

On the Statistical Mechanics of Attention

Jerry Yao-Chieh Hu^{†‡*1} Yu-An Chen^{§*2} Maojiang Su^{†3} Han Liu^{†‡4}

[†] Center for Foundation Models and Generative AI, Northwestern University

Department of Computer Science, Northwestern University

[‡] Department of Computer Science, College of Connected Computing, Vanderbilt University

[§] Department of Physics, National Taiwan University

[‡] Department of Statistics and Data Science, Northwestern University

We revisit the statistical mechanics of Transformer attention via modern Hopfield models. We present a free-energy formalism to characterize parameter regimes and critical thresholds for robust attention inference. The resulting free-energy landscape incorporates randomness in query states and attention weights. Attention then emerges as the equation of motion of the corresponding statistical-mechanical system (i.e., modern Hopfield model) through the standard free-energy minimization principle. This allows us to identify first-order phase transitions in attention between glassy and robust-inference phases. We further extend to attentions with stochastic queries and keys. Together, our results provide a rigorous statistical-mechanics framework for characterizing phase structure and inference reliability across a family of Transformer attention mechanisms.

¹ jhu@northwestern.edu

² b09202059@ntu.edu.tw

³ smj@northwestern.edu

⁴ hanliu@northwestern.edu

*These authors contributed equally to this work.

Contents

1	Introduction	2
2	Preliminaries	4
2.1	Statistical Mechanics and Energy Formalism	4
2.2	Connection between Free Energy and Tsallis Entropy	6
3	Statistical Physics of Attention: Modern Hopfield Models	7
3.1	Modern Hopfield Model (MHM)	7
3.2	Retrieval Dynamics and Attention Mechanism	10
4	Phase Transitions and Critical Points of Retrieval Success	11
5	Numerical Results	12
5.1	Temperature Dependence	13
5.2	Memory-Load Dependence	15
6	Discussion and Conclusion	18
A	Related Work	21
B	A Minimal Introduction of Statistical Mechanics	22
B.1	From Interacting Spins to Associative Memory	22
B.2	A Concrete Hopfield Free Energy	23
B.3	Landscapes, Metastability, and Retrieval Jumps	25
C	Definitions and Proofs in the Main Text	26
C.1	Basic Tools for Convex Analysis	26
C.2	Proof of Proposition 2.1	27
C.3	Tsallis Entropy	29
C.4	Proof of Theorem 2.1	33
C.5	Proof of Proposition 3.1	35
C.6	Proof of Theorem 3.1	35
C.7	Proof of Proposition 3.3	37
C.8	Proof of Proposition 4.1	39
C.9	Proof of Proposition 4.2	40
D	Fluctuated Version and Thermal Attention	43
D.1	Physics of Temperature	43
D.2	Fluctuated Variables in Modern Hopfield Model	46
D.3	Connection between Free Energy and Tsallis Entropy in Fluctuated Case	47
D.4	Retrieval Dynamics in Fluctuated Case and Thermal Attention	49
E	Dynamical Retrieval Stability Analysis	56
F	Experiment Setup	61
F.1	Effective Support Size and Sparsity	61
F.2	Settings	63
G	Additional Experiments	66
G.1	Dynamical Stability Analysis for Temperature	66
G.2	Support and Capacity-Reliability Trade-off	68
G.3	Support, Temperature, and Reliability	72

1 Introduction

We study the statistical-mechanics of Transformer attention through modern Hopfield models and characterize parameter regimes and critical threshold for robust attention inference. Our key contribution is a statistical-mechanical formulation in which Transformer attention emerges from the standard free-energy minimization principle naturally. Through this perspective, we give a clear characterization of first-order phase transitions of attention inference between glassy and robust-inference phases. Consequently, our framework provides a rigorous way to characterize reliable inference of Transformer attention, and further allows extension to the full attention family.

During a Transformer forward pass, an input query selects and combines information from the entire context. Attention mechanism captures the core function of this operation. Concretely, for a query $q \in \mathbb{R}^{d_k}$ and stored key-value pairs $\{k_\mu, v_\mu\}_{\mu=1}^M$, we denote K and V as the matrices with rows $k_\mu \in \mathbb{R}^{d_k}$ and $v_\mu \in \mathbb{R}^{d_v}$, respectively. Then attention [Vaswani et al., 2017] computes

$$\text{Attn}(q, K, V) = \text{Softmax}\left(\beta \frac{qK^\top}{\sqrt{d_k}}\right)V = \sum_{\mu=1}^M \left(\frac{e^{\beta z_\mu}}{\sum_{\nu=1}^M e^{\beta z_\nu}}\right)v_\mu, \quad \text{with} \quad z_\mu := \frac{qk_\mu^\top}{\sqrt{d_k}}.$$

Here β is the inverse-temperature. $\text{Attn}(q, K, V)$ specifies which key k_μ does the input query q retrieve and contribute to the output via the value projection V .

Recent works link this *inference* of Transformer attention to the memory *retrieval* of modern Hopfield models [Ramsauer et al., 2020, Krotov and Hopfield, 2020, Hu et al., 2023], i.e., selecting or combining stored patterns based on the input query. This connection is decomposable into two aspects:

- First, Hu et al. [2023], Wu et al. [2023] bridge between modern Hopfield models and attention family including softmax, sparsemax, and beyond [Martins and Astudillo, 2016].
- On the other hand, Hu et al. [2023] reveal that softmax solves an entropy-regularized optimization problem via variation [Peters et al., 2019]; see [Santos et al., 2025] for a review.

Together, they view Transformer attentions as physical models and lead to useful methods [Widrich et al., 2020, Hoover et al., 2023, Auer et al., 2023, Hu et al., 2024a, Wu et al., 2024, Dehmamy et al., 2025]. However, the prior work only investigate the retrieval rule itself. The parameter regimes for robust inference in these models remains unclear. This raises a natural question in physics, with beneficial implications for Transformer inference.

In statistical mechanics, phase structure characterizes how a system’s behavior changes as its parameters vary. Amit et al. [1985, 1987], Nishimori [2001] analyze the phase structure of classical Hopfield models [Hopfield, 1982] via mean-field replica method. However, they consider i.i.d. and infinite-dimension assumptions. That is, how finite size and realistic structured patterns affect retrieval behavior remains open. Modern Hopfield models offer a natural setting for extending

this study of phase transitions: connecting the physics of associative memory to Transformer attention. Yet their phase structure and the conditions governing these transitions remain poorly understood.

To address those issues, we utilize the free-energy minimization principle from statistical mechanics [Landau and Lifshitz, 1980, Binder, 1987, Zinn-Justin, 2002] to study phase structures of modern Hopfield models and hence of Transformer attention. Within this framework, different entropy induces different sparsity and hence different attention mechanism [Martins and Astudillo, 2016, Peters et al., 2019, Ramsauer et al., 2020]. Through this perspective, we explore first-order phase transitions and hence abrupt reorganization of attention weight beyond classical Hopfield models [Hopfield, 1982, Mezard et al., 1986]. This leads to a parameter-dependent mechanism behind modern Hopfield models and hence the phase structure of Transformer attention.

Contributions. We develop a statistical-mechanical framework for Transformer attention and use it to understand its inference behavior. Our main contributions are:

- **A Statistical-Mechanical Theory of Transformer Attention.** We formulate Transformer attention as a statistical-mechanical model of associative memory through the free energy of modern Hopfield networks. In this formulation, attention emerges naturally from the equations of motion (Propositions 3.2 and D.7) and its forward pass describes how a query retrieves stored information. This framework allows the unification of attention’s temperature, entropy, sparsity, fluctuations, and retrieval dynamics into one coherent physical picture.
- **Phase Transitions in Attention Inference.** We use the resulting free-energy landscape to study the reliability of information retrieval during a Transformer attention forward pass. Our analysis reveals two retrieval regimes: a glassy-like regime with competing metastable states and a robust-inference regime with stable, target-aligned retrieval. We characterize first-order phase transitions¹ between these regimes as attention parameters vary.
- **Critical Conditions for Robust Attention Inference.** We further analyze what governs the stability of memory selection in attention inference. We characterize critical conditions for the coexistence of competing states (Proposition 4.1) and their spectral stability (Proposition 4.2) from correlations among learned keys. These conditions explain how learned key correlations and attention temperature determine whether target-memory selection remains stable or becomes vulnerable to interference from competing keys.

Organization. Section 2 presents preliminaries. Section 3 presents physical formulation of attention through modern Hopfield models. Section 4 presents our phase transition results. Section 5 presents numerical validations. Section 6 presents concluding remarks and practical implications. We defer related work discussion and statistical-mechanics tutorial in Sections A and B due to space limit. Section D extends to the stochastic version, and Section E analyzes the dynamical stability. Sections F and G show experiment setup and additional experiments.

¹A first-order phase transition occurs when the equilibrium state changes discontinuously as a control parameter varies. In our case, varying attention temperature can induce such transitions in memory retrieval. Memory load and key correlations further shape the phase structure.

Table 1: Roadmap from the technical terms to their role in attention.

Concept	Meaning for Attention	Dictionary
Fenchel Conjugate	It transforms a microscopic quantity into a macroscopic one, and constructs the free-energy objective.	Definition C.6, Proposition 2.1, and Theorem 2.1
Tsallis Entropy Ψ_α	Attention Family: softmax ($\alpha = 1$), entmax ($1 < \alpha < 2$), and sparsemax ($\alpha = 2$).	Section C.3, Theorem 2.1, and Proposition 3.3
Query-Key Similarity (Overlap $z = \frac{1}{\sqrt{d}}\xi^\top x$)	It measures the alignment between queries and memories. A larger overlap indicates a stronger retrieval preference for that memory.	Definition 3.2
Free-energy Γ_α	It defines free-energy landscape and enables assigning energy value for each local minimum. It encodes equation of motion as attention rule.	Definition 2.1, Propositions 3.2 and 3.3
Coexistence (Energy Crossing)	It occurs when two distinct stable retrieval states have the same free energy Γ_α simultaneously.	Definition 4.2 and Proposition 4.1
Spinodal	It occurs when the energy barrier of a local minimum disappears. This drives a retrieval state to the other.	Definition 4.3 and Proposition 4.2
Exponential Memory Load $\kappa = \log(M)/d$	It quantifies the growth rate of the number of memories. This affects retrieval interference and reliability.	Section 5.2

2 Preliminaries

We first discuss convex analysis and geometric properties in statistical mechanics. This induces entropy and sparsity via the free energy framework. The purpose of this section is to identify the conjugate variables that characterize the retrieval dynamics and attention mechanism.

2.1 Statistical Mechanics and Energy Formalism

We introduce partition function characterizing a system with all possible configurations, and adopt it to discuss free energy through the statistical mechanics [Zinn-Justin, 2002, Srednicki, 2007].

Definition 2.1 (Partition Function). Let $\mathcal{H} : \mathcal{X} \rightarrow \mathbb{R}$ be a given function called an energy

function. We define the partition function as the Gibbs average over all configurations.

$$Z[J] := \int_{\mathcal{X}} \exp\{-\beta\mathcal{H}[x] + \sum_{i=1}^d J_i x_i\} d\mu(x),$$

where $\beta := 1/T$ is the inverse temperature, $J \in \mathbb{R}^d$ is the generating source, and $d\mu(x)$ denotes the Lebesgue measure in $\mathcal{X}$. The partition function is positive. For simplicity, we consider $\mathcal{X} = \mathbb{R}^d$.

Remark 2.1. The generating source term represents the external magnetic field in a spin system. Furthermore, if the model is discrete, then we replace $\int_{\mathcal{X}} d\mu(x)$ by $\sum_{x \in \mathcal{X}}$ with $\mathcal{X} = \{1, \dots, K\}$.

Definition 2.2 (Stationary Equations and Stationary Points). The stationary equations are

$$\nabla_x \{\beta\mathcal{H}[x] - \sum_{i=1}^d J_i x_i\}|_{x=x_J} = 0, \quad \nabla_x^2 \{\beta\mathcal{H}[x] - \sum_{i=1}^d J_i x_i\}|_{x=x_J} \succ 0, \quad (2.1)$$

where $\succ$ means positive definite. We say that the solution x_J is a stationary or (local) stable point.

We obtain stable points in [Definition 2.2](#) through the energy function $\mathcal{H}[x]$. However, it contains microscopic noise components. Hence, x_J is not a well-behaved stable point. To extract the aggregated stability from the original energy function $\mathcal{H}[x]$, we should introduce the free energy (effective action) $\Gamma[\varphi]$. It later helps characterize the macroscopic averaged energy of the system.

Definition 2.3 (Log-Partition Function and Free Energy). We have two different free energies that describe the partition function. To distinguish them, we define them carefully.

- **Log-Partition Function:** It is defined as

$$W[J] := \log Z[J] = \log \left(\int_{\mathcal{X}} \exp\{-\beta\mathcal{H}[x] + \sum_{i=1}^d J_i x_i\} d\mu(x) \right). \quad (2.2)$$

- **Free Energy (Effective Action):** It is the Fenchel conjugate of the log-partition function.

$$\Gamma[\varphi] := \sup_{J \in \mathbb{R}^d} \left\{ \sum_{i=1}^d J_i \varphi_i - W[J] \right\}, \quad \text{with } \varphi \in \mathbb{R}^d. \quad (2.3)$$

Remark 2.2. From [Section C.1](#), $W[J]$ and $\Gamma[\varphi]$ are both convex, $W[J]$ is smooth, $\nabla_J W[J]$ is single-valued, and $\Gamma[\varphi]$ is closed. However, we do not assume strict convexity. Consequently, the gradient maps $\nabla_J W[J]$ and $\nabla_{\varphi} \Gamma[\varphi]$ are not guaranteed to be bijective according to [Fact 1](#). In this way, distinct microscopic configurations are able to represent the same macroscopic state.

The optimization problem becomes minimizing the free energy $\Gamma[\varphi]$ with respect to φ . Next, we describe the physical interpretation/derivation of the Fenchel-Young duality.




Figure 1: **Illustration of the procedure from the microscopic to macroscopic energy landscapes.** **Left.** “(Microscopic) Fluctuations” means x is random, has nonzero variance, and follows Gibbs distribution $p_J(x)$. Built from microscopic x , the local stable point x_J exhibits randomness. **Middle.** $W[J] = \log(\int_{\mathcal{X}} \exp\{-\beta\mathcal{H}[x] + Jx\})$ encodes macroscopic responses, i.e., $\varphi = \mathbb{E}_{x \sim p_J}[x] = \nabla_J W[J]$. Hence, the slope of $W[J]$ gives a deterministic average. **Right.** $\Gamma[\varphi] = \sup_J \{J\varphi - W[J]\}$ is the Fenchel conjugate of $W[J]$. Its response is $J = \nabla_{\varphi} \Gamma[\varphi]$, so local stable macroscopic points occurs at $J = 0$. $\Gamma[\varphi]$ describes the energy landscape built from macroscopic φ after integrating out microscopic fluctuations.

Proposition 2.1 (Fenchel-Young Duality Relation [Srednicki, 2007, Blondel et al., 2020]). Let p_J be the Gibbs measure with external source J . From the standard statistical field theory, the induced macroscopic field $\varphi = \mathbb{E}_{p_J}[x]$ reveals the following source-field relation:

$$\varphi = \mathbb{E}_{x \sim p_J}[x] \in \partial W[J] \Leftrightarrow J \in \partial \Gamma[\varphi],$$

where ∂ is the subdifferential from Definition C.5. If W and Γ are differentiable, then ∂ reduce to gradients ∇ . We call them equations of motion (EOMs) in the physics literature. It also refers to the Fenchel-Young duality relation. For φ and J , we call them the classical solutions or saddle points.

Proof. Please see Section C.2 for a detailed proof. $\square$

From Proposition 2.1, the Fenchel conjugate variable φ is macroscopic $\mathbb{E}_{x \sim p_J}[x]$ exactly. φ eliminates fluctuations and induces deterministic landscape $\Gamma[\varphi]$. Please see Figure 1 for illustration.

2.2 Connection between Free Energy and Tsallis Entropy

In this section, we connect our free energy to Tsallis entropy. The value of α determines the sparsity structure of the retrieval dynamics. As we shall see in Section 3, this choice covers the general attention mechanisms, ranging from dense ($\alpha = 1$) to the sparse case ($\infty > \alpha > 1$).

For preliminaries on Tsallis entropy, α -deformed exponential and logarithm family, please see Section C.3 for completeness. They generalize free energy to the α -deformed version.

Theorem 2.1 (Tsallis Entropy and Deformed Free Energy). If we replace the exponential functions and logarithm functions by α -deformed functions, then we have the deformed log-partition

function

$$-W_\alpha[J] = \beta \mathbb{E}_{x \sim p_J}^{(\alpha)}[\mathcal{H}[x]] - \frac{\Psi_\alpha[p_J]}{Z_\alpha^{\alpha-1}[J] \int_{\mathcal{X}} p_J^\alpha(x) d\mu(x)} - \sum_{i=1}^d J_i \mathbb{E}_{x \sim p_J}^{(\alpha)}[x_i],$$

where $\mathbb{E}_{x \sim p_J}^{(\alpha)}[\cdot] := \int_{\mathcal{X}} p_J^\alpha(x) [\cdot] d\mu(x) / \int_{\mathcal{X}} p_J^\alpha(x) d\mu(x)$ is the escort expectation value, and

$$\Gamma_\alpha = \beta \mathbb{E}_{x \sim p_{J^*}}^{(\alpha)}[\mathcal{H}[x]] - \frac{\Psi_\alpha[p_{J^*}]}{Z_\alpha^{\alpha-1}[J^*] \int_{\mathcal{X}} p_{J^*}^\alpha(x) d\mu(x)}, \quad \text{with } \varphi = \mathbb{E}_{x \sim p_{J^*}}^{(\alpha)}[x], \quad (2.4)$$

where $J^* = \arg \max_J \{\sum_{i=1}^d J_i \varphi_i - W_\alpha[J]\}$. We assume $\mathcal{X}$ is finite-volume to avoid divergence.

Proof. Please see [Section C.4](#) for a detailed proof. $\square$

Remark 2.3. If $\alpha = 1$, an escort expectation value becomes an ordinary expectation value. Hence, if we further consider the zero-temperature limit $\beta \rightarrow \infty$, then minimizing the free energy Γ_α is equivalent to minimizing the energy function $\mathcal{H}$. In the high-temperature limit, minimizing Γ_α is equivalent to maximizing the Tsallis entropy Ψ_α . The latter case indicates a disorder regime.

Remark 2.4. [Theorem 2.1](#) relates to the evidence lower bound (ELBO) in variational autoencoders (VAE), [[Kingma and Welling, 2013](#)]. Specifically, with a posterior $q_\phi(z|x)$, we have

$$\log p_\theta(x) \geq \mathbb{E}_{z \sim q_\phi(z|x)}[\log p_\theta(x, z)] - \mathbb{E}_{z \sim q_\phi(z|x)}[\log q_\phi(z|x)].$$

We recognize that the second term on the right hand side is Shannon entropy, [\(C.3\)](#). Then,

$$\text{ELBO} := \mathbb{E}_{z \sim q_\phi(z|x)}[\log p_\theta(x, z)] + S[q_\phi(\cdot|x)].$$

It is an energy-entropy decomposition, where the expected log joint probability corresponds to a negative energy. Hence, maximizing ELBO is equivalent to minimizing a free energy [\(2.4\)](#).

3 Statistical Physics of Attention: Modern Hopfield Models

In this section, we define an energy function of MHMs in [[Hu et al., 2023](#), [Wu et al., 2023](#)]. Then we express the free energy in terms of entropy and derive the emergent attention mechanism.

3.1 Modern Hopfield Model (MHM)

Inspired by the optimization objectives in [[Ramsauer et al., 2020](#), [Krotov and Hopfield, 2020](#), [Hu et al., 2023](#), [Wu et al., 2023](#)], we introduce the explicit joint energy function of MHM as follows.

Definition 3.1 (Modern Hopfield Model (MHM)). The joint energy function of the MHM is

$$\tilde{\mathcal{H}}_{\text{MHM}}[x, p, \xi] := \frac{1}{2} \sum_{i=1}^d x_i^2 - \sum_{\mu=1}^M z_\mu p_\mu, \quad (3.1)$$

where $p \in \Delta^M$ lives on a probability simplex, i.e., p is non-negative and satisfies $\sum_{\mu=1}^M p_\mu = 1$, and $z_\mu := \sum_{i=1}^d \xi_i^\mu x_i / \sqrt{d}$. Moreover, we denote $\mathcal{H}_{\text{MHM}}[x, \xi] := \inf_{p \in \Delta^M} \tilde{\mathcal{H}}_{\text{MHM}}[x, p, \xi]$.

In our setting, query states, the probability distribution, and the stored patterns are all fixed, i.e., the distributions of x and p collapse to their escort average. Specifically, a Taylor expansion of $f(x)$ gives $\mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[f(x)] = f(\varphi) + O(\text{Var}^{(\alpha)}(x))$, where $\rho_J(x, p)$ is the joint probability distribution, $\varphi = \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[x]$ is the expectation value, and $\text{Var}^{(\alpha)}$ is the variance. Then the fixed-parameter setting corresponds to the zero variance. Hence, we are able to evaluate all quantities $f(x)$ at φ under this fixed-parameter setting exactly. Thus, we identify φ with x , and similarly for p and z .²

Definition 3.2 (Statistics-MHM Mapping). The mapping is:

- **Query State Vector.** $x \in \mathbb{R}^d$: It represents continuous query states. Due to our fixed-parameter setting, it is macroscopic and parametrizes the free energy $\Gamma_\alpha[x] = \Gamma_\alpha[\varphi]$.
- **External Source.** $J \in \mathbb{R}^d$: It represents an external bias of the query states. In particular, $J = 0$ identifies local minima of $\Gamma_\alpha[x]$ due to the Fenchel-Young relation ([Proposition 2.1](#)).
- **Probability Function.** $p \in \Delta^M$: It represents the retrieval distribution over memories.
- **Entropy.** Ψ_α : It acts as a regularizer controlling how sparse the memory weights p are.
- **Score.** $z \in \mathbb{R}^M$: It represents query-key similarity between the current x and each memory pattern ξ^μ via the relation: $z_\mu := \sum_{i=1}^d \xi_i^\mu x_i / \sqrt{d}$. Higher z_μ suggests higher similarity.

In fact, the query state vector x in [Definition 3.2](#) relates to the probability weight p as follows.

Proposition 3.1 (Query-Probability Relation). A query state relates to a probability function via

$$x_i = \frac{1}{\sqrt{d}} \sum_{\mu=1}^M \xi_i^\mu p_\mu + \frac{1}{\beta} J_i. \quad (3.2)$$

Proof. Please see [Section C.5](#) for a detailed proof. □

In [Section 2.2](#), we relate free energy to Tsallis entropy via (2.4). Inspired by that relation, we are capable of expressing the free energy in terms of the variables of MHMs as follows.

²Here we restrict attention to the deterministic regime. This allows us to isolate the physics most directly relevant to the attention mechanism. In [Section D](#), we consider the more general setting in which both x and p are random variables with nonzero variance. This leads to a richer and more complete statistical-mechanics structure, and allows us to analyze fluctuations and response properties.

Theorem 3.1 (Free Energy of MHM). The joint free energy of the MHM is

$$\tilde{\Gamma}_\alpha[x, p, \xi] = \frac{\beta}{2} \sum_{i=1}^d x_i^2 - \beta \sum_{\mu=1}^M z_\mu p_\mu. \quad (3.3)$$

Hence, after taking the infimum with respect to p , we obtain the free energy

$$\Gamma_\alpha[x, \xi] = \frac{\beta}{2} \sum_{i=1}^d x_i^2 - \beta \max_{\mu \in [1, M]} z_\mu. \quad (3.4)$$

Proof. Please see [Section C.6](#) for a detailed proof. $\square$

This free energy selects the maximal component of the overlap $\max_\mu z_\mu$. The hard selection mechanism arises from zero entropy due to the collapse of p . This captures a supremum over simplex vertices rather than the full simplex Δ^M . The relaxation from a hard selection to a distribution over multiple memories becomes a natural extension. Motivated by the sparsity of the entropy regularizer in [\[Hu et al., 2023\]](#), we consider the relaxation through Fenchel conjugate as follows.

We introduce the entropy regularizer in (3.4). This yields a generalized (relaxed) model:

$$\Gamma_\alpha^{(\text{mod})}[x, \xi] = \frac{\beta}{2} \sum_{i=1}^d x_i^2 - \Omega_\alpha^*[\beta z], \quad \Omega_\alpha := -\Psi_\alpha, \quad (3.5)$$

where $\Omega_\alpha^*[\beta z] = \sup_{p \in \Delta^M} \{\sum_{\mu=1}^M \beta z_\mu p_\mu - \Omega_\alpha[p]\}$ is the Fenchel conjugate of the negative Tsallis entropy [\(C.5\)](#). It recovers the energy form in [\[Hu et al., 2023, Wu et al., 2023\]](#) exactly.

Remark 3.1. This generalization allows us to relax p to have a distribution over multiple memories on the full simplex Δ^M . This leads to several theoretical and practical advantages:

1. **Equilibrium Selection vs. Averaging.** The MHM performs a hard maximization, selecting a single dominant component of $\Gamma_\alpha^{(\text{mod})}[x, \xi]$. In contrast, expectation-based rules compute averages over configurations. The generalization (3.5) interpolates between these regimes, i.e., it reduces to the original MHM in $\beta \rightarrow \infty$ limit. Please see [Section D](#) for details.
2. **Robustness to Noise.** As we will see in [Proposition 3.3](#), $\Gamma_\alpha^{(\text{mod})}$ emphasizes dominant components of z for $\alpha > 1$. Hence, it suppresses contributions from noisy patterns. In contrast, expectation-based rules average all components, making them more sensitive to noise.
3. **Sparsity.** The entropy function Ψ_α controls the support of sparsity. In particular, for $\alpha > 1$, the maximizer p^* is sparse, so some of the components are zero. This is also a consequence of robustness to noise, which does not appear in expectation-based formulations.
4. **Computational Accessibility.** Fenchel conjugation leads to an explicit variational object and is more accessible with gradient-based methods. In contrast, expectation-based formulations require sampling or approximation of Gibbs averages, which is more demanding.

3.2 Retrieval Dynamics and Attention Mechanism

In [Section 2](#), we introduce EOMs from [Proposition 2.1](#). These equations describe the trajectory that the flow should follow. In this section, we investigate the physical meaning of the trajectories in associative memory systems. In particular, the attention emerges as an EOM naturally.

Proposition 3.2 (EOMs of Free Energies in MHMs). Let Z_α be the partition function with α -deformed Gibbs average. Then the EOM in the J -space corresponds to the Hopfield retrieval

$$\nabla_J W_\alpha[J, \xi] = x, \quad \text{with } W_\alpha[J, \xi] := \log_\alpha Z_\alpha[J, \xi]. \quad (3.6)$$

The EOM of the joint log-partition function in the z -space corresponds to

$$\nabla_z \widetilde{W}[z, p, J] = \beta p, \quad \text{with } \widetilde{W}[z, p, J] := -\beta \widetilde{\mathcal{H}}_{\text{MHM}}[x, p, \xi] + \sum_{i=1}^d J_i x_i. \quad (3.7)$$

The EOMs of the joint free energy and the generalized one are

$$\nabla_z \widetilde{\Gamma}_\alpha[x, p, \xi] = \nabla_z \widetilde{\Gamma}_\alpha^{(\text{mod})}[x, p, \xi] = -\beta p. \quad (3.8)$$

Proof. Under fixed x and p , the log-partition function is $W_\alpha[J, \xi] = -\beta/2 \sum_{i=1}^d x_i^2 + \beta \sum_{\mu=1}^M z_\mu p_\mu + \sum_{i=1}^d J_i x_i$ following from the proof of [Theorem 3.1](#). After taking the derivative of $W_\alpha[J, \xi]$ with respect to J , we obtain the first equation. For the second and the last ones, we derive them by taking the derivative of (3.1), (3.4), and (3.5) with respect to z_μ directly. $\square$

We observe the emergent attention mechanism from the EOMs in [Proposition 3.2](#). Specifically,

Proposition 3.3 (Emergent Attention Mechanism as EOMs in MHMs). Attention emerges as:

- If $\alpha \rightarrow \infty$, it corresponds to a hardmax function:

$$p_\mu = \text{hardmax}(z_\mu) = \begin{cases} 1, & \text{for maximal } z_\mu \\ 0, & \text{otherwise} \end{cases}.$$

- If $\alpha = 2$, it corresponds to a sparsemax function [[Martins and Astudillo, 2016](#)]:

$$p_\mu = \text{Sparsemax}_\beta(z_\mu) = \frac{\beta}{2} [z_\mu - \theta]_+, \quad \text{with } \frac{\beta}{2} \sum_{\mu=1}^M [z_\mu - \theta]_+ = 1.$$

- If $1 < \alpha < 2$, it corresponds to an entmax family [[Peters et al., 2019](#)].
- If $\alpha \rightarrow 1$, it corresponds to a softmax function [[Blondel et al., 2020](#)]:

$$p_\mu = \text{Softmax}_\beta(z_\mu) = \frac{\exp\{\beta z_\mu\}}{\sum_{\nu=1}^M \exp\{\beta z_\nu\}}.$$

Proof. Please see [Section C.7](#) for a detailed proof. $\square$

Proposition 3.3 provides a re-derivation via EOMs of a free energy landscape, revealing the underlying physical interpretation of the attention mechanism. In Transformer language, [Vaswani et al., 2017], we denote $\text{Attn}(Q, K, V) := \sum_{\mu=1}^M p_{\mu}(Q, K) V_{\mu}$ as the representation of x , where Q represents the microscopic states, and K and V represent the stored memory patterns. Q, K, V are all fixed inputs. This corresponds to our zero variance or deterministic setting.

