The slowness principle: SFA can detect different slow components in non-stationary time series

Wolfgang Konen* and Patrick Koch

Institute for Informatics,
Cologne University of Applied Sciences,
Steinmüllerallee 1, D-51643 Gummersbach, Germany
E-mail: wolfgang.konen@fh-koeln.de
E-mail: patrick.koch@fh-koeln.de
*Corresponding author

Abstract: Slow feature analysis (SFA) is a bioinspired method for extracting slowly varying driving forces from quickly varying non-stationary 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; driving force detection; human perception; logistic map; non-linear regression; non-stationary time series; phase transition; slow feature analysis; SFA; slowness principle; unsupervised learning.

Reference to this paper should be made as follows: Konen, W. and Koch, P. (2011) 'The slowness principle: SFA can detect different slow components in non-stationary time series', *Int. J. Innovative Computing and Applications*, Vol. 3, No. 1, pp.3–10.

Biographical notes: Wolfgang Konen received his PhD in Physics from the University of Mainz, Germany in 1990. He is currently a Professor at the Cologne University of Applied Sciences, Faculty of Computer Science and Engineering. His research interests include computational intelligence, data mining, machine learning, reinforcement learning, pattern recognition and computer vision. He is a member of the ACM SIGEVO Special Interest Group on Genetic and Evolutionary Computation and of the GMA Special Interest Group on Computational Intelligence.

Patrick Koch received his Diploma in Computer Science from the University of Paderborn, Germany in 2008. He is a Researcher at the Faculty of Computer Science and Engineering, Cologne University of Applied Sciences and is currently working on his PhD in Computer Science at the University of Leiden, The Netherlands. His research interests include computational intelligence, machine learning and optimisation.

1 Introduction

The analysis of non-stationary 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 non-stationarities can be modelled 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, see (Wiskott, 2003).

The typical scenario of driving force analysis is: Given complex signals from machines or sensors, can we find slowly varying state information (state switch, state drift)? Application areas, among others, are EEG-analysis or monitoring of complex chemical or electrical power plants. One is interested in revealing the driving forces from the raw observed time series since they show interesting aspects of the underlying dynamics.

Several methods for detecting and visualising driving forces have been developed; based on recurrence plots, see Casdagli (1997), feedforward ANNs with an extra input unit, see Verdes et al. (2001) or, as Wiskott (2003) recently proposed, by slow feature analysis (SFA), a versatile, robust, and fast algorithm. SFA works fully unsupervised, just by searching non-linear combinations of the input signals which vary as slowly as possible in time.

The main purpose of this paper is to clarify the concept of slowness which is central to SFA. What is 'slow' in the

driving forces compared to the raw observed time series? Often it might be the case that driving forces contain 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, consider a driving force 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 of which parameters influence whether the first or the second choice is taken.

This paper is structured as follows: in Section 2, we briefly introduce the principles and the application fields of SFA. In Section 3, we describe the methods and results of our computer experiments. These results are further discussed in Section 4, which also discusses some connections to the notion of slowness in human perception. Finally, Section 5 summarises the findings from our experiments, their impact on driving force analysis and on our understanding of the slowness principle.

2 Slow feature analysis

2.1 SFA: origin and applications

We base our analysis on SFA (Wiskott and Sejnowski, 2002; Wiskott, 2003) as a driving force detection algorithm since it constitutes a versatile, robust and fast algorithm.

SFA has been originally developed by Wiskott (1998) in the context of an abstract model for unsupervised learning of invariances in the visual system of vertebrates and is described in detail in Wiskott and Sejnowski (2002) and Wiskott (2003).

SFA has been successfully applied to:

- invariant object recognition, see Franzius et al. (2008a)
- unsupervised learning of invariances in the visual system, see Wiskott (1998) and Franzius et al. (2008b)
- unsupervised learning of complex cell properties, see Berkes and Wiskott (2005)
- handwritten digit recognition, see Berkes (2005)
- driving force detection in non-stationary time series, see Wiskott (2003)
- gesture recognition from accelerometer time series, see Koch et al. (2010).

