RdTools

Release 2.0.0+0.gcac6c35.dirty

Oct 20, 2020

Contents

RdTools

RdTools is an open-source library to support reproducible technical analysis of time series data from photovoltaic energy systems. The library aims to provide best practice analysis routines along with the building blocks for users to tailor their own analyses. Current applications include the evaluation of PV production over several years to obtain rates of performance degradation and soiling loss. They also include the capability to analyze systems for system- and subsystem-level availability. RdTools can handle both high frequency (hourly or better) or low frequency (daily, weekly, etc.) datasets. Best results are obtained with higher frequency data.

Full examples are worked out in the example notebooks in the *example notebook*.

To report issues, contribute code, or suggest improvements to this documentation, visit the RdTools development repository on github.

CHAPTER 1

Degradation and Soiling

Both degradation and soiling analyses are based on normalized yield, similar to performance index. Usually, this is computed at the daily level although other aggregation periods are supported. A typical analysis of soiling and degradation contains the following:

0. Import and preliminary calculations

1. Normalize data using a performance metric

2. Filter data that creates bias

3. Aggregate data

4. Analyze aggregated data to estimate the degradation rate and/or soiling loss

Steps 1 and 2 may be accomplished with the clearsky workflow (see the *example notebook*) which can help eliminate problems from irradiance sensor drift.

Normalize

Clearsky normalize

Filter

Clearsky filter

Aggregate

Degradation

Soiling
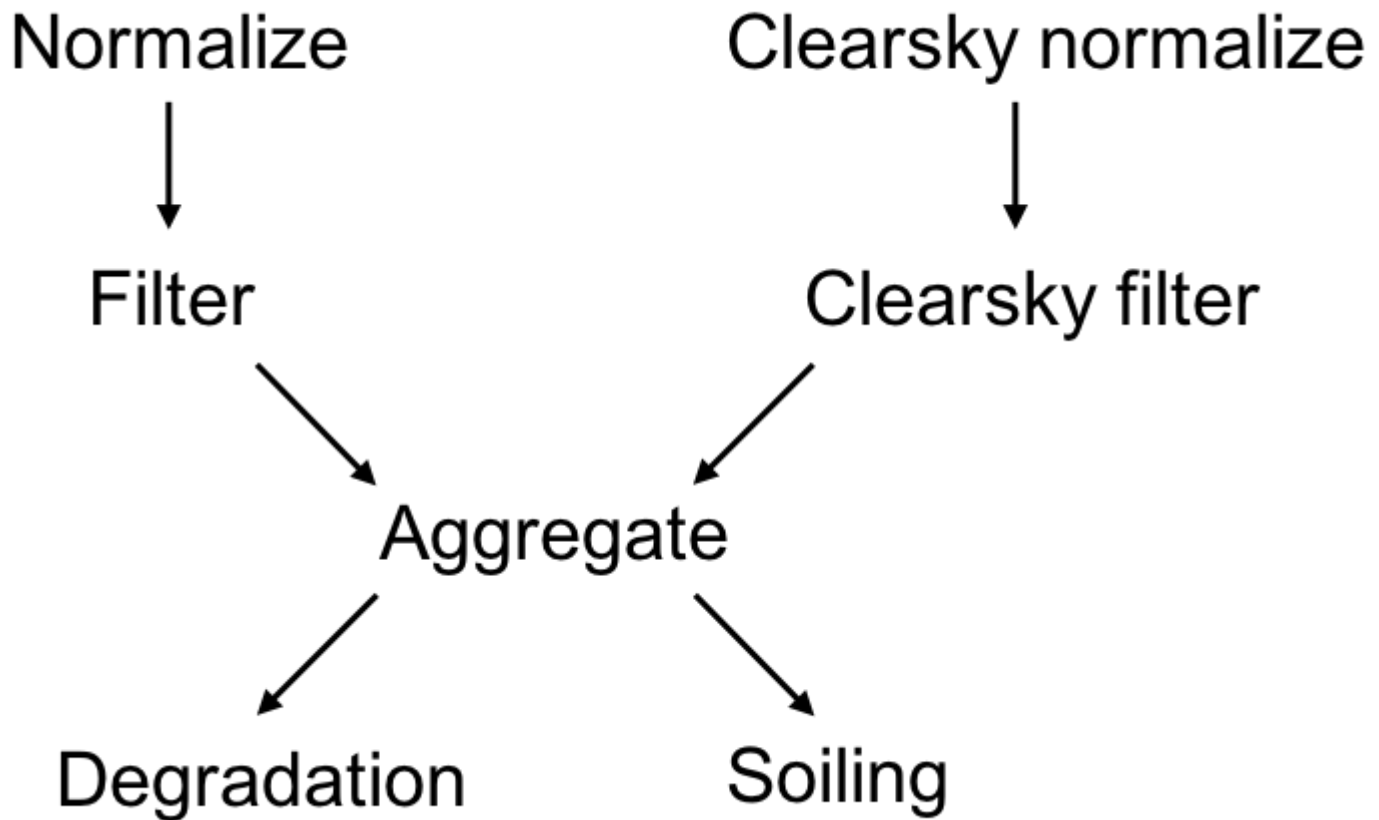
1.1 Degradation

The preferred method for degradation rate estimation is the year-on-year (YOY) approach (Jordan 2018), available in
degradation.degradation_year_on_year(). The YOY calculation yields in a distribution of degradation
rates, the central tendency of which is the most representative of the true degradation. The width of the distribution
provides information about the uncertainty in the estimate via a bootstrap calculation. The *example notebook* uses the
output of *degradation.degradation_year_on_year()* to visualize the calculation.
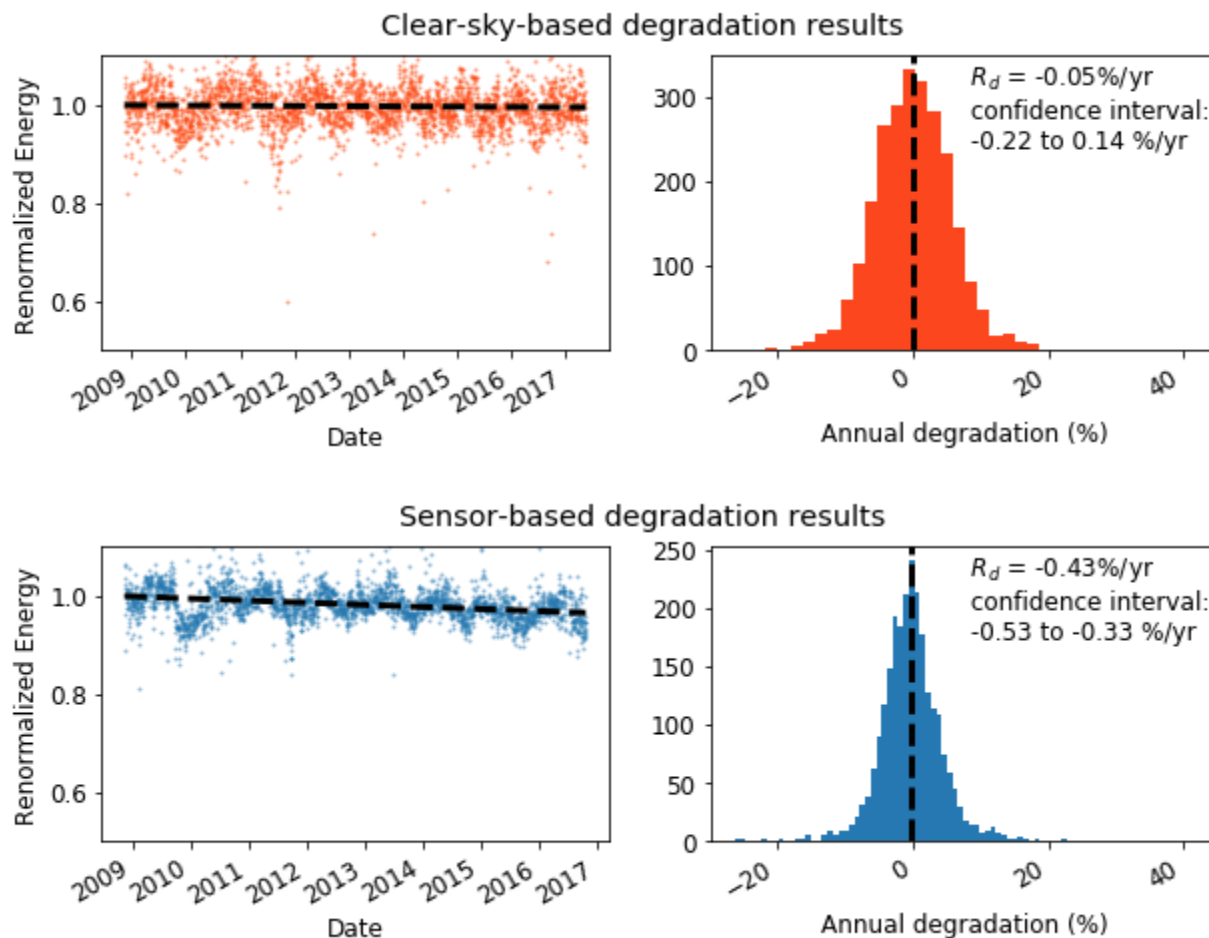
Two workflows are available for system performance ratio calculation, and illustrated in an example notebook. The sensor-based approach assumes that site irradiance and temperature sensors are calibrated and in good repair. Since this is not always the case, a 'clear-sky' workflow is provided that is based on modeled temperature and irradiance. Note that site irradiance data is still required to identify clear-sky conditions to be analyzed. In many cases, the 'clear-sky' analysis can identify conditions of instrument errors or irradiance sensor drift, such as in the above analysis.

The clear-sky analysis tends to provide less stable results than sensor-based analysis when details such as filtering are changed. We generally recommend that the clear-sky analysis be used as a check on the sensor-based results, rather than as a stand-alone analysis.

1.2 Soiling

Soiling can be estimated with the stochastic rate and recovery (SRR) method (Deceglie 2018). This method works well when soiling patterns follow a "sawtooth" pattern, a linear decline followed by a sharp recovery associated with natural or manual cleaning. `soiling.soiling_srr()` performs the calculation and returns the P50 insolation-weighted soiling ratio, confidence interval, and additional information (`soiling_info`) which includes a summary of the soiling intervals identified, `soiling_info['soiling_interval_summary']`. This summary table can, for example, be used to plot a histogram of the identified soiling rates for the dataset.
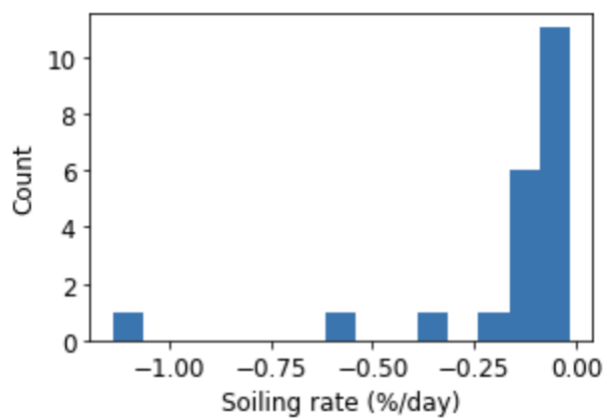
Availability

Evaluating system availability can be confounded by data loss from interrupted datalogger or system communications. RdTools implements two methods (Anderson & Blumenthal 2020) of distinguishing nuisance communication interruptions from true production outages with the *availability.AvailabilityAnalysis* class. In addition to classifying data outages, it estimates lost production and calculates energy-weighted system availability.
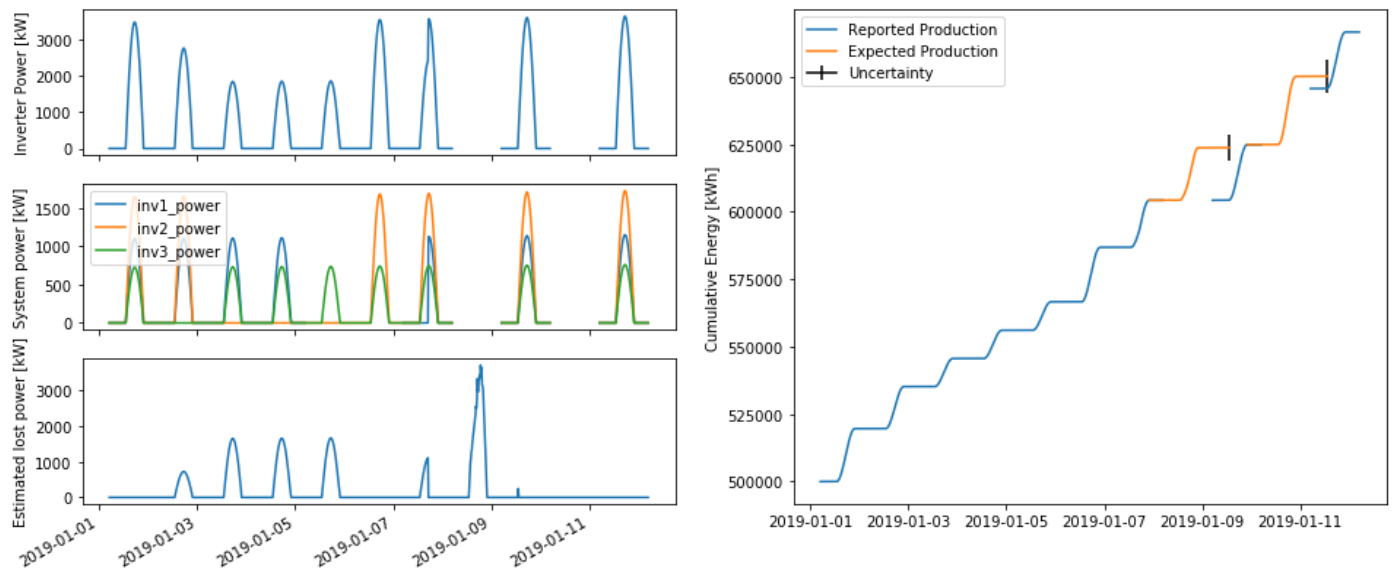
CHAPTER 3

Install RdTools using pip

RdTools can be installed automatically into Python from PyPI using the command line:

```
pip install rdtools
```

Alternatively it can be installed manually using the command line:

1. Download a release (Or to work with a development version, clone or download the rdtools repository).

2. Navigate to the repository: `cd rdtools`

3. Install via pip: `pip install .`

On some systems installation with `pip` can fail due to problems installing requirements. If this occurs, the requirements specified in `setup.py` may need to be separately installed (for example by using `conda`) before installing `rdtools`.

For more detailed instructions, see the *Developer Notes* page.

RdTools currently is tested on Python 3.6+.

CHAPTER 4

Usage and examples

Full workflow examples are found in the notebooks in *example notebook*. The examples are designed to work with python 3.7. For a consistent experience, we recommend installing the packages and versions documented in `docs/notebook_requirements.txt`. This can be achieved in your environment by first installing RdTools as described above, then running `pip install -r docs/notebook_requirements.txt` from the base directory.

The following functions are used for degradation and soiling analysis:

```python
import rdtools
```

The most frequently used functions are:

```python
normalization.normalize_with_expected_power(pv, power_expected, poa_global,
                                            pv_input='power')
  '''
  Inputs: Pandas time series of raw power or energy, expected power, and
    plane of array irradiance.
  Outputs: Pandas time series of normalized energy and POA insolation
  '''
```

```python
filtering.poa_filter(poa_global); filtering.tcell_filter(temperature_cell);
filtering.clip_filter(power_ac); filtering.normalized_filter(energy_normalized);
filtering.csi_filter(poa_global_measured, poa_global_clearsky);
  '''
  Inputs: Pandas time series of raw data to be filtered.
  Output: Boolean mask where `True` indicates acceptable data
  '''
```

```python
aggregation.aggregation_insol(energy_normalized, insolation, frequency='D')
  '''
  Inputs: Normalized energy and insolation
  Output: Aggregated data, weighted by the insolation.
  '''
```

```
degradation.degradation_year_on_year(energy_normalized)
  '''
  Inputs: Aggregated, normalized, filtered time series data
  Outputs: Tuple: `yoy_rd`: Degradation rate
    `yoy_ci`: Confidence interval `yoy_info`: associated analysis data
  '''
```

```
soiling.soiling_srr(energy_normalized_daily, insolation_daily)
  '''
  Inputs: Daily aggregated, normalized, filtered time series data for normalized␣
→performance and insolation
  Outputs: Tuple: `sr`: Insolation-weighted soiling ratio
    `sr_ci`: Confidence interval `soiling_info`: associated analysis data
  '''
```

```
availability.AvailabilityAnalysis(power_system, power_subsystem,
                                  energy_cumulative, power_expected)
  '''
  Inputs: Pandas time series system and subsystem power and energy data
  Outputs: DataFrame of production loss and availability metrics
  '''
```

Citing RdTools

The underlying workflow of RdTools has been published in several places. If you use RdTools in a published work, please cite the following as appropriate:

- D. Jordan, C. Deline, S. Kurtz, G. Kimball, M. Anderson, "Robust PV Degradation Methodology and Application", IEEE Journal of Photovoltaics, 8(2) pp. 525-531, 2018

- M. G. Deceglie, L. Micheli and M. Muller, "Quantifying Soiling Loss Directly From PV Yield," in IEEE Journal of Photovoltaics, 8(2), pp. 547-551, 2018

- K. Anderson and R. Blumenthal, "Overcoming Communications Outages in Inverter Downtime Analysis", 2020 IEEE 47th Photovoltaic Specialists Conference (PVSC)."

- RdTools, version x.x.x, https://github.com/NREL/rdtools, https://doi.org/10.5281/zenodo.1210316

    - Be sure to include the version number used in your analysis!

References

- The clear sky temperature calculation, *clearsky_temperature.get_clearsky_tamb()*, uses data from images created by Jesse Allen, NASA's Earth Observatory using data courtesy of the MODIS Land Group.

    - https://neo.sci.gsfc.nasa.gov/view.php?datasetId=MOD_LSTD_CLIM_M

    - https://neo.sci.gsfc.nasa.gov/view.php?datasetId=MOD_LSTN_CLIM_M

Other useful references which may also be consulted for degradation rate methodology include:

- D. C. Jordan, M. G. Deceglie, S. R. Kurtz, "PV degradation methodology comparison — A basis for a standard", in 43rd IEEE Photovoltaic Specialists Conference, Portland, OR, USA, 2016, DOI: 10.1109/PVSC.2016.7749593.

