

UNIVERSITY OF GRONINGEN

MASTER THESIS

Controlling Recurrent Neural Networks by Diagonal Conceptors

Author:

J.P. De Jong

Examiners:

Prof. Dr. H. Jaeger

Assistant Professor C.P. Hirsch



July 1, 2021

Contents

	Page
Abstract	4
1 Introduction	5
2 Related Work	8
3 Theory	10
3.1 Mathematical Landscape	10
3.2 Storing the Patterns and Training the Output Weights	11
3.3 Conceptors	12
3.3.1 Computing Conceptors	13
3.3.2 Morphing	14
3.3.3 Autoconceptors	14
3.4 Diagonal Conceptors	15
3.4.1 Computing Diagonal Conceptors	16
3.4.2 Morphing	18
3.4.3 Diagonal autoconceptors	18
4 Simulation Methods	21
4.1 Initializing the Parameters	21
4.2 Training Conceptors	22
4.3 Training Diagonal Conceptors	24
4.4 Pattern Comparison	26
5 Simulations	28
5.1 Periodic Patterns	28
5.1.1 Conceptors	29
5.1.2 Diagonal Conceptors	30
5.2 Chaotic Attractors	32
5.2.1 Conceptors	33
5.2.2 Diagonal Conceptors	33
5.3 Human Motion	35
5.3.1 Conceptors	36
5.3.2 Diagonal Conceptors	38
5.4 Morphing	40
5.4.1 Periodic Patterns	41
5.4.2 Chaotic Attractors	41
5.4.3 Human Motion	43
6 Online adaptation	46
7 Discussion	48
7.1 Training diagonal conceptors	48
7.1.1 Initial Diagonal Conceptors	48
7.1.2 Aperture and Regularization	50

7.1.3	Adaptation Period	51
7.2	Stability	52
8	Conclusion	54
9	Future Work	55
	Bibliography	56

Abstract

The human brain is capable of learning, memorizing, and regenerating a panoply of temporal patterns. A neuro-dynamical mechanism called *conceptors* offers a method for controlling the dynamics of a recurrent neural network by which a variety of temporal patterns can be learned and recalled. However, conceptors are matrices whose size scales quadratically with the number of neurons in the recurrent neural network, hence they quickly become impractical. In the work reported in this thesis, a variation of conceptors is introduced, called *diagonal conceptors*, which are diagonal matrices, thus reducing the computational cost drastically. It will be shown that diagonal conceptors achieve the same accuracy as conceptors, but are slightly more unstable. This instability can be improved, but requires further research. Nevertheless, diagonal conceptors show to be a promising practical alternative to the standard full matrix conceptors.

1 Introduction

In 1997, the world chess champion Garry Kasparov lost to the IBM supercomputer Deep Blue. This was the first time a world chess champion was beaten by a machine [1]. Ever since, the field of Artificial Intelligence (AI) has expanded and grown immensely. Currently, there are several books which are solely devoted to the discussion about the progress of AI, both the good and bad sides of it [1, 2]. One important open question is how close AI research is to creating AI that can outperform humans in most tasks. At present, AI can transcend human intelligence in some individual tasks, such as playing games like Atari, chess, and Go [3], or arithmetic. Furthermore, the field of neural networks has made an impressive leap forward since its creation, instigated by the perceptron [4]. At present, neural networks are researched extensively and are used for a wide range of tasks. Nevertheless, the human brain remains overall superior to AI. One of the reasons is that the human brain is able to learn a panoply of skills consisting of, for instance language, sport, and social interaction. Specifically, the ability to learn, recognize, recall, and combine a vast amount of *temporal* patterns with little effort. It is thus desirable to research a neural network that is capable of such tasks. One of the difficulties is managing long-term memory, which is the ability to permanently store temporal patterns in its synaptic weights such that it is able to recall the learned patterns whenever asked without changing those weights, also referred to as neural long-term memory. There are several researches devoted to this task, which will be discussed in Section 2. However, a neuro-computational mechanism called *conceptors*, not only offers a solution to managing long-term memory for temporal patterns, but also offers insight into a variety of other problems [5].

The conceptors architecture is a neuro-computational mechanism that can be used to control the dynamics of an *Recurrent Neural Network* (RNN), henceforth also referred to as the *reservoir*. Conceptors were first introduced in the technical report *Controlling Recurrent Neural Network by Conceptors* written by H. Jaeger, which will be referred to as the *conceptors report*. In the conceptors report many of the applications of conceptors, such as temporal pattern classification, human motion generation, de-noising and signal separation were demonstrated [5]. Consequentially, conceptors have been successfully used in several other researches [6–13]. It must be noted that conceptors can take different forms, but in this report they will only be in matrix form and will therefore be referred to as *conceptor matrices* or simply *conceptors*.

Intuitively, conceptors act as filters through which an RNN can be controlled. The crucial observation is that when an N -dimensional RNN is driven by different patterns $1, 2, \dots, p$, different areas A^1, A^2, \dots, A^p of the neural state space are being activated. Conceptors exploit this observation, which allows for retrieving a desired pattern from the neural state space. Differently stated, conceptors attempt to identify the subvolumes in which the respective patterns live. A conceptor $C^j \in \mathbb{R}^{N \times N}$ representing a pattern j constrains the neural dynamics to the volume of state space A^j such that the network will regenerate pattern j . It achieves this by leaving the area of state space associated with pattern j for the most part unchanged, yet suppressing the other areas of state space. Conceptors are robust against small parameter changes and noise. However, conceptors can become computationally inefficient. As the number of neurons in the reservoir N increases, the conceptor matrices grow quadratically in N , which poses problems with storage and computational cost. For example, a matrix with 1000 rows and columns uses approximately 8 MB of storage, so for 10 patterns, 80 MB is needed. However, if the number of neurons in the reservoir increases from $N = 1000$ to $N = 5000$, then the conceptor matrix will have 5000 rows and columns, requiring 200 MB of memory for a single conceptor. Not only storage, but also the computational cost increases drastically. If $C \in \mathbb{R}^{N \times N}$ and $\mathbf{x} \in \mathbb{R}^N$, then the complexity of the multiplication $C\mathbf{x}$ is $O(N^2)$. Therefore, conceptor matrices quickly

become unpractical as the reservoir size increases, which is unfortunate, as conceptors carry many useful qualities and it would be ideal if they could be used in practice. An alternative architecture was introduced, called the *Random Feature Conceptor* (RFC) architecture, to overcome this impracticality as well as biological implausibility [5]. The idea of the RFC architecture is to expand the state of the RNN onto a higher-dimensional space, where the state is manipulated by a diagonal conceptor matrix rather than a full conceptor matrix, and finally the manipulated state is projected back to the original space. In this architecture, the diagonal conceptors manipulate the higher-dimensional state element-wise, so it can be written in vector form. This is computationally much more efficient. During the research conducted for this thesis it was discovered that the architecture still works if the dimension of the higher-dimensional space is set equal to dimension of the original space and both projection matrices are set to the identity matrix, i.e. the same architecture as conceptors, but with a diagonal conceptor matrix. However, this only works well if a vital adjustment is made in the training scheme, which will be discussed later. This discovery led to the neuro-computational mechanism that is the subject of this thesis, called *diagonal conceptors*.

Diagonal conceptors offer a solution to the impracticality of conceptors. As can be deduced from the name, diagonal conceptors are diagonal matrices, meaning that they reduce the computational cost significantly and require much less storage. A diagonal conceptor matrix $D \in \mathbb{R}^{N \times N}$ can be written as $D = \text{diag}(\mathbf{c})$, where $\text{diag}(\cdot)$ denotes a diagonal matrix with zeros everywhere except for the diagonal. The diagonal \mathbf{c} is an N -dimensional vector called the *conception vector* and its elements c_i , for $i = 1, \dots, N$, are called the *conception weights*, where the terminology is borrowed from Section 3.14 of the conceptors report [5]. In practice, the conception vector \mathbf{c} is used, but in this report, for notational simplicity, the matrix notation D will be used. In comparison to a conceptor matrix, a conception vector of size 5000 requires only 40 KB of storage, so for 10 patterns, only 0.4 MB is required to store the conception vectors. Moreover, the complexity of computing the conception vector \mathbf{c} times a vector $\mathbf{x} \in \mathbb{R}^N$ is only $O(N)$, which is a considerable decrease compared to $O(N^2)$.

In addition, it will be shown that the conception weights can be trained individually, which provides two advantages over conceptors. First, diagonal conceptors are not biologically implausible, whereas conceptors are. Conceptors can be online adapted in a variation called *autoconceptors*, but the online adaptation requires non-local computations, making the conceptors biologically implausible. Diagonal conceptors, on the other hand, can also be adapted online, for which only local computations are required. Therefore, it can be said that diagonal conceptors are not biologically implausible, in the sense that all information for the adaptation of a synaptic weight is available as the synapse. Second, the global learning rate for conceptors is constrained by the ratio of the smallest to the largest curvature direction in the gradient space, resulting in slow convergence in areas where the curvature is small. The learning rate of the online adaptation of diagonal conceptors, on the other hand, can be chosen for each conception weight individually. Therefore, the constraint of the global learning rate for conceptors is lifted for diagonal conceptors.

Diagonal conceptors are trained differently than conceptors, which is not surprising, since the degrees of freedom of a conceptor is much higher than the degrees of freedom of a diagonal conceptor. The main change in the training scheme is that before the reservoir is driven, a randomly initialized diagonal conceptor is inserted in the update loop. These initial random diagonal conceptors randomly scale each neuron in the reservoir individually. Consequently, the different areas A^1, A^2, \dots, A^p of the neural state space that are being activated by driving the RNN with patterns $1, 2, \dots, p$ are randomly scaled, which creates new areas $\hat{A}^1, \hat{A}^2, \dots, \hat{A}^p$. The diagonal conceptors will be trained on those newly

positioned areas \hat{A}^j rather than the original areas \hat{A}^j , for all j . This is especially useful when different areas in state space overlap, since, in contrast to conceptors, diagonal conceptors have less degrees of freedom than conceptors to characterize the nuances of each area. The use of this initial random scaling is an interesting observation and it offers opportunities for more future work, as it could be researched whether the randomness of the initial random scaling can be optimized.

This report introduces the diagonal conceptors by examples that are also used in the conceptors report [5]. The examples consist of *four periodic patterns*, *chaotic attractors*, and *human motions*. Since each example highlights a different quality of conceptors, it will become evident which characteristics translate to diagonal conceptors and which do not. For each example, both conceptors and diagonal conceptors are trained, allowing a direct comparison to be made.

The objective of this report is to introduce diagonal conceptors so they can be used as a practical alternative for matrix conceptors. Furthermore, this report aims to analyse some of the properties of diagonal conceptors and how they relate the properties of conceptors. Lastly, this report is intended as a intuitive guide to using diagonal conceptors in practice.

2 Related Work

The diagonal conceptors architecture was inspired by the RFC architecture, which is described in Section 3.15 of the conceptors report [5]. Therefore, diagonal conceptors are closely related to RFCs and many of the consequences of RFCs translate to diagonal conceptors. However, it should be noted that RFCs are trained differently than diagonal conceptors. The conception weights in the RFC architecture are trained using an adaptation rule that is evaluated every time step during an adaption period, whereas the conception weights in the diagonal conceptors architecture are computed explicitly after a state collection period. Consequently, the conception weights in the RFC architecture are restricted to either 0 or the range $(0.5, 1]$, whereas the conception weights in the diagonal conceptors architecture are restricted to the broader range $(0, 1)$. This difference in the resulting conception weights changes the dynamics of the reservoir, but it does not affect the translation of the properties of RFCs to diagonal conceptors. For example, diagonal conceptors are also computationally efficient and not biologically implausible, like RFCs, which can be argued by the same reasoning as in Section 3.15 of the conceptors report [5]. Furthermore, the algebraic and logical rules for conception weights also translate to diagonal conceptors. The symbolic interpretation of conceptors is not discussed in this thesis, so an analysis on the algebraic and logical rules for the conception weights of diagonal conceptors is an opportunity for future work.

RFCs have remained mostly unexplored in literature. As of yet, there are only two articles that make use of RFCs. The first article uses RFCs in their model to learn and recognize complex, dynamic stimuli. They conclude that RFCs perform well, but are prone to instability [12]. The second used the RFC architecture to accurately model bistable perception. They conclude that the RFC architecture is a promising model for general human perception [13]. The almost nonexistence of RFCs in the literature may be attributed to the unexplored shortcomings of the model. As was pointed out in one of the articles, it can be difficult for RFCs to converge to a stable solution. In the conceptors report, this was pointed out as well [5]. In the examples in this report, diagonal conceptors are shown to yield stable reservoir dynamics, which offers a promising approach to applying conceptors in practice compared to RFCs.

Other related research can be divided up in two parts. The first is by viewing diagonal conceptors as a variation of conceptors. The scope of the conceptors architecture is much broader than learning, recognizing, and recalling temporal patterns. Conceptors were introduced as a novel perspective on the neuro-symbolic integration problem [5]. The field of neuro-symbolic integration aims to bridge the gap between two fundamentally different paradigms: symbolic systems in AI and neural network systems in AI. In turn, it intends to unite the research of different fields. There are many researches from different fields dedicated to the neuro-symbolic integration problem [14–17]. However, the field of neuro-symbolic integration is outside the scope of this report, so for an easy introduction the interested reader is referred to [18]. As diagonal conceptors are only used for storing and recalling temporal patterns in this report, the neuro-integration problem will not be discussed.

The second way of comparing diagonal conceptors to other works is by viewing the diagonal conceptors architecture as a means to store temporal patterns in neural long-term memory. The article *Using Conceptors to Manage Neural Long-Term Memories for Temporal Patterns* by H. Jaeger [19], shows how conceptors can be used to manage neural long-term memories for temporal patterns. In its introduction, a variety of different models and techniques is discussed that have attempted the same. First, associative memory is considered, which is the paradigmatic model for neural long-term memory. It was introduced by a variety of researches [20–22], only a few of which are cited here.

However, these models are mostly designed for *static* patterns, such as images, and not for temporal patterns. Second, it considers a number of approaches for managing neural memories of temporal patterns, a few of which are an extension of the classical associative memory model. It consists of, but is not limited to: hetero-associative memory [23], multiassociative memory [24], reservoir computing [25]. These approaches will not be discussed in detail here, as they are discussed in detail in [19]. However, the reason they are briefly mentioned here is that they are compared to conceptors in a later part of the introduction of [19]. In this part, a number of properties that are variously attributed to a neural memory system are listed. Each of the mentioned approaches can be attributed a number of properties, but usually they fall short in the others. Conceptors, on the other hand, facilitate all those properties except two. Now, since the argument in favor of conceptors can be translated to diagonal conceptors it follows that diagonal conceptors also facilitate all those properties except the two. Moreover, compared to conceptors, diagonal conceptors have the advantage of computational efficiency as well as the fact that diagonal conceptors are not biologically implausibility. Therefore, for all the advantages of conceptors over the discussed other methods in [19], diagonal conceptors possess the same advantages plus two more.

3 Theory

The first part of this section largely follows the conceptors report [5]. Nevertheless, this section is self-contained and introduces all the required background knowledge for this thesis. It is organized in four parts. First and second, the mathematical landscape of an RNN is introduced as well as how to store patterns in a reservoir and how to train the output weights. Third, conceptors are introduced and defined, followed by a brief introduction of autoconceptors. Fourth, diagonal conceptors are introduced and defined.

As for mathematical notation, there are a few notations that will be consistent throughout this report. A matrix is denoted by a capital letter. A vector is denoted by a bold letter. The element on the i -th row and j -th column of a matrix A is denoted by a_{ij} . The i -th element of a vector \mathbf{a} is denoted by a_i . The transpose of a matrix A is denoted by A^T . The Frobenius norm of a matrix A is denoted by $\|A\|_{fro}$ and is defined, for a real matrix, as $\|A\|_{fro} = \sqrt{\sum_i \sum_j |a_{ij}|^2}$. If the norm has no subscript, i.e. $\|A\|$, the reader can assume that the Frobenius norm is meant.

3.1 Mathematical Landscape

As mentioned in the introduction, the goal is to store a variety of temporal patterns in a single RNN such that they can later be retrieved. With *storing* it is meant that after the patterns are stored in the RNN, the RNN can autonomously regenerate any pattern. Furthermore, *retrieving* a pattern from an RNN means that the autonomous regeneration of any stored pattern can be instigated at any time such that the desired pattern is outputted from the RNN. The RNN will now be made explicit.

Let $\mathbf{P} = \{\mathbf{p}^1, \mathbf{p}^2, \dots, \mathbf{p}^p\}$ be a set of p discrete-time patterns, where pattern j at time step n is given by $\mathbf{p}^j(n) \in \mathbb{R}^M$. Furthermore, assume M input neurons as well as a bias neuron, an RNN consisting of N simple tanh neurons called *reservoir neurons*, and M' output neurons. In the more general case one can assume $M' \neq M$ output neurons, which would be the case if the driving patterns are used to produce output based on the driving patterns, e.g. classification. However, in this report the reservoir and output neurons serve as a means to self-generate the driving patterns, hence the number of input and output neurons is set equal. The input neurons drive the reservoir with $\mathbf{p}^j(n)$ via an input weights matrix $W^{in} \in \mathbb{R}^{N \times M}$ and a bias vector $\mathbf{b} \in \mathbb{R}^N$. The N neurons in the RNN are recurrently connected via a weights matrix $W^* \in \mathbb{R}^{N \times N}$. When the reservoir is driven by pattern $\mathbf{p}^j(n)$, the reservoir neurons are activated. The state of neuron i at time n is denoted by $x_i^j(n) \in (-1, 1)$ and is called the *neuron state*. The state of the reservoir at time n , denoted by $\mathbf{x}^j(n) \in (-1, 1)^N$, is given by the vector containing all the neuron states at time n and is called the *reservoir state*. Finally, the output neurons serve to read out the target signal at time step n , denoted by $\mathbf{y}^j(n) \in \mathbb{R}^{M'}$, from the reservoir state via the output weights matrix $W^{out} \in \mathbb{R}^{M' \times N}$. A visual representation of the setup can be seen in Figure 1.

Let W^* , W^{in} , and \mathbf{b} be fixed random matrices. They will not be adjusted after initialization. The reservoir state vector, or simply *state vector*, is updated according to the update equation

$$\mathbf{x}^j(n+1) = \tanh(W^* \mathbf{x}^j(n) + W^{in} \mathbf{p}^j(n+1) + \mathbf{b}) \quad (1)$$

and the output signal is given by

$$\mathbf{y}^j(n) = W^{out} \mathbf{x}^j(n), \quad (2)$$




Figure 1: Basic setup of the RNN. The input neurons drive the reservoir with pattern \mathbf{p} through an input weights matrix W^{in} as well as a bias vector \mathbf{b} . The reservoir neurons are interconnected via a reservoir weights matrix W^* . The output \mathbf{y} can be read from the reservoir state \mathbf{x} via an output weights matrix W^{out} .

where the output weights W^{out} are learned, see Section 3.2. Let the reservoir be driven for L time steps. The state vectors $\mathbf{x}^j(n)$ are collected in the *state collection matrix* X^j . The state collection matrix gives a sample of states that is representative of the volume of state space that is occupied by the driving pattern j and is given by

$$X^j = \begin{bmatrix} x_1^j(1) & x_1^j(2) & \dots & x_1^j(L) \\ x_2^j(1) & x_2^j(2) & \dots & x_2^j(L) \\ \vdots & & \ddots & \vdots \\ x_N^j(1) & x_N^j(2) & \dots & x_N^j(L) \end{bmatrix}. \quad (3)$$

Note that the information about pattern j is encoded in the X^j .

3.2 Storing the Patterns and Training the Output Weights

Up until this point, the reservoir has been driven by the input patterns. However, in order to store the patterns, the reservoir must be able to update the state vector in the absence of a driving input, which yields the following approximation:

$$\tanh(W^* \mathbf{x}^j(n) + W^{in} \mathbf{p}^j(n+1) + \mathbf{b}) \approx \tanh(W \mathbf{x}^j(n) + \mathbf{b}), \quad (4)$$

where W comprises the *recomputed reservoir weights*. W can be computed by minimizing the mean square error over all j and n , yielding

$$W = \operatorname{argmin}_{\tilde{W}} \sum_j \sum_n \|W^* \mathbf{x}^j(n) + W^{in} \mathbf{p}^j(n+1) - \tilde{W} \mathbf{x}^j(n)\|^2. \quad (5)$$

For such linear regression tasks, ridge regression will be used throughout this report. Furthermore, it should be noted that there are alternative techniques for removing the driving input from the update equations, which are discussed in Section 3.11.1 of the conceptors report [5].

