Highlights

Decomposition and Inference of Sources through Spatiotemporal Analysis of Network Signals: The DISSTANS Python Package

Tobias Köhne,Bryan Riel,Mark Simons

- Spatial and temporal coverage of GNSS Networks is significantly increasing

- Network analysis software needs to be generic, fast, open-source, and well-documented

- DISSTANS is a Python package optimized for the decomposition of GNSS timeseries

- DISSTANS can incorporate spatial correlation a priori in the analysis

- DISSTANS offers visualization tools, is easily extendable, and can be automated

Decomposition and Inference of Sources through Spatiotemporal Analysis of Network Signals: The DISSTANS Python Package

Tobias **Köhne**[a,*], Bryan **Riel**[b] and Mark **Simons**[a]

[a]*Seismological Laboratory, Division of Geological and Planetary Sciences, California Institute of Technology, 1200 E California Blvd, Pasadena, 91106, CA, USA*

[b]*Department of Earth, Atmospheric and Planetary Sciences, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, 02139, MA, USA*

ARTICLE INFO

ABSTRACT

Dense, regional-scale, continuously-operating Global Navigation Satellite System (GNSS) networks are powerful tools to monitor plate motion and surface deformation. The spatial extent and density of these networks, as well as the length of observation records, have steadily increased in the past three decades. Software to enable the efficient analysis (especially the decomposition) of the ever-increasing amount of available timeseries should have the following desirable qualities: geographic portability, computational speed, automation (minimizing the need for manual inspection of each station), use of spatial correlation (exploiting the fact that stations experience common signals), source code availability, and documentation. We introduce the DISSTANS Python package, which aims to be generic (therefore portable), parallelizable (fast), and able to exploit the spatial structure of the observation records in a user-assisted, semi-automated framework, including uncertainty propagation. The code is open-source, includes an application interface documentation as well as usage tutorials, and is easily extendable. We present two case studies that demonstrate our code, one using a synthetic dataset and one using real GNSS network timeseries.

1. Introduction

Networks of Global Navigation Satellite System (GNSS) stations enable the direct observation of surface displacement down to millimeter accuracy (e.g., Blewitt, 2015). Originally using only the Global Positioning System (GPS) and consisting of only a handful of stations, modern quasi-permanent deployments can incorporate more than 1,000 receivers and take advantage of other GNSS constellations such as the Galileo or BeiDou systems. Analyzing network position timeseries requires awareness of the many processes that affect the observations, both desired and confounding, and an ability to distinguish between them. While dominant constituents like the secular motion of a particular station can usually be inferred by simple linear regression, quantifying less prominent constituents (e.g., displacements due to low-magnitude slow slip events or small-volume magma chamber growth) requires a better understanding of the contributing processes.

Here, we present the DISSTANS Python package to facilitate the temporal and spatial decomposition of GNSS timeseries. The code is written in a generic, fully object-oriented fashion with minimal assumptions as to study location, data units, and sampling frequency. Different data loading methods are implemented that interface with common existing timeseries file formats, but are also easily adapted to new formats. All downstream processing is independent of the original format and origin. To make the code as usable and accessible as possible, it is open-source and heavily commented. The repository includes tagged versions, verbose commit messages, and full documentation. The documentation features tutorials based on synthetic and real timeseries data, a subset of which are presented here. DISSTANS already contains many common processing workflows. These workflows are usable with just a few lines of code, and more are in the planned development roadmap. DISSTANS is parallelized for the most demanding tasks — most notably the model fitting component. We also provide extensive plotting options and graphical user interfaces, simplifying interactions with the data.

Section 2 of this report introduces some key structural decisions and presents a brief literature review of previous work, placing this study in the broader scientific context. Section 3 provides an overview of the modular and flexible structure of the code. To validate our processing, Section 4 contains the analysis of a synthetic network of GNSS stations, as well as results from a real-world application using data from the Long Valley Caldera region, California, USA. Section 5 discusses key design choices. Finally, we end in Section 6 with a brief summary and some possible future avenues for extensions to DISSTANS.

2. Background

The list of scientific questions that can be addressed with GNSS networks is long, and the list of approaches that can be used is even longer. For plate motion, surface deformation, and related fields, the key data are displacement timeseries, i.e., the relative movement over time of a receiver with respect to a defined reference frame. To obtain these timeseries, processing centers start from raw receiver observables (time, pseudoranges, and phases) and take into

✉ tkoehne@caltech.edu (T. Köhne); briel@mit.edu (B. Riel);
simons@caltech.edu (M. Simons)

ORCID(s): 0000-0002-8400-7255 (T. Köhne); 0000-0003-1940-3910 (B. Riel); 0000-0003-1412-6395 (M. Simons)

account a large number of physical processes (e.g., tropospheric and ionospheric travel time delays, gravitational effects, relativistic effects) to produce the best estimate of true receiver position for any given sampling interval (e.g., daily or hourly).

With these displacement timeseries, geoscientists can now interrogate the timeseries: Is the entire signal explained by rigid or elastic plate motion? What are the causes for shortterm or longterm transients? How can we use inter-, co- and postseismic station velocities to constrain fault locking? These questions are usually addressed with geophysical model inversions. Similarly, one might also want to identify noise processes and their statistics (e.g., power-law noise contributions). All of these questions, however, require the *decomposition* of the timeseries into components that are the direct effect of specific physical processes (e.g., hydrological seasonal loading, earthquake offsets and transients, plate motion), and a residual component which is the result of noise processes and imperfect modeling.

In this study, we focus on this intermediate step, and refer to it as simply *timeseries analysis* or *timeseries decomposition*. Therefore, we will refer to the displacement timeseries as produced by GNSS network processing centers as the *raw* or *input* timeseries, and to the different timeseries constituents as *signals*.

2.1. Approaches to Timeseries Analysis

We categorize timeseries analysis tools using three main criteria.

Process-agnostic vs. process-aware

This first criterion aims to distinguish approaches that either make a priori assumptions about the physical processes affecting the data (expecting a certain structure in the data), or alternatively, assume the least possible. For example, fitting a model containing a complete set of basis functions to a timeseries is, in its most generic form, process-agnostic, but fitting a logarithmic decay function to a postseismic transient assumes a specific tectonic process.

One of the major benefits of process-agnostic approaches is that they will usually achieve the "best" fit to the observations — at least measured by the magnitude of the residuals, since that is the main available optimization criterion for such methods. However, over-reliance on the data and its residuals makes these methods susceptible to "overfitting"; i.e., interpreting noise as signal. Similarly, process-agnostic methods have difficulties determining trade-offs between different source processes, for example in the case of Principal Component Analysis (PCA) when signals manifest in multiple principal components, or a single principal component mixes different signals. By contrast, process-aware approaches might ignore parts of the observation if they either (a) do not have an appropriate way of describing the observation (e.g., an unmodeled, temporary offset caused by a heavy snow cover on the antenna), or (b) try to fit a signal with an inappropriate model (e.g., mapping postseismic deformation into the coseismic one); as these approaches

naturally prefer a decomposition that follows the assumed structural patterns.

Parametric vs. non-parametric

This second criterion assesses whether a code estimates parameters (coefficients) for predetermined models to decompose the timeseries. The models can be as complex as desired (high dimensionality, non-linear) — even machine learning algorithms are parametric at their core. Examples for non-parametric decompositions are bandpass filters or basis reprojections like Principal Component Analysis (PCA). Note that this criterion ignores the impact of hyperparameters (e.g., regularization penalties, frequency windows).

With non-parametric approaches, the assumptions, hyperparameters, and computational resources needed are minimal compared to model-based methods, thereby simplifying the problem setup immensely. Furthermore, reducing the influence of hyperparameters translates into a reduction of possible sources of errors. On the other hand, parametric approaches enable a straightforward implementation of the formal covariances between model parameters, and by extension, uncertainties in the predicted timeseries. These approaches can also deal naturally with data gaps. Crucially, a parametric approach is necessary for process-aware studies, because non-parametric approaches have no inherent knowledge about how to group different source processes into components (see above).

Station- vs. network-level

An additional criterion acknowledges the role that spatial information can play in the analysis process. For example, if the same models are fit to every timeseries in a network, regardless of where the stations are located, then the analysis code is not aware of the spatial context. These local, station-level solutions are therefore independent from another. If one recognizes, however, that geophysical signals usually have a spatially coherent signature (assuming sufficiently dense networks), then we can and should incorporate that understanding. For example, PCA makes use of the fact that all stations in the network can potentially see the same source signal (even though the network geometry is neglected). Taking advantage of potential spatial structure is usually advised, although the resulting code complexity and computational costs can become a challenge.

2.2. Previous Work

Considering the diversity of possible approaches, the selection of a certain approach (or the design of a hybrid approach) depends on one's goals and the available data. Additional factors include the ease of software implementation, or possibilities to extend the methods to include ancillary datasets (e.g., rainfall, earthquake catalogs, atmospheric pressure). We review selected published work in the field of timeseries analysis in the context of process-agnostic vs. process-aware, parametric vs. non-parametric, and degree of spatial awareness.

Before high-quality station timeseries became ubiquitous, the QOCA software (Dong et al., 1998) could be used to combine "quasi-observations" (lightly-processed input data from GPS, Electronic Distance Measurements, Satellite Laser Ranging, or Very Long Baseline Interferometry) using a Kalman filter approach. QOCA includes the popular module `analyze_tseri` to estimate linear, episodic, and stochastic motion of the different stations individually in a least-squares-based, process-aware, and parametric framework.

With an increasing number of GNSS stations, more GNSS constellations in general, and more precise understanding of the physical processes affecting GNSS positioning solutions, GNSS networks became common for monitoring surface deformation. Today, the analysis tools developed to produce GNSS displacement timeseries routinely also include simple timeseries analysis functionality. For example, the current iterations of JPL's GipsyX/RTGx (Bertiger et al., 2020) and MIT's GAMIT/GLOBK (Herring et al., 2018) software both contain methods to estimate position, velocity, seasonal variations, offsets[1], and postseismic deformations[2]. These Kalman-filter-based methods are parametric, process-aware, but in contrast to QOCA, not spatially aware (for computational reasons).

For regions where complex geophysical processes are at play (such as near a volcano or in subduction zones), more complex analysis is necessary to distinguish between different processes. A common example is the impact of an un unmodeled transient period on the estimated secular plate velocity. In the following, we present a (non-exhaustive, unordered) selection of tools that start from raw GNSS displacement timeseries to analyze stations exhibiting more complex behavior.

The Network Inversion Filter (NIF), first proposed by Segall and Matthews (1997) and subsequently expanded upon by a variety of studies (e.g., McGuire and Segall, 2003; Bekaert et al., 2016), estimates slip rates on predetermined fault structures from (GNSS or other) observations using a Kalman filter. It is therefore process-aware, and because slip on the modeled faults affect multiple stations, which are jointly used to estimate the slip coefficients, it is also spatially aware. The NIF estimates slip and therefore transient displacement signals non-parametrically, but of course the importance of the hyperparameters specifying the fault geometry and the characteristics of fault slip in time and space (e.g., smoothness) is significant.

The Median Interannual Difference Adjusted for Skewness (MIDAS, Blewitt et al., 2016) algorithm explicitly recognizes the importance of unmodeled steps and shortterm transient deformation in the raw timeseries. Not being a traditional regression scheme, it uses the median of velocities computed from data pairs separated by one year, providing a degree of insensitivity to offsets, small data gaps, and annual seasonal signals if the timeseries is sufficiently long. This process-aware, station-level method is mostly defined by its hyperparameters, although other parameters such as known maintenance and earthquake offset times are used. It is therefore a powerful, largely automated method to estimate secular plate velocities, that does not attempt to extract non-annual seasonal, transient, or decaying signals. MIDAS is at the core of UNR's Nevada Geodetic Laboratory openly-accessible global GNSS timeseries repository (Blewitt et al., 2018).

The Señales y Análisis de Ruido Interactivo (Interactive Signal and Noise Analysis, SARI, Santamaría-Gómez, 2019) software performs process-aware, parametric, station-level regression focusing on an interactive user interface. Least squares or Kalman filtering is used to fit polynomial, sinusoidal, exponential, logarithmic, and step models, allowing for a detailed analysis of functional forms underlying the timeseries. It also contains useful additional functionality such as automatic discontinuity detection, periodogram visualization, and noise characterization.

The Greedy Automatic Signal Decomposition (GrAtSiD, Bedford and Bevis, 2018) algorithm is an iterative, station-level method that focuses on detecting and modeling transient signals in the timeseries. At each iteration, a least-squares regression is performed that includes a linear trend, sinusoidal oscillations, predefined steps, as well as a selection of sparse, transient functions ("multitransients"). Only multitransients that significantly improve the data fit are kept for the next iteration, until a convergence criteria is reached. GrAtSiD can therefore be classified as a parametric approach, that is partly process-aware (for the non-multitransient parts of the regression) and partly process-agnostic (since the multitransients can have a variety of shapes and are not tied to a particular physical source).

Note that while MIDAS, SARI, and GrAtSiD are advanced timeseries analysis tools, they are limited to station-level model fit solutions, and cannot incorporate spatial awareness.

An example for a non-parametric, process-agnostic, and spatially-aware method to decompose timeseries is the variational Bayesian Independent Component Analysis (vbICA, Gualandi et al., 2016), a modern iteration of basis reprojection algorithms particularly suitable for GNSS networks. Its key distinction from traditional PCA/ICA is to recognize that probability density functions for individual components are generally not normally distributed by nature, and alleviates this problem by using mixtures of Gaussians. vbICA therefore allows for a more accurate signal separation, as well as a formal way to incorporate component uncertainties.

Finally, Riel et al. (2014) proposed a method that builds on parametric, process-aware regularized regression and adds a process-agnostic set of B-Spline functions to model transients in a spatially-aware framework. Spatial awareness is achieved by choosing the same dictionary of splines (including different center times and periods) for the transient model component at all stations in the network, and then during the regularization process, promoting coefficients that are seen by multiple nearby stations while penalizing those that are not, resulting in a sparse solution.

---

[1] Automatic detection only by GipsyX.
[2] Only GAMIT/GLOBK.

There are many other studies that have implemented or adapted codes and methods for specific study regions or purposes; an analysis and comparison of which would be beyond the scope of this work.

3. Code Overview

In the following, we present the key properties and design choices made for the DISSTANS package.

Spatial awareness. With GNSS networks becoming more widespread — and more importantly, denser — we should explicitly recognize that groups of nearby stations subject to geophysical processes may experience similar signals. If we only consider each station individually, we may miss the opportunity to identify signals around the noise floor. However, if many stations experience the same signal (with different magnitudes), a joint estimation can theoretically enhance the chance of its detection. Such a method would thereby lower the effective signal-to-noise ratio necessary for signal extraction. DISSTANS allows one to take advantage of the available spatial information by building on the spatiotemporal L0-regularized transient fitting algorithm developed by Riel et al. (2014) (also see Sections 2.2 and 3.4).

Incorporating process knowledge. Similarly, where knowledge about physical processes affecting GNSS timeseries is present (e.g., an inflating magma chamber), such information can theoretically improve model fitting. It is therefore desirable for timeseries analysis methods to both include models that best represent known physical processes, as well as methods that are flexible enough to account for unmodeled, unknown processes. DISSTANS allows for such a distinction by offering a range of process-aware, as well as three process-agnostic models (see Section 3.2).

Scalability. In order to scale well with both the number of stations, as well as the length of the observation record, it is useful to parallelize the computationally demanding parts. DISSTANS includes an option to parallelize the station-level, least squares solutions, as well as the evaluation of the predicted model timeseries including its full model covariance matrix.