- Jordan DC, Kurtz SR, VanSant KT, Newmiller J, Compendium of Photovoltaic Degradation Rates, Progress in Photovoltaics: Research and Application, 2016, 24(7), 978 - 989.

- D. Jordan, S. Kurtz, PV Degradation Rates – an Analytical Review, Progress in Photovoltaics: Research and Application, 2013, 21(1), 12 - 29.

- E. Hasselbrink, M. Anderson, Z. Defreitas, M. Mikofski, Y.-C.Shen, S. Caldwell, A. Terao, D. Kavulak, Z. Campeau, D. DeGraaff, "Validation of the PVLife model using 3 million module-years of live site data", 39th IEEE Photovoltaic Specialists Conference, Tampa, FL, USA, 2013, p. 7 – 13, DOI: 10.1109/PVSC.2013.6744087.

Documentation Contents

7.1 Degradation and soiling example with clearsky workflow

This jupyter notebook is intended to the RdTools analysis workflow. In addition, the notebook demonstrates the effects of changes in the workflow. For a consistent experience, we recommend installing the specific versions of packages used to develop this notebook. This can be achieved in your environment by running `pip install -r requirements.txt` followed by `pip install -r docs/notebook_requirements.txt` from the base directory. (RdTools must also be separately installed.) These environments and examples are tested with Python 3.7.

The calculations consist of several steps illustrated here:

Import and preliminary calculations

Normalize data using a performance metric

Filter data that creates bias

Aggregate data

Analyze aggregated data to estimate the degradation rate

Analyze aggregated data to estimate the soiling loss

After demonstrating these steps using sensor data, a modified version of the workflow is illustrated using modeled clear sky irradiance and temperature. The results from the two methods are compared at the end.

This notebook works with data from the NREL PVDAQ `[4] NREL x-Si #1` system. Note that because this system does not experience significant soiling, the dataset contains a synthesized soiling signal for use in the soiling section of the example. This notebook automatically downloads and locally caches the dataset used in this example. The data can also be found on the DuraMAT Datahub (https://datahub.duramat.org/dataset/pvdaq-time-series-with-soiling-signal).

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

(continues on next page)

```
import pvlib
import rdtools
%matplotlib inline
```

```
[2]: #Update the style of plots
import matplotlib
matplotlib.rcParams.update({'font.size': 12,
                            'figure.figsize': [4.5, 3],
                            'lines.markeredgewidth': 0,
                            'lines.markersize': 2
                            })
# Register time series plotting in pandas > 1.0
from pandas.plotting import register_matplotlib_converters
register_matplotlib_converters()
```

```
[3]: # Set the random seed for numpy to ensure consistent results
np.random.seed(0)
```

7.1.1 0: Import and preliminary calculations

This section prepares the data necessary for an `rdtools` calculation. The first step of the `rdtools` workflow is normalization, which requires a time series of energy yield, a time series of cell temperature, and a time series of irradiance, along with some metadata (see Step 1: Normalize)

The following section loads the data, adjusts units where needed, and renames the critical columns. The ambient temperature sensor data source is converted into estimated cell temperature. This dataset already has plane-of-array irradiance data, so no transposition is necessary.

A common challenge is handling datasets with and without daylight savings time. Make sure to specify a `pytz` timezone that does or does not include daylight savings time as appropriate for your dataset.

The steps of this section may change depending on your data source or the system being considered. Transposition of irradiance and modeling of cell temperature are generally outside the scope of `rdtools`. A variety of tools for these calculations are available in pvlib.

```
[4]: # Import the example data
file_url = ('https://datahub.duramat.org/dataset/a49bb656-7b36-'
            '437a-8089-1870a40c2a7d/resource/5059bc22-640d-4dd4'
            '-b7b1-1e71da15be24/download/pvdaq_system_4_2010-2016'
            '_subset_soilsignal.csv')
cache_file = 'PVDAQ_system_4_2010-2016_subset_soilsignal.pickle'

try:
    df = pd.read_pickle(cache_file)
except FileNotFoundError:
    df = pd.read_csv(file_url, index_col=0, parse_dates=True)
    df.to_pickle(cache_file)

df = df.rename(columns = {
    'ac_power':'power_ac',
    'wind_speed': 'wind_speed',
    'ambient_temp': 'Tamb',
    'poa_irradiance': 'poa',
})
```

```python
# Specify the Metadata
meta = {"latitude": 39.7406,
        "longitude": -105.1774,
        "timezone": 'Etc/GMT+7',
        "gamma_pdc": -0.005,
        "azimuth": 180,
        "tilt": 40,
        "power_dc_rated": 1000.0,
        "temp_model_params":
        pvlib.temperature.TEMPERATURE_MODEL_PARAMETERS['sapm']['open_rack_glass_
→polymer']}

df.index = df.index.tz_localize(meta['timezone'])

loc = pvlib.location.Location(meta['latitude'], meta['longitude'], tz = meta['timezone
→'])
sun = loc.get_solarposition(df.index)

# There is some missing data, but we can infer the frequency from
# the first several data points
freq = pd.infer_freq(df.index[:10])

# Then set the frequency of the dataframe.
# It is recommended not to up- or downsample at this step
# but rather to use interpolate to regularize the time series
# to its dominant or underlying frequency. Interpolate is not
# generally recommended for downsampling in this application.
df = rdtools.interpolate(df, freq)

# Calculate cell temperature
df['Tcell'] = pvlib.temperature.sapm_cell(df.poa, df.Tamb,
                                          df.wind_speed, **meta['temp_model_params'])

# plot the AC power time series
fig, ax = plt.subplots(figsize=(4,3))
ax.plot(df.index, df.power_ac, 'o', alpha=0.01)
ax.set_ylim(0,1500)
fig.autofmt_xdate()
ax.set_ylabel('AC Power (W)');
```
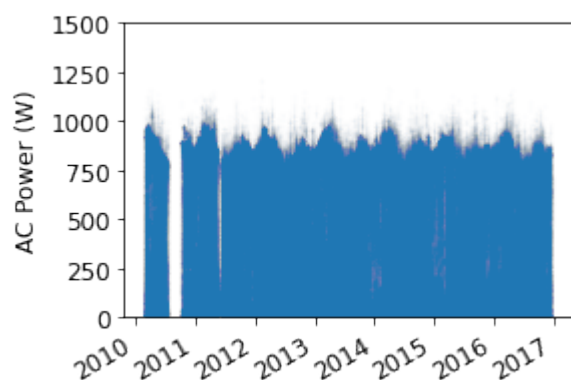
7.1.2 1: Normalize

Data normalization is achieved with `rdtools.normalize_with_expected_power()`. This function can be used to normalize to any modeled or expected power. Note that realized PV output can be given as energy, rather than power, by using an optional key word argument.

```
[5]: # Calculate the expected power with a simple PVWatts DC model
     modeled_power = pvlib.pvsystem.pvwatts_dc(df['poa'], df['Tcell'], meta['power_dc_rated
     ↪'],
                                               meta['gamma_pdc'], 25.0 )

     # Calculate the normalization, the function also returns the relevant insolation for
     # each point in the normalized PV energy timeseries
     normalized, insolation = rdtools.normalize_with_expected_power(df['power_ac'],
                                                                    modeled_power,
                                                                    df['poa'])

     df['normalized'] = normalized
     df['insolation'] = insolation

     # Plot the normalized power time series
     fig, ax = plt.subplots()
     ax.plot(normalized.index, normalized, 'o', alpha = 0.05)
     ax.set_ylim(0,2)
     fig.autofmt_xdate()
     ax.set_ylabel('Normalized energy');
```
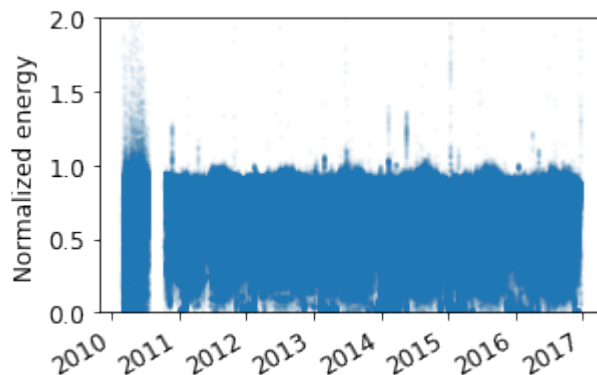

7.1.3 2: Filter

Data filtering is used to exclude data points that represent invalid data, create bias in the analysis, or introduce significant noise.

It can also be useful to remove outages and outliers. Sometimes outages appear as low but non-zero yield. Automatic functions for outage detection are not yet included in `rdtools`. However, this example does filter out data points where the normalized energy is less than 1%. System-specific filters should be implemented by the analyst if needed.

```
[6]: # Calculate a collection of boolean masks that can be used
     # to filter the time series
     normalized_mask = rdtools.normalized_filter(df['normalized'])
     poa_mask = rdtools.poa_filter(df['poa'])
     tcell_mask = rdtools.tcell_filter(df['Tcell'])
     # Note: This clipping mask may be disabled when you are sure the system is not
```

(continues on next page)

```python
# experiencing clipping due to high DC/AC ratio
clip_mask = rdtools.clip_filter(df['power_ac'])

# filter the time series and keep only the columns needed for the
# remaining steps
filtered = df[normalized_mask & poa_mask & tcell_mask & clip_mask]
filtered = filtered[['insolation', 'normalized']]

fig, ax = plt.subplots()
ax.plot(filtered.index, filtered.normalized, 'o', alpha = 0.05)
ax.set_ylim(0,2)
fig.autofmt_xdate()
ax.set_ylabel('Normalized energy');
```
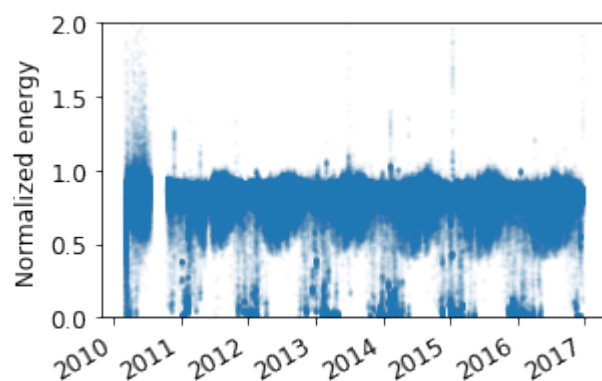

7.1.4 3: Aggregate

Data is aggregated with an irradiance weighted average. This can be useful, for example with daily aggregation, to reduce the impact of high-error data points in the morning and evening.

```python
[7]: daily = rdtools.aggregation_insol(filtered.normalized, filtered.insolation,
                                        frequency = 'D')

fig, ax = plt.subplots()
ax.plot(daily.index, daily, 'o', alpha = 0.1)
ax.set_ylim(0,2)
fig.autofmt_xdate()
ax.set_ylabel('Normalized energy');
```

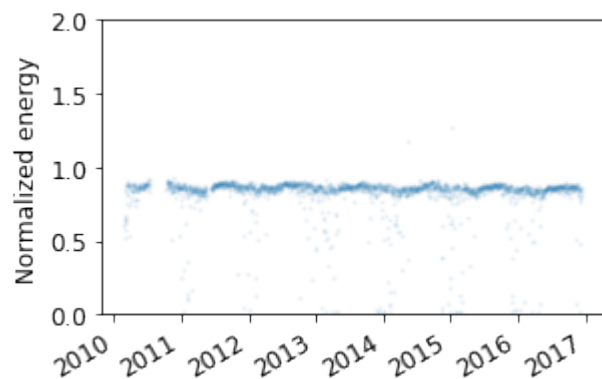### 7.1.5 4: Degradation calculation

Data is then analyzed to estimate the degradation rate representing the PV system behavior. The results are visualized and statistics are reported, including the 68.2% confidence interval, and the P95 exceedance value.

```
[8]:  # Calculate the degradation rate using the YoY method
      yoy_rd, yoy_ci, yoy_info = rdtools.degradation_year_on_year(daily, confidence_
      →level=68.2)
      # Note the default confidence_level of 68.2 is appropriate if you would like to
      # report a confidence interval analogous to the standard deviation of a normal
      # distribution. The size of the confidence interval is adjustable by setting the
      # confidence_level variable.

      # Visualize the results

      degradation_fig = rdtools.degradation_summary_plots(
          yoy_rd, yoy_ci, yoy_info, daily,
          summary_title='Sensor-based degradation results',
          scatter_ymin=0.5, scatter_ymax=1.1,
          hist_xmin=-30, hist_xmax=45, bins=100
      )
```
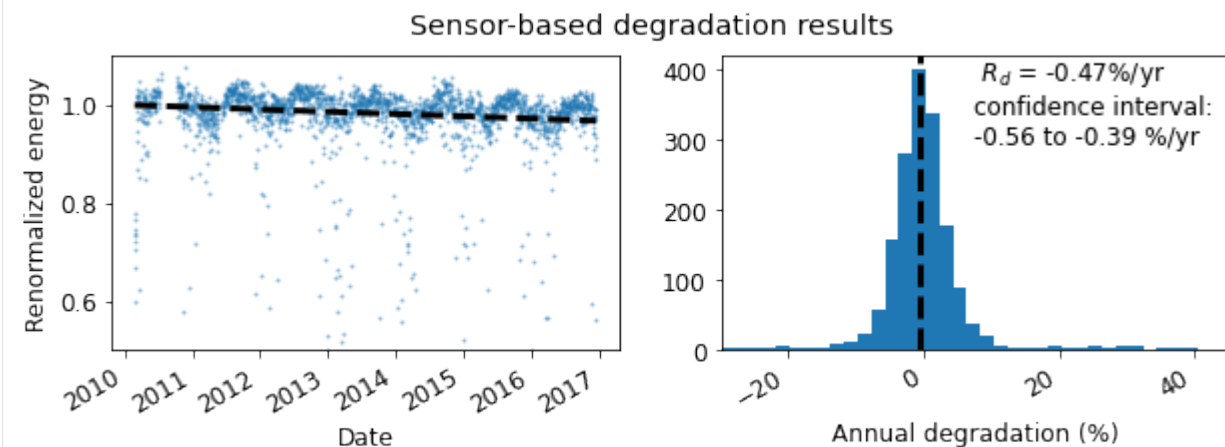

In addition to the confidence interval, the year-on-year method yields an exceedance value (e.g. P95), the degradation rate that was exceeded (slower degradation) with a given probability level. The probability level is set via the `exceedance_prob` keyword in `degradation_year_on_year`.

```
[9]:  print('The P95 exceedance level is %.2f%%/yr' % yoy_info['exceedance_level'])

      The P95 exceedance level is -0.63%/yr
```

7.1.6 5: Soiling calculations

This section illustrates how the aggregated data can be used to estimate soiling losses using the stochastic rate and recovery (SRR) method.[1] Since our example system doesn't experience much soiling, we apply an artificially generated soiling signal, just for the sake of example.

[1] M. G. Deceglie, L. Micheli and M. Muller, "Quantifying Soiling Loss Directly From PV Yield," IEEE Journal of Photovoltaics, vol. 8, no. 2, pp. 547-551, March 2018. doi: 10.1109/JPHOTOV.2017.2784682

```python
[10]: # Apply artificial soiling signal for example
      # be sure to remove this for applications on real data,
      # and proceed with analysis on `daily` instead of `soiled_daily`

      soiling = df['soiling'].resample('D').mean()
      soiled_daily = soiling*daily
```

```python
[11]: # Calculate the daily insolation, required for the SRR calculation
      daily_insolation = filtered['insolation'].resample('D').sum()

      # Perform the SRR calculation
      from rdtools.soiling import soiling_srr
      cl = 68.2
      sr, sr_ci, soiling_info = soiling_srr(soiled_daily, daily_insolation,
                                            confidence_level=cl)
```

```
/Users/mdecegli/Documents/GitHub/rdtools/rdtools/soiling.py:15: UserWarning: The
→soiling module is currently experimental. The API, results, and default behaviors
→may change in future releases (including MINOR and PATCH releases) as the code
→matures.
  'The soiling module is currently experimental. The API, results, '
```

```python
[12]: print('The P50 insolation-weighted soiling ratio is %0.3f'%sr)
```

```
The P50 insolation-weighted soiling ratio is 0.945
```

```python
[13]: print('The %0.1f confidence interval for the insolation-weighted'
            ' soiling ratio is %0.3f-%0.3f'%(cl, sr_ci[0], sr_ci[1]))
```

```
The 68.2 confidence interval for the insolation-weighted soiling ratio is 0.939-0.951
```

```python
[14]: # Plot Monte Carlo realizations of soiling profiles
      fig = rdtools.plotting.soiling_monte_carlo_plot(soiling_info, soiled_daily,
      →profiles=200);
```

```
/Users/mdecegli/Documents/GitHub/rdtools/rdtools/plotting.py:151: UserWarning: The
→soiling module is currently experimental. The API, results, and default behaviors
→may change in future releases (including MINOR and PATCH releases) as the code
→matures.
  'The soiling module is currently experimental. The API, results, '
```
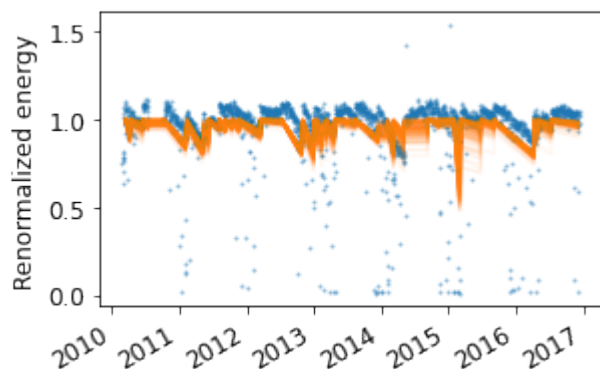


```python
[15]: # Plot the slopes for "valid" soiling intervals identified,
      # assuming perfect cleaning events
      fig = rdtools.plotting.soiling_interval_plot(soiling_info, soiled_daily);
```