2.2 The SFA algorithm

Here we briefly review the approach described in Wiskott (2003). 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 $x = x(t)$ where t indicates time and $x = [x_1, \ldots, x_N]^T$ is a vector, the task can be formalised as follows: find the input-output function $g(x)$ that generates a scalar output signal:

$$y(t) \coloneqq g\big(x(t)\big) \tag{1}$$

with its temporal variation as slow as possible, measured by the variance of the time derivative, i.e.:

$$\text{minimise } \langle \dot{y}^2 \rangle \tag{2}$$

with $\langle \cdot \rangle$ indicating the temporal mean. To avoid the trivial constant solution, the output signal has to meet the following constraints:

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

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

This is an optimisation 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 non-linear basis functions, the problem becomes tractable with the mathematical details given in Wiskott (2003). A typical choice for the non-linear basis functions are monomials of degree 2,

$$h(x) = \Big[x_1, x_2, \ldots, x_N, x_1^2, x_1 x_2, \ldots, x_N^2 \Big]^T ,$$

but other choices, e.g., monomials of higher degree or radial basis functions could be used as well. Basically, SFA consists of the following four steps:

1 expand the input signal with some set of fixed possibly non-linear functions

2 sphere the expanded signal to obtain components with zero mean and unit covariance matrix

3 compute the time derivative of the sphered expanded signal and determine the normalised eigenvector of its covariance matrix with the smallest eigenvalue

4 project the sphered expanded signal onto this eigenvector to obtain the output signal, which we denote by y_1.

Sometimes we are also interested in the second-smallest eigenvalue and the corresponding projected output signal, which we denote by y_2. As usual for eigenvectors, higher components are orthogonal to lower ones, i.e., signal y_i satisfies the $i-1$ additional conditions $\langle y_i y_j \rangle = 0$, $\forall j = 1, \ldots, i-1$.

In the remainder of this paper we work with time series $x(1), x(2), \ldots, x(t), t \in \mathbf{N}$, instead of continuous signals $x(t), t \in \mathbf{R}$, but the transfer of the algorithm described above to time series is straightforward. The time derivative is simply computed as the difference between successive data points assuming a constant sampling spacing $\Delta t = 1$.

2.3 Slowness indicator η

It is useful to have a measure for the slowness of a signal. In principle, $\langle \dot{y}^2 \rangle$ is such an indicator, but Wiskott and Sejnowski (2002) propose a more intuitive measure η defined by:

$$\eta(y) = \frac{T}{2\pi} \frac{\sqrt{\langle \dot{y}^2 \rangle}}{\sqrt{\left\langle \left(y - \langle y \rangle\right)^2 \right\rangle}}, \tag{5}$$

where $\langle \cdot \rangle$ indicates the temporal average over $[t_0, t_0 + T]$. The measurement interval T should be at least as long as one or two periods of the slowest signal component. Low η-values indicate slow signals, high η-values fast signals. Wiskott and Sejnowski (2002) show that for a pure sine wave $y(t) = \sin(n 2\pi t / T)$ the η-value counts the number of oscillations in the observation interval, i.e., $\eta(y) = n$. The denominator in equation (5) ensures that each linear transform $ay(t) + b$ has the same η as $y(t)$ (This normalisation is not necessary for the training output signals of SFA, which are already normalised by construction to zero mean and unit variance, but it allows to apply the measure η to unnormalised time series like driving forces as well. Also notice that SFA output signals derived from test data are only approximately normalised, so that the normalisation is useful to make η independent of an accidental scaling factor).

3 Experiments

3.1 The driving force

In the following, we present simulation experiments with non-stationary time series. The approach follows closely the work of Wiskott (2003) but with more systematic variations on the underlying driving force. Our 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 equation (2) than the time series $w(t)$.

