RUHR-UNIVERSITÄT BOCHUM

INSTITUT FÜR NEUROINFORMATIK

RUB

Predictable Feature Analysis (PFA)

Stefan Richthofer, Laurenz Wiskott

stefan.richthofer@ini.rub.de laurenz.wiskott@ini.rub.de

Inspired by Slow Feature Analysis (SFA, [8]), PFA extracts features by predictability rather than slowness.

This way, we prepare to handle interactive scenarios that involve notions of control, planning and decision-making. In order to perform any kind of planning or intelligent control, it is crucial to have a model that is capable of estimating the consequences of possible actions (cf. [1]). Since predictability is a desired property of such a model by definition, the PFA-setup is well suited for this purpose. While in control theory the involved models are usually formulated as a set of (partial) differential equations in a problem specific manner, with PFA we preserve the main advantages of SFA – namely its unsupervised nature and the ability to build a model in a self-organizing fashion.

Slow Feature Analysis (SFA)

On a given signal, Slow Feature Analysis (SFA, [8]) performs dimensionality reduction by selecting the slowest subsignals available under decorrelation and unit-variance constraints. Typical data-analysis and recognition tasks like regression and classification become much more feasible on the reduced signal. SFA was demonstrated to be capable of object recognition invariant to spatial transformations, self-organization of complex cell receptive fields, nonlinear blind source separation ([6, 5, 4, 2]) and other tasks of this kind.

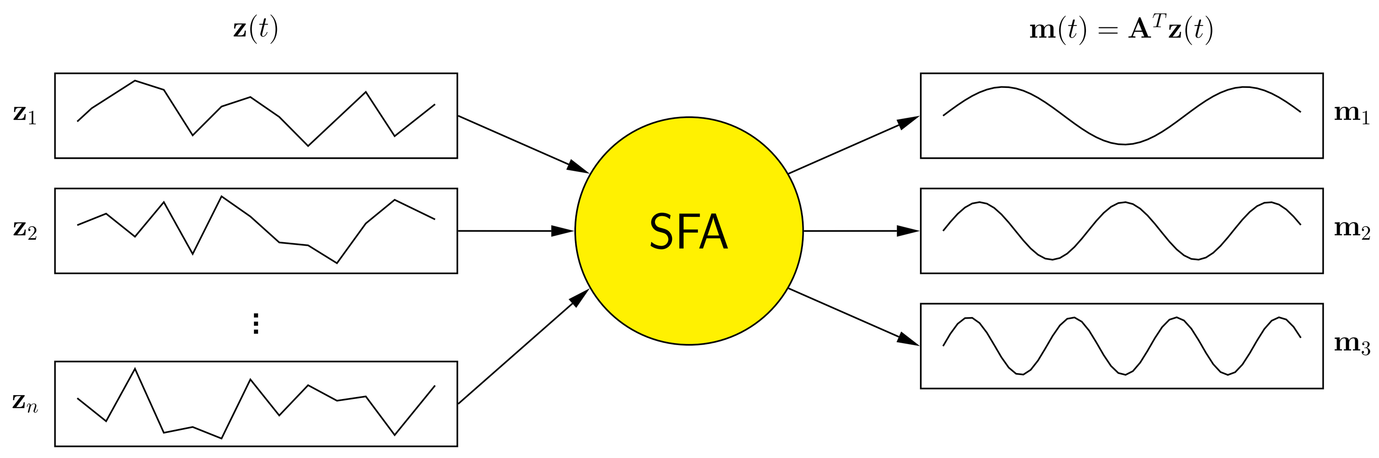

$$
\begin{aligned}
\underset{\mathbf{a}_i \in \mathbb{R}^{\dim(\mathbf{z})}}{\text{minimize}} \quad & \mathbf{a}_i^T \left\langle \dot{\mathbf{z}}\dot{\mathbf{z}}^T \right\rangle_t \mathbf{a}_i \\
\text{subject to} \quad & \mathbf{a}_i^T \left\langle \mathbf{z}\mathbf{z}^T \right\rangle_t \mathbf{a}_i = 1 && \text{(unit variance)} \\
& \mathbf{a}_i^T \left\langle \mathbf{z}\mathbf{z}^T \right\rangle_t \mathbf{a}_j = 0 \quad \forall\, j < i && \text{(uncorrelated)}
\end{aligned}
$$

Predictable Feature Analysis (PFA)

While there exist model-independent notions from information theory (cf. information bottleneck approach [3], ForeCA [7]), we consider predictability with respect to a certain prediction model. In the current setup, we consider signal components as predictable if they can be predicted by a linear autoregressive model after an optional, nonlinear preprocessing. The desired property of the extracted model function $\mathbf{m}$ can be expressed as follows:

$$
\mathbf{m}_i(t) \;:=\; \mathbf{a}_i^T \mathbf{z}(t) \;\overset{!}{\approx}\; b_{i1}\mathbf{a}_i^T \mathbf{z}(t-1) + \ldots + b_{ip}\mathbf{a}_i^T \mathbf{z}(t-p)
$$