Uncertainty estimation. Given the possible complexities of displacement timeseries, a proper interpretation of signal decomposition results can only be made if the trade-offs between and within models and east-north-up components can be quantified. The full, formal model covariances (specifically, between components, but also between models) can be estimated and propagated in the DISSTANS workflow.

Step detection. One omnipresent challenge for researchers when analyzing timeseries is the detection and subsequent estimation (or equivalently, removal) of steps in the data. While there are maintenance catalogs detailing the work being done on stations in a GNSS network (e.g., replacement of antenna, receiver, radome), there is no perfect 1:1 prediction of which type of change introduces a step in the recording and which does not. In fact, differing processing centers might even have different catalogs. Furthermore, improper step removal can significantly affect secular plate velocities as well as the character of GNSS noise (e.g., Santamaría-Gómez and Ray, 2021; Blewitt et al., 2016). However, our imperfect understanding of subsurface structure and composition does not allow us to accurately predict the coseismic offset at every network station for earthquakes of varying magnitudes. Sadly, there is also no fully-automated algorithm that would remove the need for manual inspection (e.g., Gazeaux et al., 2013). Therefore, semi-automated methods that aid the researcher in their quest to model all relevant offsets by providing both automatically estimated "guesses" as well as an interface to rapidly inspect the data are crucial for timeseries analysis codes. To this end, DISSTANS contains an automated step detector (similar to the one in GipsyX, Bertiger et al., 2020), a visualization GUI to inspect the data (see Section S1.2), and loading functions for maintenance records in multiple formats. DISSTANS also features both an empirical (following Blewitt et al., 2016) and an elastic-half-space-based method to determine whether or not to allow a coseismic offsets to be estimated at any given station and time.

Portability and extendability. As new GNSS networks are implemented, and output formats of data processing centers change, the ability to easily incorporate these changes is another desirable quality. DISSTANS separates the data loading tasks from all other analysis steps, such that the former can easily be updated without affecting the latter. Furthermore, to enable the development and integration of new approaches, DISSTANS is written as a modular, extendable framework (in contrast to single-use collections of scripts, see Sections 3.1 and S1.1).

Documentation. A detailed documentation and annotated code are crucial if the software is to be shared and used by more than one person, and the public repository of DISSTANS includes both.

In what follows, we further expand upon these points and the relevant classes, functions, and hyperparameters. DISSTANS is written in Python, and we assume the reader has a basic knowledge of Python. Furthermore, we focus on presenting the structure, methodology and results of the code, with little to no actual sample code. For sample code, please refer to the package documentation. The following typesetting will be used for clarity: classes are capitalized and typeset in bold monospace font (e.g., `Station`) and attributes, properties, variables, methods, functions as well as general code are typeset in regular monospace font (e.g., `parameters` or `import disstans`) with callables (e.g., functions and methods) additionally being trailed with parentheses (e.g., `get_mapping()`).

3.1. Structure

Fig. 1 present the modular code structure of DISSTANS. The highest level of abstraction is the `Network` class, which serves three main purposes. First, for each station in the
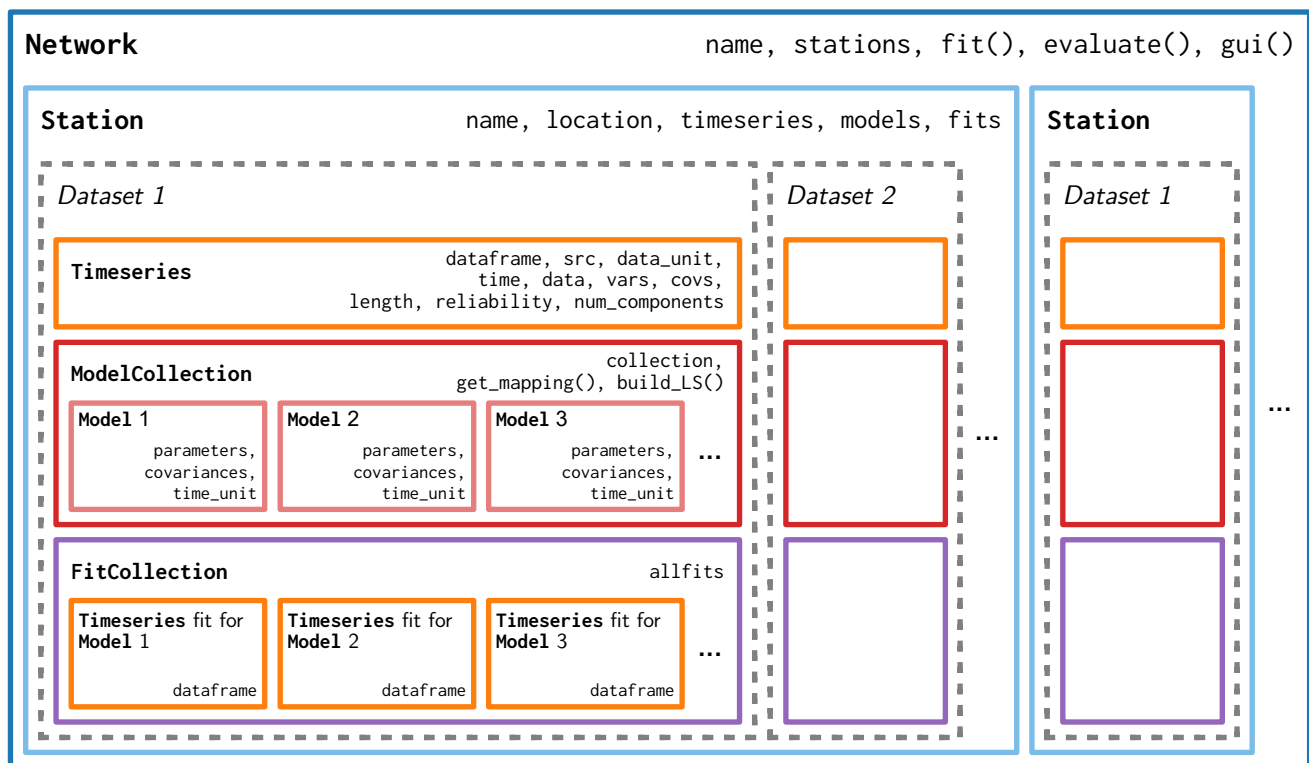
Figure 1: Code structure of DISSTANS, explained in detail in Section 3.1.

network, it contains a `Station` object in its `stations` dictionary attribute, which enables straightforward access. Its second use is to provide a suite of convenience methods that perform a certain task for each station. Without parallelization enabled, their only advantage is that a user does not have to write explicit for-loops, but crucially, `Network` methods also implement an automatic switch to parallelized execution using Python's `multiprocessing.pool` module if the configuration is set accordingly. Finally, the `Network` class contains methods that interface with all stations simultaneously; for example, the graphical user interface `gui()` and other plotting functions (more details about visualization methods in Section S1.2). Plotting functions are based on the standard Matplotlib (Hunter, 2007) and Cartopy (Met Office, 2021) packages.

One level down in the hierarchy is the `Station` class. Apart from storing the metadata information `name` and `location`, it is the container object for all datasets being assigned to the station; for example, raw or post-processed GNSS displacement timeseries (e.g., *Dataset 1* and *Dataset 2* in Fig. 1). A network can contain multiple stations, and each station can contain multiple datasets, but not all datasets have to be present at all stations. The `Station` class also provides functions that directly work on contained timeseries, such as `analyze_residuals()`.

On the third level, for each dataset, a station contains three key elements: the actual data (in the `Timeseries` object, stored in the `Station.timeseries` dictionary), the associated

models (as a `ModelCollection` object containing the individual `Model` objects, stored in the `Station.models` dictionary), and any fits to the data based on model evaluations (as a dictionary of `Timeseries` objects, one for each model, plus one for all models jointly, all stored in the `Station.fits` `FitCollection` object). Using the methods provided by the `Station` class ensures that whenever a new dataset is added (or removed), all three elements are initialized (or deleted) appropriately. While this separation might appear somewhat confusing, it is necessary to enable easy and legible access to individual objects while preserving the flexibility of real-life applications. For example, a `Timeseries` object is physically independent of whatever model one wants to apply to it, and therefore the code should reflect this (i.e., the `Timeseries` object should not change when a model is added or removed, or when an individual model is evaluated to yield a prediction). The separation into data, models, and fits also allows for the same dataset to easily have different models at different stations, or multiple models of the same class (e.g., two sets of step functions, one for maintenance and one for earthquake-induced steps). Lastly, using the `Timeseries` class also for fits (i.e., model-predicted timeseries) allows for the efficient re-use of practical `Timeseries` methods such as file storage or mathematical operations.

On the lowest level, the `Model` and `Timeseries` objects store their data using standard NumPy arrays (Harris et al., 2020) and pandas DataFrames (McKinney, 2010; The pandas development team, 2021), respectively, enabling seamless integration with existing Python-based workflows.

The open-source nature of the code, along with a defined hierarchical, object-oriented structure, allows for easy adaption and extension by the user through subclassing. For example, storing additional station metadata such as antenna information can easily be implemented by creating a Python class inheriting from the `Station` class and extending the initialization function to accept additional instance variables. Another example is the implementation of new user-defined models by subclassing `Model` which then seamlessly integrate into the rest of DISSTANS's workflow. Finally, loading timeseries data from a custom data format can be integrated into DISSTANS by subclassing the `Timeseries` class. In fact, all of the included models (see below) and timeseries file formats are subclasses of `Model` and `Timeseries`, respectively, and can be used as examples by users wishing to extend the code functionality.

3.2. Models

DISSTANS uses a linear combination of parametric models. Parametric models linear in their coefficients (i.e., not necessarily composed of linear functions) allow both simple unregularized as well as more complex L2, L1 or L0 regularized least squares fitting (more detail about regularization schemes in Section 3.4). Furthermore, estimating multiple models jointly is straightforward as their influences just get summed, and the mapping (or design) matrix is simply a horizontal stack of all the models' individual mapping matrices (everything automatically done by the `ModelCollection` class). Lastly, the formal estimated model parameter covariance matrix can usually be estimated in a closed-form way. A mathematical overview over the modeling is given in Section S1.3.

The individual `Model` classes included in DISSTANS can be separated into basic and spline models. All models can be used with one or multiple data components. The basic models included in the package are: `Polynomial`, `Step`, `Sinusoid`, `Logarithmic`, `Exponential`, `HyperbolicTangent` and `Arctangent`. They have in common that they either just have a single function (e.g., the logarithm), or that their functions form orthogonal bases within their class (e.g., polynomials). The spline modeling in `BSpline` or `ISpline` model is based on Hetland et al. (2012); Riel et al. (2014) and contains multiple cardinal B- or integrated-B-splines (respectively) of the same timescale and order but with different center times. The `SplineSet` combines several `BSpline` or `ISpline` models of different timescales into one large collection, forming a linearly-dependent (overcomplete) spanning set able to approximate arbitrary functions. The `AmpPhModulatedSinusoid` estimates a sinusoid of a given nominal frequency, but allows the instantaneous amplitude and phase to vary. Time-varying properties are enabled by modeling the linear sine and cosine coefficients of the sinusoid as being defined by a linearly-independent set of B-Spline basis functions over the given time interval. Since there is an infinite number of ways to fit a given input signal with spline models, some form of regularization is necessary to gain a meaningful result.

3.3. Synthetic Data

The creation of synthetic data is another feature directly integrated into DISSTANS. Each `Model` and `ModelCollection` object has the two methods `read_parameters()` and `evaluate()`, which integrate into existing Python workflow by accepting and returning (respectively) NumPy arrays. A typical workflow to generate datasets therefore is to instantiate `Model` objects (e.g., a polynomial of a certain order), define and read in the parameters of the model, and finally evaluate the individual models (or a `ModelCollection` containing the individual ones). If the data is then to be used within DISSTANS, a simple `Timeseries` constructor exists for NumPy arrays, otherwise one can use the regular NumPy methods for exporting the data.

3.4. Solver Functions

The provided solver functions are least squares (therefore parametric) solvers, with varying degrees of added complexity. They each

1. Build the mapping and observation matrices for a given `Timeseries` object of observations and `ModelCollection` object ($\mathbf{G}$ and $\mathbf{d}$, respectively, see Section S1.3),
2. Divide the solution process into independent subproblems if there is no data component covariance (decreasing the computational burden),
3. Call a lower-level solver to minimize the cost function $\|\mathbf{Gm} - \mathbf{d}\|_2^2$ (potentially subject to regularization; also see Section S1.4),
4. Optionally calculate the formal model parameter covariance matrix $\mathbf{C}_m$ (also see Section S1.4), and
5. Return a `Solution` object (containing the best-fitting $\mathbf{m}$).

To prevent convergence or numerical issues, the solvers and the `Solution` class keep track of model parameters that cannot be estimated (because they are not observable given the timespan of the observations) or should not be estimated (useful, for example, if some splines in a `SplineSet` are assumed to be zero). The regularized solvers additionally keep track of which model's parameters should be regularized, allowing for a flexible regularization approach.

The first, most basic solver function is `linear_regression()`, which essentially just provides the above-mentioned features as a wrapper to the least squares routine in SciPy (Virtanen et al., 2020). It can therefore be regarded as a minimal code example for new, user-defined solvers. The second provided solver, `ridge_regression()`, adds L2 regularization, and also relies on the least squares routine in SciPy.

The third solver, `lasso_regression()`, uses CVXPY (Diamond and Boyd, 2016; Agrawal et al., 2018) to provide L1 and, by means of weighted iterations, station-specific L0 regularization (Candès et al., 2008). L1 regularization is a sparsity-promoting regularization scheme that penalizes the absolute magnitudes of the estimated parameters, driving many parameters close to zero in the process. L0 regularization goes a step further by penalizing the *existence* of a parameter, not its magnitude, thereby either driving parameters
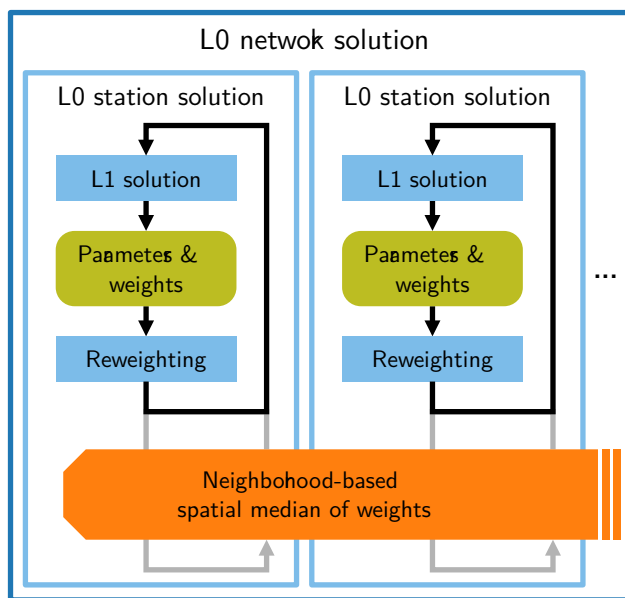
Figure 2: Flowchart of the spatiotemporal L0-regularized solver as described in Riel et al. (2014). Symbols and colors from Fig. 3. At each station, an L1-regularized least-squares fit is computed, where each parameter has an associated weight. The weight is inversely correlated to the parameter magnitude. Parameters close to zero are iteratively penalized, whereas significant parameters have their penalty gradually reduced to zero. Iterated L1 regularization effectively approximates an L0-regularized solution (see Candès et al., 2008). By combining the weights between stations with a median in an intermediate step, parameters that are significant at other nearby stations as well are promoted, and parameters that are insignificant are demoted.

to zero, or not penalizing a parameter at all. An appropriate threshold needs to be set such that the algorithm can distinguish between insignificant and significant parameters. L0 regularization is therefore more suited for physical processes which occur sporadically, are not ubiquitous, and have an "arbitrary", but significant, magnitude. Thresholds are not hard cut-offs, but rather are defined within the context of `Reweighting Function` objects. While the appropriate choice of functions and scales will vary between applications, a good (empirical) starting point are functions whose penalties close to an input value of zero are of a similar order of magnitude of the data being fitted.

