

PAPER • OPEN ACCESS

## Koopman-inspired approach for identification of exogenous anomalies in nonstationary time-series data

To cite this article: Alex Mallen *et al* 2023 *Mach. Learn.: Sci. Technol.* **4** 025033

View the [article online](#) for updates and enhancements.

### You may also like

- [Data-driven discovery of Koopman eigenfunctions for control](#)  
Eurika Kaiser, J Nathan Kutz and Steven L Brunton
- [Automatic differentiation to simultaneously identify nonlinear dynamics and extract noise probability distributions from data](#)  
Kadierdan Kaheman, Steven L Brunton and J Nathan Kutz
- [Extending the extended dynamic mode decomposition with latent observables: the latent EDMD framework](#)  
Said Ouala, Bertrand Chapron, Fabrice Collard et al.



## PAPER

## OPEN ACCESS

RECEIVED  
27 November 2022REVISED  
16 May 2023ACCEPTED FOR PUBLICATION  
9 June 2023PUBLISHED  
20 June 2023

Original Content from  
this work may be used  
under the terms of the  
[Creative Commons  
Attribution 4.0 licence](#).

Any further distribution  
of this work must  
maintain attribution to  
the author(s) and the title  
of the work, journal  
citation and DOI.



# Koopman-inspired approach for identification of exogenous anomalies in nonstationary time-series data

Alex Mallen<sup>1,4,\*</sup> , Christoph A Keller<sup>2,3</sup> and J Nathan Kutz<sup>4</sup> <sup>1</sup> Department of Computer Science, University of Washington, Seattle, WA 98195, United States of America<sup>2</sup> NASA Global Modeling and Assimilation Office, Goddard Space Flight Center, Greenbelt, MD 20771, United States of America<sup>3</sup> Morgan State University, Baltimore, MD 21251, United States of America<sup>4</sup> Department of Applied Mathematics and Electrical and Computer Engineering, University of Washington, Seattle, WA 98195, United States of America

\* Author to whom any correspondence should be addressed.

E-mail: [atmallen@cs.washington.edu](mailto:atmallen@cs.washington.edu)**Keywords:** anomaly detection, machine learning, time series, koopman theory

## Abstract

In many scenarios, it is necessary to monitor a complex system via a time-series of observations and determine when anomalous exogenous events have occurred so that relevant actions can be taken. Determining whether current observations are abnormal is challenging. It requires learning an extrapolative probabilistic model of the dynamics from historical data, and using a limited number of current observations to make a classification. We leverage recent advances in long-term probabilistic forecasting, namely *Deep Probabilistic Koopman*, to build a general method for classifying anomalies in multi-dimensional time-series data. We also show how to utilize models with domain knowledge of the dynamics to reduce type I and type II error. We demonstrate our proposed method on the important real-world task of global atmospheric pollution monitoring, integrating it with NASA's Global Earth Observing System Model. The system successfully detects localized anomalies in air quality due to events such as COVID-19 lockdowns and wildfires.

## 1. Introduction

*Time-series analysis* is used to extracting meaningful statistics and characteristics of temporal sequences of data [1], and is among the most ubiquitous mathematical methods. Indeed, time-series are universal for signal processing methods and in pattern recognition applications, dominating characterization of econometrics and finance along with almost any scientific and engineering application. Time-series methods can be broadly divided into time-domain and frequency-domain methods, the former of which uses a variety of statistical techniques to characterize a sequence, and the latter of which uses spectral (e.g. Fourier) decompositions and wavelets as the underlying representation of the signal. While the goals of time-series analysis are diverse, including signal estimation, signal classification, signal segmentation, and prediction and forecasting, time-series analysis methods almost universally aim to construct a simple and accurate model underlying the observations. Noisy, non-stationary data with anomalous exogenous events are the most difficult to characterize for any method. However, emerging machine learning methods, which are advocated here, allow for a flexible mathematical framework capable of characterizing even such difficult time-series. We leverage recent advances in long-term probabilistic forecasting, namely the Deep Probabilistic Koopman (DPK) method [2, 3], to build a general method for classifying anomalies in multi-dimensional time-series data. DPK combines deep learning with a spectral estimation approach in order to build a model that accommodates noisy, nonstationary time-series data, further allowing for the detection of anomalous events.

The importance of time-series analysis has led to a broad range of mathematical innovations characterizing sequential temporal data. In addition to traditional statistical methods, there are growing number of methods using concepts from dynamical systems [4–6] and machine learning [7, 8] for the

analysis of time-series data. Traditional statistical methods include the time-domain methods, such as the family of autoregressive (AR) models and their many variants, including AR moving average, AR integrated moving average (ARIMA), seasonal ARIMA, etc [1]. Such models use a diversity of optimization techniques to estimate parameters of a linear model with its history dependence. Traditional frequency-domain methods use the properties of short-time Fourier transforms [9] and/or wavelet transforms [10] in order to characterize the signal in a joint time-frequency representation. More recently, there have been efforts to model time-series data as from a dynamical systems perspective [4–6]. Thus the basic concept is to identify a dynamical system as a generative model that underlies the source of the temporal sequence of data. Included in such methods are the *dynamic mode decomposition* [5, 11–13], Koopman operator [14–19] and SINDy (sparse identification of nonlinear dynamics) [6, 20, 21]. A third approach is offered by machine learning, whose flexible frameworks and universal approximation properties [22], have led to a diversity of neural network architectures [23, 24] for learning time-sequences, including recurrent neural networks and long short-term memories [25]. Although often lacking in interpretability and generalization, deep learning methods have typically exhibited superior performance when provided sufficient training data.

Here we present a novel and general method for classifying anomalies from multi-dimensional time-series data which leverages recent advances in long-term probabilistic forecasting, namely DPK. While DPK is designed to be a forecasting method, it makes strong assumptions about stability which makes it useful for detecting anomalies. For the remainder of this paper, we will refer to a DPK model as a time-dependent model, rather than a forecast, because our aim is not predicting future observations. DPK is a deep learning framework which takes time-series data and infers a probabilistic model whose parameters are expressible as a function of a minimal set of sinusoids in time [3]. This fitting procedure captures both the nonstationary aspects of the time-series data, as well as a Fourier-like Koopman decomposition of the signal in time. To demonstrate the capabilities of this method, we apply it to time-series of nitrogen dioxide ( $\text{NO}_2$ ), an important air pollutant closely linked to the burning of fossil fuel. The evolution of  $\text{NO}_2$  at a given location is fundamentally non-stationary in nature with unknown exogenous influences. Finding outliers in nonstationary time-series data is a challenging task because it involves creating a long-term probabilistic model of the data that generalizes to future timesteps. Unlike most other machine learning problems which are interpolative (i.e. the inference-time data usually lies in the convex hull of the training data), this problem is extrapolative, and requires training on data from thousands of timesteps prior to inference for sufficient data size and statistical power. The presented method is able to characterize air pollution anomalies, such as changes in  $\text{NO}_2$  in the wake of COVID-19 lockdowns. Moreover, the method detects these changes only few days after the onset of the anomaly, which makes it suitable for near real-time monitoring applications. An important aspect of our demonstration application is the combined use of publicly available atmospheric observations and corresponding computer simulations. Computer models, such as the NASA Goddard Earth Observing System (GEOS) Composition Forecast system (GEOS-CF) [26], are an indispensable tool to analyze and understand the drivers of air pollution, such as pollution transport, precursor emissions, and atmospheric chemistry [27]. We show how leveraging this information with observations can help reduce type I and type II errors.