4 Phase Transitions and Critical Points of Retrieval Success

In this section, we discuss the spectral stability and thresholds for retrieval success. For the dynamical stability, please see [Section E](#). β, M, d and $\kappa := \exp(M)/d$ determine the free energy and allow multiple local minima in some ranges. This leads to our study of the phase structure. In our setting, the stored memories $\{\xi^{\mu}\}$ remain fixed, while the query state x evolves dynamically.

Definition 4.1 (Branch). Choose an initial parameter θ_0 that determines $\Gamma_{\alpha}^{(\text{mod})}[x(\theta_0), \xi]$. As $\theta \in \mathcal{I}_{\theta_0}$ varies from θ_0 , we define a branch as a continuous family of the evolution from a specific local stable point $\{x_a^*(\theta; \theta_0)\}_{\theta \in \mathcal{I}_{\theta_0}}$ satisfying the stationary condition ([Definition 2.2](#))

$$\nabla_x \Gamma_{\alpha}^{(\text{mod})}[x_a^*(\theta; \theta_0), \xi] = 0, \text{ and } \nabla_x^2 \Gamma_{\alpha}^{(\text{mod})}[x_a^*(\theta; \theta_0), \xi] \succ 0.$$

Branches help us analyze the evolution of stable points and their loss of stability. Adopting this setup, we study the phase transitions of the generalized model (3.5) following stability analysis in statistical mechanics [Binder, 1987, Zinn-Justin, 2002]. We start from defining two critical points.

Definition 4.2 (Coexistence Point). Given two distinct branches, a, b , with their stable points, $\{x_a^*(\theta; \theta_0)\}, \{x_b^*(\theta; \theta_0)\}$. A coexistence point is a parameter value θ_{coexist} such that

$$\Gamma_{\alpha}^{(\text{mod})}[x_a^*(\theta_{\text{coexist}}; \theta_0), \xi] = \Gamma_{\alpha}^{(\text{mod})}[x_b^*(\theta_{\text{coexist}}; \theta_0), \xi]. \quad (4.1)$$

However, the coexistence condition is not a necessary condition for a real jump of local minima. Even if the equality holds, there is still an energy barrier that forbids the transition.

Definition 4.3 (Spinodal Point). Given a smooth branch a , a spinodal point occurs when the branch loses its local stability. That is, the point where the stationary condition fails:

$$\lambda_{\min}(\nabla_x^2 \Gamma_{\alpha}^{(\text{mod})}[x_a^*(\theta_{\text{jump}}; \theta_0), \xi]) = 0. \quad (4.2)$$

This is also the point where the continuation along that branch breaks down.

Remark 4.1. If a branch is non-smooth, e.g., there is a sudden change in the effective support size, then jumps also occur at that discontinuous point. In this case, (4.2) does not hold necessarily.

We control the system by the parameters, $\theta = (\beta, M, d, \xi)$. For $\alpha = 1$, the retrieval map is the softmax map ([Proposition 3.3](#)), so p_{μ}^* is positive for all memories and the weights vary with z^*

smoothly. Hence, discontinuous change of retrieval only occurs if the Hessian vanishes, i.e., (4.2) holds. To investigate the transitions, we focus on the critical value of β in the following.

Proposition 4.1 (Temperature Coexistence Point for Sparse Model). Let a, b be two distinct branches and $\alpha = 2$. Denote k_a and k_b as the exact support sizes (see Lemma F.1) of branches a and b , respectively. Based on the statistical mechanics framework, the coexistence point satisfies

$$\sum_{\mu, \nu \in a} [z_\mu^* - \theta]_+ [z_\nu^* - \theta]_+ K_{\mu\nu}(\beta_{\text{coexist}}; \beta_0) = \sum_{\mu, \nu \in b} [z_\mu^* - \theta]_+ [z_\nu^* - \theta]_+ K_{\mu\nu}(\beta_{\text{coexist}}; \beta_0),$$

where $K_{\mu\nu}(\beta) := \mathbb{1}_{\{\mu, \nu\}} - \frac{\beta}{2} G_{\mu\nu}$, and $G_{\mu\nu} := \sum_{i=1}^d \xi_i^\mu \xi_i^\nu / d$ is the Gram matrix. β_{coexist} does not admit a closed-form expression. However, if the distribution is uniform on the simplex and the system satisfies the well-separation condition, i.e.,

$$G_{\mu\nu} = \begin{cases} 1 & , \mu = \nu \\ \epsilon \ll 1 & , \mu \neq \nu \end{cases}, \quad (4.3)$$

then for $k_a = 1$ and $k_b = 2$, we have the analytic solution $\beta_{\text{coexist}} = 2/(1 - \epsilon)$. At this point, $\beta = \beta_{\text{coexist}}$, p^* is degenerate and every p^* is a global minimum of the joint free energy $\tilde{\Gamma}_{\alpha=2}^{(\text{mod})}$.

Proof. Please see Section C.8 for a detailed proof. $\square$

Proposition 4.2 (Temperature Spinodal Point for Sparse Model). Consider the sparse case $\alpha = 2$ and a branch a with exact support size k . Then the spinodal critical value is

$$\beta_{\text{jump}} = \frac{2}{\lambda_{\max}(PGP)},$$

where $P_{\mu\nu} := \mathbb{1}_{\{\mu, \nu\}} - k^{-1}$, and G is the Gram matrix. If the distribution is uniform and the system satisfies the well-separation condition, (4.3), then we have the analytic solution $\beta_{\text{jump}} = 2/(1 - \epsilon)$.

Proof. Please see Section C.9 for a detailed proof. $\square$

Remark 4.2. Under the uniform distribution condition and the well-separation condition, we conclude that $\beta_{\text{coexist}} = \beta_{\text{jump}} = 2/(1 - \epsilon)$. However, we do not have these assumptions in general. Then the real jump and the coexistence point do not occur simultaneously, i.e., $\beta_{\text{jump}} \neq \beta_{\text{coexist}}$.

5 Numerical Results

We provide numerical evidence supporting the ideas in Section 4. Specifically, we display the phase structure of MHMs under temperature and memory-load scans, including both $\alpha = 1, 2$.

5.1 Temperature Dependence

We examine the temperature dependence of the phase structure in MHMs, including both the sparse ($\alpha = 2$) and dense ($\alpha = 1$) regimes. Specifically, we characterize the temperature effect on the attention map (Proposition 3.3) and the branches of free energy (3.5). For the experimental setup, please see Section F. We first start from the temperature analysis in the sparse case.




Figure 2: **Temperature analysis of the retrieval structure of MHM in the sparse regime** ($\alpha = 2$). ($M = 100, d = 64$). Results are averaged over 10 random pattern realizations, and shaded regions indicate one standard deviation across different seeds. **(a) Retrieval branches.** The maximal overlap $z_{\max}^*$ under up- and down-sweeps of β exhibits hysteresis. **(b) Support size.** We compare the exact support size (active components of the probability distribution p), the inverse participation ratio (IPR), and the fixed score-threshold count (Section F.2). **(c) Top branches.** We track the top five stable branches ranked by $z_{\max}^*$. **(d) Branch energies.** We show the corresponding free energies of those top five branches.

Sparse case ($\alpha = 2$). Here, we select four plots to represent the retrieval structure in Figure 2.

- (a) The up- and down-sweep branches do not coincide, which shows a hysteresis window in β . This indicates that the coexistence of multiple local stable retrieval states (Definition 4.2 and Proposition 4.1). The selected branch depends on initialization, analogous to first-order transitions in freezing-melting hysteresis of water [Enninfu et al., 2019]. As β increases, the maximal overlap $z_{\max}^*$ grows stepwise and the shaded regions collapse at large β . This indicates a seed-independent saturation. Specifically, it saturates at $z_{\max}^* \approx 0.0156$. Indeed,

this is the consequence of one pattern retrieval, i.e.,

$$z_\mu^* = \frac{1}{\sqrt{d}} \sum_{i=1}^d \xi_i^\mu x_i^* = \frac{1}{\sqrt{d}} \times d \times \frac{1}{\sqrt{d}} \times \frac{1}{d} = \frac{1}{d} = \frac{1}{64} \approx 0.0156.$$

- (b) As β increases, the exact support size and IPR decreases (Section F.2). This indicates a transition from distributed to sparse retrieval. The discrepancy of $|\text{supp}(p^*)|$ and IPR arises since IPR counts the components non-uniformly. However, K_Δ is zero in the high-temperature regime, where noise dominates and no overlap exceeds the threshold $\Delta = 0.01$. This suggests a disordered regime. In the low-temperature regime, all support sizes approach to one, so the system reaches single-pattern retrieval with large overlap. The collapse of the shaded regions again indicates a perfect retrieval and seed-independence.
- (c) Initialization at $\beta_0 = 290$. Multi-branch continuation indicates several competing stable retrieval branches in the intermediate- β regime. The presence of these branches indicates multiple coexisting local minima. Their crossings and discontinuous jumps in $z_{\max}^*$ show that the sparse model admits branch switching between distinct local minima (Proposition 4.2 and Remark 4.1). Hence, the selected state correspond to either stronger or weaker alignment depending on the branch we are tracking.
- (d) Initialization at $\beta_0 = 290$. The branch exhibits coexist point from the plot (Definition 4.2 and Proposition 4.1). As β varies, the global minimizer switches between branches. However, under continuation (warm-start dynamics), the system follows a given local branch rather than the instantaneous global minimizer. Hence, the transition between branches occur before reaching the exact crossing point. This leads to discontinuous changes along a branch. Hence, we interpret this as a transition within the retrieval phase rather than a transition between retrieval and a disordered phase. This simulation provides a more precise phase structure than earlier analyses, [Amit et al., 1985, 1987, Nishimori, 2001].

Dense case ($\alpha = 1$). Here, we select four plots to represent the retrieval structure in Figure 3.

- (a) The up- and down-sweep branches coincide, which indicates the absence of hysteresis. Hence, the retrieval dynamics admits one stable solution for each β . There is no metastable branch or discontinuous transition. The shaded region becomes broader as β increases. Hence, the alignment with stored pattern increases, but the variance increases as well.
- (b) The support size remains full for all β (Section F.2), since $p_\mu^* \propto \exp(\beta z_\mu^*)$ assigns positive weight to every component μ . As β increases, $K_{\text{eff}}^{\text{IPR}}$ decreases, so the distribution becomes concentrated. However, this does not induce sparsity. The shaded region of $K_{\text{eff}}^{\text{IPR}}$ becomes broader as β increases. This is due to the amplification of small differences from seeds.
- (c) Initialization at $\beta_0 = 140$. The maximal overlap $z_{\max}^*$ grows with β . This indicates a continuous sharpening of alignment with the target pattern. There is no discontinuity or spinodal instability due to the vanishing Hessian. Hence, it shows the absence of phase-transition-like behavior, based on our softmax experiment setup here. However, it does not indicate a unique local minimum necessarily. That is, other local minima may locate far away from our selected initial point so that the nearby search only displays one branch.




Figure 3: **Temperature analysis of the retrieval structure of MHM in the dense regime** ($\alpha = 1$). ($M = 100, d = 64$). Results are averaged over 10 random pattern realizations, and shaded regions indicate one standard deviation across different seeds. **(a) Retrieval branches.** The maximal overlap $z_{\max}^*$ under up- and down-sweeps of β shows no hysteresis. **(b) Support size.** We compare the exact support size (active components of the probability distribution p) and the inverse participation ratio (IPR) (Section F.2). **(c) Top branches.** We observe one single stable branch across the entire range of β . **(d) Branch energies.** We show the corresponding free energy of the branch.

(d) Initialization at $\beta_0 = 140$. The energy decreases along the branch, and the retrieval state becomes more concentrated. This reflects a continuous sharpening of the energy basin, rather than a restructuring of competing minima as in the sparse case.

Remark 5.1. Koulischer et al. [2023] analyze temperature-dependent behavior in MHM under softmax attention ($\alpha = 1$). They show that as β increases, the retrieval weights concentrate smoothly. Our results here are consistent with their findings. Besides, we further provide a characterization of this behavior by tracking the maximal overlap $z_{\max}^*$ and the support sizes.

5.2 Memory-Load Dependence

Here we study the exponential memory load $M = \exp(\kappa d)$, where M is the pattern number and d is the dimension of each pattern. Specifically, we characterize the memory-load effect on the attention map (Proposition 3.3) and the branches of free energy (3.5). For the experimental setup, please see Section F. We start from the memory-load analysis of the sparse case.

Sparse case ($\alpha = 2$). Here, we select four plots to represent the retrieval structure in Figure 4.




Figure 4: **Memory-load analysis of the retrieval structure of MHM in the sparse regime** ($\alpha = 2$). ($d = 48, \beta = 180$). Results are averaged over 10 random pattern realizations, and shaded regions indicate one standard deviation across different seeds. **(a) Retrieval branches.** The maximal overlap $z_{\max}^*$ under up- and down-sweeps of κ shows hysteresis. **(b) Support size.** We compare the exact support size $|\text{supp}(p^*)|$ (active components of the probability distribution p) and the inverse participation ratio (IPR) (Section F.2). **(c) Top branches.** We track the top five stable branches ranked by $z_{\max}^*$. **(d) Branch energies.** We plot the corresponding free energy of the five branches.

- (a) The up- and down-sweep branches do not coincide, which shows a hysteresis window in κ . This indicates the coexistence of multiple local stable retrieval states and the presence of first-order phase transitions. As κ increases, the maximal overlap $z_{\max}^*$ decreases stepwise. This suggests the delocalization when the stored patterns are too many.
- (b) The system remains in a sparse retrieval regime: although additional overlaps enter the active set as κ increases, IPR remains close to one (Section F.2). This indicates that a single dominant memory continues to control the retrieval. Besides, the variance of fixed scored-threshold support size (Section F.2) grows with κ due to the delocalization.
- (c) Initialization at $\kappa_0 = 0.12$. Multiple branches coexist and represent distinct local stable retrieval states in this κ -window. The selected state depends on initialization and the continuation path. This reflects metastability in the energy landscape, i.e., multiple stable points coexist and the dynamics remains trapped in a given basin until stability is lost.
- (d) Initialization at $\kappa_0 = 0.12$. We plot the corresponding free energies of the five branches. As κ grows, all the branch energies decrease. However, it does not indicate that the retrieval is better. Instead, panel (a) suggests the delocalization (low score) at high κ .




Figure 5: **Memory-load analysis of the retrieval structure of MHM in the dense regime** ($\alpha = 1$). ($M = \exp(\kappa d)$, $d = 48$, $\beta = 180$). Results are averaged over random pattern realizations, and shaded regions indicate one standard deviation across different seeds. **(a) Retrieval branches.** The maximal overlap $z_{\max}^*$ under up- and down-sweeps of κ shows no hysteresis. **(b) Support size.** We compare the exact support size (active components of the probability distribution p) and the inverse participation ratio (IPR) (Section F.2). **(c) Top branches.** We track and observe only a single stable branch across the entire range of κ . **(d) Branch energies.** We plot the corresponding free energy of the branch.

Dense case ($\alpha = 1$). Here, we select four plots to represent the retrieval structure in Figure 5.

- (a) The up- and down-sweep branches coincide, which shows the absence of hysteresis in κ . That is, there is no coexistence window of multiple local stable retrieval states. As κ increases, the maximal overlap $z_{\max}^*$ decreases. This suggests the delocalization at large κ .
- (b) Both the exact support size and IPR (Section F.2) grow with κ . This reflects the dense nature of softmax, i.e.,

$$p_{\mu}^* = \frac{\exp\{\beta z_{\mu}^*\}}{\sum_{\nu=1}^M \exp\{\beta z_{\nu}^*\}} > 0, \quad \forall \mu, \quad \text{with overlap } z_{\mu}^*.$$

All patterns have strictly positive weight. Hence, the support size $|\text{supp}|$ remains full, and the delocalization drives IPR to grow exponentially with κ as well.

- (c) Initialization at $\kappa_0 = 0.08$. Only a single continuation branch appear. The energy decreases with κ monotonically. This indicates a smooth deformation of the solution rather than transitions between competing states. Again, this does not indicate the uniqueness of this local minimum. Other local minima may locate far away from our initial point.

- (d) Initialization at $\kappa_0 = 0.08$. We plot the corresponding free energies of the branch. As κ grows, the branch energy decreases monotonically. However, it does not indicate that the retrieval is better. Instead, panel (a) suggests the delocalization (low κ) at high κ .

Remark 5.2. The curves show the non-smoothness as functions of κ . This is the consequence of the fact that we require M to be an integer in our simulation.

6 Discussion and Conclusion

We present a free energy framework for Transformer attention via modern Hopfield models. This viewpoint identifies attention as the equation of motion on an energy landscape. Our work shows the empirical reliable retrieval regime with respect to entropy, temperature, and exponential memory load (i.e., $\alpha/\beta/\kappa$). We study the temperature and exponential memory load effects on phase structure. In the sparse regime, discontinuous transitions occur due to support change in memories (keys). The accompanying hysteresis indicates the coexistence of multiple stable states. In contrast, the softmax regime exhibits smooth evolution without hysteresis. Consequently, the softmax attention is a mixture of all patterns, while the query-key alignment grows with inverse-temperature smoothly. These results reveal the entropy-dependent phase behavior of modern Hopfield models and hence in Transformers.

Practical Implications. We sketch some design principles from our results:

- First, the attention should match the task: sparse one is support-efficient, while dense one is more accurate when many relevant values contribute ([Proposition 3.3](#) and [Sections G.2](#) and [G.3](#)).
- Second, the retrieval depends on path under temperature changes ([Proposition 4.1](#) and [Figures 2](#) and [3](#)). Specifically, increasing versus decreasing β lead to different different memory selection in the sparse case. The outcome also depends on the initial query (i.e., input context or prompt) supplied to the retrieval update during inference.
- Third, increasing the inverse-temperature strengthens the maximal query-key similarity $z_{\max} = \max_{\mu} \sum_{i=1}^d x_i \xi_i^{\mu} / \sqrt{d}$ ([Figures 2](#) and [3](#)). In particular, sparse attention reaches single-memory retrieval saturation, while dense one sharpens smoothly with full support.
- Finally, increasing memory load reduces retrieval accuracy ([Figures 4](#) and [5](#)). This further reveals a capacity-reliability trade-off in realistic Transformer architectures ([Section G.2](#)).

Together, these provide concrete guidance for selecting attention, temperature and memory-load.

Limitations. Our analysis relies on simplified assumptions, such as prescribed correlation and noise structures, and the fixed memory setting. In contrast, realistic Transformer architectures allow interactions across heads and layers. This thereby leads to complicated noise. Furthermore, the keys and queries co-evolve through learned projection matrices and upstream layers. Hence, the theoretical analysis on phase transitions and retrieval reliability of full Transformers remain less explored. Extending our framework along these directions is crucial for future work.

We defer **more experiments** to **Sections F** and **G** due to page limit.

Impact Statement

By the theoretical nature of this work, we do not anticipate any negative social impact.

AI Usage Disclosure

We used large language models (LLMs) to aid, polish writing and debug codes, such as improving clarity, grammar, and conciseness. We also used LLMs for retrieval and discovery, for example, exhausting literature to identify potential missing related work. Finally, we used LLMs to brainstorm ideas for selected experiments in **Section G**, including suggestions for the model architecture, the number of attention heads, and the hidden dimensions [Vaswani et al., 2017, Xiong et al., 2020]. All technical content, proofs, and results are original contributions by the authors. The authors take full responsibility for the paper’s content and conclusions.

Acknowledgments

JH is partially supported by Northwestern University’s Walter P. Murphy Fellowship and Terminal Year Fellowship (Paul K. Richter Memorial Award). Han Liu is partially supported by NIH R01LM1372201, NSF AST-2421845, Simons Foundation MPS-AI-00010513, AbbVie , Dolby and Chan Zuckerberg Biohub Chicago Spoke Award. The content is solely the responsibility of the authors and does not necessarily represent the official views of the funding agencies.

Appendix

A	Related Work	21
B	A Minimal Introduction of Statistical Mechanics	22
B.1	From Interacting Spins to Associative Memory	22
B.2	A Concrete Hopfield Free Energy	23
B.3	Landscapes, Metastability, and Retrieval Jumps	25
C	Definitions and Proofs in the Main Text	26
C.1	Basic Tools for Convex Analysis	26
C.2	Proof of Proposition 2.1	27
C.3	Tsallis Entropy	29
C.4	Proof of Theorem 2.1	33
C.5	Proof of Proposition 3.1	35
C.6	Proof of Theorem 3.1	35
C.7	Proof of Proposition 3.3	37
C.8	Proof of Proposition 4.1	39
C.9	Proof of Proposition 4.2	40
D	Fluctuated Version and Thermal Attention	43
D.1	Physics of Temperature	43
D.2	Fluctuated Variables in Modern Hopfield Model	46
D.3	Connection between Free Energy and Tsallis Entropy in Fluctuated Case	47
D.4	Retrieval Dynamics in Fluctuated Case and Thermal Attention	49
E	Dynamical Retrieval Stability Analysis	56
F	Experiment Setup	61
F.1	Effective Support Size and Sparsity	61
F.2	Settings	63
G	Additional Experiments	66
G.1	Dynamical Stability Analysis for Temperature	66
G.2	Support and Capacity-Reliability Trade-off	68
G.3	Support, Temperature, and Reliability	72
G.4	Numerical Validation of Coexistence and Spinodal Formula	74

A Related Work

Classical and Modern Hopfield Models. Hopfield-type models originate from the physical interpretation of associative memory in spin systems, including the Ising model [Onsager, 1944], the Sherrington-Kirkpatrick model [Sherrington and Kirkpatrick, 1975], and the classical Hopfield model [Hopfield, 1982]. Modern Hopfield models modify the energy function so that stored associative memories become local minima on the energy landscape, leading to a higher memory capacity [Krotov and Hopfield, 2020].

Attention mechanisms [Vaswani et al., 2017] have been highly successful in large-scale architectures. Ramsauer et al. [2020] establishes the connection between softmax modern Hopfield retrieval and Transformer attention. Hu et al. [2023] further extends the softmax to sparse retrieval and sparse attention by applying convex conjugate. Peters et al. [2019] introduce the entmax family of attention maps, which interpolates between dense and sparse attention. Complementarily, Millidge et al. [2022] develop universal Hopfield networks. It formulates a retrieval dynamics through three operations: similarity, separation and projection. Santos et al. [2025] systematize this line of work into a Hopfield-Fenchel-Young framework, covering Fenchel-Young losses and the variational formulation of retrieval.

Free Energy. Our formulation adopts the variational principles used in physics. Specifically, this method appears in quantum field theory [Zinn-Justin, 2002], large- N mean-field theory [Moshe and Zinn-Justin, 2003], and low-temperature theory [Landau and Lifshitz, 1980]. This viewpoint describes the retrieval process as a variational problem over macroscopic parameters, rather than microscopic states. Related optimization-based interpretations of Transformer architectures are also explored in [Yang et al., 2022], where Transformer layers are viewed as unfolding steps of an energy minimization process. Unlike traditional analyses of Hopfield networks that focus on energy minimization alone, the free energy framework incorporates entropy, leading to a richer landscape with multiple metastable states.

Phase Transition. Physicists often characterize phase transitions in energy-based systems through spontaneous symmetry breaking [Coleman and Weinberg, 1973] and the emergence of metastable states. In classical Hopfield networks, Amit et al. [1985, 1987] and Mezard et al. [1986] analyze this behavior using replica methods. This reveals transitions between the retrieval and spin-glass phases, which contain numerous spurious states that lead to retrieval failure. Free energy provides a description of phase transitions through the deformation of energy landscapes and the emergence of local minima [Binder, 1987]. Recent work from Koulischer et al. [2023] discusses the phase transition in the softmax case of modern Hopfield models. They identify a critical inverse temperature at which the mean minima are close to individual stored patterns, rather than one common averaged minimum. However, their analysis does not include sparse retrieval map and the entmax family.

Dense Associative Memory. Dense associative memory models extend classical Hopfield models by introducing higher-order non-linear energy functions, which improved memory load storage

in energy-based models [Demircigil et al., 2017, Krotov, 2023]. These interactions further reshape the energy landscape and enhance robustness to noise [Krotov and Hopfield, 2018].

Fuzzy Memory. Modern Hopfield models admit distributed retrieval over multiple memories. Recent work studies how such retrieval behavior relates to memory capacity. Specifically, Wu et al. [2024] introduce a learnable feature map that separates stored patterns in a kernel space. This makes the distribution of memory minima more uniform. On the other hand, Hu et al. [2024b] consider spherical codes, i.e., stored memories in the kernel feature space as points on a hypersphere. This induces a tight capacity bound by analyzing how many feature vectors maintain angular separation on a hypersphere. These works specify regimes where multiple patterns are activated simultaneously. Such distributed states interpolate between sparse (single-pattern) and dense (multi-pattern) retrieval regimes.

Despite these prior works, a unified statistical framework of modern Hopfield models remains limited. Existing analyses focus on either retrieval capacity or dynamical behavior, without a systematic treatment of entropy, sparsity, and metastability within a single framework. Furthermore, replica analyses characterize transitions between retrieval, spin-glass, and disordered phases, but not capture the detailed finite-dimensional geometry for modern Hopfield retrieval. In this work, we introduce a free energy formulation that unifies these aspects and study the induced finite-dimensional retrieval landscape. In particular, we focus on phase behavior inside the retrieval regime, rather than only transitions between retrieval and disordered phases. This allows us to compare dense and sparse regime, analyze metastable states, and verify retrieval success. In addition, we analyze how exponential memory load κ changes empirical retrieval behaviors. This probes memory-dependent phase regimes which are not addressed by existing Fenchel-Young or sparse Hopfield analyses.

B A Minimal Introduction of Statistical Mechanics

This appendix gives a physical route from interacting spins to associative memory and then to attention. We set Boltzmann’s constant to one, so the inverse temperature is $\beta = 1/T$.

B.1 From Interacting Spins to Associative Memory

A Hamiltonian assigns an energy to each microscopic configuration; a choice of dynamics specifies how the system moves through that landscape.

Definition B.1 (Classical Hopfield Model [Hopfield, 1982]). Consider binary spins $x_i \in \{-1, +1\}$ with couplings $J_{ij} = d^{-1} \sum_{\mu=1}^M \xi_i^\mu \xi_j^\mu$ and $J_{ii} = 0$. This gives the classical Hopfield model

$$\mathcal{H}_{\text{Hop}}(x) = -\frac{1}{2d} \sum_{i \neq j}^d \sum_{\mu=1}^M \xi_i^\mu \xi_j^\mu x_i x_j, \quad (\text{B.1})$$

where memories are denoted as $\xi^1, \dots, \xi^M \in \{-1, +1\}^d$.

Thus, zero-temperature Hopfield dynamics is a dissipative descent toward a local energy minimum. At nonzero temperature the system instead fluctuates among configurations, as described in [Section B.2](#).

Definition B.2 (Overlap). The overlap between the current state and memory μ is

$$z_\mu(x) := \frac{1}{\sqrt{d}} \sum_{i=1}^d \xi_i^\mu x_i. \quad (\text{B.2})$$

Therefore, low energy means large overlap with one or more stored patterns.

Definition B.3 (Frustration [[Kirkpatrick, 1977](#)]). A system is frustrated if it does not satisfy $J_{ij}x_i x_j > 0$ for every pair with $J_{ij} \neq 0$.

A one-pattern Hopfield model is equivalent to an ordinary ferromagnet and retrieves without frustration. However, with several patterns, the pairwise alignments become incompatible. This creates frustration and additional metastable minima. Hence, stored patterns are attractors, but not every local minimum must be a stored memory. This is the elementary mechanism of associative retrieval shown in [Figure 6](#).




Figure 6: **Retrieval and frustration in the classical Hopfield picture.** (a) From a noisy query set, we retrieve a stored pattern by minimizing $\mathcal{H}$. Equivalently, the overlap increases during this associative retrieval. (b) Two ferromagnetic bonds and one antiferromagnetic bond on a triangle is impossible to satisfy pairwise energy minimization simultaneously. This is the simplest illustration of frustration.

B.2 A Concrete Hopfield Free Energy

In this section, we take one-pattern Hopfield model for a self-contained example of the statistical mechanics construction. That is, we choose $\mathcal{H}_{\text{1pat}}(x) = -\frac{1}{2}z^2(x)$. At temperature $T = 1/\beta$, the Gibbs distribution and partition function ([Definition 2.1](#)) are

$$p_J(x) = \frac{e^{-\beta\mathcal{H}_{\text{1pat}}(x) + J^\top x}}{Z(J)}, \quad Z(J) = \sum_{x \in \{-1, +1\}^d} e^{-\beta\mathcal{H}_{\text{1pat}}(x) + J^\top x}. \quad (\text{B.3})$$

By grouping states via Gibbs distribution and Fenchel conjugate (Definitions 2.3 and C.6),

$$\Gamma(z) = \sup_J \{J^\top \varphi - \log Z(J)\},$$

we obtain the exact form of one-pattern Hopfield free energy.

