

---

# On the role of noise in factorizers for disentangling distributed representations

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 To efficiently factorize high-dimensional distributed representations to the con-  
2 stituent atomic vectors, one can exploit the compute-in-superposition capabilities  
3 of vector-symbolic architectures (VSA). Such factorizers however suffer from the  
4 phenomenon of limit cycles. Applying noise during the iterative decoding is one  
5 mechanism to address this issue. In this paper, we explore ways to further relax  
6 the noise requirement by applying noise only at the time of VSA’s reconstruc-  
7 tion codebook initialization. While the need for noise during iterations proves  
8 analog in-memory computing systems to be a natural choice as an implementa-  
9 tion media, the adequacy of initialization noise allows digital hardware to remain  
10 equally indispensable. This broadens the implementation possibilities of factoriz-  
11 ers. Our study finds that while the best performance shifts from initialization noise  
12 to iterative noise as the number of factors increases from 2 to 4, both extend the  
13 operational capacity by at least  $50\times$  compared to the baseline factorizer resonator  
14 networks.

## 15 1 Introduction

16 Some basic Visual perception tasks, such as disentangling static elements from moving objects in  
17 a dynamic scene and enabling the understanding of object persistence, depend upon the factoriza-  
18 tion problem [1–4]. The principle of factorization extends to auditory perception, e.g. separating  
19 individual voices or instruments from a complex soundscape [5]. Factorization also plays a key  
20 role in higher-level cognitive tasks. Understanding analogies for example requires decomposing the  
21 underlying relationships between different concepts or situations [6–10]. While biological neural  
22 circuits solve the above challenges deftly, factorization remains a problem unsolvable within poly-  
23 nomial time complexity [11]. Outside the biological domain, factorization is at the core of many  
24 rapidly developing fields. In robotics, factorization can enable robots to 1) understand their environ-  
25 ment by parsing visual scenes and identifying objects, locations, and relationships [12], and 2) plan  
26 and execute tasks by decomposing complex actions into a simple sequence of steps. Factorizing  
27 semi-primes has implications for cryptography and coding theory [13]. Factorization helps develop  
28 more transparent and explainable AI systems [14], by decomposing complex decision processes into  
29 understandable components.

30 Vector-symbolic architectures (VSA) [15–18] is an emerging computing paradigm that can repre-  
31 sent and process a combination of attributes using high-dimensional holographic vectors. Unlike the  
32 traditional methods where information might be stored in specific locations (e.g., a single bit repre-  
33 senting a value), VSA distributes the information across the entire vector signifying a holographic  
34 nature. This means if some components of the vector are corrupted or lost, the overall information  
35 can still be recovered due to the distributed nature of the encoding. In high-dimensional spaces,

36 randomly generated vectors are very likely to be nearly orthogonal aka quasi-orthogonal [18]. This  
37 crucial property allows efficient and robust representation of a vast number of different concepts  
38 in VSA with minimal interference using these nearly orthogonal vectors as building blocks. More  
39 details on the background of VSA are provided in Appendix 5.1.

40 A typical algebraic operation for combining  $F$  different attributes in an object is to element-wise  
41 multiply associated  $D$ -dimensional holographic vectors. Due to the properties of high-dimensional  
42 spaces, this operation tends to produce unique representations for different attribute combinations.  
43 A deep convolutional neural network can be trained to generate these product vectors approxi-  
44 mately [14] by taking the raw image of the object. The inverse of the binding problem becomes  
45 the factorization problem which is the disentangling of an exact product vector or, as in the latter  
46 case, an inexact product vector, into its constituent attribute vectors. While binding is a straightfor-  
47 ward calculation, factorization involves searching for the correct combination of attributes among  
48 an exponentially large space of possibilities.

49 Resonator network [19, 20], a type of factorizer, built upon the VSA computing paradigm, introduces  
50 a promising approach to perform rapid and robust factorization. Resonator networks employ an  
51 iterative, parallel search strategy over high-dimensional vector data structures called *codebooks*,  
52 leveraging the properties of high-dimensional spaces to explore the search space efficiently and  
53 converge on the most likely attribute combinations. In resonator networks’ iterative search strategy  
54 *limit cycles* pose an obstacle to effective functioning.

55 As one potential remedy, the IN-MEMORY FACTORIZER (IMF) [21] harnesses the intrinsic stochas-  
56 tic noise of analog in-memory computing (IMC) to break free of limit cycles during the iterative  
57 decoding. Together with an additional nonlinearity in the form of a computationally cheap thresh-  
58 olding function on the similarity estimates, IMF solves significantly larger problem sizes compared  
59 to the original resonator networks. Indeed, more recent advancements on resonator networks [22]  
60 make also use of nonlinearities in the form of ReLU and exponentiation and noise on the similarity  
61 estimates. In this case, the noise has to be generated by a dedicated noise source (e.g., a random  
62 number generator) at every decoding iteration. IMF however does not need such an additional noise  
63 source thanks to the IMC’s intrinsic noise; yet, its performance relies on the availability of underly-  
64 ing hardware with specific noise levels across decoding iterations.