Here we consider a driving force that is made up of two frequency components:

$$\gamma(t) = \frac{1}{2}\Big(\underbrace{\sin(0.0005\nu_f t)}_{=\gamma_S(t)} + \underbrace{\sin(0.0047\nu_f t)}_{=\gamma_F(t)}\Big), \tag{6}$$

where the slow component γ_S is roughly ten times slower than the fast component 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 minimise 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 [the scale and offset of slow signals $y(t)$ extracted by SFA are fixed by the constraints and the sign is arbitrary as well). Therefore we define a *γ-aligned signal*

$$A\big(y(t);\gamma\big) = ay(t) + b, \tag{7}$$

where the constants a and b are chosen in such a way that:

$$\left\langle \left(ay(t) + b - \gamma(t)\right)^2 \right\rangle \stackrel{!}{=} \min. \tag{8}$$

Figure 1 The logistic map with driving force [see equation (9)] for different values of q (see online version for colours)

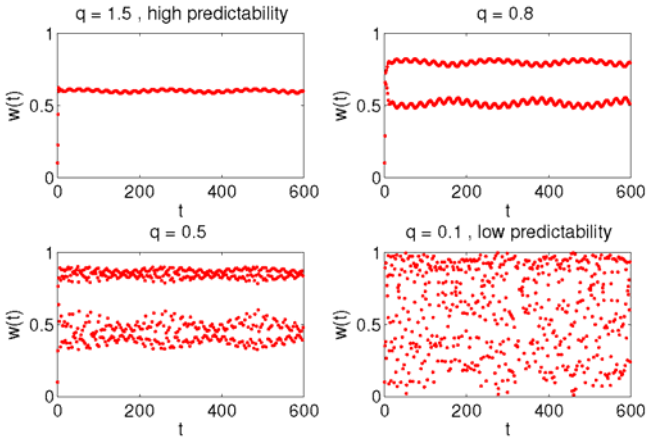

3.2 The logistic map

Now we put the slowly varying driving force γ into a fast varying logistic map. The logistic map was introduced by May (1976) as a "simple mathematical model with very complicated dynamics". We use it here because it allows capturing in one time series quite different dynamic behaviours and we are interested in the way SFA reacts on that different dynamics. The logistic map with driving force can be written as:

$$w(t+1) = \underbrace{\big(4 - q + 0.1\gamma(t)\big)}_{=r(t)} w(t)\big(1 - w(t)\big), \tag{9}$$

which maps for a constant $r(t) = 4$ the interval [0, 1] onto itself. The graph of $w(t)$ against $w(t+1)$ has the shape of an upside-down parabola crossing the abscissa at 0 and 1, its height being governed by $r(t)$. Parameter q controls the predictability of the logistic map (see Figure 1 for some examples): for $q < 0.33$ the map is fully in its chaotic regime, 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. It is this feature of the logistic map, its steerable predictability q, which makes it useful for our experiments on the slowness principle: For our first experiment we take $q = 0.1$ and since the logistic map is in

its chaotic regime, we have a highly chaotic map with no visible structure in Figure 2, top.

Figure 2 Top: logistic map time series $w(t)$ with $q = 0.1$ and with driving force for $\nu_f = 20$. Middle: the slowest SFA signal $y_1(t)$ (dots). Bottom: a detail plot for $0 < t < 500$ of the true slowly varying driving force $\gamma(t)$ (solid line), its slower component $\gamma_S(t)$ (dashed line) and the alignment of $y_1(t)$ [see equation (7)] to both components (dots) (see online version for colours)

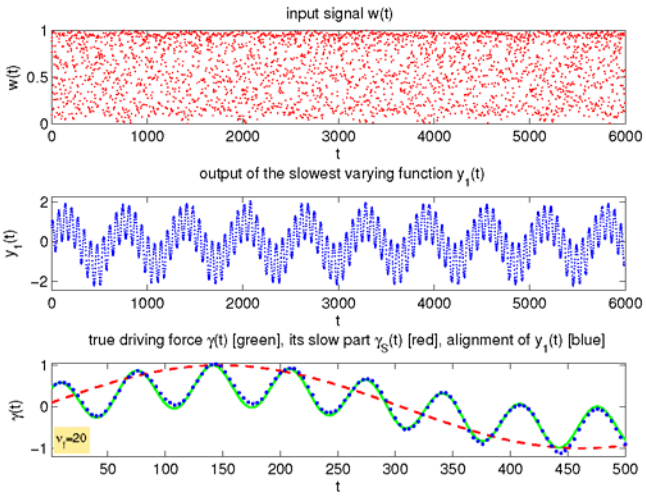

Notes: The estimated driving force $\gamma^{est}(t) = A(y_1(t);\gamma)$ is here in good agreement with $\gamma(t)$. SFA was done with embedding dimension $m = 19$.

