

Analyzing the Capacity of Distributed Vector Representations to Encode Spatial Information

1st Florian Mirus

Research, New Technologies, Innovations
BMW AG

Garching, Germany
florian.mirus@bmwgroup.com

2nd Terrence C. Stewart

Applied Brain Research Inc.
Waterloo, Ontario, Canada

terry.stewart@appliedbrainresearch.com

3rd Jörg Conradt

Dep. of Comp. Science and Technology
KTH Royal Institute of Technology

Stockholm, Sweden
conr@kth.se

Abstract—Vector Symbolic Architectures belong to a family of related cognitive modeling approaches that encode symbols and structures in high-dimensional vectors. Similar to human subjects, whose capacity to process and store information or concepts in short-term memory is subject to numerical restrictions, the capacity of information that can be encoded in such vector representations is limited and one way of modeling the numerical restrictions to cognition. In this paper, we analyze these limits regarding information capacity of distributed representations. We focus our analysis on simple superposition and more complex, structured representations involving convolutive powers to encode spatial information. In two experiments, we find upper bounds for the number of concepts that can effectively be stored in a single vector only depending on the dimensionality of the underlying vector space.

Index Terms—VSAs (Vector Symbolic Architectures), distributed representations, capacity analysis, cognitive modeling

I. INTRODUCTION

Understanding and building cognitive systems has seen extensive research over the last decades leading to the development of several cognitive architectures. A cognitive architecture is a “general proposal about the representation and processes that produce intelligent thought” [1]. On the one hand, these architectures are used to explain and better understand important aspects of human behavior and intelligence. On the other hand, they are also used to design computers and robots mimicking certain cognitive abilities of humans.

VSAs (Vector Symbolic Architectures) [2] refers to a family of related cognitive modeling approaches that represent symbols and structures by mapping them to (high-dimensional) vectors. Such vectors are one variant of distributed representations in the sense that information is captured over all dimensions of the vector instead of one single number, which allows to encode both, symbol-like and numerical structures in a similar and unified way. Additionally, the architectures’ algebraic operations allow manipulation and combination of represented entities into structured representations. There are several architectures such as MAP (Multiply-Add-Permutate) [3], BSCs (Binary Spatter Codes) [4] and HRRs (Holographic Reduced Representations) [5], which propose different compressed multiplication operations replacing the initially used tensor product [6] and resulting in vectors with the same

dimension as the input vectors. One advantage of this approach is that the number of dimensions remains fixed, independent of the number of entities combined through the architecture’s algebraic operations. Schlegel et al. [7] give an overview of eight different variants of VSAs and compare their properties and characteristics.

VSAs have been employed in a diverse variety of application domains, for instance, as one building block for implementing cognitive tasks such as RPMs (Raven’s Progressive Matrices) [8] in SNNs (Spiking Neural Networks) [9] for the Spaun (Semantic Pointer Architecture Unified Network) model [10]. Furthermore, VSAs have been used for encoding and manipulating concepts [11] as well as for human-scale knowledge representation of language vocabularies [12]. Kleyko et al. [13] used VSAs to imitate the concept learning capabilities of honey bees. In robotics [14], VSAs have been used to learn navigation policies for simple reactive behaviors to control a Braitenberg-vehicle robot [15]. In previous work, we proposed an automotive environment representation based on the SPA (Semantic Pointer Architecture), one particular VSA, and employed this representation to tasks such as context classification [16] and vehicle trajectory prediction [17]. For the latter [17], we used the convolutive power of vectors to encapsulate spatial positions of several vehicles in vectors of fixed length (cf. Fig. 1 and Sec. II). Komer et al. [18] propose a similar representation of continuous space using convolutive powers and analyze it from neural perspective.

