
orbit

Release 1.1.3

Edwin Ng, Steve Yang, Zhishi Wang, Yifeng Wu, Jing Pan

Nov 30, 2022

BASIC

1	About Orbit	1
2	Installation	3
3	Quick Start	5
4	Methods of Estimations and Predictions	9
5	Randomness Control and Reproducible Results	17
6	Using Pyro for Estimation	21
7	Damped Local Trend (DLT)	25
8	Local Global Trend (LGT)	33
9	Regression Priors in DLT	39
10	Regression Penalties in DLT	45
11	Handling Missing Response	55
12	Kernel-based Time-varying Regression - Part I	61
13	Kernel-based Time-varying Regression - Part II	69
14	Kernel-based Time-varying Regression - Part III	79
15	Kernel-based Time-varying Regression - Part IV	87
16	Prediction Decomposition	97
17	Model Diagnostics	101
18	Backtest	109
19	WBIC/BIC	119
20	EDA Utilities	129
21	Simulation Data	135
22	Other Utilities	139

23 Build Your Own Model	141
24 API Docs	149
25 Changelog	169
26 Indices and tables	175
Python Module Index	177
Index	179

ABOUT ORBIT

Orbit is a Python package for Bayesian time series modeling and inference. It provides a familiar and intuitive initialize-fit-predict interface for working with time series tasks, while utilizing probabilistic programming languages under the hood.

Currently, it supports the following models:

- Damped Local Trend (DLT)
- Exponential Smoothing (ETS)
- Local Global Trend (LGT)
- Kernel-based Time-varying Regression (KTR)

It also supports the following sampling methods for model estimation:

- Markov-Chain Monte Carlo (MCMC) as a full sampling method
- Maximum a Posteriori (MAP) as a point estimate method
- Stochastic Variational Inference (SVI) as a hybrid-sampling method on approximate distribution

Under the hood, the package is leveraging probabilistic program such as [pyro](#) and [PyStan 2.0](#).

1.1 Citation

To cite Orbit in publications, refer to the following whitepaper:

[Orbit: Probabilistic Forecast with Exponential Smoothing](#)

Bibtex:

```
@misc{
  ng2020orbit,
  title={Orbit: Probabilistic Forecast with Exponential Smoothing},
  author={Edwin Ng,
    Zhishi Wang,
    Huigang Chen,
    Steve Yang,
    Slawek Smyl
  },
  year={2020}, eprint={2004.08492}, archivePrefix={arXiv}, primaryClass={stat.CO}
}
```

1.2 Blog Post

1. Introducing Orbit, An Open Source Package for Time Series Inference and Forecasting [[Link](#)] 2. The New Version of Orbit (v1.1) is Released: The Improvements, Design Changes, and Exciting Collaborations [[Link](#)]

INSTALLATION

Install from PyPi:

```
pip install orbit-ml
```

Install from GitHub:

```
git clone https://github.com/uber/orbit.git
cd orbit
pip install -r requirements.txt
pip install .
```


QUICK START

This session covers topics:

- a forecast task on iclaims dataset
- a simple Bayesian ETS Model using PyStan
- posterior distribution extraction
- tools to visualize the forecast

3.1 Load Library

```
[1]: %matplotlib inline
import matplotlib.pyplot as plt

import orbit
from orbit.utils.dataset import load_iclaims
from orbit.models import ETS
from orbit.diagnostics.plot import plot_predicted_data
```

```
[2]: print(orbit.__version__)
```

```
1.1.3
```

3.2 Data

The *iclaims* data contains the weekly initial claims for US unemployment (obtained from [Federal Reserve Bank of St. Louis](#)) benefits against a few related Google trend queries (unemploy, filling and job) from Jan 2010 - June 2018. This aims to demo a similar dataset from the Bayesian Structural Time Series (BSTS) model ([Scott and Varian 2014](#)).

Note that the numbers are log-log transformed for fitting purpose and the discussion of **using the regressors** can be found in later chapters with the **Damped Local Trend (DLT)** model.

```
[3]: # load data
df = load_iclaims()
date_col = 'week'
response_col = 'claims'
df.dtypes
```

```
[3]: week          datetime64[ns]
      claims       float64
      trend.unemploy float64
      trend.filling float64
      trend.job     float64
      sp500         float64
      vix           float64
      dtype: object
```

```
[4]: df.head(5)
```

```
[4]:      week      claims  trend.unemploy  trend.filling  trend.job  sp500  \
0 2010-01-03  13.386595      0.219882      -0.318452   0.117500 -0.417633
1 2010-01-10  13.624218      0.219882      -0.194838   0.168794 -0.425480
2 2010-01-17  13.398741      0.236143      -0.292477   0.117500 -0.465229
3 2010-01-24  13.137549      0.203353      -0.194838   0.106918 -0.481751
4 2010-01-31  13.196760      0.134360      -0.242466   0.074483 -0.488929

      vix
0  0.122654
1  0.110445
2  0.532339
3  0.428645
4  0.487404
```

Train-test split.

```
[5]: test_size = 52
      train_df = df[:-test_size]
      test_df = df[-test_size:]
```

3.3 Forecasting Using Orbit

Orbit aims to provide an intuitive **initialize-fit-predict** interface for working with forecasting tasks. Under the hood, it utilizes probabilistic modeling API such as PyStan and Pyro. We first illustrate a Bayesian implementation of Rob Hyndman's ETS (which stands for Error, Trend, and Seasonality) Model (Hyndman et. al, 2008) using PyStan.

```
[6]: ets = ETS(
      response_col=response_col,
      date_col=date_col,
      seasonality=52,
      seed=8888,
      )
```

```
[7]: %%time
      ets.fit(df=train_df)


INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 1.000, warmups (per
↳chain): 225 and samples(per chain): 25.
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests.
↳for n_eff and Rhat.
To run all diagnostics call pystan.check_hmc_diagnostics(fit)
```

```
CPU times: user 46.7 ms, sys: 45.2 ms, total: 91.9 ms
Wall time: 862 ms
```

```
[7]: <orbit.forecaster.full_bayes.FullBayesianForecaster at 0x17b6e8e50>
```

```
[8]: predicted_df = ets.predict(df=test_df)
```

```
[9]: _ = plot_predicted_data(train_df, predicted_df, date_col, response_col, title=
↳ 'Prediction with ETS')
```



3.4 Extract and Analyze Posterior Samples

Users can use `.get_posterior_samples()` to extract posterior samples in an `OrderedDict` format.

```
[10]: posterior_samples = ets.get_posterior_samples()
posterior_samples.keys()
```

```
[10]: odict_keys(['l', 'lev_sm', 'obs_sigma', 's', 'sea_sm'])
```

The extracted parameters posteriors are pretty much compatible diagnostic with `arviz`. To do that, users can set `permute=False` to preserve chain information.

```
[11]: import arviz as az


posterior_samples = ets.get_posterior_samples(permute=False)

# example from https://arviz-devs.github.io/arviz/index.html
az.style.use("arviz-darkgrid")
az.plot_pair(
    posterior_samples,
```

(continues on next page)

(continued from previous page)

```
var_names=["sea_sm", "lev_sm", "obs_sigma"],  
kind="kde",  
marginals=True,  
textsize=15,  
)  
plt.show()
```



For more details in model diagnostics visualization, there is a subsequent section dedicated to it.

METHODS OF ESTIMATIONS AND PREDICTIONS

There are three methods supported in Orbit model parameters estimation (a.k.a posteriors in Bayesian).

1. Maximum a Posteriori (MAP)
2. Markov Chain Monte Carlo (MCMC)
3. Stochastic Variational Inference (SVI)

This session will cover the first two: **MAP** and **MCMC** which mainly uses [PyStan2.0](#) at the back end. Users can simply can leverage the args `estimator` to pick the method (`stan-map` and `stan-mcmc`). The details will be covered by the sections below. The SVI method is calling [Pyro](#) by specifying `estimator='pyro-svi'`. However, it is covered by a separate session.

4.1 Data and Libraries

```
[1]: %matplotlib inline
import matplotlib.pyplot as plt

import orbit
from orbit.utils.dataset import load_iclaims
from orbit.models import ETS
from orbit.diagnostics.plot import plot_predicted_data, plot_predicted_components
```

```
[2]: print(orbit.__version__)

1.1.3
```

```
[3]: # load data
df = load_iclaims()
test_size = 52
train_df = df[:-test_size]
test_df = df[-test_size:]
response_col = 'claims'
date_col = 'week'
```

4.2 Maximum a Posteriori (MAP)


To use MAP method, one can simply specify `estimator='stan-map'` when instantiating a model. The advantage of MAP estimation is a faster computational speed. In MAP, the uncertainty is mainly generated the noise process with bootstrapping. However, the uncertainty would not cover parameters variance as well as the credible interval from seasonality or other components.

```
[4]: %%time
ets = ETS(
    response_col=response_col,
    date_col=date_col,
    estimator='stan-map',
    seasonality=52,
    seed=8888,
)
ets.fit(df=train_df)
predicted_df = ets.predict(df=test_df)

INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.

CPU times: user 15.2 ms, sys: 4.71 ms, total: 20 ms
Wall time: 143 ms
```

```
[5]: _ = plot_predicted_data(train_df, predicted_df, date_col, response_col, title=
    ↪ 'Prediction with ETSMAP')
```



To have the uncertainty from MAP, one can specify `n_bootstrap_draws`. The default is set to be `-1` which mutes the bootstrap process. Users can also specify a particular percentiles to report prediction intervals by passing list of percentiles with args `prediction_percentiles`.

```
[6]: # default: [10, 90]
prediction_percentiles=[10, 90]
```

(continues on next page)

(continued from previous page)


```

ets = ETS(
    response_col=response_col,
    date_col=date_col,
    estimator='stan-map',
    seasonality=52,
    seed=8888,
    n_bootstrap_draws=1e4,
    prediction_percentiles=prediction_percentiles,
)
ets.fit(df=train_df)
predicted_df = ets.predict(df=test_df)

_ = plot_predicted_data(train_df, predicted_df, date_col, response_col,
                       prediction_percentiles=prediction_percentiles,
                       title='Prediction with ETS-MAP')

```

INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.



One can access the posterior estimated by calling the `.get_point_posteriors()`. The outcome from this function is a dict of dict where the top layer stores the type of point estimate while the second layer stores the parameters labels and values.

```
[7]: pt_posteriors = ets.get_point_posteriors()['map']
pt_posteriors.keys()
```

```
[7]: dict_keys(['l', 'lev_sm', 'obs_sigma', 's', 'sea_sm'])
```

In general, the first dimension is just 1 as a point estimate for each parameter. The rest of the dimension will depend on the dimension of parameter itself.

```
[8]: lev = pt_posteriors['l']
lev.shape
```

```
[8]: (1, 391)
```

4.3 MCMC

To use MCMC method, one can specify `estimator='stan-mcmc'` (the default) when instantiating a model. Compared to MAP, it usually takes longer time to fit. As the model now fitted as a **full Bayesian** model where **No-U-Turn Sampler (NUTS)** (Hoffman and Gelman 2011) is carried out under the hood. By default, a full sampling on posteriors distribution is conducted. Hence, full distribution of the predictions are always provided.


4.3.1 MCMC - Full Bayesian Sampling

```
[9]: %%time
ets = ETS(
    response_col=response_col,
    date_col=date_col,
    estimator='stan-mcmc',
    seasonality=52,
    seed=8888,
    num_warmup=400,
    num_sample=400,
)
ets.fit(df=train_df)
predicted_df = ets.predict(df=test_df)

INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 1.000, warmups (per
↪chain): 100 and samples(per chain): 100.
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests.
↪for n_eff and Rhat.
To run all diagnostics call pystan.check_hmc_diagnostics(fit)


CPU times: user 159 ms, sys: 74.6 ms, total: 234 ms
Wall time: 464 ms

[10]: _ = plot_predicted_data(train_df, predicted_df, date_col, response_col, title=
↪'Prediction with ETS-Full Bayesian')
```



Also, users can request prediction with credible intervals of each component.

```
[11]: predicted_df = ets.predict(df=df, decompose=True)
      plot_predicted_components(predicted_df, date_col=date_col,
                              plot_components=['prediction', 'trend', 'seasonality'])
```



```
[11]: array([<AxesSubplot:title={'center':'prediction'}>,
             <AxesSubplot:title={'center':'trend'}>,
             <AxesSubplot:title={'center':'seasonality'}>], dtype=object)
```

Just like the `MAPForecaster`, one can also access the posterior samples by calling the function `.get_posterior_samples()`.

```
[12]: posterior_samples = ets.get_posterior_samples()
posterior_samples.keys()
[12]: odict_keys(['l', 'lev_sm', 'obs_sigma', 's', 'sea_sm'])
```

As mentioned, in **MCMC (Full Bayesian)** models, the first dimension reflects the sample size.

```
[13]: lev = posterior_samples['l']
lev.shape
[13]: (400, 391)
```

4.3.2 MCMC - Point Estimation

Users can also choose to derive point estimates via MCMC by specifying `point_method` as `mean` or `median` via the call of `.fit`. In that case, posteriors samples are first aggregated by mean or median and store as a point estimate for final prediction. Just like other point estimate, users can specify `n_bootstrap_draws` to report uncertainties.


```
[14]: %%time
ets = ETS(
    response_col=response_col,
    date_col=date_col,
    estimator='stan-mcmc',
    seasonality=52,
    seed=8888,
    n_bootstrap_draws=1e4,
)

# specify point_method e.g. 'mean', 'median'
ets.fit(df=train_df, point_method='mean')
predicted_df = ets.predict(df=test_df)

INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 1.000, warmups (per_
↳chain): 225 and samples(per chain): 25.
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests_
↳for n_eff and Rhat.
To run all diagnostics call pystan.check_hmc_diagnostics(fit)

CPU times: user 178 ms, sys: 91.7 ms, total: 270 ms
Wall time: 779 ms

[15]: _ = plot_predicted_data(train_df, predicted_df, date_col, response_col,
                             title='Prediction with Point(Mean) Estimation')
```



One can always access the the point estimated posteriors by `.get_point_posteriors()` (including the cases fitting the parameters through MCMC).

```
[16]: ets.get_point_posteriors()['mean'].keys()
```

```
[16]: dict_keys(['l', 'lev_sm', 'obs_sigma', 's', 'sea_sm'])
```

```
[17]: ets.get_point_posteriors()['median'].keys()
```

```
[17]: dict_keys(['l', 'lev_sm', 'obs_sigma', 's', 'sea_sm'])
```


RANDOMNESS CONTROL AND REPRODUCIBLE RESULTS

There are randomness involved in the random initialization, sampling and bootstrapping process. Some of them with sufficient condition such as converging status and large amount of samples, can be fixed even without a fixed seed. However, they are not guaranteed. Two settings needed to allow fully reproducible results and will be demoed from this session:

1. fix the seed on fitting
2. fix the seed on prediction

5.1 Data and Libraries

```
[1]: import numpy as np

import orbit
from orbit.models import DLT
from orbit.utils.dataset import load_iclaims
```

```
[2]: print(orbit.__version__)

1.1.3
```

```
[3]: df = load_iclaims()
df.head(5)
```

```
[3]:
```

	week	claims	trend.unemploy	trend.filling	trend.job	sp500	\
0	2010-01-03	13.386595	0.219882	-0.318452	0.117500	-0.417633	
1	2010-01-10	13.624218	0.219882	-0.194838	0.168794	-0.425480	
2	2010-01-17	13.398741	0.236143	-0.292477	0.117500	-0.465229	
3	2010-01-24	13.137549	0.203353	-0.194838	0.106918	-0.481751	
4	2010-01-31	13.196760	0.134360	-0.242466	0.074483	-0.488929	

	vix
0	0.122654
1	0.110445
2	0.532339
3	0.428645
4	0.487404

5.2 Fixing Seed in Fitting

By default, the seed supplied during the **initialization** step is fixed. This allows fully reproducible posteriors samples by default. For other purpose, users can randomize the seed. Nonetheless, this process usually assumes stable result with or without a fixed seed. Otherwise, convergence alert should be raised.

With different seeds, results should be closed but not identical:

```
[4]: dlt1 = DLT(response_col='claims', date_col='week', seed=2021, estimator='stan-map', n_
      ↪bootstrap_draws=1e3)
      dlt2 = DLT(response_col='claims', date_col='week', seed=2020, estimator='stan-map', n_
      ↪bootstrap_draws=1e3)

      dlt1.fit(df);
      dlt2.fit(df);

      lev1 = dlt1.get_point_posteriors()['map']['l']
      lev2 = dlt2.get_point_posteriors()['map']['l']

      INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
      INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
```

```
[5]: np.all(lev1 == lev2)
```

```
[5]: False
```

```
[6]: np.allclose(lev1, lev2, rtol=1e-3)
```

```
[6]: True
```

With fixed seeds, results should be identical:

```
[7]: dlt1 = DLT(response_col='claims', date_col='week', seed=2020, estimator='stan-map', n_
      ↪bootstrap_draws=1e3)
      dlt2 = DLT(response_col='claims', date_col='week', seed=2020, estimator='stan-map', n_
      ↪bootstrap_draws=1e3)

      dlt1.fit(df);
      dlt2.fit(df);

      lev1 = dlt1.get_point_posteriors()['map']['l']
      lev2 = dlt2.get_point_posteriors()['map']['l']

      INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
      INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
```

```
[8]: np.all(lev1 == lev2)
```

```
[8]: True
```

In sampling algorithm, users should expect identical posteriors with fixed seed:

```
[9]: dlt_mcmc1 = DLT(response_col='claims', date_col='week', seed=2020, estimator='stan-mcmc')
      dlt_mcmc2 = DLT(response_col='claims', date_col='week', seed=2020, estimator='stan-mcmc')
```

(continues on next page)

(continued from previous page)

```
dlt_mcmc1.fit(df);
dlt_mcmc2.fit(df);

lev_mcmc1 = dlt_mcmc1.get_posterior_samples()['l']
lev_mcmc2 = dlt_mcmc2.get_posterior_samples()['l']
```

INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 1.000, warmups (per chain): 225 and samples(per chain): 25.
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests for n_eff and Rhat.
To run all diagnostics call pystan.check_hmc_diagnostics(fit)
INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 1.000, warmups (per chain): 225 and samples(per chain): 25.
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests for n_eff and Rhat.
To run all diagnostics call pystan.check_hmc_diagnostics(fit)

```
[10]: print(lev_mcmc1.shape)
print(lev_mcmc2.shape)
np.all(lev1 == lev2)
```

```
(100, 443)
(100, 443)
```

```
[10]: True
```

5.3 Fixing Seed in Prediction

Unlike the fitting process, the seed in prediction is set to be random by default unless users provided a fixed seed. Once a fixed seed provided through the args in `.predict()`. Users should expect identical result.

```
[11]: # check with MAP estimator
pred1 = dlt1.predict(df, seed=2020)['prediction'].values
pred2 = dlt2.predict(df, seed=2020)['prediction'].values
np.all(pred1 == pred2)
```

```
[11]: True
```

```
[12]: # check with MCMC estimator
pred1 = dlt_mcmc1.predict(df, seed=2020)['prediction'].values
pred2 = dlt_mcmc2.predict(df, seed=2020)['prediction'].values
np.all(pred1 == pred2)
```

```
[12]: True
```


USING PYRO FOR ESTIMATION

Note

Currently we are still experimenting with Pyro and support Pyro only in LGT and KTR models.

Pyro is a flexible, scalable deep probabilistic programming library built on PyTorch. **Pyro** was originally developed at Uber AI and is now actively maintained by community contributors, including a dedicated team at the Broad Institute.

```
[1]: %matplotlib inline

import pandas as pd
import numpy as np
import matplotlib.pyplot as plt

import orbit
from orbit.models import LGT
from orbit.diagnostics.plot import plot_predicted_data
from orbit.diagnostics.plot import plot_predicted_components
from orbit.utils.dataset import load_iclaims

from orbit.constants.palette import OrbitPalette

[2]: print(orbit.__version__)

1.1.3

[3]: df = load_iclaims()

[4]: test_size=52
train_df=df[:-test_size]
test_df=df[-test_size:]
```

6.1 VI Fit and Predict

Although Pyro provides a variety of ways to optimize/sample posteriors. Currently, we only support Stochastic Variational Inference (SVI). For details, please refer to this [doc](#).

To use SVI for LGT, specify estimator as `pyro-svi`.

```
[5]: lgt_vi = LGT(
      response_col='claims',
      date_col='week',
      seasonality=52,
      seed=8888,
      estimator='pyro-svi',
      num_steps=101,
      num_sample=300,
      # trigger message per 50 steps
      message=50,
      learning_rate=0.1,
    )
```

```
[6]: %%time
      lgt_vi.fit(df=train_df)


INFO:orbit:Using SVI (Pyro) with steps: 101, samples: 300, learning rate: 0.1, learning_
↳rate_total_decay: 1.0 and particles: 100.
INFO:root:Guessed max_plate_nesting = 2
INFO:orbit:step    0 loss = 658.91, scale = 0.11635
INFO:orbit:step   50 loss = -432, scale = 0.48623
INFO:orbit:step  100 loss = -444.07, scale = 0.34976

CPU times: user 4.48 s, sys: 549 ms, total: 5.03 s
Wall time: 4.58 s
```

```
[6]: <orbit.forecaster.svi.SVIForecaster at 0x15c777d30>
```

```
[7]: predicted_df = lgt_vi.predict(df=test_df)
```

```
[8]: _ = plot_predicted_data(training_actual_df=train_df, predicted_df=predicted_df,
      date_col=lgt_vi.date_col, actual_col=lgt_vi.response_col,
      test_actual_df=test_df)
```




We can also extract the ELBO loss from the training metrics.

```
[9]: loss_elbo = lgt_vi.get_training_metrics()['loss_elbo']
```

```
[10]: steps = np.arange(len(loss_elbo))
plt.subplots(1, 1, figsize=(8, 4))
plt.plot(steps, loss_elbo, color=OrbitPalette.BLUE.value)
plt.title('ELBO Loss per Step')
```

```
[10]: Text(0.5, 1.0, 'ELBO Loss per Step')
```



DAMPED LOCAL TREND (DLT)

This section covers topics including:

- DLT model structure
- DLT global trend configurations
- Adding regressors in DLT
- Other configurations

```
[1]: %matplotlib inline
import matplotlib.pyplot as plt

import orbit
from orbit.models import DLT
from orbit.diagnostics.plot import plot_predicted_data, plot_predicted_components
from orbit.utils.dataset import load_iclaims
```

```
[2]: print(orbit.__version__)
```

```
1.1.3
```

7.1 Model Structure

DLT is one of the main exponential smoothing models we support in `orbit`. Performance is benchmarked with M3 monthly, M4 weekly dataset and some Uber internal dataset (Ng and Wang et al., 2020). The model is a fusion between the classical ETS (Hyndman et. al., 2008)) with some refinement leveraging ideas from Rlgt (Smyl et al., 2019). The model has a structural forecast equations

$$\begin{aligned}y_t &= \mu_t + s_t + r_t + \epsilon_t \\ \mu_t &= g_t + l_{t-1} + \theta b_{t-1} \\ \epsilon_t &\sim \text{Student}(\nu, 0, \sigma) \\ \sigma &\sim \text{HalfCauchy}(0, \gamma_0)\end{aligned}$$

with the update process

$$\begin{aligned}g_t &= D(t) \\ l_t &= \rho_l(y_t - g_t - s_t - r_t) + (1 - \rho_l)(l_{t-1} + \theta b_{t-1}) \\ b_t &= \rho_b(l_t - l_{t-1}) + (1 - \rho_b)\theta b_{t-1} \\ s_{t+m} &= \rho_s(y_t - l_t - r_t) + (1 - \rho_s)s_t \\ r_t &= \sum_j \beta_j x_{jt}\end{aligned}$$

One important point is that using y_t as a log-transformed response usually yield better result, especially we can interpret such log-transformed model as a *multiplicative form* of the original model. Besides, there are two new additional components compared to the classical damped ETS model:

1. $D(t)$ as the deterministic trend process
2. r as the regression component with x as the regressors

```
[3]: # load log-transformed data
df = load_iclaims()
response_col = 'claims'
date_col = 'week'
```

Note

Just like LGT model, we also provide MAP and MCMC (full Bayesian) methods for DLT model (by specifying `estimator='stan-map'` or `estimator='stan-mcmc'` when instantiating a model).

MCMC is usually more robust but may take longer time to train. In this notebook, we will use the MAP method for illustration purpose.

7.2 Global Trend Configurations

There are a few choices of $D(t)$ configured by `global_trend_option`:

1. `linear` (default)
2. `loglinear`
3. `flat`
4. `logistic`

Mathematically, they are expressed as such,

1. Linear:

$$D(t) = \delta_{\text{intercept}} + \delta_{\text{slope}} \cdot t$$

2. Log-linear:

$$D(t) = \delta_{\text{intercept}} + \ln(\delta_{\text{slope}} \cdot t)$$

3. Logistic:

$$D(t) = L + \frac{U-L}{1+e^{-\delta_{\text{slope}} \cdot t}}$$

4. Flat:

$$D(t) = \delta_{\text{intercept}}$$

where $\delta_{\text{intercept}}$ and δ_{slope} are fitted parameters and t is rescaled time-step between 0 and T (=number of time steps).

To show the difference among these options, their predictions are projected in the charts below. Note that the default is set to `linear` which is also used in the benchmarking process mentioned previously. During prediction, a convenient function `make_future_df()` is called to generate future data frame (ONLY applied when you don't have any regressors!).


7.2.1 linear global trend

```
[4]: %%time

dlt = DLT(
    response_col=response_col,
    date_col=date_col,
    estimator='stan-map',
    seasonality=52,
    seed=8888,
    global_trend_option='linear',
    # for prediction uncertainty
    n_bootstrap_draws=1000,
)

dlt.fit(df)
test_df = dlt.make_future_df(periods=52 * 10)
predicted_df = dlt.predict(test_df)
_ = plot_predicted_data(df, predicted_df, date_col, response_col, title='DLT Linear
↪Global Trend')
```

INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.



CPU times: user 445 ms, sys: 43.4 ms, total: 489 ms
Wall time: 451 ms

One can use `.get_posterior_samples()` to extract the samples for all sampling parameters.


```
[5]: dlt.get_posterior_samples().keys()

[5]: dict_keys(['l', 'b', 'lev_sm', 'slp_sm', 'obs_sigma', 'nu', 'lt_sum', 's', 'sea_sm', 'gt_
↪sum', 'gb', 'gl'])
```

```
[6]: %%time
# log-linear global trend
dlt = DLT(
    response_col=response_col,
    date_col=date_col,
    seasonality=52,
    estimator='stan-map',
    seed=8888,
    global_trend_option='loglinear',
    # for prediction uncertainty
    n_bootstrap_draws=1000,
)

dlt.fit(df)
# re-use the test_df generated above
predicted_df = dlt.predict(test_df)
_ = plot_predicted_data(df, predicted_df, date_col, response_col, title='DLT Log-Linear_
↪Global Trend')
```

INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.



CPU times: user 664 ms, sys: 96.7 ms, total: 761 ms
Wall time: 392 ms

In logistic trend, users need to specify the args `global_floor` and `global_cap`. These args are with default 0 and 1.

7.2.2 logistic global trend

[7]: %%time

```
dlt = DLT(
    response_col=response_col,
    date_col=date_col,
    estimator='stan-map',
    seasonality=52,
    seed=8888,
    global_trend_option='logistic',
    global_cap=9999,
    global_floor=11.75,
    damped_factor=0.1,
    # for prediction uncertainty
    n_bootstrap_draws=1000,
)


dlt.fit(df)
predicted_df = dlt.predict(test_df)
ax = plot_predicted_data(df, predicted_df, date_col, response_col,
                        title='DLT Logistic Global Trend', is_visible=False);
ax.axhline(y=11.75, linestyle='--', color='orange')
ax.figure
```

INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.

CPU times: user 414 ms, sys: 51.2 ms, total: 466 ms

Wall time: 333 ms

[7]:




Note: Theoretically, the trend is bounded by the `global_floor` and `global_cap`. However, because of seasonality and regression, the predictions can still be slightly lower than the floor or higher than the cap.

