HOW SLOW IS SLOW? SFA DETECTS SIGNALS SLOWER THAN THE DRIVING FORCE

Wolfgang Konen and Patrick Koch
Institute for Informatics, Cologne University of Applied Sciences
Steinmüllerallee 1, D-51643 Gummersbach, Germany
{wolfgang.konen, patrick.koch}@fh-koeln.de

Abstract Slow feature analysis (SFA) is a bioinspired method for extracting slowly varying driving forces from quickly varying nonstationary time series. We show here that it is possible for SFA to detect a component which is even slower than the driving force itself (e.g. the envelope of a modulated sine wave). It depends on circumstances like the embedding dimension, the time series predictability, or the base frequency, whether the driving force itself or a slower subcomponent is detected. Interestingly, we observe a swift phase transition from one regime to another and it is the objective of this work to quantify the influence of various parameters on this phase transition. We conclude that *what* is perceived as slow by SFA varies and that a more or less fast switching from one regime to another occurs, perhaps showing some similarity to human perception.

Keywords: Driving force, phase transition, nonstationary time series, slow feature analysis

1. Introduction

The analysis of nonstationary time series plays an important role in the data understanding of various phenomena such as temperature drift in an experimental setup, global warming in climate data, or varying heart rate in cardiology. Such nonstationarities can be modeled by underlying parameters, referred to as driving forces, that change the dynamics of the system smoothly on a slow time scale or abruptly but rarely, e.g. if the dynamics switches between different discrete states [11].

Often, e.g. in EEG-analysis or in monitoring of complex chemical or electrical power plants, one is particularly interested in revealing the driving forces themselves from the raw observed time series since they show interesting aspects of the underlying dynamics, for example the switching between different dynamic regimes.

Several methods for detecting and visualizing driving forces have been developed; based on recurrence plots [2], feedforward ANNs with extra input unit [9] or, as Wiskott [11] recently proposed, by Slow Feature Analysis (SFA), a versatile, robust, and fast algorithm. SFA has been originally presented in context of a bioinspired model for unsupervised learning of invariances in the visual system of vertebrates [10] and is described in detail in [11, 12]. SFA works fully unsupervised, just by searching nonlinear combinations of the input signals which vary as slowly as possible in time.

What is 'slow' in the driving forces compared to the raw observed time series? Often it might be the case that a driving force contains components on different time scales and it is crucial to understand which time scale will be selected by the driving force algorithm. As an example we consider driving forces made up of two overlayed frequencies $f_1 < f_2$. Will the driving force detection algorithm detect the slower one of the frequencies, f_1, thus being more slow, or the combined driving force made up of f_1 and f_2, thus being more accurate? With this paper we try to deepen our understanding which parameters influence whether the first or the second choice is taken.

2. Slow Feature Analysis

We briefly review here the SFA approach described in [11]. The general objective of SFA is to extract slowly varying features from a quickly varying multidimensional signal. For a scalar output signal and an N-dimensional input signal $\boldsymbol{x} = \boldsymbol{x}(t)$ where t indicates time and $\boldsymbol{x} = [x_1, ..., x_N]^T$ is a vector, the question can be formalized as follows: Find the input-output function $g(\boldsymbol{x})$ that generates a scalar output signal

$$y(t) := g(\boldsymbol{x}(t)) \tag{1}$$

with its temporal variation as slowly as possible, measured by the variance of the time derivative:

$$\text{minimize} \quad \Delta(y) = \langle \dot{y}^2 \rangle \tag{2}$$

with $\langle \cdot \rangle$ indicating the temporal mean. Wiskott and Sejnowski [12] propose a closely related slowness indicator η proportional to $\sqrt{\Delta(y)}$. Low η-values indicate slow signals, high η-values fast signals.