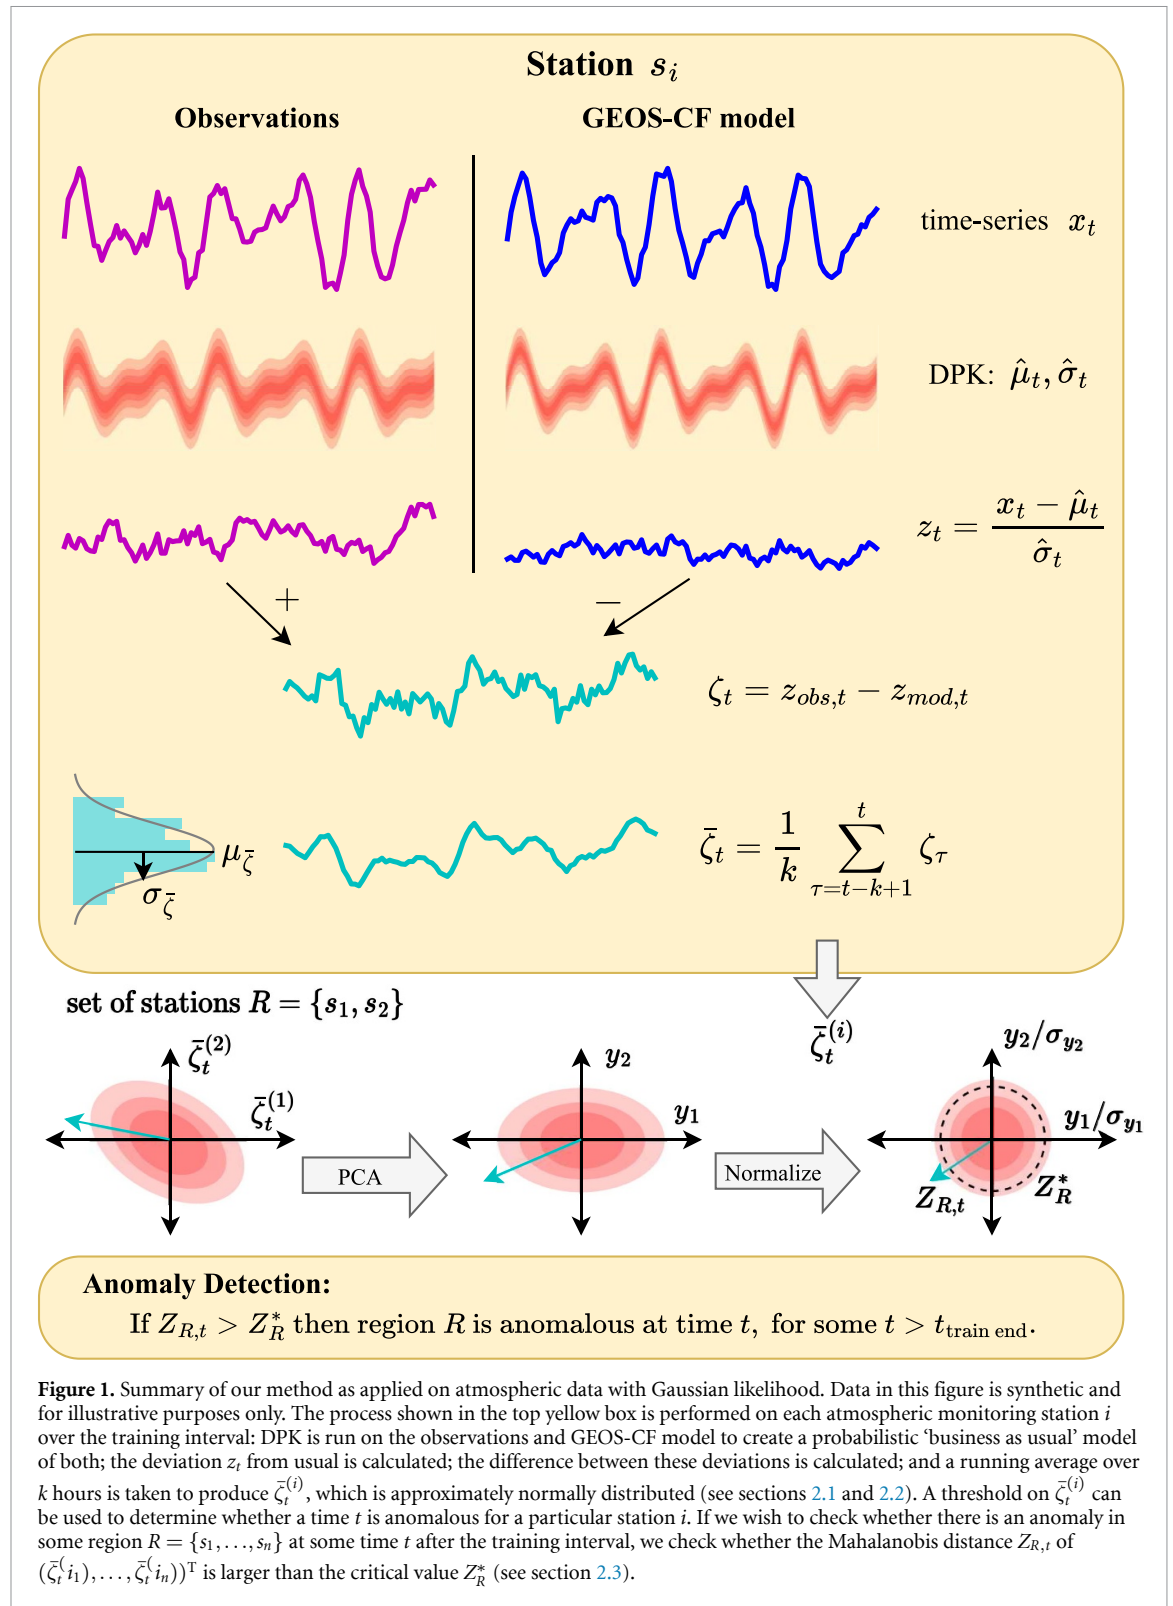
## 2. Methodology

We first present the basic methodology for classifying anomalous behavior in 1-dimensional time-series observations, then we show how this definition can optionally be expanded to leverage domain-specific models of the observations, and finally we present how to classify anomalies in multi-dimensional observations, allowing for more general and nuanced hypothesis testing. A summary of the methods can be seen in figure 1, and an example is shown in figure 2.

### 2.1. DPK for anomaly detection

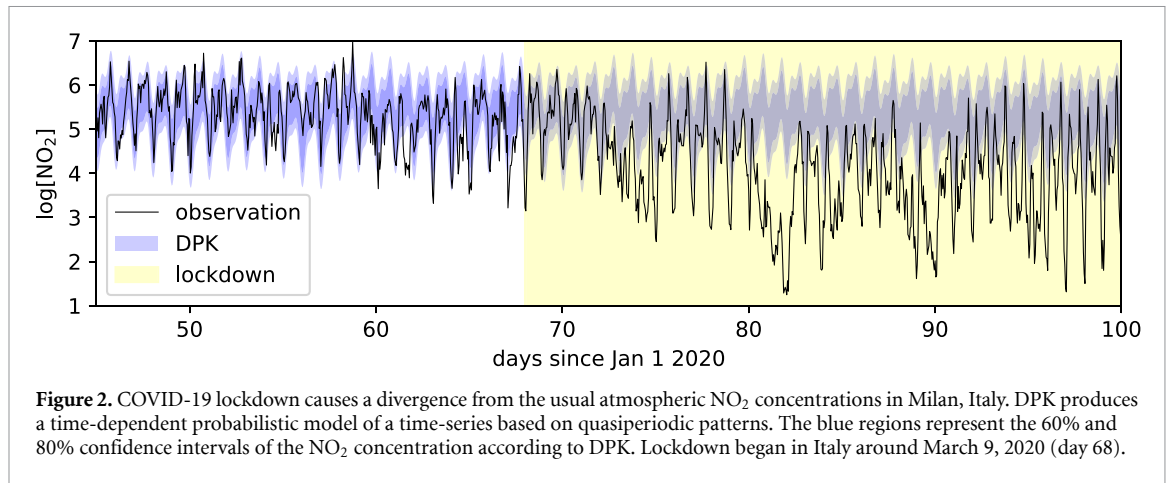
Here we present how to leverage the long-term probabilistic modeling ability of DPK to produce a calibrated anomaly classifier and monitor for such events in 1-dimensional time-series.