7.2.3 flat trend

```
[8]: %%time
```

```
dlt = DLT(  
    response_col=response_col,  
    date_col=date_col,  
    estimator='stan-map',  
    seasonality=52,  
    seed=8888,  
    global_trend_option='flat',  
    # for prediction uncertainty  
    n_bootstrap_draws=1000,  
)  
  
dlt.fit(df)  
predicted_df = dlt.predict(test_df)  
_ = plot_predicted_data(df, predicted_df, date_col, response_col, title='DLT Flat_  
↪Global Trend')
```

```
INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
```



```
CPU times: user 696 ms, sys: 71.4 ms, total: 767 ms  
Wall time: 351 ms
```

7.3 Regression

You can also add regressors into the model by specifying `regressor_col`. This serves the purpose of nowcasting or forecasting when exogenous regressors are known such as events and holidays. Without losing generality, the interface is set to be

$$\beta_j \sim \mathcal{N}(\mu_j, \sigma_j^2)$$

where $\mu_j = 0$ and $\sigma_j = 1$ by default as a non-informative prior. These two parameters are set by the arguments `regressor_beta_prior` and `regressor_sigma_prior` as a list. For example,


```
[9]: dlt = DLT(
      response_col=response_col,
      date_col=date_col,
      estimator='stan-mcmc',
      seed=8888,
      seasonality=52,
      regressor_col=['trend.unemploy', 'trend.filling'],
      regressor_beta_prior=[0.1, 0.3],
      regressor_sigma_prior=[0.5, 2.0],
    )
```

```
dlt.fit(df)
predicted_df = dlt.predict(df, decompose=True)
plot_predicted_components(predicted_df, date_col);
```

```
INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 1.000, warmups (per_
↪chain): 225 and samples(per chain): 25.
```

```
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests_
↪for n_eff and Rhat.
```

```
To run all diagnostics call pystan.check_hmc_diagnostics(fit)
```



One can also use `.get_regression_coefs` to extract the regression coefficients along with the confidence interval when posterior samples are available. The default lower and upper limits are set to be `.05` and `.95`.

```
[10]: dlt.get_regression_coefs()
```

```
[10]:      regressor regressor_sign coefficient coefficient_lower \
0 trend.unemploy      Regular    0.049229      0.022658
1 trend.filling      Regular    0.076010     -0.020334

      coefficient_upper Pr(coef >= 0) Pr(coef < 0)
0          0.072649      1.00      0.00
1          0.147020      0.91      0.09
```

There are much more configurations on regression such as the regressors sign and penalty type. They will be discussed in subsequent sections.

7.3.1 High Dimensional and Fourier Series Regression

In case of high dimensional regression, users can consider fixing the smoothness with a relatively small levels smoothing values e.g. setting `level_sm_input=0.01`. This is particularly useful in modeling higher frequency time-series such as daily and hourly data using regression on Fourier series. Check out the `examples/` folder for more details.

LOCAL GLOBAL TREND (LGT)

In this section, we will cover:

- LGT model structure
- difference between DLT and LGT
- syntax to call LGT classes with different estimation methods

LGT stands for Local and Global Trend and is a refined model from **Rlgt** (Smyl et al., 2019). The main difference is that LGT is an additive form taking log-transformation response as the modeling response. This essentially converts the model into a multiplicative with some advantages (Ng and Wang et al., 2020). **However, one drawback of this approach is that negative response values are not allowed due to the existence of the global trend term and because of that we start to deprecate the support of regression of this model.**

```
[1]: %matplotlib inline

import pandas as pd
import numpy as np
import matplotlib.pyplot as plt

import orbit
from orbit.models import LGT
from orbit.diagnostics.plot import plot_predicted_data
from orbit.diagnostics.plot import plot_predicted_components
from orbit.utils.dataset import load_iclaims
```

```
[2]: print(orbit.__version__)
```

```
1.1.3dev
```

8.1 Model Structure

$$\begin{aligned}y_t &= \mu_t + s_t + \epsilon_t \\ \mu_t &= l_{t-1} + \xi_1 b_{t-1} + \xi_2 l_{t-1}^\lambda \\ \epsilon_t &\sim \text{Student}(\nu, 0, \sigma) \\ \sigma &\sim \text{HalfCauchy}(0, \gamma_0)\end{aligned}$$

with the update process,

$$\begin{aligned}l_t &= \rho_l(y_t - s_t) + (1 - \rho_l)l_{t-1} \\ b_t &= \rho_b(l_t - l_{t-1}) + (1 - \rho_b)b_{t-1} \\ s_{t+m} &= \rho_s(y_t - l_t) + (1 - \rho_s)s_t\end{aligned}$$

Unlike **DLT** model which has a deterministic trend, **LGT** introduces a hybrid trend where it consists of

- local trend takes on a fraction ξ_1 rather than a damped factor
- global trend is with a auto-regressive term ξ_2 and a power term λ

We will continue to use the *iclaims* data with 52 weeks train-test split.

```
[3]: # load data
df = load_iclaims()
# define date and response column
date_col = 'week'
response_col = 'claims'
df.dtypes
test_size = 52
train_df = df[:-test_size]
test_df = df[-test_size:]
```

8.2 LGT Model

In orbit, we provide three methods for LGT model estimation and inferences, which are * MAP * MCMC (also providing the point estimate method, mean or median), which is also the default * SVI

Orbit follows a sklearn style model API. We can create an instance of the Orbit class and then call its fit and predict methods.

In this notebook, we will only cover MAP and MCMC methods. Refer to [this notebook](#) for the pyro estimation.

8.2.1 LGT - MAP

To use MAP, specify the estimator as `stan-map`.

```
[4]: lgt = LGT(
    response_col=response_col,
    date_col=date_col,
    estimator='stan-map',
    seasonality=52,
    seed=8888,
)
```

```
INFO:orbit:Optimizing(PyStan) with algorithm:LBFGS .
```


```
[5]: %%time
lgt.fit(df=train_df)

CPU times: user 218 ms, sys: 8.52 ms, total: 226 ms
Wall time: 405 ms
```

```
[5]: <orbit.forecaster.map.MAPForecaster at 0x142a8c650>
```

```
[6]: predicted_df = lgt.predict(df=test_df)
```

```
[7]: _ = plot_predicted_data(training_actual_df=train_df, predicted_df=predicted_df,
                           date_col=date_col, actual_col=response_col,
                           test_actual_df=test_df, title='Prediction with LGTMAP Model')
```



8.2.2 LGT - MCMC

To use MCMC sampling, specify the estimator as `stan-mcmc` (the default).

- By default, full Bayesian samples will be used for the predictions: for each set of parameter posterior samples, the prediction will be conducted once and the final predictions are aggregated over all the results. To be specific, the final predictions will be the median (aka 50th percentile) along with any additional percentiles provided. One can use `.get_posterior_samples()` to extract the samples for all sampling parameters.
- One can also specify `point_method` (either `mean` or `median`) via `.fit` to have the point estimate: the parameter posterior samples are aggregated first (mean or median) then conduct the prediction once.

LGT - full

```
[8]: lgt = LGT(
    response_col=response_col,
    date_col=date_col,
    seasonality=52,
    seed=8888,
)
```

```
[9]: %%time
lgt.fit(df=train_df)

INFO:orbit:Sampling(PyStan) with chains:4, cores:8, temperature:1.000, warmups(per_
↳chain):225 and samples(per chain):25.
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests_
↳for n_eff and Rhat.
To run all diagnostics call pystan.check_hmc_diagnostics(fit)
WARNING:pystan:2 of 100 iterations ended with a divergence (2 %).
WARNING:pystan:Try running with adapt_delta larger than 0.8 to remove the divergences.

CPU times: user 77.7 ms, sys: 63.6 ms, total: 141 ms
Wall time: 5.86 s

[9]: <orbit.forecaster.full_bayes.FullBayesianForecaster at 0x10a5b8d50>
```

```
[10]: predicted_df = lgt.predict(df=test_df)
```

```
[11]: predicted_df.tail(5)
```


```
[11]:
```

	week	prediction_5	prediction	prediction_95
47	2018-05-27	12.138702	12.236653	12.325029
48	2018-06-03	12.042856	12.164003	12.282972
49	2018-06-10	12.177416	12.271215	12.393355
50	2018-06-17	12.126099	12.228332	12.343101
51	2018-06-24	12.151659	12.278477	12.375120

```
[12]: lgt.get_posterior_samples().keys()
```

```
[12]: odict_keys(['l', 'b', 'lev_sm', 'slp_sm', 'obs_sigma', 'nu', 'lgt_sum', 'gt_pow', 'lt_
↳coef', 'gt_coef', 's', 'sea_sm'])
```

```
[13]: _ = plot_predicted_data(training_actual_df=train_df, predicted_df=predicted_df,
date_col=lgt.date_col, actual_col=lgt.response_col,
test_actual_df=test_df, title='Prediction with LGTFull Model')
```



LGT - point estimate

```
[14]: lgt = LGT(
      response_col=response_col,
      date_col=date_col,
      seasonality=52,
      seed=8888,
    )
```

```
[15]: %%time
      lgt.fit(df=train_df, point_method='mean')
```

```
INFO:orbit:Sampling(PyStan) with chains:4, cores:8, temperature:1.000, warmups(per_
↳chain):225 and samples(per chain):25.
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests_
↳for n_eff and Rhat.
To run all diagnostics call pystan.check_hmc_diagnostics(fit)
WARNING:pystan:2 of 100 iterations ended with a divergence (2 %).
WARNING:pystan:Try running with adapt_delta larger than 0.8 to remove the divergences.
```

```
CPU times: user 83 ms, sys: 71.9 ms, total: 155 ms
Wall time: 6.13 s
```

```
[15]: <orbit.forecaster.full_bayes.FullBayesianForecaster at 0x1430e8c90>
```

```
[16]: predicted_df = lgt.predict(df=test_df)
```

```
[17]: predicted_df.tail(5)
```


```
[17]:      week  prediction
      47 2018-05-27  12.198260
```

(continues on next page)

(continued from previous page)

```
48 2018-06-03 12.134447
49 2018-06-10 12.227925
50 2018-06-17 12.195388
51 2018-06-24 12.241482
```

```
[18]: _ = plot_predicted_data(training_actual_df=train_df, predicted_df=predicted_df,
                        date_col=lgt.date_col, actual_col=lgt.response_col,
                        test_actual_df=test_df, title='Predictibon with LGTAggregated Model')
```



REGRESSION PRIORS IN DLT

This notebook demonstrates usage of priors in the regression analysis. The *iclaims* data will be used in demo purpose. Examples include

1. regression with default setting
2. regression with bounded priors for regression coefficients

Generally speaking, regression coefficients are more robust under full Bayesian sampling and estimation. The default setting `estimator='stan-mcmc'` will be used in this tutorial.

```
[1]: %matplotlib inline

import matplotlib.pyplot as plt
import numpy as np

import orbit
from orbit.utils.dataset import load_iclaims
from orbit.models import DLT
from orbit.diagnostics.plot import plot_predicted_data
from orbit.constants.palette import OrbitPalette
```

```
[2]: print(orbit.__version__)

1.1.3
```

9.1 US Weekly Initial Claims

Recall the *iclaims* dataset by previous section. In order to use this data to nowcast the US unemployment claims during COVID-19 period, the dataset is extended to Jan 2021 and the S&P 500 (^GSPC) and VIX Index historical data are attached for the same period.

The data is standardized and log-transformed for the model fitting purpose.

```
[3]: # load data
df = load_iclaims(end_date='2021-01-03')
df = df[['week', 'claims', 'trend.unemploy', 'trend.job', 'sp500', 'vix']]
df = df[1:].reset_index(drop=True)

date_col = 'week'
response_col = 'claims'
df.dtypes
```

```
[3]: week          datetime64[ns]
      claims        float64
      trend.unemploy float64
      trend.job      float64
      sp500          float64
      vix            float64
      dtype: object
```

```
[4]: df.head(5)
```


```
[4]:
```

	week	claims	trend.unemploy	trend.job	sp500	vix
0	2010-01-10	13.624218	0.016351	0.181862	-0.550891	0.069878
1	2010-01-17	13.398741	0.032611	0.130569	-0.590640	0.491772
2	2010-01-24	13.137549	-0.000179	0.119987	-0.607162	0.388078
3	2010-01-31	13.196760	-0.069172	0.087552	-0.614339	0.446838
4	2010-02-07	13.146984	-0.182500	0.019344	-0.605636	0.308205

We can see from the plot below, there are seasonality, trend, and as well as a huge change point due the impact of COVID-19.

```
[5]: fig, axs = plt.subplots(2, 2, figsize=(20, 8))
      axs[0, 0].plot(df['week'], df['claims'])
      axs[0, 0].set_title('Unemployment Claims')
      axs[0, 1].plot(df['week'], df['trend.unemploy'], 'tab:orange')
      axs[0, 1].set_title('Google trend - unemploy')
      axs[1, 0].plot(df['week'], df['vix'], 'tab:green')
      axs[1, 0].set_title('VIX')
      axs[1, 1].plot(df['week'], df['sp500'], 'tab:red')
      axs[1, 1].set_title('S&P500')
```

```
[5]: Text(0.5, 1.0, 'S&P500')
```



```
[6]: # using relatively updated data
      df[['sp500']] = df[['sp500']].diff()
      df = df[1:].reset_index(drop=True)

      test_size = 12
```

(continues on next page)

(continued from previous page)

```
train_df = df[:-test_size]
test_df = df[-test_size:]
```

9.1.1 Naive Model

Here we will use DLT models to compare the model performance with vs. without regression.

```
[7]: %%time
dlt = DLT(
    response_col=response_col,
    date_col=date_col,
    seasonality=52,
    seed=8888,
    num_warmup=4000,
)
dlt.fit(df=train_df)
predicted_df = dlt.predict(df=test_df)
```

INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 1.000, warmups (per chain): 1000 and samples(per chain): 25.
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests for n_eff and Rhat.
To run all diagnostics call pystan.check_hmc_diagnostics(fit)

CPU times: user 152 ms, sys: 108 ms, total: 260 ms
Wall time: 7.01 s

9.1.2 DLT With Regression

The regressor columns can be supplied via argument `regressor_col`. Recall the regression formula in **DLT**:

$$\hat{y}_t = \mu_t + s_t + r_t$$

$$r_t = \sum_j \beta_j x_{jt}$$

$$\beta_j \sim \mathcal{N}(\mu_j, \sigma_j^2)$$

By default, $\mu_j = 0$ and $\sigma_j = 1$. In addition, we can set a *sign* constraint for each coefficient β_j . This is can be done by supplying the `regressor_sign` as a list where elements are in one of followings:

- ‘=’: $\beta_j \sim \mathcal{N}(0, \sigma_j^2)$ i.e. $\beta_j \in (-\text{inf}, \text{inf})$
- ‘+’: $\beta_j \sim \mathcal{N}^+(0, \sigma_j^2)$ i.e. $\beta_j \in [0, \text{inf})$
- ‘-’: $\beta_j \sim \mathcal{N}^-(0, \sigma_j^2)$ i.e. $\beta_j \in (-\text{inf}, 0]$

Based on some intuition, it’s reasonable to assume search terms such as “unemployment”, “filling” and **VIX** index to be positively correlated and stock index such as **SP500** to be negatively correlated to the outcome. Then we will leave whatever unspecified as a regular regressor.

```
[8]: %%time
dlt_reg = DLT(
    response_col=response_col,
```

(continues on next page)

(continued from previous page)

```

date_col=date_col,
regressor_col=['trend.unemploy', 'trend.job', 'sp500', 'vix'],
seasonality=52,
seed=8888,
num_warmup=4000,
)
dlt_reg.fit(df=train_df)
predicted_df_reg = dlt_reg.predict(test_df)

```

INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 1.000, warmups (per chain): 1000 and samples(per chain): 25.
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests for n_eff and Rhat.
To run all diagnostics call `pystan.check_hmc_diagnostics(fit)`

CPU times: user 81.1 ms, sys: 85.6 ms, total: 167 ms
Wall time: 6.96 s

The estimated regressor coefficients can be retrieved via `.get_regression_coefs()`.

```
[9]: dlt_reg.get_regression_coefs()
```

```
[9]:
```

	regressor	regressor_sign	coefficient	coefficient_lower	\
0	trend.unemploy	Regular	0.080230	0.048558	
1	trend.job	Regular	-0.039983	-0.083835	
2	sp500	Regular	-0.010453	-0.202477	
3	vix	Regular	0.010462	-0.024078	
	coefficient_upper	Pr(coef >= 0)	Pr(coef < 0)		
0	0.106878	1.00	0.00		
1	0.009523	0.13	0.87		
2	0.215410	0.42	0.58		
3	0.036222	0.71	0.29		

9.1.3 Regression with Informative Priors

Due to various reasons, users may obtain further knowledge on some of the regressors or they want to propose different regularization on different regressors. These informative priors basically means to replace the defaults (μ, σ) mentioned previously. In orbit, this process is done via the arguments `regressor_beta_prior` and `regressor_sigma_prior`. These two lists should be of the same length as `regressor_col`.

In addition, we can set a *sign* constraint for each coefficient β_j . This is can be done by supplying the `regressor_sign` as a list where elements are in one of followings:

- ‘=’: $\beta_j \sim \mathcal{N}(0, \sigma_j^2)$ i.e. $\beta_j \in (-\text{inf}, \text{inf})$
- ‘+’: $\beta_j \sim \mathcal{N}^+(0, \sigma_j^2)$ i.e. $\beta_j \in [0, \text{inf})$
- ‘-’: $\beta_j \sim \mathcal{N}^-(0, \sigma_j^2)$ i.e. $\beta_j \in (-\text{inf}, 0]$

Based on intuition, it’s reasonable to assume search terms such as “unemployment”, “filling” and **VIX** index to be positively correlated (+ sign is used in this case) and upward shock of **SP500** (- sign) to be negatively correlated to the outcome. Otherwise, an unbounded coefficient can be used (= sign).

Furthermore, regressors such as search queries may have more direct impact than stock marker indices. Hence, a smaller σ is considered.

```
[10]: dlt_reg_adjust = DLT(
      response_col=response_col,
      date_col=date_col,
      regressor_col=['trend.unemploy', 'trend.job', 'sp500', 'vix'],
      regressor_sign=['+', '=', '-', '+'],
      regressor_sigma_prior=[0.3, 0.1, 0.05, 0.1],
      num_warmup=4000,
      num_sample=1000,
      estimator='stan-mcmc',
      seed=2022,
    )
dlt_reg_adjust.fit(df=train_df)
predicted_df_reg_adjust = dlt_reg_adjust.predict(test_df)
```

```
INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 1.000, warmups (per_
↳chain): 1000 and samples(per chain): 250.
```

```
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests_
↳for n_eff and Rhat.
```

```
To run all diagnostics call pystan.check_hmc_diagnostics(fit)
```

```
[11]: dlt_reg_adjust.get_regression_coefs()
```

```
[11]:
```

	regressor	regressor_sign	coefficient	coefficient_lower \
0	trend.unemploy	Positive	0.132037	0.078566
1	vix	Positive	0.019541	0.001662
2	sp500	Negative	-0.030996	-0.095246
3	trend.job	Regular	-0.011236	-0.080350

	coefficient_upper	Pr(coef >= 0)	Pr(coef < 0)
0	0.210839	1.000	0.000
1	0.054155	1.000	0.000
2	-0.003587	0.000	1.000
3	0.060355	0.386	0.614

Let's compare the holdout performance by using the built-in function `smape()` .

```
[12]: def mae(x, y):
      return np.mean(np.abs(x - y))

naive_mae = mae(predicted_df['prediction'].values, test_df['claims'].values)
reg_mae = mae(predicted_df_reg['prediction'].values, test_df['claims'].values)
reg_adjust_mae = mae(predicted_df_reg_adjust['prediction'].values, test_df['claims'].
↳values)

print("-----Mean Absolute Error Summary-----")
print("Naive Model: {:.3f}\nRegression Model: {:.3f}\nRefined Regression Model: {:.3f}"
↳format(
      naive_mae, reg_mae, reg_adjust_mae
    ))
```

```
-----Mean Absolute Error Summary-----
Naive Model: 0.249
Regression Model: 0.249
Refined Regression Model: 0.097
```

9.2 Summary

This demo showcases a use case in nowcasting. Although this may not be applicable in real-time forecasting, it mainly introduces the regression analysis with time-series modeling in `Orbit`. For people who have concerns on the forecastability, one can consider introducing lag on regressors.

Also, `Orbit` allows informative priors where sometime can be useful in combining multiple source of insights together.

REGRESSION PENALTIES IN DLT

This notebook continues to discuss regression problems with DLT and covers various penalties:

1. fixed-ridge
2. auto-ridge
3. lasso

Generally speaking, regression coefficients are more robust under full Bayesian sampling and estimation. The default setting `estimator='stan-mcmc'` will be used in this tutorial. Besides, a fixed and small smoothing parameters are used such as `level_sm_input=0.01` and `slope_sm_input=0.01` to facilitate high dimensional regression.

```
[1]: %matplotlib inline

import matplotlib.pyplot as plt
import numpy as np

import orbit
from orbit.utils.dataset import load_iclaims
from orbit.models import DLT
from orbit.diagnostics.plot import plot_predicted_data
from orbit.constants.palette import OrbitPalette
```

```
[2]: print(orbit.__version__)

1.1.3
```

10.1 Regression on Simulated Dataset

A simulated dataset is used to demonstrate sparse regression.

```
[3]: import pandas as pd
from orbit.utils.simulation import make_trend, make_regression
from orbit.diagnostics.metrics import mse
```

A few utilities from the package is used to generate simulated data. For details, please refer to the API doc. In brief, the process generates observations y such that

$$y_t = l_t + \sum_p^P \beta_p x_{t,p}$$

for $t = 1, 2, \dots, T$

where

$$l_t = l_{t-1} + \delta_t$$

$$\delta_t \sim N(0, \sigma_\delta)$$

10.1.1 Regular Regression

To begin with, the setting $P = 10$ and $T = 100$ is used.

```
[4]: NUM_OF_REGRESSORS = 10
      SERIES_LEN = 50
      SEED = 20210101
      # sample some coefficients
      COEFS = np.random.default_rng(SEED).uniform(-1, 1, NUM_OF_REGRESSORS)
      trend = make_trend(SERIES_LEN, rw_loc=0.01, rw_scale=0.1)
      x, regression, coefs = make_regression(series_len=SERIES_LEN, coefs=COEFS)
      print(regression.shape, x.shape)
```

```
(50,) (50, 10)
```

```
[5]: # combine trend and the regression
      y = trend + regression
```

```
[6]: x_cols = [f"x{x}" for x in range(1, NUM_OF_REGRESSORS + 1)]
      response_col = "y"
      dt_col = "date"
      obs_matrix = np.concatenate([y.reshape(-1, 1), x], axis=1)
      # make a data frame for orbit inputs
      df = pd.DataFrame(obs_matrix, columns=[response_col] + x_cols)
      # make some dummy date stamp
      dt = pd.date_range(start='2016-01-04', periods=SERIES_LEN, freq="1W")
      df['date'] = dt
      df.shape
```

```
[6]: (50, 12)
```

Here is a peek on the coefficients.

```
[7]: coefs
```

```
[7]: array([ 0.38372743, -0.21084054,  0.5404565 , -0.21864409,  0.85529298,
         -0.83838077, -0.54550632,  0.80367924, -0.74643654, -0.26626975])
```

By default, `regression_penalty` is set as `fixed-ridge` i.e.

$$\beta_j \sim \mathcal{N}(\mu_j, \sigma_j^2)$$

with a default $\mu_j = 0$ and $\sigma_j = 1$

Fixed Ridge Penalty

```
[8]: %%time
dlt_fridge = DLT(
    response_col=response_col,
    date_col=dt_col,
    regressor_col=x_cols,
    seed=SEED,
    # this is default
    regression_penalty='fixed_ridge',
    # fixing the smoothing parameters to learn regression coefficients more effectively
    level_sm_input=0.01,
    slope_sm_input=0.01,
    num_warmup=4000,
)
dlt_fridge.fit(df=df)
```

```
INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 1.000, warmups (per_
↳chain): 1000 and samples(per chain): 25.
```

```
WARNING:pystan:n_eff / iter below 0.001 indicates that the effective sample size has_
↳likely been overestimated
```


```
WARNING:pystan:Rhat above 1.1 or below 0.9 indicates that the chains very likely have_
↳not mixed
```

```
CPU times: user 46.3 ms, sys: 43 ms, total: 89.3 ms
```

```
Wall time: 1.4 s
```

```
[8]: <orbit.forecaster.full_bayes.FullBayesianForecaster at 0x177f1f610>
```

```
[9]: coef_fridge = np.quantile(dlt_fridge._posterior_samples['beta'], q=[0.05, 0.5, 0.95],_
↳axis=0 )
lw=3
idx = np.arange(NUM_OF_REGRESSORS)
plt.figure(figsize=(20, 8))
plt.title("Weights of the model", fontsize=20)
plt.plot(idx, coef_fridge[1], color=OrbitPalette.GREEN.value, linewidth=lw, drawstyle=
↳'steps', label='Fixed-Ridge', alpha=0.5, linestyle='--')
plt.fill_between(idx, coef_fridge[0], coef_fridge[2], step='pre', alpha=0.3,_
↳color=OrbitPalette.GREEN.value)
plt.plot(coefs, color="black", linewidth=lw, drawstyle='steps', label="Ground truth")
plt.ylim(1, -1)
plt.legend(prop={'size': 20})
plt.grid()
```



Auto-Ridge Penalty

Users can also set the `regression_penalty` to be `auto_ridge` in case users are not sure what to set for the `regressor_sigma_prior`.

Instead of using fixed scale in the coefficients prior, a prior can be assigned to them, i.e.

$$\sigma_j \sim \text{Cauchy}^+(0, \alpha)$$

This can be done by setting `regression_penalty="auto_ridge"` with the argument `auto_ridge_scale` (default of 0.5) set the prior α . A higher `adapt_delta` is recommend to reduce divergence. Check [here](#) for details of `adapt_delta`.

```
[10]: %%time
dlt_auto_ridge = DLT(
  response_col=response_col,
  date_col=dt_col,
  regressor_col=x_cols,
  seed=SEED,
  # this is default
  regression_penalty='auto_ridge',
  # fixing the smoothing parameters to learn regression coefficients more effectively
  level_sm_input=0.01,
  slope_sm_input=0.01,
  num_warmup=4000,
  # reduce divergence
  stan_mcmc_control={'adapt_delta':0.9},
)
dlt_auto_ridge.fit(df=df)
```

```
INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 1.000, warmups (per
↳chain): 1000 and samples(per chain): 25.
```


```
WARNING:pystan:n_eff / iter below 0.001 indicates that the effective sample size has
↳likely been overestimated
```

```
WARNING:pystan:Rhat above 1.1 or below 0.9 indicates that the chains very likely have
↳not mixed
```

```
CPU times: user 210 ms, sys: 84.7 ms, total: 295 ms
Wall time: 2.53 s
```

```
[10]: <orbit.forecaster.full_bayes.FullBayesianForecaster at 0x178041f70>
```

```
[11]: coef_auto_ridge = np.quantile(dlt_auto_ridge._posterior_samples['beta'], q=[0.05, 0.5, 0.
↳95], axis=0 )
lw=3
idx = np.arange(NUM_OF_REGRESSORS)
plt.figure(figsize=(20, 8))
plt.title("Weights of the model", fontsize=24)
plt.plot(idx, coef_auto_ridge[1], color=OrbitPalette.GREEN.value, linewidth=lw,
↳drawstyle='steps', label='Auto-Ridge', alpha=0.5, linestyle='--')
plt.fill_between(idx, coef_auto_ridge[0], coef_auto_ridge[2], step='pre', alpha=0.3,
↳color=OrbitPalette.GREEN.value)
plt.plot(coefs, color="black", linewidth=lw, drawstyle='steps', label="Ground truth")
plt.ylim(1, -1)
plt.legend(prop={'size': 20})
plt.grid();
```



```
[12]: print('Fixed Ridge MSE: {:.3f}\nAuto Ridge MSE: {:.3f}'.format(
mse(coef_fridge[1], coefs), mse(coef_auto_ridge[1], coefs)
))
```

```
Fixed Ridge MSE:0.103
Auto Ridge MSE:0.075
```

10.1.2 Sparse Regrsson

In reality, users usually faces a more challenging problem with a much higher P to N ratio with a sparsity specified by the parameter `relevance=0.5` under the simulation process.

```
[13]: NUM_OF_REGRESSORS = 50
      SERIES_LEN = 50
      SEED = 20210101
      COEFS = np.random.default_rng(SEED).uniform(0.3, 0.5, NUM_OF_REGRESSORS)
      SIGNS = np.random.default_rng(SEED).choice([1, -1], NUM_OF_REGRESSORS)
      # to mimic a either zero or relative observable coefficients
      COEFS = COEFS * SIGNS
      trend = make_trend(SERIES_LEN, rw_loc=0.01, rw_scale=0.1)
      x, regression, coefs = make_regression(series_len=SERIES_LEN, coefs=COEFS, relevance=0.5)
      print(regression.shape, x.shape)

(50,) (50, 50)
```

```
[14]: # generated sparsed coefficients
      coefs
```

```
[14]: array([[ 0.          ,  0.          , -0.45404565,  0.37813559,  0.          ,
          0.          ,  0.          ,  0.48036792, -0.32535635, -0.37337302,
        -0.42474576,  0.          , -0.37000755,  0.44887456,  0.47082836,
          0.          ,  0.32678039,  0.37436121,  0.38932392,  0.40216056,
          0.          ,  0.          , -0.3076828 , -0.35036047,  0.          ,
          0.          ,  0.          ,  0.          ,  0.          ,  0.          ,
        -0.45838674,  0.3171478 ,  0.          ,  0.          ,  0.          ,
          0.          ,  0.          ,  0.41599814,  0.          , -0.30964341,
        -0.42072894,  0.36255583,  0.          , -0.39326337,  0.44455655,
          0.          ,  0.          ,  0.30064161, -0.46083203,  0.          ]])
```

```
[15]: # combine trend and the regression
      y = trend + regression
```

```
[16]: x_cols = [f"x{x}" for x in range(1, NUM_OF_REGRESSORS + 1)]
      response_col = "y"
      dt_col = "date"
      obs_matrix = np.concatenate([y.reshape(-1, 1), x], axis=1)
      # make a data frame for orbit inputs
      df = pd.DataFrame(obs_matrix, columns=[response_col] + x_cols)
      # make some dummy date stamp
      dt = pd.date_range(start='2016-01-04', periods=SERIES_LEN, freq="1W")
      df['date'] = dt
      df.shape
```

```
[16]: (50, 52)
```


10.1.3 Fixed Ridge Penalty

```
[17]: dlt_fridge = DLT(
        response_col=response_col,
        date_col=dt_col,
        regressor_col=x_cols,
        seed=SEED,
        level_sm_input=0.01,
        slope_sm_input=0.01,
        num_warmup=8000,
    )
dlt_fridge.fit(df=df)
```

```
INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 1.000, warmups (per
↳chain): 2000 and samples(per chain): 25.
WARNING:pystan:n_eff / iter below 0.001 indicates that the effective sample size has
↳likely been overestimated
WARNING:pystan:Rhat above 1.1 or below 0.9 indicates that the chains very likely have
↳not mixed
```

```
[17]: <orbit.forecaster.full_bayes.FullBayesianForecaster at 0x1783c0340>
```

```
[18]: coef_fridge = np.quantile(dlt_fridge._posterior_samples['beta'], q=[0.05, 0.5, 0.95],
↳axis=0 )
lw=3
idx = np.arange(NUM_OF_REGRESSORS)
plt.figure(figsize=(20, 8))
plt.title("Weights of the model", fontsize=24)
plt.plot(coef_fridge[1], color=OrbitPalette.GREEN.value, linewidth=lw, drawstyle='steps
↳', label="Ridge", alpha=0.5, linestyle='--')
plt.fill_between(idx, coef_fridge[0], coef_fridge[2], step='pre', alpha=0.3,
↳color=OrbitPalette.GREEN.value)
plt.plot(coefs, color="black", linewidth=lw, drawstyle='steps', label="Ground truth")
plt.legend(prop={'size': 20})
plt.grid();
```



LASSO Penalty


In high P to N problems, *LASSO* penalty usually shines compared to *Ridge* penalty.

```
[19]: dlt_lasso = DLT(
    response_col=response_col,
    date_col=dt_col,
    regressor_col=x_cols,
    seed=SEED,
    regression_penalty='lasso',
    level_sm_input=0.01,
    slope_sm_input=0.01,
    num_warmup=8000,
)
dlt_lasso.fit(df=df)
```

```
INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 1.000, warmups (per
↳chain): 2000 and samples(per chain): 25.
WARNING:pystan:n_eff / iter below 0.001 indicates that the effective sample size has
↳likely been overestimated
WARNING:pystan:Rhat above 1.1 or below 0.9 indicates that the chains very likely have
↳not mixed
```

```
[19]: <orbit.forecaster.full_bayes.FullBayesianForecaster at 0x17a5489d0>
```

```
[20]: coef_lasso = np.quantile(dlt_lasso._posterior_samples['beta'], q=[0.05, 0.5, 0.95],
↳axis=0 )
lw=3
idx = np.arange(NUM_OF_REGRESSORS)
plt.figure(figsize=(20, 8))
plt.title("Weights of the model", fontsize=24)
plt.plot(coef_lasso[1], color=OrbitPalette.GREEN.value, linewidth=lw, drawstyle='steps',
↳label="Lasso", alpha=0.5, linestyle='--')
plt.fill_between(idx, coef_lasso[0], coef_lasso[2], step='pre', alpha=0.3,
↳color=OrbitPalette.GREEN.value)
plt.plot(coefs, color="black", linewidth=lw, drawstyle='steps', label="Ground truth")
plt.legend(prop={'size': 20})
plt.grid();
```



```
[21]: print('Fixed Ridge MSE: {:.3f}\nLASSO MSE: {:.3f}'.format(
        mse(coef_fridge[1], coefs), mse(coef_lasso[1], coefs)
    ))
```

```
Fixed Ridge MSE:0.174
LASSO MSE:0.098
```

10.2 Summary

This notebook covers a few choices of penalty in regression regularization. A lasso and auto-ridge can be considered in highly sparse data.

