

Dictionary Learning with Equiprobable Matching Pursuit

Fredrik Sandin

EISLAB, Luleå University of Technology, Sweden
Email: fredrik.sandin@ltu.se

Sergio Martin-del-Campo

SKF-LTU University Technology Center and
EISLAB, Luleå University of Technology, Sweden

Abstract—Sparse signal representations based on linear combinations of learned atoms have been used to obtain state-of-the-art results in several practical signal processing applications. Approximation methods are needed to process high-dimensional signals in this way because the problem to calculate optimal atoms for sparse coding is NP-hard. Here we study greedy algorithms for unsupervised learning of dictionaries of shift-invariant atoms and propose a new method where each atom is selected with the same probability on average, which corresponds to the homeostatic regulation of a recurrent convolutional neural network. Equiprobable selection can be used with several greedy algorithms for dictionary learning to ensure that all atoms adapt during training and that no particular atom is more likely to take part in the linear combination on average. We demonstrate via simulation experiments that dictionary learning with equiprobable selection results in higher entropy of the sparse representation and lower reconstruction and denoising errors, both in the case of ordinary matching pursuit and orthogonal matching pursuit with shift-invariant dictionaries. Furthermore, we show that the computational costs of the matching pursuits are lower with equiprobable selection, leading to faster and more accurate dictionary learning algorithms.

I. INTRODUCTION

Sensors like cameras, microphones and accelerometers typically generate redundant information because the resulting data have an underlying structure. That is why most observed phenomena can be accurately approximated with mathematical relationships in the form of physical laws, and also why such data can be compressed using algorithms that identifies and approximates patterns. Redundancies in the data make the processes of communicating, analyzing and storing information inefficient, thus constraining the domain of feasible applications. The problem to identify the structure of signals and thereby derive succinct representations that are both compact and informative is a computational challenge, which limits the efficacy of sensor systems.

Sparse representation models [1], [2], [3], [4] and related machine learning algorithms [5], [6] have proven remarkably successful at extracting useful information from complex high-dimensional signals, for example in the context of denoising [7], under-determined source separation [8], [9], compressed sensing [10], super-resolution sensing [11], and classification [12]. In a sparse model, the signal is typically described as a linear combination of elementary functions called atoms, which can either be predefined (if an appropriate generative model for the signal class is known) or learned from a training

signal. The goal is to select or learn the atoms so that the model residual is minimum for a certain sparsity of the representation, or to maximize the sparsity for a certain tolerance on the model residual. Atoms typically have unit norm and the set of atoms defines a dictionary, which can be subject to further constraints like temporal or spatial translation invariance.

Fourier and wavelet analysis are two examples where predefined dictionaries are used. Models based on predefined dictionaries have enabled derivation of closed-form mathematical results and fast algorithms that are widely used, but such approaches are simplistic compared to the complexity of the underlying natural phenomena. The dictionary learning approach is based on the hypothesis that complex signals can be more accurately modeled by extracting environmentally matched atoms from the signal. However, the dictionary learning problem is NP-hard and it is also hard to find approximate solutions near the optimal sparsity level [13].

The development of dictionary learning methods was stimulated by results presented in the mid '90s by Olshausen and Field [14], [15], which demonstrate that atoms similar to the receptive fields of cells in visual cortex can be learned from natural images by imposing a few general optimization conditions, including sparsity and statistical independence of atoms. This demonstrates that some aspects of the low-level functions in the visual system can be explained by a few general computational principles, and that elementary structures of such complex signals can be automatically uncovered from examples. Since that time several probabilistic dictionary learning and sparse coding methods have been developed [2], [5], aiming for a dictionary that either maximizes the likelihood of the data, as for example in [16], or the posterior probability of the dictionary, as in [17]. Recent developments include extensions of dictionary learning methods to distributed systems [18] and low-power hardware [19].

There is a knowledge gap between receptive field models and the observed function of neural networks in biological sensory systems. For instance, neurons demonstrate a form of homeostatic adaptation of the functional properties of the network to the ongoing changes in the statistical structure of the sensory input [20], [21], [22], which may be related to optimal encoding [23] of the input by dynamic adaptation of the receptive fields. Furthermore, there is a notion that homeostatic mechanisms serve to maintain the dynamics of cortical networks at a critical point [24], [25] where the dynamic

range and information processing capacity are optimal [26], [27]. Can we learn something new about dictionary learning by incorporating and studying such homeostatic regulation mechanisms?

In this paper we focus on greedy algorithms for dictionary learning, where the approximation error is reduced iteratively by sparse decomposition of the signal followed by gradient optimization of the atoms in the dictionary. Our aim is to study a basic implementation of homeostatic regulation where all atoms are enforced to occur with the same probability on average, which in terms of neurons and receptive fields mean that neurons fire with the same average probability.