One final ingredient is missing before we can start our simulations: taking the time series $w(t)$ directly as an input signal would not give SFA enough information to estimate the driving force, because SFA considers only data (and its derivative) from one time point at a time. Thus, it is necessary to generate an embedding-vector time series as an input. An embedding vector at time point t with delay τ and dimension m is defined as:

$$x(t) := \left[w(t - s_\tau), w\left(t - (s_\tau - 1)\right), \ldots, w(t + s_\tau) \right]^T \qquad (10)$$

for scalar $w(t)$, odd dimension m and $s_\tau := \tau(m-1)/2$. The definition can be easily extended to even m, which requires an extra shift of the indices by $\tau/2$ or its next lower integer to centre the used data points at t. Centring the embedding vectors results in an optimal temporal alignment between estimated and true driving force.

The following simulations are based on 6,000 data points each and were done with MATLAB 7.0.1 (Release 14) using the SFA toolkit sfa-tk of Berkes (2003).

3.3 The phase transition as a function of q and m

Figure 2 shows the time series for $q = 0.1$, the estimated driving force (from SFA with $m = 19$, $\tau = 1$ and second order monomials), and the true driving force. The estimated driving force is at $\nu_f = 20$ in good alignment with the true driving force. If we use a three times higher base frequency $\nu_f = 60$ in Figure 3, bottom, we see that the estimated driving force is now in nearly perfect 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 an envelope around the green curve (see detail plot in Figure 3, bottom). Table 1 shows the slowness indicator η for various signals and some ν_f.

Figure 3 The bottom panel of Figure 2 for three base frequencies $\nu_f = 20, 40, 60$ clearly shows the phase transition from the complete driving force $\gamma(t)$ (solid line) to its slower subcomponent $\gamma_S(t)$ (dashed line) (see online version for colours)

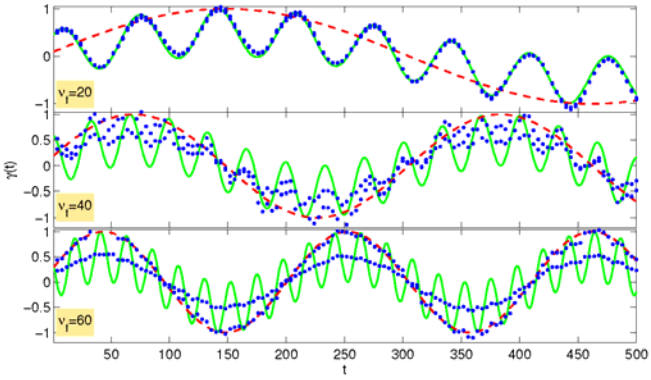

Note: We see two dotted curves since we align the slowest SFA signal once with $\gamma(t)$ and once with $\gamma_S(t)$.

Table 1 Slowness indicator η for various signals in our experiments as a function of ν_f

		Slowness η			
	Signal	$\nu_f = 20$	$\nu_f = 40$	$\nu_f = 60$	$\nu_f = 80$
Logistic map	$w(t)$	1,613	1,610	1,613	1,602
Driving force	$\gamma(t)$	*64*	127	190	235
Its slow part	$\gamma_S(t)$	9	19	29	*38*
SFA result	$y_1(t)$	*61*	69	42	*41*

Notes: Large η indicate fast, small η indicate slow signals. The logistic map has a tremendously high η. SFA detects the slower driving force for low ν_f (numbers in italics in first column). For high ν_f, SFA detects the slow component $\gamma_S(t)$ of the driving force (numbers in italics in last column).

Figure 3 shows quite clearly that there is a phase transition occurring around $\nu_f = 40$. To localise the phase transition and to study its dependence on other parameters of the system, we vary in Figure 4 the base frequency $\nu_f \in [5, 80]$

and we define the phase transition frequency $\nu(P.T.)$ as the lowest ν_f with $|C(y_1(t),\gamma(t))| < |C(y_1(t),\gamma_S(t))|$, where $C(\cdot,\cdot)$ denotes the usual correlation. Figure 4(a), shows that for small $q = 0.1$ (fully chaotic w) a phase transition occurs at $\nu_f = 33.9$ while for larger $q = 0.4$ (mix of chaotic and non-chaotic periods in w) the phase transition happens earlier at $\nu_f = 17.3$ [Figure 4(b)].