HANDLING MISSING RESPONSE

Because of the generative nature of the exponential smoothing models, they can naturally deal with missing response during the training process. It simply replaces observations by prediction during the 1-step ahead generating process. Below users can find the simple examples of how those exponential smoothing models handle missing responses.

```
[1]: import pandas as pd
import numpy as np
import orbit
import matplotlib.pyplot as plt

from orbit.utils.dataset import load_iclaims
from orbit.diagnostics.plot import plot_predicted_data, plot_predicted_components
from orbit.utils.plot import get_orbit_style
from orbit.models import ETS, LGT, DLT
from orbit.diagnostics.metrics import smape

plt.style.use(get_orbit_style())

%load_ext autoreload
%autoreload 2

%matplotlib inline
```

```
[2]: orbit.__version__
```

```
[2]: '1.1.3'
```

11.1 Data

```
[3]: # can also consider transform=False
raw_df = load_iclaims(transform=True)
raw_df.dtypes

df = raw_df.copy()
df.head()
```

```
[3]:
```

	week	claims	trend.unemploy	trend.filling	trend.job	sp500	\
0	2010-01-03	13.386595	0.219882	-0.318452	0.117500	-0.417633	
1	2010-01-10	13.624218	0.219882	-0.194838	0.168794	-0.425480	
2	2010-01-17	13.398741	0.236143	-0.292477	0.117500	-0.465229	

(continues on next page)

(continued from previous page)

```

3 2010-01-24 13.137549      0.203353      -0.194838      0.106918 -0.481751
4 2010-01-31 13.196760      0.134360      -0.242466      0.074483 -0.488929

      vix
0  0.122654
1  0.110445
2  0.532339
3  0.428645
4  0.487404

```

```
[4]: test_size=52
```

```

train_df=df[:-test_size]
test_df=df[-test_size:]

```

11.1.1 Define Missing Data

Now, we manually created a dataset with a few missing values in the response variable.

```
[5]: na_idx = np.arange(53, 100, 1)
na_idx
```

```
[5]: array([53, 54, 55, 56, 57, 58, 59, 60, 61, 62, 63, 64, 65, 66, 67, 68, 69,
        70, 71, 72, 73, 74, 75, 76, 77, 78, 79, 80, 81, 82, 83, 84, 85, 86,
        87, 88, 89, 90, 91, 92, 93, 94, 95, 96, 97, 98, 99])
```

```
[6]: train_df_na = train_df.copy()
train_df_na.iloc[na_idx, 1] = np.nan
```

11.2 Exponential Smoothing Examples

11.2.1 ETS

```
[7]: ets = ETS(
      response_col='claims',
      date_col='week',
      seasonality=52,
      seed=2022,
      estimator='stan-mcmc'
    )
ets.fit(train_df_na)
ets_predicted = ets.predict(df=train_df_na)
```

```

INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 1.000, warmups (per
↳chain): 225 and samples(per chain): 25.
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests
↳for n_eff and Rhat.
To run all diagnostics call pystan.check_hmc_diagnostics(fit)

```

11.2.2 LGT

```
[8]: lgt = LGT(
      response_col='claims',
      date_col='week',
      estimator='stan-mcmc',
      seasonality=52,
      seed=2022
    )
lgt.fit(df=train_df_na)
lgt_predicted = lgt.predict(df=train_df_na)
```

INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 1.000, warmups (per chain): 225 and samples(per chain): 25.
 WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests for n_eff and Rhat.
 To run all diagnostics call `pystan.check_hmc_diagnostics(fit)`

11.2.3 DLT

```
[9]: dlt = DLT(
      response_col='claims',
      date_col='week',
      estimator='stan-mcmc',
      seasonality=52,
      seed=2022
    )
dlt.fit(df=train_df_na)
dlt_predicted = dlt.predict(df=train_df_na)
```

INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 1.000, warmups (per chain): 225 and samples(per chain): 25.
 WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests for n_eff and Rhat.
 To run all diagnostics call `pystan.check_hmc_diagnostics(fit)`

11.2.4 Summary

Users can verify this behavior with a table and visualization of the actuals and predicted.

```
[10]: train_df_na['ets-predict'] = ets_predicted['prediction']
      train_df_na['lgt-predict'] = lgt_predicted['prediction']
      train_df_na['dlt-predict'] = dlt_predicted['prediction']
```

```
[11]: # table summary of prediction during absence of observations
      train_df_na.iloc[na_idx, :].head(5)
```

```
[11]:
```

	week	claims	trend.unemploy	trend.filling	trend.job	sp500	\
53	2011-01-09	NaN	0.152060	-0.127397	0.085412	-0.295869	
54	2011-01-16	NaN	0.186546	-0.044015	0.074483	-0.303546	
55	2011-01-23	NaN	0.169451	-0.004795	0.074483	-0.309024	

(continues on next page)


(continued from previous page)

56	2011-01-30	NaN	0.079300	0.032946	-0.005560	-0.282329
57	2011-02-06	NaN	0.060252	-0.024213	0.006275	-0.268480
	vix	ets-predict	lgt-predict	dlt-predict		
53	-0.036658	13.512120	13.507156	13.509457		
54	0.141233	13.272154	13.281752	13.286773		
55	0.222816	13.003021	13.012511	13.010722		
56	-0.006710	13.050246	13.069293	13.066893		
57	-0.021891	13.000305	13.010235	13.016641		

```
[12]: from orbit.constants.palette import OrbitPalette
```

```
# just to get some color from orbit palette
orbit_palette = [
    OrbitPalette.BLACK.value,
    OrbitPalette.BLUE.value,
    OrbitPalette.GREEN.value,
    OrbitPalette.YELLOW.value,
]
```

```
[13]: pred_list = ['ets-predict', 'lgt-predict', 'dlt-predict']
fig, axes = plt.subplots(len(pred_list), 1, figsize=(16, 16))
for idx, p in enumerate(pred_list):
    axes[idx].scatter(train_df_na['week'], train_df_na['claims'].values,
                     label='actuals' if idx == 0 else '', color=orbit_palette[0],
                     ↪alpha=0.5)
    axes[idx].plot(train_df_na['week'], train_df_na[p].values,
                  label=p, color=orbit_palette[idx + 1], lw=2.5)
fig.legend()
fig.tight_layout()
```



11.3 First Observation Exception

It is worth pointing out that the very first value of the response variable cannot be missing since this is the starting point of the time series fitting. **An error message will be raised when the first observation in response is missing.**

```
[14]: na_idx2 = list(na_idx) + [0]
      train_df_na2 = train_df.copy()
      train_df_na2.iloc[na_idx2, 1] = np.nan
      ets.fit(train_df_na2)
```

```

-----
DataInputException                                Traceback (most recent call last)
Input In [14], in <cell line: 4>()
      2 train_df_na2 = train_df.copy()
      3 train_df_na2.iloc[na_idx2, 1] = np.nan
----> 4 ets.fit(train_df_na2)

File ~/edwinnglabs/orbit/orbit/forecaster/full_bayes.py:36, in FullBayesianForecaster.
-> fit(self, df, point_method, keep_samples, sampling_temperature, **kwargs)
      28 def fit(
      29     self,
      30     df,
      (...)
      34     **kwargs,
      35 ):
----> 36     super().fit(df, sampling_temperature=sampling_temperature, **kwargs)
      37     self._point_method = point_method
      39     if point_method is not None:

File ~/edwinnglabs/orbit/orbit/forecaster/forecaster.py:147, in Forecaster.fit(self, df,
-> **kwargs)
      145 self._set_training_meta(df)
      146 # customize module
--> 147 self._model.set_dynamic_attributes(
      148     df=df, training_meta=self.get_training_meta()
      149 )
      150 # based on the model and df, set training input
      151 self.set_training_data_input()

File ~/edwinnglabs/orbit/orbit/template/ets.py:149, in ETSModel.set_dynamic_
-> attributes(self, df, training_meta)
      145 else:
      146     # should not be used anyway; just a placeholder
      147     self.seasonality_sd = training_meta[TrainingMetaKeys.RESPONSE_SD.value]
--> 149 self._set_valid_response(training_meta)

File ~/edwinnglabs/orbit/orbit/template/ets.py:156, in ETSModel._set_valid_response(self,
-> training_meta)
      154 # raise exception if the first response value is missing
      155 if self.is_valid_response[0] == 0:
--> 156     raise DataInputException(
      157         "The first value of response column {} cannot be missing..".format(
      158             training_meta[TrainingMetaKeys.RESPONSE_COL.value]
      159         )
      160     )

DataInputException: The first value of response column claims cannot be missing..

```

KERNEL-BASED TIME-VARYING REGRESSION - PART I

Kernel-based time-varying regression (KTR) is a time series model to address

1. time-varying regression coefficients
2. complex seasonality pattern

The full details of the model structure with an application in marketing media mix modeling can be found in Ng, Wang and Dai (2021). The core of KTR is the use of latent variables to define a smooth time varying representation of model coefficients, which bears a similar ideas in [Kernel Smoothing](#). The KTR approach sharply reduces the number of parameters compared to typical dynamic linear models such as Harvey (1989), and Durbin and Koopman (2002). The reduced number of parameters improves the computation speed, and allows for handling of high dimensional data and detecting small variances.

To topics covered here in **Part I**, are

1. KTR model structure
2. syntax to initialize, fit and predict a model
3. fit a model with complex seasonality
4. visualization of prediction and decomposed components

```
[1]: import numpy as np
import pandas as pd
import matplotlib.pyplot as plt

import orbit
from orbit.models import KTR
from orbit.diagnostics.plot import plot_predicted_data, plot_predicted_components
from orbit.utils.dataset import load_electricity_demand

%matplotlib inline
pd.set_option('display.float_format', lambda x: '%.5f' % x)
```

```
[2]: print(orbit.__version__)
```

```
1.1.3
```

12.1 Model Structure

This section gives the mathematical structure of the KTR model. In short, it considers a time-series (y_t) as the linear combination of three parts which are the local-trend (l_t), seasonality (s_t), and regression (r_t) terms. Mathematically,

$$y_t = l_t + s_t + r_t + \epsilon_t, \quad t = 1, \dots, T,$$

where the ϵ_t comprise a stationary random error process.

In **KTR**, the distinction between the local-trend, seasonality, and regressors while useful is semi-arbitrary and the time-series can also be considered as

$$y_t = X_t^T \beta_t + \epsilon_t, \quad t = 1, \dots, T,$$

where β_t is a P -dimensional vector of coefficients that vary over time (i.e., β_i is almost certainly different from β_j for $i \neq j$) and X_t P -dimensional covariate vector (i.e., the t th row of X , the design matrix).

To reduce the total number of parameters in the model (potentially $P \times T$) the β_t are parameterized with a weighted sum of J local latent variables (b_1, \dots, b_J). That is

$$B = Kb^T$$

where - *coefficient matrix* B has size $T \times P$ with rows equal to the β_t . - *knot matrix* b with size $P \times J$; each entry is a latent variable $b_{p,j}$. The b_j can be viewed as the “knots” from the perspective of spline regression and j is a time index such that $t_j \in [1, \dots, T]$. - *kernel matrix* K with size $T \times J$ where the i th row and j th element can be viewed as the normalized weight $k(t_j, t) / \sum_{j=1}^J k(t_j, t)$

For the level/trend,

$$l_t = \beta_{t,\text{lev}}$$

It can also be viewed as a dynamic intercept (where the regressor is a vector of ones).

For the seasonality,

$$B_{\text{seas}} = K_{\text{seas}} b_{\text{seas}}^T$$

$$s_t = X_{t,\text{seas}} \beta_{t,\text{seas}}$$

We use Fourier series to handle the seasonality; i.e., sin waves with varying periods are used for the columns of X_{seas} .

The details for the additional regressors are given in **Part II**, as they are not used in this tutorial. Note this includes different choices of kernel function (which determines the kernel matrix K) and prior for matrix b .

12.2 Data

To illustrate the usage of KTR, consider the daily series of electricity demand in Turkey from the 2000 - 2008.

```
[3]: # from 2000-01-01 to 2008-12-31
df = load_electricity_demand()
date_col = 'date'
response_col = 'electricity'
df[response_col] = np.log(df[response_col])
print(df.shape)
df.head()
```

```
(3288, 2)
```

```
/var/folders/g0/77v1jc9s455cj14mkqcht0m00000gn/T/ipykernel_1700/982097363.py:2:
↳ UserWarning: Parsing '31/12/2008' in DD/MM/YYYY format. Provide format or specify
↳ infer_datetime_format=True for consistent parsing.
df = load_electricity_demand()
```

```
[3]:
```

	date	electricity
0	2000-01-01	9.43760
1	2000-01-02	9.50130
2	2000-01-03	9.63565
3	2000-01-04	9.65392
4	2000-01-05	9.66089

```
[4]: print(f'starts with {df[date_col].min()} \nends with {df[date_col].max()} \nshape: {df.
↳ shape}')

```

```
starts with 2000-01-01 00:00:00
ends with 2008-12-31 00:00:00
shape: (3288, 2)
```

12.2.1 Train / Test Split

Split the data into a training set and test set for model validation.

```
[5]: test_size=365
train_df=df[:-test_size]
test_df=df[-test_size:]
```

12.3 A Quick Start on KTR

Here the Similar to other model types in Orbit, KTR follows sklearn model API style. First an instance of the Orbit class KTR is created. Second fit and predict methods are called for that instance. Note that unlike version $\leq 1.0.15$, the fitting API arg are within the function; thus, KTR is called directly.

```
[6]: ktr = KTR(
    response_col=response_col,
    date_col=date_col,
    seed=2021,
    estimator='pyro-svi',
    # bootstrap sampling to capture uncertainties
    n_bootstrap_draws=1e4,
    # pyro training config
    num_steps=301,
    message=100,
)
```

```
[7]: ktr.fit(train_df)
```

```
INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
INFO:orbit:Using SVI (Pyro) with steps: 301, samples: 100, learning rate: 0.1, learning_
↳ rate_total_decay: 1.0 and particles: 100.
```

(continues on next page)

(continued from previous page)

```
INFO:root:Guessed max_plate_nesting = 1
INFO:orbit:step 0 loss = -1946.6, scale = 0.093118
INFO:orbit:step 100 loss = -3136.6, scale = 0.01007
INFO:orbit:step 200 loss = -3144.1, scale = 0.0097119
INFO:orbit:step 300 loss = -3132.8, scale = 0.0098468
```

```
[7]: <orbit.forecaster.svi.SVIForecaster at 0x1640be100>
```

We can take a look how the level is fitted with the data.

```
[8]: predicted_df = ktr.predict(df=df)
predicted_df.head()
```

```
[8]:
```


	date	prediction_5	prediction	prediction_95
0	2000-01-01	9.47117	9.60187	9.73470
1	2000-01-02	9.46718	9.59950	9.73052
2	2000-01-03	9.46743	9.60012	9.73023
3	2000-01-04	9.46891	9.59984	9.73081
4	2000-01-05	9.46597	9.59948	9.73270

One can use `.get_posterior_samples()` to extract the samples for all sampling parameters.

```
[9]: ktr.get_posterior_samples().keys()
```


```
[9]: dict_keys(['lev_knot', 'lev', 'yhat', 'obs_scale'])
```

```
[10]: _ = plot_predicted_data(training_actual_df=train_df, predicted_df=predicted_df,
                             date_col=date_col, actual_col=response_col,
                             test_actual_df=test_df, markersize=20, lw=.5)
```



It can also be helpful to see the trend knot locations and levels. This is done with the `plot_lev_knots` function.

```
[11]: _ = ktr.plot_lev_knots()
```



12.4 Fitting with Complex Seasonality

The previous model fit is not satisfactory as there is clear seasonality in the electrical demand time-series that is not accounted for. In this modelling example the electrical demand data is fit with a dual seasonality for weekly and yearly patterns. Since the data is daily, the seasonality periods are 7 and 365.25. These are added into the KTR object as a list through the `seasonality` arg. Otherwise the process is the same as the previous example.

```
[12]: ktr_with_seas = KTR(
    response_col=response_col,
    date_col=date_col,
    seed=2021,
    seasonality=[7, 365.25],
    estimator='pyro-svi',
    n_bootstrap_draws=1e4,
    # pyro training config
    num_steps=301,
    message=100,
)
```

```
[13]: ktr_with_seas.fit(train_df)
```

```
INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
INFO:orbit:Using SVI (Pyro) with steps: 301, samples: 100, learning_rate: 0.1, learning_
↪rate_total_decay: 1.0 and particles: 100.
INFO:root:Guessed max_plate_nesting = 1
INFO:orbit:step    0 loss = -2190.9, scale = 0.093667
INFO:orbit:step  100 loss = -4379.9, scale = 0.0069612
INFO:orbit:step  200 loss = -4311.1, scale = 0.0071341
INFO:orbit:step  300 loss = -4294.1, scale = 0.0071724
```

```
[13]: <orbit.forecaster.svi.SVIForecaster at 0x1659933a0>
```

```
[14]: predicted_df = ktr_with_seas.predict(df=df, decompose=True)
```

```
[15]: predicted_df.head(5)
```

```
[15]:
```

	date	prediction_5	prediction	prediction_95	trend_5	trend	\
0	2000-01-01	9.54219	9.62491	9.70913	9.51294	9.59566	
1	2000-01-02	9.48945	9.57346	9.65672	9.51005	9.59405	
2	2000-01-03	9.55081	9.63458	9.71768	9.51056	9.59432	
3	2000-01-04	9.61004	9.69301	9.77611	9.51109	9.59406	
4	2000-01-05	9.59257	9.67726	9.76164	9.50927	9.59396	

	trend_95	regression_5	regression	regression_95	seasonality_7_5	\
0	9.67988	0.00000	0.00000	0.00000	-0.02771	
1	9.67731	0.00000	0.00000	0.00000	-0.07843	
2	9.67743	0.00000	0.00000	0.00000	-0.01842	
3	9.67717	0.00000	0.00000	0.00000	0.03947	
4	9.67833	0.00000	0.00000	0.00000	0.02308	

	seasonality_7	seasonality_7_95	seasonality_365.25_5	seasonality_365.25	\
0	-0.02771	-0.02771	0.05696	0.05696	
1	-0.07843	-0.07843	0.05784	0.05784	
2	-0.01842	-0.01842	0.05868	0.05868	
3	0.03947	0.03947	0.05947	0.05947	
4	0.02308	0.02308	0.06023	0.06023	


	seasonality_365.25_95
0	0.05696
1	0.05784
2	0.05868
3	0.05947
4	0.06023

Tips: there is an additional arg `seasonality_fs_order` to control the number of orders in `fourier series terms` we want to approximate the seasonality. In general, they cannot violate the condition that $2 \times \text{fourier series order} < \text{seasonality}$ since each order represents adding a pair of sine and cosine regressors.


12.5 More Diagnostic and Visualization

Here are a few more diagnostic and visualization. The fit is decomposed into components, the local trend and both periods of seasonality.

```
[16]: _ = plot_predicted_data(training_actual_df=train_df, predicted_df=predicted_df,
                           date_col=date_col, actual_col=response_col,
                           test_actual_df=test_df, markersize=10, lw=.5)
```



```
[17]: _ = plot_predicted_components(predicted_df=predicted_df, date_col=date_col, plot_
components=['trend', 'seasonality_7', 'seasonality_365.25'])
```



12.6 References

1. Ng, Wang and Dai (2021) Bayesian Time Varying Coefficient Model with Applications to Marketing Mix Modeling, arXiv preprint arXiv:2106.03322
2. Hastie, Trevor and Tibshirani, Robert. (1990), Generalized Additive Models, New York: Chapman and Hall.
3. Wood, S. N. (2006), Generalized Additive Models: an introduction with R, Boca Raton: Chapman & Hall/CRC
4. Harvey, C. A. (1989). Forecasting, Structural Time Series and the Kalman Filter, Cambridge University Press.
5. Durbin, J., Koopman, S. J.. (2001). Time Series Analysis by State Space Methods, Oxford Statistical Science Series

KERNEL-BASED TIME-VARYING REGRESSION - PART II

The previous tutorial covered the basic syntax and structure of **KTR** (or so called **BTVC**); time-series data was fitted with a KTR model accounting for trend and seasonality. In this tutorial a KTR model is fit with trend, seasonality, and additional regressors. To summarize part 1, **KTR** considers a time-series as an additive combination of local-trend, seasonality, and additional regressors. The coefficients for all three components are allowed to vary over time. The time-varying of the coefficients is modeled using kernel smoothing of latent variables. This can also be an advantage of picking this model over other static regression coefficients models.

This tutorial covers:

1. KTR model structure with regression
2. syntax to initialize, fit and predict a model with regressors
3. visualization of regression coefficients

```
[1]: import pandas as pd
import numpy as np
from math import pi
import matplotlib.pyplot as plt

import orbit
from orbit.models import KTR
from orbit.diagnostics.plot import plot_predicted_components
from orbit.utils.plot import get_orbit_style
from orbit.constants.palette import OrbitPalette

%matplotlib inline
pd.set_option('display.float_format', lambda x: '%.5f' % x)
orbit_style = get_orbit_style()
plt.style.use(orbit_style);
```

```
/Users/towinazure/opt/miniconda3/envs/jupyter/lib/python3.9/site-packages/tqdm/auto.py:
↳22: TqdmWarning: IProgress not found. Please update jupyter and ipywidgets. See https://
↳/ipywidgets.readthedocs.io/en/stable/user_install.html
from .autonotebook import tqdm as notebook_tqdm
```

```
[2]: print(orbit.__version__)
```

```
1.1.3dev
```

13.1 Model Structure

This section gives the mathematical structure of the KTR model. In short, it considers a time-series (y_t) as the linear combination of three parts. These are the local-trend (l_t), seasonality (s_t), and regression (r_t) terms at time t . That is

$$y_t = l_t + s_t + r_t + \epsilon_t, \quad t = 1, \dots, T,$$

where

- ϵ_t s comprise a stationary random error process.
- r_t is the regression component which can be further expressed as $\sum_{i=1}^J x_{i,t} \beta_{i,t}$ with covariate x and coefficient β on indexes i, t

For details of how on l_t and s_t , please refer to **Part I**.

Recall in **KTR**, we express coefficients as

$$B = Kb^T$$

where - *coefficient matrix* B has size $t \times P$ with rows equal to the β_t - *knot matrix* b with size $P \times J$; each entry is a latent variable $b_{p,j}$. The b_j can be viewed as the “knots” from the perspective of spline regression and j is a time index such that $t_j \in [1, \dots, T]$. - *kernel matrix* K with size $T \times J$ where the i th row and j th element can be viewed as the normalized weight $k(t_j, t) / \sum_{j=1}^J k(t_j, t)$

In regression, we generate the matrix K with Gaussian kernel k_{reg} as such:

$$k_{\text{reg}}(t, t_j; \rho) = \exp\left(-\frac{(t-t_j)^2}{2\rho^2}\right),$$

where ρ is the scale hyper-parameter.

13.2 Data Simulation Module

In this example, we will use simulated data in order to have true regression coefficients for comparison. We propose two set of simulation data with three predictors each:

The two data sets are: - random walk - sine-cosine like

Note the data are random so it may be worthwhile to repeat the next few sets a few times to see how different data sets work.

