



HAL
open science

Uncovering differential equations from data with hidden variables

Agustín Somacal, Yamila Barrera, Leonardo Boechi, Matthieu Jonckheere,
Vincent Lefieux, Dominique Picard, Ezequiel Smucler

► **To cite this version:**

Agustín Somacal, Yamila Barrera, Leonardo Boechi, Matthieu Jonckheere, Vincent Lefieux, et al..
Uncovering differential equations from data with hidden variables. *Physical Review E* , 2022, 105 (5),
pp.054209. 10.1103/physreve.105.054209 . hal-03945967

HAL Id: hal-03945967

<https://hal.science/hal-03945967>

Submitted on 5 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Uncovering differential equations from data with hidden variables

Agustín Somacal¹, Yamila Barrera¹

Leonardo Boechi², Matthieu Jonckheere², Vincent Lefieux³, Dominique Picard⁴, Ezequiel Smucler^{5,1}

Abstract

SINDy is a method for learning system of differential equations from data by solving a sparse linear regression optimization problem [Brunton et al., 2016]. In this article, we propose an extension of the SINDy method that learns systems of differential equations in cases where some of the variables are not observed. Our extension is based on regressing a higher order time derivative of a target variable onto a dictionary of functions that includes lower order time derivatives of the target variable. We evaluate our method by measuring the prediction accuracy of the learned dynamical systems on synthetic data and on a real data-set of temperature time series provided by the Réseau de Transport d'Électricité (RTE). Our method provides high quality short-term forecasts and it is orders of magnitude faster than competing methods for learning differential equations with latent variables.

Keywords: Differential equations, Dynamical systems, Lasso, Latent variables, Machine learning

1. Introduction

Many branches of science are based on the study of dynamical systems. Examples include meteorology, biology and physics. The usual way to model deterministic dynamical systems is by using (partial) differential equations. Typically, differential equations models for a given dynamical system are derived using apriori insights into the problem at hand; then the model is validated using empirical observations. In an era in which massive data-sets pertaining to different fields of science are widely

Email addresses: a.somacal@aristas.com.ar (Agustín Somacal), y.barrera@aristas.com.ar (Yamila Barrera), lboechi@ic.fcen.uba.ar (Leonardo Boechi), mjonckhe@dm.uba.ar (Matthieu Jonckheere), vincent.lefieux@rte-france.com (Vincent Lefieux), picard@math.univ-paris-diderot.fr (Dominique Picard), e.smucler@aristas.com.ar (Ezequiel Smucler)

¹Aristas S.R.L., Dorrego 1940, Torre A, 2do Piso, dpto. N (1425), Ciudad Autónoma de Buenos Aires, Argentina

²Instituto de Calculo-CONICET, Intendente Güiraldes 2160, Ciudad Universitaria, Pabellón II, 2do. piso, (C1428EGA), Buenos Aires, Argentina

³Réseau de Transport d'Electricité (RTE), France.

⁴Université de Paris, LPSM, UFR Mathématiques Batiment Sophie Germain, 75013 Paris.

⁵Universidad Torcuato Di Tella, Av. Figueroa Alcorta 7350 (C1428BCW) Sáenz Valiente 1010 (C1428BIJ) Ciudad de Buenos Aires, Argentina

available, an interesting problem is whether it is possible for a useful differential equations model to be learned directly from data, without any major modeling effort required by the researcher.

The SINDy method is an approach of sparse identification of nonlinear dynamical systems that consists of linking the dynamical system discovery problem to a statistical regression problem ([Brunton et al., 2016], [Rudy et al., 2017] and [Quade et al., 2018]). The main idea of the SINDy method is to consider a set of differential operators (possibly in both time and space if appropriate), discretize them, for example by using finite differences, and then regress the outcome of interest on the discretized differential operators. By solving the regression problem using an ad-hoc thresholded least-squares algorithm, they can build sparse, interpretable models, that use mostly low order derivatives. The authors explored the applicability of their method on simulated data, but only in situations in which all the variables of the simulated models are observed.

Our goal is to extend the SINDy model for the case in which not all relevant variables are observed, that is, in cases in which the main variable of interest depends on other variables of which no measurements are available. As an example of application we consider the climate time series of the Réseau de Transport d'Électricité (RTE). RTE is the main electricity network operational manager in France, who is interested in understanding the behavior of climate time series because of their impact on energy consumption. RTE uses high-level simulations of hourly temperature series to study the impact different climate scenarios have on electricity consumption, and hence on the French electrical power grid. The underlying simulations are based on the Navier-Stokes equations and include variables as wind velocity, density, pressure, etc. The resulting dynamic system is known to be chaotic, see [Kida et al., 1989]. For that reason, our goal is to learn a system of differential equations that adequately models the dynamics of the temperature time series if the only observable variable is temperature, that is, if pressure, wind velocity, etc, are hidden variables.

