

# Learning in Compressed Space

Alexander Fabisch<sup>a</sup>, Yohannes Kassahun<sup>a</sup>, Hendrik Wöhrle<sup>b</sup>, Frank  
Kirchner<sup>a,b</sup>

<sup>a</sup>*University of Bremen, Fachbereich 3 - Mathematik und Informatik, Postfach 330 440,  
28334 Bremen, Germany*

<sup>b</sup>*Robotics Innovation Center, German Research Center for Artificial Intelligence  
(DFKI), Robert-Hooke-Str. 5, 28359 Bremen, Germany*

---

## Abstract

We examine two methods which are used to deal with complex machine learning problems: compressed sensing and model compression. We discuss both methods in the context of feedforward artificial neural networks. We develop the backpropagation method for multilayer neural networks in compressed parameter space. We further show that compressing the weights of a layer of a multilayer neural network is equivalent to compressing the input of the layer. We perform experiments in supervised and reinforcement learning to demonstrate that these methods reduce training time significantly.

*Keywords:* model compression, compressed sensing, artificial neural networks, supervised learning, reinforcement learning

---

## 1. Introduction

Artificial intelligence is facing real world problems and thus machine learning problems become more and more complex. Bengio and Lecun [4] state that “a long-term goal of machine learning research is to produce methods that will enable artificially intelligent agents capable of learning complex behaviors with minimal human intervention and prior knowledge”. As a result, even more effort is shifted from human to machine.

---

*Email addresses:* [afabisch@informatik.uni-bremen.de](mailto:afabisch@informatik.uni-bremen.de) (Alexander Fabisch),  
[kassahun@informatik.uni-bremen.de](mailto:kassahun@informatik.uni-bremen.de) (Yohannes Kassahun),  
[Hendrik.Woehrle@dfki.de](mailto:Hendrik.Woehrle@dfki.de) (Hendrik Wöhrle), [Frank.Kirchner@dfki.de](mailto:Frank.Kirchner@dfki.de) (Frank  
Kirchner)

Complex machine learning problems comprise learning complex behaviors like visual and auditory perception and natural language processing as mentioned by Bengio and Lecun [4] as well as dealing with highly noisy data such as electroencephalography (EEG) signals. A common characteristic of these problems is an associated large amount of data, which includes a large input space dimension and/or a large training set. For example, vision problems usually have hundreds to millions of input components as well as thousands of training examples to cover all possible distortions and a typical brain-computer interface (BCI) [43] can sample 128 channels with 5 kHz.

Complex problems usually require long training times. We will now consider ways that reduce the training time. There are numerous ways of dealing with large input spaces like feature selection [20], dimensionality reduction or feature extraction, which in most of the cases require at least some domain specific knowledge and a manually designed preprocessing flow. But developments that reduce the need for human expertise and intervention such as convolutional neural networks (CNNs) [31] or deep belief neural networks (DBNs) [3] exist as well. The first layers of these kinds of neural networks can be seen as trainable feature extractors, which can be used to cope with large input spaces. This way of dealing with machine learning problems is better because information cannot be lost unintentionally when (almost) unprocessed data is used as input for the learning algorithm.

Another way to avoid this is using conventional machine learning algorithms without built-in preprocessing on raw data. But this is computationally more challenging and requires lots of optimizations, in particular during the training. In this article we want to present novel approaches that we can use to deal with complex machine learning problems. In particular, we want to show how to reduce the time needed to train feedforward neural networks.

We will only consider multilayer perceptrons (MLPs) here to demonstrate two ways of simplifying machine learning optimization problems: model compression and data compression. Advantages of MLPs are that they are universal function approximators with only one hidden layer [23] and they are simple in concept, well-known and widely used. Since MLP layers are contained in CNNs and the backpropagation method is also used in DBNs, we think that this restriction is justifiable. In addition, the presented concepts can be extended to many other learning algorithms. One of these learning algorithms is the support vector machine (SVM) [8, 41]. We will use SVMs and compare them to MLPs in this article.

The source code for most of the experiments we present here is available

online at <https://github.com/AlexanderFabisch/OpenANN>.

## 2. Related Work

We will first give a short overview of the foundations of our work, which include two research branches. In addition, we will discuss related ideas.

We call the first research branch “model compression”. Schmidhuber [38, 39] developed a universal *network encoding language* (NEL) to represent neural networks in a compressed form for supervised learning. This was motivated by the search for the best generalizing network, because the simplest hypothesis that fits the training instances should give the best generalization for the latent function. Since the NEL is not continuous it was not possible to use efficient optimization algorithms for this compressed representation. This was the reason for Koutník et al. [27, 28] to develop a continuous representation, where the weights of a recurrent neural network with fixed topology are represented by coefficients of an inverse discrete cosine transform. The focus of their work was reinforcement learning [40]. Therefore, they presented experiments with the evolutionary optimization algorithm CoSyNE [19]. The main goal in this research is the increase of sample efficiency, that is the reduction of episodes needed to learn a good or successful policy. We will present a very similar approach in this article. However, there will be some differences:

- We will focus on feedforward neural networks, and hence extend the standard backpropagation procedure for learning in compressed space.
- Not only a combination of orthogonal cosine functions, but any kind of orthogonal functions and even randomly generated values can be used to generate the weights.
- Koutník et al. [27, 28] did compress all weights of a neural network with the same coefficients, although they already mentioned that it would be better to have a distinct set of coefficients for each neuron. We will use a distinct set of coefficients for all incoming weights of each neuron.

The second research branch is the so called “compressed sensing” [11, 17]. In compressed sensing, sparse or compressible data is examined. It is possible to compress compressible data through random projections and reconstruct it with high probability by solving an optimization problem. In combination with machine learning it can be used as a preprocessing method that

requires almost no prior knowledge. In particular, the method does not need to compute any features from the training set. Compressed sensing has for example been combined with support vector machines by Calderbank et al. [10] and least squares regression by Maillard and Munos [34]. In these cases the reconstruction is not necessary. It is instead assumed that it is possible to distinguish the instances in compressed space because we could reconstruct the original data with high probability. This is true because random projections approximately preserve distances between instances [5]. As a contribution to the state of the art, we will show that compressed sensing and compressing the first layer of a neural network are mathematically equivalent. This is always the case when we use a multilayer perceptron that has no hidden layer (single layer perceptron, SLP). We have already shown that in a previous paper [26]. For MLP, we show that compressing the weight matrix of a layer is equivalent to compressing the input to that layer.

Reduction of training time without reducing the number of parameters can also be achieved. Recent hardware developments enable cheap parallelization, especially because of general-purpose computing on graphics processing units (GPGPU). For example, Ciresan et al. [14] achieved outstanding results when they trained huge neural networks with millions of weights on GPUs.

Our goal is to reduce training time by reducing the number of parameters that we have to optimize for a neural network. There are other approaches in machine learning that have the same goals and could be combined with the methods we present here. Two ways to do this with neural networks are weight sharing and sparse connections, which are actually a form of model compression. Prominent examples that combine both methods are CNNs [31]. These have been particularly successful in recognizing objects in images such as traffic signs [15].

### 3. Basics of Artificial Neural Networks

Before we describe the actual compression, we will summarize the basics of artificial neural networks that we need in this article. In supervised or reinforcement learning we want to build a model of an unknown function  $\mathbf{y} = f(\mathbf{x})$  with  $\mathbf{x} \in \mathbb{R}^D$ ,  $\mathbf{y} \in \mathbb{R}^F$ . When we talk about a *model* in this article, we mean a feedforward artificial neural network, which is a parametrizable function  $\mathbf{y} = \hat{f}(\mathbf{x}, \mathbf{w})$ , where  $\mathbf{w} \in \mathbb{R}^K$  is a *weight vector*.

