

Article

In Search of Dispersed Memories: Generative Diffusion Models Are Associative Memory Networks

Luca Ambrogioni

Donders Institute for Brain, Cognition, and Behaviour, Radboud University, 6525 XZ Nijmegen, The Netherlands; luca.ambrogioni@donders.ru.nl

Abstract: Uncovering the mechanisms behind long-term memory is one of the most fascinating open problems in neuroscience and artificial intelligence. Artificial associative memory networks have been used to formalize important aspects of biological memory. Generative diffusion models are a type of generative machine learning techniques that have shown great performance in many tasks. Similar to associative memory systems, these networks define a dynamical system that converges to a set of target states. In this work, we show that generative diffusion models can be interpreted as energy-based models and that, when trained on discrete patterns, their energy function is (asymptotically) identical to that of modern Hopfield networks. This equivalence allows us to interpret the supervised training of diffusion models as a synaptic learning process that encodes the associative dynamics of a modern Hopfield network in the weight structure of a deep neural network. Leveraging this connection, we formulate a generalized framework for understanding the formation of long-term memory, where creative generation and memory recall can be seen as parts of a unified continuum.

Keywords: generative diffusion models; associative memory networks; hopfield networks

1. Introduction

Memory is a mysterious thing. Human beings can form lifelong memories from fleeting events and effortlessly recall them decades later as vivid multi-sensory experiences. On the other hand, in spite of their impressive capabilities, deep learning systems tend to require extensive training sessions to encode new information, which severely limit their adaptivity and consequently their capacity to develop general intelligence. Nevertheless, the field of machine learning has a long history of research on biologically plausible memory [1–6]. Arising from the pioneering work of John Hopfield, associative memory networks have been proposed as computational models of biological memory [1,7]. These networks encode memories as stable fixed-points in an “energy landscape” defined on the space of neural activations. Interestingly, this energy function can be encoded into a pattern of synaptic connections trained using biologically plausible synaptic learning rules [2,8]. In recent years, associative memory networks have been generalized in order to greatly scale their encoding capacity [9,10]. However, these modern Hopfield networks have a much more tenuous connection with known forms of synaptic learning since their energy function cannot be straightforwardly captured by learned pairwise synaptic couplings [11]. While biologically plausible implementations of modern Hopfield networks have been proposed, they do not provide strong insights on how memories are encoded in synaptic patterns since they require hard storage of the memorized pattern [11].

While human-like long-term memory is still outside of the capabilities of artificial intelligence systems, great progress has been made in approximating several forms of human creativity. Modern generative models are nowadays capable of generating beautiful visual art and insightful written text. In fact, these generative models have attracted great attention in neuroscience since they can provide an internal model of the world [12–14] and form the mechanistic basis for imagination and top-down predictive perception [15].



Citation: Ambrogioni, L. In Search of Dispersed Memories: Generative Diffusion Models Are Associative Memory Networks. *Entropy* **2024**, *26*, 381. <https://doi.org/10.3390/e26050381>

Academic Editor: Jakub Tomczak

Received: 12 March 2024

Revised: 19 April 2024

Accepted: 25 April 2024

Published: 29 April 2024



Copyright: © 2024 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Interestingly, modern research in psychology and neuroscience suggests that there is no sharp distinction between memory and spontaneous imagination [16]. In particular, the theory of reconstructive memory suggests that most sensory aspects of our memories are not encoded but are instead reconstructed from contextual information [16–19]. Memory-related brain areas such as the hippocampus have been shown to respond to imagination, future prediction and counterfactual reasoning tasks [16]. Moreover, hippocampal replays can be seen as a form of spontaneous generation, which is considered to be vital for memory consolidation and planning [20]. Given these deep connections, it is natural to expect that generative models and associative memory models are two faces of the same coin.

In this paper, we will show that this is indeed the case for generative diffusion models, a relatively new class of generative models that have achieved state-of-the-art performance in most computer vision and audio generation tasks [21–23]. Specifically, we show that the asymptotic low-time energy landscape of a large class of generative diffusion models trained on discrete patterns is identical to the asymptotic low-temperature landscape of modern Hopfield networks. Furthermore, in our experiments, we show that this equivalence holds almost exactly in the standard numerical implementations.

Using these results, we offer a new theoretical conceptualization of associative memory that can incorporate semantic, episodic, and reconstructive memory as the result of the action of the same synaptic learning rule.

2. Preliminaries

In this section, we will review the basic theory of associative memory and diffusion modeling. To keep the text readable to a large audience, we will focus on the intuitive aspects and keep the level of mathematical formalism at a minimum. For example, we will write stochastic differential equations in terms of (infinitesimal) update equations. For a more formal treatment of these topics, we recommend the reader to refer to SDE texts such as to Kloeden et al. [24].

2.1. Associative Memory Networks

Hopfield networks have been developed to formalize the concept of associative memory in a simplified artificial neural system [1]. We will denote the activity of the D memory units with a vector $x(t) = (x_1(t), \dots, x_D(t))$. In the original paper, these neural activities were assumed to be binary variables ($x_j(t) \in \{-1, 1\}$), respectively, denoting states of rest and states of firing. The dynamic of a Hopfield network is regulated by the update equation:

$$x_j(t + dt) = [\text{sign}(Wx)]_j, \quad (1)$$

where W is a real-valued symmetric matrix of synaptic couplings (weights) with null diagonal ($W_{jj} = 0, \forall j$). The matrix W encodes pairwise associations between the different components of the pattern vectors. It can be shown that this update rule decreases monotonically the following energy function [1]:

$$u_H(x) = x^T W x, \quad (2)$$

and that it will, therefore, converge to one of its local minima. These minima can then be used to encode memories, which can be retrieved through the Hopfield dynamics when initialized in an incomplete or perturbed version of the memory state. The simplest way to encode memories in the coupling matrix is to use the Hebb's rule of association, which in its simplest form is:

$$W_{j,k} = \sum_{n=1}^N y_{n,j} y_{n,k} \quad (3)$$

where the set of vectors $\{y_1, \dots, y_n, \dots, y_N\}$ represents N "experienced" patterns of neural activity. A pattern is considered to be successfully stored if it is a stable fixed-point of the discrete dynamics. If the patterns are random, it can be proven that the storage capacity of Hopfield nets scales as $D/4 \log_2 D$ [7].