13.2.1 Random Walk Simulated Dataset

```
[3]: def sim_data_seasonal(n, RS):
      """ coefficients curve are sine-cosine like
      """
      np.random.seed(RS)
      # make the time varying coeffs
      tau = np.arange(1, n+1)/n
      data = pd.DataFrame({
          'tau': tau,
          'date': pd.date_range(start='1/1/2018', periods=n),
          'beta1': 2 * tau,
          'beta2': 1.01 + np.sin(2*pi*tau),
```

(continues on next page)

(continued from previous page)

```

    'beta3': 1.01 + np.sin(4*pi*(tau-1/8)),
    'x1': np.random.normal(0, 10, size=n),
    'x2': np.random.normal(0, 10, size=n),
    'x3': np.random.normal(0, 10, size=n),
    'trend': np.cumsum(np.concatenate((np.array([1]), np.random.normal(0, 0.1, n-
→1))))),
    'error': np.random.normal(0, 1, size=n) #stats.t.rvs(30, size=n),#
    })

    data['y'] = data.x1 * data.beta1 + data.x2 * data.beta2 + data.x3 * data.beta3 +
→data.error
    return data

```

```

[4]: def sim_data_rw(n, RS, p=3):
    """ coefficients curve are random walk like
    """
    np.random.seed(RS)

    # initializing coefficients at zeros, simulate all coefficient values
    lev = np.cumsum(np.concatenate((np.array([5.0]), np.random.normal(0, 0.01, n-1))))
    beta = np.concatenate(
        [np.random.uniform(0.05, 0.12, size=(1,p)),
         np.random.normal(0.0, 0.01, size=(n-1,p))],
         axis=0)
    beta = np.cumsum(beta, 0)

    # simulate regressors
    covariates = np.random.normal(0, 10, (n, p))

    # observation with noise
    y = lev + (covariates * beta).sum(-1) + 0.3 * np.random.normal(0, 1, n)

    regressor_col = ['x{}'.format(pp) for pp in range(1, p+1)]
    data = pd.DataFrame(covariates, columns=regressor_col)
    beta_col = ['beta{}'.format(pp) for pp in range(1, p+1)]
    beta_data = pd.DataFrame(beta, columns=beta_col)
    data = pd.concat([data, beta_data], axis=1)

    data['y'] = y
    data['date'] = pd.date_range(start='1/1/2018', periods=len(y))

    return data

```

```

[5]: rw_data = sim_data_rw(n=300, RS=2021, p=3)
    rw_data.head(10)

```

```

[5]:

```

	x1	x2	x3	beta1	beta2	beta3	y	date
0	14.02970	-2.55469	4.93759	0.07288	0.06251	0.09662	6.11704	2018-01-01
1	6.23970	0.57014	-6.99700	0.06669	0.05440	0.10476	5.35784	2018-01-02
2	9.91810	-6.68728	-3.68957	0.06755	0.04487	0.11624	4.82567	2018-01-03
3	-1.17724	8.88090	-16.02765	0.05849	0.04305	0.12294	3.63605	2018-01-04
4	11.61065	1.95306	0.19901	0.06604	0.03281	0.11897	5.85913	2018-01-05

(continues on next page)

(continued from previous page)

```

5  7.31929  3.36017 -6.09933 0.07825 0.03448 0.10836 5.08805 2018-01-06
6  0.53405  8.80412 -1.83692 0.07467 0.01847 0.10507 4.59303 2018-01-07
7 -16.03947 0.27562 -22.00964 0.06887 0.00865 0.10749 1.26651 2018-01-08
8 -17.72238 2.65195  0.22571 0.07007 0.01008 0.10432 4.10629 2018-01-09
9  -7.39895 -7.63162  3.25535 0.07715 0.01498 0.09356 4.30788 2018-01-10

```

13.2.2 Sine-Cosine Like Simulated Dataset

```
[6]: sc_data = sim_data_seasonal(n=80, RS=2021)
      sc_data.head(10)
```

```
[6]:
      tau      date  beta1  beta2  beta3      x1      x2      x3  \
0  0.01250 2018-01-01 0.02500 1.08846 0.02231 14.88609  1.56556 -14.69399
1  0.02500 2018-01-02 0.05000 1.16643 0.05894  6.76011 -0.56861  4.93157
2  0.03750 2018-01-03 0.07500 1.24345 0.11899 -4.18451 -5.38234 -13.90578
3  0.05000 2018-01-04 0.10000 1.31902 0.20098 -8.06521  9.01387 -0.75244
4  0.06250 2018-01-05 0.12500 1.39268 0.30289  5.55876  2.24944 -2.53510
5  0.07500 2018-01-06 0.15000 1.46399 0.42221 -7.05504 12.77788 14.25841
6  0.08750 2018-01-07 0.17500 1.53250 0.55601 11.30858  6.29269  7.82098
7  0.10000 2018-01-08 0.20000 1.59779 0.70098  6.45002  3.61891 16.28098
8  0.11250 2018-01-09 0.22500 1.65945 0.85357  1.06414 36.38726  8.80457
9  0.12500 2018-01-10 0.25000 1.71711 1.01000  4.22155 -12.01221  8.43176

      trend  error      y
0  1.00000 -0.73476  1.01359
1  1.07746 -0.97007 -1.00463
2  1.19201 -0.13891 -8.80009
3  1.22883  0.66550 11.59721
4  1.31341 -1.58259  1.47715
5  1.25911 -0.98049 22.68806
6  1.23484 -0.53751 15.43357
7  1.13237 -1.32858 17.15636
8  1.02834  0.87859 69.01607
9  1.00649 -0.22055 -11.27534

```

13.3 Fitting a Model with Regressors

The metadata for simulated data sets.

```
[7]: # num of predictors
      p = 3
      regressor_col = ['x{}'.format(pp) for pp in range(1, p + 1)]
      response_col = 'y'
      date_col='date'
```

As in **Part I** KTR follows sklearn model API style. First an instance of the Orbit class KTR is created. Second fit and predict methods are called for that instance. Besides providing meta data such `response_col`, `date_col` and `regressor_col`, there are additional args to provide to specify the estimator and the setting of the estimator. For details, please refer to other tutorials of the **Orbit** site.

```
[8]: ktr = KTR(
      response_col=response_col,
      date_col=date_col,
      regressor_col=regressor_col,
      prediction_percentiles=[2.5, 97.5],
      seed=2021,
      estimator='pyro-svi',
    )
```

Here `predict` has the additional argument `decompose=True`. This returns the components (l_t , s_t , and r_t) of the regression along with the prediction.

```
[9]: ktr.fit(df=rw_data)
      ktr.predict(df=rw_data, decompose=True).head(5)
```

```
INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
INFO:orbit:Using SVI (Pyro) with steps: 301, samples: 100, learning rate: 0.1, learning_
↪rate_total_decay: 1.0 and particles: 100.
INFO:root:Gussed max_plate_nesting = 1
INFO:orbit:step    0 loss = 3107.7, scale = 0.091353
INFO:orbit:step   100 loss = 302.15, scale = 0.047045
INFO:orbit:step   200 loss = 311.94, scale = 0.051746
INFO:orbit:step   300 loss = 311.42, scale = 0.053282
```

```
[9]:
```

	date	prediction_2.5	prediction	prediction_97.5	trend_2.5	trend \
0	2018-01-01	4.88856	6.00374	7.12610	4.08780	5.22386
1	2018-01-02	3.16076	4.19009	5.16861	4.17244	5.18190
2	2018-01-03	3.52741	4.78682	6.02640	4.14188	5.28647
3	2018-01-04	1.78070	2.92761	4.04994	4.05188	5.18402
4	2018-01-05	4.35078	5.28555	6.20752	4.21571	5.18052

	trend_97.5	regression_2.5	regression	regression_97.5
0	6.44394	0.47508	0.80353	1.08875
1	6.21267	-1.20094	-1.00036	-0.73915
2	6.58336	-0.80359	-0.54748	-0.16197
3	6.32235	-2.66193	-2.23109	-1.72535
4	6.14344	-0.10424	0.11233	0.33019

13.4 Visualization of Regression Coefficient Curves

The function `get_regression_coefs` to extract coefficients (they will have central credibility intervals if the argument `include_ci=True` is used).

```
[10]: coef_mid, coef_lower, coef_upper = ktr.get_regression_coefs(include_ci=True)
```

```
[11]: coef_mid.head(5)
```

```
[11]:
```


	date	x1	x2	x3
0	2018-01-01	0.00383	0.00787	0.15039
1	2018-01-02	0.00393	0.00795	0.14729
2	2018-01-03	0.00403	0.00802	0.14417
3	2018-01-04	0.00413	0.00808	0.14102
4	2018-01-05	0.00423	0.00814	0.13785

Because this is simulated data it is possible to overlay the estimate with the true coefficients.

```
[12]: fig, axes = plt.subplots(p, 1, figsize=(12, 12), sharex=True)


x = np.arange(coef_mid.shape[0])
for idx in range(p):
    axes[idx].plot(x, coef_mid['x{}'.format(idx + 1)], label='est' if idx == 0 else "",
                  color=OrbitPalette.BLUE.value)
    axes[idx].fill_between(x, coef_lower['x{}'.format(idx + 1)], coef_upper['x{}'.
    format(idx + 1)], alpha=0.2, color=OrbitPalette.BLUE.value)
    axes[idx].scatter(x, rw_data['beta{}'.format(idx + 1)], label='truth' if idx == 0
                    else "", s=10, alpha=0.6, color=OrbitPalette.BLACK.value)
    axes[idx].set_title('beta{}'.format(idx + 1))

fig.legend(bbox_to_anchor = (1,0.5));
```



To plot coefficients use the function `plot_regression_coefs` from the `KTR` class.

```
[13]: ktr.plot_regression_coefs(figsize=(10, 5), include_ci=True);
```



These type of time-varying coefficients detection problems are not new. Bayesian approach such as the R packages Bayesian Structural Time Series (a.k.a **BSTS**) by Scott and Varian (2014) and **tvReg** Isabel Casas and Ruben Fernandez-Casal (2021). Other frequentist approach such as Wu and Chiang (2000).

For further studies on benchmarking coefficients detection, Ng, Wang and Dai (2021) provides a detailed comparison of **KTR** with other popular time-varying coefficients methods; **KTR** demonstrates superior performance in the random walk data simulation.

13.5 Customizing Priors and Number of Knot Segments

To demonstrate how to specify the number of knots and priors consider the sine-cosine like simulated dataset. In this dataset, the fitting is more tricky since there could be some better way to define the number and position of the knots. There are obvious “change points” within the sine-cosine like curves. In **KTR** there are a few arguments that can leveraged to assign a priori knot attributes:

1. `regressor_init_knot_loc` is used to define the prior mean of the knot value. e.g. in this case, there is not a lot of prior knowledge so zeros are used.
2. The `regressor_init_knot_scale` and `regressor_knot_scale` are used to tune the prior sd of the global mean of the knot and the sd of each knot from the global mean respectively. These create a plausible range for the knot values.
3. The `regression_segments` defines the number of between knot segments (the number of knots - 1). The higher the number of segments the more change points are possible.

```
[14]: ktr = KTR(
  response_col=response_col,
  date_col=date_col,

  regressor_col=regressor_col,
  regressor_init_knot_loc=[0] * len(regressor_col),
  regressor_init_knot_scale=[10.0] * len(regressor_col),
```

(continues on next page)

(continued from previous page)

```

regressor_knot_scale=[2.0] * len(regressor_col),
regression_segments=6,

prediction_percentiles=[2.5, 97.5],
seed=2021,
estimator='pyro-svi',
)
ktr.fit(df=sc_data)

```

```

INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
INFO:orbit:Using SVI (Pyro) with steps: 301, samples: 100, learning_
↪rate_total_decay: 1.0 and particles: 100.
INFO:root:Guessed max_plate_nesting = 1
INFO:orbit:step    0 loss = 828.02, scale = 0.10882
INFO:orbit:step   100 loss = 340.58, scale = 0.87797
INFO:orbit:step   200 loss = 266.67, scale = 0.37411
INFO:orbit:step   300 loss = 261.21, scale = 0.43775

```

```
[14]: <orbit.forecaster.svi.SVIForecaster at 0x29c316460>
```


```

[15]: coef_mid, coef_lower, coef_upper = ktr.get_regression_coefs(include_ci=True)
fig, axes = plt.subplots(p, 1, figsize=(12, 12), sharex=True)

x = np.arange(coef_mid.shape[0])
for idx in range(p):
    axes[idx].plot(x, coef_mid['x{}'.format(idx + 1)], label='est' if idx == 0 else "",
↪color=OrbitPalette.BLUE.value)
    axes[idx].fill_between(x, coef_lower['x{}'.format(idx + 1)], coef_upper['x{}'.
↪format(idx + 1)], alpha=0.2, color=OrbitPalette.BLUE.value)
    axes[idx].scatter(x, sc_data['beta{}'.format(idx + 1)], label='truth' if idx == 0
↪else "", s=10, alpha=0.6, color=OrbitPalette.BLACK.value)
    axes[idx].set_title('beta{}'.format(idx + 1))


fig.legend(bbox_to_anchor = (1, 0.5));

```



Visualize the knots using the `plot_regression_coefs` function with `with_knot=True`.

```
[16]: ktr.plot_regression_coefs(with_knot=True, figsize=(10, 5), include_ci=True);
```



There are more ways to define knots for regression as well as seasonality and trend (a.k.a levels). These are described in **Part III**

13.6 References

1. Ng, Wang and Dai (2021). Bayesian Time Varying Coefficient Model with Applications to Marketing Mix Modeling, arXiv preprint arXiv:2106.03322
2. Isabel Casas and Ruben Fernandez-Casal (2021). tvReg: Time-Varying Coefficients Linear Regression for Single and Multi-Equations. <https://CRAN.R-project.org/package=tvReg> R package version 0.5.4.
3. Steven L Scott and Hal R Varian (2014). Predicting the present with bayesian structural time series. International Journal of Mathematical Modelling and Numerical Optimisation 5, 1-2 (2014), 4–23.

KERNEL-BASED TIME-VARYING REGRESSION - PART III

The tutorials **I** and **II** described the **KTR** model, its fitting procedure, visualizations and diagnostics / validation methods . This tutorial covers more **KTR** configurations for advanced users. In particular, it describes how to use knots to model change points in the seasonality and regression coefficients.

For more detail on this see Ng, Wang and Dai (2021)., which describes how **KTR** knots can be thought of as change points. This highlights a similarity between **KTR** and **Facebook's Prophet** package which introduces the change point detection on levels.

Part III covers different **KTR** arguments to specify knots position:

- level_segements
- level_knot_distance
- level_knot_dates

```
[1]: import pandas as pd
import numpy as np
from math import pi
import matplotlib.pyplot as plt

import orbit
from orbit.models import KTR
from orbit.diagnostics.plot import plot_predicted_data
from orbit.utils.plot import get_orbit_style
from orbit.utils.dataset import load_iclaims

%matplotlib inline
pd.set_option('display.float_format', lambda x: '%.5f' % x)
```

```
[2]: print(orbit.__version__)
```

```
1.1.3
```

14.1 Fitting with iClaims Data

The iClaims data set gives the weekly log number of claims and several regressors.

```
[3]: # without the enddate, we would get end date='2018-06-24' to make our tutorial consistent,
      ↪with the older version
df = load_iclaims(end_date='2020-11-29')

DATE_COL = 'week'
RESPONSE_COL = 'claims'

print(df.shape)
df.head()
```

```
(570, 7)
```

```
[3]:
```

	week	claims	trend.unemploy	trend.filling	trend.job	sp500	\
0	2010-01-03	13.38660	0.03493	-0.34414	0.12802	-0.53745	
1	2010-01-10	13.62422	0.03493	-0.22053	0.17932	-0.54529	
2	2010-01-17	13.39874	0.05119	-0.31817	0.12802	-0.58504	
3	2010-01-24	13.13755	0.01840	-0.22053	0.11744	-0.60156	
4	2010-01-31	13.19676	-0.05059	-0.26816	0.08501	-0.60874	

```

      vix
0 0.08456
1 0.07235
2 0.49424
3 0.39055
4 0.44931
```

14.1.1 Specifying Levels Segments

The first way to specify the knot locations and number is the `level_segments` argument. This gives the number of between knot segments; since there is a knot on each end of each the total number of knots would be the number of segments plus one. To illustrate that, try `level_segments=10` (line 5).


```
[4]: response_col = 'claims'
      date_col='week'
```

```
[5]: ktr = KTR(
      response_col=response_col,
      date_col=date_col,

      level_segments=10,
      prediction_percentiles=[2.5, 97.5],
      seed=2020,
      estimator='pyro-svi'
    )
```

```
[6]: ktr.fit(df=df)
      _ = ktr.plot_lev_knots()
```

```
INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
INFO:orbit:Using SVI (Pyro) with steps: 301, samples: 100, learning rate: 0.1, learning_
↪rate_total_decay: 1.0 and particles: 100.
INFO:root:Guessed max_plate_nesting = 1
INFO:orbit:step    0 loss = 176.47, scale = 0.083093
INFO:orbit:step   100 loss = 113.12, scale = 0.046384
INFO:orbit:step   200 loss = 113.16, scale = 0.04608
INFO:orbit:step   300 loss = 113.27, scale = 0.04625
```



Note that there are precisely there are 11 knots (triangles) evenly spaced in the above chart.

14.1.2 Specifying Knots Distance


An alternative way of specifying the number of knots is the `level_knot_distance` argument. This argument gives the distance between knots. It can be useful as number of knots grows with the length of the time-series. Note that if the total length of the time-series is not a multiple of `level_knot_distance` the first segment will have a different length. For example, in a weekly data, by putting `level_knot_distance=104` roughly means putting a knot once in two years.

```
[7]: ktr = KTR(
    response_col=response_col,
    date_col=date_col,
    level_knot_distance=104,

    # fit a weekly seasonality
    seasonality=52,
    # high order for sharp turns on each week
    seasonality_fs_order=12,
    prediction_percentiles=[2.5, 97.5],
    seed=2020,
    estimator='pyro-svi'
)
```

```
[8]: ktr.fit(df=df)
_ = ktr.plot_lev_knots()


INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
INFO:orbit:Using SVI (Pyro) with steps: 301, samples: 100, learning rate: 0.1, learning_
↳rate_total_decay: 1.0 and particles: 100.
INFO:root:Guessed max_plate_nesting = 1
INFO:orbit:step    0 loss = 145.63, scale = 0.088976
INFO:orbit:step   100 loss = -5.2717, scale = 0.036808
INFO:orbit:step   200 loss = -5.2295, scale = 0.037277
INFO:orbit:step   300 loss = -5.4426, scale = 0.037506
```



In the above chart, the knots are located about every 2-years.

To highlight the value of the next method of configuring knot position, consider the prediction for this model show below.

```
[9]: predicted_df = ktr.predict(df=df)
_ = plot_predicted_data(training_actual_df=df, predicted_df=predicted_df, prediction_
↳percentiles=[2.5, 97.5],
                    date_col=date_col, actual_col=response_col)
```



As the knots are placed evenly the model can not adequately describe the change point in early 2020. The model fit can potentially be improved by inserting knots around the sharp change points (e.g., 2020-03-15). This insertion can be done with the `level_knot_dates` argument described below.

14.1.3 Specifying Knots Dates

The `level_knot_dates` argument allows for the explicit placement of knots. It needs a string of dates; see line 4.

```
[10]: ktr = KTR(
    response_col=response_col,
    date_col=date_col,
    level_knot_dates = ['2010-01-03', '2020-03-15', '2020-03-22', '2020-11-29'],

    # fit a weekly seasonality
    seasonality=52,
    # high order for sharp turns on each week
    seasonality_fs_order=12,
    prediction_percentiles=[2.5, 97.5],
    seed=2020,
    estimator='pyro-svi'
)
```

```
[11]: ktr.fit(df=df)
INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
INFO:orbit:Using SVI (Pyro) with steps: 301, samples: 100, learning rate: 0.1, learning_
↪rate_total_decay: 1.0 and particles: 100.
INFO:root:Guessed max_plate_nesting = 1
INFO:orbit:step    0 loss = 99.358, scale = 0.096314
INFO:orbit:step  100 loss = -444.52, scale = 0.026774
```


(continues on next page)

(continued from previous page)


```
INFO:orbit:step 200 loss = -445.8, scale = 0.027848
INFO:orbit:step 300 loss = -441.83, scale = 0.029342
```

```
[11]: <orbit.forecaster.svi.SVIForecaster at 0x177d3c8e0>
```

```
[12]: _ = ktr.plot_lev_knots()
```



```
[13]: predicted_df = ktr.predict(df=df)
_ = plot_predicted_data(training_actual_df=df, predicted_df=predicted_df, prediction_
↳percentiles=[2.5, 97.5],
date_col=date_col, actual_col=response_col)
```



Note this fit is even better than the previous one using less knots. Of course, the case here is trivial because the pandemic

onset is treated as known. In other cases, there may not be an obvious way to find the optimal knots dates.

14.2 Conclusion

This tutorial demonstrates multiple ways to customize the knots location for levels. In **KTR**, there are similar arguments for seasonality and regression such as `seasonality_segments` and `regression_knot_dates` and `regression_segments`. Due to their similarities with their knots location equivalent arguments they are not demonstrated here. However it is encouraged for **KTR** users to explore them.

14.3 References

1. Ng, Wang and Dai (2021). Bayesian Time Varying Coefficient Model with Applications to Marketing Mix Modeling, arXiv preprint arXiv:2106.03322
2. Sean J Taylor and Benjamin Letham. 2018. Forecasting at scale. *The American Statistician* 72, 1 (2018), 37–45. Package version 0.7.1.

KERNEL-BASED TIME-VARYING REGRESSION - PART IV

This is final tutorial on **KTR**. It continues from **Part III** with additional details on some of the advanced arguments. For other details on **KTR** see either the previous three tutorials or the original paper Ng, Wang and Dai (2021).

In **Part IV** covers advance inputs for regression including

- regressors signs
- time-point coefficients priors

```
[1]: import pandas as pd
import numpy as np
from math import pi
import matplotlib.pyplot as plt

import orbit
from orbit.models import KTR
from orbit.diagnostics.plot import plot_predicted_components
from orbit.utils.plot import get_orbit_style
from orbit.utils.kernels import gauss_kernel
from orbit.constants.palette import OrbitPalette

%matplotlib inline
pd.set_option('display.float_format', lambda x: '%.5f' % x)
orbit_style = get_orbit_style()
plt.style.use(orbit_style);
```

```
[2]: print(orbit.__version__)
```

```
1.1.3
```

15.1 Data

To demonstrate the effect of specifying regressors coefficients sign, it is helpful to modify the data simulation code from **part II**. The simulation is altered to impose strictly positive regression coefficients.

In the **KTR** model below, the coefficient curves are approximated with **Gaussian kernels** having positive values of knots. The levels are also included in the process with vector of ones as the covariates.

The parameters used to setup the data simulation are:

- **n** : number of time steps
- **p** : number of predictors

```
[3]: np.random.seed(2021)

n = 300
p = 2
tp = np.arange(1, 301) / 300
knot_tp = np.array([1, 100, 200, 300]) / 300
beta_knot = np.array(
    [[1.0, 0.1, 0.15],
     [3.0, 0.01, 0.05],
     [3.0, 0.01, 0.05],
     [2.0, 0.05, 0.02]]
)

gk = gauss_kernel(tp, knot_tp, rho=0.2)
beta = np.matmul(gk, beta_knot)
covar_lev = np.ones((n, 1))
covar = np.concatenate((covar_lev, np.random.normal(0, 1.0, (n, p))), axis=1)\

# observation with noise
y = (covar * beta).sum(-1) + np.random.normal(0, 0.1, n)

regressor_col = ['x{}'.format(pp) for pp in range(1, p+1)]
data = pd.DataFrame(covar[:,1:], columns=regressor_col)
data['y'] = y
data['date'] = pd.date_range(start='1/1/2018', periods=len(y))
data = data[['date', 'y'] + regressor_col]
beta_col = ['beta{}'.format(pp) for pp in range(1, p+1)]
beta_data = pd.DataFrame(beta[:,1:], columns=beta_col)

data = pd.concat([data, beta_data], axis=1)
```

```
[4]: data.tail(10)
```

```
[4]:
```

	date	y	x1	x2	beta1	beta2
290	2018-10-18	2.15947	-0.62762	0.17840	0.04015	0.02739
291	2018-10-19	2.25871	-0.92975	0.81415	0.04036	0.02723
292	2018-10-20	2.18356	0.82438	-0.92705	0.04057	0.02707
293	2018-10-21	2.26948	1.57181	-0.78098	0.04077	0.02692
294	2018-10-22	2.26375	-1.07504	-0.86523	0.04097	0.02677
295	2018-10-23	2.21349	0.24637	-0.98398	0.04117	0.02663
296	2018-10-24	2.13297	-0.58716	0.59911	0.04136	0.02648
297	2018-10-25	2.00949	-2.01610	0.08618	0.04155	0.02634
298	2018-10-26	2.14302	0.33863	-0.37912	0.04173	0.02620
299	2018-10-27	2.10795	-0.96160	-0.42383	0.04192	0.02606

Just like previous tutorials in regression, some additional args are used to describe the regressors and the scale parameters for the knots.

```
[5]: ktr = KTR(
    response_col='y',
    date_col='date',
    regressor_col=regressor_col,
```

(continues on next page)

(continued from previous page)

```

regressor_init_knot_scale=[0.1] * p,
regressor_knot_scale=[0.1] * p,
prediction_percentiles=[2.5, 97.5],
seed=2021,
estimator='pyro-svi',
)
ktr.fit(df=data)

```

```

INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
INFO:orbit:Using SVI (Pyro) with steps: 301, samples: 100, learning rate: 0.1, learning_
↪rate_total_decay: 1.0 and particles: 100.
INFO:root:Gussed max_plate_nesting = 1
INFO:orbit:step    0 loss = -3.5592, scale = 0.085307
INFO:orbit:step   100 loss = -229.96, scale = 0.036306
INFO:orbit:step   200 loss = -227.35, scale = 0.038053
INFO:orbit:step   300 loss = -226.72, scale = 0.037578

```

```
[5]: <orbit.forecaster.svi.SVIForecaster at 0x14e790d90>
```

15.2 Visualization of Regression Coefficient Curves

The `get_regression_coefs` argument is used to extract coefficients with intervals by supplying the argument `include_ci=True`.

```
[6]: coef_mid, coef_lower, coef_upper = ktr.get_regression_coefs(include_ci=True)
```


The next figure shows the overlay of the estimate on the true coefficients. Since the lower bound is below zero some of the coefficient posterior samples are negative.

```
[7]: fig, axes = plt.subplots(p, 1, figsize=(12, 12), sharex=True)

x = np.arange(coef_mid.shape[0])
for idx in range(p):
    axes[idx].plot(x, coef_mid['x{}'.format(idx + 1)], label='est' if idx == 0 else "",
↪alpha=0.8, color=OrbitPalette.BLUE.value)
    axes[idx].fill_between(x, coef_lower['x{}'.format(idx + 1)], coef_upper['x{}'.
↪format(idx + 1)], alpha=0.15, color=OrbitPalette.BLUE.value)
    axes[idx].plot(x, data['beta{}'.format(idx + 1)], label='truth' if idx == 0 else "",
↪alpha=0.6, color = OrbitPalette.BLACK.value)
    axes[idx].set_title('beta{}'.format(idx + 1))

fig.legend(bbox_to_anchor = (1,0.5));

```



15.3 Regressor Sign

Strictly positive coefficients can be imposed by using the `regressor_sign` arg. It can have values “=”, “-”, or “+” which denote no restriction, strictly negative, strictly positive. Note that it is possible to have a mixture by providing a list of strings one for each regressor.

```
[8]: ktr = KTR(
    response_col='y',
    date_col='date',
    regressor_col=regressor_col,
    regressor_init_knot_scale=[0.1] * p,
    regressor_knot_scale=[0.1] * p,
    regressor_sign=['+'] * p,
    prediction_percentiles=[2.5, 97.5],
    seed=2021,
```

(continues on next page)

(continued from previous page)

```

    estimator='pyro-svi',
)
ktr.fit(df=data)

```

```

INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
INFO:orbit:Using SVI (Pyro) with steps: 301, samples: 100, learning_
↪rate_total_decay: 1.0 and particles: 100.
INFO:root:Gussed max_plate_nesting = 1
INFO:orbit:step    0 loss = 9.7375, scale = 0.10482
INFO:orbit:step   100 loss = -229.85, scale = 0.42108
INFO:orbit:step   200 loss = -230.68, scale = 0.42784
INFO:orbit:step   300 loss = -230.77, scale = 0.4182

```

```
[8]: <orbit.forecaster.svi.SVIForecaster at 0x14f045cd0>
```

```
[9]: coef_mid, coef_lower, coef_upper = ktr.get_regression_coefs(include_ci=True)
```


```
[10]: fig, axes = plt.subplots(p, 1, figsize=(12, 12), sharex=True)
```

```

x = np.arange(coef_mid.shape[0])
for idx in range(p):
    axes[idx].plot(x, coef_mid['x{}'.format(idx + 1)], label='est' if idx == 0 else "",
↪alpha=0.8, color=OrbitPalette.BLUE.value)
    axes[idx].fill_between(x, coef_lower['x{}'.format(idx + 1)], coef_upper['x{}'.
↪format(idx + 1)], alpha=0.15, color=OrbitPalette.BLUE.value)
    axes[idx].plot(x, data['beta{}'.format(idx + 1)], label='truth' if idx == 0 else "",
↪ alpha=0.6, color = OrbitPalette.BLACK.value)
    axes[idx].set_title('beta{}'.format(idx + 1))

fig.legend(bbox_to_anchor = (1,0.5));

```



Observe the curves lie in the positive range with a slightly improved fit relative to the last model.

To conclude, it is useful to have a strictly positive range of regression coefficients if that range is known a priori. **KTR** allows these priors to be specified. For regression scenarios where there is no a priori knowledge of the coefficient sign it is recommended to use the default which contains both sides of the range.

15.4 Time-point coefficient priors

Users can incorporate coefficient priors for any regressor and any time period. This feature is quite useful when users have some prior knowledge or beliefs on regressor coefficients. For example, if an A/B test is conducted for a certain regressor over a specific time range, then users can ingest the priors derived from such A/B test.

This can be done by supplying a list of dictionaries via `coef_prior_list`. Each dict in the list should have keys as `name`, `prior_start_tp_idx` (inclusive), `prior_end_tp_idx` (not inclusive), `prior_mean`, `prior_sd`, and `prior_regressor_col`.