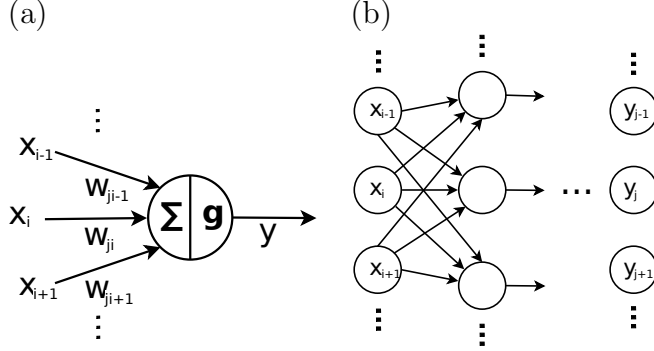


Figure 1: (a) Artificial neuron. (b) Multilayer perceptron.

The basic module of an artificial neural network (ANN) is the artificial neuron (see Figure 1 (a)). It computes a weighted sum of its inputs  $x_i$ , which is called *activation*, and applies an activation function  $g$  to generate the output  $y$ . These neurons can be connected and arranged in layers to form a multilayer perceptron, where consecutive layers are fully connected (see Figure 1 (b)). The output of the model  $\mathbf{y}$  will be calculated through forward propagation, that is the neurons will sequentially calculate their outputs and the outputs of the last layer will generate the model's output.

In order to fit the function  $f$ , we usually modify the model parameters  $\mathbf{w}$  directly, such that they minimize the error on the training set in supervised learning or maximize the return in reinforcement learning. That is, in supervised learning we use an optimization algorithm to minimize an error function  $E(\mathbf{w})$ . Fast optimization algorithms (see for example Table 1) usually require the gradient of the error function  $\nabla E(\mathbf{w}) = \sum_{n=1}^N \nabla E_n(\mathbf{w})$ , where  $E_n$  is the error of instance  $n$ .

In order to calculate the gradient, we use the backpropagation procedure [37]. This requires the calculation of a  $\delta_j$  for each neuron  $j$ . This quantity can be interpreted as the contribution of neuron  $j$  to the error. For each neuron  $j$  in the output layer this is calculated as

$$\delta_j = \frac{\partial E_n}{\partial a_j} = g'(a_j) \frac{\partial E_n}{\partial y_j}, \quad (1)$$

where  $g$  is the activation function of neuron  $j$ ,  $a_j$  is its activation and  $y_j$  its output. Then we can *backpropagate* these  $\delta$ 's in order to calculate the  $\delta$ 's of

lower layers through

$$\delta_j = \frac{\partial E_n}{\partial a_j} = g'(a_j) \sum_k w_{kj} \delta_k, \quad (2)$$

where the neurons  $k$  are in the layer above  $j$ . Finally, we can compute the derivative of the error with respect to each weight

$$\frac{\partial E_n}{\partial w_{ji}} = \delta_j z_i, \quad (3)$$

where  $z_i$  is the output of neuron  $i$ .

#### 4. Learning in Compressed Space

In this section we first describe model compression and present its advantages. Afterwards we give an introduction to data compression in the context of compressed sensing. Finally, we discuss the equivalence of model compression and data compression.

##### 4.1. Model Compression

We can compress the model  $\hat{f}(\mathbf{x}, \mathbf{w})$  by representing  $\mathbf{w} \in \mathbb{R}^K$  with less than  $K$  parameters.

**Definition 1.** Let  $w_{ji}$ ,  $i \in \{1, \dots, I_j\}$  be the weights of a neuron  $j$  as depicted in Figure 1 (a). The weight  $w_{ji}$  is generated from compressed parameter space by

$$w_{ji} = \sum_{m=1}^{M_j} \alpha_{jm} \phi_{mi}, \quad (4)$$

where  $M_j$  is the number of parameters  $\alpha_{j1}, \dots, \alpha_{jM_j}$  to be optimized and  $\phi_{mi}$  can be

- sampled from pairwise orthogonal functions, that is  $\phi_{mi} = \phi_m(t_i)$  where for all  $k, l \in \{1, \dots, M_j\}$

$$\int_0^1 \phi_k(t) \phi_l(t) dt = 0, \quad (5)$$

and  $t_i$  is defined as

$$t_i = \frac{i-1}{I_j-1}, \quad (6)$$

for example, the set  $\{\phi_0(t_i) = \cos(0\pi t_i), \phi_1(t_i) = \cos(1\pi t_i), \phi_2(t_i) = \cos(2\pi t_i), \dots\}$  contains orthogonal cosine functions,

- or sampled from a random distribution that we discuss in Section 4.3.

In order to compress the model,  $\sum_j M_j = L < K$  must be satisfied. That is, the weights of at least one neuron have to be represented by less parameters than the number of weights of the neuron.

$\alpha$  is a vector that contains all parameters that are required to generate the weights of an ANN. For supervised learning, we can extend the standard backpropagation procedure to calculate the gradient of  $E(\alpha)$ .

**Theorem 1.** *Given the partial derivative of the error  $E_n$  with respect to the weight  $w_{ji}$ , the partial derivative of the error  $E_n$  with respect to the parameter  $\alpha_{jm}$  can be calculated by*

$$\frac{\partial E_n}{\partial \alpha_{jm}} = \sum_i \frac{\partial E_n}{\partial w_{ji}} \phi_{mi}. \quad (7)$$

PROOF. We apply the chain rule for partial derivatives.

$$\frac{\partial E_n}{\partial \alpha_{jm}} = \sum_i \frac{\partial E_n}{\partial w_{ji}} \frac{\partial w_{ji}}{\partial \alpha_{jm}} \quad (8)$$

Using Equation (4), we can write

$$\frac{\partial E_n}{\partial \alpha_{jm}} = \sum_i \frac{\partial E_n}{\partial w_{ji}} \frac{\partial}{\partial \alpha_{jm}} \left( \sum_{m'=1}^{M_j} \alpha_{jm'} \phi_{m'i} \right) \quad (9)$$

$$= \sum_i \frac{\partial E_n}{\partial w_{ji}} \phi_{mi}. \quad \square \quad (10)$$

Since we focus on MLPs, the calculation of the activation for each layer  $l$  is usually done by a matrix-vector multiplication

$$\mathbf{W}^{(l)} \mathbf{x}^{(l)} = \mathbf{a}^{(l)}, \quad (11)$$

where  $\mathbf{W}^{(l)} \in \mathbb{R}^{J^{(l)} \times I^{(l)}}$  is the weight matrix of layer  $l$  with its components being  $(\mathbf{W}^{(l)})_{ji} = w_{ji}^{(l)}$ ,  $\mathbf{x}^{(l)} \in \mathbb{R}^{I^{(l)}}$  is the input of layer  $l$  and  $\mathbf{a}^{(l)} \in \mathbb{R}^{J^{(l)}}$  contains the activations of each neuron  $j \in \{1, \dots, J^{(l)}\}$ . If we use a bias in this layer,  $\mathbf{x}^{(l)}$  will be the output of the previous layer with an additional component that is always 1. For MLPs we can simplify the model compression.

