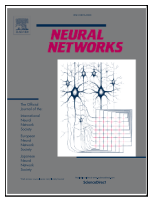
Neural Networks 206 (2027) 109729



Contents lists available at ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet



Full Length Article



Do Hopfield networks dream of stored patterns? A statistical–mechanical theory of dreaming in multidirectional associative memories

Adriano Barra [a,b,c,d], Fabrizio Durante [e], Andrea Ladiana [a,c,*], Michela Marra Solazzo [c,e,f]

[a] Dipartimento di Scienze di Base e Applicate all'Ingegneria, Sapienza Università di Roma, Rome, Italy
[b] Istituto Nazionale di Fisica Nucleare, Sezione di Roma 1, Rome, Italy
[c] Istituto Nazionale d'Alta Matematica, GNFM, Rome, Italy
[d] CNR-Nanotec, Unità di Lecce, Lecce, Italy
[e] Dipartimento di Matematica e Fisica, Università del Salento, Lecce, Italy
[f] Istituto Nazionale di Fisica Nucleare, Sezione di Lecce, Lecce, Italy

ARTICLE INFO

Keywords:
Hopfield networks
Heteroassociative memory
Hebbian dreaming
Associative memory
Replica–symmetric theory
Statistical mechanics
Energy–based models

ABSTRACT

We introduce the Dreaming L-directional Associative Memory (DLAM), a multi-layer Hebbian architecture in which off-line dreaming and supervised heteroassociative coupling coexist within a single cost function, placing our approach within the framework of energy-based models (EBMs). The replica-symmetric free energy, derived via the Guerra interpolation scheme, yields self-consistency equations governing the order parameters across the control-parameter space. The effective local field decomposes into signal, intra-layer dreaming noise, and inter-layer noise. Dreaming improves retrieval by differentially attenuating high-eigenvalue interference modes of the empirical correlation matrix, suppressing inter-pattern crosstalk while preserving the signal. Dreaming and inter-layer coupling prove synergistic, opening retrieval regions unreachable by either mechanism alone, as confirmed by Monte Carlo simulations for $L = 3$. Their interplay is most pronounced on pattern disentanglement: given a mixture state as input, the network splits the constituent patterns one-per-layer, recovering each modality-specific pattern from a common cue that simultaneously blends noisy evidence from all sensory channels. Phase diagrams are planar projections of the hyperspace (α, T, ρ, t)—where α is the storage load, T the fast-noise temperature, ρ the dataset entropy, and t the sleeping time. In the (ρ, t)-plane, the diagrams reveal a data–computation trade-off: off-line consolidation substitutes for additional training data, extending to heteroassociative architectures a phenomenon previously established for autoassociative networks. Enriching the standard Hopfield model with heteroassociativity and dreaming gives rise to EBMs capable of complex tasks beyond classical pattern recognition, contributing to a modern theory of neural information processing.

1. Introduction

Among the most provocative ideas to emerge from the intersection of neuroscience and statistical mechanics is the hypothesis that sleep serves a computational function. Crick and Mitchison (1983) proposed that rapid-eye-movement sleep operates as a form of reverse learning, selectively weakening parasitic synaptic configurations accumulated during waking experience that would otherwise corrupt stored memories. This early proposal is now situated within a broad experimental literature documenting the active role of sleep in memory consolidation, replay, and synaptic reorganisation (Diekelmann & Born, 2010; Klinzing et al., 2019; Rasch & Born, 2013). Translated into the framework of Hopfield networks (Hopfield, 1982), this amounts to an off-line adjustment of the coupling matrix that suppresses inter-pattern crosstalk and sharpens the basins of attraction around genuine stored patterns. The idea is noteworthy because the role it assigns to sleep is at its core *algorithmic*: the sleeping network performs a spectral filtering of its own synaptic matrix, a computation that the waking network, constrained to process incoming data, cannot easily carry out on-line (Agliari et al., 2024b, 2026). This interpretation is broadly consonant with synaptic-homeostasis accounts of sleep (Tononi & Cirelli, 2006), although the present kernel is a mathematical idealisation rather than a biological claim. The analogy between this biological picture and the mathematical operation of regularising a noisy empirical covariance matrix provides both a conceptual bridge and a technical starting point for the present work.

* Corresponding author at: Dipartimento di Scienze di Base e Applicate all'Ingegneria, Sapienza Università di Roma, Rome, Italy.
E-mail address: andrea.ladiana@uniroma1.it (A. Ladiana).

https://doi.org/10.1016/j.neunet.2026.109729
Received 14 May 2026; Received in revised form 1 September 2026; Accepted 5 October 2026
Available online 9 October 2026
0893-6080/© 2026 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

The analytical confirmation of this picture has unfolded over several decades. The storage capacity of the classical Hopfield network, limited to $\alpha_c \simeq 0.138$ by inter-pattern interference (Amit et al., 1985; Hopfield et al., 1983), can be raised substantially when the coupling matrix is subjected to a dreaming-type correction, whether through Hebbian unlearning, pseudo-inverse regularisation, or the spectral kernel introduced in Fachechi et al. (2019) and Alemanno et al. (2023). In the supervised learning-from-examples framework, where the network observes noisy copies of each archetype rather than the archetypes themselves, dreaming carries a dual benefit: it suppresses inter-pattern crosstalk *and* compensates for finite-sample noise, thereby acting as both a capacity booster and a data-saving mechanism (Agliari et al., 2024a,b; Benedetti et al., 2024; Serricchio et al., 2025). These results, however, have been obtained exclusively in single-layer architectures (autoassociative memories). Even there the benefit is not unbounded, and it is worth being precise about how single-layer dreaming breaks down under scarce or noisy data. Dreaming acts on the empirical correlation matrix Σ_l, whose signal eigenvalue stays detached from the noise bulk only while the examples are sufficiently informative relative to the load; in the learning-from-examples picture this is governed by the dataset entropy ρ_l (see Eq. (2)), which aggregates label corruption (small r_l) and small sample size (small M). Once ρ_l exceeds a critical value, the signal direction sinks into the bulk spectrum of Σ_l (a Baik–Ben Arous–Péché-type spectral transition) and no amount of sleep can recover it, because the kernel $\mathbb{A}_l(t_l)$ can only reweight existing eigenmodes, not regenerate a signal that the data no longer resolve (Agliari et al., 2024a,b; Alemanno et al., 2023). A single dreaming layer therefore fails *sharply* once its own data drop below this threshold, which is exactly the regime in which the cross-modal support supplied by a multi-layer architecture becomes indispensable, and which motivates the construction below. When the task is intrinsically multimodal, that is, when a memory comprises coordinated information arriving through several sensory channels, a single layer cannot represent the relational structure between modalities, and dreaming within that layer alone cannot exploit cross-modal regularities.[1] Recent work has shown that an assembly of Hopfield networks can display emergent capabilities that are absent at the level of a single network (Agliari et al., 2025b): in particular, a layered architecture can spontaneously perform *pattern disentanglement*, recovering the individual components of a composite signal, in analogy with identifying the notes composing a chord, a task that is by construction impossible for any single-layer model. This “more is different” principle, in the spirit of P. W. Anderson, motivates the multi-layer architecture developed in the present work.

The natural formalism for multimodal memory is the L–directional associative memory (LAM), illustrated in Fig. 1 (left panel), in which L layers of neurons, each encoding a distinct modality, interact through Hebbian heteroassociative couplings that bind co-occurring patterns across layers (Agliari et al., 2025a; Alessandrelli et al., 2025a,b; Kobayashi et al., 2002). The architecture has a principled biological motivation: the mammalian cortex stores multimodal representations in anatomically distributed areas and retrieves them through associative completion (McClelland et al., 1995; Rissman & Wagner, 2012): a visual cue can evoke an auditory memory, and a partial stimulus can elicit its full multisensory context (Barker & Warburton, 2020). Pattern completion across modalities and pattern separation within a modality are widely regarded as complementary operations subserved by hippocampal–cortical circuits (Rolls, 2013; Yassa & Stark, 2011); the LAM provides a tractable mathematical abstraction of this organisation. To keep this analogy disciplined rather than loose, we anchor each circuit-level operation to a specific term of the DLAM cost function. Cross-modal pattern completion –the CA3-like reactivation of a full stored representation from a minimal or partial cue, supported in the hippocampus by strong recurrent connectivity– is carried by the heteroassociative couplings g_{lm} in $\mathcal{H}_{\mathrm{LAM}}$ (6), which propagate retrieval from informed to uninformed layers; within-modality pattern separation – the dentate-gyrus-like “expansion recoding” that maps overlapping or noisy inputs onto well-separated representations – is carried by the dreaming kernel $\mathbb{A}_l(t_l)$ (3), which decorrelates overlapping representations inside each layer; and the off-line (slow-wave) consolidation budget is the sleeping time t_l, while the dataset entropy ρ_l (2) quantifies the quality of waking experience. The hippocampal–cortical reading is therefore tied to concrete control parameters (g_{lm}, t_l, ρ_l) that appear explicitly in the phase diagrams, not to an informal metaphor. From the statistical-mechanics viewpoint, inter-layer coupling extends the retrieval region relative to independent layers, but the capacity remains limited by inter-pattern interference, and the problem becomes especially acute when the training data is noisy or scarce.

These two lines of research, single-layer dreaming and multi-layer heteroassociation, have developed largely in parallel, yet the regime where their combination matters most is precisely the one that neither addresses on its own: when the training data is noisy or scarce, when the cues are ambiguous, and when successful recall requires both intra-modal noise suppression and cross-modal support. A layer whose local data is too poor for standalone retrieval might be rescued by information relayed from better-informed partner layers, but only if its own intra-layer landscape is sufficiently clean for the heteroassociative signal to take effect. The interplay between the two mechanisms is non-trivial and cannot be anticipated by studying them in isolation.

In this paper we construct the *Dreaming L–directional Associative Memory* (DLAM), a single cost function (Hamiltonian) that integrates intra-layer spectral dreaming with inter-layer Hebbian coupling, and we develop a closed replica-symmetric description of its thermodynamics. The analytical framework rests on a Hubbard–Stratonovich linearisation of both quadratic sectors, followed by a Guerra interpolation that yields the quenched free energy under the replica-symmetric (RS) ansatz. The central analytical result is a closed system of self-consistency equations governing the fixed points of the order parameters (listed in Table 1) whose effective local field decomposes into a signal term, enhanced by dreaming; a dreaming-noise term, suppressed by dreaming; and an inter-layer noise term. This decomposition provides a clear field-theoretic mechanism for the dreaming benefit and yields the phase diagram as a function of storage load, dataset entropy, temperature, and inter-layer coupling, revealing synergistic dreaming–cooperation regions not reached by either mechanism alone in the explored regime. Monte Carlo simulations on three-layer networks confirm the analytical predictions across the explored parameter space. The finding with the most immediate significance concerns the task of *disentanglement*: presenting all layers simultaneously with a common cue that blends the noisy modal evidence for a single target archetype across all sensory channels, and requiring every layer to converge to its own modality-specific version of that archetype. Without dreaming the cross-modal mixture cue excites inter-pattern correlations that trap each layer in a spurious state; with moderate dreaming the spectral filtering reorganises the landscape and joint disentanglement success rises substantially. In the (ρ, t)-plane, the numerical phase diagrams further establish a data–computation trade-off: off-line consolidation substitutes for additional training examples, extending to the multi-layer setting results recently obtained in single-layer networks (Agliari et al., 2024a,b).

The remainder of the paper is organised as follows. Section 2 defines the DLAM architecture, introduces the dreaming kernel and the inter-layer coupling, and formulates the disentanglement task that serves as

[1] The obstruction is structural rather than a matter of degree. A unipartite cost function acting on a single layer of N spins contains only intra-modal couplings J_{ij}; once the modalities are concatenated into one layer, no term of the energy can carry a cross-channel correlation $\mathbb{E}[s_{i,l}\, s_{j,m}]$ with $l \neq m$, so the corresponding Gibbs measure factorises over the channel marginals and assigns equal weight to all joint configurations compatible with them. Cross-modal regularities are thus invisible to any single-layer model and can be encoded only by an explicitly relational coupling tensor J_{ij}^{lm} that binds distinct layers, which is exactly the ingredient the multi-layer construction supplies (Agliari et al., 2025b).

the central test of the theory. Section 3 develops the replica-symmetric free energy, states the self-consistency equations, and analyses the phase structure. Section 4 presents Monte Carlo results, beginning with disentanglement and data saving, and proceeding through the full phase structure, finite-size scaling, design rules, dynamics, and robustness to heterogeneous data quality. Section 6 discusses connections to related work, implications, and limitations. Detailed derivations are collected in the Appendix.

2. Model

2.1. Architecture

The DLAM comprises L interacting layers, each populated by N binary neurons $s_{i,l} \in \{-1,+1\}$, $i = 1,\dots,N$, $l = 1,\dots,L$. Each layer represents a distinct sensory modality or data channel; the multi-layer organisation is what allows the network to bind information across modalities, in contrast with single-layer architectures that treat each modality independently. The network stores K archetypes, each of which is *multimodal* in the sense that it consists of an L-tuple of independently drawn Rademacher vectors, one per layer: $\xi_l^\mu \in \{-1,+1\}^N$ for $\mu = 1,\dots,K$ and $l = 1,\dots,L$. Within each layer, the component ξ_l^μ is a standard Hopfield pattern; the qualifier *multimodal* simply records that a single index μ simultaneously addresses one Rademacher pattern in every layer. All archetypes are quenched and hidden from the network.

It is useful to spell out the operational anatomy of this construction. Each of the L layers is a fully connected Hopfield module of N neurons that models one sensory channel or cortical area.[2] *Within* a layer, neurons interact through the (dreaming-regularised) Hebbian couplings of Section 2.2, which implement intra-modal storage and noise suppression. *Across* layers, the relational dependencies between channels are carried by the heteroassociative couplings of Section 2.3: since a stored archetype is an L-tuple of co-occurring patterns sharing the index μ, the inter-layer term binds these components, so that driving one layer toward its component of archetype μ biases the partner layers toward their own components of the *same* μ. The multipartite coupling structure made explicit in Fig. 1 is precisely what encodes the cross-channel regularities that a single module cannot represent.

This relational design is best read as the multipartite completion of the classical heteroassociative hierarchy. Kosko's bidirectional associative memory (BAM) couples *two* layers through a single generalised Hebbian matrix and can bind only ordered pairs of patterns; the three-directional associative memory (TAM) raises this to a tripartite graph in which every *pair* of three layers is coupled, enabling the storage and disentanglement of pattern triplets (Agliari et al., 2025a; Alessandrelli et al., 2025a,b; Kobayashi et al., 2002). The L-directional memory is the natural limit of this construction: the L layers form the vertices of a complete graph whose $\binom{L}{2}$ edges each carry one symmetric heteroassociative coupling g_{lm} (6), so that every cortical area is relationally tied to every other through the shared archetype index μ. Cross-modal completion is thus not delegated to a separate read-out stage but is wired into the energy itself, edge by edge, which is what allows a cue presented to one area to complete the co-occurring representations stored in the remaining $L-1$ areas. The biological reading is anchored term-by-term in the Introduction (CA3-like pattern completion $\leftrightarrow$ the couplings g_{lm}; dentate-gyrus-like pattern separation $\leftrightarrow$ the dreaming kernel $\mathbb{A}_l(t_l)$), so that the architecture mirrors the distributed, associatively completed organisation of multimodal cortical memory rather than merely evoking it.

The network learns from data, not from the archetypes directly: it never experiences the true vectors ξ_l^μ. For each archetype μ and layer l, it observes M independent noisy copies generated by flipping each bit with probability $(1-r_l)/2$:

$$\eta_{i,l}^{\mu,k} = \chi_{i,l}^{\mu,k}\,\xi_{i,l}^{\mu}, \qquad \mathbb{E}[\chi_{i,l}^{\mu,k}] = r_l \in (0,1], \tag{1}$$

where r_l controls data fidelity. The empirical archetype $\hat{\xi}_{i,l}^{\mu} := M^{-1}\sum_{k=1}^{M}\eta_{i,l}^{\mu,k}$ is the network's best estimate of the true pattern, and its fluctuations at finite sample size are governed by the dataset entropy (Alemanno et al., 2023)

$$\rho_l := \frac{1-r_l^2}{M\,r_l^2}. \tag{2}$$

The control and order parameters used throughout the analysis are collected in Table 1. When $\rho_l = 0$ the empirical archetypes coincide with the true ones; as ρ_l grows, the dataset becomes noisier, either because the observations are corrupted (r_l small) or because fewer examples are available (M small). The parameter ρ_l thus conflates two distinct sources of degradation into a single effective noise measure, and the data–computation trade-off explored later amounts to the observation that dreaming can compensate for large ρ_l. We also define $\Gamma_l := r_l^2(1+\rho_l)$ and $\Gamma_{lm} := \sqrt{\Gamma_l\Gamma_m}$, which serve as normalisation constants throughout.

Retrieval quality is measured by the Mattis magnetisation $m_l^\mu := N^{-1}\sum_i \xi_{i,l}^\mu s_{i,l}$, which quantifies the alignment of the network state with the true archetype. We also define $\hat{m}_l^\mu := N^{-1}\sum_i \hat{\xi}_{i,l}^\mu s_{i,l}$ (alignment with the empirical archetype) and the rescaled variant $m_{\eta,l} := \hat{m}_l^1/[r_l(1+\rho_l)]$, which factors out the trivial dependence on dataset quality.[3]

2.2. Dreaming kernel

Throughout this work "dreaming" denotes an *off-line modification of the synaptic couplings*, and not a relaxation of the neural state at frozen weights (free association). Concretely, it is the deterministic map $\mathbb{I} \mapsto \mathbb{A}_l(t_l)$ applied to the intra-layer coupling matrix, Eq. (3), controlled by a single scalar sleeping time $t_l \geq 0$. It is executed once, before retrieval; it revisits no stored examples (there is no replay buffer) and it does not alter the dataset. In this sense dreaming is a form of Hebbian consolidation or synaptic rescaling, after which the on-line retrieval dynamics, the single-spin-flip Glauber updates of Section 4, run on the already-consolidated couplings $\mathbb{A}_l(t_l)$ at fixed t_l.

The dreaming mechanism operates within each layer through the empirical correlation matrix $(\Sigma_l)_{\mu\nu} := (N\Gamma_l)^{-1}\sum_i \hat{\xi}_{i,l}^\mu \hat{\xi}_{i,l}^\nu$. As the sleep time t_l grows, the effective coupling matrix

$$\mathbb{A}_l(t_l) := (1+t_l)\big(\mathbb{I} + t_l\,\Sigma_l\big)^{-1} \tag{3}$$

progressively suppresses the eigenmodes of Σ_l responsible for inter-pattern crosstalk (Agliari et al., 2019; Fachechi et al., 2019). The spectral interpretation is immediate: in the eigenbasis of Σ_l, an eigenvalue λ of Σ_l is rescaled to $(1+t_l)/(1+t_l\lambda)$, which is a strictly decreasing function of λ. Consequently, eigenmodes with larger λ are suppressed more strongly than those with smaller λ. As $t_l \to +\infty$, eigenvalue λ saturates at $1/\lambda$, so the weight assigned to each mode scales inversely with its original magnitude. Crucially, the modes that dominate the inter-pattern crosstalk (the eigenvectors of Σ_l carrying the largest eigenvalues) are differentially attenuated relative to the signal direction, which contributes

[2] For analytical transparency we develop the theory for a *symmetric* architecture, in which all layers share a common size ($N_l \equiv N$) and a common neuron-index space. For heteroassociative memories, the maximal storage load $\gamma_{\max}$ is itself attained at the symmetric point $N_1 = \cdots = N_L$, and degrades monotonically as the layer sizes are made unequal (Alessandrelli et al., 2025a,b). The symmetric, common-index setting is therefore the natural *and* optimal operating regime, which is why we adopt it throughout. Relaxing it to asymmetric neuron counts $\{N_l\}$ is conceptually straightforward: it amounts to propagating layer-dependent normalisations N_l through the same steps, at the price of heavier bookkeeping and the loss of the single-layer reduction that we use to display the results.

[3] The same macroscopic order parameters are used throughout for both the standard retrieval problem and the disentanglement task. What changes between the two settings is the cueing/initialisation protocol, not the order-parameter structure of the theory.

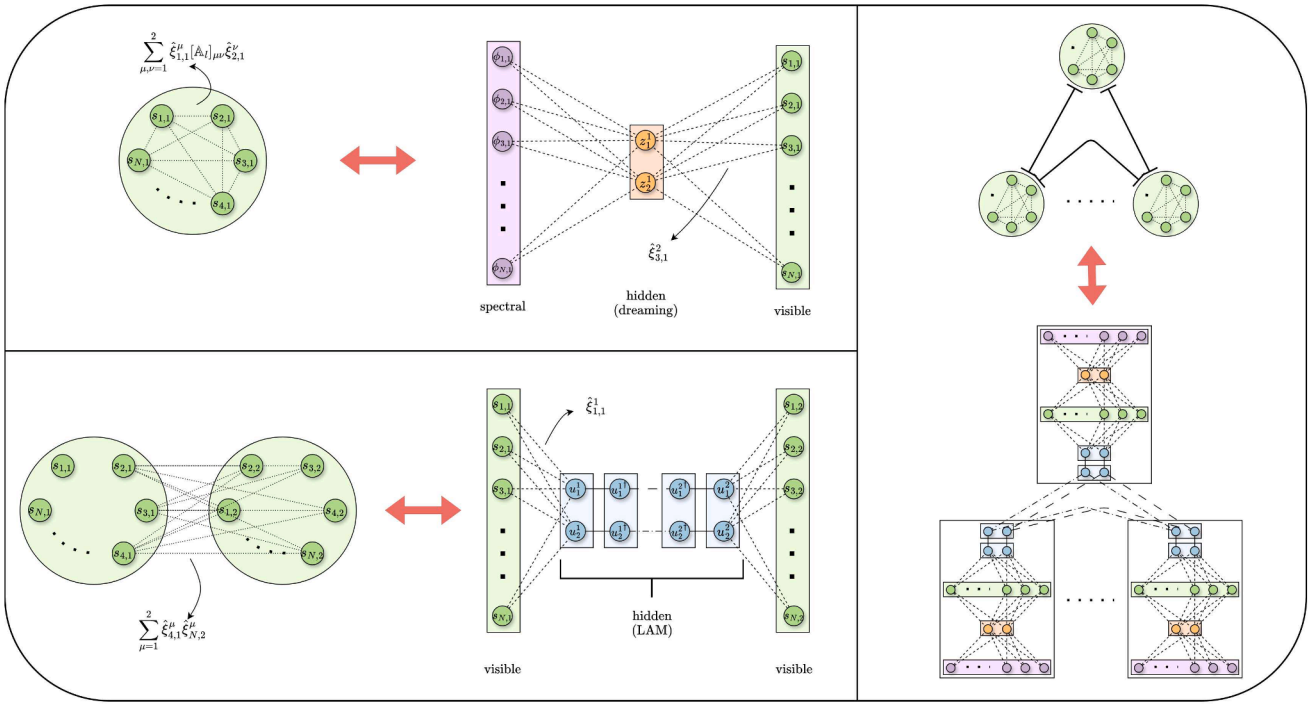


Fig. 1. Architectural equivalence of the Dreaming L-directional Associative Memory (DLAM). Top Left: Integral representation of a simple single-layer dreaming model. Through Hubbard–Stratonovich linearisation (Barra et al., 2012), the intra-layer dreaming interactions among visible neurons ($s_{i,l}$, green nodes) are decoupled into a Restricted Boltzmann Machine (RBM)-like structure mediated by site-indexed ($\phi_{i,l}$, purple nodes) and pattern-indexed (z^l_μ, orange nodes) Gaussian fields. **Bottom Left:** Integral representation of a LAM model with $L = 2$ layers, where inter-layer heteroassociative bindings are mediated by complex conjugate auxiliary fields ($u^l_\mu, u^{l\dagger}_\mu$, blue nodes). **Right:** A stylised illustration of the complete, fully coupled DLAM architecture integrating both mechanisms across L interacting layers. This unified bipartite representation forms the analytical foundation for the Guerra interpolation and the replica-symmetric theory.

across all eigenmodes but is concentrated in the condensed spin-space eigenvector. The net effect is a flattening of the effective coupling spectrum that suppresses spurious interference without destroying the signal: the kernel acts as a *spectral regulariser* that reduces the variance of the effective couplings rather than zeroing individual eigenmodes.

The intra-layer cost function constructed with this kernel reads

$$\mathcal{H}_{\text{Dream}}(\boldsymbol{s}) = -\sum_{l=1}^{L} \frac{1}{N\Gamma_l} \sum_{i,j} \sum_{\mu,\nu} \hat{\xi}^{\mu}_{i,l} \big[\mathbb{A}_l(t_l)\big]_{\mu\nu} \hat{\xi}^{\nu}_{j,l}\, s_{i,l}\, s_{j,l}. \tag{4}$$

The regularising action of the dreaming kernel admits a direct geometric interpretation in the macroscopic space of empirical overlaps. Let $\hat{\boldsymbol{m}}_l(\boldsymbol{s}) \in \mathbb{R}^K$ denote the vector with components $\hat{m}^{\mu}_l = N^{-1} \sum_i \hat{\xi}^{\mu}_{i,l} s_{i,l}$. For $t_l = 0$, the intra-layer energy reduces to the negative squared Euclidean norm, $\mathcal{H} \propto -\hat{\boldsymbol{m}}_l^{\top} \hat{\boldsymbol{m}}_l$. This isotropic metric preserves the correlated noise embedded in the empirical data.

Conversely, the kernel $\mathbb{A}_l(t_l) = (1+t_l)(\mathbb{I} + t_l \Sigma_l)^{-1}$ introduces a non-trivial metric tensor. Since the empirical covariance Σ_l is positive semi-definite, the operator $\mathbb{I} + t_l \Sigma_l$ is strictly positive definite for any $t_l > 0$, ensuring that its inverse is a well-defined precision matrix. Consequently, the intra-layer cost function $\mathcal{H}_{\text{Dream}} \propto -\hat{\boldsymbol{m}}_l^{\top} \mathbb{A}_l(t_l) \hat{\boldsymbol{m}}_l$ computes a Tikhonov-regularised Mahalanobis squared norm of the state vector. In the asymptotic limit $t_l \to +\infty$, the tensor $\mathbb{A}_l(t_l)$ behaves as Σ_l^{-1} on the non-null eigenspace of the patterns. The dreaming mechanism thus performs a spatial whitening of the overlap space, replacing the Euclidean geometry with a Mahalanobis metric that natively decorrelates the macroscopic state variables.

Remark 1. To exhibit the transition between the two metrics, and its invariance, diagonalise the symmetric positive semi-definite covariance as $\Sigma_l = O_l \Lambda_l O_l^{\top}$, with $O_l \in \mathrm{O}(K)$ and $\Lambda_l = \mathrm{diag}(\lambda_1, \dots, \lambda_K)$, $\lambda_a \geq 0$. By the functional calculus the kernel is diagonal in the same eigenframe, $\mathbb{A}_l(t_l) = O_l\, \mathrm{diag}\big(\frac{1+t_l}{1+t_l\lambda_a}\big)\, O_l^{\top}$, so that, writing $\tilde{\boldsymbol{m}}_l = O_l^{\top} \hat{\boldsymbol{m}}_l$ for the overlaps expressed in the eigenbasis,

$$\hat{\boldsymbol{m}}_l^{\top}\, \mathbb{A}_l(t_l)\, \hat{\boldsymbol{m}}_l \;=\; \sum_{a=1}^{K} \frac{1+t_l}{1+t_l\lambda_a}\, \tilde{m}^2_{l,a}. \tag{5}$$

At $t_l = 0$ every weight equals unity and (5) collapses to the Euclidean norm $\|\hat{\boldsymbol{m}}_l\|^2$; for $t_l > 0$ each mode is reweighted by the strictly positive factor $(1+t_l)/(1+t_l\lambda_a)$, i.e. by a Tikhonov-regularised inverse-covariance (Mahalanobis) weight, which tends to λ_a^{-1} on the pattern eigenspace as $t_l \to +\infty$. Crucially, the form (5) is invariant under any orthogonal relabelling of the pattern frame $\hat{\boldsymbol{m}}_l \mapsto R\,\hat{\boldsymbol{m}}_l$, $\Sigma_l \mapsto R\,\Sigma_l\,R^{\top}$ ($R \in \mathrm{O}(K)$), because $\mathbb{A}_l$ transforms by the same congruence $\mathbb{A}_l \mapsto R\,\mathbb{A}_l\,R^{\top}$ and the scalar $\hat{\boldsymbol{m}}_l^{\top}\mathbb{A}_l\hat{\boldsymbol{m}}_l$ is left unchanged. The Mahalanobis geometry is therefore intrinsic to the empirical covariance and independent of the arbitrary ordering of the stored patterns; this is the metric-invariance that makes the replacement of the Euclidean norm by the dreaming-induced norm well posed.

2.3. Inter-layer coupling

The dreaming kernel operates within each layer independently; the second ingredient of the DLAM is a cross-modal binding that couples layers through their shared archetype structure. When archetype μ is stored simultaneously in layers l and m, the heteroassociative term rewards spin configurations in which both layers align with their respective empirical estimates of μ. This is the mathematical counterpart of the biological observation that multimodal memories are stored in distributed cortical areas and retrieved through associative completion.[4]

[4] We stress that in Xiang et al. (2021) the authors use multiplex biological networks to capture impulsive dynamics, offering strong supporting evidence for the growing necessity of multilayer network paradigms in modelling complex, interconnected biological data.

Formally, the inter-layer coupling reads

$$\mathcal{H}_{\mathrm{LAM}}(\boldsymbol{s}) = -\sum_{l<m} \frac{g_{lm}}{N\Gamma_{lm}} \sum_{\mu=1}^{K} \sum_{i,j} \hat{\xi}^{\mu}_{i,l}\, \hat{\xi}^{\mu}_{j,m}\, s_{i,l}\, s_{j,m}. \tag{6}$$

The coupling is *symmetric and bidirectional*: each unordered pair $\{l, m\}$ contributes once (the sum runs over $l < m$) through a single scalar strength $g_{lm} = g_{ml} \geq 0$, so layer l acts on layer m exactly as m acts on l and no direction of cross-modal flow is privileged. This symmetry is what makes $\mathcal{H}_{\mathrm{DLAM}}$ a bona fide (Lyapunov) energy function: the single-spin-flip dynamics satisfy detailed balance and relax to the Boltzmann–Gibbs measure, just as in the single-layer Hopfield model. Consistency across layers is then enforced energetically rather than by an external constraint: as shown in Section 2.5, Eq. (11), the LAM term is equivalent, up to state-independent constants, to the contrastive penalty $\frac{1}{2}(\hat{m}^{\mu}_{l} - \hat{m}^{\mu}_{m})^2$, so configurations in which the coupled layers agree on the retrieved archetype are energetically favoured and disagreement is penalised.

2.4. Full cost function

The complete DLAM cost function (Hamiltonian) is the sum

$$\mathcal{H}_{\mathrm{DLAM}}(\boldsymbol{s}) := \mathcal{H}_{\mathrm{Dream}}(\boldsymbol{s}) + \mathcal{H}_{\mathrm{LAM}}(\boldsymbol{s}). \tag{7}$$

The retrieval behaviour is governed by four classes of control parameters: the storage load $\alpha := K/N$, the inverse temperature $\beta = \frac{1}{T}$, the dataset entropies $\{\rho_l\}$, and the sleep times $\{t_l\}$.[5]

The cost function (7) contains several known models as limiting cases: at L=1 it reduces to a single-layer dreaming network; at t_l=0 for all l it yields a multidirectional associative memory built from noisy empirical patterns; with the inter-layer couplings removed, it factorises into L independent dreaming problems (Agliari et al., 2019; Fachechi et al., 2019). The key architectural point is that, unlike a two-stage pipeline in which dreaming and heteroassociation are applied sequentially, both mechanisms live inside the same energy landscape and interact through the shared neural dynamics.

2.5. Cost function vs. loss function: a machine learning perspective

It is instructive to recast the DLAM cost function, traditionally viewed as an energy function in statistical mechanics, into a Loss function typical of the Machine Learning framework. This establishes a rigorous bridge between the physical dynamics of retrieval and the minimisation of a Mean Squared Error (MSE) objective.

Let us rewrite the empirical Mattis magnetisation as $\hat{m}^{\mu}_{l}(\boldsymbol{s}) := \frac{1}{N}\sum_i \hat{\xi}^{\mu}_{i,l} s_{i,l}$. The total energy $\mathcal{H}_{\mathrm{DLAM}}$ can be expressed entirely in terms of these macroscopic overlaps:

$$\mathcal{H}_{\mathrm{DLAM}} = -N \sum_{l=1}^{L} \frac{1}{\Gamma_l} \sum_{\mu,\nu} \hat{m}^{\mu}_{l} [\mathbb{A}_l(t_l)]_{\mu\nu} \hat{m}^{\nu}_{l} - N \sum_{l<m} \frac{g_{lm}}{\Gamma_{lm}} \sum_{\mu=1}^{K} \hat{m}^{\mu}_{l} \hat{m}^{\mu}_{m}. \tag{8}$$

To frame the dreaming sector as a reconstruction task, we introduce the rotated empirical archetypes $\tilde{\xi}^{\mu}_{l} := \sum_{\nu} \left[\mathbb{A}_l(t_l)^{1/2}\right]_{\mu\nu} \hat{\xi}^{\nu}_{l}$, with corresponding rotated overlaps $\tilde{m}^{\mu}_{l}(\boldsymbol{s}) = \frac{1}{N}\sum_i \tilde{\xi}^{\mu}_{i,l} s_{i,l}$. This rotation exactly absorbs the dreaming kernel, allowing us to define a natural L_2 Loss function measuring the distance between the network state and the target (or its reflection):

$$\tilde{\mathcal{L}}^{\mu}_{l,\pm}(\boldsymbol{s}) := \frac{1}{2N}\left\|\tilde{\xi}^{\mu}_{l} \pm \boldsymbol{s}_l\right\|^2 = \frac{1}{2N}\left\|\tilde{\xi}^{\mu}_{l}\right\|^2 + \frac{1}{2} \pm \tilde{m}^{\mu}_{l}(\boldsymbol{s}). \tag{9}$$

The transformation $\tilde{\xi}^{\mu}_{l} = \mathbb{A}_l(t_l)^{1/2}\hat{\xi}^{\mu}_{l}$ is rigorously well-defined: because the empirical correlation matrix Σ_l is a Gram matrix, it is symmetric positive semi-definite. Consequently, the operator $\mathbb{I} + t_l\Sigma_l$ has eigenvalues bounded below by 1, ensuring that the dreaming kernel $\mathbb{A}_l(t_l)$ is strictly positive definite for any $t_l \geq 0$. Thus, its principal square root exists, is unique, and maintains the geometry of the pattern space.