To avoid the trivial constant solution, the output signal has to meet the following constraints:

$$\langle y \rangle \quad = \quad 0 \quad \text{(zero mean)}\,, \tag{3}$$

$$\langle y^2 \rangle \quad = \quad 1 \quad \text{(unit variance)}\,. \tag{4}$$

This is an optimization problem of variational calculus and as such difficult to solve. But if we constrain the input-output function to be a linear combination of some fixed and possibly nonlinear basis functions, the problem becomes tractable with the mathematical details given in [11]. A typical choice for the nonlinear basis functions are monomials of degree 2, but other choices, e. g. monomials of higher degree or radial basis functions could be used as well. Basically, SFA searches the eigenvector in the expanded space with the smallest eigenvalue and projects the expanded signal onto this eigenvector to obtain the output signal, which we denote here by y or y_1.

3. Experiments

In the following we present examples with time series $w(t)$ derived from the well-known logistic map [7, 11] to illustrate the properties of SFA. The underlying driving force is always denoted by γ and may vary between -1 and 1 smoothly and considerably slower (as defined by the variance of its time derivative (2)) than the time series $w(t)$. The approach follows closely the work of Wiskott [11] but with more systematic variations in the driving force.

We consider here a driving force that is made up of two frequency components

$$\gamma(t) = \frac{1}{2}(\underbrace{\sin(0.0005\nu_f t)}_{=\gamma_S(t)} + \underbrace{\sin(0.0047\nu_f t)}_{=\gamma_F(t)}) \quad \in \quad [-1, 1], \tag{5}$$

where the first component γ_S is roughly ten times slower than γ_F. The question is whether SFA as the driving force detector detects solely the slower component γ_S of the driving force (in an attempt to minimize η) or the full driving force γ (in an attempt to extract the underlying system dynamics as accurately as possible). A second question is whether a phase transition between the two choices might occur as we vary the base frequency ν_f.

In order to visually inspect the agreement between a slow SFA-signal and the driving force γ we must bring the SFA-signal into alignment with γ (since the scale and offset of the slow signal $y(t)$ formed by SFA is fixed by the constraints and the sign is arbitrary). Therefore we define a γ-*aligned signal*

$$A_\gamma(y(t)) = ay(t) + b \tag{6}$$

where the free parameters a and b are chosen in such a way that the signal $A_\gamma(y(t))$ is in best possible alignment with $\gamma(t)$.

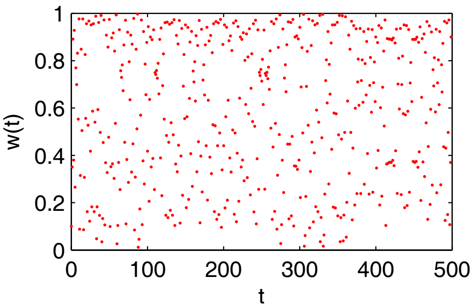

Fig. 1. Time series $w(t)$ derived from the logistic map with driving force according to Eq. (7) for $\nu_f = 20$ and $q = 0.1$. For all $q < 0.33$, no structure from the driving force is directly visible in the map.

The following simulations are based on 6000 data points each and were done with MATLAB 7.0.1 using the SFA toolkit `sfa-tk` [1].

3.1 Logistic map in chaotic and in predictable regime

We consider a time series derived from a logistic map

$$w(t+1) = (4.0 - q + 0.1\gamma(t))w(t)(1 - w(t)), \quad (7)$$

which maps the interval $[0, 1]$ onto itself and has the shape of an upside-down parabola crossing the abscissa at 0 and 1. The logistic map exhibits an interesting and complex dynamic behaviour, since its parameter $q \in [0.1, 3.9]$ controls different forms of predictabilty: For $q < 0.33$ the map is fully in its chaotic regime (a map with no visible structure, see Fig. 1), for $0.33 < q < 0.53$ we have a mixture of chaotic and predictable periods and for $0.53 < q < 3.9$ it is long-term predictable.