Proposition B.1 (Free Energy of One-Pattern Hopfield Model [Friedli and Velenik, 2017]). The free energy of a one-pattern Hopfield model is

$$\Gamma(z) = -\frac{1}{2}z^2 - \frac{1}{\beta} \log \left(\frac{d}{\frac{d+\sqrt{d}z}{2}} \right).$$

In the large- d limit, it approaches to an explicit expression

$$\Gamma(z) = \beta(U - TS) \xrightarrow{d \rightarrow \infty} \Gamma_\infty(q) = -\frac{\beta}{2}q^2 - s(q), \quad (\text{B.4})$$

where $U = -\frac{1}{2}\mathbb{E}_{p_J}[z^2]$ is the internal energy, S is the entropy (Definition C.8) in the statistical mechanics language, $q := \lim_{d \rightarrow \infty} \mathbb{E}_{p_J}[z]$ and

$$s(q) := -\frac{d+q\sqrt{d}}{2} \log \frac{1+q/\sqrt{d}}{2} - \frac{d-q\sqrt{d}}{2} \log \frac{1-q/\sqrt{d}}{2}. \quad (\text{B.5})$$

Proof. Please see [Friedli and Velenik, 2017] for the detailed proof. $\square$

Proposition B.2 (Stationary Equation of One-Pattern Hopfield Model [Friedli and Velenik, 2017]). The stationary condition of the large- d one-pattern Hopfield free energy is

$$q^* = \sqrt{d} \tanh \left(\frac{\beta q^*}{\sqrt{d}} \right). \quad (\text{B.6})$$

This is an implicit equation. The solutions depend on inverse-temperature β .

Proof. Taking the derivative, we have

$$0 = \Gamma'_\infty(q^*) = -\beta q^* + \frac{\sqrt{d}}{2} \log \frac{1+q^*/\sqrt{d}}{1-q^*/\sqrt{d}}.$$

This hence gives the stationary points q^* and proves the proposition. $\square$

At the origin, $\Gamma''_\infty(q^* = 0) = -\beta + 1$. Therefore, as illustrated in Figure 7:

- for $\beta < 1$ (high temperature), $q^* = 0$ is the unique minimum. However, $q^* = \lim_{d \rightarrow \infty} \mathbb{E}_{p_J}[z] = 0$ indicates disorderness, i.e., zero overlap.
- at $\beta_c = 1$, the Hessian at the origin vanishes.
- for $\beta > 1$ (low temperature), $q^* = 0$ becomes unstable and two minima $q_\pm^* \neq 0$ appear. Since the local minima are symmetric, the retrieval is still unique and stable.




Figure 7: **One-pattern free-energy landscape in the large- d limit.** The curves show the one-pattern Hopfield model free energy and its local minima. Cooling the system turns the minimum at $q^* = 0$ into two retrieval minima.

The deformation of the energy landscape moves the local minimum. It leads to a phase transition.

B.3 Landscapes, Metastability, and Retrieval Jumps

The one-pattern transition above is continuous: the nonzero minima emerge from $q^* = 0$ continuously. More complicated retrieval landscapes contain several separated local minima. Specifically, two minima coexist when they have equal free energy. If one is higher, then the minimum is metastable since a barrier prevents movement to the lower state. A system in one basin remains there until that basin loses local stability. It moves only when the energy barrier vanishes, which is called a *spinodal* (Definition 4.3). We illustrate this mechanism in Figure 8.




Figure 8: **Phase-transition picture.** (a) A metastable minimum is protected by a barrier. (b) We plot two smooth branches. In one smooth branch, the minimal Hessian eigenvalue vanishes at a spinodal. This leads to a jump onto another branch.

C Definitions and Proofs in the Main Text

C.1 Basic Tools for Convex Analysis

Here we define several fundamental terminologies from standard convex analysis [Rockafellar, 1997].

Definition C.1 (Proper). A function f is proper if the support domain of f satisfies

$$\text{dom}(f) := \{x \in \mathbb{R}^d : f(x) < \infty\} \neq \emptyset, \quad \text{and } f \neq +\infty.$$

Definition C.2 (Lower Semi-Continuous (lsc)). A function f is lower semi-continuous (lsc) if

$$f(x) \leq \liminf_{\tau \rightarrow \infty} f(x_\tau),$$

for every sequence $(x_i)_\tau \rightarrow x_i$.

Definition C.3 (Closed Convex). A function f is closed convex if f is proper, lsc and convex.

Definition C.4 (Essentially Smooth). A function f is essentially smooth if f is differentiable on $\text{int dom}(f)$, and $\|\nabla f(x)\| \rightarrow \infty$ whenever $x_i \rightarrow \bar{x}_i \in \partial(\text{dom}(f))$.

Definition C.5 (Subgradient and Subdifferential). Let $f : \mathbb{R}^d \rightarrow [-\infty, +\infty)$ be a proper convex function. A vector $g \in \mathbb{R}^d$ is called a subgradient of f at $x \in \text{dom}(f)$ if

$$f(y) \geq f(x) + \sum_{i=1}^d g_i(y_i - x_i) \quad \forall y \in \mathbb{R}^d.$$

The set of all subgradients of f at x is called the subdifferential of f at x , and is denoted by

$$\partial f(x) = \{g \in \mathbb{R}^d : f(y) \geq f(x) + \sum_{i=1}^d g_i(y_i - x_i) \forall y \in \mathbb{R}^d\}.$$

If f is differentiable at x , then the subdifferential reduces to $\partial f(x) = \{\nabla f(x)\}$.

Definition C.6 (Fenchel Conjugation [Rockafellar, 1997]). Consider a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ and an input $x \in \mathbb{R}^d$. Then we define the Fenchel conjugate of $f(x)$ by the following variational formula,

$$f^*(k) := \sup_{x \in \mathbb{R}^d} \left\{ \sum_{i=1}^d k_i x_i - f(x) \right\} \in \mathbb{R}, \quad k \in \mathbb{R}^d.$$

Fact 1 (Fenchel Conjugation Properties [Rockafellar, 1997]). **Definition C.6** has several properties:

1. **Convexity.** A Fenchel conjugate $f^*(x)$ is by definition a convex function.
2. **Involutivity.** If a function f is closed convex, then applying Fenchel conjugation twice yields the identity map, i.e., $f^{**} = f$ for all $x \in \mathbb{R}^d$. This is the well-known Fenchel-Moreau Theorem.
3. **Bijectiveness.** If a function f is proper, lsc, strictly convex and essentially smooth, then $\nabla f^* : \text{int dom}(f^*) \leftrightarrow \text{int dom}(f)$ is a bijection, where int and dom denote interior and domain.

Definition C.7 (Legendre Transformation). A function f is called a Legendre function if it is proper, lsc, strictly convex and essentially smooth. In other words, Legendre functions satisfy involutivity and bijectiveness. In this case, we refer to the Fenchel conjugate f^* as the Legendre transformation.

Remark C.1. A Legendre function f satisfies the Fenchel-Moreau involutivity $f^{**} = f$, and its gradient map ∇f is a bijection between the interior domains from **Fact 1** automatically.

Remark C.2. Here we adopt the definition of the Legendre transformation from [Rockafellar, 1997]. However, in the physics and machine learning literature, for example, in [Krotov and Hopfield, 2020] and [Ramsauer et al., 2020], people denote the same conjugation without verifying the conditions. Specifically, in many cases, the domain is $\mathbb{R}^d$, so $\partial\mathbb{R}^d = \emptyset$ leads to the essential smoothness condition being vacuous. The real issue is the non-bijectiveness of gradient maps. We should understand the transformations in these works as the Fenchel conjugations rather than the true Legendre transforms.

C.2 Proof of **Proposition 2.1**

Proposition C.1 (**Proposition 2.1** Restated: Fenchel-Young Duality Relation [Srednicki, 2007, Blondel et al., 2020]). Let p_J be the Gibbs measure with external source J . From the standard statistical field theory, the induced macroscopic field $\varphi = \mathbb{E}_{p_J}[x]$ reveals the following source-field relation:

$$\varphi = \mathbb{E}_{x \sim p_J}[x] \in \partial W[J] \Leftrightarrow J \in \partial \Gamma[\varphi],$$

where ∂ is the subdifferential from **Definition C.5**. If W and Γ are differentiable, then ∂ reduce to gradients ∇ . We call them equations of motion (EOMs) in the physics literature. It also refers to the Fenchel-Young duality relation. For φ and J , we call them the classical solutions or saddle points.

Proof. We repeat the proof from [Blondel et al., 2020, Rockafellar, 1997] for completeness. By definition, $\Gamma[\varphi] = \sup_{J \in \mathbb{R}^d} \{\sum_{i=1}^d J_i \varphi_i - W[J]\}$. Therefore, for any J , $\Gamma[\varphi] \geq \sum_{i=1}^d J_i \varphi_i - W[J]$.

By the definition of subgradient [Definition C.5](#), we have

$$W[J] \geq W[J^*] + \sum_{i=1}^d \varphi_i (J_i - J_i^*), \quad \forall J \in \mathbb{R}^d \forall \varphi \in \partial W[J^*].$$

Then

$$\sum_{i=1}^d J_i^* \varphi_i - W[J^*] \geq \sum_{i=1}^d J_i \varphi_i - W[J] \quad \forall J \in \mathbb{R}^d.$$

Hence, J^* is a maximizer, and therefore $J^* \in \partial \Gamma[\varphi^*]$. Conversely, if we assume that $J^* \in \partial \Gamma[\varphi^*]$, then

$$\Gamma[\varphi] \geq \Gamma[\varphi^*] + \sum_{i=1}^d J_i^* (\varphi_i - \varphi_i^*) \quad \forall \varphi \in \mathbb{R}^d.$$

In this way,

$$\sum_{i=1}^d J_i^* \varphi_i^* - \Gamma[\varphi^*] \geq \sum_{i=1}^d J_i^* \varphi_i - \Gamma[\varphi] \quad \forall \varphi \in \mathbb{R}^d.$$

Note that $W[J]$ is closed due to the form [\(2.2\)](#). Then by involutivity in [Fact 1](#), we obtain

$$W[J^*] = W^{**}[J^*] = \sup_{\varphi \in \mathbb{R}^d} \left\{ \sum_{i=1}^d \varphi_i J_i^* - \Gamma[\varphi] \right\} \geq \sum_{i=1}^d \varphi_i J_i^* - \Gamma[\varphi] \quad \forall \varphi \in \mathbb{R}^d.$$

Hence, φ^* is a maximizer, and therefore $\varphi^* \in \partial W[J^*]$. After removing the redundant notation $\star$, we arrive the Fenchel-Young duality relation. $\square$

Remark C.3. The non-zero equation, $\nabla_J W(J)$, reflects a symmetry breaking of a free theory due to the external input. For example, consider a spin system. Every spin has a random orientation, $S \in \{-1, +1\}^d$, and there exists a $\mathbb{Z}_2$ symmetry that flips the entire system upside down. Then,

$$\nabla_J W[J]|_{J=0} = \frac{S \exp\{-\beta \mathcal{H}[S] + \sum_{i=1}^d J_i S_i\}}{Z[J]}|_{J=0} = \mathbb{E}_{S \sim p(S)}[S] = 0, \quad (\text{C.1})$$

where $p(S) \propto \exp\{-\beta \mathcal{H}[S]\}$ up to a normalization factor. This gives the macroscopic variable $\varphi = 0$. However, if we input the external source, $J \neq 0$, then the spin system favors one orientation. This corresponds to a response in a ferromagnet. After this, we approach $J \rightarrow 0^+$ as a stationary condition for $\Gamma[\varphi]$. The system stays at one specific local minima. See [Figure 9](#) for the illustration. We shall see that in the sparse modern Hopfield models ([Section 3](#)), we always require this response, $\varphi = \nabla_J W[J] \neq 0$, to realize the retrieval states.




Figure 9: **Selection of a stationary state by an external source.** This figure illustrates how a source J selects one state from multiple degenerate stationary states of the free energy $\Gamma[\varphi]$. **Left.** When $J = 0$, $\Gamma[\varphi]$ is symmetric and has two degenerate local minima $\pm\varphi^*$. **Middle.** Turning on a small external source $J > 0$ tilts the landscape to $\Gamma[\varphi] - J\varphi$. This breaks the symmetry so that $+\varphi^*$ becomes preferred. Then, the system flows toward the selected basin. **Right.** After the selection, we take the limit $J \rightarrow 0^+$. Although the explicit source is removed, the system remains in the chosen local minimum $+\varphi^*$ satisfying $\nabla_\varphi \Gamma[+\varphi^*] = 0$. This source-selection procedure resolves the degeneracy between degenerate stationary states and specifies the retrieval state.

Remark C.4. From the Fenchel-Young duality relation, [Proposition 2.1](#), we know that a gradient descent trajectory in J -space has one-to-one mapping to a unique trajectory in φ -space. However, the induced unique trajectory is not a gradient flow. Specifically, suppose we define gradient descent as $\dot{\varphi} = \dot{J} = -\nabla_J W[J]$, where $\dot{J} := dJ/dt$ if we view update steps as infinitesimal. Then

$$\dot{\varphi} = \nabla_J^2 W[J](-\nabla_J W[J]) = -\varphi \nabla_J^2 W[J],$$

which is not the gradient flow, $\dot{\varphi} = -\nabla_\varphi \Gamma[\varphi]$.

C.3 Tsallis Entropy

In this section, we introduce the Tsallis entropy and its role in statistical mechanics.

Definition C.8 (Tsallis Entropy [[Tsallis, 1988](#)]). We define the Tsallis entropy, [[Tsallis, 1988](#)], of the form

$$\Psi_\alpha[p] := \frac{1}{\alpha - 1} \left(1 - \sum_{\mu=1}^M p_\mu^\alpha \right), \quad (\text{C.2})$$

where $\sum_{\mu=1}^M p_\mu = 1$, $\alpha \in \mathbb{R}_+ / \{1\}$.

Remark C.5. The maximal Tsallis entropy occurs when the probability function is uniform, i.e., $p_\mu = M^{-1}$ for every μ . Hence, if the system is random, then the entropy becomes maximal.

Proposition C.2 (Types of Entropy [[Tsallis, 1988](#)]). Specific values of α in Tsallis entropy reduce to other types of entropy.

- If $\alpha \rightarrow 1$, it reduces to the Shannon entropy [Shannon, 1948].

$$S[p] := - \sum_{\mu=1}^M p_{\mu} \log(p_{\mu}). \quad (\text{C.3})$$

- If $\alpha \rightarrow 2$, it reduces to the Gini entropy [Gini, 1912].

$$G[p] := 1 - \sum_{\mu=1}^M p_{\mu}^2. \quad (\text{C.4})$$

- If $\alpha \rightarrow \infty$, it is zero.

Proof. We repeat the proof from [Tsallis, 1988] for completeness. We prove them one-by-one as follows.

- Applying the l'Hôpital's rule for 0/0-form, we have

$$\Psi_{\alpha \rightarrow 1}[p] = - \lim_{\alpha \rightarrow 1} \sum_{\mu=1}^M \log(p_{\mu}) p_{\mu}^{\alpha} = - \sum_{\mu=1}^M p_{\mu} \log(p_{\mu}) := S[p].$$

This is the form of Shannon entropy.

- For $\alpha \rightarrow 2$, the Tsallis entropy function is by definition the Gini entropy.
- Let $p_{\max} := \max_{\mu} p_{\mu}$ and $\mu_{\max}$ be the corresponding maximal index. Then we separate the Tsallis entropy as

$$\Psi_{\alpha}[p] = \frac{1}{\alpha - 1} \left[1 - p_{\max}^{\alpha} \left(1 + \sum_{\mu \neq \mu_{\max}}^M \frac{p_{\mu}^{\alpha}}{p_{\max}^{\alpha}} \right) \right].$$

After taking $\alpha \rightarrow \infty$ and applying the l'Hôpital's rule for ∞/∞ -form, we obtain

$$\Psi_{\alpha \rightarrow \infty}[p] = - \lim_{\alpha \rightarrow \infty} \log(p_{\max}) p_{\max}^{\alpha},$$

where $(p_{\mu}/p_{\max})^{\alpha} \rightarrow 0$ as $\alpha \rightarrow \infty$. If $p_{\max} < 1$, then $\Psi_{\alpha \rightarrow \infty} = 0$. If $p_{\max} = 1$, then $\Psi_{\alpha \rightarrow \infty} \rightarrow 0$ due to the logarithm. Hence, $\Psi_{\alpha \rightarrow \infty}[p]$ for all p .

This completes the proof. □

Remark C.6. If $\alpha > 0$, then the Tsallis entropy is strictly concave since its second derivative respect to p is negative on the simplex interior. This includes Shannon entropy ($\alpha \rightarrow 1$), which is also strictly concave.

Remark C.7. Although the Tsallis entropy is not defined at $\alpha = 1$ from Definition C.8, we use the continuous extension $\Psi_1[p] := \lim_{\alpha \rightarrow 1} \Psi_{\alpha}[p]$. It is equal to the Shannon entropy according to Proposition C.2. Therefore, throughout the paper, the $\alpha = 1$ refers to this Shannon/softmax limit.

Proposition C.3 (Fenchel Conjugate of Entropy [Blondel et al., 2020]). The Fenchel conjugate of the negative Tsallis entropy is

$$\Omega_\alpha^*[z] := (-\Psi_\alpha)^*[z] = \sup_{p \in \Delta^M} \left\{ \sum_{\mu=1}^M z_\mu p_\mu + \Psi_\alpha[p] \right\}, \quad (\text{C.5})$$

where $\Delta^M := \{p_\mu \in \mathbb{R}_+^M \mid \sum_{\mu=1}^M p_\mu = 1\}$ is a $(M-1)$ -dimensional unit simplex. Then

- In the Shannon case, $\alpha \rightarrow 1$, the negative entropy Ω_α is a Legendre function.
- For $\alpha > 1$, Ω_α satisfies involutivity, but neither Legendre condition nor bijectiveness.
- For $1 > \alpha > 0$, Ω_α is a Legendre function.

Proof. First, we look at the Shannon case.

- Due to (C.3) and Remark C.6, the Shannon entropy is closed and strictly concave. By Fact 1, it satisfies involutivity, $S^{**}[p] = S[p]$. Furthermore, $\|\nabla_p S[p]\|_{p \rightarrow 0} = \infty$. Thus, it is essentially smooth. The negative Shannon entropy Ω_α is thereby a Legendre function and satisfies bijectiveness.
- For $\alpha > 1$, the Tsallis entropy is closed and strictly concave on the simplex. Therefore, $\Omega_\alpha = -\Psi_\alpha$ satisfies involutivity condition $\Omega_{\alpha>1}^{**}[p] = \Omega_{\alpha>1}[p]$. However, $\|\nabla_p \Omega_{\alpha>1}[p]\|_{p \rightarrow 0}$ does not diverge, so it is not essentially smooth. As a result, it is not a Legendre function nor satisfies bijectiveness.
- For $1 > \alpha > 0$, the Tsallis entropy is strictly concave. Hence, Ω_α is strictly convex. Furthermore, $\partial \Omega_\alpha / \partial p_\mu \rightarrow -\infty$ near the boundary, so it satisfies the essential smoothness. Therefore, it is a Legendre function satisfying the involutivity and bijectiveness properties.

This completes the proof. $\square$

Remark C.8. We are able to express unit simplex condition in two different ways: $1 = \sum_{\mu=1}^M p_\mu = \int_{\mathcal{X}} p_\mu(x) \, d\mu(x)$. That is, we transform the x configuration space, $\mathcal{X}$, onto a simplex, Δ^M .

Definition C.9 (α -Deformed Exponential Family and Logarithm [Tsallis, 1988]). Let $\alpha > 0$. We define the α -deformed exponential function as

$$\exp_\alpha(x) := \begin{cases} \exp(x) & , \text{ if } \alpha = 1 \\ [1 + (1 - \alpha)x]_+^{\frac{1}{1-\alpha}} & , \text{ if } \alpha \neq 1 \end{cases},$$

and its inverse as

$$\log_\alpha(y) := \begin{cases} \log(y) & , \text{ if } \alpha = 1 \\ \frac{y^{1-\alpha} - 1}{1-\alpha} & , \text{ if } \alpha \neq 1 \end{cases},$$

where $x \in \mathbb{R}$, $y \in \mathbb{R}_+$ for well-definedness, and $[z]_+ := \max\{z, 0\}$ for every $z \in \mathbb{R}$.

This α -deformed family has several properties, and some of which differ from the ordinary ones.

Lemma C.1 (Properties of α -Deformed Functions [Tsallis, 1988]). α -deformed functions satisfy

- $\lim_{\alpha \rightarrow 1} \exp_\alpha(x) = \exp(x)$ and $\lim_{\alpha \rightarrow 1} \log_\alpha(x) = \log(x)$.
- $\exp_\alpha(x)/\log_\alpha(y)$ are monotonic and convex/concave.
- α -deformed exponential/logarithm does not have multiplicity/additivity unless $\alpha = 1$.
- $\frac{d}{dx} \exp_\alpha(x) = [\exp_\alpha(x)]^\alpha$ in the domain $[1 + (1 - \alpha)x] > 0$ and $\frac{d}{dy} \log_\alpha(y) = y^{-\alpha}$.

Proof. We repeat the proof from [Tsallis, 1988] for completeness. We prove those properties one-by-one as follows.

- Let $u = 1 - \alpha$. By Definition C.9,

$$\lim_{\alpha \rightarrow 1} \exp_\alpha(x) = \lim_{\alpha \rightarrow 1} [1 + (1 - \alpha)x]_+^{\frac{1}{1-\alpha}} = \lim_{u \rightarrow 0} [1 + ux]^{1/u} = \exp(x) .$$

For the logarithm, we apply the l'Hôpital's rule for 0/0-form.

$$\lim_{\alpha \rightarrow 1} \log_\alpha(y) = \lim_{\alpha \rightarrow 1} \frac{y^{1-\alpha} - 1}{1 - \alpha} = - \lim_{\alpha \rightarrow 1} \log(y) y^{1-\alpha} = \log(y) .$$

- If $[1 + (1 - \alpha)x] > 0$, by differentiating $\exp_\alpha(x)$ with respect to x , we have

$$\frac{d}{dx} \exp_\alpha(x) = [1 + (1 - \alpha)x]^{\frac{\alpha}{1-\alpha}} > 0 , \quad \frac{d^2}{dx^2} \exp_\alpha(x) = \alpha [1 + (1 - \alpha)x]^{\frac{2\alpha-1}{1-\alpha}} > 0 .$$

Hence, it is strictly increasing (monotonic) and strictly convex. For $\log_\alpha$,

$$\frac{d}{dy} \log_\alpha(y) = y^{-\alpha} > 0 , \quad \frac{d^2}{dy^2} \log_\alpha(y) = -\alpha y^{-\alpha-1} < 0 .$$

Hence, it is strictly increasing (monotonic) and strictly concave.

- Consider $x_1, x_2 \in \mathbb{R}$.

$$\begin{aligned} \exp_\alpha(x_1) \exp_\alpha(x_2) &= [1 + (1 - \alpha)(x_1 + x_2) + (1 - \alpha)^2 x_1 x_2]_+^{\frac{1}{1-\alpha}} \\ &= \exp_\alpha(x_1 + x_2) \cdot \left[1 + \frac{(1 - \alpha)^2 x_1 x_2}{1 + (1 - \alpha)(x_1 + x_2)}\right]_+^{\frac{1}{1-\alpha}} . \end{aligned} \quad (\text{C.6})$$

If $\alpha \rightarrow 1$, then $[\cdot \cdot \cdot]$ in (C.6) is 1. Consider $y_1, y_2 \in \mathbb{R}_+$.

$$\begin{aligned} \log_\alpha(y_1 y_2) &= \frac{y_1^{1-\alpha} y_2^{1-\alpha} - 1}{1 - \alpha} = \frac{(y_1^{1-\alpha} - 1) + (y_2^{1-\alpha} - 1) + (y_1^{1-\alpha} - 1)(y_2^{1-\alpha} - 1)}{1 - \alpha} \\ &= \log_\alpha(y_1) + \log_\alpha(y_2) + (1 - \alpha) \log_\alpha(y_1) \log_\alpha(y_2) . \end{aligned} \quad (\text{C.7})$$

If $\alpha \rightarrow 1$, the last term in (C.7) is zero.

- For $[1 + (1 - \alpha)x] > 0$, we have

$$\frac{d}{dx} \exp_\alpha(x) = [1 + (1 - \alpha)x]^{\frac{1}{1-\alpha}-1} = [\exp_\alpha(x)]^\alpha.$$

For $\log_\alpha(y)$, we obtain the result by differentiation easily. For $\alpha \rightarrow 1$, they hold trivially.

This completes the proof. $\square$

Remark C.9. Since $p_\mu^\alpha \log_\alpha(p_\mu) = (1 - \alpha)^{-1}(p_\mu - p_\mu^\alpha)$, using $\sum_{\mu=1}^M p_\mu = 1$, we obtain

$$\Psi_\alpha[p] = - \sum_{\mu=1}^M p_\mu^\alpha \log_\alpha(p_\mu). \quad (\text{C.8})$$

C.4 Proof of Theorem 2.1

Theorem C.1 (Theorem 2.1 Restated: Tsallis Entropy and Deformed Free Energy). If we replace the exponential functions and logarithm functions by α -deformed functions, then we have the deformed log-partition function

$$-W_\alpha[J] = \beta \mathbb{E}_{x \sim p_J}^{(\alpha)}[\mathcal{H}[x]] - \frac{\Psi_\alpha[p_J]}{Z_\alpha^{\alpha-1}[J] \int_{\mathcal{X}} p_J^\alpha(x) d\mu(x)} - \sum_{i=1}^d J_i \mathbb{E}_{x \sim p_J}^{(\alpha)}[x_i],$$

where $\mathbb{E}_{x \sim p_J}^{(\alpha)}[\cdot] := \int_{\mathcal{X}} p_J^\alpha(x) [\cdot] d\mu(x) / \int_{\mathcal{X}} p_J^\alpha(x) d\mu(x)$ is the escort expectation value, and

$$\Gamma_\alpha = \beta \mathbb{E}_{x \sim p_{J^*}}^{(\alpha)}[\mathcal{H}[x]] - \frac{\Psi_\alpha[p_{J^*}]}{Z_\alpha^{\alpha-1}[J^*] \int_{\mathcal{X}} p_{J^*}^\alpha(x) d\mu(x)}, \quad \text{with } \varphi = \mathbb{E}_{x \sim p_{J^*}}^{(\alpha)}[x], \quad (\text{C.9})$$

where $J^* = \arg \max_J \{\sum_{i=1}^d J_i \varphi_i - W_\alpha[J]\}$. We assume $\mathcal{X}$ is finite-volume to avoid divergence.

