

Heuristic Evaluation of Expansions for Non-Linear Hierarchical Slow Feature Analysis

Alberto N. Escalante-B.
Institut für Neuroinformatik
Ruhr-University of Bochum, Germany
Email: alberto.escalante@ini.rub.de

Laurenz Wiskott
Institut für Neuroinformatik
Ruhr-University of Bochum, Germany
Email: laurenz.wiskott@ini.rub.de

Abstract—Slow Feature Analysis (SFA) is a feature extraction algorithm based on the slowness principle with applications to both supervised and unsupervised learning. When implemented hierarchically, it allows for efficient processing of high-dimensional data, such as images. Expansion plays a crucial role in the implementation of non-linear SFA. In this paper, a fast heuristic method for the evaluation of expansions is proposed, consisting of tests on 7 problems and two metrics. Several expansions with different complexities are evaluated. It is shown that the method allows predictions of the performance of SFA on a concrete data set, and the use of normalized expansions is justified. The proposed method is useful for the design of powerful expansions that allow the extraction of complex high-level features and provide better generalization.

I. INTRODUCTION

One important problem in machine learning is the development of algorithms for dimensionality reduction (cf. [1]). Noticeable examples of algorithms for *linear* dimensionality reduction are Principal Component Analysis (PCA), Independent Component Analysis (ICA), and Linear Discriminant Analysis (LDA) [2]. These algorithms can be applied to several other tasks besides dimensionality reduction, such as data compression, signal unmixing, prediction, classification, and regression. Their popularity and effectiveness is remarkable given that they all rely on a simple linear model, and all of them can be regarded as feature extraction algorithms that *optimally* solve a particular feature extraction problem when restricted to a *linear* feature space.

In many realistic scenarios, feature extraction problems (e.g., dimensionality reduction) cannot be satisfactorily solved by linear models. Thus, several non-linear variations of the algorithms above have been proposed [1]. For instance, in the field of image processing, feature extraction from an object is generally a non-linear problem if non-homogeneous backgrounds are used and if invariance to the position of the object is desired. Pose-normalization and other types of preprocessing are helpful to reduce the complexity of the underlying problem, but do not guarantee linearization.

Slow Feature Analysis (SFA) [3] is a promising learning algorithm for feature extraction inspired by biology (e.g., by the visual system) based on the slowness principle (cf. [4]). Given a multi-dimensional signal as a function of time, the objective of SFA is to extract those features that change as little as possible over time, while feature extraction is

done instantaneously. These features are called *slow features* and additionally satisfy a few normalization and decorrelation restrictions (see Section III-A).

Hierarchical SFA is one form of implementing SFA that allows processing high-dimensional data such as images. In this paper, we study the design of various types of non-linear expansions that are particularly suitable for hierarchical SFA. We concentrate on the heuristic evaluation of non-linear expansions that allow complex and effective feature extraction (dimensionality reduction) while at the same time offering resistance against the frequent problems of overfitting and outlier amplification (see Section III-E). For this purpose, we make use of a theoretical analysis on SFA [5] and of empirical observations gained after long lasting experimentation with the algorithm at solving several feature extraction problems. The proposed method might be also applicable to non-linear hierarchical implementations of PCA, ICA, or LDA.

II. PREVIOUS WORK

SFA is typically a non-linear algorithm that can be implemented in two ways. The most sophisticated one relies on the introduction of kernels [6], [7] and is possible due to the kernel trick. This allows the consideration of large or even infinite feature spaces. In practice, kernel SFA has been, however, less used. This might be explained by the difficulty of choosing a good kernel or by high computation costs for large data sets. The second way, which is the one followed in this paper, consists on the application of linear SFA after the data has been transformed by an explicit non-linear expansion. An efficient optimal algorithm for linear SFA is employed [3]. A widely used expansion is the Quadratic Expansion (*QExp*), as follows: given some data sample $\mathbf{x} = (x_1, x_2, \dots, x_I)$, $QExp(\mathbf{x}) \stackrel{\text{def}}{=} (I(\mathbf{x}), QM(\mathbf{x}))$, that is, *QExp* is the vectorial concatenation of $I(\mathbf{x}) \stackrel{\text{def}}{=} (x_1, x_2, \dots, x_I)$, the identity function, and $QM(\mathbf{x}) \stackrel{\text{def}}{=} (x_1^2, x_1x_2, \dots, x_1x_I, x_2^2, \dots, x_2x_I, x_3^2, \dots, x_I^2)$, all the quadratic monomials, (i.e., all the products of pairs of elements of \mathbf{x}). Higher degree expansions can be defined analogously. For example, $CExp(\mathbf{x}) \stackrel{\text{def}}{=} (I(\mathbf{x}), QM(\mathbf{x}), CM(\mathbf{x}))$ is a cubic expansion, where *CM* are third degree monomials.