However, given the mathematical properties of VSAs, there are systematical limitations to the amount of information that can be encoded in such a vector representation. These limitations are strongly connected to the chosen dimension of the underlying vector space and are a feature of such modeling architectures for being able to model limitations of cognitive functions of living beings, who are also not able to store unlimited amounts of information. Considering human subjects for instance, the capacity to process and store information or concepts in short-term memory as well as other cognitive tasks is subject to numerical restrictions [19]. Hence, numerical limitations of cognitive architectures like VSAs are one way of modeling the numerical restrictions to cognition observed in human subjects. In our context of interest, i.e., automated driving [17], however, we need to analyze these




Fig. 1. Visualization of the convolutive vector power encoding one particular driving scene in a 512-dimensional vector. The left plot depicts a scene from a real-world driving data set, while the middle and right plots visualize the similarity between the representation vector of that scene and auxiliary comparison vectors created from a sequence of discrete values as heat map for the target vehicle (middle) and other cars (right).

restrictions imposed by VSAs in general and the SPA in particular to provide upper borders regarding the amount of information that can be stored in our vector representation.

A. Contribution

In this paper, we analyze the limits regarding information capacity of distributed representations with the goal of finding bounds for, e.g., the number of concepts that can effectively be stored in a single vector before the accumulation of noise makes it impossible to retrieve the original individual vectors. Therefore, our contribution is a two-stage analysis: First, we analyze the amount of information that can effectively be stored in a single vector through superposition (i.e., addition) of several concept vectors. A similar but slightly different experiment has been conducted in [20]: the atomic vocabulary vectors, referred to as elemental vectors in [20], are sparse in the sense, that they mostly contain 0 elements, and the superposed vectors are normalized after adding them. Furthermore, [20] only compares the similarity between the superposition and the original vector with the similarity between the original vector and the most recently added random vector as baseline for the expected similarity between randomly chosen vectors. In contrast, we calculate the similarity between the superposition vector and n other random vectors for reference.

Secondly, we analyze the information capacity of vector representations involving the convolutive vector power for encapsulating spatial information. Given our scene representation proposed in [17] (cf. Fig. 1 and Sec. II), we are primarily interested in representing two-dimensional values in vectors, which is why we focus our analysis of the convolutive power encoding scheme on this case. In our analysis, we show that the information capacity is tightly linked to the dimension of the underlying vector space and furthermore, we give upper bounds for the capacity of superposition and convolutive power representations for three different vector dimensionalities.

II. MATERIALS AND METHODS

A. Convolutive vector-power

The SPA [9] is based on Plate’s HRRs [5], which is one special case of a Vector Symbolic Architecture [2]. Here,

atomic vectors are picked from the real-valued unit sphere, the dot product serves as a measure of similarity, which we denote by ϕ , and the algebraic operations are component-wise vector addition \oplus and circular convolution \otimes . In this work, we make use of the fact that for any two vectors \mathbf{v}, \mathbf{w} , we can write

$$\mathbf{v} \otimes \mathbf{w} = \text{IDFT}(\text{DFT}(\mathbf{v}) \odot \text{DFT}(\mathbf{w})), \quad (1)$$

where \odot denotes element-wise multiplication, DFT and IDFT denote the Discrete Fourier Transform and Inverse Discrete Fourier Transform respectively. Using Eq. (1), we define the *convolutive power* of a vector \mathbf{v} by an exponent $p \in \mathbb{R}$ as

$$\mathbf{v}^p := \Re\left(\text{IDFT}\left(\left(\text{DFT}_j(\mathbf{v})\right)_{j=0}^{D-1}\right)\right), \quad (2)$$

where \Re denotes the real part of a complex number. Furthermore, we call a vector \mathbf{u} *unitary*, if $\|\mathbf{v}\| = \|\mathbf{v} \otimes \mathbf{u}\|$ for any other \mathbf{v} (see [5, Sec. 3.6.3 and 3.6.5] for more details on the convolutive power and unitary vectors).