**Theorem 2.** *If we use the same number of parameters  $M^{(l)}$  for all neurons in layer  $l$ , the weight matrix of layer  $l$  is generated by*

$$\mathbf{W}^{(l)} = \boldsymbol{\alpha}^{(l)} \boldsymbol{\Phi}^{(l)}, \quad (12)$$

where  $\boldsymbol{\alpha}^{(l)} \in \mathbb{R}^{J^{(l)} \times M^{(l)}}$  is a parameter matrix with the components  $(\boldsymbol{\alpha}^{(l)})_{jm} = \alpha_{jm}^{(l)}$  and  $\boldsymbol{\Phi}^{(l)} \in \mathbb{R}^{M^{(l)} \times I^{(l)}}$  is a constant matrix whose components are generated by  $(\boldsymbol{\Phi}^{(l)})_{mi} = \phi_{mi}$ .

PROOF. When we look at the components of the matrices in Equation (12)

$$\begin{pmatrix} \alpha_{11}^{(l)} & \dots & \alpha_{1M^{(l)}}^{(l)} \\ \vdots & \ddots & \vdots \\ \alpha_{J^{(l)}1}^{(l)} & \dots & \alpha_{J^{(l)}M^{(l)}}^{(l)} \end{pmatrix} \cdot \begin{pmatrix} \phi_{11} & \dots & \phi_{1I^{(l)}} \\ \vdots & \ddots & \vdots \\ \phi_{M^{(l)}1} & \dots & \phi_{M^{(l)}I^{(l)}} \end{pmatrix} \quad (13)$$

$$= \begin{pmatrix} w_{11}^{(l)} & \dots & w_{1I^{(l)}}^{(l)} \\ \vdots & \ddots & \vdots \\ w_{J^{(l)}1}^{(l)} & \dots & w_{J^{(l)}I^{(l)}}^{(l)} \end{pmatrix}, \quad (14)$$

we can see that the weight  $w_{ji}^{(l)}$  is generated by Equation (4), since

$$\begin{pmatrix} \alpha_{j1}^{(l)} & \dots & \alpha_{jM^{(l)}}^{(l)} \end{pmatrix} \cdot \begin{pmatrix} \phi_{1i} \\ \vdots \\ \phi_{M^{(l)}i} \end{pmatrix} = \sum_{m=1}^{M^{(l)}} \alpha_{jm}^{(l)} \phi_{mi}. \quad \square \quad (15)$$

**Theorem 3.** *If we use backpropagation, we can construct a matrix  $\mathbf{V}^{(l)} \in \mathbb{R}^{J^{(l)} \times I^{(l)}}$ , where  $(\mathbf{V}^{(l)})_{ji} = \frac{\partial E_n}{\partial (\mathbf{W}^{(l)})_{ji}}$  and use this matrix to form another matrix  $\boldsymbol{\Omega}^{(l)} \in \mathbb{R}^{J^{(l)} \times M^{(l)}}$ , where  $(\boldsymbol{\Omega}^{(l)})_{jm} = \frac{\partial E_n}{\partial (\boldsymbol{\alpha}^{(l)})_{jm}}$  with*

$$\boldsymbol{\Omega}^{(l)} = \mathbf{V}^{(l)} (\boldsymbol{\Phi}^{(l)})^T. \quad (16)$$

The proof for this theorem is straight forward and only requires to write Equation (7) in matrix form. Thus, we will omit it here.

In this article, we use the following notation to describe the architecture of an MLP: 2-20-20-1 means that there are 2 inputs, 20 hidden nodes in the first hidden layer, 20 hidden nodes in the second hidden layer and 1 output node and there is an additional bias in every layer. Similarly, we describe the compression: 3-6-3 means that we compress the weights of the first hidden layer with 3 parameters per neuron, in the second hidden layer with 6 parameters per neuron and in the output layer with 3 parameters per neuron.

Algorithm	Time complexity
Newton Algorithm	$O(L^3)$
Gauß-Newton	$O(L^3)$
Levenberg-Marquardt	$O(L^3)$
BFGS	$O(L^2)$
Conjugate Gradient	$O(L)$

Table 1: Time complexity of each iteration for several gradient based optimization algorithms according to LeCun et al. [32]. The number of parameters that will be optimized is denoted by  $L$ .

#### 4.2. Advantages of Model Compression

There are at least three ways to reduce the training time of neural networks:

- Option 1: We can reduce the time that is needed for forward propagation and backpropagation.
- Option 2: We can reduce the time that the optimization algorithm needs for each iteration. Here we have to exclude the time of forward propagation and backpropagation.
- Option 3: We can reduce the number of iterations an optimization algorithm needs to converge.

When we talk about reinforcement learning in this article, we will focus on direct policy search. Usually, reinforcement learning requires the approximation of value functions that indicate how good it is for an agent to be in a given state and then we can infer policies from these value functions [40]. In direct policy search, the policies are learned directly. Reducing the number of model parameters has different effects in supervised learning and reinforcement learning with direct policy search.

When we look at the extended backpropagation that is required when we use this model compression method for supervised learning, we can see that it only adds more computational cost. In addition, we have to generate the weights anew, when we adjust the parameters. Thus, we clearly do not use option 1. We even cannot assure that the number of iterations is reduced (option 3), because an indirect weight optimization can even make the optimization harder. Nevertheless, the overall learning process will speed up, when we use complex optimization algorithms that are in  $O(L^2)$  or even  $O(L^3)$ , where  $L$  is the number of parameters that have to be optimized,

and we can reduce the number of parameters in comparison to the number of weights sufficiently. Then each iteration of the optimization algorithm requires only a fraction of time. Interestingly, some of the best optimization algorithms for batch learning have high time complexity in each optimization step (see Table 1).

We think that the computational cost is not important for reinforcement learning. This is because the goal of reinforcement learning is to deal with complex simulation environments that are computationally expensive or with the real world. In the real world we want to minimize the possibility to wear or break real systems. Thus, we have to minimize the number of function evaluations an optimization algorithm needs. Each function evaluation in this setup is called *episode*. The derivative-free optimization algorithm we use here for direct policy search is CMA-ES [21]. To our knowledge there is no convergence proof for CMA-ES yet, but our empirical results show that the convergence speed depends on the search space dimension  $L$ . This is the reason why the reduction of parameters will speed up the learning process.

#### 4.3. Data Compression

Compressed sensing can be used for data compression. Baraniuk et al. [2] summarize

“In *Compressed Sensing* (CS) [...], a random projection of a high-dimensional but sparse or compressible signal vector onto a lower-dimensional space has been shown, with high probability, to contain enough information to enable signal reconstruction with small or zero error.”

In specific terms that means one of the key equations in CS is

$$\mathbf{y} = \mathbf{\Phi} \mathbf{x} \quad [11, 17], \quad (17)$$

where  $\mathbf{x} \in \mathbb{R}^N$  is sparse or compressible,  $\mathbf{y} \in \mathbb{R}^M$  and  $M < N$ .  $\mathbf{\Phi} \in \mathbb{R}^{M \times N}$  is generated randomly and has to satisfy the *restricted isometry property* (RIP) [12]. In this setting, we could approximately reconstruct  $\mathbf{x}$  from  $\mathbf{y}$  with high probability. However, this is not required when we use CS as a preprocessing method for machine learning.

The random distributions we use to generate the matrix  $\mathbf{\Phi}$  have to satisfy the RIP. Candès and Tao [12] originally define the restricted isometry property of a matrix  $\mathbf{\Phi}$  that has *restricted isometry constants*  $\delta_M$ . These

constants limit the deviation of the norm, when we project any vector  $\mathbf{x}$  on any  $M$ -dimensional subspace by multiplying it with a submatrix  $\Phi_M$  of  $\Phi$ , that is

$$(1 - \delta_M) \|\mathbf{x}\|^2 \leq \|\Phi_M \mathbf{x}\|^2 \leq (1 + \delta_M) \|\mathbf{x}\|^2. \quad (18)$$

This definition is then used to prove that sparse data can be reconstructed after compression. In addition, they show that there exist matrices with good restricted isometry constants.