Below is an illustrative example by using the simulated data above.

```
[11]: from copy import deepcopy
```

```
[12]: prior_duration = 50
coef_list_dict = []
prior_idx=[
    np.arange(150, 150 + prior_duration),
    np.arange(200, 200 + prior_duration),
]
regressor_idx = range(1, p + 1)
plot_dict = {}
for i in regressor_idx:
    plot_dict[i] = {'idx': [], 'val': []}
```

```
[13]: for idx, idx2, regressor in zip(prior_idx, regressor_idx, regressor_col):
    prior_dict = {}
    prior_dict['name'] = f'prior_{regressor}'
    prior_dict['prior_start_tp_idx'] = idx[0]
    prior_dict['prior_end_tp_idx'] = idx[-1] + 1
    prior_dict['prior_mean'] = beta[idx, idx2]
    prior_dict['prior_sd'] = [0.1] * len(idx)
    prior_dict['prior_regressor_col'] = [regressor] * len(idx)

    plot_dict[idx2]['idx'].extend(idx)
    plot_dict[idx2]['val'].extend(beta[idx, idx2])

    coef_list_dict.append(deepcopy(prior_dict))
```

```
[14]: ktr = KTR(
    response_col='y',
    date_col='date',
    regressor_col=regressor_col,
    regressor_init_knot_scale=[0.1] * p,
    regressor_knot_scale=[0.1] * p,
    regressor_sign=['+'] * p,
    coef_prior_list = coef_list_dict,
    prediction_percentiles=[2.5, 97.5],
    seed=2021,
    estimator='pyro-svi',
)
ktr.fit(df=data)
```

```
INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
INFO:orbit:Using SVI (Pyro) with steps: 301, samples: 100, learning rate: 0.1, learning_
↪rate_total_decay: 1.0 and particles: 100.
INFO:root:Guessed max_plate_nesting = 1
INFO:orbit:step    0 loss = -5741.9, scale = 0.094521
INFO:orbit:step   100 loss = -7140.2, scale = 0.31433
INFO:orbit:step   200 loss = -7141.1, scale = 0.31618
INFO:orbit:step   300 loss = -7141.6, scale = 0.33042
```

```
[14]: <orbit.forecaster.svi.SVIForecaster at 0x14efcf130>
```


```
[15]: coef_mid, coef_lower, coef_upper = ktr.get_regression_coefs(include_ci=True)
```

```
[16]: fig, axes = plt.subplots(p, 1, figsize=(10, 8), sharex=True)
```

```
x = np.arange(coef_mid.shape[0])
for idx in range(p):
    axes[idx].plot(x, coef_mid['x{}'.format(idx + 1)], label='est', alpha=0.8,
    ↪ color=OrbitPalette.BLUE.value)
    axes[idx].fill_between(x, coef_lower['x{}'.format(idx + 1)], coef_upper['x{}'.
    ↪ format(idx + 1)], alpha=0.15, color=OrbitPalette.BLUE.value)
    axes[idx].plot(x, data['beta{}'.format(idx + 1)], label='truth', alpha=0.6, color =
    ↪ OrbitPalette.BLACK.value)
    axes[idx].set_title('beta{}'.format(idx + 1))
    axes[idx].scatter(plot_dict[idx + 1]['idx'], plot_dict[idx + 1]['val'],
    ↪ s=5, color=OrbitPalette.RED.value, alpha=.6, label='ingested priors')

handles, labels = axes[0].get_legend_handles_labels()
fig.legend(handles, labels, loc='upper center', ncol=3, bbox_to_anchor=(.5, 1.05))

plt.tight_layout()
```



As seen above, for the ingested prior time window, the estimation is aligned better with the truth and the resulting confidence interval also becomes narrower compared to other periods.

15.5 References

1. Ng, Wang and Dai (2021). Bayesian Time Varying Coefficient Model with Applications to Marketing Mix Modeling, arXiv preprint arXiv:2106.03322

PREDICTION DECOMPOSITION

In this section, we will demonstrate how to visualize

- time series forecasting
- predicted components

by using the plotting utilities that come with the Orbit package.

```
[1]: %matplotlib inline

import pandas as pd
import numpy as np
import matplotlib.pyplot as plt

import orbit
from orbit.models import DLT
from orbit.diagnostics.plot import plot_predicted_data, plot_predicted_components
from orbit.utils.dataset import load_iclaims

import warnings
warnings.filterwarnings('ignore')
```

```
[2]: print(orbit.__version__)

1.1.3
```

```
[3]: # load log-transformed data
df = load_iclaims()
train_df = df[df['week'] < '2017-01-01']
test_df = df[df['week'] >= '2017-01-01']

response_col = 'claims'
date_col = 'week'
regressor_col = ['trend.unemploy', 'trend.filling', 'trend.job']
```

16.1 Fit a model

Here we use the DLTFull model as example.

```
[4]: dlt = DLT(
      response_col=response_col,
      regressor_col=regressor_col,
      date_col=date_col,
      seasonality=52,
      prediction_percentiles=[5, 95],
    )

dlt.fit(train_df)

INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 1.000, warmups (per_
↳chain): 225 and samples(per chain): 25.
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests_
↳for n_eff and Rhat.
To run all diagnostics call pystan.check_hmc_diagnostics(fit)


[4]: <orbit.forecaster.full_bayes.FullBayesianForecaster at 0x16bc45220>
```

16.2 Plot Predictions

First, we do the prediction on the training data before the year 2017.

```
[5]: predicted_df = dlt.predict(df=train_df, decompose=True)


_ = plot_predicted_data(train_df, predicted_df,
                       date_col=dlt.date_col, actual_col=dlt.response_col)
```



Next, we do the predictions on the test data after the year 2017. This plot is useful to help check the overall model performance on the out-of-sample period.

```
[6]: predicted_df = dlt.predict(df=test_df, decompose=True)

_ = plot_predicted_data(training_actual_df=train_df, predicted_df=predicted_df,
                        date_col=dlt.date_col, actual_col=dlt.response_col,
                        test_actual_df=test_df)
```




16.3 Plot Predicted Components

`plot_predicted_components` is the utility to plot each component separately. This is useful when one wants to look into the model prediction results and inspect each component separately.


```
[7]: predicted_df = dlt.predict(df=train_df, decompose=True)

_ = plot_predicted_components(predicted_df, date_col)
```



One can use `plot_components` to have more components to be plotted if they are available in the supplied `predicted_df`.

```
[8]: _ = plot_predicted_components(predicted_df, date_col,
    plot_components=['prediction', 'trend', 'seasonality',
    ↪ 'regression'])
```



MODEL DIAGNOSTICS

In this section, we introduce to a few recommended diagnostic plots to diagnostic Orbit models. The posterior samples in **SVI** and **Full Bayesian** i.e. `FullBayesianForecaster` and `SVIForecaster`.

The plots are created by `arviz` for the plots. **ArviZ** is a Python package for exploratory analysis of Bayesian models, includes functions for posterior analysis, data storage, model checking, comparison and diagnostics.

- Trace plot
- Pair plot
- Density plot

```
[1]: import pandas as pd
import numpy as np
import matplotlib.pyplot as plt
import arviz as az
import seaborn as sns

%matplotlib inline

import orbit
from orbit.models import LGT, DLT
from orbit.utils.dataset import load_iclaims

import warnings
warnings.filterwarnings('ignore')

from orbit.diagnostics.plot import params_comparison_boxplot
from orbit.constants import palette
```

```
[2]: print(orbit.__version__)

1.1.3
```

17.1 Load data

```
[3]: df = load_iclaims()
df.dtypes
```

```
[3]: week          datetime64[ns]
claims           float64
trend.unemploy   float64
trend.filling    float64
trend.job        float64
sp500            float64
vix              float64
dtype: object
```

```
[4]: df.head(5)
```

```
[4]:
```

	week	claims	trend.unemploy	trend.filling	trend.job	sp500	vix
0	2010-01-03	13.386595	0.219882	-0.318452	0.117500	-0.417633	0.122654
1	2010-01-10	13.624218	0.219882	-0.194838	0.168794	-0.425480	0.110445
2	2010-01-17	13.398741	0.236143	-0.292477	0.117500	-0.465229	0.532339
3	2010-01-24	13.137549	0.203353	-0.194838	0.106918	-0.481751	0.428645
4	2010-01-31	13.196760	0.134360	-0.242466	0.074483	-0.488929	0.487404

17.2 Fit a Model

```
[5]: DATE_COL = 'week'
RESPONSE_COL = 'claims'
REGRESSOR_COL = ['trend.unemploy', 'trend.filling', 'trend.job']
```

```
[6]: dlt = DLT(response_col=RESPONSE_COL,
              date_col=DATE_COL,
              regressor_col=REGRESSOR_COL,
              seasonality=52,
              num_warmup=2000,
              num_sample=2000,
              chains=4)
```

```
[7]: dlt.fit(df=df)
```

```
INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 1.000, warmups (per_
↪chain): 500 and samples(per chain): 500.
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests_
↪for n_eff and Rhat.
To run all diagnostics call pystan.check_hmc_diagnostics(fit)
```

```
[7]: <orbit.forecaster.full_bayes.FullBayesianForecaster at 0x17961e670>
```

We can use `.get_posterior_samples()` to extract posteriors. Note that we need `permute=False` to retrieve additional information such as chains when we extract posterior samples for posteriors plotting. For regression, in order to collapse and relabel regression from parameters (usually called as `beta`), we use `relabel=True`.

```
[8]: ps = dlt.get_posterior_samples(relabel=True, permute=False)
      ps.keys()
```

```
[8]: odict_keys(['l', 'b', 'lev_sm', 'slp_sm', 'obs_sigma', 'nu', 'lt_sum', 's', 'sea_sm',
               ↪ 'gt_sum', 'gb', 'gl', 'trend.unemploy', 'trend.filling', 'trend.job'])
```


17.3 Diagnostics Visualization

In the following section, we are going to use the regression coefficients as an example. In practice, you could check other parameters extracted from the model. For now, it only supports 1-D parameter which in generally capture the most important parameters of the model (e.g. `obs_sigma`, `lev_sm` etc.)

17.3.1 Convergence Status

Trace plots help us verify the convergence of model. In general, a largely overlapped distribution across samples from different chains indicates the convergence.

```
[9]: az.style.use('arviz-darkgrid')
      az.plot_trace(
          ps,
          var_names=['trend.unemploy', 'trend.filling', 'trend.job'],
          chain_prop={"color": ['r', 'b', 'g', 'y']},
          figsize=(10, 8),
      );
```




Note that this is only applicable for FullBayesianForecaster using sampling method such as MCMC.

17.3.2 Samples Density

We can also check the density of samples by pair plot.

```
[10]: az.plot_pair(
    ps,
    var_names=['trend.unemploy', 'trend.filling', 'trend.job'],
    kind=["scatter", "kde"],
    marginals=True,
    point_estimate="median",
    textsize=18.5,
);
```



17.3.3 Compare Models

You can also compare posteriors across different models with the same parameters. You can use plots such as density plot and forest plot to do so.


```
[11]: dlt_smaller_prior = DLT(
    response_col=RESPONSE_COL,
    date_col=DATE_COL,
    regressor_col=REGRESSOR_COL,
    regressor_sigma_prior=[0.05, 0.05, 0.05],
    seasonality=52,
    num_warmup=2000,
    num_sample=2000,
    chains=4
)
dlt_smaller_prior.fit(df=df)
ps_smaller_prior = dlt_smaller_prior.get_posterior_samples(relabel=True, permute=False)
```

```
INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 1.000, warmups (per
↳ chain): 500 and samples(per chain): 500.
```


```
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests.
↳ for n_eff and Rhat.
```

```
To run all diagnostics call pystan.check_hmc_diagnostics(fit)
```


```
[12]: az.plot_density(  
    [ps, ps_smaller_prior],  
    var_names=['trend.unemploy', 'trend.filling', 'trend.job'],  
    data_labels=["Default", "Smaller Sigma"],  
    shade=0.1,  
    textsize=18.5,  
);
```



```
[13]: az.plot_forest(  
    [ps, ps_smaller_prior],  
    var_names=['trend.unemploy'],  
    model_names=["Default", "Smaller Sigma"],  
);
```



```
[14]: params_comparison_boxplot(
    [ps, ps_smaller_prior],
    var_names=['trend.unemploy', 'trend.filling', 'trend.job'],
    model_names=["Default", "Smaller Sigma"],
    box_width = .1, box_distance=0.1,
    showfliers=True
);
```



17.4 Conclusion

Orbit models allow multiple visualization to diagnostics models and compare different models. We briefly introduce some basic syntax and usage of `arviz`. There is an [example gallery](#) built by the original team. Users can learn the details and more advance usage there. Meanwhile, the Orbit team aims to continue expand the scope to leverage more work done from the `arviz` project.

BACKTEST

This section will cover following topics:

- How to create a TimeSeriesSplitter
- How to create a BackTester and retrieve the backtesting results
- How to leverage the backtesting to tune the hyper-parameters for orbit models

```
[1]: %matplotlib inline

import pandas as pd
import numpy as np
import matplotlib.pyplot as plt

import orbit
from orbit.models import LGT, DLT
from orbit.diagnostics.backtest import BackTester, TimeSeriesSplitter
from orbit.diagnostics.plot import plot_bt_predictions
from orbit.diagnostics.metrics import smape, wmape
from orbit.utils.dataset import load_iclaims

import warnings
warnings.filterwarnings('ignore')
```

```
[2]: print(orbit.__version__)

1.1.3
```

```
[3]: # load log-transformed data
data = load_iclaims()
```

```
[4]: data.shape
```

```
[4]: (443, 7)
```

The way to gauge the performance of a time-series model is through re-training models with different historic periods and check their forecast within certain steps. This is similar to a time-based style cross-validation. More often, it is called `backtest` in time-series modeling.

The purpose of this notebook is to illustrate how to `backtest` a single model using `BackTester`

`BackTester` will compose a `TimeSeriesSplitter` within it, but `TimeSeriesSplitter` is useful as a standalone, in case there are other tasks to perform that requires splitting but not backtesting. `TimeSeriesSplitter` implemented

each 'slices' as generator, i.e it can be used in a for loop. You can also retrieve the composed `TimeSeriesSplitter` object from `BackTester` to utilize the additional methods in `TimeSeriesSplitter`

Currently, there are two schemes supported for the back-testing engine: expanding window and rolling window.

- **expanding window**: for each back-testing model training, the train start date is fixed, while the train end date is extended forward.
- **rolling window**: for each back-testing model training, the training window length is fixed but the window is moving forward.

18.1 Create a TimeSeriesSplitter

There two main way to splitting a time series: expanding and rolling. Expanding window has a fixed starting point, and the window length grows as users move forward in time series. It is useful when users want to incorporate all historical information. On the other hand, rolling window has a fixed window length, and the starting point of the window moves forward as users move forward in time series. Below section illustrates how users can use `TimeSeriesSplitter` to split the claims time series.

18.1.1 Expanding window

```
[5]: # configs
min_train_len = 380 # minimal length of window length
forecast_len = 20 # length forecast window
incremental_len = 20 # step length for moving forward
```

```
[6]: ex_splitter = TimeSeriesSplitter(df=data,
                                     min_train_len=min_train_len,
                                     incremental_len=incremental_len,
                                     forecast_len=forecast_len,
                                     window_type='expanding',
                                     date_col='week')
```

```
[7]: print(ex_splitter)
```


```
----- Fold: (1 / 3)-----
Train start date: 2010-01-03 00:00:00 Train end date: 2017-04-09 00:00:00
Test start date: 2017-04-16 00:00:00 Test end date: 2017-08-27 00:00:00

----- Fold: (2 / 3)-----
Train start date: 2010-01-03 00:00:00 Train end date: 2017-08-27 00:00:00
Test start date: 2017-09-03 00:00:00 Test end date: 2018-01-14 00:00:00

----- Fold: (3 / 3)-----
Train start date: 2010-01-03 00:00:00 Train end date: 2018-01-14 00:00:00
Test start date: 2018-01-21 00:00:00 Test end date: 2018-06-03 00:00:00
```


Users can visualize the splits using the internal `plot()` function. One may notice that the last few data points may not be included in the last split, which is expected when `min_train_len` is specified.

```
[8]: _ = ex_splitter.plot()
```



If users want to visualize the scheme in terms of indices, one can do the following.

```
[9]: _ = ex_splitter.plot(show_index=True)
```




18.1.2 Rolling window

```
[10]: # configs
min_train_len = 380 # in case of rolling window, this specify the length of window length
forecast_len = 20 # length forecast window
incremental_len = 20 # step length for moving forward
```

```
[11]: roll_splitter = TimeSeriesSplitter(data,
                                         min_train_len=min_train_len,
                                         incremental_len=incremental_len,
                                         forecast_len=forecast_len,
                                         window_type='rolling', date_col='week')
```

Users can visualize the splits, green is training window and yellow it the forecasting window. The window length is always 380, while the starting point moves forward 20 weeks each steps.

```
[12]: _ = roll_splitter.plot()
```




18.1.3 Specifying number of splits

User can also define number of splits using `n_splits` instead of specifying minimum training length. That way, minimum training length will be automatically calculated.

```
[13]: ex_splitter2 = TimeSeriesSplitter(data,
                                     min_train_len=min_train_len,
                                     incremental_len=incremental_len,
                                     forecast_len=forecast_len,
                                     n_splits=5,
                                     window_type='expanding', date_col='week')
```

```
[14]: _ = ex_splitter2.plot()
```



18.1.4 TimeSeriesSplitter as generator

`TimeSeriesSplitter` is implemented as a generator, therefore users can call `split()` to loop through it. It comes handy even for tasks other than backtest.

```
[15]: for train_df, test_df, scheme, key in roll_splitter.split():
       print('Initial Claim slice {} rolling mean: {:.3f}'.format(key, train_df['claims
       ↪'].mean()))
```

```
Initial Claim slice 0 rolling mean:12.712
Initial Claim slice 1 rolling mean:12.671
Initial Claim slice 2 rolling mean:12.647
```

18.2 Create a BackTester

To actually run backtest, first let's initialize a DLT model and a `BackTester`. You pass in `TimeSeriesSplitter` parameters to `BackTester`.

```
[16]: # instantiate a model
dlt = DLT(
    date_col='week',
    response_col='claims',
    regressor_col=['trend.unemploy', 'trend.filling', 'trend.job'],
    seasonality=52,
```

(continues on next page)

(continued from previous page)

```

estimator='stan-map',
# reduce number of messages
verbose=False,
)

```

```

[17]: # configs
min_train_len = 100
forecast_len = 20
incremental_len = 100
window_type = 'expanding'

bt = BackTester(
    model=dlt,
    df=data,
    min_train_len=min_train_len,
    incremental_len=incremental_len,
    forecast_len=forecast_len,
    window_type=window_type,
)

```

18.3 Backtest fit and predict

The most expensive portion of backtesting is fitting the model iteratively. Thus, users can separate the API calls for `fit_predict` and `score` to avoid redundant computation for multiple metrics or scoring methods

```
[18]: bt.fit_predict()
```

Once `fit_predict()` is called, the fitted models and predictions can be easily retrieved from `BackTester`. Here the data is grouped by the date, `split_key`, and whether or not that observation is part of the training or test data


```
[19]: predicted_df = bt.get_predicted_df()
predicted_df.head()
```

```
[19]:
```

	date	actual	prediction	training_data	split_key
0	2010-01-03	13.386595	13.386595	True	0
1	2010-01-10	13.624218	13.649157	True	0
2	2010-01-17	13.398741	13.373225	True	0
3	2010-01-24	13.137549	13.210564	True	0
4	2010-01-31	13.196760	13.187855	True	0

A plotting utility is also provided to visualize the predictions against the actuals for each split.

```
[20]: plot_bt_predictions(predicted_df, metrics=smape, ncol=2, include_vline=True);
```



Users might find this useful for any custom computations that may need to be performed on the set of predicted data. Note that the columns are renamed to generic and consistent names.

Sometimes, it might be useful to match the data back to the original dataset for ad-hoc diagnostics. This can easily be done by merging back to the original dataset

```
[21]: predicted_df.merge(data, left_on='date', right_on='week')
```

```
[21]:
```

	date	actual	prediction	training_data	split_key	week	\
0	2010-01-03	13.386595	13.386595	True	0	2010-01-03	
1	2010-01-03	13.386595	13.386595	True	1	2010-01-03	
2	2010-01-03	13.386595	13.386595	True	2	2010-01-03	
3	2010-01-03	13.386595	13.386595	True	3	2010-01-03	
4	2010-01-10	13.624218	13.649157	True	0	2010-01-10	
...
1075	2017-12-17	12.568616	12.566429	False	3	2017-12-17	
1076	2017-12-24	12.691451	12.675786	False	3	2017-12-24	
1077	2017-12-31	12.769532	12.783315	False	3	2017-12-31	
1078	2018-01-07	12.908227	12.800098	False	3	2018-01-07	
1079	2018-01-14	12.777193	12.536719	False	3	2018-01-14	
	claims	trend.unemploy	trend.filling	trend.job	sp500	vix	
0	13.386595	0.219882	-0.318452	0.117500	-0.417633	0.122654	
1	13.386595	0.219882	-0.318452	0.117500	-0.417633	0.122654	
2	13.386595	0.219882	-0.318452	0.117500	-0.417633	0.122654	
3	13.386595	0.219882	-0.318452	0.117500	-0.417633	0.122654	
4	13.624218	0.219882	-0.194838	0.168794	-0.425480	0.110445	
...
1075	12.568616	0.298663	0.248654	-0.216869	0.434042	-0.482380	
1076	12.691451	0.328516	0.233616	-0.358839	0.430410	-0.373389	
1077	12.769532	0.503457	0.069313	-0.092571	0.456087	-0.553539	
1078	12.908227	0.527849	0.051295	0.029532	0.471673	-0.456456	
1079	12.777193	0.465717	0.032946	0.006275	0.480271	-0.352770	

(continues on next page)

(continued from previous page)

[1080 rows x 12 columns]

18.4 Backtest Scoring

The main purpose of `BackTester` are the evaluation metrics. Some of the most widely used metrics are implemented and built into the `BackTester` API.

The default metric list is **smape**, **wmape**, **mape**, **mse**, **mae**, **rmsse**.

```
[22]: bt.score()
```

```
[22]:  metric_name  metric_values  is_training_metric
0      smape      0.006787          False
1      wmape      0.006785          False
2       mape      0.006770          False
3       mse      0.013098          False
4       mae      0.086193          False
5      rmsse      0.816876          False
```

It is possible to filter for only specific metrics of interest, or even implement your own callable and pass into the `score()` method. For example, see this function that uses last observed value as a predictor and computes the `mse`. Or `naive_error` which computes the error as the delta between predicted values and the training period mean.

Note these are not really useful error metrics, just showing some examples of callables you can use ;)

```
[23]: def mse_naive(test_actual):
      actual = test_actual[1:]
      prediction = test_actual[:-1]
      return np.mean(np.square(actual - prediction))

      def naive_error(train_actual, test_prediction):
          train_mean = np.mean(train_actual)
          return np.mean(np.abs(test_prediction - train_mean))
```

```
[24]: bt.score(metrics=[mse_naive, naive_error])
```

```
[24]:  metric_name  metric_values  is_training_metric
0   mse_naive      0.019628          False
1 naive_error      0.230727          False
```

It doesn't take additional time to refit and predict the model, since the results are stored when `fit_predict()` is called. Check docstrings for function criteria that is required for it to be supported with this api.

In some cases, users may want to evaluate our metrics on both train and test data. To do this you can call `score` again with the following indicator

```
[25]: bt.score(include_training_metrics=True)
```

```
[25]:  metric_name  metric_values  is_training_metric
0      smape      0.006787          False
1      wmape      0.006785          False
2       mape      0.006770          False
3       mse      0.013098          False
```

(continues on next page)

(continued from previous page)

4	mae	0.086193	False
5	rmsse	0.816876	False
6	smape	0.002739	True
7	wmape	0.002742	True
8	mape	0.002738	True
9	mse	0.003116	True
10	mae	0.035042	True

18.5 Backtest Get Models

In cases where `BackTester` doesn't cut it or for more custom use-cases, there's an interface to export the `TimeSeriesSplitter` and predicted data, as shown earlier. It's also possible to get each of the fitted models for deeper diving

```
[26]: fitted_models = bt.get_fitted_models()
```


```
[27]: model_1 = fitted_models[0]
      model_1.get_regression_coefs()
```

```
[27]:      regressor regressor_sign coefficient
0  trend.unemploy      Regular    -0.048311
1  trend.filling      Regular    -0.120104
2   trend.job         Regular     0.028545
```

`BackTester` composes a `TimeSeriesSplitter` within it, but `TimeSeriesSplitter` can also be created on its own as a standalone object. See section below on `TimeSeriesSplitter` for more details on how to use the splitter.

All of the additional `TimeSeriesSplitter` args can also be passed into `BackTester` on instantiation

```
[28]: ts_splitter = bt.get_splitter()
      _ = ts_splitter.plot()
```



18.6 Hyperparameter Tunning

After seeing the results from the backtest, users may wish to fine tune the hyperparameters. Orbit also provide a `grid_search_orbit` utilities for parameter searching. It uses `Backtester` under the hood so users can compare backtest metrics for different parameters combination.

To have a consistent outlook here, users need to make sure an updated version of `ipywidgets` is installed.

```
[29]: from orbit.utils.params_tuning import grid_search_orbit
```

```
[30]: # defining the search space for level smoothing paramter and seasonality smooth paramter
param_grid = {
    'level_sm_input': [0.3, 0.5, 0.8],
    'seasonality_sm_input': [0.3, 0.5, 0.8],
}
```

```
[31]: # configs
min_train_len = 380 # in case of rolling window, this specify the length of window length
forecast_len = 20 # length forecast window
incremental_len = 20 # step length for moving forward
best_params, tuned_df = grid_search_orbit(
    param_grid,
    model=dlt,
    df=data,
    min_train_len=min_train_len,
    incremental_len=incremental_len,
    forecast_len=forecast_len,
    metrics=None,
    criteria="min",
    verbose=False,
)

0%|          | 0/9 [00:00<?, ?it/s]
```

```
[32]: tuned_df.head() # backtest output for each parameter searched
```

```
[32]:
```

	level_sm_input	seasonality_sm_input	metrics
0	0.3	0.3	0.004909
1	0.3	0.5	0.004058
2	0.3	0.8	0.003609
3	0.5	0.3	0.007908
4	0.5	0.5	0.006306

```
[33]: best_params # output best parameters
```

```
[33]: [{'level_sm_input': 0.3, 'seasonality_sm_input': 0.8}]
```


WBIC/BIC

This notebook gives a tutorial on how to use Watanabe-Bayesian information criterion (WBIC) and Bayesian information criterion (BIC) for feature selection (Watanabe[2010], McElreath[2015], and Vehtari[2016]). The WBIC or BIC is an information criterion. Similar to other criteria (AIC, DIC), the WBIC/BIC endeavors to find the most parsimonious model, i.e., the model that balances fit with complexity. In other words a model (or set of features) that optimizes WBIC/BIC should neither over nor under fit the available data.

In this tutorial a data set is simulated using the damped linear trend (DLT) model. This data set is then used to fit DLT models with varying number of features as well as a global local trend model (GLT), and a Error-Trend-Seasonal (ETS) model. The WBIC/BIC criteria is then show to find the true model.

Note that we recommend the use of WBIC for full Bayesian and SVI estimators and BIC for MAP estimator.

```
[1]: from datetime import timedelta

import pandas as pd
import numpy as np

import matplotlib.pyplot as plt
%matplotlib inline

import orbit
from orbit.models import DLT, ETS, KTRLite, LGT
from orbit.utils.simulation import make_trend, make_regression
```

```
[2]: print(orbit.__version__)
```

```
1.1.3dev
```

19.1 Data Simulation

This block of code creates random data set (365 observations with 10 features) assuming a DLT model. Of the 10 features 5 are effective regressors; i.e., they are used in the true model to create the data.