Finally, we consider any two vectors similar, if their similarity is significantly higher than what we would expect from two randomly chosen vectors. For growing dimension D , the cosine similarity follows approximately a normal distribution $\mathcal{N}_{\mu, \sigma}$, with $\mu = 0$ and $\sigma = \frac{1}{\sqrt{D}}$ [21]. Using the three-sigma-rule, we denote $\epsilon_{\text{weak}} = \frac{2}{\sqrt{D}}$ as *weak similarity threshold* and $\epsilon_{\text{strong}} = \frac{3}{\sqrt{D}}$ as *strong similarity threshold*.

B. Vector representation

We are primarily interested in representing two-dimensional values in vectors, which is why we focus our analysis of the convolutive power encoding scheme on this case. Hence, we encode two numerical values x, y , i.e., a two-dimensional entity, by generating two random, unitary vectors \mathbf{X}, \mathbf{Y} representing the corresponding units and applying Equation (2)

$$\mathbf{V} = \mathbf{X}^x \otimes \mathbf{Y}^y. \quad (3)$$

To encode a sequence (x_i, y_i) for $i = 1, \dots, n$ of two-dimensional numerical values all sharing the same units, we

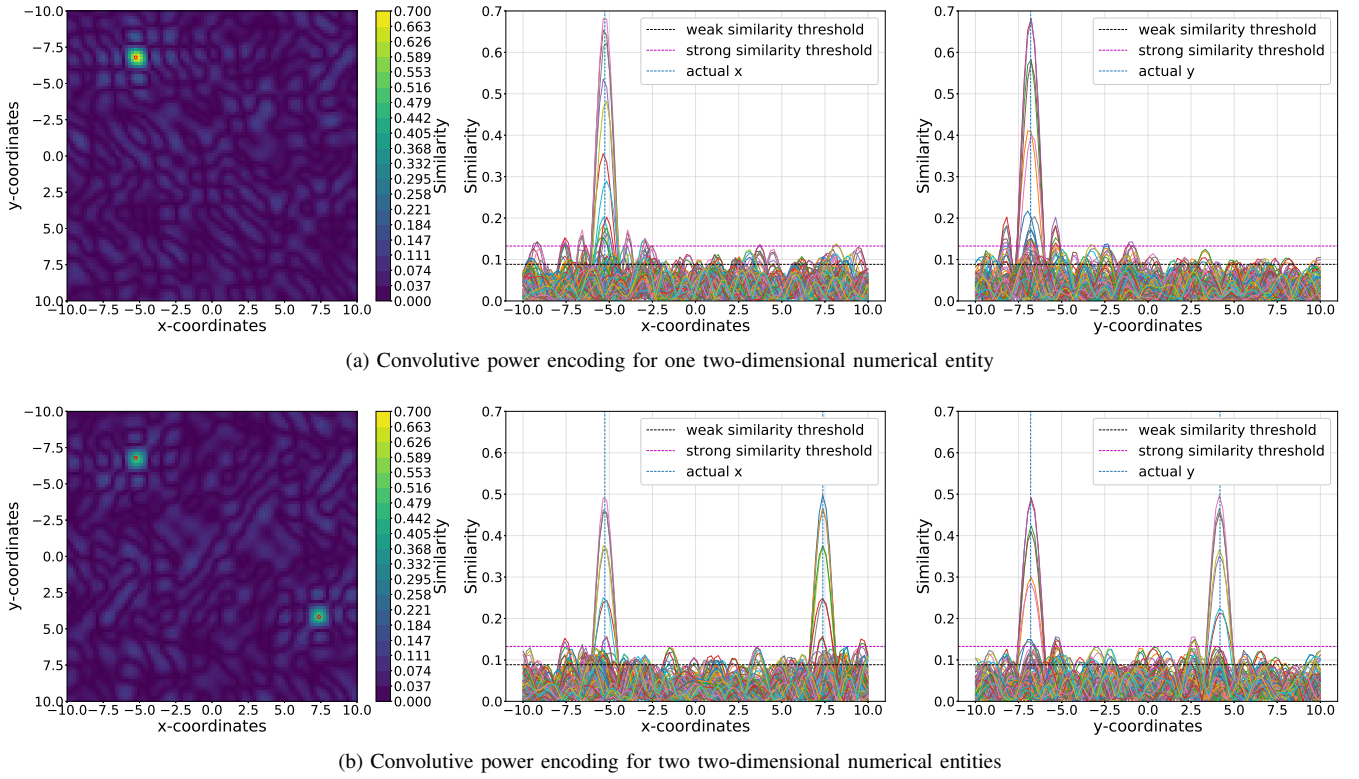


Fig. 2. Visualization of the convolutive power encoding scheme for 512-dimensional representation vectors depicting the similarity between the representation vector and auxiliary comparison vectors created from a sequence of discrete values. The left plot in both rows shows a two-dimensional grid of the similarities, while the middle and right plot show the individual entities respectively. The red circles in the left plot and the dashed blue lines in the middle and right plots indicate the actual encoded values.

simply sum up a their individual encoding vectors generated via Equation (3), which leads to

$$\mathbf{V} = \sum_{i=1}^n \mathbf{X}^{x_i} \otimes \mathbf{Y}^{y_i}. \quad (4)$$

Figure 2 visualizes vectors encoding one (cf. Fig. 2a and Equation (3)) and two numerical entities (cf. Fig. 2b and Equation (4)) given by two units within one vector. To generate the similarities shown in Fig. 2, we calculate the dot product between the vectors actually representing the encoded values and vectors $\tilde{\mathbf{V}}_i = \mathbf{X}^{\tilde{x}_i} \otimes \mathbf{Y}^{\tilde{y}_i}$ encoding a sequence of discrete sample values $(\tilde{x}_i, \tilde{y}_i)$ for $i = 1, \dots, m$. The left plot in each row of Fig. 2 depicts the similarities as heat map over a two-dimensional grid. The middle and