Figure 4 Absolute correlation of the slowest SFA-signal y_1 with the driving force γ (dash-dotted) and with its slow component γ_S (dashed lines). (a) $q = 0.1$ (logistic map in its chaotic regime)* (b) the same for $q = 0.4$ (logistic map, only partly chaotic) gives the same qualitative results, only the phase transition is shifted to the lower frequency $\nu_f = 17$ ** (see online version for colours)

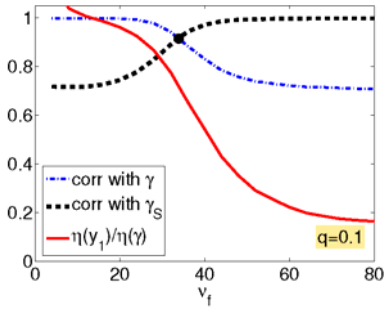

(a)

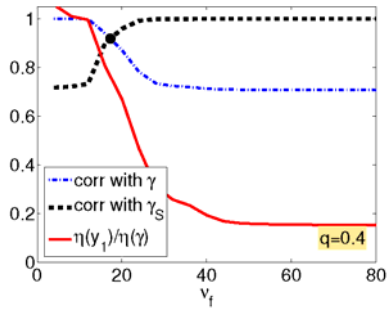

(b)

Notes: *For sufficiently large ν_f the correlation with the slow component γ_S is stronger than the correlation with γ. The solid line is the quotient $\eta(y_1)/\eta(\gamma)$ and it drops largely for $\nu_f = 34$ and finally approaches $\eta(\gamma_S)/\eta(\gamma) \approx 0.1$.

**SFA was carried out with monomials of degree 2, embedding dimension $m = 19$ and delay $\tau = 1$.

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 Figure 5. 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.

Before we discuss the results further in Section 4, a second remark is in order concerning the SFA implementation sfa-tk of Berkes (2003): while it worked well for small embedding dimensions m, larger m led quite inevitably to wrong 'slow' signals y_1 which were neither slow nor did they respect the unit variance condition $\langle y^2 \rangle = 1$. This is reported in more detail in Konen (2009) where it is also shown that this behaviour can be traced back to numeric instabilities of the implementation and a slightly modified implementation based on SVD is presented – closer along the lines of Wiskott and Sejnowski (2002) – 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 envelope. 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 hypothesise 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 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 Figure 5. 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 Figure 5, 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.

Figure 5 (a) Phase transition frequency $\nu(P.T.)$ as a function of q and m * (b) the absolute correlation values at fixed $\nu_f = 40$ show a phase transition towards the slower component as we increase the embedding dimension m (top) or the predictability q of the logistic map (bottom) (see online version for colours)

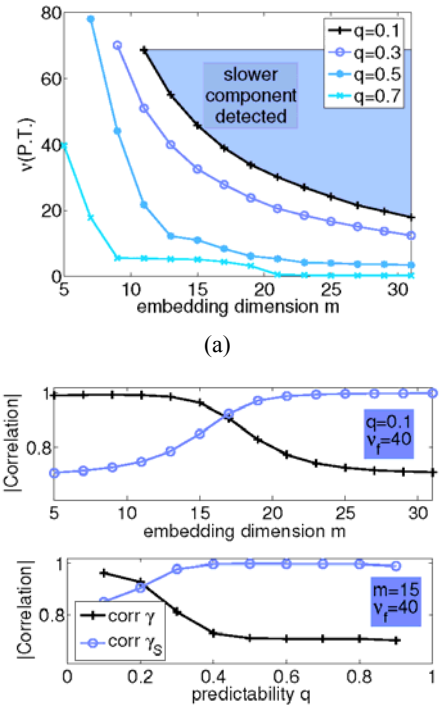

Note: In the area above each phase transition curve the detection of the slower component γ_S of the driving force is stronger.

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 it is so slow that 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.

4.1 Non-linear regression

Hypotheses 1 and 2 are also supported by the following non-linear regression experiment: for the set of non-linear basis function 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 Figure 4 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 non-linear regression experiment.