7.1. Degradation and soiling example with clearsky workflow   **23**
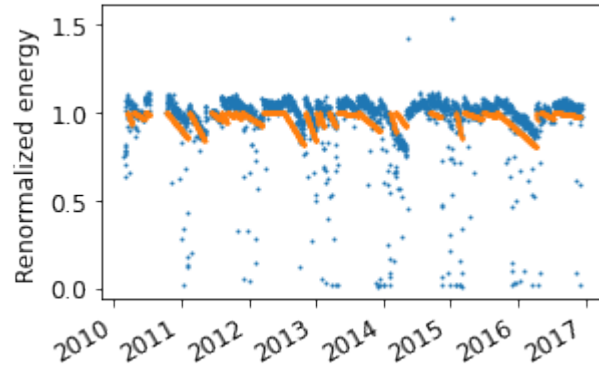
```
/Users/mdecegli/Documents/GitHub/rdtools/rdtools/plotting.py:211: UserWarning: The
↪soiling module is currently experimental. The API, results, and default behaviors
↪may change in future releases (including MINOR and PATCH releases) as the code
↪matures.
  'The soiling module is currently experimental. The API, results, '
```



```
[16]:  # View the first several rows of the soiling interval summary table
       soiling_summary = soiling_info['soiling_interval_summary']
       soiling_summary.head()
```

```
[16]:                     start                        end   soiling_rate  \
       0 2010-02-25 00:00:00-07:00  2010-03-06 00:00:00-07:00       0.000000
       1 2010-03-07 00:00:00-07:00  2010-03-11 00:00:00-07:00       0.000000
       2 2010-03-12 00:00:00-07:00  2010-04-08 00:00:00-07:00      -0.002505
       3 2010-04-09 00:00:00-07:00  2010-04-11 00:00:00-07:00       0.000000
       4 2010-04-12 00:00:00-07:00  2010-06-15 00:00:00-07:00      -0.000594

          soiling_rate_low  soiling_rate_high  inferred_start_loss  \
       0          0.000000           0.000000             0.685379
       1          0.000000           0.000000             1.053439
       2         -0.005069           0.000000             1.058785
       3          0.000000           0.000000             1.044975
       4         -0.000997          -0.000174             1.011211

          inferred_end_loss  length  valid
       0           0.863517       9  False
       1           1.003025       4  False
       2           0.991144      27   True
       3           1.044975       2  False
       4           0.973207      64   True
```

```
[17]:  # View a histogram of the valid soiling rates found for the data set
       fig = rdtools.plotting.soiling_rate_histogram(soiling_info, bins=50)
```

```
/Users/mdecegli/Documents/GitHub/rdtools/rdtools/plotting.py:251: UserWarning: The
↪soiling module is currently experimental. The API, results, and default behaviors
↪may change in future releases (including MINOR and PATCH releases) as the code
↪matures.
  'The soiling module is currently experimental. The API, results, '
```
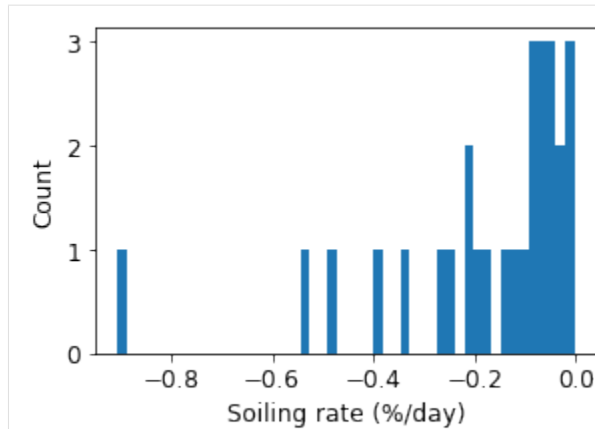
These plots show generally good results from the SRR method. In this example, we have slightly overesti-mated the soiling loss because we used the default behavior of the `method` key word argument in `rdtools.soiling_srr()`, which does not assume that every cleaning is perfect but the example artificial soiling signal did include perfect cleaning. We encourage you to adjust the options of `rdtools.soiling_srr()` for your application.

```python
[18]: # Calculate and view a monthly soiling rate summary
      from rdtools.soiling import monthly_soiling_rates
      monthly_soiling_rates(soiling_info['soiling_interval_summary'],
                            confidence_level=cl)
```

```
[18]:     month  soiling_rate_median  soiling_rate_low  soiling_rate_high  \
      0       1            -0.000942         -0.001954          -0.000692
      1       2            -0.001794         -0.006180          -0.000752
      2       3            -0.001096         -0.002230          -0.000394
      3       4            -0.000924         -0.001899          -0.000122
      4       5            -0.000305         -0.000733          -0.000086
      5       6            -0.000331         -0.000777          -0.000091
      6       7            -0.000404         -0.001342          -0.000140
      7       8            -0.000674         -0.001779          -0.000182
      8       9            -0.000856         -0.001572          -0.000191
      9      10            -0.000881         -0.001413          -0.000203
      10     11            -0.000920         -0.001894          -0.000229
      11     12            -0.000947         -0.002455          -0.000691

          interval_count
      0                6
      1                7
      2               10
      3                9
      4                7
      5                8
      6                8
      7                7
      8                8
      9                8
      10               8
      11               6
```

```python
[19]: # Calculate and view annual insolation-weighted soiling ratios and their confidence
      # intervals based on the Monte Carlo simulation. Note that these losses include the
```

```python
# assumptions of the cleaning assumptions associated with the method parameter
# of rdtools.soiling_srr(). For anything but 'perfect_clean', each year's soiling
# ratio may be impacted by prior years' soiling profiles. The default behavior of
# rdtools.soiling_srr uses method='half_norm_clean'

from rdtools.soiling import annual_soiling_ratios
annual_soiling_ratios(soiling_info['stochastic_soiling_profiles'],
                      daily_insolation,
                      confidence_level=cl)
```

```
[19]:   year  soiling_ratio_median  soiling_ratio_low  soiling_ratio_high
     0  2010              0.961769           0.950512            0.969079
     1  2011              0.944563           0.937086            0.950570
     2  2012              0.939465           0.931211            0.945439
     3  2013              0.954355           0.944595            0.961878
     4  2014              0.949834           0.929179            0.965085
     5  2015              0.950557           0.921117            0.966028
     6  2016              0.937150           0.925213            0.944815
```

7.1.7 Clear sky workflow

The clear sky workflow is useful in that it avoids problems due to drift or recalibration of ground-based sensors. We use `pvlib` to model the clear sky irradiance. This is renormalized to align it with ground-based measurements. Finally we use `rdtools.get_clearsky_tamb()` to model the ambient temperature on clear sky days. This modeled ambient temperature is used to model cell temperature with `pvlib`. If high quality ambient temperature data is available, that can be used instead of the modeled ambient; we proceed with the modeled ambient temperature here for illustrative purposes.

In this example, note that we have omitted wind data in the cell temperature calculations for illustrative purposes. Wind data can also be included when the data source is trusted for improved results

We generally recommend that the clear sky workflow be used as a check on the sensor workflow. It tends to be more sensitive than the sensor workflow, and thus we don't recommend it as a stand-alone analysis.

Note that the calculations below rely on some objects from the steps above

7.1.8 Clear Sky 0: Preliminary Calculations

```python
# Calculate the clear sky POA irradiance
clearsky = loc.get_clearsky(df.index, solar_position=sun)

cs_sky = pvlib.irradiance.isotropic(meta['tilt'], clearsky.dhi)
cs_beam = pvlib.irradiance.beam_component(meta['tilt'], meta['azimuth'],
                                          sun.zenith, sun.azimuth, clearsky.dni)
df['clearsky_poa'] = cs_beam + cs_sky

# Renormalize the clear sky POA irradiance
df['clearsky_poa'] = rdtools.irradiance_rescale(df.poa, df.clearsky_poa,
                                                method='iterative')

# Calculate the clearsky temperature
df['clearsky_Tamb'] = rdtools.get_clearsky_tamb(df.index, meta['latitude'],
                                                meta['longitude'])
df['clearsky_Tcell'] = pvlib.temperature.sapm_cell(df.clearsky_poa, df.clearsky_Tamb,
                                                   0, **meta['temp_model_params'])
```

7.1.9 Clear Sky 1: Normalize

Normalize as in step 1 above, but this time using clearsky modeled irradiance and cell temperature

```
[21]: # Calculate the expected power with a simple PVWatts DC model
clearsky_modeled_power = pvlib.pvsystem.pvwatts_dc(df['clearsky_poa'],
                                                   df['clearsky_Tcell'],
                                                   meta['power_dc_rated'], meta[
↪'gamma_pdc'], 25.0 )

# Calculate the normalization, the function also returns the relevant insolation for
# each point in the normalized PV energy timeseries
clearsky_normalized, clearsky_insolation = rdtools.normalize_with_expected_power(
    df['power_ac'],
    clearsky_modeled_power,
    df['clearsky_poa']
)

df['clearsky_normalized'] = clearsky_normalized
df['clearsky_insolation'] = clearsky_insolation
```

7.1.10 Clear Sky 2: Filter

Filter as in step 2 above, but with the addition of a clear sky index (csi) filter so we consider only points well modeled by the clear sky irradiance model.

```
[22]: # Perform clearsky filter
cs_normalized_mask = rdtools.normalized_filter(df['clearsky_normalized'])
cs_poa_mask = rdtools.poa_filter(df['clearsky_poa'])
cs_tcell_mask = rdtools.tcell_filter(df['clearsky_Tcell'])

csi_mask = rdtools.csi_filter(df.insolation, df.clearsky_insolation)

clearsky_filtered = df[cs_normalized_mask & cs_poa_mask & cs_tcell_mask &
                       clip_mask & csi_mask]
clearsky_filtered = clearsky_filtered[['clearsky_insolation', 'clearsky_normalized']]
```

7.1.11 Clear Sky 3: Aggregate

Aggregate the clear sky version of of the filtered data

```
[23]: clearsky_daily = rdtools.aggregation_insol(clearsky_filtered.clearsky_normalized,
                                           clearsky_filtered.clearsky_insolation)
```

7.1.12 Clear Sky 4: Degradation Calculation

Estimate the degradation rate and compare to the results obtained with sensors. In this case, we see that the degradation rate estimated with the clearsky methodology is not far off from the sensor-based estimate.

```
[24]: # Calculate the degradation rate using the YoY method
cs_yoy_rd, cs_yoy_ci, cs_yoy_info = rdtools.degradation_year_on_year(
    clearsky_daily,
```

```
    confidence_level=68.2
)

# Note the default confidence_level of 68.2 is appropriate if you would like to
# report a confidence interval analogous to the standard deviation of a normal
# distribution. The size of the confidence interval is adjustable by setting the
# confidence_level variable.

# Visualize the results
clearsky_fig = rdtools.degradation_summary_plots(
    cs_yoy_rd, cs_yoy_ci, cs_yoy_info, clearsky_daily,
    summary_title='Clear-sky-based degradation results',
    scatter_ymin=0.5, scatter_ymax=1.1,
    hist_xmin=-30, hist_xmax=45, plot_color='orangered',
    bins=100);

print('The P95 exceedance level with the clear sky analysis is %.2f%%/yr' %
      cs_yoy_info['exceedance_level'])
```

```
The P95 exceedance level with the clear sky analysis is -0.81%/yr
```
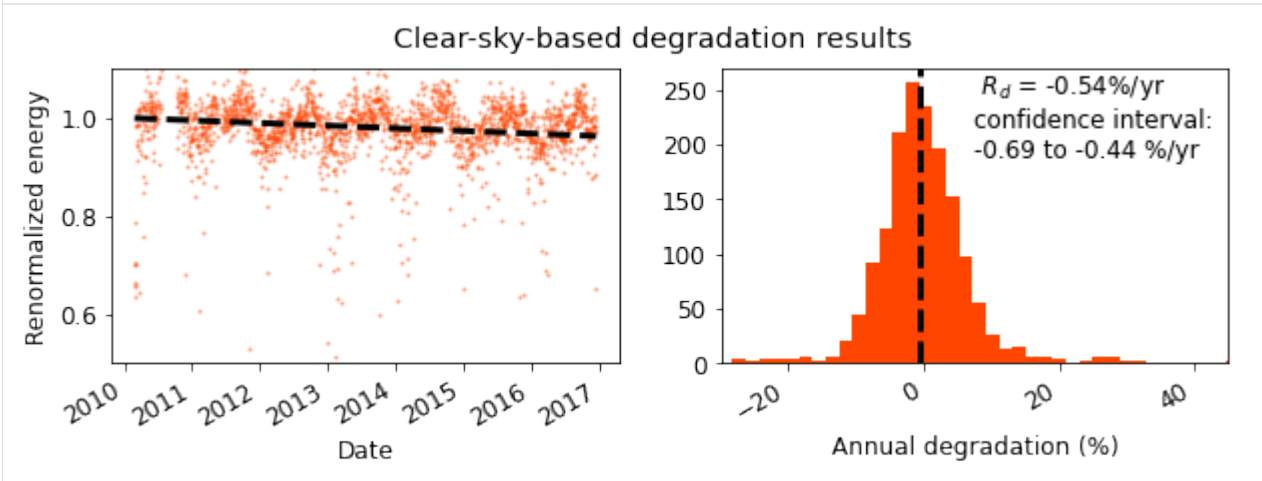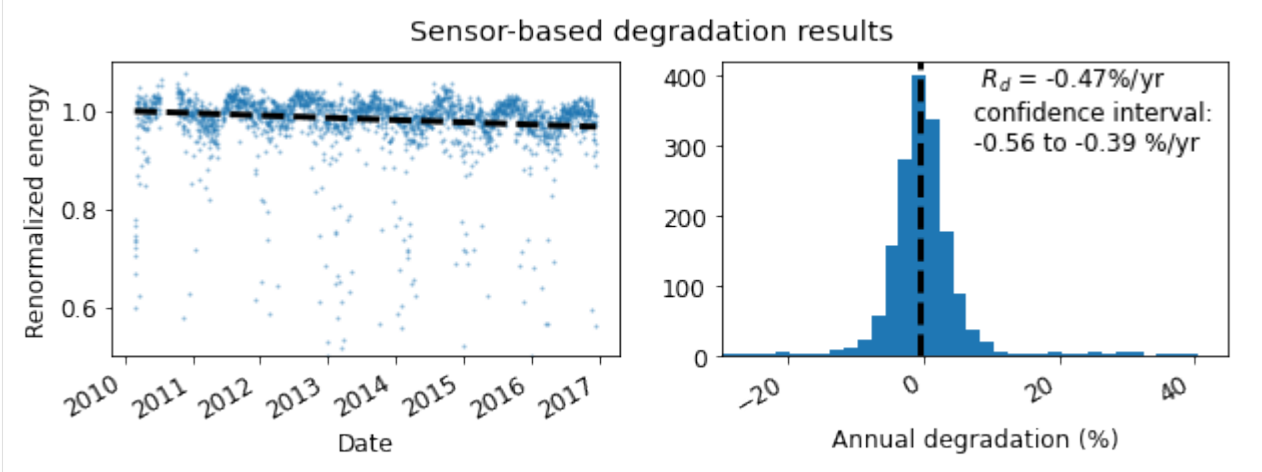


```
[25]: # Compare to previous sensor results
      degradation_fig
```

```
[25]:
```

7.2 System availability example

This notebook shows example usage of the inverter availability functions. As with the degradation and soiling example, we recommend installing the specific versions of packages used to develop this notebook. This can be achieved in your environment by running `pip install -r requirements.txt` followed by `pip install -r docs/notebook_requirements.txt` from the base directory. (RdTools must also be separately installed.) These environments and examples are tested with Python 3.7.

RdTools currently implements two methods of quantifying system availability. The first method compares power measurements from inverters and the system meter to distinguish subsystem communication interruptions from true outage events. The second method determines the uncertainty bounds around an energy estimate of a total system outage and compares with true production calculated from a meter's cumulative production measurements. The RdTools `AvailabilityAnalysis` class uses both methods to quantify downtime loss.

These methods are described in K. Anderson and R. Blumenthal, "Overcoming Communications Outages in Inverter Downtime Analysis", 2020 IEEE 47th Photovoltaic Specialists Conference (PVSC).

```python
[1]: import rdtools
     import pvlib

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

Quantifying the production impact of inverter downtime events is complicated by gaps in a system's historical data caused by communication interruptions. Although communication interruptions may prevent remote operation, they usually do not result in production loss. Accurate production loss estimates require the ability to distinguish true outages from communication interruptions.

The first method focuses on partial outages where some of a system's inverters are reporting production and some are not. In these cases, the method examines the AC power measurements at the inverter and system meter level to classify each timestamp individually and estimate timeseries production loss. This level of granularity is made possible by comparing timeseries power measurements between inverters and the meter.

7.2.1 Create a test dataset

First we'll generate a test dataset to demonstrate the method. This code block just puts together an artificial dataset to use for the analysis – feel free to skip ahead to where it gets plotted.

```python
[2]: def make_dataset():
         """
         Make an example dataset with several types of data outages for availability
     ↪analysis.