To accommodate the possibility of hidden variables we note that, for a large class of dynamical systems, it is possible to reconstruct a trajectory (equivalent to the original one) given only one of the model variables, using its higher order derivatives ([Takens, 1981]).

Related to our approach is the Generalized Polynomial Modeling method (known as GPoMo), that addresses the recovery problem via a combinatorial search among a predefined set of polynomial functions of the observable variables [Mangiarotti et al., 2012] and [Bongard and Lipson, 2007a]. The GPoMo method proceeds by choosing iteratively a family of combination of terms that minimize the Akaike or the Bayesian information criterion. Finally, it returns the set of best models. The authors also discuss the ability of their algorithm to find equations able to capture the dynamics when only some variables are observed. However, we will show that unfortunately this approach does not scale to large problems. Other approaches for learning dynamical systems from data available in the literature,

such as those based on symbolic regression ([Bongard and Lipson, 2007b]), also have the drawback of being too computationally expensive.

The article is organized as follows. In Section 2 we review the GPoMo and SINDy methods and afterwards, describe our methodology (named L-ODEfind) in detail. Section 3 presents the results of our experiments. In particular, in Section 3.2, we compare the performance of GPoMo and L-ODEfind in recovering differential equations using empirical data in the case in which all relevant variables are observed. In Section 3.3 we compare them in the harder case in which at least one relevant variable driving the dynamical system is latent. We apply our proposed method to real world temperature times series in Section 3.4. Finally, in Section 4 we discuss future work and possible extensions.

2. Methods

2.1. GPoMo: Differential equations recovery as a combinatorial search problem

GPoMo is a method proposed and implemented by [Mangiarotti et al., 2012] that addresses the differential equations recovery problem via a combinatorial search in the space of differential equations that can be expressed as polynomial functions of the observed variables. The method uses a genetic algorithm in which at each step new test models are generated by randomly choosing some polynomial terms to be included in the equations. This choice is made by making small variations (take or add a few terms) over the best previously seen models. Then, to select the winning models at each step, they are integrated and compared to the original data.

This combination of combinatorial search and integration steps makes the method slow (as we will see in the experiments). Moreover, except for polynomial combinations of the variables, it does not allow other types of regressors to be included, such as functions of the time variable. To overcome the aforementioned limitations of GPoMo, we based our approach on SINDy, which uses sparse regression to discover governing physical equations from measurement data. We briefly review the SINDy method in the following section.

2.2. SINDy: Differential equations recovery as a linear regression problem

In our work we focus on ordinary differential equations even though SINDy can also tackle partial differential equations. In particular, consider a dynamical system represented by functions $f_1(t) \dots f_H(t)$ satisfying a set of differential equations of the form

$$\mathbb{D}\mathbf{f} = \mathbf{U}(\mathbf{f}, \mathbb{E}\mathbf{f}), \tag{1}$$

where $\mathbf{f} = (f_1, \dots, f_H)$, \mathbb{D}, \mathbb{E} are differential operators in the temporal variables (t) and $\mathbf{U} : \mathbb{R}^J \rightarrow \mathbb{R}^H$ is an unknown map. Suppose we have a series of T equally spaced in time measurements, that is, we

observe $f_h(t_i)$ $i \in \{1, \dots, T\}, h \in \{1, \dots, H\}$. An example of a dynamical system we will study in this paper is the classical Rossler system (2) ([Rossler, 1976]). This system was originally designed to have similar properties and be simpler than the Lorenz system ([Lorenz, 1963]) which was, at the same time, a simplified model for atmospheric convection. The system is given by

$$\begin{aligned}\frac{df_1(t)}{dt} &= -f_2(t) - f_3(t), \\ \frac{df_2(t)}{dt} &= f_1(t) + \alpha f_2(t), \\ \frac{df_3(t)}{dt} &= \beta + f_3(f_1(t) - \gamma),\end{aligned}\tag{2}$$

for constants α, β, γ . This system can be written in the form (1) by taking $\mathbf{f} = (f_1, f_2, f_3)$, $\mathbb{D} = (d/dt, d/dt, d/dt)$ and $\mathbf{U} = (U_1, U_2, U_3)$ where $U_1(v_1, v_2, v_3) = -v_2 - v_3$, $U_2(v_1, v_2, v_3) = v_1 + \alpha v_2$ and $U_3(v_1, v_2, v_3) = \beta + v_3(v_1 - \gamma)$. For certain values of the parameters α, β, ρ , the system is known to have chaotic solutions [Amaral et al., 2006, Ibrahim et al., 2018].