Lastly, the `Network.spatialfit()` method extends the possibilities of station-specific L0 regularization to also take into account the weights of a given model at nearby stations. The approach implemented here follows Riel et al. (2014) closely, with the goal to identify signals close to the noise floor, suppress local noise, and promote sparse models in both time *and* space. A visual summary of the method is given in Fig. 2. DISSTANS is able to perform the station-specific fits in parallel, resulting in a large runtime improvement.

3.5. Example Workflow

Even though DISSTANS is modular and therefore highly flexible, we propose the workflow presented in Fig. 3 for the following real-world example and as a general starting point for timeseries analysis with DISSTANS.

The first step is the acquisition and preparation of the raw input datasets: in this case, GNSS network station displacement timeseries (and, if available, associated maintenance and seismic catalogs). Applying quality metrics such as requiring a minimum number of observations or station reliability (through their respective `Timeseries` attributes `num_observations` and `reliability`) ensures that the fitting process is not hindered by bad data.

We view the second step as a "preprocessing" one, where we identify and remove statistical outliers and the common mode error (see Dong et al., 2006; Huang et al., 2012) from the observations (see Fig. 4 for more detail). The relevant functions are `median()`, `clean()` and `common_mode()`, which are called on the entire network (respecting parallelization) through the `Network` methods `call_func_ts_return()` and `call_func_no_return()`. The common mode is an error that is systematic for the entire network and exists even if all geophysical processes have been estimated and removed perfectly, as it is the result of the imperfections during the production of the raw input displacement timeseries in the first place. It manifests itself as a high-frequency noise realization that should be estimated independently of model fits (which could create additional systematic errors). To estimate the common mode error (e.g., using PCA/ICA), we first remove empirically all the potentially interesting, low-frequency signal using a running median. A median filter, compared to general high-pass filtering, has the benefit of being more robust when handling large steps in the data (which are of course omnipresent before any step removal is performed). Outlier removal is performed on the residual between low-pass and input signal, based on the residual's variance.

The biggest obstacle for model fitting are offsets (or steps) in the data. Left unaccounted for, they will influence every other model component (e.g., the secular plate velocity). While big jumps in the data can easily be spotted by comparing a measurement with the variance around the mean of previous observations, smaller offsets that are either below or similar to the data variance, and/or are accompanied by transient motion, are more challenging to detect. Ideally, all occuring offsets are known in advance based on ancillary catalogs, and could be categorized into equipment changes and physical processes.

Maintenance events (e.g., antenna replacements, software changes, receiver upgrades) usually are well-recorded and accessible. Functions like `parse_maintenance_table()` and `parse_unr_steps()` are useful for these purposes. However, not all maintenance events automatically have a visible effect in the data, and therefore there are "grey zones" where the addition of a modeled step may be more harmful than beneficial. In these cases, we can perform an iterative
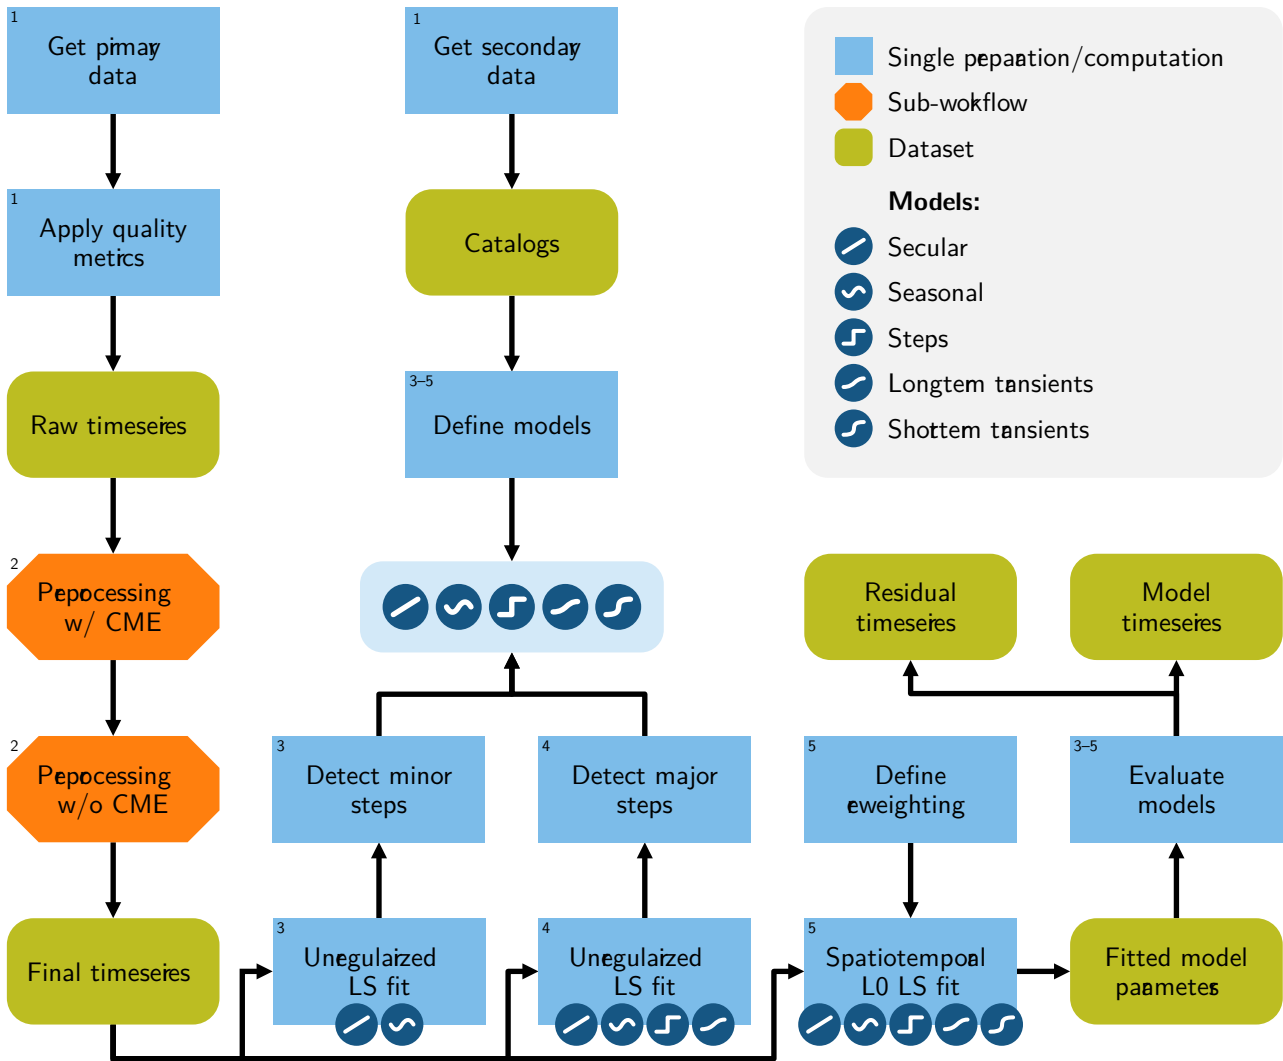
Figure 3: Example workflow for using DISSTANS, explained in detail in Section 3.5. Blue rectangles represent single computational steps, orange rectangles with cut corners represent sub-workflows discussed in more detail elsewhere, and green, rounded rectangles represent datasets at their different stages of processing. The numbered steps in the text correspond to the numbering in the top left corners of the rectangles.

process between fitting larger signals, and then checking again for evidence of smaller offsets.

A similar case can be made for the presence of coseismic displacements. Large, nearby earthquakes produce offsets that can be predicted from seismic catalogs and simple forward modeling of the expected displacement at any given station. (The `earthquakes` module provides this functionality in DISSTANS.) However, smaller events might not necessarily warrant an additional modeled step, and very fast transients would be better fit by transient models. Therefore, we recommend an iterative approach here as well.

The next steps in the proposed workflow are therefore iterations of step-detection and model-fitting. In the third step, an unregularized least-squares fit with only a polynomial and some sinusoidal models is performed at each station individually. Using the `StepDetector` class, extremely

prominent offsets in the data are well resolved, and are added to a list of offsets to be fit (with the `Step` model class).

In the fourth step, using the initial simple models, the defined list of offsets, and a `SplineSet` dictionary of longterm transient splines, another unregularized least-squares fit is computed. Together with external maintenance and seismic catalogs, a second run of the `StepDetector` then aims to identify smaller steps that are to be estimated.

For the fifth step of the proposed workflow, the `spatialfit()` method is used to perform a network-wide, spatiotemporal, L0-regularized fit using the aforementioned polynomial, sinusoidal, and step models, as well as an expanded spline dictionary that includes also shorter-term transients. (Only the spline parameters are included in the regularization process, although this setting can be changed.) The definition of an appropriate `ReweightingFunction` is necessary to ensure a sparse, yet well-fitting solution. When seasonal effects are
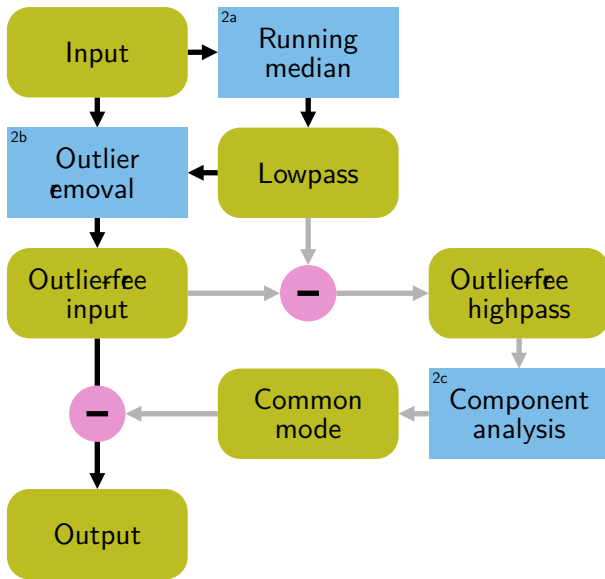
4. Validation

We present two validation datasets and results. The first, in Section 4.1, is a synthetic dataset of 16 stations exhibiting some commonly seen patterns in GNSS network timeseries. Using this synthetic network, we demonstrate key capabilities of this code in estimating complex signals, all while being able to compare fitted models to the true underlying timeseries. The second dataset, in Section 4.3, is a collection of GNSS stations in the Long Valley Caldera region in California, USA. Here, the main goal is to illustrate the example workflow, and discuss some subtleties when dealing with imperfect, real-world data.

4.1. Synthetic Dataset

The code for this analysis, the synthetic model parameters, quantitative fitting statistics, convergence measurements, as well as additional discussion, can be found in Tutorial 3 of the online documentation.

One of the main features of DISSTANS is its ability to use spatial coherence as an additional source of information and constraint. In general, signals like earthquakes, slow fault slip events, or seasonal loading are spatially correlated, as the processes affecting each station have the same underlying sources. By using this knowledge in combination with the enforcement of sparsity, we can make sure that the estimated models are consistent between stations.

On the flipside, processes that only affect a single station are considered noise for the purposes of this study (e.g., antenna maintenance or strongly localized displacements). By making it harder for a station to fit a model when no station around it sees the same signal (as done by the spatial L0 regularization), its residuals will become very prominent around the time in question, and it is therefore easier to detect processes that need to be fitted, or removed, for the station in question.

Setup

The synthetic dataset is comprised of 16 stations randomly positioned on an elongated, rectangular grid (see Fig. 5). Each two-component station is affected by a secular, linear trend, one annual seasonal signal, an earthquake (with both co- and postseismic components), two shortterm slow slip events, one longterm transient, common mode error, and measurement error (correlated between the components). The linear trend, coseismic and postseismic signals are all equal in direction and magnitude, whereas the seasonal signal is random at each station. The three transients are all equal in onset time, duration, and direction, but differ in magnitude. Furthermore, one station ("Cylon") experiences significant powerlaw noise, and a different station is affected by an unmodeled maintenance step. Lastly, the amplitudes of the three transients decrease exponentially towards the east.

The processing follows a simplified version of the example workflow presented in Section 3.5 and Fig. 3. Because the data is synthetic, no quality metrics need to be applied, nor is the second preprocessing step without CME or the

Figure 4: Preprocessing sub-workflow, following the same symbolic and coloring as Fig. 3 (step 2), with rose circles representing mathematical operations. First, a running median of the input is calculated, which results in a lowpass filtered timeseries. The variance of the input around the lowpass timeseries is used to detect outliers. Removing them from the input yields the outlier-free input. Without common mode estimation, this is also the final output. To remove the common mode, the difference between the lowpassed input and the outlier-free input is calculated, which yields an outlier-free, highpassed input. The dominant component of this timeseries is the best estimate of the common mode error. Removing this from the outlier-free input yields the outlier-free, common-mode-removed output.

found to be strongly varying, allowing the seasonal signal to vary amplitude over time (using `AmpPhModulatedSinusoid` models), can also greatly improve the fit to all other models.

The final (as well as all intermediate) results are therefore a set of model parameters, in each data component, together with a complete parameter covariance matrix (if desired). They can be evaluated at all stages to yield the overall model-predicted timeseries (including its predicted uncertainty), as well as the individual model contributions. The residuals can always be computed using the `Network.math()` methods and analyzed using the `Network.gui()` method to assure no systematic misfit is present.

There are many variations to this example workflow. For example, one might already start the first unregularized least-squares computation with some predefined, well-known steps modeled or removed. Or, one might want to add a cleaning postprocessing stage at the end, or have a user-defined process act on the data between steps. Because of its modular and standard-Python architecture, this is all easily implemented with DISSTANS.
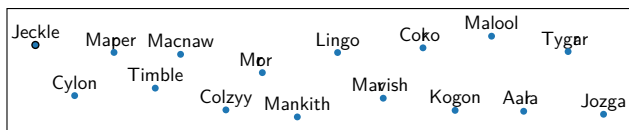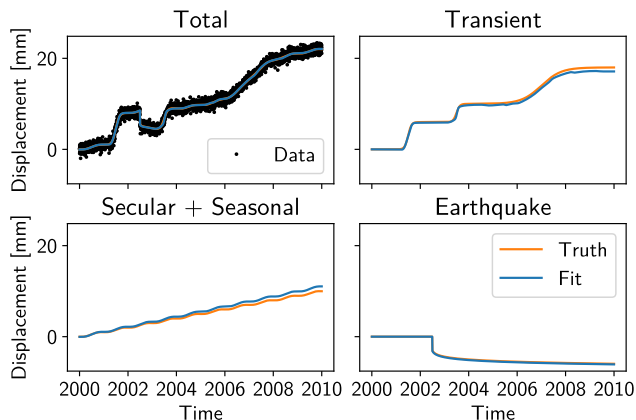
Figure 5: Map view of the synthetic network.