         Returns
         -------
         df_reported : pd.DataFrame
             Simulated data as a data acquisition system would report it, including the
             effect of communication interruptions.
         df_secret : pd.DataFrame
             The secret true data of the system, not affected by communication
             interruptions.  Only used for comparison with the analysis output.
         expected_power : pd.Series
             An "expected" power signal for this hypothetical PV system, simulating a
             modeled power from satellite weather data or some other method.
```

(continues on next page)

```python
    (This function creates instananeous data. SystemAvailability is technically␣
→designed
    to work with right-labeled averages. However, for the purposes of the example, the
    approximation is suitable.)
    """

    # generate a plausible clear-sky power signal
    times = pd.date_range('2019-01-01', '2019-01-12', freq='15min', tz='US/Eastern',
                          closed='left')
    location = pvlib.location.Location(40, -80)
    clearsky = location.get_clearsky(times, model='haurwitz')
    # just scale GHI to power for simplicity
    base_power = 2.5*clearsky['ghi']
    # but require a minimum irradiance to turn on, simulating start-up voltage
    base_power[clearsky['ghi'] < 20] = 0

    df_secret = pd.DataFrame({
        'inv1_power': base_power,
        'inv2_power': base_power * 1.5,
        'inv3_power': base_power * 0.66,
    })

    # set the expected_power to be pretty close to actual power,
    # but with some autocorrelated noise and a bias:
    expected_power = df_secret.sum(axis=1)
    np.random.seed(2020)
    N = len(times)
    expected_power *= 0.9 - (0.3 * np.sin(np.arange(0, N)/7 +
                             np.random.normal(0, 0.2, size=N)))

    # Add a few days of individual inverter outages:
    df_secret.loc['2019-01-03':'2019-01-05', 'inv2_power'] = 0
    df_secret.loc['2019-01-02', 'inv3_power'] = 0
    df_secret.loc['2019-01-07 00:00':'2019-01-07 12:00', 'inv1_power'] = 0

    # and a full system outage:
    full_outage_date = '2019-01-08'
    df_secret.loc[full_outage_date, :] = 0

    # calculate the system meter power and cumulative production,
    # including the effect of the outages:
    df_secret['meter_power'] = df_secret.sum(axis=1)
    interval_energy = rdtools.energy_from_power(df_secret['meter_power'])
    df_secret['meter_energy'] = interval_energy.cumsum()
    # fill the first NaN from the cumsum with 0
    df_secret['meter_energy'] = df_secret['meter_energy'].fillna(0)
    # add an offset to reflect previous production:
    df_secret['meter_energy'] += 5e5
    # calculate cumulative energy for an inverter as well:
    inv2_energy = rdtools.energy_from_power(df_secret['inv2_power'])
    df_secret['inv2_energy'] = inv2_energy.cumsum().fillna(0)

    # now that the "true" data is in place, let's add some communications␣
→interruptions:
    df_reported = df_secret.copy()
    # in full outages, we lose all the data:
```

```
    df_reported.loc[full_outage_date, :] = np.nan
    # add a communications interruption that overlaps with an inverter outage:
    df_reported.loc['2019-01-05':'2019-01-06', 'inv1_power'] = np.nan
    # and a communication outage that affects everything:
    df_reported.loc['2019-01-10', :] = np.nan

    return df_reported, df_secret, expected_power
```

Let's visualize the dataset before analyzing it with RdTools. The dotted lines show the "true" data that wasn't recorded by the datalogger because of interrupted communications.

```python
[3]: df, df_secret, expected_power = make_dataset()

fig, axes = plt.subplots(3, 1, sharex=True, figsize=(8,6))
colors = plt.rcParams['axes.prop_cycle'].by_key()['color'][:3]

# inverter power
df_secret[['inv1_power', 'inv2_power', 'inv3_power']].plot(ax=axes[0],
                                                legend=False, ls=':',
                                                color=colors)
df[['inv1_power', 'inv2_power', 'inv3_power']].plot(ax=axes[0], legend=False)
# meter power
df_secret['meter_power'].plot(ax=axes[1], ls=':', color=colors[0])
df['meter_power'].plot(ax=axes[1])
# meter cumulative energy
df_secret['meter_energy'].plot(ax=axes[2], ls=':', color=colors[0])
df['meter_energy'].plot(ax=axes[2])

axes[0].set_ylabel('Inverter Power [kW]')
axes[1].set_ylabel('Meter Power [kW]')
axes[2].set_ylabel('Cumulative\nMeter Energy [kWh]')
plt.show()
```
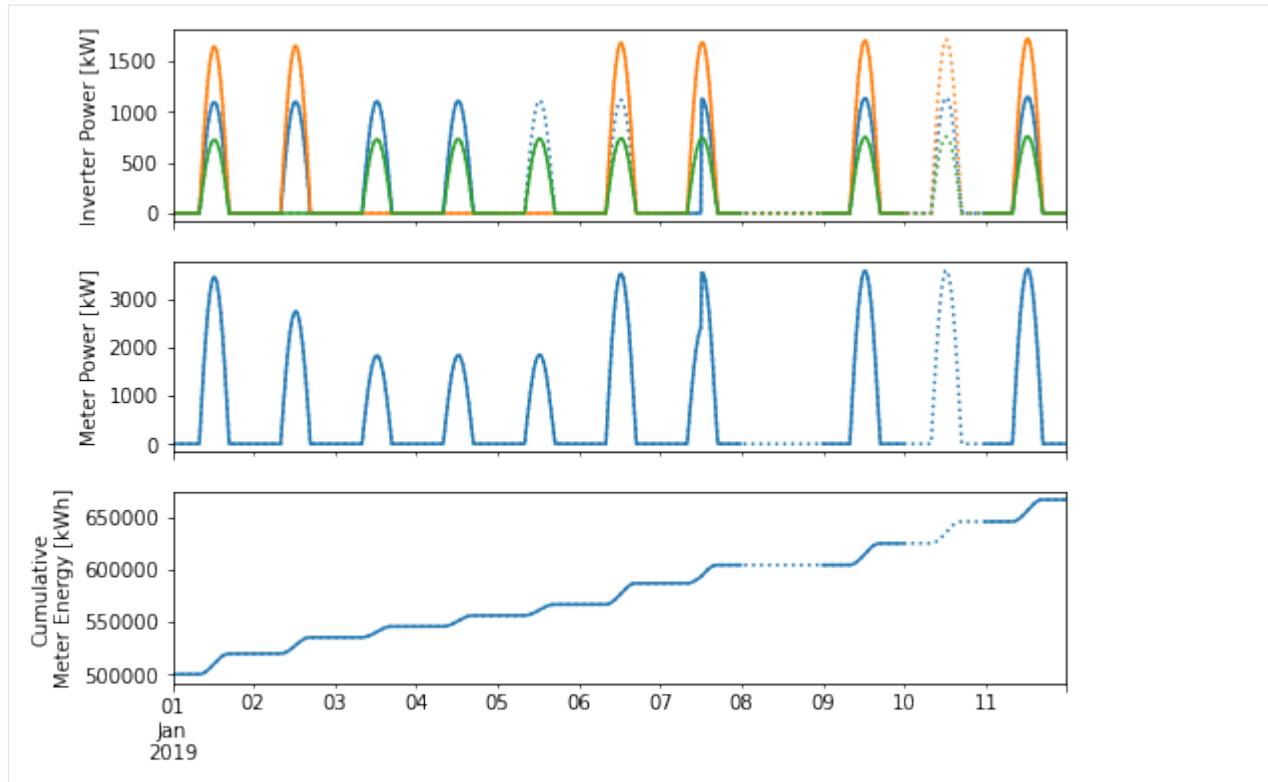
Note that the solid lines show the data that would be available in our example while the dotted lines show the true underlying behavior that we normally wouldn't know.

If we hadn't created this dataset ourselves, it wouldn't necessarily be obvious why the meter shows low or no production on some days – maybe it was just cloudy weather, maybe it was a nuisance communication outage (broken cell modem power supply, for example), or maybe it was a true power outage. This example also shows how an inverter can appear to be offline while actually producing normally. For example, just looking at inverter power on the 5th, it appears that only the small inverter is producing. However, the meter shows two inverters' worth of production. Similarly, the 6th shows full meter production despite one inverter not reporting power. Using only the inverter-reported power would overestimate the production loss because of the communication interruption.

7.2.2 System availability analysis

Now we'll hand this data off to RdTools for analysis:

```python
from rdtools.availability import AvailabilityAnalysis
aa = AvailabilityAnalysis(
    power_system=df['meter_power'],
    power_subsystem=df[['inv1_power', 'inv2_power', 'inv3_power']],
    energy_cumulative=df['meter_energy'],
    power_expected=expected_power,
)
# identify and classify outages, rolling up to daily metrics for this short dataset:
aa.run(rollup_period='D')
```

```
/Users/mdecegli/opt/anaconda3/envs/final_release_test/lib/python3.7/site-packages/
↪rdtools/availability.py:18: UserWarning: The availability module is currently␣
↪experimental. The API, results, and default behaviors may change in future releases␣
↪(including MINOR and PATCH releases) as the code matures.
  'The availability module is currently experimental. The API, results, '
```

First, we can visualize the estimated power loss and outage information:

```
[5]: fig = aa.plot()
fig.set_size_inches(16, 7)
fig.axes[1].legend(loc='upper left');
```

```
/Users/mdecegli/opt/anaconda3/envs/final_release_test/lib/python3.7/site-packages/
↪rdtools/plotting.py:320: UserWarning: The availability module is currently␣
↪experimental. The API, results, and default behaviors may change in future releases␣
↪(including MINOR and PATCH releases) as the code matures.
  'The availability module is currently experimental. The API, results, '
```
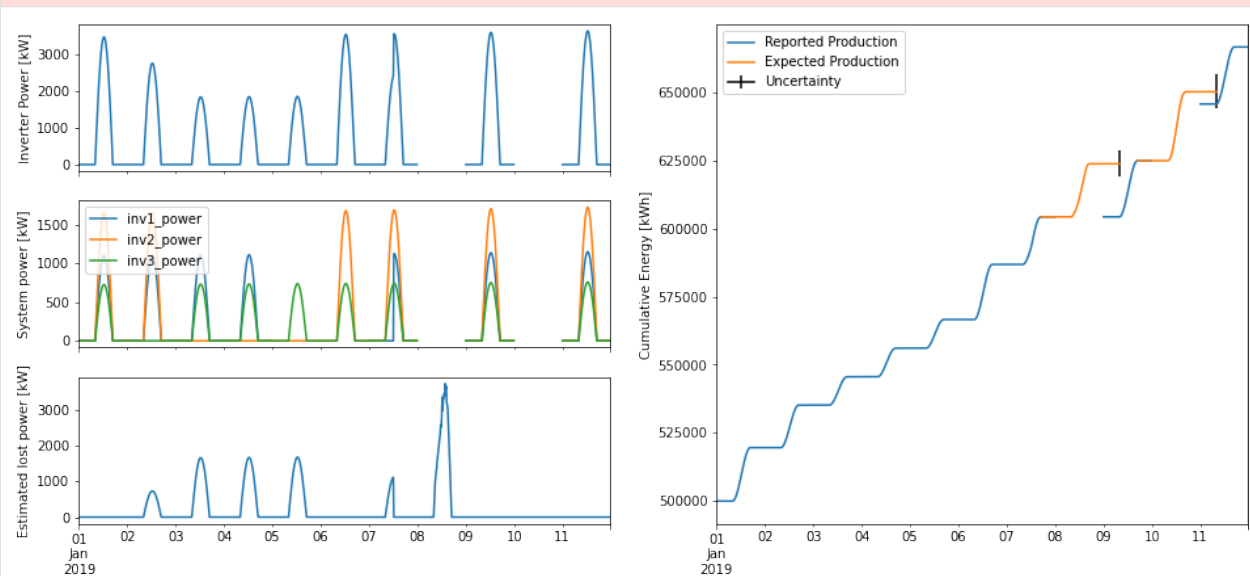

Examining the plot of estimated lost power, we can see that the estimated loss is roughly in proportion to the amount of offline capacity. In particular, the loss estimate is robust to mixed outage and communication interruption like on the 5th when only the smallest inverter is reporting production but the analysis correctly inferred that one of the other inverters is producing but not communicating.

RdTools also reports rolled-up production and availability metrics:

```
[6]: pd.set_option('precision', 3)
aa.results
```

```
[6]:                            lost_production  actual_production  availability
2019-01-01 00:00:00-05:00            0.000          19606.785         1.000
2019-01-02 00:00:00-05:00         4114.031          15583.450         0.791
2019-01-03 00:00:00-05:00         9396.788          10399.112         0.525
2019-01-04 00:00:00-05:00         9466.477          10476.235         0.525
2019-01-05 00:00:00-05:00         9522.325          10538.040         0.525
2019-01-06 00:00:00-05:00            0.000          20185.784         1.000
2019-01-07 00:00:00-05:00         2859.565          17459.339         0.859
2019-01-08 00:00:00-05:00        19448.084              0.000         0.000
2019-01-09 00:00:00-05:00            0.000          20607.950         1.000
```

(continues on next page)

<div style="text-align: right">(continued from previous page)</div>

```
2019-01-10 00:00:00-05:00            0.000        20763.718            1.000
2019-01-11 00:00:00-05:00            0.000        20926.869            1.000
```

The `AvailabilityAnalysis` object has other attributes that may be useful to inspect as well. The `outage_info` dataframe has one row for each full system outage with several columns, perhaps the most interesting of which are `type` and `loss`.

See `AvailabilityAnalysis?` or `help(AvailabilityAnalysis)` for full descriptions of the available attributes.

```python
[7]: pd.set_option('precision', 2)
     # Show the first half of the dataframe
     N = len(aa.outage_info.columns)
     aa.outage_info.iloc[:, :N//2]
```

```
[7]:                     start                       end        duration  \
     0 2019-01-07 17:00:00-05:00 2019-01-09 08:00:00-05:00 1 days 15:00:00
     1 2019-01-09 17:00:00-05:00 2019-01-11 08:00:00-05:00 1 days 15:00:00

        intervals  daylight_intervals  error_lower  error_upper
     0        157                  35        -0.24         0.25
     1        157                  35        -0.24         0.25
```

```python
[8]: # Show the second half
     aa.outage_info.iloc[:, N//2:]
```

```
[8]:    energy_expected  energy_start  energy_end  energy_actual  ci_lower  \
     0         19448.08     604248.74   604248.74           0.00  14819.33
     1         25284.75     624856.69   645620.41       20763.72  19266.84

        ci_upper   type       loss
     0  24271.15   real  19448.08
     1  31555.29  comms      0.00
```

7.2.3 Other use cases

Although this demo applies the methods for an entire PV system (comparing inverters against the meter and comparing the meter against expected power), it can also be used at the individual inverter level. Because there are no subsystems to compare against, the "full outage" analysis branch is used for every outage. That means that instead of basing the loss off of the other inverters, it relies on the expected power time series being accurate, which in this example causes the loss estimates to lose some accuracy. In this case, because the expected power signal is somewhat inaccurate, it causes the loss estimate to be overestimated:

```python
[9]: # make a new analysis object:
     aa2 = rdtools.availability.AvailabilityAnalysis(
         power_system=df['inv2_power'],
         power_subsystem=df['inv2_power'].to_frame(),
         energy_cumulative=df['inv2_energy'],
         # okay to use the system-level expected power here because it gets rescaled anyway
         power_expected=expected_power,
     )
     # identify and classify outages, rolling up to daily metrics for this short dataset:
     aa2.run(rollup_period='D')
     print(aa2.results['lost_production'])
```

```
2019-01-01 00:00:00-05:00        0.00
2019-01-02 00:00:00-05:00        0.00
2019-01-03 00:00:00-05:00     9931.24
2019-01-04 00:00:00-05:00    11453.27
2019-01-05 00:00:00-05:00    11238.57
2019-01-06 00:00:00-05:00        0.00
2019-01-07 00:00:00-05:00        0.00
2019-01-08 00:00:00-05:00     9505.33
2019-01-09 00:00:00-05:00        0.00
2019-01-10 00:00:00-05:00        0.00
2019-01-11 00:00:00-05:00        0.00
Freq: D, Name: lost_production, dtype: float64
```