Suppose now that we have access to a particular time series that was generated by this system. Our objective is to find some system of differential equations that can explain the behaviour of the measurements. The SINDy method works by choosing a large dictionary of functions and regressing discretisations of $\frac{\partial f_1}{\partial t}, \dots, \frac{\partial f_H}{\partial t}$ on the dictionary. The dictionary in question can be formed, for example, by collecting polynomial powers of f_h , $h = 1, \dots, H$, spatial derivatives of f_1, \dots, f_H and trigonometric functions of t . A concrete simple example of such a dictionary in the case in which $H = 1$ is the following:

$$\mathcal{A} = \{t, t^2, \sin(t), f_1, f_1^2, t f_1\}.$$

Of course in practice all derivatives are replaced by the corresponding finite differences taken from the measurements represented in f . Having chosen a dictionary, we let $\mathcal{A} = (A_1, \dots, A_p)$ be the vector collecting all members of the dictionary.

Using the observations of the dynamical system, a regression model can be fitted to find the combination of the elements of the dictionary of functions that adequately explains the behavior of $\frac{\partial f_1}{\partial t}, \dots, \frac{\partial f_H}{\partial t}$. That is, we look for a vector of regression coefficients $\mathbf{c} = (c_1, \dots, c_p)$ such that for all $h = 1, \dots, H$

$$\frac{\partial f_h(t)}{\partial t} \approx \sum_{i=1}^p c_{i,h} A_i(t).\tag{3}$$

The regression model has to be learned using the available data. This regression problem could be solved in principle using least-squares. However, the ordinary least-squares regression estimator is ill-defined in cases in which the number of predictor variables p is larger than the number of observations. Since the analyst is usually uncertain about the number of elements in the dictionary needed to adequately model the system of interest, the method used to solve the regression problem

at hand should allow for large number of predictor variables (possibly larger than the number of observations) and automatically estimate sparse models, that is, generate accurate models that only use a relatively small fraction of predictor variables. The Lasso regression technique is perfectly suited for this task. The Lasso is an ℓ_1 -regularised least-squares regression estimator, defined as follows. For $h = 1, \dots, H$ such that f_h is observable we let

$$\mathbf{c}_h^* = \arg \min_{\mathbf{c}_h \in \mathbb{R}^p} \sum_t \left(\frac{\partial^n f_h}{\partial t^n}(t_i) - \sum_{i=1}^p c_{i,h} A_i(t_i) \right)^2 + \lambda \|\mathbf{c}_h\|_1, \quad (4)$$

where $\lambda > 0$ is a tuning constant, measuring the amount of regularization. It can be shown ([Hastie et al., 2015]) that the ℓ_1 penalty encourages sparse solutions and that the larger $\lambda > 0$ the sparser the solution vector \mathbf{c}_h^* will be. In practice, λ is usually chosen by cross-validation. Note that any other sparse regression technique could have been used to estimate the coefficients. We prefer the Lasso due to its simplicity and the wide availability of efficient algorithms to compute it. See for example [Friedman et al., 2007].

The main assumption behind this methodology is that the dynamical system that generated the data at hand can, in reality, be at least approximated using a sparse model, that is, that the vectors \mathbf{c}_h^* in (4) are either exactly or approximately sparse. This hypothesis is known to hold for several dynamical systems of interest in different fields of science. See [Quade et al., 2018]. If the hypothesis holds, we can expect the Lasso estimates to select only a few elements of the dictionary, namely, those that do a good job at explaining variations in the response variable ([Hastie et al., 2015]).

In [Brunton et al., 2016], [Rudy et al., 2017] and [Quade et al., 2018] the authors propose to use an ad-hoc linear regression estimator based on iteratively thresholding the least-squares estimator and applied this method only to first order systems. Through extensive numerical experiments, they show that this methodology is able to learn systems of partial differential equations that adequately model the dynamical system that generated the data. Unfortunately, if some variables are latent, that is, if one is unable to measure at least one of f_1, \dots, f_H , the approach described above cannot be used directly. Next, we describe a way of extending this methodology to deal with the case in which some variables are latent.