The storage capacity of Hopfield-like networks can be increased by including non-linear mappings $F(\cdot)$ in its energy function [9,10,25]. The general form for the energy of a (discrete) modern Hopfield network can be written as:

$$u_F(\mathbf{x}) = h\left(\sum_{n=1}^N F(\mathbf{x}^T \mathbf{y}_n)\right), \quad (4)$$

where $h(\cdot)$ is an arbitrary differentiable and strictly monotonic function, which does not affect the location and stability of the local minima. This expression reduces to the standard Hopfield energy for $F(x) = x^2$ and $h(x) = x$. However, it is possible to achieve much higher theoretical capacity by using more complex functions. For example, a modern Hopfield network with $F(x) = e^x$ can store up to $2^{D/2}$ binary patterns [10]. These associative networks have been recently generalized to have continuous dynamics. If the activation vector is continuous, the energy function needs to include a regularization term to enforce stability. For example, [26] proposed the use of an energy function of the form:

$$u_{\text{MH}}(\mathbf{x}, \beta) = -\beta^{-1} \log\left(\sum_{n=1}^N e^{\beta \mathbf{x}^T \mathbf{y}_n}\right) + \|\mathbf{x}\|_2^2/2 \quad (5)$$

where β is a positive-valued parameter. We omitted the terms that are additive constant in \mathbf{x} since they do not change the fixed-points and the resulting dynamics. In these models, the variable \mathbf{x} is assumed to be real-valued and the resulting energy-minimization dynamics can be expressed either through update rule or through a system of ordinary differential equations. Moreover, recent work [11] showed that this dynamics can be expressed in terms of biologically plausible binary association between latent neurons, in a way that is similar to the architecture of a restricted Boltzmann machine [27]. Unfortunately, in these models, the numerical values of the patterns directly determine the synaptic weights between latent and observable neurons. This implies that the patterns need to be stored in memory instead of being converted into distributed synaptic patterns. In this sense, modern Hopfield networks offer a model of memory recall but do not provide insight into learning and memory storage in the brain.

2.2. Generative Diffusion Models

Consider a target distribution $\phi(\mathbf{y})$ that is only available through a training dataset $\mathcal{D} = \{\mathbf{y}_1, \dots, \mathbf{y}_N\}$ of independently sampled data-points. Our goal is to learn the structure of the training set in order to generate new samples from $\phi(\mathbf{y})$. To this aim, we first define a noise-injective process that turns the training samples into random noise states. We will then ‘invert’ this process to turn random noise into new samples. To be consistent with the notation used in the continuous Hopfield model, we deviate from the diffusion modeling literature by writing this process in reversed time, with the noiseless data corresponding to a final time T . In this notation, the diffusive dynamics can be determined by the following backward recursive Equation (corresponding to the forward process described in the standard diffusion literature):

$$\mathbf{x}(t - dt) = \mathbf{x}(t) + \sigma\sqrt{dt}\delta(t), \quad (6)$$

where σ determines the standard deviation of the noise injected at time t and $\delta(t)$ follows a standard normal distribution. This is known as the variance exploding equation in the generative diffusion modeling literature [23], which corresponds to a mathematical Brownian motion. We can sample from this process by sampling an initial state \mathbf{y} from the dataset and using it as the final state in the recursion defined by Equation (6). Note that any other Itô diffusion process defined by a stochastic differential equation can equivalently be used to define a generative diffusion model [23], we will cover the general case in

Section 3.2. It can be shown that, if Equation (6) is initialized with the target distribution $\phi(\mathbf{y})$, the ‘inverse’ equation is given by:

$$\mathbf{x}(t + dt) = \mathbf{x}(t) + \sigma^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x}(t)) dt + \sigma \sqrt{dt} \delta(t). \tag{7}$$

where $p_t(\mathbf{x})$ is the marginal distribution of the noise-injection process at time t . In the case of the variance-exploding process, the marginal can be computed analytically and it can be expressed as:

$$p_t(\mathbf{x}) = \mathbb{E}_{\mathbf{y} \sim \phi(\mathbf{y})} \left[\frac{1}{\sqrt{2\pi(T-t)\sigma^2}} e^{-\frac{\|\mathbf{x}-\mathbf{y}\|_2^2}{2(T-t)\sigma^2}} \right]. \tag{8}$$

This formula involves an average over the distribution of the data, which is usually not available in a generative modeling task. However, we can approximate the drift term of the dynamics (i.e., the so-called score $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$) using a parameterized deep network $\mathbf{s}(\mathbf{x}(t), t; W)$ trained using the denoising loss [23,28,29]:

$$\mathcal{L}(W) = \frac{1}{2} \mathbb{E}_{\mathbf{y} \sim \mathcal{D}, t} \left[\mathbb{E}_{\mathbf{x}(t)|\mathbf{y}} \left[\|\delta(t) - \mathbf{s}(\mathbf{x}(t), t; W)\|_2^2 \right] \right] \tag{9}$$

where $\delta(t) = (\mathbf{x}(t) - \mathbf{y})/(\sigma t)$ is the total noise added to the pattern \mathbf{y} up to time t , and t is uniformly sampled over $[0, T]$. The score can then be recovered from the network using the formula:

$$\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) \approx -\sigma^{-1} \mathbf{s}(\mathbf{x}(t), t; W) \tag{10}$$

which becomes exact when the dataset is infinitely large and the loss is minimized globally.

3. The Equivalence between Diffusion Models and Modern Hopfield Networks

We can now show that, when used for storing discrete patterns, the dynamics of generative diffusion models minimizes the energy function of continuous modern Hopfield networks. We will start from the simpler case of variance-exploding diffusion models as their analysis provides all the main insights and results. We will then generalize the result to arbitrary diffusion models defined by a large class of SDEs.

3.1. Variance-Exploding Models

The first step is to formulate the deterministic dynamics of the model as the negative gradient of an energy function: $\mathbf{x}(t + dt) = \mathbf{x}(t) - \nabla_{\mathbf{x}} u(\mathbf{x}, t) dt + \sigma \sqrt{dt} \delta(t)$. As noted in [30], the energy function is:

$$u_{\text{DM}}(\mathbf{x}, t) = -\sigma^2 \log p_t(\mathbf{x}) = -\sigma^2 \log \mathbb{E}_{\mathbf{y} \sim \phi(\mathbf{y})} \left[e^{-\frac{\|\mathbf{x}-\mathbf{y}\|_2^2}{2(T-t)\sigma^2}} \right] + c, \tag{11}$$

where c does not depend on \mathbf{x} and can, therefore, be omitted without affecting the dynamics. In order to establish a link between the diffusion model and Hopfield networks, we can now assume that the data source is a finite collection of N patterns that we wish to store as memories. This led to the energy:

$$u_{\text{DM}}(\mathbf{x}, t) = -\sigma^2 \log \left(\frac{1}{N} \sum_{n=1}^N e^{-\frac{\|\mathbf{x}-\mathbf{y}_n\|_2^2}{2(T-t)\sigma^2}} \right). \tag{12}$$

If we now assume that the patterns are normalized ($\|\mathbf{y}\|_2^2 = 1$), by expanding the square in the exponent and omitting constant additive terms, we obtain:

$$u_{\text{DM}}(\mathbf{x}, t)/\sigma^2 = -\log \left(\sum_{n=1}^N e^{\frac{\mathbf{x}^T \mathbf{y}_n}{(T-t)\sigma^2}} \right) + \frac{\|\mathbf{x}\|_2^2}{(T-t)^2}. \tag{13}$$

Finally, if we define $\beta(t)^{-1} = (T - t) \cdot \sigma^2$ and we multiply both sides by $\beta(t)^{-1}$, we obtain:

$$\beta^{-1}(t)u_{DM}(x, t)/\sigma^2 = -\beta(t)^{-1} \log \left(\sum_{n=1}^N e^{\beta(t)x^T y_n} \right) + \frac{\|x\|_2^2}{2} = u_{MH}(x, \beta(t)), \quad (14)$$

which for a fixed t is identical to the continuous Hopfield network energy in Equation (5), and it, therefore, has the same fixed-point structure at the limit $\beta \rightarrow \infty$. Note that the scaling factor $\beta^{-1}(t)/\sigma^2$ does not change the fixed-points and their stability as it is a positive constant of x . While we derived this result assuming normalized patterns, this is not actually necessary since, for $\beta(t) \rightarrow \infty$, we have that $\beta^{-1}(t)u_{DM}(x, t)/\sigma^2 \sim -\frac{1}{2} \left(x^T y^* + \|y^*\|_2^2 + \|x\|_2^2 \right)$, where y^* is the pattern that maximizes the quadratic form $x^T y + \|y\|_2^2$. As shown by this expression, at this limit the norm y^* only adds an irrelevant constant shift in the energy.