65 In this article, we explore novel factorizers with alternative noise requirements to mitigate limit cy-  
66 cles. In particular, we propose ASYMMETRIC CODEBOOK FACTORIZER (ACF), where codebooks  
67 are initialized with some noisy perturbations and the same instance of noise used across iterations  
68 providing a relaxed noise requirement to circumvent limit cycles. We conceptualize a purely digital  
69 factorizer design that can provide energy efficiency gains by eliminating data conversions. We find  
70 that compared to the baseline resonator network [19], the variants embracing noise, IMF, and ACF,  
71 always perform better in terms of operational capacity and require fewer iterations to converge at  
72 different sizes of search spaces.

## 73 2 Background: Resonator Networks

74 The resonator network is an iterative algorithm designed to factorize high-dimensional vectors [19].  
75 Essential to the operation of resonator networks are *codebooks*, which serve as repositories of quasi-  
76 orthogonal high-dimensional vectors called *codevectors*, with each codevector representing a dis-  
77 tinct attribute value. For instance, in a system designed for visual object recognition, 3 separate  
78 codebooks might store representations for shapes (like circles, squares, triangles), colors (red, green,  
79 blue), and positions (top left, bottom right, center) each having 3 distinct possibilities for the attribute  
80 values. See Fig. 1. The number of potential combinations of these codevectors to form product vec-  
81 tors grows  $M^F$  with respect to the number of attributes (i.e., factors  $F$ ) and attribute values (i.e.,  
82 codebook size  $M$ ) exponentially. However, the dimensionality ( $D$ ) of these vectors is usually fixed  
83 to several hundred leading to ( $D \ll M^F$ ), searching for the specific combination of codevectors  
84 that compose a given product vector becomes computationally expensive, presenting a hard combi-  
85 natorial search problem. Thanks to the quasi-orthogonal property of codevectors however, resonator  
86 networks can rapidly navigate the vast search space.

87 Starting with an initial estimate for each factor, and the product vector, it updates the estimates one  
88 factor at a time. It achieves this by first “*unbinding*” (UB) or removing the influence of the estimated  
89 values of all but the selected factor from the product vector. In the context of bipolar spaces, where

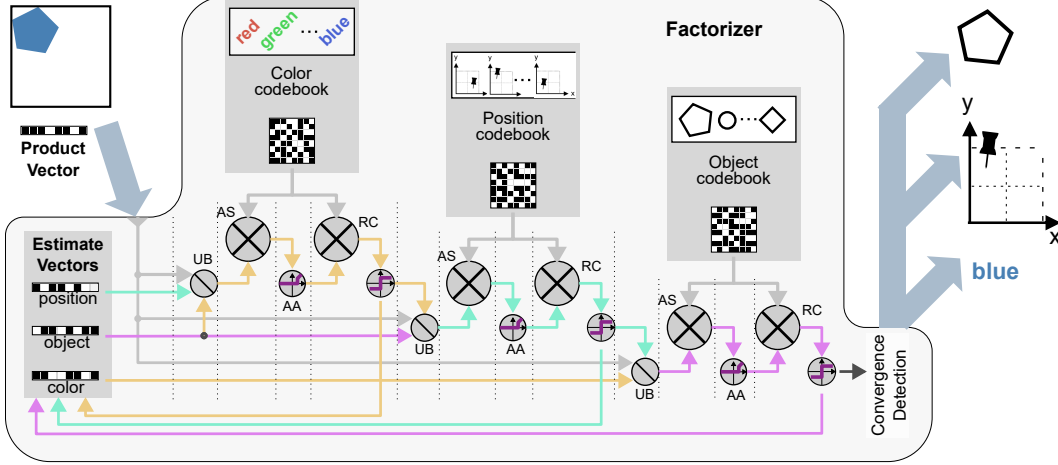


Fig. 1: Factorization of perceptual representations: color position and object type using factorizers. Taking the product vector as input, and starting from initial estimate vectors the factorizer undergoes unbinding (UB), associative search (AS), attention activation (AA), and reconstruction (RC) phases to iteratively refine the estimate. AS and RC take the biggest share of the computing and memory. They map to a predominantly MVM operation.

90 codevector components are +1 or -1, unbinding is accomplished through element-wise multiplication.  
 91 Secondly, the unbound vector is compared against all codevectors within the corresponding  
 92 codebook using an associative search (AS), typically a series of cosine similarity or dot products  
 93 that can be formulated as a matrix-vector-multiplication (MVM) operation. The associative search  
 94 outputs an attention vector of length  $M$ , measuring how the unbound vector aligns with the codevec-  
 95 tors. Thirdly, the attention vector is passed through an optional attention activation (AA) function.  
 96 This allows us to filter out uninteresting codevectors. Fourthly, using activated attentions as weights,  
 97 codevectors are superimposed to reconstruct a new estimate vector for the selected factor. This re-  
 98 construction (RC) operation can also be formulated by an MVM operation with the transpose of the  
 99 codebook itself acting as the reconstruction matrix and the activated attention acting as the vector.  
 100 The above four phases are repeated, refining the estimates for each factor with each iteration until  
 101 the resonator network converges to a stable solution representing the most probable factorization of  
 102 the product vector.