Below we review previous architectures based on SFA, the expansions that were used, and the measures employed to

reduce overfitting. Franzius et al. [8] developed a hierarchical SFA network capable of learning pose and object identity from synthetic fish images. The expansion *QExp* was used, and a clipping step was introduced to keep signal amplitudes within the range $(-4, 4)$. Moreover, Gaussian noise was added to improve generalization.

Berkes et al. [4] analyzed sequences of pairs of patches taken from natural images with a single SFA node. It was demonstrated that the features extracted by SFA show the same properties described by complex cells found in the visual cortex (e.g., direction selectivity). The structure of their system was: dimensionality reduction with PCA, followed by *QExp* and SFA. Overfitting was avoided through the preprocessing step with PCA and through the use of a large amount of training data (about 400,000 samples).

Escalante et al. [9] developed a linear and a non-linear hierarchical network for SFA. Both networks were capable of accurately estimating the age and gender of artificial subjects generated from 3-D face models under constant pose and expression. The non-linear network was based on a different expansion: $0.8Exp(\mathbf{x}) \stackrel{\text{def}}{=} (I(\mathbf{x}), |\mathbf{x}|^{0.8}) = (x_1, x_2, \dots, x_I, |x_1|^{0.8}, |x_2|^{0.8}, \dots, |x_I|^{0.8})$. The linear network outperformed the non-linear one on test data due to insufficient training data. This simple expansion overcame the problem of uncontrolled outlier amplification without resorting to clipping.

Later, Escalante et al. (cf. [10, Appendix A]) developed SFA networks with expansions similar to *0.8Exp* and applied them to various feature extraction problems on frontal face photographs, such as the estimation of position and size of a face in an image patch. This was consolidated in a system for face detection. For the problems addressed, linear SFA gave poor performance. Non-linear SFA, however, performed much better confirming that these problems are highly non-linear and showing that even a modest expansion repeated in several layers is capable of untangling the underlying problem.

III. LINEAR AND NON-LINEAR SFA

A. SFA Optimization Problem

The SFA problem [3], [4] can be stated as follows. Given an I -dimensional signal $\mathbf{x}(t) = (x_1(t), \dots, x_I(t))^T$, find a set of real-valued instantaneous functions $\mathbf{g}(\mathbf{x}) = (g_1(\mathbf{x}), \dots, g_J(\mathbf{x}))$ lying in a function space \mathcal{F} so that for each component $y_j(t)$ of the output signal $\mathbf{y}(t) \stackrel{\text{def}}{=} \mathbf{g}(\mathbf{x}(t))$, for $1 \leq j \leq J$, the following holds:

$$\Delta(y_j) \stackrel{\text{def}}{=} \langle \dot{y}_j(t)^2 \rangle \text{ is minimal (objective function)} \quad (1)$$

under the constraints

$$\langle y_j(t) \rangle = 0, \text{ (zero mean)} \quad (2)$$

$$\langle y_j(t)^2 \rangle = 1, \text{ (unit variance)} \quad (3)$$

$$\langle y_j(t)y_{j'}(t) \rangle = 0, \forall j' < j \text{ (decorrelation)}. \quad (4)$$

The conditions (2)–(4) assure that the output signals are not constant and code different information of the input signal.

Notice that the delta value $\Delta(y_j)$ is defined as the average energy of the temporal derivative of y_j , and is therefore a measure of the temporal variation of such signal. The problem is solved iteratively beginning with y_1 and finishing with y_J .

In some theoretical analysis (e.g., [5]), it is common to let \mathbf{g} belong to an unrestricted function space. In practice, this function space is usually limited to the function space of all quadratic functions.

B. Linear SFA Algorithm

The SFA algorithm [3] takes advantage of the fact that the optimal *linear* solutions to the optimization problem only depend on the covariance matrix $\mathbf{C} \stackrel{\text{def}}{=} \langle \mathbf{x}\mathbf{x}^T \rangle$ of the training signal $\mathbf{x}(t)$, where the mean of $\mathbf{x}(t)$ was removed, and on the second-moment matrix $\dot{\mathbf{C}} \stackrel{\text{def}}{=} \langle \dot{\mathbf{x}}\dot{\mathbf{x}}^T \rangle$ of the time derivative $\dot{\mathbf{x}}(t)$ of the training signal. In practice, discrete time is used. Therefore, the training signal becomes a sequence of samples, and $\dot{\mathbf{x}}$ is usually approximated by the difference of consecutive samples: $\dot{\mathbf{x}}(t) \stackrel{\text{def}}{=} \mathbf{x}(t+1) - \mathbf{x}(t)$. The matrix \mathbf{C} is consequently approximated by its corresponding sample covariance matrix: $\mathbf{C} = \frac{1}{N-1} \sum_{t=1}^N (\mathbf{x}(t) - \bar{\mathbf{x}})(\mathbf{x}(t) - \bar{\mathbf{x}})^T$, and $\dot{\mathbf{C}}$ is computed as $\dot{\mathbf{C}} = \frac{1}{N-1} \sum_{t=1}^{N-1} (\mathbf{x}(t+1) - \mathbf{x}(t))(\mathbf{x}(t+1) - \mathbf{x}(t))^T$.