The main differences between the two approaches are (1) in diffusion models, $\beta(t)$ tends toward this divergent limit as part of the denoising dynamics, while the denoising iterations of modern Hopfield networks keep β fixed and (2) in Hopfield networks, the energy function is minimized deterministically, whereas in diffusion models, there is an additional stochastic term. However, these two differences ‘cancel each other out’ since the divergence of $\beta(t)$ leads to the suppression of the stochastic fluctuations and to exact convergence on the same patterns that minimize the modern Hopfield energy for $\beta \rightarrow 0$. As shown in our experiments, there is no meaningful difference as far as β is large enough.

3.2. The General Case

We can now prove the equivalence in a much more general case where the noise injection dynamics follows the equation:

$$x(t - dt) = x(t) + \nabla v(x)dt + \sigma(t)\sqrt{dt}\delta(t), \quad (15)$$

where $v(x)$ is a differentiable scalar potential function and $\sigma(t)$ is a continuous function of time. This form covers most of the equations used in the diffusion modeling literature but it excludes non-conservative dynamics and state-dependent noise models. The generative dynamics corresponding to this more general noise-injection model is given by the equation:

$$x(t + dt) = x(t) + \left(\sigma(t)^2 \nabla_x \log p_t(x(t)) - \nabla v(x) \right) dt + \sigma(t)\sqrt{dt}\delta(t). \quad (16)$$

where the conditional marginal distributions of the stochastic differential equation is given by the expression:

$$p_t(x) = \mathbb{E}_{y \sim \phi(y)} [k(x(t), t; y, T)]. \quad (17)$$

In this formula, the solution kernel $k(x'', t''; x', t')$ gives the conditional probability density of a state x' at t' to be moved to x'' at time t'' under the noise-injecting dynamics. Unfortunately, this solution kernel cannot be expressed analytically in the general case. We define the function $\psi(x, t; y) = \log k(x, t; y, T)$ as the logarithm of the solution kernel, which allows us to express the energy function of the model in the following form:

$$u_{vDM}(x, t) = -\sigma^2(t) \log \left(\sum_{n=1}^N e^{\psi(x, t; y_n)} \right) + v(x) \quad (18)$$

As we showed in the previous section, the equivalence with the Hopfield model can be shown at the limit $t \rightarrow T$, which corresponds to $\beta \rightarrow \infty$. At this limit, the diffusion dynamics of Equation (17) around a pattern y^* is well-approximated by the linearized equation:

$$x(t - dt) - x(t) = \nabla v(y^*)dt + H(y^*)(x - y^*)dt + \sigma(t)\sqrt{dt}\delta(t), \quad (19)$$

with $H(\mathbf{y}^*)$ denoting the Hessian matrix of the potential at \mathbf{y}^* , with $H_{ij}(\mathbf{y}^*) = \frac{\partial^2 u(\mathbf{y}^*)}{\partial x_i \partial x_j}$. Since the equation is linear, its solution kernel is given by a Gaussian density whose mean vector and covariance matrix depend on the potential $v(x)$ and on $\sigma(t)$. For $t \rightarrow T$, the drift term is negligible compared to the diffusion term as the former scales linearly while the latter scales as a squared root. For this reason, the kernel simplifies further into the following asymptotic expression:

$$\psi(\mathbf{x}, t; \mathbf{y}) \sim -\frac{1}{2\sigma^2(T)(T-t)} \|\mathbf{x} - \mathbf{y}\|_2^2 + c. \quad (20)$$

where c does not depend on x . Up to multiplicative and additive constants, this leads to the asymptotic energy function:

$$u_{v\text{DM}}(\mathbf{x}, t) \sim -\log \left(\sum_{n=1}^N e^{-\frac{\|\mathbf{x} - \mathbf{y}_n\|_2^2}{2(T-t)\sigma(T)^2}} \right), \quad (21)$$

which is identical to the energy function in Equation (12). Note that we neglected the “regularization term” $v(x)$, since it does not diverge for $t \rightarrow T$ and it is, therefore, negligible at this limit. From this result, we can conclude that the fixed-point structure of the diffusion model at $t \rightarrow T$ does not depend on the specific form of the SDE and it is therefore fully characterized by the Brownian motion model analyzed in Section 3.1.

4. Encoding Memories by Denoising Neural State

To date, we considered exact diffusion models by analyzing the theoretically optimal score function. However, in real-world applications it is often not possible to compute the score analytically. Instead, as outlined in Section 2.2, the score function is approximated using a denoising deep neural network $s(\mathbf{x}(t), t; W)$ parameterized by a large array W of synaptic weights. Therefore, in real applications the memories are ultimately encoded in the pattern of synaptic weights updated by SGD. Specifically, for a single data-point, the loss function in Equation (9) results to a SGD weight updates of the form:

$$W_{t+1} = W_t - \eta \frac{\partial s}{\partial W} \left(\delta(t) - s(\mathbf{x}, t; W) \right). \quad (22)$$

where $\delta(t)$ can be interpreted as the residual between a past low-noise state \mathbf{y} and a current noisy-state \mathbf{x} . In the brain, we conjecture that the residual might be obtained through very short-term sensory memory, which can buffer recent states and then compare them with the current activity. In psychology and neuroscience, this form of ‘buffering’ is often referred to as iconic memory, which is estimated to last approximately 1000 ms [31–33]. Interestingly, recent studies show that the backpropagation term $\partial s / \partial W$ can also be obtained through noise injection [34], which suggests that noise-injecting processes may play a major role in learning. While the activity of the brain is ongoing and cannot be neatly separated into forward and reversed dynamics, there is tantalizing evidence that the time-evolution of cortical networks oscillate between feed-forward and feed-back phases [35], which corresponds to states with different signal-to-noise ratios [36]. This phenomenon might offer a mechanistic implementation of generative diffusion training in the brain, possibly in relation with the hippocampus theta cycle, which is known to play a crucial role in memory consolidation [37]. Coordination between the frontal cortex is likely to play a central role in the consolidation of memory traces [38,39,39] and in the modulation of the oscillatory cycles [40,41]. This suggests that a denoising training might be carried out through frontal-hippocampal feedback process, where cycles of forward and reversed processes are alternated in order to learn both from real and simulated experiences (replays). However, these are speculative suggestions that need to be tested through detailed biophysical theoretical modeling and experimental studies.