To allow SFA to reconstruct the driving force, it is necessary to generate from the scalar $w(t)$ a time series of embedding vectors $\boldsymbol{x}(t)$ as input to SFA. The embedding vector at time point t is defined as

$$\boldsymbol{x}(t) := [w(t - s_\tau),\, w(t - (s_\tau - 1)),\, ..., w(t + s_\tau)]^T \quad (8)$$

with delay τ, odd dimension m and $s_\tau := \tau(m-1)/2$. Centering the embedding vectors results in an optimal temporal alignment between estimated and true driving force.

Fig. 2 shows the estimated driving force (from SFA with $m = 19$, $q = 0.1$, $\tau = 1$ and second order monomials) and the true driving force. At the higher frequency $\nu_f = 60$ the estimated driving force is in alignment with the slower component $\gamma_S(t)$. This is remarkable since the slower component is not directly visible in the driving force, only indirectly as envelope of the solid curve. Quite clearly there is a phase transition occuring around $\nu_f = 40$.

In Fig. 3 we vary the base frequency $\nu_f \in [4, 80]$ and we see a swift phase transition. The transition frequency $\nu(P.T.)$ is the crossover point of the two correlation curves, shown in Fig. 3 as black dot. For small $q = 0.1$ (fully chaotic w; left part of Fig. 3) a phase transition occurs at $\nu(P.T.) = 34$ (black dot) while for larger $q = 0.4$ (mix of chaotic and non-chaotic periods in w; right part of Fig. 3) the phase transition happens earlier and occurs swifter at $\nu(P.T.) = 17$.

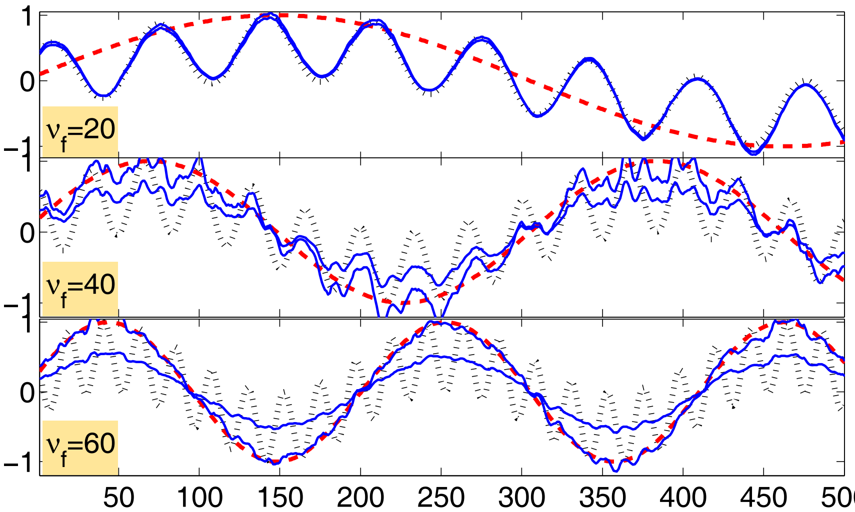

Figure 2. *SFA outputs $y_1(t)$ (solid lines) aligned to the driving forces (see Eq. (6)) for base frequencies $\nu_f = 20, 40, 60$ clearly show a phase transition from the complete driving force $\gamma(t)$ (dotted line) to its slower subcomponent $\gamma_S(t)$ (dashed line). We see two solid curves since we align the slowest SFA signal once with $\gamma(t)$ and once with $\gamma_S(t)$. For clarity only the first 500 time steps out of 6000 are shown.*