right plots in Fig. 2 visualize the similarities of each unit, which is similar to plotting the heat map in three dimensions as ridges and slicing them through one of the ground axes. In both rows, we observe high similarity peaks way above both similarity thresholds at the actual encoded values and significantly lower similarity values everywhere else. However, comparing the similarities at the positions of the encoded values, we observe a drop of similarity values from roughly 0.7 to 0.5 when encoding two two-dimensional numerical values instead of only one. This gives a graphic explanation that there is a limit of how many values can effectively be encoded in such a representation before the actual values can not be properly

recovered anymore. Such limitations regarding the number of concepts that can be represented in one vector depending on its dimension are a recurrent theme in the field of VSAs and the subject of our analysis in the following section.

III. EXPERIMENTS

In this section, we conduct our analysis regarding the capacity of distributed representations for simple superposition in Sec. III-A and representations employing the convolutive power in Sec. III-B. The code for reproducing our analysis, results and figures can be found online ¹.

A. Superposition capacity

First, we evaluate the capacity of superposition, i.e., the addition operation. Superposition is used to store and combine several concept vectors \mathbf{v}_i for $i = 0, \dots, n$ in an unordered set

$$\mathbf{s} = \sum_{i=0}^n \mathbf{v}_i. \quad (5)$$

We can determine if a vector of interest \mathbf{w} belongs to that ordered set by calculating the similarity $\phi(\mathbf{s}, \mathbf{w})$ between the superposition vector and the vector of interest. For sufficiently high-dimensional vectors, the similarity $\phi(\mathbf{s}, \mathbf{w})$ will be close to 0 in case the vector \mathbf{w} is not part of the sum. However, the more vectors we add to the superposition vector \mathbf{s} , the more

¹https://github.com/fmirus/spa_capacity_analysis




Fig. 3. Visualization of the SPA’s superposition capacity for vector dimensions 256, 512 and 1024. The blue boxes indicate the similarity between the superposition vector and its summands, the orange boxes illustrate the similarity between the superposition vector and other randomly generated vectors. The dotted lines visualize the similarity threshold based on the vector dimensionality for reference.