Baraniuk et al. [2] show for example that

$$\Phi_{mi} \sim \mathcal{N}(0, \frac{1}{M_j}), \quad (19)$$

$$\Phi_{mi} = \begin{cases} \frac{+1}{\sqrt{M_j}} & \text{with probability } \frac{1}{2} \\ \frac{-1}{\sqrt{M_j}} & \text{with probability } \frac{1}{2} \end{cases}, \quad (20)$$

$$\text{and } \Phi_{mi} = \begin{cases} \frac{+1}{\sqrt{M_j}} & \text{with probability } \frac{1}{6} \\ 0 & \text{with probability } \frac{2}{3} \\ \frac{-1}{\sqrt{M_j}} & \text{with probability } \frac{1}{6} \end{cases} \quad (21)$$

satisfy the RIP.

Another property of these random projections is that they approximately preserve distances between instances because they satisfy the Johnson-Lindenstrauss lemma [24] according to Baraniuk et al. [2]. In other words, random projections preserve similarities of data vectors. This property is especially useful for machine learning algorithms, because their ability to generalize is based upon the fact that similar inputs have similar outputs. We can regard compressed sensing as a preprocessing method that requires almost no domain specific knowledge. We only have to know that the data is compressible.

#### 4.4. Equivalence of Model Compression and Data Compression

**Theorem 4.** *Compressing the weight matrix  $\mathbf{W}^{(l)}$  of layer  $l$  is equivalent to compressing the input  $\mathbf{x}^{(l)}$  of this layer.*

PROOF. Let us assume that we generate  $\mathbf{W}^{(l)}$  according to Equation (12).

Then we can write the calculation of the activations in layer  $l$  as

$$\mathbf{W}^{(l)} \mathbf{x}^{(l)} = (\boldsymbol{\alpha}^{(l)} \boldsymbol{\Phi}^{(l)}) \mathbf{x}^{(l)} \quad (22)$$

$$= \boldsymbol{\alpha}^{(l)} (\boldsymbol{\Phi}^{(l)} \mathbf{x}^{(l)}) \quad (23)$$

$$= \boldsymbol{\alpha}^{(l)} \mathbf{x}'^{(l)}. \quad \square \quad (24)$$

Hence, we can interpret  $\boldsymbol{\alpha}^{(l)}$  as the new weight matrix and  $\mathbf{x}'^{(l)}$  as the compressed input of layer  $l$ .

Compressing the weight matrix of layer  $l = 0$  of the ANN is equivalent to compressing its input  $\mathbf{x}^{(0)}$  (with a bias). Compressing the input can be done using Equation (17). In particular, for single layer networks compressed sensing and model compression are mathematically identical. Therefore, we can transfer the ideas from compressed sensing to model compression.

## 5. Supervised Learning

One of the key contributions of this article is the backpropagation algorithm for compressed MLPs. Thus, we want to show the advantages of compression by means of supervised learning. We want to demonstrate that a reduced number of parameters accelerates the optimization process. The main reason for this is the reduced time for each iteration of the optimization algorithm because some of the best optimization algorithms for ANNs have quadratic or cubic time complexity.

In the following experiments we use the Levenberg-Marquardt algorithm (LMA) [33, 35] to optimize the parameters of the ANN. The implementation we use is from the library ALGLIB [7]. This algorithm has cubic complexity in every iteration. But it usually needs significantly less iterations than other optimization algorithms (for example conjugate gradient).

### 5.1. Two Spirals

In this experiment we use MLPs to compare the influence of model compression on the training time. The two spirals data set, which has been developed by Lang and Witbrock [30], is a standard benchmark for classification algorithms. The goal is to separate two classes of points that are located on a two-dimensional surface (see Figure 2). They are arranged as interlocked spirals. The data set we use here is taken from the fast artificial neural network library (FANN) [36]. It consists of 193 training instances and 193 test instances.

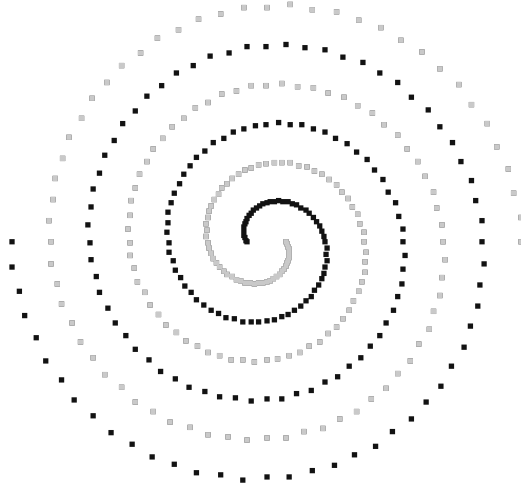


Figure 2: Two interlocked spirals. Different colors indicate different classes.

We use MLPs with two hidden layers to solve this problem. The instances of this data set are not compressible and an indirect weight representation actually makes the optimization more difficult. Here, we only want to demonstrate that the mere reduction of parameters accelerates the optimization.

In this experiment we use MLPs with two hidden layers to compare the influence of compression on the training time. We use hyperbolic tangent in every layer as activation function. The classes are encoded as -1 and 1. We compress the weights with orthogonal cosine functions. We draw the initial parameters from  $\mathcal{N}(0, 0.05)$ . The error function is the sum of squared errors (SSE). In this experiment we use the following stopping criteria:

- the maximum number of iterations is 1,000,
- the difference of the error between consecutive iterations must not be too small, hence we stop the optimization if

$$|E^{t+1} - E^t| \leq 10^{-8} \cdot \max\{|E^{t+1}|, |E^t|, 1\},$$

- the gradient must not be too small, therefore we stop if

$$|\mathbf{g}| \leq 10^{-8},$$

where  $t$  is the current iteration,  $E^t$  is the error at iteration  $t$ , and  $\mathbf{g}$  is the current gradient.

<i>Architecture</i>	<i>Compression</i>	<i>L</i>	<i>Correct predictions</i>	<i>Iterations</i>	<i>Training time</i>
			$\mu \pm \sigma$	$\mu \pm \sigma$	$\mu$
2-20-10-1	-	281	$192.49 \pm 0.795$	$313 \pm 8.560$	2575 ms
2-20-20-1	-	501	$192.66 \pm 0.705$	$251 \pm 6.060$	7176 ms
2-20-20-1	3-21-21	501	$192.33 \pm 0.843$	$286 \pm 6.599$	8344 ms
2-20-20-1	3-12-12	312	$192.46 \pm 0.805$	$323 \pm 7.802$	3430 ms
2-20-20-1	3-6-6	186	$191.97 \pm 0.941$	$401 \pm 9.261$	1740 ms
2-20-20-1	3-6-3	183	$192.22 \pm 0.883$	$484 \pm 11.672$	2042 ms

Table 2: Results of experiments with the two spirals data set. The values are averaged over 100 runs. The training time is measured with an Intel Core i7-2600K.

The results are shown in Table 2. The number of parameters to be optimized is denoted as  $L$ . It is either the number of weights for uncompressed ANNs or the number of parameters  $\sum_j M_j$  for compressed ANNs. For each configuration we did 100 test runs.

The differences in the classification performance are negligible. As we expected, the optimization problem became more difficult due to indirect weight representation. The indicator for this is the number of iterations. This number even increases when we reduce the number of parameters. However, the training time also decreases because the time for each optimization step decreases significantly. For example, the shortest average training time (1740 ms) is achieved with a compressed MLP that has only 186 parameters but needs 401 iterations on average, that is each iteration takes less than 5 ms. In comparison, the uncompressed MLP with the same topology has an average training time of 7176 ms although only 251 iterations are required on average because each iteration takes more than 28 ms. This shows that the additional complexity of the extended backpropagation is by far not as crucial as the reduced complexity of the optimization step.