5. Beyond Classical Associative Memories

In the previous section, we showed that the asymptotic energy landscape of generative diffusion models trained on a finite set of discrete patterns matches the energy function of modern Hopfield networks for $\beta \rightarrow \infty$. This implies that, in this regime, the two models have the same fixed-point structure and, consequently, that they have the same memory capacity. Nevertheless, generative diffusion models are more general than associative memory networks in two respects: (1) they provide a framework for probabilistic memory recall and (2) they can be used to encode non-discrete memory structures of arbitrary dimensionality. In this sense, generative diffusion model theory can be seen as a wide generalization of the classical theory of associative memory.

5.1. Probabilistic Recall

At its core, the classical theory of associative memory is centered around the idea of deterministic energy minimization (however, see [42] for related probabilistic methods). On the other hand, the theory of generative diffusion modeling is based on partially random processes. This is not a coincidence since these models have been devised for the purpose of generation and statistical variability is central to the generative task. However, probability theory is also central to the problem of memory retrieval, since in general there are many possible patterns \mathbf{y} that could have generated a given corrupted/incomplete state \mathbf{x} . This is particularly central when the level of corruption is high (i.e., for large values of $(T - t)$), since in this regime it is often impossible to distinguish between several possible patterns and the best hope is to cover their (posterior) probability distribution. When initialized in one of these partially corrupted states, the stochastic generative dynamics of diffusion models can be interpreted in this fashion, with each denoising trajectory eventually reaching one of the possible fixed-points that are compatible with the initial state. Probabilistic recall can be used to encode Bayesian uncertainty, with several possible memories being simultaneously co-activated by parallel denoising processes. This form of encoding can be used to solve information gathering, exploration and other meta-learning problems [43] and has been suggested to be central to the functioning of the mammalian hippocampus [44].

5.2. Higher-Dimensional Memory Structures

In the classical associative memory literature, a memory is encoded as a single (fixed-) point in a m -dimensional space of possible neural activities. In modern mathematics, a set of points is often interpreted as a 0-dimensional space as there are no local degrees of freedom, meaning that it is impossible to locally perturb a point while remaining inside its ‘space’. The energy landscape and associated vector field of this form of memories is visualized in Figure 1a. This presents a very strict definition of memory that is incompatible with its cognitive reality in most biological and artificial forms of intelligence. In fact, even the sharpest human memory only encodes a minority of the details of the original neural state, with many of these details being reconstructed during recall based on contextual information [19]. This suggests that a more realistic description of memory should include “internal” degrees of freedom that allow for partial encoding. In mathematical terms, this can be done by defining a memory as a connected d -dimensional sub-space embedded in the larger m -dimensional space of possible neural activities. A diffusion model trained on N of these “extended memories”, each represented by a sub-space \mathcal{S}_n has the energy function:

$$\log \left(\sum_{n=1}^N \int_{\mathcal{S}_n} e^{\beta(t)x^T \mathbf{y}_n} d\mathbf{y}_n \right), \quad (23)$$

where the integral is taken over the whole sub-space \mathcal{S}_n , whose points correspond equally valid interpretations of the same memory. An example of an energy landscape for this form of extended localized memory is given in Figure 1b. In general, each sub-space \mathcal{S}_n

can potentially have a different dimensionality or even be a more complex geometrical structure with fractional (fractal) or variable dimensionality.

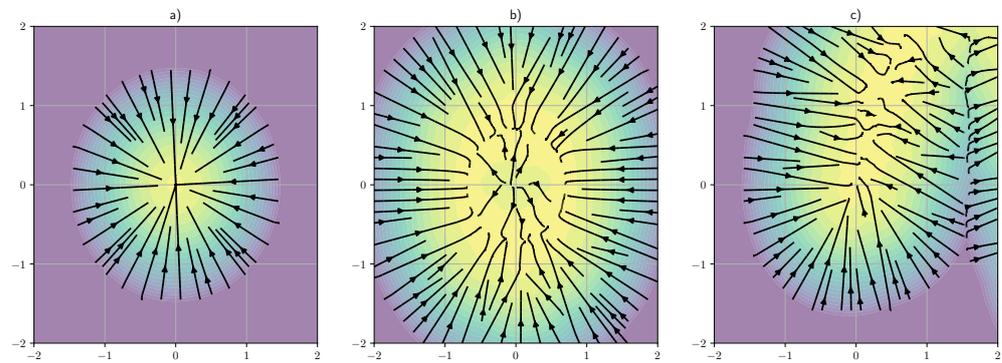


Figure 1. Visualization of different kinds of energy landscape and gradient vector fields corresponding to different forms of memory (in a two-dimensional space): (a) classical point-like memory; (b) extended localized memory; (c) non-localized (semantic) memory structure. The color denotes the probability density of the learned distribution while the lines represent the integral trajectories of the vector field induced by the score function.

6. A Theoretical Framework for Memory in Biological and Artificial Intelligent Systems

In this section, we will leverage the connection we established between associative memory and diffusion models in order to outline a theoretical framework for biologically plausible memory. Our goal is both to provide conceptual tools to theoretical and experimental neuroscientists and to promote developments in naturalistic machine intelligence systems.

6.1. Semantic, Episodic, and Reconstructive Memory

In psychology and neuroscience, the semantic (or structural) memory system is thought to learn the general structure of the sensory input, discarding the idiosyncratic details of individual events. This can be modeled with an energy function where the sum over patterns is replaced by an integral over a continuous density $\phi(\mathbf{y})$:

$$\sigma^2 \log \int e^{\beta(t)\mathbf{x}^T \mathbf{y}} \phi(\mathbf{y}) d\mathbf{y}. \tag{24}$$

In practice, the distribution $\phi(\mathbf{y})$ is often defined on a manifold or some other lower dimensional structure, leading to dynamics similar to what is visualized in Figure 1c. This is exactly the kind of behavior we expect from a generative model, initial perturbed states are gradually pushed towards a point on the manifold of possible patterns. In a sense, it can also be seen as a probabilistic generalization of the high dimensional memory structures described in Equation (23), where the sub-space \mathcal{S}_n is no longer localized in a small sub-region of the space of neural states. In cognitive terms, roughly speaking, we can say that a memory is “episodic” if its sub-space is localized, whereas a memory is considered semantic or structural if it is part of a large, non-localized sub-space.