noise accumulates in the representation and thus decreases the similarity between the superposition vector \mathbf{s} and its individual ingredients \mathbf{v}_i . In order to analyze how many vectors can be added together by superposition before individual vectors become irretrievable, we conducted the following experiment: assuming we want to add n vectors \mathbf{v}_i for $i = 1, \dots, n$ into a superposition vector \mathbf{s} as in Equation (5), we randomly generate a vocabulary of $2n$ vectors \mathbf{v}_i for $i = 1, \dots, 2n$ and sum up the first n members to create our superposition vector \mathbf{s} . Then we calculate the cosine similarity $\phi(\mathbf{s}, \mathbf{v}_i)$ between the superposition vector \mathbf{s} and every vector \mathbf{v}_i for $i = 1, \dots, 2n$ in the vocabulary.

Figure 3 shows the result of our experiment for 3 random vocabularies per superposition length containing vectors of dimension 256, 512 and 1024. The blue boxes in each figure illustrate the similarity between the superposition vector \mathbf{s} and each of the individual vectors \mathbf{v}_i for $i = 1, \dots, n$ it contains, i.e., the members of the unordered superposition set. The orange boxes depict the similarity between \mathbf{s} and the other vocabulary vectors \mathbf{v}_i for $i = n + 1, \dots, 2n$ it does not contain, i.e., the non-members. The dotted red and green lines indicate the SPA’s weak and strong similarity threshold depending on the dimension of the vector space. Considering the weak similarity threshold $\epsilon_{weak} = \frac{2}{\sqrt{D}}$, we observe that for a vector dimension of 256 the SPA allows roughly 50 items to be stored in a superposition vector. For higher vector dimensions 512 and 1024, the number of items that can be superposed increases to roughly 100 and 200 respectively. Considering the strong similarity threshold $\epsilon_{strong} = \frac{3}{\sqrt{D}}$, the upper borders for the

number of items being stored in a superposition vector are slightly more conservative with 25, 50 and 100 for vector space dimensions of 256, 512 and 1024 respectively. We also observe in our experiments that the similarity between the superposition vector and non-member random vectors is consistently below the weak similarity threshold ϵ_{weak} for the majority of the samples. However, once the similarity between the superposition vector and its members drops below either of the similarity thresholds for the majority of the samples, we can not distinguish between members and non-members with a sufficiently high probability. For 256 dimensional vectors for instance, we even observe that the members and non-members become nearly indistinguishably when adding more than 80 vectors. Thus, we have to choose rather conservative bounds for the number of items to be encoded in a superposition vector.

B. Capacity of structured representations involving convolutive powers

In the previous section, we have analyzed the SPA’s capacity regarding the number of items that can be stored in an unordered set using superposition. For encoding more complex information, such as driving situations [17], in a semantic vector substrate, we employ more complex representation than superposition of single items alone. In this section, we analyze the capacity of structured vector representations involving the convolutive vector-power (see Eq. (2)). Figure 2 illustrates that we can unbind positions by querying the representation vector with sample vectors encoding discrete position examples, but also that encoding several entities of the same type in one




Fig. 4. Capacity analysis for the superposition of vectors encoding spatial positions using the convolutive vector-power for varying vector dimensions.

position vector as in Fig. 2b yields lower similarities of the true positive positions compared to the encoding of only one item as in Fig. 2a. Hence, our capacity analysis has to cover the amount of objects encoded in one vector, but also the number of items per object class.