The starting point is the dictionary learning method introduced by Smith and Lewicki [28], where the sparse code is generated with Matching Pursuit (MP) [29], [30] and the shift-invariant dictionary is updated with probabilistic gradient ascent on the likelihood of the data [16]. In principle this model corresponds to a particular type of recurrent convolutional neural network with max pooling (cf. Fig. 5 in [15]), which we extend here with a homeostatic regularization mechanism. Typically a subset of the atoms are selected in the MP for one particular signal, while some atoms are rarely selected and do not take significant part in the gradient-based dictionary update process. We implement homeostatic regulation by enforcing equiprobable atom selection in the MP, which implies that each atom is equally likely to occur in the linear combination on average, see Figure 1. In addition to MP we consider dictionary learning with Orthogonal MP (OMP) [30], [31], which in the case of shift-invariant dictionaries is applicable to high-dimensional signals in the form of local OMP [32].

We study the effects of equiprobable selection on the learned dictionary and sparse representation of music and birdsong. With equiprobable selection all atoms adapt to the training signal, see Figure 2. Our main finding is that the entropy and reconstruction accuracy are higher for equiprobable sparse representations, and that the computational cost is reduced by the equiprobability constraint. Equiprobable selection can be generalized to other greedy dictionary learning methods and these results motivate further investigations of dictionary learning models incorporating more realistic homeostatic regulation mechanisms observed by neuroscientists.




Figure 1. Two equally sparse representations of 0.15 seconds of 44.1 kHz rock music obtained with two different matching pursuit algorithms and learned dictionaries. Discs represent onsets of atoms and disc sizes are proportional to the atom coefficients. The panel on the (left-hand) right-hand side is obtained with (non-equiprobable) equiprobable MP. With equiprobable MP all atoms adapt to the signal and contributes to the sparse representation.

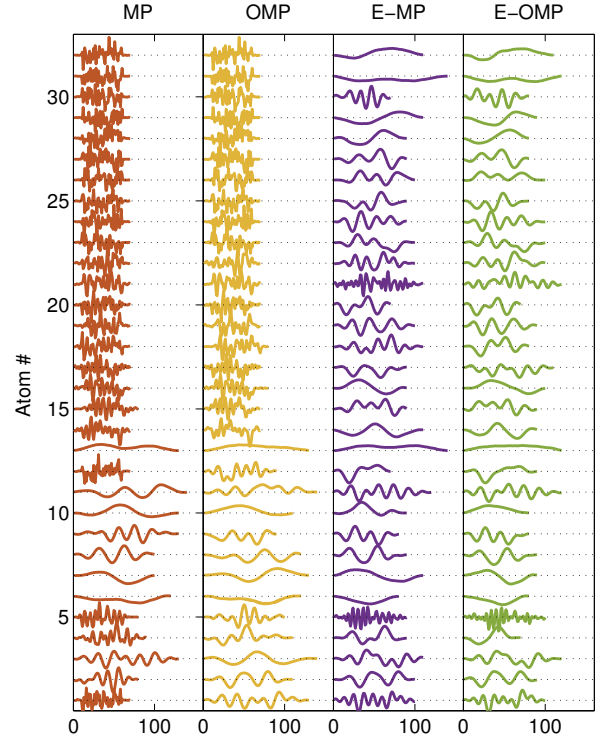


Figure 2. Dictionaries of 32 atoms learned from 44.1 kHz rock music using the dictionary learning methods described below, which are based on four different matching pursuit algorithms (MP, OMP, E-MP and E-OMP). Initially the four dictionaries are identical and include atoms that are sampled from a Gaussian distribution with zero mean.

II. DICTIONARY LEARNING METHOD

The signal, $x(t)$, is modelled as a linear superposition of waveforms, φ_i , with compact support and additive noise

$$x(t) = \sum_{i=1}^M \sum_{j=1}^{N_i} a_{i,j} \varphi_i(t - \tau_{i,j}) + \epsilon(t). \quad (1)$$

The functions φ_i are shift-invariant *atoms* that represent elementary waveforms of the signal and M is the number of different atoms in the model. The variable $\epsilon(t)$ represents the model residual, including noise. The variable N_i refers to the number of instances of atom φ_i , and the temporal position and amplitude of the j -th instance of atom φ_i are denoted by $\tau_{i,j}$ and $a_{i,j}$, respectively. The set of M atoms defines a dictionary

$$\Phi = \{\varphi_1, \dots, \varphi_M\}, \quad (2)$$

which we want to adapt to the signal. In principle $x(t)$ and φ_i can be multidimensional, but here we limit the numerical experiments to scalar signals.

Eq. (1) defines a sparse approximation of $x(t)$ if the number of terms are few compared to the number of samples of $x(t)$ and $\epsilon(t)$ is small compared to $x(t)$, which requires that the dictionary Φ is adapted to the signal. Thus, the dictionary learning problem is to calculate a set of φ_i that minimizes $\epsilon(t)$ under ℓ_0 regularization of the coefficient matrix. This problem is NP-hard [13] and cannot be solved explicitly. Instead, a greedy algorithm that reduces the approximation error in a

two-step iteration process is used [33]: *A) Encoding step*; optimize the sparse representation of the signal $x(t)$ with MP or OMP and a constant dictionary Φ . *B) Learning step*; update the dictionary Φ using the sparse representation and model residual so that the approximation error is reduced.