However, modern memory research suggests that real-life episodic memories are thought to be largely reconstructive, meaning that most of the sensory details are re-created during recall based on contextual information [18,45]. For example, the memory of a car crash may evoke the memory of broken glass, although the windshield was not actually broken during the real event. This suggests that human episodic memory has a stored lower-dimensional “representational core” that does not fully constrains the dynamics of the system. This can be formalized using a mixture of discrete (or localized) and continuous distributions:

$$\log \left(\sum_n^N e^{\beta(t)f_n(x)^T \xi_n} + \int e^{-\beta(t)\mathbf{x}^T \mathbf{y}} \phi(\mathbf{y}) d\mathbf{y} \right), \tag{25}$$

where the function $f_n : \mathbb{R}^D \rightarrow \mathbb{R}^W$ with $W < D$ is a lower dimensional encoding of the state and $\zeta^n \in \mathbb{R}^W$ is a stored lower-dimensional pattern that encode the “core” of the memory. Since $W < D$, this energy will only constrain W degrees of freedom to converge to the pattern, while the other degrees of freedom are left free to evolve under the dynamics determined by the corpus of semantic memory. This leads to a form of memory recall where some features are recovered faithfully while others are reconstructed based on learned contextual associations.

6.2. Consolidation and Replays

All the different forms of memory outlined in the previous section can be learned through the same synaptic update rule given in Equation (22). In this sense, our model provides a possible unification for most of the known-form of long-term memory in humans and other animals. During training, the only difference between learning a density $\phi(\mathbf{y})$ or a discrete set of patterns is that in the former case each pattern \mathbf{y} is (almost surely) sampled only once while in the latter case each pattern is re-sampled with finite probability throughout training. In generative machine learning, re-sampling is often performed in practice to maximize data efficiency, although it can lead to overfitting since the generative model could memorize the patterns themselves instead of extrapolating the underlying density $\phi(\mathbf{y})$.

In biological systems, each data point can only be experienced once and this would likely not result in the formation of localized episodic memories, since their encoding requires some form of re-sampling. However, it is well known that the formation of episodic memories in humans requires an extensive phase of consolidation that depends on the activity of the hippocampus. Under our theoretical framework, a possible explanation is that new simulated experiences are generated in the hippocampal network and are then used for synaptic training. In fact, it is well known that these forms of replays can be observed in the hippocampus both during sleep and wakefulness and that their disruption compromises memory consolidation [46]. Training on these self-generated samples can result in a positive feedback, with each replay increasing the probability of the same event being re-sampled in the future. Through this process, a single event can be ‘bootstrapped’ into a self-reinforcing process of activity that can eventually form a localized episodic memory. Furthermore, this would result in the formation of a low-dimensional episodic “core”, as described in the previous sub-section, if only a sub-set of features are re-sampled faithfully, with the other being generated from structural memory.

7. Experiments

Figure 2 visualizes the qualitative behavior of both learned and exact score models compared with modern Hopfield iterations. We consider a five-dimensional associative memory problem with four randomly sampled binary patterns. For the forward diffusion dynamics, we used a variance-preserving model as they are numerically more stable and more widely used [23]. We considered both exact score models and trained diffusion models parameterized by three layer perceptrons $s(x, t; W)$ with 750 hidden units in each hidden layer and ReLu activation functions in both hidden layers. The time variable t was embedded by appending the value $\theta_t = e^{-t/2}$ to the activations of each layer, including the input, hidden, and output layers. This means that each weight matrix acted on a $n + 1$ -dimensional feature space, where n is the number of activations. We used the training approach described in [23] but with a simpler constant noise schedule ($\sigma = 1$). The noise-injection model is given by the equation:

$$x(t - dt) = (1 - \alpha dt)x(t) + \sigma\sqrt{dt}\delta(t), \quad (26)$$

where $\alpha = 0.5$ and $\delta(t) \sim \mathcal{N}(0, I)$. Accordingly, the network was trained using the denoiser autoencoder loss:

$$\mathcal{L}(W) = \mathbb{E}_{\theta_t \sim \mathcal{U}(0.6, 1), \delta \sim \mathcal{N}(0, I)} \left[\sum_n \left\| \delta - s \left(\theta_t \mathbf{y}_n + \sqrt{1 - \theta_t^2} \delta, t; W \right) \right\|_2^2 \right]. \quad (27)$$

Using this loss, the network is trained to recover the added noise from noise-corrupted patterns. The loss was optimized using Adam with fixed batches including all patterns to be memorized. The networks were trained for 5000 iterations. Diffusion trajectories were integrated using a simple Euler approach for the deterministic ODE samples (see [23]):

$$\mathbf{x}(t + dt) = \mathbf{x}(t) + \frac{1}{2} \left(\alpha \mathbf{x}(t) - \frac{s(\mathbf{x}(t), t; W_{\text{trained}})}{\sqrt{1 - e^{-2\alpha t}}} \right) dt, \quad (28)$$

while we used a Euler–Maruyama method for the SDE sampler:

$$\mathbf{x}(t + dt) = \alpha \mathbf{x}(t) - \frac{s(\mathbf{x}(t), t; W_{\text{trained}})}{\sqrt{1 - e^{-2\alpha t}}} dt + \sigma \sqrt{dt} \delta(t). \quad (29)$$

Both samplers used 300 time steps discretized linearly from $t_{\text{star}} = 10^{-3}$ to a t corresponding to the noise level. The modern Hopfield iterations were implemented as specified in [26], with $\beta = 5$ and 4 updates. Consistent with our analysis, all algorithms approximately converge to the same target points. In line with previous work, the modern Hopfield iteration converges after a single iteration, whereas the diffusion trajectories converge smoothly to the target. However, note that, in this setting, diffusion denoising can also be performed in one step simply by re-scaling the score.

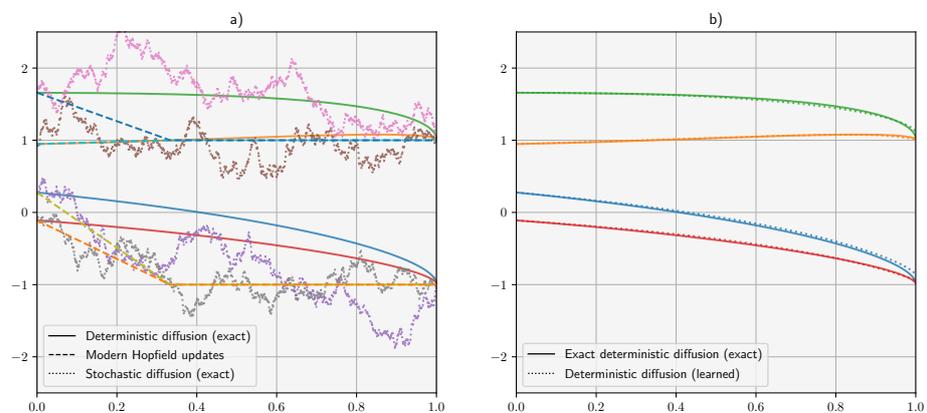


Figure 2. Qualitative analysis of the (marginal) denoising trajectories of a binary associative memory problem with four patterns in a five-dimensional space. (a) Comparison between denoising trajectories of diffusion models and modern Hopfield updates. The diffusion curves are integrated using the Euler method with 2000 steps. The trajectories are overlaid to four modern Hopfield updates. (b) Comparison between exact and learned deterministic denoising trajectories. The colors are used to identify individual trajectories.