For computing W^{out} , a similar approach is employed, where the output $\mathbf{y}^j(n) = W^{out} \mathbf{x}^j(n)$ must approximate the input pattern $\mathbf{p}^j(n)$. This gives the approximation

$$\mathbf{p}^j(n) \approx \mathbf{y}^j(n) = W^{out} \mathbf{x}^j(n), \quad (6)$$

where W^{out} is then computed by minimizing the mean square error over all j and n , which yields

$$W^{out} = \operatorname{argmin}_{\tilde{W}^{out}} \sum_j \sum_n \|\mathbf{p}^j(n) - \tilde{W}^{out} \mathbf{x}^j(n)\|^2. \quad (7)$$

The process of recomputing the reservoir weights W and computing the output weights W^{out} is referred to as *storing* the patterns in the reservoir. The reservoir is called *loaded* after the patterns have been stored.

3.3 Conceptors

After the patterns are stored in the reservoir, there exists a superposition of patterns in the reservoir. If the network were to simply start updating the state vector by $\mathbf{x}(n+1) = \tanh(W\mathbf{x}(n) + \mathbf{b})$ it would exhibit unpredictable behavior, since the reservoir does not know which pattern to engage in. So, how is a specific pattern retrieved from the reservoir? Note that driving the reservoir with pattern j creates a cloud of points in state space, which is characteristic for pattern j . This point cloud is given by the columns of the state collection matrix X^j . Ideally, the point cloud associated with pattern j is confined to a proper subspace $S^j \subset \mathbb{R}^N$ such that the state vector can be projected on S^j via a projection matrix \mathbf{P}_{S^j} . However, the state vectors usually span the entire space \mathbb{R}^N , so $\mathbf{P}_{\mathbb{R}^N}$ would be the identity matrix, which would not be helpful in singling out an individual pattern. Therefore, instead of applying the projection $\mathbf{P}_{\mathbb{R}^N}$, the state vector is projected on a subspace that is spanned by a number of leading principal components of the point cloud associated with pattern j . How many of the leading principal components should be used in this soft projection is adjustable by means of a control parameter called *aperture*¹. Each pattern j will have its own projection matrix, which is called the *conceptor* associated with pattern j and is denoted by C^j . The conceptor matrix C^j should behave such that it leaves states associated with pattern j intact, but suppresses the states of other patterns. This trade-off is mitigated by the aperture.

Applying C^j to the state vector can be viewed as a neural network consisting of two layers, where the first layer is projected to the second layer via C^j and is projected back via W . Therefore, the update equation is written in two parts

$$\begin{aligned} \mathbf{r}(n+1) &= \tanh(W\mathbf{z}(n) + \mathbf{b}), \\ \mathbf{z}(n+1) &= C^j \mathbf{r}(n+1), \end{aligned} \quad (8)$$

where the output is given by

$$\mathbf{y}(n) = W^{out} \mathbf{z}(n). \quad (9)$$

A visual representation is depicted in Figure 2.

¹The name aperture has its roots in optics, where it means the diameter of the effective lens opening, hence it negotiates the amount of light energy that reaches the film.




Figure 2: A visual representation of Equations 8 and 9. The state \mathbf{r} is projected to a state \mathbf{z} via a conceptor matrix C and \mathbf{z} is inserted back in the loop via W . The output is read from the state \mathbf{z} by the output weights W^{out} .

3.3.1 Computing Conceptors

Conceptor C^j must leave states $\mathbf{z}^j(n)$ unchanged, but also suppress components that are not typical for $\mathbf{z}^j(n)$. This essentially means that C^j should act as the identity matrix for states $\mathbf{z}^j(n)$, but as the null matrix for the unwanted components. This leads to the following quadratic loss function:

$$\begin{aligned} \mathcal{L}(C^j) &= \mathbb{E}[\|\mathbf{z}^j - C^j \mathbf{z}^j\|^2] + \alpha^{-2} \|C^j\|^2 \\ &= \text{Tr}[(I - (C^j)^T)(I - C^j)] \mathbb{E}[\mathbf{z}^j (\mathbf{z}^j)^T] + (\alpha^j)^{-2} \text{Tr}[(C^j)^T C^j], \end{aligned} \quad (10)$$

where α^j is the aperture associated with pattern j and $\|\cdot\|$ is the Frobenius norm.

The first component of $\mathcal{L}(C^j)$ is the time-average of the difference between the projected state vectors $C^j \mathbf{z}^j$ and the state vector \mathbf{z}^j . This reflects that C^j must leave \mathbf{z}^j unchanged. This component is optimal for $C^j = I$. The second part of the loss function represents the suppression of the outlying components of the state vector. Note that it is minimal for C^j equals the null matrix. Therefore, a large aperture results in a conceptor matrix close to the identity matrix. Conversely, a small aperture will shrink the conceptor matrix towards the null matrix.

To find the matrix C^j for which $\mathcal{L}(C^j)$ is minimal, the gradient of $\mathcal{L}(C^j)$ with respect to C^j is computed. The gradient is given by

$$\frac{\partial}{\partial C^j} \left(\mathbb{E}[\|\mathbf{z}^j - C^j \mathbf{z}^j\|^2] + (\alpha^j)^{-2} \|C^j\|^2 \right) = -2(I - C^j) \mathbb{E}[\mathbf{z}^j (\mathbf{z}^j)^T] + 2(\alpha^j)^{-2} C^j. \quad (11)$$

It is then straightforward to see that the minimization of $\mathcal{L}(C^j)$ leads to the following solution

$$C^j = R^j (R^j + (\alpha^j)^{-2} I)^{-1}, \quad (12)$$

where $R^j = \mathbb{E}[\mathbf{z}^j (\mathbf{z}^j)^T]$ is the *state correlation matrix* associated with pattern j . For $\alpha^j \in (0, \infty)$, the solution is well-defined, but for $\alpha^j = 0$ it is not. The cases of $\alpha^j = 0$ and $\alpha^j = \infty$ will not be

discussed here, but they are discussed in Section 3.8.1 of the conceptors report [5]. In simulations, the state correlation matrix R^j is estimated by $\hat{R}^j = Z^j(Z^j)^T / L$, where $Z^j = [\mathbf{z}^j(1) \mathbf{z}^j(2) \dots \mathbf{z}^j(L)] \in \mathbb{R}^{N \times L}$ is the state collection matrix from Section 3.1.

A few properties of C^j and R^j can be inferred and are worth mentioning:

1. C^j and R^j have the same N eigenvectors, meaning that if $R^j = U\Sigma U^T$ is the singular value decomposition of R^j then $C^j = USU^T$, where Σ and S contain the singular values of R^j and C^j , respectively.
2. The singular values of C^j , denoted by s_i^j , are the normalized singular values of R^j , denoted by σ_i^j , and relate by $s_i^j = \sigma_i^j(\sigma_i^j + (\alpha^j)^{-2})^{-1}$.
3. From the previous property it can be concluded that $0 \leq s_i^j \leq 1$.

3.3.2 Morphing

The dynamics of the reservoir can be morphed by conceptors. The term morphing is usually used to describe the smooth transition from one image to another by slow gradual interpolation, but here it is used to describe a smooth transition between reservoir dynamics. Conceptors can be used to morph different patterns that are stored in the reservoir. Let the reservoir be loaded with patterns j_1 and j_2 and let the corresponding conceptors C^{j_1} and C^{j_2} be given. Then, a mixture of the reservoir dynamics associated with \mathbf{p}^{j_1} and \mathbf{p}^{j_2} can be obtained by a linear combination of C^{j_1} and C^{j_2} , given by

$$\begin{aligned} \mathbf{r}(n+1) &= \tanh(W\mathbf{z}(n) + \mathbf{b}), \\ \mathbf{z}(n+1) &= ((1-\mu)C^{j_1} + \mu C^{j_2})\mathbf{r}(n+1), \end{aligned} \quad (13)$$

where $\mu \in \mathbb{R}$ is the *mixture parameter*. Notice that the mixture parameter μ is not constrained to $[0, 1]$. As will be shown in the simulations in Section 5.4, conceptors are not only able to interpolate, but also extrapolate. In the conceptors report, it was shown that for two sine waves with different periods that conceptors are capable of capturing the period of the sine waves and then not only interpolate between them, but also extrapolate [5]. The periods in the example were ≈ 8.83 and 9.83 , but with the extrapolation, the conceptors were able to create sine waves of periods between ≈ 7.5 and ≈ 11.9 .

Furthermore, it must be noted that a morph is not restricted to two patterns. In the more general case, a morph can be performed for p patterns, where the term $(1-\mu)C^{j_1} + \mu C^{j_2}$ would be substituted by $\sum_{i=1}^p \mu_i C^{j_i}$. In this sum, the mixing parameter μ_i determines how much the outputted pattern is influenced by the dynamics of pattern j_i . The simulations in this report only make use of the case where $p = 2$.

3.3.3 Autoconceptors

Thus far, conceptors are computed with Equation 12, after which they are stored so the dynamics of the reservoir can be constrained at a later time. This is practical and achievable for machine learning applications, where the conceptors can be written to a file. However, this is not always the case. Specifically, from a neuroscience point of view, it seems unlikely that a filter is created for every new pattern that needs to be learned, where the filter has the same size as the reservoir. This compels for a different architecture, where the conceptors are created while the reservoir is being driven and the conceptors need not be stored.

The idea is to make a conceptor matrix time-dependent, so it can be adapted while the reservoir is being driven. The update equations in Equation 8 then change to

$$\begin{aligned}\mathbf{r}(n+1) &= \tanh(W\mathbf{z}(n) + \mathbf{b}), \\ \mathbf{z}(n+1) &= C(n)\mathbf{r}(n+1),\end{aligned}\tag{14}$$

where the only difference is that C is now time-dependent. The adaptation rule for the online adaptation of $C(n)$ can be read directly from Equation 11, yielding

$$C(n+1) = C(n) + \lambda \left((I - C(n))\mathbf{z}(n)\mathbf{z}^T(n) - \alpha^{-2}C(n) \right),\tag{15}$$

where $\lambda > 0$ is the learning rate. Note that λ is constrained by the ratio of the smallest to the largest curvature direction in the gradient space. This leads to slow convergence in areas where the local curvature is lower.

It was shown that if the autoconceptor $C(n)$ converges, it will possess the same algebraic properties as a conceptor [5]. Autoconceptors comprise a large part of the conceptors report, so for more details the reader is referred there.

3.4 Diagonal Conceptors

There are two main areas in which conceptors fall short. First, conceptors are computationally expensive. As mentioned before, each pattern that needs to be stored in the reservoir, also requires a conceptor matrix, which has the same size as the reservoir. Furthermore, the conceptor matrices quickly increase in size as they scale quadratically with the number of reservoir neurons N . Therefore, conceptors could hardly be used in real-world applications where the dimension of the reservoir is large. Second, conceptors are biologically implausible. The online adaptation of an element of the conceptor matrix, denoted by C_{ij} , requires information that would biologically not all be available at the synapse of C_{ij} .

In Section 3.15 of the conceptors, an architecture, called *Random Feature Conceptors* (RFC), is introduced that solves the shortcomings of conceptors [5]. In short, the idea of random feature conceptors is to project the reservoir state to a higher-dimensional space, where it is manipulated by a diagonal conceptor matrix rather than a full conceptor matrix, and finally the manipulated state is projected back to the original space. In this architecture, the diagonal matrix conceptors manipulate the higher-dimensional state element-wise, so it can be written in vector form. This is computationally much more efficient and is not biologically implausible, which is explained in more detail in Section 3.15 of the conceptors report [5].

Here, an architecture, called *diagonal conceptors*, is proposed that is closely related to RFC. In contrast to RFC, in the diagonal conceptors architecture, the projection to the higher-dimensional space is unnecessary. Therefore, the conceptor matrix C^j in Equation 8 can be simply substituted with a *diagonal conceptor matrix* D^j . It will be shown that this architecture only works well if an adjustment is made in the training algorithm of conceptors, which will be made clear in Section 4.3. This section discusses how the diagonal conceptors would theoretically be computed.

A diagonal matrix is a matrix in which all the off-diagonal entries are zero. Let $\mathbf{c} = [c_1 \ c_2 \ \dots \ c_N] \in \mathbb{R}^N$,

then a diagonal matrix $D \in \mathbb{R}^{N \times N}$ is often denoted by

$$D = \text{diag}(\mathbf{c}) = \begin{bmatrix} c_1 & 0 & \dots & 0 \\ 0 & c_2 & \dots & 0 \\ \vdots & 0 & \ddots & \vdots \\ 0 & 0 & \dots & c_N \end{bmatrix}. \quad (16)$$

Let $D^j = \text{diag}(\mathbf{c}^j) \in \mathbb{R}^{N \times N}$ denote the *diagonal conceptror matrix*, or simply *diagonal conceptror*, associated with pattern j , where $\mathbf{c}^j = [c_1^j \ c_2^j \ \dots \ c_N^j] \in \mathbb{R}^N$ is called the *conception vector* and its elements c_i^j , for $i = 1, 2, \dots, N$, are called the *conception weights*. Conception vector and conception weights are terminology that is borrowed from the RFC architecture, described in Section 3.15 of the conceptrors report [5]. In Equation 8, the conceptror matrix C^j is substituted for the diagonal conceptror matrix D^j , which yields

$$\begin{aligned} \mathbf{r}(n+1) &= \tanh(W\mathbf{z}(n) + \mathbf{b}) \\ \mathbf{z}(n+1) &= D^j \mathbf{r}(n+1) \end{aligned} \quad (17)$$

and the output can be read from the reservoir according to

$$\mathbf{y}(n) = W^{\text{out}} \mathbf{z}(n). \quad (18)$$

Note that $\mathbf{z}(n)$ is now a scaling of $\mathbf{r}(n)$, where the scaling parameters are given by \mathbf{c}^j . A convenient way to write Equation 17 is by means of the conception weights. This gives

$$\begin{aligned} \mathbf{r}(n+1) &= \tanh(W\mathbf{z}(n) + \mathbf{b}) \\ \mathbf{z}(n+1) &= \mathbf{c}^j \circ \mathbf{r}(n+1), \end{aligned} \quad (19)$$

where \circ is the Hadamard, or element-wise, multiplication operator. A visualization of Equations 18 and 19 is shown in Figure 3, where a dashed line depicts a simple scaling operation. Simulations make use of Equation 19 rather than Equation 17, because $\mathbf{z}(n)$ can then be computed by vector-vector multiplication rather than matrix-vector multiplication.

Diagonal conceptrors are computationally cheaper than conceptrors, as they allow for vector storage and vector-vector multiplication. Furthermore, diagonal conceptrors are not biologically implausible, which will be discussed further in Section 3.4.3. However, it must be noted that the diagonal conceptrors architecture does not imply biological plausibility. It merely suggest that it is not biologically implausible.

3.4.1 Computing Diagonal Conceptrors

The derivation of a conceptror matrix can be translated directly to diagonal conceptrors. Moreover, instead of deriving an expression for the diagonal conceptror D^j , an expression for the individual conception weights c_i^j is derived. Similar to conceptror matrices, the diagonal conceptror associated with pattern j must leave the states associated with pattern j untouched, while suppressing the states associated with other patterns. Again, this balance is mitigated by the parameter *aperture*, denoted by α^j . This leads to a loss function similar to the loss function for a conceptror C^j . The loss function for diagonal conceptror D^j is given by

$$\mathcal{L}(D^j) = \mathbb{E}[\|\mathbf{z}^j - D^j \mathbf{z}^j\|^2] + (\alpha^j)^{-2} \|D^j\|^2, \quad (20)$$




Figure 3: A visual representation of Equations 18 and 19. The state \mathbf{r} is element-wise scaled to the state \mathbf{z} . This is denoted by the dashed lines, where each dashed line depicts a scaling by a conception weight c_i . The state \mathbf{z} is inserted back in the loop via W . The output is read from the state \mathbf{z} by the output weights W^{out} .

Since D^j is a diagonal matrix, Equation 20 can be more conveniently written element-wise, which gives

$$\begin{aligned} \mathcal{L}(c_i^j) &= \mathbb{E}[(z_i^j - c_i^j z_i^j)^2] + (\alpha^j)^{-2} (c_i^j)^2 \\ &= (1 - c_i^j)^2 \mathbb{E}[(z_i^j)^2] + (\alpha^j)^{-2} (c_i^j)^2. \end{aligned} \quad (21)$$

An expression for c_i^j is found by minimizing the loss function $\mathcal{L}(c_i^j)$, for which the derivative with respect to c_i^j must be computed. This gives

$$\frac{\partial}{\partial c_i^j} \left((1 - c_i^j)^2 \mathbb{E}[(z_i^j)^2] + (\alpha^j)^{-2} (c_i^j)^2 \right) = -2(1 - c_i^j) \mathbb{E}[(z_i^j)^2] + 2(\alpha^j)^{-2} c_i^j. \quad (22)$$

Setting this equal to 0 and solving for c_i^j immediately yields

$$c_i^j = \frac{\mathbb{E}[(z_i^j)^2]}{\mathbb{E}[(z_i^j)^2] + (\alpha^j)^{-2}}. \quad (23)$$

This expression indicates that $0 < c_i^j < 1$, where the boundaries 0 and 1 require a bit more attention.

First, consider the value $\mathbb{E}[(z_i^j)^2]$. Note that $\mathbb{E}[(z_i^j)^2]$ can be written as $(c_i^j)^2 \mathbb{E}[(r_i^j)^2]$, where $0 < \mathbb{E}[(r_i^j)^2] < 1$, because of the hyperbolic tangent function. Therefore, $\mathbb{E}[(z_i^j)^2]$ will come close to 1, but it will never reach 1. In addition, $\mathbb{E}[(z_i^j)^2]$ will only be 0 if $c_i^j = 0$. Therefore, $0 \leq \mathbb{E}[(z_i^j)^2] < 1$.

Second, consider the value α^j . In simulations, it is assumed that $0 < \alpha^j < \infty$, for which c_i^j is always well-defined. However, for completeness, one can look at the boundary values $\alpha^j = 0$ and $\alpha^j = \infty$.

by considering the behavior of c_i^j in the limits $\alpha^j \downarrow 0$ and $\alpha^j \rightarrow \infty$. Let $0 \leq \mathbb{E}[(z_i^j)^2] < 1$ be constant. Then, c_i^j in the limiting values of $\alpha^j \downarrow 0$ and $\alpha^j \rightarrow \infty$ is given by

$$\lim_{\alpha^j \downarrow 0} c_i^j = \lim_{\alpha^j \downarrow 0} \frac{\mathbb{E}[(z_i^j)^2]}{\mathbb{E}[(z_i^j)^2] + (\alpha^j)^{-2}} = 0 \quad \text{and} \quad \lim_{\alpha^j \rightarrow \infty} c_i^j = \lim_{\alpha^j \rightarrow \infty} \frac{\mathbb{E}[(z_i^j)^2]}{\mathbb{E}[(z_i^j)^2] + (\alpha^j)^{-2}} = 1. \quad (24)$$

There is an intuitive relation between α^j and $\mathbb{E}[(z_i^j)^2]$, which is clearer if Equation 23 is written as

$$c_i^j = \frac{1}{1 + \frac{1}{(\alpha^j)^2 \mathbb{E}[(z_i^j)^2]}}. \quad (25)$$

Notice that excited neurons, i.e. $\mathbb{E}[(z_i^j)^2]$ close to 1, yield larger values of c_i^j , hence those neuron states remain mostly untouched. Conversely, less excited neurons, i.e., $\mathbb{E}[(z_i^j)^2]$ close to 0, yield smaller values of c_i^j , hence those neuron states are suppressed. How much the neurons are untouched or suppressed depends on α^j . Increasing α^j also increases c_i^j and, conversely, decreasing α^j also decreases c_i^j . However, this behavior should not come as a surprise, as this is exactly how the loss function was designed.

3.4.2 Morphing

Morphing conceptors can be directly translated to diagonal conceptors. Let the reservoir be loaded with patterns j_1 and j_2 and let the associated diagonal conceptors D^{j_1} and D^{j_2} be given. Then, similar to morphing conceptors, a mixture of the reservoir dynamics associated with patterns j_1 and j_2 is obtained by taking a linear combination of D^{j_1} and D^{j_2} as follows:

$$\begin{aligned} \mathbf{r}(n+1) &= \tanh(W\mathbf{z}(n) + \mathbf{b}), \\ \mathbf{z}(n+1) &= ((1-\mu)D^{j_1} + \mu D^{j_2})\mathbf{r}(n+1), \end{aligned} \quad (26)$$

where $\mu \in \mathbb{R}$ is the mixture parameter. As mentioned in Section 3.3.2, the mixture parameter for morphing conceptors could be used to not only interpolate between reservoir states, but also extrapolate. However, in the examples shown in Section 5.4, interpolation was possible, but extrapolation was not.