The encoding step is by itself an iterative process that terminates at some predefined sparsity or tolerance on the model residual. Each iteration of the encoding process includes two steps:

- $A_1)$ *Atom selection* – Find the atom φ_i and offset $\tau_{i,j}$ that maximizes the cross-correlation with the model residual, $r_k(t)$, and calculate the corresponding $a_{i,j}$.
- $A_2)$ *Residual update* – Update the residual by subtracting the contribution from the selected atom, $r_{k+1}(t) = r_k(t) - a_{i,j}\varphi_i(t - \tau_{i,j})$ with $r_0(t) = x(t)$.

The accuracy of the resulting signal approximation depends on the complex interplay between the atom selection rule and the dictionary learning process.

MP optimizes the parameters $\tau_{i,j}$ and $a_{i,j}$ of the most recently selected atom, while local OMP [32] re-optimizes all $a_{i,j}$ for selected atoms with overlapping support in each iteration (OMP compensates for the interference between atom instances). In the following we consider both MP and OMP when introducing homeostatic regulation. Our implementation of MP and OMP are based on efficient computational methods like those described in [32] and [34].

The atom selection rule defined above is optimal for sparse decomposition of a signal with a constant dictionary [30], but it does *not* imply optimal dictionary learning. For example, if some atoms are more frequently selected than others the entropy of the resulting sparse representation is expected to be suboptimal, potentially leading to information loss and lower reconstruction accuracy. Furthermore, atoms that are rarely selected mainly contribute to the model complexity. Therefore, we introduce a modified atom selection rule that regulates the average probability for each atom to be selected:

- $A_1^*)$ *Equiprobable atom selection* – Find the atom φ_i and offset $\tau_{i,j}$ that maximizes the cross-correlation with the model residual, $r_k(t)$, under the constraint $P(\varphi_i) < p$ and calculate the corresponding $a_{i,j}$.

The atom selection constraint $P(\varphi_i) < p$ can be applied to both MP and OMP for dictionary learning purposes. The constrained matching pursuit terminates when no atom is selected, at which point all atoms occur with the same probability p . We refer to the resulting “equiprobable” MP and OMP as E-MP and E-OMP, respectively. MP and OMP are defined by steps A_1 , A_2 and B above, while E-MP and E-OMP are defined by A_1^* , A_2 and B . The resulting four matching pursuits are summarized in Table I and Algorithm 1, which is a straightforward extension of the local OMP algorithm presented in [32] to equiprobable atom selection.

The function $P(\varphi_i, \Phi_k)$ is the probability for each shift-invariant atom to be selected, which is defined here by the relative number of occurrences of each atom in the subdictionary Φ_k of selected atoms and offsets. For E-MP and E-OMP the

Algorithm 1 Matching pursuit

```

1: function match( $\Phi, x, p$ )
2:    $r_0 = x$ 
3:    $a_0 = 0$ 
4:    $\tau_0 = 0$ 
5:    $\Phi_0 = \emptyset$ 
6:   for  $k = 1$  to  $I$  do
7:      $\{\varphi_k, t_k\} = \operatorname{argmax}_{\varphi_i, \tau_{i,j}} \langle r_{k-1}, \Phi \rangle$ 
       subject to constraint
8:      $\Phi_k = \Phi_{k-1} \cup \varphi_k$ 
9:      $\Psi_k = \text{neighborhood}$ 
10:     $\chi_k = (\Psi_k^* \Psi_k)^{-1} \Psi_k^* r_{k-1}$ 
11:     $a_k = a_{k-1} + \chi_k$ 
12:     $\tau_k = \tau_{k-1} + t_k$ 
13:     $r_k = r_{k-1} - \Psi_k \chi_k$ 
14:  end for
15:  return  $\{a_I, \tau_I, r_I\}$ 
16: end function

```

total number of iterations, I , is defined so that $P(\varphi_i, \Phi_I) \equiv p$ for each shift-invariant atom, φ_i , which implies that each atom is included in the sparse approximation with the same probability, p , on average. For MP and OMP the total number of iterations, I , is defined to be equivalent to the number of E-MP and E-OMP iterations. Thus, the resulting MP, OMP, E-MP and E-OMP approximations are equally sparse and can be compared in terms of reconstruction error, denoising error etc. With MP and OMP the shift-invariant atoms typically occur with different probabilities, resulting in a learned dictionary where only a subset of the atoms adapt to the training signal.

The dictionary learning problem is to infer the set of atomic waveforms, φ_i , in the dictionary, Φ , so that the matching pursuit results in a sparse representation with low residual. A computationally feasible formulation of this problem can be obtained by rewriting Eq. (1) in probabilistic form

$$p(x|\Phi) = \int p(x|a, \Phi) p(a) da \quad (3)$$

$$\approx p(x|\hat{a}, \Phi) p(\hat{a}), \quad (4)$$

where \hat{a} is the maximum a posteriori (MAP) estimation of a ,

$$\hat{a} = \arg \max_a p(a|x, \Phi) = \arg \max_a p(x|a, \Phi) p(a), \quad (5)$$

that is generated by the matching pursuit [16], [28], [29]. Furthermore, we assume that the noise term, $\epsilon(t)$, in Eq. (1) is Gaussian. Thus, the data likelihood, $p(x|a, \Phi)$, is also

Table I
LIST OF MATCHING PURSUITS DEFINED BY ALGORITHM 1.