DPK produces a stable long-term probabilistic model of time-series data by expressing it as a parameterized probability distribution (in this case Gaussian) whose parameters vary as a function of sinusoids with different frequencies, as well as a limiting zero frequency for aperiodic trends. It is a result of Koopman operator theory [12, 15, 17, 18, 28] that any deterministic stable dynamical system can be expressed as a function of sinusoids with various frequencies [2, 14, 29]. This is extended to the case of stochastic stable dynamical systems by letting our mapping from the sinusoidal Koopman observables be stochastic. Mathematically,



$$X_t \sim P\left(\vec{\theta}_t = g_\Theta\left(\begin{bmatrix} \cos(\vec{\omega}t) \\ \sin(\vec{\omega}t) \end{bmatrix}\right)\right), \tag{1}$$

where  $\vec{\omega}$  is a vector of frequencies,  $P(\vec{\theta}_t)$  is a parameterized probability distribution such as a normal distribution (denoted  $\mathcal{N}(\vec{\theta}_t)$ ),  $\vec{\theta}_t$  are the distribution parameters, and  $g_\Theta$  is a feedforward neural network parameterized by  $\Theta$ . In this model  $x_t$  is drawn from  $X_t$  for all  $t$ . The neural network parameters  $\Theta$  are trained



through maximum likelihood. DPK can be understood as providing a ‘business as usual’ model of a time-series based on seasonal patterns.

As described in [3], DPK can be applied only on quasi-periodic time-series, those expressible as a nonlinear function of sinusoids with different frequencies, modulo a long-term trend.

Once a DPK model has been trained on the time-series of observations, it is used for anomaly detection. First consider the case of univariate time-series for which there is no domain-specific model (such as the GEOS-CF model). Consider the time-series of observed values  $x_{obs,t}$ . For a choice of time-varying distribution  $P(\vec{\theta}_t)$ , we define  $z_{obs,t} = \text{ICDF}_{\mathcal{N}(0,1)}(\text{CDF}_{P(\vec{\theta}_t)}(x_{obs,t}))$ , where CDF denotes the cumulative distribution function and ICDF denotes the inverse CDF. This can be understood as transforming the observations such that we model their distribution as a standard normal distribution at each point in time. When  $P$  is a Gaussian this is simply the z-score  $(x_{obs,t} - \hat{\mu}_{obs,t}) / \hat{\sigma}_{obs,t}$ . If our choice of  $P$  is accurate and the DPK model is calibrated, then the distribution of  $z_{obs,t}$  will empirically be Gaussian. Thus, a time is classified as anomalous based on the average of  $z_{obs,t}$  over  $k$  consecutive timesteps,  $\bar{z}_{obs,t} = \frac{1}{k} \sum_{\tau=t-k+1}^t z_{obs,\tau}$ . Consecutive values of  $z_{obs,t}$  are positively correlated due to physical continuity over time, so we cannot apply the central limit theorem to obtain the sampling distribution of  $\bar{z}_{obs,t}$ . Instead, the empirical sampling distribution of  $\bar{z}_{obs,t}$  is used, and fit with a Gaussian distribution  $\mathcal{N}(\mu_{\bar{z}_{obs}}, \sigma_{\bar{z}_{obs}})$ . The sampling distribution of  $\bar{z}_{obs,t}$  is Gaussian on the training distribution, as can be observed in figure 3(a), which is expected because the distribution of  $z_{obs,t}$  is Gaussian.

**Univariate anomaly detection without a domain-specific model:** A time  $t$  is considered anomalous if and only if  $\bar{z}_{obs,t}$  is as extreme or more extreme than the two-tailed critical value  $\bar{z}_{obs}^*$  for  $\mathcal{N}(\mu_{\bar{z}_{obs}}, \sigma_{\bar{z}_{obs}})$  at a certain  $\alpha$  level. That is, when the probability of observing a  $\bar{z}_{obs,t}$  as extreme or more extreme than that observed, assuming our business-as-usual model, is less than  $\alpha$ , then time  $t$  is considered anomalous.

These anomalous events represent times in which the recent values of  $x_{obs,t}$  cannot be explained by usual variation as modeled by DPK.

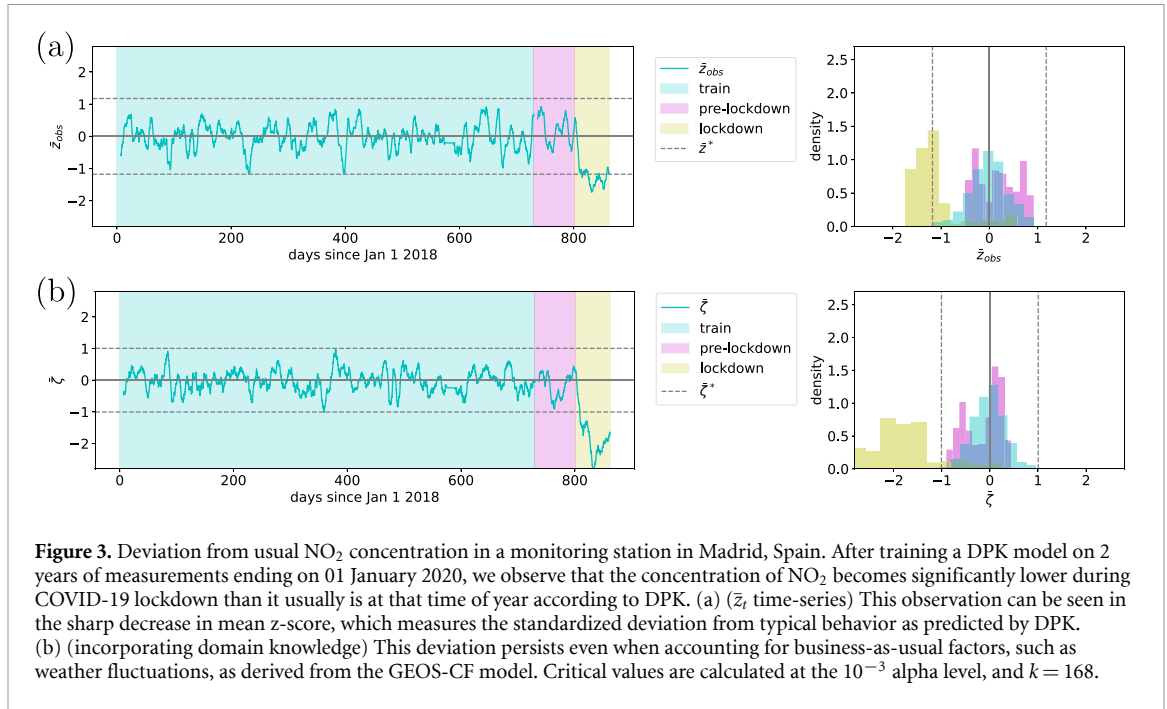
Because the training distribution is supposed to represent observations which are not anomalous, we remove outliers from the training distribution before computing  $\mu_{\bar{z}_{obs}}$  and  $\sigma_{\bar{z}_{obs}}$ . We discard  $\bar{z}_{obs}$  values not in  $[Q_1 - \lambda \text{IQR}, Q_3 + \lambda \text{IQR}]$  where  $Q_1$  and  $Q_3$  are the first and third quartiles, IQR is the interquartile range, and we choose  $\lambda = 2$ .