103 This can be mathematically formulated as follows:

104 Let us consider a three-factor ( $F = 3$ ) problem with the factors originating from three codebooks:

105  $\mathbf{A} = \{\mathbf{a}^{(1)}, \mathbf{a}^{(2)}, \dots, \mathbf{a}^{(M)}\}$ ,  $\mathbf{B} = \{\mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \dots, \mathbf{b}^{(M)}\}$ ,  $\mathbf{C} = \{\mathbf{c}^{(1)}, \mathbf{c}^{(2)}, \dots, \mathbf{c}^{(M)}\}$ ,  
 106 each with  $D$  dimensional bipolar codevectors  $\mathbf{A}, \mathbf{B}, \mathbf{C} \in \{-1, +1\}^{D \times M}$ . Let the estimate for  
 107 each factor be represented using  $\hat{\mathbf{a}}, \hat{\mathbf{b}}, \hat{\mathbf{c}}$  respectively, and the product vector denoted with  $\mathbf{x}$ , then for  
 108 the first factor, unbinding of other factors is given by:  $\tilde{\mathbf{a}} = \mathbf{x} \oslash \hat{\mathbf{b}} \oslash \hat{\mathbf{c}}$ . Based on the unbound vector  $\tilde{\mathbf{a}}$   
 109 of the first factor, AS, AA, and RC phases can be written as in the following equations respectively.

110 
$$\alpha_a(t) = \tilde{\mathbf{a}}(t)\mathbf{A}^T \quad (1) \quad , \quad \alpha'_a(t) = f(\alpha_a(t)) \quad (2) \quad , \quad \hat{\mathbf{a}}(t+1) = \text{sign}(\alpha'_a(t)\mathbf{A}) \quad (3)$$

111 Where  $t, \alpha, f, \text{sign}$  stands for iteration number, attention vector, activation function, and signum  
 112 function respectively. A similar approach can be followed to compute the next estimate vector of  
 113 the other factors  $\hat{\mathbf{b}}, \hat{\mathbf{c}}$ .

114 **3 Breaking Free from Limit Cycles with Noise**

115 As we discussed, resonator networks operate iteratively, progressively refining their estimates for  
 116 each factor by comparing them to codebook entries. However, during this iterative process, the  
 117 network can sometimes get trapped in a repeating sequence of states. Then the network’s estimate  
 118 for a factor oscillates between a small set of codevectors without ever settling on the true factor.  
 119 This phenomenon, referred to as a *limit cycle*, can prevent the network from reaching the optimal  
 120 solution forever.

121 The emergence of limit cycles can be attributed to the symmetric and deterministic nature of the  
 122 codebooks used in the AS and RC phases of the baseline resonator network’s (BRN) [19, 20] search  
 123 procedure. This deterministic behavior is particularly problematic when the search space is large  
 124 and contains many local minima, which can trap the network’s updates in a repeating pattern. When  
 125 a resonator network gets stuck in a limit cycle, it fails to converge to the correct factorization, even  
 126 if given an unlimited number of iterations. This lack of convergence can significantly impact the  
 127 network’s accuracy and efficiency, rendering it ineffective for tasks that require precise factorization.

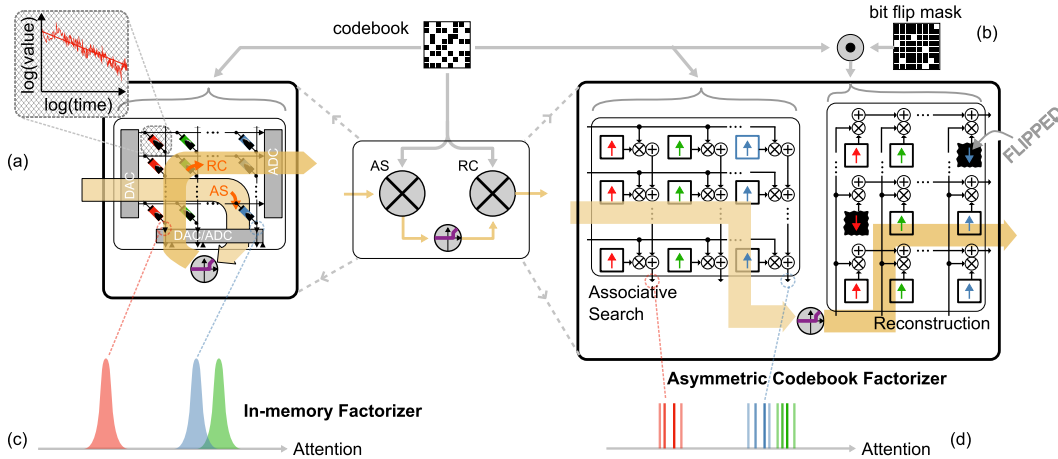


Fig. 2: Implementing noise during a decoding iteration of a single factor of factorizer. (a) The codebooks are implemented on an analog memory device crossbar array which introduce intrinsic noise to each iteration of both the AS and RC phases. (b) The codebooks are implemented on digital memory devices. The second codebook used in the RC phase is made asymmetric from the first codebook in the AS phase by a bit flip mask perturbation. (c) The resulting attention using analog IMC. It can have a continuous distribution. (d) The resulting attention using asymmetric codebooks. It follows a discrete distribution.

128 Introducing stochasticity into the network’s update rules can help it break free from deterministic  
 129 limit cycles. This is a crucial finding that not only pushes the factorizers’ operational capacity but  
 130 also shifts the hardware landscape they thrive. In particular, there is a potential for leveraging the  
 131 intrinsic randomness associated with memristive devices in IMC implementations of factorizers as  
 132 prescribed in IMF [21]. An example implementation is illustrated in Fig. 2(a). The codevectors of a  
 133 codebook are programmed along the columns of the crossbar arrays. In the first crossbar used for  
 134 the AS phase, the inputs are passed through the west side digital to analog converters (DACs), and  
 135 the resulting currents/charges are collected on the south periphery and converted back to the digital  
 136 domain using the analog to digital converters (ADCs). After activating the resulting attentions,  
 137 the sparse attention vector is input through the south-side DACs for the RC phase. The resulting  
 138 currents/charges are collected through the east side periphery and converted to digital using ADCs  
 139 before converting via signum function to the next estimate vector.