METHOD	constraint	neighborhood
MP	$\varphi_i \in \Phi$	φ_k
OMP	$\varphi_i \in \Phi$	Φ_k
E-MP	$\varphi_i \in \Phi, P(\varphi_i, \Phi_k) < p$	φ_k
E-OMP	$\varphi_i \in \Phi, P(\varphi_i, \Phi_k) < p$	Φ_k

Gaussian and takes the form

$$p(x|a, \Phi) \approx \exp\left(-\frac{\|x - a\Phi\|^2}{2\sigma_\epsilon^2}\right), \quad (6)$$

where

$$\|x - a\Phi\|^2 = \left\|x - \sum_{i=1}^M \sum_{j=1}^{N_i} a_{i,j} \varphi_i(t - \tau_{i,j})\right\|^2, \quad (7)$$

and σ_ϵ^2 is the variance of the noise. Note that x , a and Φ are matrices in these probabilistic expressions, and that the dictionary, Φ , includes all possible shifts of each atom φ_i .

Under these assumptions the atoms can be optimized by performing gradient ascent on the approximate log data probability [28]. It follows from Eq. (4) that

$$\frac{\partial}{\partial \varphi_i} \log(p(x|\Phi)) = \frac{\partial}{\partial \varphi_i} [\log(p(x|\hat{a}, \Phi)) + \log(p(\hat{a}))]. \quad (8)$$

By taking the derivative and substituting the likelihood term with Eq. (6) this becomes

$$\begin{aligned} & \frac{\partial}{\partial \varphi_i} \log(p(x|\Phi)) \\ &= \frac{-1}{2\sigma_\epsilon^2} \frac{\partial}{\partial \varphi_i} \left\|x - \sum_{i=1}^M \sum_{j=1}^{N_i} a_{i,j} \varphi_i(t - \tau_{i,j})\right\|^2 \end{aligned} \quad (9)$$

$$= \frac{1}{\sigma_\epsilon^2} \sum_j a_{i,j} [r_I]_{\tau_{i,j}}. \quad (10)$$

The factor $[r_I]_{\tau_{i,j}}$ represents the model residual coinciding with atom φ_i at temporal position $\tau_{i,j}$. In other words, the shape of each atom is adapted with a weighted average of the residual elements coinciding with the matches identified by the matching pursuit. This is a form of nonlinear Hebbian learning because the atoms are adapted to patterns in the signal that they correlate with.

In order to use the gradient for dictionary learning we introduce a relative steplength parameter, η , and define the gradient ascent update of atom φ_i as

$$\varphi_i \rightarrow \varphi_i + \frac{\eta}{\sigma_\epsilon^2} \sum_j a_{i,j} [r_I]_{\tau_{i,j}}. \quad (11)$$

This implies that the dictionary adaptation rate depends on the activation rate of atoms. We zero-pad all atoms with ten elements and grow each tail in length with ten additional elements if the RMS of the tail exceeds 0.1 of the atom RMS. The resulting dictionary learning method is summarized in Algorithm 2.

The function `randdict(M)` generates a random dictionary of M normalized atoms, where each atom includes fifty elements sampled from a Gaussian distribution with zero mean and two vanishing tails that are ten elements long. Thus, the M different atoms are seventy elements long initially. The `extnorm(φ)` function extends the length of an atom φ with ten vanishing elements whenever the RMS of a tail exceeds the predefined threshold mentioned above, and it also normalizes the atom. Training data is fetched with the `getdata()` function, which allows block-based processing of

Algorithm 2 Dictionary learning

```

1: function dlearn( $M, p, \eta$ )
2:    $n = 0$ 
3:    $r_0 = 0$ 
4:    $\Phi^{(0)} = \text{randdict}(M)$ 
5:   while  $x_{n+1} = \text{getdata}(n, r_n)$  is not empty do
6:      $n = n + 1$ 
7:      $\{a_n, \tau_n, r_n\} = \text{match}(\Phi^{(n-1)}, x_n, p)$ 
8:     for each  $\xi_i$  in  $a_n$  do
9:        $\delta_i = \eta \xi_i[r_n]_{\tau_n} / \text{var}(r_n)$ 
10:       $\varphi_i^{(n)} = \text{extnorm}(\varphi_i^{(n-1)} + \delta_i), \quad \varphi_i^{(n)} \in \Phi^{(n)}$ 
11:    end for
12:  end while
13:  return  $\Phi^{(n)}$ 
14: end function

```

sampled signals using overlapping window functions and part of the former residual, r_{i-1} , of the matching pursuit.

III. RESULTS

We investigate the effects of equiprobable atom selection on the learned dictionary and sparse approximation accuracy with numerical experiments using two different signals. One signal is a 155 seconds long 44.1 kHz rock music track with lyrics [35] and the second signal is a 26 seconds long 48 kHz recording of Zebra Finch song phrases [36]. The dictionary learning method is defined by Eq. (11) with steplength $\eta = 10^{-6}$. We do not observe significant improvements in the resulting model accuracy using other values of η and our qualitative analysis does not depend on finetuning of this hyperparameter. The `getdata()` function is defined so that each block of data (x_n in Algorithm 2) is five seconds long sampled from one random location in the data set. The progression of the dictionary learning protocol is thereby quantified in terms of time rather than epochs.