In the linear function space, the outputs can be written as $y_j(t) = g_j(\mathbf{x}(t)) = \mathbf{w}_j^T \mathbf{x}$, and the SFA problem can be reduced to a generalized eigenvalue problem (cf. [4]). That is, weights $\mathbf{W} = \{\mathbf{w}_j\}$ are computed such that $\mathbf{W}^T \mathbf{C} \mathbf{W} = \mathbf{I}$ and $\dot{\mathbf{C}} \mathbf{W} = \mathbf{W} \mathbf{\Lambda}$, where $\mathbf{\Lambda}$ is a diagonal matrix with diagonal $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_J$. Hence, efficient algorithms for solving the latter problem can be used, and the SFA algorithm has a complexity similar to PCA.

C. Hierarchical Slow Feature Analysis

For high-dimensional data, the direct use of linear SFA might be too expensive since it has a computational complexity of $\mathcal{O}(ND^2 + D^3)$ where N is the number of samples and D is the input dimensionality. An expansion would make this problem even more severe. Hierarchical SFA is a greedy implementation of SFA that allows to cope with high-dimensional data by dividing it in lower-dimensional chunks (e.g., 2 chunks of N samples of dimensionality $D/2$) that are separately processed by SFA nodes in a layer, cf. [11]. Afterwards, the resulting slow signals computed by each node are grouped together and further processed by an SFA node in the next layer. This process can be repeated recursively and the nodes can be organized in a multi-layer network, where each layer is trained separately from the bottom to the top, and the top layer contains a single node.

D. Non-Linear Hierarchical SFA Networks

In a non-linear SFA network each (linear) SFA node is preceded by an expansion node. The increase in dimensionality of the data chunk depends on the expansion that was applied, which might be different for each SFA node. The number of slow signals preserved by the SFA nodes is kept similar as in the linear case to avoid an explosion in their number.




Fig. 1. Example of a hierarchical network, where the inputs are 135 by 135-pixel images with an output of 40 slow features.

In principle, the larger the expansion is, the more complex the features learned by SFA might be and, thus, a large expansion might allow the extraction of slower features. Ideally, it would be desirable to allow full polynomial expansion of the data, and let linear SFA find an appropriate linear combination of the expanded data. The feature space available to SFA would contain all polynomial functions, and could, thus, extract extremely complex non-linear features. On the other hand, expansions drastically increase the number of dimensions of the data, and consequently, increase the number of free parameters in the model to be learned, which under normal conditions gives rise to overfitting.

E. Overfitting and Outlier Amplification

Overfitting is the over-adaptation of a model to the peculiarities of the training data as opposed to the true distribution of the data, resulting in poor performance on new test data. This problem is more evident when a large expansion is used or when the number of training samples is small. In the context of SFA, the performance can be measured in terms of the delta values $\Delta(y_j)$, for $1 \leq j \leq J$, of the extracted features.

Another recurrent problem found when using some expansions is *outlier amplification*, in which test samples less similar to the samples seen during training are mapped into even more atypical features, for example, by the quadratic terms of $QExp$. In hierarchical implementations, this problem accumulates across layers, resulting in an uncontrolled explosion of the magnitude of the signals for particular samples. For example, on test data we have frequently seen features with amplitudes of 10^{60} , whereas on training data the features have zero mean and unit variance. We remark that this problem only occurs in a fraction of the samples of the test data, and is directly related, but different from the problem of overfitting. Clipping (e.g. [8]) can reduce this problem, but is a suboptimal solution because feature extraction is impaired due to information loss.

Overfitting and outlier amplification in the context of expansions can be measured with heuristics proposed in Section V.

IV. EXPANSIONS AND BASIS FUNCTIONS

Let us assume that an SFA node is trained with data that can be described by an underlying slow parameter that monotonically increases as a function of time, and that the parameter can be solely computed from each data sample (with

an unknown arbitrarily complex function). Then, the theory of optimal free SFA responses [5] predicts that the slowest feature learned by SFA with an unrestricted function space is a half-cosine function (with arbitrary sign) $\pm\sqrt{2}\cos(\pi t/T)$, for $0 \leq t \leq T$. In principle, since such a signal is also monotonically increasing (or decreasing) in this interval, a mapping from the slowest signal to the slow parameter can be found. When restricted function spaces are used, such as in quadratic SFA, we cannot expect the extraction of the half-cosine function exactly, but we can expect a distorted and noisy version of this function. In other words, a restricted version of SFA can be seen as an approximation, in a restricted function space, to an optimally slow function.