```
[10]: aa2.plot();
```

```
/Users/mdecegli/opt/anaconda3/envs/final_release_test/lib/python3.7/site-packages/
→rdtools/plotting.py:320: UserWarning: The availability module is currently␣
→experimental. The API, results, and default behaviors may change in future releases␣
→(including MINOR and PATCH releases) as the code matures.
  'The availability module is currently experimental. The API, results, '
```
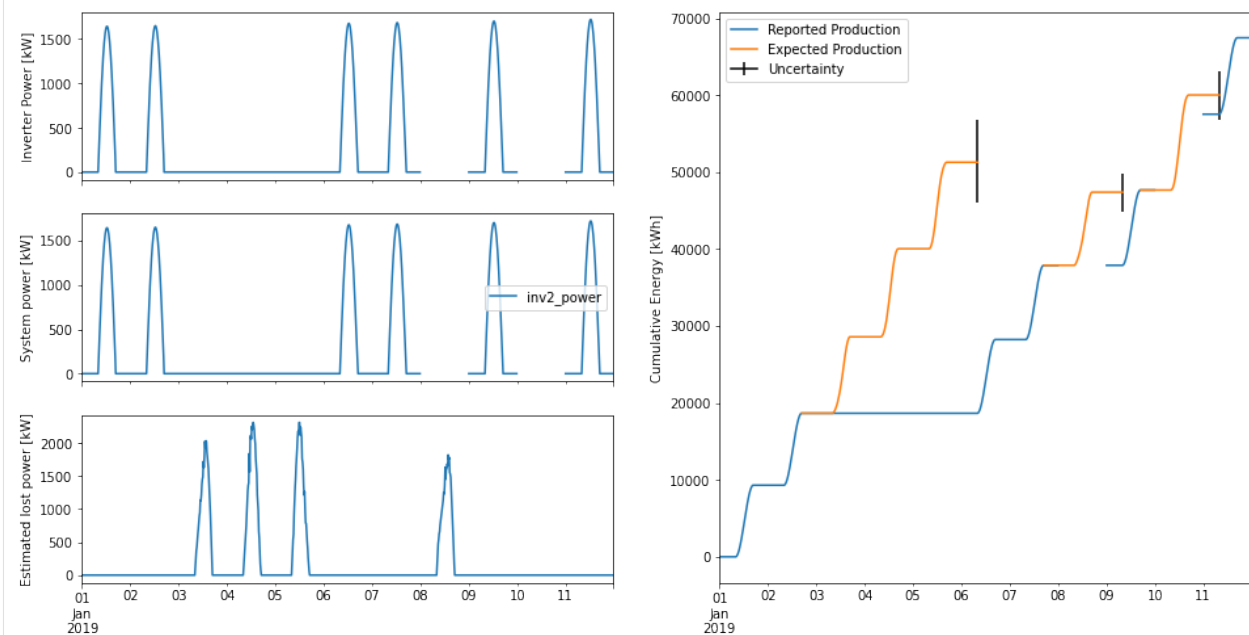

7.3 API reference

7.3.1 Submodules

RdTools is organized into submodules focused on different parts of the data analysis workflow.

degradation	Functions for calculating the degradation rate of photovoltaic systems.
soiling	Functions for calculating soiling metrics from photovoltaic system data.

Table 1 – continued from previous page

availability	Functions for detecting and quantifying production loss from photovoltaic system downtime events.
filtering	Functions for filtering and subsetting PV system data.
normalization	Functions for normalizing, rescaling, and regularizing PV system data.
aggregation	Functions for calculating weighted aggregates of PV system data.
clearsky_temperature	Functions for estimating clear-sky ambient temperature.
plotting	Functions for plotting degradation and soiling analysis results.

rdtools.degradation

Functions for calculating the degradation rate of photovoltaic systems.

Functions

degradation_classical_decomposition(...[, ...])	Estimate the trend of a timeseries using a classical decomposition approach (moving average) and calculate various statistics, including the result of a Mann-Kendall test and a Monte Carlo-derived confidence interval of slope.
degradation_ols(energy_normalized[, ...])	Estimate the trend of a timeseries using ordinary least-squares regression and calculate various statistics including a Monte Carlo-derived confidence interval of slope.
degradation_year_on_year(energy_normalized)	Estimate the trend of a timeseries using the year-on-year decomposition approach and calculate a Monte Carlo-derived confidence interval of slope.

rdtools.soiling

Functions for calculating soiling metrics from photovoltaic system data.

The soiling module is currently experimental. The API, results, and default behaviors may change in future releases (including MINOR and PATCH releases) as the code matures.

Functions

annual_soiling_ratios(...[, confidence_level])	Return annualized soiling ratios and associated confidence intervals based on stochastic soiling profiles from SRR.
monthly_soiling_rates(soiling_interval_summary)	Use Monte Carlo to calculate typical monthly soiling rates.
soiling_srr(energy_normalized_daily, ...[, ...])	Functional wrapper for *SRRAnalysis*.

Classes

SRRAnalysis(energy_normalized_daily, ...[, ...])	Class for running the stochastic rate and recovery (SRR) photovoltaic soiling loss analysis presented in Deceglie et al.

Exceptions

NoValidIntervalError	raised when no valid rows appear in the result dataframe

rdtools.availability

Functions for detecting and quantifying production loss from photovoltaic system downtime events.

The availability module is currently experimental. The API, results, and default behaviors may change in future releases (including MINOR and PATCH releases) as the code matures.

Classes

AvailabilityAnalysis(power_system, ...)	A class to perform system availability and loss analysis.

rdtools.filtering

Functions for filtering and subsetting PV system data.

Functions

clip_filter(power_ac[, quantile])	Filter data points likely to be affected by clipping with power greater than or equal to 99% of the *quant* quantile.
csi_filter(poa_global_measured, ...[, threshold])	Filtering based on clear-sky index (csi)
normalized_filter(energy_normalized[, ...])	Select normalized yield between low_cutoff and high_cutoff
poa_filter(poa_global[, poa_global_low, ...])	Filter POA irradiance readings outside acceptable measurement bounds.
tcell_filter(temperature_cell[, ...])	Filter temperature readings outside acceptable measurement bounds.

rdtools.normalization

Functions for normalizing, rescaling, and regularizing PV system data.

Functions

check_series_frequency(series, ...)	Deprecated since version 2.0.0.
delta_index(series)	Deprecated since version 2.0.0.

Continued on next page

Table 8 – continued from previous page

energy_from_power(power[, target_frequency, ...])	Returns a regular right-labeled energy time series in units of Wh per interval from a power time series.
interpolate(time_series, target[, ...])	Returns an interpolation of time_series, excluding times associated with gaps in each column of time_series longer than max_timedelta; NaNs are returned within those gaps.
irradiance_rescale(irrad, irrad_sim[, ...])	Attempt to rescale modeled irradiance to match measured irradiance on clear days.
normalize_with_expected_power(pv, ...[, ...])	Normalize PV power or energy based on expected PV power.
normalize_with_pvwatts(energy, pvwatts_kws)	Deprecated since version 2.0.0.
normalize_with_sapm(energy, sapm_kws)	Deprecated since version 2.0.0.
pvwatts_dc_power(poa_global, power_dc_rated)	Deprecated since version 2.0.0.
sapm_dc_power(pvlib_pvsystem, met_data)	Deprecated since version 2.0.0.

Exceptions

ConvergenceError	Rescale optimization did not converge

rdtools.aggregation

Functions for calculating weighted aggregates of PV system data.

Functions

aggregation_insol(energy_normalized, insolation)	Insolation weighted aggregation

rdtools.clearsky_temperature

Functions for estimating clear-sky ambient temperature.

Functions

get_clearsky_tamb(times, latitude, longitude)	Estimates the ambient temperature at latitude and longitude for the given times using a Gaussian rolling window.

rdtools.plotting

Functions for plotting degradation and soiling analysis results.

Functions

availability_summary_plots(power_system, ...)	Create a figure summarizing the availability analysis results.
degradation_summary_plots(yoy_rd, yoy_ci, ...)	Create plots (scatter plot and histogram) that summarize degradation analysis results.
soiling_interval_plot(soiling_info, ...[, ...])	Create figure to visualize valid soiling profiles used in the SRR analysis.
soiling_monte_carlo_plot(soiling_info, ...)	Create figure to visualize Monte Carlo of soiling profiles used in the SRR analysis.
soiling_rate_histogram(soiling_info[, bins])	Create histogram of soiling rates found in the SRR analysis.

7.3.2 Degradation

Functions for calculating the degradation rate of photovoltaic systems.

degradation_classical_decomposition(...[, ...])	Estimate the trend of a timeseries using a classical decomposition approach (moving average) and calculate various statistics, including the result of a Mann-Kendall test and a Monte Carlo-derived confidence interval of slope.
degradation_ols(energy_normalized[, ...])	Estimate the trend of a timeseries using ordinary least-squares regression and calculate various statistics including a Monte Carlo-derived confidence interval of slope.
degradation_year_on_year(energy_normalized)	Estimate the trend of a timeseries using the year-on-year decomposition approach and calculate a Monte Carlo-derived confidence interval of slope.

rdtools.degradation.degradation_classical_decomposition

rdtools.degradation.**degradation_classical_decomposition**(*energy_normalized*, *confidence_level=68.2*)

Estimate the trend of a timeseries using a classical decomposition approach (moving average) and calculate various statistics, including the result of a Mann-Kendall test and a Monte Carlo-derived confidence interval of slope.

> **Parameters**
>
> - **energy_normalized** (*pd.Series*) – Daily or lower frequency time series of normalized system ouput. Must be regular time series.
>
> - **confidence_level** (*float, default 68.2*) – The size of the confidence interval to return, in percent.
>
> **Returns**
>
> - **Rd_pct** (*float*) – Estimated degradation relative to the year 0 system capacity [%/year]
>
> - **Rd_CI** (*np.array*) – The calculated confidence interval bounds.
>
> - **calc_info** (*dict*) – A dict that contains slope, intercept, root mean square error of regression ('rmse'), standard error of the slope ('slope_stderr'), intercept ('intercept_stderr'), and least

squares RegressionResults object ('ols_results'), pandas series for the annual rolling mean ('series'), and Mann-Kendall test trend ('mk_test_trend')

rdtools.degradation.degradation_ols

rdtools.degradation.**degradation_ols**(*energy_normalized*, *confidence_level=68.2*)

Estimate the trend of a timeseries using ordinary least-squares regression and calculate various statistics including a Monte Carlo-derived confidence interval of slope.

Parameters

- **energy_normalized** (*pd.Series*) – Daily or lower frequency time series of normalized system ouput.

- **confidence_level** (*float, default 68.2*) – The size of the confidence interval to return, in percent.

Returns

- **Rd_pct** (*float*) – Estimated degradation relative to the year 0 system capacity [%/year]

- **Rd_CI** (*np.array*) – The calculated confidence interval bounds.

- **calc_info** (*dict*) – A dict that contains slope, intercept, root mean square error of regression ('rmse'), standard error of the slope ('slope_stderr'), intercept ('intercept_stderr'), and least squares RegressionResults object ('ols_results')

rdtools.degradation.degradation_year_on_year

rdtools.degradation.**degradation_year_on_year**(*energy_normalized*, *recenter=True*, *exceedance_prob=95*, *confidence_level=68.2*)

Estimate the trend of a timeseries using the year-on-year decomposition approach and calculate a Monte Carlo-derived confidence interval of slope.

Parameters

- **energy_normalized** (*pd.Series*) – Daily or lower frequency time series of normalized system ouput.

- **recenter** (*bool, default True*) – Specify whether data is internally recentered to normalized yield of 1 based on first year median. If False, Rd_pct is calculated assuming energy_normalized is passed already normalized to the year 0 system capacity.

- **exceedance_prob** (*float, default 95*) – The probability level to use for exceedance value calculation, in percent.

- **confidence_level** (*float, default 68.2*) – The size of the confidence interval to return, in percent.

Returns

- **Rd_pct** (*float*) – Estimated degradation relative to the year 0 median system capacity [%/year]

- **confidence_interval** (*np.array*) – confidence interval (size specified by confidence_level) of degradation rate estimate

- **calc_info** (*dict*) –

  – *YoY_values* - pandas series of right-labeled year on year slopes

- *renormalizing_factor* - float of value used to recenter data

- *exceedance_level* - the degradation rate that was outperformed with probability of *exceedance_prob*

7.3.3 Soiling

Functions for calculating soiling metrics from photovoltaic system data.

The soiling module is currently experimental. The API, results, and default behaviors may change in future releases (including MINOR and PATCH releases) as the code matures.

soiling_srr(energy_normalized_daily, ...[, ...])	Functional wrapper for *SRRAnalysis*.
monthly_soiling_rates(soiling_interval_summary)	Use Monte Carlo to calculate typical monthly soiling rates.
annual_soiling_ratios(...[, confidence_level])	Return annualized soiling ratios and associated confidence intervals based on stochastic soiling profiles from SRR.
SRRAnalysis(energy_normalized_daily, ...[, ...])	Class for running the stochastic rate and recovery (SRR) photovoltaic soiling loss analysis presented in Deceglie et al.
SRRAnalysis.run([reps, day_scale, ...])	Run the SRR method from beginning to end.

rdtools.soiling.soiling_srr

rdtools.soiling.**soiling_srr**(*energy_normalized_daily*, *insolation_daily*, *reps=1000*, *precipitation_daily=None*, *day_scale=14*, *clean_threshold='infer'*, *trim=False*, *method='half_norm_clean'*, *clean_criterion='shift'*, *precip_threshold=0.01*, *min_interval_length=2*, *exceedance_prob=95.0*, *confidence_level=68.2*, *recenter=True*, *max_relative_slope_error=500.0*, *max_negative_step=0.05*)

Functional wrapper for *SRRAnalysis*. Perform the stochastic rate and recovery soiling loss calculation. Based on the methods presented in Deceglie et al. JPV 8(2) p547 2018.

Parameters

- **energy_normalized_daily** (*pd.Series*) – Daily performance metric (i.e. performance index, yield, etc.) Alternatively, the soiling ratio output of a soiling sensor (e.g. the photocurrent ratio between matched dirty and clean PV reference cells). In either case, data should be insolation-weighted daily aggregates.

- **insolation_daily** (*pd.Series*) – Daily plane-of-array insolation corresponding to *energy_normalized_daily*. Arbitrary units.

- **reps** (*int, default 1000*) – number of Monte Carlo realizations to calculate

- **precipitation_daily** (*pd.Series, default None*) – Daily total precipitation. Units ambiguous but should be the same as precip_threshold. Note default behavior of precip_threshold. (Ignored if clean_criterion='shift'.)

- **day_scale** (*int, default 14*) – The number of days to use in rolling median for cleaning detection, and the maximum number of days of missing data to tolerate in a valid interval

- **clean_threshold** (*float or 'infer', default 'infer'*) – The fractional positive shift in rolling median for cleaning detection. Or specify 'infer' to automatically

use outliers in the shift as the threshold.

- **trim** (*bool, default False*) – Whether to trim (remove) the first and last soiling intervals to avoid inclusion of partial intervals

- **method**(*str, default 'half_norm_clean'*) – How to treat the recovery of each cleaning event

  - 'random_clean' - a random recovery between 0-100%

  - 'perfect_clean' - each cleaning event returns the performance metric to 1

  - **'half_norm_clean'(default) - The starting point of each interval is taken** randomly from a half normal distribution with its mode (mu) at 1 and its sigma equal to 1/3 * (1-b) where b is the intercept of the fit to the interval.

- **clean_criterion** (*{'precip_and_shift', 'precip_or_shift', 'precip', 'shift'} default 'shift'*) – The method of partitioning the dataset into soiling intervals. If 'precip_and_shift', rolling median shifts must coincide with precipitation to be a valid cleaning event. If 'precip_or_shift', rolling median shifts and precipitation events are each sufficient on their own to be a cleaning event. If 'shift', only rolling median shifts are treated as cleaning events. If 'precip', only precipitation events are treated as cleaning events.

- **precip_threshold** (*float, default 0.01*) – The daily precipitation threshold for defining precipitation cleaning events. Units must be consistent with precip.

- **min_interval_length** (*int, default 2*) – The minimum duration, in days, for an interval to be considered valid. Cannot be less than 2 (days).

- **exceedance_prob** (*float, default 95.0*) – the probability level to use for exceedance value calculation in percent

- **confidence_level** (*float, default 68.2*) – the size of the confidence interval to return, in percent

- **recenter** (*bool, default True*) – specify whether data is centered to normalized yield of 1 based on first year median

- **max_relative_slope_error** (*float, default 500.0*) – the maximum relative size of the slope confidence interval for an interval to be considered valid (percentage).

- **max_negative_step** (*float, default 0.05*) – The maximum magnitude of negative discrete steps allowed in an interval for the interval to be considered valid (units of normalized performance metric).

Returns

- **insolation_weighted_soiling_ratio** (*float*) – P50 insolation weighted soiling ratio based on stochastic rate and recovery analysis

- **confidence_interval** (*np.array*) – confidence interval (size specified by confidence_level) of degradation rate estimate

- **calc_info** (*dict*) –

  - 'renormalizing_factor' - value used to recenter data

  - 'exceedance_level' - the insolation-weighted soiling ratio that was outperformed with probability of exceedance_prob

  - 'stochastic_soiling_profiles' - List of Pandas series corresponding to the Monte Carlo realizations of soiling ratio profiles

- – 'soiling_ratio_perfect_clean' - Pandas series of the soiling ratio during valid soiling intervals assuming perfect cleaning and P50 slopes

- – 'soiling_interval_summary' - Pandas dataframe summarizing the soiling intervals identified. The columns of the dataframe are as follows:

Column Name	Description
'start'	Start timestamp of the soiling interval
'end'	End timestamp of the soiling interval
'soiling_rate'	P50 Soiling rate for interval, in day^1 Negative value indicates soiling is occurring. E.g. a rate of 0.01 indicates 1% soiling loss per day.
'soiling_rate_low'	Low edge of confidence interval for soiling rate for interval, in day^1
'soiling_rate_high'	High edge of confidence interval for soiling rate for interval, in day^1
'inferred_start_loss'	Estimated performance metric at the start of the interval
'inferred_end_loss'	Estimated performance metric at the end of the interval
'length'	Number of days in the interval
'valid'	Whether the interval meets the criteria to be treated as a valid soiling interval

rdtools.soiling.monthly_soiling_rates

rdtools.soiling.**monthly_soiling_rates**(*soiling_interval_summary*, *min_interval_length=14*, *max_relative_slope_error=500.0*, *reps=100000*, *confidence_level=68.2*)

Use Monte Carlo to calculate typical monthly soiling rates. Samples possible soiling rates from soiling rate confidence intervals associated with soiling intervals assuming a uniform distribution. Soiling intervals get samples proportionally to their overlap with each calendar month.

Parameters

- **soiling_interval_summary** (`pd.DataFrame`) – DataFrame describing soiling intervals. Typically from `soiling_info['soiling_interval_summary']` obtained with *rdtools.soiling.soiling_srr()* or *rdtools.soiling.SRRAnalysis.run()* Must have columns `soiling_rate_high`, `soiling_rate_low`, `soiling_rate`, `length`, `valid`, `start`, and `end`.

- **min_interval_length** (`int, default 14`) – The minimum number of days a soiling interval must contain to be included in the calculation. Similar to the same parameter in soiling_srr() and SRRAnalysis.run() but with a more conservative default value as a starting point for monthly soiling rate analyses.

- **max_relative_slope_error** (`float, default 500.0`) – The maximum relative size of the slope confidence interval for an interval to be included in the calculation (percentage).

- **reps** (`int, default 100000`) – The number of Monte Carlo samples to take for each month.

- **confidence_level** (`float, default 68.2`) – The size of the confidence interval, as a percentage, to use in determining the upper and lower quantiles reported in the returned DataFrame. (The median is always included in the result.)

Returns

DataFrame describing monthly soiling rates.

Column Name	Description
'month'	Integer month, January (1) to December (12)
'soiling_rate_median'	The median soiling rate for the month over the entire dataset, in units of day^1. Negative value indicates soiling is occurring. E.g. a rate of 0.01 indicates 1% soiling loss per day.
'soiling_rate_low'	The lower edge of the confidence interval for the monthly soiling rate in units of day^1
'soiling_rate_high'	The upper edge of the confidence interval for the monthly soiling rate in units of day^1
'interval_count'	The number of soiling intervals contributing to the monthly calculation. If only a few intervals contribute, the confidence interval is likely to underestimate the true uncertainty.

Return type pd.DataFrame

rdtools.soiling.annual_soiling_ratios

rdtools.soiling.**annual_soiling_ratios**(*stochastic_soiling_profiles*, *insolation_daily*, *confidence_level=68.2*)

Return annualized soiling ratios and associated confidence intervals based on stochastic soiling profiles from SRR. Note that each year may be affected by previous years' profiles for all SRR cleaning assumptions (i.e. method) except perfect_clean.

Parameters

- **stochastic_soiling_profiles** (`list`) – List of pd.Series representing profile realizations from the SRR monte carlo. Typically `soiling_interval_summary['stochastic_soiling_profiles']` obtained with *rdtools.soiling.soiling_srr()* or *rdtools.soiling.SRRAnalysis.run()*

- **insolation_daily** (`pd.Series`) – Daily plane-of-array insolation with DatetimeIndex. Arbitrary units.

- **confidence_level** (`float, default 68.2`) – The size of the confidence interval to use in determining the upper and lower quantiles reported in the returned DataFrame. (The median is always included in the result.)

Returns

DataFrame describing annual soiling rates.

Column Name	Description
'year'	Calendar year
'soiling_ratio_median'	The median insolation-weighted soiling ratio for the year
'soiling_ratio_low'	The lower edge of the confidence interval for insolation-weighted soiling ratio for the year
'soiling_ratio_high'	The upper edge of the confidence interval for insolation-weighted soiling ratio for the year

Return type pd.DataFrame

rdtools.soiling.SRRAnalysis

class rdtools.soiling.**SRRAnalysis**(*energy_normalized_daily*, *insolation_daily*, *precipitation_daily=None*)

Class for running the stochastic rate and recovery (SRR) photovoltaic soiling loss analysis presented in Deceglie et al. JPV 8(2) p547 2018

> **Parameters**
>
> - **energy_normalized_daily** (*pd.Series*) – Daily performance metric (i.e. performance index, yield, etc.) Alternatively, the soiling ratio output of a soiling sensor (e.g. the photocurrent ratio between matched dirty and clean PV reference cells). In either case, data should be insolation-weighted daily aggregates.
>
> - **insolation_daily** (*pd.Series*) – Daily plane-of-array insolation corresponding to *energy_normalized_daily*. Arbitrary units.
>
> - **precipitation_daily** (*pd.Series, default None*) – Daily total precipitation. (Ignored if clean_criterion='shift' in subsequent calculations.)

__init__(*energy_normalized_daily*, *insolation_daily*, *precipitation_daily=None*)
Initialize self. See help(type(self)) for accurate signature.

Methods

__init__(energy_normalized_daily, ...[, ...])	Initialize self.
run([reps, day_scale, clean_threshold, ...])	Run the SRR method from beginning to end.

rdtools.soiling.SRRAnalysis.run

SRRAnalysis.**run**(*reps=1000*, *day_scale=14*, *clean_threshold='infer'*, *trim=False*, *method='half_norm_clean'*, *clean_criterion='shift'*, *precip_threshold=0.01*, *min_interval_length=2*, *exceedance_prob=95.0*, *confidence_level=68.2*, *recenter=True*, *max_relative_slope_error=500.0*, *max_negative_step=0.05*)

Run the SRR method from beginning to end. Perform the stochastic rate and recovery soiling loss calculation. Based on the methods presented in Deceglie et al. "Quantifying Soiling Loss Directly From PV Yield" JPV 8(2) p547 2018.

> **Parameters**
>
> - **reps** (*int, default 1000*) – number of Monte Carlo realizations to calculate
>
> - **day_scale** (*int, default 14*) – The number of days to use in rolling median for cleaning detection, and the maximum number of days of missing data to tolerate in a valid interval
>
> - **clean_threshold** (*float or 'infer', default 'infer'*) – The fractional positive shift in rolling median for cleaning detection. Or specify 'infer' to automatically use outliers in the shift as the threshold.
>
> - **trim** (*bool, default False*) – Whether to trim (remove) the first and last soiling intervals to avoid inclusion of partial intervals
>
> - **method** (*str, default 'half_norm_clean'*) – How to treat the recovery of each cleaning event:

- – 'random_clean' - a random recovery between 0-100%

- – 'perfect_clean' - each cleaning event returns the performance metric to 1

- – 'half_norm_clean' - The starting point of each interval is taken randomly from a half normal distribution with its mode (mu) at 1 and its sigma equal to 1/3 * (1-b) where b is the intercept of the fit to the interval.

- **clean_criterion** (*{'precip_and_shift', 'precip_or_shift', 'precip', 'shift'} default 'shift'*) – The method of partitioning the dataset into soiling intervals. If 'precip_and_shift', rolling median shifts must coincide with precipitation to be a valid cleaning event. If 'precip_or_shift', rolling median shifts and precipitation events are each sufficient on their own to be a cleaning event. If 'shift', only rolling median shifts are treated as cleaning events. If 'precip', only precipitation events are treated as cleaning events.

- **precip_threshold** (*float, default 0.01*) – The daily precipitation threshold for defining precipitation cleaning events. Units must be consistent with `self.precipitation_daily`

- **min_interval_length** (*int, default 2*) – The minimum duration for an interval to be considered valid. Cannot be less than 2 (days).

- **exceedance_prob** (*float, default 95.0*) – The probability level to use for exceedance value calculation in percent

- **confidence_level** (*float, default 68.2*) – The size of the confidence interval to return, in percent

- **recenter** (*bool, default True*) – Specify whether data is centered to normalized yield of 1 based on first year median

- **max_relative_slope_error** (*float, default 500*) – the maximum relative size of the slope confidence interval for an interval to be considered valid (percentage).

- **max_negative_step** (*float, default 0.05*) – The maximum magnitude of negative discrete steps allowed in an interval for the interval to be considered valid (units of normalized performance metric).

Returns

- **insolation_weighted_soiling_ratio** (*float*) – P50 insolation-weighted soiling ratio based on stochastic rate and recovery analysis

- **confidence_interval** (*np.array*) – confidence interval (size specified by confidence_level) of insolation-weighted soiling ratio

- **calc_info** (*dict*) –

  - – 'renormalizing_factor' - value used to recenter data

  - – 'exceedance_level' - the insolation-weighted soiling ratio that was outperformed with probability of exceedance_prob

  - – 'stochastic_soiling_profiles' - List of Pandas series corresponding to the Monte Carlo realizations of soiling ratio profiles

  - – 'soiling_ratio_perfect_clean' - Pandas series of the soiling ratio during valid soiling intervals assuming perfect cleaning and P50 slopes

  - – 'soiling_interval_summary' - Pandas dataframe summarizing the soiling intervals identified. The columns of the dataframe are as follows:

Column Name	Description
'start'	Start timestamp of the soiling interval
'end'	End timestamp of the soiling interval
'soiling_rate'	P50 Soiling rate for interval, in day^1 Negative value indicates soiling is occurring. E.g. a rate of 0.01 indicates 1% soiling loss per day.
'soiling_rate_low'	Low edge of confidence interval for soiling rate for interval, in day^1
'soiling_rate_high'	High edge of confidence interval for soiling rate for interval, in day^1
'inferred_start_loss'	Estimated performance metric at the start of the interval
'inferred_end_loss'	Estimated performance metric at the end of the interval
'length'	Number of days in the interval
'valid'	Whether the interval meets the criteria to be treated as a valid soiling interval

7.3.4 System Availability

Functions for detecting and quantifying production loss from photovoltaic system downtime events.

The availability module is currently experimental. The API, results, and default behaviors may change in future releases (including MINOR and PATCH releases) as the code matures.

AvailabilityAnalysis(power_system, ...)	A class to perform system availability and loss analysis.
AvailabilityAnalysis.run([low_threshold, ...])	Run the availability analysis.
AvailabilityAnalysis.plot()	Create a figure summarizing the availability analysis results.

rdtools.availability.AvailabilityAnalysis

class rdtools.availability.**AvailabilityAnalysis**(*power_system*, *power_subsystem*, *energy_cumulative*, *power_expected*)

A class to perform system availability and loss analysis.

This class follows the analysis procedure described in[1], and implements two distinct algorithms. One for partial (subsystem) outages and one for system-wide outages. The *AvailabilityAnalysis.run()* method executes both algorithms and combines their results.

The input timeseries don't need to be in any particular set of units as long as all power and energy units are consistent, with energy units being the hourly-integrated power (e.g., kW and kWh). The units of the analysis outputs will match the inputs.

> **Parameters**
>
> - **power_system** (*pd.Series*) – Timeseries total system power. In the typical case, this is meter power data. Should be a right-labeled interval average (this is what is typically recorded in many DAS).

---

[1] Anderson K. and Blumenthal R. "Overcoming communications outages in inverter downtime analysis", 2020 IEEE 47th Photovoltaic Specialists Conference (PVSC).

- **power_subsystem** (*pd.DataFrame*) – Timeseries power data, one column per subsystem. In the typical case, this is inverter AC power data. Each column is assumed to represent a subsystem, so no extra columns may be included. The index must match power_system. Should be a right-labeled interval average.

- **energy_cumulative** (*pd.Series*) – Timeseries cumulative energy data for the entire system (e.g. meter). These values must be recorded at the device itself (rather than summed by a downstream device like a datalogger or DAS provider) to preserve its integrity across communication interruptions. Units must match power integrated to hourly energy (e.g. if power is in kW then energy must be in kWh).

- **power_expected** (*pd.Series*) – Expected system power data with the same index as the measured data. This can be modeled from on-site weather measurements if instruments are well calibrated and there is no risk of data gaps. However, because full system outages often cause weather data to be lost as well, it may be more useful to use data from an independent weather station or satellite-based weather provider. Should be a right-labeled interval average.

results

Rolled-up production, loss, and availability metrics. The index is a datetime index of the period passed to *AvailabilityAnalysis.run()*. The columns of the dataframe are as follows:

Column Name	Description
'lost_production'	Production loss from outages. Units match the input power units (e.g. if power is given in kW, 'lost_production' will be in kWh).
'actual_production'	System energy production. Same units as 'lost_production'.
'availability'	Energy-weighted system availability as a fraction (0-1).

Type pd.DataFrame

loss_system

Estimated timeseries lost power from system outages.

Type pd.Series

loss_subsystem

Estimated timeseries lost power from subsystem outages.

Type pd.Series

loss_total

Estimated total lost power from outages.

Type pd.Series

reporting_mask

Boolean mask indicating whether subsystems appear online or not.

Type pd.DataFrame

power_expected_rescaled

Expected power rescaled to better match system power during periods where the system is performing normally.

Type pd.Series

energy_expected_rescaled

Interval expected energy calculated from *power_expected_rescaled*.

> **Type** pd.Series

energy_cumulative_corrected
> Cumulative system production after filling in data gaps from outages with estimated production.
>
> > **Type** pd.Series