A. Model accuracy and rate of convergence

In the first experiment we study the signal-to-noise ratio (SNR) of the sparse approximation of rock music with an average atom selection probability of $p = 0.05$, see Figure 3. The resulting dictionaries of 32 atoms are displayed in Figure 2. Equiprobable selection reduces the initial dictionary learning rate, but it also leads to an improved convergence time and accuracy of the sparse model. For example, the accuracy of the E-MP-based approximation exceeds that of the OMP-based approximation after about 1000 seconds of learning. This is remarkable considering that the computational cost of E-MP is lower than both MP and OMP (further details below). With 64 atoms a longer learning time is needed to reach a comparable SNR, but after about 2000 seconds of learning the SNR exceeds that of models with 32 atoms. Note that atoms are shift invariant. Thus, the number of possible sparse representations of a five-second long window is astronomical even if there are only a few atoms in the dictionary.




Figure 3. Reconstruction SNR versus learning time for rock music with an average atom selection probability of $p = 0.05$. Bold (thin) lines are obtained with a dictionary of 32 (64) learned atoms.




Figure 4. Reconstruction SNR versus the average atom selection probability. Dictionaries are learned from 1500 seconds of rock music with an average event probability of 0.05 (dashed line).

B. Effect of varying atom selection probability

Next we study the sensitivity of the model to variations in the average atom selection probability, p . Using dictionaries learned from 1500 seconds of rock music we calculate the SNR of the sparse approximation of rock music for different values of p , see Figure 4. This result demonstrates that the models based on dictionaries learned with E-MP and E-OMP degrades gracefully when the sparsity of the model changes.




Figure 5. Reconstruction SNR versus the relative standard deviation of additive Gaussian noise, σ_n/σ_s , for a track of 44.1 kHz rock music with standard deviation σ_s .

The learned dictionaries generalize to such varying conditions and are not “overfitted” to one particular value of p . Thus, equiprobable atom selection/activation can be feasible also with low-power computing substrates like neuromorphic chips where the average value of p cannot be precisely defined due to device mismatch and noise.

C. Denoising

Sparse approximations of signals based on learned dictionaries are useful for solving denoising problems because atoms represent repeating additive structure, not noise-like components of the signal which mostly end up in the model residual, $\epsilon(t)$. We investigate the denoising capability of E-MP- and E-OMP-based models by adding Gaussian noise to the rock music signal and comparing the SNR of the model for different noise levels with $p = 0.05$, see Figure 5. The noise level is quantified in terms of the ratio of the standard deviation of additive Gaussian noise, σ_n , to the standard deviation of the signal, σ_s . As expected the OMP-based model produces denoising results that are far better than MP, at least for moderate levels of noise, $\sigma_n/\sigma_s \leq 0.1$. The denoising accuracy obtained with E-MP and E-OMP is comparable to that of OMP for moderate levels of noise, $\sigma_n/\sigma_s \leq 0.1$, and it is superior to OMP for high levels of noise $\sigma_n/\sigma_s > 0.1$. This is another remarkable consequence of equiprobable atom selection, which indicates that the selection constraint prevents overfitting of noise-like signal components.

D. Computational cost

Next we investigate the computational cost of the four different dictionary learning methods. The experiments are done using one 2.3 GHz Intel Core i7 processor core and a

C++ implementation of the algorithms executed in Matlab. We calculate the average core time per signal sample and iteration of the matching pursuit, see Figure 6. This definition




Figure 6. Processing time versus the window length for dictionaries with 32 atoms using one 2.3 GHz Intel Core i7 processor core.

implies that the displayed processing time should be multiplied with the number of samples squared times the average atom selection probability, p , to get the actual computing time. We find that equiprobable selection is beneficial in terms of computational cost. In particular, MP is more costly than E-MP, and OMP is more costly than E-OMP. Furthermore, with a window length of 200 kilosamples the processing time of E-OMP is comparable to that of MP-based dictionary learning. Note that these results are obtained with the efficient implementations of MP and local OMP introduced above, which is the reason why the computational cost per sample decreases with the window length. This is not the case for straightforward implementations of the algorithms, which are significantly more costly.

E. Entropy of sparse representation

Sparse representation with matching pursuit and a learned dictionary is a form of lossy compression where a signal is approximated using prior information encoded in the learned atoms. Next we study the information entropy of the sparse representations generated by the matching pursuits introduced above and the corresponding learned dictionaries. The number of selected atoms per second depends on the signal being processed and the method and dictionary used.

The top-25 most frequently MP/OMP-selected atoms obtained using the learned dictionaries are illustrated in Figure 7 for both rock music and birdsong. Illustrated in the figure are also the constant rate of atom selection events for E-MP and E-OMP with an average selection probability of $p = 0.05$. For MP and OMP some atoms are typically more likely to be selected than others, while some atoms may not be selected

Table II
ENTROPY IN BITS PER ATOM SELECTED BY THE DIFFERENT ALGORITHMS.