Similar to morphing with conceptors, the morph is not restricted to two patterns. Again, in the more general case, a morph could be performed for p patterns, where the term $(1-\mu)D^{j_1} + \mu D^{j_2}$ would be substituted for $\sum_{i=1}^p \mu_i D^{j_i}$. The morphing simulations in this report only make use of the case where $p = 2$.

3.4.3 Diagonal autoconceptors

The arguments for autoconceptors can be translated directly to diagonal conceptors, creating diagonal autoconceptors. If the diagonal conceptor matrix D is made time-dependent, it yields the following update equations:

$$\begin{aligned} \mathbf{r}(n+1) &= \tanh(W\mathbf{z}(n) + \mathbf{b}), \\ \mathbf{z}(n+1) &= D(n)\mathbf{r}(n+1). \end{aligned} \quad (27)$$

The adaptation rule for $D(n) = \text{diag}(\mathbf{c}(n))$, where $\mathbf{c}(n)$ is the time-dependent conception vector, can be written element-wise, and is read directly from Equation 22. This yields

$$c_i(n+1) = c_i(n) + \lambda_i \left((1 - c_i(n)) z_i^2(n) - \alpha^{-2} c_i(n) \right), \quad (28)$$

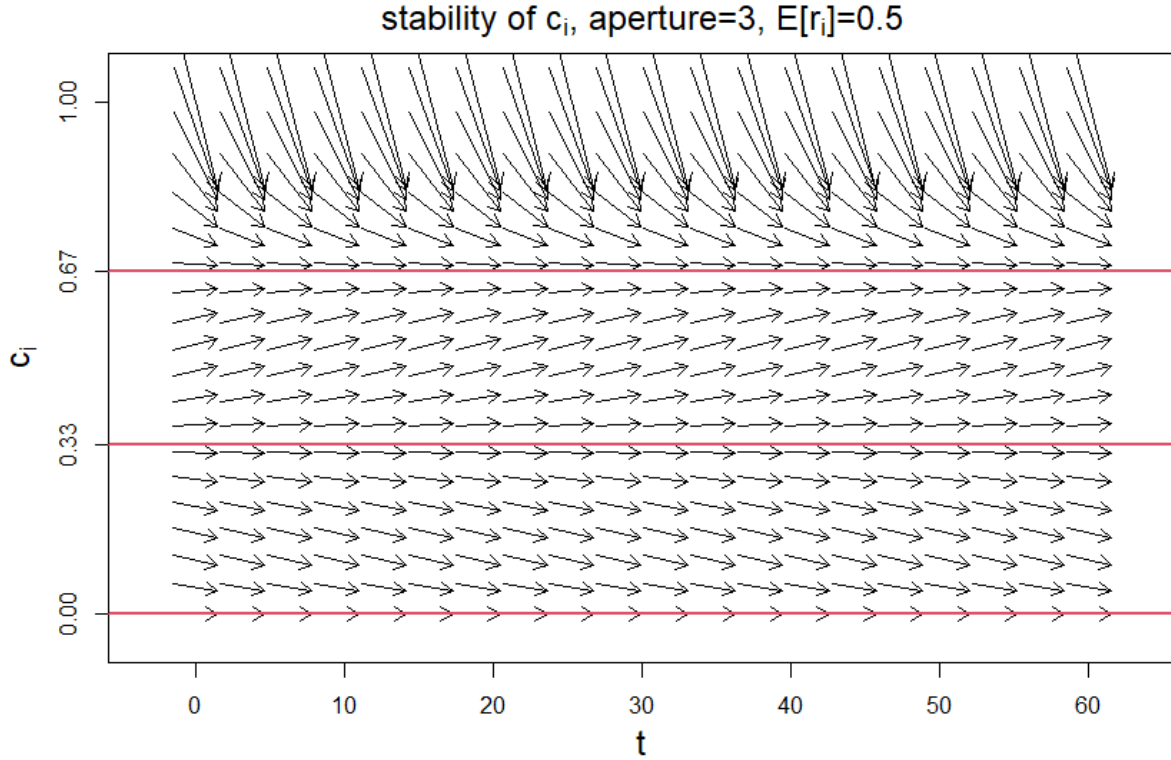


Figure 4: A direction field of Equation 29. The solutions, as computed in Equation 30, are plotted as red lines.

where $c_i(n)$ and $z_i(n)$ are the i -th elements of $\mathbf{c}(n)$ and $\mathbf{z}(n)$, respectively, and λ_i is the learning rate associated with neuron i . Note that the learning rate can be set individually for each neuron, hence lifting the constraint of a global learning rate for autoconceptors. Furthermore, the biological implausibility of autoconceptors is lifted as well, since each neuron state is updated independently of each other. Therefore, in contrast to conceptors, all the information required for updating a conception weight at a synapse is available at that synapse. This does not directly imply biological plausibility, but it certainly does not deny it. This is argued in more detail in Section 3.15 of the conceptors report and this argumentation can be used for diagonal autoconceptors as well [5], so for more information the reader is referred there.

Furthermore, the dynamics of diagonal autoconceptors are easier to analyse than the dynamics of autoconceptors, because it comprises scalars instead of matrices. A short, intuitive analysis is given here. It will only grasp a small portion of the full dynamics of the reservoir, but it will give some insight nonetheless. Writing the adaptation rule given in Equation 28 as a continuous-time differential equation and letting $z_i = c_i r_i$ gives

$$\dot{c}_i(t) = (1 - c_i(t))c_i^2(t)r_i^2(t) - \alpha^{-2}c_i(t). \quad (29)$$

Setting this equal to 0 and solving for $c_i(t)$ gives

$$c_{i,\pm} = \frac{1}{2} \pm \frac{1}{2} \sqrt{1 - 4\alpha^{-2}(\mathbb{E}[r_i^2])^{-1}} \quad \text{and} \quad c_i = 0. \quad (30)$$

This is a special case of what is more generally shown in Section 3.14.4 of the conceptors report [5].

Upon closer inspection it turns out that only $c_{i,+}$ and $c_i = 0$ are stable solutions, which is neatly shown in Figure 4.

Figure 4 shows that over time the value of c_i will either go towards the upper red line, corresponding to $c_{i,+}$ in Equation 30, or $c_i = 0$. It can also be seen that $c_{i,-}$ is not a stable solution. Furthermore, the curvature is much larger for values above $c_{i,+}$ than for values below $c_{i,+}$, hence the convergence rate in those areas will most likely differ. Therefore, in simulations one could, for example, increase the learning rate for values of c_i smaller than 0.5 to increase the rate of convergence.

At the present moment, diagonal autoconceptors remain mostly unexplored. A more comprehensive analysis is required to fully understand the dynamics of this system.

4 Simulation Methods

Several simulations will be discussed in Section 5 to demonstrate the capabilities of diagonal conceptors compared to conceptors. This section describes how those simulations were set up. In addition, this section describes how to train conceptors and diagonal conceptors. The last part of this section describes how the accuracy of the conceptors and diagonal conceptors is determined.

4.1 Initializing the Parameters

When setting up computer simulations, the RNN must be initialized appropriately. How many connections between reservoir neurons are optimal and how should they be weighted? From what distribution should they be drawn? The same questions can be asked for the input weights and the bias vector. This section introduces some of the rules of thumb for initializing the simulation parameters.

Remember that the system is given by the equations

$$\mathbf{x}^j(n+1) = \tanh(W^* \mathbf{x}^j(n) + W^{in} \mathbf{p}^j(n+1) + \mathbf{b}) \quad (31)$$

where $\mathbf{x}^j(n) \in \mathbb{R}^N$ is the reservoir state associated with pattern j at time step n , $\mathbf{p}^j(n) \in \mathbb{R}^M$ is the n -th time step of pattern j , $\mathbf{b} \in \mathbb{R}^N$ is the bias vector, $W^* \in \mathbb{R}^{N \times N}$ are the reservoir weights, and $W^{in} \in \mathbb{R}^{N \times M}$. The simulations in this report will use a slightly different system consisting of so-called *leaky-integrated* neurons, which changes Equation 31 to

$$\begin{aligned} \tilde{\mathbf{x}}^j(n+1) &= \tanh(W^* \mathbf{x}^j(n) + W^{in} \mathbf{p}^j(n+1) + \mathbf{b}) \\ \mathbf{x}^j(n+1) &= (1 - \alpha_l) \mathbf{x}^j(n) + \alpha_l \tilde{\mathbf{x}}^j(n+1) \end{aligned} \quad (32)$$

where $0 \leq \alpha_l \leq 1$ is the leaking rate. Leaky-integrated neurons are an extension of the non-leaky integrated neurons that were described so far, since Equation 31 is recovered for $\alpha_l = 1$, and they offer better fine-tuning of the reservoir dynamics [26]. The leaking rate is a so-called *exponential smoother* and it determines how much of the previous state is still remembered in computing the next state. Adding the leaking rate into the update equation does not change anything that has been discussed up until this point. Intuitively, lowering the leaking rate results in a smoother output pattern, since big deviations are slightly suppressed.

The elements of W^{in} , \mathbf{b} and W^* are all drawn from the standard normal distribution and scaled appropriately. The input weights matrix and bias vector are full matrices, whereas the reservoir weights matrix is a sparse matrix. The sparsity of the reservoir weights matrix is determined by the number of neurons and can be chosen in different manners [27]. For simulations discussed in this report, the sparsity was set to 1 if $N < 20$ and $10/N$ if $N \geq 20$, where N is the number of reservoir neurons.

The scaling of W^{in} and \mathbf{b} mostly determines how much the neurons in the reservoir are excited. For example, lower scalings of W^{in} and \mathbf{b} result in less excited states. Intuitively, if the neurons are too excited their values will be close to -1 or 1 , which will make the neurons act more in a binary switching way. This is not ideal for smooth patterns. On the other hand, if the neurons are not excited enough, their values will be close to 0 , which is desired for very linear tasks, but not for non-linear tasks.

Furthermore, the scalings of W^{in} , \mathbf{b} , and the spectral radius of W^* must be appropriate to ensure that the RNN possesses the so-called *echo state property* [28]. The echo state property ensures that the

RNN "forgets" the initial state. Often, the initial state of a network is not known and will simply be set to the all-zero vector. Therefore, the network must be given some time to *washout* the initial state of the system. This is ensured via a washout period in the simulation. However, the initial state of the RNN must also be forgotten and the echo state property makes sure this happens. The length of the washout period, denoted by $n_{washout}$, can be found by driving the reservoir from two initial states, e.g. the all-zero vector and the all-ones vector, and then determining how long it takes for the two reservoir states to converge, up to a residual error, to the same state. Note that $n_{washout}$ depends on the other parameters of the simulation, e.g., a lower leaking rate will require a longer washout period.

Setting the parameters mentioned above is mainly a matter of trial and error. In order to get some intuition into the parameters, it is advised to manually adjust a single parameter and see how it changes the dynamics of the reservoir or the outputs. It is possible to implement parameter optimization algorithms, but when the output requires visual inspection it may be better to try manually adjusting the parameters one-by-one. Much research has been conducted in finding good scaling parameters, which has been concisely summarized in the paper by M. Lukoševičius, *A Practical Guide to Applying Echo State Networks* [27]. For more intuition on setting up the parameters, the reader is referred to that paper.

4.2 Training Conceptors

This section describes how conceptors are trained. Let there be a set of patterns $\mathbf{P} = \{\mathbf{p}^1, \mathbf{p}^2, \dots, \mathbf{p}^p\}$, where $\mathbf{p}^j \in \mathbb{R}^M$ for all j . Initialize the following:

- number of neurons in the reservoir (N),
- input weights ($W^{in} \in \mathbb{R}^{N \times M}$): Random values drawn from the standard normal distribution and scaled appropriately,
- bias ($\mathbf{b} \in \mathbb{R}^N$): Random values drawn from the standard normal distribution and scaled appropriately,
- initial reservoir weights ($W^* \in \mathbb{R}^{N \times N}$): Sparse matrix with random values drawn from the standard normal distribution and scaled appropriately,
- leaking rate α_l .

Washout

Drive the reservoir with pattern j for $n_{washout}$ time steps according to Equation 32. All information during this period is discarded. The initial state can be initialized as any random vector in $[-1, 1]^N$, but usually $\mathbf{x}^j(0) = \mathbf{0}$, for all j , for simplicity.

Learning

The reservoir is continued to be driven for another n_{learn} time steps. During the learning phase, the aim is to compute:

- $R^j = X^j(X^j)^T / n_{learn}$, where X^j is the state collection matrix,
- W , such that $W\mathbf{x}^j(n) \approx W^*\mathbf{x}^j(n) + W^{in}\mathbf{p}^j(n)$ for all n, j ,
- W^{out} , such that $W^{out}\mathbf{x}^j(n) \approx \mathbf{p}^j(n)$ for all n, j ,

so the following vectors are collected:

- $\mathbf{x}^j(n)$, to create: X^j , where $(X^j)_{*,n} = \mathbf{x}^j(n_{washout} + n)$,
- $\mathbf{x}^j(n-1)$, to create: \tilde{X}^j , where $(\tilde{X}^j)_{*,n} = \mathbf{x}^j(n_{washout} + n - 1)$,
- $\mathbf{p}^j(n)$, to create: P^j , where $(P^j)_{*,n} = \mathbf{p}^j(n_{washout} + n)$
- $W^*\mathbf{x}^j(n) + W^{in}\mathbf{p}^j(n)$, to create: W_{target}^j , where $(W_{target}^j)_{*,n} = W^*\mathbf{x}^j(n_{washout} + n) + W^{in}\mathbf{p}^j(n_{washout} + n)$.

In the above, $(\cdot)_{*,n}$ is used to denote the n -th column of a matrix.

Computing Conceptors

The washout and learning period is repeated for all patterns $j = 1, 2, \dots, P$, during which the state correlation matrices R^j are collected. Furthermore, the aperture α^j is set for each pattern j . The conceptor matrix C^j is computed with

$$C^j = R^j(R^j + (\alpha^j)^{-2}I)^{-1}. \quad (33)$$

The aperture is chosen heuristically, but there exist optimization techniques, see Section 3.8 in [5]. These techniques will not be employed in this report.

Computing W and W^{out}

The collected matrices X^j , \tilde{X}^j , P^j , and W_{target}^j are concatenated, which gives $X = [X^1|X^2|\dots|X^P]$, $\tilde{X} = [\tilde{X}^1|\tilde{X}^2|\dots|\tilde{X}^P]$, $P = [P^1|P^2|\dots|P^P]$, and $W_{target} = [W_{target}^1|W_{target}^2|\dots|W_{target}^P]$. Following Section 3.2, the patterns are stored and the output weights are computed via ridge regression as follows:

$$W = (\tilde{X}\tilde{X}^T + \varrho^W I)^{-1}\tilde{X}W_{target}^T, \quad W^{out} = (XX^T + \varrho^{W^{out}} I)^{-1}XP^T, \quad (34)$$

where ϱ^W and $\varrho^{W^{out}}$ are the regularization constants for W and W^{out} respectively, and I is the identity matrix. The regularization constants have a big influence on the output of the simulation. A small regularization constant may result in overfitting, whereas a large regularization constant may lead to underfitting. The regularization constants are essentially another two parameters in the simulation and they need to be chosen appropriately. Again, this is done heuristically.

Self-generating reservoir

After W , W^{out} , and C^j have been computed, the reservoir should be able to self-generate any pattern j by inserting C^j in the update equation, yielding

$$\begin{aligned} \mathbf{r}(n+1) &= (1 - \alpha_l)\mathbf{r}(n) + \alpha_l \tanh(W\mathbf{z}(n) + \mathbf{b}) \\ \mathbf{z}(n+1) &= C^j\mathbf{r}(n+1) \end{aligned} \quad (35)$$

and

$$\mathbf{y}(n) = W^{out}\mathbf{z}(n), \quad (36)$$

where \mathbf{y} is expected to behave like pattern j .

4.3 Training Diagonal Conceptors

The training scheme of diagonal conceptors was inspired by the training scheme of RFC, described in Section 3.15 of the conceptors report [5]. The main contrast between the two training schemes is that in the case of diagonal conceptors, the conception weights are initialized to randomly drawn values from the uniform distribution on $[0, 1]$ instead of 1. This random initialization is an important feature of the training scheme of diagonal conceptors and it requires more attention than can be given in this section. Therefore, the consequences of the random initialization will be discussed in more detail in Section 7.1.1.

The training commences with initializing the following:

- number of neurons in reservoir (N),
- input weights ($W^{in} \in \mathbb{R}^{N \times M}$): Random values drawn from the standard normal distribution and scaled appropriately,
- bias ($\mathbf{b} \in \mathbb{R}^N$): Random values drawn from the standard normal distribution and scaled appropriately,
- initial reservoir weights ($W^* \in \mathbb{R}^{N \times N}$): Sparse matrix with random values drawn from the standard normal distribution and scaled appropriately,
- leaking rate α_l ,
- initial diagonal conceptors ($D_0^j = \text{diag}(\mathbf{c}_0^j)$, where $\mathbf{c}_0^j \in [0, 1]^N$): Elements of \mathbf{c}_0^j are randomly drawn from the uniform distribution on the interval $[0, 1]$,

Washout

The reservoir is driven with pattern j for $n_{washout}$ time steps according to the following update equations:

$$\begin{aligned} \mathbf{r}^j(n+1) &= (1 - \alpha_l)\mathbf{r}^j(n) + \alpha_l \tanh(W^* \mathbf{z}^j(n) + W^{in} \mathbf{p}^j(n+1) + \mathbf{b}), \\ \mathbf{z}^j(n+1) &= D_0^j \mathbf{r}^j(n+1), \end{aligned} \quad (37)$$

where the initial state $\mathbf{z}^j(0)$ is, again, simply set to the all-zero vector.

Adaptation

After the washout period, the reservoir is continued to be driven for n_{adapt} time steps according to Equation 37, but now the states $\mathbf{z}^j(n)$ are collected in the state collection matrix

$$Z^j = \begin{bmatrix} \mathbf{z}^j(n_{washout} + 1) & \mathbf{z}^j(n_{washout} + 2) & \dots & \mathbf{z}^j(n_{washout} + n_{adapt}) \end{bmatrix} \in \mathbb{R}^{N \times n_{adapt}}.$$

As shown in Section 3.4, the conception weights c_i^j can be computed element-wise

$$c_i^j = \frac{\mathbb{E}[(z_i^j)^2]}{\mathbb{E}[(z_i^j)^2] + (\alpha^j)^{-2}}, \quad (38)$$

where z_i^j is the i -th element of \mathbf{z}^j and α^j is the aperture for pattern j . Note that in simulations, \mathbf{c}^j is computed in a vectorized manner, which gives

$$\mathbf{c}^j = \mathbb{E}[\mathbf{z}^j \circ \mathbf{z}^j] \oslash (\mathbb{E}[\mathbf{z}^j \circ \mathbf{z}^j] + (\alpha^j)^{-2}), \quad (39)$$

where \circ is the Hadamard product, or the element-wise multiplication operator, and \oslash is the element-wise division operator. Furthermore, $\mathbb{E}[\mathbf{z}^j \circ \mathbf{z}^j]$ is approximated by

$$\mathbb{E}[\mathbf{z}^j \circ \mathbf{z}^j] \approx \frac{1}{n_{adapt}} \sum_l^{n_{adapt}} (Z^j \circ Z^j)_{*,l}, \quad (40)$$

which states that Z^j is multiplied with itself element-wise and then the mean of each row is taken. $(Z^j \circ Z^j)_{*,l}$ denotes the l -th column of $Z^j \circ Z^j$.

Note that this period is called the *adaptation* period, which may be a bit misleading as the diagonal conceptors are not adapted, but rather they are learned. The main reason it is referred to as the adaptation period regardless is because this is the terminology in the RFC architecture in the conceptors report [5], where the conception weights are adapted rather than learned. The second reason is that it allows for the incorporation of diagonal autoconceptors, which are described in Section 3.4.3 and briefly simulated in Section 6.

Learning

The reservoir is continued to be driven for another n_{learn} time steps, but now with the newly computed diagonal conceptors $D^j = \text{diag}(\mathbf{c}^j)$ in the loop rather than D_0^j . This gives the new, slightly different update equations:

$$\begin{aligned} \mathbf{r}^j(n+1) &= (1 - \alpha_l)\mathbf{r}^j(n) + \alpha_l \tanh(W^* \mathbf{z}^j(n) + W^{in} \mathbf{p}^j(n+1) + \mathbf{b}), \\ \mathbf{z}^j(n+1) &= D^j \mathbf{r}(n+1). \end{aligned} \quad (41)$$

During the learning period, the goal is to compute:

- W , such that $W \mathbf{z}^j(n) \approx W^* \mathbf{z}^j(n) + W^{in} \mathbf{p}^j(n)$ for all n, j ,
- W^{out} , such that $W^{out} \mathbf{z}^j(n) \approx \mathbf{p}^j(n)$ for all n, j ,

so the following vectors are collected:

- $\mathbf{r}^j(n)$, to create: X^j , where $(X^j)_{*,n} = \mathbf{r}^j(n_0 + n)$,
- $\mathbf{z}^j(n-1)$, to create: \tilde{Z}^j , where $(\tilde{Z}^j)_{*,n} = \mathbf{z}^j(n_0 + n - 1)$,
- $\mathbf{p}^j(n)$, to create: P^j , where $(P^j)_{*,n} = \mathbf{p}^j(n_0 + n)$,
- $W^* \mathbf{z}^j(n) + W^{in} \mathbf{p}^j(n)$, to create: W_{target}^j , where $(W_{target}^j)_{*,n} = W^* \mathbf{z}^j(n_0 + n) + W^{in} \mathbf{p}^j(n_0 + n)$,

where $n_0 = n_{washout} + n_{adapt}$ and $(\cdot)_{*,n}$ denotes the n -th column of a matrix.