Therefore, we conduct the following experiment: assuming we want to encode n spatial entities, i.e., objects o_i with two-dimensional location information (x_i, y_i) for $i = 1, \dots, n$ as shown in Fig. 2 for $n = 1$ (Fig. 2a) and $n = 2$ (Fig. 2b), into a single representation vector \mathbf{s} , we generate a vocabulary of random vectors \mathbf{v}_i for $i = 1, \dots, n$ encoding object class labels and random unitary vectors \mathbf{X}, \mathbf{Y} to encode the units of the spatial information. In contrast to the experiment in III-A, where we simply summed up a certain number of random vectors, we are interested in a more specific analysis, since there are several possibilities to distribute the positional values (x_i, y_i) over the available object class vectors \mathbf{v}_i . For instance, for a total number of two superpositions, i.e., $n = 2$, there are two possibilities to generate our representation vector, namely

$$\mathbf{s}_1 = \mathbf{v}_1 \otimes \mathbf{X}^{x_1} \otimes \mathbf{Y}^{y_1} + \mathbf{v}_1 \otimes \mathbf{X}^{x_2} \otimes \mathbf{Y}^{y_2}, \quad (6)$$

$$\mathbf{s}_2 = \mathbf{v}_1 \otimes \mathbf{X}^{x_1} \otimes \mathbf{Y}^{y_1} + \mathbf{v}_2 \otimes \mathbf{X}^{x_2} \otimes \mathbf{Y}^{y_2}. \quad (7)$$

The vector \mathbf{s}_1 in Equation (6) encodes two objects of the same type, while the vector \mathbf{s}_2 encodes occurrences of two different object types at the given locations. As we are working with random vectors in this experiment, we can, without loss of generality, skip the vector encoding two objects of type represented by the vector \mathbf{v}_2 , which would yield a result equivalent to Equation (6). More generally, we are interested

in all sets

$$C_{m,j} = \left\{ 0 < k_1, \dots, k_m \leq n \mid m \leq n \text{ and } \sum_{i=1}^m k_i = n \right\} \quad (8)$$

of natural numbers k_i summing up to the total number of objects n ignoring permutations of the k_i . We index the sets with j , since there potentially exist several possibilities to decompose n into sums of m natural numbers.

In our experiments, for each number n of total objects to be encoded in the vector representation, we calculate all possible sets $C_{m,j}$ (ignoring permutations) and generate random position values (x_i, y_i) for $i = 1, \dots, n$ and a random vocabulary as described above. For each set $C_{m,j}$, we generate a representation vector

$$\mathbf{s}_{m,j} = \sum_{i=1}^m \sum_{l=1}^{k_i} \mathbf{v}_i \otimes \mathbf{X}^{x_i} \otimes \mathbf{Y}^{y_i} \quad (9)$$

as well as query vectors $\mathbf{P}_i = \mathbf{X}^{\tilde{x}_i} \otimes \mathbf{Y}^{\tilde{y}_i}$ encoding a sequence of discrete sample values $(\tilde{x}_i, \tilde{y}_i)$ for $i = 1, \dots, M$ evenly distributed over the length of the positional encoding grid. In other words, Equation (9) states, that each class label \mathbf{v}_i appears k_i times yielding a sum of n objects. We query the representation vector for the position of each class by binding it to the pseudo-inverse (see also [5], [9]) element $\bar{\mathbf{v}}_i$ for each class label vector, i.e.,

$$\mathbf{s}_{m,j} \otimes \bar{\mathbf{v}}_i \approx \sum_{l=1}^{k_i} \mathbf{X}^{x_i} \otimes \mathbf{Y}^{y_i}, \quad (10)$$

and calculate the similarity with the discretized position vectors \mathbf{P}_k to get

$$s_{i,k} = |\phi(\mathbf{s}_{m,j} \otimes \bar{\mathbf{v}}_i, \mathbf{P}_k)|. \quad (11)$$




Fig. 5. Capacity analysis for the superposition of vectors encoding spatial positions using the convolutive vector-power for varying vector dimensions. In contrast to Fig. 4, this figure illustrates the similarity for vectors containing spatial information for several objects of the same class.

For samples close to the originally encoded positions, i.e., $|x_i - \tilde{x}_i| < \varepsilon$ and $|y_i - \tilde{y}_i| < \varepsilon$ for a certain threshold ε (here, we use $\varepsilon = 0.4$), we label $s_{i,k}$ as positive similarity denoting a member of the representation vector. Otherwise, we consider the similarity $s_{i,k}$ at position $(\tilde{x}_i, \tilde{y}_i)$ not a member of the representation vector $s_{m,j}$.