Figure 6: Overall model fit to the data at station Jeckle (see Fig. 5), including the decomposition into the primary constituents (North component only).
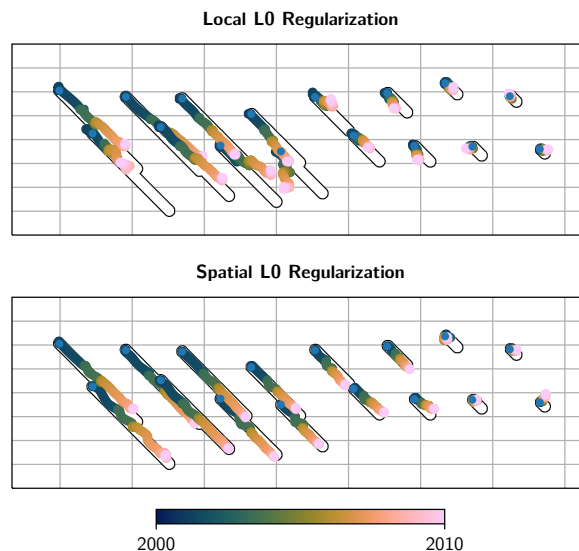

Figure 7: Map view of the transient motion of the synthetic network over the entire timespan. The top panel shows the result without spatial regularization, the bottom one with. Markers correspond to the position of a station relative to its initial position, with colors corresponding to time. The white background curves with black outlines are the true synthetic transient, which is clearly better matched by the solution incorporating spatial awareness.

multiple step detection steps necessary. Instead, to demonstrate the benefits of using the spatially-aware L0 regularization, its output will be compared to results obtained by the station-specific L1 regularization, as well as the local L0 regularization.

Fig. 6 shows the north component of a representative station. The inferred model fit the synthetic data well. A small tradeoff can be observed between the secular and transient models, although it should be noted that in real world applications, such a conclusion is frequently difficult.

Regularization effect

To analyze the sparsity and spatial coherence of the transient signals, scalograms for all three regularization schemes at stations Jeckle (see Fig. 6) and Cylon are presented in Fig. S4 (respective timeseries are shown in Fig. S3). Using the L1 solver, the transients (two shortterm, one longterm afterwards) are sparsely fitted in time. However, for the same transients (i.e., same shape, centertime, and period), different splines end up being "chosen" by the solver (as a result of different noise realizations) at the two stations. For signals as significant as the ones present at the two stations, this is not problematic; but in the context of identifying signals closer to the noise level that are appearing at multiple stations (and conversely, identifying shortterm noise processes that are not seen by other stations), it is preferable that a given signal is fit by the same splines (in time and period).

Going from the L1 to the (local) L0 regularization, the situation does not improve, even though there is a slight reduction in the number of non-zero splines overall. Adding the spatial component, however, accomplishes two things:

(1) the same transients are fit by the same splines (with slightly differing magnitudes of course), and (2) some colored noise at station Cylon (previously present as shortterm splines) is suppressed since it is not seen by the nearby stations Jeckle and others. The result is therefore a dictionary of splines that are both sparse in time *and* space.

Fig. 7 shows the improvement from local to spatial L0 regularization in map view for all stations: the transient components are smoother (therefore fitting less noise) and more closely follow the true signal (shown in the background). Importantly, the homogenous displacement field is obtained without degrading the fit to the data. Fig. S3 shows timeseries plots for both (east and north) data components at two neighboring stations (Jeckle and Cylon) for all three regularization schemes, and they are virtually undistinguishable (except for the overfitting of some colored noise as discussed above).

Model parameter correlations

Since the elemental solver procedures follow least-squares principles, it is possible to compute model correlations: between components (component covariance), within a model (model covariance), and between models (component/model cross-covariance). Fig. 8 shows the correlation for the spatial L0-regularized solution, i.e., the normalized covariances. Tradeoffs between models and within splines are clearly identifiable, thereby giving a closer insight into the fitting process and allowing for statements about parameter uncertainties.
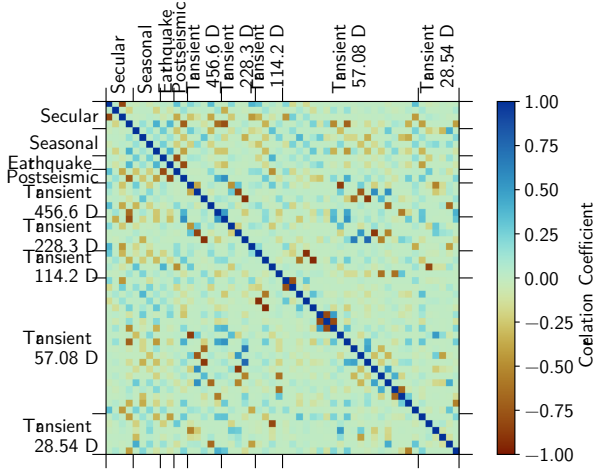
Figure 8: Model parameter correlation matrix for both data components. The covariances for spline parameters that are estimated to be close to zero are set to zero as well, and not shown.
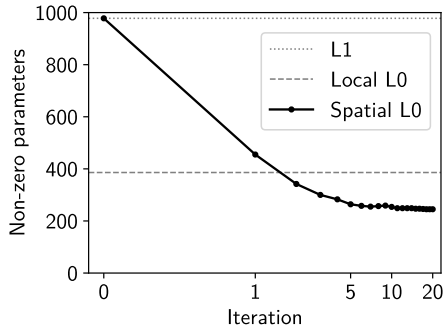

Figure 9: Convergence of the iterative, spatiotemporal L0-regularized solver. Shown are the decreasing and stabilizing total number of unique parameters fitted (curve), as well as the number of unique parameters for the station-specific L0- and L1-regularized solutions (horizontal lines). In total, there are 8416 parameters.

Convergence

Finally, to validate that the spatiotemporal L0 regularization is also numerically stable and converges, Fig. 9 shows how the total number of non-zero parameters declines asymptotically. Fig. S5 shows that this decrease can both be seen in the total number of non-zero parameters (i.e., the sum over all stations and components of the number of non-zero coefficients) as well as the number of *unique* non-zero parameters (i.e., the number of all splines that are non-zero at least at one station, per component). The latter number specifically demonstrates the effect of spatial sparsity.

4.2. Influence of Number of Stations

The code for this analysis, the synthetic model parameters, as well as the exploration of additional explored hyperparameters, can be found in Tutorial 5 of the online documentation.

To further validate our claim that incorporating information from nearby stations improves the quality of the

resulting model fit, we here explore the dependence of the model error on the number of stations used. In this section, we use a new synthetic network of $N = 20$ stations, distributed randomly, that is only affected by a single transient process and white noise. The noise level relative to the maximum amplitude of the transient signal, σ, is one of the hyperparameters we vary. The other variable is the number of stations $2 \leq n \leq N$ used by the spatial L0-regularized solver. For each test case, we therefore subsample the original network to create a subnetwork of smaller size n, comprised of randomly selected stations. (We also calculate the result of using a local L0-regularized solver for comparison, where by construction $n = 1$). The number of samples m, for each n to test, is given by the maximum of either the amount of possible permutations, or a defined maximum value M based on computational considerations ($M = 50$ in our case).

For each σ and each n, we therefore have m samples to test. The metric we choose to compare is the root-mean-squared true model error (RMSE), calculated from the final fit of each sampled subnetwork (ensuring the solvers iterate long enough to converge). For each n, we therefore compute the double mean of the RMSE, ϵ, first across the subnetwork, and then across samples. We also compute the standard deviation σ_ϵ of the samples of the subnetwork-wide mean RMSEs.

Fig. 10 shows the results of our experiment. For all of the cases, the mean RMSE ϵ decreases with increasing number of stations used in the fitting process (approximately by $1/\sqrt{n}$). Furthermore, the variance of the errors decreases as well. Importantly, for the case of $\sigma = 3$ (i.e., the white noise standard deviation is three times the maximum magnitude of the transient signal), the local L0-regularized solution has a high error variance centered close to the maximum allowable error (defined as not fitting a transient at all). Including multiple stations in the estimation process, however, decreases the mean error and error variance significantly — with 20 stations, as low as the mean error for the local L0-regularized solution for $\sigma = 1$. In the highest noise case presented here, $\sigma = 10$, most local L0-regularized solutions actually overfit the data. Incorporating spatial awareness prevents the solver to do so. Overall, as shown by the reduction of error, error variance, and susceptibility to overfitting, the importance of using spatial awareness for transient model fitting becomes clear.

4.3. Long Valley Caldera

The code for this analysis, processing settings, model- and catalog details, as well as additional discussion, can be found in Example 1 of the online documentation.

To validate DISSTANS with real data, we consider time-series from the Long Valley Caldera (LVC) region in California, USA. Because of the geophysical interest into the magmatic, seismic, and hydrological processes at work there, the LVC has been monitored by an ever-extending network of GNSS stations since the late 1990s (e.g., J et al., 2013; Montgomery-Brown et al., 2015; Silverii et al., 2020).
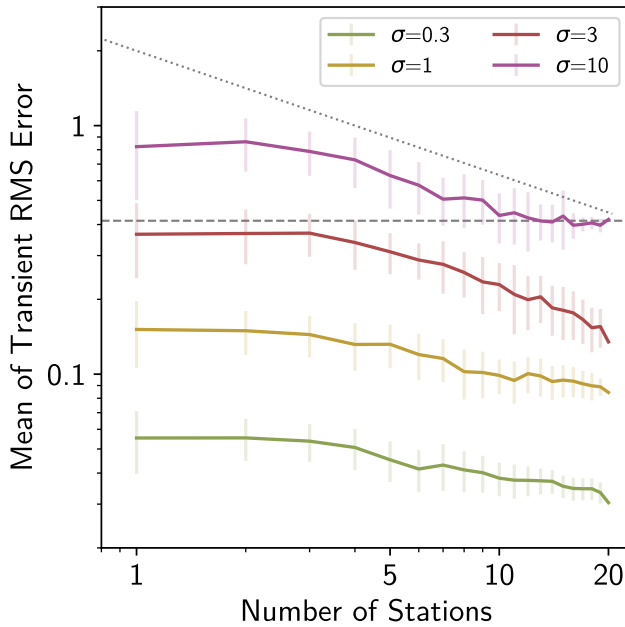
Figure 10: Sample mean (ϵ, colored lines) and sample standard deviation (σ_ϵ, vertical errorbars) for all the sampled subnetworks as a function of the number of stations used in the solution process (n, horizontal axis), and noise level ratio (σ, different colors). The dashed, horizontal grey line corresponds to the maximum allowable error if no transient signal is fitted at all. The dotted grey line is a reference line parallel to $1/\sqrt{n}$.

The displacement timeseries are complemented by detailed maintenance and seismic catalogues, which are crucial for determining the best set of steps to include in the fitting process. In this example, the goals are threefold: (1) to illustrate the example workflow proposed in Section 3, (2) to present the best-fit transient model to the periods of unrest in the Long Valley Caldera, and (3) to showcase the importance of allowing the seasonal signal models to vary in amplitude over time. Any in-depth physical modeling of the extracted signals is beyond the scope of this study.

Setup

The processing follows the example workflow presented in Fig. 3. The data and corresponding maintenance and seismic events catalog are downloaded with DISSTANS-included tools from the GNSS timeseries repository maintained by the University of Nevada at Reno's Nevada Geodetic Laboratory (Blewitt et al., 2018). Only stations with a reliability of over 50% and an observation record at least one year long are considered, and outliers in each timeseries (more than 10 standard deviations away from the median) as well as the common mode error are removed.

With help of the available maintenance catalog, we iteratively identify steps in the data. This process is aided significantly by DISSTANS's included step detector and visualization routines. Finally, we compute the spatiotemporal L0-regularized fit.

Transient Signals

The timespan between 2012 and 2015 (approximately) is dominated by a significant expansion of the caldera's dome, as observed by both the GNSS network and Interferometric Synthetic Aperture Radar (InSAR) timeseries (Montgomery-Brown et al., 2015; Silverii et al., 2020).

Fig. 11 shows the horizontal component estimated by DISSTANS in map view: the radial extension of the network from the center of the dome is clearly visible. Fig. 12 shows the extracted transient signal of selected stations in and around the resurgent dome, projected along the direction of maximum transient displacement during the period of maximum unrest, for the entirety of the available data, and ordered by azimuth. Two periods of significant expansion can clearly be distinguished: around 2002–2003, and from 2011–2021. A smaller period of unrest can be seen throughout the network between approximately 2008–2010, and the station CASA allows us to see a period of extremely rapid expansion around 1998. These results are comparable to Silverii et al. (2020, Fig. S3a), where transients were recovered using non-parametric multiyear filters, even though the directions of maximum displacements are different. Crucially, however, we did not enforce the secular long-term motion to be zero during a specific timespan. As a result, many stations appear to never reach a "steady-state" matching the general plate motion, because the transient motion, even when regularized, is dominant for large parts of each timeseries. The a priori removal of a secular trend is easily done with DISSTANS, and additionally, given the computation of the full model parameter covariance matrix, the adjustment could even be made analytically after the model fit.

Seasonal Signals

Traditional least-squares model fitting for GNSS timeseries usually either approximate the seasonal signal as having a constant amplitude and phase over the entire timespan considered (or piecewise within that timeseries) (e.g., Heflin et al., 2020), or estimate a more accurate seasonal deformation signal from filtering or component-analysis methods (e.g., Silverii et al., 2020). The two approaches are usually acceptable, as either the resulting residual signals are small enough not to affect any other significant model, or are not prone to producing large seasonal residuals in the first place (respectively). Given our transient modeling of even small timescales (down to the order of less than 100 days), our method does suffer from these seasonal residuals, as annual rain- and snowfall can vary widely, especially in the Sierra Nevada. In fact, because seasonal residual are highly correlated between stations, they are not even removed by our spatial L0 regularization. Modeling the seasonal signal as the sum of both an unregularized, constant, nominal signal (using `Sinusoid`), and a simple, L1-regularized, station-specific deviation model of the same nominal frequencies that is allowed to vary amplitude (and by construction, instantaneous phase) over time (using `AmpPh`
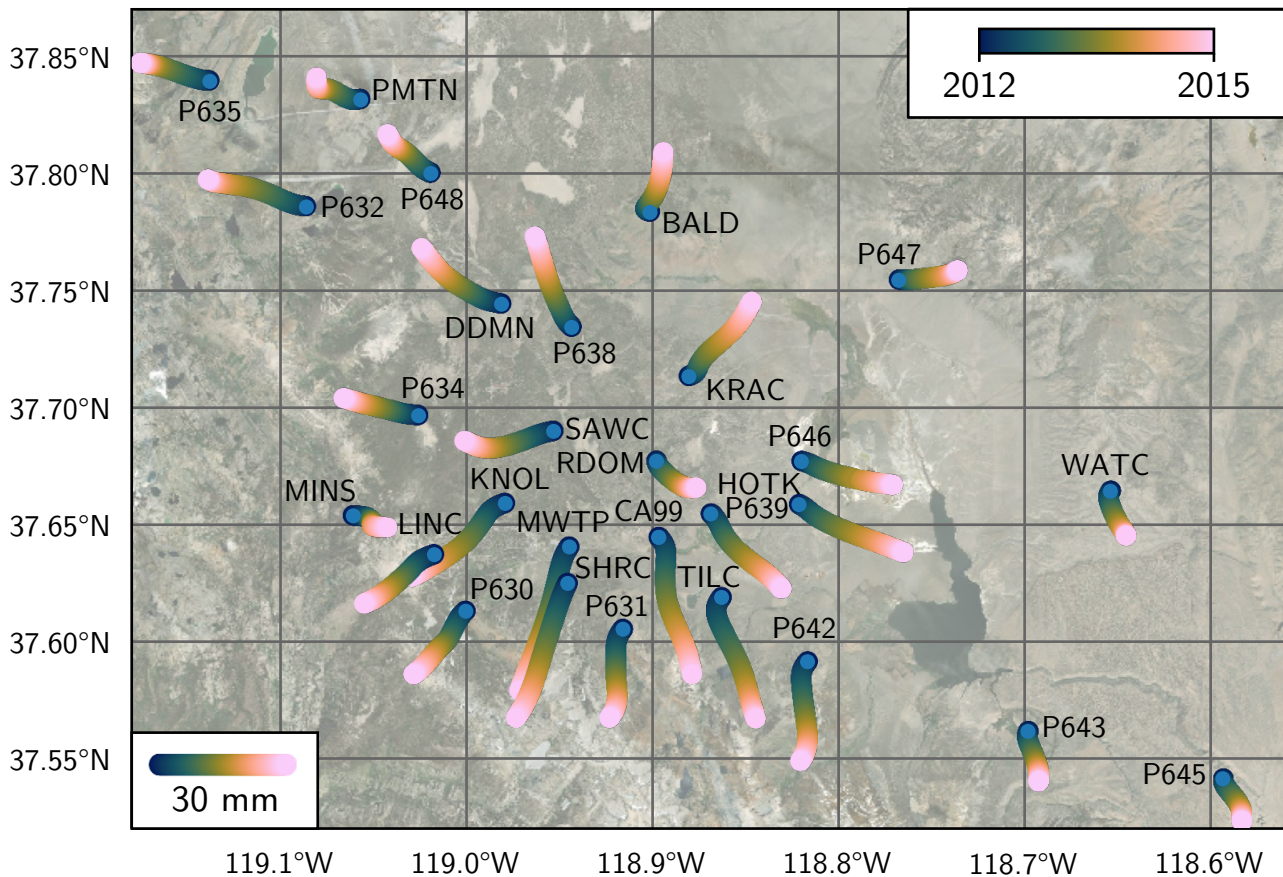
Figure 11: Modeled horizontal transient displacements of selected stations inside the Long Valley Caldera during the period between 2012 and 2015. The curves begin at the nominal location of each station, with the color of the markers corresponding to the time. Background satellite imagery by Earthstar Geographics & Esri.