As an exercise left to the user once you have run the code once try changing the NUM_OF_EFFECTIVE_REGRESSORS (line 2), the SERIES_LEN (line 3), and the SEED (line 4) to see how it effects the results.

```
[3]: NUM_OF_REGRESSORS = 10
NUM_OF_EFFECTIVE_REGRESSORS = 4
SERIES_LEN = 365
SEED = 1
```

(continues on next page)

(continued from previous page)

```

# sample some coefficients
COEFS = np.random.default_rng(SEED).uniform(-1, 1, NUM_OF_EFFECTIVE_REGRESSORS)
trend = make_trend(SERIES_LEN, rw_loc=0.01, rw_scale=0.1)
x, regression, coefs = make_regression(series_len=SERIES_LEN, coefs=COEFS)

# combine trend and the regression
y = trend + regression
y = y - y.min()

x_extra = np.random.normal(0, 1, (SERIES_LEN, NUM_OF_REGRESSORS - NUM_OF_EFFECTIVE_
↳REGRESSORS))
x = np.concatenate([x, x_extra], axis=-1)

x_cols = [f"x{x}" for x in range(1, NUM_OF_REGRESSORS + 1)]
response_col = "y"
dt_col = "date"
obs_matrix = np.concatenate([y.reshape(-1, 1), x], axis=1)
# make a data frame for orbit inputs
df = pd.DataFrame(obs_matrix, columns=[response_col] + x_cols)
# make some dummy date stamp
dt = pd.date_range(start='2016-01-04', periods=SERIES_LEN, freq="1W")
df['date'] = dt

```

```

[4]: print(df.shape)
print(df.head())

```

```

(365, 12)

```

	y	x1	x2	x3	x4	x5	x6	\
0	4.426242	0.172792	0.000000	0.165219	-0.000000	-0.171564	0.646560	
1	5.580432	0.452678	0.223187	-0.000000	0.290559	1.760342	1.809456	
2	5.031773	0.182286	0.147066	0.014211	0.273356	-1.536165	-1.029521	
3	3.264027	-0.368227	-0.081455	-0.241060	0.299423	0.591493	0.505090	
4	5.246511	0.019861	-0.146228	-0.390954	-0.128596	-1.447045	-1.678974	

	x7	x8	x9	x10	date
0	-0.140274	2.024782	-0.334100	-1.179794	2016-01-10
1	-1.191844	-0.990802	-2.894241	1.757803	2016-01-17
2	0.572564	-0.750318	0.168999	0.524423	2016-01-24
3	1.348104	0.464408	1.263463	-1.705559	2016-01-31
4	0.292021	0.534810	1.463977	-0.137525	2016-02-07

19.2 WBIC

In this section, we use DLT model as an example. Different DLT models (the number of features used changes) are fitted and their WBIC values are calculated respectively.

```
[5]: %%time
WBIC_ls = []
for k in range(1, NUM_OF_REGRESSORS + 1):
    regressor_col = x_cols[:k]
    dlt_mod = DLT(
        response_col=response_col,
        date_col=dt_col,
        regressor_col=regressor_col,
        seed=2022,
        # fixing the smoothing parameters to learn regression coefficients more
        ↪effectively
        level_sm_input=0.01,
        slope_sm_input=0.01,
        num_warmup=4000,
        num_sample=4000,
    )
    WBIC_temp = dlt_mod.fit_wbic(df=df)
    print("WBIC value with {:d} regressors: {:.3f}".format(k, WBIC_temp))
    print('-----')
    WBIC_ls.append(WBIC_temp)
```

```
INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 5.900, warmups (per
↪chain): 1000 and samples(per chain): 1000.
```

```
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests
↪for n_eff and Rhat.
```

```
To run all diagnostics call pystan.check_hmc_diagnostics(fit)
```

```
INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 5.900, warmups (per
↪chain): 1000 and samples(per chain): 1000.
```

```
WBIC value with 1 regressors: 1202.366
-----
```

```
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests
↪for n_eff and Rhat.
```

```
To run all diagnostics call pystan.check_hmc_diagnostics(fit)
```

```
INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 5.900, warmups (per
↪chain): 1000 and samples(per chain): 1000.
```

```
WBIC value with 2 regressors: 1150.291
-----
```

```
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests
↪for n_eff and Rhat.
```

```
To run all diagnostics call pystan.check_hmc_diagnostics(fit)
```

```
INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 5.900, warmups (per
↪chain): 1000 and samples(per chain): 1000.
```

```
WBIC value with 3 regressors: 1104.491
-----
```

```
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests.
↳for n_eff and Rhat.
```

```
To run all diagnostics call pystan.check_hmc_diagnostics(fit)
```

```
INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 5.900, warmups (per
↳chain): 1000 and samples(per chain): 1000.
```

```
WBIC value with 4 regressors: 1054.450
-----
```

```
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests.
↳for n_eff and Rhat.
```

```
To run all diagnostics call pystan.check_hmc_diagnostics(fit)
```

```
INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 5.900, warmups (per
↳chain): 1000 and samples(per chain): 1000.
```

```
WBIC value with 5 regressors: 1060.473
-----
```

```
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests.
↳for n_eff and Rhat.
```

```
To run all diagnostics call pystan.check_hmc_diagnostics(fit)
```

```
INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 5.900, warmups (per
↳chain): 1000 and samples(per chain): 1000.
```

```
WBIC value with 6 regressors: 1066.991
-----
```

```
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests.
↳for n_eff and Rhat.
```

```
To run all diagnostics call pystan.check_hmc_diagnostics(fit)
```

```
INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 5.900, warmups (per
↳chain): 1000 and samples(per chain): 1000.
```

```
WBIC value with 7 regressors: 1073.808
-----
```

```
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests.
↳for n_eff and Rhat.
```

```
To run all diagnostics call pystan.check_hmc_diagnostics(fit)
```

```
INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 5.900, warmups (per
↳chain): 1000 and samples(per chain): 1000.
```

```
WBIC value with 8 regressors: 1080.625
-----
```

```
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests.
↳for n_eff and Rhat.
```

```
To run all diagnostics call pystan.check_hmc_diagnostics(fit)
```

```
INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 5.900, warmups (per
↳chain): 1000 and samples(per chain): 1000.
```

```
WBIC value with 9 regressors: 1084.251
-----
```

```
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests.
↳for n_eff and Rhat.
```

```
To run all diagnostics call pystan.check_hmc_diagnostics(fit)
```

```
WBIC value with 10 regressors: 1088.111
```

```
-----
CPU times: user 3.12 s, sys: 4.08 s, total: 7.2 s
Wall time: 4min 37s
```

It is also interesting to see if WBIC can distinguish between model types; not just do feature selection for a given type of model. To that end the next block fits an LGT and ETS model to the data; the WBIC values for both models are then calculated.

Note that WBIC is supported for both the ‘stan-mcmc’ and ‘pyro-svi’ estimators. Currently only the LGT model has both. Thus WBIC is calculated for LGT for both estimators.

```
[6]: %%time
lgt = LGT(response_col=response_col,
          date_col=dt_col,
          regressor_col=regressor_col,
          seasonality=52,
          estimator='stan-mcmc',
          seed=8888)
WBIC_lgt_mcmc = lgt.fit_wbic(df=df)
print("WBIC value for LGT model (stan MCMC): {:.3f}".format(WBIC_lgt_mcmc))

lgt = LGT(response_col=response_col,
          date_col=dt_col,
          regressor_col=regressor_col,
          seasonality=52,
          estimator='pyro-svi',
          seed=8888)
WBIC_lgt_pyro = lgt.fit_wbic(df=df)
print("WBIC value for LGT model (pyro SVI): {:.3f}".format(WBIC_lgt_pyro))

ets = ETS(
    response_col=response_col,
    date_col=dt_col,
    seed=2020,
    # fixing the smoothing parameters to learn regression coefficients more
    ↪ effectively
    level_sm_input=0.01,
)

WBIC_ets = ets.fit_wbic(df=df)
print("WBIC value for ETS model: {:.3f}".format(WBIC_ets))

WBIC_ls.append(WBIC_lgt_mcmc)
WBIC_ls.append(WBIC_lgt_pyro)
WBIC_ls.append(WBIC_ets)

INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 5.900, warmups (per
↪ chain): 225 and samples(per chain): 25.
WARNING:pystan:Maximum (flat) parameter count (1000) exceeded: skipping diagnostic tests
↪ for n_eff and Rhat.
To run all diagnostics call pystan.check_hmc_diagnostics(fit)
INFO:orbit:Using SVI (Pyro) with steps: 301, samples: 100, learning rate: 0.1, learning
↪ rate_total_decay: 1.0 and particles: 100.
```


(continues on next page)

(continued from previous page)

```
INFO:root:Guessed max_plate_nesting = 2
WBIC value for LGT model (stan MCMC): 1144.606
INFO:orbit:step    0 loss = 303.3, scale = 0.11508
INFO:orbit:step  100 loss = 116.73, scale = 0.50351
INFO:orbit:step  200 loss = 116.58, scale = 0.5112
INFO:orbit:step  300 loss = 116.84, scale = 0.50453
INFO:orbit:Sampling (PyStan) with chains: 4, cores: 8, temperature: 5.900, warmups (per_
↳chain): 225 and samples(per chain): 25.
WBIC value for LGT model (pyro SVI): 1124.428
WARNING:pystan:n_eff / iter below 0.001 indicates that the effective sample size has_
↳likely been overestimated
WARNING:pystan:Rhat above 1.1 or below 0.9 indicates that the chains very likely have_
↳not mixed
WBIC value for ETS model: 1203.031
CPU times: user 50.9 s, sys: 8.06 s, total: 59 s
Wall time: 57.9 s
```

The plot below shows the WBIC vs the number of features / model type (blue line). The true model is indicated by the vertical red line. The horizontal gray line shows the minimum (optimal) value. The minimum is at the true value.

```
[7]: labels = ["DLT_{}".format(x) for x in range(1, NUM_OF_REGRESSORS + 1)] + ['LGT_MCMC',
↳'LGT_SVI', 'ETS']
fig, ax = plt.subplots(1, 1, figsize=(12, 6), dpi=80)
markerline, stemlines, baseline = ax.stem(
    np.arange(len(labels)), np.array(WBIC_ls), label='WBIC', markerfmt='D')
baseline.set_color('none')
markerline.set_markersize(12)
ax.set_ylim(1020, )
ax.set_xticks(np.arange(len(labels)))
ax.set_xticklabels(labels)
# because list type is mixed index from 1;
ax.axvline(x=NUM_OF_EFFECTIVE_REGRESSORS - 1, color='red', linewidth=3, alpha=0.5,
↳linestyle='-', label='truth')
ax.set_ylabel("WBIC")
ax.set_xlabel("# of Features / Model Type")
ax.legend();
```



19.3 BIC

In this section, we use DLT model as an example. Different DLT models (the number of features used changes) are fitted and their BIC values are calculated respectively.

```
[8]: %%time
BIC_ls = []
for k in range(0, NUM_OF_REGRESSORS):
    regressor_col = x_cols[:k + 1]
    dlt_mod = DLT(
        estimator='stan-map',
        response_col=response_col,
        date_col=dt_col,
        regressor_col=regressor_col,
        seed=2022,
        # fixing the smoothing parameters to learn regression coefficients more
        ↪effectively
        level_sm_input=0.01,
        slope_sm_input=0.01,
    )
    dlt_mod.fit(df=df)
    BIC_temp = dlt_mod.get_bic()
    print("BIC value with {:d} regressors: {:.3f}".format(k + 1, BIC_temp))
    print('-----')
    BIC_ls.append(BIC_temp)
```

```
INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
```

```
BIC value with 1 regressors: 1247.445
```

```
-----
BIC value with 2 regressors: 1191.889
```

```
-----
BIC value with 3 regressors: 1139.408
```

```
INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
```

```
INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
```

```
BIC value with 4 regressors: 1079.639
```

```
-----
BIC value with 5 regressors: 1082.445
```

```
INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
```

```
INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
```

```
BIC value with 6 regressors: 1082.482
```

```
-----
BIC value with 7 regressors: 1082.217
```

```
INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
```

```
INFO:orbit:Optimizing (PyStan) with algorithm: LBFGS.
```

```
BIC value with 8 regressors: 1081.698
```

```
-----
BIC value with 9 regressors: 1078.876
```


```
-----
BIC value with 10 regressors: 1187.269
```

```
-----
CPU times: user 935 ms, sys: 58.7 ms, total: 994 ms
```

```
Wall time: 995 ms
```

The plot below shows the BIC vs the number of features (blue line). The true model is indicated by the vertical red line. The horizontal gray line shows the minimum (optimal) value. The minimum is at the true value.

```
[9]: labels = ["DLT_{}".format(x) for x in range(1, NUM_OF_REGRESSORS + 1)]
fig, ax = plt.subplots(1, 1, figsize=(12, 6), dpi=80)
markerline, stemlines, baseline = ax.stem(
    np.arange(len(labels)), np.array(BIC_ls), label='BIC', markerfmt='D')
baseline.set_color('none')
markerline.set_markersize(12)
ax.set_ylim(1020, )
ax.set_xticks(np.arange(len(labels)))
ax.set_xticklabels(labels)
# because list type is mixed index from 1;
ax.axvline(x=NUM_OF_EFFECTIVE_REGRESSORS - 1, color='red', linewidth=3, alpha=0.5,
↵linestyle='-', label='truth')
ax.set_ylabel("BIC")
ax.set_xlabel("# of Features")
ax.legend();
```



19.4 References

1. Watanabe Sumio (2010). “Asymptotic Equivalence of Bayes Cross Validation and Widely Applicable Information Criterion in Singular Learning Theory”. *Journal of Machine Learning Research*. 11: 3571–3594.
2. McElreath Richard (2015). “Statistical Rethinking: A Bayesian course with examples in R and Stan” Second Ed. 193-221.
3. Vehtari Aki, Gelman Andrew, Gabry Jonah (2016) “Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC”

EDA UTILITIES

In this section, we will introduce a rich set of plotting functions in orbit for the EDA (exploratory data analysis) purpose. The plots include

- Time series heatmap
- Correlation heatmap
- Dual axis time series plot
- Wrap plot

```
[1]: import seaborn as sns
      from matplotlib import pyplot as plt
      import pandas as pd
      import numpy as np

      import orbit
      from orbit.utils.dataset import load_iclaims
      from orbit.eda import eda_plot
```

```
[2]: print(orbit.__version__)
```

```
1.1.3
```

```
[3]: df = load_iclaims()
      df['week'] = pd.to_datetime(df['week'])
```

```
[4]: df.head()
```


```
[4]:
```

	week	claims	trend.unemploy	trend.filling	trend.job	sp500	\
0	2010-01-03	13.386595	0.219882	-0.318452	0.117500	-0.417633	
1	2010-01-10	13.624218	0.219882	-0.194838	0.168794	-0.425480	
2	2010-01-17	13.398741	0.236143	-0.292477	0.117500	-0.465229	
3	2010-01-24	13.137549	0.203353	-0.194838	0.106918	-0.481751	
4	2010-01-31	13.196760	0.134360	-0.242466	0.074483	-0.488929	


	vix
0	0.122654
1	0.110445
2	0.532339
3	0.428645
4	0.487404

20.1 Time series heat map

```
[5]: _ = eda_plot.ts_heatmap(df = df, date_col = 'week', seasonal_interval=52, value_col=  
    ↪ 'claims')
```




```
[6]: _ = eda_plot.ts_heatmap(df = df, date_col = 'week', seasonal_interval=52, value_col=  
    ↪ 'claims', normalization=True)
```




20.2 Correlation heatmap

```
[7]: var_list = ['trend.unemploy', 'trend.filling', 'trend.job', 'sp500', 'vix']
_ = eda_plot.correlation_heatmap(df, var_list = var_list,
                                fig_width=10, fig_height=6)
```



20.3 Dual axis time series plot

```
[8]: _ = eda_plot.dual_axis_ts_plot(df=df, var1='trend.unemploy', var2='claims', date_col='week')
```



20.4 Wrap plots for quick glance of data patterns


```
[9]: var_list=['week', 'trend.unemploy', 'trend.filling', 'trend.job', 'sp500', 'vix']
df[var_list].melt(id_vars = ['week'])
```

```
[9]:
```

	week	variable	value
0	2010-01-03	trend.unemploy	0.219882
1	2010-01-10	trend.unemploy	0.219882
2	2010-01-17	trend.unemploy	0.236143
3	2010-01-24	trend.unemploy	0.203353
4	2010-01-31	trend.unemploy	0.134360
...
2210	2018-05-27	vix	-0.175192
2211	2018-06-03	vix	-0.275119
2212	2018-06-10	vix	-0.291676
2213	2018-06-17	vix	-0.152422
2214	2018-06-24	vix	0.003284

```
[2215 rows x 3 columns]
```

```
[10]: _ = eda_plot.wrap_plot_ts(df, 'week', var_list)
```



SIMULATION DATA

Orbit provide the functions to generate the simulation data including:

1. Generate the data with time-series trend:
 - random walk
 - arima
2. Generate the data with seasonality
 - discrete
 - fourier series
3. Generate regression data

```
[1]: import numpy as np
import matplotlib.pyplot as plt
from sklearn.linear_model import LinearRegression

import orbit
from orbit.utils.simulation import make_trend, make_seasonality, make_regression
from orbit.utils.plot import get_orbit_style
plt.style.use(get_orbit_style())
from orbit.constants.palette import OrbitPalette

%matplotlib inline
```


```
[2]: print(orbit.__version__)

1.1.3dev
```

21.1 Trend

21.1.1 Random Walk


```
[3]: rw = make_trend(200, rw_loc=0.02, rw_scale=0.1, seed=2020)
_ = plt.plot(rw, color = OrbitPalette.BLUE.value)
```



21.1.2 ARMA

reference for the ARMA process: https://www.statsmodels.org/stable/generated/statsmodels.tsa.arima_process.ArmaProcess.html

```
[4]: arma_trend = make_trend(200, method='arma', arma=[.8, -.1], seed=2020)
_ = plt.plot(arma_trend, color = OrbitPalette.BLUE.value)
```



21.2 Seasonality

21.2.1 Discrete

generating a weekly seasonality(=7) where seasonality within a day is constant(duration=24) on an hourly time-series

```
[5]: ds = make_seasonality(500, seasonality=7, duration=24, method='discrete', seed=2020)
_ = plt.plot(ds, color = OrbitPalette.BLUE.value)
```



21.3 Regression

generating multiplicative time-series with trend, seasonality and regression components

```
[8]: # define the regression coefficients
     coefs = [0.1, -.33, 0.8]
```

```
[9]: x, y, coefs = make_regression(200, coefs, scale=2.0, seed=2020)
```

```
[10]: _ = plt.plot(y, color = OrbitPalette.BLUE.value)
```



```
[11]: # check if get the coefficients as set up
     reg = LinearRegression().fit(x, y)
     print(reg.coef_)
```

```
[ 0.1586677  -0.33126796  0.7974205 ]
```

OTHER UTILITIES

22.1 Generating Full Span of multiple time-series

```
[1]: import pandas as pd
import numpy as np
from orbit.utils.general import expand_grid, regenerate_base_df
```

Define the series keys and datetime array.

```
[2]: dt = pd.date_range('2020-01-31', '2022-12-31', freq='M')
keys = ['x' + str(x) for x in range(10)]
print(keys)
print(dt)

['x0', 'x1', 'x2', 'x3', 'x4', 'x5', 'x6', 'x7', 'x8', 'x9']
DatetimeIndex(['2020-01-31', '2020-02-29', '2020-03-31', '2020-04-30',
               '2020-05-31', '2020-06-30', '2020-07-31', '2020-08-31',
               '2020-09-30', '2020-10-31', '2020-11-30', '2020-12-31',
               '2021-01-31', '2021-02-28', '2021-03-31', '2021-04-30',
               '2021-05-31', '2021-06-30', '2021-07-31', '2021-08-31',
               '2021-09-30', '2021-10-31', '2021-11-30', '2021-12-31',
               '2022-01-31', '2022-02-28', '2022-03-31', '2022-04-30',
               '2022-05-31', '2022-06-30', '2022-07-31', '2022-08-31',
               '2022-09-30', '2022-10-31', '2022-11-30', '2022-12-31'],
              dtype='datetime64[ns]', freq='M')
```

Users can use `expand_grid` to generate dataframe with observations in key and dt levels.

```
[3]: df_base = expand_grid({
    'key': keys,
    'dt': dt,
})
x = np.random.normal(0, 1, 10 * 36)
df_base['x'] = x
print(df_base.shape)
df_base.head(5)
```

```
(360, 3)
```

```
[3]:   key      dt      x
0  x0 2020-01-31  0.833432
1  x0 2020-02-29 -0.524211
```

(continues on next page)

(continued from previous page)

```
2  x0 2020-03-31 -0.443138
3  x0 2020-04-30 -1.522778
4  x0 2020-05-31 -1.982342
```

22.2 Regenerate Multiple Timeseries with Missing rows

Create missing rows.

```
[4]: np.random.seed(2022)
drop_idx = np.random.choice(df_base.index, 5, replace=False)
df_missing = df_base.drop(drop_idx).reset_index(drop=True)
print(df_missing.shape)
df_missing.head(5)
```

```
(355, 3)
```

```
[4]:   key      dt      x
0  x0 2020-01-31  0.833432
1  x0 2020-02-29 -0.524211
2  x0 2020-03-31 -0.443138
3  x0 2020-04-30 -1.522778
4  x0 2020-05-31 -1.982342
```

Use `regenerate_base_df` to regenerate the base dataframe.

```
[5]: time_col = "dt"
key_col = "key"
new_df_base = regenerate_base_df(df_missing, time_col, key_col, val_cols=['x'])
```

By default, the missing entries regenerated come with a null value.

```
[6]: new_df_base.iloc[drop_idx]
```

```
[6]:      dt key  x
286 2022-11-30  x7 NaN
274 2021-11-30  x7 NaN
75  2020-04-30  x2 NaN
135 2022-04-30  x3 NaN
43  2020-08-31  x1 NaN
```

Users can also use `fill_na` option to fill the missing values.

```
[7]: new_df_base = regenerate_base_df(df_missing, time_col, key_col, val_cols=['x'], fill_
↪na=0)
```

```
[8]: new_df_base.iloc[drop_idx]
```

```
[8]:      dt key  x
286 2022-11-30  x7  0.0
274 2021-11-30  x7  0.0
75  2020-04-30  x2  0.0
135 2022-04-30  x3  0.0
43  2020-08-31  x1  0.0
```

BUILD YOUR OWN MODEL

One important feature of `orbit` is to allow developers to build their own models in a relatively flexible manner to serve their own purpose. This tutorial will go over a demo on how to build up a simple Bayesian linear regression model using `Pyro API` in the backend with `orbit` interface.

23.1 Orbit Class Design

In version `1.1.0`, the classes within `Orbit` are re-designed as such:

1. Forecaster
2. Model
3. Estimator

23.1.1 Forecaster

Forecaster provides general interface for users to perform `fit` and `predict` task. It is further inherited to provide different types of forecasting methodology:

1. [Maximum a posterior \(MAP\)](#)
2. [\[Stochastic Variational Inference \(SVI\)\]](#)
3. [Full Bayesian](#)

The discrepancy on these three methods mainly lie on the posteriors estimation where **MAP** will yield point posterior estimate and can be extracted through the method `get_point_posterior()`. Meanwhile, **SVI** and **Full Bayesian** allow posterior sample extraction through the method `get_posteriors()`. Alternatively, you can also approximate point estimate by passing through additional arg such as `point_method='median'` in the `.fit()` process.

To make use of a **Forecaster**, one must provide these two objects:

1. Model
2. Estimator

Theses two objects are prototyped as abstract and next subsections will cover how they work.

23.1.2 Model

Model is an object defined by a class inherited from `BaseTemplate` a.k.a **Model Template** in the diagram below. It mainly turns the logic of `fit()` and `predict()` concrete by supplying the `fitter` as a file (**PyStan**) or a callable class (**Pyro**) and the internal `predict()` method. This object defines the overall inputs, model structure, parameters and likelihoods.

23.1.3 Estimator

Meanwhile, there are different APIs implement slightly different ways of sampling and optimization (for **MAP**). `orbit` is designed to support various APIs such as **PyStan** and **Pyro** (hopefully PyMC3, Numpyro in the future!). The logic separating the call of different APIs with different interface is done by the **Estimator** class which is further inherited in `PyroEstimator` and `StanEstimator`.




Diagram above shows the interaction across classes under the Orbit package design.

23.2 Creating a Bayesian Linear Regression Model

The plan here is to build a classical regression model with the formula below:

$$y = \alpha + X\beta + \epsilon$$

where α is the intercept, β is the coefficients matrix and ϵ is the random noise.

To start with let's load the libraries.

```
[1]: import pandas as pd
import numpy as np
import torch
import pyro
import pyro.distributions as dist
from copy import deepcopy
import matplotlib.pyplot as plt

import orbit
from orbit.template.model_template import ModelTemplate
from orbit.forecaster import SVIForecaster
from orbit.estimated.pyro_estimator import PyroEstimatorSVI

from orbit.utils.simulation import make_regression
from orbit.diagnostics.plot import plot_predicted_data
from orbit.utils.plot import get_orbit_style
plt.style.use(get_orbit_style())

%matplotlib inline
```

```
[2]: print(orbit.__version__)

1.1.3
```

Since the **Forecaster** and **Estimator** are already built inside `orbit`, the rest of the ingredients to construct a new model will be a **Model** object that contains the follow:

- a callable class as a fitter
- a predict method

23.2.1 Define a Fitter

For **Pyro** users, you should find the code below familiar. All it does is to put a Bayesian linear regression (**BLR**) model code in a callable class. Details of **BLR** will not be covered here. Note that the parameters here need to be consistent .

```
[3]: class MyFitter:
    max_plate_nesting = 1 # max number of plates nested in model

    def __init__(self, data):
        for key, value in data.items():
            key = key.lower()
            if isinstance(value, (list, np.ndarray)):
                value = torch.tensor(value, dtype=torch.float)
```

(continues on next page)

```

        self.__dict__[key] = value

def __call__(self):
    extra_out = {}

    p = self.regressor.shape[1]
    bias = pyro.sample("bias", dist.Normal(0, 1))
    weight = pyro.sample("weight", dist.Normal(0, 1).expand([p]).to_event(1))
    yhat = bias + weight @ self.regressor.transpose(-1, -2)
    obs_sigma = pyro.sample("obs_sigma", dist.HalfCauchy(self.response_sd))

    with pyro.plate("response_plate", self.num_of_obs):
        pyro.sample("response", dist.Normal(yhat, obs_sigma), obs=self.response)

    log_prob = dist.Normal(yhat[...], 1:), obs_sigma).log_prob(self.response[1:])
    extra_out.update(
        {"log_prob": log_prob}
    )

    return extra_out

```

23.2.2 Define the Model Class

This is the part requires the knowledge of `orbit` most. First we construct a class by plugging in the `fitter` callable. Users need to let the `orbit` estimators know the required input in addition to the defaults (e.g. `response`, `response_sd` etc.). In this case, it takes `regressor` as the matrix input from the data frame. That is why there are lines of code to provide this information in

1. `_data_input_mapper` - a list or Enum to let estimator keep tracking required data input
2. `set_dynamic_attributes` - the logic define the actual inputs i.e. `regressor` from the data frame. This is a **reserved function** being called inside **Forecaster**.

Finally, we code the logic in `predict()` to define how we utilize posteriors to perform in-sample / out-of-sample prediction. Note that the output needs to be a dictionary where it supports **components decomposition**.

```

[4]: class BayesLinearRegression(ModelTemplate):
    _fitter = MyFitter
    _data_input_mapper = ['regressor']
    _supported_estimator_types = [PyroEstimatorSVI]

    def __init__(self, regressor_col, **kwargs):
        super().__init__(**kwargs)
        self.regressor_col = regressor_col
        self.regressor = None
        self._model_param_names = ['bias', 'weight', 'obs_sigma']

    def set_dynamic_attributes(self, df, training_meta):
        self.regressor = df[self.regressor_col].values

    def predict(self, posterior_estimates, df, training_meta, prediction_meta, include_
    ↪error=False, **kwargs):

```

(continues on next page)

(continued from previous page)

```

model = deepcopy(posterior_estimates)
new_regressor = df[self.regressor_col].values.T
bias = np.expand_dims(model.get('bias'), -1)
obs_sigma = np.expand_dims(model.get('obs_sigma'), -1)
weight = model.get('weight')

pred_len = df.shape[0]
batch_size = weight.shape[0]

prediction = bias + np.matmul(weight, new_regressor) + \
    np.random.normal(0, obs_sigma, size=(batch_size, pred_len))
return {'prediction': prediction}

```

23.3 Test the New Model with Forecaster

Once the model class is defined. User can initialize an object and build a forecaster for fit and predict purpose. Before doing that, the demo provides a simulated dataset here.