error_info
> Records about the error between expected power and actual power.
>
> > **Type** pd.DataFrame

interp_lower, interp_upper
> Functions to estimate the uncertainty interval bounds of expected production based on outage length.
>
> > **Type** function

outage_info
> Records about each detected system outage, one row per outage. The primary columns of interest are `type`, which can be either `'real'` or `'comms'` and reports whether the outage was determined to be a real outage with lost production or just a communications interruption with no production impact; and `loss` which reports the estimated production loss for the outage. The columns are as follows:

Column Name	Description
'start'	Timestamp of the outage start.
'end'	Timestamp of the outage end.
'duration'	Length of the outage (*i.e.* `outage_info['end'] - outage_info['start']`).
'intervals'	Total count of data intervals contained in the outage.
'daylight_intervals'	Count of data intervals contained in the outage occurring during the day.
'error_lower'	Lower error bound as a fraction of expected energy.
'error_upper'	Upper error bound as a fraction of expected energy.
'energy_expected'	Total expected production for the outage duration.
'energy_start'	System cumulative production at the outage start.
'energy_end'	System cumulative production at the outage end.
'energy_actual'	System production during the outage (*i.e.*, `outage_info['energy_end'] - outage_info['energy_start']`).
'ci_lower'	Lower bound for the expected energy confidence interval.
'ci_upper'	Lower bound for the expected energy confidence interval.
'type'	Type of the outage ('real or 'comms').
'loss'	Estimated production loss.

> > **Type** pd.DataFrame

Notes

This class's ability to detect short-duration outages is limited by the resolution of the system data. For instance, 15-minute averages would not be able to resolve the rapid power cycling of an intermittent inverter. Additionally, the loss at the edges of an outage may be underestimated because of masking by the interval averages.

This class expects outages to be represented in the timeseries by NaN, zero, or very low values. If your DAS does not record data from outages (e.g., a three-hour outage results in three hours of omitted timestamps), you should insert those missing rows before using this analysis.

References

__init__(*power_system*, *power_subsystem*, *energy_cumulative*, *power_expected*)
   Initialize self. See help(type(self)) for accurate signature.

Methods

init(power_system, power_subsystem, ...)	Initialize self.
plot()	Create a figure summarizing the availability analysis results.
run([low_threshold, relative_sizes, ...])	Run the availability analysis.

rdtools.availability.AvailabilityAnalysis.run

AvailabilityAnalysis.**run**(*low_threshold=None*, *relative_sizes=None*, *power_system_limit=None*, *quantiles=(0.01, 0.99)*, *rollup_period='M'*)
   Run the availability analysis.

   Parameters

   - **low_threshold** (*float or pd.Series, optional*) – An optional threshold used to naively classify subsystems as online. If the threshold is a scalar, it will be used for all subsystems. For subsystems with different capacities, a pandas Series may be passed with index values matching the columns in power_subsystem. Units must match power_subsystem and power_system. If omitted, the limit is calculated for each subsystem independently as 0.001 times the 99th percentile of its power data.

   - **relative_sizes** (*dict or pd.Series, optional*) – The production capacity of each subsystem, normalized by the mean subsystem capacity. If not specified, it will be estimated from power data.

   - **power_system_limit** (*float or pd.Series, optional*) – Maximum allowable system power in the same units as the input power timeseries. This parameter is used to account for cases where online subsystems can partially mitigate the loss of an offline subsystem, for example a system with a plant controller and dynamic inverter setpoints. This constraint is only applied to the subsystem loss calculation.

   - **quantiles** (*2-element tuple, default (0.01, 0.99)*) – The quantiles of the error distribution used for the expected energy confidence interval. The lower bound is used to classify outages as either (1) a simple communication interruption with no production loss or (2) a power outage with an associated production loss estimate.

   - **rollup_period** (*pandas DateOffset or alias, default 'M'*) – The period on which to roll up losses and calculate availability.

rdtools.availability.AvailabilityAnalysis.plot

AvailabilityAnalysis.**plot**()
   Create a figure summarizing the availability analysis results. The analysis must be run using the run() method before using this method.

>
> **Returns fig**
>
> **Return type** matplotlib Figure

7.3.5 Filtering

Functions for filtering and subsetting PV system data.

clip_filter(power_ac[, quantile])	Filter data points likely to be affected by clipping with power greater than or equal to 99% of the *quant* quantile.
csi_filter(poa_global_measured, ...[, threshold])	Filtering based on clear-sky index (csi)
poa_filter(poa_global[, poa_global_low, ...])	Filter POA irradiance readings outside acceptable measurement bounds.
tcell_filter(temperature_cell[, ...])	Filter temperature readings outside acceptable measurement bounds.
normalized_filter(energy_normalized[, ...])	Select normalized yield between `low_cutoff` and `high_cutoff`

rdtools.filtering.clip_filter

rdtools.filtering.**clip_filter**(*power_ac*, *quantile=0.98*)
> Filter data points likely to be affected by clipping with power greater than or equal to 99% of the *quant* quantile.
>
> **Parameters**
>
> - **power_ac** (*pd.Series*) – AC power in Watts
>
> - **quantile** (*float, default 0.98*) – Value for upper threshold quantile
>
> **Returns** Boolean Series of whether the given measurement is below 99% of the quantile filter.
>
> **Return type** pd.Series

rdtools.filtering.csi_filter

rdtools.filtering.**csi_filter**(*poa_global_measured*, *poa_global_clearsky*, *threshold=0.15*)
> Filtering based on clear-sky index (csi)
>
> **Parameters**
>
> - **poa_global_measured** (*pd.Series*) – Plane of array irradiance based on measurments
>
> - **poa_global_clearsky** (*pd.Series*) – Plane of array irradiance based on a clear sky model
>
> - **threshold** (*float, default 0.15*) – threshold for filter
>
> **Returns** Boolean Series of whether the clear-sky index is within the threshold around 1.
>
> **Return type** pd.Series

rdtools.filtering.poa_filter

rdtools.filtering.**poa_filter**(*poa_global*, *poa_global_low=200*, *poa_global_high=1200*)

    Filter POA irradiance readings outside acceptable measurement bounds.

        Parameters

- **poa_global** (*pd.Series*) – POA irradiance measurements.
- **poa_global_low** (*float, default 200*) – The lower bound of acceptable values.
- **poa_global_high** (*float, default 1200*) – The upper bound of acceptable values.

        Returns Boolean Series of whether the given measurement is within acceptable bounds.

        Return type pd.Series

rdtools.filtering.tcell_filter

rdtools.filtering.**tcell_filter**(*temperature_cell*, *temperature_cell_low=-50*, *temperature_cell_high=110*)

    Filter temperature readings outside acceptable measurement bounds.

        Parameters

- **temperature_cell** (*pd.Series*) – Cell temperature measurements.
- **temperature_cell_low** (*float, default -50*) – The lower bound of acceptable values.
- **temperature_cell_high** (*float, default 110*) – The upper bound of acceptable values.

        Returns Boolean Series of whether the given measurement is within acceptable bounds.

        Return type pd.Series

rdtools.filtering.normalized_filter

rdtools.filtering.**normalized_filter**(*energy_normalized*, *energy_normalized_low=0.01*, *energy_normalized_high=None*)

    Select normalized yield between `low_cutoff` and `high_cutoff`

        Parameters

- **energy_normalized** (*pd.Series*) – Normalized energy measurements.
- **energy_normalized_low** (*float, default 0.01*) – The lower bound of acceptable values.
- **energy_normalized_high** (*float, optional*) – The upper bound of acceptable values.

        Returns Boolean Series of whether the given measurement is within acceptable bounds.

        Return type pd.Series

7.3.6 Normalization

Functions for normalizing, rescaling, and regularizing PV system data.

`energy_from_power`(power[, target_frequency, ...])	Returns a regular right-labeled energy time series in units of Wh per interval from a power time series.
`interpolate`(time_series, target[, ...])	Returns an interpolation of time_series, excluding times associated with gaps in each column of time_series longer than max_timedelta; NaNs are returned within those gaps.
`irradiance_rescale`(irrad, irrad_sim[, ...])	Attempt to rescale modeled irradiance to match measured irradiance on clear days.
`normalize_with_expected_power`(pv, ...[, ...])	Normalize PV power or energy based on expected PV power.
`normalize_with_pvwatts`(energy, pvwatts_kws)	Deprecated since version 2.0.0.
`normalize_with_sapm`(energy, sapm_kws)	Deprecated since version 2.0.0.
`pvwatts_dc_power`(poa_global, power_dc_rated)	Deprecated since version 2.0.0.
`sapm_dc_power`(pvlib_pvsystem, met_data)	Deprecated since version 2.0.0.
`delta_index`(series)	Deprecated since version 2.0.0.
`check_series_frequency`(series, ...)	Deprecated since version 2.0.0.

rdtools.normalization.energy_from_power

rdtools.normalization.**energy_from_power**(*power*, *target_frequency=None*, *max_timedelta=None*, *power_type='right_labeled'*)

Returns a regular right-labeled energy time series in units of Wh per interval from a power time series. For instantaneous timeseries, a trapezoidal sum is used. For right labeled time series, a rectangular sum is used. NaN is filled where the gap between input data points exceeds `max_timedelta`. Power_series should be given in Watts.

> **Parameters**
>
> - **power** (`pd.Series`) – Time series of power in Watts
>
> - **target_frequency** (`DatetimeOffset or frequency string, default None`) – The frequency of the energy time series to be returned. If omitted, use the median timestep of `power`, or if `power` has fewer than two elements, use ``power.index.freq`.
>
> - **max_timedelta** (`pd.Timedelta, default None`) – The maximum allowed gap between power measurements. If the gap between consecutive power measurements exceeds `max_timedelta`, NaN will be returned for that interval. If omitted, `max_timedelta` is set internally to the median time delta in `power`. Ignored when `power` has fewer than two elements.
>
> - **power_type** (`{'right_labeled', 'instantaneous'}`) – The labeling convention used in power. Default: 'right_labeled'
>
> **Returns** right-labeled energy in Wh per interval

> **Return type** pd.Series

rdtools.normalization.interpolate

rdtools.normalization.**interpolate**(*time_series*, *target*, *max_timedelta=None*, *warning_threshold=0.1*)

Returns an interpolation of time_series, excluding times associated with gaps in each column of time_series longer than max_timedelta; NaNs are returned within those gaps.

> **Parameters**
>
> - **time_series** (`pd.Series, pd.DataFrame`) – Original values to be used in generating the interpolation
>
> - **target** (`pd.DatetimeIndex, DatetimeOffset, or frequency string`) –
>
>   - If DatetimeIndex: the index onto which the interpolation is to be made
>
>   - If DatetimeOffset or frequency string: the frequency at which to resample and interpolate
>
> - **max_timedelta** (`pd.Timedelta, default None`) – The maximum allowed gap between values in `time_series`. Times associated with gaps longer than `max_timedelta` are excluded from the output. If omitted, `max_timedelta` is set internally to two times the median time delta in `time_series`.
>
> - **warning_threshold** (`float, default 0.1`) – The fraction of data exclusion above which a warning is raised. With the default value of 0.1, a warning will be raised if the fraction of data excluded because of data gaps longer than `max_timedelta` is above than 10%.
>
> **Returns**
>
> **Return type** pd.Series or pd.DataFrame (matching type of time_series) with DatetimeIndex

---

Note: Timezone information in the DatetimeIndexes is handled automatically, however both `time_series` and `target` should be time zone aware or they should both be time zone naive.

---

rdtools.normalization.irradiance_rescale

rdtools.normalization.**irradiance_rescale**(*irrad*, *irrad_sim*, *max_iterations=100*, *method='iterative'*, *convergence_threshold=1e-06*)

Attempt to rescale modeled irradiance to match measured irradiance on clear days.

> **Parameters**
>
> - **irrad** (`pd.Series`) – measured irradiance time series
>
> - **irrad_sim** (`pd.Series`) – modeled/simulated irradiance time series
>
> - **max_iterations** (`int, default 100`) – The maximum number of times to attempt rescale optimization. Ignored if `method = 'single_opt'`
>
> - **method** (`str, default 'iterative'`) – The calculation method to use. 'single_opt' implements the irradiance_rescale of rdtools v1.1.3 and earlier. 'iterative' implements a more stable calculation that may yield different results from the single_opt method.

- **convergence_threshold** (*float, default 1e-6*) – The accept-
  able iteration-to-iteration scaling factor difference to determine convergence.
  If the threshold is not reached after `max_iterations`, raise `rdtools.`
  `normalization.ConvergenceError`. Must be greater than zero. Only used
  if `method=='iterative'`.

Returns Rescaled modeled irradiance time series

Return type pd.Series

rdtools.normalization.normalize_with_expected_power

rdtools.normalization.**normalize_with_expected_power**(*pv*, *power_expected*, *poa_global*,
pv_input='power')

Normalize PV power or energy based on expected PV power.

Parameters

- **pv** (*pd.Series*) – Right-labeled time series PV energy or power. If energy, should *not* be
  cumulative, but only for preceding time step. Type (energy or power) must be specified in
  the `pv_input` parameter.

- **power_expected** (*pd.Series*) – Right-labeled time series of expected PV power.
  (Note: Expected energy is not supported.)

- **poa_global** (*pd.Series*) – Right-labeled time series of plane-of-array irradiance as-
  sociated with `expected_power`

- **pv_input** (*{'power' or 'energy'}*) – Specifies the type of input used for `pv` pa-
  rameter. Default: 'power'

Returns

- **energy_normalized** (*pd.Series*) – Energy normalized based on `power_expected`

- **insolation** (*pd.Series*) – Insolation associated with each normalized point

rdtools.normalization.normalize_with_pvwatts

rdtools.normalization.**normalize_with_pvwatts**(*energy*, *pvwatts_kws*)

Deprecated since version 2.0.0: The normalize_with_pvwatts function was deprecated in rdtools 2.0.0 and will
be removed in 3.0.0. Use normalize_with_expected_power instead.

Normalize system AC energy output given measured poa_global and meteorological data. This method uses the
PVWatts V5 module model.

Energy timeseries and poa_global timeseries can be different granularities.

Parameters

- **energy** (*pd.Series*) – Energy time series to be normalized in watt hours. Must be a
  right-labeled regular time series.

- **pvwatts_kws** (*dict*) – Dictionary of parameters used in the pvwatts_dc_power function.
  See Other Parameters.

Other Parameters

- **poa_global** (*pd.Series*) – Total effective plane of array irradiance.

- **power_dc_rated** (*float*) – Rated DC power of array in watts

- **temperature_cell** (*pd.Series, optional*) – Measured or derived cell temperature [degrees Celsius]. Time series assumed to be same frequency as *poa_global*. If omitted, the temperature term will be ignored.

- **poa_global_ref** (*float, default 1000*) – Reference irradiance at standard test condition [W/m**2].

- **temperature_cell_ref** (*float, default 25*) – Reference temperature at standard test condition [degrees Celsius].

- **gamma_pdc** (*float, default None*) – Linear array efficiency temperature coefficient [1 / degree Celsius]. If omitted, the temperature term will be ignored.

Note: All series are assumed to be right-labeled, meaning that the recorded value at a given timestamp refers to the previous time interval

---

Returns

- **energy_normalized** (*pd.Series*) – Energy divided by PVWatts DC energy.

- **insolation** (*pd.Series*) – Insolation associated with each normalized point

rdtools.normalization.normalize_with_sapm

rdtools.normalization.**normalize_with_sapm**(*energy*, *sapm_kws*)

Deprecated since version 2.0.0: The normalize_with_sapm function was deprecated in rdtools 2.0.0 and will be removed in 3.0.0. Use normalize_with_expected_power instead.

Normalize system AC energy output given measured met_data and meteorological data. This method relies on the Sandia Array Performance Model (SAPM) to compute the effective DC energy using measured irradiance, ambient temperature, and wind speed.

Energy timeseries and met_data timeseries can be different granularities.

Parameters

- **energy** (`pd.Series`) – Energy time series to be normalized in watt hours. Must be a right-labeled regular time series.

- **sapm_kws** (`dict`) – Dictionary of parameters required for sapm_dc_power function. See Other Parameters.

Other Parameters

- **pvlib_pvsystem** (*pvlib-python LocalizedPVSystem object*) – Object contains orientation, geographic coordinates, equipment constants (including DC rated power in watts). The object must also specify either the `temperature_model_parameters` attribute or both `racking_model` and `module_type` to infer the model parameters.

- **met_data** (*pd.DataFrame*) – Measured met_data, ambient temperature, and wind speed. Expected column names are ['DNI', 'GHI', 'DHI', 'Temperature', 'Wind Speed']

Note: All series are assumed to be right-labeled, meaning that the recorded value at a given timestamp refers to the previous time interval

---

Returns

- **energy_normalized** (*pd.Series*) – Energy divided by Sandia Model DC energy.

- **insolation** (*pd.Series*) – Insolation associated with each normalized point

rdtools.normalization.pvwatts_dc_power

rdtools.normalization.**pvwatts_dc_power**(*poa_global*, *power_dc_rated*, *temperature_cell=None*, *poa_global_ref=1000*, *temperature_cell_ref=25*, *gamma_pdc=None*)

Deprecated since version 2.0.0: The pvwatts_dc_power function was deprecated in rdtools 2.0.0 and will be removed in 3.0.0. Use normalize_with_expected_power instead.

PVWatts v5 Module Model: DC power given effective poa poa_global, module nameplate power, and cell temperature. This function differs from the PVLIB implementation by allowing cell temperature to be an optional parameter.

Parameters

- **poa_global** (`pd.Series`) – Total effective plane of array irradiance.

- **power_dc_rated** (`float`) – Rated DC power of array in watts

- **temperature_cell** (`pd.Series, optional`) – Measured or derived cell temperature [degrees Celsius]. Time series assumed to be same frequency as `poa_global`. If omitted, the temperature term will be ignored.

- **poa_global_ref** (`float, default 1000`) – Reference irradiance at standard test condition [W/m**2].

- **temperature_cell_ref** (`float, default 25`) – Reference temperature at standard test condition [degrees Celsius].

- **gamma_pdc** (`float, default None`) – Linear array efficiency temperature coefficient [1 / degree Celsius]. If omitted, the temperature term will be ignored.

Note: All series are assumed to be right-labeled, meaning that the recorded value at a given timestamp refers to the previous time interval

---

Returns power_dc – DC power in watts determined by PVWatts v5 equation.

Return type pd.Series

rdtools.normalization.sapm_dc_power

rdtools.normalization.**sapm_dc_power**(*pvlib_pvsystem*, *met_data*)

Deprecated since version 2.0.0: The sapm_dc_power function was deprecated in rdtools 2.0.0 and will be removed in 3.0.0. Use normalize_with_expected_power instead.

Use Sandia Array Performance Model (SAPM) and PVWatts to compute the effective DC power using measured irradiance, ambient temperature, and wind speed. Effective irradiance and cell temperature are calculated with SAPM, and DC power with PVWatts.

Parameters

- **pvlib_pvsystem** (`pvlib-python LocalizedPVSystem object`) – Object contains orientation, geographic coordinates, equipment constants (including DC rated power in watts). The object must also specify either the

---

> `temperature_model_parameters` attribute or both `racking_model` and
> `module_type` attributes to infer the temperature model parameters.

- **met_data** (*pd.DataFrame*) – Measured irradiance components, ambient temperature,
  and wind speed. Expected met_data DataFrame column names: ['DNI', 'GHI', 'DHI',
  'Temperature', 'Wind Speed']

Note: All series are assumed to be right-labeled, meaning that the recorded value at a given timestamp refers
to the previous time interval

Returns

- **power_dc** (*pd.Series*) – DC power in watts derived using Sandia Array Performance Model
  and PVWatts.

- **effective_poa** (*pd.Series*) – Effective irradiance calculated with SAPM

rdtools.normalization.delta_index

rdtools.normalization.**delta_index**(*series*)

Deprecated since version 2.0.0: The _delta_index function was deprecated in rdtools 2.0.0 and will be removed
in 3.0.0.

Takes a pandas series with a DatetimeIndex as input and returns (time step sizes, average time step size) in hours

Parameters **series** (*pd.Series*) – A pandas timeseries

Returns

- **deltas** (*pd.Series*) – A timeseries representing the timestep sizes of `series`

- **mean** (*float*) – The average timestep

rdtools.normalization.check_series_frequency

rdtools.normalization.**check_series_frequency**(*series*, *series_description*)

Deprecated since version 2.0.0: The _check_series_frequency function was deprecated in rdtools 2.0.0 and will
be removed in 3.0.0.

Returns the inferred frequency of a pandas series, raises ValueError using `series_description` if it can't.

Parameters

- **series** (*pd.Series*) – The timeseries to infer the frequency of.

- **series_description** (*str*) – The description to use when raising an error.

Returns **freq** – The inferred index frequency

Return type pandas Offsets string

7.3.7 Aggregation

Functions for calculating weighted aggregates of PV system data.

---

aggregation_insol(energy_normalized, insolation)	Insolation weighted aggregation

rdtools.aggregation.aggregation_insol

rdtools.aggregation.**aggregation_insol**(*energy_normalized*, *insolation*, *frequency='D'*)

> Insolation weighted aggregation

> **Parameters**
>
> - **energy_normalized** (*pd.Series*) – Normalized energy time series
> - **insolation** (*pd.Series*) – Time series of insolation associated with each *energy_normalized* point
> - **frequency** (*Pandas offset string, default 'D'*) – Target frequency at which to aggregate
>
> **Returns** **aggregated** – Insolation weighted average, aggregated at frequency
>
> **Return type** pd.Series

7.3.8 Clear-Sky Temperature

Functions for estimating clear-sky ambient temperature.

get_clearsky_tamb(times, latitude, longitude)	Estimates the ambient temperature at latitude and longitude for the given times using a Gaussian rolling window.

rdtools.clearsky_temperature.get_clearsky_tamb

rdtools.clearsky_temperature.**get_clearsky_tamb**(*times*, *latitude*, *longitude*, *window_size=40*, *gauss_std=20*)

> Estimates the ambient temperature at latitude and longitude for the given times using a Gaussian rolling window.

> **Parameters**
>
> - **times** (*pd.DatetimeIndex*) – A pandas DatetimeIndex, localized to local time
> - **latitude** (*float*) – Coordinates in decimal degrees.
> - **longitude** (*float*) – Coordinates in decimal degrees. Positive is east of the prime meridian.
> - **window_size** (*int, default 40*) – The window size in days to use when calculating rolling averages.
> - **gauss_std** (*int, default 20*) – The standard deviation in days to use for the Gaussian rolling window.
>
> **Returns** clear sky ambient temperature
>
> **Return type** pd.Series

Notes

Uses data from images created by Jesse Allen, NASA's Earth Observatory using data courtesy of the MODIS Land Group.

- https://neo.sci.gsfc.nasa.gov/view.php?datasetId=MOD_LSTD_CLIM_M

- https://neo.sci.gsfc.nasa.gov/view.php?datasetId=MOD_LSTN_CLIM_M

7.3.9 Plotting

Functions for plotting degradation and soiling analysis results.

degradation_summary_plots(yoy_rd, yoy_ci, ...)	Create plots (scatter plot and histogram) that summarize degradation analysis results.
soiling_monte_carlo_plot(soiling_info, ...)	Create figure to visualize Monte Carlo of soiling profiles used in the SRR analysis.
soiling_interval_plot(soiling_info, ...[, ...])	Create figure to visualize valid soiling profiles used in the SRR analysis.
soiling_rate_histogram(soiling_info[, bins])	Create histogram of soiling rates found in the SRR analysis.
availability_summary_plots(power_system, ...)	Create a figure summarizing the availability analysis results.

rdtools.plotting.degradation_summary_plots

rdtools.plotting.**degradation_summary_plots**(*yoy_rd*, *yoy_ci*, *yoy_info*, *normalized_yield*, *hist_xmin=None*, *hist_xmax=None*, *bins=None*, *scatter_ymin=None*, *scatter_ymax=None*, *plot_color=None*, *summary_title=None*, *scatter_alpha=0.5*)

Create plots (scatter plot and histogram) that summarize degradation analysis results.

Parameters

- **yoy_rd** (*float*) – rate of relative performance change in %/yr

- **yoy_ci** (*float*) – one-sigma confidence interval of degradation rate estimate

- **yoy_info** (*dict*) – a dictionary with keys:

  - YoY_values - pandas series of right-labeled year on year slopes

  - renormalizing_factor - float value used to recenter data

  - exceedance_level - the degradation rate that was outperformed with a probability given by the exceedance_prob parameter in the *degradation. degradation_year_on_year()*

- **normalized_yield** (*pd.Series*) – PV yield data that is normalized, filtered and aggregated

- **hist_xmin** (*float, optional*) – lower limit of x-axis for the histogram

- **hist_xmax** (*float, optional*) – upper limit of x-axis for the histogram

- **bins** (*int, optional*) – Number of bins in the histogram distribution. If omitted, len(yoy_values) // 40 will be used

- **scatter_ymin** (*float, optional*) – lower limit of y-axis for the scatter plot
- **scatter_ymax** (*float, optional*) – upper limit of y-axis for the scatter plot
- **plot_color** (*str, optional*) – color of the summary plots
- **summary_title** (*str, optional*) – overall title for summary plots
- **scatter_alpha** (*float, default 0.5*) – Transparency of the scatter plot

---

> **Note:** It should be noted that the yoy_rd, yoy_ci and yoy_info are the outputs from *degradation.degradation_year_on_year()*.

> **Returns fig** – Figure with two axes
>
> **Return type** matplotlib Figure

rdtools.plotting.soiling_monte_carlo_plot

rdtools.plotting.**soiling_monte_carlo_plot**(*soiling_info*, *normalized_yield*, *point_alpha=0.5*, *profile_alpha=0.05*, *ymin=None*, *ymax=None*, *profiles=None*, *point_color=None*, *profile_color='C1'*)

Create figure to visualize Monte Carlo of soiling profiles used in the SRR analysis.

> **Warning:** The soiling module is currently experimental. The API, results, and default behaviors may change in future releases (including MINOR and PATCH releases) as the code matures.

> **Parameters**
>
> - **soiling_info** (*dict*) – soiling_info returned by *soiling.SRRAnalysis.run()* or *soiling.soiling_srr()*.
> - **normalized_yield** (*pd.Series*) – PV yield data that is normalized, filtered and aggregated.
> - **point_alpha** (*float, default 0.5*) – tranparency of the `normalized_yield` points
> - **profile_alpha** (*float, default 0.05*) – transparency of each profile
> - **ymin** (*float, optional*) – minimum y coordinate
> - **ymax** (*float, optional*) – maximum y coordinate
> - **profiles** (*int, optional*) – the number of stochastic profiles to plot. If not specified, plot all profiles.
> - **point_color** (*str, optional*) – color of the normalized_yield points
> - **profile_color** (*str, default 'C1'*) – color of the stochastic profiles
>
> **Returns fig**
>
> **Return type** matplotlib Figure

rdtools.plotting.soiling_interval_plot

rdtools.plotting.**soiling_interval_plot**(*soiling_info*, *normalized_yield*, *point_alpha=0.5*, *profile_alpha=1*, *ymin=None*, *ymax=None*, *point_color=None*, *profile_color=None*)
Create figure to visualize valid soiling profiles used in the SRR analysis.

> **Warning:** The soiling module is currently experimental. The API, results, and default behaviors may change in future releases (including MINOR and PATCH releases) as the code matures.

Parameters

- **soiling_info** (`dict`) – soiling_info returned by *soiling.SRRAnalysis.run()* or *soiling.soiling_srr()*.

- **normalized_yield** (`pd.Series`) – PV yield data that is normalized, filtered and aggregated.

- **point_alpha** (`float, default 0.5`) – tranparency of the `normalized_yield` points

- **profile_alpha** (`float, default 1`) – transparency of soiling profile

- **ymin** (`float, optional`) – minimum y coordinate

- **ymax** (`float, optional`) – maximum y coordinate

- **point_color** (`str, optional`) – color of the `normalized_yield` points

- **profile_color** (`str, optional`) – color of the soiling intervals

Returns fig

Return type matplotlib Figure

rdtools.plotting.soiling_rate_histogram

rdtools.plotting.**soiling_rate_histogram**(*soiling_info*, *bins=None*)
Create histogram of soiling rates found in the SRR analysis.

> **Warning:** The soiling module is currently experimental. The API, results, and default behaviors may change in future releases (including MINOR and PATCH releases) as the code matures.

Parameters

- **soiling_info** (`dict`) – soiling_info returned by *soiling.SRRAnalysis.run()* or *soiling.soiling_srr()*.

- **bins** (`int`) – number of histogram bins to use

Returns fig

Return type matplotlib Figure

rdtools.plotting.availability_summary_plots

rdtools.plotting.**availability_summary_plots**(*power_system*, *power_subsystem*, *loss_total*, *energy_cumulative*, *energy_expected_rescaled*, *outage_info*)

Create a figure summarizing the availability analysis results.

Because all of the parameters to this function are products of an AvailabilityAnalysis object, it is usually easier to use *availability.AvailabilityAnalysis.plot()* instead of running this function manually.

> **Warning:** The availability module is currently experimental. The API, results, and default behaviors may change in future releases (including MINOR and PATCH releases) as the code matures.

Parameters

- **power_system** (*pd.Series*) – Timeseries total system power.
- **power_subsystem** (*pd.DataFrame*) – Timeseries power data, one column per subsystem.
- **loss_total** (*pd.Series*) – Timeseries system lost power.
- **energy_cumulative** (*pd.Series*) – Timeseries system cumulative energy.
- **energy_expected_rescaled** (*pd.Series*) – Timeseries expected energy, rescaled to match actual energy. This reflects interval energy, not cumulative.
- **outage_info** (*pd.DataFrame*) – A dataframe with information about system outages.

Returns fig

Return type matplotlib Figure

See also:

rdtools.availability.AvailabilityAnalysis.plot()

Examples