`ModulatedSinusoid`), the solver is once again able to separate seasonal (i.e., periodic) signals from (aperiodic) transient motion. Different regularization penalties for the horizontal and vertical components furthermore allow to reflect the fact that horizontal observation precision is usually much higher than in the vertical direction. One example of the resulting seasonal fit in the vertical direction is shown in Fig. 13 for station P647. Variations in the amplitude, and sometimes instantaneous phase, are clearly visible, which showcases the importance of properly removing or estimating the full seasonal signal at stations that are affected by major hydrological processes. Fig. S6 shows the annual model's vertical amplitude and phase in map view, and Fig. S7 shows the overall vertical seasonal signal for the stations in Fig. 12 for the entire timespan.

5. Discussion

The choice to incorporate process-agnostic, spatial awareness into the timeseries decomposition problem by means of a parametric, spline-based model that requires regularization and iteration may possibly appear odd — after all, vbICA and comparable methods already have an inherent sense of space. However, even though basis decompositions have a spatial component, the geometry of the network is neglected (e.g., relative distance between stations). Network geometry and extent becomes relevant when networks are large, and some signals are spatially confined: different processes at different locations may be mapped into the same component, complicating its interpretation. Furthermore, in order to obtain a clean decomposition using vbICA or similar methods, maintenance and earthquake coseismic offsets still have to be removed ahead of time, as well as the linear secular trends. Therefore, not only do these non-parametric decomposition approaches require a significant amount of preprocessing in the first place, the separation of preprocessing and actual decomposition precludes a straightforward way to quantify the covariance between the constituents. Using parametric models that are both process-aware (such as secular, seasonal, and maintenance offset models) and process-agnostic (using a dictionary of splines for transients and seasonal variations), by contrast, offers this correlation by design, while the spatial L0 regularization accomplishes the goals for sparsity and spatial awareness. The significant trade-off between secular, linear velocity, and longterm transients has been discussed in Section 4.1. Lastly, we note that DISSTANS can still be used for purely data
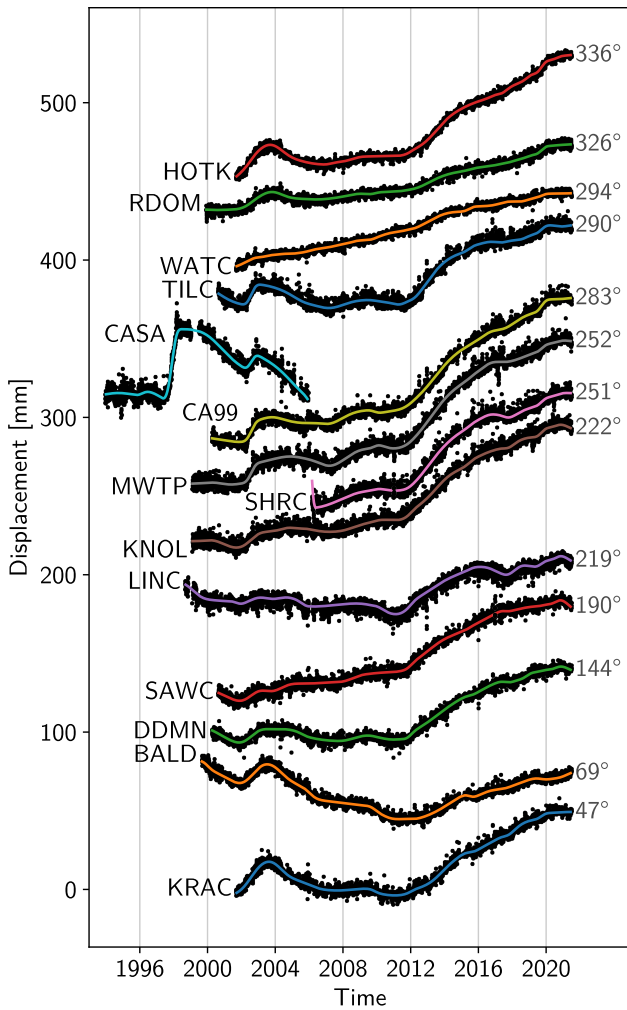
Figure 12: Modeled horizontal transient displacement (colored lines) of selected stations (names on the left) from Fig. 11, projected along the direction of maximum displacement during the period of 2012–2015. The directions (in grey to the right of the timeseries) are measured counterclockwise from east. CA99's direction is used for CASA. Black dots are the joint model's residuals, centered on the transient model.

Figure 13: Modeled seasonal vertical displacement timeseries for station P647 (see Fig. 11 for location). Blue and orange lines correspond to the annual and biannual signals, respectively, and the black line is their sum. The deviation component only includes the annual frequency.

management and preprocessing purposes, and may therefore still be of value to studies relying on basis decompositions — in fact, PCA/ICA is included in DISSTANS for the common mode removal.

Another benefit of using a parametric approach for this study is the potential to include prior knowledge beyond the preprocessing steps. Incorporating such knowledge is already partly possible through the very definition of the models (e.g., inserting a postseismic displacement model after a large earthquake), but least-squares-based methods such as the one used by DISSTANS also allow analytic inclusion of a priori model parameter knowledge, which may be added in future versions.

Hyperparameters such as the regularization penalties, number of iterations, the choice (shape and strength) of the reweighting functions, have so far only been introduced,
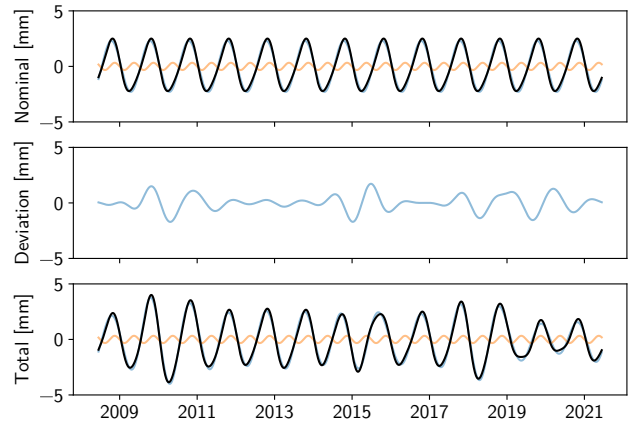
but not examined closely. We omit a detailed look here, as differing scientific goals, as well as different underlying data, have a large impact on what the "best" choice is, and general assertions are therefore not possible. The code presented here therefore does not relieve the user of the task of finding the best set of parameter for their data and problem formulation, although DISSTANS's online documentation does include the specific choices for the cases presented in the previous section (based on both analytic and empirical considerations) and may provide a good starting point.

An important caveat of using an overcomplete, sparse dictionary of splines for the modeling of transient signals is that such fits are not phase-invariant. Not accounting for phase invariance means that processes that move both in space and time (e.g., slow slip events in Cascadia) are "discretized" by the onset times of individual splines, and that multiple splines (of possibly different periods) are necessary to capture a potentially simple signal moving in time. Failing to account for different onset times throughout the network could negatively impact the quality of fit, as well as reduce the sparsity of the solution. However, experience shows that phase invariance is not as crucial as it may seem: First, observation noise makes exact onset times of transient signals hard to determine, and simultaneously allows the solver to fit splines that are adjacent in time when the "best" onset time would be somewhere in between the splines' onset times. Second, if the problem persists, more splines of different periods or new onset times can be easily inserted into the models (with the main drawback being higher computational costs). In neither the synthetic nor real validation datasets presented here did the splines' periods or onset times have to be adjusted from an initial, default configuration to obtain a high-quality decomposition solution.

While DISSTANS includes the estimation of the formal covariance matrix between parameters and components in a least-squares sense, it is not a fully probabilistic or Bayesian

method. However, since the necessary data to perform a Markov Chain Monte Carlo (MCMC) calculation, for example, is chiefly the forward model (i.e. the mapping matrix from model parameters to fitted timeseries, which is a core part of DISSTANS's computations), a future extension to include probabilistic methods is viable. In these cases, the standard spatial-L0-regularized solution could be used for the problem initialization.

Probabilistic methods may also be useful when extending the solution process to include nonlinear models — currently, only models linear in their coefficients are possible. Alternatively, nonlinear least squares (i.e iterative least squares linearized about a prior) may provide useful enough to incorporate into future versions of the code. However, the necessity for nonlinear models has so far been limited (especially beyond what an overcomplete dictionary of basis functions could accommodate).

Finally, DISSTANS's computation and parallelization are CPU-based because of its relative simplicity of implementation thanks to Python's core modules. GPU subroutines, as well as parallelization based on arrays of GPUs, will warrant a closer look in the future to assess possible performance upgrades. Indeed, in the core computational step when performing a local L0-regularized fit, CUDA-based algorithms may already provide a speed-up of about an order of magnitude (Schubiger et al., 2020, Fig. 2).

6. Conclusion

Displacement timeseries of regional GNSS networks have been used for over two decades now to monitor surface deformation, plate motion, as well as transient signals such as hydrological loading or aseismic slip events. A crucial step in these analyses is the decomposition of the input (raw) timeseries into its constituents: (assumed) linear plate motion, periodic seasonal variations, step offsets due to earthquakes, etc. As networks continue to grow in number and size, so does the need for software that aides researchers in efficiently analyzing timeseries. We aim to combine the accomplishments of previously published methods into a single, generic, open-source code. The DISSTANS Python package includes the following key features: (1) incorporation of spatial information through the use of a spatial L0-regularized least-squares solver, (2) CPU-based parallelization to provide scalability for large networks, (3) formal uncertainty quantification with covariance matrices between components and models, (4) a suite of supporting tools including timeseries files data management, common mode estimation, and simple, automated step detection, as well as (5) visualization methods to accelerate data and model inspection by the user.

Validation with synthetic GNSS network timeseries shows the beneficial effect of fitting transient signals with the spatial, L0-regularized solver: transients in the data are fit sparsely both in time and space, and are able to recover the true underlying motion better than comparable solutions without spatial awareness. An analysis of GNSS displacement timeseries from the Long Valley Caldera region in the Sierra Nevada, California, USA furthermore demonstrates the viability of our approach using real-world data, jointly estimating a number of models including step offset, secular long-term motion, transient signals, as well as time-varying seasonal displacements. Finally, some avenues of future improvements to the code, including GPU utilization and the adaptation for probabilistic methods, are presented.

Computer Code Availability

DISSTANS is available at `https://github.com/tobiscode/disstans` under the GPL-3.0 License.

Acknowledgements

CRediT authorship contribution statement

Tobias Köhne: Methodology, Software, Validation, Writing: Original Draft, Visualization. **Bryan Riel:** Methodology, Writing: Review & Editing. **Mark Simons:** Conceptualization, Resources, Writing: Review & Editing, Supervision, Funding acquisition.

References

Agrawal, A., Verschueren, R., Diamond, S., Boyd, S., 2018. A rewriting system for convex optimization problems. Journal of Control and Decision 5, 42–60. doi:`10.1080/23307706.2017.1397554`.

Bedford, J, Bevis, M., 2018. Greedy Automatic Signal Decomposition and Its Application to Daily GPS Time Series. Journal of Geophysical Research: Solid Earth 123, 6992–7003. doi:`10.1029/2017JB014765`.

Bekaert, D.P.S., Segall, P., Wright, T.J, Hooper, A.J, 2016. A Network Inversion Filter combining GNSS and InSAR for tectonic slip modeling. Journal of Geophysical Research: Solid Earth 121, 2069–2086. doi: `10.1002/2015JB012638`.

Bertiger, W., Bar-Sever, Y., Dorsey, A., Haines, B., Harvey, N., Hemberger, D., Heflin, M., Lu, W., Miller, M., Moore, A.W., Murphy, D., Ries, P., Romans, L., Sibois, A., Sibthorpe, A., Szilagyi, B., Vallisneri, M., Willis, P., 2020. GipsyX/RTGx, a new tool set for space geodetic operations and research. Advances in Space Research 66, 469–489. doi:`10.1016/j.asr.2020.04.015`.

Blewitt, G., 2015. GPS and Space-Based Geodetic Methods, in: Schubert, G. (Ed.), Treatise on Geophysics (Second Edition). Elsevier, Oxford, pp. 307–338. doi:`10.1016/B978-0-444-53802-4.00060-9`.

Blewitt, G., Hammond, W., Kreemer, C., 2018. Harnessing the GPS Data Explosion for Interdisciplinary Science. Eos 99. doi:`10.1029/2018EO104623`.

Blewitt, G., Kreemer, C., Hammond, W.C., Gazeaux, J, 2016. MIDAS robust trend estimator for accurate GPS station velocities without step detection. Journal of Geophysical Research: Solid Earth 121, 2054–2068. doi:`10.1002/2015JB012552`.

Candès, E.J, Wakin, M.B., Boyd, S.P., 2008. Enhancing Sparsity by Reweighted ℓ1 Minimization. Journal of Fourier Analysis and Applications 14, 877–905. doi:`10.1007/s00041-008-9045-x`.

Diamond, S., Boyd, S., 2016. CVXPY: A Python-Embedded Modeling Language for Convex Optimization. Journal of machine learning research: JMLR 17, 83.

Dong, D., Fang, P., Bock, Y., Webb, F., Prawirodirdjo, L., Kedar, S., Jamason, P., 2006. Spatiotemporal filtering using principal component analysis and Karhunen-Loeve expansion approaches for regional GPS network analysis. Journal of Geophysical Research: Solid Earth 111. doi:10.1029/2005JB003806.

Dong, D., Herring, T.A., King, R.W., 1998. Estimating regional deformation from a combination of space and terrestrial geodetic data. Journal of Geodesy 72, 200–214. doi:10.1007/s001900050161.