When SFA is implemented as an expansion followed by linear SFA, linear SFA finds an optimal linear combination of the expanded signal, such that the slowness is minimized. Basically, the expansion is playing the role of a basis that defines a feature space, from which SFA linearly chooses the slowest possible function. If the expansion is rich enough, the slowest output of SFA will closely resemble the optimally slow function.

Not all expansions are equally useful. This is a direct consequence of the richness of the different feature spaces defined by different expansions, and of the fact that in hierarchical networks the input to each node has a particular distribution. Recall that the input to an SFA node is the concatenation of the outputs of previous SFA nodes, which must satisfy normalization and decorrelation conditions (Equations (2)–(4)). Thus, the definition of an expansion is mostly relevant where the probability density of the input is noticeable, and not on the whole possible input domain \mathbb{R}^I .

The polynomial expansion (polynomial basis), $PExp(\mathbf{x}) \stackrel{\text{def}}{=} (1, I(\mathbf{x}), QM(\mathbf{x}), CM(\mathbf{x}), \dots)$, describes a great function space that allows fitting to any function at an arbitrary number of points. In practice, infinite expansions cannot be used, and they should be truncated. In this paper, we truncate by keeping functions up to a given degree. Clearly, $QExp$ and $CExp$ are just truncated versions of $PExp$. The main disadvantage of $PExp$ is that its functions (e.g. $x_1^2 x_2^3$) typically grow at a fast pace resulting in bad function fitting for test samples that are too different from training samples.

A. Generation of Expansions

Polynomial expansions ($PExp$, $QExp$, $CExp$, etc.) are very useful because they allow the derivation of various families of expansions with normalization properties that improve generalization. Here we propose two of them:

- 1) Let $q(\mathbf{x})$ be a positive function. Then, $PExp(\mathbf{x})/q(\mathbf{x}) \stackrel{\text{def}}{=} (1/q(\mathbf{x}), x_1/q(\mathbf{x}), \dots, x_1^2/q(\mathbf{x}), \dots)$ is a normalized polynomial expansion. In contrast to $PExp$, in this expansion more functions do not grow more than linearly (e.g. $x_1^2/(1 + \|\mathbf{x}\|^2)$).
- 2) Let $m(\mathbf{x})$ be a bijective mapping of vectors. Then, $PExp(m(\mathbf{x}))$ is a useful expansion. An example of such expansion is $PExp(\mathbf{x}^{*k})$, where $\mathbf{x}^{*k} \stackrel{\text{def}}{=} x_1^{*k}, \dots, x_I^{*k}$,

and $x_i^{*k} \stackrel{\text{def}}{=} \text{sgn}(x_i)|x_i|^k$ is an invertible sign-preserving exponential function of x_i , for $k > 0$.

Of course, more expansions can be defined by starting from an expansion other than $PExp$ or by using other functions for $q(\mathbf{x})$ or $m(\mathbf{x})$. We introduce each one of the expansions evaluated in this paper in Table I, including truncated versions of $PExp(\mathbf{x})/q(\mathbf{x})$ and $PExp(\mathbf{x}^{*k})$.

V. HEURISTIC EVALUATION OF EXPANSIONS

For concrete data, a good choice of an expansion for SFA (one achieving high slowness) depends on the peculiarities of the data (number of samples, dimensionality, type of data, etc.). If the slow parameter can be linearly extracted from the input data, the identity expansion $I(\mathbf{x})$ is already optimal, otherwise a more complex expansion can be useful. Different expansions could be tested for each SFA node of a network to locally achieve top slowness. However, such an optimization procedure would be computationally expensive because it requires training each node for every particular expansion. Therefore, here we propose a heuristic method to evaluate expansions without requiring training of SFA for any particular data. This method is intended to aid in the design of good expansions for a great variety of natural problems.

In a concrete execution of SFA, ideally the outputs ordered in decreasing slowness are equal to the optimal free responses (ignoring polarity w.l.o.g.): $c_1(t) = \sqrt{2} \cos(t)$, $c_2(t) = \sqrt{2} \cos(2t)$, $c_3(t) = \sqrt{2} \cos(3t)$, etc., for $0 \leq t \leq \pi$. That is, in an ideal case SFA should compute cosine functions of increasing frequencies solely from the samples $\mathbf{x}(t)$. When appropriate expansions are employed, the computation of noisy versions of these signals has been empirically verified, particularly at the top nodes [8], [9].