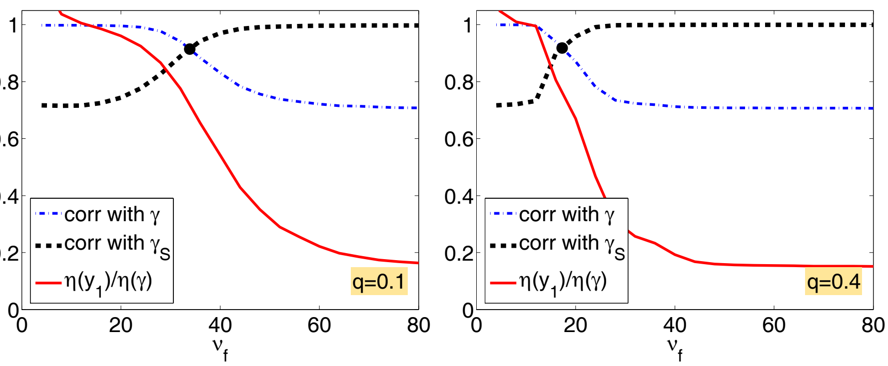

Figure 3. *We show the correlation of the SFA-output y_1 with the driving force γ (dash-dotted line) and with its slow component γ_S (thick dashed line). The black dot indicates the phase transition at $\nu(P.T.)$. The slowness quotient $\eta(y_1)/\eta(\gamma)$ (solid line) drops largely near the phase transition. Left: $q = 0.1$, phase transition at $\nu(P.T.) = 34$. Right: $q = 0.4$, phase transition at $\nu(P.T.) = 17$.*

3.2 The phase transition as a function of q and m

How does the phase transition frequency $\nu(P.T.)$ vary as a function of the predictability q and the embedding dimension m of the SFA-input signal? Both parameters are varied systematically over a broad range and the results are depicted in Fig. 4. First of all it is interesting to note that the SFA algorithm, being basically parameter-free, works very well over this broadly varying input material, which makes SFA a robust and versatile algorithm.

A second remark is necessary concerning the SFA implementation `sfa-tk` [1]: While it worked well for small embedding dimensions m, larger m led quite inevitably to numerical instabilities resulting in wrong "slow" signals y_1 which were neither slow nor did they respect the unit variance condition $\langle y^2 \rangle = 1$. We presented in [6] a slightly modified implementation (closer along the lines of [12]) and based on SVD which successfully avoids these numeric instabilities. This modified implementation is used throughout the experiments in this paper.

4. Discussion

It is important for driving force analysis with SFA to understand the mechanisms by which the slowest signal is selected. If the driving force contains two components of different frequencies, two interesting things might happen: If the base frequency ν_f is large enough then SFA will return the slower component as the slowest signal. This is quite remarkable, since SFA detects a signal with a smaller η than the driving force itself. Recall that this slower component is not directly visible in the driving force, only indirectly as the modulation. But after all, it is also quite understandable: If we view the dynamical system as a two-stage process where the slow component γ_S is considered as a modulating force acting on the other (faster) component γ_F with the output of this stage acting on the dynamical system, then in such a system description, the slower component γ_S becomes directly visible.

Surprisingly, if we lower the base frequency ν_f, we reach the point where the slow component comes "out of sight" and the slowest signal returned by SFA is well-aligned with the driving force itself (slow plus fast component). Why is the slow component alone no longer detected by SFA? We hypothesize that two reasons are responsible for this:

1. If we lower the base frequency ν_f, the fast component γ_F becomes slower and thus contains less information within a given embedding horizon m. This makes the reconstruction of the slow component γ_S more and more noisy. We finally reach the point where for a given embedding dimension m the smoother reconstruction of γ

gets a smaller η (becomes slower) than the noisy reconstruction of γ_S. Increasing m should make the reconstruction of γ_S smoother, thus making γ_S again detectable as the slow component.

2. Another reason might be the chaotic nature of the logistic map. In the chaotic region of the map $w(t)$, noise is amplified and makes the reconstruction of the slow component γ_S noisier until it again comes to the point where the noisy reconstruction has a larger η than the (smoother) reconstruction of γ. If this is true, then moving to a better predictable region of the logistic map (increasing q) should make the slow component again detectable.

Both hypotheses are well-supported by the results shown in Fig. 4. On the left-hand side we see the location of the phase transition. For most input signals which are a function of q and ν_f there seems to be a sufficiently large m so that the slow component becomes detectable. For $q = 0.7$ this occurs already at very low frequencies. The curve for $q = 0.6$ (not shown) is for $m > 10$ very similar to $q = 0.7$, which is well-understandable if we recall that all $q > 0.53$ make the time series long-term predictable, thus even a very slow subcomponent becomes detectable. On the right-hand side of Fig. 4 we see that both methods, increasing m or increasing q, finally lead to a reliable detection of the slow subcomponent as it is claimed by our hypotheses.