Proof. By Definition C.9, we express $W_\alpha[J] = \log_\alpha Z_\alpha[J]$ and

$$p_J(x) = \frac{\exp_\alpha\{-\beta \mathcal{H}[x] + \sum_{i=1}^d J_i x_i\}}{Z_\alpha[J]}, \quad Z_\alpha[J] = \int_{\mathcal{X}} \exp_\alpha\{-\beta \mathcal{H}[x] + \sum_{i=1}^d J_i x_i\} d\mu(x).$$

Notice that since $\mathcal{X}$ is finite-volume, $Z_\alpha[J]$ is finite for appropriate J and the following computation is reliable. Then we have

$$\begin{aligned} \log_\alpha p_J(x) &= \log_\alpha \left(\frac{\exp_\alpha\{-\beta \mathcal{H}[x] + \sum_{i=1}^d J_i x_i\}}{Z_\alpha[J]} \right) \\ &= \frac{Z_\alpha^{\alpha-1}[J] \exp_\alpha^{1-\alpha}\{-\beta \mathcal{H}[x] + \sum_{i=1}^d J_i x_i\} - 1}{1 - \alpha} && (\text{By Definition C.9}) \\ &= \frac{Z_\alpha^{\alpha-1}[J] \{1 + (1 - \alpha)(-\beta \mathcal{H}[x] + \sum_{i=1}^d J_i x_i)\} - 1}{1 - \alpha} && (\text{By Definition C.9}) \end{aligned}$$

$$\begin{aligned}
&= Z_\alpha^{\alpha-1}[J](-\beta\mathcal{H}[x] + \sum_{i=1}^d J_i x_i) + \frac{Z_\alpha^{\alpha-1}[J] - 1}{1 - \alpha} \\
&= Z_\alpha^{\alpha-1}[J](-\beta\mathcal{H}[x] + \sum_{i=1}^d J_i x_i) - Z_\alpha^{\alpha-1}[J] \cdot \left(\frac{Z_\alpha^{1-\alpha}[J] - 1}{1 - \alpha} \right) \\
&= Z_\alpha^{\alpha-1}[J](-\beta\mathcal{H}[x] + \sum_{i=1}^d J_i x_i - W_\alpha[J]). \quad (\text{By Definition C.9 and } W_\alpha[J] = \log_\alpha Z_\alpha[J])
\end{aligned}$$

By (C.8), we express the Tsaliis entropy as

$$\begin{aligned}
\Psi_\alpha[p_J] &= - \int_{\mathcal{X}} p_J^\alpha(x) \log_\alpha(p_J(x)) \, d\mu(x) \\
&= - \int_{\mathcal{X}} p_J^\alpha(x) Z_\alpha^{\alpha-1}[J](-\beta\mathcal{H}[x] + \sum_{i=1}^d J_i x_i - W_\alpha[J]) \, d\mu(x). \tag{C.10}
\end{aligned}$$

This gives

$$\begin{aligned}
-W_\alpha[J] &= \frac{-\Psi_\alpha[p_J] + \beta \int_{\mathcal{X}} p_J^\alpha(x) Z_\alpha^{\alpha-1}[J] \mathcal{H}[x] \, d\mu(x) - \int_{\mathcal{X}} \sum_{i=1}^d p_J^\alpha(x) Z_\alpha^{\alpha-1}[J] J_i x_i \, d\mu(x)}{Z_\alpha^{\alpha-1}[J] \int_{\mathcal{X}} p_J^\alpha(x) \, d\mu(x)} \\
&= \beta \frac{\int_{\mathcal{X}} p_J^\alpha(x) \mathcal{H}[x] \, d\mu(x)}{\int_{\mathcal{X}} p_J^\alpha(x) \, d\mu(x)} - \frac{\Psi_\alpha[p_J]}{Z_\alpha^{\alpha-1}[J] \cdot \int_{\mathcal{X}} p_J^\alpha(x) \, d\mu(x)} - \sum_{i=1}^d J_i \frac{\int_{\mathcal{X}} p_J^\alpha(x) x_i \, d\mu(x)}{\int_{\mathcal{X}} p_J^\alpha(x) \, d\mu(x)}.
\end{aligned}$$

The first and last term are the escort expectation values $\mathbb{E}_{x \sim p_J}^{(\alpha)}[\cdot]$. Hence, it is the form of our first equation. After applying the Fenchel conjugation, we have

$$\begin{aligned}
\Gamma_\alpha &= \sup_J \left\{ \sum_{i=1}^d J_i \varphi_i - W_\alpha[J] \right\} \\
&= \sup_J \left\{ \sum_{i=1}^d J_i (\varphi_i - \mathbb{E}_{x \sim p_J}^{(\alpha)}[x_i]) + \beta \mathbb{E}_{x \sim p_J}^{(\alpha)}[\mathcal{H}[x]] - \frac{\Psi_\alpha[p_J]}{Z_\alpha^{\alpha-1}[J] \cdot \int_{\mathcal{X}} p_J^\alpha(x) \, d\mu(x)} \right\} \\
&= \beta \mathbb{E}_{x \sim p_{J^*}}^{(\alpha)}[\mathcal{H}[x]] - \frac{\Psi_\alpha[p_{J^*}]}{Z_\alpha^{\alpha-1}[J^*] \cdot \int_{\mathcal{X}} p_{J^*}^\alpha(x) \, d\mu(x)},
\end{aligned}$$

where $\varphi = \mathbb{E}_{x \sim p_{J^*}}^{(\alpha)}[x]$. This proves the theorem. $\square$

Remark C.10. The finite-volume assumption is only to ensure that the α -deformed partition function is well-defined. Specifically, if we consider whole $\mathbb{R}^d$ with $\alpha = 2$ and $a > 0$, then

$$\int_{\mathbb{R}^d} \exp_2\{-a\|x\|_2^2\} \, d\mu(x) = \int_{\mathbb{R}^d} \frac{d\mu(x)}{1 + a\|x\|_2^2} = S_{d-1} \int_0^\infty \frac{r^{d-1}}{1 + ar^2} \, dr,$$

where S_{d-1} is the surface area of a unity sphere S^{d-1} . Hence, only if $d < 2$ is the integral

finite. This leads to the consideration of a finite-volume domain $\mathcal{X}$. Passing to a finite-volume domain $\mathcal{X} \subset \mathbb{R}^d$ introduces a boundary. It removes the divergence but does not guarantee strict convexity of $W_\alpha[J]$ nor the essential smoothness at the boundary of $\text{dom} W_\alpha$. Hence, according to [Definition C.7](#), the transformation here should be understood as a Fenchel conjugate on $\text{dom} W_\alpha$, not as a Legendre transform necessarily.

C.5 Proof of [Proposition 3.1](#)

Proposition C.4 ([Proposition 3.1](#) Restated: Query-Probability Relation). A query state vector relates to a probability function via

$$x_i = \frac{1}{\sqrt{d}} \sum_{\mu=1}^M \xi_i^\mu p_\mu + \frac{1}{\beta} J_i.$$

Proof. In the deterministic limit, we have

$$Z_\alpha[J] = \int_{\mathcal{X}} \exp_\alpha \left\{ -\beta \tilde{\mathcal{H}}_{\text{MHM}}[x, p, \xi] + \sum_{i=1}^d J_i x_i \right\} d\mu(x) = \exp_\alpha \left\{ -\beta \tilde{\mathcal{H}}_{\text{MHM}}[x, p, \xi] + \sum_{i=1}^d J_i x_i \right\},$$

since we identify x with the mean. The query state vector that maximizes $Z_\alpha[J]$ is the deterministic x itself. Hence, we have

$$\begin{aligned} 0 &= \frac{\partial}{\partial x_i} \left\{ \beta \tilde{\mathcal{H}}_{\text{MHM}}[x, p, \xi] - \sum_{j=1}^d J_j x_j \right\} \\ &= \beta x_i - \frac{\beta}{\sqrt{d}} \sum_{\mu=1}^M \xi_i^\mu p_\mu - J_i. \end{aligned}$$

This proves the proposition. □

C.6 Proof of [Theorem 3.1](#)

Proposition C.5 ([Theorem 3.1](#) Restated: Free Energy of MHM). The joint free energy of the MHM is

$$\tilde{\Gamma}_\alpha[x, p, \xi] = \frac{\beta}{2} \sum_{i=1}^d x_i^2 - \beta \sum_{\mu=1}^M z_\mu p_\mu.$$

Hence,

$$\Gamma_\alpha[x, \xi] = \frac{\beta}{2} \sum_{i=1}^d x_i^2 - \beta \max_{\mu \in [1, M]} z_\mu.$$

Proof. The form of the partition function is

$$Z_\alpha[J, \xi] = \exp_\alpha\{W_\alpha[J, \xi]\} = \int_{\mathcal{X}} \int_{\Delta^M} \exp_\alpha\left\{-\frac{\beta}{2} \sum_{i=1}^d x_i^2 + \beta \sum_{\mu=1}^M z_\mu p_\mu + \sum_{i=1}^d J_i x_i\right\} d\mu(p) d\mu(x).$$

According to the form of the Tsallis entropy (C.8) and (C.10), we have

$$\begin{aligned} \Psi_\alpha[\rho_J] &= - \int_{\mathcal{X}} \int_{\Delta^M} \rho_J^\alpha(x, p) \log_\alpha[\rho_J(x, p)] d\mu(p) d\mu(x) \\ &= \frac{\beta}{2} Z_\alpha^{\alpha-1}[J, \xi] \sum_{i=1}^d \int_{\mathcal{X}} \int_{\Delta^M} \rho_J^\alpha(x, p) x_i^2 d\mu(p) d\mu(x) \\ &\quad - \beta Z_\alpha^{\alpha-1}[J, \xi] \sum_{\mu=1}^M \int_{\mathcal{X}} \int_{\Delta^M} \rho_J^\alpha(x, p) p_\mu z_\mu(x) d\mu(p) d\mu(x) \\ &\quad - Z_\alpha^{\alpha-1}[J, \xi] \sum_{i=1}^d J_i \int_{\mathcal{X}} \int_{\Delta^M} \rho_J^\alpha(x, p) x_i d\mu(p) d\mu(x) \\ &\quad + Z_\alpha^{\alpha-1}[J, \xi] W_\alpha[J, \xi] \int_{\mathcal{X}} \int_{\Delta^M} \rho_J^\alpha(x, p) d\mu(p) d\mu(x). \end{aligned}$$

It yields

$$\begin{aligned} -W_\alpha[J, \xi] &= \frac{\beta}{2} \sum_{i=1}^d \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[x_i^2] - \beta \sum_{\mu=1}^M \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[z_\mu p_\mu] \\ &\quad - \frac{\Psi_\alpha[\rho_J]}{Z_\alpha^{\alpha-1}[J, \xi] \int_{\mathcal{X}} \int_{\Delta^M} \rho_J^\alpha(x, p) d\mu(p) d\mu(x)} - \sum_{i=1}^d J_i \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[x_i]. \end{aligned} \quad (\text{C.11})$$

Since x and p are fixed, so the probability density function concentrates to an indicator, i.e.,

$$\rho_J(x, p) \rightarrow \mathbb{1}_{\{x, \tilde{\mathbb{E}}_{x \sim p_J}^{(\alpha)}[x]\}} \mathbb{1}_{\{p, \tilde{\mathbb{E}}_{p \sim p_{\Delta, J}}^{(\alpha)}[p]\}},$$

where

$$\tilde{\mathbb{E}}_{x \sim p_J}^{(\alpha)}[x] := \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[x] = \frac{\int_{\mathcal{X}} C(x) x p_J^\alpha(x) d\mu(x)}{\int_{\mathcal{X}} C(x) p_J^\alpha(x) d\mu(x)}, \quad \text{with } C(x) := \int_{\Delta^M} \rho_J^\alpha(p|x) d\mu(p), \quad (\text{C.12})$$

and

$$\tilde{\mathbb{E}}_{p \sim p_{\Delta, J}}^{(\alpha)}[p_\mu] := \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[p_\mu] = \frac{\int_{\Delta^M} C(p) p_\mu p_{\Delta, J}^\alpha(p) d\mu(p)}{\int_{\Delta^M} C(p) p_{\Delta, J}^\alpha(p) d\mu(p)}, \quad \text{with } C(p) := \int_{\mathcal{X}} \rho_J^\alpha(x|p) d\mu(x). \quad (\text{C.13})$$

Then

$$Z_\alpha[J, \xi] \rightarrow \exp_\alpha \{ -\beta \tilde{\mathcal{H}}_{\text{MHM}}[\tilde{\mathbb{E}}_{x \sim p_J}^{(\alpha)}[x], \tilde{\mathbb{E}}_{p \sim p_{\Delta, J}}^{(\alpha)}[p], \xi] + \sum_{i=1}^d J_i \tilde{\mathbb{E}}_{x \sim p_J}^{(\alpha)}[x_i] \}.$$

Hence, by the Fenchel conjugate, we obtain

$$\begin{aligned} \tilde{\Gamma}_\alpha &= \sup_{J \in \text{dom} W_\alpha} \left\{ \sum_{i=1}^d J_i \varphi_i - W_\alpha[J, \xi] \right\} \\ &= \sup_{J \in \text{dom} W_\alpha} \left\{ \sum_{i=1}^d J_i (\varphi_i - \tilde{\mathbb{E}}_{x \sim p_J}^{(\alpha)}[x_i]) + \frac{\beta}{2} \sum_{i=1}^d (\tilde{\mathbb{E}}_{x \sim p_J}^{(\alpha)}[x_i^2]) - \beta \sum_{\mu=1}^M \tilde{\mathbb{E}}_{x \sim p_J}^{(\alpha)}[z_\mu] \tilde{\mathbb{E}}_{p \sim p_{\Delta, J}}^{(\alpha)}[p_\mu] \right. \\ &\quad \left. - \frac{\Psi_\alpha[\mathbf{1}]}{\exp_\alpha^{\alpha-1} \{ -\beta \tilde{\mathcal{H}}_{\text{MHM}}[\tilde{\mathbb{E}}_{x \sim p_J}^{(\alpha)}[x], \tilde{\mathbb{E}}_{p \sim p_{\Delta, J}}^{(\alpha)}[p], \xi] + \sum_{i=1}^d J_i \tilde{\mathbb{E}}_{x \sim p_J}^{(\alpha)}[x_i] \}} \right\} \\ &= \frac{\beta}{2} \sum_{i=1}^d (\tilde{\mathbb{E}}_{x \sim p_{J^*}}^{(\alpha)}[x_i^2]) - \beta \sum_{\mu=1}^M \tilde{\mathbb{E}}_{x \sim p_{J^*}}^{(\alpha)}[z_\mu] \tilde{\mathbb{E}}_{p \sim p_{\Delta, J^*}}^{(\alpha)}[p_\mu], \end{aligned} \quad (\text{By } \Psi_\alpha[\mathbf{1}] = 0)$$

where J^* is the maximizing source. Since we identify $\tilde{\mathbb{E}}_{x \sim p_{J^*}}^{(\alpha)}[x]$ with x , $\tilde{\mathbb{E}}_{x \sim p_{J^*}}^{(\alpha)}[x^2]$ with x^2 , and $\tilde{\mathbb{E}}_{p \sim p_{\Delta, J^*}}^{(\alpha)}[p]$ with p , we have

$$\tilde{\Gamma}_\alpha[x, p, \xi] = \frac{\beta}{2} \sum_{i=1}^d x_i^2 - \beta \sum_{\mu=1}^M z_\mu p_\mu.$$

Notice that

$$\inf_{p \in \Delta^M} \left\{ - \sum_{\mu=1}^M z_\mu p_\mu \right\} = - \max_{\mu \in [1, M]} z_\mu.$$

After substituting into the joint free energy to eliminate the dependence of p , we prove the second equation in this proposition. $\square$

C.7 Proof of Proposition 3.3

Proposition C.6 (Proposition 3.3 Revisited: Emergent Attention Mechanism as EOMs in MHMs). Attention emerges as follows.

- If $\alpha \rightarrow \infty$, it corresponds to a hardmax function.

$$p_\mu = \text{hardmax}(z_\mu) = \begin{cases} 1, & \text{for maximal } z_\mu \\ 0, & \text{otherwise} \end{cases}.$$

- If $\alpha = 2$, it corresponds to a sparsemax function [Martins and Astudillo, 2016].

$$p_\mu = \text{Sparsemax}_\beta(z_\mu) = \frac{\beta}{2}[z_\mu - \theta]_+, \quad \text{with } \frac{\beta}{2} \sum_{\mu=1}^M [z_\mu - \theta]_+ = 1.$$

- If $1 < \alpha < 2$, it corresponds to an entmax family [Peters et al., 2019].
- If $\alpha \rightarrow 1$, it corresponds to a softmax function [Blondel et al., 2020].

$$p_\mu = \text{Softmax}_\beta(z_\mu) = \frac{\exp\{\beta z_\mu\}}{\sum_{\nu=1}^M \exp\{\beta z_\nu\}}.$$

Proof. We include the proof for completeness, adapting the standard argument from [Martins and Astudillo, 2016, Peters et al., 2019, Blondel et al., 2020] to our notation. The training trajectory follows $p_\mu = \arg \max_{p \in \Delta^M} \{\beta \sum_{\mu=1}^M z_\mu p_\mu + \Psi_\alpha[p]\}$. We derive the physical meanings of this optimization problem via the Lagrangian function:

$$\frac{\partial}{\partial p_\mu} \left\{ \sum_{\nu=1}^M z_\nu p_\nu + \frac{1}{\beta} \cdot \frac{1}{\alpha - 1} (1 - \sum_{\nu=1}^M p_\nu^\alpha) + \theta (1 - \sum_{\nu=1}^M p_\nu) + \sum_{\nu=1}^M \lambda_\nu p_\nu \right\} = 0,$$

where we use the form of Tsallis entropy, (C.2), and the Lagrangian multipliers, λ and $\lambda_\nu \geq 0$, to include the simplex constraints. This gives the complementary slackness, $\lambda_\mu p_\mu = 0$, and the stationary condition,

$$z_\mu - \frac{\alpha}{\beta(\alpha - 1)} p_\mu^{\alpha-1} - \theta + \lambda_\mu = 0 \implies \begin{cases} p_\mu = \{\beta(1 - \frac{1}{\alpha})[z_\mu - \theta]_+\}^{\frac{1}{\alpha-1}} & , \text{ if } \lambda_\mu = 0 \\ p_\mu = 0 & , \text{ if } \lambda_\mu > 0 \end{cases},$$

where $[u]_+ := \max(u, 0)$. For $\alpha \rightarrow \infty$ and $\alpha = 2$, we obtain the forms of p_μ by inserting α directly. For $\alpha = 1$, we use the Shannon form, (C.3). Hence, for non-zero p_μ , we have

$$z_\mu + \frac{1}{\beta} \frac{\partial}{\partial p_\mu} \left(- \sum_{\nu=1}^M p_\nu \log p_\nu \right) - \theta = 0.$$

This gives

$$p_\mu = \exp\{\beta(z_\mu - \theta) - 1\}.$$

The normalization condition gives $1 = \sum_{\mu=1}^M p_\mu = \exp\{-\beta\theta\} \sum_{\mu=1}^M \exp\{\beta z_\mu - 1\}$. This yields $\theta = \beta^{-1} \log \sum_{\mu=1}^M \exp\{\beta z_\mu - 1\}$. After plugging this into p_μ , we obtain the softmax function. For $2 > \alpha > 1$, $(\alpha - 1)^{-1} > 1$. Hence, p_μ is the entmax family interpolating between softmax and sparsemax. $\square$

C.8 Proof of Proposition 4.1

Proposition C.7 (Proposition 4.1 Restated: Temperature Coexistence Point for Sparse Model). Let a, b be two distinct branches and $\alpha = 2$. Denote k_a and k_b as the exact support sizes (see Lemma F.1) of branches a and b , respectively. Based on the statistical mechanics framework, the coexistence point satisfies

$$\sum_{\mu, \nu \in a} [z_\mu^* - \theta]_+ [z_\nu^* - \theta]_+ K_{\mu\nu}(\beta_{\text{coexist}}; \beta_0) = \sum_{\mu, \nu \in b} [z_\mu^* - \theta]_+ [z_\nu^* - \theta]_+ K_{\mu\nu}(\beta_{\text{coexist}}; \beta_0),$$

where $K_{\mu\nu}(\beta) := \mathbb{1}_{\{\mu, \nu\}} - \frac{\beta}{2} G_{\mu\nu}$, and $G_{\mu\nu} := \sum_{i=1}^d \xi_i^\mu \xi_i^\nu / d$ is the Gram matrix. β_{coexist} does not admit a closed-form expression. However, if the distribution is uniform on the simplex and the system satisfies the well-separation condition, i.e.,

$$G_{\mu\nu} = \begin{cases} 1 & , \mu = \nu \\ \epsilon \ll 1 & , \mu \neq \nu \end{cases},$$

then for $k_a = 1$ and $k_b = 2$, we have the analytic solution $\beta_{\text{coexist}} = 2/(1 - \epsilon)$. At this point, $\beta = \beta_{\text{coexist}}$, p^* is degenerate and every p^* is a global minimum of the joint free energy $\tilde{\Gamma}_{\alpha=2}^{(\text{mod})}$.

Proof. Consider a system with the minimizer, p_μ^* . Using that at a stable point, $J_i = 0$ for all $i = 1, \dots, d$, we have

$$z_\mu^* = \frac{1}{\sqrt{d}} \sum_{i=1}^d \xi_i^\mu x_i^* = \frac{1}{\sqrt{d}} \sum_{i=1}^d \xi_i^\mu \left(\frac{1}{\sqrt{d}} \sum_{\nu=1}^M \xi_i^\nu p_\nu^* + \frac{1}{\beta} J_i \right) = \sum_{\nu=1}^M G_{\mu\nu} p_\nu^*.$$

Moreover,

$$\sum_{i=1}^d (x_i^*)^2 = \frac{1}{d} \sum_{i=1}^d \sum_{\mu, \nu=1}^M \xi_i^\mu \xi_i^\nu p_\mu^* p_\nu^* = \sum_{\mu, \nu=1}^M G_{\mu\nu} p_\mu^* p_\nu^*.$$

Suppose the support is k , we denote the indices of the non-zero components as $\tilde{\mu} = 1, \dots, k$. Then we have

$$\begin{aligned} \tilde{\Gamma}_\alpha^{(\text{mod})}[x^*, p^*, \xi] &= \frac{\beta}{2} \sum_{i=1}^d (x_i^*)^2 - \beta \sum_{\tilde{\mu}=1}^k z_{\tilde{\mu}}^* p_{\tilde{\mu}}^* - \left(1 - \sum_{\tilde{\mu}=1}^k (p_{\tilde{\mu}}^*)^2\right) \\ &= -1 - \sum_{\tilde{\mu}, \tilde{\nu}=1}^k p_{\tilde{\mu}}^* \left(\frac{\beta}{2} G_{\tilde{\mu}\tilde{\nu}} - \mathbb{1}_{\{\tilde{\mu}, \tilde{\nu}\}} \right) p_{\tilde{\nu}}^* \end{aligned}$$

From the uniform and well-separation condition, we write $G_{\tilde{\mu}\tilde{\nu}} = (1 - \epsilon) \mathbb{1}_{\{\tilde{\mu}, \tilde{\nu}\}} + \epsilon$. Then this

gives $\sum_{\tilde{\mu}, \tilde{\nu}=1}^k p_{\tilde{\mu}}^* G_{\tilde{\mu}\tilde{\nu}} p_{\tilde{\nu}}^* = (1 - \epsilon) \|p^*\|_2^2 + \epsilon$. In this way, we have

$$\tilde{\Gamma}_{\alpha}^{(\text{mod})}[x^*, p^*, \xi] = \left(1 - \frac{\beta}{2}(1 - \epsilon)\right) \|p^*\|^2 - 1 - \frac{\beta}{2}\epsilon. \quad (\text{C.14})$$

By **Lemma F.1**, we express the minimizer as $p_{\tilde{\mu}}^* = \frac{\beta}{2}(z_{\tilde{\mu}}^* - \theta)$ for $\alpha = 2$. By matching the free energy on branch a and branch b , we prove the first equation. For the well-separation and uniform distribution conditions, we have

$$\tilde{\Gamma}_{\alpha}^{(\text{mod})}[x_a^*, p^*, \xi] = -1 - \frac{\beta_{\text{coexist}}}{2} \times 1 + 1 = -\frac{\beta_{\text{coexist}}}{2}.$$

For branch b ,

$$\tilde{\Gamma}_{\alpha}^{(\text{mod})}[x_b^*, p^*, \xi] = -1 - \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \end{pmatrix} \begin{pmatrix} \frac{\beta_{\text{coexist}}}{2} - 1 & \frac{\beta_{\text{coexist}}\epsilon}{2} \\ \frac{\beta_{\text{coexist}}\epsilon}{2} & \frac{\beta_{\text{coexist}}}{2} - 1 \end{pmatrix} \begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix} = -\frac{1}{2} - \frac{\beta_{\text{coexist}}}{4}(1 + \epsilon).$$

After equating these two energies, we obtain $\beta_{\text{coexist}} = 2/(1 - \epsilon)$. Notice that when $\beta = \beta_{\text{coexist}}$, all p^* are degenerate, i.e., they are all global minima in free energy (C.14). This proves the proposition. $\square$

C.9 Proof of Proposition 4.2

Proposition C.8 (**Proposition 4.2** Restated: Temperature Spinodal Point for Sparse Model). Consider the sparse case $\alpha = 2$ and a branch a with exact support size k . Then the spinodal critical value is

$$\beta_{\text{jump}} = \frac{2}{\lambda_{\max}(PGP)},$$

where $P_{\mu\nu} := \mathbb{1}_{\{\mu, \nu\}} - k^{-1}$, and G is the Gram matrix. If the distribution is uniform and the system satisfies the well-separation condition, (4.3), then we have the analytic solution $\beta_{\text{jump}} = 2/(1 - \epsilon)$.

Proof. Since for $\alpha = 2$, the minimizer satisfies $p^* = \frac{\beta}{2}(z_{\mu}^* - \theta)$, where $\mu = 1, \dots, k$, we have

$$\frac{\partial p_{\mu}^*}{\partial z_{\nu}^*} = \frac{\beta}{2} \mathbb{1}_{\{\mu, \nu\}} - \frac{\beta}{2} \frac{\partial}{\partial z_{\nu}^*} \left(\frac{\sum_{\rho=1}^k z_{\rho}^* - 1}{k} \right) = \frac{\beta}{2} (\mathbb{1}_{\{\mu, \nu\}} - \frac{1}{k}) = \frac{\beta}{2} P_{\mu\nu}.$$

Then the gradient of the free energy at the local stable point, x^* , is

$$\begin{aligned} \frac{\partial}{\partial x_i} \Gamma_{\alpha}^{(\text{mod})}[x^*, \xi] &= \beta x_i^* - \frac{\partial}{\partial x_i} \sup_{p \in \Delta^M} \left\{ \beta \sum_{\mu=1}^k z_{\mu} p_{\mu} + 1 - \sum_{\mu=1}^k p_{\mu}^2 \right\} \Big|_{x=x^*} \\ &= \beta x_i^* - \frac{\beta}{\sqrt{d}} \sum_{\mu=1}^k \xi_i^{\mu} p_{\mu}^*. \end{aligned}$$

The Hessian is

$$\begin{aligned} \frac{\partial^2}{\partial x_i \partial x_j} \Gamma_\alpha^{(\text{mod})}[x^*, \xi] &= \beta \mathbf{1}_{\{i,j\}} - \frac{\beta}{\sqrt{d}} \sum_{\mu, \nu=1}^k \xi_i^\mu \frac{\partial p_\mu^*}{\partial z_\nu^*} \frac{\partial z_\nu^*}{\partial x_j} \Big|_{x=x^*} \\ &= \beta \mathbf{1}_{\{i,j\}} - \frac{\beta^2}{2d} \sum_{\mu, \nu=1}^k \xi_i^\mu \xi_j^\nu P_{\mu\nu}. \end{aligned} \quad (\text{C.15})$$

The smallest Hessian eigenvalue vanishes when $\frac{\beta_{\text{jump}}}{2d} \lambda_{\max}(\sum_{\mu, \nu=1}^k \xi_i^\mu P_{\mu\nu} \xi_j^\nu) = 1$. Notice that $P_{\mu\nu}$ is a projector, i.e.,

$$\sum_{\nu=1}^k P_{\mu\nu} P_{\nu\rho} = P_{\mu\rho}.$$

Suppose u is the eigenvector such that $\sum_{j=1}^d \sum_{\mu, \nu=1}^k \xi_i^\mu P_{\mu\nu} \xi_j^\nu u_j = \lambda u_i$. We construct

$$C_{\mu\nu} := \sum_{\rho, \sigma=1}^k P_{\mu\rho} \left(\sum_{i=1}^d \xi_i^\rho \xi_i^\sigma \right) P_{\sigma\nu} = d \sum_{\rho, \sigma=1}^k P_{\mu\rho} G_{\rho\sigma} P_{\sigma\nu}.$$

In fact, its eigenvector is $v_\mu = \sum_{\nu=1}^k \sum_{i=1}^d P_{\mu\nu} \xi_i^\nu u_i$. In this way,

$$\begin{aligned} (Cv)_\mu &= d \sum_{\nu, \rho, \sigma=1}^k \sum_{\kappa=1}^k \sum_{i=1}^d P_{\mu\rho} G_{\rho\sigma} P_{\sigma\nu} P_{\nu\kappa} \xi_i^\kappa u_i \\ &= \sum_{\rho, \sigma=1}^k \sum_{\kappa=1}^k \sum_{i,j=1}^d P_{\mu\rho} \xi_j^\rho \xi_j^\sigma P_{\sigma\kappa} \xi_i^\kappa u_i && (\text{By } P \text{ being a projector}) \\ &= \sum_{\rho=1}^k \sum_{j=1}^d P_{\mu\rho} \xi_j^\rho \lambda u_j && (\text{By } \sum_{j=1}^d \sum_{\mu, \nu=1}^k \xi_i^\mu P_{\mu\nu} \xi_j^\nu u_j = \lambda u_i) \\ &= \lambda v_\mu. \end{aligned}$$

Hence, its eigenvalue is the same as which of $\sum_{\mu, \nu=1}^k \xi_i^\mu P_{\mu\nu} \xi_j^\nu$. That is,

$$\lambda_{\max} \left(\sum_{\mu, \nu=1}^k \xi_i^\mu P_{\mu\nu} \xi_j^\nu \right) = d \cdot \lambda_{\max} \left(\sum_{\rho, \sigma=1}^k P_{\mu\rho} G_{\rho\sigma} P_{\sigma\nu} \right).$$

Finally, we obtain

$$\beta_{\text{jump}} = \frac{2d}{\lambda_{\max}(\sum_{\mu, \nu=1}^k \xi_i^\mu P_{\mu\nu} \xi_j^\nu)} = \frac{2}{\lambda_{\max}(\sum_{\rho, \sigma=1}^k P_{\mu\rho} G_{\rho\sigma} P_{\sigma\nu})}.$$

If the system satisfies the well-separation condition, $G_{\mu\nu} = (1 - \epsilon)\mathbb{1}_{\{\mu,\nu\}} + \epsilon$, with $\epsilon \ll 1$. Then

$$\begin{aligned} \sum_{\rho,\sigma=1}^k P_{\mu\rho} G_{\rho\sigma} P_{\sigma\nu} &= (1 - \epsilon) \sum_{\rho,\sigma=1}^k P_{\mu\rho} \mathbb{1}_{\{\rho,\sigma\}} P_{\sigma\nu} + \epsilon \sum_{\rho,\sigma=1}^k P_{\mu\rho} P_{\sigma\nu} \\ &= (1 - \epsilon) P_{\mu\nu} \quad (\text{By } P \text{ being a projector and } P(1 \ 1 \ \dots \ 1)^\top = 0) \end{aligned}$$

Let the vector space be $I \oplus I^\perp$, where $I := \text{span}\{(1, \dots, 1)^\top\}$. Since the eigenvalue is $(1 - \epsilon)$ on $I^\perp$ and zero on I , the largest eigenvalue is $(1 - \epsilon)$. After substituting into the expression of β_{jump} , we prove the proposition. $\square$

D Fluctuated Version and Thermal Attention

In this section, we revisit the modern Hopfield models but with the fluctuated variables, x, p , and z . First, we study how the temperature influences the fluctuation property. Then we move on to discuss the modern Hopfield models.