2.3. Our proposal: *L-ODEfind*

To accommodate the possibility of latent variables we note that, for a large class of dynamical systems, it is possible to reconstruct a trajectory (equivalent to the original one) given only one of the model variables, using its higher order derivatives ([Takens, 1981]). Moreover, we recall that in the case of a linear system of n ordinary differential equations there is an equivalence between this multidimensional system and a single differential equation of order n , which we can interpret as

latently including the information of the other $n - 1$ unobserved variables. [Mangiarotti et al., 2012] also makes use of higher order time derivatives to deal with unobserved variables.

Based on this ideas, we propose to augment the methodology developed in [Brunton et al., 2016], [Rudy et al., 2017] and [Quade et al., 2018] by choosing the target variable to be a higher-order time derivative, to tackle situations in which not all relevant variables are observed. We estimate the coefficients of the dynamical system using the Lasso estimator (4). As mentioned earlier, we chose to use the Lasso due to its simplicity, the abundance of theoretical guarantees on its performance ([Hastie et al., 2015]) and the availability of efficient algorithms to solve the convex optimization problem that defines the estimator. The choice of the tuning constant λ in (4) is done by 10-fold cross validation, using the LassoCV method from sklearn [Pedregosa et al., 2011] with a maximum number of steps equal to 10000, and 100 candidate λ s. After estimating the regression coefficients, we build a forecasting method by integrating the retrieved differential equation. We call this method Latent ODE find (L-ODEfind).

For instance, suppose we have observed a single time series f ($H = 1$). If we choose as a target variable the third time derivative and we use polynomial combinations up to degree 2 of the series and derivatives up to order 2 as regressors, equation (3) becomes

$$\frac{\partial^3 f(t)}{\partial t^3} \approx \sum_{i=1}^p c_i A_i(t), \quad (5)$$

and

$$\mathcal{A} = \left\{ 1, f, \frac{\partial f(t)}{\partial t}, \frac{\partial^2 f(t)}{\partial t^2}, f^2, \left(\frac{\partial f(t)}{\partial t} \right)^2, \left(\frac{\partial^2 f(t)}{\partial t^2} \right)^2, f \frac{\partial f(t)}{\partial t}, f \frac{\partial^2 f(t)}{\partial t^2}, \frac{\partial f(t)}{\partial t} \frac{\partial^2 f(t)}{\partial t^2} \right\}.$$

This regression model is then fitted using the Lasso (4), as described earlier.

2.4. Evaluation with predictions

Evaluating the equations found by these methods can be challenging. The evaluation of a method that aims at recovering the differential equation behind the observed data needs to be done in different ways depending on the information available.

In the case of a simulation where the coefficients of the differential equation are known and the variables are fully observed, the adjusted coefficients and the true ones can be compared using the mean squared error. If not all variables are observed, the adjusted coefficients refer to a different differential equation that in most cases cannot be obtained analytically (exceptions are, for instance, linear ODEs). In this case, the coefficients comparison cannot be done and another way of evaluating the method is needed. The same happens when working with real world data where the differential equation behind is not known. In order to evaluate a method in this context, we propose to integrate

the fitted differential equation to make predictions for different time horizons and compare them to the observed data by using the symmetric mean absolute percentage error (SMAPE). Using a time horizon of n , the SMAPE is defined as:

$$SMAPE(A, F) = \frac{1}{n} \sum_{t=1}^n \frac{|F_t - A_t|}{(|A_t| + |F_t|)/2},$$

where A_t is the real value and F_t the forecasted value. If the predictions made by method A are better (lower values of SMAPE) than the ones made by method B, the differential equation found by the method A is deemed better than the one found by method B.

2.5. Implementation details

All the experiments in this paper were performed using Python 3.8, except for GPoMo , which was performed using R 3.6 [R Core Team, 2020] and the GPoM package [Mangiarotti et al., 2020]. For L-ODEfind we use our own Python implementation (available in <https://github.com/agussomacal/L-ODEfind>). The integration of the differential equations was done using odeint from python Scipy library [Virtanen et al., 2020].

3. Results

In this section we first compare L-ODEfind with GPoMo for the problem of learning an ordinary differential equation with no hidden variables. Then, we compare the performance of these methods in simulated systems with hidden variables and finally in temperatures series provided by RTE.