Gazeaux, J., Williams, S., King, M., Bos, M., Dach, R., Deo, M., Moore, A.W., Ostini, L., Petrie, E., Roggero, M., Teferle, F.N., Olivares, G., Webb, F.H., 2013. Detecting offsets in GPS time series: First results from the detection of offsets in GPS experiment. Journal of Geophysical Research: Solid Earth 118, 2397–2407. doi:10.1002/jgrb.50152.

Gualandi, A., Serpelloni, E., Belardinelli, M.E., 2016. Blind source separation problem in GPS time series. Journal of Geodesy 90, 323–341. doi:10.1007/s00190-015-0875-4.

Harris, C.R., Millman, K.J., van der Walt, S.J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N.J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M.H., Brett, M., Haldane, A., del Río, J.F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., Oliphant, T.E., 2020. Array programming with NumPy. Nature 585, 357–362. doi:10.1038/s41586-020-2649-2.

Heflin, M., Donnellan, A., Parker, J., Lyzenga, G., Moore, A., Ludwig, L.G., Rundle, J., Wang, J., Pierce, M., 2020. Automated Estimation and Tools to Extract Positions, Velocities, Breaks, and Seasonal Terms From Daily GNSS Measurements: Illuminating Nonlinear Salton Trough Deformation. Earth and Space Science 7, e2019EA000644. doi:10.1029/2019EA000644.

Herring, T.A., King, R.W., Floyd, M.A., McClusky, S.C., 2018. Introduction to GAMIT/GLOBK.

Hetland, E.A., Musé, P., Simons, M., Lin, Y.N., Agram, P.S., DiCaprio, C.J., 2012. Multiscale InSAR Time Series (MInTS) analysis of surface deformation. Journal of Geophysical Research: Solid Earth 117. doi:10.1029/2011JB008731.

Huang, D.W., Dai, W.J., Luo, F.X., 2012. ICA Spatiotemporal Filtering Method and Its Application in GPS Deformation Monitoring. Applied Mechanics and Materials 204–208, 2806–2812. doi:10.4028/www.scientific.net/AMM.204-208.2806.

Hunter, J.D., 2007. Matplotlib: A 2D Graphics Environment. Computing in Science Engineering 9, 90–95. doi:10.1109/MCSE.2007.55.

Ji, K.H., Herring, T.A., Llenos, A.L., 2013. Near real-time monitoring of volcanic surface deformation from GPS measurements at Long Valley Caldera, California. Geophysical Research Letters 40, 1054–1058. doi:10.1002/grl.50258.

McGuire, J.J., Segall, P., 2003. Imaging of aseismic fault slip transients recorded by dense geodetic networks. Geophysical Journal International 155, 778–788. doi:10.1111/j.1365-246X.2003.02022.x.

McKinney, W., 2010. Data Structures for Statistical Computing in Python. Proceedings of the 9th Python in Science Conference , 56–61doi:10.25080/Majora-92bf1922-00a.

Met Office, 2021. Cartopy: A cartographic python library with a Matplotlib interface.

Montgomery-Brown, E.K., Wicks, C.W., Cervelli, P.F., Langbein, J.O., Svarc, J.L., Shelly, D.R., Hill, D.P., Lisowski, M., 2015. Renewed inflation of Long Valley Caldera, California (2011 to 2014). Geophysical Research Letters 42, 5250–5257. doi:10.1002/2015GL064338.

Riel, B., Simons, M., Agram, P., Zhan, Z., 2014. Detecting transient signals in geodetic time series using sparse estimation techniques. Journal of Geophysical Research: Solid Earth 119, 5140–5160. doi:10.1002/2014JB011077.

Santamaría-Gómez, A., 2019. SARI: Interactive GNSS position time series analysis software. GPS Solutions 23, 52. doi:10.1007/s10291-019-0846-y.

Santamaría-Gómez, A., Ray, J., 2021. Chameleonic Noise in GPS Position Time Series. Journal of Geophysical Research: Solid Earth 126, e2020JB019541. doi:10.1029/2020JB019541.

Schubiger, M., Banjac, G., Lygeros, J., 2020. GPU acceleration of ADMM for large-scale quadratic programming. Journal of Parallel and Distributed Computing 144, 55–67. doi:10.1016/j.jpdc.2020.05.021.

Segall, P., Matthews, M., 1997. Time dependent inversion of geodetic data. Journal of Geophysical Research: Solid Earth 102, 22391–22409. doi:10.1029/97JB01795.

Silverii, F., Montgomery-Brown, E.K., Borsa, A.A., Barbour, A.J., 2020. Hydrologically Induced Deformation in Long Valley Caldera and Adjacent Sierra Nevada. Journal of Geophysical Research: Solid Earth 125, e2020JB019495. doi:10.1029/2020JB019495.

The pandas development team, 2021. Pandas-dev/pandas: Pandas 1.3.0. Zenodo. doi:10.5281/zenodo.5060318.

Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S.J., Brett, M., Wilson, J., Millman, K.J., Mayorov, N., Nelson, A.R.J., Jones, E., Kern, R., Larson, E., Carey, C.J., Polat, İ., Feng, Y., Moore, E.W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E.A., Harris, C.R., Archibald, A.M., Ribeiro, A.H., Pedregosa, F., van Mulbregt, P., 2020. SciPy 1.0: Fundamental algorithms for scientific computing in Python. Nature Methods 17, 261–272. doi:10.1038/s41592-019-0686-2.

Tobias Köhne received his B.S. (2016) in Mechanical Engineering from the Technical University of Munich, Germany and his M.S. (2018) in Aerospace Engineering from the University of Texas at Austin, USA focusing on orbital mechanics. He is currently a Ph.D. candidate at the California Institute of Technology, where he is working on modern tools for geodetic timeseries analysis of large-scale, continuously-operating GNSS regional networks. At Caltech, he was also part of an investigation into how machine learning approaches can aid to extract more information out of the currently available InSAR timeseries datasets, and worked on a small study on the dynamics and origins of retrograde Jupiter Trojans.

Bryan Riel received a B.Sc. degree in 2008 and M.Sc. degree in 2010 in Aerospace Engineering from the University of Texas at Austin and a Ph.D. in Geophysics from the California Institute of Technology in 2016. His graduate research was focused on geodetic time series analysis, geophysical modeling, and inverse theory, as well as radiometric calibration of synthetic aperture radar (SAR) data and quantification of elevation errors in backscatter products. In 2016, he joined the NASA Jet Propulsion Laboratory (JPL) as a signal analysis engineer in the Radar Science and Engineering Section. In 2019, he joined the Massachusetts Institute of Technology (MIT) as a research scientist, studying glacier dynamics using geodetic data and machine learning. Dr. Riel's research interests are broad and include geodetic time series analysis and decomposition, glaciology, volcanology, machine learning, physical modeling, inverse theory, and uncertainty quantification. Dr. Riel received the National Aeronautics and Space Administration Earth and Space Sciences Fellowship in 2012.

Mark Simons received the B.Sc. degree in geophysics and space physics from the University of California at Los Angeles, Los Angeles, CA, USA, in 1989, and the Ph.D. degree in geophysics from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 1996. He has been with the California Institute of Technology, Pasadena, CA, USA, since 1996, where he is currently the John W. and Herberta M. Miles Professor of Geophysics at the Seismological Laboratory, Division of Geological and Planetary Science. He is also the Chief Scientist of the Jet Propulsion Laboratory. His research interests include processes associated with the seismic cycle, migration of magma and water in the subsurface, tides, and glacial rebound; tectonics and the relationship between short and long time-scale processes; glaciology, particularly basal mechanics and ice rheology; tools and applications using space geodesy, particularly GNSS and SAR; Bayesian methods for large geophysical inverse problems; and application of space geodesy for monitoring and rapid response to natural disasters. Prof. Simons is a fellow of the American Geophysical Union.

Decomposition and Inference of Sources through Spatiotemporal Analysis of Network Signals: The DISSTANS Python Package

Supplementary Material

Tobias Köhne   Bryan Riel   Mark Simons

January 28, 2022

S1 Code Overview: Additional Details

S1.1 Data Formats

All timeseries datasets are stored as objects of `Timeseries` subclasses. The `Timeseries` parent class defines an internal data structure that all further processing done by DISSTANS methods of all levels rely on. It also implements properties such as the calculation of a timeseries length or reliability, the possibilty to use Python's in-built mathematical operators to create new timeseries, and convenience functions such as cutting the timeseries or building covariance matrices at a particular timestep.

Subclasses, in turn, define how any particular input file gets loaded to match the common structure. The two provided subclasses are `GipsyTimeseries` (for $\mathbf{P}$L's GipsyX `.tseries` files) and `UNRTimeseries` (for UNR's `.tenv3` files). User-defined classes can easily be created by adhering to the format of the two existing subclasses, and checking the documentation of `Timeseries`.

S1.2 Visualization

Because the raw data contained within `Timeseries` objects are standard pandas DataFrames, they can be plotted using standard Matplotlib code using their `Timeseries.time` and `Timeseries.data` attributes. Utilizing commonly-used Python object formats enables uncomplicated inspections of a particular station, timeseries, or fit; and allow for non-standard user-desired plotting (e.g., for publication-quality figures). Model parameter values and covariances (accessed through their `Model.parameters` and `Model.covariances` NumPy array attributes) are also directly plottable with Matplotlib.

There are, however, high-level visualization routines already included in DISSTANS. The core functionality is contained within the `Network.gui()` method. In its simplest call signature, it provides a clickable map of the network (optionally with satellite imagery background), and a separate figure with all the timeseries contained by a station. If a timeseries contains fitted models, the overall model prediction is plotted, and optionally, can be split up into the different model components, and if there are `SplineSet` models present at a station, a scalogram can be shown. All figures can also be saved directly to files.

Furthermore, to visualize station motion in a map view, the `Network.wormplot()` method can produce still maps and animated video files of the stations' displacement (or individual model components of them). Lastly, the `Network.graphical_cme()` method performs common mode estimation (see Section 3.5) and presents the temporal and spatial components separately for validation purposes.

S1.3 Mathematical Model Formulation

In DISSTANS, the joint mathematical formulation $g(t)$ is the sum of all `num_models` individual models contained in a `ModelCollection`. Each individual model g_{Model} (described by `Model` objects) can again be a linear superposition of spanning functions g_j and corresponding coefficients m_j:

$$g(t) = \sum g_{\text{Model}}(t) = \sum_{j=1}^{\text{num\_parameters}} m_j g_j(t)$$

Here, `num_parameters` is the total number of all individual spanning functions, and therefore also the number of all coefficients to be estimated. The basic models included in DISSTANS are:

$$g_{\text{Arctangent}}(t) = m' \left(\frac{1}{\pi} \arctan \left(\frac{t}{\tau} \right) + 0.5 \right)$$

$$g_{\text{Exponential}}(t) = m' \left(1 - \exp \left(-\frac{t}{\tau} \right) \right)$$

$$g_{\text{HyperbolicTangent}}(t) = m' \left(\frac{1}{2} \tanh \left(\frac{t}{\tau} \right) + 0.5 \right)$$

$$g_{\text{Logarithmic}}(t) = m' \log \left(1 + \frac{t}{\tau} \right)$$

$$g_{\text{Polynomial}}(t) = \sum_l m'_l t^l$$

$$g_{\text{Step}}(t) = \sum_l m'_l H \left(t - t_l^{\text{step}} \right)$$

$$g_{\text{Sinusoid}}(t) = m'_0 \cos \left(\omega t \right) + m'_1 \sin \left(\omega t \right)$$

where all m' are just stand-ins for the overall set of m_j, the τ can of course vary between models, the t_l^{step} are step times, and $H(t)$ is the Heaviside function.

While the models are continuous in time t, timeseries analysis inherently works on discrete observations d_i at times t_i. Using matrix notation, the least squares problem can be formulated as follows:

$$\mathbf{d} = \mathbf{Gm} + \epsilon$$

Where

$$\mathbf{d} = \left(d_i\right) \in \mathbb{R}^{\texttt{num\_observations} \times 1}$$

$$\mathbf{G} = \left(G_{i,j}\right) \in \mathbb{R}^{\texttt{num\_observations} \times \texttt{num\_parameters}} = \left(g_j\left(t_i\right)\right) \in \mathbb{R}^{\texttt{num\_observations} \times \texttt{num\_parameters}}$$

$$\mathbf{m} = \left(m_j\right) \in \mathbb{R}^{\texttt{num\_parameters} \times 1}$$

and $\epsilon \in \mathbb{R}^{\texttt{num\_observations} \times 1}$ is the column vector of residuals. All solvers start from this formulation to find the best set of m_j that minimizes a given cost function dependent on ϵ (potentially including regularization criteria, see Section S1.4). The choice of the data misfit loss function implicitly defines the assumed distribution from which ϵ is drawn (e.g., a Normal distribution in the case of unregularized least squares). In DISSTANS, the mapping (design) matrices $\mathbf{G}$ are computed by the get_mapping() methods, $\mathbf{d}$ is represented by Timeseries objects, and $\mathbf{m}$ is returned by the solver in Solution objects and added to each Model object.

In the following two subsections, we go more into the details of the spline-based models.

Linearly dependent, overcomplete dictionary of splines: SplineSet

We start with the formulation of a single cardinal B-spline basis function (*spline function*) of reference time t_{ref}. "Normalized" timestamps t' can be calculated as follows:

$$t' = \frac{t - t_{\text{ref}}}{\rho}$$

By default, this single spline function is then shifted to multiple center times by using its timescale ρ, leading to different normalized timevectors for each spline function:

$$t'_j = \frac{\left(t - t_{\text{ref}}\right) - j \cdot \rho}{\rho}$$

(Here, $j = 0 \ldots$ num_splines only considers the spline functions.)

To create the spline functions of a certain degree p (with order $n = p + 1$), we can then use the following relation (Butzer et al., 1988; Schoenberg, 1973):

$$g_j(t'_j) = \sum_{k=0}^{n} \frac{(-1)^k}{p!} \cdot \binom{n}{k} \cdot \left(t'_j + \frac{n}{2} - k\right)^p$$

This is the model represented by BSpline. Based on Riel et al. (2014), this study uses the integrated form of this spline function to represent transients. Its mathematical representation is:

$$g_j(t'_j) = \sum_{k=0}^{n} \frac{(-1)^k}{(p+1)!} \cdot \binom{n}{k} \cdot \left(t'_j + \frac{n}{2} - k\right)^{p+1}$$

The final spline model (a single BSpline or ISpline object) over all the available center times is therefore

$$g_{\{\texttt{B,I}\}\texttt{Spline}}(t) = \sum_{j=0}^{\texttt{num\_splines}} m_j g_j(t)$$

This model can then be repeated again for different timescales ρ, which is the purpose of the SplineSet class. Fig. S1 shows example spline functions.