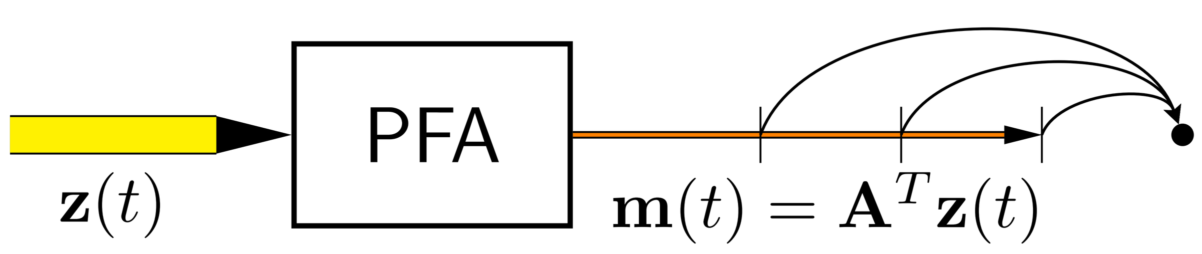

The problem of finding optimal vectors $\mathbf{a}_i$ and $\mathbf{b}_i$ arises. We formalize it in least-squares-sense and include the constraints from SFA to avoid trivial or repeated solutions:

$$
\begin{aligned}
\underset{\mathbf{a}_i \in \mathbb{R}^{\dim(\mathbf{z})},\, \mathbf{b}_i \in \mathbb{R}^p}{\text{minimize}} \quad & \left\langle \left(\mathbf{a}_i^T \mathbf{z}(t) - \sum_{k=1}^{p} b_{ik} \mathbf{a}_i^T \mathbf{z}(t-k) \right)^2 \right\rangle_t \\
\text{subject to} \quad & \mathbf{a}_i^T \left\langle \mathbf{z}\mathbf{z}^T \right\rangle_t \mathbf{a}_i = 1 && \text{(unit variance)} \\
& \mathbf{a}_i^T \left\langle \mathbf{z}\mathbf{z}^T \right\rangle_t \mathbf{a}_j = 0 \quad \forall\, j < i && \text{(uncorrelated)}
\end{aligned}
$$

The extracted features must be optimized for predictability, but judging their predictability is an optimization problem by itself: If $\mathbf{a}_i$ or $\mathbf{b}_i$ is fixed, the other one can be obtained efficiently, but optimizing in turns runs into local optima. The following sketch illustrates a relaxation that allows efficient computation of near-optimal $\mathbf{a}_i$ and $\mathbf{b}_i$:

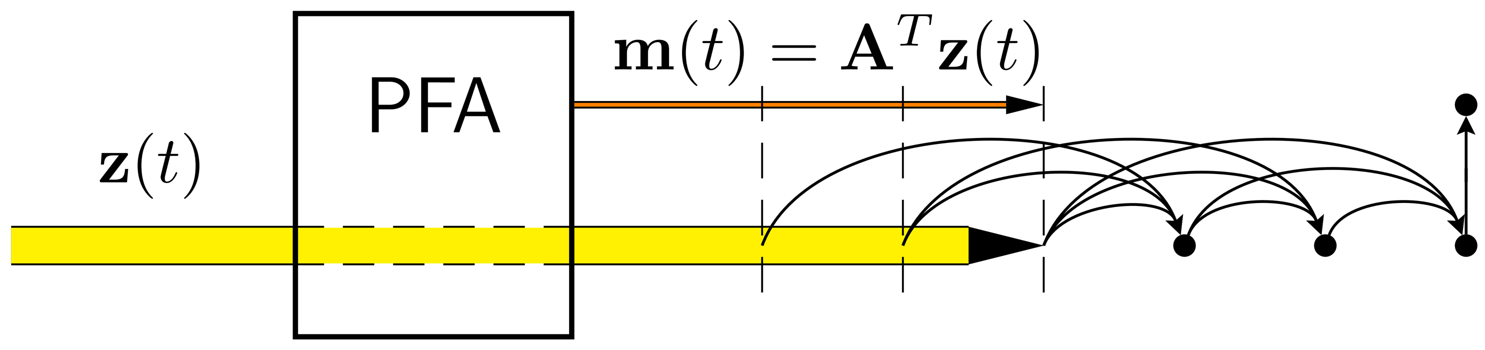

We extract those signals that have the lowest error after an iterated linear extrapolation of the entire input signal. As a side-effect, $\mathbf{m}$ results in a signal that can be well predicted independently from $\mathbf{z}$:

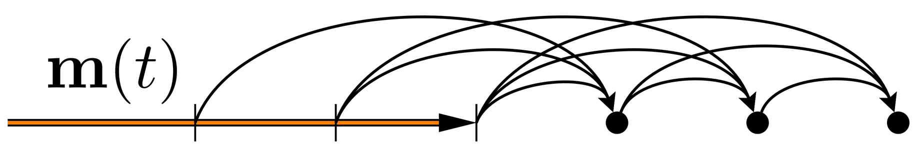

Future work: PFA using external information

An experimental variant of PFA allows to optimize predictability regarding to an additional stream of information.