3.1. Names abbreviation

The names abbreviation used in this section can be found in Table 1. Target time derivative is the degree of the fitted differential equation and poly degree is the maximum degree of the polynomial combinations of the derivatives used as regressors. An example of the regression problem with target derivative 3 and poly degree 2 can be found in equation (5). Notice that the number in the model name refers to the degree of the differential equation.

| Method | Parameters | Name |
|-----------|--------------------------------------|------------|
| L-ODEfind | Target derivative: 1, poly degree: 3 | L-odefind1 |
| L-ODEfind | Target derivative: 2, poly degree: 3 | L-odefind2 |
| L-ODEfind | Target derivative: 3, poly degree: 3 | L-odefind3 |
| GPoMo | Target derivative: 2, poly degree: 3 | GPoMo2 |
| GPoMo | Target derivative: 3, poly degree: 3 | GPoMo3 |

Table 1: Models names abbreviations used in the graphics above.

3.2. Simulated data with fully observed variables

We compare the performance of L-ODEfind with that of GPoMo for the task of learning an ordinary differential equation with no hidden variables. Note that since in this case there are no hidden variables, our method coincides with SINDy. The goal of this comparison is to highlight the fact that, because L-ODEfind solves a continuous optimization problem and GPoMo approximately solves a combinatorial optimization problem, L-ODEfind can be orders of magnitude faster than GPoMo.

We generated data using the Lorenz attractor equations:

$$\begin{aligned}\frac{dx}{dt} &= \sigma(y - x), \\ \frac{dy}{dt} &= x(\rho - z) - y, \\ \frac{dz}{dt} &= xy - \beta z,\end{aligned}\tag{6}$$

using the coefficients $\sigma = 10$, $\rho = 28$ and $\beta = \frac{8}{3}$, for twenty different random starting conditions, with distribution $N(0, 1)$. The differential equation was integrated using a discretization step of 0.01. Then, we fitted the resulting datasets using GPoMo and L-ODEfind. For GPoMo we set a maximum number of steps of 10240. We compare the true coefficients of the differential equations with the ones found by each model using as a metric the mean square error (MSE). All the coefficients were taken into account for computing the MSE, including the ones that are supposed to be zero. We average the MSE corresponding to different random starting conditions and report this as our goodness of fit metric.

Figure 1 shows that L-ODEfind method is nearly two orders of magnitude more accurate than the best GPoMo case. On the other hand, L-ODEfind is also orders of magnitude faster, taking less than 10 seconds to compute accurate approximations of the true coefficients while GPoMo takes hours.




Figure 1: Comparison between L-ODEfind and GPoMo when finding the coefficients of data generated by the Lorenz attractor differential equations. We plot the average MSE across different random initial conditions for the system, and the time (in seconds) needed by the algorithm to fit the model.

3.3. Simulated data with hidden variables

In this section we use three ODE systems, an oscillator, and the Rossler and Lorenz systems, as examples to evaluate and compare, using the methodology explained in 2.4, the accuracy of several L-ODEfind and GPoMo models. For a given differential equation, we generated time series of length 5000 points and integration step of $dt = 0.01$ for 20 different initial conditions. Following the methodology described in 2.4 we fit each series with GPoMo or L-ODEfind and then integrate the equation found in each case to predict the values of the series for several time horizons and compute the corresponding SMAPE. The number of maximum iterations for GPoMo is set to 5120.

3.3.1. Example 1: Oscillator

We start by considering an oscillator, which is a first order linear system of two variables. An equation of order two involving only one of the variables can be derived. So, in this case, the problem of having hidden variables can be effectively tackled by choosing a higher order time derivative (order two) as the target.


We used the two variables oscillator equation which can be written in general form as:

$$\frac{d}{dt} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot \begin{pmatrix} x \\ y \end{pmatrix} \quad (7)$$

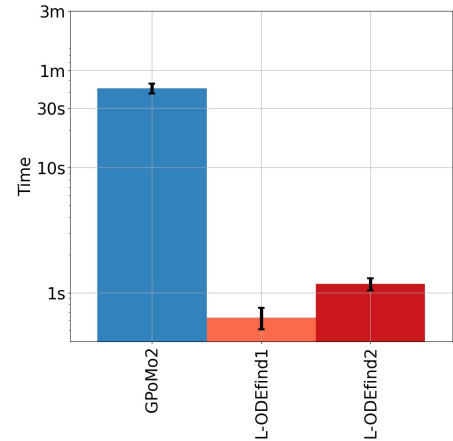
where x and y are the variables and a, b, c, d are the coefficients of the linear equation that links the variable with its derivatives. If we only had access to the variable x , this system could be rewritten in a second order differential equation taking the form:

$$\frac{d^2x}{dt^2} = (a+d)\frac{dx}{dt} - (ad-bc)x = \beta\frac{dx}{dt} + \alpha x. \quad (8)$$