### 5.2. P300 Speller

The two spirals data set is a difficult problem for neural networks. Nevertheless, it is not a complex problem because the input dimension is low and the training set is small. In contrast, brain-computer interfaces (BCIs) [43] generate a huge amount of data. Ordinary systems measure brain activity with 64 or more electrodes and sample signals with a high frequency. The relevant signal usually has a significantly lower frequency. Therefore the data is highly compressible. In the following experiment we use data from a BCI as an example for a complex problem and show that we can use compressed sensing in combination with model compression to speed up the training time and even improve the performance of the classifier.

A brain-computer interface can be used to spell characters. The P300 speller [16, 18] is based on the *oddball paradigm*: whenever a user encounters an important rare stimulus in a sequence of stimuli a P300 potential is elicited. The P300 potential is a so-called event-related potential (ERP) that occurs around 300 ms after the corresponding stimulus.

In this setup the user concentrates on the character he wants to spell at the character matrix that is shown in Figure 3. All columns and rows are intensified 15 times in random order. Each of these repetitions is called a trial. A classifier is used to recognize P300 potentials. After each character epoch the classifier’s score for each row and column will be accumulated and the result will be the character which is contained in the column and the row with the highest scores. A low classifier accuracy does not necessarily lead to a low character accuracy because we accumulate 15 trials. Here we use the the BCI competition III data set II [6, 29]. We chose this data set because there are many published results that we can compare to our results.

An instance consists of 15,360 components, that is one second sampled with 240 Hz from 64 channels. We use a single layer perceptron (SLP) with hyperbolic tangent activation function for classification and allow two preprocessing methods:

- Lowpass filter with cut-off frequency 10 Hz.
- Downsampling with a factor of eleven.

The combination of both is called decimation. After downsampling, an instance consists of 1,344 components.

(W)O R D						
A	B	C	D	E	F	
G	H	I	J	K	L	
M	N	O	P	Q	R	
S	T	U	V	W	X	
Y	Z	1	2	3	4	
5	6	7	8	9	_	

Figure 3: Character matrix.

The compression we use here is random compression, where each component of  $\Phi$  is drawn from the following distribution:

$$\Phi_{mi} = \begin{cases} +\sqrt{\frac{3}{M_j}} & \text{with probability } \frac{1}{6} \\ 0 & \text{with probability } \frac{2}{3} \\ -\sqrt{\frac{3}{M_j}} & \text{with probability } \frac{1}{6} \end{cases} \quad (25)$$

The result is a matrix with approximately  $\frac{2}{3}$  entries of 0, that is not all parameters affect all weights. We have shown that compressing the input of a layer is equivalent to compressing the weights of a layer (see Theorem 4). In an SLP we only have one layer. Thus compressing the data is equivalent to compressing all weights. There is only one difference: we do not compress the weight of the bias, when we compress the data. Nevertheless we do compress the data here. This can also be regarded as a preprocessing method.

The most successful setups are listed in Table 3. We tried several configurations and noticed that neither higher compression nor lower compressions yield better results. The results are averaged over two subjects and ten runs per subject. The compression matrix  $\Phi$  and the initial parameter vector  $\alpha_0$  varied during the ten runs. An uncompressed SLP needed more than one day of training and did not reach an accuracy better than 20 %.

The best results of the BCI competition III are shown in Table 4. Three of our results would have been ranked second in this competition. Note that we neither did use preprocessing methods like channel selection, detrending, PCA, etc. nor did we use an ensemble of classifiers. Almost no domain

<i>Parameters</i>	<i>Compression</i>	<i>Lowpass filter</i>	<i>Downsampling</i>	<i>Accuracy</i>		<i>Training time</i>	<i>Iterations</i>
				15 trials	5 trials		
801	✓	✓	✓	93.9 %	63.8 %	59.6 s	12.4
801	✓	✓	-	94.1 %	64.2 %	85.3 s	16.1
1201	✓	-	-	87.8 %	55.3 %	143.8 s	16.1
1345	-	✓	✓	93.5 %	64.1 %	145.5 s	16.3
15360*	-	-	-	20.0 %	-	25 h	-

Table 3: P300 speller results. The values are averaged over two subjects and ten runs per subject. (\*) We only did one run with one subject for this configuration.

Accuracy		Classifier	Preprocessing
15 trials	5 trials		
96,5 %	73,5 %	Ensemble of SVM	Channel selection, bandpass filter, downsampling
90,5 %	55,0 %	Bagging with SVM	Channel selection, bandpass filter, eye movement artifacts removal, downsampling
90 %	59,5 %	-	Lowpass filter, downsampling, detrending, PCA, t-statistic
89,5 %	53,5 %	Gradient Boosting	Detrending, decimation
87,5 %	57,5 %	Bagging with LDA	Channel selection, bandpass filter, downsampling

Table 4: Best results of the BCI Competition III [6].

specific knowledge is required for our method. Only a lowpass filter is really required in order to achieve comparable results. So we had to make a rough assumption about the frequency of the P300 potential. We have no information about the training time used to generate the competition results. Thus, we cannot certainly say that we developed a significantly faster method.

### 5.3. Single Trial P300 Detection

A more challenging task in BCI is P300 detection with only a single trial. Here we use data from the project IMMI (Intelligent Man-Machine Interface)<sup>1</sup> of the German Research Center for Artificial Intelligence (DFKI) Robotics Innovation Center and the University of Bremen. We have 24 data sets, that is three independent sessions from eight subjects.

In the experiment two kinds of visual stimuli were presented to the test person: irrelevant “standards” and relevant “targets”. The targets are supposed to elicit P300 potentials. When a target was presented the test person had to react with a movement of the right arm. The ratio between standards and targets was 8:1. The voltage is recorded at 5 kHz sampling rate with 124 electrodes. For the experiments we used 62 EEG electrodes. The data was acquired using an actiCap system (Brain Products GmbH, Munich, Germany) and amplified by four 32 channel BrainAmp DC amplifiers. We preprocessed the data as follows:

1. Standardization: We subtracted the mean of each channel and divided the data by the standard deviation.
2. Decimation to 25 Hz: The data was filtered with a lowpass filter and subsampled.
3. Lowpass filter with cut-off frequency 4 Hz.

After this preprocessing an instance consists of 1,550 components: 25 samples from 62 channels.