Note that non-linear regression is not a substitute for SFA, since it requires the output signal (the driving force) to be known beforehand while SFA finds the driving force unsupervisedly. Non-linear regression serves here as a diagnostic tool to clarify why SFA decides for this or the other component of the driving force.

4.2 Robustness of SFA

SFA as tested in this paper works robustly over a large range of parameters ν_f, m, q. It is, however, necessary to deal carefully with zero eigenvalues which occur frequently when the embedding dimension m is large and/or the noise is low. If zero eigenvalues are not handled properly, it is likely to see numerical instabilities. A numerically robust way to handle them is based on SVD as described by Konen (2009).

4.3 Accuracy of the estimated driving force

The driving force is estimated with high accuracy, although the estimation is undetermined up to any invertible transformation, see Wiskott (2003). We found that this is true even for large embedding dimensions, e.g., $m = 51$, in contrast to the hypothesis of Wiskott (2003) that only small embedding dimensions would avoid more complicated invertible transformations.

4.4 Noise sensitivity

The results described in this paper were obtained with noise-free data. In some simulations we tested the effects of adding Gaussian noise to the data before embedding. For $m = 19$, $\nu_f = 56$ and $q = 0.4$, the effect of adding 1%, 2% and 5% noise brought the correlation

between slow component γ_S and slowest SFA-signal y_1 from $|C| = 0.999$ down to 0.85, 0.75 and 0.60, respectively. Thus for small noise levels < 1% the main effects persist, but in contrast to Wiskott (2003) the noise sensitivity is somewhat higher, which means that 5% noise destroys most of the correlation with the slow component. This might be due to the more complex nature of the driving force build up from multiple components. On the other hand, we found that larger embedding dimensions, e.g., $m = 51$ stabilise the system and bring up the correlation again to 0.963, 0.893 and 0.804, respectively, but further experiments are needed to investigate this systematically.

4.5 Hierarchical SFA

As the preceding paragraph has shown, higher dimensional input data usually improve the SFA performance. However, since the number of monomials grows quadratically with the embedding dimension, the requirements in computer memory and computing time quickly increase. It is, therefore, interesting to investigate whether similar results as with high m can be obtained by hierarchical approaches where first smaller parts of the embedding vector are analysed by a first SFA, whose outputs are then combined by a second SFA, as has been demonstrated by Wiskott and Sejnowski (2002).

4.6 Connection to human perception

Since SFA was originally developed by Wiskott (1998) as a model for neural information processing, it might be natural to ask whether the observed switch between components and its phase transition has any parallel in human perception or human motion coordination. Several phenomena with switching effects are discussed in the literature:

- The well-known backward spinning-wheel illusion described by Purves et al. (1996) occurs frequently in movies or under stroboscopic lighting conditions and it shows the transition from a fast forward rotation percept to a slow backward rotation percept. 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, see Kline et al. (2004) and Purves et al. (1996). No snapshot-like explanation is possible here; the percept is continuous having a greater resemblance to the smoothly varying driving force of our SFA experiments. A possible explanation of the sunlight spinning-wheel illusion is according to Kline et al. (2004) that rivalry between different motion detectors in the brain occurs.
- Kelso (1981) describes another well-known phase transition which occurs in bimanual motion coordination when performing certain movements with the index fingers of both hands. A theoretical model for this phenomenon exists, the so-called Haken-Kelso-Bunz model, see Haken et al. (1985), which predicts a phase transition and certain hysteresis effects.

SFA has shown similar capabilities in the sense that the same setup can learn to synchronise 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 Conclusions

In this paper we have investigated the notion of *slowness* in SFA. It has been verified that SFA can reliably detect either slow driving forces or their subcomponents over a broad range of parameters in non-stationary 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 that moment when its reconstruction in the expanded function space used by SFA has more temporal variation than the reconstruction of the complete driving force.

There are still a number of open questions. Does hierarchical SFA, which achieves larger embedding dimensions with a smaller computing budget, allow for the detection of very slow components which are 'out-of-reach' for plain SFA? Can an extended SFA system model more closely certain switching effects known from human perception [as for example the backward spinning-wheel illusion of Purves et al. (1996)]? Finally, it is necessary to apply SFA to real world data and to see whether the results reported in this study have some similar parallel there.