It should also be noted that long-term probabilistic models are subject to overconfidence due to imperfect extrapolation and overfitting. DPK is typically not more than 10 to 30% overconfident, as measured by the root mean square of test-time z-scores, but this overconfidence is accounted for in our selection of the  $\alpha$  level via a hyperparameter sweep in appendix section ‘Hyperparameter selection’.

## 2.2. Incorporating domain knowledge

We can incorporate domain-specific models into our anomaly detection system such that known events are not considered anomalies. This is an optional modification to the technique, and the remaining contributions in section 2.3 should be of general applicability if this step is skipped because no such model is available for the application.

When a domain-specific model’s predictions  $x_{mod,t}$  are available, we let  $\zeta_t = z_{obs,t} - z_{mod,t}$  and use  $\zeta_t$  in place of  $z_{obs,t}$ .  $z_{mod,t}$  is calculated in the same way as  $z_{obs,t}$ , except that it uses the model time-series  $x_{mod,t}$



rather than observations. This implies that we train a DPK model on the domain-specific model’s predictions in order to obtain a time-dependent model of the domain-specific model’s usual predictions. Deviations from this DPK model correspond to differences that can be attributed to known factors. As before, an anomaly is flagged when the current  $\bar{z}_t$  is extreme in the sampling distribution  $\mathcal{N}(\mu_{\bar{z}}, \sigma_{\bar{z}})$  of  $\bar{z}_t$ . This process is shown in the main box of figure 1. Figure 3 shows a comparison between using  $\bar{z}_{obs}$  and using  $\bar{z}$ .

### 2.3. Extending to multidimensional time-series

Given a set of observables  $R = \{s_1, \dots, s_n\}$ , we wish to determine whether the entire system those observables measure is anomalous, as shown in figure 5. Let  $\bar{\zeta}_{R,t} = (\bar{\zeta}_t^{(s_1)}, \dots, \bar{\zeta}_t^{(s_n)})^T$ , where  $\bar{\zeta}_t^{(s_i)}$  is the  $\bar{z}_t$  described in section 2.2 for station  $s_i$ .

We employ principal component analysis (PCA) to detect anomalies in multivariate data. PCA finds linearly independent basis dimensions in a dataset, sorted by the extent to which they explain the variation in all of the observables, enabling us to remove the least explanatory (noisiest) directions. Let our multivariate Gaussian estimate of the distribution from which  $\bar{\zeta}_{R,t}$  is drawn be denoted  $\mathcal{N}(\mu_R, \Sigma_R)$ . Assuming a multivariate normal distribution allows us to estimate the joint distribution using the product of the marginal distributions along each principal component, which provides a practical way to estimate the likelihood of an event.

While PCA produces linearly independent random variables, it does not guarantee statistically independent random variables, which is required in order to factor the joint distribution in terms of the marginal distributions. However, in most applications it seems reasonable to assume that any statistical dependence between the residual errors of different observables is approximately linear, such that PCA produces approximately statistically independent random variables. This can be supported by visually verifying random pairs of PCA RVs are independent (granted, pairwise independence does not imply mutual independence). We empirically verify this method’s calibration (see section 3) as a test of all assumptions.

We use Mahalanobis distance to determine whether an observation is anomalous. Usually, the Mahalanobis distance between an observation  $\mathbf{d} \in \mathbb{R}^n$  and a distribution  $D$  with mean  $\mu_d$  and covariance matrix  $S_d$  is given by

$$\delta_M(\mathbf{d}, D) = \sqrt{(\mathbf{d} - \mu_d) S_d^{-1} (\mathbf{d} - \mu_d)}. \tag{2}$$

However, we rewrite this in terms of our principal component projection so that we can remove noisy dimensions. After performing PCA so that  $y = V^T(\mathbf{d} - \mu_d)$ , where the rows of  $V^T$  are the eigenvectors of the



empirical covariance matrix, we have  $\mu_y = \mathbf{0}$  and  $S_y = \text{diag}(\sigma_1^2, \dots, \sigma_n^2)$  is a diagonal matrix of the variance along each principal component. Then the Mahalanobis distance simplifies to

$$\delta_M(\mathbf{d}, D) = \left\| \begin{bmatrix} y_1/\sigma_1 \\ \vdots \\ y_n/\sigma_n \end{bmatrix} \right\|_2. \quad (3)$$

Unlike the usual Mahalanobis distance, we do not use all  $n$  principal components shown in equation (3), since the components with the smallest variance  $\sigma_i^2$  are mostly noise. Instead, we truncate to the smallest dimensionality that explains at least 90% of the variance [9, 30], although more sophisticated methods for rank truncation exist [31, 32]. See the appendix figure 8 for a hyperparameter sweep demonstrating this.

$\mathbf{d}$  is drawn from a multivariate Gaussian, so  $(y_1/\sigma_1, \dots, y_n/\sigma_n)^T$  is drawn from a spherical Gaussian with mean  $\mathbf{0}$  and radius 1, and the  $p$ -value of  $\mathbf{d}$  can be calculated by integrating the probability density over the volume lying outside of the  $n$ -ball of radius  $\delta_M(\mathbf{d}, D)$  centered at the origin. Let  $Z_{R,t} = \delta_M(\bar{\zeta}_{R,t}, \mathcal{N}(\mu_R, \Sigma_R))$ . Then the  $p$ -value, the probability of making an observation as extreme or more extreme than a given observation  $Z_{R,t}$  assuming 'business as usual,' is given by

$$p = P(Z_R \geq Z_{R,t}) \propto \int_{Z_{R,t}}^{\infty} Z^{n-1} e^{-Z^2} dZ, \quad (4)$$

for  $Z_R, t \geq 0$ . We numerically integrate and normalize this to calculate the  $p$ -value.

#### 2.4. Application to atmospheric pollution

Here we demonstrate our method using ground-based NO<sub>2</sub> observations [33] obtained from the OpenAQ platform<sup>5</sup> and the air quality data portal of the European Environment Agency<sup>6</sup>. The time-series  $x_{obs,t}$  is the log of the observed concentration.

The DPK model is trained using observations from years 2018–2019, and evaluated over 2020. We choose to provide the DPK model with four frequencies  $\omega_i$  observed in our data: daily, weekly, annual, and the long-term trend. We choose  $P$  to be a Gaussian because residuals are approximately normally distributed.

We employ GEOS-CF to incorporate domain knowledge into our anomaly detection pipeline. The GEOS-CF model is a global atmospheric chemistry model that provides global hourly model estimates of atmospheric composition at  $25 \times 25 \text{ km}^2$  spatial resolution [26], available in near real-time at <sup>7</sup>. It is important to note that GEOS-CF uses historical estimates as anthropogenic emissions input, and thus does not capture changes in anthropogenic emission reductions related to COVID-19 restrictions. However, GEOS-CF captures other events that are known to impact air pollution, most notably meteorology and wildfires.