Linearly independent spline basis for time-varying sinusoids: AmpPhModulatedSinusoid

The simple Sinusoid class models a seasonal signal, given a certain frequency ω, as the linear combination of a sine and cosine combination, allowing to estimate both phase ϕ and amplitude A as a linear problem:

$$g_{\texttt{Sinusoid}}(t) = A \cos\left(\omega t - \phi\right) = a \cos\left(\omega t\right) + b \sin\left(\omega t\right)$$
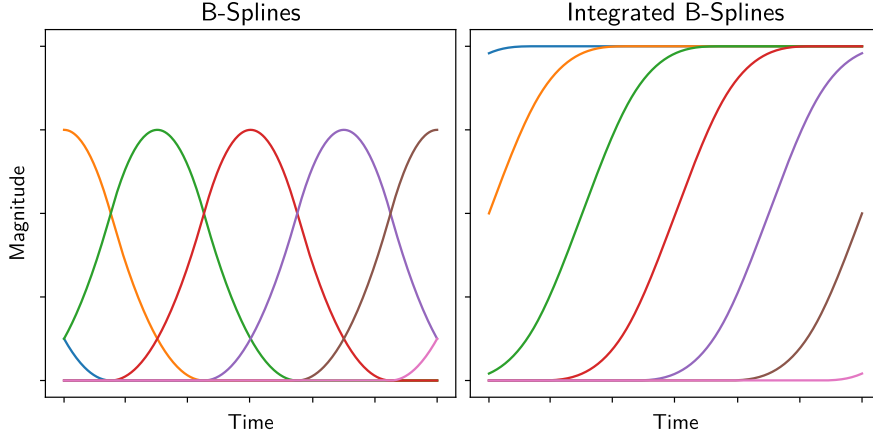
3

Figure S1: Example of `BSpline` (left) and `ISpline` (right) spline functions for a single timescale and order at seven different center times. Each curve on the right is the integral of the curve on the left of the same color.

If we want to allow the overall amplitude A to change over time, we can extend the definition of a (and similarly, b):

$$a \rightarrow a(t) = \bar{a} + \Delta a(t)$$

To keep the problem linear, we can use a spline representation for $\Delta a(t), \Delta b(t)$:

$$a(t) = \bar{a} + \sum_j a_j h_j(t)$$

Where the a_j (and respectively, b_j) are the parameters m_j to estimate, and h_j are the spline basis functions (more on h_j below). Expanding $g_{\text{Sinusoid}}(t)$ with the extended definition leads to a natural separation of components:

$$\begin{aligned}
a(t)\cos(\omega t) + b(t)\sin(\omega t) &= (\bar{a} + \Delta a(t))\cos(\omega t) + \left(\bar{b} + \Delta b(t)\right)\sin(\omega t) \\
&= \left(\bar{a}\cos(\omega t) + \bar{b}\sin(\omega t)\right) + (\Delta a(t)\cos(\omega t) + \Delta b(t)\sin(\omega t)) \\
&= \left(\bar{a}\cos(\omega t) + \bar{b}\sin(\omega t)\right) + \sum_j \left(a_j h_j(t)\cos(\omega t) + b_j h_j(t)\sin(\omega t)\right) \\
&= g_{\text{Sinusoid}}(t) + g_{\text{AmpPhModulatedSinusoid}}(t)
\end{aligned}$$

Here, the first term represents the *nominal* component, and the second term the *deviation* component. In DISSTANS, the terms correspond to the `Sinusoid` and `AmpPhModulatedSinusoid`, respectively.

Note that the h_j are not the same as for the dictionary of splines defined above. The dictionary is comprised of a single (cardinal) spline, that is of a defined length scale (i.e., period), and centered at specified timestamps. Here, for `AmpPhModulatedSinusoid`, we do not need the spline to be the same one shifted and scaled, instead we can default to the more general notion of B-Splines: the one of a complete basis for polynomials of a given degree on a given interval. Furthermore, this relaxation allows us to use SciPy's basis function implementation directly (Virtanen et al., 2020). Fig. S2 shows an example set of spline basis functions h_j, as well as the resulting modulated cosine and sine terms used as the spanning functions for `AmpPhModulatedSinusoid`.

Although it is not strictly necessary to include $\bar{a}$ and $\bar{b}$ explicitly in $a(t)$ and $b(t)$ (splines can also represent any constant function), the separation is useful because it allows a regularized solver to not penalize the nominal component.

S1.4 Solver Equations

For the `linear_regression()` solver, the cost function to be minimized is (based on the notation introduced in Section S1.3):

$$\|\mathbf{Gm} - \mathbf{d}\|_2^2$$

and the posterior covariance matrix $\mathbf{C}_m$ given the data covariance matrix $\mathbf{C}_d$ is

$$\mathbf{C}_m = \left(\mathbf{G}^T \mathbf{C}_d^{-1} \mathbf{G}\right)^g$$
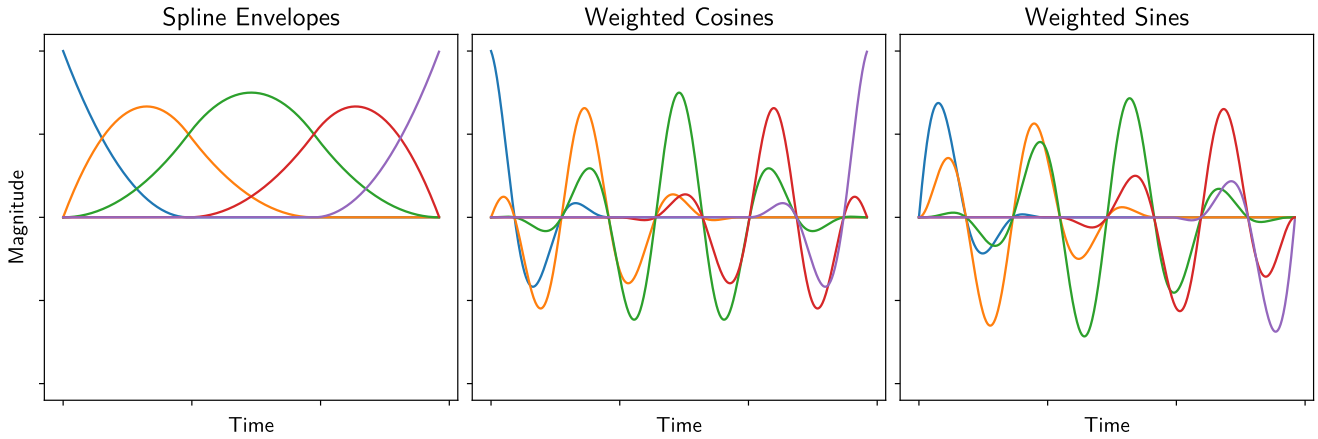
4

Figure S2: Visualization of the intermediate functions used by `AmpPhModulatedSinusoid`. The spline basis functions h_j over the considered time interval are in the left panel. Multiplying the spline functions with the cosine and sine of a given period then yields the modulated cosines and sines in the center and right panel, respectively.

where ^{g} is the generalized pseudo-inverse.

The `ridge_regression()` solver introduces L2-norm regularization of the model coefficients, minimizing

$$\|\mathbf{Gm} - \mathbf{d}\|_2^2 + \lambda \|\mathbf{m}_{\mathrm{reg}}\|_2^2$$

where λ is a chosen regularization penalty hyperparameter, and $\mathbf{m}_{\mathrm{reg}}$ is the subset of $\mathbf{m}$ that should be regularized. The posterior covariance matrix takes the regularization into account:

$$\mathbf{C}_m = \left(\mathbf{G}^T \mathbf{C}_d^{-1} \mathbf{G} + \lambda \mathbf{I}_{\mathrm{reg}}\right)^g$$

The `lasso_regression()` solver instead uses L1-norm regularization of the model coefficients, minimizing

$$\|\mathbf{Gm} - \mathbf{d}\|_2^2 + \lambda \|\mathbf{m}_{\mathrm{reg}}\|_1$$

By iterating on the L1-regularized solution, `lasso_regression()` can also approximate the solution for the L0-regularized [1] least-squares problem, minimizing

$$\|\mathbf{Gm} - \mathbf{d}\|_2^2 + \|\mathbf{m}_{\mathrm{reg}}\|_0$$

Because the result of an L0-regularized solution is approximately the same as if an unregularized problem was solved with only a subset of model parameters to be estimated, the posterior covariance matrix for `lasso_regression()` is the same as for `linear_regression()`, but setting to zero the covariances which were not estimated. Lastly, when `Network.spatialfit()` is used, DISSTANS can also minimize the jointly L1- and L0-regularized problem:

$$\|\mathbf{Gm} - \mathbf{d}\|_2^2 + \lambda \|\mathbf{m}_{\mathrm{reg,L1}}\|_1 + \|\mathbf{m}_{\mathrm{reg,L0}}\|_0$$

---

[1]Here, we use the definition of Candès et al. (2008): $\|\mathbf{x}\|_0 = |\{i \,:\, x_i \neq 0\}|$, i.e., the number of non-zero elements.
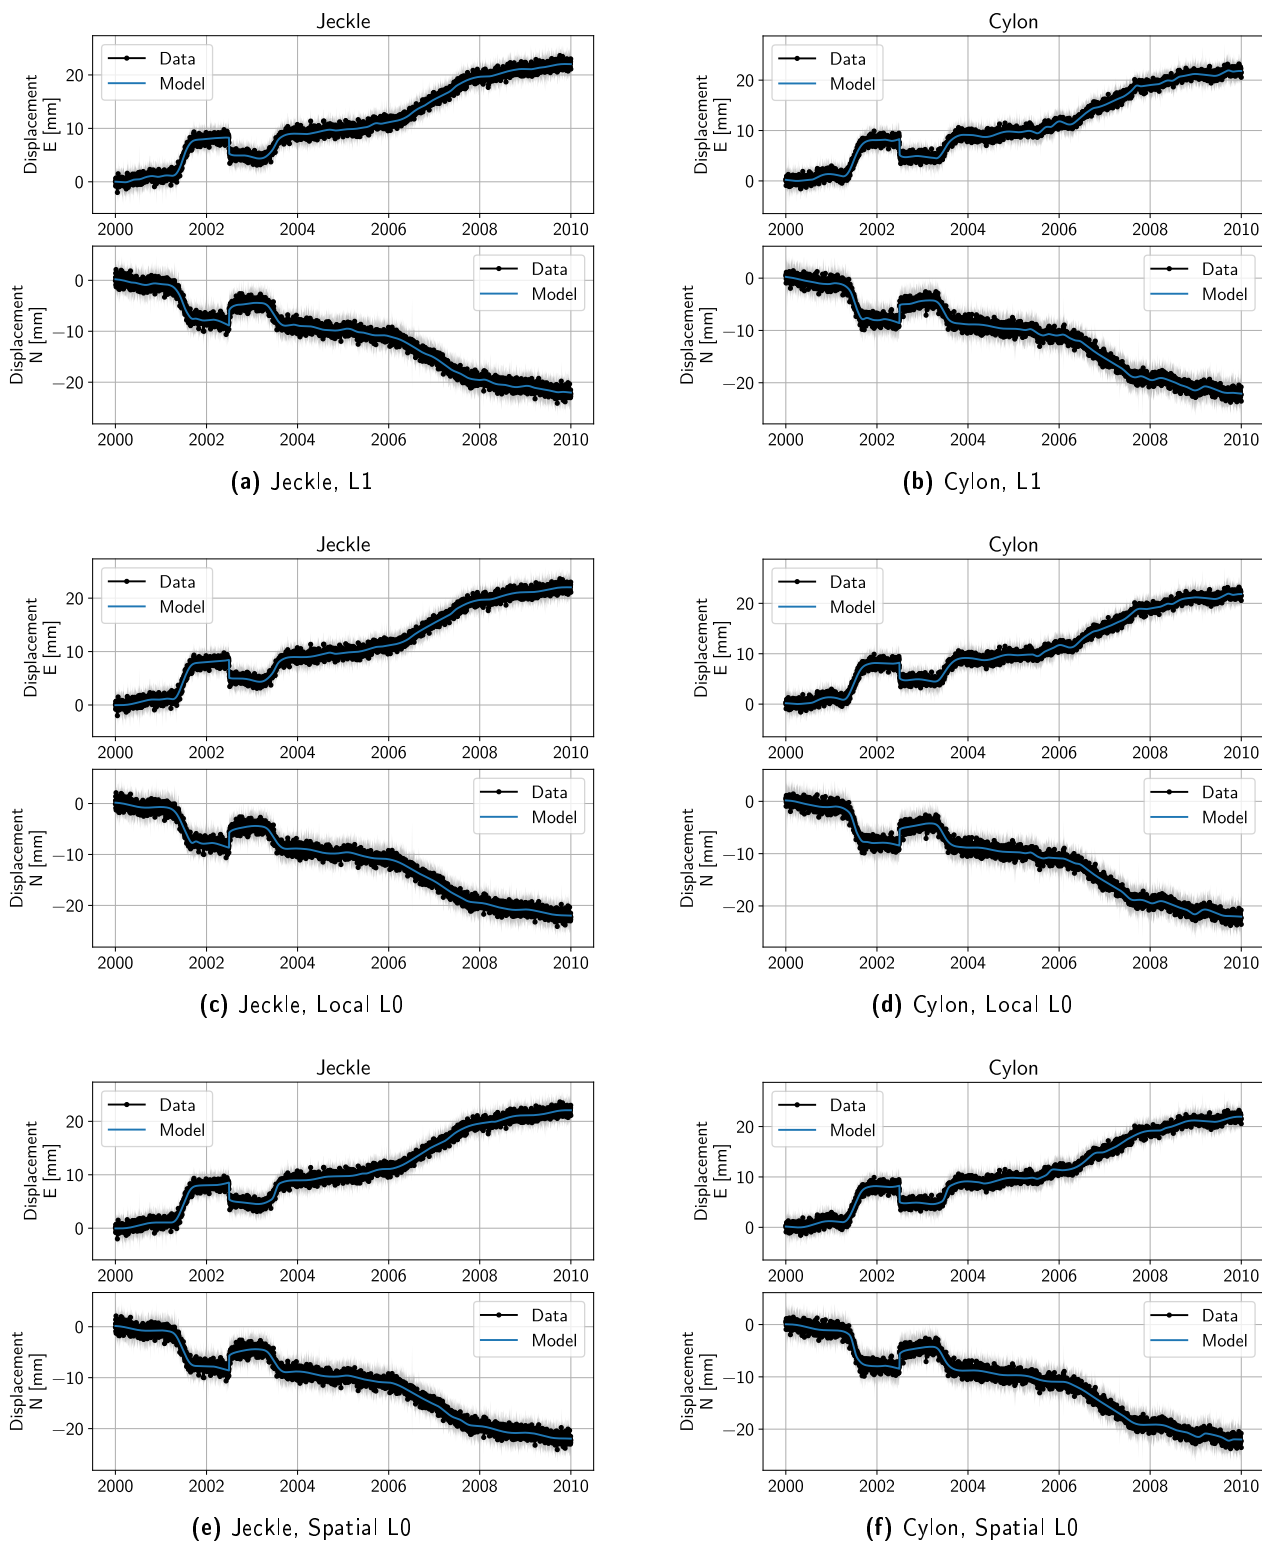
S2 Synthetic data timeseries fits & scalograms