Here we compare two types of machine learning algorithms: SLPs and SVMs. We use the library LIBSVM [13] with a linear kernel. We use compressed sensing for both classifiers. We will measure generalization performance and training time. In order to determine the generalization performance, we use the data of the first two sessions as training set and the data

---

<sup>1</sup>Website of the project: <http://robotik.dfki-bremen.de/en/research/projects/space-robotics/immi.html>

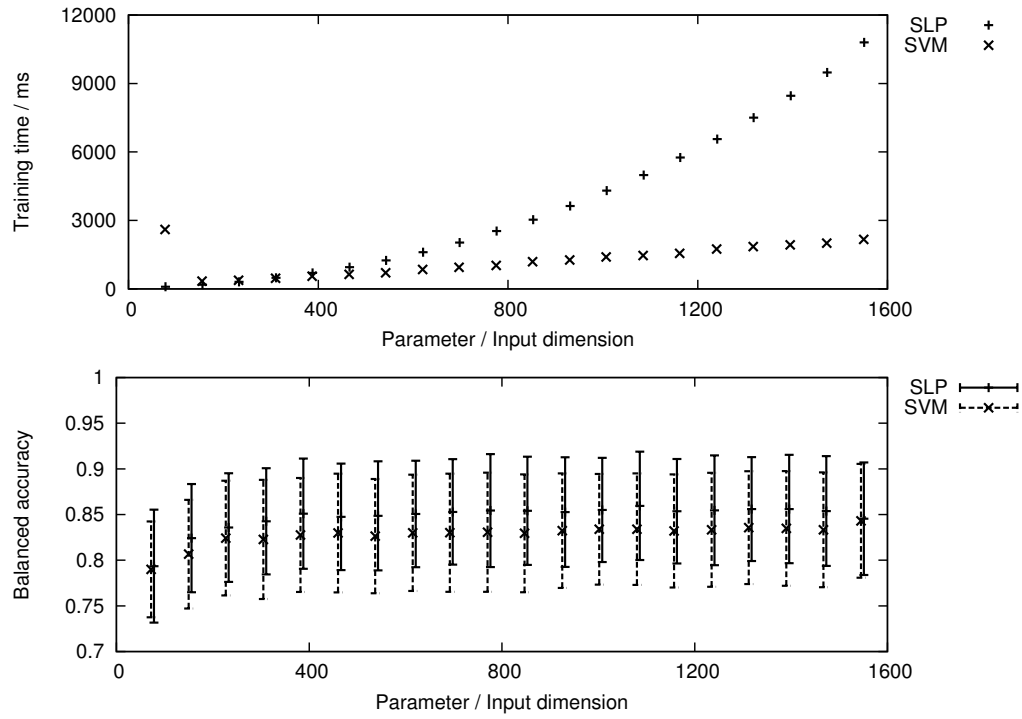


Figure 4: Balanced accuracy and training time of compressed SLP and SVM. The values are averaged over 80 experiments (ten per subject). The interval  $[\mu - \sigma, \mu + \sigma]$  is shown for the balanced accuracy.

of the third session as test set and calculate the *balanced accuracy* [9]

$$\text{balanced accuracy} = \frac{\frac{TP}{TP+FN} + \frac{TN}{TN+FP}}{2}, \quad (26)$$

which is suitable for unbalanced class distributions. Here, TP are true positives, FN are false negatives, TN are true negatives and FP are false positives.

The results averaged over all subjects are shown in Figure 4. The training time of an SVM is usually much shorter and the reduction of parameters is more effective for SLPs in terms of training time. However, we can approximately preserve the accuracy when we reduce the number of parameters to 15 % of the number of weights (233 parameters) while at the same time the training time is reduced to less than 20 % for SVMs (380 ms) and less than 3 % for SLPs (308 ms). The training time of this compressed SLP is just a little bit shorter than the training time of the compressed SVM but it is significantly shorter than the training time of an uncompressed SVM. The balanced accuracy of SVMs and SLPs is approximately the same, but SLPs seem to be slightly better for compressed data.

When the compression of the data or the model is too high, the balanced accuracy will be low and for SVMs the optimization problem will become harder, that is the training time might become longer as we can see in the plot.

## 6. Reinforcement Learning

An area of application for model compression that has already been explored by Koutník et al. [27] is reinforcement learning. Here the computational complexity of an optimization step is negligible. The only thing that counts is the number of episodes required to learn a successful policy because we usually deal with simulation environments that are computationally expensive or with the real world.

In these experiments we evolve the parameters or weights of MLPs that represent the agent’s policy. This is a form of neuroevolution.

### 6.1. Pole Balancing

Our first reinforcement learning experiment is a very simple benchmark. The pole balancing environment is depicted in Figure 5. The goal is to balance two poles mounted on a cart by applying a force  $F$ . The dynamics are described by Wieland [42]. We will analyze single and double pole balancing

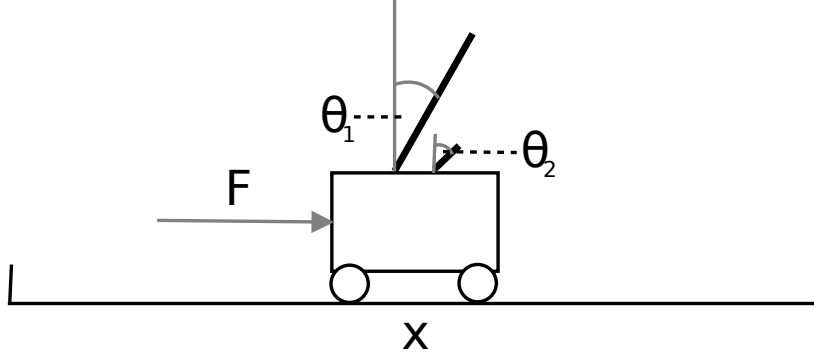


Figure 5: Pole balancing environment.

with and without velocities. This has already been done by Koutník et al. [27] with the optimization algorithm CoSyNE [19].

We perform direct policy search to solve this problem. The policy  $\pi : S \rightarrow A$  is represented by a single layer perceptron without bias, where  $S$  is the state space and  $A$  is the action space. We use the optimization algorithm IPOP-CMA-ES [1, 21] to adjust the weights. A similar approach has been examined by Heidrich-Meisner and Igel [22]. The fitness function we want to maximize is the number of time steps per episode, that is the weights will be updated after each episode. When we do not know the velocities, we have to reconstruct them somehow from the current and the last state. One way to do this is to realize a kind of memory within the neural network. This is possible with recurrent neural networks. The option we chose here is an augmented neural network with Kalman filters (ANKF) [25]. ANKFs are neural networks with a specialized recurrent neural network ( $\alpha$ - $\beta$  filter), which is used to estimate the velocities. We will not compress the parameters of the  $\alpha$ - $\beta$  filters. Thus, we have to optimize  $L$  parameters for the single layer perceptron and one parameter for each  $\alpha$ - $\beta$  filter (we will call this number  $K$ ).

We compare compressed and uncompressed SLPs in terms of the number of episodes required to learn a successful policy. The results are summarized in Table 5. SLPs with less parameters tend to converge faster if they are able to represent a successful policy. So, we only listed the best compression setups. None of the 1,000 experiments for each configuration failed. It was always possible to learn a successful policy with the listed configurations. In three out of four environments it was possible to reduce the

Environment		Setup		Episodes					Time
<i>Velocities</i>	<i>Poles</i>	<i>Compression</i>	<i>L + K</i>	$\mu$	$\sigma$	<i>min</i>	<i>max</i>	<i>median</i>	
✓	1	-	4	33.1	20.8	1	142	28	119 ms
✓	1	✓	1	<b>2.5</b>	2.4	1	40	2	104 ms
✓	2	-	6	261.2	174.3	28	1410	224	210 ms
✓	2	✓	5	<b>201.4</b>	229.8	10	1336	139	160 ms
-	1	-	4 + 2	31.4	15.4	1	102	30	117 ms
-	1	✓	3 + 2	<b>14.3</b>	9.6	1	57	12	114 ms
-	2	-	6 + 3	<b>425.5</b>	220.9	3	1714	388	229 ms
-	2	✓	5 + 3	434.3	318.4	25	1909	352	195 ms

Table 5: Pole balancing results. The values are averaged over 1000 experiments.

number of episodes required to learn a successful policy with model compression. Double pole balancing without velocities is the exception. This result is consistent with the result of Koutník et al. [27]. They even slowed down the learning speed significantly by compressing the weights for double pole balancing without velocities. Another similarity is that single pole balancing can be solved in 2.5 episodes on average. Koutník et al. [27] explained that a solution only requires the weights of an SLP to be equal and positive. When we represent all weights with only one parameter they always have the same value. Therefore, the optimization algorithm only has to find a positive value.