Figure 4 shows the results of this capacity analysis regarding the total number of superposed objects within the representation vector for varying vector dimensions. In Sec. III-A, we have already analyzed the SPA’s capacity regarding the number of items that can be stored in an unordered set using superposition. Similar to Fig. 2, we observe that the similarity of the non-members is in the order of magnitude of the similarity thresholds while the similarity for the member position decreases with a growing number of spatial items encoded in the vector. However, for encoding more complex information, like automotive scenes [17] in a semantic vector substrate, we employ more complex representation than superposition of single items alone.

Therefore, Fig. 5 shows a different evaluation of the same data showing the number of addition operations per class on the x -axis. In contrast to Fig. 4, Fig. 5 illustrates the similarity for vectors containing spatial information for several objects of the same class. That is, Fig. 5 illustrates the similarity of vectors containing a specific number k of superpositions per class on its x -axis independent of the total number of superpositions. We observe, that not only the similarity of the members decreases with growing number of superpositions per class, but the similarity of the non-member increases beyond the weak similarity threshold. Similar to the simple

superposition capacity analysis, we consider the point in the plots where the member similarities fall below the strong similarity threshold the upper border for the maximal number of spatial objects per class to be encoded in this representational substrate. For instance, this upper bound for 256 dimensional vectors is 10 superpositions per class, which is roughly half of the upper bound for the number of simple superpositions. For higher-dimensional vectors (here 512 and 1024), these limits are beyond the evaluated number of superpositions per class. Therefore, using at least 512 dimensional vectors for automotive scene representation yields a sufficiently high information capacity. On the other hand, we expect upper bounds for the higher dimensions similar to 256-dimensional vectors, i.e., roughly half the number of superpositions as stated in Sec. III-A.

IV. CONCLUSION

In this paper, we analyzed the capacity of structured vector representations in VSAs based on simple superposition and superposition combined with the convolutive power encoding of spatial information. We provided a more detailed analysis of the superposition capacity compared to those available in the literature, e.g. in [20]. Furthermore, we evaluated the capacity of structured representations involving convolutive vector powers to encode spatial information, which, to the best of our knowledge, is the first of its kind. Thereby, we found upper bounds for the amount of information that can effectively be encoded in such representations depending on the dimension of the underlying vector space. These bounds have to be considered in future work regarding distributed rep-

representations, for instance, of automotive scenes like in [17] to evaluate if the amount of information to be encoded in the vector representation is compliant with these bounds. This will allow a conclusive assessment of the limits of structured vector representations in general, whereas our particular focus is on automotive context.