Computing W and W^{out}

The washout, adaptation, and learning period is repeated for all patterns $j = 1, 2, \dots, p$ and the collected matrices are concatenated. This gives $X = [X^1 | X^2 | \dots | X^p]$, $\tilde{Z} = [\tilde{Z}^1 | \tilde{Z}^2 | \dots | \tilde{Z}^p]$, $P = [P^1 | P^2 | \dots | P^p]$, and $W_{target} = [W_{target}^1 | W_{target}^2 | \dots | W_{target}^p]$. Using ridge regression, the new reservoir weights W and the output weights W^{out} are then computed as follows:

$$W = (\tilde{Z}\tilde{Z}^T + \varrho^W I)^{-1} \tilde{Z} W_{target}^T, \quad W^{out} = (X X^T + \varrho^{W^{out}} I)^{-1} X P^T, \quad (42)$$

where ρ^W and $\rho^{W^{out}}$ are the regularization constants for W and W^{out} respectively, and I is the identity matrix.

Self-generating reservoir

After computing W and W^{out} , the reservoir is expected to self-generate pattern j by letting the reservoir run according to the update equations

$$\begin{aligned}\mathbf{r}^j(n+1) &= (1 - \alpha_l)\mathbf{r}^j(n) + \alpha_l \tanh(W\mathbf{z}^j(n) + \mathbf{b}), \\ \mathbf{z}^j(n+1) &= D^j\mathbf{r}(n+1),\end{aligned}\tag{43}$$

where the output can be read by

$$\mathbf{y}(n+1) = W^{out}\mathbf{r}(n).\tag{44}$$

The output \mathbf{y} is expected to behave like pattern j .

4.4 Pattern Comparison

Let $P = [\mathbf{p}(1) \mathbf{p}(2) \dots \mathbf{p}(L)]^T \in \mathbb{R}^{L \times M}$ and $Y = [\mathbf{y}(1) \mathbf{y}(2) \dots \mathbf{y}(L')]^T \in \mathbb{R}^{L' \times M}$ denote the input pattern and the pattern that is self-generated by the reservoir, respectively. When comparing Y and P , the first step is always visual inspection. Especially when choosing the simulation parameters, it is useful to see how the self-generated pattern compares to the input pattern. The second step is quantifying the comparison. This is practical in case the patterns are high-dimensional or many patterns are compared simultaneously. In this report, the accuracy of the match between a self-generated pattern and an input pattern will be compared with the Normalized Root Mean Square Error (NRMSE).

Let $\mathbf{p}_{obs} \in \mathbb{R}^K$ be an *observed* pattern and let $\mathbf{p}_{target} \in \mathbb{R}^K$ be a *target* pattern. The NRMSE of \mathbf{p}_{obs} and \mathbf{p}_{target} is given by

$$NRMSE(\mathbf{p}_{obs}, \mathbf{p}_{target}) = \sqrt{\frac{\frac{1}{K} \sum_{n=1}^K (\mathbf{p}_{obs}(n) - \mathbf{p}_{target}(n))^2}{\text{Var}(\mathbf{p}_{target})}},\tag{45}$$

where $\text{Var}(\mathbf{p}_{target})$ is the variance of \mathbf{p}_{target} . The NRMSE is a normalized version of the Root Mean Square Error (RMSE):

$$RMSE(\mathbf{p}_{obs}, \mathbf{p}_{target}) = \sqrt{\frac{1}{K} \sum_{n=1}^K (\mathbf{p}_{obs}(n) - \mathbf{p}_{target}(n))^2}.\tag{46}$$

Note that $NRMSE = 1$ if $\mathbf{p}_{obs}(n)$ is equal to the mean of \mathbf{p}_{target} for all n . This indicates that a reasonable model should achieve a NRMSE of between 0 and 1. In this report, a simulation has achieved good accuracy if the NRMSE is 0.1 or lower.

If the patterns are higher-dimensional, so $P_{obs} = [\mathbf{p}_{obs}(1) \mathbf{p}_{obs}(2) \dots \mathbf{p}_{obs}(K)]^T \in \mathbb{R}^{K \times M}$ and $P_{target} = [\mathbf{p}_{target}(1) \mathbf{p}_{target}(2) \dots \mathbf{p}_{target}(K)]^T \in \mathbb{R}^{K \times M}$, where $M > 1$, then the NRMSE is taken for each column. This yields an M -dimensional vector of NRMSEs. One could simply take the mean over all values. However, if the observed pattern has a small NRMSE in one column, but a large NRMSE in another column, this information would get lost in the mean. Therefore, when comparing higher-dimensional patterns, it is more informative to not only take the mean, but also show the smallest and largest

NRMSE values, and the standard deviation among all NRMSE values. In this case, a simulation has achieved good accuracy if the mean over all NRMSEs is 0.1 or lower and the standard deviation is 0.1 or lower.

The length of the input pattern P , denoted by L , is usually given and cannot be easily altered, unless the data is synthetically generated. The length of the self-generated pattern Y , denoted by L' , on the other hand, is usually regulated by the simulation and can easily be adjusted. Therefore, it is assumed that $L' \geq L$, where, in the simulations in this report, L' is simply set equal to L . In order to compare the input pattern and the self-generated pattern, the comparison length, denoted by K , must be chosen appropriately.

There is a nontrivial problem to measuring the accuracy between an input pattern and a self-generated pattern. If the observed pattern is matched against the target pattern for a long time, i.e., K is large, small phase shifts can lead to large NRMSEs. For example, if the target pattern is a sine wave with period 1 and the observed pattern is a sine wave with period 1.001. Then, after some time, the two patterns will have zero correlation and the NRMSE will be 1. Consequently, the NRMSE is more reliable if the patterns are compared for a short time, i.e., K is small. However, in that case, one cannot assess whether the generating RNN is long-term stable. Therefore, in this report, the long-term stability is assessed by computing the NRMSE for the first K time steps and the last K time steps. If the NRMSE in both cases is below 0.1, the match between the input pattern and the self-generated pattern is deemed good and long-term stable.

Lastly, the autonomous run of the reservoir is usually started from a random initial state, so the observed pattern will most likely not be phase-aligned with the target pattern. The observed pattern \mathbf{p}_{obs} and the target pattern \mathbf{p}_{target} are phase-aligned by finding a shifting parameter $d \in \mathbb{N}$, with $d \geq 1$, such that

$$NRMSE\left(\mathbf{p}_{obs}[d : K], \mathbf{p}_{target}[1 : (K - d + 1)]\right) \quad (47)$$

is minimum, where $\mathbf{p}_{obs}[d : K]$ denotes the subset of \mathbf{p}_{obs} containing only the elements from time step d up to and including K , which is the same notation used for indexing in the R programming language. The upper bound of d is set larger than the period of the target pattern to ensure optimal phase alignment. After the optimal d has been found, the shifted patterns are stored as well as their NRMSE. Note that in this setup, the NRMSE is computed over $K - d + 1$ time steps. This phase alignment is only performed for 1-dimensional patterns in this report.

5 Simulations

As mentioned in the previous sections, conceptors were first introduced in the conceptor report [5]. In that report, several examples with different patterns are used to highlight the numerous qualities of conceptors. It is for that reason that the patterns used in this section are taken from those examples. This allows for a direct comparison between conceptors and diagonal conceptors. Therefore, this section comprises three examples. A simulation with conceptors and a simulation with diagonal conceptors was conducted for each example. The findings are reported in the first three parts of this section. In addition, a morphing simulation with conceptors as well as a morphing simulation with diagonal conceptors was conducted. This is reported in the last part of this section. Simulation findings will be discussed in Section 7.

The patterns used in the three examples increase in difficulty. The first example consists of four periodic patterns, two irrational-periodic sine waves and two 5-periodic random points. This example serves as a introductory demonstration, which also shows how similar patterns can be distinguished by both conceptors and diagonal conceptors. The second example comprises four patterns sampled from the well-known chaotic attractors, the Lorenz, Rössler, Mackey-Glass, and Hénon attractor. Chaotic attractors are intrinsically unpredictable, which makes them a fitting challenge for conceptors as well as diagonal conceptors. In this example, the change in dynamics due to the aperture is highlighted. The last example models a collection of human motions and it is the most difficult out of the three examples. It comprises 15 61-dimensional patterns of lengths between roughly 200 and 900 time steps.

All the simulations in this section were implemented in R and the source code can be found [here](#) [29].

5.1 Periodic Patterns

A function f is periodic if $f(t+k) = f(t)$ for some constant k . For discrete-time patterns this means that the pattern is either *integer-periodic* or *irrational-periodic*. If a pattern is integer-periodic, there exists a positive integer k such that $\mathbf{p}(n+k) = \mathbf{p}(n)$. If a pattern is irrational-periodic, the pattern is sampled with a sample interval $d \in \mathbb{Z}_{>0}$ from a continuous-time signal with periodicity $k \in \mathbb{R}_{>0}$ and the ratio d/k is irrational. An integer-periodic pattern with period k will occupy k points in state space, as it will simply re-visit those points after every cycle. An irrational-periodic pattern, on the other hand, will occupy a more complex set of points in state space. For example, a one-dimensional irrational-periodic pattern will occupy a set of points that can be characterized topologically as a one-dimensional cycle that is homeomorphic to the unit cycle in \mathbb{R}^2 [5].

Let \mathbf{p}^1 and \mathbf{p}^2 be sampled from a sine wave with period ≈ 8.83 and ≈ 9.83 , respectively. They are irrational-periodic. Let \mathbf{p}^3 be a random 5-periodic pattern and let \mathbf{p}^4 be a 5-periodic pattern that is a slight variation of \mathbf{p}^3 . The set of patterns is given by $\mathbf{P} = \{\mathbf{p}^1, \mathbf{p}^2, \mathbf{p}^3, \mathbf{p}^4\}$. The patterns are 1-dimensional and are initialized to a length of $L = 5000$ time steps. An interval of 30 time steps is shown in the left column of Figure 5.

For both the conceptors and diagonal conceptors simulation, the reservoir is initialized with the number of reservoir neurons $N = 100$, leaking rate $\alpha_l = 1$, and regularization constants $\varrho^W = 0.001$ and $\varrho^{W^{out}} = 0$. The reservoir weights W^* and the input weights W^{in} are both scaled by 1 and the bias \mathbf{b} is scaled by 0.2. The length of the washout period is $n_{washout} = 200$. In the case of conceptors, the apertures was set to $\alpha^j = 100$ and for diagonal conceptors it was set to $\alpha^j = 8$, for all j . The length of

the learning period for conceptors was set to $n_{learn} = L - n_{washout} = 4800$ time steps. For the diagonal conceptors, the length of the adaptation period was set to $n_{adapt} = 500$, and the length of the learning period was set to $n_{learn} = L - n_{washout} - n_{adapt} = 4400$ time steps. For comparability, the network size N is set equal to the the network size in the example of the conceptors report [5].

The goal of this example is to demonstrate that diagonal conceptors can distinguish similar patterns with just as little effort as conceptors.

5.1.1 Conceptors

Lets start with the conceptors simulation. The reservoir was driven with the four patterns and the conceptors were trained according to the training algorithm described in Section 4.2. The conceptor matrices were stored and the reservoir was autonomously run for $n_{learn} = 4800$ time steps with C^j , for all j , yielding the self-generated patterns, denoted by \mathbf{y}^j . The self-generated patterns were phase-aligned with the input patterns, as described in Section 4.4.

Both the self-generated patterns and the target patterns are plotted on a small interval, which can be seen in the first column of Figure 5, where the target pattern (driver) is depicted by the black line, the conceptors-generated pattern (\mathbf{y}) is depicted by the red dashed line, and the NRMSE is shown in the bottom left corner. In the second column of Figure 5, the neuron state of four randomly chosen neurons, in the same interval as the first column, is plotted. From the similarity between the neuron state and the input pattern as well as the periodicity of the neuron state, it can be inferred that the dynamics of the patterns have been captured by the reservoir.

To illustrate the difference in the dynamics of the reservoir between the integer- and irrational-periodic patterns, the \log_{10} of the singular values of R^j and the 10 largest singular values of R^j are plotted in the third and fourth columns of Figure 5, respectively. Let $R^j = U\Sigma U^T$ be the singular value decomposition of R^j . The singular values of R^j are given by the diagonal of the diagonal matrix Σ and the principal components of R^j are given by U . The third column of Figure 5 neatly shows that for the irrational-periodic patterns, the principal components of R^j span all of R^N , whereas the principal components of the 5-periodic patterns are only non-zero in 5 directions, as the reservoir periodically visits the same 5 states. The fourth column shows that a large contribution of the singular values is concentrated on a few principal directions for all four patterns. This means that the point cloud created by the reservoir states, whose characteristics are captured by R^j , has a negligible variance in most directions. The small singular values for the 5-periodic patterns are due to numerical errors.

Adjusting the simulation parameters may result in slightly better or worse performance. Specifically, the conceptors still achieve good accuracy, as defined in Section 4.4, when the parameters are individually modified in the following ranges:

- W^{in} scaling: $[0.6, 2]$,
- W^* scaling: $[0.8, 1]$,
- \mathbf{b} scaling: $[0.2, 0.6]$,
- α_t : $[0.9, 1]$,
- $\varrho^{W^{out}}$: $[0, 1000]$,

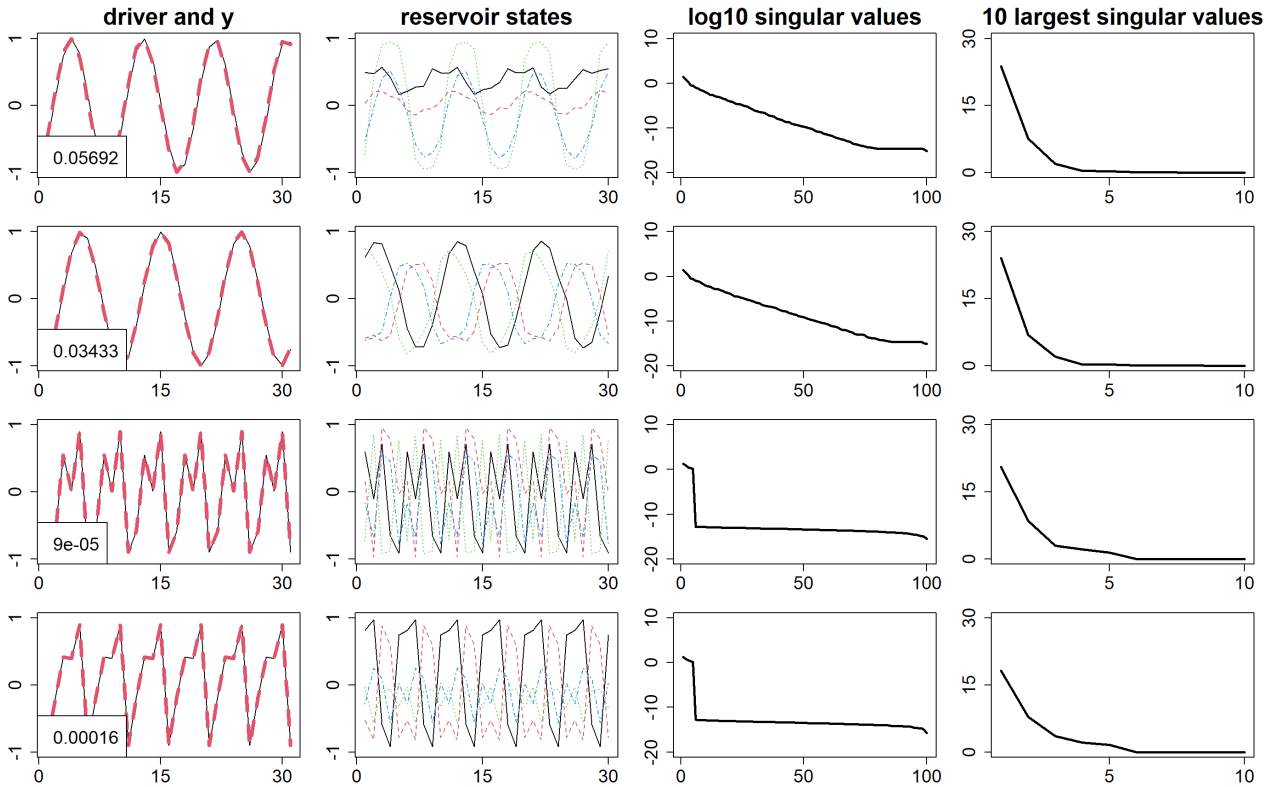


Figure 5: Row j depicts the results for pattern j . The left column shows the first 30 time steps of the driving patterns (black lines) and the conceptor-generated patterns (red dashed lines) and the NRMSE between the patterns over all time steps n . The second column shows the first 30 time steps of 4 randomly chosen neuron states. The third shows the \log_{10} of the singular values of R^j and the fourth column shows the 10 largest singular values of R^j .

- $\varrho^W: [0, 0.01]$,
- $\alpha^j: [10, 10000]$, for all j .

It must be noted that these ranges indicate that good accuracy is still achieved if a single parameter is adjusted to anywhere in its given range, while the other parameters remain unchanged. Regardless, there are combinations of parameters possible that lie outside of these ranges, e.g., the following parameters also yield good accuracy: W^{in} scaling = 0.5, W^* scaling = 1, \mathbf{b} scaling = 0.7, $\alpha_l = 1$, $\varrho^{W^{out}} = 0.000001$, $\varrho^W = 0.001$, $\alpha^j = 100$ for all j .

5.1.2 Diagonal Conceptors

Similar to the conceptors simulations, the reservoir is driven with the four periodic patterns. The diagonal conceptors are computed and the patterns are stored in the reservoir. Afterwards, the reservoir is able to self-generate the patterns with the diagonal conceptors. The self-generated patterns are phase-aligned the same as in the conceptors simulations and compared to the original patterns using NRMSE, see Section 4.4. The results are depicted in Figure 6 and are depicted in a similar way as Figure 5, which allows for direct comparison. Since the meaning of every column has already been discussed in Section 5.1.1, the interpretation of Figure 6 will not be described again.

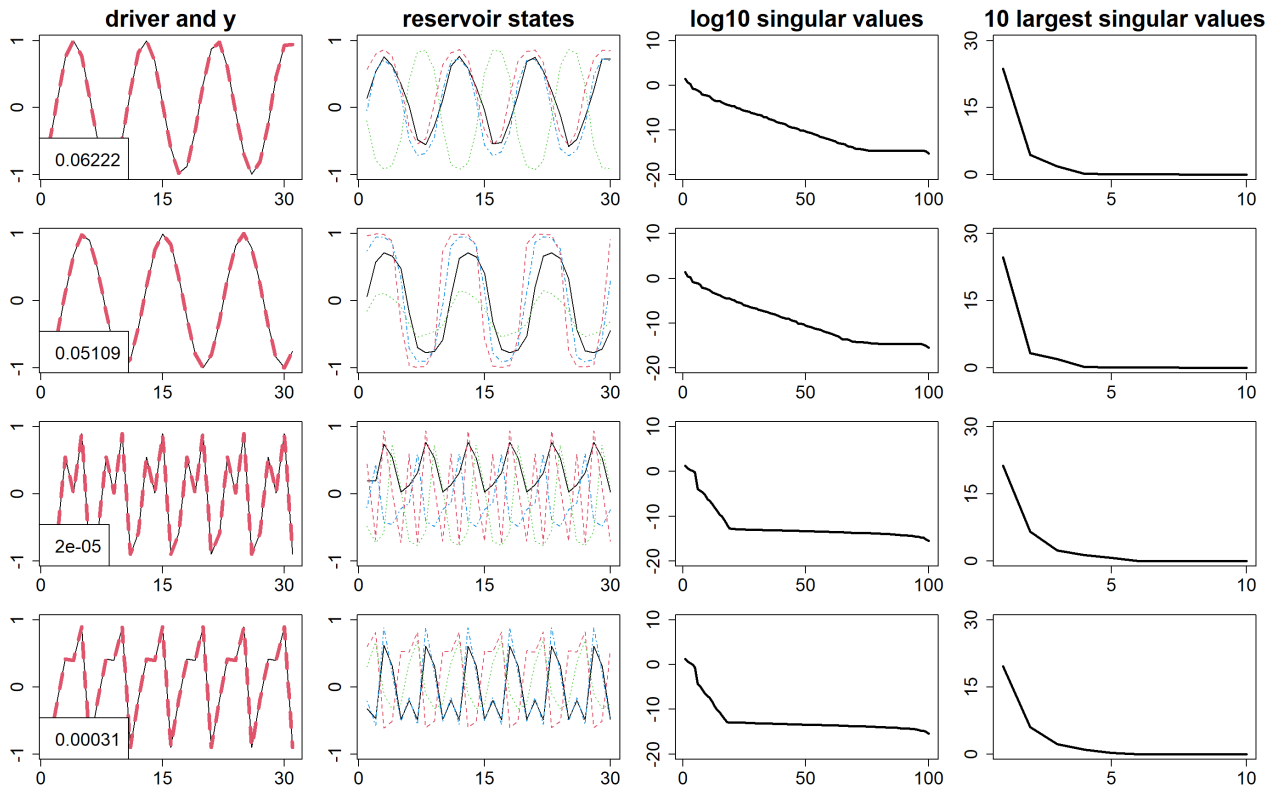


Figure 6: For explanation see Figure 5.