23.3.1 Data Simulation

```
[5]: x, y, coefs = make_regression(120, [3.0, -1.0], bias=1.0, scale=1.0)
```

```
[6]: df = pd.DataFrame(
    np.concatenate([y.reshape(-1, 1), x], axis=1), columns=['y', 'x1', 'x2']
)
df['week'] = pd.date_range(start='2016-01-04', periods=len(y), freq='7D')
```

```
[7]: df.head(5)
```

```
[7]:
```

	y	x1	x2	week
0	2.382337	0.345584	0.000000	2016-01-04
1	2.812929	0.330437	-0.000000	2016-01-11
2	3.600130	0.905356	0.446375	2016-01-18
3	-0.884275	-0.000000	0.581118	2016-01-25
4	2.704941	0.364572	0.294132	2016-02-01

```
[8]: test_size = 20
train_df = df[:-test_size]
test_df = df[-test_size:]
```

23.3.2 Create the Forecaster

As mentioned previously, model is the inner object to control the math. To use it for fit and predict purpose, we need a **Forecaster**. Since the model is written in **Pyro**, the pick here should be **SVIForecaster**.

```
[9]: model = BayesLinearRegression(
      regressor_col=['x1', 'x2'],
    )
```

```
[10]: blr = SVIForecaster(
        model=model,
        response_col='y',
        date_col='week',
        estimator_type=PyroEstimatorSVI,
        verbose=True,
        num_steps=501,
        seed=2021,
    )
```

```
[11]: blr
```

```
[11]: <orbit.forecaster.svi.SVIForecaster at 0x166027070>
```

Now, an object `blr` is instantiated as a `SVIForecaster` object and is ready for fit and predict.

```
[12]: blr.fit(train_df)
```

```
INFO:orbit:Using SVI (Pyro) with steps: 501, samples: 100, learning rate: 0.1, learning_
↪rate_total_decay: 1.0 and particles: 100.
INFO:root:Guessed max_plate_nesting = 2
INFO:orbit:step    0 loss = 27333, scale = 0.077497
INFO:orbit:step  100 loss = 12594, scale = 0.0092399
INFO:orbit:step  200 loss = 12596, scale = 0.0095782
INFO:orbit:step  300 loss = 12592, scale = 0.0097062
INFO:orbit:step  400 loss = 12595, scale = 0.0091866
INFO:orbit:step  500 loss = 12592, scale = 0.0095684
```

```
[12]: <orbit.forecaster.svi.SVIForecaster at 0x166027070>
```

23.3.3 Compare Coefficients with Truth

```
[13]: estimated_weights = blr.get_posterior_samples()['weight']
```

The code below is to compare the median of coefficients posteriors which is labeled as `weight` with the truth.


```
[14]: print("True Coef: {:.3f}, {:.3f}".format(coefs[0], coefs[1]) )
      estimated_coef = np.median(estimated_weights, axis=0)
      print("Estimated Coef: {:.3f}, {:.3f}".format(estimated_coef[0], estimated_coef[1]))
```

```
True Coef: 3.000, -1.000
Estimated Coef: 2.947, -0.965
```

23.3.4 Examine Forecast Accuracy

```
[15]: predicted_df = blr.predict(df)
```

```
[16]: _ = plot_predicted_data(train_df, predicted_df, 'week', 'y', test_actual_df=test_df,
↳ prediction_percentiles=[5, 95])
```



23.3.5 Additional Notes

In general, most of the **diagnostic tools** in orbit such as posteriors checking and plotting is applicable in the model created in this style. Also, users can provide `point_method='median'` in the `fit()` under the `SVIForecaster` to extract median of posteriors directly.

24.1 orbit package

24.1.1 Subpackages

`orbit.constants` package

Submodules

`orbit.constants.constants` module

`class orbit.constants.constants.BacktestFitKeys`(*value*)

Bases: `enum.Enum`

column names of the dataframe used in the output from the `backtest.BackTester.fit_predict()` or any labels of the intermediate variables to generate such outcome dataframe

`ACTUAL = 'actual'`

`DATE = 'date'`

`METRIC_NAME = 'metric_name'`

`METRIC_VALUES = 'metric_values'`

`PREDICTED = 'prediction'`

`SPLIT_KEY = 'split_key'`

`TEST_ACTUAL = 'test_actual'`

`TEST_PREDICTED = 'test_prediction'`

`TRAIN_ACTUAL = 'train_actual'`

`TRAIN_FLAG = 'training_data'`

`TRAIN_METRIC_FLAG = 'is_training_metric'`

`TRAIN_PREDICTED = 'train_prediction'`

`class orbit.constants.constants.CompiledStanModelPath`

Bases: `object`

the directory path for compiled stan models

`CHILD = 'stan_compiled'`

```
PARENT = 'orbit'
```

```
class orbit.constants.constants.EstimatorsKeys(value)
```

```
Bases: enum.Enum
```

alias for all available estimator types when they are called under model wrapper functions

```
PyroSVI = 'pyro-svi'
```

```
StanMAP = 'stan-map'
```

```
StanMCMC = 'stan-mcmc'
```

```
class orbit.constants.constants.KTRTimePointPriorKeys(value)
```

```
Bases: enum.Enum
```

hash table keys for the dictionary of back-test aggregation analysis result

```
NAME = 'name'
```

```
PRIOR_END_TP_IDX = 'prior_end_tp_idx'
```

```
PRIOR_MEAN = 'prior_mean'
```

```
PRIOR_REGRESSOR_COL = 'prior_regressor_col'
```

```
PRIOR_SD = 'prior_sd'
```

```
PRIOR_START_TP_IDX = 'prior_start_tp_idx'
```

```
class orbit.constants.constants.PlotLabels(value)
```

```
Bases: enum.Enum
```

used in multiple prediction plots

```
ACTUAL_RESPONSE = 'actual_response'
```

```
PREDICTED_RESPONSE = 'predicted_response'
```

```
TRAINING_ACTUAL_RESPONSE = 'training_actual_response'
```

```
class orbit.constants.constants.PredictMethod(value)
```

```
Bases: enum.Enum
```

The predict method for all of the stan template. Often used are mean and median.

```
FULL_SAMPLING = 'full'
```

```
MAP = 'map'
```

```
MEAN = 'mean'
```

```
MEDIAN = 'median'
```

```
class orbit.constants.constants.PredictionKeys(value)
```

```
Bases: enum.Enum
```

column names for the data frame of predicted result with decomposed components

```
PREDICTION = 'prediction'
```

```
REGRESSION = 'regression'
```

```
REGRESSOR = 'regressor'
```

```
SEASONALITY = 'seasonality'
```

```
TREND = 'trend'
```

```
class orbit.constants.constants.PredictionMetaKeys(value)
    Bases: enum.Enum

    prediction input meta data dictionary processed under Forecaster.predict()

    DATE_ARRAY = 'date_array'
    END = 'prediction_end'
    END_INDEX = 'end'
    FUTURE_STEPS = 'n_forecast_steps'
    PREDICTION_DF_LEN = 'df_length'
    START = 'prediction_start'
    START_INDEX = 'start'

class orbit.constants.constants.TimeSeriesSplitSchemeKeys(value)
    Bases: enum.Enum

    hash table keys for the dictionary of back-test meta data

    MODEL = 'model'
    SPLIT_TYPE_EXPANDING = 'expanding'
    SPLIT_TYPE_ROLLING = 'rolling'
    TEST_IDX = 'test_idx'
    TRAIN_END_DATE = 'train_end_date'
    TRAIN_IDX = 'train_idx'
    TRAIN_START_DATE = 'train_start_date'

class orbit.constants.constants.TrainingMetaKeys(value)
    Bases: enum.Enum

    training meta data dictionary processed under Forecaster.fit()

    DATE_ARRAY = 'date_array'
    DATE_COL = 'date_col'
    END = 'training_end'
    NUM_OF_OBS = 'num_of_obs'
    RESPONSE = 'response'
    RESPONSE_COL = 'response_col'
    RESPONSE_MEAN = 'response_mean'
    RESPONSE_SD = 'response_sd'
    START = 'training_start'
```

orbit.constants.dlt module

orbit.constants.lgt module

orbit.constants.palette module

```
class orbit.constants.palette.KTRPalette(value)
```

```
    Bases: enum.Enum
```

```
    str
```

```
    KNOTS_REGION = '#05A357'
```

```
    KNOTS_SEGMENT = '#276ef1'
```

```
class orbit.constants.palette.OrbitColorMap(value)
```

```
    Bases: enum.Enum
```

```
    matplotlib.ColorMap
```

```
    BLACK_GRADIENT = <matplotlib.colors.LinearSegmentedColormap object>
```

```
    BLUE_GRADIENT = <matplotlib.colors.LinearSegmentedColormap object>
```

```
    GREEN_GRADIENT = <matplotlib.colors.LinearSegmentedColormap object>
```

```
    PURPLE_GRADIENT = <matplotlib.colors.LinearSegmentedColormap object>
```

```
    RAINBOW = <matplotlib.colors.LinearSegmentedColormap object>
```

```
    RED_GRADIENT = <matplotlib.colors.LinearSegmentedColormap object>
```

```
    YELLOW_GRADIENT = <matplotlib.colors.LinearSegmentedColormap object>
```

```
class orbit.constants.palette.OrbitPalette(value)
```

```
    Bases: enum.Enum
```

```
    str
```

```
    BLACK = '#000000'
```

```
    BLUE = '#276EF1'
```

```
    BLUE600 = '#174291'
```

```
    BROWN = '#99644C'
```

```
    GREEN = '#05A357'
```

```
    GREEN600 = '#03582F'
```

```
    ORANGE = '#ED6E33'
```

```
    PURPLE = '#7356BF'
```

```
    RED = '#E11900'
```

```
    WHITE = '#FFFFFF'
```

```
    YELLOW = '#FFC043'
```

```
    YELLOW400 = '#FFC043'
```

```
class orbit.constants.palette.PredictionPaletteClassic(value)
```

```
    Bases: enum.Enum
```

```
    str
```

```

ACTUAL_OBS = '#000000'
HOLDOUT_VERTICAL_LINE = '#000000'
PREDICTION_INTERVAL = '#276EF1'
PREDICTION_LINE = '#276EF1'
TEST_OBS = '#FFC043'

```

Module contents

orbit.diagnostics package

Submodules

orbit.diagnostics.plot module

```

orbit.diagnostics.plot.metric_horizon_barplot(df, model_col='model',
                                              pred_horizon_col='pred_horizon', metric_col='smape',
                                              bar_width=0.1, path=None, figsize=None,
                                              fontsize=None, is_visible=False)

```

```

orbit.diagnostics.plot.params_comparison_boxplot(data, var_names, model_names,
                                                color_list=[(0.12156862745098039,
0.4666666666666667, 0.7058823529411765), (1.0,
0.4980392156862745, 0.054901960784313725),
(0.17254901960784313, 0.6274509803921569,
0.17254901960784313), (0.8392156862745098,
0.15294117647058825, 0.1568627450980392),
(0.5803921568627451, 0.403921568627451,
0.7411764705882353), (0.5490196078431373,
0.33725490196078434, 0.29411764705882354),
(0.8901960784313725, 0.4666666666666667,
0.7607843137254902), (0.4980392156862745,
0.4980392156862745, 0.4980392156862745),
(0.7372549019607844, 0.7411764705882353,
0.13333333333333333), (0.09019607843137255,
0.7450980392156863, 0.8117647058823529)],
                                              title='Params Comparison', fig_size=(10, 6),
                                              box_width=0.1, box_distance=0.2,
                                              showfliers=False)

```

compare the distribution of parameters from different models using a boxplot. :param data: a list of dict with keys as the parameters of interest :param var_names: a list of strings, the labels of the parameters to compare :param model_names: a list of strings, the names of models to compare :param color_list: a list of strings, the color to use for differentiating models :param title: string

the title of the chart

Parameters

- **fig_size** – tuple figure size
- **box_width** – float width of the boxes in the boxplot
- **box_distance** – float the distance between each boxes in the boxplot

- **showfliers** – boolean show outliers in the chart if set as True

Returns a boxplot comparing parameter distributions from different models side by side

```
orbit.diagnostics.plot.plot_bt_predictions(bt_pred_df, metrics=<function smape>,
                                          split_key_list=None, ncol=2, figsize=None,
                                          include_vline=True, title="", fontsize=20, path=None,
                                          is_visible=True)
```

function to plot and visualize the prediction results from back testing.

bt_pred_df [data frame] the output of `orbit.diagnostics.backtest.BackTester.fit_predict()`, which includes the actuals/predictions for all the splits

metrics [callable] the metric function

split_key_list: list; default None with given model, which split keys to plot. If None, all the splits will be plotted

ncol [int] number of columns of the panel; number of rows will be decided accordingly

figsize [tuple] figure size

include_vline [bool] if plotting the vertical line to cut the in-sample and out-of-sample predictions for each split

title [str] title of the plot

fontsize: int; optional fontsize of the title

path [string] path to save the figure

is_visible [bool] if displaying the figure

```
orbit.diagnostics.plot.plot_bt_predictions2(bt_pred_df, metrics=<function smape>,
                                           split_key_list=None, figsize=None, include_vline=True,
                                           title="", fontsize=20, markersize=50, lw=2, fig_dir=None,
                                           is_visible=True, fix_xylim=True, export_gif=False)
```

a different style backtest plot compare to `plot_bt_prediction` where it writes separate plot for each split; this is also used to produce an animation to summarize every split

```
orbit.diagnostics.plot.plot_predicted_components(predicted_df, date_col,
                                                prediction_percentiles=None,
                                                plot_components=None, title="", figsize=None,
                                                path=None, fontsize=None, is_visible=True)
```

Plot predicted components with the data frame of decomposed prediction where components has been pre-defined as *trend*, *seasonality* and *regression*.

predicted_df [pd.DataFrame] predicted data response data frame. two columns required: `actual_col` and `pred_col`. If user provide `pred_percentiles_col`, it needs to include them as well.

date_col [str] the date column name

prediction_percentiles [list] a list should consist exact two elements which will be used to plot as lower and upper bound of confidence interval

plot_components [list] a list of strings to show the label of components to be plotted; by default, it uses values in `orbit.constants.constants.PredictedComponents`.

title [str; optional] title of the plot

figsize [tuple; optional] figsize pass through to `matplotlib.pyplot.figure()`

path [str; optional] path to save the figure

fontsize [int; optional] fontsize of the title

is_visible [boolean] whether we want to show the plot. If called from unittest, is_visible might = False.

Returns

Return type matplotlib axes object

```
orbit.diagnostics.plot.plot_predicted_data(training_actual_df, predicted_df, date_col, actual_col,
                                          pred_col='prediction', prediction_percentiles=None, title="",
                                          test_actual_df=None, is_visible=True, figsize=None,
                                          path=None, fontsize=None, line_plot=False,
                                          markersize=50, lw=2, linestyle='-')
```

plot training actual response together with predicted data; if actual response of predicted data is there, plot it too.

Parameters

- **training_actual_df** (*pd.DataFrame*) – training actual response data frame. two columns required: actual_col and date_col
- **predicted_df** (*pd.DataFrame*) – predicted data response data frame. two columns required: actual_col and pred_col. If user provide prediction_percentiles, it needs to include them as well in such *prediction_{x}* where x is the correspondent percentiles
- **prediction_percentiles** (*list*) – list of two elements indicates the lower and upper percentiles
- **date_col** (*str*) – the date column name
- **actual_col** (*str*) –
- **pred_col** (*str*) –
- **title** (*str*) – title of the plot
- **test_actual_df** (*pd.DataFrame*) – test actual response dataframe. two columns required: actual_col and date_col
- **is_visible** (*boolean*) – whether we want to show the plot. If called from unittest, is_visible might = False.
- **figsize** (*tuple*) – figsize pass through to *matplotlib.pyplot.figure()*
- **path** (*str*) – path to save the figure
- **fontsize** (*int; optional*) – fontsize of the title
- **line_plot** (*bool; default False*) – if True, make line plot for observations; otherwise, make scatter plot for observations
- **markersize** (*int; optional*) – point marker size
- **lw** (*int; optional*) – out-of-sample prediction line width
- **linestyle** (*str*) – linestyle of prediction plot

Returns

Return type matplotlib axes object

```
orbit.diagnostics.plot.residual_diagnostic_plot(df, dist='norm', date_col='week',
                                               residual_col='residual', fitted_col='prediction',
                                               sparams=None)
```

Parameters

- **df** (*pd.DataFrame*) –
- **dist** (*str*) –
- **date_col** (*str*) – column name of date
- **residual_col** (*str*) – column name of residual
- **fitted_col** (*str*) – column name of fitted value from model
- **sparams** (*float or list*) – extra parameters used in distribution such as t-dist

Notes

1. residual by time
2. residual vs fitted
3. residual histogram with vertical line as mean
4. residuals qq plot
5. residual ACF
6. residual PACF

Module contents

orbit.estimated package

Submodules

orbit.estimated.base_estimator module

class orbit.estimated.base_estimator.**BaseEstimator**(*seed=8888, verbose=True*)

Bases: object

Base Estimator class for both Stan and Pyro Estimator

Parameters

- **seed** (*int*) – seed number for initial random values
- **verbose** (*bool*) – If True (default), output all diagnostics messages from estimators

abstract fit(*model_name, model_param_names, data_input, fitter=None, init_values=None*)

Parameters

- **model_name** (*str*) – name of model - used in mapping the right sampling file (stan/pyro/...)
- **model_param_names** (*list*) – list of strings of model parameters names to extract
- **data_input** (*dict*) – key-value pairs of data input as required by definition in samplers (stan/pyro/...)
- **fitter** – model object used for fitting; this will be used instead of model_name if supplied to search for model object

- **init_values** (*float* or *np.array*) – initial sampler value. If None, ‘random’ is used

Returns

- **posteriors** (*dict*) – key value pairs where key is the model parameter name and value is *num_sample* x posterior values
- **training_metrics** (*dict*) – metrics and meta data related to the training process

orbit.estimated.pyro_estimator module

```
class orbit.estimated.pyro_estimator.PyroEstimator(num_steps=301, learning_rate=0.1,
                                                    learning_rate_total_decay=1.0, message=100,
                                                    **kwargs)
```

Bases: *orbit.estimated.base_estimator.BaseEstimator*

Abstract PyroEstimator with shared args for all PyroEstimator child classes

Parameters

- **num_steps** (*int*) – Number of estimator steps in optimization
- **learning_rate** (*float*) – Estimator learning rate
- **learning_rate_total_decay** (*float*) – A config re-parameterized from lrd in ClippedAdam. For example, 0.1 means a 90% reduction of the final step as of original learning rate where linear decay is implied along the steps. In the case of 1.0, no decay is applied. All steps will have the constant learning rate specified by *learning_rate*.
- **seed** (*int*) – Seed int
- **message** (*int*) – Print to console every *message* number of steps
- **kwargs** – Additional BaseEstimator args

Notes

See http://docs.pyro.ai/en/stable/_modules/pyro/optim/clipped_adam.html for optimizer details

```
abstract fit(model_name, model_param_names, data_input, fitter=None, init_values=None)
```

Parameters

- **model_name** (*str*) – name of model - used in mapping the right sampling file (stan/pyro/...)
- **model_param_names** (*list*) – list of strings of model parameters names to extract
- **data_input** (*dict*) – key-value pairs of data input as required by definition in samplers (stan/pyro/...)
- **fitter** – model object used for fitting; this will be used instead of model_name if supplied to search for model object
- **init_values** (*float* or *np.array*) – initial sampler value. If None, ‘random’ is used

Returns

- **posteriors** (*dict*) – key value pairs where key is the model parameter name and value is *num_sample* x posterior values
- **training_metrics** (*dict*) – metrics and meta data related to the training process

class orbit.estimatedors.pyro_estimator.**PyroEstimatorSVI**(*num_sample=100, num_particles=100, init_scale=0.1, **kwargs*)

Bases: *orbit.estimatedors.pyro_estimator.PyroEstimator*

Pyro Estimator for VI Sampling

Parameters

- **num_sample** (*int*) – Number of samples ot draw for inference, default 100
- **num_particles** (*int*) – Number of particles used in :class: *~pyro.infer.Trace_ELBO* for SVI optimization
- **init_scale** (*float*) – Parameter used in *pyro.infer.autoguide*; recommend a larger number of small dataset
- **kwargs** – Additional *PyroEstimator* class args

fit(*model_name, model_param_names, data_input, sampling_temperature, fitter=None, init_values=None*)

Parameters

- **model_name** (*str*) – name of model - used in mapping the right sampling file (*stan/pyro/...*)
- **model_param_names** (*list*) – list of strings of model parameters names to extract
- **data_input** (*dict*) – key-value pairs of data input as required by definition in samplers (*stan/pyro/...*)
- **fitter** – model object used for fitting; this will be used instead of *model_name* if supplied to search for model object
- **init_values** (*float or np.array*) – initial sampler value. If None, ‘random’ is used

Returns

- **posteriors** (*dict*) – key value pairs where key is the model parameter name and value is *num_sample* x posterior values
- **training_metrics** (*dict*) – metrics and meta data related to the training process

orbit.estimatedors.stan_estimator module

class orbit.estimatedors.stan_estimator.**StanEstimator**(*num_warmup=900, num_sample=100, chains=4, cores=8, algorithm=None, **kwargs*)

Bases: *orbit.estimatedors.base_estimator.BaseEstimator*

Abstract StanEstimator with shared args for all StanEstimator child classes

Parameters

- **num_warmup** (*int*) – Number of samples to warm up and to be discarded, default 900
- **num_sample** (*int*) – Number of samples to return, default 100
- **chains** (*int*) – Number of chains in stan sampler, default 4
- **cores** (*int*) – Number of cores for parallel processing, default *max(cores, multiprocessing.cpu_count())*
- **algorithm** (*str*) – If None, default to Stan defaults
- **kwargs** – Additional *BaseEstimator* class args

abstract fit(*model_name*, *model_param_names*, *data_input*, *fitter=None*, *init_values=None*)

Parameters

- **model_name** (*str*) – name of model - used in mapping the right sampling file (stan/pyro/...)
- **model_param_names** (*list*) – list of strings of model parameters names to extract
- **data_input** (*dict*) – key-value pairs of data input as required by definition in samplers (stan/pyro/...)
- **fitter** – model object used for fitting; this will be used instead of *model_name* if supplied to search for model object
- **init_values** (*float or np.array*) – initial sampler value. If None, ‘random’ is used

Returns

- **posteriors** (*dict*) – key value pairs where key is the model parameter name and value is *num_sample* x posterior values
- **training_metrics** (*dict*) – metrics and meta data related to the training process

class orbit.estimators.stan_estimator.**StanEstimatorMAP**(*stan_map_args=None*, ***kwargs*)

Bases: [orbit.estimators.stan_estimator.StanEstimator](#)

Stan Estimator for MAP Posteriors

Parameters **stan_map_args** (*dict*) – Supplemental stan vi args to pass to PyStan.optimizing()

fit(*model_name*, *model_param_names*, *data_input*, *fitter=None*, *init_values=None*)

Parameters

- **model_name** (*str*) – name of model - used in mapping the right sampling file (stan/pyro/...)
- **model_param_names** (*list*) – list of strings of model parameters names to extract
- **data_input** (*dict*) – key-value pairs of data input as required by definition in samplers (stan/pyro/...)
- **fitter** – model object used for fitting; this will be used instead of *model_name* if supplied to search for model object
- **init_values** (*float or np.array*) – initial sampler value. If None, ‘random’ is used

Returns

- **posteriors** (*dict*) – key value pairs where key is the model parameter name and value is *num_sample* x posterior values
- **training_metrics** (*dict*) – metrics and meta data related to the training process

class orbit.estimators.stan_estimator.**StanEstimatorMCMC**(*stan_mcmc_control=None*,
stan_mcmc_args=None, ***kwargs*)

Bases: [orbit.estimators.stan_estimator.StanEstimator](#)

Stan Estimator for MCMC Sampling

Parameters

- **stan_mcmc_control** (*dict*) – Supplemental stan control parameters to pass to PyStan.sampling()

- **stan_mcmc_args** (*dict*) – Supplemental stan mcmc args to pass to PyStan.sampling()

fit(*model_name*, *model_param_names*, *sampling_temperature*, *data_input*, *fitter=None*, *init_values=None*)

Parameters

- **model_name** (*str*) – name of model - used in mapping the right sampling file (stan/pyro/...)
- **model_param_names** (*list*) – list of strings of model parameters names to extract
- **data_input** (*dict*) – key-value pairs of data input as required by definition in samplers (stan/pyro/...)
- **fitter** – model object used for fitting; this will be used instead of model_name if supplied to search for model object
- **init_values** (*float or np.array*) – initial sampler value. If None, ‘random’ is used

Returns

- **posteriors** (*dict*) – key value pairs where key is the model parameter name and value is *num_sample* x posterior values
- **training_metrics** (*dict*) – metrics and meta data related to the training process

Module contents

orbit.models package

Submodules

orbit.models.ets module

`orbit.models.ets.ETS`(*seasonality=None*, *seasonality_sm_input=None*, *level_sm_input=None*, *estimator='stan-mcmc'*, ***kwargs*)

Parameters

- **seasonality** (*int*) – Length of seasonality
- **seasonality_sm_input** (*float*) – float value between [0, 1], applicable only if *seasonality* > 1. A larger value puts more weight on the current seasonality. If None, the model will estimate this value.
- **level_sm_input** (*float*) – float value between [0.0001, 1]. A larger value puts more weight on the current level. If None, the model will estimate this value.
- **estimator** (*string; {'stan-mcmc', 'stan-map'}*) – default to be ‘stan-mcmc’.
- **response_col** (*str*) – Name of response variable column, default ‘y’
- **date_col** (*str*) – Name of date variable column, default ‘ds’
- **n_bootstrap_draws** (*int*) – Number of samples to bootstrap in order to generate the prediction interval. For full Bayesian and variational inference forecasters, samples are drawn directly from original posteriors. For point-estimated posteriors, it will be used to sample

noise parameters. When -1 or None supplied, full Bayesian and variational inference forecasters will assume number of draws equal the size of original samples while point-estimated posteriors will mute the draw and output prediction without interval.

- **prediction_percentiles** (*list*) – List of integers of prediction percentiles that should be returned on prediction. To avoid reporting any confident intervals, pass an empty list
- ****kwargs** – additional arguments passed into `orbit.estimated.stan_estimator`

orbit.models.lgt module

```
orbit.models.lgt.LGT(seasonality=None, seasonality_sm_input=None, level_sm_input=None,
                    regressor_col=None, regressor_sign=None, regressor_beta_prior=None,
                    regressor_sigma_prior=None, regression_penalty='fixed_ridge', lasso_scale=0.5,
                    auto_ridge_scale=0.5, slope_sm_input=None, estimator='stan-mcmc', **kwargs)
```

Parameters

- **seasonality** (*int*) – Length of seasonality
- **seasonality_sm_input** (*float*) – float value between [0, 1], applicable only if *seasonality* > 1. A larger value puts more weight on the current seasonality. If None, the model will estimate this value.
- **level_sm_input** (*float*) – float value between [0.0001, 1]. A larger value puts more weight on the current level. If None, the model will estimate this value.
- **regressor_col** (*list*) – Names of regressor columns, if any
- **regressor_sign** (*list*) – list with values { '+', '-', '=' } such that '+' indicates regressor coefficient estimates are constrained to [0, inf). '-' indicates regressor coefficient estimates are constrained to (-inf, 0]. '=' indicates regressor coefficient estimates can be any value between (-inf, inf). The length of *regressor_sign* must be the same length as *regressor_col*. If None, all elements of list will be set to '='.
- **regressor_beta_prior** (*list*) – list of prior float values for regressor coefficient betas. The length of *regressor_beta_prior* must be the same length as *regressor_col*. If None, use non-informative priors.
- **regressor_sigma_prior** (*list*) – list of prior float values for regressor coefficient sigmas. The length of *regressor_sigma_prior* must be the same length as *regressor_col*. If None, use non-informative priors.
- **regression_penalty** ({ 'fixed_ridge', 'lasso', 'auto_ridge' }) – regression penalty method
- **lasso_scale** (*float*) – float value between [0, 1], applicable only if *regression_penalty* == 'lasso'
- **auto_ridge_scale** (*float*) – float value between [0, 1], applicable only if *regression_penalty* == 'auto_ridge'
- **slope_sm_input** (*float*) – float value between [0, 1]. A larger value puts more weight on the current slope. If None, the model will estimate this value.
- **estimator** (*string*; {'stan-mcmc', 'stan-map', 'pyro-svi'}) – default to be 'stan-mcmc'.
- **response_col** (*str*) – Name of response variable column, default 'y'

- **date_col** (*str*) – Name of date variable column, default ‘ds’
- **n_bootstrap_draws** (*int*) – Number of samples to bootstrap in order to generate the prediction interval. For full Bayesian and variational inference forecasters, samples are drawn directly from original posteriors. For point-estimated posteriors, it will be used to sample noise parameters. When -1 or None supplied, full Bayesian and variational inference forecasters will assume number of draws equal the size of original samples while point-estimated posteriors will mute the draw and output prediction without interval.
- **prediction_percentiles** (*list*) – List of integers of prediction percentiles that should be returned on prediction. To avoid reporting any confident intervals, pass an empty list
- ****kwargs** – additional arguments passed into `orbit.estimated.stan_estimator` or `orbit.estimated.pyro_estimator`

orbit.models.dlt module

```
orbit.models.dlt.DLT(seasonality=None, seasonality_sm_input=None, level_sm_input=None,
                    regressor_col=None, regressor_sign=None, regressor_beta_prior=None,
                    regressor_sigma_prior=None, regression_penalty='fixed_ridge', lasso_scale=0.5,
                    auto_ridge_scale=0.5, slope_sm_input=None, period=1, damped_factor=0.8,
                    global_trend_option='linear', global_cap=1.0, global_floor=0.0,
                    global_trend_sigma_prior=None, forecast_horizon=1, estimator='stan-mcmc',
                    **kwargs)
```

Parameters

- **seasonality** (*int*) – Length of seasonality
- **seasonality_sm_input** (*float*) – float value between [0, 1], applicable only if *seasonality* > 1. A larger value puts more weight on the current seasonality. If None, the model will estimate this value.
- **level_sm_input** (*float*) – float value between [0.0001, 1]. A larger value puts more weight on the current level. If None, the model will estimate this value.
- **regressor_col** (*list*) – Names of regressor columns, if any
- **regressor_sign** (*list*) – list with values { ‘+’, ‘-’, ‘=’ } such that ‘+’ indicates regressor coefficient estimates are constrained to [0, inf). ‘-’ indicates regressor coefficient estimates are constrained to (-inf, 0]. ‘=’ indicates regressor coefficient estimates can be any value between (-inf, inf). The length of *regressor_sign* must be the same length as *regressor_col*. If None, all elements of list will be set to ‘=’.
- **regressor_beta_prior** (*list*) – list of prior float values for regressor coefficient betas. The length of *regressor_beta_prior* must be the same length as *regressor_col*. If None, use non-informative priors.
- **regressor_sigma_prior** (*list*) – list of prior float values for regressor coefficient sigmas. The length of *regressor_sigma_prior* must be the same length as *regressor_col*. If None, use non-informative priors.
- **regression_penalty** (*{ 'fixed_ridge', 'lasso', 'auto_ridge' }*) – regression penalty method
- **lasso_scale** (*float*) – float value between [0, 1], applicable only if *regression_penalty* == ‘lasso’