Figure S3: Timeseries comparison for two stations and three regularization schemes in the two east and north components. Black dots are the synthetic observations, with the grey shading corresponding to three standard deviations of simulated observation uncertainty. The blue line represents the final overall model fit.
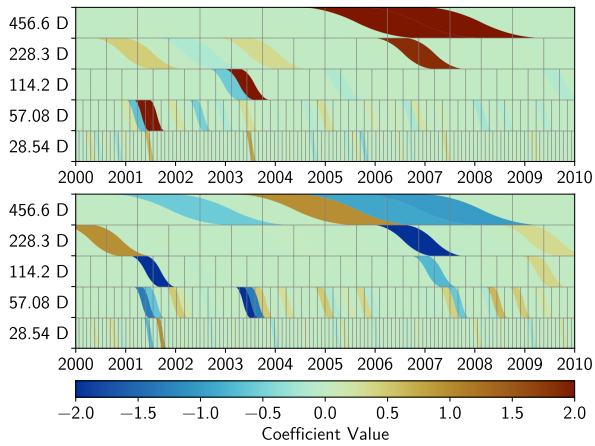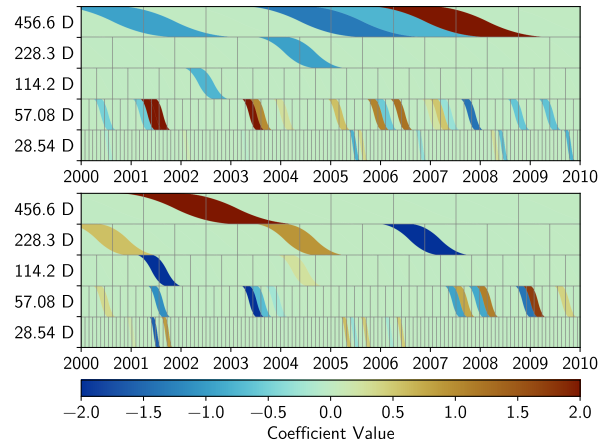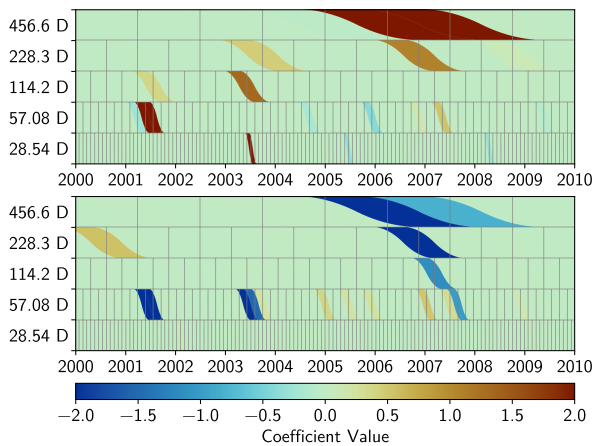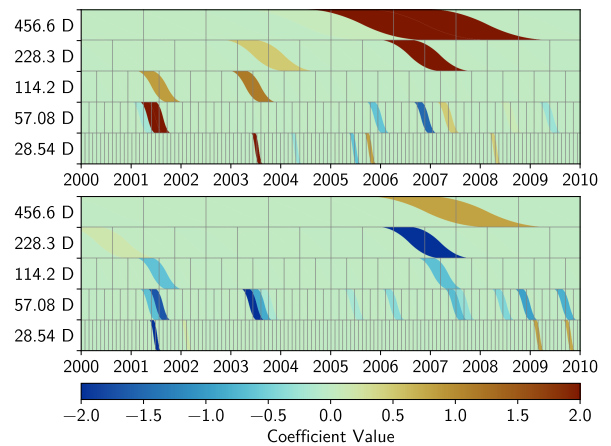
(a) Jeckle, L1
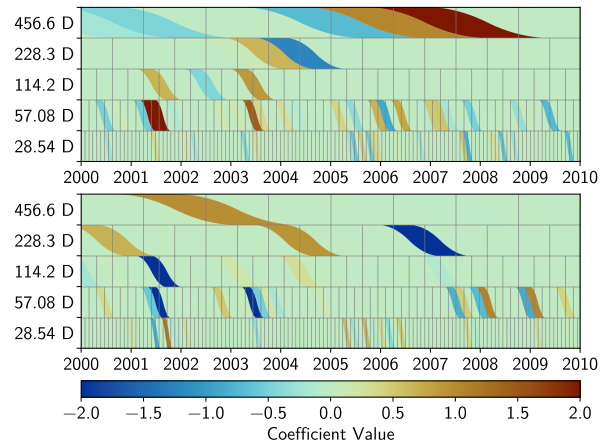
(b) Cylon, L1
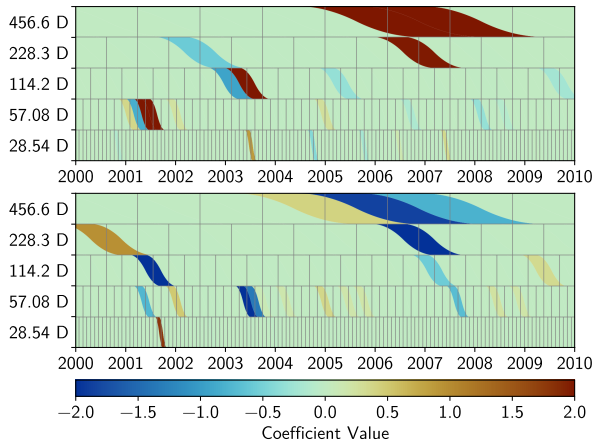
(c) Jeckle, Local L0

(d) Cylon, Local L0

(e) Jeckle, Spatial L0

(f) Cylon, Spatial L0

Figure S4: Scalograms of the transient model for the stations and regularization in Fig. S3. The horizontal and vertical axes correspond to time and the discrete periods of the splines, respectively. Patches (colored by the spline coefficient's value) in this time-period-space represent a single spline in the dictionary, with their extent in time defined as the active period of the spline (i.e., having non-zero gradient), and their height defined by the relative magnitude of the particular spline compared to all splines active at that time.

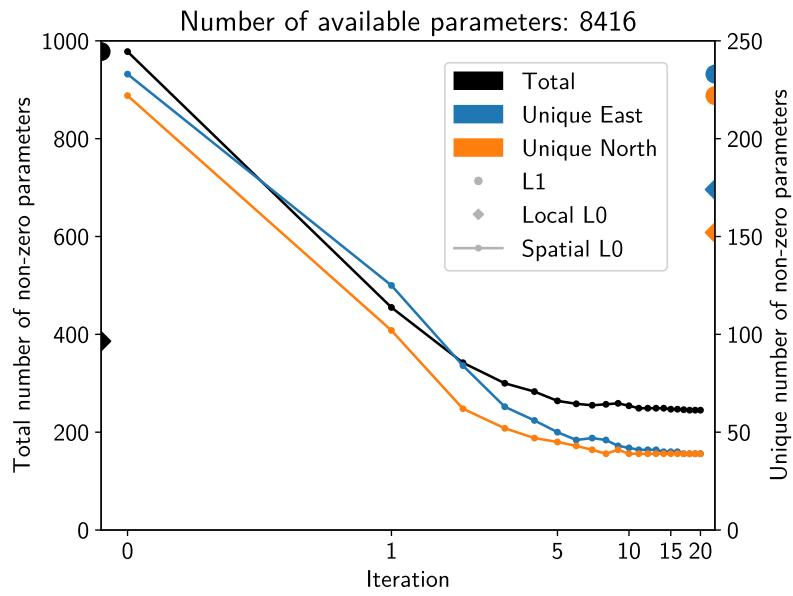# S3 Synthetic network solver convergence

Figure S5: Convergence of the iterative, spatiotemporal L0-regularized solver (line). For reference, values for the local L1-regularized (half circles) and local L0-regularized (triangles) solutions are also plotted on the axes. The results are shown both for the individual components (blue and orange), as well as the overall solution (black).

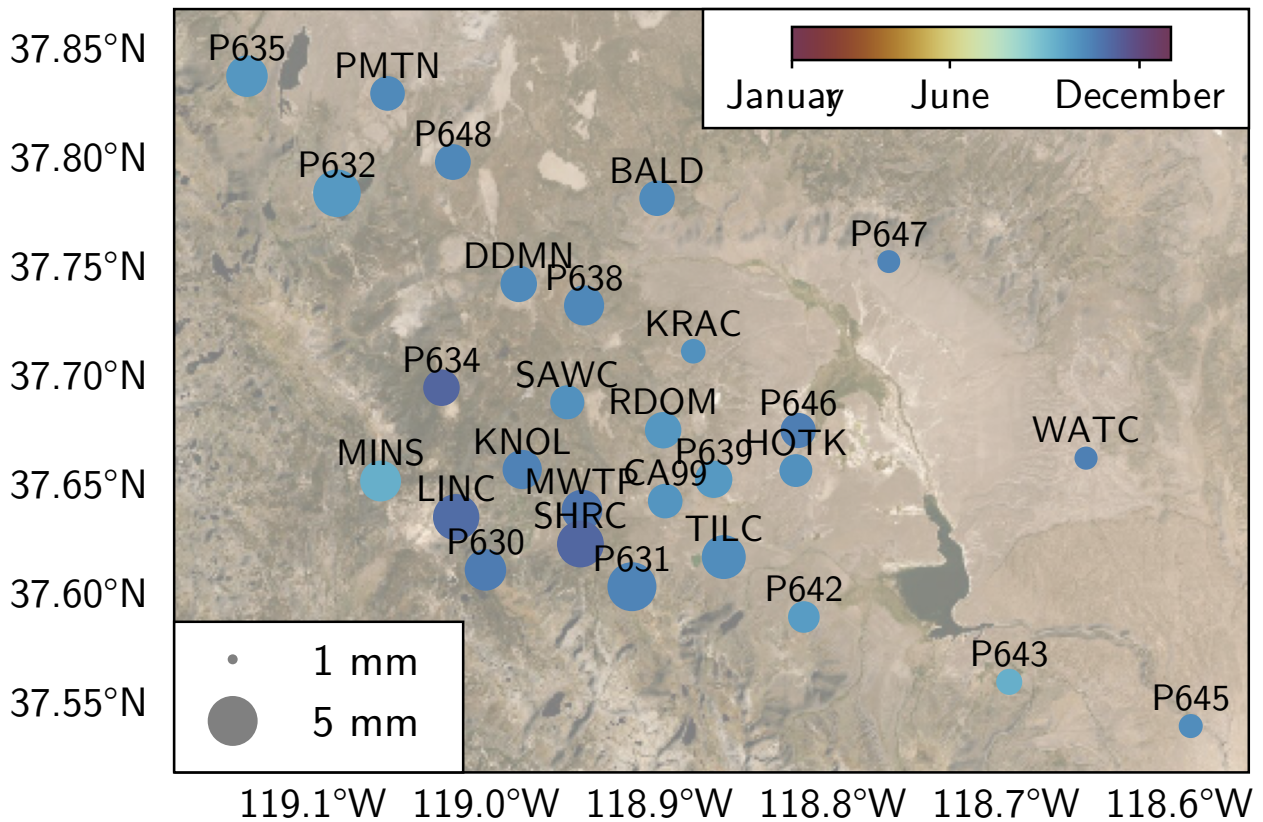## S4 Long Valley Caldera seasonal signals

Figure S6: Unregularized (average) component of the vertical seasonal fitted annual model. The marker size corresponds to the amplitude, and the color corresponds to the time of the sinusoid's maximum during the year. Background satellite imagery by Earthstar Geographics & Esri.
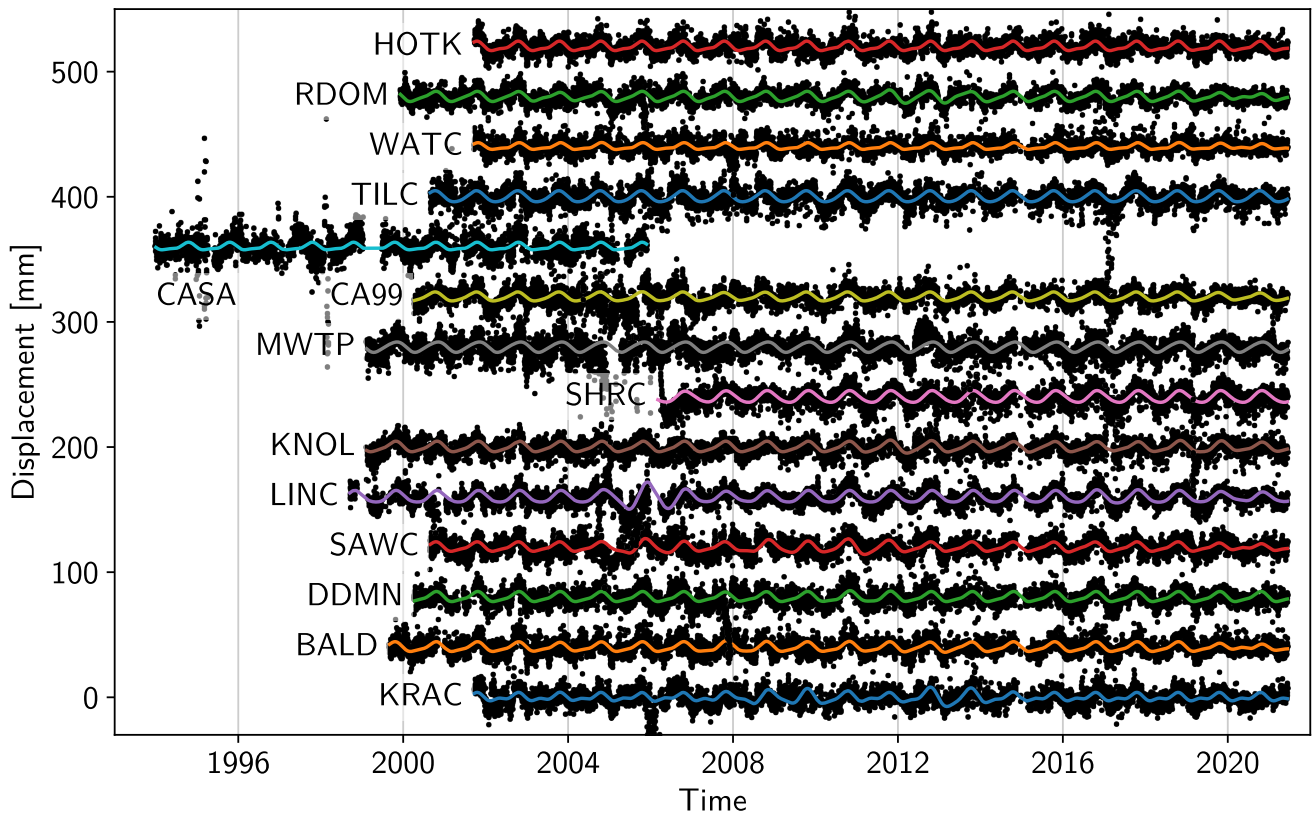
Figure S7: Full seasonal model (unregularized and regularized, annual and semi-annual) for the selected stations from Fig. 12, in the same order. Black dots are the overall model's residuals, centered on the seasonal model.

References

Butzer, P.L., Schmidt, M., Stark, E.L., 1988. Observations on the History of Central B-Splines. Archive for History of Exact Sciences 39, 137–156.

Candès, E.J., Wakin, M.B., Boyd, S.P., 2008. Enhancing Sparsity by Reweighted $\ell 1$ Minimization. Journal of Fourier Analysis and Applications 14, 877–905. doi:10.1007/s00041-008-9045-x.

Riel, B., Simons, M., Agram, P., Zhan, Z., 2014. Detecting transient signals in geodetic time series using sparse estimation techniques. Journal of Geophysical Research: Solid Earth 119, 5140–5160. doi: 10.1002/2014JB011077.

Schoenberg, I.J, 1973. The Basis Property of B-Splines, in: Cardinal Spline Interpolation. Society for Industrial and Applied Mathematics. CBMS-NSF Regional Conference Series in Applied Mathematics, pp. 11–19. doi:10.1137/1.9781611970555.ch2.

Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S.J, Brett, M., Wilson, J., Millman, K.J, Mayorov, N., Nelson, A.R.J, Jones, E., Kern, R., Larson, E., Carey, C.J, Polat, İ., Feng, Y., Moore, E.W., VanderPlas, J, Laxalde, D., Perktold, J, Cimrman, R., Henriksen, I., Quintero, E.A., Harris, C.R., Archibald, A.M., Ribeiro, A.H., Pedregosa, F., van Mulbregt, P., 2020. SciPy 1.0: Fundamental algorithms for scientific computing in Python. Nature Methods 17, 261–272. doi:10.1038/s41592-019-0686-2.