From the first column of Figure 6 it is evident that diagonal conceptors achieve equally good accuracy as conceptors. Additionally, the second, third, and fourth columns of Figure 6 are approximately identical to the second, third, and fourth columns of Figure 5.

Similar to the conceptors simulation, individual parameters were varied to see how it would affect the accuracy. The diagonal conceptors still achieve good accuracy, as defined in Section 4.4, when the parameters are individually modified in the following ranges:

- W^{in} scaling: $[0.8, 1.05]$ and $[1.15, 1.4]$,
- W^* scaling: $[0.9, 1]$,
- \mathbf{b} scaling: $[0.15, 0.35]$,
- α_l : $[0.85, 1]$,
- $\varrho^{W^{out}}$: $[0, 10000]$,
- ϱ^W : $(0, 0.01]$,
- α^j : $[6, 9]$, for all j .

For ϱ^W it should be noted that 0 is not included in the interval due to the singularity in computing the inverse matrix in ridge regression. In addition, it is remarkable that the input weights scaling is split up into two intervals. For values 1.05 – 1.15, the simulation yielded diagonal conceptors that were unstable.

Again, the above intervals only indicate that the simulation still yields good results if one parameter at a time is modified. Other parameter combinations that yield good accuracy are possible, e.g., W^{in} scaling = 0.6, W^* scaling = 1, \mathbf{b} scaling = 0.2, $\alpha_l = 1$, $\varrho^{W^{out}} = 0$, $\varrho^W = 0.000001$, $\alpha^j = 25$ for all j .

5.2 Chaotic Attractors

This example comprises four chaotic patterns sampled from the well-known Rössler, Lorenz, Mackey-Glass and Hénon attractors. The patterns are chaotic, hence fragile, and it will require careful fine-tuning of the dynamics of the reservoir in combination with the aperture of the conceptors and diagonal conceptors. For the defining equations of these well-known chaotic attractors, the reader is referred to Section 4.2 of the conceptors report [5]. The four chaotic patterns are shown in Figure 7.

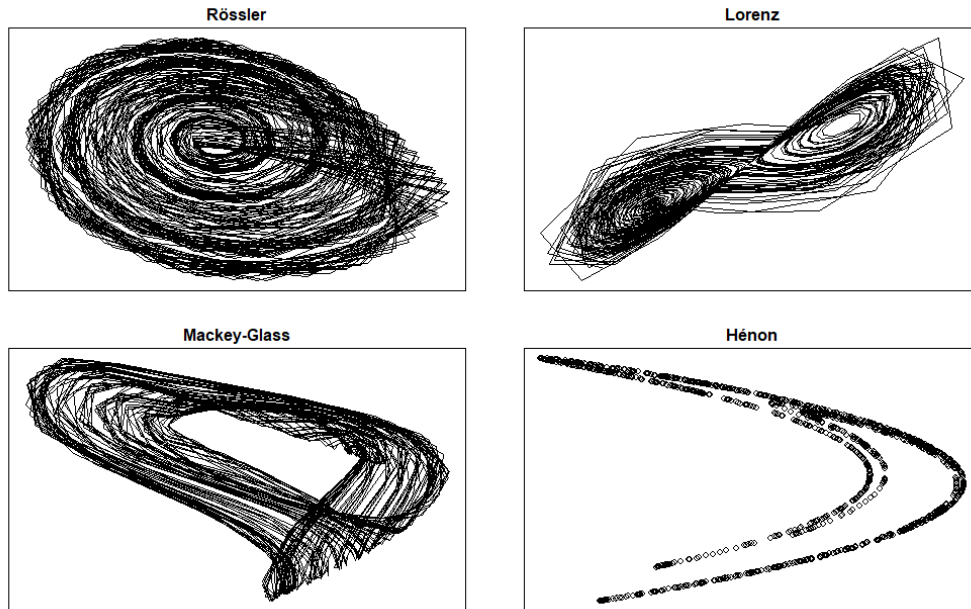


Figure 7: Four chaotic patterns, sampled from the Rössler (top left), Lorenz (top right), Mackey-Glass (bottom left) and Hénon (bottom right) attractors and plotted with time-delay embedding of 1500 time steps on the interval $[0, 1]^2$.

Let the set of patterns be denoted by $\mathbf{P} = \{\mathbf{p}^1, \mathbf{p}^2, \mathbf{p}^3, \mathbf{p}^4\}$, where $\mathbf{p}^j(n) \in \mathbb{R}^2$, for all j . Patterns 1, 2, 3, and 4 were sampled from the *Rössler*, *Lorenz*, *Mackey-Glass*, and *Hénon* attractors, respectively, and projected to the 2-dimensional plane via the same projections as in the conceptors report [5]. Each pattern has length $L = 1500$. The number of neurons in the reservoir was set to $N = 500$. The input weights W^{in} and the reservoir weights W^* were both scaled by 1.5, the bias \mathbf{b} was scaled by 1. The leaking rate was set to $\alpha_l = 1$. The length of the washout period was set to $n_{washout} = 100$, and for diagonal conceptors the length of the adaptation period was set to $n_{adapt} = 400$, yielding a learning period length of $n_{learn} = L - n_{washout} - n_{adapt} = 1000$ for diagonal conceptors and $n_{learn} = L - n_{washout} = 1400$ for conceptors. For diagonal conceptors, the apertures were set to $\alpha^1 = 10$, $\alpha^2 = 6$, $\alpha^3 = 9$, $\alpha^4 = 5$, where α^j corresponds to pattern j . For conceptors, the apertures were set to $\alpha^1 = 140$, $\alpha^2 = 50$, $\alpha^3 = 140$ and $\alpha^4 = 20$. The regularization constant for computing the output weights was set to $\varrho^{W^{out}} = 0$ for both conceptors and diagonal conceptors and the regularization constant for the

recomputing the reservoir weights was set to $\rho^W = 0.1$ for conceptors and $\rho^W = 0.01$ for diagonal conceptors.

The chaotic behavior of the patterns requires a carefully chosen aperture. Therefore, the goal of this example is to demonstrate how the role of the aperture for diagonal conceptors differs from the role of the aperture for conceptors. The control that the aperture offers for conceptors is slightly lost for diagonal conceptors. The reason is as follows. On the one hand, in the training scheme of conceptors, the conceptor matrices are computed independently of W and W^{out} . Therefore, after the reservoir is loaded, an adjustment in the aperture will not affect W nor W^{out} . Adjusting the aperture will only affect the conceptor matrices. On the other hand, in the training scheme of diagonal conceptors, the diagonal conceptor matrices are computed and inserted in the update equations *before* the patterns are stored. However, since an adjustment in the aperture changes the diagonal conceptor matrices, that adjustment in the aperture also affects the matrices W and W^{out} and thus the resulting dynamics of the reservoir. Consequently, changing the aperture associated with pattern j may affect the self-generated pattern associated with diagonal conceptor $i \neq j$. As a result, for diagonal conceptors there is a delicate interplay between the aperture and the regularization constants for W and W^{out} .

5.2.1 Conceptors

The conceptors were trained according to the procedure described in Section 4.2. The resulting conceptor matrices were then used to self-generate the stored patterns. The self-generated patterns were compared to the original patterns by visual inspection, because they are not expected to align perfectly with their target patterns. The target patterns are chaotic in nature, so if the self-generated patterns would be able to predict the chaotic patterns it would defy the unpredictability of the chaotic patterns. There are methods for comparing two chaotic patterns, for example by computing and comparing the Lyapunov spectrums, however, these were not employed in this thesis. Instead, the characteristics of the self-generated patterns and the target patterns were visually inspected to assess the accuracy of the match. In Figure 8 it can be seen that the characteristics are indeed captured. The black patterns are the target patterns and the red patterns are the conceptor-generated patterns. Furthermore, they are stable, meaning that after perturbation, the dynamics of the reservoir will return to the dynamics associated with the chaotic attractor.

In the periodic patterns simulation, the individual parameters were slightly varied to find in which ranges the self-generated patterns would still be satisfactory. However, in the case of chaotic attractors, this is a bit more complicated. The reason is that the aperture must be handled with care, so changing a parameter such as the bias scaling leads to a necessary adjustment in the apertures. For example, adjusting the input weights scaling from 1.5 to 1.4 yields unstable self-generated patterns, but if the apertures of patterns 3 and 4 are then also adjusted to $\alpha^2 = 35$ and $\alpha^3 = 65$, the self-generated patterns are stable again.

5.2.2 Diagonal Conceptors

The diagonal conceptors were trained according to the scheme described in Section 4.3. The resulting diagonal conceptors were then used to self-generate the patterns. Similar to conceptors, the self-generated patterns were compared to the target patterns by visual inspection, and both the target patterns and the self-generated patterns are shown in Figure 9. The diagonal conceptor-generated patterns resemble the target patterns, just as well as in the conceptors simulation. Furthermore, they also yield stable reservoir dynamics like conceptors.




Figure 8: Conceptor-generated patterns (red) plotted next to their target patterns (black) in a time-delay embedding of 1500 time steps on the interval $[0, 1]^2$.

Similar to the parameters in the conceptors simulation, the parameters in the diagonal conceptors simulation could not simply be adjusted individually. Moreover, the aperture adjustment for the chaotic patterns was more sensitive than for the periodic patterns, i.e., a slight modification could yield unstable results. The main reason is that adjusting the aperture for j may affect the outcome for pattern $i \neq j$, as mentioned in the introduction of this example. For example, adjusting the aperture of pattern 1 from $\alpha^1 = 10$ to $\alpha^1 = 11$ resulted in instability for pattern 2, whereas pattern 1, 3 and 4 remained stable. For this reason, it was slightly more difficult to find stable solutions. Once a stable solution has been found for some of the patterns, one cannot simply adjust the aperture of the unstable patterns, because it could disrupt the stability of the stable patterns. In contrast to conceptors, one must start by adjusting the apertures to overall acceptable solutions and then gradually fine-tune until satisfaction. Unfortunately, in the current training procedure, the diagonal conceptors cannot be adjusted one by one. Furthermore, the stability of the solutions is not only dependent on the aperture, but also on the regularization constants ϱ^W and $\varrho^{W^{out}}$. There is a subtle balance between ϱ^W and the aperture. Intuitively, if the diagonal conceptor-generated patterns seem to be chaotic, either the aperture must be decreased or the regularization constant must be increased.

However, regardless of the increased difficulty of finding the appropriate parameters, there were still many different parameter configurations that yielded stable solutions. Similar to the periodic patterns simulations, some of the parameters could be individually varied while maintaining good accuracy, as described in Section 4.4, in the following ranges:

- W^{in} scaling: $[1.5, 1.7]$,
- W^* scaling: $[1.4, 1.5]$,
- \mathbf{b} scaling: $[0.8, 1]$,
- α_j : $[0.9, 1]$,




Figure 9: Diagonal conceptor-generated patterns (red) plotted next to their target patterns (black) for a period of $L = 1500$ time steps.

- $\varrho^{W^{out}} : [0, 1]$.

First, notice that the ranges are smaller compared to the periodic patterns simulations, which made it initially more difficult to find parameters that yield stable solutions. Second, note that the ranges for ϱ^W and α^j are omitted. This is due to the fact that even a slight adjustment in either ϱ^W or α^j would cause the self-generated patterns to be unstable. Lastly, similar to the periodic patterns simulations, other combinations of parameters also yielded good solutions, e.g., W^{in} scaling = 1.1, W^* scaling = 1.2, \mathbf{b} scaling = 0.8, $\alpha_l = 1$, $\varrho^{W^{out}} = 0.001$, $\varrho^W = 0.01$, $\alpha^1 = 30$, $\alpha^2 = 30$, $\alpha^3 = 30$, $\alpha^4 = 20$.

5.3 Human Motion

The patterns used in the simulations in this section are taken from the supplementary code of the paper *Using Conceptors to Manage Long-Term Memory For Temporal Patterns* by H. Jaeger [19]. In this paper, it is shown how conceptors can be used to manage long-term memory for temporal patterns by the use of a few examples. The human motion example is a real-world example demonstrating that conceptors are capable of storing a larger number of patterns of different lengths and timescales. Therefore, this example is ideal for demonstrating that diagonal conceptors are also capable of that. Even though the patterns in this example were taken from the paper, it must be noted that the simulations were also performed with a different set of human motions, which yielded similar results. These simulations can be found at the aforementioned GitHub repository [29].

The patterns are derived from real human motion, recorded by a motion capture lab consisting of 12 infrared camera capable of recording 120 Hz with images of 4 megapixel resolution. The data is taken from the open source Carnegie Mellon University Motion Capture Database (MoCap) [30]. A test subject wears a black jumpsuit and has 41 markers taped to the suit. The cameras pick up the markers and the different images are then used to triangulate the positions of the markers to output 3D data.

The data can be used in two ways. Either the data is given in the form of a .c3d file, which contains the marker positions, or the data is given by a pair of .asf and .amc files, where the .asf file describes the skeleton and its joints and the .amc file describes the movement data. Since the data is in a specific format, it is not so straightforward to extract the human motion from the data. Fortunately, code has been written in different programming languages to accommodate this. The patterns in the supplementary code of [19] are already processed for the most part, but in the case of data taken directly from the MoCap database, an R package called *mocap* (<https://github.com/gsimchoni/mocap>) was used for parsing the .asf and .amc files. For a more detailed explanation about the structure of the data, the reader is referred to the MoCap website (<http://mocap.cs.cmu.edu/>). Since the data processing is rather involved, it will not be discussed here, but in the case of the patterns that are used in this section, the reader can read more about it in the appendix of [19].

Let $\mathbf{P} = \{\mathbf{p}^1, \mathbf{p}^2, \dots, \mathbf{p}^{15}\}$ be the set of patterns where pattern j at time step n is given by $\mathbf{p}^j(n) \in \mathbb{R}^{61}$, for all j . Patterns 1 until 15 represent the following human motions, respectively: three different types of boxing, doing a cartwheel, crawling, striding, getting down on the knees, getting seated, jogging, sitting, walking slowly, standing up, standing up from a stool, walking, waltzing. The patterns are diverse in their length L , which ranges between roughly 200 and 900 time steps. Some are transient (sitting down), some are periodic (walking, running), and some are irregular stochastic (boxing). The columns of each pattern are scaled such that they lie in the range $[-1, 1]$. In addition, the patterns were smoothed to remove some of the rough edges of the data. The first dimension of four patterns is shown in Figure 10.




Figure 10: The first of 61 dimensions is plotted for four different human motion patterns: boxing 1, standing up, cartwheel, and boxing 2.

The goal of this example is to demonstrate the ability of diagonal conceptors for a larger, more diverse set of patterns. The fact that this is a real-world example makes it a fitting challenge for diagonal conceptors.

5.3.1 Conceptors

The number of neurons in the reservoir was set to $N = 600$. The input weights W^{in} were scaled by 0.2, the bias \mathbf{b} was scaled by 0.8 and the reservoir weights W^* were scaled by 1. The leaking rate was set to $\alpha_l = 0.3$. The length of the washout period was set to $n_{washout} = 50$, yielding a learning period of length $n_{learn} = L - 50$, which was different for each patterns. The apertures were set to $\alpha^1 = 25$, $\alpha^2 = 55$, $\alpha^3 = 17$, $\alpha^4 = 49$, $\alpha^5 = 89$, $\alpha^6 = 49$, $\alpha^7 = 9$, $\alpha^8 = 65$, $\alpha^9 = 49$, $\alpha^{10} = 9$, $\alpha^{11} = 25$, $\alpha^{12} = 25$, $\alpha^{13} = 17$, $\alpha^{14} = 9$, and $\alpha^{15} = 7$, where each α^j corresponds to the aperture for pattern j . The apertures were found by grid search where the grid was set from 1 to 100 with intervals of

9, where the aperture for pattern 15 was adjusted after the grid search. The regularization constants were set to $\varrho^{W^{out}} = 0$ and $\varrho^W = 0.001$.

The patterns were stored in the reservoir and the concepton matrices were computed. With the concepton matrices, the reservoir was able to self-generate the patterns from a given starting reservoir state. This is different from the periodic patterns and chaotic attractors simulations, where the reservoir could start from any reservoir state. The concepton-generated patterns would sometimes show unpredictable behavior if the reservoir would start from a random initial reservoir state, which is expected from, e.g., the transient patterns, as they must start from the correct starting state. For example, going from a sitting position to a standing position requires that the starting point is the sitting position. Therefore, the starting state of the learning period was saved and used as the starting state of the reservoir for the self-generation of the patterns. The reservoir was run for a period 4 times longer than the target pattern and the first dimensions of the outputted patterns are shown in Figure 11.




Figure 11: The first of 61 dimensions is plotted for the concepton-generated patterns (red) as well as the target pattern (black), where the concepton-generated pattern is run for a period that is 4 times longer than its target pattern. The human motions associated with the four different patterns are: boxing 1, standing up, cartwheel, and boxing 2.

For each target pattern and concepton-generated pattern, the NRMSE is computed along each dimension for the overlapping period, i.e., the first n_{learn} time steps. This yields a vector of 61 NRMSEs per pattern, of which the minimum, maximum, mean and standard deviation are shown in Table 1. The table shows that the mean NRMSE of each pattern, except for pattern 15, is below 0.1, which is desired. Pattern 15 corresponds to the waltz human motion, which is the longest pattern out of the 15 patterns with a length of 900 time steps. As mentioned in Section 4.1, the rule of thumb for initializing N is to set it equal to the length of the longest pattern. However, in this case, $N = 600$ and the longest pattern has length 900, which is assumed to be the reason that the NRMSE for pattern 15 is slightly larger. Furthermore, the small standard deviation row indicates that the NRMSE stay roughly around the mean NRMSE across all 61 dimensions, i.e., the mean is representative of the overall NRMSE.

The admissible ranges of the parameters were less broad than the admissible ranges of the parameters in previous simulations. For example, changing the input weights scaling from 0.2 to 0.3 disrupts the conceptors associated with pattern 2 and 15. However, the conceptors associated with the disrupted patterns can often be recomputed with different apertures, without changing the conceptors of the other patterns. This way, if the balance is disrupted, it can be restored by adjusting the apertures accordingly.

Furthermore, the conceptors are moderately robust to unknown reservoir states, which is why it has no

Pattern	1	2	3	4	5	6	7	8	9
Min	0.058	0.061	0.030	0.021	0.041	0.039	0.012	0.009	0.070
Max	0.123	0.183	0.099	0.074	0.141	0.108	0.077	0.096	0.202
Mean	0.076	0.092	0.047	0.036	0.057	0.058	0.030	0.023	0.083
Std	0.017	0.025	0.016	0.012	0.023	0.018	0.017	0.016	0.030

10	11	12	13	14	15
0.031	0.041	0.023	0.008	0.058	0.123
0.237	0.107	0.117	0.161	0.289	0.466
0.060	0.056	0.039	0.025	0.084	0.234
0.039	0.017	0.022	0.023	0.042	0.084

Table 1: The NRMSE of the target patterns and the conceptor-generated patterns. The NRMSE is computed per dimension, which yields a vector of 61 NRMSEs. The Mean row shows the average NRMSE. The Min row shows the smallest NRMSE. The Max row shows the largest NRMSE. The Std row shows the standard deviation of the NRMSEs vector.

trouble repeating a periodic pattern, as is seen Figure 11. This is because the starting reservoir state for a periodic human motion will be roughly the same as the ending reservoir state, e.g., a cartwheel starts from a standing position and ends in an ending position. However, in the case of transient patterns, the long-term behavior is undefined, which is why the self-generated pattern will show unpredictable behavior after the pattern has been regenerated. The long-term behavior of the self-generated patterns for both the periodic and transient patterns can be seen in Figure 11. Notice how the periodic human motion repeats after the target pattern has been regenerated, but the transient pattern stagnates.