Simple algebra reveals that the intra-layer dreaming energy is exactly proportional to the product of these symmetric losses:

$$-\sum_{\mu,\nu} \hat{m}^{\mu}_{l} [\mathbb{A}_l(t_l)]_{\mu\nu} \hat{m}^{\nu}_{l} = \sum_{\mu=1}^{K} \left(\tilde{\mathcal{L}}^{\mu}_{l,+}(\boldsymbol{s})\, \tilde{\mathcal{L}}^{\mu}_{l,-}(\boldsymbol{s}) - C^{\mu}_{l}(t_l)^2\right), \tag{10}$$

where $C^{\mu}_{l}(t_l) = \frac{1}{2N}\|\tilde{\xi}^{\mu}_{l}\|^2 + \frac{1}{2}$ is a state-independent structural constant.

Simultaneously, the heteroassociative LAM sector acts as a multi-view *alignment loss*. By rewriting the coupling term as

$$\hat{m}^{\mu}_{l}\hat{m}^{\mu}_{m} = \frac{1}{2}(\hat{m}^{\mu}_{l})^2 + \frac{1}{2}(\hat{m}^{\mu}_{m})^2 - \frac{1}{2}(\hat{m}^{\mu}_{l} - \hat{m}^{\mu}_{m})^2, \tag{11}$$

the LAM energy explicitly minimises the cross-modal discrepancy:

$$\mathcal{L}^{\mu,l,m}_{\mathrm{cross}}(\boldsymbol{s}) := \frac{1}{2}\left(\hat{m}^{\mu}_{l}(\boldsymbol{s}) - \hat{m}^{\mu}_{m}(\boldsymbol{s})\right)^2. \tag{12}$$

The decomposition of the heteroassociative LAM energy establishes a formal link with modern self-supervised learning. The identity in Eq. (11) illustrates that inter-layer dynamics extend beyond simple cross-modal alignment (i.e., the minimisation of the discrepancy $\frac{1}{2}(\hat{m}^{\mu}_{l} - \hat{m}^{\mu}_{m})^2$). Specifically, the positive auto-overlap terms $\frac{1}{2}(\hat{m}^{\mu}_{l})^2$ function as variance-maximising regularisers. In the context of self-supervised frameworks, these terms mitigate representation collapse, preventing the network from converging to uninformative trivial solutions in the pursuit of cross-layer synchronisation (Bardes et al., 2022).

Thus, the DLAM thermal relaxation inherently performs a dual optimisation: the dreaming kernel minimises the intra-layer MSE (effectively whitening the empirical noise), while the LAM couplings minimise the inter-layer contrastive loss, forcing distinct sensory channels to converge onto a unified, coherent multisensory representation.

2.6. Disentanglement task

The preceding definitions equip the DLAM with the machinery for both intra-layer noise suppression and cross-modal binding. The most stringent test of whether these mechanisms work together is not standard pattern retrieval, which probes only the depth of a single attractor, but *disentanglement*: the simultaneous recovery of each layer's modality-specific representation from a single common cue that blends noisy evidence from all sensory channels at once.

Concretely, a target archetype $\mu^\star$ is drawn at random. For each layer $l \in \{1, 2, 3\}$,[6] a single noisy training example $\eta^{\mu^\star, k_l}_{i,l}$ is drawn independently (example index $k_l \in \{1, \ldots, M\}$ drawn uniformly at random) from the layer-l dataset for archetype $\mu^\star$ (using the notation of Section 2.1). A weighted superposition of these three modal examples is formed and used as the initial state of *all* layers simultaneously:

$$s^{\mathrm{init}}_{i,l} = \mathrm{sign}\left(\alpha_1\, \eta^{\mu^\star,k_1}_{i,1} + \alpha_2\, \eta^{\mu^\star,k_2}_{i,2} + \alpha_3\, \eta^{\mu^\star,k_3}_{i,3}\right), \qquad l = 1, 2, 3, \tag{13}$$

with $(\alpha_1, \alpha_2, \alpha_3) \geq 0$ and $\alpha_1 + \alpha_2 + \alpha_3 = 1$. Ties (i.e. $\alpha_1\, \eta^{\mu^\star,k_1}_{i,1} + \alpha_2\, \eta^{\mu^\star,k_2}_{i,2} + \alpha_3\, \eta^{\mu^\star,k_3}_{i,3} = 0$) are broken by independently assigning +1 or −1 with equal probability $\frac{1}{2}$; this event has positive probability on simplex edges and

[5] For L-layer networks with symmetric architecture it is natural to adopt the following parametrisation of the coupling tensors. One assigns a common *auto-coupling* strength $g_{\mathrm{auto}} := g_{ll}$ to all diagonal entries (l=m) and a common *hetero-coupling* strength $g_{\mathrm{hetero}} := g_{lm}$ (l≠m) to all off-diagonal entries of the coupling matrix $\{g_{lm}\}$. Concretely, this amounts to replacing $1 \mapsto g_{\mathrm{auto}}$ in $\mathcal{H}_{\mathrm{Dream}}$ and $g_{lm} \mapsto g_{\mathrm{hetero}}$ uniformly in $\mathcal{H}_{\mathrm{LAM}}$. All numerical experiments in Section 4 use this parametrisation, and all caption entries labelled g_{auto} and g_{hetero} refer to these two scalars. This parametrisation is a modelling choice: the general theory in Sections 3–4 holds for arbitrary $\{g_{lm}\}$.

[6] Throughout the numerical experiments we instantiate the model with $L = 3$ layers, but the disentanglement framework generalises naturally to an arbitrary number of layers L.

must be handled consistently to keep the initial state in the binary space $\{-1,+1\}^N$. The construction implicitly assumes a *shared neuron-index space*: the ith component of each modal sample $\eta_{i,l}^{\mu^\star}$ refers to the same index position across layers, so that the coordinatewise sum is geometrically well-defined. This common-basis convention is stated here explicitly: it is not a structural requirement of the DLAM (the layers may in general have independent index sets), but it is the natural choice for the present benchmark. The weight vector $(\alpha_1,\alpha_2,\alpha_3)$ lives on the 2-simplex: its vertices, e.g. $(1,0,0)$, correspond to pure single-view cues (cue derived entirely from layer 1's training data), the edges to two-view mixtures, and the barycentre $(\frac{1}{3},\frac{1}{3},\frac{1}{3})$ to equal weighting of all three modal views. To systematically test the network's resilience, the initial state can be further corrupted by an explicit cue noise $\varepsilon \in [0,0.5]$, which represents the fraction of randomly flipped spins in $s_{i,l}^{\text{init}}$ before the onset of the retrieval dynamics.

A trial is successful only if the tail-averaged magnetisation $m_l = N^{-1}\langle \boldsymbol{s}_l \cdot \boldsymbol{\xi}_l^{\mu^\star}\rangle$ exceeds a threshold m_{th} on *all* layers simultaneously. This is a genuine disentanglement requirement: all three layers start from the exact same mixed vector, one that belongs to none of the individual layers' pattern spaces, and each must recover its own modality-specific target archetype.

To make the claim precise, "disentanglement success" is the event $\{\, m_l \geq m_{\text{th}} \text{ for all } l = 1,\dots,L \,\}$ evaluated on the post-relaxation state, and the quantity reported throughout is its empirical probability P_{success} over independent realisations of the dataset, the cue indices $\{k_l\}$, and the stochastic dynamics. We do not claim a deterministic guarantee: the result is the observed behaviour that, above a finite sleeping time t and at the parameters stated in each caption, $P_{\text{success}} \to 1$ essentially uniformly over the cue simplex. The specific value $m_{\text{th}} = 0.80$ is a conventional high-overlap criterion, and the classification is insensitive to its precise choice: the tail-averaged magnetisation is sharply bimodal –successful trials concentrate near $m \simeq 0.9$–1 and failures near the mixture/spin-glass band $m \simeq 0$– so any threshold in the wide gap $m_{\text{th}} \in [0.7,0.9]$ produces the same success/failure partition.

The difficulty varies systematically with $(\alpha_1,\alpha_2,\alpha_3)$. Near the barycentre every layer receives a diluted but nonzero signal from its own modal view, so that inter-layer cooperation can bootstrap all three layers simultaneously. At a simplex vertex, by contrast, only one layer receives useful signal while the other two start from essentially zero overlap with their targets; joint retrieval must then rely entirely on the heteroassociative coupling to rescue the uninformed layers, a far harder task when the intra-layer basins are contaminated by noise. The central prediction of the DLAM theory is that, by suppressing the inter-pattern correlations encoded in each Σ_l, the dreaming kernel $\mathbb{A}_l(t_l)$ cleans the basins sufficiently for the heteroassociative signal to propagate and rescue even uninformed layers, expanding the region of joint success from the barycentre toward the vertices. This prediction is confirmed quantitatively in Section 4.

It is worth making explicit why this task is, by construction, the proper arena for a *heteroassociative* architecture, and why a baseline of L independent (autoassociative) dreaming networks would not be an informative comparison. The cue (13) is a superposition of modal patterns, that is, exactly a spurious mixture state of the kind that constitutes the structural Achilles' heel of single-layer Hopfield memories. A lone autoassociative layer fed such a mixture either locks onto it as a metastable attractor or, once dreaming has destabilised that mixture, falls back toward its own single target; in neither case does it possess any mechanism to assign the remaining modal components to other representations, for it has none. Disentanglement therefore cannot even be *posed* for a collection of decoupled layers: splitting one common mixed cue into L distinct, modality-specific patterns (one per layer) intrinsically requires the inter-layer couplings g_{lm} to route each component to its layer. Dreaming alone removes spurious mixtures, but only its combination with heteroassociation (dreaming to clean the intra-layer basins, the LAM couplings to redistribute the components) realises the task. This is precisely why we benchmark dreaming+LAM against LAM-only in Fig. 2, rather than against independent dreaming networks, whose joint success rate is identically zero except at the trivial single-view vertices.

3. Statistical–mechanical analysis

This section develops the replica-symmetric theory for the DLAM. Rather than reproducing every algebraic step, we present the derivation strategy, state the main results, and focus on the physical interpretation of the self-consistency equations and the effective-field decomposition. Full details (Hubbard–Stratonovich decoupling, Guerra interpolation, one-body calculations, and explicit LAM scalar forms for L=3) are in the Appendix.

A word on the role of $L = 3$ is in order, since it is a presentational choice rather than a limitation of the theory. The derivation in this section and in the Appendix is carried out for *generic* L up to the LAM Gaussian sector, whose closure requires the determinant of an $L \times L$ coupling matrix; this determinant exists and is computable for every L, but its expansion produces a number of distinct terms that grows combinatorially, so a *compact* closed-form free energy can be displayed cleanly only for small L, and we present it for the representative case $L = 3$. Beyond its didactic convenience, $L = 3$ is moreover a case of genuine practical interest: a large class of real multimodal signals decomposes naturally into three components –the three RGB channels of a colour image, the three notes of an elementary chord– so the tripartite memory is itself a natural operating point and not merely a toy reduction. This restriction is, in any case, purely a matter of readability: we provide, as supplementary material together with the public code, a short symbolic routine that generates the LAM determinant and the associated self-consistency closure for arbitrary L. Moreover, because the Monte Carlo simulations and the self-consistency equations are shown to agree throughout the explored parameter space (Section 4), the two can be used interchangeably; for $L > 3$, where the analytic free energy is unwieldy to typeset, the validated simulator provides the same phenomenology at no loss of rigour. We therefore frame $L \leq 3$ as a transparent illustrative setting, not as the boundary of applicability of the model.

3.1. Hubbard–Stratonovich decoupling and partition function

The goal is to compute the intensive quenched free energy

$$f(\beta) := -\lim_{N\to\infty} \frac{1}{\beta N}\, \mathbb{E}\big[\log \mathcal{Z}_{N,L}(\beta \mid D)\big] \tag{14}$$

where $\mathcal{Z}_{N,L}$ is the partition function and $\mathbb{E}$ denotes expectation over the quenched dataset D. Both the dreaming and the LAM sectors of (7) are quadratic in the spins; both are linearised by independent Hubbard–Stratonovich (HS) transformations. Set

$$J_{i,l}^{\mu} := \frac{\hat{\xi}_{i,l}^{\mu}}{\sqrt{\Gamma_l}}, \qquad p_l^{\mu}(\boldsymbol{s}) := \frac{1}{\sqrt{N}} \sum_{i=1}^{N} J_{i,l}^{\mu}\, s_{i,l}, \tag{15}$$

so that $\hat{\xi}_{i,l}^{\mu}\hat{\xi}_{j,m}^{\mu}/\Gamma_{lm} = J_{i,l}^{\mu} J_{j,m}^{\mu}$ and the LAM energy takes the form $-\beta \sum_{l<m} g_{lm} \sum_{\mu} p_l^{\mu} p_m^{\mu}$. Under the replica-symmetric condensation ansatz, mode $\mu = 1$ is macroscopic (signal) and modes $\mu \geq 2$ are noise; the signal–noise split reads

$$\sum_{\mu=1}^{K} p_l^{\mu} p_m^{\mu} = p_l^1 p_m^1 + \sum_{\mu\geq 2} p_l^{\mu} p_m^{\mu}. \tag{16}$$

The restriction to a *single* condensed mode corresponds to the investigation of *pure* retrieval states. Among the K overlaps p_l^{μ}, promoting one to a macroscopic $\mathcal{O}(1)$ value describes the successful retrieval of a target archetype, whereas the remaining $K-1$ overlaps stay of order $\mathcal{O}(N^{-1/2})$ and act collectively as a Gaussian noisy channel (Agliari et al., 2017). This collective interference is exactly what the Hubbard–Stratonovich linearisation of the non-condensed sector encodes below. While mixed

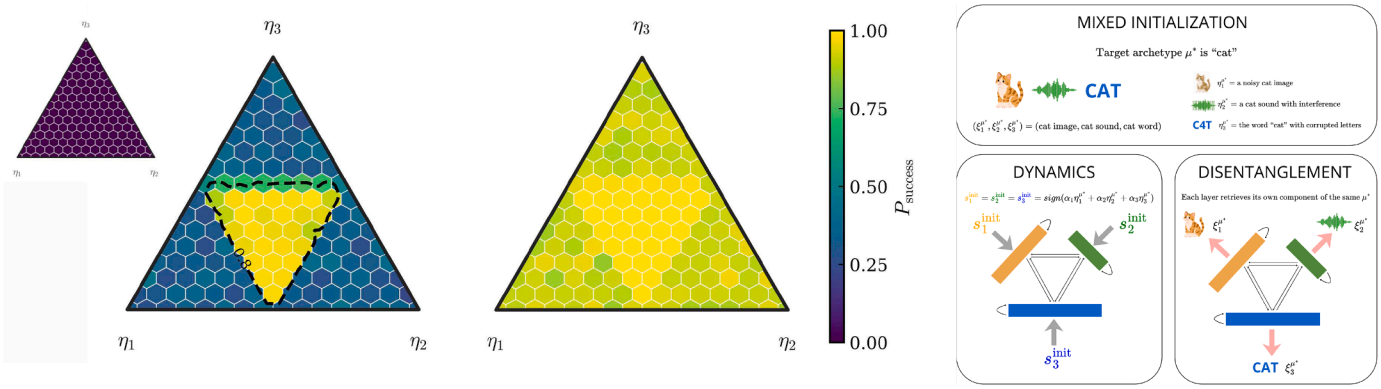


Fig. 2. Disentanglement under dreaming. *Left and centre*: Joint success rate on the cue-weight simplex (13). **Top-left inset:** Independent baseline (g_{hetero}=0): Joint success rate is identically zero ($P_{\text{success}} \equiv 0$) across the entire simplex, demonstrating that decoupled single-layer dreaming networks cannot perform cross-modal disentanglement. **Left:** LAM only (t=0): Retrieval succeeds only near the barycentre, where every layer receives some signal from its own modal view and inter-layer cooperation can bootstrap joint recovery; near the vertices, where a single view dominates and two layers start uninformed, joint convergence fails. **Centre:** Dreaming + LAM (t=64): The high-success region expands from the barycentre toward the vertices, covering nearly the entire simplex; even vertex cues (single-view evidence) now yield reliable joint retrieval. Dashed black contour at P=0.8. Colour encodes P_{success} on the *viridis* scale (dark = failure, bright = success). **Right:** Conceptual illustration of the disentanglement task: A common cross-modal cue blending noisy samples from all three sensory layers is presented simultaneously to every layer as the initial state; each layer must then converge independently to its own modality-specific archetype. Without dreaming, inter-pattern correlations trap each layer in a spurious mixed attractor; with dreaming, the spectral kernel $\mathbb{A}_l(t)$ suppresses off-diagonal correlations and carves out well-separated per-modality basins, enabling every layer to jointly recover its correct target from the shared cue. Simplex panels (a,b): N=400, K=72 (α=0.18), M=30, r=0.75, β=10, g_{auto}=0.82, g_{hetero}=0.41, cue noise ε=0.10, m_{th}=0.80; 200 Glauber steps, 40 independent replicates.

states with multiple macroscopic modes do exist as valid metastable solutions (in direct analogy with the symmetric mixture states of the Amit–Gutfreund–Sompolinsky solution for the Hopfield model; Amit et al., 1985), they correspond to a different phenomenological regime. Therefore, focusing on a single macroscopic mode (labelled as $\mu = 1$ by convention) is the standard and necessary assumption to evaluate the network's pristine retrieval capabilities.

The condensed empirical Mattis magnetisation is

$$\widehat{M}_l(\boldsymbol{s}) := \widehat{m}_l^1(\boldsymbol{s}) := \frac{1}{N}\sum_{i=1}^{N} \hat{\xi}_{i,l}^1 \, s_{i,l}. \tag{17}$$

Dreaming sector. The intra-layer dreaming term is linearised by N site-indexed Gaussian fields $\phi_{i,l} \sim \mathcal{N}(0,1)$, collected in the product measure $D\phi_l := \prod_i (2\pi)^{-1/2} e^{-\phi_{i,l}^2/2} d\phi_{i,l}$, and packaged into the complexified multi-spin

$$k_{i,l} := s_{i,l} + i\sqrt{\frac{t_l}{\beta(1+t_l)}}\,\phi_{i,l}. \tag{18}$$

In this HS representation, the purely spin overlap $\widehat{M}_l(\boldsymbol{s})$ is lifted to the complexified condensed dreaming overlap

$$\eta_l(\boldsymbol{s},\boldsymbol{\phi}) := \frac{1}{NM}\sum_{i=1}^{N}\sum_{k=1}^{M} \eta_{i,l}^{1,k}\, k_{i,l} = \frac{1}{N}\sum_{i=1}^{N} \hat{\xi}_{i,l}^1\, k_{i,l}. \tag{19}$$

Here $\eta_{i,l}^{1,k}$ denotes the kth noisy example of the condensed archetype in layer l.

The weight $\sqrt{t_l/[\beta(1+t_l)]}$ in the imaginary part is chosen so that Gaussian integration over $\boldsymbol{\phi}$ reproduces the dressed dreaming kernel $\mathbb{A}_l(t_l) = (1+t_l)(\mathbb{I} + t_l\Sigma_l)^{-1}$ rather than the bare identity (verified by completing the square; Appendix A). The non-condensed dreaming modes are further linearised by pattern-indexed fields z_μ^l drawn from the measure

$$Dz_l := \prod_{\mu\geq 2} \frac{dz_\mu^l}{\sqrt{2\pi}} \exp\Big(-\frac{(z_\mu^l)^2}{2(1+t_l)}\Big), \tag{20}$$

so that $z_\mu^l \sim \mathcal{N}(0, 1+t_l)$.

Inter-layer sector. For the non-condensed LAM modes, complex auxiliary fields $u_\mu^l, u_\mu^{l\dagger} \in \mathbb{C}$ ($l = 1,\dots,L$, $\mu \geq 2$) are introduced with a non-diagonal Gaussian measure $Duu^\dagger$ whose cross-covariance encodes the off-diagonal connectivity matrix $\mathbf{G}_0 \in \mathbb{R}^{L\times L}$ defined by

$$(\mathbf{G}_0)_{lm} := \begin{cases} g_{lm}, & l \neq m, \\ 0, & l = m. \end{cases} \tag{21}$$

Integrating out $u, u^\dagger$ against $Duu^\dagger$ regenerates $\exp\big(\frac{\beta}{2}\sum_\mu \boldsymbol{p}^{\mu\top}\mathbf{G}_0\boldsymbol{p}^\mu\big)$, recovering the non-condensed inter-layer coupling exactly (Lemma 1 of the appendix).

Exact decoupled partition function. Assembling the condensation split (16) with the three HS linearisations yields the exact representation

$$\begin{aligned}\mathcal{Z}(\beta \mid \mathcal{D}) = \sum_{\boldsymbol{s}} e^{S_{\text{LAM}}(\boldsymbol{s})} \prod_{l=1}^{L}\int D\phi_l\, e^{\beta S_{\text{Dream}}(\boldsymbol{s},\boldsymbol{\phi})} \\ \times \mathcal{N}_{\text{Dream}}(\boldsymbol{s},\boldsymbol{\phi})\,\mathcal{N}_{\text{LAM}}(\boldsymbol{s}),\end{aligned} \tag{22}$$

with signal factors

$$S_{\text{LAM}}(\boldsymbol{s}) := \beta \sum_{l<m} g_{lm}\, p_l^1(\boldsymbol{s})\, p_m^1(\boldsymbol{s}), \tag{23}$$

$$\beta S_{\text{Dream}}(\boldsymbol{s},\boldsymbol{\phi}) := \sum_{l=1}^{L} \frac{\beta N(1+t_l)}{2\Gamma_l}\,\eta_l(\boldsymbol{s})^2, \tag{24}$$

and noise factors

$$\begin{aligned}\mathcal{N}_{\text{Dream}}(\boldsymbol{s},\boldsymbol{\phi}) := \prod_{l=1}^{L}\int Dz_l \times \\ \times \exp\Big(\sqrt{\frac{\beta}{\Gamma_l N}}\sum_{\mu\geq 2}\sum_{i=1}^{N} \hat{\xi}_{i,l}^\mu z_\mu^l k_{i,l}\Big),\end{aligned} \tag{25}$$

$$\begin{aligned}\mathcal{N}_{\text{LAM}}(\boldsymbol{s}) := \int Duu^\dagger \times \\ \times \exp\Big(\sqrt{\frac{\beta}{N}}\sum_{\mu\geq 2}\sum_{l=1}^{L}\sum_{i=1}^{N} J_{i,l}^\mu s_{i,l} u_\mu^l\Big).\end{aligned} \tag{26}$$

The factorisation (22) is exact; no approximation has been made. Its three hidden-variable sectors correspond to the three families of auxiliary nodes in Fig. 1: the $\phi_{i,l}$ (purple) mediate intra-layer dreaming corrections; the z_μ^l (orange) absorb non-condensed dreaming fluctuations; and the complex pair $u_\mu^l, u_\mu^{l\dagger}$ (blue) implement the inter-layer LAM connectivity.

Table 1
Control parameters and order parameters of the DLAM. Layer indices $l, m \in \{1, \dots, L\}$ with $l < m$ for coupling parameters.

Symbol	Description
Control parameters	
$\alpha := K/N$	storage load
$\beta = 1/T$	inverse temperature
ρ_l	dataset entropy, Eq. (2)
t_l	sleep time (dreaming budget, per layer)
g_{lm}	inter-layer coupling strength
Order parameters (RS ansatz)	
$\hat{M}_l$	empirical Mattis magnetisation
$\bar{\eta}_l$	auxiliary signal variable (dreaming sector)
$\bar{q}_l$	Edwards–Anderson overlap (off-diagonal, dreaming)
$\bar{Q}_l$	diagonal k-variable overlap (dreaming)
$\bar{p}_l, \bar{P}_l$	conjugate dreaming-noise overlaps
H_l	spin-glass overlap (LAM sector)
$\bar{\Pi}_l$	LAM noise order parameter

The quenched average is handled by a Guerra-type interpolation: an auxiliary parameter $x \in [0,1]$ continuously deforms the interacting system (x=1) into a factorised one-body problem (x=0), and the fundamental theorem of calculus yields $f(x=1) = f(x=0) + \int_0^1 dx\,(df/dx)$. The streaming term df/dx is evaluated via Gaussian integration by parts (the Stein–Nishimori identity) and, under the RS self-averaging hypothesis, reduces to a deterministic expression involving the order parameters listed in Table 1.

3.2. Replica-symmetric (RS) free energy

Assembling the one-body contribution at x=0 and the streaming contributions from the three sectors yields the RS free energy:

$$\begin{aligned}-\beta f^{\mathrm{RS}}_{\mathrm{DLAM}} = &-\sum_{l<m}\frac{\beta\, g_{lm}}{\Gamma_{lm}}\hat{M}_l\hat{M}_m - \frac{\alpha}{2}\log\det \mathbf{C}(H) + \\ &+\frac{\alpha}{2}\mathbb{D}_{\mathrm{LAM}}(H) - \frac{\alpha\beta}{4}\sum_{l=1}^{L}\bar{\Pi}_l(1-H_l) + \\ &+\sum_{l=1}^{L}\Bigg[\log 2 + \mathbb{E}\big[\log\cosh\Psi_l\big] - \tfrac{1}{2}\log\bar{D}_l + \tfrac{\alpha}{2}\log\tfrac{1+t_l}{S_l} + \\ &+\frac{\alpha\beta\bar{q}_l(1+t_l)}{2S_l} - \frac{\alpha\bar{p}_l\, t_l}{2\bar{D}_l(1+t_l)} + \frac{\beta(\bar{D}_l-1)(1+t_l)}{2\bar{D}_l\, t_l} + \\ &-\frac{\beta(1+t_l)\bar{\eta}_l^2(t_l+\bar{D}_l)}{2\Gamma_l\,\bar{D}_l} + \frac{\alpha\beta}{2}(\bar{p}_l\bar{q}_l - \bar{P}_l\bar{Q}_l)\Bigg].\end{aligned} \tag{27}$$

Here $\mathbf{C}(H)$ is the $2L\times 2L$ LAM covariance matrix, $\mathbb{D}_{\mathrm{LAM}}(H)$ its quadratic correction, and $\bar{D}_l := 1 + \alpha\frac{t_l}{1+t_l}(\bar{P}_l - \bar{p}_l)$, $S_l := 1 - \beta(1+t_l)(\bar{Q}_l - \bar{q}_l)$. It is worth noting that for an arbitrary number of layers L the terms governing the inter-layer noise sector do not admit a general compact-form algebraic representation.[7] We provide in the Appendix the formulas for the specialisation ($L = 3$).

3.3. Self-consistency equations and the effective field

Extremising (27) with respect to each order parameter yields the closed RS self-consistency system. Its structure is most perspicuous when written in terms of the effective local field

$$\Psi_l = \underbrace{I_l\,\iota}_{\text{LAM noise}} + \underbrace{\sigma_l\,\varphi}_{\text{dreaming noise}} + \underbrace{\hat{X}_l\,(u_l + v_l)}_{\text{signal}}, \tag{28}$$

where $\iota, \varphi \sim \mathcal{N}(0,1)$ are independent standard Gaussians, $\hat{X}_l \overset{d}{=} \hat{\xi}^1_{i,l}$ with $\mathbb{E}[\hat{X}_l^2] = \Gamma_l$, and

$$\begin{aligned} I_l &= \sqrt{\tfrac{\alpha\beta}{2}\bar{\Pi}_l}, \quad \sigma_l = \frac{\sqrt{\alpha\beta\,\bar{p}_l}}{\bar{D}_l}, \\ u_l &= \sum_{m\neq l}\frac{\beta g_{lm}}{\Gamma_{lm}}\hat{M}_m, \quad v_l = \frac{\beta(1+t_l)}{\Gamma_l}\frac{\bar{\eta}_l}{\bar{D}_l}. \end{aligned} \tag{29}$$

This decomposition is the central object of the theory. The signal reaching neuron i on layer l has two sources: v_l, the autoassociative signal enhanced by dreaming, and u_l, the heteroassociative signal relayed from neighbouring layers. These compete against two independent noise channels: σ_l, the residual interference from the $K-1$ non-recalled patterns within the same layer (dreaming noise), and I_l, the noise generated by cross-talk in the inter-layer sector (LAM noise). Pattern recall succeeds when the signal dominates both noise terms.

The self-consistency equations, obtained by setting the derivatives of f^{RS} to zero, are:

Signal sector.

$$\hat{M}_l = \mathbb{E}\big[\hat{X}_l\tanh\Psi_l\big], \qquad \bar{\eta}_l = \frac{\hat{M}_l}{t_l + \bar{D}_l}. \tag{30}$$

Here $\hat{M}_l$ is the *empirical* Mattis magnetisation $N^{-1}\sum_i \hat{\xi}^1_{i,l}s_{i,l}$, i.e. the overlap with the stored empirical archetype rather than the hidden true archetype ξ^1_l. The two are related by $m^1_l = \hat{M}_l\, r_l/\Gamma_l + \mathcal{O}(N^{-1/2}) = \hat{M}_l/[r_l(1+\rho_l)] + \mathcal{O}(N^{-1/2})$ in probability (see Section 2.1 and the rescaled overlap $m_{\eta,l}$, which satisfies $m_{\eta,l} \approx m^1_l$ in the thermodynamic limit); the true-archetype overlap m^μ_l, which is the quantity plotted in simulation figures, is recovered from $\hat{M}_l$ via this rescaling.

Dreaming overlap sector.

$$\bar{q}_l = \mathbb{E}\big[\tanh^2\Psi_l\big], \tag{31}$$

$$\bar{Q}_l = \bar{q}_l + \frac{1-\bar{q}_l}{\bar{D}_l^2} - \frac{t_l}{\beta(1+t_l)\bar{D}_l}, \tag{32}$$

$$\bar{p}_l = \beta\frac{(1+t_l)^2}{S_l^2}\bar{q}_l, \qquad \bar{P}_l = \frac{1+t_l}{S_l} + \beta\frac{(1+t_l)^2}{S_l^2}\bar{q}_l. \tag{33}$$

LAM sector.

$$H_l = \mathbb{E}\big[\tanh^2\Psi_l\big], \qquad \bar{\Pi}_l = \frac{2}{\beta}\Big[\partial_{H_l}\log\det\mathbf{C}(H) - \partial_{H_l}\mathbb{D}_{\mathrm{LAM}}(H)\Big]. \tag{34}$$

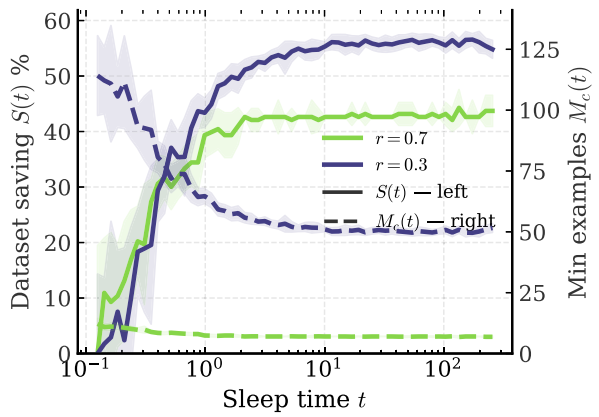
Remark 2. The overlap $\bar{q}_l$ (dreaming sector) and the spin-glass Edwards–Anderson parameter H_l (LAM sector) satisfy the same functional equation $\mathbb{E}[\tanh^2\Psi_l]$, but are closed by different companion equations; they are not redundant. For L=3 the LAM derivatives in (34) admit fully explicit scalar form; the expressions are given in the Appendix.

3.4. Phase boundaries and the data-saving mechanism

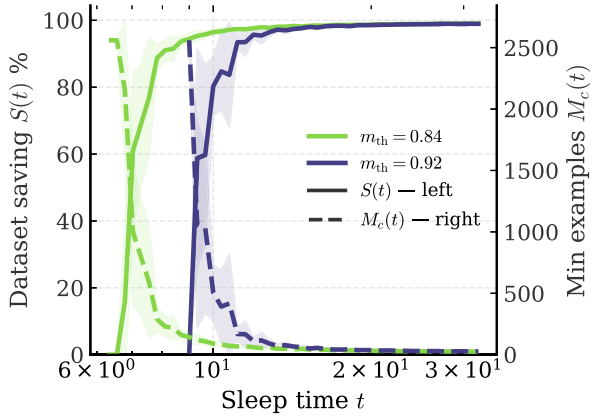
The system (30)–(34) admits three classes of fixed points:

- *Retrieval*: $\hat{M}_l > 0$, $\bar{q}_l > 0$, the network has recalled a stored pattern (i.e., the signal prevails).
- *Spin glass*: $\hat{M}_l = 0$, $\bar{q}_l > 0$, pattern information is lost and the system retains frozen disorder.
- *Paramagnetic*: $\hat{M}_l = 0$, $\bar{q}_l = 0$, where the noise prevails.

[7] This absence of a compact scalar form does not, however, affect the scalability of the self-consistency iteration: the only L-dependent objects it requires – $\det\mathbf{C}(H)$, the log-determinant gradient $\partial_{H_l}\log\det\mathbf{C}(H)$, and the quadratic correction $\mathbb{D}_{\mathrm{LAM}}(H)$ – are evaluated numerically at each step at the cost of an $L\times L$ determinant/inverse, i.e. $\mathcal{O}(L^3)$ per iteration, negligible against the $\mathcal{O}(N)$ Gaussian quadratures of the one-body sector. A symbolic routine generating these objects for arbitrary L is released with the code (see the Reproducibility statement), so that the fixed point can be closed at any L even where the free energy is too unwieldy to typeset.



(a) Unstructured random data.



(b) MNIST structured data.