140 The memory devices used in these arrays are fabricated using phase-change memory (PCM) tech-  
 141 nology and exhibit natural variations in their behavior forming a near-Gaussian distribution of  
 142 the attention result as seen in Fig. 2(c). This can be expressed in an approximated form as  
 143  $\alpha_a(t) = \tilde{\mathbf{a}}(t)\mathbf{A}^T + n$ . Where  $n$  is the  $M$ -dimensional noise vector sampled from i.i.d. Gaus-  
 144 sian distribution  $n \sim \mathcal{N}(0, \sigma \cdot \mathbf{I}_M)$  and  $\sigma$  denotes the standard deviation of the noise, which ranges  
 145 around 0.01 in a recent large-scale chip based on phase-change memory devices [23]. As seen in

146 Sec. 4, this  $\sigma$  value falls in the *useful noise* range enabling escaping repeating patterns and exploring  
 147 a wider range of solutions.

148 Even if an arbitrary number of iterations are allowed, a deterministic digital design with unperturbed  
 149 symmetric codebooks fails to achieve the accuracy of the IMF due to its susceptibility to limit cycles.  
 150 Stochasticity, as a key operation to eliminate limit cycles, has significant added costs in terms of  
 151 energy and area in mature digital hardware implementation. For example, generating Gaussian noise  
 152 involves several expensive floating point operations such as exponentiation and multiplication.

153 We explore the possibility of inserting the noise into the codebooks at the time of initialization. If  
 154 the same noise is added to both copies of the codebook it would still keep the same quasi-orthogonal  
 155 relationship of the codevectors and would not change the dynamics of the factorizer. Instead, we  
 156 propose perturbing only a single copy of the codebook using a randomly generated bitflip mask.  
 157 Fig. 2(b) shows perturbing the codebook used in the RC phase by applying a bit flip mask  $BFM \in$   
 158  $\{-1, +1\}^{D \times M}$  of certain sparsity as shown in Eq. 4. We call this type of model an Asymmetric  
 159 Codebook Factorizer (ACF).

$$BFM(r) = \begin{cases} +1 & \text{if } \mathbf{u} + r > 1 \\ -1 & \text{otherwise} \end{cases} \quad \begin{aligned} \mathbf{A}_{RC} &= \mathbf{A} \odot BFM(r) \\ \hat{\mathbf{a}}(t+1) &= \text{sign}(\alpha'_a(t)\mathbf{A}_{RC}) \end{aligned} \quad (4)$$

160 Where  $\mathbf{u} \sim \mathcal{U}(0, 1) \in [0, 1]^{D \times M}$  is a noise matrix sampled from the uniform distribution. The  
 161 sparsity  $r$  is a hyperparameter that can be optimally obtained using a hyperparameter search scheme.  
 162 The perturbed codebook for the RC phase  $\mathbf{A}_{RC}$  is calculated by element-wise multiplication between  
 163  $BFM$  and the codebook as shown in Eq. 4. As  $r$  increases the factorizer starts losing its ability  
 164 to converge and iterate on the correct solution. However, with the presence of the convergence  
 165 detection circuit detailed in Appendix 5.2, the need to *resonate* is not essential. Similarly, codebooks  
 166 that are not bipolar in nature [24] can be perturbed using an appropriate random function.

167 Apart from noise, another known approach to solving limit cycles and converge faster to the right  
 168 solution involves the use of non-linear activation functions. One such activation function, employed  
 169 both in IMF and ACF, uses a threshold to sparsify the attention vector. It is explained further in  
 170 Appendix 5.3.

## 171 4 Results and Discussion

172 We conduct experiments to measure the operational capacity and other behaviors of different variants  
 173 of factorizers, namely the BRN, IMF, and ACF. Operational capacity is defined as the maximum size  
 174 of the search space that can be handled with more than 99% accuracy while requiring fewer iterations  
 175 than what a brute force approach would have taken. A brute force approach would in the worst case  
 176 find the correct factors in  $M^F$  steps.

177 The results are presented in Fig. 3. We consider 3 cases for the number of factors  $F = \{2, 3, 4\}$ . The  
 178 dimensions of the codevectors for these cases are set based on the minimum dimensions reported in  
 179 the BRN, namely  $D = 1000, 1500, 2000$  respectively for  $F = 2, 3, 4$  respectively. For each case, we  
 180 span the search space size starting from  $10^4$  up to a maximum of  $10^{11}$  until the operational capacity  
 181 is reached. At each selected search space size, we conduct 5000 trials factorizing randomly sampled  
 182 product vectors. From this, we calculate the average accuracy and the average number of iterations  
 183 as reported in Fig. 3.