- **auto_ridge_scale** (*float*) – float value between [0, 1], applicable only if *regression_penalty* == ‘auto_ridge’
- **slope_sm_input** (*float*) – float value between [0, 1]. A larger value puts more weight on the current slope. If None, the model will estimate this value.
- **period** (*int*) – Used to set *time_delta* as $1 / \max(\text{period}, \text{seasonality})$. If None and no seasonality, then *time_delta* == 1
- **damped_factor** (*float*) – Hyperparameter float value between [0, 1]. A smaller value further dampens the previous global trend value. Default, 0.8
- **global_trend_option** (*{ 'linear', 'loglinear', 'logistic', 'flat' }*) – Transformation function for the shape of the forecasted global trend.
- **global_cap** (*float*) – Maximum value of global logistic trend. Default is set to 1.0. This value is used only when *global_trend_option* = ‘logistic’
- **global_floor** (*float*) – Minimum value of global logistic trend. Default is set to 0.0. This value is used only when *global_trend_option* = ‘logistic’
- **global_trend_sigma_prior** (*sigma prior of the global trend; default uses 1 standard deviation of response*) –
- **forecast_horizon** (*int*) – forecast_horizon will be used only when users want to specify optimization forecast horizon > 1
- **estimator** (*string; { 'stan-mcmc', 'stan-map' }*) – default to be ‘stan-mcmc’.
- **response_col** (*str*) – Name of response variable column, default ‘y’
- **date_col** (*str*) – Name of date variable column, default ‘ds’
- **n_bootstrap_draws** (*int*) – Number of samples to bootstrap in order to generate the prediction interval. For full Bayesian and variational inference forecasters, samples are drawn directly from original posteriors. For point-estimated posteriors, it will be used to sample noise parameters. When -1 or None supplied, full Bayesian and variational inference forecasters will assume number of draws equal the size of original samples while point-estimated posteriors will mute the draw and output prediction without interval.
- **prediction_percentiles** (*list*) – List of integers of prediction percentiles that should be returned on prediction. To avoid reporting any confident intervals, pass an empty list
- ****kwargs** – additional arguments passed into `orbit.estimated.stan_estimator`

orbit.models.ktrlite module

```
orbit.models.ktrlite.KTRLite(level_knot_scale=0.1, level_segments=10, level_knot_distance=None,
                             level_knot_dates=None, seasonality=None, seasonality_fs_order=None,
                             seasonality_segments=2, seasonal_initial_knot_scale=1.0,
                             seasonal_knot_scale=0.1, degree_of_freedom=30, date_freq=None,
                             estimator='stan-map', **kwargs)
```

Parameters

- **level_knot_scale** (*float*) – sigma for level; default to be .1
- **level_segments** (*int*) – the number of segments partitioned by the knots of level (trend)
- **level_knot_distance** (*int*) – the distance between every two knots of level (trend)

- **level_knot_dates** (*array like*) – list of pre-specified dates for the level knots
- **seasonality** (*int, or list of int*) – multiple seasonality
- **seasonality_fs_order** (*int, or list of int*) – fourier series order for seasonality
- **seasonality_segments** (*int*) – the number of segments partitioned by the knots of seasonality
- **seasonal_initial_knot_scale** (*float*) – scale parameter for seasonal regressors initial coefficient knots; default to be 1
- **seasonal_knot_scale** (*float*) – scale parameter for seasonal regressors drift of coefficient knots; default to be 0.1.
- **degree_of_freedom** (*int*) – degree of freedom for error t-distribution
- **date_freq** (*str*) – date frequency; if not supplied, `pd.infer_freq` will be used to imply the date frequency.
- **estimator** (*string; {'stan-map'}*) –
- **response_col** (*str*) – Name of response variable column, default 'y'
- **date_col** (*str*) – Name of date variable column, default 'ds'
- **n_bootstrap_draws** (*int*) – Number of samples to bootstrap in order to generate the prediction interval. For full Bayesian and variational inference forecasters, samples are drawn directly from original posteriors. For point-estimated posteriors, it will be used to sample noise parameters. When -1 or None supplied, full Bayesian and variational inference forecasters will assume number of draws equal the size of original samples while point-estimated posteriors will mute the draw and output prediction without interval.
- **prediction_percentiles** (*list*) – List of integers of prediction percentiles that should be returned on prediction. To avoid reporting any confident intervals, pass an empty list
- ****kwargs** – additional arguments passed into `orbit.estimated.stan_estimator`

Module contents

orbit.pyro package

Submodules

orbit.pyro.lgt module

```
class orbit.pyro.lgt.Model(data)
    Bases: object
    max_plate_nesting = 1
```

Module contents

orbit.utils package

Submodules

orbit.utils.general module

orbit.utils.general.**expand_grid**(*base*)

Given a base key values span, expand them into a dataframe covering all combinations :param base: dictionary with keys equal columns name and value equals key values :type base: dict

Returns `pd.DataFrame`

Return type dataframe generate based on user specified base

orbit.utils.general.**get_parent_path**(*current_file_path*)

Parameters

- **current_file_path** (*str*) – The given file path, should be an absolute path
- **Returns** –
- ----- – *str* : The parent path of give file path

orbit.utils.general.**is_empty_dataframe**(*df*)

A simple function to tell whether the passed in df is an empty dataframe or not. :param df: given input dataframe :type df: `pd.DataFrame`

Returns `bool`

Return type True if df is none, or if df is an empty dataframe; False otherwise.

orbit.utils.general.**is_even_gap_datetime**(*array*)

Returns True if array is evenly distributed

orbit.utils.general.**is_ordered_datetime**(*array*)

Returns True if array is ordered and non-repetitive

orbit.utils.general.**regenerate_base_df**(*df, time_col, key_col, val_cols=[], fill_na=None*)

Given a dataframe, key column, time column and value column, re-generate multiple time-series to cover full range date-time with all the keys. This can be a useful utils for working multiple time-series.

Parameters

- **df** (*pd.DataFrame*) –
- **time_col** (*str*) –
- **key_col** (*str*) –
- **val_cols** (*List[str]*; values column considered to be imputed) –
- **fill_na** (*Optional[float]*; values to fill when there are missing values of the row) –

orbit.utils.general.**update_dict**(*original_dict, append_dict*)

orbit.utils.pyro module

`orbit.utils.pyro.get_pyro_model(model_name)`

orbit.utils.stan module

`orbit.utils.stan.compile_stan_model(stan_model_name)`

Compile stan model and save as pkl

`orbit.utils.stan.compile_stan_model_simplified(path)`

A more flexible way to load compile stan model with a path provided :param path:

`orbit.utils.stan.get_compiled_stan_model(stan_model_name)`

Load compiled Stan model

`orbit.utils.stan.get_compiled_stan_model_simplified(path)`

A more flexible way to load pre-compiled model :param path:

`orbit.utils.stan.set_compiled_stan_path(parent, child='stan_compiled')`

Set the path for compiled stan models.

parent: the primary directory level child: the secondary directory level

class `orbit.utils.stan.suppress_stdout_stderr`

Bases: object

A context manager for doing a “deep suppression” of stdout and stderr in Python, i.e. will suppress all print, even if the print originates in a compiled C/Fortran sub-function.

This will not suppress raised exceptions, since exceptions are printed

to stderr just before a script exits, and after the context manager has exited (at least, I think that is why it lets exceptions through).

Module contents

24.1.2 Submodules

24.1.3 orbit.exceptions module

exception `orbit.exceptions.AbstractMethodException`

Bases: Exception

exception `orbit.exceptions.BacktestException`

Bases: Exception

exception `orbit.exceptions.DataInputException`

Bases: Exception

exception `orbit.exceptions.EstimatorException`

Bases: Exception

exception `orbit.exceptions.ForecasterException`

Bases: Exception

exception `orbit.exceptions.IllegalArgument`

Bases: Exception

exception orbit.exceptions.**ModelException**
Bases: Exception

exception orbit.exceptions.**PlotException**
Bases: Exception

exception orbit.exceptions.**PredictionException**
Bases: Exception

24.1.4 orbit.orbit module

Top level Orbit class

class orbit.orbit.**Orbit**
Bases: object

24.1.5 Module contents

25.1 1.1.2 (2022-04-28) (release notes)

Core changes

- Add Conda installation option (#679)
- Suppress the lengthy Stan logging message (#696)
- WBIC for pyro SVI sampling and BIC for MAP optimization (#719, #710)
- Backtest module to include confidence intervals (#724)
- Allow configuration for compiled Stan model path (#713)
- Box plot for regression coefficient comparison (#737)
- Bounded logistic growth for DLT model (#712)
- Enhance regression output reporting (#739)

Documentation

- Add blacking linting to Github action workflow (#708)
- Tutorial enhancement

Utilities

- Add a new method *make_future_df* to prepare data frame for forecasting (#695)

25.2 1.1.2alpha (2022-04-06) (release notes)

Core changes

- Add Conda installation option (#679)
- Suppress the lengthy Stan logging message (#696)
- WBIC for pyro SVI sampling and BIC for MAP optimization (#719, #710)
- Backtest module to include confidence intervals (#724)
- Allow configuration for compiled Stan model path (#713)
- Box plot for regression coefficient comparison (#737)
- Bounded logistic growth for DLT model (#712)
- Enhance regression output reporting (#739)

Documentation

- Add blacking linting to Github action workflow (#708)
- Tutorial enhancement

Utilities

- Add a new method *make_future_df* to prepare data frame for forecasting (#695)

25.3 1.1.1 (2022-03-03) (release notes)

Core changes

- fix the mplstyle file path bug (#714)

25.4 1.1.0 (2022-01-11) (release notes)

Core changes

- Redesign the model class structure with three core components: model template, estimator, and forecaster (#506, #507, #508, #513)
- Introduce the Kernel-based Time-varying Regression (KTR) model (#515)
- Implement the negative coefficient for LGT and KTR (#600, #601, #609)
- Allow to handle missing values in response for LGT and DLT (#645)
- Implement WBIC value for model candidate selection (#654)

Documentation

- A new series of tutorials for KTR (#558, #559)
- Migrate the CI from TravisCI to Github Actions (#556)
- Missing value handle tutorial (#645)
- WBIC tutorial (#663)

Utilities

- New Plotting Palette (#571, #589)
- Redesign the diagnostic plotting (#581, #607)
- Raise a warning when date index is not evenly distributed (#639)

25.5 1.0.17 (2021-08-30) (release notes)

Core changes

- Use global mean instead of median in ktrx model before next major release

25.6 1.0.16 (2021-08-27) (release notes)

Core changes

- Bug fix and code improvement before next major release (#540, #541, #546)

25.7 1.0.15 (2021-08-02) (release notes)

Core changes

- Prediction functionality refactoring (#430)
- KTRLite model enhancement and interface cleanup (#440)
- More flexible scheduling config in Backtester (#447)
- Allow extraction of training related metrics (e.g. ELBO loss) in Pyro SVI (#443)
- Add a flag to keep the posterior samples or not in aggregated model (#465)
- Bug fix and code improvement (#428, #438, #459, #470)

Documentation

- Clean up and standardize example notebooks (#462)
- Tutorial update and enhancement (#431, #474)

Utilities

- Diagnostic plot with Arviz (#433)
- Refine plotting palette (#434, #473)
- Create an orbit-featured plotting style (#434)

25.8 1.0.13 (2021-04-02) (release notes)

Core changes

- Implement a new model KTRLite (#380)
- Refactoring of BaseTemplate (#382, #384)
- Add MAPTemplate, FullBayesianTemplate, and AggregatedPosteriorTemplate (#394)
- Remove dependency of scikit-learn (#379, #381)

Documentation

- Add changelogs, release process, and contribution guidance (#363, #369, #370, #372)
- Setup documentation deployment via TravisCI (#291)
- New tutorial of making your own model (#389)
- Tutorial enhancement (#383, #388)

Utilities

- New EDA plot utilities (#403, #407, #408)
- More options for existing plot utilities (#396)

25.9 1.0.12 (2021-02-19) (release notes)

- Documentation update (#354, #362)
- Providing prediction intervals for point posteriors such as AggregatedPosterior and MAP (#357, #359)
- Abstract classes created to refactor posteriors estimation as templates (#360)
- Automating documentation and tutorials; migrating docs to readthedocs (#291)

25.10 1.0.11 (2021-02-18) (release notes)

Core changes

- a simple ETS class is created (#280, #296)
- DLT is replacing LGT as the model used in the quick start and general demos (#305)
- DLT and LGT are refactored to inherit from ETS (#280)
- DLT now supports regression with strictly positive/negative signs (#296)
- deprecation on regression with LGT (#305)
- dependency update; remove enum34 and update other dependencies versions (#301)
- fixed pickle error (#342)

Documentation

- updated tutorials (#309, #329, #332)
- docstring cleanup with inherited classes (#350)

Utilities

- include the provide hyper-parameters tuning (#288)
- include dataloader with a few standard datasets (#352, #337, #277, #248)
- plotting functions now returns the plot object (#327, #325, #287, #279)

25.11 1.0.10 (2020-11-15) (Initial Release)

- dpl v2 for travis config (#295)

25.12 1.0.9 (2020-11-15)

- debug travis pypi deployment (#293)
- Debug travis package deployment (#294)

25.13 1.0.8 (2020-11-15)

- debug travis pypi deployment (#293)

25.14 1.0.7 (2020-11-14)

- #279
- reorder fourier series calculation to match the df (#286)
- plot utility enhancement (#287)
- Setup TravisCI deployment for PyPI (#292)

25.15 1.0.6 (2020-11-13)

- #251
- #257
- #259
- #263
- #248
- #264
- #265
- #270
- #273
- #277
- #281
- #282

INDICES AND TABLES

- [genindex](#)
- [modindex](#)
- [search](#)

PYTHON MODULE INDEX

O

- orbit, 167
- orbit.constants, 153
- orbit.constants.constants, 149
- orbit.constants.palette, 152
- orbit.diagnostics, 156
- orbit.diagnostics.plot, 153
- orbit.estimated, 160
- orbit.estimated.base_estimator, 156
- orbit.estimated.pyro_estimator, 157
- orbit.estimated.stan_estimator, 158
- orbit.exceptions, 166
- orbit.models, 164
- orbit.models.dlt, 162
- orbit.models.ets, 160
- orbit.models.ktrlite, 163
- orbit.models.lgt, 161
- orbit.orbit, 167
- orbit.pyro, 165
- orbit.pyro.lgt, 164
- orbit.utils, 166
- orbit.utils.general, 165
- orbit.utils.pyro, 166
- orbit.utils.stan, 166

A

AbstractMethodException, 166
 ACTUAL (*orbit.constants.constants.BacktestFitKeys* attribute), 149
 ACTUAL_OBS (*orbit.constants.palette.PredictionPaletteClassic* attribute), 152
 ACTUAL_RESPONSE (*orbit.constants.constants.PlotLabels* attribute), 150

B

BacktestException, 166
 BacktestFitKeys (*class in orbit.constants.constants*), 149
 BaseEstimator (*class in orbit.estimators.base_estimator*), 156
 BLACK (*orbit.constants.palette.OrbitPalette* attribute), 152
 BLACK_GRADIENT (*orbit.constants.palette.OrbitColorMap* attribute), 152
 BLUE (*orbit.constants.palette.OrbitPalette* attribute), 152
 BLUE600 (*orbit.constants.palette.OrbitPalette* attribute), 152
 BLUE_GRADIENT (*orbit.constants.palette.OrbitColorMap* attribute), 152
 BROWN (*orbit.constants.palette.OrbitPalette* attribute), 152

C

CHILD (*orbit.constants.constants.CompiledStanModelPath* attribute), 149
 compile_stan_model() (*in module orbit.utils.stan*), 166
 compile_stan_model_simplified() (*in module orbit.utils.stan*), 166
 CompiledStanModelPath (*class in orbit.constants.constants*), 149

D

DataInputException, 166
 DATE (*orbit.constants.constants.BacktestFitKeys* attribute), 149

DATE_ARRAY (*orbit.constants.constants.PredictionMetaKeys* attribute), 151
 DATE_ARRAY (*orbit.constants.constants.TrainingMetaKeys* attribute), 151
 DATE_COL (*orbit.constants.constants.TrainingMetaKeys* attribute), 151
 DLT() (*in module orbit.models.dlt*), 162

E

END (*orbit.constants.constants.PredictionMetaKeys* attribute), 151
 END (*orbit.constants.constants.TrainingMetaKeys* attribute), 151
 END_INDEX (*orbit.constants.constants.PredictionMetaKeys* attribute), 151
 EstimatorException, 166
 EstimatorsKeys (*class in orbit.constants.constants*), 150
 ETS() (*in module orbit.models.ets*), 160
 expand_grid() (*in module orbit.utils.general*), 165

F

fit() (*orbit.estimators.base_estimator.BaseEstimator* method), 156
 fit() (*orbit.estimators.pyro_estimator.PyroEstimator* method), 157
 fit() (*orbit.estimators.pyro_estimator.PyroEstimatorSVI* method), 158
 fit() (*orbit.estimators.stan_estimator.StanEstimator* method), 158
 fit() (*orbit.estimators.stan_estimator.StanEstimatorMAP* method), 159
 fit() (*orbit.estimators.stan_estimator.StanEstimatorMCMC* method), 160
 ForecasterException, 166
 FULL_SAMPLING (*orbit.constants.constants.PredictMethod* attribute), 150
 FUTURE_STEPS (*orbit.constants.constants.PredictionMetaKeys* attribute), 151

G

get_compiled_stan_model() (*in module or-*

- bit.utils.stan*), 166
- `get_compiled_stan_model_simplified()` (in module *orbit.utils.stan*), 166
- `get_parent_path()` (in module *orbit.utils.general*), 165
- `get_pyro_model()` (in module *orbit.utils.pyro*), 166
- GREEN (*orbit.constants.palette.OrbitPalette* attribute), 152
- GREEN600 (*orbit.constants.palette.OrbitPalette* attribute), 152
- GREEN_GRADIENT (*orbit.constants.palette.OrbitColorMap* attribute), 152
- ## H
- HOLDOUT_VERTICAL_LINE (*orbit.constants.palette.PredictionPaletteClassic* attribute), 153
- ## I
- IllegalArgument, 166
- `is_empty_dataframe()` (in module *orbit.utils.general*), 165
- `is_even_gap_datetime()` (in module *orbit.utils.general*), 165
- `is_ordered_datetime()` (in module *orbit.utils.general*), 165
- ## K
- KNOTS_REGION (*orbit.constants.palette.KTRPalette* attribute), 152
- KNOTS_SEGMENT (*orbit.constants.palette.KTRPalette* attribute), 152
- KTRLite() (in module *orbit.models.ktrlite*), 163
- KTRPalette (class in *orbit.constants.palette*), 152
- KTRTimePointPriorKeys (class in *orbit.constants.constants*), 150
- ## L
- LGT() (in module *orbit.models.lgt*), 161
- ## M
- MAP (*orbit.constants.constants.PredictMethod* attribute), 150
- `max_plate_nesting` (*orbit.pyro.lgt.Model* attribute), 164
- MEAN (*orbit.constants.constants.PredictMethod* attribute), 150
- MEDIAN (*orbit.constants.constants.PredictMethod* attribute), 150
- `metric_horizon_barplot()` (in module *orbit.diagnostics.plot*), 153
- METRIC_NAME (*orbit.constants.constants.BacktestFitKeys* attribute), 149
- METRIC_VALUES (*orbit.constants.constants.BacktestFitKeys* attribute), 149
- Model (class in *orbit.pyro.lgt*), 164
- MODEL (*orbit.constants.constants.TimeSeriesSplitSchemeKeys* attribute), 151
- ModelError, 166
- module
- orbit*, 167
 - orbit.constants*, 153
 - orbit.constants.constants*, 149
 - orbit.constants.palette*, 152
 - orbit.diagnostics*, 156
 - orbit.diagnostics.plot*, 153
 - orbit.estimateors*, 160
 - orbit.estimateors.base_estimator*, 156
 - orbit.estimateors.pyro_estimator*, 157
 - orbit.estimateors.stan_estimator*, 158
 - orbit.exceptions*, 166
 - orbit.models*, 164
 - orbit.models.dlt*, 162
 - orbit.models.ets*, 160
 - orbit.models.ktrlite*, 163
 - orbit.models.lgt*, 161
 - orbit.orbit*, 167
 - orbit.pyro*, 165
 - orbit.pyro.lgt*, 164
 - orbit.utils*, 166
 - orbit.utils.general*, 165
 - orbit.utils.pyro*, 166
 - orbit.utils.stan*, 166
- ## N
- NAME (*orbit.constants.constants.KTRTimePointPriorKeys* attribute), 150
- NUM_OF_OBS (*orbit.constants.constants.TrainingMetaKeys* attribute), 151
- ## O
- ORANGE (*orbit.constants.palette.OrbitPalette* attribute), 152
- orbit*
- module, 167
- Orbit (class in *orbit.orbit*), 167
- orbit.constants*
- module, 153
 - orbit.constants.constants* module, 149
 - orbit.constants.palette* module, 152
 - orbit.diagnostics* module, 156
 - orbit.diagnostics.plot* module, 153
 - orbit.estimateors* module, 160
 - orbit.estimateors.base_estimator*

- module, 156
 - orbit. estimators. pyro_estimator
 - module, 157
 - orbit. estimators. stan_estimator
 - module, 158
 - orbit. exceptions
 - module, 166
 - orbit. models
 - module, 164
 - orbit. models. dlt
 - module, 162
 - orbit. models. ets
 - module, 160
 - orbit. models. ktrlite
 - module, 163
 - orbit. models. lgt
 - module, 161
 - orbit. orbit
 - module, 167
 - orbit. pyro
 - module, 165
 - orbit. pyro. lgt
 - module, 164
 - orbit. utils
 - module, 166
 - orbit. utils. general
 - module, 165
 - orbit. utils. pyro
 - module, 166
 - orbit. utils. stan
 - module, 166
 - OrbitColorMap (class in orbit.constants.palette), 152
 - OrbitPalette (class in orbit.constants.palette), 152
- ## P
- params_comparison_boxplot() (in module orbit.diagnostics.plot), 153
 - PARENT (orbit.constants.constants.CompiledStanModelPath attribute), 149
 - plot_bt_predictions() (in module orbit.diagnostics.plot), 154
 - plot_bt_predictions2() (in module orbit.diagnostics.plot), 154
 - plot_predicted_components() (in module orbit.diagnostics.plot), 154
 - plot_predicted_data() (in module orbit.diagnostics.plot), 155
 - PlotException, 167
 - PlotLabels (class in orbit.constants.constants), 150
 - PREDICTED (orbit.constants.constants.BacktestFitKeys attribute), 149
 - PREDICTED_RESPONSE (orbit.constants.constants.PlotLabels attribute), 150
 - PREDICTION (orbit.constants.constants.PredictionKeys attribute), 150
 - PREDICTION_DF_LEN (orbit.constants.constants.PredictionMetaKeys attribute), 151
 - PREDICTION_INTERVAL (orbit.constants.palette.PredictionPaletteClassic attribute), 153
 - PREDICTION_LINE (orbit.constants.palette.PredictionPaletteClassic attribute), 153
 - PredictionException, 167
 - PredictionKeys (class in orbit.constants.constants), 150
 - PredictionMetaKeys (class in orbit.constants.constants), 150
 - PredictionPaletteClassic (class in orbit.constants.palette), 152
 - PredictMethod (class in orbit.constants.constants), 150
 - PRIOR_END_TP_IDX (orbit.constants.constants.KTRTimePointPriorKeys attribute), 150
 - PRIOR_MEAN (orbit.constants.constants.KTRTimePointPriorKeys attribute), 150
 - PRIOR_REGRESSOR_COL (orbit.constants.constants.KTRTimePointPriorKeys attribute), 150
 - PRIOR_SD (orbit.constants.constants.KTRTimePointPriorKeys attribute), 150
 - PRIOR_START_TP_IDX (orbit.constants.constants.KTRTimePointPriorKeys attribute), 150
 - PURPLE (orbit.constants.palette.OrbitPalette attribute), 152
 - PURPLE_GRADIENT (orbit.constants.palette.OrbitColorMap attribute), 152
 - PyroEstimator (class in orbit. estimators. pyro_estimator), 157
 - PyroEstimatorSVI (class in orbit. estimators. pyro_estimator), 157
 - PyroSVI (orbit.constants.constants.EstimatorsKeys attribute), 150
- ## R
- RAINBOW (orbit.constants.palette.OrbitColorMap attribute), 152
 - RED (orbit.constants.palette.OrbitPalette attribute), 152
 - RED_GRADIENT (orbit.constants.palette.OrbitColorMap attribute), 152
 - regenerate_base_df() (in module orbit.utils.general), 165
 - REGRESSION (orbit.constants.constants.PredictionKeys attribute), 150

- REGRESSOR (*orbit.constants.constants.PredictionKeys* attribute), 150
- residual_diagnostic_plot() (in module *orbit.diagnostics.plot*), 155
- RESPONSE (*orbit.constants.constants.TrainingMetaKeys* attribute), 151
- RESPONSE_COL (*orbit.constants.constants.TrainingMetaKeys* attribute), 151
- RESPONSE_MEAN (*orbit.constants.constants.TrainingMetaKeys* attribute), 151
- RESPONSE_SD (*orbit.constants.constants.TrainingMetaKeys* attribute), 151
- S**
- SEASONALITY (*orbit.constants.constants.PredictionKeys* attribute), 150
- set_compiled_stan_path() (in module *orbit.utils.stan*), 166
- SPLIT_KEY (*orbit.constants.constants.BacktestFitKeys* attribute), 149
- SPLIT_TYPE_EXPANDING (or-
bit.constants.constants.TimeSeriesSplitSchemeKeys attribute), 151
- SPLIT_TYPE_ROLLING (or-
bit.constants.constants.TimeSeriesSplitSchemeKeys attribute), 151
- StanEstimator (class in or-
bit.estimators.stan_estimator), 158
- StanEstimatorMAP (class in or-
bit.estimators.stan_estimator), 159
- StanEstimatorMCMC (class in or-
bit.estimators.stan_estimator), 159
- StanMAP (*orbit.constants.constants.EstimatorsKeys* attribute), 150
- StanMCMC (*orbit.constants.constants.EstimatorsKeys* attribute), 150
- START (*orbit.constants.constants.PredictionMetaKeys* attribute), 151
- START (*orbit.constants.constants.TrainingMetaKeys* attribute), 151
- START_INDEX (*orbit.constants.constants.PredictionMetaKeys* attribute), 151
- suppress_stdout_stderr (class in *orbit.utils.stan*), 166
- T**
- TEST_ACTUAL (*orbit.constants.constants.BacktestFitKeys* attribute), 149
- TEST_IDX (*orbit.constants.constants.TimeSeriesSplitSchemeKeys* attribute), 151
- TEST_OBS (*orbit.constants.palette.PredictionPaletteClassic* attribute), 153
- TEST_PREDICTED (*orbit.constants.constants.BacktestFitKeys* attribute), 149
- TimeSeriesSplitSchemeKeys (class in or-
bit.constants.constants), 151
- TRAIN_ACTUAL (*orbit.constants.constants.BacktestFitKeys* attribute), 149
- TRAIN_END_DATE (*orbit.constants.constants.TimeSeriesSplitSchemeKeys* attribute), 151
- TRAIN_FLAG (*orbit.constants.constants.BacktestFitKeys* attribute), 149
- TRAIN_IDX (*orbit.constants.constants.TimeSeriesSplitSchemeKeys* attribute), 151
- TRAIN_METRIC_FLAG (or-
bit.constants.constants.BacktestFitKeys attribute), 149
- TRAIN_PREDICTED (or-
bit.constants.constants.BacktestFitKeys attribute), 149
- TRAIN_START_DATE (or-
bit.constants.constants.TimeSeriesSplitSchemeKeys attribute), 151
- TRAINING_ACTUAL_RESPONSE (or-
bit.constants.constants.PlotLabels attribute), 150
- TrainingMetaKeys (class in *orbit.constants.constants*), 151
- TREND (*orbit.constants.constants.PredictionKeys* attribute), 150
- U**
- update_dict() (in module *orbit.utils.general*), 165
- W**
- WHITE (*orbit.constants.palette.OrbitPalette* attribute), 152
- Y**
- YELLOW (*orbit.constants.palette.OrbitPalette* attribute), 152
- YELLOW400 (*orbit.constants.palette.OrbitPalette* attribute), 152
- YELLOW_GRADIENT (or-
bit.constants.palette.OrbitColorMap attribute), 152