Fig. 3. Data saving under dreaming. In both panels the bright (yellow-green) curve corresponds to the high-quality data condition and the dark (blue-purple) curve to the low-quality condition. **Left:** Critical number of training examples $M_c(t)$ (dashed, right axis) and dataset saving $S(t) = 1 - M_c(t)/M_c(0)$ (solid, left axis) versus sleep time, for synthetic data with r=0.70 (bright) and r=0.30 (dark). Moderate dreaming consistently reduces M_c, demonstrating the data–computation trade-off in its most direct form: off-line consolidation substitutes for a significant fraction of the required training examples, with savings reaching 40–60% at moderate sleep times. The effect is even more pronounced at lower r, highlighting the increased effectiveness of consolidation in the presence of higher initial data corruption. **Right:** Same dual-axis display on structured synthetic data (MNIST-like patterns, N=784, three horizontal bands used as independent views) at two retrieval thresholds ($m_{\rm th}$=0.84, bright; $m_{\rm th}$=0.92, dark): a stricter success criterion requires more examples at any fixed t, but dreaming erases most of that gap, confirming that the trade-off is robust to the choice of evaluation metric and to structured (non-random) pattern statistics. *Dataset pipeline for panel (b):* MNIST digits (28×28 pixels) are binarised by thresholding at the per-image median intensity (pixels above median set to +1, below to −1); K=10 randomly chosen digit classes serve as archetypes (one archetype per class, selected as the image closest to the class centroid); train/test sets are split 8:2 within each class, and M denotes the number of binarised training images used per class per view. Panel (a): N=300, K=30, $r\in\{0.70, 0.30\}$, $g_{\rm auto}$=1.0, $g_{\rm hetero}$=0.3, β=10; $M_c(t)$ determined via binary search (see text) with n=15 replicates and success criterion $P(\langle m\rangle \geq 0.8) \geq 0.8$.

The phase boundaries, namely the critical load α_c, the critical temperature T_c, and the critical dataset entropy ρ_c, are the loci where the retrieval fixed point ceases to exist or becomes thermodynamically subdominant, yielding a phase diagram whose boundaries delimit the retrieval, paramagnetic, and spin-glass states.

The role of dreaming is visible directly in the field coefficients (29). As the sleep time t_l increases:

- the dreaming-noise coefficient σ_l decreases, because the kernel $\mathbb{A}_l$ differentially attenuates the high-eigenvalue modes of Σ_l (which carry most of the inter-pattern crosstalk), reducing the effective noise overlap $\bar{p}_l$;
- the signal coefficient v_l increases, because the condensed retrieval direction is preserved (its effective coupling saturates at a finite value) while $\bar{D}_l$ stabilises.

This heuristic argument, based on simultaneous noise suppression and signal amplification, predicts that the critical boundary should shift outward with increasing t. In the single-layer limit this mechanism has been established rigorously (Alemanno et al., 2023; Fachechi et al., 2019); in the multi-layer model, the self-consistency Eqs. (30)–(34), which are exact under the RS ansatz, provide a closed characterisation of the equilibrium phase structure. Numerical solutions via fixed-point iterations confirm the predicted outward shift of the retrieval boundary across all parameter planes. Phase diagrams in the (ρ, t) and (α, t) planes, both from the self-consistency equations and from direct Monte Carlo simulations on finite-size networks, are presented in Section 4 and confirm the phenomenology quantitatively. In the (ρ, t) plane, the shift has a direct operational interpretation: at fixed load, the maximum tolerable dataset entropy $\rho_c(t)$ increases with t, so that fewer training examples are needed: the data–computation trade-off recently explored in the single-layer setting by Agliari et al. (2024a) and Agliari et al. (2024b).

The phase diagrams in the (α, T)-plane, obtained by iterating the self-consistency equations for varying inter-layer coupling g and presented in Section 4.3, display a weak re-entrant structure at intermediate temperatures. This feature is not specific to the DLAM architecture, but corresponds to the usual re-entrant behaviour already observed in replica-symmetric Hopfield-type models when the retrieval boundary approaches the spin-glass region. We therefore do not interpret this branch as a new DLAM-specific mechanism, but as a standard RS-level feature close to the glassy regime.

3.5. Consistency reductions

The RS closure recovers the expected limiting theories in the appropriate regimes. Removing the heteroassociative sector yields L independent single-layer supervised-dreaming problems; taking the limit $t_l \to 0$ for all l recovers a noisy-data multidirectional associative memory (note that a literal substitution $t_l = 0$ is not defined from the displayed formula because of the explicit $1/t_l$ term; the limit is understood in the sense $t_l \to 0^+$, which regularises the expression); removing the dreaming sector recovers the supervised LAM closure. The explicit reduction dictionaries are reported in the Appendix.

4. Numerical results

The analytical theory in Section 3 provides closed-form self-consistency equations describing the evolution of order parameters within the DLAM control parameter space in the thermodynamic limit. This framework suffices to derive the phase diagram of the network. Nevertheless, the Monte Carlo experiments that follow serve three purposes: first, to confirm the predicted phenomenology at finite N; second, to complement the analytical phase diagrams with finite-size scaling and quantitative tests where the RS solution serves as a benchmark; and finally, to probe dynamical phenomena—such as convergence acceleration and the temporal structure of disentanglement—that lie beyond the scope of the present theory.

All simulations use L=3 symmetric layers with N neurons each.[8] Two initialisation protocols are employed, matched to the tasks of:

- *Pattern retrieval.* Layer 1 (the cued layer) is initialised at overlap $m_0 = 1 - 2\varepsilon$ with the target archetype; layers 2 and 3 start from random configurations ($m_0 \simeq 0$). This protocol isolates the dynamics of cross-modal completion and is used in the temperature-dependence experiments.
- *Pattern disentanglement.* All layers are initialised with the same weighted cross-modal cue (13), blending single noisy examples drawn from each layer's dataset for a shared target archetype. The network must recover the correct modality-specific pattern on every layer simultaneously. This protocol is used in the majority of the experiments.

Retrieval is declared successful when the tail-averaged Mattis magnetisation exceeds a threshold m_{th} (typically 0.8).

Dataset-saving metric. In the experiments of Section 4.1 we measure the critical training-set size $M_c(t)$, the minimum number of examples per archetype such that retrieval is reliable at sleep time t. Operationally, $M_c(t)$ is found by binary search: starting from a bracket $[M_{\min}, M_{\max}]$ that is expanded geometrically until the upper bound yields reliable retrieval, the search bisects the interval until the gap is smaller than one example. At each candidate M, n=15 independent trials are run; retrieval is deemed reliable if at least 80% of trials achieve tail-averaged magnetisation $\langle m \rangle \geq 0.8$. The dataset saving is then $S(t) = 1 - M_c(t)/M_c(0)$.

4.1. Disentanglement under dreaming

We begin with the experiment that most directly tests the central prediction of the DLAM theory, as storage grows: can the dreaming kernel destabilise the glassy attractor that traps ambiguous mixture cues and promote the emergence of sharp, individual retrieval basins? The disentanglement task is the right stress test because it probes not just the depth of a single basin but the global organisation of the free-energy landscape under competing attractors.

To probe this prediction we deploy the simplex initialisation (13) defined in Section 2.6, sweeping all cue compositions $(\alpha_1, \alpha_2, \alpha_3)$ on a 30×30 simplex grid and recording the joint success rate averaged over 40 independent replicates per point.

Fig. 2(a, b) display the central visual result of the paper. Without dreaming, the joint success rate on the simplex is appreciable only near the barycentre, where every layer receives a diluted but nonzero signal from its own modal view and inter-layer cooperation can bootstrap all three layers. Moving toward the vertices, where the cue is dominated by a single view and two layers start essentially uninformed, the uninformed layers cannot recover and joint retrieval fails. With dreaming (e.g., $t = 64$), the picture changes qualitatively: the high-success region expands from the barycentre toward the vertices, covering nearly the entire simplex. The spectral filtering by $\mathbb{A}_l$ has cleaned the intra-layer basins so that the heteroassociative signal can rescue even uninformed layers; the transition is sharp in t, and above a threshold disentanglement becomes the generic outcome across the full simplex.

Alongside the benefits observed in the disentanglement task, the regularising action of dreaming provides a complementary and equally crucial advantage concerning data efficiency, an aspect that we quantify in terms of data saving (Fig. 3). While the Hebbian dreaming mechanism was recently shown to reduce the volume of training examples needed for retrieval in single-layer networks (Agliari et al., 2024a,b), our current simulations structurally extend and reaffirm this evidence within a considerably more complex multimodal architecture. As sleep time progresses, the critical number of training examples required to learn a pattern, $M_c(t)$, reliably decreases; the resulting relative dataset saving $S(t) = 1 - M_c(t)/M_c(0)$ scales rapidly up to a significant fraction of the required training data. Equivalently, $M_c(t)$ gives a direct quantitative comparison between a dream-consolidated network trained with fewer examples and an undreamed network trained with more examples: both are evaluated at the same retrieval criterion, and the ratio $M_c(t)/M_c(0)$ measures the amount of data replaced by off-line consolidation. This outlines a concrete, operational *data–computation trade-off*: the network compensates for a sparse or highly noisy training set by expending computational effort during the off-line consolidation phase. The importance of this trade-off holds acute practical value since, in many real-world multimodal scenarios, collecting large batches of clean, aligned data across multiple sensors is far more cost-prohibitive and bottleneck-prone than allocating extra cycles for internal parameter tuning.

4.2. Robustness and phase structure

The disentanglement results demonstrate a qualitative transformation of the free-energy landscape at a single operating point. A natural follow-up question is whether this effect is robust across the parameter space explored, or whether it is a fragile anecdote tied to a particular choice of load and noise. The answer, explored across three complementary projections of the phase diagram, is that the dreaming benefit is broad and systematic.

Fig. 4(a) quantifies the basin expansion. The critical cue corruption ε_c, the noise fraction at which retrieval drops below 50%, shifts markedly rightward as t increases: dreaming enlarges the retrieval basins in configuration space (equivalently, it broadens the set of corrupted or partial input cues from which the network still converges to the correct stored pattern), making the network robust to substantially noisier inputs.[9]

The cooperative contribution is visible in Fig. 4(b). In the (α, ε)-plane three retrieval boundaries are compared: the baseline without dreaming, dreaming alone, and the full dreaming + cooperation regime. Inter-layer coupling shifts the phase boundary outward by a further margin, so that the combined regime encloses a retrieval region substantially larger than either mechanism alone. Dreaming and inter-layer coupling are thus synergistic: they open operating regions not reached by either mechanism in isolation within the explored regime.

Fig. 4(c) maps the retrieval boundary in the (ρ, t)-plane for several values of the storage load α, the operational content of the data–computation trade-off. As argued in Section 3.4, the noise-suppression mechanism of the dreaming kernel increases ρ_c with t: the boundary shifts rightward (to larger tolerable ρ_c), so that at fixed load the same retrieval quality is achieved with fewer training examples. The picture is consistent with the analytical predictions and with the data-saving results of Agliari et al. (2024a,b) established in the single-layer setting. At higher α the boundary moves inward, because more patterns demand cleaner data, but the relative gain from dreaming is preserved, confirming that the spectral filtering mechanism operates independently of the storage load.

[8] Concretely, each Monte Carlo run samples the Boltzmann–Gibbs equilibrium of the DLAM cost function (7) at inverse temperature $\beta = 1/T$ by asynchronous single-spin-flip (Glauber/Metropolis) dynamics acting on the already-consolidated couplings $\mathbb{A}_l(t_l)$, at finite system sizes $N \leq 500$ neurons per layer and $L = 3$ layers. The observables extracted are the tail-averaged Mattis magnetisations m_l (the overlap with the true, hidden archetype) and the success probabilities derived from them, each averaged over the number of independent realisations stated in the corresponding caption.

[9] Note that ρ and ε operate at distinct stages: ρ represents static synaptic noise, modifying the energy landscape during consolidation, whereas ε denotes dynamic state noise, quantifying the initial distance from the attractor. While dreaming improves resilience to both, they are not analytically interchangeable, as ρ alters the basin topology, while ε measures the initial position within that topology.

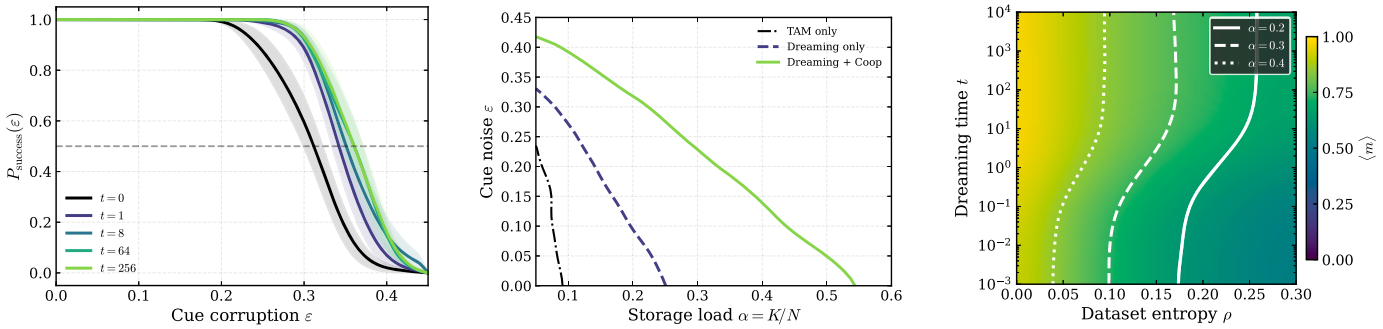


Fig. 4. Robustness and phase structure. Left: Noise robustness: retrieval probability $P_{\text{success}}(\varepsilon)$ versus cue corruption level for increasing sleep times t. Curves are coloured with a viridis gradient (dark = short dreaming, bright = long dreaming); black marks the undreamed baseline t=0. The critical corruption ε_c at which retrieval drops to 50% shifts markedly rightward with t, quantifying basin expansion. Shaded bands: $\pm 1\sigma$ over independent realisations. **Centre:** Two-dimensional phase boundary in the (α, ε)-plane, comparing three regimes: LAM only (black dash-dotted), dreaming alone (dark viridis, dashed), and dreaming combined with inter-layer cooperation (bright viridis, solid). The combined boundary encloses a substantially larger retrieval region, demonstrating genuine synergy. The viridis heatmap shows P_{success} for the cooperative condition. **Right:** Phase diagram in the (ρ, t)-plane for three values of the storage load α. The retrieval boundary (iso-contour at $\langle m \rangle$=0.85) separates the *retrieval phase* (upper-left region: low dataset entropy ρ and large dreaming time t, shown bright in the viridis heatmap) from the *non-retrieval phase* (lower-right region, dark in the heatmap: high entropy or insufficient sleep). Moving along the vertical axis at fixed α, increasing t pushes the boundary rightward, meaning that a larger ρ (poorer or scarcer data) is tolerated at fixed retrieval quality, equivalently, fewer training examples suffice for the same performance. At higher load the boundary moves inward (lower ρ_c at fixed t), but the relative gain from dreaming is preserved. Background: Gaussian-smoothed $\langle m \rangle$ for α=0.4. Panel (a): N=500, M=30, r_1=0.50, $r_{2,3}$=0.80, α=0.15, $t \in \{0, 1, 8, 64, 256\}$, β=10, 50 Glauber steps, 20 replicates. Panel (c): N=500, M=20, α=0.40 (background), β=100, 20 steps, 10 runs per cell on an 80×80 $(\rho, \log t)$ grid, retrieval threshold $\langle m \rangle$=0.85.

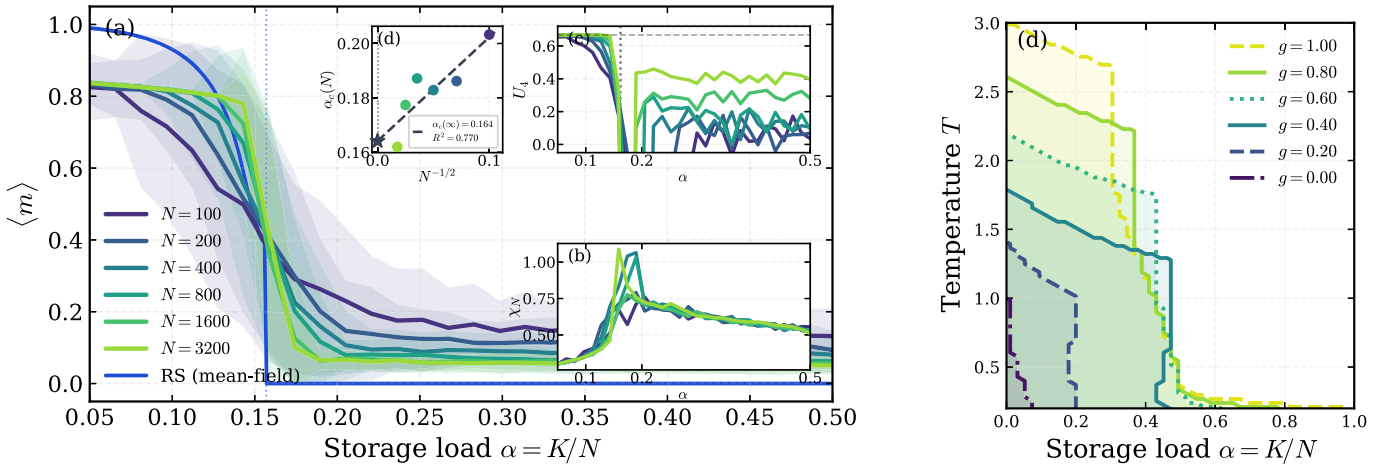


Fig. 5. Critical behaviour and finite-size scaling. (a) Finite-size scaling of the retrieval transition: mean magnetisation $\langle m \rangle$ versus storage load α for increasing system sizes $N \in \{100, 200, 400, 800, 1600, 3200\}$. The blue solid line is the theoretical prediction obtained from the numerical solution of the Replica Symmetric (RS) equations. The transition steepens progressively and provides the main visual summary of the critical crossover. **(b)** Inset: susceptibility $\chi_N = N\,\mathrm{Var}(m)$, showing the growth of the critical peak with N. **(c)** Inset: Binder cumulant U_4, whose crossing identifies a size-independent estimate of the critical load α_c. **(d)** Inset: Amit–Gutfreund–Sompolinsky (AGS) finite-size-scaling extrapolation of the pseudo-critical load. The per-size estimate $\alpha_c(N)$, read off from the susceptibility peak in panel (b), is plotted against $N^{-1/2} - \nu = \frac{1}{2}$ being the classical mean-field exponent of the AGS power-law fit – and fitted with $\alpha_c(N) = \alpha_c(\infty) + c\,N^{-1/2}$. The extrapolation to the thermodynamic limit ($N^{-1/2} \to 0$, star) yields $\alpha_c(\infty) \simeq 0.164$ ($R^2 \simeq 0.77$), in agreement with the RS spinodal of panel (a) and confirming that the finite-N transition extrapolates to a genuine thermodynamic critical load rather than drifting to zero. **(e)** Analytical retrieval regions in the (α, T)-plane for six representative values of the inter-layer coupling, g=0.0, 0.2, 0.4, 0.6, 0.8, 1.0, obtained by numerical solution of the RS self-consistency Eqs. (30)–(34). The silhouettes are overlaid on the same axes, showing the expansion of the retrieval phase as coupling increases. Parameters for (a–d): Monte Carlo at r=0.5, M=40, dreaming time t=8, β=1 (150 replicates per (N, α)). Parameters for (e): ρ=0.10, retrieval regions identified via multi-start fixed-point iteration with threshold $\epsilon_M = 10^{-3}$.

4.3. *Critical behaviour and finite-size scaling*

A natural question is whether the retrieval-to-noise transition visible in the phase diagrams is a genuine phase transition or a finite-size artefact.

The retrieval-to-noise transition visible in the preceding phase diagrams must be distinguished from a finite-size artefact before it can be interpreted as evidence of a genuine phase transition. Standard finite-size scaling analysis provides the necessary test.

The magnetisation curves (Fig. 5a) steepen progressively with N: as the system grows, the crossover from high to low magnetisation near the critical load sharpens markedly with system size, consistent with an increasingly abrupt thermodynamic transition. The susceptibility $\chi_N = N\,\mathrm{Var}(m)$ (Fig. 5b) develops a peak whose height grows with system size in a manner consistent with power-law divergence, $\chi_{\max} \sim N^{\gamma/\nu}$, the hallmark of a genuine critical point. In the present setting, the Binder cumulant $U_4 = 1 - \langle m^4 \rangle / (3\langle m^2 \rangle^2)$ (Fig. 5c) provides a convenient finite-size diagnostic of the retrieval transition: the curves for different N display an approximately size-independent crossing near the critical load α_c. Taken together, these three diagnostics confirm that the boundary observed in the Monte Carlo experiments corresponds to a thermodynamic phase transition, rather than a gradual degradation caused by finite system size.

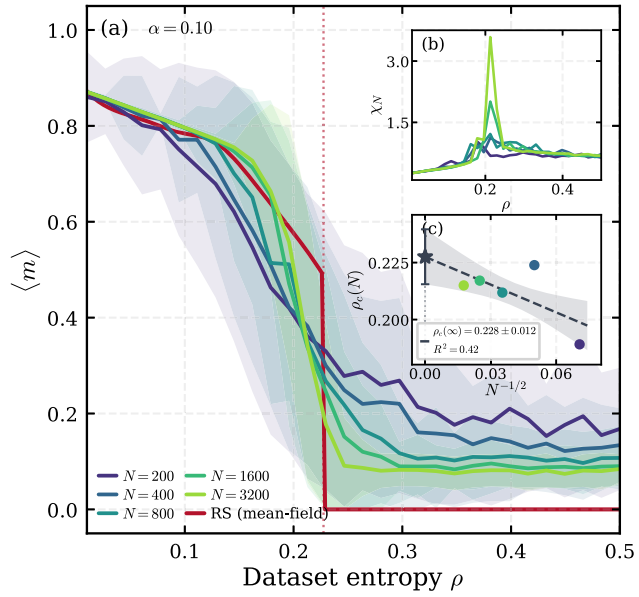


Fig. 6. Retrieval transition along the dataset-entropy axis, and comparison with the replica-symmetric prediction. (a) Mean magnetisation $\langle m \rangle$ versus the dataset entropy $\rho = (1 - r^2)/(Mr^2)$ at *fixed* storage load $\alpha = 0.10$, for increasing system sizes $N \in \{200, 400, 800, 1600, 3200\}$ (shaded bands: one std across disorder realisations). The same ρ is imposed on all three layers, and is tuned at fixed $M = 40$ through the example quality $r(\rho) = (1 + M\rho)^{-1/2}$; since the RS equations depend on the dataset only through ρ – the quality r cancels identically once the overlap is written in terms of the rescaled magnetisation $m_\eta = \hat{M}/[r(1 + \rho)]$, cf. Section 2.1 – tuning r at fixed M or M at fixed r traces the same theoretical curve. The solid line is the retrieval branch of the RS self-consistency Eqs. (30)–(34). **(b)** Inset: susceptibility $\chi_N = N\,\mathrm{Var}(m)$, whose peak sharpens and grows with N and defines the pseudo-critical entropy $\rho_c(N)$. **(c)** Inset: finite-size-scaling extrapolation of $\rho_c(N)$ against $N^{-1/2}$, fitted with the mean-field law $\rho_c(N) = \rho_c(\infty) + c\,N^{-1/2}$ ($\nu = \frac{1}{2}$, as in Fig. 5d; shaded band: 1σ confidence interval). The extrapolation to the thermodynamic limit (star) gives $\rho_c(\infty) \simeq 0.228 \pm 0.012$ ($R^2 \simeq 0.42$), which coincides with the spinodal of the RS branch in panel (a): the collapse of retrieval driven by data degradation occurs, in the thermodynamic limit, at the entropy predicted by the replica-symmetric theory. Other parameters: Monte Carlo at dreaming time $t = 8$, $\beta = 5$, 150 replicates per (N, ρ) cell, 600 sweeps with tail average over the last 150.

To turn this qualitative picture into a quantitative comparison with the mean-field prediction, we extend the simulations up to $N = 3200$ and extract, for each size, the pseudo-critical load $\alpha_c(N)$ from the peak of the susceptibility χ_N (Fig. 5b). Within the Amit–Gutfreund–Sompolinsky finite-size-scaling framework these estimates are expected to approach their thermodynamic value with the classical mean-field exponent $\nu = \frac{1}{2}$, $\alpha_c(N) = \alpha_c(\infty) + c\,N^{-1/2}$. Plotting $\alpha_c(N)$ against $N^{-1/2}$ (Fig. 5d) indeed yields a clean linear trend whose extrapolation to the thermodynamic limit gives $\alpha_c(\infty) \simeq 0.164$ (coefficient of determination $R^2 \simeq 0.77$), in quantitative agreement with the spinodal of the RS curve overlaid in panel (a). The finite-size estimates therefore converge to a strictly positive critical load rather than drifting toward zero, which is precisely the signature distinguishing a genuine thermodynamic transition from a finite-size crossover; the residual gap between the $N \leq 800$ data and the RS curve visible above the transition is the expected $\mathcal{O}(N^{-1/2})$ correction and is reabsorbed by the extrapolation.

The same analysis can be repeated along the second control axis: Fig. 6 fixes the storage load at $\alpha = 0.10$, comfortably inside the retrieval phase, since $\alpha_c(\infty) \simeq 0.164$, and degrades the training set instead, sweeping the dataset entropy ρ at fixed M through the example quality. The phenomenology mirrors that of the load transition: the magnetisation curves steepen monotonically with N, the susceptibility develops a peak that grows and sharpens with system size (inset b), and the pseudo-critical entropies $\rho_c(N)$ extracted from those peaks fall on a straight line when plotted against $N^{-1/2}$ (inset c), extrapolating to a strictly positive $\rho_c(\infty) \simeq 0.228$. The retrieval collapse driven by data degradation is therefore a genuine thermodynamic transition and not a finite-size crossover, exactly as its load-driven counterpart.

Crucially, the RS retrieval branch computed along the same axis terminates at a spinodal that coincides with $\rho_c(\infty)$, and tracks the Monte Carlo data throughout the retrieval region. That the two independent sections of the phase diagram are both reproduced quantitatively indicates that the replica-symmetric solution captures the location of the retrieval boundary faithfully over the region of parameter space explored here, notwithstanding the fact that a full de Almeida–Thouless stability analysis (which would delimit the regime in which the ansatz is exact rather than merely accurate) remains outside the scope of the present work and is left to a dedicated study.

The analytical panel (Fig. 5e), obtained by numerically iterating the RS self-consistency Eqs. (30)–(34) to convergence, complements the finite-size analysis with a direct comparison among six representative couplings g. Overlaying the six retrieval silhouettes on the same (α, T) axes makes the structural effect of cooperation immediate: as the inter-layer coupling g increases, the region of stable retrieval expands outward substantially. Moreover, the qualitative agreement between these purely analytical phase maps and the threshold dynamics extrapolated from our finite-size Monte Carlo simulations supplies a strong, direct validation of the replica-symmetric framework.

4.4. Design rules and trade-offs

The preceding sections establish the qualitative and quantitative effects of dreaming on the retrieval region; we now turn to the complementary engineering question: *given a finite sleep budget, how should it be allocated among the competing cognitive tasks?* This perspective translates the physics of the preceding sections into actionable design principles for multi-layer architectures.

Fig. 7(a) maps the retrieval boundary in the (α, t)-plane for several values of the data quality r. The phase boundary shifts rightward with increasing sleep time: at fixed r, moderate dreaming nearly doubles the viable storage load. Lower data quality pushes the boundary inward, but the relative gain from dreaming is preserved, confirming the generality of the spectral filtering mechanism. The picture is consistent with the (ρ, t) diagrams of Section 4.2 and provides a complementary view centred on storage capacity rather than dataset entropy.

Fig. 7(b) addresses network heterogeneity from an architectural standpoint. The marginal gain ΔP measures the increase in joint retrieval success when additional helper layers are introduced alongside the weak cued layer (fixed at r_1=0.50), as a function of the helper data quality $r_{2,3}$. At high $r_{2,3}$ (high helper quality), the first helper provides a dramatic boost: cross-modal support from a clean partner injects the heteroassociative signal u_l into the effective field (28) and can rescue a layer that would otherwise stall in the non-retrieval phase on its own. As $r_{2,3}$ decreases toward the weak layer's quality, the helper becomes progressively less informative and the marginal benefit declines monotonically. The second helper always contributes less than the first, consistent with diminishing returns, but its incremental gain peaks at intermediate $r_{2,3}$ values, precisely the critical regime where the first helper partially rescues the weak layer but the additional signal channel is still consequential. Crucially, both marginal gains remain non-negative throughout the explored range (at fixed r_1=0.50, varying helper quality $r_{2,3} \in [0.30, 0.90]$): additional helper layers never degrade performance within this regime, a practical consequence of the non-negative definiteness of the heteroassociative coupling term u_l in the effective field (28). This observation leads to a practical design rule: in a constrained architecture, invest in one high-quality helper layer before adding further redundancy; the second helper is worth adding only when the first has been fully exploited.

Finally, the Pareto front (Fig. 7(c)) crystallises the fundamental auto/hetero trade-off in the stressed regime (α=0.20, near the capac-

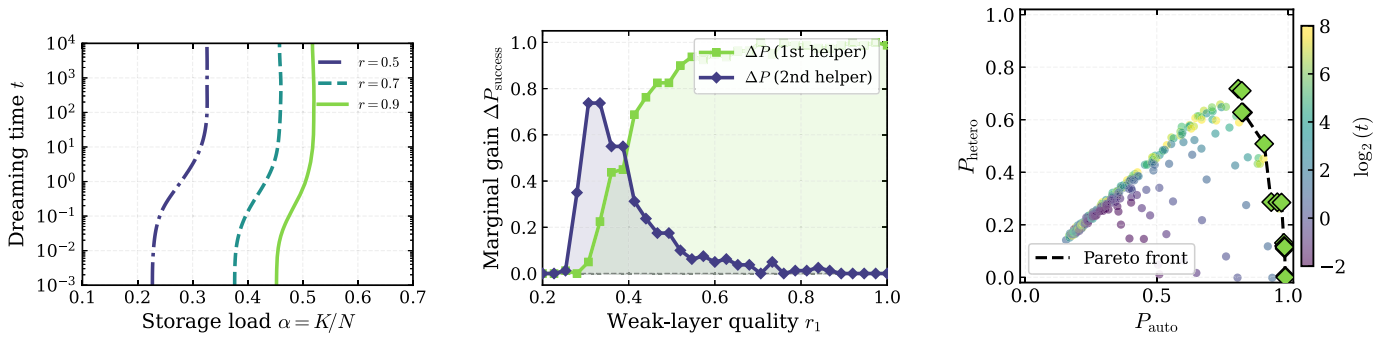


Fig. 7. Design rules for cooperative dreaming networks. Left: Phase diagram in the (α, t)-plane for three values of the data quality r. The retrieval boundary (iso-contour at $\langle m \rangle$=0.85) separates the *retrieval phase* (upper-left region: low storage load α and large dreaming time t) from the *non-retrieval phase* (lower-right region, where the load exceeds the network's effective capacity at the given sleep budget). The boundary shifts rightward with increasing sleep time, meaning that higher storage loads are tolerated: moderate dreaming can nearly double the viable α at fixed r. At lower r (noisier or scarcer data) the boundary moves inward, but the relative gain from dreaming is preserved, confirming that spectral filtering remains effective regardless of raw data quality. Background: $\langle m \rangle$ heatmap for r=0.5. **Centre:** Marginal gain ΔP from additional helper layers as a function of the helper-layer data quality $r_{2,3}$ (weak layer fixed at r_1=0.50). Bright (yellow-green) curve: $\Delta P = P_{\text{success}}(1 \text{ helper}) - P_{\text{success}}(0 \text{ helpers})$; dark (blue-purple) curve: incremental gain of a second helper, $P_{\text{success}}(2 \text{ helpers}) - P_{\text{success}}(1 \text{ helper})$. The first helper provides the dominant boost at high $r_{2,3}$ (high-quality helpers), where cross-modal support is essential to rescue the weak layer; as helper quality decreases the marginal benefit diminishes. The second helper peaks at intermediate $r_{2,3}$ (the critical regime where the first helper partially rescues the weak layer but additional signal channels are still consequential) and remains positive throughout the explored range, confirming that richer cross-modal coupling never degrades performance. Shaded bands indicate the signed contribution of each helper relative to the zero-helper baseline. **Right:** Pareto front in the stressed regime (α=0.20, near capacity without dreaming). Each scatter point represents a specific auto/hetero coupling allocation and sleep time t; colour encodes $\log_2 t$ on the viridis scale (dark = little dreaming, bright = long dreaming). The horizontal axis P_{auto} measures retrieval success when the network is initialised from a partial within-layer cue (standard pattern reconstruction); the vertical axis P_{hetero} measures cross-modal completion success when the cue is provided in a different layer (heteroassociative propagation). Without dreaming (dark cluster near the origin) both tasks fail in the high-load regime; increasing t expands the jointly achievable performance region toward the upper-right. Diamond markers trace the Pareto-optimal frontier: no configuration above this curve exists, so it defines the efficiency boundary for the simultaneous auto/hetero allocation problem. Its convexity implies that a single intermediate sleep budget serves both tasks nearly as well as two specialised budgets. Both axes are success probabilities in $[0, 1]$, and a configuration is preferable the closer it lies to the upper-right corner, where *both* the autoassociative (P_{auto}) and the heteroassociative (P_{hetero}) tasks succeed; a configuration is *Pareto-optimal* when neither coordinate can be improved without sacrificing the other, and the diamond frontier collects exactly these non-dominated points, i.e. the best trade-offs jointly attainable in this stressed regime. Panel (a): N=500, M=20, r=0.5 (background), β=100, $\alpha \in [0.1, 0.7]$, $t \in [10^{-3}, 10^4]$ (log-sampled), 20 steps, 10 runs per cell. Panel (b): N=500, α=0.28, M=30, r_1=0.50, $r_{2,3} \in [0.30, 0.90]$, t=4, g_{auto}=0.82, g_{hetero}=0.41, β=10, 50 steps, 30 replicates. Panel (c): N=500, α=0.20, M=20, r=0.6, β=8, ε=0.15, $t = 2^k$ ($k \in \{-2, \ldots, 8\}$), 200 steps, 50 replicates.

ity limit without dreaming). Without dreaming (dark points clustered near the origin of the P_{auto}–P_{hetero} plane) neither task is reliably achievable: the high storage load overwhelms both intra-layer pattern completion and cross-modal signal propagation. Increasing t (progressively brighter points) lifts the cloud toward the upper-right corner, jointly expanding the achievable performance space for both tasks. The Pareto-optimal frontier (diamond markers) traces the efficiency boundary: no configuration can improve P_{auto} without degrading P_{hetero}, or vice versa, above this curve. The convexity of the frontier implies that intermediate allocations of the auto/hetero couplings and sleep budget t are Pareto-efficient: a single cooperative network can simultaneously achieve reliable autoassociative recall and reliable cross-modal completion without severe penalty on either objective.

4.5. Dynamics and convergence

The equilibrium analysis captures the asymptotic properties of the DLAM; a natural complementary question is whether dreaming also affects how rapidly the network reaches equilibrium and how the layers coordinate in time. The answer is affirmative on both counts. All simulations reported in this section employ the synchronous Glauber rule detailed in Algorithm 1; the full implementation is publicly available at https://github.com/andrea-ladiana/dlam-mc-core.