Hypothesis 1 is also supported by the following experiment: If we lower the frequency of the slow component γ_S but keep the fast component γ_F the same, then SFA will always reliably detect the slow component γ_S, even if only a quarter of its wave length appears in the time series data. This is because the same γ_F allows a reconstruction of γ_S at always the same smoothness level.

Nonlinear Regression. Hypotheses 1 and 2 are also supported by the following nonlinear regression experiment: For the set of nonlinear basis functions used by SFA (e.g. monomials of degree 2) and for a given output signal (e.g. γ and γ_S) we seek the best reconstruction in the least-square sense. Decreasing m or q leads to more and more noisy reconstructions of γ_S. We find empirically that quite precisely at the same phase transition points as in Fig. 3 the reconstruction of γ_S gets a higher η (becomes less slow) than the reconstruction of γ. This is remarkable since the slowness principle was not used at all in this nonlinear regression experiment.

Connection to human perception. Since SFA has been originally developed as a model for neural information processing [10], it might be natural to ask, whether the observed switch between components and

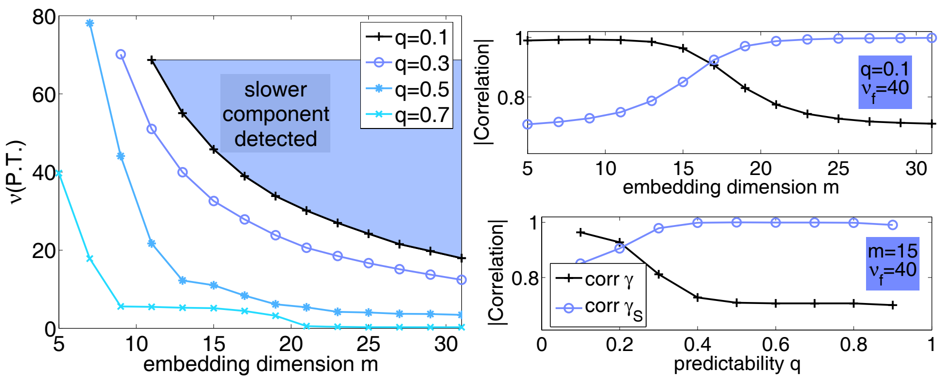

Figure 4. *Left: Phase transition frequency* $\nu(P.T.)$ *as a function of* q *and* m*. Right: Absolute values of correlation at fixed* $\nu_f = 40$ *when varying either* m *or* q*.*

its phase transition has any parallel to human perception and motion coordination. Several phenomena with switching effects are discussed in the literature:

The well-known backward spinning-wheel illusion [8] occurs frequently in movies or under stroboscopic lighting conditions and it shows the transition from a fast forward rotation detection to a slow backward rotation detection. This effect is usually explained by the snapshot-like presentation of the percept which has ambiguous motion interpretations. Somewhat less known is that a similar, although harder to perceive effect can occur under plain sunlight and direct view with the eye [5, 8]. No snapshot-like explanation is possible here, the percept is continous having a greater resemblance to the smoothly varying driving force of our SFA experiments. A possible explanation of the sunlight spinning-wheel illusion is that rivalry between different motion detectors in the brain occurs [5].

Another well-known phase transition occurs in bimanual motion coordination when performing certain movements with the index fingers of both hands [4]. For the observed phenomena there exists a theoretical model, the Haken-Kelso-Bunz model [3], which describes the phase transition and certain hysteresis effects.

SFA has shown similar capabilities in the sense that the same setup can learn to synchronize with different components of a driving force, depending on the experimental conditions. It remains however to be studied, whether *one* trained SFA system can (without further learning) switch between different components when applied to signals with smoothly varying base frequency and whether a hysteresis effect can be observed.