One final advice for real world driving force detection might be drawn from our study: In real world data it will often not be possible to vary the base frequency or the degree of non-linearity in the observed dynamic system systematically (as we did here in our simulations). Therefore, the advice is to vary the embedding dimension m over a broad range in order to detect potentially different 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.

Acknowledgements

The authors are grateful to Laurenz Wiskott for helpful discussions on SFA and to Pietro Berkes for providing the MATLAB code for sfa-tk, see Berkes (2003). This work has been supported by the Bundesministerium für Forschung und Bildung (BMBF) under the grant SOMA (AIF FKZ 17N1009, 'Systematic optimization of models in IT and automation') and by the Cologne University of Applied Sciences under the research focus grant COSA.

References

Berkes, P. (2003) 'sfa-tk: slow feature analysis toolkit for MATLAB (v.1.0.1)', available at http://people.brandeis.edu/~berkes/software/sfa-tk (accessed on June 2010).

Berkes, P. (2005) 'Pattern recognition with slow feature analysis', Cognitive Sciences EPrint Archive (CogPrint) 4104, available at http://cogprints.org/4104/ (accessed on June 2010).

Berkes, P. and Wiskott, L. (2005) 'Slow feature analysis yields a rich repertoire of complex cells', *Journal of Vision*, Vol. 5, No. 6, pp.579–602.

Casdagli, M.C. (1997) 'Recurrence plots revisited', *Physica D: Nonlinear Phenomena*, Vol. 108, Nos. 1–2, pp.12–44.

Franzius, M., Wilbert, N. and Wiskott, L. (2008a) 'Invariant object recognition with slow feature analysis', *Lecture Notes in Computer Science*, Vol. 5163, pp.961–970.

Franzius, M., Wilbert, N. and Wiskott, L. (2008b) 'Unsupervised learning of invariant 3D-object and pose representations with slow feature analysis', *Proceedings of the 6th FENS Forum of European Neuroscience*, Geneva, Switzerland, 12–16 July.

Haken, H., Kelso, J. and Bunz, H. (1985) 'A theoretical model of phase transitions in human hand movements', *Biological Cybernetics*, Vol. 51, No. 5, pp.347–356.

Kelso, J. (1981) 'On the oscillatory basis of movement', *Bulletin of the Psychonomic Society*, Vol. 18, p.63.

Kline, K., Holcombe, A. and Eagleman, D. (2004) 'Illusory motion reversal is caused by rivalry, not by perceptual snapshots of the visual field', *Vision Research*, Vol. 44, No. 23, pp.2653–2658.

Koch, P., Konen, W. and Hein, K. (2010) 'Gesture recognition on few training data using slow feature analysis and parametric bootstrap', *Proceedings of IEEE World Congress on Computational Intelligence (WCCI)*, Barcelona, Spain, 18–23 July.

Konen, W. (2009) 'On the numeric stability of the SFA implementation sfa-tk', arXiv.org e-Print archive, available at http://arxiv.org/abs/0912.1064 (accessed on June 2010).

May, R.M. (1976) 'Simple mathematical models with very complicated dynamics', *Nature*, Vol. 261, pp.459–467.

Purves, D., Paydarfar, J.A. and Andrews, T.J. (1996) 'The wagon-wheel illusion in movies and reality', *Proceedings of National Academy of Sciences USA*, Vol. 93, pp.3693–3697.

Verdes, P.F., Granitto, P.M., Navone, H.D. and Ceccatto, H.A. (2001) 'Nonstationary timeseries analysis: accurate reconstruction of driving forces', *Physical Review Letters*, Vol. 87, No. 12, pp.124101-1–4.

Wiskott, L. (1998) 'Learning invariance manifolds', *Proceedings of the 5th Joint Symposium on Neural Computation*, San Diego, CA, 16 May, Vol. 8, pp.196–203.

Wiskott, L. (2003) 'Estimating driving forces of nonstationary time series with slow feature analysis', arXiv.org e-Print archive, available at http://arxiv.org/abs/cond-mat/0312317/ (accessed on June 2010).

Wiskott, L. and Sejnowski, T. (2002) 'Slow feature analysis: unsupervised learning of invariances', *Neural Computation*, Vol. 14, No. 4, pp.715–770.