The magnetisation trajectories (Fig. 8a) provide direct evidence: without dreaming, layer 1 reaches a metastable plateau at intermediate magnetisation, a signature of the rugged free-energy landscape characteristic of the high-load regime. With moderate dreaming ($t \geq 2$), these local minima are smoothed out, the plateau disappears, and all three layers converge rapidly to high magnetisation. The inter-layer dynamics displays a clear causal structure: layers 2 and 3 lag behind layer 1 by 1–2 Monte Carlo sweeps, reflecting the time needed for the heteroassociative signal to propagate across layers. This lag has a simple dynamical origin. The cue first builds a finite magnetisation in the initially informed layer; only then does this magnetisation enter the effective fields of the other layers through the heteroassociative contribution $u_l = \sum_{m \neq l} (\beta g_{lm} / \Gamma_{lm}) \hat{M}_m$ in Eq. (29). The observed 1–2 sweep delay is therefore the finite-time propagation of the cooperative signal under the Monte Carlo update rule.

The convergence time $T_{0.8}$, operationally defined as the number of Monte Carlo sweeps required for the cued layer's magnetisation to reach $m_1 \geq 0.8$ (Fig. 8b), quantifies the acceleration: it decreases monotonically with dreaming time, while the fraction of trials that successfully converge increases in parallel. Dreaming thus improves not only the asymptotic success rate but also the speed of retrieval, a distinction that matters in practice, since a network that converges slowly may be operationally useless even if it eventually reaches the correct attractor.

4.6. Sleep-budget allocation under heterogeneous data quality

A realistic multimodal system will have layers of unequal quality: one sensory channel may be clean and abundant while another is noisy or data-poor. In that setting, the practically relevant question is how a finite dreaming budget should be distributed across weak and strong layers. This is, structurally, a resource-allocation problem under a fixed total budget, formally analogous to duty-cycle allocation problems studied in other disciplines (for instance the almost-blank-subframe allocation devised for LTE-U/Wi-Fi coexistence; Alhulayil & López-Benítez, 2019) in which a shared, finite "on-time" must be partitioned among competing users (here, layers) so as to optimise a global objective.

Fig. 9 yields a practical design rule with non-trivial structure. The experiment fixes the total sleep budget τ and sweeps the allocation frac-

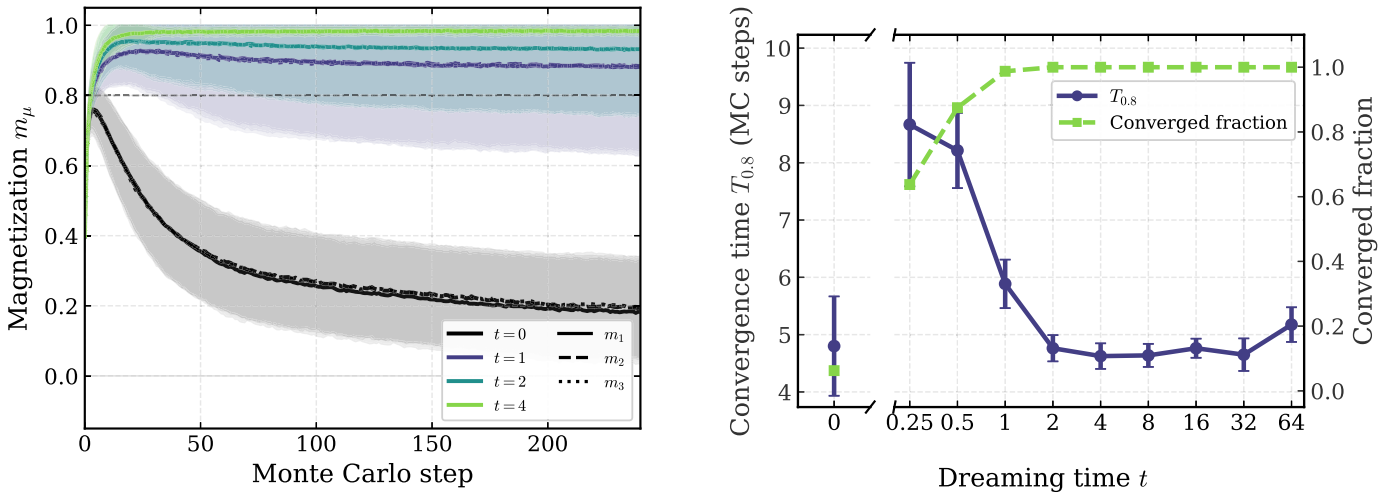


Fig. 8. Dynamics and convergence acceleration. Left: Magnetisation traces m_1, m_2, m_3 (solid, dashed, dotted) under the disentanglement protocol for selected dreaming times $t \in \{0, 1, 2, 4\}$. Without dreaming (t=0, black) the system stalls at a metastable plateau; increasing t smooths the landscape and accelerates convergence. Shaded bands: $\pm 1\sigma$. **Right:** Convergence acceleration: convergence time $T_{0.8}$ (steps to reach $m_1 \geq 0.8$, left axis) and converged-trial fraction (right axis) versus dreaming time t. Both metrics improve monotonically, demonstrating that dreaming accelerates retrieval. Note the broken x-axis in the right panel, which combines the always-awake baseline at $t = 0$ with a logarithmic scale for $t > 0$. Both panels: N=500, M=30, r=0.70, 80 replicates. Panel (a): α=0.30, β=2, ϵ=0.10, g_{auto}=0.82, g_{hetero}=0.41, $t_{\text{dream}} \in \{0, 1, 2, 4\}$, 240 Glauber steps. Panel (b): α=0.20, β=5, ϵ=0.20, g_{auto}=0.98, g_{hetero}=0.12, $t_{\text{dream}} \in [0, 64]$ (log-sampled), 40 steps.

Algorithm 1: Synchronous Glauber dynamics for the 3-layer DLAM.

Input : Empirical patterns $\{\hat{\boldsymbol{\xi}}_l^\mu\}$; dreaming kernels $\mathbb{A}_l(t_l)$ (Eq. (3)); couplings $\{g_{lm}\}$, normalisations $\{\Gamma_l, \Gamma_{lm}\}$; inverse temperature $\beta = 1/T$; initial states $\boldsymbol{s}_l^{(0)} \in \{-1, +1\}^N$; number of steps T

Output: Magnetisation trajectories $\{m_l^{(t)}\}_{t=0}^T$, final states $\{\boldsymbol{s}_l^{(T)}\}$

for $t = 1, \ldots, T$ **do**
  for $l = 1, 2, 3$ (independently) **do**
    // Pattern-space empirical overlaps (K-vector)
    $\hat{m}_l^\mu \leftarrow \frac{1}{N} \sum_{j=1}^N \hat{\xi}_{j,l}^\mu \, s_{j,l}^{(t-1)}, \quad \mu = 1, \ldots, K$
    // Mix intra- and inter-layer signals
    $v_l^\mu \leftarrow \frac{g_{ll}}{\Gamma_l} \left[\mathbb{A}_l(t_l) \hat{\boldsymbol{m}}_l\right]_\mu + \sum_{m \neq l} \frac{g_{lm}}{\Gamma_{lm}} \hat{m}_m^\mu, \quad \mu = 1, \ldots, K$
    // Back-project to site space: effective local field
    $h_{i,l} \leftarrow \sum_{\mu=1}^K \hat{\xi}_{i,l}^\mu \, v_l^\mu, \quad i = 1, \ldots, N$
    // Synchronous Glauber flip
    draw $u_i \overset{\text{iid}}{\sim} \mathcal{U}[0, 1]$ for $i = 1, \ldots, N$
    $s_{i,l}^{(t)} \leftarrow +1$ if $u_i < \frac{1}{2}\left(1 + \tanh(\beta h_{i,l})\right)$, else -1
  end
  $m_l^{(t)} \leftarrow \frac{1}{N} \sum_{i=1}^N \xi_{i,l}^{\mu\star} \, s_{i,l}^{(t)}, \quad l = 1, 2, 3$
end

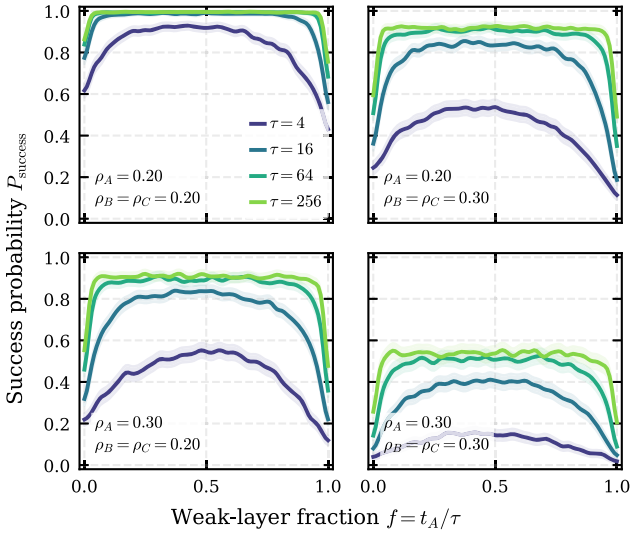


Fig. 9. Cooperative sleep-budget allocation under variable data-quality gaps. A total budget τ is split between the target layer (fraction f) and the helpers (fraction $(1-f)/2$ each). The four panels explore different combinations of dataset entropies for the target layer ($\rho_A \in \{0.20, 0.30\}$) and the helpers ($\rho_B = \rho_C \in \{0.20, 0.30\}$). Across all configurations, a broad interior optimum emerges and shifts leftward as τ grows, reflecting the saturation of the target layer's own dreaming. Parameters: N=500, α=0.18, M=35, β=10; $\tau \in \{4, 16, 64, 256\}$, g_{auto}=0.82, g_{hetero}=0.41, cue noise ϵ=0.15, m_{th}=0.80; 100 Glauber steps, 200 replicates.

tion $f = t_A/\tau$ (fraction directed to the target layer, so the two helpers each receive $(1 - f)/2$), recording P_{success} at four representative values of $\tau \in \{4, 16, 64, 256\}$. To probe the sensitivity of this mechanism to the heterogeneity of the datasets, and specifically to observe the effect of reducing the quality gap, we examine a 2×2 grid of configurations. By varying both the target layer's entropy ($\rho_A \in \{0.20, 0.30\}$) and the helpers' entropy ($\rho_B = \rho_C \in \{0.20, 0.30\}$), the diagonal panels ($\rho_A = \rho_B$) naturally represent the limiting case of a zero gap (homogeneous datasets), serving as a baseline for the heterogeneous cases.

Several features of the curves are noteworthy and robust across all parameter combinations. First, concentrating all dreaming on the target layer ($f \to 1$) consistently underperforms: at f=1 the helpers receive no sleep at all, so their noise-dominant basins cannot relay a clean heteroassociative signal u_l to the target layer, even if the latter's own intra-layer landscape has been perfectly cleaned. Conversely, allocating all budget to the helpers ($f \to 0$) also underperforms: the target layer is left undreamed, its non-condensed eigenmodes remain at their initial magnitude, and the inter-layer signal cannot overcome the strong intra-layer noise even with clean helpers. The optimum $f^\star$ lies in the interior, robustly between 0.2 and 0.7 across all budgets and quality gaps. Naturally, when the gap is smaller or absent (e.g., $\rho_A = \rho_B$, right column), the optimum for small budgets is closer to a symmetric split $f \approx 1/3$,

whereas a larger gap (left column) initially commands a larger share of the budget for the weaker layer.

Second, as τ grows (darker to brighter curves) the optimal fraction shifts leftward, from $f^\star \approx 0.5$–0.7 at small τ to $f^\star \approx 0.2$–0.3 at τ=256. The physical interpretation is saturation: once t_A is large enough to suppress all non-condensed modes in the target layer (approximately $t_A \sim 1/\alpha$), additional sleep on that layer yields no further noise reduction, so the marginal return on target-layer investment drops to nearly zero. At this point, budget is best redirected to the helpers, whose cleaner representations amplify the heteroassociative coefficient u_l and provide a net lift to the target layer through the coupling. This shift is consistent with the effective-field decomposition (29): at saturation $\bar{D}_l \approx 1$ and $\sigma_l \approx 0$ for the target layer, so further increases in t_A have negligible effect on either the dreaming-noise coefficient or the signal coefficient v_l, and the dominant lever becomes the inter-layer term u_l, which benefits from helper dreaming.

Third, the success probability at the optimum ($P^\star \approx 0.8$–0.95 depending on τ) substantially exceeds both extremes f=0 and f=1, confirming that a cooperative sleep allocation strategy consistently and significantly outperforms naive alternatives across all budget sizes.

It is worth stating precisely what can, and what cannot, be proved about the interiority and the stability of this optimum across arbitrary total budgets τ. That the maximiser $f^\star(\tau)$ lies strictly inside $(0, 1)$ for every finite τ is structural rather than numerical: the two endpoints are controlled by the effective-field decomposition (29). At f=1 the helpers receive no sleep, their non-condensed modes are unattenuated and the relayed signal u_l they can supply is throttled by their residual crosstalk; at f=0 the target layer is left undreamed, its dreaming-noise coefficient σ_l stays at its undreamed value and the clean helper signal cannot overcome it. Since $P_{\text{success}}(f)$ is continuous on the compact interval $[0, 1]$ and both boundary values are strictly dominated, it necessarily attains its maximum in the interior; this much is rigorous and budget-independent. A fully closed-form characterisation of $f^\star(\tau)$ for arbitrary τ, by contrast, would require the joint retrieval-region functional in closed form, which – as noted for the LAM noise sector (Section 3) – is not available beyond the L=3 specialisation, so we do not claim an explicit analytic proof of global stability. What *does* explain the robustness of the interior optimum, and its smooth leftward drift with τ, is the cooperativeness intrinsic to multidirectional memories (Alessandrelli et al., 2025a). There, providing layers with entropy-heterogeneous datasets produces a redistribution effect: the retrieval regions do not simply track each layer's own dataset quality but settle, through the inter-layer couplings, to a common amplitude – the stronger layer cedes part of its retrieval region so that the weaker layers may enlarge theirs. Sleep-budget allocation is the dynamical, dual face of the same conservation law: dreaming time, like dataset quality, controls the size of a layer's retrieval basin, and the heteroassociative couplings couple these basins so that what one layer gains in effective capacity another tends to relinquish. The allocation problem is therefore essentially a *zero-sum redistribution* of retrieval-region area across layers under the fixed budget τ, and a balanced interior split is the natural fixed point: away from it, the cooperative coupling restores balance, because the over-dreamed layer hits the saturation $\bar{D}_l \to 1$, $\sigma_l \to 0$ (diminishing returns) and the marginal budget is automatically best spent on the layer whose basin is still expanding. As τ grows the target layer saturates first (at $t_A \sim 1/\alpha$), so the conserved "area" is increasingly cheaply bought on the helper side, which is exactly the observed leftward migration of $f^\star$.

5. Signal-to-noise picture and the optimal dreaming time

The replica-symmetric theory characterises the equilibrium phases; we close the analysis with a complementary one-step (signal-to-noise) probe in the noiseless limit $\beta \to +\infty$ (equally, $T \to 0^+$), which returns the learning threshold in closed form for arbitrary sleeping time and exposes, layer by layer, how consolidation reshapes retrieval. The construction, developed in full in Appendix F, prepares every layer on its own, never observed, archetype $\boldsymbol{\xi}^1$ and applies a single Hinton update. The decisive simplification is that the dreaming coupling, transported from the K-dimensional pattern space to the N-dimensional neuron space, becomes a resolvent of the empirical overlap matrix $G_l = \frac{1}{N} J_l J_l^\top$,

$$\mathcal{J}^{(l)} = (1 + t_l)\, G_l\, (\mathbb{I} + t_l G_l)^{-1}, \qquad \mathbb{M}_l := (\mathbb{I} + t_l G_l)^{-1}, \tag{35}$$

so that the field aligned with the archetype is controlled by three self-averaging spectral functionals of $\mathbb{M}_l$,

$$m_l = \tfrac{1}{N} \mathrm{Tr}\, \mathbb{M}_l, \quad q_l = \tfrac{1}{N} \boldsymbol{\xi}_l^{1\top} \mathbb{M}_l\, \boldsymbol{\xi}_l^1, \quad q_l^{(2)} = \tfrac{1}{N} \boldsymbol{\xi}_l^{1\top} \mathbb{M}_l^2\, \boldsymbol{\xi}_l^1. \tag{36}$$

From these the first moment of the internal field, the one-step magnetisation and the signal-to-noise ratio follow in closed form, the temperature cancelling in the ratio:

$$\mu_1^{(l)} = \beta\, \tfrac{1+t_l}{t_l}\, (m_l - q_l), \qquad m_l^{(2)} = \mathrm{erf}\left(\frac{m_l - q_l}{\sqrt{2\,(q_l^{(2)} - q_l^2)}} \right),$$

$$\mathrm{SNR}_l^2 = \frac{(m_l - q_l)^2}{q_l^{(2)} - q_l^2}. \tag{37}$$

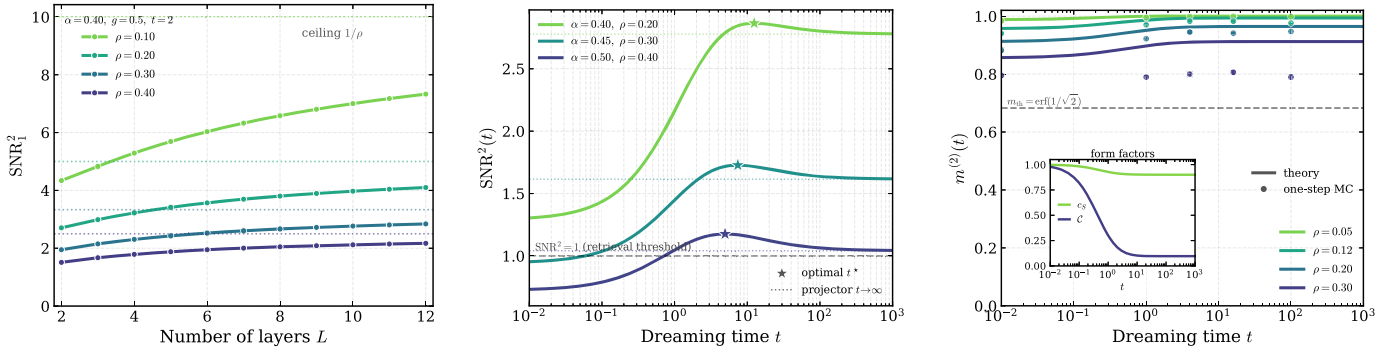
Under the Gaussian equivalence of the non-condensed modes, m_l is fixed by the Marchenko–Pastur self-consistency $t_l\, m_l(m_l + \alpha - 1) = 1 - m_l$, while the rank-one signal spike yields $q_l = m_l\,(1 + t_l \delta_l m_l)/(1 + t_l m_l)$ with sampling-noise fraction $\delta_l = \rho_l/(1 + \rho_l)$ and $q_l^{(2)} = \mathrm{d}(t_l q_l)/\mathrm{d}t_l$. The dreaming signal is then modulated by a single form factor, $\mu_1^{\text{Dream}} = \frac{\beta}{1+\rho_l}\, c_S(t_l)$ with $c_S(t_l) = (1 + t_l) m_l^2/(1 + t_l m_l)$, and the single-layer learning threshold ($\theta = 1/\sqrt{2}$) interpolates between two transparent endpoints,

$$\underbrace{\alpha(1+\rho_l)^2 + \rho_l < 1}_{t_l = 0 \text{ (Hebbian)}} \longrightarrow \underbrace{\alpha(1+\rho_l^2) + \rho_l < 1}_{t_l \to +\infty \text{ (projector)}}, \tag{38}$$

so that consolidation removes the linear example-interference penalty $2\alpha\rho_l$ from the load budget. These closed forms reproduce direct one-step Monte Carlo sampling across the whole range of t_l (Fig. 10c).

Two structural consequences follow. In the multi-layer architecture the heteroassociative channels supply a cooperative signal that grows with the number of coupled layers; the coherent finite-sample noise, which scales as $\rho_l (\mu_1^{(l)})^2$, however caps the gain, so that SNR_l^2 increases monotonically with L yet remains bounded above by $1/\rho_l$ (Fig. 10a). Adding modalities therefore never degrades retrieval, and the cooperative benefit saturates once the relayed signal dominates the residual sampling noise.

This same decomposition makes transparent why *disentanglement* is intrinsically a heteroassociative task and cannot be accomplished by dreaming alone. The aligning signal in Ψ_l splits additively into the autoassociative term v_l and the relayed heteroassociative term $u_l = \sum_{m \neq l} (\beta g_{lm}/\Gamma_{lm}) \hat{M}_m$ of Eq. (29). When the cue is a mixture that is uninformative for layer l, as at the simplex vertices of Section 2.6, the layer starts with $\hat{M}_l \approx 0$, so its autoassociative drive $v_l \propto \tilde{\eta}_l \propto \hat{M}_l$ is itself negligible: the one-step magnetisation $m_l^{(2)} = \mathrm{erf}\big((m_l - q_l)/\sqrt{2(q_l^{(2)} - q_l^2)}\big)$ stays near zero, and *no* amount of consolidation can raise it, because dreaming only reweights the eigenmodes of Σ_l and cannot manufacture an overlap that the cue does not contain. The only term that can lift such a layer out of the $m_l \approx 0$ band is u_l, which injects the overlap retrieved by the informed partner layers; setting $g_{lm} = 0$ (decoupled, purely autoassociative dreaming networks) removes this term identically and freezes the uninformed layers at zero magnetisation, so the joint success probability $\{m_l \geq m_{\text{th}}\ \forall l\}$ collapses to zero away from the trivial single-view vertices. Disentanglement is therefore realised only by the combination of the two mechanisms – dreaming to clean the intra-layer basins so the relayed signal is not reburied by crosstalk, and the heteroassociative couplings g_{lm} to route each modal component to its layer – which



(a) (b) (c)

Fig. 10. Signal-to-noise picture of the dreaming kernel. (a) One-step signal-to-noise ratio SNR_1^2 versus the number of coupled layers L (fixed per-bond coupling g, load α and sleeping time t), for several dataset entropies ρ. Inter-layer cooperation increases the signal-to-noise ratio monotonically with L, but it is bounded by the coherent sampling-noise ceiling $1/\rho$ (dotted lines): additional modalities never degrade retrieval and the cooperative gain saturates. **(b)** $\mathrm{SNR}^2(t)$ versus dreaming time for three (α, ρ) pairs. The ratio rises, attains a maximum at a finite optimal sleeping time $t^\star$ (stars), and then relaxes to the projector value (dotted): the over-consolidation effect. The horizontal dashed line is the retrieval threshold $\mathrm{SNR}^2 = 1$. **(c)** One-step magnetisation $m^{(2)}(t)$: closed-form theory (solid lines) against direct one-step Monte Carlo (symbols) for several ρ at $\alpha = 0.1$; the inset shows the signal form factor $c_S(t)$ and the intra-layer crosstalk form factor $C(t)$. All closed forms are derived in Appendix F.

is exactly the cooperative gain quantified in Fig. 10a and the reason the meaningful ablation is dreaming+LAM versus LAM-only.

The dependence on the sleeping time is more delicate, and it answers a question the numerics naturally raise. Increasing t_l suppresses the inter-pattern crosstalk, but the projector limit $\mathbb{A}_l \to \Sigma_l^{-1}$ inverts, and thereby amplifies, the finite-sample fluctuation of the empirical archetypes. The two effects balance at a finite *optimal sleeping time* $t_l^\star(\alpha, \rho_l)$: the signal-to-noise ratio rises, peaks at $t_l^\star$, and decreases beyond it (Fig. 10b) -the signal-to-noise face of the over-consolidation (overfitting) regime familiar from regularised, early-stopped learning (Agliari et al., 2024b). Why, then, do the Monte Carlo phase diagrams of Section 4 show retrieval improving monotonically with t and merely saturating? The reconciliation is quantitative (Fig. 11). The overshoot of SNR_l^2 over its projector value scales steeply with the load (from a fraction of a percent for $\alpha \lesssim 0.2$ to order unity only as $\alpha \to 1$) while the optimum itself sits at $t_l^\star \sim \mathcal{O}(1/\rho_l)$, which the clean datasets employed here ($\rho_l \lesssim 0.1$) push well beyond the explored window $t \lesssim 10^2$. The erf map compresses even a substantial overshoot of the ratio into a magnetisation bump of order 10^{-3}, below the Monte Carlo resolution; and the one-step probe, initialised exactly on the archetype, is intrinsically sharper than the full relaxation, which can re-enter the (shifted) retrieval basin. In the low-load, well-sampled regime of our simulations the optimum is thus at once out of range and below resolution, and dreaming behaves as a practically monotone benefit –precisely as observed– whereas the over-consolidation predicted here emerges, and is confirmed by full Monte Carlo, only at high load and degraded data (Fig. 11b).

6. Discussion

The preceding sections have established that the DLAM, a single cost function (Hamiltonian) combining off-line dreaming with heteroassociative coupling, produces a qualitative enlargement of the retrieval region, a measurable data-saving effect, and a robust improvement on the demanding task of disentanglement. We now place these findings in their broader context, interpret them from the three complementary viewpoints of statistical mechanics, machine learning, and biologically inspired computation, and state the limitations of the present analysis.

The mechanism: spectral filtering meets cross-modal support

The effective-field decomposition (28) provides a compact explanation of why the two mechanisms are synergistic rather than merely additive. Dreaming operates on the intra-layer sector: by differentially attenuating the high-eigenvalue modes of the empirical correlation matrix (which carry the inter-pattern interference), it reduces the dreaming-noise coefficient σ_l while preserving the autoassociative signal v_l. Inter-layer coupling, by contrast, introduces a second signal channel u_l that is independent of the quality of layer l's own data and depends only on the retrieval state of the neighbouring layers. The two improvements are complementary: one reduces noise, the other injects signal, and their joint effect exceeds the sum of their individual contributions, as confirmed by the synergistic retrieval regions documented in Fig. 4(b).

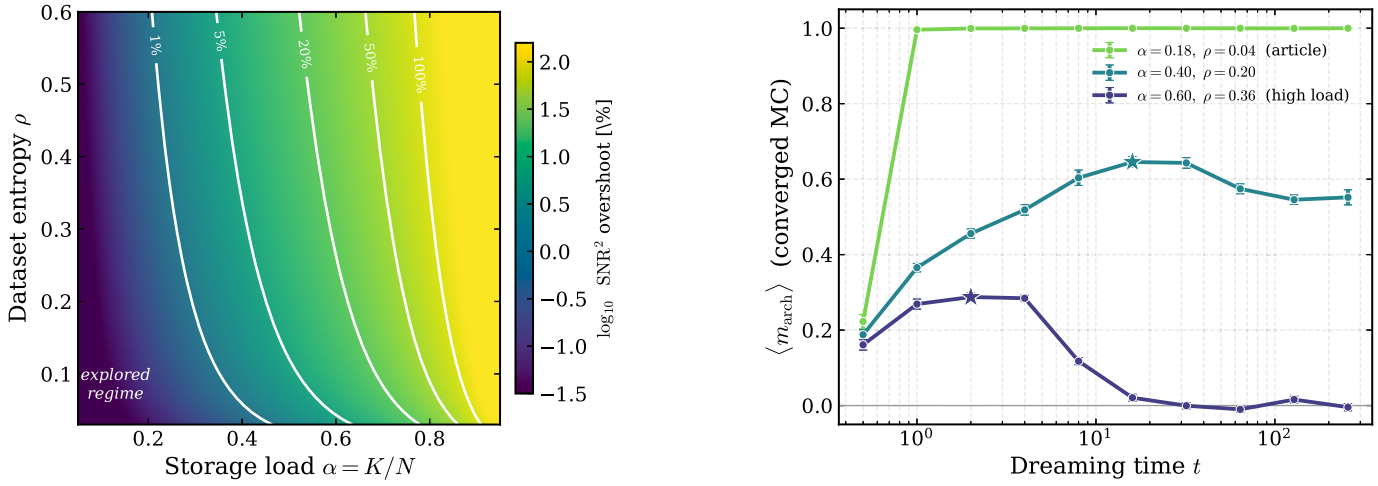
An important distinction separates the DLAM's dreaming mechanism from static decorrelation methods such as Kanter–Sompolinsky pseudo-inverse couplings (Kanter & Sompolinsky, 1987) or Gram–Schmidt orthogonalisation. In the single-layer limit the two approaches are equivalent, but in the multidirectional setting the equivalence breaks down. The cooperative effects arise because the dreaming kernel and the heteroassociative couplings share a common thermal dynamics: dreaming cleans the intra-layer representation, which in turn sharpens the inter-layer signal, which further stabilises the intra-layer state. This feedback loop is absent when each layer is decorrelated independently before coupling, and it is responsible for the re-entrant phase structure visible in the corresponding (α, T) phase diagrams.

A natural theoretical extension of this framework concerns the investigation of multi-condensed ansatzes. While the present work focuses on the thermodynamics of pure retrieval states, extending the replica-symmetric analysis to account for mixture states –where two or more macroscopic overlaps coexist– would allow us to explore the complete taxonomy of the model's metastable fixed points and properly quantify the interference across extensively loaded multi-layer architectures.

Data–computation trade-off and the role of disentanglement

The data–computation trade-off revealed by the (ρ, t) phase diagrams (Fig. 4c) has a direct practical reading from the machine-learning viewpoint: at fixed retrieval quality, off-line consolidation reduces the number of training examples required, so that computation substitutes for data acquisition. This phenomenon was first identified in single-layer dreaming networks (Agliari et al., 2024a,b); the present work shows that it persists, and is amplified, in the multi-layer architecture.

The effective-field decomposition (29) also provides a principled criterion for *when* dreaming is worth the computational investment: the benefit is largest when the dreaming-noise coefficient σ_l is the domi-



(a) (b)

Fig. 11. Why the optimal dreaming time is invisible at low load. (a) Relative overshoot of the signal-to-noise ratio over its projector value, $\max_t \mathrm{SNR}^2(t)/\mathrm{SNR}^2(\infty) - 1$, across the (α, ρ) plane (logarithmic colour scale; white contours at $1, 5, 20, 50, 100\%$). The over-consolidation maximum is negligible in the low-load, well-sampled corner explored by the simulations and becomes order-one only as the load approaches the projector capacity $\alpha \to 1$. **(b)** Converged Monte Carlo archetype overlap $\langle m_{\mathrm{arch}} \rangle$ versus dreaming time (averaged over 30 realisations, $N = 500$, $\beta = 10$, relaxation from a held-out test example). In the clean, low-load regime (article parameters) retrieval rises monotonically and saturates; at high load it peaks at the optimal $t^\star$ (stars) and then collapses, the macroscopic confirmation of the over-consolidation predicted by the signal-to-noise analysis.

nant source of interference and the signal v_l still has room to grow, that is, when the load is moderate and the data quality is low. In the opposite low-load, high-quality-data regime, the benefit of dreaming naturally saturates. This should not be interpreted as an additional phase boundary. Rather, the system is already deep inside the retrieval phase: the non-condensed interference is weak, the dreaming-noise coefficient σ_l is small, and the signal term v_l has little room to increase. Since the main role of dreaming is to suppress residual interference, further increasing t only produces a negligible marginal improvement. This reading is consistent with Fig. 4(c), where the retrieval boundary in the (ρ, t)-plane is shifted outward by dreaming, but has little operational relevance when ρ is already small. It is also reflected in Fig. 10, where the low-load, well-sampled regime displays a monotone rise followed by a plateau rather than a sharp saturation transition.

Disentanglement provides the most challenging stress test for this trade-off. The cross-modal mixture cue (13) blends independent modal samples of a single target archetype from all three layers; at the barycentre every layer receives a diluted but nonzero starting signal, allowing bootstrap through inter-layer cooperation even without dreaming. Toward the simplex vertices, however, one layer receives all the signal while the other two are essentially uninformed; joint convergence then requires the heteroassociative coupling to rescue the blind layers, which is impossible when the basins are contaminated by inter-pattern noise. With dreaming, the spectral filter suppresses precisely these off-diagonal correlations, cleaning the basins so that even a single-view cue can propagate through the coupling and achieve joint retrieval. The simplex visualisation (Fig. 2a, b) makes this transition vivid: the region of high joint success expands from the barycentre outward toward the vertices as dreaming is activated. We note that disentanglement is not merely a technical benchmark: it captures the computational problem of multi-source memory integration, namely fusing partial, channel-specific evidence into separate, coherent per-modality representations, which is directly relevant to episodic memory consolidation, and the setting where the interaction between dreaming and heteroassociation is most consequential.

Biological resonance

The biological reading of the DLAM is disciplined but substantive. The architecture captures two features of cortical memory organisation, namely distributed multimodal storage and off-line consolidation (Frankland & Bontempi, 2005; Klinzing et al., 2019), within a unified energy-based framework. The heteroassociative couplings model the associative links between cortical areas that support cross-modal pattern completion, while the dreaming kernel models the synaptic rescaling that occurs during sleep. The quantitative prediction that emerges, namely that sleep is most valuable when waking data is scarce, noisy, or ambiguous, resonates with experimental findings on the role of slow-wave sleep in memory consolidation under conditions of interference and weak encoding (Denis et al., 2021). To address directly the question of biological correspondence: the dreaming step maps most naturally onto systems-level consolidation during slow-wave sleep, in which hippocampal replay reinstates recent traces and drives an off-line, experience-independent reorganisation of cortical synapses (Ji & Wilson, 2007; Klinzing et al., 2019). In this reading the consolidation time t_l is the analogue of the sleep budget, and the spectral filtering by $\mathbb{A}_l(t_l)$ plays the role of the synaptic down-selection of parasitic configurations attributed to that stage, rather than of REM-style free association at frozen synapses. The disentanglement task of Section 4.1 provides a concrete instantiation: the cross-modal mixture cue (13) can be read as the kind of multi-source, partially encoded memory trace that sleep is hypothesised to separate into clearly distinct representations (Crick & Mitchison, 1983), an interpretation supported by evidence of coordinated replay across hippocampal and cortical areas during sleep (Ji & Wilson, 2007).

This framing is a principled analogy, not a claim of neural realism: the DLAM assumes binary neurons, symmetric couplings, and substantial simplifications of cortical circuitry, yet it captures the essential computational logic –off-line spectral noise suppression and Hebbian cross-modal stabilisation– with enough precision to yield quantitative phase-diagram predictions. Recent evidence that dreaming optimises memory even under strict, biologically realistic bounds on synaptic strengths (Marinari et al., 2026) supports the plausibility of this

paradigm; whether its specific spectral mechanism has a counterpart in biological synaptic homeostasis remains an open empirical question.