184 The BRN reaches its capacity at approximately  $10^5, 10^6,$  and  $10^7$  for  $F = 2, 3, 4$  cases respectively.  
 185 Although the accuracy momentarily dropped slightly below 99% in a few search space sizes between  
 186  $10^5$  and  $10^6$  at  $F = 2$ , IMF does not reach the operational capacity for all search space sizes tested.  
 187 For ACF, we observe the reaching of the operational capacity at  $> 5 \times 10^9$  for  $F = 4$ . In the  
 188 other two cases, ACF did not reach the operational capacity point for the search space sizes tested.  
 189 The momentary drop in accuracy and rise in iterations in certain search spaces can be attributed to  
 190 inadequate hyperparameter exploration. The optimum hyperparameter setting we achieved during  
 191 our experiments is further detailed in Appendix 5.4

192 Both IMF and ACF exhibit better performance in terms of operational capacity and the number of  
 193 iterations compared to the BRN. In theory, the IMF has better control over noise as it is applied  
 194 over the iterations. This becomes clear in the  $F = 4$  case where it outperforms ACF by achieving

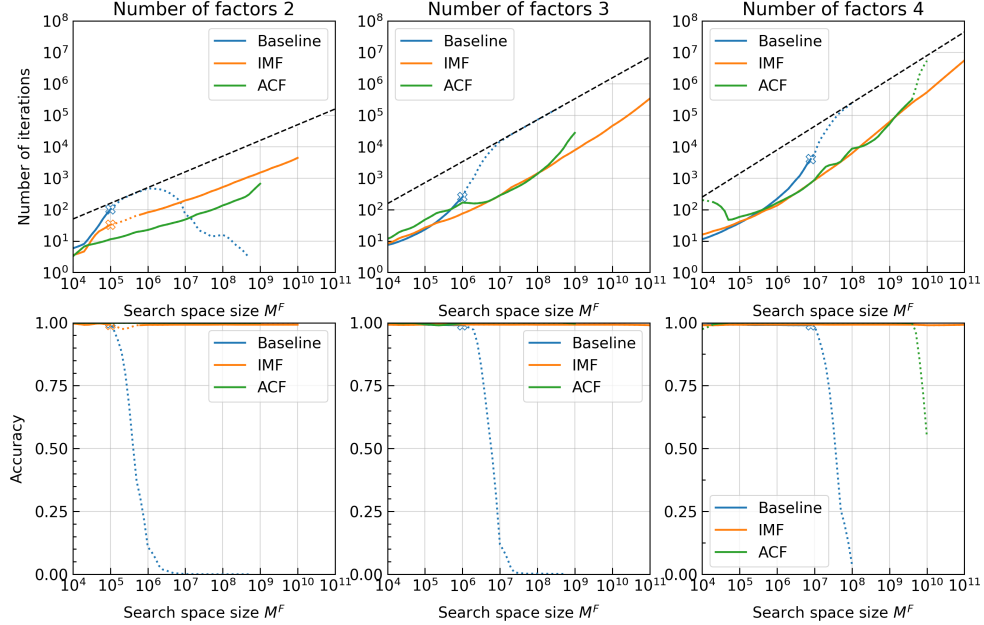


Fig. 3: The number of iterations (top row) and accuracy (bottom row) for three variants of the factorizer: baseline resonator (in blue) [19], in-memory factorizer (in orange) [21], and our asymmetric codebook factorizer (in green). The left, center, and right column results correspond to 2,3, and 4-factor scenarios, respectively. The regions that meet operational capacity criteria (i.e.  $\geq 99\%$  accuracy) are plotted with solid lines while other regions are plotted with dotted lines.

195 greater operational capacity. ACF however edges over IMF in the  $F = 2$  case where there are fewer  
 196 interactions among factors.

197 Another aspect to consider is the hardware design costs. As shown in Fig. 2, IMF achieves area  
 198 efficiency with a single copy of codebooks in a crossbar array and achieves energy efficiency by  
 199 performing arithmetic operations implicitly using device physics. ACF on the other hand has explicit  
 200 multipliers and adders but saves the bulk of the energy spent on converting data from digital to analog  
 201 and vice versa several times per decoding iteration. As a consequence of the converter-less design,  
 202 ACF can operate faster, with several nanoseconds per iteration as opposed to several microseconds.  
 203 Thus ACF can achieve more iterations per unit period of time.

204 The principles used in the IMF and ACF are not mutually exclusive. While in this work we study  
 205 and compare their standalone performance, there is no reason that prevents them from being em-  
 206 ployed in unison. One possible realization of this involves perturbing the target conductance values  
 207 corresponding to the codebook values before they get programmed into the analog memory devices  
 208 in the IMF. Incorporating both sources of noise may result in a synergistic effect enabling more  
 209 performant factorizers.

## 210 5 Conclusion

211 In conventional wisdom, stochasticity and noise are considered a bane in computing. We demon-  
 212 strate that factorizers grounded on VSAs empower a new computing paradigm that embraces noise  
 213 to push the limits of operational capacity. We discuss two variants, the first harnessing intrinsic noise  
 214 in analog in-memory computing during MVM operation, the second initializing the codebooks with  
 215 noisy perturbations yielding a model that widely appeals to deterministic digital design systems.  
 216 While there are tradeoffs, both these variants empirically outperform the baseline resonator networks  
 217 in multiple facets. When combined with appropriate hardware designs, they provide promising di-  
 218 rections to solve factorization problems in large-scale search spaces within reasonable timescales.