SIGNAL	NO ATOMS	MP	OMP	E-MP	E-OMP
Birdsong	32	3.5	3.6	5.0	5.0
Music	32	3.5	4.1	5.0	5.0
Music	64	4.7	5.2	6.0	6.0

and learned at all. In contrast to that the average probability for selecting each atom is constant for E-MP and E-OMP. For example, for the equiprobable pursuits $0.05 \times 48000/32 = 75$ events per second are expected for each atom on average. Thus, the entropy of the sequence of selected atom numbers is expected to be higher for E-MP and E-OMP than for the conventional matching pursuits, which is part of the motivation of this study outlined in the Introduction.

We calculate the Shannon entropy, $-\sum_i p_i \log_2(p_i)$, of the atom number sequence generated by the four matching pursuits and the two different signals, see Table II. We find numerically that OMP events have higher entropy than MP events for the signals considered here, and that E-MP and E-OMP have a maximum entropy of $\log_2(M)$ as expected. For example, with a dictionary of 32 atoms the entropy is $\log_2(32) = 5$ bits per selected atom with E-MP and E-OMP. For comparison, the entropy is 3.5 (4.1) bits per selected atom for MP (OMP) in the case of rock music, and 3.5 (3.6) bits per selected atom for birdsong.

The distribution of atom coefficients, $a_{m,i}$, also depend on the method used and the signal being processed. Histograms of the atom coefficients calculated for rock music with 32 learned atoms are illustrated in Figure 8. An explicit calculation of the




Figure 7. Number of atom selection events per second for rock music and birdsong with an average selection probability of $p = 0.05$. Bold (thin) lines are obtained with a dictionary of 32 (64) learned atoms. For clarity only the twentyfive most frequently selected atoms are included in each case.




Figure 8. Histograms of atom coefficients generated from rock music with 32 learned atoms.

Table III
COEFFICIENT ENTROPIES RELATIVE TO THE OMP COEFFICIENT ENTROPY.

BINS	OMP	E-OMP	MP	E-MP
16	1	1.04	1.07	1.10
32	1	1.02	1.06	1.08
64	1	1.00	1.04	1.05

Shannon entropy of the event coefficients is not meaningful due to the continuous (floating point) nature of these numbers. However, we calculate the entropies of multiple histograms with different bin counts (16, 32 and 64) and note a systematic difference between the four models, see Table III. Quantized OMP coefficients have the lowest entropy, followed by E-OMP, MP and E-MP in order of increasing entropy. The maximum difference is 10% when comparing the coefficient entropies of E-MP and OMP with 16 bins (about 4-bits of precision), which can be compared to the difference of about 20% in Table II. Thus, by adding the entropies of atom indices and coefficients we conclude that E-MP has the highest event entropy followed by E-OMP for the signals considered here.

In these numerical experiments the average atom selection probability is a constant with value 0.05, which implies that on average there is one selected atom in the sparse approximation for every 20 samples of the signal. This implies that the Shannon information of the sparse approximations generated by the different matching pursuits are directly proportional to the entropy per event. Therefore, for the signals considered here the sparse approximations calculated with equiprobable selection have higher Shannon information, in line with the higher reconstruction accuracy obtained with these methods in the former subsections.

IV. DISCUSSION

We extend a well-known dictionary learning and sparse representation model [28] with a basic homeostatic regulation mechanism. The extension is motivated by the observation that the information entropy of such sparse representations is

sub-optimal by construction, and partially also by the central role of homeostatic regulation in cortical networks and spiking neural network models of sensory areas, see for example [37], [38].

The sparse representations are generated with Matching Pursuit (MP) [30], Local Orthogonal MP (OMP) [32] and the homeostatic extensions Equiprobable MP (E-MP) and Equiprobable OMP (E-OMP) introduced here. Dictionaries of shift-invariant atoms are learned using probabilistic gradient ascent for two different signals (rock music and Zebra Finch song phrases). With dictionary learning based on E-MP- and E-OMP we obtain an improved rate of convergence and sparse representation SNR (Figure 3), improved denoising results (Figure 5), a lower computational cost (Figure 6) and higher information entropy of the sparse representation (Table II and Table III).

Note that E-MP and E-OMP only make sense in a dictionary learning setting where the atoms are adapted to the signal, otherwise the regular OMP/MP methods should be used. Furthermore, dictionary learning with equiprobable selection is only sensible with complex signals when the number of shift-invariant atoms in the dictionary is lower than the number of independent and unique signal components, which typically is the case in practical applications. Otherwise the preferred solution is to learn one atom for each independent component, which in general is not expected with equiprobable selection.

The low coefficient entropy of E-OMP in combination with maximum atom selection entropy is an interesting property of E-OMP-based dictionary learning that could improve the accuracy of dictionary learning implementations with quantized coefficients. Furthermore, with E-MP and E-OMP atoms are enforced to occur with the same probability on average and all atoms adapt to the training signal, which is not the case with MP/OMP. In principle the equiprobable selection mechanism resembles a dropout [39] mechanism where the probability of dropout dynamically depends on the selection rate of each atom. We study the sensitivity of the method to variations in the average selection probability, p , and find that the learned dictionaries are useful with other values of p and that the model degrades gracefully for lower values of p . Thus, low-power neuromorphic implementations of the proposed equiprobable selection and dictionary learning methods could be feasible regardless of uncertainties associated with for example device mismatch and noise.