In a simple model, we approximate the input signals of a node (e.g., extracted by a previous node) as $o_i(t) = a_i \cdot (c_i(t) + n_i(t))$ where a_i is a normalization constant assuring unit variance and $n_i(t)$ is Gaussian noise with variance σ_i^2 and zero mean. Similarly, the outputs are approximated as $o'_j(t) = a'_j \cdot (c_j(t) + n'_j(t))$. The input to a node in a network is typically the concatenation of the outputs of many nodes, but for simplicity we view the input as if it only originated from a single one. In practice, noise variance is smaller for the slowest changing signals and larger for the fastest changing signals. For simplicity, we consider here all variances σ_i^2 to be constant and equal to σ^2 . It has been also observed that in hierarchical networks the extracted signals at the lower layers are typically much more noisy and distorted than at the top layers, thus one should consider different variances σ^2 for different layers.

Roughly speaking, it can be conceived that an SFA network attempts to compute $c_j(t)$, for $1 \leq j \leq J$, from the training data. This computation is performed on each layer yielding very crude estimations at the first layers and more accurate estimations at the top node. Let us focus on a single node generating outputs o'_j , for $1 \leq j \leq J$. Notice that, although the input o_i is decorrelated with the input $o_{i'}$ for $i \neq i'$, these signals are not independent. The signal c_1 (or a noisy version of it) can be seen as distributedly coded in the signals

o_1, \dots, o_I . Of course, o_1 provides more information about c_1 than any other extracted signal, however, the other components still contain useful information that might be exploited after a non-linear transformation to better approximate c_1 . Thus, in hierarchical processing, the non-linearities should be powerful enough to allow the computation of versions of c_j , for $1 \leq j \leq J$, and particularly of c_1 , as faithfully as possible from the input signals o_1, \dots, o_I .

Based on previous empirical observations, we postulate that a good expansion Exp should be capable of allowing the approximate computation of a series of functions. By \mathbf{P} : $(x_1, \dots, x_I) \mapsto y$ we denote the problem of approximating a function $y(t)$ from the input data $\mathbf{x}(t) \stackrel{\text{def}}{=} (x_1(t), \dots, x_I(t))$. In the context of a particular expansion Exp for SFA, we approximate $y(t)$ linearly as $\ell(t) \stackrel{\text{def}}{=} \mathbf{w}^T(\mathbf{1}, Exp(\mathbf{x}(t)))^T$, where \mathbf{w} is a weight vector computed through linear regression to best fit y and the constant 1 compensates for feature means different from zero. Depending on the expansion, the approximation will be more or less accurate, and this can be measured in terms of the Root Mean Squared Error (RMSE) between $y(t)$ and $\ell(t)$. We propose the following function approximation problems to test the computational richness of an expansion.

- 1) **P1**: $o_1 \mapsto o_1$; (Identity function) That is, linearly approximate/compute o_1 from (the expansion of) o_1 .
- 2) **P2**: $o_1 \mapsto c_2$; (Frequency doubler) Compute c_2 from o_1 .
- 3) **P3**: $o_1, o_2 \mapsto c_3$; Approximate c_3 from o_1 and o_2 .
- 4) **P4**: $o_1 \mapsto \sin(t)$; Approximate $\sin(t)$ from a noisy version of $\cos(t)$.
- 5) **P5**: $o_2, o_3 \mapsto c_1$; Approximate of c_1 from higher frequency noisy harmonics o_2 and o_3 .
- 6) **P6**: $o_2, o_3, o_4 \mapsto c_1$. Similar to **P5**.
- 7) **P7**: $o_3, o_4, o_5 \mapsto c_1$. Also similar to **P5**.

Each one of these problems is motivated by specific reasons. Suppose that a particular node extracts a very slow signal, then a good performance on **P1** ensures that a subsequent node in the network is at least able to preserve such a signal after expansion. A good performance on **P2** and **P3** indicates that a node is capable of generating (or reducing the noise of) higher frequency harmonics given access to lower frequency ones. **P4** ensures that the output signals could be freed from sine components if needed. A good performance on **P5**, **P6**, and **P7** indicates that lower frequency harmonics might be generated from higher frequency ones.

The problems above are representative and subsume a larger number of problems, for instance $o_1 \mapsto o_1$ is equivalent to $o_i \mapsto o_i$, whereas $o_2, o_3 \mapsto c_1$ is equivalent to $o_{2i}, o_{3i} \mapsto c_i$, and $o_1 \mapsto \sin(t)$ is equivalent to $o_i \mapsto \sin(it)$, where $i \in \mathbb{N}$.

In the context of SFA, *overfitting* might be defined as the difference (or ratio) between the delta values of the slowest signals extracted from training and test data, but computing this requires training of SFA. We propose a simple and practical way of measuring overfitting for non-linear expansions. The basic idea is to use the expanded data to learn a random zero-mean signal through linear regression.