```
>>> aa = AvailabilityAnalysis(...)
>>> aa.run()
>>> fig = rdtools.plotting.availability_summary_plots(aa.power_system,
...     aa.power_subsystem, aa.loss_total, aa.energy_cumulative,
...     aa.energy_expected_rescaled, aa.outage_info)
```

7.4 RdTools Change Log

7.4.1 v2.0.0 (October 20, 2020)

Version 2.0.0 adds experimental soiling and availability modules, plotting capability, and includes updates to normalization work flow. This major release introduces some breaking changes to the API. Details below.

API Changes

- The calculations internal to `normalize_with_pvwatts()` and `normalize_with_sapm()` have changed. Generally, when working with raw power data it should be converted to right-labeled energy with `energy_from_power()` before being used with these normalization functions (GH #105, GH #108).

- Remove `low_power_cutoff` parameter in `clip_filter()` (GH #84).

- Many kwargs have changed name (but not input order) to bring nomenclature into closer alignment with the DuraMAT pv-terms project: (GH #185)

  - `aggregation_insol()` first kwarg is now `energy_normalized`.

  - `degradation_year_on_year()`, `degradation_ols()` and `degradation_classical_decomposition()` first kwarg is now `energy_normalized`.

  - `normalized_filter()` input kwargs are now `energy_normalized`, `energy_normalized_low` and `energy_normalized_high`.

  - `poa_filter()` input kwargs are now `poa_global`, `poa_global_low` and `poa_global_high`.

  - `tcell_filter()` input kwargs are now `temperature_cell`, `temperature_cell_low` and `temperature_cell_high`.

  - `clip_filter()` input kwargs are now `power_ac` and `quantile`.

  - `csi_filter()` first two kwargs are now `poa_global_measured`, `poa_global_clearsky`.

  - `normalize_with_pvwatts()` pvwatts_kws dictionary keys have been renamed.

  - `pvwatts_dc_power()` input kwargs are now `poa_global`, `power_dc_rated`, `temperature_cell`, `poa_global_ref`, `temperature_cell_ref`, `gamma_pdc`.

  - `irradiance_rescale()` second kwarg is now `irrad_sim`

Deprecations

- The functions `pvwatts_dc_power()`, `sapm_dc_power()`, `normalize_with_pvwatts()`, and `normalize_with_sapm()` have been deprecated in favor of `normalize_with_expected_power()`. (GH #215)

- `delta_index()` and `check_series_frequency()` (GH #222)

Enhancements

- Add new `soiling` module to implement the stochastic rate and recovery method:

  - Create new class `SRRAnalysis` and helper function `soiling_srr()` (GH #112, GH #168, GH #169, GH #176, GH #208, GH #213)

  - Create functions `monthly_soiling_rates()` and `annual_soiling_ratios()` (GH #193, GH #207)

- Create new module `availability` with the class `AvailabilityAnalysis` for estimating timeseries system availability (GH #131)

- Add new function `normalize_with_expected_power()` (GH #173).

- Add new functions `energy_from_power()` and `interpolate()` (GH #105, GH #108, GH #182, GH #212).

- Add new function `normalized_filter()` (GH #139)

- Add new *plotting* module for generating standard plots (GH #138, GH #131)

- Add parameter `convergence_threshold` to *irradiance_rescale()* (GH #152).

Bug fixes

- Allow `max_iterations=0` in *irradiance_rescale()* (GH #152).

Testing

- Add Python 3.7 and 3.8 to CI testing (GH #135).

- Add CI configuration based on the minimum dependency versions (GH #197)

Documentation

- Create sphinx documentation and set up ReadTheDocs (GH #125).

- Add guides on running tests and building sphinx docs (GH #136).

- Improve module-level docstrings (GH #137).

- Update landing page and add new "Inverter Downtime" documentation page based on the availability notebook (GH #131)

Requirements

- Drop support for Python 2.7, minimum supported version is now 3.6 (GH #135).

- Increase minimum pvlib version to 0.7.0 (GH #170)

- Update requirements.txt and notebook_requirements.txt to avoid conflicting specifications. Taken together, they represent the complete environment for the notebook example (GH #164).

- Add minimum matplotlib requirement of 3.0.0 (released September 18, 2018) (GH #197)

- Increase minimum numpy version from 1.12 (released January 15, 2017) to 1.15 (released July 23, 2018) (GH #197)

Example Updates

- Seed `numpy.random` to ensure repeatable results (GH #164).

- Use *normalized_filter()* instead of manually filtering the normalized energy timeseries. Also updated the associated mask variable names (GH #139).

- Add soiling section to the original example notebook.

- Add a new example notebook that analyzes data from a PV system located at NREL's South Table Mountain campus (PVDAQ system #4) (GH #171).

- Explicitly register pandas datetime converters which were deprecated.

- Add new `system_availability_example.ipynb` notebook (GH #131)

Contributors

- Mike Deceglie (@mdeceglie)
- Kevin Anderson (@kanderso-nrel)
- Chris Deline (@cdeline)
- Will Vining (@wfvining)

7.4.2 v1.2.3 (April 12, 2020)

- Updates dependencies
- Versioneer bug fix
- Licence update

Contributors

- Mike Deceglie (@mdeceglie)

7.4.3 v1.2.2 (October 12, 2018)

Patch that adds author email to enable pypi deployment

Contributors

- Mike Deceglie (@mdeceglie)

7.4.4 v1.2.1 (October 12, 2018)

This update includes automated testing and deployment to support development along with some bug fixes to the library itself, a documented environment for the example notebook, and new example results to reflect changes in the example dataset. It addresses GH #49, GH #76, GH #78, GH #79, GH #80, GH #85, GH #86, and GH #92.

Contributors

- Mike Deceglie (@mdeceglie)
- Adam Shinn (@abshinn)
- Chris Deline (@cdeline)
- nb137 (@nb137)

7.4.5 v1.2.0 (March 30, 2018)

This incorporates changes including:

- Enables users to control confidence intervals reported in degradation calculations (GH #59)
- Adds python 3 support (GH #56 and GH #67)

- Fixes bugs (GH #61 GH #57)
- Improvements/typo fixes to docstrings
- Fixes error in check for two years of data in degradation_year_on_year
- Improves the calculations underlying irradiance_rescale

Contributors

- Mike Deceglie (@mdeceglie)
- Ambarish Nag (@ambarishnag)
- Gregory Kimball (@GregoryKimball)
- Chris Deline (@cdeline)
- Mark Mikofski (@mikofski)

7.4.6 v1.1.3 (December 6, 2017)

This patch includes the following changes:

1. Update the notebook for improved plotting with Pandas v.0.21.0
2. Fix installation bug related to package data

Contributors

- Mike Deceglie (@mdeceglie)
- Chris Deline (@cdeline)

7.4.7 v1.1.2 (November 6, 2017)

This patch includes the following changes:

1. Fix bugs in installation
2. Update requirements
3. Notebook plots made compatible with pandas v.0.21.0

Contributors

- Mike Deceglie (@mdeceglie)

7.4.8 v1.1.1 (November 1, 2017)

This patch:

1. Improves documentation
2. Fixes installation requirements

Contributors

- Mike Deceglie (@mdeceglie)

- Adam Shinn (@abshinn)

- Chris Deline (@cdeline)

7.4.9 v1.1.0 (September 30, 2017)

This update includes the addition of filters, functions to support a clear-sky workflow, and updates to the example notebook.

Contributors

- Mike Deceglie (@mdeceglie)

- Adam Shinn (@abshinn)

- Ambarish Nag (@ambarishnag)

- Gregory Kimball (@GregoryKimball)

- Chris Deline (@cdeline)

- Jiyang Yan (@yjy1663)

7.5 Developer Notes

This page documents some of the workflows specific to RdTools development.

7.5.1 Installing RdTools source code

To make changes to RdTools, run the test suite, or build the documentation locally, you'll need to have a local copy of the git repository. Installing RdTools using pip will install a condensed version that doesn't include the full source code. To get the full source code, you'll need to clone the RdTools source repository from Github with e.g.