Connections to modern associative memories

The DLAM bears a structural resemblance to modern Hopfield networks (Ramsauer et al., 2020) and dense associative memories (Krotov & Hopfield, 2016), in that all three families manipulate the energy landscape to sharpen retrieval basins. The mechanism differs fundamentally: where modern Hopfield networks achieve sharper basins through higher-order (polynomial or continuous) energy functions, a pathway recently shown to profoundly expand associative capacities across both neurobiology and modern machine learning (Krotov & Hopfield, 2021), the DLAM achieves a comparable effect through off-line spectral filtering of a *quadratic* energy, followed by on-line thermal relaxation. The two approaches are not mutually exclusive: one could in principle apply dreaming-type regularisation to a dense associative memory, and a formal comparison of their capacity scaling and phase structure is an interesting direction for future work. In this respect, the nonlinear-PDE approach to the statistical mechanics of dense associative memories (Agliari et al., 2022) provides a complementary analytical framework for the same capacity and energy-landscape questions addressed here, and a natural bridge between the quadratic, dreaming-regularised energy of the DLAM and the higher-order energies of dense models.

Limitations

Several limitations should be stated explicitly. The replica-symmetric ansatz is expected to break down near the spin-glass transition, where full replica-symmetry breaking (RSB) may modify the critical thresholds quantitatively. Qualitatively, the RS description should therefore be treated with greater caution close to the glassy regime. In the (α, T) plane, the potentially unstable region is expected to be analogous to that below the de Almeida–Thouless (AT) stability line of the standard Hopfield model. Conversely, consistently with the rigorous results available for the dreaming Hopfield model, we expect the spin-glass region to be progressively suppressed as t increases and to disappear in the $t \to \infty$ limit. The extension of this stability analysis to the full (α, T, ρ, t) parameter space is currently under investigation and is left for future work. In particular, in the long-dreaming-time limit the (multidirectional extension of the) network is expected to approach the Kanter–Sompolinsky perfect-memory network (Kanter & Sompolinsky, 1987), which is replica-symmetric because orthogonalising the archetypes removes spin-glass frustration already at finite N; consistently, the single-layer dreaming Hopfield model is rigorously known to retain only the retrieval and ergodic phases, the spin-glass phase disappearing, as $t \to \infty$ (Agliari et al., 2019), an effect further elucidated by the Hopfield–RBM duality of Leonelli et al. (2021). The present model assumes i.i.d. Rademacher archetypes and symmetric layer sizes; correlated patterns, continuous-valued neurons, and heterogeneous layer dimensions are natural extensions that may introduce qualitatively new phenomena. The dreaming mechanism is parametrised by a single scalar t_l per layer; a dynamical model of the dreaming process itself, in which t_l evolves according to its own learning rule, would be both more biologically motivated and potentially richer in phenomenology. Finally, while the dedicated finite-size-scaling campaign of Section 4.3 reaches $N = 3200$ and confirms the thermodynamic-limit critical load, the remaining Monte Carlo experiments are run at moderate system sizes ($N \leq 500$) and at equilibrium or near-equilibrium dynamics; larger-scale studies of the full phase structure and of non-equilibrium protocols remain to be explored.

7. Conclusion

We have introduced the Dreaming L–directional Associative Memory, a multi-layer Hopfield architecture that unifies off-line Hebbian dreaming with supervised heteroassociative coupling within a single energy function. The replica-symmetric theory yields a closed system of self-consistency equations for the order parameters as the control parameters are made to vary. The resulting effective local field decomposes into signal, dreaming noise, and inter-layer noise, a structure that identifies, at the level of individual neurons, why dreaming improves retrieval: it differentially attenuates the high-eigenvalue interference modes of the empirical correlation matrix, suppressing inter-pattern crosstalk while preserving the signal direction. This intra-layer cleaning, in turn, amplifies the effectiveness of the inter-layer coupling.

The two mechanisms are synergistic. The phase diagrams reveal retrieval regions not reached by either dreaming or heteroassociation alone, and the numerics confirm this prediction across the explored parameter space. The most consequential manifestation is on the task of disentanglement, where the dreaming kernel destabilises the glassy attractor that traps the dynamics near ambiguous mixture cues and promotes sharp, well-separated retrieval basins, raising the joint success rate substantially. Importantly, this work provides the *first* treatment of the disentanglement problem with a fully general, heterogeneous cue-weight vector $(\alpha_1, \alpha_2, \alpha_3)$ on the 2-simplex, rather than restricted to the symmetric barycentric cue. This generalisation is both welcome and non-trivial: it is welcome because mixtures are rarely used with equal intensity among the base patterns–e.g. in audio and/or video processing – non-trivial because the asymmetry of the cue weights breaks the symmetry of the initial condition across layers, so that different layers start with vastly different overlaps with their respective targets. Establishing that the dreaming benefit extends throughout the entire simplex (not just near the symmetric barycentre but all the way to the vertices, where a single-view cue provides signal to only one layer while the other two start essentially uninformed) is a substantively stronger result than the homogeneous case, and is the main empirical finding of the disentanglement experiments.

The data–computation trade-off documented in the (ρ, t) phase diagrams (Fig. 4c) and the data-saving experiments (Fig. 3) deserves special emphasis. Off-line consolidation substitutes for additional training data in a quantitatively significant way: moderate dreaming consistently reduces the critical dataset size $M_c(t)$ by 40–60% across a broad range of data-quality conditions and pattern statistics, and this saving is robust to both the choice of evaluation metric and the structure of the patterns (random Rademacher or MNIST-derived). In the multimodal setting these savings are of acute practical value, because collecting large batches of clean, aligned data across multiple sensors is far more cost-prohibitive than allocating extra off-line computation. The extension of this data-saving phenomenon from single-layer to multi-layer architectures is therefore not a minor quantitative extension but a qualitative broadening of its applicability.

From a broader perspective, the DLAM demonstrates that a principled combination of off-line noise suppression and cross-modal binding can extend the operating envelope of associative memories into regimes of high load, poor data quality, and ambiguous cues that are inaccessible to conventional architectures. The effective-field decomposition provides not only a mechanistic explanation but a practical design tool: it identifies when dreaming is most beneficial and how a finite sleep budget should be allocated across layers of unequal data quality.

Several directions remain open: a full de Almeida–Thouless stability analysis to delineate the regime of validity of the replica-symmetric predictions; extensions to correlated, structured, and continuous-valued patterns; dynamical models of the dreaming process with their own evolution equations; and applications to real multimodal datasets where data acquisition is expensive but off-line computation is cheap.

Furthermore, while not addressed here, potential implementations in neuromorphic hardware represent a promising direction: modern

memristors could eventually support the intensive off-line computing demands of dreaming networks (see, e.g., Kamsma et al., 2024; Yu et al., 2025).

The central message of this work, namely that off-line computation can substitute for part of the burden of data in multimodal memory systems, suggests that the classical unlearning hypothesis of Crick and Mitchison (1983) has richer implications for learning architectures than have yet been fully explored.

CRediT authorship contribution statement

Adriano Barra: Writing – original draft, Supervision, Project administration, Funding acquisition, Conceptualization; **Fabrizio Durante:** Writing – original draft, Project administration, Funding acquisition, Conceptualization; **Andrea Ladiana:** Writing – original draft, Software, Investigation, Formal analysis, Data curation, Conceptualization; **Michela Marra Solazzo:** Writing – original draft, Investigation, Formal analysis.

Reproducibility

All code and simulation data accompanying this work are available at https://github.com/andrea-ladiana/dlam-mc-core. The repository also includes, as supplementary material, the symbolic (MATHEMATICA) routine `LAM_symbolic_arbitraryL.wl` that generates the LAM coupling determinant and the associated replica-symmetric self-consistency closure (the kernel $\mathbf{C}(H)$, $\det \mathbf{C}(H)$, the log-determinant gradient, the quadratic correction $\mathbb{D}_{\mathrm{LAM}}$, and the conjugate overlaps $\bar{\Pi}_l$) for an arbitrary number of layers L. The script is self-validating: on evaluation it reproduces every explicit $L = 3$ formula reported here, and it yields in closed form the homogeneous determinant $\det \mathbf{C}(H) = (1-(L-1)u)(1+u)^{L-1}$ with $u = \beta g(1-H)$ for general L, so that the equations used here for $L = 3$ can be extended to higher L on demand.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Fabrizio Durante reports financial support was provided by Ministry of Education and Merit. Michela Marra Solazzo reports financial support was provided by Ministry of Education and Merit. The submitting author (Prof. A. Barra) belongs to the Editorial Board Member of the Journal. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

F.D. and M.M.S. acknowledge support from the European Union–NextGenerationEU under the Italian National Recovery and Resilience Plan (PNRR), Mission 4, Component 2, Investment 1.3, via the Italian Ministry of University and Research (MUR), through the extended partnership "Future Artificial Intelligence Research (FAIR)" (Project Code PE00000013, CUP H97G22000210007, Spoke 6 "Symbiotic AI") and the specific cascade grant project "DELTA-CORE" (CUP F88H24002810007).

A.B. acknowledges support from Sapienza University of Rome, Prot. n. RM12519999AB8CA9, *Neural Networks and Learning Machines: asymptotic behaviors on structured datasets*.

A.B., A.L., and M.M.S. acknowledge financial support from the Italian National Institute for Higher Mathematics (INdAM) through the GNFM group. Additionally, M.M.S. acknowledges support from the National Institute for Nuclear Physics (INFN), Sezione di Lecce, and A.B. acknowledges support from INFN, Sezione di Roma 1.

A.B. also acknowledges the grant NEEDLE from Sapienza Università di Roma.

The authors acknowledge the use of the Lagrange Multi-GPU Server at the Department of Mathematics, Sapienza University of Rome, for computational resources supporting this work.

Appendix A. Derivation of the replica-symmetric free energy

This section constructs the replica-symmetric representation of the DLAM partition function. Starting from the Boltzmann–Gibbs measure associated with the full cost function (Hamiltonian) (7), we isolate the signal and noise contributions in each coupling sector, perform Hubbard–Stratonovich linearisations tailored to the symmetry of each quadratic form, and embed the result in a Guerra-type interpolating action. The fundamental theorem of calculus then decomposes the quenched pressure into a solvable one-body term and a streaming remainder.

A.1. Boltzmann–Gibbs measure, partition function and free energy

The model depends on the following control parameters: the inverse temperature $\beta = 1/T$; the vector of layer-wise sleep times $\boldsymbol{t} := (t_1, \dots, t_L)$; the intra-layer load ratios $\boldsymbol{\rho} := (\rho_1, \dots, \rho_L)$, which enter through $\Gamma_l := r_l^2(1+\rho_l)$; the inter-layer coupling matrix $\mathbf{G} := (g_{lm})_{1\le l<m\le L}$; and the pattern load $\alpha := (K-1)/N$. All thermodynamic functions depend on this full parameter set; for compactness, we suppress its explicit appearance in argument lists whenever no ambiguity arises.

For a fixed realisation of the quenched disorder $\mathcal{D}$, the Boltzmann–Gibbs measure at inverse temperature $\beta = 1/T$ takes the form

$$\mathbb{P}_{N,L}(\boldsymbol{s}) := \frac{\exp\big(-\beta\,\mathcal{H}_{\mathrm{DLAM}}(\boldsymbol{s})\big)}{\mathcal{Z}_{N,L}}, \tag{A.1}$$

where the normalisation

$$\mathcal{Z}_{N,L} := \sum_{\boldsymbol{s}\in\{-1,+1\}^{NL}} \exp\big(-\beta\,\mathcal{H}_{\mathrm{DLAM}}(\boldsymbol{s})\big) \tag{A.2}$$

defines the finite-volume partition function. The thermal average of an observable $F : \{-1,+1\}^{NL} \to \mathbb{R}$ under this measure is denoted by

$$\omega_N(F) := \frac{1}{\mathcal{Z}_{N,L}} \sum_{\boldsymbol{s}} F(\boldsymbol{s})\, e^{-\beta \mathcal{H}_{\mathrm{DLAM}}(\boldsymbol{s})}. \tag{A.3}$$

Because the disorder enters only through the dataset, it is natural to condition on the full collection of examples

$$\mathcal{D} := \Big\{\eta_{i,l}^{\mu,k}\Big\}_{\substack{i\le N,\, l\le L\\ \mu\le K,\, k\le M}} \quad \text{(equivalently, on } \{\hat{\xi}_{i,l}^{\mu}\} \text{ and } \{\Sigma_l\}\text{)}, \tag{A.4}$$

and to express the finite-N partition function as

$$\mathcal{Z}_{N,L}(\beta \mid \mathcal{D}) := \sum_{\boldsymbol{s}} \exp\big\{-\beta \mathcal{H}_{\mathrm{DLAM}}(\boldsymbol{s})\big\}. \tag{A.5}$$

Substituting the explicit forms (4)–(6), the Boltzmann weight splits cleanly into its inter-layer and intra-layer contributions:

$$\begin{aligned}\mathcal{Z}_{N,L}(\beta \mid \mathcal{D}) = \sum_{\boldsymbol{s}} \exp\Bigg\{ &\beta \sum_{1\le l<m\le L} \frac{g_{lm}}{N\Gamma_{lm}} \sum_{\mu=1}^{K} \sum_{i,j=1}^{N} \hat{\xi}_{i,l}^{\mu} \hat{\xi}_{j,m}^{\mu} s_{i,l} s_{j,m} \\ &+ \beta \sum_{l=1}^{L} \frac{1}{N\Gamma_l} \sum_{\mu,\nu=1}^{K} \sum_{i,j=1}^{N} \hat{\xi}_{i,l}^{\mu} (\mathbb{A}_l)_{\mu\nu} \hat{\xi}_{j,l}^{\nu} s_{i,l} s_{j,l} \Bigg\}.\end{aligned} \tag{A.6}$$

The Boltzmann weight in (A.6) is already a function of the empirical Mattis magnetisations alone: every microscopic sum over site indices is absorbed into the collective overlaps $\hat{m}_l^{\mu}$, yielding

$$\begin{aligned}\mathcal{Z}_{N,L}(\beta \mid \mathcal{D}) = \sum_{\boldsymbol{s}} \exp\Bigg\{ &\beta N \sum_{1\le l<m\le L} \frac{g_{lm}}{\Gamma_{lm}} \sum_{\mu=1}^{K} \hat{m}_l^{\mu} \hat{m}_m^{\mu} \\ &+ \beta N \sum_{l=1}^{L} \frac{1}{\Gamma_l} \sum_{\mu,\nu=1}^{K} \hat{m}_l^{\mu} (\mathbb{A}_l)_{\mu\nu} \hat{m}_l^{\nu} \Bigg\}.\end{aligned} \tag{A.7}$$

The identity is purely algebraic and requires no thermodynamic limit.

Letting $\mathbb{E}$ denote the average over all realisations of the quenched disorder, we define the intensive quenched pressure

$$\mathcal{A}_N(\beta) := \frac{1}{N}\,\mathbb{E}\big[\log \mathcal{Z}_{N,L}\big], \tag{A.8}$$

and, in the thermodynamic limit,

$$\mathcal{A}(\beta) := \lim_{N\to\infty} \mathcal{A}_N(\beta), \tag{A.9}$$

from which the quenched free-energy density follows as

$$f(\beta) := -\frac{1}{\beta}\,\mathcal{A}(\beta). \tag{A.10}$$

A.2. Partition function and Hubbard–Stratonovich decouplings

The partition function (A.7) contains both inter-layer and intra-layer quadratic forms in the spins; a tractable mean-field theory is obtained through Hubbard–Stratonovich (HS) linearisation of both sectors. The starting point is

$$\mathcal{Z}(\beta \mid D) = \sum_{\boldsymbol{s}} \exp\Big(-\beta \mathcal{H}_{\mathrm{LAM}}(\boldsymbol{s}) - \beta \mathcal{H}_{\mathrm{Dream}}(\boldsymbol{s})\Big), \tag{A.11}$$

with the LAM and dreaming Hamiltonians given by

$$\mathcal{H}_{\mathrm{LAM}}(\boldsymbol{s}) := -\sum_{1\le l<m\le L} \frac{g_{lm}}{N\Gamma_{lm}} \sum_{\mu=1}^{K}\sum_{i,j=1}^{N} \hat{\xi}^{\mu}_{i,l}\hat{\xi}^{\mu}_{j,m} s_{i,l}s_{j,m}, \tag{A.12}$$

$$\mathcal{H}_{\mathrm{Dream}}(\boldsymbol{s}) := -\sum_{l=1}^{L} \frac{1}{N\Gamma_l} \sum_{i,j=1}^{N}\sum_{\mu,\nu=1}^{K} \hat{\xi}^{\mu}_{i,l}[\mathbb{A}_l(t_l)]_{\mu\nu}\hat{\xi}^{\nu}_{j,l} s_{i,l}s_{j,l},$$

$$\mathbb{A}_l(t_l) = (1+t_l)(\mathbb{I}+t_l\Sigma_l)^{-1}, \tag{A.13}$$

so that the full Hamiltonian is

$$\mathcal{H}_{\mathrm{DLAM}}(\boldsymbol{s}) := \mathcal{H}_{\mathrm{Dream}}(\boldsymbol{s}) + \mathcal{H}_{\mathrm{LAM}}(\boldsymbol{s}), \tag{A.14}$$

with the normalisation constants

$$(\Sigma_l)_{\mu\nu} := \frac{1}{N\Gamma_l}\sum_{i=1}^{N} \hat{\xi}^{\mu}_{i,l}\hat{\xi}^{\nu}_{i,l}, \qquad \Gamma_l := r_l^2(1+\rho_l), \quad \Gamma_{lm} := \sqrt{\Gamma_l\Gamma_m}. \tag{A.15}$$

A.2.1. LAM sector

The inter-layer couplings are bilinear in spins belonging to different layers, and their HS decoupling becomes most transparent once the layer-wise normalisation is absorbed into the patterns. To this end we define the rescaled archetypes

$$J^{\mu}_{i,l} := \frac{\hat{\xi}^{\mu}_{i,l}}{\sqrt{\Gamma_l}} \qquad \Rightarrow \qquad \frac{1}{\Gamma_{lm}}\hat{\xi}^{\mu}_{i,l}\hat{\xi}^{\mu}_{j,m} = J^{\mu}_{i,l}J^{\mu}_{j,m}, \tag{A.16}$$

and the corresponding layer-resolved overlaps

$$p^{\mu}_l(\boldsymbol{s}) := \frac{1}{\sqrt{N}}\sum_{i=1}^{N} J^{\mu}_{i,l}s_{i,l}, \tag{A.17}$$

the LAM energy assumes the manifestly quadratic form

$$\beta\mathcal{H}_{\mathrm{LAM}}(\boldsymbol{s}) = -\beta \sum_{1\le l<m\le L} g_{lm}\sum_{\mu=1}^{K} p^{\mu}_l(\boldsymbol{s})p^{\mu}_m(\boldsymbol{s}). \tag{A.18}$$

Within the standard replica-symmetric condensation ansatz, a single class (conventionally $\mu = 1$) acquires a macroscopic overlap with the spin configuration and thus carries the signal, while the remaining $K-1$ non-condensed modes generate the noise. Isolating these two contributions as

$$\sum_{\mu=1}^{K} p^{\mu}_l p^{\mu}_m = p^1_l p^1_m + \sum_{\mu=2}^{K} p^{\mu}_l p^{\mu}_m, \tag{A.19}$$

one defines the signal and noise factors

$$\mathcal{S}_{\mathrm{LAM}}(\boldsymbol{s}) := \beta \sum_{l<m} g_{lm}\, p^1_l(\boldsymbol{s})p^1_m(\boldsymbol{s}),$$

$$\mathcal{N}_{\mathrm{LAM}}(\boldsymbol{s}) := \exp\Big(\beta\sum_{l<m} g_{lm}\sum_{\mu=2}^{K} p^{\mu}_l(\boldsymbol{s})p^{\mu}_m(\boldsymbol{s})\Big) = \exp\Big(\frac{\beta}{2}\sum_{\mu=2}^{K} \boldsymbol{p}^{\mu\top}\mathbf{G}_0\,\boldsymbol{p}^{\mu}\Big), \tag{A.20}$$

where the noise factor involves the overlap vectors $\boldsymbol{p}^{\mu} = (p^{\mu}_1, \dots, p^{\mu}_L)^{\top}$ associated with all non-condensed modes, mutually coupled through the symmetric off-diagonal connectivity matrix

$$\mathbf{G}_0 \in \mathbb{R}^{L\times L}, \qquad (\mathbf{G}_0)_{lm} := \begin{cases} g_{lm}, & l \neq m, \\ 0, & l = m, \end{cases} \qquad \Rightarrow$$

$$\sum_{l<m} g_{lm}p^{\mu}_l p^{\mu}_m = \frac{1}{2}\boldsymbol{p}^{\mu\top}\mathbf{G}_0\boldsymbol{p}^{\mu}. \tag{A.21}$$

Lemma 1 (Complex HS for heteroassociative couplings). *Let $L \ge 2$ and let $\mathbf{G}_0 \in \mathbb{R}^{L\times L}$ be the weighted off-diagonal matrix defined in (A.21). For each non-condensed index $\mu \ge 2$ introduce complex auxiliary fields $\boldsymbol{u}_\mu = (u^1_\mu, \dots, u^L_\mu) \in \mathbb{C}^L$ and their complex conjugates $\boldsymbol{u}^{\dagger}_\mu = (u^{1\dagger}_\mu, \dots, u^{L\dagger}_\mu) \in \mathbb{C}^L$. We endow $(\boldsymbol{u}_\mu, \boldsymbol{u}^{\dagger}_\mu)$ with the Gaussian measure*[10]

$$Duu^{\dagger} := \prod_{\mu=2}^{K}\Bigg[\prod_{\ell=1}^{L}\frac{du^{\ell}_\mu\, du^{\ell\dagger}_\mu}{2\pi}\Bigg]\exp\Bigg\{-\frac{1}{2}\sum_{\mu=2}^{K}\Big(\boldsymbol{u}^{\top}_\mu\boldsymbol{u}^{\dagger}_\mu - \boldsymbol{u}^{\dagger\top}_\mu\mathbf{G}_0\,\boldsymbol{u}_\mu\Big)\Bigg\}. \tag{A.22}$$

This measure is non-diagonal: the cross-covariance term $\boldsymbol{u}^{\dagger\top}_\mu\mathbf{G}_0\,\boldsymbol{u}_\mu$ correlates hidden variables attached to distinct visible layers, and it is precisely this correlation that reconstructs the original inter-layer Hebbian couplings upon integration.

For every fixed $\boldsymbol{p} \in \mathbb{R}^L$, the (selective) complex HS representation is chosen so that

$$\exp\Big(\frac{\beta}{2}\boldsymbol{p}^{\top}\mathbf{G}_0\,\boldsymbol{p}\Big) = \int Duu^{\dagger}\,\exp\Big(\sqrt{\beta}\,\boldsymbol{p}^{\top}\boldsymbol{u}\Big), \tag{A.23}$$

where the symbol $Duu^{\dagger}$ in (A.23) denotes the one-mode factor (single μ) of the measure (A.22), and the conjugate variables $\boldsymbol{u}^{\dagger}$ enter only through the non-diagonal Gaussian measure rather than through an additional linear source.

Applying the identity (A.23) mode by mode to each $\mu \ge 2$ yields the linearised representation of the noise factor,

$$\mathcal{N}_{\mathrm{LAM}}(\boldsymbol{s}) = \int Duu^{\dagger}\,\exp\Bigg[\sqrt{\beta}\sum_{\mu=2}^{K}\boldsymbol{p}^{\mu\top}\boldsymbol{u}_\mu\Bigg], \tag{A.24}$$

which, after expanding p^{μ}_l in terms of the individual spins, becomes manifestly linear in each spin variable:

$$\mathcal{N}_{\mathrm{LAM}}(\boldsymbol{s}) = \int Duu^{\dagger}\,\exp\Bigg[\sqrt{\frac{\beta}{N}}\sum_{\mu=2}^{K}\sum_{\ell=1}^{L}\sum_{i=1}^{N} J^{\mu}_{i,\ell}s_{i,\ell}\,u^{\ell}_\mu\Bigg]. \tag{A.25}$$

A.2.2. Dreaming sector

The intra-layer dreaming couplings are linearised by introducing, for each layer l, an array of N site-indexed Gaussian auxiliary fields $\{\phi_{i,l}\}_{i=1}^{N}$ distributed as

$$D\phi_l := \prod_{i=1}^{N}\frac{d\phi_{i,l}}{\sqrt{2\pi}}\exp\Big(-\frac{1}{2}\sum_{i=1}^{N}\phi^2_{i,l}\Big), \tag{A.26}$$

and by combining the spin and noise degrees of freedom into the complexified multi-spin

$$k_{i,l} := s_{i,l} + i\sqrt{\frac{t_l}{\beta(1+t_l)}}\,\phi_{i,l}, \tag{A.27}$$

whose imaginary component, weighted by the sleep time t_l, is calibrated so that Gaussian integration regenerates the dressed interaction $\mathbb{A}_l$. In

[10] By a standard abuse of notation, when a single non-condensed mode μ is considered (with no product over μ written explicitly), the same symbol $Duu^{\dagger}$ denotes the corresponding one-mode factor of (A.22).

this representation, the dreaming signal for the condensed class reduces to

$$\exp\Big(\beta S_{\text{Dream}}(\boldsymbol{s},\boldsymbol{\phi})\Big) := \prod_{l=1}^{L} \exp\left(\frac{\beta N(1+t_l)}{2\Gamma_l}\eta_l^2\right), \tag{A.28}$$

while the non-condensed modes are linearised through a second family of auxiliary fields $\{z_\mu^l\}_{\mu=2}^K$, drawn from a Gaussian with variance $(1+t_l)$:

$$Dz_l := \prod_{\mu=2}^{K} \frac{dz_\mu^l}{\sqrt{2\pi}} \exp\Big(-\frac{1}{2(1+t_l)}\sum_{\mu=2}^{K}(z_\mu^l)^2\Big). \tag{A.29}$$

The dreaming noise factor then takes the form

$$\mathcal{N}_{\text{Dream}}(\boldsymbol{s},\boldsymbol{\phi}) := \prod_{l=1}^{L}\int Dz_l \exp\left[\sqrt{\frac{\beta}{\Gamma_l N}}\sum_{\mu=2}^{K}\sum_{i=1}^{N}\hat{\xi}_{i,l}^{\mu} z_\mu^l k_{i,l}\right]. \tag{A.30}$$

A.2.3. Full HS-decoupled representation

Combining the LAM signal–noise decomposition (A.20), the complex HS representation of the inter-layer noise (A.24)–(A.25), and the dreaming-native formulation of the intra-layer noise (A.28)–(A.30), we arrive at the exact HS-decoupled partition function

$$\begin{aligned}\mathcal{Z}(\beta \mid D) = \sum_{\boldsymbol{s}} \exp\big(\mathcal{S}_{\text{LAM}}(\boldsymbol{s})\big)\left[\prod_{l=1}^{L}\int D\phi_l\right] \\ \times \exp\big(\beta \mathcal{S}_{\text{Dream}}(\boldsymbol{s},\boldsymbol{\phi})\big)\mathcal{N}_{\text{Dream}}(\boldsymbol{s},\boldsymbol{\phi})\mathcal{N}_{\text{LAM}}(\boldsymbol{s}).\end{aligned} \tag{A.31}$$

The representation is exact: integrating out the complex auxiliary fields $\boldsymbol{u}_\mu, \boldsymbol{u}_\mu^\dagger$ in (A.24) regenerates the inter-layer quadratic form $\exp\big(\frac{\beta}{2}\boldsymbol{p}^{\mu\top}\mathbf{G}_0\boldsymbol{p}^\mu\big) = \exp\big(\beta\sum_{l<m} g_{lm}p_l^\mu p_m^\mu\big)$ for each non-condensed mode $\mu \geq 2$, while integrating out z_μ^l and $\phi_{i,l}$ in (A.30) reconstructs the dressed dreaming kernel $\mathbb{A}_l(t_l)$ of (4).

A.3. Guerra interpolation for the coupled LAM+Dreaming model

The quenched pressure is extracted through a Guerra-type interpolation, which continuously deforms the fully coupled system into a solvable reference model while tracking the free-energy cost at each step. A preliminary ingredient is the Gaussian approximation for the non-condensed modes $\mu \geq 2$, in which the rescaled empirical archetypes are replaced by i.i.d. standard Gaussians,

$$J_{i,l}^{\mu} = \frac{\hat{\xi}_{i,l}^{\mu}}{\sqrt{\Gamma_l}} \approx \lambda_{i,l}^{\mu}, \qquad \lambda_{i,l}^{\mu} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1), \qquad \mu \geq 2, \tag{A.32}$$

with the family $\{\lambda_{i,l}^{\mu}\}_{i,l,\mu\geq 2}$ taken independent of the condensed class and of every other random variable in play. By the central limit theorem this replacement is asymptotically exact for the non-condensed sector in the thermodynamic limit; its effect is to turn both the LAM noise factor (A.25) and the dreaming noise factor (A.30) into linear couplings of the spins ($s_{i,l}$ and $k_{i,l}$) to i.i.d. Gaussian random fields. In the dreaming sector the coupling reads $\sqrt{\beta/N}\,\tilde{\lambda}_{i,l}^{\mu} z_\mu^l k_{i,l}$, where $\{\tilde{\lambda}_{i,l}^{\mu}\}_{i,l,\mu\geq 2}$ is an independent copy of $\{\lambda_{i,l}^{\mu}\}$, kept distinct so as to preserve the statistical independence of the LAM and dreaming Gaussian sectors.

The interpolating partition function $\mathcal{Z}_{N,L}(x)$ depends on an auxiliary parameter $x \in [0,1]$, and the associated quenched pressure

$$\mathcal{A}_{N,L}(x) := \frac{1}{N}\,\mathbb{E}\log\mathcal{Z}_{N,L}(x), \qquad x \in [0,1], \tag{A.33}$$

in which $\mathbb{E}$ now denotes the joint average over the quenched dataset and all auxiliary Gaussian fields. The fundamental theorem of calculus then decomposes the physical pressure (at $x=1$) into a solvable boundary value and a streaming integral,

$$\mathcal{A}_{N,L}(1) = \mathcal{A}_{N,L}(0) + \int_0^1 dx\,\frac{d}{dx}\mathcal{A}_{N,L}(x), \tag{A.34}$$

where the derivative $\frac{d}{dx}\mathcal{A}$ is evaluated by Gaussian integration by parts via the Nishimori–Guerra identity.

The interpolating action is assembled from several independent families of i.i.d. $\mathcal{N}(0,1)$ auxiliary variables: site-indexed fields $\{\varphi_{i,l}\}_{i,l}$ and $\{\iota_{i,l}\}_{i,l}$, pattern-indexed fields $\{\theta_\mu^l\}_{l,\mu\geq 2}$, and mode-indexed fields $\{\delta_\mu^l\}_{\mu\geq 2,l}$ and $\{\tilde{\delta}_\mu^l\}_{\mu\geq 2,l}$. All of these are mutually independent and independent of the dataset, of $\{\lambda_{i,l}^{\mu}\}$, and of $\{\tilde{\lambda}_{i,l}^{\mu}\}_{i,l,\mu\geq 2}$–the latter being the independent copy of $\lambda_{i,l}^{\mu}$ employed exclusively in the dreaming sector. Paired with these random fields are the variational parameters $\{\psi_l^{(m)}\}_{l,m=1}^L$ controlling the LAM signal sector, the layer-specific constants $B_l, C_l, D_l, E_l, F_l, I_l \in \mathbb{R}$ governing the dreaming sector, and $C_u, C_{u^\dagger}, F_u, F_{u^\dagger} \in \mathbb{R}$ for the LAM auxiliary-field sector. Expressed through $k_{i,l}$ and η_l as in (18) and in the main text, the full interpolating action reads

$$\begin{aligned}S_x :=& x\,\beta \sum_{1\leq l<m\leq L} g_{lm}\,p_l^1(\boldsymbol{s})p_m^1(\boldsymbol{s}) + (1-x)\,N \\ &\times \sum_{1\leq l<m\leq L}\Big(\psi_l^{(m)}\hat{m}_l^1(\boldsymbol{s}) + \psi_m^{(l)}\hat{m}_m^1(\boldsymbol{s})\Big) \\ &+ \sqrt{\frac{x\beta}{2N}}\sum_{\mu=2}^{K}\sum_{l=1}^{L}\sum_{i=1}^{N}\lambda_{i,l}^{\mu}s_{i,l}\,u_\mu^l + \sqrt{1-x}\sum_{l=1}^{L}I_l\sum_{i=1}^{N}\iota_{i,l}s_{i,l} \\ &+ (1-x)\sum_{\mu=2}^{K}\sum_{l=1}^{L}\Big[C_u(u_\mu^l)^2 + C_{u^\dagger}({u_\mu^l}^\dagger)^2\Big] \\ &+ \sqrt{1-x}\sum_{\mu=2}^{K}\sum_{l=1}^{L}\Big[F_u\,\delta_\mu^l u_\mu^l + F_{u^\dagger}\,\tilde{\delta}_\mu^l\,{u_\mu^l}^\dagger\Big] \\ &+ \sum_{l=1}^{L}\Bigg[x\,\frac{\beta N(1+t_l)}{2\Gamma_l}\eta_l^2 - \frac{B_l(1-x)}{2(1+t_l)}\sum_{\mu=2}^{K}(z_\mu^l)^2 + \sqrt{\frac{x\beta}{N}}\sum_{\mu=2}^{K}\sum_{i=1}^{N}\tilde{\lambda}_{i,l}^{\mu}z_\mu^l k_{i,l} \\ &\quad + (1-x)C_l\sum_{i=1}^{N}k_{i,l}^2 + D_l\sqrt{1-x}\sum_{i=1}^{N}\varphi_{i,l}k_{i,l} \\ &+ (1-x)E_l N\eta_l + F_l\sqrt{1-x}\sum_{\mu=2}^{K}\theta_\mu^l z_\mu^l\Bigg].\end{aligned} \tag{A.35}$$

At $x=1$ all trial contributions vanish and the partition function collapses to the physical HS-decoupled form (A.31) under the Gaussian replacement (A.32). At $x=0$, conversely, the Gaussian couplings to the quenched patterns disappear and the remaining action splits into independent site- and mode-factors, rendering the system one-body and exactly soluble. The boundary value at $x=0$ thus supplies the Cauchy datum for the sum rule (A.34).