As a first quantitative analysis, we evaluated the Pearson correlation coefficient between the output of modern Hopfield iteration and (a) a diffusion model, (b) a classical Hopfield network, and (c) the ground truth pattern. For a given dimensionality d , n binary patterns \mathbf{y} were randomly generated, and subsequently, corrupted with noise using the formula $\tilde{\mathbf{y}} = \theta \mathbf{y} + \sqrt{1 - \theta^2} \epsilon$, where ϵ is a standard Gaussian noise vector. We used a noise level of $\theta = 0.68$. We kept the dimensionality equal to 10 and we evaluated the correlation for 10, 20, and 30 stored patterns. The simulation was repeated 100 times in order to reliably compute the correlations. In order to avoid to have to re-train a neural model hundreds of times, for the diffusion models we used the exact score formula (see Appendix A). For the

modern Hopfield model, we used 150 updates in order to maximize performance. The final output of all methods was binarized using the sign function. Table 1 shows the estimated Pearson correlations between methods. As expected from our analysis, the correlation between the modern Hopfield iterations and the diffusion model is extremely close to one.

Table 1. Pearson correlation between output of modern Hopfield network and other models (plus ground truth pattern) in both denoising and completion experiments.

Denoising Task	Number of Patterns	Diffusion Model	Classic Hopfield	True Patterns
	10	0.995	0.732	0.893
	20	0.991	0.704	0.822
	30	0.991	0.715	0.81
Completion task				
	10	0.996	0.741	0.897
	20	0.991	0.707	0.838
	30	0.989	0.700	0.795

Average Pearson correlations between the output of (exact) diffusion models, modern Hopfield networks and classical Hopfield networks in denoising and completion tasks performed on random binary patterns.

We also performed the same experiments in a completion task, where the patterns were partially zero-masked instead of being corrupted by white noise. The binary masks were sampled randomly from a Bernoulli distribution with $p = 0.5$. Again, the output of the exact diffusion models correlates almost perfectly with the output of the modern Hopfield iterations.

Next, we estimated the error and capacity of the models. The corrupted patterns were fed to the algorithms and the results were compared with the original pattern using the Hamming error. The patterns were considered to be correctly recovered if the error was smaller than 3%. Figure 3a shows the error of an exact diffusion model for different numbers of patterns as function of the dimensionality. For a given noise level and threshold, the capacity was defined as the maximum number of patterns that can on average be recovered. Figure 3b shows the estimated capacity of the exact diffusion model (blue), modern Hopfield network (green), classical Hopfield network (red), and trained diffusion model (black dots). The details of the experiments are given in Appendix A.

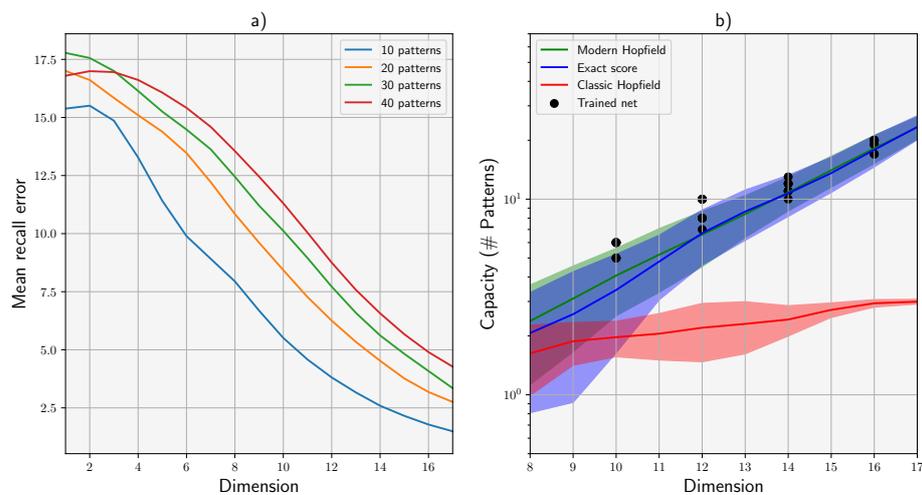


Figure 3. (a) Median error of exact diffusion model as function of the dimensionality. (b) Capacity of diffusion models and Hopfield networks in log scale. The shaded area denotes the estimated 95% intervals.

8. Discussion

In this paper, we demonstrated that a popular class of modern Hopfield networks with exponential non-linearities is mathematically equivalent to a large class of continuous diffusion models at the limit of $\beta \rightarrow \infty$. In our experiments, we showed that this equivalence holds almost exactly for a finite β value and for the more stable variance-preserving models, both with exact and trained score. The equivalence depends on the fact that the diffusion models are trained on a finite number of discrete patterns, and in fact the diffusion models can generalize the modern Hopfield energy function to a setting where both episodic memory and semantic memory (i.e., generative manifolds) are jointly encoded by the dynamics of the same network. From the point of view of theoretical neuroscience, this may be used to model the different forms of long-term memory as a result of the same learning mechanism, in a single distributed neural system.

While generative diffusion models offer an attractive paradigm for modeling memory, imagery, and even perception in the brain, more work needs to be done in order for its components to be implemented in a biologically plausible way. In particular, it is unlikely that biological neural networks implement pure noise-injection dynamics. However, the mathematics of generative diffusion models can be written in term of any stochastic differential Equation [23], which can be used to implement more plausible transformation such as, for example, the feedforward perceptual feature as implemented in the hierarchy of sensory cortical areas [47]. More fundamentally, it is unclear how the reverse and forward dynamics can be simultaneously implemented in the brain networks. However, there is tantalizing evidence that this could correspond to the theta rhythm driving periodic cholinergic modulations [37].

The denoising loss given by Equation (9) is somewhat similar to the self-prediction errors used in predictive coding models [48–50]. This opens the door for potentially fruitful connections with existing predictive models of memory [51,52]. Particularly interesting is the recent use of generative predictive coding models for associative memory storage [53], which is also based on a generative model and has some similarities with the approach discussed here.

The main biological issue with episodic memory as conceptualized in this paper is that it seems to require the re-sampling of the same events, while in the real world, each event can only happen once. This can be potentially implemented with some form of bootstrap re-sampling, which might correspond to the replays observed in the hippocampus [20].

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Data Availability Statement: Data are contained within the article.

Acknowledgments: We would like to thank Gabriel Raya for reviewing the manuscript and providing valuable feedback.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Details of the Experiments

Binary patterns were generated independently by applying the sign function to standard Gaussian vectors. To test recovery, we corrupted the patterns using the formula $\tilde{\mathbf{y}} = \theta \mathbf{y} + \sqrt{1 - \theta^2} \delta$, where δ is a standard Gaussian vector. In all the experiments, we used $\theta = 0.68$.

We compared the storage and recovery performance of several models. Modern Hopfield networks were implemented with the iterative update given in [26]. We set β to be equal to 5 and we applied 150 iterations to each noisy pattern.