## References

- 219
- 220 [1] Roland Memisevic and Geoffrey E. Hinton. Learning to represent spatial transformations with  
221 factored higher-order boltzmann machines. *Neural Computation*, 22(6):1473–1492, 06 2010.
- 222 [2] Yoram Burak, Uri Rokni, Markus Meister, and Haim Sompolinsky. Bayesian model of dy-  
223 namic image stabilization in the visual system. *Proceedings of the National Academy of Sci-  
224 ences*, 107(45):19525–19530, 2010.
- 225 [3] Charles F. Cadieu and Bruno A. Olshausen. Learning intermediate-level representations of  
226 form and motion from natural movies. *Neural Computation*, 24(4):827–866, 04 2012.
- 227 [4] Alexander G. Anderson, Kavitha Ratnam, Austin Roorda, and Bruno A. Olshausen. High-  
228 acuity vision from retinal image motion. *Journal of Vision*, 20(7):34–34, 07 2020.
- 229 [5] Ying Hu and Guizhong Liu. Separation of singing voice using nonnegative matrix partial  
230 co-factorization for singer identification. *IEEE/ACM Transactions on Audio, Speech, and Lan-  
231 guage Processing*, 23(4):643–653, 2015.
- 232 [6] J. E. Hummel and K. J. Holyoak. Distributed Representations of Structure: A Theory of  
233 Analogical Access and Mapping. *Psychological Review*, 104(3):427–466, 1997.
- 234 [7] P. Kanerva. Dual Role of Analogy in the Design of a Cognitive Computer. In *Advances in  
235 Analogy Research: Integration of Theory and Data from the Cognitive, Computational, and  
236 Neural Sciences*, pages 164–170, 1998.
- 237 [8] P. Kanerva. Pattern Completion with Distributed Representation. In *International Joint Con-  
238 ference on Neural Networks*, pages 1416–1421, 1998.
- 239 [9] T. A. Plate. Analogy Retrieval and Processing with Distributed Vector Representations. *Expert  
240 Systems: The International Journal of Knowledge Engineering and Neural Networks*, 17(1):  
241 29–40, 2000.
- 242 [10] R. W. Gayler and S. D. Levy. A Distributed Basis for Analogical Mapping: New frontiers in  
243 Analogy Research. In *New frontiers in Analogy Research, Second International Conference  
244 on the Analogy*, pages 165–174, 2009.
- 245 [11] Steven G Krantz. *The proof is in the pudding: The changing nature of mathematical proof*.  
246 Springer, 2011.
- 247 [12] Alpha Renner, Lazar Supic, Andreea Danielescu, Giacomo Indiveri, E Paxon Frady,  
248 Friedrich T Sommer, and Yulia Sandamirskaya. Visual odometry with neuromorphic resonator  
249 networks. *Nature Machine Intelligence*, 6(6):653–663, 2024.
- 250 [13] Denis Kleyko, Connor Bybee, Christopher J Kymn, Bruno A Olshausen, Amir Khosrowshahi,  
251 Dmitri E Nikonov, Friedrich T Sommer, and E Paxon Frady. Integer factorization with compo-  
252 sitional distributed representations. In *Proceedings of the 2022 Annual Neuro-Inspired Com-  
253 putational Elements Conference*, pages 73–80, 2022.
- 254 [14] Michael Hersche, Mustafa Zeqiri, Luca Benini, Abu Sebastian, and Abbas Rahimi. A neuro-  
255 vector-symbolic architecture for solving raven’s progressive matrices. *Nature Machine Intelli-  
256 gence*, 5(4):363–375, 2023.
- 257 [15] R. W. Gayler. Vector Symbolic Architectures Answer Jackendoff’s Challenges for Cognitive  
258 Neuroscience. In *Joint International Conference on Cognitive Science*, pages 133–138, 2003.
- 259 [16] T. A. Plate. Holographic Reduced Representations. *IEEE Transactions on Neural Networks*, 6  
260 (3):623–641, 1995.
- 261 [17] T. A. Plate. *Holographic Reduced Representations: Distributed Representation for Cognitive  
262 Structures*. Stanford University, 2003.
- 263 [18] P. Kanerva. Hyperdimensional Computing: An Introduction to Computing in Distributed Rep-  
264 resentation with High-Dimensional Random Vectors. *Cognitive Computation*, 1(2):139–159,  
265 2009.

- 266 [19] E. P. Frady, S. J. Kent, B. A. Olshausen, and F. T. Sommer. Resonator Networks, 1: An  
267 Efficient Solution for Factoring High-Dimensional, Distributed Representations of Data Structures. *Neural Computation*, 32(12):2311–2331, 2020.  
268
- 269 [20] S. J. Kent, E. P. Frady, F. T. Sommer, and B. A. Olshausen. Resonator Networks, 2: Factoriza-  
270 tion Performance and Capacity Compared to Optimization-Based Methods. *Neural Computa-  
271 tion*, 32(12):2332–2388, 2020.
- 272 [21] Jovin Langenegger, Geethan Karunaratne, Michael Hersche, Luca Benini, Abu Sebastian, and  
273 Abbas Rahimi. In-memory factorization of holographic perceptual representations. *Nature  
274 Nanotechnology*, 18(5):479–485, 2023.
- 275 [22] Alpha Renner, Lazar Supic, Andreea Danielescu, Giacomo Indiveri, Bruno A Olshausen, Yulia  
276 Sandamirskaya, Friedrich T Sommer, and E Paxon Frady. Neuromorphic visual scene under-  
277 standing with resonator networks. *Nature Machine Intelligence*, 6(6):641–652, 2024.
- 278 [23] Riduan Khaddam-Aljameh, Milos Stanisavljevic, et al. Hermes core—a 14nm cmos and PCM-  
279 based in-memory compute core using an array of 300ps/LSB linearized CCO-based ADCs and  
280 local digital processing. In *2021 Symposium on VLSI Circuits*, pages 1–2. IEEE, 2021.
- 281 [24] Michael Hersche, Aleksandar Terzic, Geethan Karunaratne, Jovin Langenegger, Angéline  
282 Pouget, Giovanni Cherubini, Luca Benini, Abu Sebastian, and Abbas Rahimi. Factorizers  
283 for distributed sparse block codes. *Neurosymbolic Artificial Intelligence*, 2024.
- 284 [25] Denis Kleyko, Dmitri A. Rachkovskij, Evgeny Osipov, and Abbas Rahimi. A survey on hy-  
285 perdimensional computing aka vector symbolic architectures, part I: models and data transfor-  
286 mations. *ACM Comput. Surv.*, may 2022.
- 287 [26] Denis Kleyko, Dmitri A. Rachkovskij, Evgeny Osipov, and Abbas Rahimi. A survey on hy-  
288 perdimensional computing aka vector symbolic architectures, part II: applications, cognitive  
289 models, and challenges. *ACM Comput. Surv.*, June 2022.