We note that despite having a domain-specific model for this application, GEOS-CF cannot be directly used for anomaly detection because (1) it is not probabilistic, (2) it does not necessarily capture all long-term patterns, and (3) its comparison to observations is complicated by representation errors ( $25 \times 25 \text{ km}^2$  model averages vs. point source observations) and systematic model biases. DPK is uniquely suited to solve all of these problems because it (1) is probabilistic, (2) solves a global (frequency domain) optimization problem which enables long-term modeling ability, and (3) converges to an unbiased model.

### 3. Results

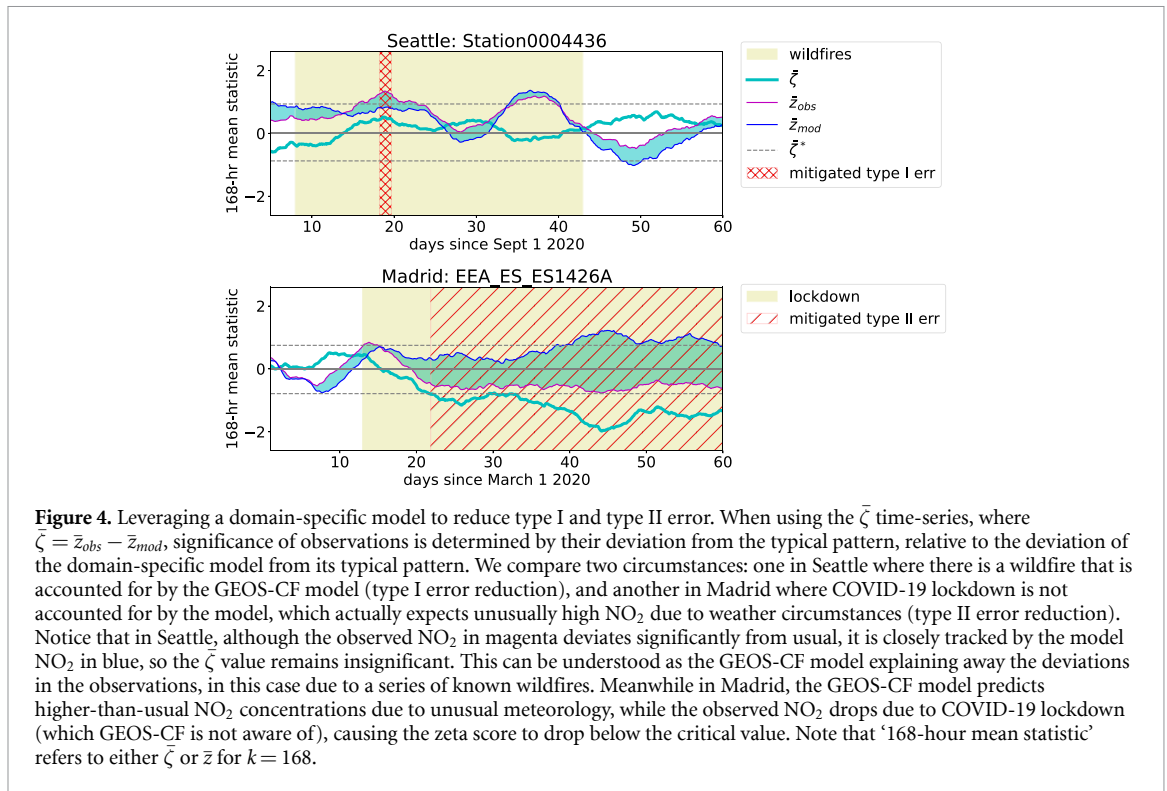
Figure 2 shows observed NO<sub>2</sub> concentrations over Milan, Italy (black) for February–May 2020 compared against the corresponding time-dependent DPK model (red). The DPK-predicted NO<sub>2</sub> concentrations agree well with the observations for all of February and the first days of March 2020, but the observations start to diverge from the DPK values afterwards, reflecting the observed decline of NO<sub>2</sub> in the wake of the COVID-19 pandemic [34]. By design, this decrease is not captured by DPK, which allows for detection of anomalies.

Figure 3(a) shows the time-series of the  $\bar{z}_{obs,t}$  used to make anomaly classifications. The training distribution is normally distributed as expected. The pre-lockdown control distribution does not contain any events classified as anomalous (as desired), though there is insufficient data to determine whether  $\bar{z}_{obs,t}$

<sup>5</sup> <https://openaq.org/>.

<sup>6</sup> <https://discomap.eea.europa.eu/map/fme/AirQualityExport.htm>.

<sup>7</sup> [https://gmao.gsfc.nasa.gov/weather\\_prediction/GEOS-CF/data\\_access/](https://gmao.gsfc.nasa.gov/weather_prediction/GEOS-CF/data_access/).



continues to be normally distributed. Meanwhile, soon after lockdown begins most  $\bar{z}_{obs,t}$  values cross the threshold to be classified as anomalous.

### 3.1. Incorporating the GEOS-CF model

We only wish to be alerted of anomalies that are not already known by the GEOS-CF model. As can be seen in figure 4, incorporating GEOS-CF reduces both type I error and type II error. Type I error is mitigated when domain knowledge introduced by GEOS-CF explains away the deviation in  $\text{NO}_2$ , so that no anomaly is flagged. Type II error is mitigated when an anomaly is veiled by other (e.g. meteorological) factors that push observations in the opposite direction of the anomaly, and the anomaly is only noticed when comparing the observations and the domain-specific model.

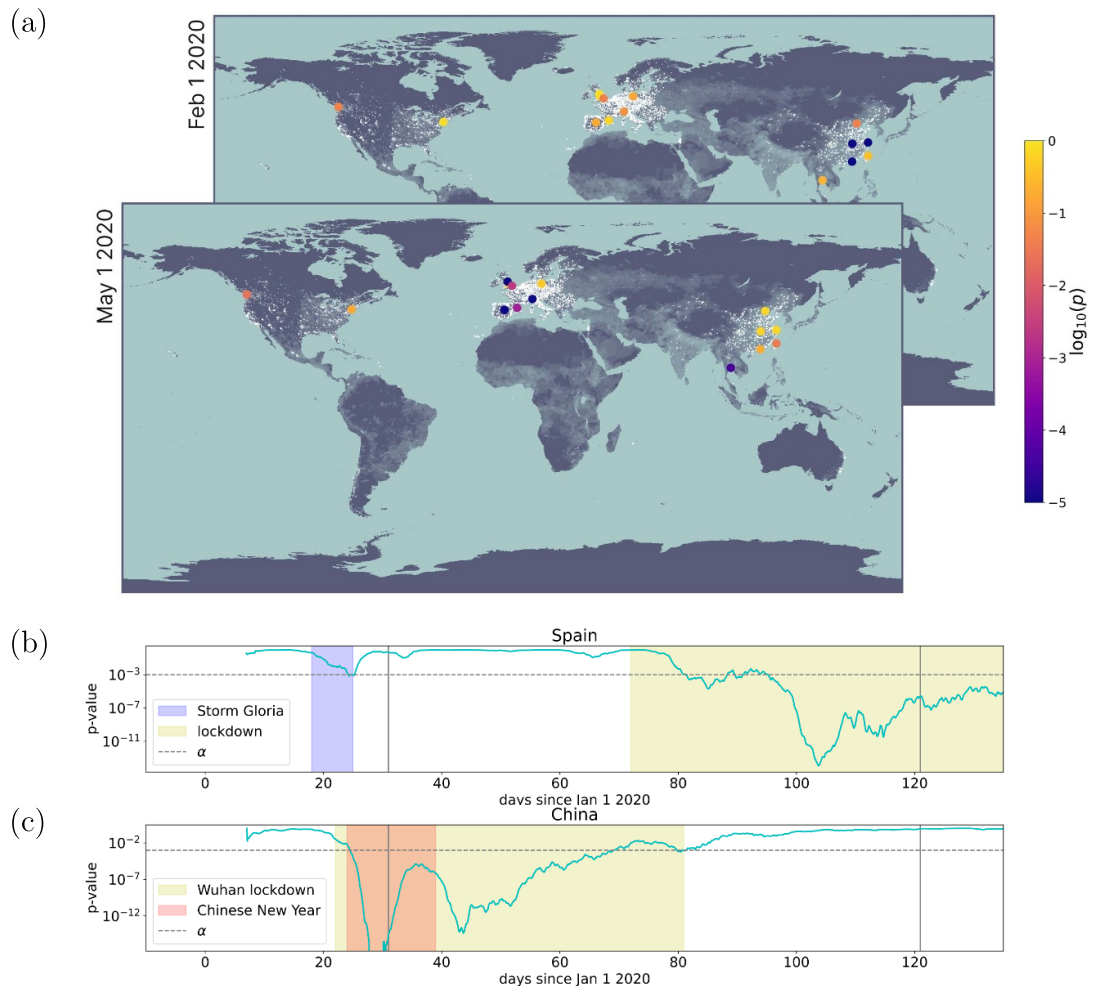
### 3.2. Global and regional atmospheric anomaly detection