Appendix B. One-body calculations at $x=0$

At the endpoint $x=0$ of the Guerra interpolation, the couplings mediated by the quenched disorder switch off entirely and the partition function reduces to a product of elementary Gaussian and spin-trace factors–the solvable Cauchy datum for the sum rule (A.34). Evaluating the interpolating action (A.35) at $x=0$ yields

$$\mathcal{Z}_{N,L}(0) := \sum_{\boldsymbol{s}}\left[\prod_{l=1}^{L}\int D\phi_l\right]\left[\prod_{l=1}^{L}\int Dz_l\right]\int Duu^\dagger\,\exp\Big(S_0\Big), \tag{B.1}$$

$$\begin{aligned}S_0 =& N\sum_{1\leq l<m\leq L}\Big(\psi_l^{(m)}\hat{m}_l^1(\boldsymbol{s}) + \psi_m^{(l)}\hat{m}_m^1(\boldsymbol{s})\Big) + \sum_{l=1}^{L}I_l\sum_{i=1}^{N}\iota_{i,l}\,s_{i,l} \\ &+ \sum_{\mu=2}^{K}\sum_{l=1}^{L}\Big[C_u(u_\mu^l)^2 + C_{u^\dagger}({u_\mu^l}^\dagger)^2 + F_u\delta_\mu^l u_\mu^l + F_{u^\dagger}\tilde{\delta}_\mu^l{u_\mu^l}^\dagger\Big] \\ &+ \sum_{l=1}^{L}\Bigg[-\frac{B_l}{2(1+t_l)}\sum_{\mu=2}^{K}(z_\mu^l)^2 + C_l\sum_{i=1}^{N}k_{i,l}^2 \\ &+ D_l\sum_{i=1}^{N}\varphi_{i,l}\,k_{i,l} + E_l N\eta_l + F_l\sum_{\mu=2}^{K}\theta_\mu^l\,z_\mu^l\Bigg].\end{aligned} \tag{B.2}$$

The complexified multi-spin $k_{i,l}$, the empirical overlap η_l, and the Mattis magnetisation $\hat{m}^1_l(s)$ retain their definitions from the interacting theory:

$$k_{i,l} = s_{i,l} + i\sqrt{\frac{t_l}{\beta(1+t_l)}}\phi_{i,l}, \qquad \eta_l = \frac{1}{NM}\sum_{i=1}^{N}\sum_{k=1}^{M}\eta^{1,k}_{i,l}\,k_{i,l},$$

$$\hat{m}^1_l(\boldsymbol{s}) = \frac{1}{N}\sum_{i=1}^{N}\hat{\xi}^1_{i,l}s_{i,l}.$$

At $x = 0$ the three families of auxiliary variables–the complex LAM fields $(u, u^\dagger)$, the dreaming-noise fields z, and the spin–site-noise pair (s, ϕ)–couple only to their own independent Gaussian sources. The partition function therefore factorises:

$$\mathcal{Z}_{N,L}(0) = \mathcal{Z}^{(0)}_{\text{LAM}}\,\mathcal{Z}^{(0)}_{\text{Dream}}\,\mathcal{Z}^{(0)}_{s-\phi} \tag{B.3}$$

$$\mathcal{Z}^{(0)}_{\text{LAM}} := \int Duu^\dagger\,\exp\Bigg\{\sum_{\mu=2}^{K}\sum_{l=1}^{L}\Big[C_u(u^l_\mu)^2 + C_{u^\dagger}(u^{l\,\dagger}_\mu)^2 + F_u\delta^l_\mu u^l_\mu + F_{u^\dagger}\bar{\delta}^l_\mu u^{l\,\dagger}_\mu\Big]\Bigg\}. \tag{B.4}$$

The dreaming-noise sector gathers the pattern-indexed Gaussian fields z^l_μ, each carrying the layer-specific variance $(1 + t_l)$:

$$\mathcal{Z}^{(0)}_{\text{Dream}} := \prod_{l=1}^{L}\int Dz_l\,\exp\Bigg\{-\frac{B_l}{2(1+t_l)}\sum_{\mu=2}^{K}(z^l_\mu)^2 + F_l\sum_{\mu=2}^{K}\theta^l_\mu\,z^l_\mu\Bigg\}. \tag{B.5}$$

The local (s, ϕ) sector collects the spin traces together with the site-indexed Gaussian fields that encode the dreaming-kernel dressing:

$$\mathcal{Z}^{(0)}_{s-\phi} := \sum_{\boldsymbol{s}}\Bigg[\prod_{l=1}^{L}\int D\phi_l\Bigg]\exp\Bigg\{N\sum_{1\le l<m\le L}\Big(\psi^{(m)}_l\hat{m}^1_l + \psi^{(l)}_m\hat{m}^1_m\Big) + \sum_{l=1}^{L}\Big[I_l\sum_{i=1}^{N}\iota_{i,l}s_{i,l} + C_l\sum_{i=1}^{N}k^2_{i,l} + D_l\sum_{i=1}^{N}\varphi_{i,l}k_{i,l} + E_lN\eta_l\Big]\Bigg\}. \tag{B.6}$$

Because the logarithm maps products to sums, the quenched pressure at $x = 0$ decomposes additively:

$$\mathcal{A}_{N,L}(0) = \frac{1}{N}\mathbb{E}\log\mathcal{Z}^{(0)}_{\text{LAM}} + \frac{1}{N}\mathbb{E}\log\mathcal{Z}^{(0)}_{\text{Dream}} + \frac{1}{N}\mathbb{E}\log\mathcal{Z}^{(0)}_{s-\phi}. \tag{B.7}$$

B.1. Dreaming-noise and local sectors at $x = 0$

The two sectors evaluated below–the dreaming-noise Gaussian partition function $\mathcal{Z}^{\text{Dream-}z}$ and the local (s, ϕ) factor $\mathcal{Z}^{s,\phi}$–share the functional structure of their counterparts in the classical Hebbian dreaming theory. External signal or clamping fields could be reinstated by including the appropriate linear Mattis sources; they are set to zero here, in line with the protocol adopted throughout.

B.1.1. Dreaming-noise Gaussian sector $\mathcal{Z}^{\text{Dream-}z}$

The Gaussian measure governing each layer l is

$$Dz_l := \prod_{\mu=2}^{K}\frac{dz^l_\mu}{\sqrt{2\pi}}\,\exp\Big(-\frac{1}{2(1+t_l)}\sum_{\mu=2}^{K}(z^l_\mu)^2\Big), \tag{B.8}$$

with $\{\theta^l_\mu\}$ i.i.d. standard normals. At $x = 0$ the dreaming-noise partition function reduces to

$$\mathcal{Z}^{\text{Dream-}z}(0) := \prod_{l=1}^{L}\int Dz_l\,\exp\Bigg\{-\frac{B_l}{2(1+t_l)}\sum_{\mu=2}^{K}(z^l_\mu)^2 + F_l\sum_{\mu=2}^{K}\theta^l_\mu\,z^l_\mu\Bigg\}. \tag{B.9}$$

Because the integrand factorises over the mode index $\mu \ge 2$, the entire calculation reduces to a product of independent one-dimensional Gaussian integrals. Combining the quadratic pieces in (B.8) and (B.9), the effective single-mode exponent reads

$$-\frac{1}{2(1+t_l)}z^2 - \frac{B_l}{2(1+t_l)}z^2 + F_l\theta^l_\mu z = -\frac{1+B_l}{2(1+t_l)}z^2 + F_l\theta^l_\mu z. \tag{B.10}$$

Completing the square and performing the resulting Gaussian integral yields

$$\int\frac{dz}{\sqrt{2\pi}}\exp\Big(-\frac{1+B_l}{2(1+t_l)}z^2 + F_l\theta\,z\Big) = \sqrt{\frac{1+t_l}{1+B_l}}\,\exp\Big(\frac{(1+t_l)F_l^2}{2(1+B_l)}\,\theta^2\Big). \tag{B.11}$$

Taking the product over $\mu = 2, \dots, K$ delivers the closed-form expression

$$\mathcal{Z}^{\text{Dream-}z}(0) = \prod_{l=1}^{L}\Big(\sqrt{\frac{1+t_l}{1+B_l}}\Big)^{K-1}\exp\Bigg(\frac{(1+t_l)F_l^2}{2(1+B_l)}\sum_{\mu=2}^{K}(\theta^l_\mu)^2\Bigg). \tag{B.12}$$

Averaging over θ (using $\mathbb{E}[(\theta^l_\mu)^2] = 1$) and invoking the thermodynamic scaling $(K-1)/N \to \alpha$, we extract the quenched-pressure contribution

$$\frac{1}{N}\,\mathbb{E}\log\mathcal{Z}^{\text{Dream-}z}(0) = \sum_{l=1}^{L}\Bigg[\frac{\alpha}{2}\log\Big(\frac{1+t_l}{1+B_l}\Big) + \alpha\,\frac{(1+t_l)F_l^2}{2(1+B_l)}\Bigg]. \tag{B.13}$$

B.1.2. Local (s, ϕ) sector $\mathcal{Z}^{s,\phi}$

The last one-body factor collects the spin sums and the site-indexed Gaussian fluctuations. Writing $D\phi_l$ for the standard Gaussian measure over the analogue variables $\{\phi_{i,l}\}_{i=1}^N$,

$$D\phi_l := \prod_{i=1}^{N}\frac{d\phi_{i,l}}{\sqrt{2\pi}}\,e^{-\frac{1}{2}\phi^2_{i,l}} \tag{B.14}$$

and recalling the complexified local fields from the dreaming notation,

$$k_{i,l} := s_{i,l} + i\,a_l\,\phi_{i,l}, \qquad a_l := \sqrt{\frac{t_l}{\beta(1+t_l)}} \tag{B.15}$$

the (s, ϕ) partition function at $x = 0$ takes the form

$$\mathcal{Z}^{s,\phi}(0) := \sum_{\boldsymbol{s}}\prod_{l=1}^{L}\int D\phi_l\,\exp\Bigg\{N\sum_{1\le l<m\le L}\Big(\psi^{(m)}_l\hat{m}^1_l(\boldsymbol{s}) + \psi^{(l)}_m\hat{m}^1_m(\boldsymbol{s})\Big) + \sum_{l=1}^{L}\sum_{i=1}^{N}\Big[I_l\,\iota_{i,l}\,s_{i,l} + C_lk^2_{i,l} + D_l\,\varphi_{i,l}\,k_{i,l} + E_l\,\hat{\xi}^1_{i,l}\,k_{i,l}\Big]\Bigg\}. \tag{B.16}$$

The inter-layer Mattis coupling can be absorbed into a site-local effective field by inserting $\hat{m}^1_l(\boldsymbol{s}) = \frac{1}{N}\sum_{i=1}^N\hat{\xi}^1_{i,l}s_{i,l}$:

$$N\sum_{1\le l<m\le L}\Big(\psi^{(m)}_l\hat{m}^1_l + \psi^{(l)}_m\hat{m}^1_m\Big) = \sum_{l=1}^{L}\sum_{i=1}^{N}\Big(\hat{\xi}^1_{i,l}\,h_l\Big)s_{i,l}, \qquad h_l := \sum_{m\ne l}\psi^{(m)}_l. \tag{B.17}$$

Once this reduction is performed, the entire exponent becomes a sum of independent site-indexed terms and the partition function factorises:

$$\mathcal{Z}^{s,\phi}(0) = \prod_{l=1}^{L}\prod_{i=1}^{N}\mathcal{Z}_{i,l}, \tag{B.18}$$

where each single-site factor has the form

$$\mathcal{Z}_{i,l} := \sum_{s=\pm1}\int\frac{d\phi}{\sqrt{2\pi}}e^{-\phi^2/2}\exp\Big\{C_lk^2 + k\,\Omega_{i,l} + s\,h^{\text{ext}}_{i,l}\Big\}, \tag{B.19}$$

with the shorthand

$$\Omega_{i,l} := D_l\,\varphi_{i,l} + E_l\,\hat{\xi}^1_{i,l}, \qquad h^{\text{ext}}_{i,l} := I_l\,\iota_{i,l} + \hat{\xi}^1_{i,l}\,h_l, \qquad k = s + ia_l\phi. \tag{B.20}$$

The ϕ-integral in (B.19) is Gaussian and can be carried out in closed form. Expanding $k^2 = (s + ia_l\phi)^2 = 1 - a_l^2\phi^2 + 2ia_l s\phi$, we collect the exponent into terms quadratic, linear, and independent of ϕ:

$$-\frac{1}{2}\phi^2 + C_l(1 - a_l^2\phi^2 + 2ia_l s\phi) + (s + ia_l\phi)\Omega_{i,l} + sh^{\text{ext}}_{i,l}. \tag{B.21}$$

The quadratic, linear, and ϕ-independent contributions read

$$-\frac{1}{2}\Big(1 + 2C_la_l^2\Big)\phi^2 + ia_l\Big(2C_ls + \Omega_{i,l}\Big)\phi + \Big(C_l + s(\Omega_{i,l} + h^{\text{ext}}_{i,l})\Big). \tag{B.22}$$

Introducing the effective precision and its inverse,

$$A_l^{\rm eff} := 1 + 2C_l a_l^2, \qquad \Delta_l := (A_l^{\rm eff})^{-1} = \frac{1}{1+2C_l a_l^2}, \tag{B.23}$$

and invoking the standard Gaussian identity

$$\int \frac{d\phi}{\sqrt{2\pi}} \exp\Big(-\tfrac{1}{2}A\phi^2 + J\phi\Big) = \frac{1}{\sqrt{A}} \exp\Big(\frac{J^2}{2A}\Big), \qquad (\Re A > 0), \tag{B.24}$$

with $A = A_l^{\rm eff}$ and $J = ia_l(2C_l s + \Omega_{i,l})$, the single-site partition function reduces to

$$\mathcal{Z}_{i,l} = \sum_{s=\pm1} \sqrt{\Delta_l}\ \exp\left\{ C_l + s(\Omega_{i,l} + h_{i,l}^{\rm ext}) - \frac{\Delta_l a_l^2}{2}\big(2C_l s + \Omega_{i,l}\big)^2 \right\}. \tag{B.25}$$

Expanding the squared term and recalling $s^2 = 1$ for Ising spins,

$$\begin{aligned} -\frac{\Delta_l a_l^2}{2}(2C_l s + \Omega)^2 &= -\frac{\Delta_l a_l^2}{2}\Big(4C_l^2 s^2 + \Omega^2 + 4C_l \Omega\, s\Big) \\ &= -\frac{\Delta_l a_l^2}{2}\Big(4C_l^2 + \Omega^2 + 4C_l \Omega\, s\Big), \end{aligned} \tag{B.26}$$

the effective coefficient of s in the exponent of (B.25) simplifies via the identity $1 - 2\Delta_l a_l^2 C_l = \Delta_l$, which follows directly from (B.23):

$$\Omega_{i,l} - 2\Delta_l a_l^2 C_l\, \Omega_{i,l} = \Omega_{i,l}\Big(1 - 2\Delta_l a_l^2 C_l\Big) = \Delta_l\, \Omega_{i,l}. \tag{B.27}$$

Carrying out the remaining binary spin trace $\sum_{s=\pm1}$ yields the compact closed form

$$\mathcal{Z}_{i,l} = 2\sqrt{\Delta_l}\ \exp\left\{ C_l - \frac{\Delta_l a_l^2}{2}\Big(4C_l^2 + \Omega_{i,l}^2\Big) \right\} \cosh\Big(h_{i,l}^{\rm ext} + \Delta_l \Omega_{i,l}\Big). \tag{B.28}$$

By site exchangeability, the quenched pressure of this sector follows from (B.18):

$$\frac{1}{N}\,\mathbb{E}\log \mathcal{Z}^{s,\phi}(0) = \sum_{l=1}^{L} \mathbb{E}\log \mathcal{Z}_{1,l} \tag{B.29}$$

Spelling this out gives

$$\begin{aligned} \frac{1}{N}\,\mathbb{E}\log \mathcal{Z}^{s,\phi}(0) = \sum_{l=1}^{L} \Bigg[&\log 2 + \frac{1}{2}\log \Delta_l + C_l - \frac{\Delta_l a_l^2}{2}\Big(4C_l^2 + \mathbb{E}[\Omega_{1,l}^2]\Big) \\ &+ \mathbb{E}\log\cosh\Big(h_{1,l}^{\rm ext} + \Delta_l \Omega_{1,l}\Big)\Bigg]. \end{aligned} \tag{B.30}$$

Here the expectation $\mathbb{E}$ runs over the auxiliary Gaussians (ι, φ) and, through $\hat{\xi}^1$, over the dataset.

B.2. LAM sector at $x = 0$

The third and final one-body factor at $x = 0$ is the LAM Gaussian sector, which governs the non-condensed complex auxiliary fields. Decomposing each complex field into real and imaginary parts reduces the problem to a $2L$-dimensional real Gaussian integral.

For each non-condensed mode $\mu \geq 2$, the complex HS fields $u_\mu^l \in \mathbb{C}$ $(l = 1, \dots, L)$ are stacked into the real vector

$$\boldsymbol{x}_\mu := \big(\Re u_\mu^1, \Im u_\mu^1, \dots, \Re u_\mu^L, \Im u_\mu^L\big)^\top \in \mathbb{R}^{2L}. \tag{B.31}$$

At the replica-symmetric level the Gaussian sector is parametrised by the overlaps $H_l := \overline{\langle h_{12}^{(l)} \rangle_x}$. The covariance is the $2L \times 2L$ matrix $\mathbf{C}(H)$ with 2×2 diagonal blocks

$$\mathbf{C}(H) = \big(\mathbf{C}^{lm}\big)_{l,m=1}^{L}, \qquad \mathbf{C}^{ll} = \begin{pmatrix} 1 - \frac{\beta}{2}(1 - H_l) & 0 \\ 0 & 1 + \frac{\beta}{2}(1 - H_l) \end{pmatrix}, \tag{B.32}$$

and off-diagonal $(l \neq m)$ blocks encoding the inter-layer coupling:

$$\mathbf{C}^{lm} = g_{lm}\begin{pmatrix} 1 & -i \\ -i & -1 \end{pmatrix}, \qquad g_{lm} = g_{ml}, \qquad g_{ll} = 0. \tag{B.33}$$

As dictated by the heteroassociative architecture, each visible layer couples to a single Gaussian degree of freedom per conjugate pair rather than to both. This constraint is enforced by the i.i.d. $\delta_\mu^l \sim \mathcal{N}(0,1)$ together with the $2L$-source vector

$$\boldsymbol{b}_\mu := \sqrt{\frac{\beta}{2}}\big(\sqrt{H_1}\,\delta_\mu^1,\ i\sqrt{H_1}\,\delta_\mu^1,\ \dots,\ \sqrt{H_L}\,\delta_\mu^L,\ i\sqrt{H_L}\,\delta_\mu^L\big)^\top. \tag{B.34}$$

Since the modes decouple at $x = 0$, the LAM partition function factorises as

$$\begin{aligned} \mathcal{Z}^{\rm LAM}(0) &= \prod_{\mu=2}^{K} \mathcal{Z}_\mu^{\rm LAM}(0), \\ \mathcal{Z}_\mu^{\rm LAM}(0) &:= \int_{\mathbb{R}^{2L}} \frac{d^{2L}x}{(2\pi)^L} \exp\Big(-\tfrac{1}{2}\boldsymbol{x}^\top \mathbf{C}(H)\boldsymbol{x} + \boldsymbol{b}_\mu^\top \boldsymbol{x}\Big). \end{aligned} \tag{B.35}$$

Provided that $\Re\, \mathbf{C}(H) \succ 0$, the standard $2L$-dimensional Gaussian identity yields

$$\mathcal{Z}_\mu^{\rm LAM}(0) = (\det \mathbf{C}(H))^{-1/2} \exp\Big(\tfrac{1}{2}\,\boldsymbol{b}_\mu^\top \mathbf{C}(H)^{-1}\boldsymbol{b}_\mu\Big). \tag{B.36}$$

Summing over μ and passing to the intensive quenched pressure, we obtain

$$\frac{1}{N}\,\mathbb{E}\log \mathcal{Z}^{\rm LAM}(0) = \alpha\left[-\frac{1}{2}\log\det \mathbf{C}(H) + \frac{1}{2}\,\mathbb{E}_\delta\big[\boldsymbol{b}_\mu^\top \mathbf{C}(H)^{-1}\boldsymbol{b}_\mu\big] \right]. \tag{B.37}$$

The quadratic correction decomposes as

$$\mathbb{E}_\delta\big[\boldsymbol{b}_\mu^\top \mathbf{C}(H)^{-1}\boldsymbol{b}_\mu\big] = \frac{\beta}{2}\sum_{l=1}^{L} H_l\, \mathcal{Q}_l(H), \qquad \mathcal{Q}_l(H) := \boldsymbol{v}_l^\top \mathbf{C}(H)^{-1}\boldsymbol{v}_l, \tag{B.38}$$

where $\boldsymbol{v}_l \in \mathbb{C}^{2L}$ is the $2L$-vector that selects the $(\Re, \Im)$ components of layer l along $(1, i)$.

Remark 3 (Exact correspondence with the classical Hebbian-dreaming one-body)**.** As a consistency check, we verify that the DLAM one-body term at $x = 0$ reproduces, term by term, the standard Hebbian-dreaming one-body formula of Agliari et al. (2024a) whenever the external sources are switched off ($\psi \equiv 0$, $I_l = 0$); the two expressions then coincide up to a relabelling of the trial parameters.

Local (s, ϕ) sector. In the sourceless protocol $h_{i,l}^{\rm ext} = 0$, so that $\Theta''_{i,l} = \Delta_l \Omega_{i,l}$ with $\Omega_{i,l} := D_l\, \varphi_{i,l} + E_l\, \hat{\xi}_{i,l}^1$, and the site-factorised contribution (B.28) at $x = 0$ becomes

$$\begin{aligned} \frac{1}{N}\,\mathbb{E}\log \mathcal{Z}^{s,\phi}(0) = \sum_{l=1}^{L} \Bigg[&\log 2 + \frac{1}{2}\log \Delta_l + C_l \\ &- \frac{\Delta_l a_l^2}{2}\Big(4C_l^2 + \mathbb{E}[\Omega_{1,l}^2]\Big) + \mathbb{E}\log\cosh\Big(\Delta_l \Omega_{1,l}\Big)\Bigg], \end{aligned} \tag{B.39}$$

where $a_l^2 := t_l/[\beta(1+t_l)]$ and $\Delta_l := 1/(1 + 2C_l a_l^2) = \beta(1+t_l)/[\beta(1+t_l) + 2t_l C_l]$. Identifying the prefactors

$$\frac{1}{2}\log \Delta_l = -\frac{1}{2}\log\Big(1 + \frac{2t_l C_l}{\beta(1+t_l)}\Big), \qquad \frac{\Delta_l a_l^2}{2} = \frac{1}{2}\,\frac{t_l}{2t_l C_l + \beta(1+t_l)}, \tag{B.40}$$

we recognise the two characteristic classical dreaming terms. Restricted to a single layer, (B.39) reduces to

$$\begin{aligned} &\log 2 + C - \frac{1}{2}\log\Big(1 + \frac{2tC}{\beta(1+t)}\Big) \\ &\quad - \frac{1}{2}\,\frac{t}{2tC + \beta(1+t)}\Big(4C^2 + \mathbb{E}[\Omega^2]\Big) + \mathbb{E}\log\cosh(\Delta\Omega), \end{aligned} \tag{B.41}$$

which is precisely the classical dreaming one-body structure. The variance $\mathbb{E}[\Omega_{1,l}^2] = D_l^2 + E_l^2\,\Gamma_l$ (since $\mathbb{E}[\varphi_{1,l}^2] = 1$ and $\mathbb{E}[(\hat{\xi}_{1,l}^1)^2] = \Gamma_l$) matches the classical contribution $D^2 + E_{\rm app}^2\, r^2/\Gamma$ under the identification

$$E_l \;=\; E_{l,\rm app}\,\frac{r_l}{\Gamma_l}, \tag{B.42}$$

which leaves $E_l^2\,\Gamma_l = E_{l,\rm app}^2\, r_l^2/\Gamma_l$ invariant.

Gaussian dreaming-noise sector. The z-sector at $x = 0$ (cf. (B.13)) evaluates to

$$\frac{1}{N}\,\mathbb{E}\log \mathcal{Z}^{\text{Dream-}z}(0) = \sum_{l=1}^{L}\left[\frac{\alpha}{2}\log\Big(\frac{1+t_l}{1+B_l}\Big) + \frac{\alpha}{2}\,\frac{F_l^2(1+t_l)}{1+B_l}\right]. \tag{B.43}$$

The classical dreaming convention employs $(1 - B_{\text{app}})$ in the denominator; the two expressions agree under the sign convention

$$B_{l,\text{app}} := -B_l, \tag{B.44}$$

which leaves $1 - B_{l,\text{app}} = 1 + B_l$.

Combining (B.39) and (B.43) under the identifications (B.42)–(B.44) recovers the classical Hebbian-dreaming one-body expression of Agliari et al. (2024a) verbatim–including the logarithmic prefactor, the quadratic correction with denominator $2tC + \beta(1 + t)$, and the z-sector contributions. Reactivating the LAM coupling introduces a single modification: an additional deterministic Mattis shift $\hat{\xi}^1_{i,l}\, h_l$ inside the effective field $h^{\text{ext}}_{i,l} := I_l \iota_{i,l} + \hat{\xi}^1_{i,l}\, h_l$ (with $h_l := \sum_{m\neq l} \psi^{(m)}_l$), which displaces $\Theta''_{i,l}$ while leaving the functional form of the one-body invariant.

Appendix C. Streaming equation and trial-parameter fixing

The second ingredient of the Guerra sum rule is the streaming term $\frac{d}{dx}\mathcal{A}_{N,L}(x)$. Differentiating the interpolating action (A.35) and applying Gaussian integration by parts converts the derivative into thermal expectations of replica overlaps. Once the trial parameters are tuned so as to enforce replica-symmetric self-averaging, the streaming reduces to a deterministic expression supplemented by centred fluctuations that vanish in the thermodynamic limit.

C.1. Gibbs state, replicas, and order parameters

For a fixed realisation of the disorder, the partition function at interpolation parameter x reads

$$\mathcal{Z}_{N,L}(x) := \sum_{\boldsymbol{s}}\left[\prod_{l=1}^{L}\int D\phi_l\right]\left[\prod_{l=1}^{L}\int Dz_l\right]\int Duu^{\dagger}\ \exp\big(\mathcal{S}_x\big) \tag{C.1}$$

with $\mathcal{S}_x$ as in (A.35). The thermal average at fixed disorder and its quenched counterpart are defined by

$$\omega_x(F) := \frac{1}{\mathcal{Z}_{N,L}(x)}\sum_{\boldsymbol{s}}\left[\prod_{l=1}^{L}\int D\phi_l\right]\left[\prod_{l=1}^{L}\int Dz_l\right]\int Duu^{\dagger}\ F\ \exp\big(\mathcal{S}_x\big), \tag{C.2}$$

with $\langle\cdot\rangle_x := \mathbb{E}\,\omega_x(\cdot)$ denoting the quenched bracket. Throughout this section we make use of two independent replicas $a, b \in \{1, 2\}$ drawn from $\omega_x^{\otimes 2}$.

The layer-resolved replica overlaps are

$$\text{(spins)}\qquad h^{(l)}_{ab} := \frac{1}{N}\sum_{i=1}^{N} s^{(a)}_{i,l}\, s^{(b)}_{i,l}, \tag{C.3}$$

$$\text{(dreaming } k\text{)}\qquad q^{(l)}_{ab} := \frac{1}{N}\sum_{i=1}^{N} k^{(a)}_{i,l}\, k^{(b)}_{i,l}, \tag{C.4}$$

$$\text{(dreaming } z\text{)}\qquad p^{(l)}_{ab} := \frac{1}{K-1}\sum_{\mu=2}^{K} z^{l,(a)}_{\mu}\, z^{l,(b)}_{\mu}, \tag{C.5}$$

$$\text{(LAM } u\text{)}\qquad P^{(l)}_{ab} := \frac{1}{K-1}\sum_{\mu=2}^{K} u^{l,(a)}_{\mu}\, u^{l,(b)}_{\mu}. \tag{C.6}$$

The empirical observable $\eta_l = \frac{1}{NM}\sum_{i,k}\eta^{1,k}_{i,l}\, k_{i,l}$ (cf. Section 2 of the main text) completes the set. As elsewhere, $\alpha := (K-1)/N$ is held at $\mathcal{O}(1)$ as $N \to \infty$.

The workhorse identity for differentiating through Gaussian disorder is the Stein lemma: for $X \sim \mathcal{N}(0, 1)$ and a smooth function F with suitable integrability,

$$\mathbb{E}\,[X\,F(X)] = \mathbb{E}\,[F'(X)]. \tag{C.7}$$

When a Gaussian variable X enters the interpolating action linearly as $\sqrt{x}\,c\,X\,U$, with U an observable independent of X, the Stein lemma generates the two-replica structure

$$\begin{aligned}\mathbb{E}\big[X\,\omega_x(U)\big] &= \sqrt{x}\,c\,\mathbb{E}\Big[\omega_x(U^2) - \omega_x(U)^2\Big]\\ &= \sqrt{x}\,c\,\mathbb{E}\Big[\omega_x^{\otimes 2}\big(U^{(1)}U^{(1)} - U^{(1)}U^{(2)}\big)\Big].\end{aligned} \tag{C.8}$$

Two-replica overlaps thus emerge naturally from single-replica thermal averages.

C.2. Streaming equation: explicit computation of $\frac{d}{dx}\mathcal{A}_{N,L}$

Differentiating the quenched pressure with respect to the interpolation parameter gives

$$\frac{d}{dx}\,\mathcal{A}_{N,L}(x) = \frac{1}{N}\,\mathbb{E}\,\omega_x\big(\partial_x \mathcal{S}_x\big). \tag{C.9}$$

This derivative splits naturally into three additive contributions, one per sector of the Hamiltonian:

$$\frac{d}{dx}\mathcal{A}_{N,L} = \mathcal{D}_{\text{sig}}(x) + \mathcal{D}_{\text{LAM}}(x) + \mathcal{D}_{\text{Dream}}(x). \tag{C.10}$$

C.2.1. LAM signal part

The inter-layer signal terms in (A.35) pit the full-strength condensed overlaps $x\beta\sum_{l<m} p^1_l p^1_m$ against the one-body field $(1 - x)N\sum_{l<m}[\psi^{(m)}_l \hat{m}^1_l + \psi^{(l)}_m \hat{m}^1_m]$, which absorbs the corresponding energy at $x = 0$. Differentiating with respect to x, dividing by N, and rewriting $p^1_l = \frac{\sqrt{N}}{\sqrt{\Gamma_l}}\,\hat{m}^1_l$ through (A.16)–(A.17), we obtain

$$\mathcal{D}_{\text{sig}}(x) = \sum_{1\le l<m\le L}\left[\frac{g_{lm}\beta}{\Gamma_{lm}}\big\langle \hat{m}^1_l \hat{m}^1_m\big\rangle_x - \psi^{(m)}_l\big\langle \hat{m}^1_l\big\rangle_x - \psi^{(l)}_m\big\langle \hat{m}^1_m\big\rangle_x\right]. \tag{C.11}$$

C.2.2. LAM part

The LAM sector is governed by the Gaussian coupling $\sqrt{\frac{x\beta}{2N}}\sum_{\mu\ge 2,l,i}\lambda^{\mu}_{i,l}\, s_{i,l}u^l_{\mu}$, whose x-derivative produces a factor $(2\sqrt{x})^{-1}$ that the subsequent Stein integration by parts absorbs. Applying the identity to each $\lambda^{\mu}_{i,l}$ and rewriting the result in two-replica form yields

$$\mathcal{D}^{(\lambda)}_{\text{LAM}}(x) = \frac{\alpha\beta}{4}\sum_{l=1}^{L}\Big[\big\langle P^{(l)}_{11}\big\rangle_x - \big\langle h^{(l)}_{12}\,P^{(l)}_{12}\big\rangle_x\Big] + \mathcal{O}\Big(\frac{1}{N}\Big). \tag{C.12}$$

The reduction to two-replica overlaps proceeds as usual: integration by parts generates $\omega_x\big((s_{i,l}u^l_{\mu})^2\big) - \omega_x(s_{i,l}u^l_{\mu})^2$, which, upon summing over (i, μ) and dividing by N, yields $\langle P^{(l)}_{11}\rangle_x$ and $\langle h^{(l)}_{12}P^{(l)}_{12}\rangle_x$ respectively, up to $\mathcal{O}(1/N)$ corrections of diagonal and exchangeability origin.