In our experiments we set $a = 0.1, b = -1, c = 1, d = 0$ ($\alpha = -1, \beta = 0.1$) and we only observed the variable x so the corresponding second order equation derived as in 8 is $0.1\frac{dx}{dt} - x$. We can see in Figure 2 (a) that using the second derivative in time as target gives the lowest prediction error for both GPoMo and L-ODEfind, as expected from equation (8). We also see that the model found by L-ODEfind manages to approach the true model much better than GPoMo as it maintains an SMAPE below 0.02 in horizons were GPoMo has already arrived to 0.1. Moreover, when looking at the coefficients found by GPoMo ($\alpha = -0.99506 \pm 0.00372$ and $\beta = 0 \pm 0$) and L-ODEfind ($\alpha = -0.99965 \pm 0.00030$ and $\beta = 0.09963 \pm 0.00013$) are within a small tolerance the expected from equation (8). Finally, when looking at the fitting time (Fig. 2 (b)), we find that L-ODEfind is around 50 times faster than GPoMo.



(a) Mean SMAPE versus time horizon.



(b) Fitting time.

Figure 2: Prediction accuracy and fitting time for L-ODEfind and GPoMo models when data comes from the oscillator system with x as observed variable (y hidden).

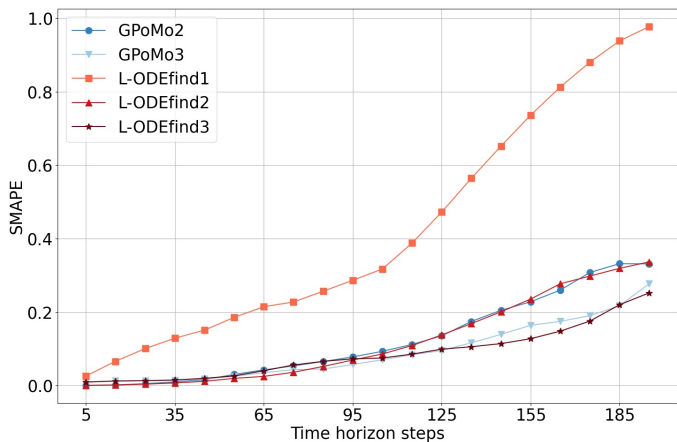
3.3.2. Example 2: Rossler

Next, we consider a more complicated case, the Rossler system, which is a non-linear (quadratic) system:


$$\begin{aligned}\frac{dx}{dt} &= -x - z, \\ \frac{dy}{dt} &= x + ay, \\ \frac{dz}{dt} &= b + z(x - c).\end{aligned}\tag{9}$$

If only variable y is observed (x and z hidden) an equation of order 3 and polynomial degree 2 can be deduced for y [Letellier et al., 2005].

In our experiments, we set $a = 0.52, b = 2, c = 4$ and only observe y (x and z latent). L-ODEfind is consistently faster than GPoMo although, in this example, GPoMo3 and L-ODEfind3 have almost the same prediction SMAPE. Notice that the SMAPE is less than 0.2 for all models with target time derivative 2 or 3 up to 125 time horizon steps.



(a) Mean SMAPE versus time horizon.



(b) Fitting time.

Figure 3: Prediction accuracy and fitting time for L-ODEfind and GPoMo models when data comes from the Rossler system with y as observed variable (x and z hidden).

Therefore, when considering systems where an analytical solution can be deduced, such as the oscillator (x observed, y hidden) and Rossler system (y observed, x and z hidden), both methods perform very well in terms of SMAPE prediction.

Next, we consider the Rossler system in the case of variable x observed (y and z hidden). In this scenario, there is no analytical solution using only polynomials [Letellier et al., 2005]. Interestingly, we can see that using higher order time derivatives as target, both GPoMo and L-ODEfind find a

differential equation that is able to provide predictions whose accuracy is comparable (for short-term horizons) to the previous case (y observed) where there was an analytical solution (Figure 4). The fitting times continue to show that L-ODEfind is consistently faster than GPoMo.




Figure 4: Prediction accuracy and fitting time for L-ODEfind and GPoMo models when data comes from the Rossler system with x as observed variable (y and z hidden).

3.3.3. Example 3: Lorenz attractor

Here we consider the Lorenz system [Lorenz, 1963], discussed in Section 3.2, where no differential equation using only polynomials can be derived for x as the only observed variable.