D.1 Physics of Temperature

In this section, we discuss the relation of φ and x_J first. Then we introduce the role that temperature plays in the perturbative analysis.

Proposition D.1 (Saddle Points as Stationary Points in the Zero-Temperature Limit [Landau and Lifshitz, 1980]). Consider the zero-temperature limit, $\beta \rightarrow \infty$. If $(-\beta\mathcal{H}[x] + \sum_{i=1}^d J_i x_i)$ has a unique maximizer, then the saddle points approach to the stationary points, $\varphi_i = (x_J)_i$.

Proof. We repeat the proof from [Landau and Lifshitz, 1980, Dembo and Zeitouni, 2010] for completeness. We denote the maximizer as x_J . Since it is unique, for any $x_i \neq (x_J)_i$, we have

$$\frac{p_J(x)}{p_J(x_J)} = \exp\{-\beta(\mathcal{H}[x] - \mathcal{H}[x_J]) + \sum_{i=1}^d J_i[x_i - (x_J)_i]\} := e^{-\beta\Delta+C},$$

where $\Delta > 0$, and C is a finite value. Then in the zero-temperature limit, $p_J(x) \rightarrow p_J(x_J)$. Thus, $\varphi_i = \mathbb{E}_{p_J}[x_i] \rightarrow (x_J)_i$. $\square$

Proposition D.2 (Fluctuation and Loop Expansion [Thouless et al., 1977]). The free energy captures noise and fluctuations. The fluctuation term behaves like a data-dependent regularizer. Specifically, if we expand x around φ , then

$$\Gamma[\varphi] = \beta\mathcal{H}[\varphi] - \frac{d}{2} \log(2\pi) + \frac{1}{2} \log \det(\beta \nabla_x^2 \mathcal{H}[\varphi]) + O(\beta^{-1}) \quad (\text{D.1})$$

$$:= \Gamma_{\text{tree}}[\varphi] + \Gamma_{1\text{-loop}}[\varphi] + \Gamma_{2\text{-loop}}[\varphi] + \dots \quad (\text{D.2})$$

This free energy admits higher-order (higher-loop) corrections arising from fluctuations, which we denote as $O(\beta^{-1})$. Specifically, the two-loop is of order β^{-1} , three-loop is of order β^{-2} , and so forth. Conventionally, we ignore the constant term, $d \log(2\pi)/2$, since it is independent of φ and does not affect the dynamics.

Proof. We include the proof for completeness, adapting the standard argument from [Thouless et al., 1977, Zinn-Justin, 2002]. By Definition 2.3, we write

$$\begin{aligned} \exp\{-\Gamma[\varphi]\} &= \exp\{W[J] - \sum_{i=1}^d J_i \varphi_i\} \\ &= \int_{\mathcal{X}} \exp\{-\beta\mathcal{H}[x] + \sum_{i=1}^d J_i(x_i - \varphi_i)\} d\mu(x). \end{aligned}$$

Then we expand the right hand side around the stationary point, φ , perturbatively:

$$\begin{aligned}
\exp\{-\Gamma[\varphi]\} &= \int_{\mathcal{X}} \exp\left\{-\beta\mathcal{H}[\varphi] - \beta \sum_{i=1}^d \frac{\partial \mathcal{H}}{\partial x_i} \Big|_{x=\varphi} (x_i - \varphi_i)\right. \\
&\quad - \frac{\beta}{2} \sum_{i,j=1}^d \frac{\partial^2 \mathcal{H}}{\partial x_i \partial x_j} \Big|_{x=\varphi} (x_i - \varphi_i)(x_j - \varphi_j) \\
&\quad - \frac{\beta}{3!} \sum_{i,j,k=1}^d \frac{\partial^3 \mathcal{H}}{\partial x_i \partial x_j \partial x_k} \Big|_{x=\varphi} (x_i - \varphi_i)(x_j - \varphi_j)(x_k - \varphi_k) \\
&\quad \left. + O(\|x - \varphi\|_2^4)\right\} \exp\left\{\sum_{i=1}^d J_i(x_i - \varphi_i)\right\} d\mu(\mathbf{x}) \\
&= \exp\{-\beta\mathcal{H}[\varphi]\} \int_{\mathcal{X}} \exp\left\{-\frac{\beta}{2} \sum_{i,j=1}^d (x_i - \varphi_i)(\nabla_x^2 \mathcal{H}[\varphi])_{ij}(x_j - \varphi_j)\right\} \cdot \left(1 - \frac{\beta}{3!} \times \right. \\
&\quad \left. \sum_{k,\ell,m=1}^d \frac{\partial^3 \mathcal{H}}{\partial x_k \partial x_\ell \partial x_m} \Big|_{x=\varphi} (x_k - \varphi_k)(x_\ell - \varphi_\ell)(x_m - \varphi_m) + O(\beta\|x - \varphi\|_2^4)\right) d\mu(x) \\
&\hspace{15em} \text{(By Definition 2.2)}
\end{aligned}$$

where $\|x' - \varphi\|_2^2 := \sum_{i=1}^d (x'_i - \varphi_i)^2$. Firstly,

$$\int_{\mathcal{X}} \exp\left\{-\frac{\beta}{2} \sum_{i,j=1}^d (x_i - \varphi_i)(\nabla_x^2 \mathcal{H}[\varphi])_{ij}(x_j - \varphi_j)\right\} d\mu(x) = (2\pi)^{d/2} (\det(\beta \nabla_x^2 \mathcal{H}[\varphi]))^{-1/2}.$$

(By Gaussian Formula)

Since the odd function vanishes after the integration, the first leading correction is of $O(\|x - \varphi\|_2^4)$. We focus on this integration:

$$\begin{aligned}
&\int_{\mathcal{X}} \exp\left\{-\frac{\beta}{2} \sum_{i,j=1}^d (x_i - \varphi_i)(\nabla_x^2 \mathcal{H}[\varphi])_{ij}(x_j - \varphi_j)\right\} O(\beta\|x - \varphi\|_2^4) d\mu(x) \\
&= \int_{\mathcal{Y}} \exp\left\{-\frac{1}{2} \sum_{i,j=1}^d (\nabla_x^2 \mathcal{H}[\varphi])_{i,j} y_i y_j\right\} O(\beta\|y\|_2^4) \cdot (\beta^{-1/2})^d d\mu(y) \quad \text{(By } x \rightarrow y = \sqrt{\beta}(x - \varphi)\text{)} \\
&= O(\beta^{-1}).
\end{aligned}$$

Putting everything together, we obtain

$$\exp\{-\Gamma[\varphi]\} = \exp\left\{-\beta\mathcal{H}[\varphi] - \frac{1}{2} \log \det(\beta \nabla_x^2 \mathcal{H}[\varphi]) + \frac{d}{2} \log(2\pi) + O(\beta^{-1})\right\}.$$

This proves the proposition. $\square$

Proposition D.3 (Low-Temperature Limit and Laplace Approximation [Zinn-Justin, 2002, Dembo and Zeitouni, 2010]). If $(-\beta\mathcal{H}[x] + \sum_{i=1}^d J_i x_i)$ has unique maximizer x_J with positive-definite Hessian, then for continuous cases,

$$\frac{1}{\beta} \log \int_{\mathcal{X}} \exp\{-\beta\mathcal{H}[x] + \sum_{i=1}^d J_i x_i\} d\mu(x) = \sup_{x \in \mathcal{X}} \{-\mathcal{H}[x] + \frac{1}{\beta} \sum_{i=1}^d J_i x_i\} + O\left(\frac{\log \beta}{\beta}\right).$$

For discrete cases,

$$\frac{1}{\beta} \log \sum_{x \in \mathcal{X}} \exp\{-\beta\mathcal{H}[x] + \sum_{i=1}^d J_i x_i\} = \sup_{x \in \mathcal{X}} \{-\mathcal{H}[x] + \frac{1}{\beta} \sum_{i=1}^d J_i x_i\} + O\left(\frac{\log |\mathcal{X}|}{\beta}\right). \quad (\text{D.3})$$

Proof. We repeat the proof from [Zinn-Justin, 2002, Dembo and Zeitouni, 2010] for completeness. We start from the continuous case. Since the maximizer is unique, we Taylor expand up to the first leading order as

$$-\beta\mathcal{H}[x] + \sum_{i=1}^d J_i x_i = -\beta\mathcal{H}[\varphi] + \sum_{i=1}^d J_i \varphi_i - \frac{1}{2} \beta [x_i - \varphi_i] (\nabla_x^2 \mathcal{H}[\varphi])_{ij} [x_j - \varphi_j].$$

(By Definition 2.2)

This yields

$$\begin{aligned} & \frac{1}{\beta} \log \int_{\mathcal{X}=\mathbb{R}^d} \exp\{-\beta\mathcal{H}[x] + \sum_{i=1}^d J_i x_i\} d\mu(x) \\ &= \frac{1}{\beta} \log \left(\exp\{-\beta\mathcal{H}[\varphi] + \sum_{i=1}^d J_i \varphi_i\} \right) + \frac{1}{\beta} \log \left(\left(\frac{2\pi}{\beta}\right)^{d/2} (\det(\beta \nabla_x^2 \mathcal{H}[\varphi]))^{-1/2} \right) \\ &= -\mathcal{H}[\varphi] + \frac{1}{\beta} \sum_{i=1}^d J_i \varphi_i + \mathcal{O}\left(\frac{\log \beta}{\beta}\right). \end{aligned}$$

This proves the first equation. For discrete cases,

$$\begin{aligned} \max_{x \in \mathcal{X}} (-\beta\mathcal{H}[x] + \sum_{i=1}^d J_i x_i) &\leq \log \sum_{x \in \mathcal{X}} \exp\{-\beta\mathcal{H}[x] + \sum_{i=1}^d J_i x_i\} \\ &\leq \max_{x \in \mathcal{X}} (-\beta\mathcal{H}[x] + \sum_{i=1}^d J_i x_i) + \log |\mathcal{X}|. \end{aligned}$$

After plugging this into (D.3), we complete the proof. $\square$

Remark D.1. In the zero-temperature limit, Proposition D.3 shows that $W[J]/\beta$ is a Fenchel conjugate of the Hamiltonian by Definition C.6.

Remark D.2. From the proof in [Proposition D.2](#) and [Proposition D.3](#), we notice that the perturbation only works in the low-temperature regime. When the temperature increases, the strength of fluctuations increases and the higher-loop corrections starts to grow. There exists a level where the perturbative method breaks down.

D.2 Fluctuated Variables in Modern Hopfield Model

In the following, we start to consider the fluctuated version of MHMs.

Definition D.1 (Statistics-MHM (Fluctuated) Mapping). The mapping between the variables in statistical mechanics and which in the fluctuated version of MHMs is:

- **Microscopic State** $x \in \mathbb{R}^d$: It represents continuous query. It is also the dynamical variable whose energy landscape defines retrieval basins.
- **External Source** $J \in \mathbb{R}^d$: It represents an external bias of the query.
- **Classical Solution** $\varphi \in \mathbb{R}^d$: It represents macroscopic state. Retrieval corresponds to φ_i aligning with some pattern ξ_i^μ .
- **Probability Function** $p \in \Delta^M$: It represents the retrieval distribution over memories.
- **Tsallis Entropy** Ψ_α : It represents a regularizer controlling how sparse the memory weights p are.
- **Conjugate Field** $z \in \mathbb{R}^M$: It represents the query-key similarity between current state and each memory, i.e., $z_\mu := \sum_{i=1}^d \xi_i^\mu x_i / \sqrt{d}$.

Notice that we distinguish the microscopic state, x , from the macroscopic one, φ , and so are p and z in this case.

Proposition D.4 (Expectation Values and Macroscopic States). Let $\mathcal{X}$ be a finite-volume domain in $\mathbb{R}^d$ and $J \in \text{dom} W_\alpha$. The macroscopic states are determined by the mean of microscopic states. Specifically,

$$\mathbb{E}_{x \sim \rho_J(\cdot|p)}^{(\alpha)}[x_i|p] = \frac{1}{\sqrt{d}} \sum_{\mu=1}^M \xi_i^\mu p_\mu + \frac{1}{\beta} J_i + B_i^{\mathcal{X}}(p, J) := m_i(p, J) + B_i^{\mathcal{X}}(p, J), \quad (\text{D.4})$$

$$\mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[x_i] = \varphi_i = \frac{1}{\sqrt{d}} \sum_{\mu=1}^M \xi_i^\mu \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[p_\mu] + \frac{1}{\beta} J_i + \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[B_i^{\mathcal{X}}(p, J)], \quad (\text{D.5})$$

where $\rho_J(x, p)$ is the probability density function in $\mathcal{X} \times \Delta^M$, $p_J(x)$ is which in $\mathcal{X}$, and $p_{\Delta, J}(p)$ is which in Δ^M . Furthermore, there is a boundary correction

$$B_i^{\mathcal{X}}(p, J) := \frac{\int_{\mathcal{Y}} y_i \exp_\alpha^\alpha \left\{ -\frac{\beta}{2} (\|y\|_2^2 - \|m\|_2^2) \right\} d\mu(y)}{\int_{\mathcal{Y}} \exp_\alpha^\alpha \left\{ -\frac{\beta}{2} (\|y\|_2^2 - \|m\|_2^2) \right\} d\mu(y)}, \quad \text{with } \mathcal{Y} = \mathcal{X} - m.$$

Proof. We consider x and p follow the distribution $\rho_J(x, p)$ with the joint energy function, [\(3.1\)](#).

Then

$$\begin{aligned}
\mathbb{E}_{x \sim \rho_J(\cdot|p)}^{(\alpha)}[x_i|p] &= \frac{\int_{\mathcal{X}} x_i \rho_J^\alpha(x|p) \, d\mu(x)}{\int_{\mathcal{X}} \rho_J^\alpha(x|p) \, d\mu(x)} \\
&= \frac{\int_{\mathcal{X}} x_i \exp_\alpha^\alpha \left\{ -\frac{\beta}{2} \sum_{j=1}^d x_j^2 + \frac{\beta}{\sqrt{d}} \sum_{\mu=1}^M \sum_{j=1}^d \xi_j^\mu x_j p_\mu + \sum_{j=1}^d J_j x_j \right\} \, d\mu(x)}{\int_{\mathcal{X}} \exp_\alpha^\alpha \left\{ -\frac{\beta}{2} \sum_{j=1}^d x_j^2 + \frac{\beta}{\sqrt{d}} \sum_{\mu=1}^M \sum_{j=1}^d \xi_j^\mu x_j p_\mu + \sum_{j=1}^d J_j x_j \right\} \, d\mu(x)} \\
&= \frac{\int_{\mathcal{Y}} (y_i + m_i) \exp_\alpha^\alpha \left\{ -\frac{\beta}{2} \sum_{j=1}^d y_j^2 + \frac{\beta}{2} \sum_{j=1}^d m_j^2 \right\} \, d\mu(y)}{\int_{\mathcal{Y}} \exp_\alpha^\alpha \left\{ -\frac{\beta}{2} \sum_{j=1}^d y_j^2 + \frac{\beta}{2} \sum_{j=1}^d m_j^2 \right\} \, d\mu(y)} \\
&= m_i(p, J) + B_i^{\mathcal{X}}(p, J).
\end{aligned}$$

This proves the first equation. In this way,

$$\begin{aligned}
\mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[x_i] &= \frac{\int_{\Delta^M} \int_{\mathcal{X}} x_i \rho_J^\alpha(x, p) \, d\mu(x) d\mu(p)}{\int_{\Delta^M} \int_{\mathcal{X}} \rho_J^\alpha(x, p) \, d\mu(x) d\mu(p)} \\
&= \frac{\int_{\Delta^M} (\int_{\mathcal{X}} x_i \rho_J^\alpha(x|p) \, d\mu(x)) p_{\Delta, J}^\alpha(p) \, d\mu(p)}{\int_{\Delta^M} (\int_{\mathcal{X}} \rho_J^\alpha(x|p) \, d\mu(x)) p_{\Delta, J}^\alpha(p) \, d\mu(p)} \quad (\text{By } p_{\Delta, J}(p) \rho_J(x|p) := \rho_J(x, p)) \\
&= \frac{\int_{\Delta^M} (\int_{\mathcal{X}} \rho_J^\alpha(x'|p) \, d\mu(x')) \mathbb{E}_{x \sim \rho_J(\cdot|p)}^{(\alpha)}[x_i|p] p_{\Delta, J}^\alpha(p) \, d\mu(p)}{\int_{\Delta^M} (\int_{\mathcal{X}} \rho_J^\alpha(x|p) \, d\mu(x)) p_{\Delta, J}^\alpha(p) \, d\mu(p)} \\
&= \frac{1}{\sqrt{d}} \sum_{\mu=1}^M \xi_i^\mu \tilde{\mathbb{E}}_{p \sim p_{\Delta, J}}^{(\alpha)}[p_\mu] + \frac{1}{\beta} J_i + \tilde{\mathbb{E}}_{p \sim p_{\Delta, J}}^{(\alpha)}[B_i^{\mathcal{X}}(p, J)],
\end{aligned}$$

where

$$C(p) := \int_{\mathcal{X}} \rho_J^\alpha(x|p) \, d\mu(x), \quad \text{and} \quad \tilde{\mathbb{E}}_{p \sim p_{\Delta, J}}^{(\alpha)}[p_\mu] := \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[p_\mu] = \frac{\int_{\Delta^M} C(p) p_\mu p_{\Delta, J}^\alpha(p) \, d\mu(p)}{\int_{\Delta^M} C(p) p_{\Delta, J}^\alpha(p) \, d\mu(p)}.$$

This completes the proof. $\square$

Remark D.3. The correction $B_i^{\mathcal{X}}$ vanishes on translation-invariant domains (e.g., $\mathbb{R}^d$ if the integral is normalizable). Furthermore, it is negligible in the large-volume or low-temperature regime when the dominant integrand is localized far from the boundary of $\mathcal{X}$.

D.3 Connection between Free Energy and Tsallis Entropy in Fluctuated Case

In [Section 2.2](#), we relate the log-partition function to the Tsallis entropy via (2.4). Here, we express the log-partition function using the variables in fluctuated version of MHMs, [Definition D.1](#). To make $Z_\alpha[J]$ finite, we assume that the integration region of x is finite volume, i.e., $\mathcal{X} \subset \mathbb{R}^d$ such that $|\mathcal{X}| < +\infty$.

Proposition D.5. Let $\rho_J(x, p)$ and $p_J(x)$ be the probability density function in $\mathcal{X} \times \Delta^M$ and $\mathcal{X}$, respectively. Inspired by [Theorem 2.1](#), we have

$$\begin{aligned} -W_\alpha[J, \xi] &= \frac{\beta}{2} \sum_{i=1}^d \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[x_i^2] - \beta \sum_{\mu=1}^M \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[z_\mu p_\mu] \\ &\quad - \frac{\Psi_\alpha[\rho_J]}{Z_\alpha^{\alpha-1}[J, \xi] \int_{\mathcal{X}} \int_{\Delta^M} \rho_J^\alpha(x, p) \, d\mu(p) d\mu(x)} - \sum_{i=1}^d J_i \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[x_i]. \end{aligned} \quad (\text{D.6})$$

Then the free energy is

$$\begin{aligned} \tilde{\Gamma}_\alpha &= \frac{\beta}{2} \sum_{i=1}^d \mathbb{E}_{(x,p) \sim \rho_{J^*}}^{(\alpha)}[x_i^2] - \beta \sum_{\mu=1}^M \mathbb{E}_{(x,p) \sim \rho_{J^*}}^{(\alpha)}[z_\mu p_\mu] \\ &\quad - \frac{\Psi_\alpha[\rho_{J^*}]}{Z_\alpha^{\alpha-1}[J^*, \xi] \int_{\mathcal{X}} \int_{\Delta^M} \rho_{J^*}^\alpha(x, p) \, d\mu(p) d\mu(x)}, \end{aligned} \quad (\text{D.7})$$

where J^* is the maximizing source, i.e., $J^* = \arg \max_J \{\sum_{i=1}^d J_i \varphi_i - W_\alpha[J, \xi]\}$, and $\varphi = \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[x]$.

Proof. Following the proof of [Theorem 3.1](#), we have (C.11). This is exactly the first equation we would like to prove. For the free energy, we consider the Fenchel conjugation [Definition C.6](#) of the log-partition function.

$$\begin{aligned} \tilde{\Gamma}_\alpha &= \sup_{J \in \text{dom} W_\alpha} \left\{ \sum_{i=1}^d J_i \varphi_i - W_\alpha[J, \xi] \right\} \\ &= \sup_{J \in \text{dom} W_\alpha} \left\{ \sum_{i=1}^d J_i (\varphi_i - \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[x_i]) + \frac{\beta}{2} \sum_{i=1}^d \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[x_i^2] - \beta \sum_{\mu=1}^M \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[z_\mu p_\mu] \right. \\ &\quad \left. - \frac{\Psi_\alpha[\rho_J]}{Z_\alpha^{\alpha-1}[J, \xi] \int_{\mathcal{X}} \int_{\Delta^M} \rho_J^\alpha(x, p) \, d\mu(p) d\mu(x)} \right\} \\ &= \frac{\beta}{2} \sum_{i=1}^d \mathbb{E}_{(x,p) \sim \rho_{J^*}}^{(\alpha)}[x_i^2] - \beta \sum_{\mu=1}^M \mathbb{E}_{(x,p) \sim \rho_{J^*}}^{(\alpha)}[z_\mu p_\mu] \\ &\quad - \frac{\Psi_\alpha[\rho_{J^*}]}{Z_\alpha^{\alpha-1}[J^*, \xi] \int_{\mathcal{X}} \int_{\Delta^M} \rho_{J^*}^\alpha(x, p) \, d\mu(p) d\mu(x)}, \end{aligned}$$

where $\varphi = \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[x]$. This completes the proof. $\square$

Proposition D.6 (MHMs in the Zero-Temperature Limit). In the zero-temperature limit, $\beta \rightarrow$

∞ , minimizing the free energy becomes minimizing the following quantity:

$$\mathcal{F}[\varphi, \xi] = \frac{1}{2} \sum_{i=1}^d \varphi_i^2 - \max_{\mu} \bar{z}_{\mu}, \quad (\text{D.8})$$

where $\bar{z}_{\mu} := \sum_{i=1}^d \xi_i^{\mu} \varphi_i / \sqrt{d}$. Notably, this is the free energy in the fixed case, (3.4), up to a scaling, β .

Proof. In the zero-temperature limit, the Gibbs probability density function concentrates to an indicator, i.e.,

$$\rho_J(x, p) \rightarrow \mathbb{1}_{\{x, \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[x]\}} \mathbb{1}_{\{p, \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[p]\}},$$

as $\beta \rightarrow \infty$. Hence,

$$\mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[x^2] \rightarrow (\mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[x])^2 = (\tilde{\mathbb{E}}_{x \sim p_J}^{(\alpha)}[x])^2 = \varphi^2,$$

and

$$\mathbb{E}_{(x,p) \sim \rho_J}[z_{\mu} p_{\mu}] \rightarrow \frac{1}{\sqrt{d}} \sum_{i=1}^d \xi_i^{\mu} \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[x_i] \mathbb{E}_{(x,p) \sim \rho_J}^{(\alpha)}[p_{\mu}] = \frac{1}{\sqrt{d}} \sum_{i=1}^d \xi_i^{\mu} \varphi_i \tilde{\mathbb{E}}_{p \sim p_{\Delta, J}}^{(\alpha)}[p_{\mu}] := \bar{z}_{\mu} \bar{p}_{\mu},$$

where $\tilde{\mathbb{E}}_{x \sim p_J}^{(\alpha)}[x]$ and $\tilde{\mathbb{E}}_{p \sim p_{\Delta, J}}^{(\alpha)}[p]$ are defined in (C.12) and (C.13), respectively. The free energy (D.7) becomes

$$\Gamma_{\alpha} \rightarrow \frac{\beta}{2} \sum_{i=1}^d \varphi_i^2 - \beta \sum_{\mu=1}^M \bar{z}_{\mu} \bar{p}_{\mu} - \Psi_{\alpha}[\mathbb{1}].$$

Since $\Psi_{\alpha}[\mathbb{1}] = 0$, we have

$$\inf_{\bar{p} \in \Delta^M} \left\{ - \sum_{\mu=1}^M \bar{z}_{\mu} \bar{p}_{\mu} \right\} = - \max_{\mu} \bar{z}_{\mu}.$$

Substituting it into the free energy to eliminate the dependence of $\bar{p}_{\mu}$, we prove this proposition. $\square$

Remark D.4. Proposition D.6 is consistent with Proposition D.2 in the zero-temperature limit. Specifically, $\Gamma_{\alpha}[\varphi] \rightarrow (\beta \mathcal{H}[\varphi] + \text{constant})$ as $\beta \rightarrow \infty$. Following the proof of Theorem 3.1, we conclude the same result.