Monitoring for anomalies in a hierarchy of stations has multiple advantages: (1) being able to ask a wider variety of questions about  $\text{NO}_2$  anomalies. (2) being able to leverage more data to evaluate questions. In figure 7, we demonstrate hierarchical anomaly detection in Spain.

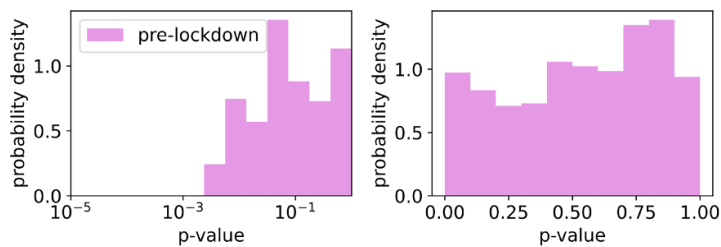
We find that our method successfully detects many known anomalies, and otherwise has a uniform distribution of  $p$ -values on the interval  $[0, 1]$ . We ran our analysis on various of the largest cities from around the world (filtered for data quality) to validate our method. We chose to look at cities because we expect their air quality to significantly improve when COVID-19 lockdown reduces traffic volume, and because cities have anywhere from a few to several dozen monitoring stations, allowing us to evaluate on a wide variety of time-series dimensionalities. The training interval is from 01 January 2018 to 01 January 2020, the explained variance threshold is set to 90%, and  $k = 168$ . In figure 5(a) the effects of COVID-19 are statistically significant in Chinese cities in February, when the outbreak was mostly limited to China; then by May China has largely resumed normal activity while other major cities are in lockdown, which is reflected in the  $p$ -values of  $\text{NO}_2$  concentrations. When looking at the national level in China and Spain (5(b) and (c)), it can be seen that low  $p$ -values generally correspond to known exceptional events, including Storm Gloria, COVID-19 lockdown, and Chinese New Year (in which entire cities are largely shut down for days).

Our statistical method is calibrated. In figure 6, we see that  $p$ -values are roughly uniformly distributed pre-lockdown for various cities. Selection criteria for these cities is described in the appendix. In figure 7(d) we see that our method is calibrated for individual stations in Spain.

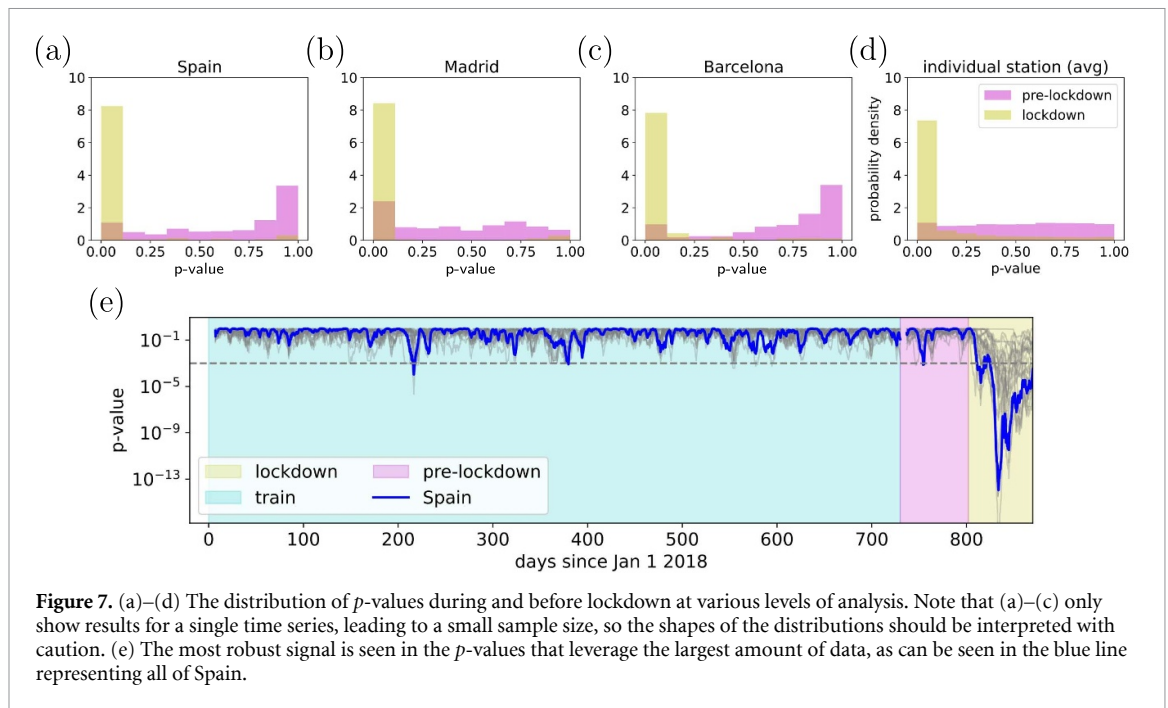




**Figure 5.** A global monitoring system. (a) Of the major cities shown in colored dots, with darker colors indicating more anomalous  $\text{NO}_2$  values, Chinese cities can be seen experiencing the effects of lockdown first, followed later by cities outside of China. By May 1, Chinese activity is roughly back to normal while European cities and Bangkok still show the effects of lockdown. (b), (c) The multivariate anomaly detection method allows for anomalies to be detected at the level of countries. For example, (b) the effects of Storm Gloria and COVID-19 lockdown can be seen in Spanish cities, and (c) the effects of Chinese New Year and COVID-19 lockdown can be seen in Chinese cities.



**Figure 6.** Distribution of  $p$ -values pre-lockdown, averaged over several cities, shown in both log and linear scale. During periods without known anomalies,  $p$ -values are approximately uniformly distributed, indicating a calibrated method. While figure 7 shows that the method is calibrated for univariate time-series, this figure shows that the method is also calibrated for high dimensional time-series (typically a few dozen stations), and does not produce excess false positives. As can be seen in the logarithmic scaled plot, none of the observations in this control setting have a  $p$ -value of less than  $\alpha = 10^{-3}$ , the threshold value for statistical significance we use.