An overview of the best results in pole balancing at the moment is shown in Table 6. We could at least reach or even outperform most of the results. Again, the exception is double pole balancing without velocities. Kassahun et al. [25] achieved a significantly better result even though almost the same methods were used. The reason is that they did not learn the parameters of the  $\alpha$ - $\beta$  filters. These parameters were fixed.

Koutník et al. [27] already achieved similar results with a similar approach. The differences are

- they used an optimization algorithm that does not work as good as IPOP-CMA-ES in this case,

Pole balancing with velocities		Episodes	
Method	Publication	Single	Double
CoSyNE	Gomez et al. [19]	98	954
CMA-NeuroES	Heidrich-Meisner and Igel [22]	91	585
CoSyNE and DCT	Koutník et al. [27]	2	258

Pole balancing without velocities		Episodes	
Method	Publication	Single	Double
CMA-NeuroES	Heidrich-Meisner and Igel [22]	192	860
CoSyNE and DCT	Koutník et al. [27]	151	3,421
CoSyNE	Gomez et al. [19]	127	1,249
CMA-ES and ANKF	Kassahun et al. [25]	-	302

Table 6: State of the art results in pole balancing. The results are compared in terms of episodes that were required to learn a successful policy.

- they did use a recurrent neural network instead of  $\alpha$ - $\beta$  filters, which also was disadvantageous.

## 6.2. Octopus Arm

The octopus arm environment (available at <http://www.cs.mcgill.ca/~dprecup/workshops/ICML06/octopus.html>) is a complex reinforcement learning problem. The dimensions of the state and action space are greater than those of pole balancing. The environment is depicted in Figure 6. In the environment configuration we use here, the arm consists of 12 compartments and each compartment consists of four point masses that are shared with adjacent compartments and has three muscles. The position and velocity of the point masses are components of the state as well as the positions and velocities of the food. The agent has to control the muscles that can be contracted with a continuous degree of strength. A state consists of 106 components and an action consists of 36 components. The values of the action’s components have to be in  $[0, 1]$ .

In these experiments, the return of an episode is the sum of all rewards. The reward for each step is usually -0.01. When the agent feeds the mouth with the left piece of food it receives the reward 5 and when it feeds the right piece of food it receives the reward 7. An episode stops if either both pieces of food are in the mouth or 1,000 steps have passed. Hence, the minimal return is -10 and the upper bound is 12.

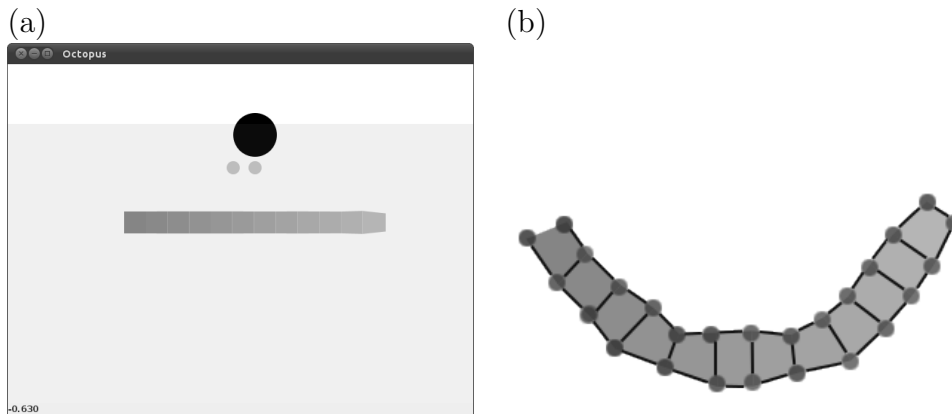


Figure 6: (a) Octopus arm environment. The dark gray line is the straightened arm, the light gray circles are pieces of food and the black circle is a mouth. (b) Octopus arm. The black lines mark muscles and the circles designate point masses of the octopus arm.

We use almost the same approach to solve this problem as for pole balancing. But the policy is represented by an MLP with ten hidden units and a bias, that is the MLP has 1466 weights. We assume that the task is simple enough to be solved with less parameters. We compress the weights with orthogonal cosine functions. We calculate the average return in each experiment as an indicator for the speed of the learning algorithm and the maximum return as an indicator for the best representable policy.

The results are listed in Table 7. One can conclude from the maximum return that each configuration is in principle able to represent a very good policy. The uncompressed MLP yields by far the worst result. It is possible to learn a good policy but this does not happen fast. When we represent the weights with orthogonal cosine functions and do not reduce the number of parameters, the result is already better. Thus, this representation is advantageous for this problem. But a significant reduction of parameters further increases the mean average return and thus improves the learning speed.

## 7. Conclusion

We developed a method to compress the weights of a neural network and extended the backpropagation method. We have shown that compressing the input of a neural network is the same as compressing the weights of the network. Therefore, we are now able to transfer some ideas from compressed

Compression	$L$	Return in 900 episodes (average of 20 experiments)		
		$\mu$	$\sigma$	max
-	1466	-4,52	1,02	11,71
107-11	1466	2,49	0,82	11,73
80-11	1196	1,89	1,32	11,72
40-11	796	2,06	0,86	11,80
20-11	596	2,28	0,60	11,72
10-11	496	2,73	1,30	11,73
5-11	446	<b>3,60</b>	1,08	11,72

Table 7: Octopus arm benchmark results. The maximum return is an indicator for the best policy that is representable with these setups and the mean average return indicates how fast a good policy is learned.

sensing. For example, we can use randomly generated matrices to compress the weights of a neural network, when we know that the data is compressible.

We have applied model compression to both supervised learning and reinforcement learning. We found that model compression

- reduces the computational cost of an optimization iteration and
- can reduce the number of episodes required to learn good policies in a complex reinforcement learning problem.

We have examined only fully connected layers of neural networks here. However, fully connected layers exist in convolutional neural networks or deep belief networks. Therefore it is simple to transfer this method to these types of neural networks. But it is even possible to adapt this method to layers that are not fully connected or share weights, for example convolutional layers in CNNs.

## Acknowledgement

This work was supported through a grant of the Federal Ministry of Economics and Technology (BMW, FKZ 50 RA 1011).

## 8. References

- [1] Auger, A., Hansen, N., 2005. A restart CMA evolution strategy with increasing population size. In: Proceedings of the IEEE Congress on Evolutionary Computation. Vol. 2. IEEE Press, pp. 1769–1776.
- [2] Baraniuk, R. G., Davenport, M. A., DeVore, R. A., Wakin, M. B., 2008. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation* (3), 253–263.
- [3] Bengio, Y., 2007. Learning deep architectures for AI. Tech. Rep. 1312, Dept. IRO, Universite de Montreal.
- [4] Bengio, Y., Lecun, Y., 2007. Scaling Learning Algorithms towards AI. MIT Press, pp. 321–360.
- [5] Bingham, E., Mannila, H., 2001. Random projection in dimensionality reduction: applications to image and text data. In: Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining. ACM, pp. 245–250.
- [6] Blankertz, B., 2005. BCI competition III webpage. Webpage.  
URL <http://www.bbci.de/competition/iii/>
- [7] Bochkhanov, S., Bystritsky, V., 2011. Levenberg-marquardt algorithm for multivariate optimization. Webpage.  
URL <http://www.alglib.net/optimization/levenbergmarquardt.php>
- [8] Boser, B. E., Guyon, I. M., Vapnik, V. N., 1992. A training algorithm for optimal margin classifiers. In: Proceedings of the 5th Annual ACM Workshop on Computational Learning Theory. ACM Press, pp. 144–152.
- [9] Brodersen, K. H., Ong, C. S., Stephan, K. E., Buhmann, J. M., 2010. The balanced accuracy and its posterior distribution. In: Proceedings of the 2010 20th International Conference on Pattern Recognition. IEEE Computer Society, pp. 3121–3124.
- [10] Calderbank, R., Jafarpour, S., Schapire, R., 2009. Compressed learning: Universal sparse dimensionality reduction and learning in the measurement domain. Tech. rep., Princeton University.  
URL <http://dsp.rice.edu/files/cs/cl.pdf>