Name	Definition	Example of functions contained	Expanded dimensionality
Identity	$I(\mathbf{x})$	x_1, x_2, x_I	I
$QExp$	$I(\mathbf{x}), QM(\mathbf{x})$	$x_1, x_2, x_I, x_1^2, x_1x_2, x_2^2, x_I^2$	$I(I+3)/2$
$CExp$	$I(\mathbf{x}), QM(\mathbf{x}), CM(\mathbf{x})$	$x_1, x_I, x_1^2, x_1x_2, x_I^2, x_1^3, x_1x_2x_3, x_2^2x_4, x_I^3$	$I(I^2+6I+11)/6$
$QExp/q_2$	$q_2(\mathbf{x}) = (1 + \ \mathbf{x}\ ^2)$	$x_1/q_2(\mathbf{x}), x_1x_2/q_2(\mathbf{x}), x_I^2/q_2(\mathbf{x})$	$I(I+3)/2$
$CExp/q_3$	$q_3(\mathbf{x}) = (1 + \ \mathbf{x}\ ^2)$	$x_1/q_3(\mathbf{x}), x_1x_2/q_3(\mathbf{x}), x_I^2/q_3(\mathbf{x}), x_I^3/q_3(\mathbf{x})$	$I(I^2+6I+11)/6$
$QExp^{*0.4}$	$QExp(\mathbf{x}^{*0.4})$	$x_1^{*0.4}, x_2^{*0.4}, x_1 ^{0.8}, (x_1x_2)^{*0.4}, x_I ^{0.8}$	$I(I+3)/2$
$CExp^{*0.3}$	$CExp(\mathbf{x}^{*0.3})$	$x_1^{*0.3}, x_2^{*0.3}, x_1 ^{0.6}, (x_1x_2)^{*0.3}, x_1^{*0.9}, (x_1x_2x_3)^{*0.3}, x_I^{*0.9}$	$I(I^2+6I+11)/6$
$0.8Exp$	$I(\mathbf{x}), \mathbf{x} ^{0.8}$	$x_1, x_2, x_I, x_1 ^{0.8}, x_2 ^{0.8}, x_I ^{0.8}$	$2I$
$SExp$	$I(\mathbf{x}), \mathbf{x}^2$	$x_1, x_2, x_I, x_1^2, x_2^2, x_I^2$	$2I$

TABLE I
LIST OF EXPANSIONS EVALUATED IN THIS PAPER.

Since the samples and the random labels are unrelated, the best possible fit is to always output the constant 0 for any data sample. Given that a zero complexity algorithm optimally solves this problem (by always outputting 0), any deviation of the learned output from 0 for test data is purely due to overfitting. More concretely, let $\mathbf{x}(t) = (o_1(t), \dots, o_I(t))$ be the simulated input to SFA (before expansion), and let $s(t) = \Pi(c_1(t))$ be a randomly time-shuffled version of $c_1(t)$ (shuffling is possible because time is discrete when using samples). We learn a function f approximating $s(t)$ from $\mathbf{x}(t)$, as $f(\mathbf{x}) \stackrel{\text{def}}{=} \mathbf{w}^T Exp(1, s(t)) \approx s(t)$, where \mathbf{w} is found using linear regression. Let $\mathbf{x}'(t)$ be the test data (with the same distribution as $\mathbf{x}(t)$), and $y'(t) \stackrel{\text{def}}{=} f(\mathbf{x}'(t))$ be the estimated labels for the test data. Then, overfitting is defined as:

$$Ov(Exp, \mathbf{x}(t), \mathbf{x}'(t)) \stackrel{\text{def}}{=} \sqrt{\langle (y')^2 \rangle}, \quad (5)$$

where $\langle \cdot \rangle$ indicates temporal averaging.

We present a metric to quantify *outlier amplification*, where extreme/atypical outputs computed at different layers of the network are magnified, which is particularly notorious for test data (cf. Section III-E). Roughly speaking, the metric compares the rareness of extreme samples in the data before and after expansion. Let $\mathbf{X} \stackrel{\text{def}}{=} (\mathbf{x}(1), \dots, \mathbf{x}(T))$ be the training samples, and Exp the expansion we want to evaluate. Let $\mathbf{X}' \stackrel{\text{def}}{=} \mathcal{W}(X)$ and $\mathbf{Y}' \stackrel{\text{def}}{=} \mathcal{W}(Exp(X))$ be whitened versions of \mathbf{X} and the expanded samples, respectively. Let \mathbf{S} be a set of vectors (samples or expanded samples), and $A(\mathbf{v}) \stackrel{\text{def}}{=} \sqrt{\sum_{i=1}^k (v_i)^2} / N$ be a scaled norm of a vector $\mathbf{v} \stackrel{\text{def}}{=} (v_1, \dots, v_k)$. A vector $\mathbf{v} \in \mathbf{S}$ is called an *outlier* if $A(\mathbf{v})$ is in the top 5% among all vectors $\mathbf{v}' \in \mathbf{S}$, and $\mathbf{O}_{\mathbf{S}} \subset \mathbf{S}$ denotes the set of all outliers of \mathbf{S} . Outlier amplification is defined as

$$O_{amp}(Exp, \mathbf{X}) \stackrel{\text{def}}{=} \frac{\sum_{\mathbf{y} \in \mathbf{O}_{\mathbf{Y}'}} A(\mathbf{y})}{\sum_{\mathbf{x} \in \mathbf{O}_{\mathbf{X}'}} A(\mathbf{x})} \quad (6)$$

and can be interpreted as the fraction between the scaled norm of the samples after and before expansion, in both cases including a whitening step.