## 4. Conclusion

We introduced a novel and general method for anomaly detection in complex high-dimensional stochastic time-series data. Our method makes use of recent advances in long-term probabilistic forecasting based on Koopman theory to create a robust ‘business as usual’ model of observations, against which current observations can be compared. We make use of multidimensional time-series signals to enable online detection of anomalies within several dozen timesteps, or a few days in our application example. The statistical power of the anomaly detection can be improved by incorporating domain-specific numerical models. We demonstrate this using time series of  $\text{NO}_2$  in conjunction with the GEOS-CF numerical atmospheric composition model. Our results show how this technique can be used for near real-time detection of atmospheric anomalies. While our example used ground observations, the presented method is equally applicable to satellite observations.

When applied to other time-series, such as electricity demand, we could see our methodological contributions aiding in feature discovery by indicating when observations deviated significantly from the long-term trend. A potential future direction is automatically selecting subsets of the observation vector that might be anomalous, rather than having to check exponentially-many possible subsets or hand-select them.

Given the ubiquity of *Time-series analysis*, the advocated Koopman based methods provides an powerful mathematical method for extracting meaningful statistics and characteristics of temporal sequences of data. Indeed, Koopman theory provides an interpretable framework from which a variety of signal processing goals and pattern recognition applications can be characterized in diverse engineering and scientific fields.

## Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://gmao.gsfc.nasa.gov/gmaoftp/geoscf/UW/>.

## Acknowledgments

J N K acknowledges funding from the National Science Foundation AI Institute in Dynamic Systems Grant Number 2112085.

## Appendix

### Experimental details

#### *DPK training details*

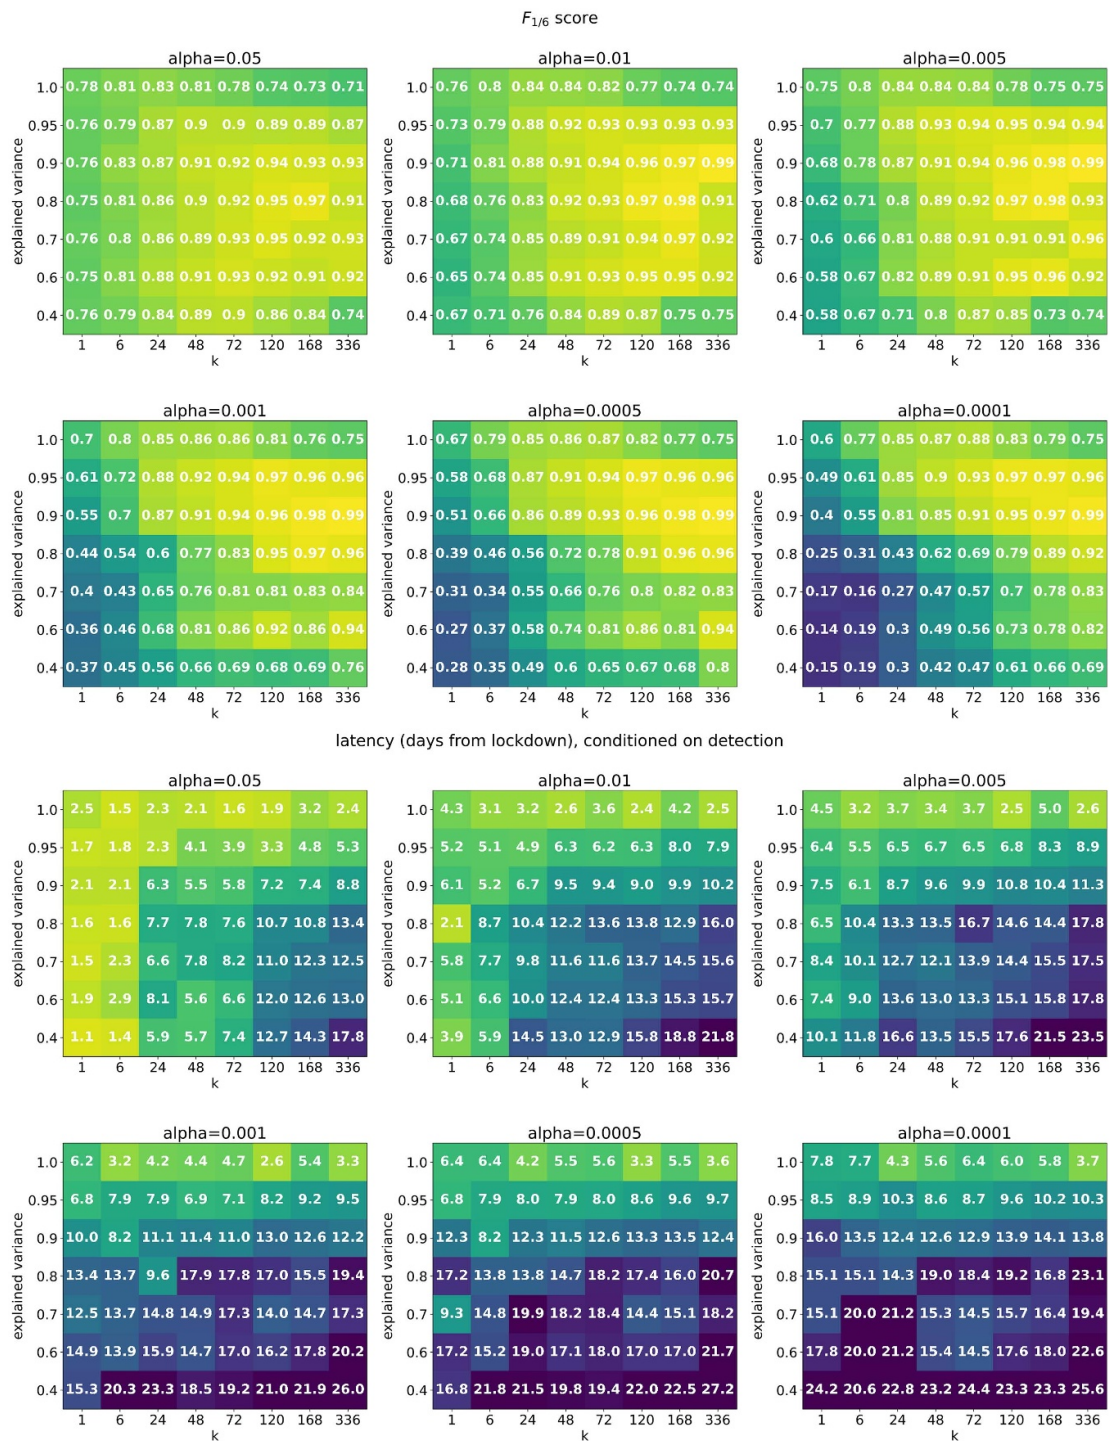
The code is available at <https://github.com/AlexTMallen/DPK-anomaly-detection>. The DPK model for each station is trained with a learning rate of 0.0001, a weight decay regularization of 0.001, for 400 epochs, on hourly observations from 01 January 2018 to 01 January 2020. Each model is initialized using the parameters of a model trained on another station's observations to aid efficiency through transfer learning.