5.3.2 Diagonal Conceptors

The number of neurons in the reservoir was set to $N = 1000$. The input weights W^{in} were scaled by 0.1, the bias \mathbf{b} was scaled by 0.1 and the reservoir weights W^* were scaled by 1. The leaking rate was set to $\alpha_l = 0.15$. The length of the washout period was set to $n_{washout} = 50$ and the length of the adaptation period n_{adapt} was set to the ceiling of one-third of the total length of the patterns, except for pattern 15, where a slightly larger period of $n_{adapt} = 400$ was required. This resulted in a learning period of length $n_{learn} = L - n_{adapt} - n_{washout}$, which was different for every pattern. The apertures were set to $\alpha^1 = 8$, $\alpha^2 = 8$, $\alpha^3 = 20$, $\alpha^4 = 8$, $\alpha^5 = 35$, $\alpha^6 = 35$, $\alpha^7 = 35$, $\alpha^8 = 8$, $\alpha^9 = 35$, $\alpha^{10} = 35$, $\alpha^{11} = 35$, $\alpha^{12} = 35$, $\alpha^{13} = 30$, $\alpha^{14} = 35$, $\alpha^{15} = 8$, where each α^j corresponds to the aperture for pattern j . The regularization constants were set to $\rho^{W^{out}} = 0.1$ and $\rho^W = 0.05$.

After the patterns were stored in the reservoir and the diagonal conceptors were computed, the reservoir was run freely with diagonal conceptors in the loop, where, again, the starting reservoir state was given. The diagonal conceptor-generated patterns are run for a period of length $4 * n_{learn}$, similar to the conceptors simulation, and the first dimension of a few of the diagonal conceptor-generated patterns are shown in Figure 12. The black patterns show the training patterns and the red patterns show the diagonal conceptor-generated patterns.

First of all, notice that the adaptation period shortens the training patterns. For the periodic patterns

and the chaotic attractors, this was not a problem, because the data was synthetically generated and an adaptation period could simply be generated to fit the needs of the simulation. However, for these real-world patterns this becomes a bit of an issue, because an adaptation period is required to train the diagonal conceptors. Ideally, the characteristics of a pattern are fully captured in the adaptation period and then repeated in the learning period. Unfortunately, this is sometimes not the case for real-world examples. One way to overcome this would be to reuse the pattern from the adaptation period in the learning period. This is discussed in more detail in Section 7.1.3.




Figure 12: The first of 61 dimensions is plotted for the diagonal conceptor-generated output (red) as well as the target pattern (black), where the conceptor-generated patterns are run for a period that is 4 times longer than the target pattern. The human motion associated with the four different patterns are: boxing 1, standing up, cartwheel, and boxing 2.

Second of all, from Figure 12 it seems that, after the target pattern has been matched, the diagonal conceptors-generated patterns are not fully behaving like they are expected to. The diagonal conceptors are trained on n_{learn} reservoir states, so starting from a given reservoir state the diagonal conceptors can constrain the dynamics reservoir as they were taught for n_{learn} time steps. However, after those n_{learn} time steps, the reservoir state is unknown to the diagonal conceptors. At that point, the diagonal conceptors should force the reservoir state back to the reservoir state it knows. In the conceptors simulation it was seen that the conceptors were more robust to unknown reservoir states. The diagonal conceptors are more vulnerable to unknown reservoir states than conceptors, as is seen by the fact that the diagonal conceptors-generated patterns show unpredictable behavior after the reservoir has autonomously run for n_{learn} time steps. In other words, diagonal conceptors are less powerful than conceptors to guide the reservoir back from unknown states to states that it knows.

Similar to the conceptors simulation, the patterns were compared for the first n_{learn} steps, corresponding to the black line in Figure 12, and the NRMSEs are shown in Table 2. The results are similar to the results for conceptors. The mean of the NRMSE of each pattern is smaller than 0.1, which is desired, and the standard deviation is also smaller than 0.1, indicating that the mean is representative of the performance over most dimensions. Nonetheless, compared to Table 1, there are a few more outliers, as can be seen in the Max row. Especially pattern 8, which has one dimension in which the NRMSE is incredibly large. However, upon closer inspection, this concerns an individual outlier, which only diverges slightly towards the end. Overall, the diagonal conceptors achieve good accuracy.

The admissible range of the parameters is similar to, but slightly more fragile than, the admissible range of the parameters of the conceptors simulation. For example, changing the input weights scaling from 0.1 to 0.2 disrupts the diagonal conceptors associated with patterns 2 and 15, similar to conceptors. Again, this disruption can then be countered by adjusting the apertures α^2 and α^{15} . How-

Pattern	1	2	3	4	5	6	7	8	9
Min	0.013	0.009	0.020	0.005	0.005	0.023	0.009	0.015	0.004
Max	0.058	0.064	0.254	0.037	0.136	0.473	0.364	0.534	0.051
Mean	0.025	0.020	0.060	0.011	0.013	0.090	0.026	0.061	0.010
Std	0.011	0.009	0.044	0.005	0.022	0.067	0.052	0.127	0.011

10	11	12	13	14	15
0.010	0.011	0.008	0.019	0.045	0.017
0.339	0.146	0.342	1.308	0.175	0.107
0.036	0.043	0.012	0.091	0.086	0.054
0.075	0.032	0.061	0.184	0.033	0.021

Table 2: The NRMSE of the target patterns and the diagonal conceptor-generated patterns. The NRMSE is computed per dimension, which yields a vector of 61 NRMSEs. The Mean row shows the average NRMSE. The Min row shows the smallest NRMSE. The Max row shows the largest NRMSE. The Std row shows the standard deviation of the NRMSEs vector.

ever, unlike in the case of conceptors, the aperture adjustment of the diagonal conceptors is a bit more delicate, as described in the chaotic attractors simulation. Therefore, adjusting apertures α^2 and α^{15} may lead to disruption in other diagonal conceptors as well. It is for this reason that it can be difficult to find the right parameters for diagonal conceptors when it involves a larger number of patterns with diverse lengths and higher dimensionality. Regardless, other parameter combinations were found that achieved good accuracy.

5.4 Morphing

In this section, the conceptors and diagonal conceptors that were computed in the previous examples were used to morph between different patterns. All examples of morphing in this section are first carried out with conceptors and then with diagonal conceptors to allow for direct comparison.

The first example (periodic patterns) is also demonstrated in the conceptors report and it shows three types of morphing: frequency, shape, and heterogeneous [5]. The first type morphs two sine waves with different periods, the second type morphs two 5-periodic patterns and the third type morphs a 5-periodic pattern into a sine wave. The second example (chaotic attractors) shows a morph between two chaotic patterns, specifically, the Rössler and Mackey-Glass attractors. The third example (human motions) demonstrates a morph between two human motions, boxing and a cartwheel.

As discussed in Sections 3.3.2 and 3.4.2, a morph between two or more patterns is governed by the mixture parameter $\mu \in \mathbb{R}$, which is typically between 0 and 1. In this example, a morph will be only between two patterns, however this is not necessary in general. Let the patterns be denoted by j_1 and j_2 with their associated conceptors C^{j_1} and C^{j_2} , or diagonal conceptors D^{j_1} and D^{j_2} . Pattern j_1 can be linearly morphed into pattern j_2 by conceptor matrices C^{j_1} and C^{j_2} according to

$$\begin{aligned} \mathbf{r}(n+1) &= \tanh(W\mathbf{z}(n) + \mathbf{b}), \\ \mathbf{z}(n+1) &= ((1-\mu)C^{j_1} + \mu C^{j_2})\mathbf{r}(n+1), \end{aligned} \tag{48}$$

or by diagonal conceptor matrices D^{j_1} and D^{j_2} according to

$$\begin{aligned}\mathbf{r}(n+1) &= \tanh(W\mathbf{z}(n) + \mathbf{b}), \\ \mathbf{z}(n+1) &= ((1-\mu)D^{j_1} + \mu D^{j_2})\mathbf{r}(n+1),\end{aligned}\tag{49}$$

where the leaking rate is omitted from the equations for notational ease. Note that $0 \leq \mu \leq 1$ indicates interpolation and $\mu < 0$ and $\mu > 1$ indicate extrapolation. The simulations in this section start by driving the reservoir with a constant $\mu = \mu_{min}$ for a period of time. Then the reservoir is continued to be driven while μ is gradually increased for a period of n_{morph} time steps until $\mu = \mu_{max}$.

5.4.1 Periodic Patterns

For morphing conceptors, the boundaries of the mixture parameter were set to $\mu_{min} = -2$ and $\mu_{max} = 3$ in the case of for the frequency morphing, and $\mu_{min} = 0$ and $\mu_{max} = 1$ for the other two types. For morphing diagonal conceptors, the boundaries were set to $\mu_{min} = 0$ and $\mu_{max} = 1$ for all three types. The reservoir was run according to Equation 48 for conceptors and Equation 49 for diagonal conceptors, where μ was initialized as $\mu = \mu_{min}$. The first $n_{washout} = 100$ steps were discarded. Then, the reservoir was continued to be driven for 50 steps. Afterwards, μ was linearly increased until μ_{max} for n_{morph} time steps. The morphing length was set to $n_{morph} = 200$ for the frequency morphing and $n_{morph} = 50$ for the others. Finally, the reservoir was run for another 50 steps with $\mu = \mu_{max}$. The resulting morphed patterns for conceptors and diagonal conceptors are shown in Figure 13 and 14, respectively.

The results for conceptors are similar to the results in the conceptors report, so for more details, the reader is referred to Section 3.7 of the conceptors report [5]. In short, in the case of the two sine waves, the conceptors are not only able to interpolate between $\mu = 0$ and $\mu = 1$, but also extrapolate for values $\mu < 0$ and $\mu > 1$. Remarkably, the information about the patterns that is extrapolated is the period of the sine waves. Therefore, for $\mu = -2$, the period of the sine wave is approximately 7.34 and for $\mu = 3$, the period is approximately 12.6, whereas the periods of the original patterns are approximately 8.83 and 9.83. This extrapolation feature does not translate to diagonal conceptors. The morphed pattern shows unpredictable behavior for $\mu_{min} < 0$ and $\mu_{max} > 1$. Furthermore, morphing the two 5-periodic patterns, as seen in the middle plot of Figures 13 and 14, shows that diagonal conceptors are just as good as conceptors to morph between similar shapes. Lastly, as seen in the bottom plot of Figures 13 and 14, diagonal conceptors can smoothly transition between dissimilar shapes.

Overall, diagonal conceptors morph between these periodic patterns almost as well as conceptors. However, in contrast to conceptors, they do not facilitate extrapolation.

5.4.2 Chaotic Attractors

The reservoir was run according to Equation 48 for conceptors and Equation 49 for diagonal conceptors, where the mixing parameter was initialized as $\mu = 0$ and step-wise increased by 0.125 until $\mu = 1$, yielding 9 values for μ . For every value of μ , the first $n_{washout} = 100$ steps were discarded and the reservoir was continued to be run for another $n_{morph} = 500$ time steps. For each value of μ , a time-delay embedding of 500 time steps is shown, resulting in 9 plots. The resulting plots for conceptors and diagonal conceptors are shown in Figures 15 and 16, respectively.

Both Figures 15 and 16 must be read from the top left to the bottom right plot. The value of μ corresponding to the plot is shown in the bottom left corner. Notice that $\mu = 0$ means that the reservoir

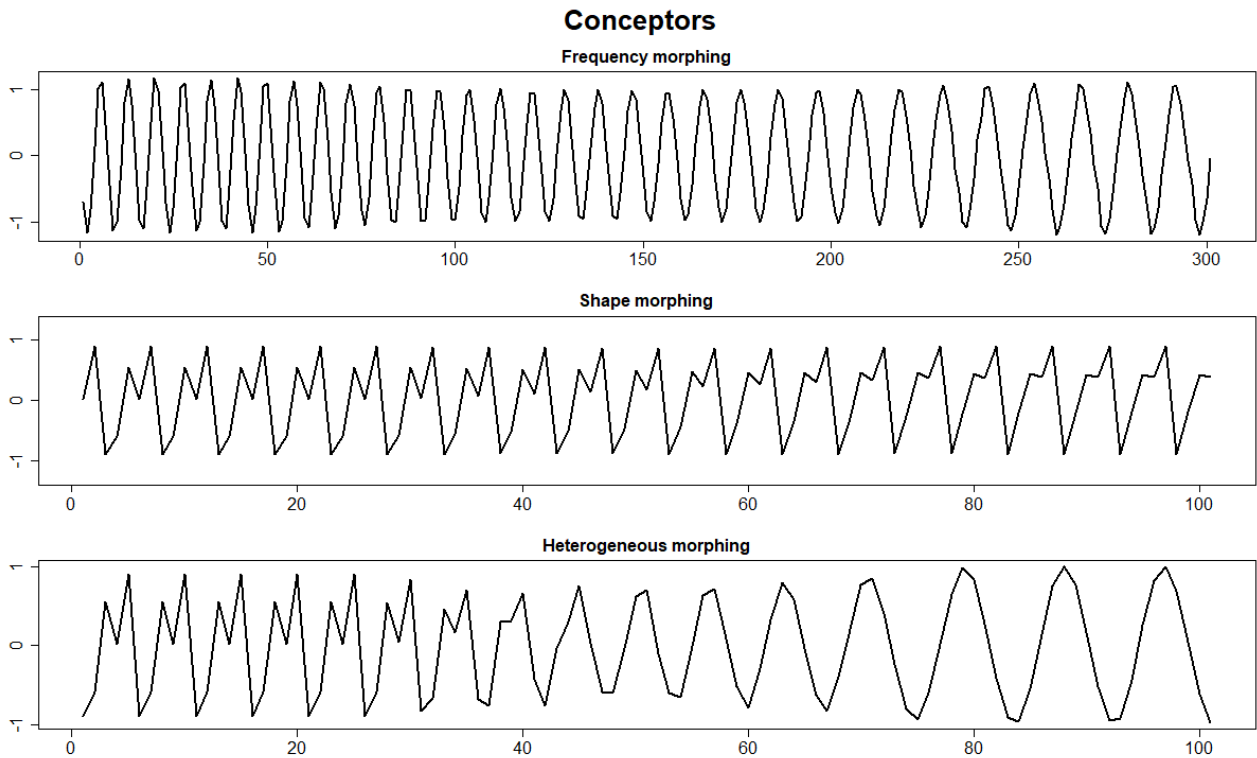


Figure 13: The upper plot shows a morph between two sine waves with periods of approximately 8.83 and 9.93, where $-2 \leq \mu \leq 3$. The middle plot shows a morph between two similar 5-periodic patterns, where $0 \leq \mu \leq 1$. The bottom plot shows a morph between a sine wave and a 5-periodic pattern, where $0 \leq \mu \leq 1$.




Figure 14: The upper plot shows a morph between two sine waves with periods of approximately 8.83 and 9.93, where $0 \leq \mu \leq 1$. The middle plot shows a morph between two similar 5-periodic patterns, where $0 \leq \mu \leq 1$. The bottom plot shows a morph between a sine wave and a 5-periodic pattern, where $0 \leq \mu \leq 1$.

is run with only the (diagonal) conceptor corresponding to the Rössler pattern and for $\mu = 1$ the reservoir is run with only the (diagonal) conceptor corresponding to the Mackey-Glass pattern. The plots show that interpolation, i.e., $0 < \mu < 1$, of conceptor matrices yields much more predictable behavior than the interpolation of diagonal conceptor matrices. The diagonal conceptors seem to output unpredictable behavior for $0 < \mu < 1$.

5.4.3 Human Motion

Again, the reservoir was run according to Equation 48 for conceptors and Equation 49 for diagonal conceptors. The (diagonal) conceptor matrices $(D^{j_1}) C^{j_1}$ and $(D^{j_2}) C^{j_2}$ correspond to a boxing pattern and a cartwheel pattern, respectively. The starting reservoir state was given, so there was no washout period required. The boundaries of the mixture parameter were given by $\mu_{min} = 0$ and $\mu_{max} = 1$. The reservoir was run with $\mu = 0$ for the length of the training pattern j_1 . Then the pattern was morphed for $n_{morph} = 150$ time steps. In reality, this would amount to a little more than 1 second, which is a fair amount of time needed between boxing jabs and doing a cartwheel. After the morphing period, the reservoir was run for another period with the length equal to that of the training pattern j_2 .

Remember that the diagonal conceptors were unstable to unknown reservoir states. Therefore, in the case of diagonal conceptors, during the morphing period, the reservoir state was slightly nudged to the correct starting state. This was implemented by adjusting the reservoir state $\mathbf{z}(n)$ as follows during the morphing period:

$$\mathbf{z}(n) = (1 - \mu)\mathbf{z}(n) + \mu\mathbf{z}^{j_2}(0), \quad (50)$$

where $\mathbf{z}^{j_2}(0)$ is the starting state of pattern j_2 , and μ is the same mixing parameter as for morphing. This nudging procedure was only executed for diagonal conceptors. For conceptors, this was not necessary, however, there was a washout period required after the morphing period.

The first three dimensions of the 61-dimensional output were plotted and are shown in Figures 17 and 18 for conceptors and diagonal conceptors, respectively. The left and right columns in both figures show patterns j_1 (boxing) and j_2 (cartwheel). The middle column in both figures shows how pattern j_1 is morphed to pattern j_2 . In the case of conceptors, it seems that the morph combines the characteristics of both patterns, where the beginning of the morph exhibits more characteristics of pattern j_1 and the end of the morph shows more characteristics of pattern j_2 , as is expected. In the case of diagonal conceptors, the morphing period shows less characteristics of both patterns. This is due to the fact that the reservoir dynamics are nudged towards the correct starting state of pattern j_2 , as described earlier.

At present, a more in-depth analysis of a morph with diagonal conceptors remains mostly unexplored. For example, it may be interesting to research how the geometric properties of the diagonal conceptors are carried through the morph. Therefore, any conclusions that are drawn are based on visual inspection rather than mathematical rigour.

Diagonal conceptors are less suitable for morphing patterns than conceptors. In the case of the periodic patterns, they performed well, as the morphing period did not seem to contain any sudden changes in reservoir dynamics. However, this was not the case for the chaotic attractors. The self-generated pattern seemed to exhibit unpredictable behavior during the morphing period for diagonal conceptors, whereas for conceptors the self-generated pattern showed less sudden changes in the reservoir dynamics. The same was true for the human motions, where the reservoir required to be nudged during the morphing period in the case of diagonal conceptors.

Conceptors

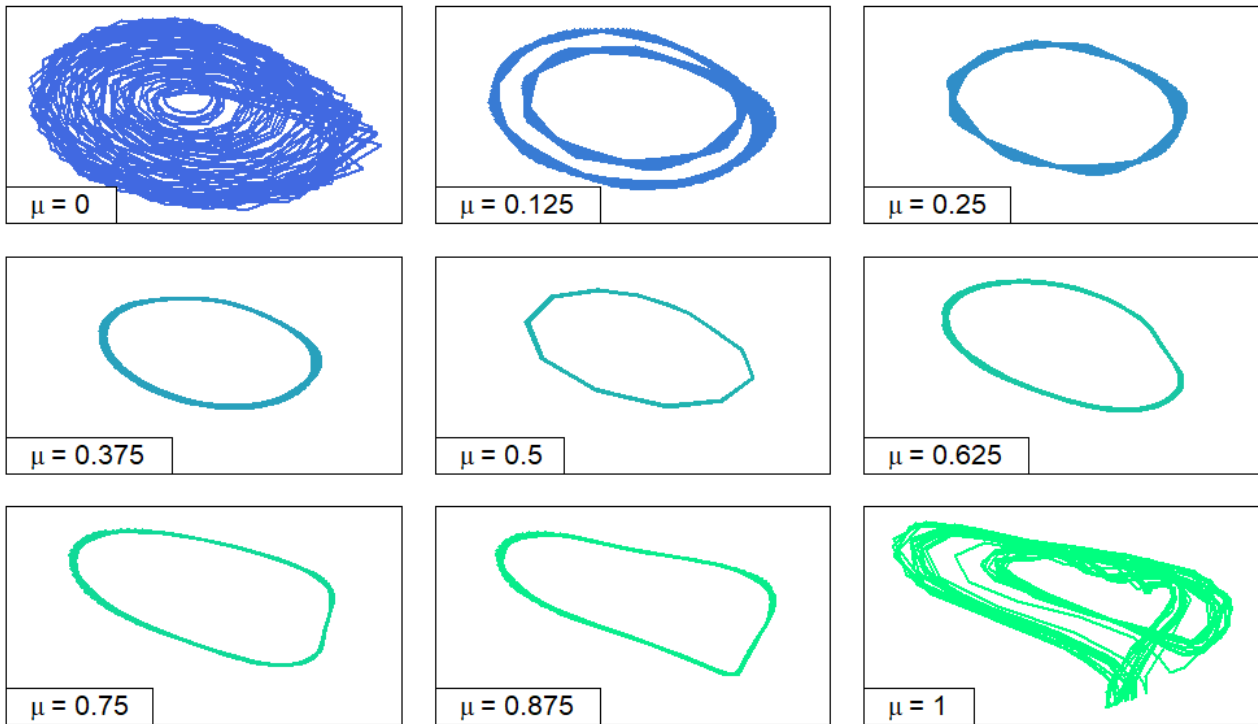


Figure 15: Time-delay embedding of 500 data points of a morph between the Rössler pattern and the Mackey-Glass pattern, where the mixing parameter μ is held still at 9 equidistant values from $\mu = 0$ to $\mu = 1$. The top left plot shows the Rössler pattern and the bottom right plot shows the Mackey-Glass pattern.