D.4 Retrieval Dynamics in Fluctuated Case and Thermal Attention

In Section 2, we introduce the equation of motions according to Proposition 2.1 and investigate the attention mechanism from the trajectories in associative memory systems. In this section,

we consider the fluctuated version, which gives an alternative thermal attention mechanism. Let $\widetilde{W}_\alpha := \log_\alpha \int_{\Delta^M} \exp_\alpha \{ -\beta \widetilde{\mathcal{H}}_{\text{MHM}} + \sum_{i=1}^d J_i x_i \} d\mu(p)$. Notice that in the fixed setting, it reduces to the simpler form, (3.7). Then,

Proposition D.7 (EOMs of Free Energies in MHMs). The EOM of the log-partition function in the x -sector corresponds to the Hopfield retrieval

$$\nabla_J W_\alpha[J, \xi] = \int_{\mathcal{X}} \int_{\Delta^M} x_i \rho_J^\alpha(x, p) d\mu(p) d\mu(x). \quad (\text{D.9})$$

The EOM of the log-partition function in the p -sector corresponds to the unnormalized escort moment [Naudts, 2004]

$$\nabla_z \widetilde{W}_\alpha[z, J, \xi] = \int_{\Delta^M} \beta p \rho_J^\alpha(p|x) d\mu(p) := \beta p^{(\text{esc})}. \quad (\text{D.10})$$

Proof. By Definition 2.3, $W_\alpha[J, \xi] = \log_\alpha Z_\alpha[J, \xi]$. After taking the derivative with respect to J , we obtain

$$\begin{aligned} \frac{\partial}{\partial J_i} W_\alpha[J, \xi] &= \frac{\partial}{\partial J_i} \log_\alpha Z_\alpha[J, \xi] \\ &= Z_\alpha^{-\alpha}[J, \xi] \frac{\partial Z_\alpha[J, \xi]}{\partial J_i} \quad (\text{By Lemma C.1}) \\ &= Z_\alpha^{-\alpha}[J, \xi] \frac{\partial}{\partial J_i} \int_{\mathcal{X}} \int_{\Delta^M} \exp_\alpha \{ -\beta \widetilde{\mathcal{H}}_{\text{MHM}}[x, p, \xi] + \sum_{j=1}^d J_j x_j \} d\mu(p) d\mu(x) \\ &= Z_\alpha^{-\alpha}[J, \xi] \int_{\mathcal{X}} \int_{\Delta^M} \exp_\alpha^\alpha \{ -\beta \widetilde{\mathcal{H}}_{\text{MHM}}[x, p, \xi] + \sum_{j=1}^d J_j x_j \} \quad (\text{By Lemma C.1}) \\ &\quad \times \frac{\partial}{\partial J_i} (-\beta \widetilde{\mathcal{H}}_{\text{MHM}}[x, p, \xi] + \sum_{k=1}^d J_k x_k) d\mu(p) d\mu(x) \\ &= \int_{\mathcal{X}} \int_{\Delta^M} x_i \frac{\exp_\alpha^\alpha \{ -\beta \widetilde{\mathcal{H}}_{\text{MHM}}[x, p, \xi] + \sum_{j=1}^d J_j x_j \}}{Z_\alpha^\alpha[J, \xi]} d\mu(p) d\mu(x) \\ &= \int_{\mathcal{X}} \int_{\Delta^M} x_i \rho_J^\alpha(x, p) d\mu(p) d\mu(x). \end{aligned}$$

For the second equation, we substitute the definition of $\widetilde{W}_\alpha$ directly:

$$\begin{aligned} \frac{\partial}{\partial z_\mu} \widetilde{W}_\alpha[z, J, \xi] &= \frac{\partial}{\partial z_\mu} \log_\alpha \int_{\Delta^M} \exp_\alpha \{ -\frac{\beta}{2} \sum_{i=1}^d x_i^2 + \beta \sum_{\mu=1}^M z_\mu p_\mu + \sum_{i=1}^d J_i x_i \} d\mu(p) \\ &= \left(\int_{\Delta^M} \exp_\alpha \{ -\frac{\beta}{2} \sum_{i=1}^d x_i^2 + \beta \sum_{\mu=1}^M z_\mu p_\mu + \sum_{i=1}^d J_i x_i \} d\mu(p) \right)^{-\alpha} \end{aligned}$$

$$\begin{aligned}
& \times \int_{\Delta^M} \beta p_\mu \exp_\alpha^\alpha \left\{ -\frac{\beta}{2} \sum_{i=1}^d x_i^2 + \beta \sum_{\mu=1}^M z_\mu p_\mu + \sum_{i=1}^d J_i x_i \right\} d\mu(p) \\
& \quad \quad \quad \text{(By Lemma C.1)} \\
& = \int_{\Delta^M} \beta p_\mu \rho_J^\alpha(p|x) d\mu(p) .
\end{aligned}$$

This completes the proof. $\square$

Remark D.5. If $\alpha \rightarrow 1$, then everything simplifies to the ordinary Gibbs case. In this way, the Hopfield retrieval becomes

$$\begin{aligned}
\nabla_J W_{\alpha \rightarrow 1}[J, \xi] &= \int_{\mathcal{X}} \int_{\Delta^M} x \rho_J(x, p) d\mu(p) d\mu(x) \\
&= \mathbb{E}_{(x,p) \sim \rho_J}[x] \\
&= \int_{\mathcal{X}} x p_J(x) d\mu(x) \\
&= \mathbb{E}_{x \sim p_J}[x].
\end{aligned}$$

Furthermore, the escort moment becomes the conditional expectation value of the probability function for general β .

$$\nabla_z \widetilde{W}_{\alpha \rightarrow 1}[z, J, \xi] = \beta \mathbb{E}_{p \sim \rho_J(\cdot|x)}[p] .$$

Remark D.6. If we do the similar analysis of [Proposition 3.3](#) for the fluctuated case, then we have the thermal attention mechanism. In particular, if $\alpha \rightarrow 1$, then the learning trajectory follows the ordinary Gibbs expectation of the conditional distribution function, i.e., $\nabla_z \widetilde{W}_{\alpha \rightarrow 1} = \beta \mathbb{E}_{p \sim \rho_J(\cdot|x)}[p]$, due to [Remark D.5](#). If $\alpha > 1$, then the system amplifies high-density regions. This leads to a sharper attention mechanism. If $1 > \alpha > 0$, then it corresponds to a softer attention mechanism.

Remark D.7. From [\(D.9\)](#), we obtain a macroscopic retrieval state depending on the external field at microscopic queries, $\{x_i\}$. It is similar to a spin system in an external magnetic field. After turning on the magnetic field, the probability distribution of spins starts to deform and yields an expectation value $\mathbb{E}[S_i]$, where S_i is the notation of a spin at site i . We call it a Landau response. Hence, the Hopfield retrieval is a Landau response in the x -sector (J -space). For the attention mechanism, it is an inner Landau response in the p -sector (z -space). In this way, we view the query-to-key score, $z_\mu = \sum_{i=1}^d \xi_i^\mu x_i / \sqrt{d}$, as an external field at $\{p_\mu\}$ in Δ^M space. This formalism gives a physical interpretation of the retrieval dynamics.

Lemma D.1 (Linear Function and its Minimizer on Simplex [[Rockafellar, 1997](#)]). Consider a

function $h(p)$ that is linear in p , i.e.,

$$h(p) = \sum_{\mu=1}^M a_{\mu} p_{\mu} + b,$$

for some constant $a_{\mu} \in \mathbb{R}^M$. Then $h(p)$ cannot have a unique minimizer p^* in the interior of the simplex.

Proof. This is a standard convex analysis [Rockafellar, 1997]. We repeat the proof for the completeness. Assume that a minimizer lies in the interior of Δ^M . Then $p^* > 0$. Since p^* is a minimizer, we have

$$\frac{d}{d\varepsilon} h(p^* + \varepsilon v)|_{\varepsilon=0} = 0,$$

for all tangent vectors $v \in T_{p^*} \Delta^M = \{v \in \mathbb{R}^M \mid \sum_{\mu=1}^M v_{\mu} = 0\} \simeq \mathbb{R}^{M-1}$. Then, after expanding the linear function h , it yields

$$0 = \frac{d}{d\varepsilon} \left(\sum_{\mu=1}^M a_{\mu} p_{\mu}^* + \varepsilon \sum_{\mu=1}^M a_{\mu} v_{\mu} + b \right) |_{\varepsilon=0} = \sum_{\mu=1}^M a_{\mu} v_{\mu}.$$

Since the tangent vector satisfies $\sum_{\mu=1}^M v_{\mu} = 0$, we conclude that $a_{\mu} = \lambda$ for all μ and some constant λ . Hence,

$$h(p) = \lambda \sum_{\mu=1}^M p_{\mu} + b = \lambda + b,$$

which is constant on Δ^M . This contradicts the assumption that the minimizer p^* is unique. This proves the lemma. $\square$

Lemma D.2 (Normalization Factor in the Zero-Temperature Limit). Consider the zero-temperature limit $\beta \rightarrow \infty$ and $\alpha > 1$. Assume that $\tilde{\mathcal{H}}_{\text{MHM}} \geq 0$, and $\rho_J(p|x)$ has a unique maximizer p^* such that $\tilde{\mathcal{H}}_{\text{MHM}}[x, p^*, \xi] = 0$. If the p is continuous on the simplex Δ^M , where $M > 1$, then the normalization function has the following limit

$$\lim_{\beta \rightarrow \infty} C(x) = \lim_{\beta \rightarrow \infty} \int_{\Delta^M} \rho_J^{\alpha}(p|x) d\mu(p) = \infty.$$

If $\alpha \rightarrow 1$ for general β or if p is discrete on the simplex Δ^M for general α , then

$$\lim_{\alpha \rightarrow 1} C(x) = \lim_{\beta \rightarrow \infty} C^{\text{dis}}(x) = \lim_{\beta \rightarrow \infty} \sum_{p \in \Delta^M} (\rho_J^{\text{dis}})^{\alpha}(p|x) = 1.$$

Proof. We first consider the continuous case. Since $\tilde{\mathcal{H}}_{\text{MHM}}$ is linear in p , its unique minimizer p^*

must be at a vertex of Δ^M according to [Lemma D.1](#). Hence, we denote

$$p_\mu^\star = \begin{cases} 1, & \text{if } \mu = \mu^\star \\ 0, & \text{if } \mu \neq \mu^\star \end{cases}. \quad (\text{D.11})$$

In this way, for $p \neq p^\star$, we have

$$\begin{aligned} \tilde{\mathcal{H}}_{\text{MHM}}[x, p, \xi] &= \tilde{\mathcal{H}}_{\text{MHM}}[x, p, \xi] - \tilde{\mathcal{H}}_{\text{MHM}}[x, p^\star, \xi] && (\text{By } \tilde{\mathcal{H}}_{\text{MHM}}[x, p^\star, \xi] = 0) \\ &= \frac{1}{2} \sum_{i=1}^d x_i^2 + \sum_{\mu=1}^M z_\mu p_\mu - \left(\frac{1}{2} \sum_{i=1}^d x_i^2 + \sum_{\mu=1}^M z_\mu p_\mu^\star \right) \\ &= \sum_{\mu=1}^M z_\mu (p_\mu - p_\mu^\star) \\ &= \sum_{\mu=\mu^\star} z_\mu (p_\mu - p_\mu^\star) + \sum_{\substack{\mu=1 \\ \mu \neq \mu^\star}}^M z_\mu (p_\mu - p_\mu^\star) \\ &= z_{\mu^\star} p_{\mu^\star} - z_{\mu^\star} p_{\mu^\star}^\star + \sum_{\substack{\mu=1 \\ \mu \neq \mu^\star}}^M z_\mu (p_\mu - p_\mu^\star) \\ &= z_{\mu^\star} p_{\mu^\star} - z_{\mu^\star} + \sum_{\substack{\mu=1 \\ \mu \neq \mu^\star}}^M z_\mu p_\mu && (\text{By (D.11)}) \\ &= \sum_{\substack{\mu=1 \\ \mu \neq \mu^\star}}^M (z_\mu - z_{\mu^\star}) p_\mu. && (\text{By } 1 = p_{\mu^\star} + \sum_{\mu \neq \mu^\star}^M p_\mu) \end{aligned}$$

Since $p^\star$ is a unique minimizer of $\tilde{\mathcal{H}}_{\text{MHM}}$, $(z_\mu - z_{\mu^\star}) > 0$ for all $\mu \neq \mu^\star$. From [Definition C.9](#), we obtain

$$\begin{aligned} \exp_\alpha \left\{ -\beta \tilde{\mathcal{H}}_{\text{MHM}}[x, p, \xi] + \sum_{i=1}^d J_i x_i \right\} &= \left\{ 1 + (\alpha - 1) (\beta \tilde{\mathcal{H}}_{\text{MHM}}[x, p, \xi] - \sum_{i=1}^d J_i x_i) \right\}^{1/(1-\alpha)} \\ &= \left\{ 1 + (\alpha - 1) \beta \sum_{\mu \neq \mu^\star}^M (z_\mu - z_{\mu^\star}) p_\mu - (\alpha - 1) \sum_{i=1}^d J_i x_i \right\}^{1/(1-\alpha)} \end{aligned}$$

for $\alpha > 1$. Then

$$C(x) = \frac{\int_{\Delta^M} \exp_\alpha^\alpha \left\{ -\beta \tilde{\mathcal{H}}_{\text{MHM}}[x, p, \xi] + \sum_{i=1}^d J_i x_i \right\} d\mu(p)}{\left(\int_{\Delta^M} \exp_\alpha \left\{ -\beta \tilde{\mathcal{H}}_{\text{MHM}}[x, p', \xi] + \sum_{i=1}^d J_i x_i \right\} d\mu(p') \right)^\alpha}$$

$$= \frac{\int_{\Delta^M} \left\{ [1 - (\alpha - 1) \sum_{i=1}^d J_i x_i] + (\alpha - 1) \beta \sum_{\mu \neq \mu^*}^M (z_\mu - z_{\mu^*}) \right\}^{\alpha/(1-\alpha)} d\mu(p)}{\left(\int_{\Delta^M} \left\{ [1 - (\alpha - 1) \sum_{i=1}^d J_i x_i] + (\alpha - 1) \beta \sum_{\mu \neq \mu^*}^M (z_\mu - z_{\mu^*}) \right\}^{1/(1-\alpha)} d\mu(p') \right)^\alpha}$$

Since the integrands only depend on $\mu \neq \mu^*$, we change the variable to $u_\mu = \beta q_\mu$, where $q_\mu := p_\mu$ with $\mu \neq \mu^*$. There are $(M - 1)$ such coordinates, so $d\mu(p) \sim \beta^{-(M-1)} d\mu(u)$. Then, for $\alpha > 1$ and $M > 1$, we have

$$C(x) \sim \frac{\beta^{-(M-1)}}{\left(\beta^{-(M-1)}\right)^\alpha} = \beta^{(\alpha-1)(M-1)} \rightarrow \infty$$

as $\beta \rightarrow \infty$. For $\alpha \rightarrow 1$, it is simply $C(x) = \int_{\Delta^M} \rho_J(p|x) d\mu(p) = 1$ for any β . For the discrete case, the conditional distribution is

$$\rho_J^{\text{dis}}(p|x) = \frac{\exp_\alpha\{-\beta\tilde{\mathcal{H}}_{\text{MHM}}[x, p, \xi] + \sum_{i=1}^d J_i x_i\}}{\sum_{p' \in \Delta^M} \exp_\alpha\{-\beta\tilde{\mathcal{H}}_{\text{MHM}}[x, p', \xi] + \sum_{i=1}^d J_i x_i\}},$$

with discrete p . Notice that as $\beta \rightarrow \infty$, $\exp_\alpha\{-\beta\tilde{\mathcal{H}}_{\text{MHM}}[x, p, \xi]\}$ is suppressed to zero except for the maximizer p^* , i.e., $\exp_\alpha\{-\beta\tilde{\mathcal{H}}_{\text{MHM}}[x, p^*, \xi]\} = 1$. Hence, as $\beta \rightarrow \infty$, the denominator of $\rho_J^{\text{dis}}(p|x)$ approaches to

$$\lim_{\beta \rightarrow \infty} \sum_{p' \in \Delta^M} \exp_\alpha\{-\beta\tilde{\mathcal{H}}_{\text{MHM}}[x, p', \xi] + \sum_{i=1}^d J_i x_i\} = \exp_\alpha\left\{\sum_{i=1}^d J_i x_i\right\}.$$

Then

$$\begin{aligned} \lim_{\beta \rightarrow \infty} C^{\text{dis}}(x) &= \lim_{\beta \rightarrow \infty} \sum_{p \in \Delta^M} (\rho_J^{\text{dis}})^\alpha(p|x) \\ &= \lim_{\beta \rightarrow \infty} \left\{ \sum_{p=p^*} (\rho_J^{\text{dis}})^\alpha(p|x) + \sum_{\substack{p \in \Delta^M \\ p \neq p^*}} (\rho_J^{\text{dis}})^\alpha(p|x) \right\} \\ &= \lim_{\beta \rightarrow \infty} \left\{ \frac{\exp_\alpha^\alpha\{\sum_{i=1}^d J_i x_i\}}{(\exp_\alpha\{\sum_{i=1}^d J_i x_i\})^\alpha} \quad (\text{By } \tilde{\mathcal{H}}_{\text{MHM}}[x, p^*, \xi] = 0) \right. \\ &\quad \left. + \sum_{\substack{p \in \Delta^M \\ p \neq p^*}} \frac{\exp_\alpha^\alpha\{-\beta\tilde{\mathcal{H}}_{\text{MHM}}[x, p, \xi] + \sum_{i=1}^d J_i x_i\}}{(\exp_\alpha\{\sum_{i=1}^d J_i x_i\})^\alpha} \right\} \\ &= \lim_{\beta \rightarrow \infty} \{1 + 0\} \\ &= 1. \end{aligned}$$

This completes the proof. □

Remark D.8. In practical attention mechanisms, the retrieved memories are a finite set. The attention weights are finite-dimensional vectors. Therefore, the discrete normalization behavior in [Lemma D.2](#) occurs in practice. In contrast, the divergence of $C(x)$ in the continuous case is mainly a theoretical issue.

Proposition D.8 (Emergence of Hard Argmax from Escort Moments). Assume that $\tilde{\mathcal{H}}_{\text{MHM}} \geq 0$ has a unique minimizer $p^* \in \Delta^M$. In the zero-temperature limit, $\beta \rightarrow \infty$, the following escort variables become deterministic hard argmax selections:

$$\begin{cases} \alpha \rightarrow 1 & : p^{(\text{esc})} \\ \alpha > 1, \text{ Continuous} & : p^{(\text{esc})}/C(x) \xrightarrow{\beta \rightarrow \infty} p^* = \arg \max_{p \in \Delta^M} \left\{ \sum_{\mu=1}^M z_\mu p_\mu \right\}. \\ \alpha > 1, \text{ Discrete} & : p^{(\text{esc})} \end{cases} \quad (\text{D.12})$$

Proof. For $\alpha \rightarrow 1$, the escort moment reduces to the ordinary Gibbs expectation,

$$p^{(\text{esc})} = \int_{\Delta^M} p \rho_J(p|x) \, d\mu(p).$$

By [Proposition D.3](#), the Gibbs distribution collapses at the unique minimizer in the zero-temperature limit,

$$\rho_J(p|x) \rightarrow \mathbb{1}_{\{p, p^*\}},$$

as $\beta \rightarrow \infty$. Hence, $p^{(\text{esc})} \rightarrow p^*$ as $\beta \rightarrow \infty$. Next, we consider $\alpha > 1$ with continuous p . By [Lemma D.2](#), $\lim_{\beta \rightarrow \infty} C(x)$ diverges. However, since the escort measure still concentrates at the unique minimizer p^* , we obtain

$$\lim_{\beta \rightarrow \infty} \frac{p^{(\text{esc})}}{C(x)} = \lim_{\beta \rightarrow \infty} \frac{\int_{\Delta^M} p \rho_J^\alpha(p|x) \, d\mu(p)}{\int_{\Delta^M} \rho_J^\alpha(p'|x) \, d\mu(p')} = p^*.$$

Finally, for $\alpha > 1$ with discrete p , we have $\lim_{\beta \rightarrow \infty} C^{\text{dis}}(x) = 1$ according to [Lemma D.2](#). Hence,

$$\lim_{\beta \rightarrow \infty} \frac{p^{(\text{esc})}}{C^{\text{dis}}(x)} = \lim_{\beta \rightarrow \infty} p^{(\text{esc})} = p^*.$$

This proves the proposition. □

Remark D.9. [Proposition D.8](#) shows that escort-weighted retrieval ([D.10](#)) reduces to a deterministic hard argmax selection rule in the zero-temperature limit. This statistical-mechanics framework demonstrates how sparse retrieval behavior is able to emerge from escort concentration in the general fluctuated (stochastic) setting.

E Dynamical Retrieval Stability Analysis

In this section, we investigate the dynamical retrieval stability. Then we discuss how the convergences of x and p relate to each other. To do this, we apply gradient-based optimization and perturb around the fixed point. The stability analysis verifies whether the training trajectory flows back to the fixed-point or not.

Definition E.1 (Retrieval Dynamics). The retrieval update rules are

$$\frac{d}{dt}x_i(t) = -(\nabla_x \Gamma_\alpha^{(\text{mod})}[x(t), \xi])_i \quad (\text{Continuous}), \quad (\text{E.1})$$

$$x_i(t+1) = x_i(t) - \eta(\nabla_x \Gamma_\alpha^{(\text{mod})}[x(t), \xi])_i \quad (\text{Discrete}), \quad (\text{E.2})$$

where η is the learning rate.

In [Section 2](#), we note that $\nabla_\varphi \Gamma_\alpha[\varphi]$ is not bijective. Hence, it converges to a set rather than a point in general. In our setting, we consider retrieval dynamics of the generalized model $\Gamma_\alpha^{(\text{mod})}$ in the query space (x -space), while the memory patterns remain fixed. The projection mapping $z_\mu = \sum_{i=1}^d \xi_i^\mu x_i / \sqrt{d}$ is not injective generally, so distinct x^* correspond to one specific fixed-point z^* . However, due to $\frac{\beta}{2} \sum_{i=1}^d x_i^2$ term, such projection degeneracy does not imply the degeneracy of the full $\Gamma_\alpha^{(\text{mod})}[x, \xi]$.

By the proposition below, we are capable of studying the stability and refer the stability to the stability modulo the full degeneracy. This interpretation is consistent with the analysis of gradient-based optimization developed in [\[Lee et al., 2016\]](#), where the local Hessian structure governs the stability of gradient flows near critical points. Furthermore, we only apply the gradient-based optimization in x -space. Hence, the stable fixed-point in x -space must correspond to one specific stable fixed-point in z -space.

Proposition E.1 (State-Retrieval Stability Modulo Degeneracy). Let the macroscopic state space be $\mathcal{M}_x$, $\mathcal{K}(x) := \ker(\nabla_x^2 \Gamma_\alpha^{(\text{mod})}[x, \xi])$ be the kernel space, and $\mathcal{K}^\perp(x)$ be the orthogonal complement of $\mathcal{K}(x)$. Assume there exists a local stable point $x^* \in \mathbb{R}^d$ such that

$$\nabla_x \Gamma_\alpha^{(\text{mod})}[x^*, \xi] = 0, \quad (\text{E.3})$$

and

$$\text{Hessian } \nabla_x^2 \Gamma_\alpha^{(\text{mod})}[x^*, \xi] \text{ is positive definite on } \mathcal{K}^\perp(x^*). \quad (\text{E.4})$$

Let $\mathcal{N}_{x^*} \subset \mathcal{M}_\varphi$ be the degeneracy space through x^* such that $\Gamma_\alpha^{(\text{mod})}[x, \xi]$ is constant on it, and $T_{x^*} \mathcal{N}_{x^*} = \mathcal{K}(x^*)$. Then x^* is a stable local minimizer modulo degeneracy. That is, the gradient flow in the x -space converges to $(x^* + v)$ for some $v \in \mathcal{K}(x^*)$. Additionally, for discrete gradient

descent, we require

$$0 < \eta < \frac{2}{\lambda_{\max}(\nabla_x^2 \Gamma_{\alpha}^{(\text{mod})}[x^*, \xi])|_{\mathcal{K}^{\perp}(x^*)}} , \quad (\text{E.5})$$

where $\lambda_{\max}(\cdot)|_{\mathcal{K}^{\perp}(x^*)}$ is the maximum eigenvalue with the Hessian being on $\mathcal{K}^{\perp}(x^*)$.

Proof. Let $x_i - x_i^* := \delta_i := \delta_i^{\parallel} + \delta_i^{\perp}$ be infinitesimal, where $\delta^{\parallel} \in \mathcal{K}(x^*)$ and $\delta^{\perp} \in \mathcal{K}^{\perp}(x^*)$. Then for non-zero $\delta^{\perp}$,

$$\begin{aligned} \Gamma_{\alpha}^{(\text{mod})}[x, \xi] &= \Gamma_{\alpha}^{(\text{mod})}[x^*, \xi] + \sum_{i=1}^d (\nabla_x \Gamma_{\alpha}^{(\text{mod})}[x^*, \xi])_i \delta_i + \frac{1}{2} \sum_{i,j=1}^d (\nabla_x^2 \Gamma_{\alpha}^{(\text{mod})}[x^*, \xi])_{ij} \delta_i \delta_j + O(\delta^3) \\ &= \Gamma_{\alpha}^{(\text{mod})}[x^*, \xi] + \frac{1}{2} \sum_{i,j=1}^d (\nabla_x^2 \Gamma_{\alpha}^{(\text{mod})}[x^*, \xi])_{ij} \delta_i^{\perp} \delta_j^{\perp} + O(\delta^3) . \end{aligned} \quad (\text{By (E.3) and (E.4)})$$

Since the Hessian is positive definite on $\mathcal{K}^{\perp}(x^*)$, we obtain

$$\Gamma_{\alpha}^{(\text{mod})}[x, \xi] = \Gamma_{\alpha}^{(\text{mod})}[x^* + \delta, \xi] > \Gamma_{\alpha}^{(\text{mod})}[x^*, \xi] .$$

Hence, the system flows back to the fixed-point, x^* . If the infinitesimal change does not contain the orthogonal complementary component, i.e., $\delta_i = \delta_i^{\parallel}$, then $x^* + \delta^{\parallel} \in \mathcal{N}_{x^*}$. By the assumption of $\mathcal{N}_{x^*}$, $\Gamma_{\alpha}^{(\text{mod})}[x, \xi]$ remains unchanged on $\mathcal{N}_{x^*}$. Hence, the higher-order correction vanishes, i.e., $O(\delta^3) = 0$. Consider the gradient flow:

$$\begin{aligned} \frac{d}{dt} \{ \Gamma_{\alpha}^{(\text{mod})}[x(t), \xi] - \Gamma_{\alpha}^{(\text{mod})}[x^*, \xi] \} &= \sum_{i=1}^d (\nabla_x \Gamma_{\alpha}^{(\text{mod})}[x(t), \xi])_i \frac{dx_i(t)}{dt} \\ &= - \sum_{i=1}^d (\nabla_x \Gamma_{\alpha}^{(\text{mod})}[x(t), \xi])_i^2 \leq 0 , \end{aligned}$$

where we use (E.1). The gradient flow approaches to $(x^* + v)$ for some $v \in \mathcal{K}(x^*)$. For gradient descent,

$$\begin{aligned} \delta_i(t+1) &= x_i(t+1) - x_i^* \\ &= \delta_i(t) - \eta (\nabla_x \Gamma_{\alpha}[x(t), \xi])_i \quad (\text{By (E.2)}) \\ &= \sum_{j=1}^d \{ \mathbb{1}_{\{i,j\}} - \eta (\nabla_x^2 \Gamma_{\alpha}^{(\text{mod})}[x^*, \xi])_{ij} \} \delta_j(t) . \end{aligned} \quad (\text{E.6})$$

To guarantee that the gradient descent converges for $\mathcal{K}^{\perp}(x^*)$, we require that the maximal absolute eigenvalue of $\{\cdot \cdot \cdot\}$ above should be smaller than one. According to (E.4), the eigenvalues are all

positive along $\mathcal{K}^\perp(x^*)$. Hence,

$$|1 - \eta \lambda_{\max}(\nabla_x^2 \Gamma_\alpha^{(\text{mod})}[x^*, \xi])|_{\mathcal{K}^\perp(x^*)}| < 1 \implies 0 < \eta < \frac{2}{\lambda_{\max}(\nabla_x^2 \Gamma_\alpha^{(\text{mod})}[x^*, \xi])|_{\mathcal{K}^\perp(x^*)}}.$$

If we do not restrict on $\mathcal{K}^\perp(x^*)$, then the $\lambda_{\max}(\cdot) = 0$ along $\mathcal{K}(x^*)$. Hence, the gradient descent approaches to $(x^* + v)$ for some $v \in \mathcal{K}(x^*)$. This proves (E.5). $\square$

Remark E.1. For symmetry and projection degeneracy, [Proposition E.1](#) holds since symmetries are always global, i.e., the group element does not differ for different φ , and the projection is affine. However, for constrained systems, we require the condition that Γ_α remains the same on the degeneracy manifold of the constrained space. This does not always hold.

Proposition E.2 (State-Retrieval Stability as Query-Key-Retrieval Stability). The dynamical retrieval stability in the macroscopic state space, $\mathcal{M}_x$, corresponds to which in the memory score z -space, $\mathcal{M}_z$. Specifically, let x^* be the stable fixed-point in $\mathcal{M}_x$, and $\mathcal{D} := \ker(\xi_i^\mu) \subset \mathbb{R}^d$.

$$x(t) \rightarrow x^* \implies z(t) \rightarrow z^* := \frac{1}{\sqrt{d}} \sum_{i=1}^d \xi_i^\mu x_i^*. \quad (\text{E.7})$$

Conversely,

$$z(t) \rightarrow z^* \implies \Pi_{\mathcal{D}^\perp} x(t) \rightarrow \Pi_{\mathcal{D}^\perp} x^*, \quad (\text{E.8})$$

where $\Pi_{\mathcal{D}^\perp}$ is the projection map to $\mathcal{D}^\perp$. Moreover, the convergence rates in z imply convergence rates in x via

$$\|\delta x\|_2 \leq \frac{d}{\sigma_{\min}} \|\delta z\|_2, \quad \forall \delta x \in \mathcal{D}^\perp, \quad (\text{E.9})$$

where $\sigma_{\min} > 0$ denotes the smallest non-zero singular value of (ξ_i^μ) .

Proof. Since $\mathbb{R}^d = \mathcal{D} \oplus \mathcal{D}^\perp$, we decompose $x_i = x_i^\parallel + x_i^\perp$ with $\sum_{i=1}^d \xi_i^\mu x_i^\parallel = 0$. For (E.7), it is trivial since $\bar{z}$ is deterministic due to $z_\mu = \sum_{i=1}^d \xi_i^\mu x_i / \sqrt{d}$. For (E.8), we consider $\delta x(t) = \Pi_{\mathcal{D}^\perp}(x(t) - x^*) \in \mathcal{D}^\perp$. Then

$$\frac{1}{\sqrt{d}} \sum_{i=1}^d \xi_i^\mu \delta x_i(t) = \frac{1}{\sqrt{d}} \sum_{i=1}^d \xi_i^\mu (x_i(t) - x_i^*) = z_\mu(t) - z^* \rightarrow 0.$$

If $\delta x \in \mathcal{D}$, then $\sum_{i=1}^d \xi_i^\mu \delta x_i(t) = \sum_{i=1}^d \xi_i^\mu \Pi_{\mathcal{D}}(x_i(t) - x_i^*) = 0$ does not give any information of $x(t)$ and x^* . We restrict to $\mathcal{D}^\perp$ now. Since on $\mathcal{D}^\perp$, (ξ_i^μ) is injective and $\delta x \in \mathcal{D}^\perp$ satisfies

$$\sum_{\mu=1}^M (z_\mu(t) - z_\mu^*)^2 = \frac{1}{d} \sum_{\mu=1}^M \left(\sum_{i=1}^d \xi_i^\mu \delta x_i(t) \right)^2 \geq \frac{\sigma_{\min}^2}{d} \sum_{i=1}^d \delta x_i^2(t),$$

we thereby prove (E.9). $\square$

Proposition E.3 (Score-space Attention-Retrieval Stability). Consider the fixed-point in the z -space be z_μ^* . Then the Jacobian of the update map is

$$\mathcal{J}_{\mu\nu}(z^*) = -\frac{1}{\beta} \sum_{\rho=1}^M G_{\mu\rho} (\nabla_z^2 \tilde{\Gamma}_\alpha^{(\text{mod})}[x^*, p, \xi])_{\rho\nu}, \quad (\text{E.10})$$

where $z_\mu^* = \sum_{i=1}^d x_i^* \xi_i^\mu / \sqrt{d}$. The dynamical retrieval update is stable at z_μ^* if

$$\rho(\mathcal{J}(z^*)) < 1,$$

where $\rho(\cdot)$ is the spectral radius, i.e., the maximal absolute value of eigenvalues.