The one-body noise field $\sqrt{1-x}\,I_l\sum_i \iota_{i,l}s_{i,l}$ stores the LAM fluctuation energy that is released as $x \to 1$. Its x-derivative equals $-\frac{I_l}{2\sqrt{1-x}}\sum_i \iota_{i,l}s_{i,l}$; applying the Stein lemma to $\iota_{i,l}$ then converts it into the spin-overlap variance

$$\mathcal{D}^{(\iota)}_{\text{LAM}}(x) = -\frac{1}{2}\sum_{l=1}^{L} I_l^2\Big[1 - \big\langle h^{(l)}_{12}\big\rangle_x\Big]. \tag{C.13}$$

Finally, the quadratic and linear trial couplings present in the interpolating action contribute

$$\begin{aligned}\mathcal{D}^{(\text{trial})}_{\text{LAM}}(x) = &-\alpha\sum_{l=1}^{L}\Big[C_u\,\big\langle P^{(l)}_{11}\big\rangle_x + C_{u^\dagger}\,\big\langle P^{\dagger,(l)}_{11}\big\rangle_x\Big]\\ &-\frac{\alpha}{2}\sum_{l=1}^{L}\Big[F_u^2\big(\langle P^{(l)}_{11}\rangle_x - \langle P^{(l)}_{12}\rangle_x\big) + F_{u^\dagger}^2\big(\langle P^{\dagger,(l)}_{11}\rangle_x - \langle P^{\dagger,(l)}_{12}\rangle_x\big)\Big],\end{aligned} \tag{C.14}$$

with $P^{\dagger,(l)}_{ab} := \frac{1}{K-1}\sum_{\mu\ge 2} u^{l,\dagger(a)}_{\mu}\, u^{l,\dagger(b)}_{\mu}$. Adding the three contributions (C.12)–(C.14) produces the total LAM streaming

$$\mathcal{D}_{\text{LAM}}(x) = \mathcal{D}^{(\lambda)}_{\text{LAM}}(x) + \mathcal{D}^{(\iota)}_{\text{LAM}}(x) + \mathcal{D}^{(\text{trial})}_{\text{LAM}}(x). \tag{C.15}$$

C.2.3. Dreaming part

Each layer contributes independently, so the dreaming streaming is a sum over l of three pieces with distinct origin. The Mattis-like overlap η_l enters the action quadratically through $x\,\frac{\beta N(1+t_l)}{2\Gamma_l}\,\eta_l^2$, whose x-derivative is immediate and yields

$$\mathcal{D}^{(\eta^2)}_{\text{Dream}}(x) = \sum_{l=1}^{L} \frac{\beta(1+t_l)}{2\Gamma_l} \left\langle \eta_l^2 \right\rangle_x. \tag{C.16}$$

The same Stein mechanism deployed for the LAM λ-coupling applies verbatim to the independent Gaussian copy $\sqrt{\frac{x\beta}{N}}\sum_{\mu\geq 2,i}\tilde{\lambda}^{\mu}_{i,l}\,z^{l}_{\mu}k_{i,l}$: the family $\{\tilde{\lambda}\}$ shares the marginal law of λ but is statistically independent of it, ensuring that the LAM and dreaming Gaussian sectors remain decoupled. The resulting two-replica expression reads

$$\mathcal{D}^{(\tilde{\lambda})}_{\text{Dream}}(x) = \frac{\alpha\beta}{2} \sum_{l=1}^{L} \left[\left\langle p^{(l)}_{11} q^{(l)}_{11} \right\rangle_x - \left\langle p^{(l)}_{12} q^{(l)}_{12} \right\rangle_x \right] + \mathcal{O}\Big(\frac{1}{N}\Big). \tag{C.17}$$

The remaining one-body trial couplings, once differentiated and processed through Stein's identity acting on the Gaussian fields φ and θ, give

$$\begin{aligned} \mathcal{D}^{(\text{trial})}_{\text{Dream}}(x) = \sum_{l=1}^{L} \Bigg[& \frac{\alpha B_l}{2(1+t_l)} \left\langle p^{(l)}_{11} \right\rangle_x - C_l \left\langle q^{(l)}_{11} \right\rangle_x - E_l \left\langle \eta_l \right\rangle_x \\ & - \frac{D_l^2}{2} \left(\left\langle q^{(l)}_{11} \right\rangle_x - \left\langle q^{(l)}_{12} \right\rangle_x \right) \\ & - \frac{\alpha F_l^2}{2} \left(\left\langle p^{(l)}_{11} \right\rangle_x - \left\langle p^{(l)}_{12} \right\rangle_x \right) \Bigg]. \end{aligned} \tag{C.18}$$

Assembling the three dreaming contributions (C.16)–(C.18) delivers

$$\mathcal{D}_{\text{Dream}}(x) = \mathcal{D}^{(\eta^2)}_{\text{Dream}}(x) + \mathcal{D}^{(\tilde{\lambda})}_{\text{Dream}}(x) + \mathcal{D}^{(\text{trial})}_{\text{Dream}}(x). \tag{C.19}$$

C.3. Full streaming equation

Combining the LAM contributions (C.11)–(C.15) with the dreaming sector (C.19), the complete streaming takes the form

$$\frac{d}{dx}\mathcal{A}_{N,L}(x) = \mathcal{D}_{\text{sig}}(x) + \mathcal{D}_{\text{LAM}}(x) + \mathcal{D}_{\text{Dream}}(x) + \mathcal{O}\Big(\frac{1}{N}\Big). \tag{C.20}$$

C.4. Replica–symmetric self–averaging and exact fixing of the trial parameters

The trial parameters introduced in the interpolating action remain free thus far. We now fix them by requiring that every streaming term involving random overlaps decompose into a deterministic RS contribution plus centred fluctuations that vanish as $N \to \infty$. Concretely, we impose the RS self-averaging scenario, in which each relevant overlap concentrates on a deterministic limit,

$$\langle \mathcal{O} \rangle_x \xrightarrow[N\to\infty]{} \overline{\mathcal{O}}(x) \qquad \text{for} \qquad \mathcal{O} \in \left\{ \hat{m}^1_l,\ h^{(l)}_{12},\ p^{(l)}_{ab},\ q^{(l)}_{ab},\ P^{(l)}_{ab},\ \eta_l \right\}. \tag{C.21}$$

For each layer l we introduce the corresponding RS scalars:

$$\begin{aligned} \hat{M}_l &:= \overline{\langle \hat{m}^1_l \rangle_x}, & H_l &:= \overline{\langle h^{(l)}_{12} \rangle_x}, & \bar{\eta}_l &:= \overline{\langle \eta_l \rangle_x}, \\ \bar{q}_l &:= \overline{\langle q^{(l)}_{12} \rangle_x}, & \bar{Q}_l &:= \overline{\langle q^{(l)}_{11} \rangle_x}, & \bar{p}_l &:= \overline{\langle p^{(l)}_{12} \rangle_x}, \\ \bar{P}_l &:= \overline{\langle p^{(l)}_{11} \rangle_x}, & \bar{\Pi}_l &:= \overline{\langle P^{(l)}_{12} \rangle_x}, & \overline{\mathcal{P}}_l &:= \overline{\langle P^{(l)}_{11} \rangle_x}. \end{aligned} \tag{C.22}$$

Within each sector, the trial parameters are calibrated to absorb the deterministic cost of the streaming, leaving only centred fluctuations–that is, differences between the empirical overlaps and their RS limits.

LAM signal fields. Inspecting (C.11), the requirement that the deterministic component vanish for each inter-layer pair (l, m) immediately gives

$$\psi^{(m)}_l = \frac{g_{lm}\beta}{\Gamma_{lm}}\,\hat{M}_m, \qquad \psi^{(l)}_m = \frac{g_{lm}\beta}{\Gamma_{lm}}\,\hat{M}_l. \tag{C.23}$$

LAM parameters. Matching the LAM streaming (C.12)–(C.14) under the RS ansatz is more delicate, since the trial sector accommodates both the physical channel u^l_μ and its conjugate $u^{l,\dagger}_\mu$. The trial coefficients must be chosen so that the deterministic part of the streaming reproduces the standard "RS plus centred fluctuations" structure; any residual deterministic term left uncompensated by the interaction would unbalance the decomposition. By construction of the interpolating action, the λ-term couples the spins only to u^l_μ and not to its conjugate, so that the integration-by-parts contribution (C.12) generates overlaps of type $P^{(l)}_{ab}$ but no analogue producing $P^{\dagger,(l)}_{ab}$. Any deterministic contribution from the conjugate trial sector, $-\alpha\, C_{u^\dagger}\langle P^{\dagger,(l)}_{11}\rangle_x - \frac{\alpha}{2}F^2_{u^\dagger}(\langle P^{\dagger,(l)}_{11}\rangle_x - \langle P^{\dagger,(l)}_{12}\rangle_x)$, would therefore lack a counterpart to cancel against and would spoil the structure. RS fixing accordingly enforces

$$C_{u^\dagger} = 0, \qquad F_{u^\dagger} = 0, \tag{C.24}$$

in agreement with the supervised LAM convention, where the conjugate coefficients associated with the non-excited linear channel are likewise set to zero.

The remaining LAM coefficients F^2_u and C_u are determined by deterministic cancellation. Under RS self-averaging one has $\langle P^{(l)}_{11}\rangle_x = \overline{\mathcal{P}}_l$ + fluct., $\langle P^{(l)}_{12}\rangle_x = \bar{\Pi}_l$ + fluct., and $\langle h^{(l)}_{12}P^{(l)}_{12}\rangle_x = H_l\bar{\Pi}_l$ + fluct.. Gathering the deterministic parts of (C.12) and (C.14) layer by layer leaves

$$\frac{\alpha\beta}{4}\overline{\mathcal{P}}_l - \frac{\alpha\beta}{4}H_l\bar{\Pi}_l - \alpha C_u\,\overline{\mathcal{P}}_l - \frac{\alpha}{2}F^2_u(\overline{\mathcal{P}}_l - \bar{\Pi}_l).$$

Demanding cancellation for arbitrary $(\overline{\mathcal{P}}_l, \bar{\Pi}_l)$ produces the two scalar conditions

$$\frac{\beta}{4} - C_u - \frac{F^2_u}{2} = 0, \qquad \frac{F^2_u}{2} - \frac{\beta}{4}H_l = 0,$$

hence

$$F^2_u = \frac{\beta}{2}\,H_l, \qquad C_u = \frac{\beta}{4}(1 - H_l). \tag{C.25}$$

The final LAM contribution is the ι-term; matching (C.13) against the RS noise level fixes

$$I^2_l = \frac{\alpha\beta}{2}\,\bar{\Pi}_l. \tag{C.26}$$

Dreaming parameters. Applying the same deterministic-cancellation strategy to (C.16)–(C.18) yields

$$\begin{aligned} B_l &= -\beta(1+t_l)\,(\bar{Q}_l - \bar{q}_l), & C_l &= \frac{\alpha\beta}{2}\,(\bar{P}_l - \bar{p}_l), & D^2_l &= \alpha\beta\,\bar{p}_l, \\ E_l &= \beta\,\frac{1+t_l}{\Gamma_l}\,\bar{\eta}_l, & & & F^2_l &= \beta\,\bar{q}_l. \end{aligned} \tag{C.27}$$

C.5. Final simplified form of the streaming after fixing trials

Inserting these RS choices into the exact streaming identity (C.20) and centring every empirical overlap around its RS limit (e.g. $\hat{m}^1_l - \hat{M}_l$, $p^{(l)}_{ab} - \bar{p}_l$, and so on), we arrive at the canonical decomposition

$$\frac{d}{dx}\,\mathcal{A}_{N,L}(x) = \mathcal{D}^{\text{RS}}(x) + \mathcal{R}^{\text{RS}}_{N,L}(x) + \mathcal{O}\Big(\frac{1}{N}\Big). \tag{C.28}$$

Here $\mathcal{D}^{\text{RS}}(x)$ collects all terms that survive the $N \to \infty$ limit, whereas $\mathcal{R}^{\text{RS}}_{N,L}(x)$ gathers the centred fluctuation terms, which vanish under the RS concentration hypothesis. The deterministic piece splits naturally by sector,

$$\mathcal{D}^{\text{RS}}(x) = \mathcal{D}^{\text{RS}}_{\text{sig}}(x) + \mathcal{D}^{\text{RS}}_{\text{LAM}}(x) + \mathcal{D}^{\text{RS}}_{\text{Dream}}(x), \tag{C.29}$$

with the three contributions obtained as follows. Substituting $\psi^{(m)}_l = \frac{g_{lm}\beta}{\Gamma_{lm}}\hat{M}_m$ and $\psi^{(l)}_m = \frac{g_{lm}\beta}{\Gamma_{lm}}\hat{M}_l$ into (C.11) produces the condensed-signal identity

$$\mathcal{D}_{\text{sig}}(x) = -\sum_{1\leq l<m\leq L} \frac{g_{lm}\beta}{\Gamma_{lm}}\,\hat{M}_l\hat{M}_m + \sum_{1\leq l<m\leq L} \frac{g_{lm}\beta}{\Gamma_{lm}} \left\langle (\hat{m}^1_l - \hat{M}_l)(\hat{m}^1_m - \hat{M}_m) \right\rangle_x, \tag{C.30}$$

so that the deterministic signal streaming reads

$$\mathcal{D}^{\mathrm{RS}}_{\mathrm{sig}}(x) = -\sum_{1\le l<m\le L} \frac{g_{lm}\beta}{\Gamma_{lm}}\,\hat{M}_l\hat{M}_m. \tag{C.31}$$

For the LAM sector, once the trial parameters have been fixed, the streaming admits the same decomposition

$$\mathcal{D}_{\mathrm{LAM}}(x) = \mathcal{D}^{\mathrm{RS}}_{\mathrm{LAM}}(x) + \mathcal{R}^{\mathrm{RS,LAM}}_{N,L}(x) + \mathcal{O}\Big(\frac{1}{N}\Big), \tag{C.32}$$

whose deterministic piece reads

$$\mathcal{D}^{\mathrm{RS}}_{\mathrm{LAM}} = -\frac{\alpha\beta}{4}\sum_{l=1}^{L}\bar{\Pi}_l\,(1-H_l). \tag{C.33}$$

With B_l, C_l, D_l, E_l, F_l expressed in terms of the RS limits $\bar{Q}_l, \bar{q}_l, \bar{P}_l, \bar{p}_l, \bar{\eta}_l$, the dreaming sector exhibits the same deterministic-plus-fluctuation structure,

$$\mathcal{D}_{\mathrm{Dream}}(x) = \mathcal{D}^{\mathrm{RS}}_{\mathrm{Dream}}(x) + \mathcal{R}^{\mathrm{RS,Dream}}_{N,L}(x) + \mathcal{O}\Big(\frac{1}{N}\Big), \tag{C.34}$$

with deterministic RS piece

$$\mathcal{D}^{\mathrm{RS}}_{\mathrm{Dream}}(x) = \sum_{l=1}^{L}\left[-\frac{\alpha\beta}{2}(\bar{P}_l\bar{Q}_l - \bar{p}_l\bar{q}_l) - \frac{\beta(1+t_l)}{2\Gamma_l}\,\bar{\eta}_l^{\,2}\right], \tag{C.35}$$

which is the multi-layer extension of the classical Hebbian–dreaming RS streaming. Collecting (C.31), (C.33) and (C.35), we obtain the full RS deterministic streaming

$$\begin{aligned}\mathcal{D}^{\mathrm{RS}}(x) = &-\sum_{1\le l<m\le L}\frac{g_{lm}\beta}{\Gamma_{lm}}\,\hat{M}_l\hat{M}_m - \frac{\alpha\beta}{4}\sum_{l=1}^{L}\bar{\Pi}_l\,(1-H_l)\\ &+\sum_{l=1}^{L}\left[-\frac{\alpha\beta}{2}(\bar{P}_l\bar{Q}_l - \bar{p}_l\bar{q}_l) - \frac{\beta(1+t_l)}{2\Gamma_l}\,\bar{\eta}_l^{\,2}\right].\end{aligned} \tag{C.36}$$

The remainder $\mathcal{R}^{\mathrm{RS}}_{N,L}(x)$ in (C.28) collects all centred fluctuation terms generated by (C.30) and by the analogous centring of the LAM and dreaming sectors. Explicitly,

$$\begin{aligned}\mathcal{R}^{\mathrm{RS}}_{N,L}(x) = &\underbrace{\sum_{1\le l<m\le L}\frac{g_{lm}\beta}{\Gamma_{lm}}\left\langle(\hat{m}^1_l - \hat{M}_l)(\hat{m}^1_m - \hat{M}_m)\right\rangle_x}_{\mathcal{R}^{\mathrm{RS,sig}}_{N,L}(x)}\\ &+\underbrace{\frac{\alpha\beta}{4}\sum_{l=1}^{L}\left\langle(P^{(l)}_{11}-\overline{P}_l) - (h^{(l)}_{12}-H_l)(P^{(l)}_{12}-\bar{\Pi}_l)\right\rangle_x}_{\mathcal{R}^{\mathrm{RS,LAM}}_{N,L}(x)}\\ &+\underbrace{\frac{\alpha\beta}{2}\sum_{l=1}^{L}\left\langle(p^{(l)}_{11}-\bar{P}_l)(q^{(l)}_{11}-\bar{Q}_l) - (p^{(l)}_{12}-\bar{p}_l)(q^{(l)}_{12}-\bar{q}_l)\right\rangle_x + \sum_{l=1}^{L}\frac{\beta(1+t_l)}{2\Gamma_l}\left\langle(\eta_l-\bar{\eta}_l)^2\right\rangle_x}_{\mathcal{R}^{\mathrm{RS,Dream}}_{N,L}(x)}.\end{aligned} \tag{C.37}$$

Under the RS self-averaging hypothesis, each centred term vanishes as $N\to\infty$; consequently, $\mathcal{R}^{\mathrm{RS}}_{N,L}(x)\to 0$ and the streaming equation reduces to the deterministic expression $\mathcal{D}^{\mathrm{RS}}(x)$.

Appendix D. Replica-symmetric statistical pressure

The one-body partition function at $x=0$ and the RS-fixed streaming equation jointly determine the quenched pressure through the Guerra sum rule. All remaining quantities then reduce to RS order parameters and layer-wise auxiliary expectations.

D.1. Guerra sum rule and RS decomposition

The interpolating pressure

$$\mathcal{A}_N(x) := \frac{1}{N}\mathbb{E}\log\mathcal{Z}(x), \qquad x\in[0,1], \tag{D.1}$$

recovers the full quenched pressure at $x=1$. The fundamental theorem of calculus delivers the Guerra sum rule

$$\mathcal{A}_N(1) = \mathcal{A}_N(0) + \int_0^1 dx\,\frac{d\mathcal{A}_N(x)}{dx}. \tag{D.2}$$

The streaming decomposes by sector,

$$\frac{d\mathcal{A}_N(x)}{dx} = \mathcal{D}^{\mathrm{sig}}_N(x) + \mathcal{D}^{\mathrm{LAM}}_N(x) + \sum_{l=1}^{L}\mathcal{D}^{\mathrm{Dream}}_{N,l}(x), \tag{D.3}$$

and, once the trial parameters are set to their RS values, each sector factorises as

$$\mathcal{D}^{\bullet}_N(x) = \mathcal{D}^{\mathrm{RS}}_{\bullet} + \mathcal{R}^{\bullet}_N(x), \tag{D.4}$$

where $\mathcal{D}^{\mathrm{RS}}_{\bullet}$ depends only on the RS order parameters–and is therefore x-independent–while the remainders $\mathcal{R}^{\bullet}_N(x)$ are centred fluctuations. Under the RS self-averaging hypothesis, these fluctuations integrate to zero in the thermodynamic limit,

$$\lim_{N\to\infty}\int_0^1 dx\,\mathcal{R}^{\bullet}_N(x) = 0, \qquad \bullet\in\{\mathrm{sig},\mathrm{LAM},\mathrm{Dream}\}, \tag{D.5}$$

so that (D.2)–(D.4) collapse to the RS pressure

$$\mathcal{A}^{\mathrm{RS}} := \lim_{N\to\infty}\mathcal{A}_N(1) = \mathcal{A}^{(0)} + \mathcal{D}^{\mathrm{RS}}_{\mathrm{sig}} + \mathcal{D}^{\mathrm{RS}}_{\mathrm{LAM}} + \sum_{l=1}^{L}\mathcal{D}^{\mathrm{RS}}_{\mathrm{Dream},l}, \tag{D.6}$$

where $\mathcal{A}^{(0)} := \lim_{N\to\infty}\mathcal{A}_N(0)$ denotes the one-body contribution.

D.2. Deterministic RS streaming terms

The three deterministic RS streaming contributions, derived in the preceding section, read

$$\mathcal{D}^{\mathrm{RS}}_{\mathrm{sig}} = -\sum_{1\le l<m\le L}\frac{g_{lm}\beta}{\Gamma_{lm}}\,\hat{M}_l\hat{M}_m, \tag{D.7}$$

$$\mathcal{D}^{\mathrm{RS}}_{\mathrm{LAM}} = -\frac{\alpha\beta}{4}\sum_{l=1}^{L}\bar{\Pi}_l\,(1-H_l), \tag{D.8}$$

$$\mathcal{D}^{\mathrm{RS}}_{\mathrm{Dream},l} = -\frac{\alpha\beta}{2}(\bar{P}_l\bar{Q}_l - \bar{p}_l\bar{q}_l) - \frac{\beta(1+t_l)}{2\Gamma_l}\,\bar{\eta}_l^{\,2}. \tag{D.9}$$

Here $\hat{M}_l := \overline{\langle\hat{m}^1_l\rangle}$ denotes the Mattis magnetisation on layer l, H_l is the RS diagonal overlap in the LAM Gaussian sector, and $(\bar{P}_l,\bar{p}_l)$, $(\bar{Q}_l,\bar{q}_l)$ are the diagonal/off-diagonal RS overlaps for the dreaming z- and k-variables, respectively.

D.3. One–body term at $x=0$

Setting $x=0$ decouples all interaction terms, and the partition function factorises as

$$\mathcal{Z}(0) = \mathcal{Z}^{(0)}_{\mathrm{LAM}}\mathcal{Z}^{(0)}_{\mathrm{Dream}}\mathcal{Z}^{(0)}_{s,\phi}, \tag{D.10}$$

and the quenched pressure at $x=0$ becomes

$$\mathcal{A}^{(0)} = \lim_{N\to\infty}\frac{1}{N}\mathbb{E}\log\mathcal{Z}^{(0)}_{\mathrm{LAM}} + \lim_{N\to\infty}\frac{1}{N}\mathbb{E}\log\mathcal{Z}^{(0)}_{\mathrm{Dream}} + \lim_{N\to\infty}\frac{1}{N}\mathbb{E}\log\mathcal{Z}^{(0)}_{s,\phi}. \tag{D.11}$$

D.4. Final RS pressure

Inserting the RS trial values into $\mathcal{A}^{(0)}$ and adding the deterministic streaming contributions (D.7)–(D.9) yields the complete RS pressure.

For each layer l we define

$$\bar{D}_l := 1 + \alpha\frac{t_l}{1+t_l}(\bar{P}_l - \bar{p}_l), \qquad \Delta_l := \frac{1}{\bar{D}_l}, \tag{D.12}$$

together with the effective RS local field

$$\Psi_l := I_l\iota + \frac{\sqrt{\alpha\beta\bar{p}_l}}{\bar{D}_l}\varphi + \xi_l\left(\sum_{m\neq l}\frac{g_{lm}\beta}{\Gamma_{lm}}\hat{M}_m + \frac{\beta(1+t_l)}{\Gamma_l}\,\frac{\bar{\eta}_l}{\bar{D}_l}\right), \tag{D.13}$$

where $\iota,\varphi\sim\mathcal{N}(0,1)$ are independent, ξ_l is a quenched random variable encoding the typical empirical archetype component (with variance Γ_l), and $I_l^2 = \frac{\alpha\beta}{2}\bar{\Pi}_l$.

Dreaming contribution (layer-wise). Combining the z-Gaussian piece, the local (s,ϕ) piece, and the deterministic streaming (D.9), we obtain, for each layer l,

$$\begin{aligned}\mathcal{A}^{\mathrm{RS}}_{\mathrm{Dream},l} = &-\frac{1}{2}\log\bar{D}_l - \frac{\alpha\bar{p}_l}{2\bar{D}_l}\frac{t_l}{1+t_l} + \mathbb{E}\log\cosh(\Psi_l) + \log 2\\ &+ \frac{\alpha}{2}\log\left(\frac{1+t_l}{1-\beta(\bar{Q}_l-\bar{q}_l)(1+t_l)}\right) + \frac{\alpha}{2}\frac{\beta\bar{q}_l(1+t_l)}{1-\beta(\bar{Q}_l-\bar{q}_l)(1+t_l)}\\ &+ \frac{\beta}{2}\frac{\bar{D}_l-1}{\bar{D}_l}\frac{1+t_l}{t_l} - \frac{\beta(1+t_l)}{2\Gamma_l}\bar{\eta}_l^2\frac{t_l+\bar{D}_l}{\bar{D}_l}\\ &+ \frac{\alpha\beta}{2}(\bar{p}_l\bar{q}_l - \bar{P}_l\bar{Q}_l).\end{aligned} \tag{D.14}$$

LAM contribution. The RS LAM sector combines the one-body Gaussian term at $x=0$ with the deterministic LAM streaming remainder. For each non-condensed mode $\mu\geq 2$ we introduce the real vector $\boldsymbol{x}_\mu = (\Re u^1_\mu, \Im u^1_\mu, \dots, \Re u^L_\mu, \Im u^L_\mu)\in\mathbb{R}^{2L}$, together with the associated $2L\times 2L$ Gaussian kernel $\mathbf{C}(H)$ built from the RS spin overlaps $H_l = \overline{\langle h^{(l)}_{12}\rangle}$. The standard Gaussian identity then produces the structure

$$-\frac{\alpha}{2}\log\det\mathbf{C}(H) \;+\; \frac{\alpha}{2}\,\mathbb{E}_\delta\left[\boldsymbol{b}^\top\mathbf{C}(H)^{-1}\boldsymbol{b}\right],$$

in which $\boldsymbol{b}$ is the RS linear source vector encoding the selective couplings. Introducing the quadratic correction

$$\mathbb{D}_{\mathrm{LAM}} := \mathbb{E}_\delta\left[\boldsymbol{b}^\top\mathbf{C}(H)^{-1}\boldsymbol{b}\right], \tag{D.15}$$

the RS LAM contribution takes the form

$$\mathcal{A}^{\mathrm{RS}}_{\mathrm{LAM}} = -\frac{\alpha}{2}\log\det\mathbf{C}(H) + \frac{\alpha}{2}\mathbb{D}_{\mathrm{LAM}} - \frac{\alpha\beta}{4}\sum_{l=1}^{L}\bar{\Pi}_l(1-H_l). \tag{D.16}$$

The last term is the deterministic RS streaming contribution fixed in Appendix C. Explicit closed forms for $\det\mathbf{C}(H)$ and $\mathbb{D}_{\mathrm{LAM}}$ are provided below in the specialisation to $L=3$.

Finally, the cross-layer coupling generated by the condensed LAM interaction contributes

$$\mathcal{A}^{\mathrm{RS}}_{\mathrm{sig}} = -\sum_{1\leq l<m\leq L}\frac{g_{lm}\beta}{\Gamma_{lm}}\hat{M}_l\hat{M}_m. \tag{D.17}$$

Full RS pressure. Combining the three sectors, the total quenched free-energy density $f(\beta) = -\frac{1}{\beta}\mathcal{A}^{\mathrm{RS}}_{\mathrm{DLAM}}$ is governed by

$$\begin{aligned}\mathcal{A}^{\mathrm{RS}}_{\mathrm{DLAM}} = &-\sum_{1\leq l<m\leq L}\frac{\beta\, g_{lm}}{\Gamma_{lm}}\hat{M}_l\hat{M}_m - \frac{\alpha}{2}\log\det\mathbf{C}(H) + \frac{\alpha}{2}\mathbb{D}_{\mathrm{LAM}}\\ &- \frac{\alpha\beta}{4}\sum_{l=1}^{L}\bar{\Pi}_l(1-H_l)\\ &+ \sum_{l=1}^{L}\Bigg[\log 2 + \mathbb{E}\left[\log\cosh\Psi_l\right] - \tfrac{1}{2}\log\bar{D}_l\\ &+ \tfrac{\alpha}{2}\log\left(\frac{1+t_l}{1-\beta(\bar{Q}_l-\bar{q}_l)(1+t_l)}\right)\\ &+ \frac{\alpha}{2}\frac{\beta\bar{q}_l(1+t_l)}{1-\beta(\bar{Q}_l-\bar{q}_l)(1+t_l)} - \frac{\alpha\bar{p}_l}{2\bar{D}_l}\frac{t_l}{1+t_l} + \frac{\beta}{2}\frac{\bar{D}_l-1}{\bar{D}_l}\frac{1+t_l}{t_l}\\ &- \frac{\beta(1+t_l)}{2\Gamma_l}\bar{\eta}_l^2\frac{t_l+\bar{D}_l}{\bar{D}_l} + \frac{\alpha\beta}{2}(\bar{p}_l\bar{q}_l - \bar{P}_l\bar{Q}_l)\Bigg].\end{aligned} \tag{D.18}$$

D.5. Explicit $L=3$ specialisation of the LAM contribution

In the general RS pressure (D.18), the cross-layer LAM signal and the LAM one-body Gaussian are given by (D.17) and (D.16), respectively. Specialising to $L=3$ amounts to evaluating the LAM one-body contribution through a *real* Hubbard–Stratonovich decoupling over the three inter-layer modes; this produces a 3×3 real symmetric kernel for which closed-form expressions are available. An equivalent derivation for the three-directional architecture can be found in Alessandrelli et al. (2025b).

The condensed signal term (D.17) reduces immediately to a sum over the three distinct layer pairs $(1,2)$, $(1,3)$, $(2,3)$,

$$\mathcal{A}^{\mathrm{RS}}_{\mathrm{sig}}\Big|_{L=3} = -\beta\left(\frac{g_{12}\hat{M}_1\hat{M}_2}{\Gamma_{12}} + \frac{g_{13}\hat{M}_1\hat{M}_3}{\Gamma_{13}} + \frac{g_{23}\hat{M}_2\hat{M}_3}{\Gamma_{23}}\right). \tag{D.19}$$

For the LAM Gaussian kernel, each non-condensed mode $\mu\geq 2$ is decoupled by a triplet of real auxiliary fields $\phi_l\in\mathbb{R}$, one per layer. After annealing over the random patterns and taking the replica-symmetric saddle point, the single-mode LAM partition function reduces to a standard three-dimensional real Gaussian integral $\int \mathcal{D}\phi\, e^{-\frac{1}{2}\phi^\top\mathbf{C}\phi+\mathbf{b}^\top\phi}$, with 3×3 real symmetric kernel

$$[\mathbf{C}(H)]_{lm} = \delta_{lm} \;-\; \beta\, g_{lm}\sqrt{(1-H_l)(1-H_m)}, \qquad g_{ll} := 0, \tag{D.20}$$

whose diagonal entries equal one and whose off-diagonal entries are $-\beta g_{lm}\sqrt{(1-H_l)(1-H_m)}$. In the standard three-layer architecture one has $g_{12}=g_{13}=g_{23}=g$, and at the homogeneous RS fixed point $H_l=H$ the kernel simplifies to

$$\mathbf{C}(H) = (1+u)\,I_3 - u\,J_3, \qquad u := \beta g(1-H), \tag{D.21}$$

where J_3 denotes the 3×3 all-ones matrix, so that $\mathbf{C}(H)$ has diagonal entries 1 and off-diagonal entries $-u$. This matrix commutes with every permutation of the layer indices; its spectrum is therefore determined by the two irreducible representations of S_3: the uniform mode $\mathbf{e}_0 = (1,1,1)^\top/\sqrt{3}$ carries the eigenvalue $\lambda_0 = 1-2u$, while the orthogonal shear doublet $(\mathbf{e}_1,\mathbf{e}_2)$ carries the doubly degenerate eigenvalue $\lambda_1 = 1+u$. The kernel is positive-definite–and the LAM Gaussian integral well-defined–if and only if

$$u < \tfrac{1}{2}, \qquad \text{i.e.} \quad T > 2g(1-H); \tag{D.22}$$

since the RS retrieval branch enforces a high overlap $H=q$ at low T, the product $\beta g(1-q)$ remains bounded, and condition (D.22) holds throughout the physical retrieval region.

The eigenvalue structure above gives the homogeneous determinant

$$\det\mathbf{C}(H) = (1-2u)(1+u)^2, \qquad u=\beta g(1-H), \tag{D.23}$$

whose only singularity, $u=\frac{1}{2}$, corresponds to $T=2g(1-H)$ (and, for $g=1$ and $H\to 0$, to the critical value $T^*=2$), well outside the physical retrieval region. In the heterogeneous case, with arbitrary g_{lm} and H_l, the 3×3 determinant expands as

$$\begin{aligned}\det\mathbf{C}(H) = 1 &- \beta^2\left[g_{12}^2(1-H_1)(1-H_2) + g_{13}^2(1-H_1)(1-H_3)\right.\\ &\left. + g_{23}^2(1-H_2)(1-H_3)\right]\\ &- 2\beta^3 g_{12}g_{13}g_{23}(1-H_1)(1-H_2)(1-H_3),\end{aligned} \tag{D.24}$$

which is the determinant that enters the LAM one-body free energy. The corresponding contribution to the RS free-energy density, computed per non-condensed pattern and arising from the real-HS Gaussian integral over $\mathbf{C}(\mathbf{H})$, takes the form

$$f_{\mathrm{LAM}}(\mathbf{H}) = -\frac{\alpha}{2}\log\det\mathbf{C}(\mathbf{H}) + \frac{\alpha\,\mathcal{D}^{\mathrm{num}}_{\mathrm{LAM}}(\mathbf{H})}{\det\mathbf{C}(\mathbf{H})}, \tag{D.25}$$

where $\mathcal{D}^{\mathrm{num}}_{\mathrm{LAM}}$ collects the quadratic overlap correction. At the homogeneous RS fixed point $H_l=H=q$ (the Edwards–Anderson overlap), with $u=\beta g(1-q)$, this correction admits the closed form

$$\mathcal{D}^{\mathrm{num}}_{\mathrm{LAM}}\Big|^{\mathrm{hom}}_{L=3} = 3q(1-q)\,\beta^2 g^2\left[1+\beta g(1-q)\right] = 3q(1-q)\,\beta^2 g^2(1+u), \tag{D.26}$$

while for the heterogeneous case $\mathcal{D}^{\mathrm{num}}_{\mathrm{LAM}}(\mathbf{H})$ follows by differentiating (D.24) with respect to the individual overlaps H_l, producing a polynomial in β, g_{lm} and $(1-H_l)$ whose explicit form can be read off from the 3×3 matrix-adjugate formula.

The derivatives required by the LAM saddle-point equations follow from Jacobi's formula for the derivative of a log-determinant, $\partial_\theta\log\det A = \mathrm{Tr}[A^{-1}\partial_\theta A]$. The stationarity condition $\partial_{H_l} f_{\mathrm{LAM}} = 0$ then determines the conjugate LAM overlaps $\bar{\Pi}_l$ via

$$\frac{\partial}{\partial H_l}\log\det\mathbf{C}(H) = \mathrm{Tr}\left[\mathbf{C}(H)^{-1}\,\frac{\partial\mathbf{C}}{\partial H_l}(H)\right], \qquad l=1,2,3, \tag{D.27}$$

where the only non-zero entries of $\partial_{H_l}\mathbf{C}$ are, in view of (D.20),

$$(\partial_{H_l}\mathbf{C})_{lm} = (\partial_{H_l}\mathbf{C})_{ml} = \frac{\beta\, g_{lm}}{2}\sqrt{\frac{1-H_m}{1-H_l}}, \qquad m \neq l,$$

with every other entry vanishing. In the homogeneous case this reduces to $\partial_{H_l}\mathbf{C} = \frac{\beta g}{2(1-H)}\big[\mathbf{C}(H) - (1+u)I_3\big]$. Combining (D.19) with (D.16), (D.23) and (D.26), we obtain

$$\Big(\mathcal{A}^{\mathrm{RS}}_{\mathrm{sig}} + \mathcal{A}^{\mathrm{RS}}_{\mathrm{LAM}}\Big)\Big|_{L=3} = -\beta\left(\frac{g_{12}\hat{M}_1\hat{M}_2}{\Gamma_{12}} + \frac{g_{13}\hat{M}_1\hat{M}_3}{\Gamma_{13}} + \frac{g_{23}\hat{M}_2\hat{M}_3}{\Gamma_{23}}\right) - \frac{\alpha}{2}\log\det\mathbf{C}(H) + \frac{\alpha\beta}{2}\,\frac{\mathcal{D}^{\mathrm{num}}_{\mathrm{LAM}}(\mathbf{H})}{\det\mathbf{C}(H)} - \frac{\alpha\beta}{4}\sum_{\ell=1}^{3}\bar{\Pi}_\ell(1-H_\ell) + \mathrm{const.} \tag{D.28}$$

Appendix E. Self-consistency equations

Extremising the RS pressure $\mathcal{A}^{\mathrm{RS}}_{\mathrm{DLAM}}$ with respect to every order parameter and its conjugate produces the self-consistency (saddle-point) equations that close the theory. The derivations are carried out at generic L wherever possible; specialisation to $\boldsymbol{L=3}$ is required only in the LAM sector, where the one-body Gaussian contribution demands an explicit scalar form.