Diagonal conceptors

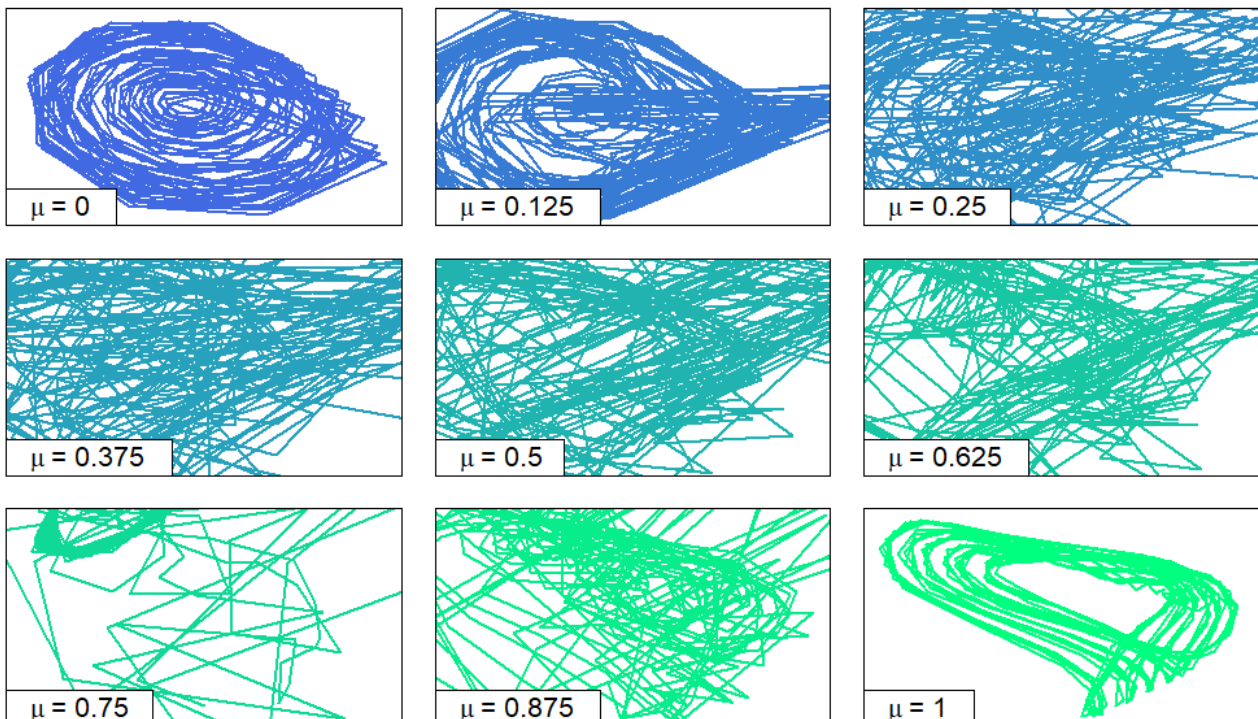


Figure 16: For explanation see Figure 15.

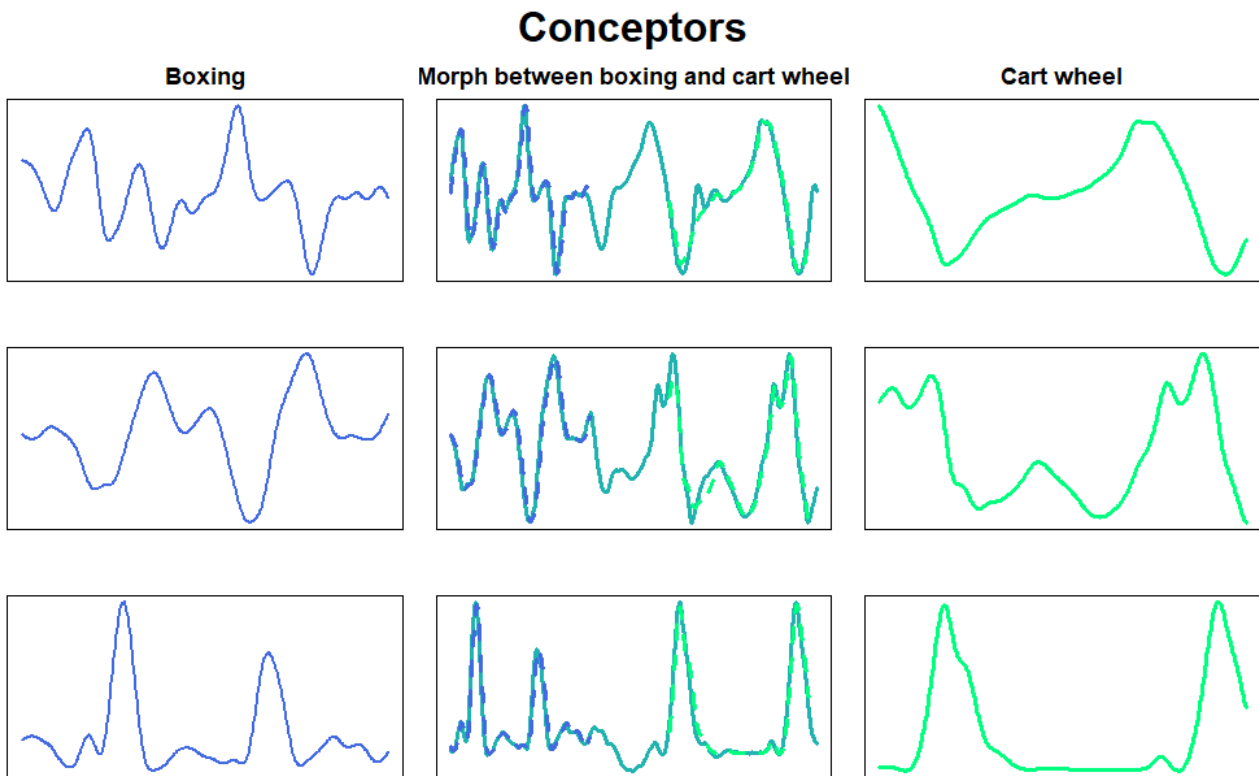


Figure 17: Morphing between two human motions, a boxing motion and a cartwheel motion. Since the patterns are 61-dimensional, only the first 3 dimensions are shown, each corresponding to a row in the figure. The left column shows the boxing patterns and the right column shows the cartwheel pattern. The middle column shows how the boxing pattern is linearly morphed into the cartwheel pattern and the dashed lines show the training patterns, which are added to compare the outputs.




Figure 18: For explanation see Figure 17.

6 Online adaptation

As mentioned in Section 3.4.3, diagonal conceptors can be trained by online adaptation, similar to autoconceptors. The section dedicated to autoconceptors in the conceptors report comprises various examples and insights into the workings of these autoconceptors [5]. Diagonal autoconceptors were briefly introduced in Section 3.4.3 for the sake of completeness, but they are not analysed to the same extent as autoconceptors in the conceptors report. Nevertheless, this section demonstrates how diagonal autoconceptors could be implemented through an easy example using the online adaptation rule from Equation 28.

The patterns in this example are the same as in Section 5.1, initialized with all the corresponding parameters. For notational simplicity, the superscript j is omitted. The simulation is conducted almost identically, except that during the adaptation period, the conception weights are adapted every time step according to the adaptation rule

$$c_i(n+1) = c_i(n) + \lambda_i((1 - c_i(n))z_i^2(n) - \alpha^{-2}c_i(n)), \quad (51)$$

where c_i is the conception weight of neuron i , λ_i is the learning rate associated with conception weight i , α is the aperture, and $z_i(n)$ is the i -th element of $\mathbf{z}(n)$ at time step n . The learning rate was set to $\lambda_i = 0.5$ for all i and the length of the adaptation period was set to $n_{adapt} = 3000$ to ensure convergence. Note that true convergence would only be achieved in the limit of $\lambda_i \rightarrow 0$. However, for the purpose of this simulation, true convergence is not necessary. The conception weights were collected over the whole adaptation period and plotted as a function of time, which is shown in Figure 19. It can be seen that the conception weights converge to either 0 or a value in the range $[0.5, 1]$, which confirms what was derived in Section 3.4.3.

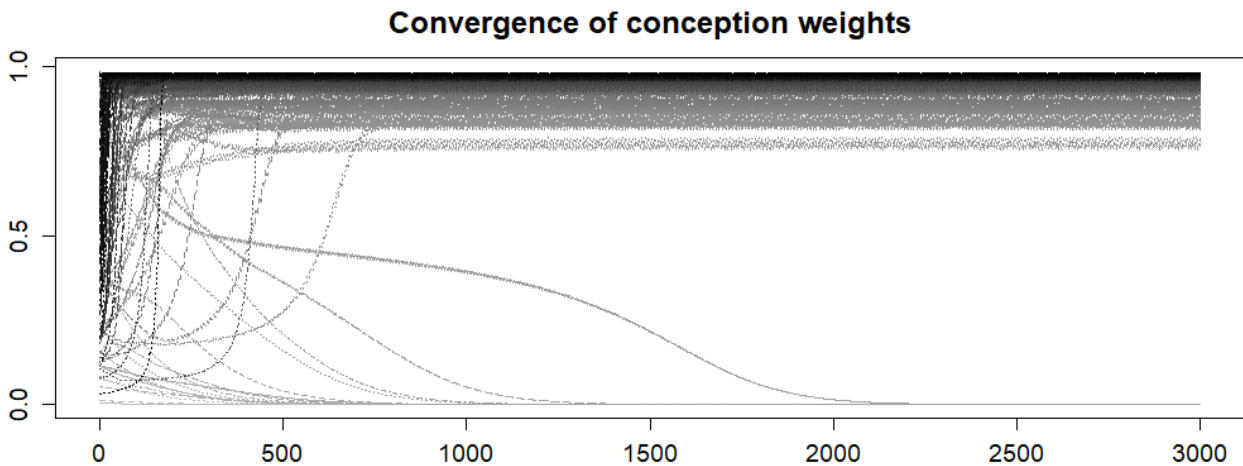


Figure 19: Convergence of the conception weights c_i over the span of $n = 3000$ time steps.

The simulation was also conducted without the adaptation rule, i.e., the diagonal conceptors were computed at the end of the adaptation period according to Equation 25, to allow for comparison. Computing the diagonal conceptors with an adaptation rule will be referred to as computing the diagonal conceptors *iteratively* and without adaptation rule will be referred to as *explicitly*. At this point it should be noted that both the iterative and explicit method yielded good accuracy for this

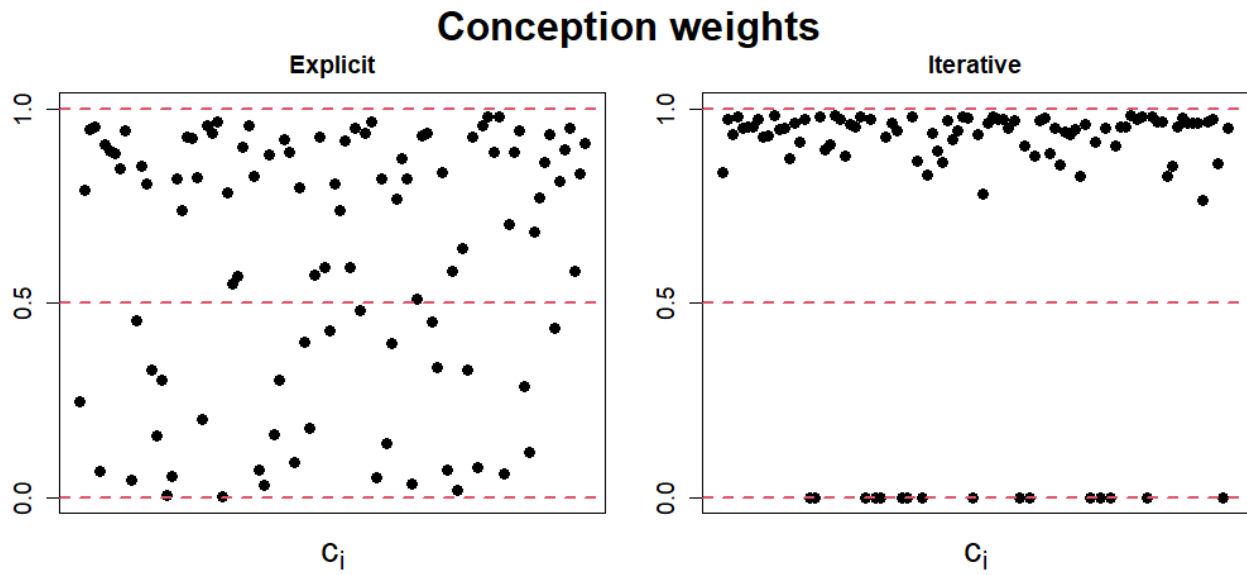


Figure 20: The conception weights after the adaptation period for both the iterative and explicit simulation.

example. The conception weights after the adaptation period for both the iterative and explicit method simulation are shown in Figure 20.

The explicitly computed conception weights are distributed over the range $[0, 1]$, where it appears slightly denser towards 1. Increasing the aperture will shift more conception weights towards 1. The iteratively computed conception weights are either 0 or in the range $[0.5, 1]$, as discussed in Section 3.4.3. Note that some conception weights are 0, which means that some neurons will be completely suppressed. For this reason, it can be hypothesized that the iterative method fails for more complex problems. However, this remains unexplored.

7 Discussion

Diagonal conceptors remain mostly unexplored at the present moment. However, this section dives a bit deeper into some of the features.

From the simulations in the previous section, it was clear that, for the most part, diagonal conceptors work just as well as conceptors. Consequently, diagonal conceptors offer an alternative, computationally efficient architecture for controlling a recurrent neural network. Aside from their advantages, diagonal conceptors also have their disadvantages. Both will be discussed in this section.

This section will attempt to analyse the dynamics of diagonal conceptors and how they relate to, and are different than, conceptors. Some of the steps in the training scheme will be analyzed, specifically, the randomly initialized diagonal conceptors. Furthermore, the advantages and disadvantages will be considered as well as some guidelines on how to manually train diagonal conceptors.

7.1 Training diagonal conceptors

This section collects several characteristics of diagonal conceptors, all related to the training algorithm. First, the use of randomly initialized diagonal conceptors is analyzed. Second, the influence of the aperture and the regularization constants $\rho^{W^{out}}$ and ρ^W is explored. Lastly, the necessary adaptation period of the training algorithm is discussed.

7.1.1 Initial Diagonal Conceptors

Let the reservoir comprise N neurons. For simplicity, it is assumed that the neurons are non-leaky, i.e., the leaking rate is set to $\alpha_l = 1$. The training of diagonal conceptors starts with initializing a number of parameters, specifically, and of interest, the initial diagonal conceptors $D_0^j = \text{diag}(\mathbf{c}_0^j)$, where $\mathbf{c}_0^j = [c_1^j \ c_2^j \ \dots \ c_N^j]^T$ is the conception vector and c_i^j are the conception weights, for all $i = 1, 2, \dots, N$. In contrast to the conceptors training scheme, randomly initialized diagonal conceptors are inserted in the update equations from the start, yielding

$$\begin{aligned} \mathbf{r}(n+1) &= \tanh(W^* \mathbf{z}(n) + W^{in} \mathbf{p}(n+1) + \mathbf{b}), \\ \mathbf{z}(n+1) &= D_0 \mathbf{r}(n+1). \end{aligned} \tag{52}$$

The reservoir is run according to Equation 52 for an initial washout period of $n_{washout}$ time steps and then for another n_{adapt} time steps. To understand the necessity of these initial diagonal conceptors in the loop, it is helpful to recap the idea behind conceptors.

If a reservoir is driven by pattern j , it creates a point cloud of reservoir states. The geometry of this point cloud is captured by the conceptor associated with pattern j . The conceptor can then be used to constrain the dynamics of the reservoir by projecting reservoir states onto the point cloud. A conceptor can distinguish similar patterns well, because their high degrees of freedom allows for minor adjustments in the projection onto the point cloud. However, diagonal conceptors have significantly fewer degrees of freedom and will have more trouble distinguishing similar patterns. This is where the initial diagonal conceptors come in. The initial diagonal conceptors create a random scaling of the reservoir states $\mathbf{r}(n)$. Intuitively, the initial diagonal conceptor shrinks the area in state space associated with pattern j differently for each pattern j . Therefore, the reservoir states, $\mathbf{z}(n) = D_0 \mathbf{r}(n)$, that are created are still characteristic for each pattern. Similar patterns, like the two sine waves in

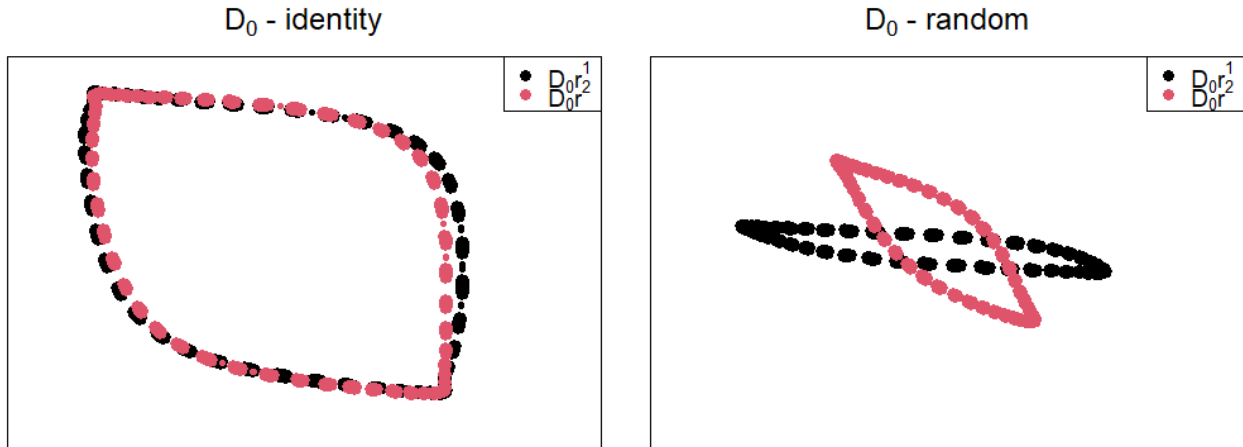


Figure 21: The neuron states during the adaptation period of two neurons are depicted. In the left plot, the initial diagonal conceptors were initialized as identity. In the right plot, the initial diagonal conceptors were initialized randomly. The black dots show states representing pattern 1. The red dots show states associated with pattern 2.

Section 5.1, are "pulled apart", making it easier for the diagonal conceptors to determine which state to follow.

Interestingly, this initial random scaling strategy can be extended to matrix conceptors, which provides a novel way of working with conceptors. This is worth exploring further as it allows for a different training scheme for conceptors, which can be compared to the current training scheme described in Section 4.2. This comparison was not explored in this thesis. However, training conceptors with the random scaling was briefly attempted for the periodic patterns that are described in Section 5.1 and the simulation yielded good results, as defined in Section 4.4. The simulation can be found in the aforementioned GitHub repository [29].

Returning to diagonal conceptors, the two sine waves from the periodic patterns example from Section 5.1 are taken as an example to show in more detail how similar patterns are "pulled apart". The sine wave with period ≈ 8.83 is called pattern 1 and the sine wave with period ≈ 9.83 is called pattern 2. The simulation is conducted twice. The first time with the diagonal conceptors initialized to the identity matrix and the second time with randomly initialized diagonal conceptors. Regarding terminology it should be noted that initializing the diagonal conceptors to the identity matrix or the conception vector to the all-ones vector or initializing no diagonal conceptors all mean the same thing. In the first simulation, $c_i = 1$ and in the second run c_i is drawn from the uniform distribution on $[0, 1]$, for all i . In both simulations, the reservoir states $\mathbf{z}^j = D^0 \mathbf{r}^j$, for $j = 1, 2$, are collected during the adaptation period of length $n_{adapt} = 500$. The results show the neuron state of two selected neurons for both pattern 1 and pattern 2 in both cases. Figure 21 depicts the findings.