Due to whitening, this metric is invariant to linear transformations, and due to the use of a scaled norm, it provides some robustness to the expanded dimensionality. Various definitions of outliers exist, such as distance- or density-based [12], which

might lead to alternative definitions of $A(\mathbf{v})$. For instance, $A(\mathbf{v})$ could be the distance to the nearest neighbour of \mathbf{v} . However this might be less efficient for a high number of samples and dimensions.

VI. EXPERIMENTAL RESULTS

To evaluate the expansions, we considered three scenarios that resemble the conditions of the inputs to SFA nodes at different stages of the network. In each scenario, the input signals before expansion $o_i(t)$ are modelled as cosine functions of increasing frequency including additive Gaussian noise with variance ranging from a large variance $\sigma^2 = 0.75^2$ for the bottom layers, to $\sigma^2 = 0.5^2$ for the middle layers, and a small variance $\sigma^2 = 0.25^2$ for the top layers.

The performance for problems **P2** to **P7** and Ov was measured with respect to test data generated with the same distribution as the training data. For **P1** and O_{amp} , only training data was used because the corresponding problem does not demand test data. In order to obtain realistic results, parameters were set similarly to those used in previous experiments with real data and SFA networks; 10,000 simulated samples were used and the data dimensionality before expansion was set to 60.

To validate the heuristics, the expansions were tested in a hierarchical network to learn and estimate the horizontal position of a face in face photographs. Similarly as in [10], the resulting estimation was robust to the face's size, vertical position, expression, illumination, background, etc.

Table II presents the performances of each expansion on each problem, metric, and on real data. Results were averaged over $\sigma^2 \in \{0.25^2, 0.5^2, 0.75^2\}$, except for O_{amp} , where only $\sigma^2 = 0.75^2$ was relevant. The expansions $QExp/q_2$ and $QExp^{*0.4}$ gave a poor performance on **P1**. This could be corrected by including the identity function in them. For problems **P2–P7**, quadratic expansions provided good performance in general, whereas cubic expansions performed excellent, however the latter are inappropriate in this setup because the expanded dimensionality is too high compared to the number of samples (consequently, they would be strongly penalized by Ov). In contrast, expansions $0.8Exp$ and $SExp$ are only good at **P2** and **P4**. For the metrics Ov and O_{amp} , the expanded dimension for the cubic expansions (39,711) was larger than the number of samples and would cause extreme

	Identity	$QExp$	$CExp$	$QExp/q_2$	$CExp/q_3$	$QExp^{*0.4}$	$CExp^{*0.3}$	$0.8Exp$	$SExp$
P1	0.0000	0.0000	0.0000	0.3777	0.1122	0.2648	0.0122	0.0000	0.0000
P2	1.0000	0.7810	0.7810	0.7348	0.7815	0.7431	0.7498	0.7431	0.7810
P3	1.0000	0.7071	0.7032	0.6706	0.6421	0.7225	0.6659	1.0001	1.0001
P4	1.0000	0.7933	0.7934	0.7725	0.8259	0.7766	0.7853	0.7766	0.7933
P5	1.0000	0.8209	0.7755	0.8084	0.7118	0.8202	0.7251	1.0000	1.0000
P6	1.0000	0.7137	0.6962	0.6603	0.6134	0.6673	0.6099	1.0000	1.0000
P7	1.0000	0.7163	0.6599	0.6608	0.5518	0.6607	0.5574	1.0001	1.0000
O_v	0.9969	1.0959	$\gg 1$	1.0946	$\gg 1$	1.0938	$\gg 1$	0.9970	0.9977
O_{amp}	1.0000	1.0520	—	0.9631	—	0.9793	—	0.9854	1.0872
$\Delta(y_1)$ training	1.0460	0.1844	—	0.0588	—	0.0283	—	0.2829	0.5176
$\Delta(y_1)$ test	1.0701	275.3278	—	0.1555	—	0.1154	—	0.2855	12.7479
Estimation error [pixels]	17.96	27.54	—	5.85	—	6.18	—	6.16	16.44

TABLE II

PERFORMANCE OF EACH EXPANSION ON EACH PROBLEM, METRIC, AND ON REAL DATA. THE PERFORMANCE FOR PROBLEMS **P1–P7** AND THE ESTIMATION ERROR ARE REPORTED IN TERMS OF THE RMSE. FOR ALL TESTS, SMALLER VALUES ARE BETTER.

overfitting and computational complications. Thus, O_v and O_{amp} were not evaluated for cubic expansions. O_v strongly depended on the expanded dimensionality and the particular expansion played a less important role. In the case of outlier amplification, the opposite was observed: The concrete shapes of the expansion functions are determinant, while the expanded dimensionality is less important. Regarding real data, it was verified that a good performance on problems **P1–P7** is related to slower extracted features for training data. It was also verified that the metric O_{amp} is a good predictor of outlier amplification for test data (which occurred for $QExp$ and $SExp$). For $O_{amp} \leq 1$, O_v showed correlation with overfitting when seen as the ratio between delta values for test and training data.