For reasons of stability, we used variance-preserving diffusion models defined by the following noise-injection SDE:

$$d\mathbf{x}_t = -\gamma \mathbf{x}_t + \gamma dW_t, \quad (\text{A1})$$

where W_t is a standard Wiener process. This is slightly different from the variance-exploding models discussed in the main text, but it delivers nearly identical results. The diffusion models were used as denoiser by applying the deterministic ODE dynamics given in [23]:

$$\frac{dx_t}{dt} = \frac{1}{2} \left(x_t + \sigma^2 \nabla_x \log p_t(x(t)) \right), \quad (\text{A2})$$

which exactly reproduces the marginal densities of the stochastic dynamics. We integrated the dynamics using 300 steps of Euler integration. The initial time was set to match the match level using the following formula: $t_{\text{start}} = -\gamma^{-1} 2 \log(\theta)$, with $\gamma = 0.8$. This formula implies the use of a constant noise scheduling with variance equal to γ .

In an exact-score-based diffusion model, we used the exact formula for the score given the marginal density:

$$\nabla \log p_t(x) = (x - Yh_t(x)) / (1 - \theta_t^2), \quad (\text{A3})$$

where Y is a matrix having the patterns y^j as columns. The weight vector h_t is obtained using the softmax function:

$$h_{t,j}(x) = \text{softmax} \left(\dots, -\|x - y^k\|_2^2 / (2(1 - \theta_t^2)), \dots \right)_j, \quad (\text{A4})$$

which depends on the correlation between the state x and each of the patterns. This is very similar to the softmax formula given in the update of modern Hopfield networks, which is unsurprising given the equivalence of their energy functions.

Learned generative diffusion models used a three layers fully connected architecture with reLu non-linearities, d input and output units and $80d$ hidden units in each layer. The time index was embedded by converting it to $\theta_t = \exp(-0.5\gamma t)$ and then by concatenating this value to each layer. They were trained using the optimizer Adam with base rate 0.001 and fixed batches containing all the generated patterns.

For the Hopfield and exact models, the capacity was estimated by evaluating the average reconstruction error in a grid of dimensionalities and number of patterns. Given each combination, the error was computed using 140 randomly sampled patterns. The capacity was then estimated by finding the maximum number of patterns such that, for a given dimensionality, the average Hamming error did not exceed a threshold value of 3%. Error intervals were obtained by bootstrap re-sampling of the error. In order to increase the snd, we smoother both the mean capacity and the bounds by convolving them with a Gaussian kernel with $\sigma = 1.5$.

This estimation method would have been too expensive for evaluating the capacity of learned models, since the network needs to be fully re-trained for any given set of patterns.

Instead, for a given dimensionality d , we trained the network with n_d patterns, where n_d is the estimated capacity of the exact score model plus 4. After training, we evaluated the error on 30 batches of unseen noise corrupted versions of the training patterns. If the error was below threshold, we returned n_d as the estimated capacity, otherwise, we reduced n_d by one and repeated the training until the error was below threshold. This procedure was repeated 8 times for d equal to 10, 12, 14, and 16.

References

1. Hopfield, J.J. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. USA* **1982**, *79*, 2554–2558. [[CrossRef](#)]
2. Michel, A.N.; Farrell, J.A. Associative memories via artificial neural networks. *IEEE Control Syst. Mag.* **1990**, *10*, 6–17. [[CrossRef](#)]
3. Graves, A.; Wayne, G.; Reynolds, M.; Harley, T.; Danihelka, I.; Grabska-Barwińska, A.; Colmenarejo, S.G.; Grefenstette, E.; Ramalho, T.; Agapiou, J.; et al. Hybrid computing using a neural network with dynamic external memory. *Nature* **2016**, *538*, 471–476. [[CrossRef](#)]
4. Lopez-Paz, D.; Ranzato, M. Gradient episodic memory for continual learning. *Adv. Neural Inf. Process. Syst.* **2017**, *30*. [[CrossRef](#)]
5. Guo, Y.; Liu, M.; Yang, T.; Rosing, T. Improved schemes for episodic memory-based lifelong learning. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 1023–1035.

