl.lower_upper_config.MetaEmbedding method), 98`
`to () (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFeedForward method), 83`
`to () (flood_forecast.transformer_xl.multi_head_base.MultiAttnHeadSimple method), 109`
`to () (flood_forecast.transformer_xl.transformer_basic.CustomTransformer method), 125`
`to () (flood_forecast.transformer_xl.transformer_basic.SimplePositionalEncoding method), 133`
`to () (flood_forecast.transformer_xl.transformer_decoder.CUSTOMTransformerDecoder method), 150`
`to () (flood_forecast.transformer_xl.transformer_decoder.MultiHeadAttention method), 118`
`to () (flood_forecast.transformer_xl.transformer_decoder.PositionalEncoding method), 166`
`to () (flood_forecast.transformer_xl.transformer_decoder.PositionalEmbedding method), 151`
`to () (flood_forecast.transformer_xl.transformer_decoder.StandardWordEmbedding method), 174`
`to () (flood_forecast.transformer_xl.TransformerXL method), 137`
`torch_single_train() (in module flood_forecast.pytorch_training), 13`
`train() (flood_forecast.custom.custom_opt.GaussianLoss method), 50`
`train() (flood_forecast.custom.custom_opt.MAPELoss method), 51`
`train() (flood_forecast.custom.custom_opt.QuantileLoss method), 67`
`train() (flood_forecast.custom.custom_opt.RMSELoss method), 44`
`train() (flood_forecast.da_rnn.model.DARNN method), 200`
`train() (flood_forecast.da_rnn.modules.Decoder method), 218`

train() (flood_forecast.da_rnn.modules.Encoder type () (flood_forecast.custom.custom_opt.QuantileLoss method), 210
 train() (flood_forecast.transformer_xl.dummy_torch.DummyTorchModel_forecast.custom.custom_opt.RMSELoss method), 44
 train() (flood_forecast.transformer_xl.lower_upper_config.AR () (flood_forecast.da_rnn.model.DARNN method), 200
 train() (flood_forecast.transformer_xl.lower_upper_config.MetaEmbedding (flood_forecast.da_rnn.modules.Decoder method), 218
 train() (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFF (flood_forecast.da_rnn.modules.Encoder method), 210
 train() (flood_forecast.transformer_xl.multi_head_base.MultiAttnHeadSimple (flood_forecast.transformer_xl.dummy_torch.DummyTorchModel method), 76
 train() (flood_forecast.transformer_xl.transformer_basic.CustomTransformerXL.lower_upper_config.AR method), 92
 train() (flood_forecast.transformer_xl.transformer_basic.SimplePositionwiseFF (flood_forecast.transformer_xl.lower_upper_config.MetaEmbedding method), 100
 train() (flood_forecast.transformer_xl.transformer_basic.SimplePositionwiseFF (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFF method), 84
 train() (flood_forecast.transformer_xl.decoderBlock (flood_forecast.transformer_xl.multi_head_base.MultiAttnHeadSimple method), 109
 train() (flood_forecast.transformer_xl.decoderBlock (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFF method), 127
 train() (flood_forecast.transformer_xl.decoderBlock (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFF method), 134
 train() (flood_forecast.transformer_xl.decoderBlock (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFF method), 119
 train() (flood_forecast.transformer_xl.decoderBlock (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFF method), 160
 train() (flood_forecast.transformer_xl.decoderBlock (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFF method), 144
 train() (flood_forecast.transformer_xl.decoderBlock (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFF method), 127
 train() (flood_forecast.transformer_xl.decoderBlock (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFF method), 167
 train() (flood_forecast.transformer_xl.decoderBlock (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFF method), 152
 train() (flood_forecast.transformer_xl.decoderBlock (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFF method), 175
 train() (flood_forecast.transformer_xl.decoderBlock (flood_forecast.transformer_xl.lower_upper_config.PositionwiseFF method), 183
 train() (in module flood_forecast.da_rnn.train_da), type () (flood_forecast.transformer_xl.transformer_basic.SimplePositionwiseFF method), 187
 train_function() (in module flood_forecast.trainer), 17
 train_iteration() (in module flood_forecast.da_rnn.train_da), 187
 train_size(flood_forecast.da_rnn.custom_types.TrainConfig attribute), 191
 train_transformer_style() (in module flood_forecast.pytorch_training), 13
 TrainConfig (class in module flood_forecast.da_rnn.custom_types), 191
 TrainData (class in module flood_forecast.da_rnn.custom_types), 191
 TransformerXL (class in module flood_forecast.transformer_xl.transformer_xl), 175
 TransformerXL.to() (in module flood_forecast.transformer_xl.transformer_xl), 182
 type () (flood_forecast.custom.custom_opt.GaussianLoss method), 60
 type () (flood_forecast.custom.custom_opt.MAPELoss method), 52
 type () (flood_forecast.time_model.PyTorchForecast variable_forecast_layer() (in module flood_forecast.transformer_xl.lower_upper_config), 77

W

zero_grad() (*flood_forecast.transformer_xl.transformer_xl.StandardWordEmbedding*.
wandb_init () (*flood_forecast.time_model.PyTorchForecast* method), 175
method), 15 zero_grad() (*flood_forecast.transformer_xl.transformer_xl.TransformerXL*.
wandb_init () (*flood_forecast.time_model.TimeSeriesModel* method), 183
method), 15
warmup_constant () (in module
flood_forecast.custom.custom_opt), 37
warmup_cosine () (in module
flood_forecast.custom.custom_opt), 37
warmup_linear () (in module
flood_forecast.custom.custom_opt), 37

Z

zero_grad() (*flood_forecast.custom.custom_opt.BertAdam*
method), 68 zero_grad() (*flood_forecast.custom.custom_opt.GaussianLoss*
method), 60 zero_grad() (*flood_forecast.custom.custom_opt.MAPELoss*
method), 52 zero_grad() (*flood_forecast.custom.custom_opt.QuantileLoss*
method), 67 zero_grad() (*flood_forecast.custom.custom_opt.RMSELoss*
method), 45 zero_grad() (*flood_forecast.da_rnn.model.DARNN*
method), 201 zero_grad() (*flood_forecast.da_rnn.modules.Decoder*
method), 218 zero_grad() (*flood_forecast.da_rnn.modules.Encoder*
method), 211 zero_grad() (*flood_forecast.transformer_xl.dummy_torch.DummyTorchModel*
method), 76 zero_grad() (*flood_forecast.transformer_xl.lower_upper_config.AR*
method), 92 zero_grad() (*flood_forecast.transformer_xl.lower_upper_config.MetaEmbedding*
method), 100 zero_grad() (*flood_forecast.transformer_xl.lower_upper_config.PositionwiseFeedForward*
method), 85 zero_grad() (*flood_forecast.transformer_xl.multi_head_base.MultiAttnHeadSimple*
method), 109 zero_grad() (*flood_forecast.transformer_xl.transformer_basic.CustomTransformerDecoder*
method), 127 zero_grad() (*flood_forecast.transformer_xl.transformer_basic.SimplePositionalEncoding*
method), 135 zero_grad() (*flood_forecast.transformer_xl.transformer_basic.SimpleTransformer*
method), 119 zero_grad() (*flood_forecast.transformer_xl.transformer_xl.DecoderBlock*
method), 160 zero_grad() (*flood_forecast.transformer_xl.transformer_xl.MultiHeadAttention*
method), 145 zero_grad() (*flood_forecast.transformer_xl.transformer_xl.PositionalEmbedding*
method), 168 zero_grad() (*flood_forecast.transformer_xl.transformer_xl.PositionwiseFF*
method), 152