Following the same methodology as before, we tried different target derivatives for both GPoMo and L-ODEfind to fit observed variable x (y and z served). We found that using higher order time derivatives (in particular second order) helps to find models that can approximate better the observed time series when integrated, although the prediction accuracy suffers from the added complexity of the problem (Fig. 5). As a consequence, the accuracy degrades faster reaching an $SMAPE = 0.5$ as soon as 35 time steps while in the previous cases this was attained around 185 steps. For a narrow difference again L-ODEfind outperformed GPoMo while also keeping fitting times 4 to 40 times faster than GPoMo.

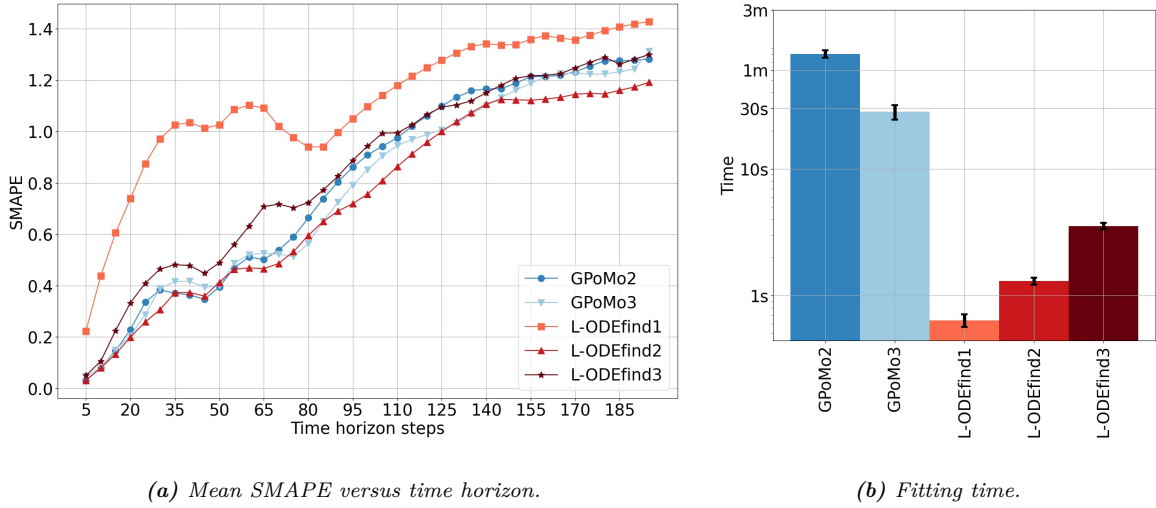


Figure 5: Prediction accuracy and fitting time for L-ODEfind and GPoMo models when data comes from the Lorenz system with x as observed variable (y and z hidden).

3.4. Temperature series provided by RTE

In this section, we fit different models to temperatures series provided by the Réseau de Transport d’Électricité (RTE). The data consist of 200 hourly measured temperatures time series. These time series correspond to temperatures in Paris along a year for 200 different possible years or scenarios. The time series are not measured temperatures nor the output of a simulation, but rather the result of a reanalysis process. Also, some relevant variables for modeling the atmospheric system are not available to us, for instance, wind and pressure. Our temperature time-series are not historical measurements but can be thought as a possible realization of the temperature in Paris. These temperature time series have $365 \times 24 = 8760$ time points.


In order to evaluate the methods, we select 39 out of the 200 temperature time series (due to GPoMo’s computation time), fit the different models and use them to predict for short time horizons. The fit was done with the first 8560 time points and the prediction was evaluated with the following time points (up to a time horizon of 15 time points).

In tackling this complex problem we want to analyze the performance of **L-ODEfind** and **GPoMo** in the forecasting task and compare their behavior to classical forecasting methods **naive predictor** (for every time horizon, predicts the average of the last 24 time points) and **exponential smoothing (ES)** (triple exponential smoothing with an additive seasonal component⁶).

In Figure 6 the SMAPE of the prediction for different time horizons in hours is displayed. L-

⁶We used the implementation available in the Python library sktime [Löning et al.]

ODEfind2 gives the lowest SMAPE for time horizons lower or equal to 6 hours, whereas GPoMo3 is better for time horizons greater than 6 hours. Notice that for time horizons greater than 6 hours, the naive predictor performs better than exponential smoothing, giving a rough idea of the reasonable predictability horizon that forecasting methods can give. In any case, both L-ODEfind and GPoMo give better forecasts than both ES and naive when the target derivative is 3. As can be seen in figure 6 (b), the fitting times for GPoMo are much greater than L-ODEfind, whereas L-ODEfind and exponential smoothing have similar fitting times.