REFERENCES

- [1] P. Thagard, "Cognitive Architectures," in *The Cambridge Handbook of Cognitive Science*, K. Frankish and W. Ramsey, Eds. Cambridge University Press, 2012, pp. 50–70.
- [2] R. Gayler, "Vector Symbolic Architectures answer Jackendoff's challenges for cognitive neuroscience," in *ICCS/ASCS International Conference on Cognitive Science*, P. Slezak, Ed., University of New South Wales. CogPrints, 2003, pp. 133–138.
- [3] —, "Multiplicative Binding, Representation Operators and Analogy," in *Advances in analogy research: Integration of theory and data from the cognitive, computational, and neural sciences*, B. N. K. D. Gentner, K. J. Holyoak, Ed. Sofia, Bulgaria: New Bulgarian University, 1998, pp. 1–4. [Online]. Available: <http://cogprints.org/502>
- [4] P. Kanerva, *Sparse Distributed Memory*. MIT Press, 1988.
- [5] T. Plate, "Distributed Representations and Nested Compositional Structure," Ph.D. dissertation, University of Toronto, 1994.
- [6] P. Smolensky, "Tensor Product Variable Binding and the Representation of Symbolic Structures in Connectionist Systems," *Artificial Intelligence*, vol. 46, no. 1-2, pp. 159–216, 1990.
- [7] K. Schlegel, P. Neubert, and P. Protzel, "A comparison of Vector Symbolic Architectures," p. arXiv:2001.11797, Jan. 2020. [Online]. Available: <https://arxiv.org/abs/2001.11797>
- [8] D. Rasmussen and C. Eliasmith, "A Neural Model of Rule Generation in Inductive Reasoning," *Topics in Cognitive Science*, vol. 3, no. 1, pp. 140–153, 2011.
- [9] C. Eliasmith, *How to build a brain: A neural architecture for biological cognition*. Oxford University Press, 2013.
- [10] C. Eliasmith, T. C. Stewart, X. Choo, T. Bekolay, T. DeWolf, Y. Tang, and D. Rasmussen, "A Large-Scale Model of the Functioning Brain," *Science*, vol. 338, no. 6111, pp. 1202–1205, 2012. [Online]. Available: <http://science.sciencemag.org/content/338/6111/1202>
- [11] P. Blouw, E. Solodkin, P. Thagard, and C. Eliasmith, "Concepts as Semantic Pointers: A Framework and Computational Model," *Cognitive Science*, vol. 40, no. 5, pp. 1128–1162, 2016.
- [12] E. Crawford, M. Gingerich, and C. Eliasmith, "Biologically Plausible, Human-Scale Knowledge Representation," *Cognitive Science*, vol. 40, no. 4, pp. 782–821.
- [13] D. Kleyko, E. Osipov, R. W. Gayler, A. I. Khan, and A. G. Dyer, "Imitation of honey bees' concept learning processes using Vector Symbolic Architectures," *Biologically Inspired Cognitive Architectures*, vol. 14, pp. 57–72, 2015. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S2212683X15000456>
- [14] P. Neubert, S. Schubert, and P. Protzel, "An introduction to hyperdimensional computing for robotics," *KI - Künstliche Intelligenz*, vol. 33, no. 4, pp. 319–330, 2019. [Online]. Available: <https://doi.org/10.1007/s13218-019-00623-z>
- [15] —, "Learning Vector Symbolic Architectures for Reactive Robot Behaviours," in *Proc. of International Conference on Intelligent Robots and Systems (IROS), Workshop on Machine Learning Methods for High-Level Cognitive Capabilities in Robotics*, 2016.
- [16] F. Mirus, T. C. Stewart, and J. Conrads, "Towards cognitive automotive environment modelling: reasoning based on vector representations," in *26th European Symposium on Artificial Neural Networks, ESANN 2018, Bruges, Belgium*, 2018, pp. 55–60.
- [17] F. Mirus, P. Blouw, T. C. Stewart, and J. Conrads, "An investigation of vehicle behavior prediction using a vector power representation to encode spatial positions of multiple objects and neural networks," *Frontiers in Neurobotics*, vol. 13, p. 84, oct 2019. [Online]. Available: <https://www.frontiersin.org/article/10.3389/fnbot.2019.00084>
- [18] B. Komer, T. C. Stewart, A. R. Voelker, and C. Eliasmith, "A neural representation of continuous space using fractional binding," in *41st Annual Meeting of the Cognitive Science Society*. Cognitive Science Society, 2019.
- [19] G. A. Miller, "The magical number seven, plus or minus two: some limits on our capacity for processing information," *Psychological Review*, vol. 63, no. 2, pp. 81–97, 1956.
- [20] M. Wahle, D. Widdows, J. R. Herskovic, E. V. Bernstam, and T. Cohen, "Deterministic binary vectors for efficient automated indexing of medline/pubmed abstracts," vol. 2012, no. PMC3540485, pp. 940–949. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3540485/>
- [21] D. Widdows and T. Cohen, "Reasoning with vectors: A continuous model for fast robust inference," *Logic Journal of IGPL*, vol. 23, no. 2, pp. 141–173, 2014. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC4646228/>