In summary, our main finding is that the accuracy and learning rate of a well-known dictionary learning method are improved by the equiprobable atom selection constraint introduced here, and that the computational cost of the resulting E-MP and E-OMP methods are lower than of MP and local OMP, respectively. Equiprobable selection can be applied to other greedy methods for dictionary learning and there are also opportunities to further develop the basic homeostatic regulation mechanism outlined here in search for more efficient approaches to address this NP-hard problem.

ACKNOWLEDGMENT

This work was stimulated by discussions at the CapoCaccia Neuromorphic Engineering workshops in 2015 and 2016. In particular F.S. acknowledge the discussions with Prof. Christopher Kello about cortical criticality and the lectures by Prof. Yves Frégnac on receptive fields where a knowledge gap between models and observed properties of cells in the visual cortex were highlighted. F.S. is funded by a Gunnar Öquist Fellowship from the Kempe Foundations and S.M.C. is funded by the SKF-LTU University Technology Center. We acknowledge travel support from the Swedish Foundation for International Cooperation in Research and Higher Education (STINT), grant number IG2011-2025.

REFERENCES

- [1] S. Mallat, *A Wavelet Tour of Signal Processing: The Sparse Way*, 3rd ed. Academic Press, 2008.
- [2] A. Bruckstein, D. Donoho, and M. Elad, "From sparse solutions of systems of equations to sparse modeling of signals and images," *SIAM Review*, vol. 51, no. 1, pp. 34–81, 2009.
- [3] M. Elad, *Sparse and redundant representations: from theory to applications in signal and image processing*. Springer, 2010.
- [4] S. Nam, M. Davies, M. Elad, and R. Gribonval, "The cosparse analysis model and algorithms," *Applied and Computational Harmonic Analysis*, vol. 34, no. 1, pp. 30–56, 2013.
- [5] R. Rubinstein, A. Bruckstein, and M. Elad, "Dictionaries for sparse representation modeling," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 1045–1057, June 2010.
- [6] I. Tosić and P. Frossard, "Dictionary learning," *Signal Processing Magazine, IEEE*, vol. 28, no. 2, pp. 27–38, March 2011.
- [7] T. Blumensath and M. Davies, "Gradient pursuits," *IEEE Transactions on Signal Processing*, vol. 56, no. 6, pp. 2370–2382, 2008.
- [8] M. Zibulevsky and B. A. Pearlmutter, "Blind source separation by sparse decomposition in a signal dictionary," *Neural Computation*, vol. 13, no. 4, pp. 863–882, 2016/02/04 2001. [Online]. Available: <http://dx.doi.org/10.1162/089976601300014385>
- [9] P. Bofill and M. Zibulevsky, "Underdetermined blind source separation using sparse representations," *Signal Processing*, vol. 81, no. 11, pp. 2353 – 2362, 2001. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0165168401001207>
- [10] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inform. Theory*, vol. 52, pp. 1289–1306, 2006.
- [11] E. J. Cands and C. Fernandez-Granda, "Towards a mathematical theory of super-resolution," *Communications on Pure and Applied Mathematics*, vol. 67, no. 6, pp. 906–956, 2014. [Online]. Available: <http://dx.doi.org/10.1002/cpa.21455>
- [12] J. Mairal, J. Ponce, G. Sapiro, A. Zisserman, and F. R. Bach, "Supervised dictionary learning," in *Advances in Neural Information Processing Systems 21*, D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, Eds. Curran Associates, Inc., 2009, pp. 1033–1040. [Online]. Available: <http://papers.nips.cc/paper/3448-supervised-dictionary-learning.pdf>
- [13] A. M. Tillmann, "On the computational intractability of exact and approximate dictionary learning," *IEEE Signal Processing Letters*, vol. 22, no. 1, pp. 45–49, Jan 2015.
- [14] B. Olshausen and D. Field, "Emergence of simple-cell receptive field properties by learning a sparse code for natural images," *Nature*, vol. 381, pp. 607–609, 1996.
- [15] B. A. Olshausen and D. J. Field, "Sparse coding with an overcomplete basis set: A strategy employed by v1?" *Vision Research*, vol. 37, pp. 3311–3325, 1997.
- [16] M. S. Lewicki and T. J. Sejnowski, "Learning overcomplete representations," *Neural Computation*, vol. 12, no. 2, pp. 337–365, March 2000.
- [17] K. Kreutz-Delgado, J. F. Murray, B. D. Rao, K. Egan, T.-W. Lee, and T. J. Sejnowski, "Dictionary learning algorithms for sparse representation," *Neural Computation*, vol. 15, no. 2, pp. 349–396, 2016/02/03 2003. [Online]. Available: <http://dx.doi.org/10.1162/089976603762552951>