## 290 Appendix

### 291 5.1 Vector-symbolic architectures

292 Here, we provide a brief overview of vector-symbolic architectures (VSAs) [15–18] of which the  
293 resonator networks [19, 20] are based on. VSA is a powerful computing framework that is built on  
294 an algebra in which all representations are high-dimensional holographic vectors of the same, fixed  
295 dimensionality denoted by  $D$ . This is attributed to modeling the representation of information in the  
296 brain as distributed over many neurons. In this work, we consider a VSA model based on bipolar  
297 vector space [15], i.e.,  $\{-1, +1\}^D$ . The similarity between two vectors is defined as the cosine  
298 similarity:

$$\text{sim}(\mathbf{x}_1, \mathbf{x}_2) = \frac{\langle \mathbf{x}_1, \mathbf{x}_2 \rangle}{\|\mathbf{x}_1\| \|\mathbf{x}_2\|} = \frac{\langle \mathbf{x}_1, \mathbf{x}_2 \rangle}{D} \quad (5)$$

299 As one of the main property of the high-dimensional vector space, any two randomly drawn vectors  
300 lie close to quasi-orthogonality to each other, i.e., their expected similarity is close to zero with a  
301 high probability [18]. The vectors can represent symbols, and can be manipulated by a rich set of  
302 dimensionality-preserving algebraic operations:

- 303 • **Binding:** Denoted by  $\odot$ , the Hadamard (i.e., element-wise) product of two input vectors  
304 implements the binding operation. It is useful to represent a hierarchical structure whereby  
305 the resulting vector lies quasi-orthogonal to all the input vectors. The binding operation  
306 follows the commutative law  $\mathbf{x}_1 \odot \mathbf{x}_2 = \mathbf{x}_2 \odot \mathbf{x}_1 = \mathbf{p}$ .
- 307 • **Unbinding:** The unbinding operation reverses the binding operation. As the element-wise  
308 multiplication in the bipolar space is self-inverse, the same operation as for the binding can  
309 be used. Using the unbinding operator  $\oslash$  the operation is defined as  $\mathbf{p} \oslash \mathbf{x}_1 = \mathbf{x}_2$ .
- 310 • **Bundling:** The superposition of two vectors is calculated by the bundling operation  $\oplus$ . The  
311 operation is defined by an element-wise sum with consecutive bipolarization. In case of an  
312 element-wise sum equal to zero, we randomly bipolarize.
- 313 • **Clean-up:** The clean-up operation maps a noisy vector to its noise-free representation by  
314 an associative memory lookup.
- 315 • **Permutation:** Permutation is a unary operation on a vector that yields a quasi-orthogonal  
316 vector of its input. This operation rotates the coordinates of the vector. A simple way to  
317 implement this is as a cyclic shift by one position.

318 Interested readers can refer to a detailed survey [25, 26] about VSAs.

### 319 5.2 Detection of convergence

320 The iterative factorization problem is said to be converged if, for two consecutive time steps, all  
321 the estimates are constant, i.e.,  $\hat{\mathbf{x}}_f(t+1) = \hat{\mathbf{x}}_f(t)$  for  $f \in [1, F]$ . We define an early convergence  
322 detection algorithm since it avoids unnecessary iterations and in the case of Asymmetric Codebook  
323 Factorizer (ACF), the legacy definition of convergence no longer holds.

324 In the new definition, the factorizer is said to be converged if a single similarity value across all the  
325 factors surpasses a convergence detection threshold:

$$\text{converged} = \begin{cases} \text{true,} & \text{if } \max(\alpha_f(t)[i]) > T_{\text{convergence}} \\ \text{false,} & \text{otherwise,} \end{cases} \quad (6)$$

326 where  $i \in [1, M]$  for  $f \in [1, F]$ . Upon convergence, the predicted factorization is given by the  
327 codevector associated with the largest similarity value per each codebook. This algorithm also  
328 eliminates the need to store the history of prior estimates. In the legacy definition of convergence, the  
329 previous estimate for each factor had to be stored to be able to compare it to the current estimate and  
330 detect convergence, resulting in a total of  $F \cdot D$  stored bits. Our experiments show that the optimal  
331 convergence detection threshold stays at a fixed ratio of  $D$  for any given set of hyperparameters and  
332 problem sizes.