Proof. If we perturb around the fixed-point, $\delta z_\mu(t) := z_\mu(t) - z_\mu^*$ gives

$$\begin{aligned} \delta z_\mu(t+1) &= \sum_{\nu=1}^M G_{\mu\nu} p_\nu(z(t)) + b_\mu(J) - \sum_{\nu=1}^M G_{\mu\nu} p_\nu(z^*) - b_\mu(J) \\ &= \sum_{\nu,\rho=1}^M G_{\mu\nu} \frac{\partial p_\nu}{\partial z_\rho} \Big|_{z=z^*} \delta z_\rho(t). \end{aligned}$$

According to (3.8), we express p in terms of the gradient of $\tilde{\Gamma}_\alpha^{(\text{mod})}$. In this way,

$$\delta z_\mu(t+1) = -\frac{1}{\beta} \sum_{\nu,\rho=1}^M G_{\mu\nu} (\nabla_z^2 \tilde{\Gamma}_\alpha^{(\text{mod})}[x^*, p, \xi])_{\nu\rho} \delta z_\rho(t)$$

By $\delta z_\mu(t+1) = \sum_{\nu=1}^M \mathcal{J}_{\mu\nu} \delta z_\nu(t)$, we obtain the Jacobian, (E.10). If $\rho(\mathcal{J}) < 1$, then $\mathcal{J}^t \rightarrow 0$ as $t \rightarrow \infty$. Hence, $\delta z_\mu(t) \rightarrow 0$, and the system converges to the fixed-point, z_μ^* . $\square$

Remark E.2. The update map is a non-linear discrete map. Hence, it is neither a gradient flow nor a gradient descent in general. This fixed-point perturbation analysis is consistent to the correspondence between the gradient-based optimization in $\mathcal{M}_x$.

Now we define the Tsallis divergence which helps us characterizes the training dynamics and the convergences. It generalizes KL-divergence and induces dual geometric structures on the probability simplex. [Amari, 1985]

Definition E.2 (Tsallis Divergence [Furuichi et al., 2004]). The Tsallis divergence generalizes the Kullback-Leibler (KL) divergence [Kullback and Leibler, 1951]. It has the following form:

$$D_\alpha(p||q) := -\sum_{\mu=1}^M p_\mu \log_\alpha \left(\frac{q_\mu}{p_\mu} \right) = \frac{1}{\alpha-1} \left(\sum_{\mu=1}^M p_\mu^\alpha q_\mu^{1-\alpha} - 1 \right).$$

Remark E.3. If $\alpha \rightarrow 1$, the Tsallis divergence reduces to the KL divergence. It follows from [Lemma C.1](#) in which $\log_\alpha \rightarrow \log$ as $\alpha \rightarrow 1$. Hence,

$$D_{\alpha \rightarrow 1}(p||q) = -\lim_{\alpha \rightarrow 1} \sum_{\mu=1}^M p_\mu \log_\alpha \left(\frac{q_\mu}{p_\mu} \right) = \sum_{\mu=1}^M p_\mu \log \left(\frac{p_\mu}{q_\mu} \right) := D_{\text{KL}}(p||q).$$

The dynamical stability of p relates to the attention-retrieval stability via the Tsallis divergence.

Proposition E.4 (Tsallis Divergence and Attention-Retrieval Stability). Let $\alpha > 1$, $p^* \in \Delta^M$ be the fixed maximizer in the simplex interior satisfying $z_\mu^* = (Gp)_\mu^* + b_\mu(J)$. Then there exist constants, $c_1, c_2 > 0$ such that

$$c_1 \|z(t) - z^*\|_2^2 \leq D_\alpha(p(t)||p^*) \leq c_2 \|p(t) - p^*\|_2^2.$$

Proof. Since

$$\frac{\partial}{\partial p_\mu} D_\alpha(p||q) = \frac{\alpha}{\alpha-1} q_\mu^{1-\alpha} p_\mu^{\alpha-1}, \quad \frac{\partial^2}{\partial p_\mu \partial p_\nu} D_\alpha(p||q) = \alpha q_\mu^{1-\alpha} p_\mu^{\alpha-2} \mathbf{1}_{\{\mu, \nu\}},$$

we have the Taylor expansion

$$\begin{aligned} D_\alpha(p(t)||p^*) &= D_\alpha(p^*||p^*) + \frac{\alpha}{\alpha-1} \sum_{\mu=1}^M (p_\mu(t) - p_\mu^*) + \frac{\alpha}{2} \sum_{\mu=1}^M \frac{1}{p_\mu^*} (p_\mu(t) - p_\mu^*)^2 + O(\|p(t) - p^*\|_2^3) \\ &= \frac{\alpha}{2} \sum_{\mu=1}^M \frac{1}{p_\mu^*} (p_\mu(t) - p_\mu^*)^2 + O(\|p(t) - p^*\|_2^3). \end{aligned}$$

Since p^* is in the simplex interior, $1 > p_\mu^* > 0$ for all μ . Hence, there exist $c_1, c_2 > 0$ such that

$$c_1 \|p(t) - p^*\|_2^2 \leq D_\alpha(p(t)||p^*) \leq c_2 \|p(t) - p^*\|_2^2.$$

Moreover,

$$z_\mu(t) - z_\mu^* = [G(p(t) - p^*)]_\mu \implies \|z(t) - z^*\|_2 \leq \|G\|_{\text{op}} \cdot \|p(t) - p^*\|_2.$$

Combining these two inequalities, we prove this proposition. $\square$

Remark E.4. From [Proposition E.4](#), we notice that the convergence in $\mathcal{M}_z$ does not control the convergence of the Tsallis divergence and the convergence in p -space. If the Gram matrix is invertible, then $\|p(t) - p^*\|_2 \leq \|G^{-1}\|_{\text{op}} \cdot \|z(t) - z^*\|_2$. In this way, the convergence in $\mathcal{M}_x$ controls the convergence in $\mathcal{M}_z$ and the p -space. However, in the real world, the high correlations between memory patterns make $G_{\mu\nu}$ non-invertible. Hence, we do not guarantee the convergence of p^* .

F Experiment Setup

In our simulation, [Section 5](#) shows the phase structure of modern Hopfield models by scanning the inverse temperature β , and the exponential memory load $\kappa = \log(M)/d$ in both the sparse ($\alpha = 2$) and dense ($\alpha = 1$) regimes. Here we introduce the support size to characterize the usage of memory weights. Then we provide the experimental setup and further details in [Section F.2](#).

F.1 Effective Support Size and Sparsity

To relate entropy functions to the size and the randomness of information concretely, we introduce the effective support size and the sparsity here.

Definition F.1 (Effective Support Size [\[Hill, 1973\]](#)). We define the (Tsallis-Rényi) effective support size as

$$K_{\text{eff}}^{(\alpha)}[p] := \left(\sum_{\mu=1}^M p_{\mu}^{\alpha} \right)^{\frac{1}{1-\alpha}},$$

where $p \in \Delta^M$ satisfying $\sum_{\mu=1}^M p_{\mu} = 1$ and $\alpha \in \mathbb{R}_+/\{1\}$.

Proposition F.1 (Range of Effective Support Size [\[Hill, 1973\]](#)). The effective support size satisfies the inequality:

$$1 \leq K_{\text{eff}}^{(\alpha)}[p] \leq M$$

If $K_{\text{eff}}^{(\alpha)}[p] = 1$, then p is a peak. If $K_{\text{eff}}^{(\alpha)}[p] = M$, then p is uniform.

Proof. We repeat the proof from [\[Hill, 1973\]](#) for completeness. We decompose the range of α into $\alpha > 1$ and $1 > \alpha > 0$ and focus on the former case first. For $\alpha > 1$, we have $\sum_{\mu=1}^M p_{\mu}^{\alpha} \leq 1$ and

$$\begin{aligned} 1 &\leq \left(\sum_{\mu=1}^M p_{\mu}^{\alpha} \right)^{\frac{1}{1-\alpha}} = K_{\text{eff}}^{(\alpha)}[p] \\ &= M^{\frac{1}{1-\alpha}} \cdot \left(\frac{1}{M} \sum_{\mu=1}^M p_{\mu}^{\alpha} \right)^{\frac{1}{1-\alpha}} \\ &\leq M^{\frac{1}{1-\alpha}} \cdot \left[\left(\frac{1}{M} \sum_{\mu=1}^M p_{\mu} \right)^{\alpha} \right]^{\frac{1}{1-\alpha}} \quad (\text{By the Jensen's Inequality}) \\ &= M. \end{aligned}$$

For $1 > \alpha > 0$, then $1 \leq \sum_{\mu=1}^M p_{\mu}^{\alpha}$ and the similar derivation holds. We note that it saturates 1 when p is a vertex on the simplex, i.e., $p_{\mu_{\max}} = 1$ and $p_{\mu \neq \mu_{\max}} = 0$, where $\mu_{\max} := \arg \max_{\mu} p_{\mu}$ is

non-degenerate. It saturates the Jensen's inequality when $p_\mu = M^{-1}$ for every μ . $\square$

Remark F.1. By [Definition C.8](#),

$$K_{\text{eff}}^{(\alpha)}[p] = (1 - (\alpha - 1)\Psi_\alpha[p])^{\frac{1}{1-\alpha}},$$

we differentiate it with respect to Ψ_α . It yields $(K_{\text{eff}}^{(\alpha)}[p])^{\alpha/(1-\alpha)} > 0$. Hence, the effective support size is strictly increasing (monotonic) with respect to the Tsallis entropy. The higher the entropy is, the larger the effective support size is.

Lemma F.1 (Threshold of Probability Function [[Martins and Astudillo, 2016](#), [Peters et al., 2019](#)]). The minimizer, $p \in \Delta^M$, of $\Gamma_\alpha^{(\text{mod})}[x, p, \xi]$ satisfies

$$p_\mu = \left(\frac{\alpha - 1}{\alpha}\beta[z_\mu - \theta]_+\right)^{\frac{1}{\alpha-1}}, \quad (\text{F.1})$$

where $[u]_+ := \max\{u, 0\}$. We determine θ by the normalization condition, $\sum_{\mu=1}^M p_\mu = 1$. This truncation only appears for $\alpha > 1$. The support of the minimizers is thereby

$$\text{supp}(p) = \{\mu \in [1, M] : z_\mu > \theta\}.$$

We define $|\text{supp}(p)|$ as the exact support size.

Proof. We include the proof for completeness, adapting the standard argument from [[Martins and Astudillo, 2016](#), [Peters et al., 2019](#)] to our notation. Since $0 \leq p_\mu \leq 1$ and $\sum_{\mu=1}^M p_\mu = 1$, we impose these constraints as Lagrangian multiplier terms:

$$\mathcal{L}(p, z, \theta, \lambda) = \beta \sum_{\mu=1}^M z_\mu p_\mu + \frac{1}{\alpha - 1} \left(1 - \sum_{\mu=1}^M p_\mu^\alpha\right) + \beta\theta \left(\sum_{\mu=1}^M p_\mu - 1\right) + \sum_{\mu=1}^M \lambda_\mu p_\mu.$$

Let the optimal saddles that maximizes the Lagrangian function be $\bar{p}$, $\bar{z}$, $\bar{\theta}$, and $\bar{\lambda}$. Then the Karush-Kuhn-Tucker (KKT) conditions are

$$\bar{p}_\mu \geq 0, \quad \sum_{\mu=1}^M \bar{p}_\mu = 1, \quad (\text{Primal})$$

$$\bar{\lambda}_\mu \geq 0, \quad (\text{Dual})$$

$$\bar{\lambda}_\mu \bar{p}_\mu = 0, \quad (\text{Complementary})$$

$$\frac{\partial \mathcal{L}}{\partial p_\mu} \Big|_{p=\bar{p}, z=\bar{z}, \theta=\bar{\theta}, \lambda=\bar{\lambda}} = \beta \bar{z}_\mu - \frac{\alpha}{\alpha - 1} (\bar{p}_\mu)^{\alpha-1} + \beta \bar{\theta} + \bar{\lambda}_\mu = 0, \quad (\text{Stationary})$$

for all $\mu \in [1, M]$. By solving the stationary equation, for $\bar{p}_\mu > 0$ and $\alpha > 1$,

$$\bar{p}_\mu = \left(\frac{\alpha - 1}{\alpha}\beta[\bar{z}_\mu + \bar{\theta}]\right)^{\frac{1}{\alpha-1}}.$$

For $\bar{p}_\mu = 0$ and $\alpha > 1$,

$$\bar{z}_\mu + \bar{\theta} + \bar{\lambda}_\mu = 0 \implies \bar{z}_\mu + \bar{\theta} \leq 0.$$

Hence, after replacing $\bar{\theta} \rightarrow -\bar{\theta}$, we have

$$p_\mu = \left(\frac{\alpha - 1}{\alpha} \beta [z_\mu - \theta]_+ \right)^{\frac{1}{\alpha-1}},$$

where we neglect bars to avoid redundant notations. $\square$

Remark F.2. For $0 < \alpha < 1$, the exponent of (F.1) is negative, so the base should remain positive. In this way, there is no truncation and all minimizers $p_\mu > 0$. For $\alpha \rightarrow 1$, the Shannon entropy has the stationary equation,

$$\beta z_\mu - \log(p_\mu) - 1 + \beta \theta + \lambda_\mu = 0.$$

The minimizer, p_μ , is positive everywhere and does not truncate for finite $(\theta, z_\mu, \lambda_\mu)$.

Definition F.2 (Sparsity [Martins and Astudillo, 2016, Peters et al., 2019]). We say $p \in \Delta^M$ is k -sparse if the exact support size satisfies

$$|\text{supp}(p)| = k,$$

where $0 \leq k \leq M$. The large the k is, the less the sparsity is. In other words, if more points satisfy $p = 0$, then the system is sparser.

Remark F.3. Higher entropy suggests larger effective support size but not higher sparsity. The sparsity depends on the boundary geometry of the entropy function. Lemma F.1 states that only for $\alpha > 1$, there is a truncation and several p_μ are zero. It is consistent with the argument in Proposition C.3. Due to Definition C.4, the Tsallis entropy is not essentially smooth only for $\alpha > 1$. Therefore, for $\alpha > 1$, the optimizer sit on the boundaries of the simplex when the corresponding z_μ is below the threshold. This introduces the sparsity. Otherwise, the divergence of gradient on the boundary forbids the sparsity.

F.2 Settings

Setup. We consider a system of dimension d with $M = e^{\kappa d}$ normalized stored patterns, $\{\xi^\mu\}_{\mu=1}^M$. To probe the energy landscape, $\Gamma_\alpha^{(\text{mod})}$, we perform continuous scans by varying β or κ . For each scan, we use both up and down sweeps to detect hysteresis and metastability.

1. **Up Sweep:** Initialize from low β (or low κ) and increase.
2. **Down Sweep:** Initialize from high β (or high κ) and decrease.

At each parameter value (β or κ), we initialize the optimization from the converged solution at the previous step (warm start). Then we minimize the free energy with respect to x by gradient

descent,

$$x(t+1) = x(t) - \eta(t) \nabla_x \Gamma_\alpha^{(\text{mod})}(x(t)), \quad \eta(t) = \eta_0 2^{-m(t)},$$

where $\eta_0 = 0.5$ and $m(t) := \min\{m \in \mathbb{N}_0 \mid \Gamma_\alpha^{(\text{mod})}(x(t+1)) \leq \Gamma_\alpha^{(\text{mod})}(x(t))\}$. This procedure tracks continuous branches of local minima of $\Gamma_\alpha^{(\text{mod})}$. To explore competing solutions, we perform multi-branch continuation.

Data. In our experiment, we consider clustered normalized patterns.

$$\xi_i^\mu = \sqrt{1-\delta} \eta_i^{c(\mu)} + \sqrt{\delta} \zeta_i^\mu,$$

where $c(\mu) \in \{1, \dots, C\}$ assigns pattern μ to a cluster, $\eta_i^{c(\mu)}$ are cluster centers sampled as $\eta^c \sim \mathcal{N}(0, \mathbb{1}_d)$, $\zeta^\mu \sim \mathcal{N}(0, \mathbb{1}_d)$ are i.i.d. normalized noise vectors, and $\delta \in [0, 1]$ indicates the noise strength. This construction introduces off-diagonal correlations in the Gram matrix, which makes the simulation more realistic than the i.i.d. setting. In our simulation, we choose $C = 4$ and $\delta = 0.04$.

Observables. To characterize the coexistence regimes and retrieval success, we measure the following quantities.

1. **Maximum Overlap:** To verify whether a stable point aligns with the stored patterns, we measure the maximum overlap. $z_\mu^* = \max_\mu |z_\mu^*|$.
2. **Exact Support Size:** Each stable point has its own support and exact support size. For example, $z_\mu^* = \sum_{i=1}^d x_i^* \xi_i^\mu / \sqrt{d}$ gives a specific value of p_μ^* via [Lemma F.1](#), which determines $K_{\text{eff}}^{(\alpha>1)}$. For $\alpha = 2$, we measure $|\text{supp}(p^*)| = \#\{\mu \in [1, M] : z_\mu^* > \theta\}$ according to [Lemma F.1](#), where $\#$ indicates the cardinality. For $\alpha = 1$, we measure $\{\mu \in [1, M] : p_\mu^* > \varepsilon\}$, where we choose $\varepsilon = 10^{-10}$ in our simulation. Hence, we denote the measured exact support size as

$$|\text{supp}(p^*)| := \begin{cases} \#\{\mu \in [1, M] : p_\mu^* > 10^{-10}\} & , \alpha = 1 \\ \#\{\mu \in [1, M] : z_\mu^* > \theta\} & , \alpha = 2 \end{cases}.$$

3. **Fixed Scored-Threshold Support Size:** In temperature analysis in the sparse case, we further define

$$K_\Delta := \{\mu \in [1, M] : z_\mu^* > \Delta\},$$

and we choose $\Delta = 0.01$. Hence, it serves as a diagnostic for counting how many memory components achieve large query-key alignment.

4. **Inverse Participation Ratio (IPR):** It is defined as the effective support size with $\alpha = 2$, i.e., $K_{\text{eff}}^{\text{IPR}} := K_{\text{eff}}^{(2)} = (\sum_\mu (p_\mu^*)^2)^{-1}$. If IPR is large, then there are more components of p_μ^* that are non-zero.

5. **Free Energy:** Distinct stable points have distinct free energy in general. To see which is more preferable, i.e., which one has the lowest free energy, we measure $\Gamma_{\alpha}^{(\text{mod})}[x, \xi]$.

Results. We display the results in **Figures 2 to 5**. Please see **Section 5** for detailed illustrations.

G Additional Experiments

In this section, we first show the experiment of dynamical stability in modern Hopfield models. Then we analyze the empirical retrieval behaviors in realistic end-to-end Transformer architectures. Finally, we validate the special coexistence point formula [Proposition 4.1](#) and spinodal formula [Proposition 4.2](#) numerically.

G.1 Dynamical Stability Analysis for Temperature

In this subsection, we provide the numerical validation for the theoretical dynamical retrieval stability condition; please see [Section E](#) for a detailed discussion and derivations. Specifically, we only analyze the ordinary gradient descent (discrete) algorithm and verify the bound in [Proposition E.1](#).

Setup. We consider the number of pattern be $M = 100$, pattern dimension be $d = 64$, and inverse-temperature be $\beta = 300$ as an example. For simplicity, we apply fixed-step gradient descent

$$x(t+1) = x(t) - \eta \nabla_x \Gamma_\alpha^{(\text{mod})}(x(t)),$$

where η is a constant. Let $\eta_{\text{crit}} = 2/\lambda_{\max}(\nabla^2 \Gamma_\alpha^{(\text{mod})}(x^*))$ be the theoretical critical learning rate threshold, $\lambda_{\max}(\cdot)$ be the maximal eigenvalue, and x^* be a fixed local stable point. We investigate the convergence through iteration numbers, t . We define the initial condition as $x(t=0) = x^* + Au_{\text{soft}}$, where $A = 10^{-2}$ is a small amplitude and u_{soft} is the eigen-direction of the Hessian with the smallest eigenvalue. Notice that we derive u_{soft} numerically, so it is not the exact eigenvector but exhibits tiny component in other directions.

Data. We use the same pattern set $\{\xi^\mu\}_{\mu=1}^M$ as in the previous experiments, where each pattern $\xi^\mu \in \mathbb{R}^d$ is normalized. We generate this collection through the method in [Section F.2](#).

Observables. We track two quantities along the gradient-descent trajectory: the distance to the fixed point, $\|x(t) - x^*\|_2$ and the maximal overlap, $z_{\max}(x(t)) := \max_\mu z_\mu(x(t))$.

Results. In the sparse regime ($\alpha = 2$), we demonstrate the result in [Figure 10](#).

- (a) For $\eta < \eta_{\text{crit}}$, the distance (blue curve) approaches to zero asymptotically, indicating that the training $x(t)$ converges to the fixed stable point, x^* . At $\eta \approx \eta_{\text{crit}}$, the convergence becomes marginal. For $\eta > \eta_{\text{crit}}$, the training $x(t)$ deviates from the local point as iteration increases. This indicates a dynamical instability. Therefore, this plot provides the evidence that η_{crit} separates the convergent regime.
- (b) The maximal overlap $z_{\max}(x(t))$ exhibits the same behavior. For $\eta < \eta_{\text{crit}}$, $z_{\max}(x(t))$ converges to the fixed value $z_{\max}(x^*)$. This indicates stable retrieval of the dominant memory pattern. Notably, the convergent value is $z_{\max}(x^*) \approx 1/d \approx 0.0156$, which is the one-pattern




Figure 10: **Dynamical retrieval stability of MHM gradient descent in the sparse regime** ($\alpha = 2$). ($M = 100, d = 64$). We study the gradient descent update $x(t+1) = x(t) - \eta \nabla \Gamma_{\alpha}^{(\text{mod})}(x(t))$. We denote the critical learning rate as $\eta_{\text{crit}} = 2/\lambda_{\text{max}}(\nabla^2 \Gamma_{\alpha}^{(\text{mod})}(x^*))$ with λ_{max} the largest eigenvalue at a fixed stable point x^* . **(a) Distance to the fixed point.** For $\eta < \eta_{\text{crit}}$, the training converges to the fixed point. Otherwise, the dynamics is unstable or even diverges. **Maximal overlap.** For $\eta < \eta_{\text{crit}}$, the training converges to the fixed value of overlap $z_{\text{max}} \approx 0.0156$. It is the saturation value of one-pattern perfect retrieval in Figure 2 (a). Otherwise, the dynamics is unstable or even diverges.

retrieval in Figure 2 (a). At $\eta \approx \eta_{\text{crit}}$, it exhibits persistent oscillations around $z_{\text{max}}(x^*)$. For $\eta > \eta_{\text{crit}}$, the oscillations grow, indicating a dynamical instability in retrieval.

In the dense regime ($\alpha = 1$), we demonstrate the result in Figure 11.

- (a) For $\eta < \eta_{\text{crit}}$, the iteration converges to the stable fixed point x^* rapidly. Larger the learning rate is, more rapid the convergence is. For $\eta > \eta_{\text{crit}}$, the distance declines first, and then grows up after some number of iteration. This phenomenon arises from unstable modes being small due to our initialization $x(t=0)$. In particular, from (E.6), we have $\delta_i(t+1) = (1 - \eta\lambda_i)\delta_i(t)$, where $\delta_i(t) \equiv x_i(t) - x_i^*$, and λ_i is the i -th eigenvalue of the Hessian $\nabla_x^2 \Gamma_{\alpha}^{(\text{mod})}$. The convergence condition is

$$|1 - \eta\lambda_i| < 1, \quad \text{or } 0 < \lambda_i < 2/\eta.$$

Hence, even if $\eta > \eta_{\text{crit}}$, some modes $x_i(t)$ still converge to the fixed point x_i^* if the eigenvalues satisfy $\lambda_i < 2/\eta$. Initially, we consider the dominant mode be which with λ_{min} . Hence, it converges to the fixed point in the early stage. However, for $\lambda_i > 2/\eta$, $\delta_i(t)$ grows with t . Hence, the unstable modes start to dominate, and the trajectory diverges.

- (b) For $\eta < \eta_{\text{crit}}$, the iteration converges to the stable fixed value $z_{\text{max}}(x^*)$. Larger the learning rate is, more rapid the convergence is. For $\eta > \eta_{\text{crit}}$, it exhibits the same behavior as in distance analysis. The dynamics converges first, then it diverges after certain iteration. This arises from unstable modes being small due to our initialization $x(t=0)$ as well.




Figure 11: **Dynamical retrieval stability of MHM gradient descent in the dense regime** ($\alpha = 1$). ($M = 100, d = 64$). We study the gradient descent update $x(t+1) = x(t) - \eta \nabla \Gamma_{\alpha}^{(\text{mod})}(x(t))$. We denote the critical learning rate as $\eta_{\text{crit}} = 2/\lambda_{\text{max}}(\nabla^2 \Gamma_{\alpha}^{(\text{mod})}(x^*))$ with λ_{max} the largest eigenvalue at a fixed stable point x^* . **(a) Distance to the fixed point.** For $\eta < \eta_{\text{crit}}$, the training converges to the fixed point. For $\eta > \eta_{\text{crit}}$, a divergence follows after a decay. This reflects that the iterate close to x^* initially. After this, the unstable eigen-directions start to dominate. We provide the precise explanation in this section. **(b) Maximal overlap.** For $\eta < \eta_{\text{crit}}$, the training converges to the fixed value of overlap z_{max} . For $\eta > \eta_{\text{crit}}$, a growing oscillations follow after an initial convergence. This is similar for the behavior in the distance plot.

G.2 Support and Capacity-Reliability Trade-off

We study how increasing the number of memories affects retrieval reliability and effective support (IPR) in a trained Transformer. We compare three attention maps, $\alpha \in \{1, 1.5, 2\}$ (softmax, entmax-1.5, and sparsemax), respectively.

Setup. The data generation pipeline is

$$u_Q \xrightarrow[\text{fixed}]{A_Q} q^{\text{in}}, \quad \text{and} \quad u_{\mu} \xrightarrow[\text{fixed}]{A_K} k_{\mu}^{\text{in}}.$$

We generate latent query and key variables and map them to observed inputs. After that, we input q^{in} and k_{μ}^{in} into a multi-head cross-attention Transformer, yielding the output $\hat{y} \in \mathbb{R}^{d_v}$.

Data. Each example contains a query $q^{\text{in}} \in \mathbb{R}^{d_{\text{in}}}$, a set of memory keys $k_{\mu}^{\text{in}} \in \mathbb{R}^{d_{\text{in}}}$, and independent values $v_{\mu} \sim \mathcal{N}(0, \mathbb{1}_{d_v})$ with $d_v = 32$. The latent dimension is $d_{\text{lat}} = 16$, the input dimension is $d_{\text{in}} = 24$, and the number of latent clusters is $C = 4$. For each example, we draw normalized

cluster centers independently

$$c_r = \frac{\tilde{c}_r}{\|\tilde{c}_r\|_2}, \quad \tilde{c}_r \sim \mathcal{N}(0, \mathbb{1}_{d_{\text{lat}}}), \quad r \in \{1, \dots, C\}.$$

Memory μ is assigned a cluster index $a(\mu)$, and its latent key is

$$u_\mu = \frac{\sqrt{1-\delta} c_{a(\mu)} + \sqrt{\delta} \epsilon_\mu}{\|\sqrt{1-\delta} c_{a(\mu)} + \sqrt{\delta} \epsilon_\mu\|_2}, \quad \epsilon_\mu \sim \mathcal{N}(0, \mathbb{1}_{d_{\text{lat}}}).$$

We choose $\delta = 0.12$ to control the latent key variance. The observed query and keys are then

$$q^{\text{in}} = A_Q u_Q + \sigma_{\text{obs}} \eta_Q, \quad k_\mu^{\text{in}} = A_K u_\mu + \sigma_K \eta_\mu,$$

where $A_Q, A_K \in \mathbb{R}^{d_{\text{in}} \times d_{\text{lat}}}$ with the orthogonalization $A_Q^\top A_Q = A_K^\top A_K = \mathbb{1}_{\text{lat}}$, $\eta_Q, \eta_\mu \sim \mathcal{N}(0, \mathbb{1}_{\text{in}})$, $\sigma_{\text{obs}} = 0.03$, and $\sigma_K = 0.05$. Here we divide u_Q generation into two cases, the selective retrieval case and the uniform-aggregation case.