```
git clone https://github.com/NREL/rdtools.git
```

from the command line, or using a GUI git client like Github Desktop. This will clone the entire git repository onto your computer.

7.5.2 Installing RdTools dependencies

The packages necessary to run RdTools itself can be installed with `pip`. You can install the dependencies along with RdTools itself from PyPI:

```
pip install rdtools
```

This will install the latest official release of RdTools. If you want to work with a development version and you have cloned the Github repository to your computer, you can also install RdTools and dependencies by navigating to the repository root, switching to the branch you're interested in, for instance:

```
git checkout development
```

and running:

```
pip install .
```

This will install based on whatever RdTools branch you have checked out. You can check what version is currently installed by inspecting `rdtools.__version__`:

```
>>> rdtools.__version__
'1.2.0+188.g5a96bb2'
```

The hex string at the end represents the hash of the git commit for your installed version.

Installing optional dependencies

RdTools has extra dependencies for running its test suite and building its documentation. These packages aren't necessary for running RdTools itself and are only needed if you want to contribute source code to RdTools.

Note: These will install RdTools along with other packages necessary to build its documentation and run its test suite. We recommend doing this in a virtual environment to keep package installations between projects separate!

Optional dependencies can be installed with the special syntax:

```
pip install rdtools[test]   # test suite dependencies
pip install rdtools[doc]    # documentation dependecies
```

Or, if your local repository has an updated dependencies list:

```
pip install .[test]   # test suite dependencies
pip install .[doc]    # documentation dependecies
```

7.5.3 Running the test suite

RdTools uses pytest to run its test suite. If you haven't already, install the testing depencencies (*Installing optional dependencies*).

To run the entire test suite, navigate to the git repo folder and run

```
pytest
```

For convenience, pytest lets you run tests for a single module if you don't want to wait around for the entire suite to finish:

```
pytest rdtools/test/soiling_test.py
```

And even a single test function:

```
pytest rdtools/test/soiling_test.py::test_soiling_srr
```

You can also evaluate code coverage when running the test suite using the coverage package:

```
coverage run -m pytest
coverage report
```

The first line runs the test suite and keeps track of exactly what lines of code were run during test execution. The second line then prints out a summary report showing how much much of each source file was executed in the test suite. If a percentage is below 100, that means a function isn't tested or a branch inside a function isn't tested. To get specific details, you can run `coverage html` to generate a detailed HTML report at `htmlcov/index.html` to view in a browser.

7.5.4 Building documentation locally

RdTools uses Sphinx to build its documentation. If you haven't already, install the documentation depencencies (*Installing optional dependencies*).

Once the required packages are installed, change your console's working directory to `rdtools/docs/sphinx` and run

```
make html
```

Note that on Windows, you don't actually need the `make` utility installed for this to work because there is a `make.bat` in this directory. Building the docs should result in output like this:

```
(venv)$ make html
Running Sphinx v1.8.5
making output directory...
[autosummary] generating autosummary for: api.rst, example.nblink, index.rst, readme_
→link.rst
[autosummary] generating autosummary for: C:
→\Users\KANDERSO\projects\rdtools\docs\sphinx\source\generated\rdtools.aggregation.
→aggregation_insol.rst, C:
→\Users\KANDERSO\projects\rdtools\docs\sphinx\source\generated\rdtools.aggregation.
→rst, C:\Users\KANDERSO\projects\rdtools\docs\sphinx\source\generated\rdtools.
→clearsky_temperature.get_clearsky_tamb.rst, C:
→\Users\KANDERSO\projects\rdtools\docs\sphinx\source\generated\rdtools.clearsky_
→temperature.rst, C:
→\Users\KANDERSO\projects\rdtools\docs\sphinx\source\generated\rdtools.degradation.
→degradation_classical_decomposition.rst, C:
→\Users\KANDERSO\projects\rdtools\docs\sphinx\source\generated\rdtools.degradation.
→degradation_ols.rst, C:
→\Users\KANDERSO\projects\rdtools\docs\sphinx\source\generated\rdtools.degradation.
→degradation_year_on_year.rst, C:
→\Users\KANDERSO\projects\rdtools\docs\sphinx\source\generated\rdtools.degradation.
→rst, C:\Users\KANDERSO\projects\rdtools\docs\sphinx\source\generated\rdtools.
→filtering.clip_filter.rst, C:
→\Users\KANDERSO\projects\rdtools\docs\sphinx\source\generated\rdtools.filtering.csi_
→filter.rst, ..., C:
→\Users\KANDERSO\projects\rdtools\docs\sphinx\source\generated\rdtools.normalization.
→normalize_with_pvwatts.rst, C:
→\Users\KANDERSO\projects\rdtools\docs\sphinx\source\generated\rdtools.normalization.
→normalize_with_sapm.rst, C:
→\Users\KANDERSO\projects\rdtools\docs\sphinx\source\generated\rdtools.normalization.
→pvwatts_dc_power.rst, C:
→\Users\KANDERSO\projects\rdtools\docs\sphinx\source\generated\rdtools.normalization.
→rst, C:\Users\KANDERSO\projects\rdtools\docs\sphinx\source\generated\rdtools.
→normalization.sapm_dc_power.rst, C:
→\Users\KANDERSO\projects\rdtools\docs\sphinx\source\generated\rdtools.normalization.
→t_step_nanoseconds.rst, C:
→\Users\KANDERSO\projects\rdtools\docs\sphinx\source\generated\rdtools.normalization.
→trapz_aggregate.rst, C:
→\Users\KANDERSO\projects\rdtools\docs\sphinx\source\generated\rdtools.soiling.
→C:\Users\KANDERSO\projects\rdtools\docs\sphinx\source\generated\rdtools.soiling.
→soiling_srr.rst, C:
→\Users\KANDERSO\projects\rdtools\docs\sphinx\source\generated\rdtools.soiling.srr_
→analysis.rst
```

(continues on next page)

```
building [mo]: targets for 0 po files that are out of date
building [html]: targets for 4 source files that are out of date
updating environment: 33 added, 0 changed, 0 removed
reading sources... [100%] readme_link
looking for now-outdated files... none found
pickling environment... done
checking consistency... done
preparing documents... done
writing output... [100%] readme_link
generating indices... genindex py-modindex
writing additional pages... search
copying images... [100%] ../build/doctrees/nbsphinx/example_33_2.png
copying static files... done
copying extra files... done
dumping search index in English (code: en) ... done
dumping object inventory... done
build succeeded.

The HTML pages are in build\html.
```

If you get an error like `Pandoc wasn't found`, you can install it with conda:

```
conda install -c conda-forge pandoc
```

The built documentation should be in `rdtools/docs/sphinx/build` and opening `index.html` with a web browser will display it.

7.5.5 Contributing

Community participation is welcome! New contributions should be based on the `development` branch as the `master` branch is used only for releases.

RdTools follows the PEP 8 style guide. We recommend setting up your text editor to automatically highlight style violations because it's easy to miss some isses (trailing whitespace, etc) otherwise.

Additionally, our documentation is built in part from docstrings in the source code. These docstrings must be in NumpyDoc format to be rendered correctly in the documentation.

Finally, all code should be tested. Some older tests in RdTools use the unittest module, but new tests should all use pytest.

CHAPTER 8

Indices and tables

- genindex
- modindex
- search

Python Module Index

r

Index

Symbols

P

R

S

T