- [11] Candès, E. J., Romberg, J. K., 2005. Practical signal recovery from random projections. In: *Proceedings of SPIE Computational Imaging III*. SPIE Press, pp. 76–86.
- [12] Candès, E. J., Tao, T., 2005. Decoding by linear programming. *IEEE Transactions on Information Theory* 51 (12), 4203–4215.
- [13] Chang, C.-C., Lin, C.-J., 2011. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology* 2 (3), 27:1–27:27.
- [14] Ciresan, D. C., Meier, U., Gambardella, L. M., Schmidhuber, J., 2010. Deep, big, simple neural nets for handwritten digit recognition. *Neural Computation* 22 (12), 3207–3220.
- [15] Ciresan, D. C., Meier, U., Masci, J., Schmidhuber, J., 2011. A committee of neural networks for traffic sign classification. In: *International Joint Conference on Neural Networks*. IEEE Press, pp. 1918–1921.
- [16] Donchin, E., Spencer, K. M., Wijesinghe, R., 2000. The mental prosthesis: Assessing the speed of a p300-based brain-computer interface. *IEEE Transactions on Rehabilitation Engineering* 8, 174–179.
- [17] Donoho, D. L., 2006. Compressed sensing. *IEEE Transactions on Information Theory* 52 (4), 1289–1306.
- [18] Farwell, L. A., Donchin, E., 1988. Talking off the top of your head: Toward a mental prosthesis utilizing event-related brain potentials. *Electroencephalography and Clinical Neurophysiology* 70, 510–523.
- [19] Gomez, F., Schmidhuber, J., Miikkulainen, R., 2008. Accelerated neural evolution through cooperatively coevolved synapses. *Journal of Machine Learning Research* 9, 937–965.
- [20] Guyon, I., Elisseeff, A., 2003. An introduction to variable and feature selection. *Journal of Machine Learning Research* 3, 1157–1182.
- [21] Hansen, N., Ostermeier, A., 2001. Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation* 9 (2), 159–195.

- [22] Heidrich-Meisner, V., Igel, C., 2009. Neuroevolution strategies for episodic reinforcement learning. *Journal of Algorithms* 64 (4), 152–168.
- [23] Hornik, K., Stinchcombe, M. B., White, H., 1989. Multilayer feedforward networks are universal approximators. *Neural Networks* 2 (5), 359–366.
- [24] Johnson, W. B., Lindenstrauss, J., 1984. Extensions of lipschitz mapping into hilbert space. In: *Conference in Modern Analysis and Probability*. Vol. 26 of *Contemporary Mathematics*. American Mathematical Society, pp. 189–206.
- [25] Kassahun, Y., de Gea, J., Edgington, M., Metzen, J. H., Kirchner, F., 2008. Accelerating neuroevolutionary methods using a Kalman filter. In: *GECCO '08: Proceedings of the 10th Annual Conference on Genetic and Evolutionary Computation*. ACM, pp. 1397–1404.
- [26] Kassahun, Y., Wöhrle, H., Fabisch, A., Tabie, M., 2012. Learning parameters of linear models in compressed parameter space. In: Villa, A., Duch, W., Érdi, P., Masulli, F., Palm, G. (Eds.), *Artificial Neural Networks and Machine Learning – ICANN 2012*. Vol. 7553 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, pp. 108–115.
- [27] Koutník, J., Gomez, F. J., Schmidhuber, J., 2010. Evolving neural networks in compressed weight space. In: Pelikan, M., Branke, J. (Eds.), *GECCO*. ACM, pp. 619–626.
- [28] Koutník, J., Gomez, F. J., Schmidhuber, J., 2010. Searching for minimal neural networks in fourier space. In: Baum, E., Hutter, M., Kitzelmann, E. (Eds.), *Proceedings of The Third Conference on Artificial General Intelligence (AGI 2010)*. Atlantic Press, pp. 61–66.
- [29] Krusienski, D., Schalk, G., 2004. Wadsworth BCI dataset (P300 evoked potentials). Data set description.  
URL [http://www.bbc.de/competition/iii/desc\\_II.pdf](http://www.bbc.de/competition/iii/desc_II.pdf)
- [30] Lang, K. J., Witbrock, M. J., 1988. Learning to tell two spirals apart. In: Touretzky, D., Hinton, G., Sejnowski, T. (Eds.), *Proceedings of the 1988 Connectionist Models Summer School*. Morgan Kaufmann, pp. 52–61.

- [31] LeCun, Y., Bengio, Y., 1995. Convolutional Networks for Images, Speech and Time Series. MIT Press, pp. 255–258.
- [32] LeCun, Y., Bottou, L., Orr, G. B., Müller, K.-R., 1998. Efficient back-prop. In: Orr, G. B., Müller, K.-R. (Eds.), *Neural Networks: Tricks of the Trade*. Springer, pp. 9–50.
- [33] Levenberg, K., 1944. A method for the solution of certain problems in least squares. *Quarterly of Applied Mathematics* 2, 164–168.
- [34] Maillard, O.-A., Munos, R., 2009. Compressed least-squares regression. In: Bengio, Y., Schuurmans, D., Lafferty, J., Williams, C. K. I., Culotta, A. (Eds.), *Advances in Neural Information Processing Systems* 22. pp. 1213–1221.
- [35] Marquardt, D., 1963. An algorithm for least-squares estimation of non-linear parameters. *Journal of the Society for Industrial and Applied Mathematics* 11 (2), 431–441.
- [36] Nissen, S., 2003. Implementation of a fast artificial neural network library (fann). Tech. rep., Department of Computer Science University of Copenhagen (DIKU).  
URL <http://leenissen.dk/fann/report/report.html>
- [37] Rumelhart, D. E., Hinton, G. E., Williams, R. J., 1986. Learning representations by back-propagating errors. *Nature* 323 (6088), 533–536.
- [38] Schmidhuber, J., 1995. Discovering solutions with low Kolmogorov complexity and high generalization capability. In: *International Conference on Machine Learning*. Morgan Kaufmann, pp. 488–496.
- [39] Schmidhuber, J., 1997. Discovering neural nets with low Kolmogorov complexity and high generalization capability. *Neural Networks* 10 (5), 857–873.
- [40] Sutton, R. S., Barto, A. G., 1998. *Reinforcement Learning: An Introduction*. MIT Press.
- [41] Vapnik, V. N., 1995. *The Nature of Statistical Learning Theory*. Springer-Verlag New York, Inc.

- [42] Wieland, A. P., 1990. Evolving controls for unstable systems. In: Touretzky, D. S., Elman, J. L., Sejnowski, T. J., Hinton, G. E. (Eds.), *Connectionist Models: Proceedings of the 1990 Summer School*. CA: Morgan Kaufmann, pp. 91–102.
- [43] Wolpaw, J. R., Birbaumer, N., McFarland, D. J., Pfurtscheller, G., Vaughan, T. M., 2002. Brain-computer interfaces for communication and control. *Clinical Neurophysiology* 113 (6), 767–91.