- [18] S. Chouvardas, Y. Kopsinis, and S. Theodoridis, "An online algorithm for distributed dictionary learning," in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, April 2015, pp. 3292–3296.
- [19] J. sun Seo, B. Lin, M. Kim, P.-Y. Chen, D. Kadetotad, Z. Xu, A. Mohanty, S. Vrudhula, S. Yu, J. Ye, and Y. Cao, "On-chip sparse learning acceleration with cmos and resistive synaptic devices," *Nanotechnology, IEEE Transactions on*, vol. 14, no. 6, pp. 969–979, Nov 2015.
- [20] J. Fournier, C. Monier, M. Pananceau, and Y. Fregnac, "Adaptation of the simple or complex nature of v1 receptive fields to visual statistics," *Nature Neuroscience*, vol. 14, no. 8, pp. 1053–1060, Aug 2011.
- [21] A. Ramirez, E. A. Pnevmatikakis, J. Merel, L. Paninski, K. D. Miller, and R. M. Bruno, "Spatiotemporal receptive fields of barrel cortex revealed by reverse correlation of synaptic input," *Nature Neuroscience*, vol. 17, no. 6, pp. 866–875, Jun 2014.
- [22] J. Fournier, C. Monier, M. Levy, O. Marre, K. Sri, Z. F. Kisvrdy, and Y. Frgnac, "Hidden complexity of synaptic receptive fields in cat v1," *The Journal of Neuroscience*, vol. 34, no. 16, pp. 5515–5528, 2014.
- [23] M. J. Wainwright, "Visual adaptation as optimal information transmission," *Vision Research*, vol. 39, no. 23, pp. 3960 – 3974, 1999.
- [24] J. M. Beggs, "The criticality hypothesis: how local cortical networks might optimize information processing," *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 366, no. 1864, pp. 329–343, 2008. [Online]. Available: <http://rsta.royalsocietypublishing.org/content/366/1864/329>
- [25] J. M. Beggs and N. Timme, "Being critical of criticality in the brain," *Frontiers in Physiology*, vol. 3, no. 163, 2012. [Online]. Available: http://www.frontiersin.org/fractal_physiology/10.3389/fphys.2012.00163/abstract
- [26] W. L. Shew and D. Plenz, "The functional benefits of criticality in the cortex," *The Neuroscientist*, vol. 19, no. 1, pp. 88–100, 2013. [Online]. Available: <http://nro.sagepub.com/content/19/1/88.abstract>
- [27] R. V. Williams-García, M. Moore, J. M. Beggs, and G. Ortiz, "Quasicritical brain dynamics on a nonequilibrium widom line," *Phys. Rev. E*, vol. 90, p. 062714, Dec 2014. [Online]. Available: <http://link.aps.org/doi/10.1103/PhysRevE.90.062714>
- [28] E. Smith and M. S. Lewicki, "Efficient auditory coding," *Nature*, no. 7079, pp. 978–982, 02.
- [29] E. Smith and M. Lewicki, "Efficient coding of time-relative structure using spikes," *Neural Computation*, vol. 17, no. 1, pp. 19–45, 2005.
- [30] S. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE T Signal Proces*, 1993.
- [31] Y. Pati, R. Rezaifar, and P. Krishnaprasad, "Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition," in *Signals, Systems and Computers, 1993. 1993 Conference Record of The Twenty-Seventh Asilomar Conference on*, Nov 1993, pp. 40–44 vol.1.
- [32] B. Mailhe, R. Gribonval, F. Bimbot, and P. Vandergheynst, "A low complexity orthogonal matching pursuit for sparse signal approximation with shift-invariant dictionaries," in *Acoustics, Speech and Signal Processing, 2009. ICASSP 2009. IEEE International Conference on*, April 2009, pp. 3445–3448.
- [33] M. S. Lewicki and B. A. Olshausen, "Probabilistic framework for the adaptation and comparison of image codes," *J. Opt. Soc. Am. A*, vol. 16, no. 7, pp. 1587–1601, Jul 1999. [Online]. Available: <http://josaa.osa.org/abstract.cfm?URI=josaa-16-7-1587>
- [34] S. Krstulovic and R. Gribonval, "MPTK: Matching Pursuit made tractable," in *Proc. Int. Conf. Acoust. Speech Signal Process. (ICASSP'06)*, vol. 3, May 2006, pp. III–496 – III–499.
- [35] The Red Thread, "After the last," licensed under CC BY-NC 3.0. [Online]. Available: <http://freemusicarchive.org>
- [36] M. Anderson, "Song phrases of Zebra Finch," cat. nr. XC287103. [Online]. Available: www.xeno-canto.org/287103/download
- [37] P. D. King, J. Zylberberg, and M. R. DeWeese, "Inhibitory interneurons decorrelate excitatory cells to drive sparse code formation in a spiking model of v1," *Journal of Neuroscience*, vol. 33, no. 13, pp. 5475–5485, 2013.
- [38] C. S. N. Brito and W. Gerstner, "Nonlinear hebbian learning as a unifying principle in receptive field formation," 2016. [Online]. Available: <http://arxiv.org/abs/1601.00701>
- [39] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, vol. 15, pp. 1929–1958, 2014.