6. Kadam, S.; Vaidya, V. Review and analysis of zero, one and few shot learning approaches. In *Intelligent Systems Design and Applications, Proceedings of the 18th International Conference on Intelligent Systems Design and Applications (ISDA 2018), Vellore, India, 6–8 December 2018*; Springer International Publishing: Berlin/Heidelberg, Germany, 2020.
7. Abu-Mostafa, Y.; Jacques, J.S. Information capacity of the Hopfield model. *IEEE Trans. Inf. Theory* **1985**, *31*, 461–464. [[CrossRef](#)]
8. Sejnowski, T.J.; Tesauro, G. The Hebb rule for synaptic plasticity: Algorithms and implementations. In *Neural Models of Plasticity*; Elsevier: Amsterdam, The Netherlands, 1989; pp. 94–103.
9. Krotov, D.; Hopfield, J.J. Dense associative memory for pattern recognition. *Adv. Neural Inf. Process. Syst.* **2016**, *29*. [[CrossRef](#)]
10. Demircigil, M.; Heusel, J.; Löwe, M.; Upgang, S.; Vermet, F. On a model of associative memory with huge storage capacity. *J. Stat. Phys.* **2017**, *168*, 288–299. [[CrossRef](#)]
11. Krotov, D.; Hopfield, J. Large associative memory problem in neurobiology and machine learning. In *Proceedings of the International Conference on Learning Representations, Virtual Event, Austria, 3–7 May 2021*.
12. Friston, K.; Kilner, J.; Harrison, L. A free energy principle for the brain. *J. Physiol.* **2006**, *100*, 70–87. [[CrossRef](#)]
13. Friston, K. The free-energy principle: A unified brain theory? *Nat. Rev. Neurosci.* **2010**, *11*, 127–138. [[CrossRef](#)]
14. Ororbias, A.; Kifer, D. The neural coding framework for learning generative models. *Nat. Commun.* **2022**, *13*, 2064. [[CrossRef](#)]
15. Jones, M.; Wilkinson, S. *From Prediction to Imagination*; Cambridge University Press: Cambridge, MA, USA, 2020.
16. Schacter, D.L.; Addis, D.R.; Buckner, R.L. Remembering the past to imagine the future: The prospective brain. *Nat. Rev. Neurosci.* **2007**, *8*, 657–661. [[CrossRef](#)]
17. Kolodner, J.L. Reconstructive memory: A computer model. *Cogn. Sci.* **1983**, *7*, 281–328.
18. Hemmer, P.; Steyvers, M. A Bayesian account of reconstructive memory. *Top. Cogn. Sci.* **2009**, *1*, 189–202. [[CrossRef](#)] [[PubMed](#)]
19. Nash, R.A.; Wheeler, R.L.; Hope, L. On the persuadability of memory: Is changing people’s memories no more than changing their minds? *Br. J. Psychol.* **2015**, *106*, 308–326. [[CrossRef](#)]
20. Buhry, L.; Azizi, A.H.; Cheng, S. Reactivation, replay, and preplay: How it might all fit together. *Neural Plast.* **2011**, *2011*, 203462. [[CrossRef](#)] [[PubMed](#)]
21. Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; Ganguli, S. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proceedings of the International Conference on Machine Learning, Lille, France, 6–11 July 2015*.
22. Ho, J.; Jain, A.; Abbeel, P. Denoising diffusion probabilistic models. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 6840–6851.
23. Song, Y.; Sohl-Dickstein, J.; Kingma, D.P.; Kumar, A.; Ermon, S.; Poole, B. Score-Based Generative Modeling through Stochastic Differential Equations. In *Proceedings of the International Conference on Learning Representations, Virtual, 3–7 May 2021*.
24. Kloeden, P.E.; Platen, E.; Kloeden, P.E.; Platen, E. *Stochastic Differential Equations*; Springer: Berlin/Heidelberg, Germany, 1992.
25. Krotov, D. A new frontier for Hopfield networks. *Nat. Rev. Phys.* **2023**, *5*, 366–367. [[CrossRef](#)]
26. Ramsauer, H.; Schäfl, B.; Lehner, J.; Seidl, P.; Widrich, M.; Adler, T.; Gruber, L.; Holzleitner, M.; Pavlović, M.; Sandve, G.K.; et al. Hopfield networks is all you need. In *Proceedings of the International Conference on Learning Representations, Virtual, 3–7 May 2021*.
27. Fischer, A.; Igel, C. An introduction to restricted Boltzmann machines. In *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications, Proceedings of the 17th Iberoamerican Congress, CIARP, Buenos Aires, Argentina, 3–6 September 2012*; Springer: Berlin/Heidelberg, Germany, 2012.
28. Hyvärinen, A. Estimation of Non-Normalized Statistical Models by Score Matching. *J. Mach. Learn. Res.* **2005**, *6*, 695–709.
29. Vincent, P. A connection between score matching and denoising autoencoders. *Neural Comput.* **2011**, *23*, 1661–1674. [[CrossRef](#)]
30. Raya, G.; Ambrogioni, L. Spontaneous symmetry breaking in generative diffusion models. *Adv. Neural Inf. Process. Syst.* **2024**, *36*. [[CrossRef](#)]
31. Di Lollo, V. Temporal characteristics of iconic memory. *Nature* **1977**, *267*, 241–243. [[CrossRef](#)] [[PubMed](#)]
32. Sakitt, B. Iconic memory. *Psychol. Rev.* **1976**, *83*, 257. [[CrossRef](#)] [[PubMed](#)]
33. Graziano, M.; Sigman, M. The dynamics of sensory buffers: Geometric, spatial, and experience-dependent shaping of iconic memory. *J. Vis.* **2008**, *8*, 9. [[CrossRef](#)] [[PubMed](#)]
34. Dalm, S.; van Gerven, M.; Ahmad, N. Effective Learning with Node Perturbation in Deep Neural Networks. *arXiv* **2023**, arXiv:2310.00965.
35. Hasselmo, M.E.; Cekic, M. Suppression of synaptic transmission may allow combination of associative feedback and self-organizing feedforward connections in the neocortex. *Behav. Brain Res.* **1996**, *79*, 153–161. [[CrossRef](#)] [[PubMed](#)]
36. Hasselmo, M.E.; Linster, C.; Patil, M.; Ma, D.; Cekic, M. Noradrenergic suppression of synaptic transmission may influence cortical signal-to-noise ratio. *J. Neurophysiol.* **1997**, *77*, 3326–3339. [[CrossRef](#)] [[PubMed](#)]
37. Hasselmo, M.E. Septal modulation of hippocampal dynamics: What is the function of the theta rhythm? In *The Behavioral Neuroscience of the Septal Region*; Springer: Berlin/Heidelberg, Germany, 2000; pp. 92–114.
38. Nieuwenhuis, I.L.; Takashima, A. The role of the ventromedial prefrontal cortex in memory consolidation. *Behav. Brain Res.* **2011**, *218*, 325–334. [[CrossRef](#)]
39. Tronel, S.; Feenstra, M.G.; Sara, S.J. Noradrenergic action in prefrontal cortex in the late stage of memory consolidation. *Learn. Mem.* **2004**, *11*, 453–458. [[CrossRef](#)]
40. Benchenane, K.; Tiesinga, P.H.; Battaglia, F.P. Oscillations in the prefrontal cortex: A gateway to memory and attention. *Curr. Opin. Neurobiol.* **2011**, *21*, 475–485. [[CrossRef](#)]

41. Comeaux, P.; Clark, K.; Noudoost, B. A recruitment through coherence theory of working memory. *Prog. Neurobiol.* **2023**, *228*, 102491. [[CrossRef](#)]
42. Hancock, E.R.; Kittler, J. A Bayesian interpretation for the Hopfield network. In Proceedings of the IEEE International Conference on Neural Networks, San Francisco, CA, USA, 28 March–1 April 1993.
43. Gupta, A.; Mendonca, R.; Liu, Y.; Abbeel, P.; Levine, S. Meta-reinforcement learning of structured exploration strategies. *Adv. Neural Inf. Process. Syst.* **2018**, *31*.
44. Ambrogioni, L.; Ólafsdóttir, H.F. Rethinking the hippocampal cognitive map as a meta-learning computational module. *Trends Cogn. Sci.* **2023**, *27*, 702–712. [[CrossRef](#)]
45. Roediger, H.L.; McDermott, K.B. Creating false memories: Remembering words not presented in lists. *J. Exp. Psychol. Learn. Mem. Cogn.* **1995**, *21*, 803. [[CrossRef](#)]
46. Ólafsdóttir, H.F.; Bush, D.; Caswell, B. The role of hippocampal replay in memory and planning. *Curr. Biol.* **2018**, *28*, R37–R50. [[CrossRef](#)]
47. Güçlü, U.; van Gerven, M.A.J. Deep neural networks reveal a gradient in the complexity of neural representations across the ventral stream. *J. Neurosci.* **2015**, *35*, 10005–10014. [[CrossRef](#)]
48. Rao, R.P.N.; Ballard, D.H. Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* **1999**, *2*, 79–87. [[CrossRef](#)]
49. Spratling, M.W. A review of predictive coding algorithms. *Brain Cogn.* **2017**, *112*, 92–97. [[CrossRef](#)] [[PubMed](#)]
50. Millidge, B.; Seth, A.; Buckley, C.L. Predictive coding: A theoretical and experimental review. *arXiv* **2021**, arXiv:2107.12979.
51. Stachenfeld, K.L.; Botvinick, M.M.; Gershman, S.J. The hippocampus as a predictive map. *Nat. Neurosci.* **2017**, *20*, 1643–1653. [[CrossRef](#)]
52. Barron, H.C.; Auzsztulewicz, R.; Friston, K. Prediction and memory: A predictive coding account. *Prog. Neurobiol.* **2020**, *192*, 101821. [[CrossRef](#)] [[PubMed](#)]
53. Salvatori, T.; Song, Y.; Hong, Y.; Sha, L.; Frieder, S.; Xu, Z.; Bogacz, R.; Lukasiewicz, T. Associative memories via predictive coding. *Adv. Neural Inf. Process. Syst.* **2021**, *34*, 3874–3886. [[PubMed](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.