333 **5.3 Threshold-based Activation**

334 Replacing the identity function, which acts as a linear activation in the BRN, with a nonlinear  
 335 winner-take-all approach can enhance both the convergence rate and the maximum solvable search  
 336 space[21]. This strategy sparsifies the similarity vector, essentially zeroing out weaker similarity  
 337 values and focusing the network’s attention on the most promising candidates. By suppressing less  
 338 likely solutions, sparse activation functions can help prevent the network from getting bogged down  
 339 in local minima and facilitate convergence to the global optimum. For a threshold  $T$ , the threshold-  
 340 based attention activation is given as:

$$\forall i \in (1, M) \quad \alpha'[i] = \begin{cases} \alpha[i], & \text{if } \alpha_i > T \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

341 **5.4 Hyperparameter Search**

342 For the three cases of experiments we conducted, we first set the following common hyperparam-  
 343 eters for both ACF and IMF. These include dimension ( $D$ ) and search space size ( $M^F$ ). Then the  
 344 following hyperparameters have to be tuned to achieve the best results. For ACF, they include the  
 345 sparsity parameter  $r$  and the activation threshold ( $T$ ). For IMF, they include iterative noise standard  
 346 deviation  $\sigma$  and the activation threshold ( $T$ ). Tables 1, 2, and 3 provide the optimum hyperparameter  
 347 combinations that give rise to the best accuracy and number of iterations results reported in Fig. 3.

Table 1: The hyperparameters that achieve the highest accuracy while minimizing the number of iterations for two factors

Common parameters			ACF-specific		IMF-specific	
Search space size $M^F$	Number of factors F	Dimension D	BFM Sparsity r	Activation Threshold T	Iterative Noise $\sigma$	Activation Threshold T
10000	2	1000	0.005	0.01	0.008	0.001
21609	2	1000	0.075	0	0.008	0.1
46225	2	1000	0.1	0	0.008	0.1
99856	2	1000	0.1	0	0.008	0.1
215296	2	1000	0.1	0	0.008	0
463761	2	1000	0.1	0	0.008	0
1000000	2	1000	0.1	0.075	0.008	0
2155024	2	1000	0.1	0	0.008	0
4639716	2	1000	0.1	0	0.008	0
9998244	2	1000	0.1	0	0.008	0
21548164	2	1000	0.1	0.1	0.008	0.1
46416969	2	1000	0.1	0.1	0.008	0.1
1.00E+08	2	1000	0.05	0.1	0.008	0.1
2.15E+08	2	1000	0.05	0.1	0.008	0.1
4.64E+08	2	1000	0.05	0.1	0.008	0.1
1.00E+09	2	1000	0.01	0.1	0.008	0.1

Table 2: The hyperparameters that achieve the highest accuracy while minimizing the number of iterations for three factors

Common parameters			ACF-specific		IMF-specific	
Search space size $M^F$	Number of factors F	Dimension D	BFM Sparsity r	Activation Threshold T	Iterative Noise $\sigma$	Activation Threshold T
10648	3	1500	0.1	0.01	0.007	0.001
21952	3	1500	0.1	0.01	0.007	0.01
46656	3	1500	0.05	0.01	0.007	0.01
97336	3	1500	0.05	0.01	0.007	0.01
216000	3	1500	0.005	0.01	0.007	0.05
456533	3	1500	0.1	0.05	0.007	0.05
1000000	3	1500	0.1	0.05	0.007	0.05
2146689	3	1500	0.1	0.05	0.007	0.05
4657463	3	1500	0.05	0.05	0.007	0.05
9938375	3	1500	0.05	0.05	0.007	0.05
21484952	3	1500	0.01	0.05	0.007	0.05
46268279	3	1500	0.01	0.05	0.007	0.05
99897344	3	1500	0.0005	0.05	0.007	0.05
2.15E+08	3	1500	0	0.05	0.007	0.05
4.64E+08	3	1500	0	0.05	0.007	0.05
1.00E+09	3	1500	0	0.05	0.007	0.05

Table 3: The hyperparameters that achieve the highest accuracy while minimizing the number of iterations for four factors

Common parameters			ACF-specific		IMF-specific	
Search space size $M^F$	Number of factors F	Dimension D	BFM Sparsity r	Activation Threshold T	Iterative Noise $\sigma$	Activation Threshold T
10000	4	2000	0.1	0	0.006	0.01
20736	4	2000	0.09	0	0.006	0.01
50625	4	2000	0.05	0	0.006	0.001
104976	4	2000	0.02	0.001	0.006	0.001
234256	4	2000	0.006	0.01	0.006	0.01
456976	4	2000	0.005	0	0.006	0.01
1048576	4	2000	0.001	0	0.006	0.01
2085136	4	2000	0.008	0.025	0.006	0.01
4477456	4	2000	0.006	0.025	0.006	0.01
9834496	4	2000	0.006	0.025	0.006	0.05
21381376	4	2000	0.006	0.025	0.006	0.05
47458321	4	2000	0.003	0.03	0.006	0.05
1.00E+08	4	2000	0.002	0.03	0.006	0.05
2.14E+08	4	2000	0.02	0.04	0.006	0.05
4.67E+08	4	2000	0.008	0.04	0.006	0.05
1.00E+09	4	2000	0.008	0.04	0.006	0.05