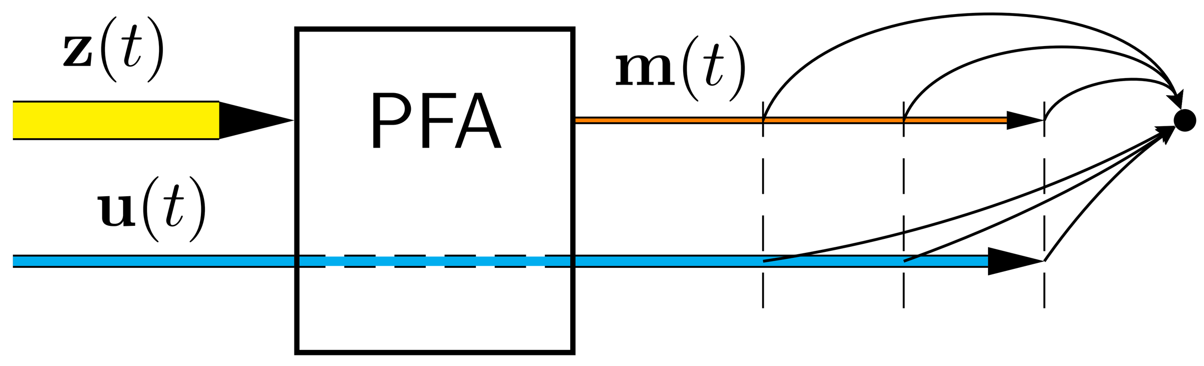

This way we attempt to model problems that involve analysis of the relationship between two signals – typically representing control and sensor components – action and outcome. Planned tasks:

- let an agent explore an environment, learning the relation between its movement commands and the resulting position
- use the learned relation to reach any desired feasible state
- apply this to various scenarios (e.g. pendulum/pole swing up, moving blocks around, multi-room navigation, visual homing)

References

[1] W. Bialek, I. Nemenman, and N. Tishby. Predictability, complexity, and learning. *Neural Comput*, 13:2409–2463, Nov 2001. [DOI:10.1162/089976601753195969] [PubMed:11674845].

[2] Tobias Blaschke, Tiziano Zito, and Laurenz Wiskott. Independent slow feature analysis and nonlinear blind source separation. *Neural Computation*, 19(4):994–1021, 2007.

[3] Felix Creutzig, Amir Globerson, and Naftali Tishby. Past-future information bottleneck in dynamical systems. *Physical Review E*, 79:041925, 2009.

[4] Sven Daehne, Niko Wilbert, and Laurenz Wiskott. Self-organization of v1 complex cells based on slow feature analysis and retinal waves. In *Proc. Bernstein Conference on Computational Neuroscience, Sep 27–Oct 1, Berlin, Germany*, 2010.

[5] Mathias Franzius, Henning Sprekeler, and Laurenz Wiskott. Slowness and sparseness lead to place-, head direction-, and spatial-view cells. In *Proc. 3rd Annual Computational Cognitive Neuroscience Conference, Nov. 1–2, San Diego, USA*, pages III–8, 2007.

[6] Mathias Franzius, Niko Wilbert, and Laurenz Wiskott. Invariant object recognition and pose estimation with slow feature analysis. *Neural Computation*, 23(9):2289–2323, 2011.

[7] Georg Goerg. Forecastable component analysis. In Sanjoy Dasgupta and David Mcallester, editors, *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, volume 28, pages 64–72. JMLR Workshop and Conference Proceedings, May 2013.

[8] L. Wiskott, P. Berkes, M. Franzius, H. Sprekeler, and N. Wilbert. Slow feature analysis. *Scholarpedia*, 6(4):5282, 2011.

This work is funded by a grant from the German Research Foundation (Deutsche Forschungsgemeinschaft, DFG) to L. Wiskott (SFB 874, TP B3) and supported by the German Federal Ministry of Education and Research within the National Network Computational Neuroscience – Bernstein Fokus: "Learning behavioral models: From human experiment to technical assistance", grant FKZ 01GQ0951.