#### *Selection criteria for cities in section 6*

From a list of major cities around the world, we look at high-pollution (greater than 10 ppb NO<sub>2</sub> multi-year average), high-population cities with clean data during the time in 2020 before COVID-19 lockdown. Most cities are excluded because of a lack of clean data, if any is available at all. We also exclude a city if it has a significant  $p$ -value that corresponds with a known anomalous event. Here are the cities we excluded due to known anomalies. Madrid and Barcelona experienced Storm Gloria ([https://en.wikipedia.org/wiki/Storm\\_Gloria](https://en.wikipedia.org/wiki/Storm_Gloria)). Birmingham, London, and Berlin experienced Storm Ciara ([https://en.wikipedia.org/wiki/Storm\\_Ciara](https://en.wikipedia.org/wiki/Storm_Ciara)).

### Hyperparameter selection

Hyperparameter tuning is critical for almost any data-driven method. In the application considered here, hyperparameter sweeps are done over the  $\alpha$  level, the number of hours  $k$  over which to average, and the explained variance threshold for calculating  $Z_{R,t}$ . Figure 8 summarizes the performance of the hyperparameter sweeps.



**Figure 8.** Hyperparameter sweep over the  $\alpha$  level, the number of hours  $k$  over which to average, and the explained variance threshold for calculating  $Z_{R,t}$ . We measure both  $F_{1/6}$  classification accuracy and latency of the anomaly detection, conditioned on detection.  $F_{1/6} = (1 + \frac{1}{6^2})(\text{precision} \cdot \text{recall}) / (\frac{1}{6^2}\text{precision} + \text{recall})$  is an accuracy measure that places 6 times as much weight on precision as on recall. We place more weight on precision than recall for two reasons. First, most times are not anomalous, so a low precision would cause an excessive quantity of false positives. Second, not all times within an anomalous event need to be recalled-as long as a single time within the event is classified as anomalous, we can be alerted of the anomaly. Based on these results we selected 90% explained variance,  $k=168$  hours (1 week) averaging, and  $\alpha=0.001$ .



## ORCID iDs

Alex Mallen  <https://orcid.org/0000-0002-5120-1350>

J Nathan Kutz  <https://orcid.org/0000-0002-6004-2275>

## References

- [1] Shumway R H, Stoffer D S and Stoffer D S 2000 *Time Series Analysis and its Applications* vol 3 (Berlin: Springer)
- [2] Lange H, Brunton S L and Nathan Kutz J 2021 From fourier to koopman: spectral methods for long-term time series prediction *J. Mach. Learn. Res.* **22** 1–38
- [3] Mallen A, Lange H and Kutz J N 2021 Deep probabilistic koopman: long-term time-series forecasting under periodic uncertainties (arXiv:2106.06033)
- [4] Hoffmann M et al 2021 Deeptime: a python library for machine learning dynamical models from time series data *Mach. Learn.: Sci. Technol.* **3** 015009
- [5] Kutz J N, Brunton S L, Brunton B W and Proctor J L 2016 *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems* (Philadelphia, PA: SIAM)
- [6] Brunton S L and Kutz J N 2019 *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems and Control* (Cambridge: Cambridge University Press)
- [7] Dieterich T G 2002 Machine learning for sequential data: a review *Joint IAPR Int. Workshops on Statistical Techniques in Pattern Recognition (SPR) and Structural And Syntactic Pattern Recognition (SSPR)* (Springer) pp 15–30
- [8] Masini R P, Medeiros M C and Mendes E F 2021 Machine learning advances for time series forecasting *J. Econ. Surv.* **37** 76–111
- [9] Kutz J N 2013 *Data-Driven Modeling & Scientific Computation: Methods for Complex Systems & big Data* (Oxford: Oxford University Press)
- [10] Mallat Sephane 1999 *A Wavelet Tour of Signal Processing* (Amsterdam: Elsevier)
- [11] Schmid P J 2010 Dynamic mode decomposition of numerical and experimental data *J. Fluid Mech.* **656** 5–28
- [12] Rowley C W, Mezic I, Bagheri S, Schlatter P and Henningson D S 2009 Spectral analysis of nonlinear flows *J. Fluid Mech.* **645** 115–27
- [13] Askham T and Nathan Kutz J 2018 Variable projection methods for an optimized dynamic mode decomposition *SIAM J. Appl. Dyn. Syst.* **17** 380–416
- [14] Koopman B O 1931 Hamiltonian systems and transformation in Hilbert space *Proc. Natl Acad. Sci. USA* **17** 315
- [15] Mezić I 2005 Spectral properties of dynamical systems, model reduction and decompositions *Nonlinear Dyn.* **41** 309–25
- [16] Brunton S L, Budišić M, Kaiser E and Nathan Kutz J 2021 Modern koopman theory for dynamical systems (arXiv:2102.12086)
- [17] Budišić M, Mohr R and Mezić I 2012 Applied koopmanism *Chaos* **22** 047510
- [18] Mezić I 2013 Analysis of fluid flows via spectral properties of the koopman operator *Annu. Rev. Fluid Mech.* **45** 357–78
- [19] Arbabi H and Mezić I 2017 Ergodic theory, dynamic mode decomposition and computation of spectral properties of the koopman operator *SIAM J. Appl. Dyn. Syst.* **16** 2096–126
- [20] Brunton S L, Proctor J L and Nathan Kutz J 2016 Discovering governing equations from data by sparse identification of nonlinear dynamical systems *Proc. Natl Acad. Sci.* **113** 3932–7
- [21] Champion K, Lusch B, Nathan Kutz J and Steven L B 2019 Data-driven discovery of coordinates and governing equations *Proc. Natl Acad. Sci.* **116** 22445–51
- [22] Hornik K 1991 Approximation capabilities of multilayer feedforward networks *Neural Netw.* **4** 251–7
- [23] Goodfellow I, Bengio Y and Courville A 2016 *Deep Learning* (Cambridge, MA: MIT Press)
- [24] Yong Y, Xiaosheng S, Changhua H and Zhang J 2019 A review of recurrent neural networks: Lstm cells and network architectures *Neural Comput.* **31** 1235–70
- [25] Hochreiter S and Schmidhuber Jurgen 1997 Long short-term memory *Neural Comput.* **9** 1735–80
- [26] Keller C A et al 2021 Description of the nasa geos composition forecast modeling system geos-cf v1.0 *J. Adv. Model. Earth Syst.* **13** e2020MS002413
- [27] Brasseur G P and Jacob D J 2017 *Modeling of Atmospheric Chemistry* (Cambridge: Cambridge University Press)
- [28] Brunton S L, Budišić M, Kaiser E and Nathan Kutz J 2021 Modern Koopman theory for dynamical systems (arXiv:2102.12086)
- [29] Koopman B O and Neumann J V 1932 Dynamical systems of continuous spectra *Proc. Natl Acad. Sci. USA* **18** 255
- [30] Berkooz G, Holmes P and Lumley J L 1993 The proper orthogonal decomposition in the analysis of turbulent flows *Annu. Rev. Fluid Mech.* **25** 539–75
- [31] Gavish M and Donoho D L 2014 The optimal hard threshold for singular values is  $4\sqrt{3}$  *IEEE Trans. Inf. Theory* **60** 5040–53
- [32] Forkman J, Josse J and Piepho H-P 2019 Hypothesis tests for principal component analysis when variables are standardized *J. Agric. Biol. Environ. Stat.* **24** 289–308
- [33] Keller C 2021 *Observations and Model Forecasts of Atmospheric NO<sub>2</sub>* (GEOS-CF) (available at: <https://gmao.gsfc.nasa.gov/gmaoftp/geoscf/UW/>)
- [34] Keller C A et al 2021 Global impact of covid-19 restrictions on the surface concentrations of nitrogen dioxide and ozone *Atmos. Chem. Phys.* **21** 3555–92