1. **Selective retrieval.** In this case, u_Q retrieves one particular memory. Consider that we sample a target memory index t from $\{1, \dots, M\}$. Therefore, we denote u_t as the latent key located at the sampled target slot t . We then generate the latent query by adding noise

$$u_Q = \frac{u_t + \sigma_Q \epsilon_Q}{\|u_t + \sigma_Q \epsilon_Q\|_2}, \quad \epsilon_Q \sim \mathcal{N}(0, \mathbb{1}_{d_{\text{lat}}}).$$

During training, the query-noise level cycles through $\sigma_Q \in \{0.05, 0.10, 0.15\}$ and the number of memories cycles through $M = \{8, 16, 32, 64\}$ simultaneously. During evaluation, the query noise is fixed at $\sigma_Q = 0.10$ while M is increasing.

2. **Uniform-aggregation.** In this case, we sample a target cluster r^* . Then the latent query is

$$u_Q = \frac{c_{r^*} + \sigma_Q \epsilon_Q}{\|c_{r^*} + \sigma_Q \epsilon_Q\|_2}, \quad \epsilon_Q \sim \mathcal{N}(0, \mathbb{1}_{d_{\text{lat}}}).$$

We denote the relevant set as $R = \{\mu : a(\mu) = r^*\}$. Since there are $C = 4$ clusters, we choose $|R| = M/4$ in this case.

Transformer. Inside the Transformer, we apply one-layer multi-head cross-attention. Specifically, we conduct the first step via

$$q_h = W_h^Q \text{LN}_Q(W_Q^{\text{in}} q^{\text{in}} + r_Q), \quad k_{h,\mu} = W_h^K \text{LN}_K(W_K^{\text{in}} k_\mu^{\text{in}} + r_K),$$

$$v_{h,\mu} = W_h^V \text{LN}_V(W_V^{\text{in}} v_\mu + r_V),$$

where $W_Q^{\text{in}}, W_K^{\text{in}} \in \mathbb{R}^{d_{\text{mod}} \times d_{\text{in}}}$, $W_V^{\text{in}} \in \mathbb{R}^{d_{\text{mod}} \times d_v}$, $W_h^Q, W_h^K, W_h^V \in \mathbb{R}^{d_h \times d_{\text{mod}}}$, LN denotes the Layer-Norm, and $h = \{1, \dots, H\}$. Then the multi-head output is followed by two residual updates

$$x_1 = W_Q^{\text{in}} q^{\text{in}} + r_Q + W^O \text{Concat} \left(\sum_{\mu=1}^M p_{1,\mu} v_{1,\mu}, \dots, \sum_{\mu=1}^M p_{H,\mu} v_{H,\mu} \right) \in \mathbb{R}^{d_{\text{mod}}},$$

and

$$x_2 = x_1 + \text{FFN}(\text{LN}_{\text{FFN}}(x_1)) \in \mathbb{R}^{d_{\text{mod}}},$$

where $W^O \in \mathbb{R}^{d_{\text{mod}} \times d_{\text{mod}}}$, $p_{h,\mu}$ is the attention with the form [Proposition 3.3](#), and $\text{FFN}(x) := W_2^{\text{FFN}} \text{GELU}(W_1^{\text{FFN}} x)$. Finally, the final predicted output is

$$\hat{y} = W^{\text{out}} \text{LN}_{\text{final}}(x_2) \in \mathbb{R}^{d_v},$$

where $W^{\text{out}} \in \mathbb{R}^{d_v \times d_{\text{mod}}}$. We choose the architectural hyperparameters as

- Model dimension $d_{\text{mod}} = d_h \times H = 96$.
- 4 attention heads ($H = 4$) with head dimension $d_h = 24$.
- FFN with hidden dimension 192, i.e., $W_1^{\text{FFN}} \in \mathbb{R}^{192 \times 96}$, and $W_2^{\text{FFN}} \in \mathbb{R}^{96 \times 192}$.

In the Training stage, we choose 2500 AdamW steps, batch size $B = 64$, learning rate $\eta = 3 \times 10^{-4}$, weight decay 10^{-4} , MSE loss, $\beta = 1$, $M \in \{8, 16, 32, 64\}$. Specifically, the MSE we evaluate is

$$\text{MSE} = \frac{1}{B} \sum_{b=1}^B \left\| \hat{y}^{(b)} - y_{\text{target}}^{(b)} \right\|_2^2,$$

where $y_{\text{target}}^{(b)} = y_{t_b}$ for the selective retrieval case, and $y_{\text{target}}^{(b)} = \sum_{\mu \in R_b} y_{\mu}^{(b)} / |R_b|$ for the uniform-aggregation case.

In inference stage, we fix $\beta = 1$ and the learned parameters, e.g., $W_h^Q, W^O, r_Q, \dots$. Then we increase the memory load from $M = 8$ to $M = 4096$ to measure how good the prediction $\hat{y}$ is.

Observables. We measure

- **Selective retrieval: Value top-1 accuracy.** We define the value top-1 accuracy as

$$A_{\text{value}} := \frac{1}{B} \sum_{b=1}^B \mathbb{1} \left\{ \arg \min_{\mu} \left\| \hat{y}^{(b)} - y_{\mu}^{(b)} \right\|_2^2 = t_b \right\}.$$

This measures end-to-end success, rather than simply measures whether attention puts its largest weight on the correct key.

- **Selective retrieval: Support size (IPR).** As introduced in [Section F.2](#), we measure the support via the inverse partition ratio (IPR). However, since we consider batch size and

multi-head here, we further average over them, i.e.,

$$K_{\text{eff}}^{\text{IPR}} := \frac{1}{HB} \sum_{b=1}^B \sum_{h=1}^H \frac{1}{\sum_{\mu=1}^M p_{b,h,\mu}^2}.$$

- **Uniform-aggregation: Error/zero-output baseline.** The target is $y_{\text{target}} = \sum_{\mu \in R} y_{\mu} / |R|$, where $y_{\mu} \sim \mathcal{N}(0, \mathbf{1}_{d_v})$. Hence, the zero-output baseline is $E_0 := \mathbb{E} \|y_{\text{target}}\|_2^2 = d_v / |R|$. In this way, we define the error/zero-output baseline E_{rel} as

$$E_{\text{rel}} := \frac{\text{MSE}}{E_0} = \frac{|R|}{Bd_v} \sum_{b=1}^B \left\| \hat{y}^{(b)} - \frac{1}{|R_b|} \sum_{\mu \in R_b} y_{\mu}^{(b)} \right\|_2^2.$$

- **Uniform-Aggregation: IPR / $|R|$.** We measure the normalized IPR with the size of the target set

$$\rho_{\text{IPR}} := \frac{K_{\text{eff}}^{\text{IPR}}}{|R|} = \frac{1}{HB|R|} \sum_{b=1}^B \sum_{h=1}^H \frac{1}{\sum_{\mu=1}^M p_{b,h,\mu}^2}.$$

Results. We present the experiments as follows.




Figure 12: Capacity-reliability trade-off for selective retrieval. We evaluate the trained Transformer at fixed $\beta_{\text{train}} = 1$ while increasing M . **(a).** We measure the value top-1 accuracy to characterize the retrieval reliability. Sparsemax retain the highest accuracy. Furthermore, their accuracy decreases as M grows, indicating the capacity-reliability trade-off. **(b).** Softmax support remains full, while sparsemax use the least supports and is the most support-efficient.

From **Figure 12 (a)**, the sparsemax arrives the best value top-1 accuracy, while the softmax arrives the worst one. Furthermore, both of their accuracy decrease as M increases. This indicates that high memory load leads to lower reliability, and reveals the capacity-reliability trade-off. From **Figure 12 (b)**, the support of softmax remains full, while which of sparsemax is low. This suggests that sparsemax is support-efficient. Those plots conclude that for a selective task, sparsemax has not only high-accuracy but also high support-efficiency.




Figure 13: **Capacity-reliability trade-off for uniform-aggregated retrieval.** We evaluate the trained Transformer at fixed $\beta_{\text{train}} = 1$ while increasing M . **(a).** We measure the ratio of the output error and the zero-output baseline to characterize the retrieval reliability. Sparsemax exhibits the worst reliability. Furthermore, their retrieval reliabilities become worse as M grows, indicating the capacity-reliability trade-off. **(b).** Softmax support remains full, while sparsemax use the least supports and is the most support-efficient.

Compared to Figure 12, Figure 13 (a) shows that softmax arrives the highest reliability among three attentions and the capacity-reliability trade-off as well. From Figure 13 (b), we still observe that the sparsemax uses the least aggregated support size, and hence the most support-efficient. In conclusion, although a large α leads to support-efficiency, the reliability depends on the task: sparsemax predicts a selective retrieval the best, while softmax predicts a uniform-aggregation retrieval the best.

G.3 Support, Temperature, and Reliability

In this section, we investigate the temperature effect on effective support size (IPR) and retrieval reliability in a realistic Transformer architecture numerically. We compare three attention maps, $\alpha \in \{1, 1.5, 2\}$ (softmax, entmax-1.5, and sparsemax), respectively.

Setup. We adopt the same setup and Transformer architecture as in Section G.2, i.e., one-layer four-head cross-attention Transformer with MSE training objective. The only different setup is that we isolate β by fixing M at $M = 128$ with query-noise level $\sigma_Q = 0.10$ in the inference stage. We vary the inverse-temperature over $\beta \in \{2^{-2}, 2^{-1}, 2^0, \dots, 2^9\}$. Furthermore, we compare two protocols:

- **Frozen- β .** In this case, we first train each model at $\beta_{\text{train}} = 1$. After training, we freeze every learned parameter (W_h^Q, W^O, r_Q , and so forth). At inference, we observe how reliable the output performance is at each $\beta \in \{2^{-2}, 2^{-1}, 2^0, \dots, 2^9\}$.
- **Matched- β .** In this case, we train a separate Transformer for every β . That is, for different β , we have different trained parameters W_h^Q, W^O, r_Q , and so forth.




Figure 14: **Temperature effect on frozen-parameter Transformer.** We train the Transformer parameters at $\beta_{\text{train}} = 1$, then we fix those parameters, set $M = 128$ and vary β during inference. **(a).** The three attention families both exhibit high accuracy for a selective retrieval task. **(b).** For selective retrieval, the sparsemax uses the least effective support size (IPR), which shows the support-efficiency and reliability of high- β sparsemax. **(c).** For a uniform-aggregation task, the three attention families have worse error relative to the zero-output baseline as β increases. **(d).** The IPR relative to target-set size decline as β grows, which indicates that the retrieval becomes selective. This leads to a mismatch for multi-memory aggregation.

Observables. We adopt the same observables as which in [Section G.2](#).

Results. We present the experiments as follows.

From [Figure 14](#) (a) and (b), we observe that sparsemax has highest accuracy and support-efficiency for frozen- β case in a selective retrieval task. Furthermore, both of the accuracy and IPR saturates at high β , arriving a single-pattern reliable retrieval for sparsemax. For softmax, the support is full, but the IPR shows that its retrieval becomes strongly concentrated. For the uniform-aggregation task, [Figure 14](#) (c) and (d) show that all three frozen models perform far worse than the zero-output baseline in the high- β regime. That is, since we train the parameters only at $\beta = 1$, the Transformer exhibits greater selectivity, but creates a mismatch for multi-memory aggregation.




Figure 15: **Temperature effect on matched-parameter Transformer.** We train the Transformer parameters at each β , then we evaluate with $M = 128$ during inference. **(a).** All three attention families maintain high accuracy in a selective retrieval task. **(b).** Sparsemax and entmax-1.5 are support-efficient in a selective retrieval task. **(c).** Softmax has the highest reliability among the three in a uniform-aggregation task. **(d).** Softmax uses numerous memories, which indicates a poor support-efficiency.

From Figure 15 (a) and (b), all three attention families maintain perfect accuracy in a selective retrieval task. Furthermore, sparsemax and entmax-1.5 are more support-efficient than softmax. For the uniform-aggregation task, Figure 15 (c) and (d) show that softmax achieves the highest reliability but its effective support remains broader than the relevant set. Again, this indicates that high β increases the selectivity, but achieves poor performance for an aggregation task.

G.4 Numerical Validation of Coexistence and Spinodal Formula

In this section, we validate the coexistence formula in Proposition 4.1 and the spinodal formula in Proposition 4.2 for $\alpha = 2$ numerically.

Setup. We test the special two-key formulas under the uniform distribution and well-separation assumptions of Proposition 4.1 and Proposition 4.2. Then we construct a second dataset to test the full spinodal formula on larger supports.

Data. We construct two experiments as follows.

- **Two-key systems.** We consider two unit-norm keys $\tilde{k}_1$ and $\tilde{k}_2$ satisfying $\tilde{k}_1^\top \tilde{k}_2 = \epsilon$. Therefore, we set

$$\tilde{k}_1 = e_1, \quad \tilde{k}_2 = \epsilon e_1 + \sqrt{1 - \epsilon^2} e_2$$

Then we apply a random orthogonal transformation Q_r to both keys

$$k_1 = Q_r \tilde{k}_1, \quad k_2 = Q_r \tilde{k}_2,$$

where the Gram matrix is preserved under this rotation, but removes the coordinate dependence. We choose the state dimension as $d = 16$, correlation as $\epsilon \in \{0, 0.05, 0.10 \dots, 0.85\}$ and use 5 independent rotation. Hence, this gives $5 \times 18 = 90$ two-key systems.

- **General k -support systems.** In this experiment, we use $k \in \{3, 4, 5, 6, 8\}$ each with 10 independently generated systems and ambient dimension $d = 20$. For each key index $\mu \in \{0, \dots, k-1\}$, we construct the feature vector

$$\tilde{k}_\mu = \left(b, \left\{ a_\ell \cos\left(\frac{2\pi\ell\mu}{k}\right), a_\ell \sin\left(\frac{2\pi\ell\mu}{k}\right) \right\}_{\ell=1}^m \right), \quad m = \lfloor \frac{k-1}{2} \rfloor,$$

where we sample $b \sim \text{Unif}(1.5, 3.0)$, $a_\ell \sim \text{Unif}(0.15, 0.80)$ independently. Then we normalize every vector and apply a random orthogonal rotation to k_μ .

Observables. We measure

- **Two-key: Special coexistence formula.** Under a uniform distribution and well-separation condition, the coexistence point occurs at

$$\beta_{\text{coexist}} = \frac{2}{1 - \epsilon}.$$

This is the special closed-form result of [Proposition 4.1](#).

- **Two-key: Numerical energy crossing.** We define the free energy difference

$$\Delta\Gamma(\beta) = \Gamma_{\alpha=2}^{(\text{mod})}(x_{\text{loc}}; \beta) - \Gamma_{\alpha=2}^{(\text{mod})}(x_{\text{mix}}; \beta),$$

where $x_{\text{loc}} := k_1$ and $x_{\text{mix}} := (k_1 + k_2)/2$. Then we compute the coexistence point from $\Delta\Gamma(\hat{\beta}_{\text{coexist}}) = 0$ numerically.

- **Two-key: Relative error.** We define the coexistence relative error as

$$E_{\text{coexist}} := \frac{|\hat{\beta}_{\text{coexist}} - \beta_{\text{coexist}}|}{\beta_{\text{coexist}}}.$$

For the spinodal formula, we have $\beta_{\text{sp}} = 2/(1 - \epsilon)$ in the two-key systems. We obtain the

empirical spinodal point by solving $\widehat{c}_h(\widehat{\beta}_{\text{sp}}) = 0$ numerically, where

$$\widehat{c}_h(\beta) = \frac{\Gamma_{\alpha=2}^{(\text{mod})}(x^* + hu; \beta) + \Gamma_{\alpha=2}^{(\text{mod})}(x^* - hu; \beta) - 2\Gamma_{\alpha=2}^{(\text{mod})}(x^*; \beta)}{h^2},$$

$h = 2 \times 10^{-5}$, and $u := (k_1 - k_2)/\|k_1 - k_2\|_2$. Then the spinodal relative error is defined as

$$E_{\text{sp}} := \frac{|\widehat{\beta}_{\text{sp}} - \beta_{\text{sp}}|}{\beta_{\text{sp}}}.$$

- **General k -support: Spinodal formula.** From [Proposition 4.2](#), the predicted critical spinodal point is

$$\beta_{\text{sp}} = \frac{2}{\lambda_{\max}(PGP)},$$

where $P_{\mu\nu} = \mathbb{1}_{\{\mu,\nu\}} - k^{-1}$, and $G_{\mu\nu} = k_{\mu}^{\top} k_{\nu}/d$. We compute the empirical spinodal point $\widehat{\beta}_{\text{sp}}$ from the equation $\widehat{c}_h(\widehat{\beta}_{\text{sp}}) = 0$ numerically.

Results. We present the experiments as follows.




Figure 16: Two-key coexistence and spinodal validation. We consider two unit keys with correlation ϵ . **(a).** The solid blue curve presents $\beta_{\text{coexist}} = 2/(1 - \epsilon)$. The numerical points are the root of the free energy difference $\Delta\Gamma(\widehat{\beta}_{\text{coexist}}) = 0$. The numerical energy crossings coincide with the theoretical curve over the full correlation range. **(b).** We plot the relative errors of energy crossing (coexistence) and curvature zero (spinodal). The low values validate the theoretical formulas in [Proposition 4.1](#) and [Proposition 4.2](#).

From [Figure 16 \(a\)](#), the blue curve represents the theoretical formula $\beta_{\text{coexist}} = 2/(1 - \epsilon)$. The numerical energy crossing points lie on the theoretical curve. This validates the special coexistence formula. For [Figure 16 \(b\)](#), we plot the relative errors of energy crossing (coexistence) and curvature zero (spinodal). The maximal relative error of the coexistence formula is 5.73×10^{-15} ,

and which of the spinodal formula is 1.72×10^{-6} . They both validate the theoretical formulas in [Proposition 4.1](#) and [Proposition 4.2](#) under the uniform distribution and well-separation conditions.




Figure 17: General k -support spinodal validation. We test the general k -support spinodal formula $\beta_{\text{jump}} = 2/\lambda_{\max}(PGP)$. We consider 50 systems with unit keys. The numerical points agree with the theoretical values, which validates [Proposition 4.2](#).

In [Figure 17](#), we compare the theoretical spinodal formula [Proposition 4.2](#) and the empirical spinodal values. Across all 50 general k -support systems, the predicted β_{sp} aligns with the empirical one closely, with the maximal relative error 1.92×10^{-6} . Hence, the agreement supports not only the qualitative existence of the transitions, but also the quantitative location of the predicted critical inverse-temperatures.

References

- Shun-ichi Amari. *α -Divergence and α -Projection in Statistical Manifold*, pages 66–103. Springer New York, New York, NY, 1985. ISBN 978-1-4612-5056-2. doi: 10.1007/978-1-4612-5056-2_3.
- Daniel J. Amit, Hanoch Gutfreund, and H. Sompolinsky. Storing infinite numbers of patterns in a spin-glass model of neural networks. *Phys. Rev. Lett.*, 55:1530–1533, Sep 1985. doi: 10.1103/PhysRevLett.55.1530.
- Daniel J Amit, Hanoch Gutfreund, and H Sompolinsky. Statistical mechanics of neural networks near saturation. *Annals of Physics*, 173(1):30–67, 1987. ISSN 0003-4916. doi: [https://doi.org/10.1016/0003-4916\(87\)90092-3](https://doi.org/10.1016/0003-4916(87)90092-3).
- Andreas Auer, Martin Gauch, Daniel Klotz, and Sepp Hochreiter. Conformal prediction for time series with modern hopfield networks. *Advances in neural information processing systems*, 36: 56027–56074, 2023.
- K Binder. Theory of first-order phase transitions. *Reports on Progress in Physics*, 50(7):783, jul 1987. doi: 10.1088/0034-4885/50/7/001.
- Mathieu Blondel, André FT Martins, and Vlad Niculae. Learning with fenchel-young losses. *Journal of Machine Learning Research*, 21(35):1–69, 2020.
- Sidney Coleman and Erick Weinberg. Radiative corrections as the origin of spontaneous symmetry breaking. *Phys. Rev. D*, 7:1888–1910, Mar 1973. doi: 10.1103/PhysRevD.7.1888.
- Nima Dehmamy, Benjamin Hoover, Bishwajit Saha, Leo Kozachkov, Jean-Jacques Slotine, and Dmitry Krotov. Nrgpt: An energy-based alternative for gpt. *arXiv preprint arXiv:2512.16762*, 2025.
- Amir. Dembo and Ofer. Zeitouni. *Large deviations techniques and applications*. Stochastic modelling and applied probability, 38. Springer, Berlin, 2nd ed. edition, 2010. ISBN 9783642033117.
- Mete Demircigil, Judith Heusel, Matthias Löwe, Sven Upgang, and Franck Vermet. On a model of associative memory with huge storage capacity. *Journal of Statistical Physics*, 168(2):288–299, May 2017. ISSN 1572-9613. doi: 10.1007/s10955-017-1806-y.
- Henry R. N. B. Enniful, Daniel Schneider, Antonia Hoppe, Sandra König, Michael Fröba, Dirk Enke, and Rustem Valiullin. Comparative gas sorption and cryoporometry study of mesoporous glass structure: Application of the serially connected pore model. *Frontiers in Chemistry*, Volume 7 - 2019, 2019. ISSN 2296-2646. doi: 10.3389/fchem.2019.00230.
- Sacha Friedli and Yvan Velenik. *Statistical Mechanics of Lattice Systems: A Concrete Mathematical Introduction*. Cambridge University Press, 2017.

- S. Furuichi, K. Yanagi, and K. Kuriyama. Fundamental properties of tsallis relative entropy. *Journal of Mathematical Physics*, 45(12):4868–4877, 12 2004. ISSN 0022-2488. doi: 10.1063/1.1805729.
- Corrado Gini. *Variabilità e mutabilità*. Tipografia di Paolo Cuppini, Bologna, 1912.
- M. O. Hill. Diversity and evenness: A unifying notation and its consequences. *Ecology*, 54(2): 427–432, 1973. doi: <https://doi.org/10.2307/1934352>.
- Benjamin Hoover, Yuchen Liang, Bao Pham, Rameswar Panda, Hendrik Strobelt, Duen Horng Chau, Mohammed Zaki, and Dmitry Krotov. Energy transformer. *Advances in neural information processing systems*, 36:27532–27559, 2023.
- J J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558, 1982. doi: 10.1073/pnas.79.8.2554.
- Jerry Yao-Chieh Hu, Donglin Yang, Dennis Wu, Chenwei Xu, Bo-Yu Chen, and Han Liu. On sparse modern hopfield model. *Advances in neural information processing systems*, 36:27594–27608, 2023.
- Jerry Yao-Chieh Hu, Pei-Hsuan Chang, Haozheng Luo, Hong-Yu Chen, Weijian Li, Wei-Po Wang, and Han Liu. Outlier-efficient hopfield layers for large transformer-based models. In *International Conference on Machine Learning*. PMLR, 2024a.
- Jerry Yao-Chieh Hu, Dennis Wu, and Han Liu. Provably optimal memory capacity for modern hopfield models: Transformer-compatible dense associative memories as spherical codes. *Advances in Neural Information Processing Systems*, 37:70693–70729, 2024b.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Scott Kirkpatrick. Frustration and ground-state degeneracy in spin glasses. *Phys. Rev. B*, 16: 4630–4641, Nov 1977. doi: 10.1103/PhysRevB.16.4630.
- Felix Koulischer, Cédric Goemaere, Tom Van Der Meersch, Johannes Deleu, and Thomas De-meester. Exploring the temperature-dependent phase transition in modern hopfield networks. *arXiv preprint arXiv:2311.18434*, 2023.
- Dmitry Krotov. A new frontier for hopfield networks. *Nature Reviews Physics*, 5, 05 2023. doi: 10.1038/s42254-023-00595-y.
- Dmitry Krotov and John Hopfield. Dense associative memory is robust to adversarial inputs. *Neural Computation*, 30(12):3151–3167, 12 2018. ISSN 0899-7667. doi: 10.1162/neco_a_01143.

- Dmitry Krotov and John Hopfield. Large associative memory problem in neurobiology and machine learning. *arXiv preprint arXiv:2008.06996*, 2020.
- Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86, 1951.
- L. D. Landau and E. M. Lifshitz. *Statistical Physics, Part 1*, volume 5 of *Course of Theoretical Physics*. Butterworth-Heinemann, Oxford, 1980. ISBN 978-0-7506-3372-7. doi: 10.1016/c2009-0-24487-4.
- Jason D Lee, Max Simchowitz, Michael I Jordan, and Benjamin Recht. Gradient descent converges to minimizers. *arXiv preprint arXiv:1602.04915*, 2016.
- Andre Martins and Ramon Astudillo. From softmax to sparsemax: A sparse model of attention and multi-label classification. In *International conference on machine learning*, pages 1614–1623. PMLR, 2016.
- M Mezard, G Parisi, and M Virasoro. *Spin Glass Theory and Beyond*. WORLD SCIENTIFIC, 1986. doi: 10.1142/0271.
- Beren Millidge, Tommaso Salvatori, Yuhang Song, Thomas Lukasiewicz, and Rafal Bogacz. Universal hopfield networks: A general framework for single-shot associative memory models. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 15561–15583. PMLR, 17–23 Jul 2022.
- Moshe Moshe and Jean Zinn-Justin. Quantum field theory in the large n limit: a review. *Physics Reports*, 385(3–6):69–228, October 2003. ISSN 0370-1573. doi: 10.1016/s0370-1573(03)00263-1.
- Jan Naudts. Generalized thermostatics and mean-field theory. *Physica A: Statistical Mechanics and its Applications*, 332:279–300, February 2004. ISSN 0378-4371. doi: 10.1016/j.physa.2003.10.013.
- Hidetoshi Nishimori. In *Statistical Physics of Spin Glasses and Information Processing: An Introduction*. Oxford University Press, 07 2001. ISBN 9780198509417. doi: 10.1093/acprof:oso/9780198509417.001.0001.
- Lars Onsager. Crystal statistics. i. a two-dimensional model with an order-disorder transition. *Phys. Rev.*, 65:117–149, Feb 1944. doi: 10.1103/PhysRev.65.117.
- Ben Peters, Vlad Niculae, and André FT Martins. Sparse sequence-to-sequence models. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1504–1519, 2019.

- Hubert Ramsauer, Bernhard Schäfl, Johannes Lehner, Philipp Seidl, Michael Widrich, Thomas Adler, Lukas Gruber, Markus Holzleitner, Milena Pavlović, Geir Kjetil Sandve, et al. Hopfield networks is all you need. *arXiv preprint arXiv:2008.02217*, 2020.
- R Tyrrell Rockafellar. *Convex analysis*, volume 28. Princeton university press, 1997.
- Saul Santos, Vlad Niculae, Daniel McNamee, and André FT Martins. Hopfield-fenchel-young networks: A unified framework for associative memory retrieval. *Journal of Machine Learning Research*, 26(265):1–51, 2025.
- Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948.
- David Sherrington and Scott Kirkpatrick. Solvable model of a spin-glass. *Phys. Rev. Lett.*, 35: 1792–1796, Dec 1975. doi: 10.1103/PhysRevLett.35.1792.
- M. Srednicki. *Quantum field theory*. Cambridge University Press, 1 2007. ISBN 978-0-521-86449-7, 978-0-511-26720-8. doi: 10.1017/CBO9780511813917.
- D. J. Thouless, P. W. Anderson, and R. G. Palmer. Solution of 'solvable model of a spin glass'. *The Philosophical Magazine: A Journal of Theoretical Experimental and Applied Physics*, 35 (3):593–601, 1977. doi: 10.1080/14786437708235992.
- Constantino Tsallis. Possible Generalization of Boltzmann-Gibbs Statistics. *J. Statist. Phys.*, 52: 479–487, 1988. doi: 10.1007/BF01016429.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Michael Widrich, Bernhard Schäfl, Milena Pavlović, Hubert Ramsauer, Lukas Gruber, Markus Holzleitner, Johannes Brandstetter, Geir Kjetil Sandve, Victor Greiff, Sepp Hochreiter, et al. Modern hopfield networks and attention for immune repertoire classification. *Advances in neural information processing systems*, 33:18832–18845, 2020.
- Dennis Wu, Jerry Yao-Chieh Hu, Weijian Li, Bo-Yu Chen, and Han Liu. Stanhop: Sparse tandem hopfield model for memory-enhanced time series prediction. *arXiv preprint arXiv:2312.17346*, 2023.
- Dennis Wu, Jerry Yao-Chieh Hu, Teng-Yun Hsiao, and Han Liu. Uniform memory retrieval with larger capacity for modern hopfield models. In *International Conference on Machine Learning*, pages 53471–53514. PMLR, 2024.
- Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tieyan Liu. On layer normalization in the transformer architecture. In *International conference on machine learning*, pages 10524–10533. PMLR, 2020.

Yongyi Yang, zengfeng Huang, and David P Wipf. Transformers from an optimization perspective. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 36958–36971. Curran Associates, Inc., 2022.

Jean Zinn-Justin. Quantum field theory and critical phenomena. *Int. Ser. Monogr. Phys.*, 113: 1–1054, 2002.