(a) SMAPE vs time horizons in hours for different predictive methods. (b) Time to fit the different models in log scale.

Figure 6: SMAPE and fitting times for different models applied on a real world data set of temperatures.

4. Discussion

In this paper we addressed the problem of recovering differential equations from data where not all variables are observed by enhancing the approach outlined in [Brunton et al., 2016]. We tested this approach in simple and complex ODE systems and consistently found that the proposed approach of using time derivatives of higher order as target regressing variables allows to find models whose future predictions are more reliable than only using first order derivatives. We also compared our method to GPoMo and found that: (i) our proposal is orders of magnitude faster than GPoMo, and (ii) our proposal learns models with comparable or even higher prediction accuracy for several dynamical systems. Finally, we faced the challenge of addressing a real world problem and found that both L-ODEfind and GPoMo used as forecasting methods gave comparable results while at the same time outperforming classical forecasting methods, exponential smoothing and naive. In summary, L-ODEfind proved to be an accurate and fast method for recovering ordinary differential equations from data with hidden variables.

5. Acknowledgements

Funding: This work was supported by Aristas S. R. L. and RTE (Réseau de Transport d'Électricité).

References

- G. Amaral, C. Letellier, and L. Aguirre. Piecewise affine models of chaotic attractors: The rossler and lorenz systems. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 16, 2006.
- J. Bongard and H. Lipson. Automated reverse engineering of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 104(24):9943–9948, 2007a.
- J. Bongard and H. Lipson. Automated reverse engineering of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 104:9943 – 9948, 2007b.
- S. L. Brunton, J. L. Proctor, and J. N. Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016.
- J. Friedman, T. Hastie, H. Höfling, and R. Tibshirani. Pathwise coordinate optimization. *Ann. Appl. Stat.*, 1(2):302–332, 12 2007.
- T. Hastie, R. Tibshirani, and M. Wainwright. *Statistical learning with sparsity: the lasso and generalizations*. Chapman and Hall/CRC, 2015.
- K. M. Ibrahim, R. K. Jamal, and F. H. Ali. Chaotic behaviour of the rossler model and its analysis by using bifurcations of limit cycles and chaotic attractors. *Journal of Physics: Conference Series*, 1003:012099, may 2018. doi: 10.1088/1742-6596/1003/1/012099. URL <https://doi.org/10.1088/1742-6596/1003/1/012099>.
- S. Kida, M. Yamada, and K. Ohkitani. Route to Chaos in a Navier-Stokes Flow. In M. Mimura and T. Nishida, editors, *Recent Topics in Nonlinear PDE IV*, volume 160 of *North-Holland Mathematics Studies*, pages 31–47. North-Holland, 1989.
- C. Letellier, L. A. Aguirre, and J. Maquet. Relation between observability and differential embeddings for nonlinear dynamics. *Physical Review*, 2005.
- M. Löning, A. Bagnall, S. Ganesh, V. Kazakov, J. Lines, and F. J. Király. sktime: A Unified Interface for Machine Learning with Time Series. In *Workshop on Systems for ML at NeurIPS 2019*.
- E. N. Lorenz. Deterministic nonperiodic flow. *Journal of the atmospheric sciences*, 20(2):130–141, 1963.

- S. Mangiarotti, R. Coudret, L. Drapeau, and L. Jarlan. Polynomial search and global modeling: Two algorithms for modeling chaos. 86(4), Oct. 2012.
- S. Mangiarotti, M. Huc, F. L. Jean, M. Chassan, L. Drapeau, I. de Recherche pour le Développement, and C. N. de la Recherche Scientifique and. *GPoM: Generalized Polynomial Modelling*, 2020. R package version 1.3.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- M. Quade, M. Abel, J. Nathan Kutz, and S. L. Brunton. Sparse identification of nonlinear dynamics for rapid model recovery. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 28(6):063116, 2018.
- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2020. URL <https://www.R-project.org/>.
- O. Rossler. Chaotic Behavior in Simple Reaction Systems. *Zeitschrift Naturforschung Teil A*, 31:259, 04 1976.
- S. H. Rudy, S. L. Brunton, J. L. Proctor, and J. N. Kutz. Data-driven discovery of partial differential equations. *Science Advances*, 3(4), 2017.
- F. Takens. Detecting strange attractors in turbulence. *Lecture Notes in Mathematics, Berlin Springer Verlag*, 898:366, 1981.
- P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17: 261–272, 2020. doi: 10.1038/s41592-019-0686-2.