5. Conclusion

In this paper we have investigated the notion of *slowness* in slow feature analysis (SFA). It has been verified that SFA can reliably detect slow driving forces or their subcomponents over a broad range of parameters in nonstationary time series, even in the presence of chaotic motion.

However it has also been seen that what is perceived as *slow* can vary for driving forces made up of components on different time scales. Depending on the embedding dimensions and the predictability of the underlying dynamical system we observe phase transitions where the slowest SFA-signal moves from alignment to a slow subcomponent to alignment with the (faster varying) complete driving force. Notably, when alignment to the slow subcomponent occurs, SFA is capable of detecting slow signals with an η-indicator considerably lower than the η-value of the true driving force. We found that the slow subcomponent is lost precisely in the moment when its reconstruction in the expanded function space used by SFA has more temporal variation than the reconstruction of the complete driving force.

In real world data it is often not possible to vary the base frequency or the degree of nonlinearity in the observed dynamical system systematically. Therefore, one advice from the present study should be to vary the embedding dimension over a broad range in order to detect possible slow signals which otherwise might be hidden. In any case, SFA has shown to be robustly working on a broad range of input data and it is able to reveal subtle components in the driving forces, thus making it a versatile tool for driving force detection.

Acknowledgments

We are grateful to Laurenz Wiskott for helpful discussions on SFA and to Pietro Berkes for providing the MATLAB code for `sfa-tk` [1]. This work has been supported by the Bundesministerium für Bildung und Forschung (BMBF) under the grant SOMA (AIF FKZ 17N1009) and by the Cologne University of Applied Sciences under research focus grant COSA.
A slightly extended preprint version of this article is available in the arXiv.org e-Print archive at http://arxiv.org/abs/0911.4397.

References

[1] P. Berkes. sfa-tk: Slow Feature Analysis Toolkit for Matlab (v.1.0.1). `http://people.brandeis.edu/~berkes/software/sfa-tk`, 2003.

[2] M. C. Casdagli. Recurrence plots revisited. *Physica D: Nonlinear Phenomena*, 108(1-2):12–44, 1997.

[3] H. Haken, J. Kelso, and H. Bunz. A theoretical model of phase transitions in human hand movements. *Biological Cybernetics*, 51:347–356, 1985.

[4] J. Kelso. On the oscillatory basis of movement. *Bulletin of the Psychonomic Society*, 18:63, 1981.

[5] K. Kline, A. Holcombe, and D. Eagleman. Illusory motion reversal is caused by rivalry, not by perceptual snapshots of the visual field. *Vision Research*, 44:2653–2658, 2004.

[6] W. Konen. On the numeric stability of the SFA implementation `sfa-tk`. arXiv.org e-Print archive, http://arxiv.org/abs/0912.1064, Dec. 2009.

[7] R. M. May. Simple mathematical models with very complicated dynamics. *Nature*, 261:459–467, 1976.

[8] D. Purves, J. A. Paydarfar, and T. J. Andrews. The wagon-wheel illusion in movies and reality. *Proceedings of National Academy of Sciences USA*, 93:3693–3697, 1996.

[9] P. F. Verdes, P. M. Granitto, H. D. Navone, and H. A. Ceccatto. Nonstationary time-series analysis: Accurate reconstruction of driving forces. *Phys. Rev. Lett.*, 87(12):124101, Sep 2001.

[10] L. Wiskott. Learning invariance manifolds. In *Proc. of the 5th Joint Symp. on Neural Computation, May 16, San Diego, CA*, volume 8, pages 196–203, San Diego, CA, 1998. Univ. of California.

[11] L. Wiskott. Estimating driving forces of nonstationary time series with slow feature analysis. arXiv.org e-Print archive, `http://arxiv.org/abs/cond-mat/0312317/`, Dec. 2003.

[12] L. Wiskott and T. Sejnowski. Slow feature analysis: Unsupervised learning of invariances. *Neural Computation*, 14(4):715–770, 2002.