It is evident from Figure 21 that the neuron states for the randomly initialized diagonal conceptors are more distinguishable than the neuron states for the identity initialized diagonal conceptors. In the left plot it is seen that the neuron states of the sine waves are very similar, hence difficult for the diagonal conceptors to tell apart. In the right plot, the scaled neuron states have been shrunk randomly and are now much easier to tell apart by the diagonal conceptors.

A few remarks can be made here. First, the randomness of the initialized diagonal conceptors is of a big influence on the training scheme. In general it is good practice to set a seed for computer simulations for reproducibility. However, in the case of diagonal conceptors this is vital, because underlying the training of diagonal conceptors lies the random initialization of the initial diagonal conceptors. Therefore, changing the seed of the simulation will most likely break the simulation. Second, there probably exists an optimal initialization. In this report, the conception weights were drawn from a uniform distribution on $[0, 1]$, but this does not have to be optimal. One could experiment with other distributions and other ranges. A few experiments were conducted where the conception weights were drawn from the uniform distribution on $[-1, 1]$ and the standard normal distribution. Both yielded unsatisfactory results. In the case of the former, this could be attributed to the fact that negative conception weights change the sign of negative neuron states, which could disrupt the reservoir dynamics. In the case of the latter, it is not guaranteed that a neuron state stays in the range $[-1, 1]$, which can also disrupt the reservoir dynamics. However, neither cases were analyzed thoroughly.

7.1.2 Aperture and Regularization

Besides the randomly initialized diagonal conceptors, the aperture α and regularization constant for recomputing the reservoir weights ϱ^W play a big role in training the diagonal conceptors. Regarding the regularization constant for computing the output weights, $\varrho^{W^{out}}$, it was found that in most cases it was of little influence on the output. Therefore, it will not be discussed here. To demonstrate how α and ϱ^W are of influence on the reservoir dynamics, the example from the previous section is continued. Four simulations were conducted where α and ϱ^W were changed. The results are shown in Figure 22 and require a bit of explaining.

The aperture was either set to $\alpha^j = 8$ or $\alpha^j = 30$, for $j = 1, 2$, and ϱ^W was either to $\varrho^W = 0.001$ or $\varrho^W = 100$. The reservoir was run with randomly initialized diagonal conceptors for $n_{adapt} = 500$ time steps, after which the diagonal conceptors D^j were computed. This computation depends heavily on α^j . Afterwards, the reservoir was continued to be driven for $n_{learn} = 500$ time steps, during which the states $\mathbf{z}^j = D^j \mathbf{r}^j$ were collected, for $j = 1, 2$. Then, the patterns were stored in the reservoir, which is regulated by ϱ^W . Finally, the reservoir was run for $n_{run} = 500$ time steps with the newly computed reservoir weights. During this period, the states $\mathbf{z}^j = D^j \mathbf{r}^j$ were again collected.

Similar to Figure 21, two neuron states were shown, yielding a two-dimensional plot for each simulation. The black color represents the neuron states associated with pattern 1 and the red color represent the neuron states associated with pattern 2. The dotted patterns are the states collected during the learning period and the filled-in areas represent the states collected during the self-generation period, where the states lie on the boundary of the area. The area within this boundary is filled in for aesthetic ease.

From Figure 22, a few observations can be made. First, increasing the aperture decreases the differences between the states. Consequently, for a large aperture, the point cloud depicted in the left left plot of Figure 21, where $D_0 = I$, is recovered. The states in the left column of Figure 22 seem to be a combination of the right and left plot of Figure 21. By adjusting the aperture it was found that the shape of the states, as seen in Figure 22, tends towards the left and right plot in Figure 21 for large and small apertures, respectively. Therefore, increasing the aperture appears to remove the advantage of the randomly initialized diagonal conceptors. This was also seen in the chosen apertures for the diagonal conceptors in the simulations in Section 5, which were consistently smaller than the apertures

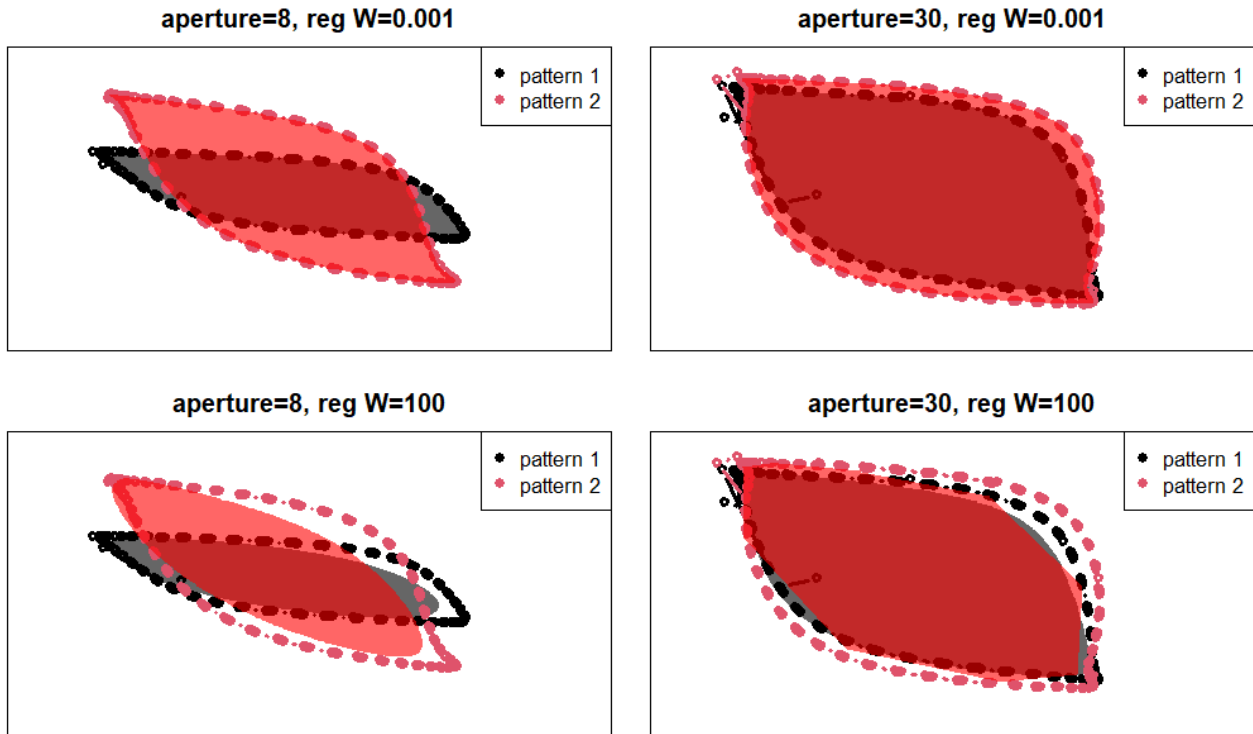


Figure 22: The neuron states during the reservoir learning period (dots) and the neuron states during the self-generation period (filled-in areas) of two neurons are shown. The black color represents the neuron states associated with pattern 1 and the red color represent the neuron states associated with pattern 2. For explanation see text.

for the conceptors. Second, the regularization constant ϱ^W behaves like a regularization constant is expected to behave. In ridge regression, the regularization constant mitigates between the bias and the variance, where an increase in the regularization constant leads to an increase in the bias, but a decrease in the variance. Therefore, a large ϱ^W will compress the shape of the states, thus decreasing the variance. What this means in practice is that if the aperture is increased it will be likely that ϱ^W should be decreased. This phenomenon can be seen most clearly in the right column of Figure 22. As the aperture is increased, the states will become more similar. However, if ϱ^W is small, it will capture the nuances that are distinct for each pattern. If ϱ^W is large, those nuances are suppressed, leading to indistinguishable state patterns. This is seen clearly in the bottom right plot of Figure 22, where the gray and red areas almost completely overlap.

7.1.3 Adaptation Period

Diagonal conceptors require a period during which the diagonal conceptors are computed, the adaptation period, as described in Section 4.3. The states in that period can not be used for storing the patterns in the reservoir, which means that all information of that period is not taken into account in the learning period. This information is lost, like in the washout period. Therefore, the learned pattern will be shorter than for conceptors. This was seen in the human motion simulation. The recalled patterns were shorter for diagonal conceptors compared to conceptors. In the case of the periodic patterns or the chaotic attractors, this was not a problem, because the patterns were generated synthetically.

As mentioned in the human motion simulations, in Section 5.3, a solution could be to first run the

adaptation period to compute the diagonal conceptors. Then, the pattern from the adaptation period would be reused in the learning period, yielding much more training data for storing the patterns. This has successfully been attempted for the periodic patterns, chaotic attractors, and human motions, in simulations that can be found in the aforementioned GitHub repository [29]. In all three cases, some of the parameters had to be adjusted to yield results with a good accuracy. For example, the aperture in the periodic patterns simulation needed to be decreased from 8 to 7, and in the chaotic attractors simulation, the aperture of pattern 2 needed to be increased from 6 to 7. After these adjustments, the periodic patterns and chaotic attractors performed equally well as in the simulations described in Sections 5.1 and 5.2. In the case of the human motions, the number of reservoir neurons had to be increased from $N = 1000$ to $N = 1500$ in order to yield results similar to the simulations in Section 5.3. From these simulations it can be concluded that reusing the pattern from the adaptation period in the learning period should be considered in future simulations of diagonal conceptors. In the case of the periodic patterns or the chaotic attractors it was not necessary, but in the case of the human motions it made a significant improvement, as the learned patterns were notable longer. Therefore, for a set of patterns such as the human motions, where patterns may be transient or irregular stochastic it is advised to apply this adjustment in the training scheme.

7.2 Stability

This section starts with the following observation. For a reservoir comprised of N neurons, a conceptor matrix C has $(N + 1)N/2$ degrees of freedom². A diagonal conceptor D has only N degrees of freedom. Therefore, as the number of reservoir neurons increase, the degrees of freedom of a conceptor matrix scales much faster than the degrees of freedom of a diagonal conceptor matrix. Because of this nonlinear scaling, one expects that the number of reservoir neurons should be significantly larger in the diagonal conceptors simulations compared to the conceptors simulations. However, the simulations in Section 5 showed that this was not necessary. In the case of the periodic patterns and the chaotic attractors, the number of reservoir neurons was set equal for conceptors and diagonal conceptors. Even in the case of the human motions, which comprised many more dimensions (61) compared to the periodic patterns (1) and the chaotic attractors (2), the number of reservoir neurons for diagonal conceptors was only $1000/600 = 1\frac{2}{3}$ times larger than for conceptors. The fact that the number of reservoir neurons for diagonal conceptors does not scale faster is mainly due to the different dynamics of diagonal conceptors imposed by the randomly initialized diagonal conceptors. Nevertheless, it must be noted that diagonal conceptors show slightly more unstable behavior than conceptors.

In the periodic patterns and chaotic attractors simulations, the diagonal conceptors were robust against perturbations. In the case of the chaotic attractors, this was especially evident in the morphing simulation, where the self-generated pattern was unpredictable during the morphing period, but afterwards it stabilized again to the desired pattern. Furthermore, self-generation with diagonal conceptors was usually initiated from the all-zeros vector, but other starting states were tried. It was found that any starting state in \mathbb{R}^N yielded stable patterns. This stable behavior was found for both the periodic patterns and the chaotic attractors. Regardless of the starting state, the correct pattern was always engaged in.

In addition, the parameters of the periodic patterns simulation were quite robust. Large variations in the parameter value still preserved stability. For the chaotic attractors simulation, the parameters were less robust, as a variation in the scaling of the weights could disrupt the stability of the self-generated

²A conceptor matrix is symmetric, so it has $N + (N - 1) + (N - 2) + \dots + 1 = N(N + 1)/2$ degrees of freedom.

patterns. Nevertheless, the parameters could still be varied to a certain extent, while stability of the self-generated patterns was retained. In the case of the human motion simulation, setting the parameters required a bit more patience. It was shown that the found parameters could not be varied much, as it would disrupt the stability of the most complex patterns. Overall, the stability of diagonal conceptors depends on the complexity of the task, but diagonal conceptors are more stable than one would initially expect.

8 Conclusion

Throughout this report, diagonal conceptors have been shown to be a practical alternative to conceptors. Surprisingly, what is gained in computation efficiency is not immediately lost in stability. Nevertheless, diagonal conceptors are more unstable than conceptors for more complex tasks.

In the simulation of the periodic patterns, in Section 5.1, the only deficit of diagonal conceptors was the ability to extrapolate the periods of two sine waves in the morphing. However, besides that, diagonal conceptors perform just as well as conceptors. The fact that diagonal conceptors can distinguish between similar sine wave patterns is mainly due to the randomly initialized diagonal conceptors. They are inserted in the update equations from the beginning to randomly scale the reservoir states. This training is different from conceptors, but proved to be vital for the performance of diagonal conceptors.

In the simulation of the chaotic attractors, in Section 5.2, diagonal conceptors show that they are also capable of capturing the dynamics of more complicated systems. Furthermore, they proved to be stable, as perturbations in the reservoir dynamics would quickly be resolved. This could be seen by the fact that they would yield stable self-generated patterns from any starting state in \mathbb{R}^N . Furthermore, the patterns could be morphed, but during the morphing period the pattern showed unpredictable behavior. Nevertheless, after the morphing period, the diagonal conceptors were able to stabilize the self-generated pattern to the desired pattern. Furthermore, in this simulation it was shown that the aperture parameter for diagonal conceptors does not offer the same amount of control as it does for conceptors. An adjustment of the aperture of a single pattern can disrupt the stability of other patterns. This is due to the fact that the diagonal conceptors are computed before the patterns are stored in the reservoir, which creates a delicate balance between the aperture and storing the patterns in the reservoir. This balance is mitigated between the aperture and the regularization constant for recomputing the reservoir weights. Conceptors, on the other hand, are computed independently of storing the patterns, hence the aperture offers much more control over the stability of the self-generated patterns.

In the simulation of the human motions, in Section 5.3, diagonal conceptors were able to recall the learned patterns. After the learned patterns had been recalled, the self-generated continuation departed from the kind of motion that the initial training data represented. Note that this is expected for transient patterns, but not for periodic patterns. The diagonal conceptors were not "strong" enough to guide unknown states back to a learned state. The adaptation period shortens the length of the training data, but an alternative approach was proposed in Section 7.1.3, where the pattern in the adaptation period is reused in the learning period.

9 Future Work

The research of diagonal conceptors is still in the early stages. Some features were already discussed in Section 7, but there is much potential for future work, for example:

- Optimizing the randomly initialized diagonal conceptors.
- Researching diagonal autoconceptors, similar to the autoconceptors that are discussed in Section 3.14 of the conceptors report [5].
- Improving the delicate balance between the aperture and the regularization constants for storing the pattern in the reservoir.
- Researching how and if the other features of conceptors translate to diagonal conceptors. They features include, but are not limited to: Boolean operations, neural memory management, recognizing patterns.

The last point is of the most importance, as the comprehensive list of capabilities of conceptors is long and a practical implementation of those features could open up further research in many fields.

Bibliography

- [1] M. Tegmark, *Life 3.0: Being Human in the Age of Artificial Intelligence*. Knopf Publishing Group, 2017.
- [2] N. Bostrom, *Superintelligence Paths, Dangers, Strategies*. Oxford, UK: Oxford University Press, 2014.
- [3] J. Schrittwieser, I. Antonoglou, T. Hubert, K. Simonyan, L. Sifre, S. Schmitt, A. Guez, E. Lockhart, D. Hassabis, T. Graepel, T. Lillicrap, and D. Silver, “Mastering Atari, Go, chess and shogi by planning with a learned model,” *Nature*, vol. 588, 12 2020.
- [4] F. Rosenblatt, “The perceptron: a probabilistic model for information storage and organization in the brain.,” *Psychological Review*, vol. 65 6, pp. 386–408, 1958.
- [5] H. Jaeger, “Controlling Recurrent Neural Networks by Conceptors,” *ArXiv*, vol. abs/1403.3369, 2014.
- [6] C. Kiefer, “Sample-level sound synthesis with recurrent neural networks and conceptors,” *PeerJ Comput. Sci.*, vol. 5, p. e205, 2019.
- [7] X. He and H. Jaeger, “Overcoming Catastrophic Interference by Conceptors,” *ArXiv*, vol. abs/1707.04853, 2017.
- [8] T. Liu, L. Ungar, and J. Sedoc, “Continual Learning for Sentence Representations Using Conceptors,” *ArXiv*, vol. abs/1904.09187, 2019.
- [9] S. Krause, O. Otto, and F. Stolzenburg, “Fast Classification Learning with Neural Networks and Conceptors for Speech Recognition and Car Driving Maneuvers,” *ArXiv*, vol. abs/2102.05588, 2021.
- [10] T. Liu, J. Sedoc, and L. H. Ungar, “Correcting the Common Discourse Bias in Linear Representation of Sentences using Conceptors,” *ArXiv*, vol. abs/1811.11002, 2018.
- [11] A. Strock, N. Rougier, and X. Hinaut, “Transfer between long-term and short-term memory using Conceptors,” *ArXiv*, vol. abs/2003.11640, 2020.
- [12] R. Gast, P. Faion, K. Standvoss, A. Suckro, B. Lewis, and G. Pipa, “Encoding and Decoding Dynamic Sensory Signals with Recurrent Neural Networks: An Application of Conceptors to Birdsongs,” *bioRxiv*, 2017.
- [13] F. Meyer zu Driehausen, R. Busche, J. Leugering, and G. Pipa, “Bistable Perception in Conceptor Networks,” in *ICANN*, 2019.
- [14] P. Hitzler, F. Bianchi, M. Ebrahimi, and M. K. Sarker, “Neural-symbolic integration and the Semantic Web,” *Semantic Web*, vol. 11, pp. 3–11, 2020.
- [15] B. Wagner and A. Garcez, “Neural-Symbolic Integration for Fairness in AI,” in *AAAI Spring Symposium: Combining Machine Learning with Knowledge Engineering*, 2021.
- [16] S. Tran, “Unsupervised Neural-Symbolic Integration,” *ArXiv*, vol. abs/1706.01991, 2017.

-
- [17] S. Bader and P. Hitzler, “Dimensions of Neural-symbolic Integration - A Structured Survey,” *ArXiv*, vol. abs/cs/0511042, 2005.
- [18] P. Hitzler and B. Hammer, *Perspectives of Neural-Symbolic Integration*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007.
- [19] H. Jaeger, “Using Conceptors to Manage Neural Long-Term Memories for Temporal Patterns,” *Journal of Machine Learning Research*, vol. 18, pp. 1–43, 2017.
- [20] D. Willshaw, O. Buneman, and H. Longuet-Higgins, “Non-Holographic Associative Memory,” *Nature*, vol. 222, pp. 960–962, 1969.
- [21] L. Cooper, “A possible organization of animal memory and learning,” in *Collective Properties of Physical Systems*, pp. 252–264, Elsevier, 1973.
- [22] J. Hopfield, “Neural networks and physical systems with emergent collective computational abilities.” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 79(8), pp. 2554–8, 1982.
- [23] S. Amari, “Learning Patterns and Pattern Sequences by Self-Organizing Nets of Threshold Elements,” *IEEE Transactions on Computers*, vol. C-21, pp. 1197–1206, 1972.
- [24] J. Kolen and J. Pollack, “Multiassociative Memory,” in *The Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*, 1991.
- [25] X. Hinaut and P. F. Dominey, “A three-layered model of primate prefrontal cortex encodes identity and abstract categorical structure of behavioral sequences,” *Journal of Physiology-Paris*, vol. 105, pp. 16–24, 2011.
- [26] H. Jaeger, M. Lukoševičius, D. Popovici, and U. Siewert, “Optimization and applications of echo state networks with leaky- integrator neurons,” *Neural Networks*, vol. 20, 4 2007.
- [27] M. Lukoševičius, “A Practical Guide to Applying Echo State Networks,” in *Neural Networks: Tricks of the Trade: Second Edition* (M. Grégoire and G. B. M. K.-R. Orr, eds.), pp. 659–686, Berlin, Heidelberg: Springer Berlin Heidelberg, 2012.
- [28] I. B. Yildiz, H. Jaeger, and S. Kiebel, “Re-visiting the echo state property,” *Neural networks : the official journal of the International Neural Network Society*, vol. 35, pp. 1–9, 2012.
- [29] J. de Jong, “Diagonal Conceptors,” 2021. <https://github.com/jorispdejong/DiagonalConceptors>.
- [30] Carnegie Mellon Graphics Lab, “CMU Graphics Lab Motion Capture Database.” Available at <http://mocap.cs.cmu.edu/>.