VII. CONCLUSION

In this paper, a method for the heuristic evaluation of expansions used in non-linear hierarchical SFA was proposed. It consists of tests on seven different problems and two metrics. It was verified experimentally with a concrete data set that a good performance on the problems is related to the richness of the spanned feature space and, thus, to the computational capabilities of SFA. The metric O_{amp} provides a simple way to predict (and avoid) the occurrence of the outlier amplification phenomenon. The metric O_v in conjunction with O_{amp} provides a rough way to predict high overfitting.

The main advantage of the method is that it allows for very fast evaluation of expansions because it does not require training of SFA or the use of real data. The method can be easily applied to efficiently tune parametrized expansions, such as $QExp/(a + ||x||^b)$ for good choices of a and b . For optimizing slowness on real data, the heuristics can be used to find an initial small set of good expansions and then refine them using a concrete hierarchical SFA network.

The method has the disadvantage that it is less realistic for the lower layers of the hierarchical network, in which the extracted signals typically do not resemble noisy cosine functions, but the results are still consistent with real data.

One open problem is the computation of a single numerical score from all individual tests, which would ease the compar-

ison of expansions. In future work, we are interested in using this method for the guided design of powerful expansions with good scores on all considered problems and metrics, but with a small expanded dimensionality. For instance, we are interested in selecting a small subset of the functions of a normalized cubic expansion (e.g., of size of about $20 \cdot I$) providing a complex feature space and allowing for efficient training.

The results on outlier amplification suggest the use of normalized expansions with functions that do not grow much more than linearly for large values of the input (e.g., in contrast to $QExp$) to achieve better performances.

REFERENCES

- [1] I. K. Fodor, "A survey of dimension reduction techniques," *Communications of The ACM*, 2002.
- [2] R. A. Fisher, "The use of multiple measurements in taxonomic problems," *Annals Eugen.*, vol. 7, pp. 179–188, 1936.
- [3] L. Wiskott and T. Sejnowski, "Slow feature analysis: Unsupervised learning of invariances," *Neural Computation*, vol. 14, no. 4, pp. 715–770, 2002.
- [4] P. Berkes and L. Wiskott, "Slow feature analysis yields a rich repertoire of complex cell properties," *J. Vis.*, vol. 5, no. 6, pp. 579–602, 7 2005.
- [5] L. Wiskott, "Slow feature analysis: A theoretical analysis of optimal free responses," *Neural Computation*, vol. 15, no. 9, pp. 2147–2177, Sep. 2003.
- [6] A. Bray and D. Martinez, "Kernel-based extraction of slow features: Complex cells learn disparity and translation invariance from natural images," in *Proc. Advances in neural information processing systems (NIPS'03)*, vol. 15. Cambridge, MA: MIT Press, 2003, p. 253260.
- [7] R. Vollgraf and K. Obermayer, "Sparse optimization for second order kernel methods," in *Proc. IJCNN'06*, 2006, pp. 145 – 152.
- [8] M. Franzius, N. Wilbert, and L. Wiskott, "Invariant object recognition with slow feature analysis," in *Proc. 18th Intl. Conf. on Artificial Neural Networks*, ser. LNCS, vol. 5163. Springer, Sep. 2008, pp. 961–970.
- [9] A. Escalante and L. Wiskott, "Gender and age estimation from synthetic face images with hierarchical slow feature analysis," in *International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems*, 2010, pp. 240–249.
- [10] N. M. Mohamed and H. Mahdi, "A simple evaluation of face detection algorithms using unpublished static images," in *10th International Conference on Intelligent Systems Design and Applications, ISDA*, 2010, pp. 1–5.
- [11] M. Franzius, H. Sprekeler, and L. Wiskott, "Slowness and sparseness lead to place, head-direction, and spatial-view cells," *PLoS Computational Biology*, vol. 3, no. 8, p. e166, August 2007.
- [12] S. Subramaniam, T. Palpanas, D. Papadopoulos, V. Kalogeraki, and D. Gunopulos, "Online outlier detection in sensor data using non-parametric models," in *Proc. 32nd international conference on Very large data bases*, ser. VLDB '06. VLDB Endowment, 2006, pp. 187–198.