E.1. Dependency map and conventions for the RS stationarity analysis

The RS order parameters to be extremised fall into three layer-wise families: the signal/LAM variable $\hat{M}_l$; the dreaming variables $\bar{\eta}_l, \bar{q}_l, \bar{Q}_l, \bar{p}_l, \bar{P}_l$; and the LAM variables $H_l, \bar{\Pi}_l$, where H_l is the RS spin overlap in the LAM channel and $\bar{\Pi}_l$ the RS off-diagonal LAM overlap entering the LAM streaming. Schematically, the RS pressure decomposes as

$$\mathcal{A}^{\mathrm{RS}}_{\mathrm{DLAM}} = \mathcal{A}^{\mathrm{RS}}_{\mathrm{sig}} + \sum_{l=1}^{L}\mathcal{A}^{\mathrm{RS}}_{\mathrm{Dream},l} + \mathcal{A}^{\mathrm{RS}}_{\mathrm{LAM}}. \tag{E.1}$$

The signal and dreaming sectors admit stationarity equations at generic L. The LAM Gaussian sector, by contrast, involves $\log\det\mathbf{C}$ and the quadratic correction $\mathbb{D}_{\mathrm{LAM}}$, whose fully explicit scalar form is tractable only after specialising to $L=3$.

To prevent any ambiguity with the microscopic archetypes $\xi^\mu_{i,l}$, the symbol ξ_l is not used in this section: for each layer l, we let $\hat{X}_l$ denote a scalar random variable that represents a generic one-site component of the condensed empirical archetype,

$$\hat{X}_l \overset{d}{=} \hat{\xi}^1_{i,l} \qquad \text{for a generic site } i, \tag{E.2}$$

with moments $\mathbb{E}[\hat{X}_l]=0$ and $\mathbb{E}[\hat{X}_l^2]=\Gamma_l$ in the benchmark Rademacher setting. With this convention the effective RS field on layer l (cf. (D.13)) takes the form

$$\Psi_l = I_l\,\iota + \frac{\sqrt{\alpha\beta\,\bar{p}_l}}{\bar{D}_l}\,\varphi + \hat{X}_l\left(\sum_{m\neq l}\frac{g_{lm}\beta}{\Gamma_{lm}}\hat{M}_m + \frac{\beta(1+t_l)}{\Gamma_l}\,\frac{\bar{\eta}_l}{\bar{D}_l}\right), \tag{E.3}$$

with $\iota,\varphi\sim\mathcal{N}(0,1)$ independent of each other and of $\hat{X}_l$. Unless stated otherwise, the layer-l expectation is taken over the joint law of the auxiliary Gaussians and of $\hat{X}_l$, i.e. $\mathbb{E}_l[\cdot] := \mathbb{E}_{\iota,\varphi,\hat{X}_l}[\cdot]$; we shall write simply $\mathbb{E}[\cdot]$ whenever no ambiguity can arise.

The stationarity equations are derived in three successive blocks: first, the signal/LAM and mixed local observables $(\hat{M}_l, \bar{\eta}_l)$ at generic L; second, the dreaming sector $(\bar{q}_l, \bar{Q}_l, \bar{p}_l, \bar{P}_l)$, again at generic L; and finally the LAM sector $(H_l, \bar{\Pi}_l)$, specialised to $L=3$ at the single point where the explicit $L=3$ LAM algebra is unavoidable.

E.2. Stationarity conditions for the signal sector: derivation of $\hat{M}_l$

The Mattis magnetisations $\hat{M}_l$ enter the RS pressure both through the condensed LAM bilinear and, via the effective fields, through the entropic terms $\mathbb{E}_l\log\cosh(\Psi_l)$. Differentiating with respect to $\hat{M}_k$ and setting the result to zero produces a vector equation that, on the physical non-degenerate branch, collapses to a componentwise self-consistency relation. Explicitly, the relevant part of the pressure reads

$$\mathcal{A}^{\mathrm{RS}}_{\mathrm{DLAM}} = -\sum_{1\le l<m\le L}\frac{g_{lm}\beta}{\Gamma_{lm}}\hat{M}_l\hat{M}_m + \sum_{l=1}^{L}\mathbb{E}_l\log\cosh(\Psi_l) + (\text{terms independent of }\hat{M}), \tag{E.4}$$

where, using the shorthand introduced in Appendix E.1,

$$\Psi_l = I_l\iota + \sigma_l\varphi + \hat{X}_l\,(u_l+v_l), \qquad u_l := \sum_{m\neq l}\frac{g_{lm}\beta}{\Gamma_{lm}}\hat{M}_m,$$
$$v_l := \frac{\beta(1+t_l)}{\Gamma_l}\,\frac{\bar{\eta}_l}{\bar{D}_l}, \tag{E.5}$$

with $\sigma_l := \frac{\sqrt{\alpha\beta\,\bar{p}_l}}{\bar{D}_l}$ and $\iota,\varphi,\hat{X}_l$ mutually independent under $\mathbb{E}_l$. Fixing $k\in\{1,\dots,L\}$, differentiation of the symmetric LAM quadratic form yields

$$\frac{\partial}{\partial\hat{M}_k}\left(-\sum_{1\le l<m\le L}\frac{g_{lm}\beta}{\Gamma_{lm}}\hat{M}_l\hat{M}_m\right) = -\sum_{m\neq k}\frac{g_{km}\beta}{\Gamma_{km}}\hat{M}_m, \tag{E.6}$$

whereas $\hat{M}_k$ enters the effective field Ψ_l only through u_l, and only for $l\neq k$: indeed $u_k = \sum_{m\neq k}\frac{g_{km}\beta}{\Gamma_{km}}\hat{M}_m$ is itself independent of $\hat{M}_k$. Hence

$$\frac{\partial u_l}{\partial\hat{M}_k} = \begin{cases}\frac{g_{lk}\beta}{\Gamma_{lk}}, & l\neq k,\\ 0, & l=k,\end{cases} \tag{E.7}$$

and the chain rule produces the entropic derivative

$$\frac{\partial}{\partial\hat{M}_k}\sum_{l=1}^{L}\mathbb{E}_l\log\cosh(\Psi_l) = \sum_{l\neq k}\mathbb{E}_l\left[\tanh(\Psi_l)\,\frac{\partial\Psi_l}{\partial\hat{M}_k}\right] = \sum_{l\neq k}\frac{g_{lk}\beta}{\Gamma_{lk}}\,\mathbb{E}_l\big[\hat{X}_l\,\tanh(\Psi_l)\big]. \tag{E.8}$$

Setting $\partial_{\hat{M}_k}\mathcal{A}^{\mathrm{RS}}_{\mathrm{DLAM}} = 0$ and combining (E.6)–(E.8) produces, for each k,

$$\sum_{m\neq k}\frac{g_{km}\beta}{\Gamma_{km}}\big(\mathbb{E}_m[\hat{X}_m\tanh(\Psi_m)] - \hat{M}_m\big) = 0. \tag{E.9}$$

Equivalently, upon introducing the weighted LAM kernel

$$(\mathbb{K}_\Gamma)_{km} := \begin{cases}\frac{g_{km}\beta}{\Gamma_{km}}, & k\neq m,\\ 0, & k=m,\end{cases} \tag{E.10}$$

and the residual vector $\Delta^{(M)}_l := \mathbb{E}_l[\hat{X}_l\tanh(\Psi_l)] - \hat{M}_l$, the stationarity condition takes the compact form

$$\mathbb{K}_\Gamma\,\boldsymbol{\Delta}^{(M)} = 0. \tag{E.11}$$

On the physically relevant branch, where $\ker\mathbb{K}_\Gamma = \{0\}$, this enforces $\boldsymbol{\Delta}^{(M)}=0$ componentwise, yielding the RS magnetisation equation

$$\hat{M}_l = \mathbb{E}_l\big[\hat{X}_l\tanh(\Psi_l)\big], \qquad l=1,\dots,L. \tag{E.12}$$

The non-degeneracy assumption holds automatically for $L=3$: the kernel is then the weighted off-diagonal 3×3 matrix

$$\begin{pmatrix} 0 & g_{12}\beta/\Gamma_{12} & g_{13}\beta/\Gamma_{13}\\ g_{12}\beta/\Gamma_{12} & 0 & g_{23}\beta/\Gamma_{23}\\ g_{13}\beta/\Gamma_{13} & g_{23}\beta/\Gamma_{23} & 0\end{pmatrix},$$

with determinant $\det\mathbb{K}_\Gamma = 2g_{12}g_{13}g_{23}\beta^3/(\Gamma_{12}\Gamma_{13}\Gamma_{23})$ (the denominator being positive, since $\Gamma_{lm}>0$). The kernel is therefore invertible whenever $g_{12}g_{13}g_{23}\neq0$, and Eq. (E.12) follows rigorously on that branch; should one or more heteroassociative channels be switched off, a separate branch analysis would be required.

E.3. Stationarity conditions for the auxiliary parameter: derivation of $\bar{\eta}_l$

The dreaming auxiliary parameter $\bar{\eta}_l$ enters only the layer-wise dreaming contribution $\mathcal{A}^{\mathrm{RS}}_{\mathrm{Dream},l}$; the derivation is therefore layer-local and applies for generic L.

From (D.14), the parameter $\bar{\eta}_l$ enters $\mathcal{A}^{\mathrm{RS}}_{\mathrm{Dream},l}$ both implicitly–through the deterministic shift $\upsilon_l = \frac{\beta(1+t_l)}{\Gamma_l}\frac{\bar{\eta}_l}{\bar{D}_l}$ inside the effective field Ψ_l–and explicitly, through the quadratic energy $-\frac{\beta(1+t_l)}{2\Gamma_l}\bar{\eta}_l^2\frac{t_l+\bar{D}_l}{\bar{D}_l}$. Retaining only the $\bar{\eta}_l$-dependent contributions,

$$\mathcal{A}^{\mathrm{RS}}_{\mathrm{Dream},l}\Big|_{\bar{\eta}_l} = \mathbb{E}_l \log\cosh(\Psi_l) - \frac{\beta(1+t_l)}{2\Gamma_l}\,\bar{\eta}_l^2\,\frac{t_l+\bar{D}_l}{\bar{D}_l}, \tag{E.13}$$

where $\bar{D}_l$ depends on $\bar{\eta}_l$ neither directly nor implicitly, since it is fixed by the combination $\bar{P}_l - \bar{p}_l$. Because Ψ_l depends on $\bar{\eta}_l$ exclusively through υ_l,

$$\frac{\partial\Psi_l}{\partial\bar{\eta}_l} = \hat{X}_l\,\frac{\partial\upsilon_l}{\partial\bar{\eta}_l} = \hat{X}_l\,\frac{\beta(1+t_l)}{\Gamma_l\bar{D}_l}, \tag{E.14}$$

so the chain rule produces the entropic derivative

$$\begin{aligned}\frac{\partial}{\partial\bar{\eta}_l}\,\mathbb{E}_l\log\cosh(\Psi_l) &= \mathbb{E}_l\left[\tanh(\Psi_l)\frac{\partial\Psi_l}{\partial\bar{\eta}_l}\right]\\ &= \frac{\beta(1+t_l)}{\Gamma_l\bar{D}_l}\,\mathbb{E}_l\left[\hat{X}_l\tanh(\Psi_l)\right],\end{aligned} \tag{E.15}$$

while differentiating the explicit quadratic term gives

$$\frac{\partial}{\partial\bar{\eta}_l}\left[-\frac{\beta(1+t_l)}{2\Gamma_l}\,\bar{\eta}_l^2\,\frac{t_l+\bar{D}_l}{\bar{D}_l}\right] = -\frac{\beta(1+t_l)}{\Gamma_l}\,\bar{\eta}_l\,\frac{t_l+\bar{D}_l}{\bar{D}_l}. \tag{E.16}$$

Imposing $\partial_{\bar{\eta}_l}\mathcal{A}^{\mathrm{RS}}_{\mathrm{DLAM}} = 0$ and combining (E.15)-(E.16) yields

$$\frac{\beta(1+t_l)}{\Gamma_l\bar{D}_l}\left[\mathbb{E}_l[\hat{X}_l\tanh(\Psi_l)] - (t_l+\bar{D}_l)\bar{\eta}_l\right] = 0, \tag{E.17}$$

which, by the strict positivity of β, $1+t_l$, Γ_l and $\bar{D}_l$ on the physical branch, reduces to $(t_l+\bar{D}_l)\bar{\eta}_l = \mathbb{E}_l[\hat{X}_l\tanh(\Psi_l)]$. Substituting the magnetisation Eq. (E.12) on the right-hand side then gives

$$\bar{\eta}_l = \frac{\hat{M}_l}{t_l+\bar{D}_l}, \qquad l = 1,\dots,L. \tag{E.18}$$

The dreaming auxiliary order parameter $\bar{\eta}_l$ is thus a rescaled Mattis magnetisation, the renormalisation being controlled by t_l and by the dreaming denominator $\bar{D}_l$.

E.4. Algebraic stationarity with respect to $\bar{q}_l$ and $\bar{Q}_l$: determination of $\bar{p}_l$ and $\bar{P}_l$

Differentiating the RS pressure with respect to $\bar{q}_l$ and $\bar{Q}_l$ produces purely algebraic relations that fix the conjugate dreaming parameters $\bar{p}_l$ and $\bar{P}_l$. The simplification rests on the observation that $\bar{q}_l$ and $\bar{Q}_l$ enter $\mathcal{A}^{\mathrm{RS}}_{\mathrm{Dream},l}$ only through the explicit rational and logarithmic terms, together with the bilinear coupling $\frac{\alpha\beta}{2}(\bar{p}_l\bar{q}_l - \bar{P}_l\bar{Q}_l)$, and not through the effective field Ψ_l; consequently, no Gaussian integration by parts is required. Defining the layer-wise denominator

$$S_l := 1 - \beta(1+t_l)(\bar{Q}_l - \bar{q}_l), \tag{E.19}$$

the $(\bar{q}_l,\bar{Q}_l)$-dependent part of $\mathcal{A}^{\mathrm{RS}}_{\mathrm{Dream},l}$ becomes

$$\mathcal{A}^{\mathrm{RS}}_{\mathrm{Dream},l}\Big|_{(\bar{q}_l,\bar{Q}_l)} = \frac{\alpha}{2}\log\left(\frac{1+t_l}{S_l}\right) + \frac{\alpha\beta(1+t_l)}{2}\,\frac{\bar{q}_l}{S_l} + \frac{\alpha\beta}{2}\,(\bar{p}_l\bar{q}_l - \bar{P}_l\bar{Q}_l), \tag{E.20}$$

together with the elementary identities

$$\frac{\partial S_l}{\partial\bar{q}_l} = \beta(1+t_l), \qquad \frac{\partial S_l}{\partial\bar{Q}_l} = -\beta(1+t_l). \tag{E.21}$$

Differentiating (E.20) with respect to $\bar{q}_l$ and using these identities, we find

$$\frac{\partial\mathcal{A}^{\mathrm{RS}}_{\mathrm{Dream},l}}{\partial\bar{q}_l} = -\frac{\alpha}{2}\,\frac{1}{S_l}\,\frac{\partial S_l}{\partial\bar{q}_l} + \frac{\alpha\beta(1+t_l)}{2}\left(\frac{1}{S_l} - \frac{\bar{q}_l}{S_l^2}\frac{\partial S_l}{\partial\bar{q}_l}\right) + \frac{\alpha\beta}{2}\,\bar{p}_l. \tag{E.22}$$

Substituting (E.21),

$$\begin{aligned}\frac{\partial\mathcal{A}^{\mathrm{RS}}_{\mathrm{Dream},l}}{\partial\bar{q}_l} &= -\frac{\alpha\beta(1+t_l)}{2S_l} + \frac{\alpha\beta(1+t_l)}{2S_l} - \frac{\alpha\beta^2(1+t_l)^2}{2S_l^2}\,\bar{q}_l + \frac{\alpha\beta}{2}\,\bar{p}_l\\ &= \frac{\alpha\beta}{2}\left(\bar{p}_l - \beta\frac{(1+t_l)^2}{S_l^2}\bar{q}_l\right).\end{aligned} \tag{E.23}$$

The stationarity condition $\partial_{\bar{q}_l}\mathcal{A}^{\mathrm{RS}}_{\mathrm{DLAM}} = 0$ thus yields

$$\bar{p}_l = \beta\,\frac{(1+t_l)^2}{S_l^2}\,\bar{q}_l. \tag{E.24}$$

An entirely analogous calculation for $\bar{Q}_l$ produces

$$\frac{\partial\mathcal{A}^{\mathrm{RS}}_{\mathrm{Dream},l}}{\partial\bar{Q}_l} = -\frac{\alpha}{2}\,\frac{1}{S_l}\,\frac{\partial S_l}{\partial\bar{Q}_l} + \frac{\alpha\beta(1+t_l)}{2}\,\bar{q}_l\left(-\frac{1}{S_l^2}\frac{\partial S_l}{\partial\bar{Q}_l}\right) - \frac{\alpha\beta}{2}\,\bar{P}_l. \tag{E.25}$$

Using (E.21),

$$\frac{\partial\mathcal{A}^{\mathrm{RS}}_{\mathrm{Dream},l}}{\partial\bar{Q}_l} = \frac{\alpha\beta(1+t_l)}{2S_l} + \frac{\alpha\beta^2(1+t_l)^2}{2S_l^2}\,\bar{q}_l - \frac{\alpha\beta}{2}\,\bar{P}_l. \tag{E.26}$$

Imposing $\partial_{\bar{Q}_l}\mathcal{A}^{\mathrm{RS}}_{\mathrm{DLAM}} = 0$,

$$\bar{P}_l = \frac{1+t_l}{S_l} + \beta\,\frac{(1+t_l)^2}{S_l^2}\,\bar{q}_l. \tag{E.27}$$

Subtracting (E.24) from (E.27) yields the identity

$$\bar{P}_l - \bar{p}_l = \frac{1+t_l}{S_l}, \tag{E.28}$$

which, combined with $\bar{D}_l = 1 + \alpha\frac{t_l}{1+t_l}(\bar{P}_l - \bar{p}_l)$ from (D.12), gives the equivalent expressions

$$\bar{D}_l = 1 + \frac{\alpha t_l}{S_l} = 1 + \frac{\alpha t_l}{1-\beta(1+t_l)(\bar{Q}_l - \bar{q}_l)}. \tag{E.29}$$

Eqs. (E.24)–(E.29) are purely algebraic and hold for each layer l at generic L; they will be combined with the Stein-derived equations for $\bar{q}_l$ and $\bar{Q}_l$ in the next subsection.

E.5. Stationarity with respect to $\bar{p}_l$, $\bar{P}_l$: exact dreaming equations (Generic L)

Differentiating the RS pressure with respect to $\bar{p}_l$ and $\bar{P}_l$ probes the entropic content of the effective field Ψ_l, with the Gaussian fluctuations of φ extracted through Stein's identity. Throughout the calculation it is convenient to fix the layer l and adopt the shorthand

$$\kappa_l := \alpha\frac{t_l}{1+t_l}, \qquad \bar{D}_l := 1 + \kappa_l(\bar{P}_l - \bar{p}_l), \tag{E.30}$$

together with the scalar quenched variable $\hat{X}_l \overset{d}{=} \hat{\xi}^1_{i,l}$, with $\mathbb{E}[\hat{X}_l^2] = \Gamma_l$ (we again write $\hat{X}_l$ to avoid clashing with the microscopic archetypes ξ_l). The effective RS field then takes the form

$$\begin{aligned}&\Psi_l = I_l\iota + \sigma_l\varphi + \hat{X}_l(u_l+\upsilon_l), \qquad \sigma_l := \frac{\sqrt{\alpha\beta\,\bar{p}_l}}{\bar{D}_l},\\ &u_l := \sum_{m\neq l}\frac{g_{lm}\beta}{\Gamma_{lm}}\hat{M}_m, \qquad \upsilon_l := \frac{\beta(1+t_l)}{\Gamma_l}\,\frac{\bar{\eta}_l}{\bar{D}_l},\end{aligned} \tag{E.31}$$

where $\iota,\varphi\sim\mathcal{N}(0,1)$ are independent of each other and of $\hat{X}_l$. Two layer-wise expectations will appear repeatedly:

$$\mathcal{T}_l := \mathbb{E}_l\left[\tanh^2(\Psi_l)\right], \qquad \mathcal{M}_l := \mathbb{E}_l\left[\hat{X}_l\tanh(\Psi_l)\right], \tag{E.32}$$

where $\mathbb{E}_l$ denotes the average over $(\iota, \varphi, \hat{X}_l)$. By the magnetisation equation derived above, $\mathcal{M}_l = \hat{M}_l$.

Since $\bar{D}_l = 1 + \kappa_l(\bar{P}_l - \bar{p}_l)$, the partial derivatives $\partial_{\bar{p}_l}\bar{D}_l = -\kappa_l$ and $\partial_{\bar{P}_l}\bar{D}_l = +\kappa_l$ propagate through σ_l and v_l according to

$$\frac{\partial \sigma_l}{\partial x} = \sigma_l\left(\frac{\mathbf{1}_{x=\bar{p}_l}}{2\bar{p}_l} - \frac{1}{\bar{D}_l}\frac{\partial \bar{D}_l}{\partial x}\right), \qquad \frac{\partial v_l}{\partial x} = -\frac{v_l}{\bar{D}_l}\frac{\partial \bar{D}_l}{\partial x}, \qquad x \in \{\bar{p}_l, \bar{P}_l\}. \tag{E.33}$$

The only non-algebraic dependence on $(\bar{p}_l, \bar{P}_l)$ resides in $\mathbb{E}_l \log\cosh(\Psi_l)$; the chain rule then yields

$$\frac{\partial}{\partial x}\mathbb{E}_l \log\cosh(\Psi_l) = \mathbb{E}_l\left[\tanh(\Psi_l)\left(\frac{\partial \sigma_l}{\partial x}\varphi + \hat{X}_l\frac{\partial v_l}{\partial x}\right)\right]. \tag{E.34}$$

The φ-term is processed by Stein's lemma, which gives $\mathbb{E}_l[\varphi\tanh(\Psi_l)] = \sigma_l\,\mathbb{E}_l[\mathrm{sech}^2(\Psi_l)] = \sigma_l(1-\mathcal{T}_l)$; the entropic derivative then becomes

$$\frac{\partial}{\partial x}\mathbb{E}_l \log\cosh(\Psi_l) = \left(\sigma_l\partial_x\sigma_l\right)(1-\mathcal{T}_l) + (\partial_x v_l)\,\mathcal{M}_l, \qquad x \in \{\bar{p}_l, \bar{P}_l\}. \tag{E.35}$$

The algebraic part of the layer-wise pressure–that is, the part of $\mathcal{A}^{\mathrm{RS}}_{\mathrm{Dream},l}$ in (D.14) other than the entropic term–depends on $\bar{D}_l$ through

$$\begin{aligned}\mathcal{A}_l^{\exp}(\bar{D}_l, \bar{p}_l, \bar{P}_l) := &-\frac{1}{2}\log\bar{D}_l - \frac{\alpha\bar{p}_l}{2\bar{D}_l}\frac{t_l}{1+t_l} + \frac{\beta}{2}\frac{\bar{D}_l - 1}{\bar{D}_l}\frac{1+t_l}{t_l}\\ &- \frac{\beta(1+t_l)}{2\Gamma_l}\bar{\eta}_l^2\frac{t_l + \bar{D}_l}{\bar{D}_l}\\ &+ \frac{\alpha\beta}{2}\left(\bar{p}_l\bar{q}_l - \bar{P}_l\bar{Q}_l\right),\end{aligned} \tag{E.36}$$

whereas the remaining terms in (D.14) are independent of $\bar{p}_l$ and $\bar{P}_l$. Differentiating $\mathcal{A}_l^{\exp}$ with respect to $\bar{D}_l$ yields

$$\frac{\partial \mathcal{A}_l^{\exp}}{\partial \bar{D}_l} = -\frac{1}{2\bar{D}_l} + \frac{\alpha\bar{p}_l}{2\bar{D}_l^2}\frac{t_l}{1+t_l} + \frac{\beta(1+t_l)}{2t_l}\frac{1}{\bar{D}_l^2} + \frac{\beta(1+t_l)t_l}{2\Gamma_l}\frac{\bar{\eta}_l^2}{\bar{D}_l^2}. \tag{E.37}$$

Combining (E.35), (E.33) and (E.37) gives the exact stationarity conditions:

$$\begin{aligned}0 &= \frac{\partial \mathcal{A}^{\mathrm{RS}}_{\mathrm{DLAM}}}{\partial \bar{p}_l}\\ &= -\kappa_l\frac{\partial \mathcal{A}_l^{\exp}}{\partial \bar{D}_l} - \frac{\alpha}{2\bar{D}_l}\frac{t_l}{1+t_l} + \frac{\alpha\beta}{2}\bar{q}_l\\ &\quad + \left[\frac{\alpha\beta}{2\bar{D}_l^2} + \frac{\kappa_l\alpha\beta\,\bar{p}_l}{\bar{D}_l^3}\right](1-\mathcal{T}_l) + \frac{\kappa_l v_l}{\bar{D}_l}\,\mathcal{M}_l.\end{aligned} \tag{E.38}$$

$$\begin{aligned}0 &= \frac{\partial \mathcal{A}^{\mathrm{RS}}_{\mathrm{DLAM}}}{\partial \bar{P}_l}\\ &= +\kappa_l\frac{\partial \mathcal{A}_l^{\exp}}{\partial \bar{D}_l} - \frac{\alpha\beta}{2}\bar{Q}_l - \frac{\kappa_l\alpha\beta\,\bar{p}_l}{\bar{D}_l^3}(1-\mathcal{T}_l) - \frac{\kappa_l v_l}{\bar{D}_l}\,\mathcal{M}_l.\end{aligned} \tag{E.39}$$

Eqs. (E.38)-(E.39) are the exact RS stationarity conditions for the dreaming variance sector. Adding them cancels every term proportional to $\partial_{\bar{D}_l}\mathcal{A}_l^{\exp}$, $\bar{p}_l$, and $v_l\mathcal{M}_l$, leaving the identity

$$\bar{Q}_l - \bar{q}_l = \frac{1}{\bar{D}_l^2}(1-\mathcal{T}_l) - \frac{t_l}{\beta(1+t_l)\bar{D}_l}, \tag{E.40}$$

which is the most convenient route to couple the physical overlap sector $(\bar{q}_l, \bar{Q}_l)$ to the effective-field statistics. Invoking the identities $\mathcal{M}_l = \hat{M}_l$ and $(t_l + \bar{D}_l)\bar{\eta}_l = \hat{M}_l$ established above, the $v_l\mathcal{M}_l$ contributions in (E.38)-(E.39) can be re-expressed entirely through $(\hat{M}_l, \bar{D}_l)$, which is the form best suited to the final closed RS system.

E.6. Closure of the dreaming overlap sector: explicit equations for $\bar{q}_l$ and $\bar{Q}_l$ (Generic L)

The dreaming-sector RS system is closed by deriving explicit equations for $\bar{q}_l$ and $\bar{Q}_l$. Throughout this subsection we keep the shorthand

$$\bar{D}_l = 1 + \alpha\frac{t_l}{1+t_l}(\bar{P}_l - \bar{p}_l), \qquad \Psi_l = I_l\iota + \sigma_l\varphi + \hat{X}_l(u_l + v_l),$$
$$\sigma_l = \frac{\sqrt{\alpha\beta\,\bar{p}_l}}{\bar{D}_l}, \qquad v_l = \frac{\beta(1+t_l)}{\Gamma_l}\frac{\bar{\eta}_l}{\bar{D}_l}, \tag{E.41}$$

in force, with $\hat{X}_l \overset{d}{=} \hat{\xi}^1_{i,l}$ and $\mathbb{E}[\hat{X}_l^2] = \Gamma_l$, together with the auxiliary expectations $\mathcal{T}_l := \mathbb{E}_l[\tanh^2(\Psi_l)]$ and $\mathcal{M}_l := \mathbb{E}_l[\hat{X}_l\tanh(\Psi_l)]$. Adding the stationarity equations $\partial_{\bar{p}_l}\mathcal{A}^{\mathrm{RS}}_{\mathrm{DLAM}} = 0$ and $\partial_{\bar{P}_l}\mathcal{A}^{\mathrm{RS}}_{\mathrm{DLAM}} = 0$ derived above produces the identity

$$\bar{Q}_l - \bar{q}_l = \frac{1}{\bar{D}_l^2}(1-\mathcal{T}_l) - \frac{t_l}{\beta(1+t_l)\bar{D}_l}, \tag{E.42}$$

valid for each layer l at generic L. Combining the same stationarity conditions with the algebraic relations extracted from the $\partial_{\bar{q}_l}$- and $\partial_{\bar{Q}_l}$-variations eliminates the $\bar{D}_l$-derivative terms; a direct (if lengthy) substitution then isolates the entropic contribution and yields the overlap equation

$$\bar{q}_l = \mathbb{E}_l\left[\tanh^2(\Psi_l)\right] = \mathcal{T}_l, \tag{E.43}$$

which identifies $\bar{q}_l$ as the Edwards–Anderson-like order parameter generated by the effective one-body field Ψ_l. Substituting (E.43) into (E.42) yields the diagonal overlap equation

$$\bar{Q}_l = \bar{q}_l + \frac{1-\bar{q}_l}{\bar{D}_l^2} - \frac{t_l}{\beta(1+t_l)\bar{D}_l}, \tag{E.44}$$

or equivalently $\bar{Q}_l = \bar{q}_l + \Delta_l^2(1-\bar{q}_l) - \frac{t_l}{\beta(1+t_l)}\Delta_l$ with $\Delta_l := 1/\bar{D}_l$. Eq. (E.43) is the standard RS Edwards–Anderson relation for the off-diagonal overlap, whereas (E.44) departs from the binary-spin identity $Q = 1$ precisely because Q_l refers to the diagonal overlap of the dreaming variables $k_{i,l} = s_{i,l} + ia_l\phi_{i,l}$ rather than to that of the Ising spins alone.

Together with the algebraic equations for $(\bar{p}_l, \bar{P}_l)$, this completes the closure of the dreaming overlap sector at generic L:

$$\begin{aligned}&\bar{p}_l = \beta\frac{(1+t_l)^2}{S_l^2}\bar{q}_l,\\ &\bar{P}_l = \frac{1+t_l}{S_l} + \beta\frac{(1+t_l)^2}{S_l^2}\bar{q}_l, \qquad S_l := 1 - \beta(1+t_l)(\bar{Q}_l - \bar{q}_l).\\ &\bar{q}_l = \mathbb{E}_l[\tanh^2(\Psi_l)],\\ &\bar{Q}_l = \bar{q}_l + \frac{1-\bar{q}_l}{\bar{D}_l^2} - \frac{t_l}{\beta(1+t_l)\bar{D}_l},\end{aligned} \tag{E.45}$$

E.7. Stationarity conditions for the LAM sector: derivation of H_l and $\bar{\Pi}_l$

The LAM-sector ($L = 3$) order parameters $\{H_l, \bar{\Pi}_l\}_{l=1}^L$ are coupled through three distinct ingredients of the RS pressure: the explicit LAM streaming term $-\frac{\alpha\beta}{4}\bar{\Pi}_l(1-H_l)$; the LAM Gaussian one-body contribution $-\frac{\alpha}{2}\log\det\mathbf{C}(H) + \frac{\alpha}{2}\mathbb{D}_{\mathrm{LAM}}(H)$; and the entropic term $\sum_m \mathbb{E}_m\log\cosh(\Psi_m)$, of which only the layer $m = l$ depends on $\bar{\Pi}_l$. Explicitly, the LAM-dependent part of the pressure reads

$$\mathcal{A}^{\mathrm{RS}}_{\mathrm{LAM}} = -\frac{\alpha}{2}\log\det\mathbf{C}(H) + \frac{\alpha}{2}\mathbb{D}_{\mathrm{LAM}}(H) - \frac{\alpha\beta}{4}\sum_{l=1}^{L}\bar{\Pi}_l(1-H_l), \tag{E.46}$$

where $\mathbf{C}(H)$ denotes the LAM complex-symmetric Gaussian kernel matrix and

$$\mathbb{D}_{\mathrm{LAM}}(H) := \mathbb{E}_{\boldsymbol{\delta}}\left[\boldsymbol{b}^{\top}\mathbf{C}(H)^{-1}\boldsymbol{b}\right]. \tag{E.47}$$

In addition, $\bar{\Pi}_l$ enters the entropic sector through the field amplitude $I_l^2 = \frac{\alpha\beta}{2}\bar{\Pi}_l$, and hence through Ψ_l (cf. (D.13)).

Stationarity with respect to $\bar{\Pi}_l$: equation for H_l (Generic L)

Fix a layer l. Since $\bar{\Pi}_l$ enters linearly in the LAM streaming and contributes to Ψ_l only through I_l, we obtain

$$\frac{\partial \mathcal{A}^{\mathrm{RS}}_{\mathrm{DLAM}}}{\partial \bar{\Pi}_l} = -\frac{\alpha\beta}{4}(1-H_l) + \frac{\partial}{\partial \bar{\Pi}_l}\,\mathbb{E}_l \log\cosh(\Psi_l), \tag{E.48}$$

where $\mathbb{E}_l$ averages over the random variables that enter Ψ_l on layer l. Since $I_l(\bar{\Pi}_l) = \sqrt{\frac{\alpha\beta}{2}\bar{\Pi}_l}$ implies $\partial I_l/\partial\bar{\Pi}_l = \alpha\beta/(4I_l)$, and since Ψ_l contains I_l only through the product $I_l\iota$ with $\iota \sim \mathcal{N}(0,1)$, the chain rule combined with Stein's lemma applied to ι (i.e. $\mathbb{E}_l[\iota\tanh(\Psi_l)] = I_l\,\mathbb{E}_l[\mathrm{sech}^2(\Psi_l)]$) collapses the entropic derivative to

$$\frac{\partial}{\partial \bar{\Pi}_l}\mathbb{E}_l \log\cosh(\Psi_l) = \frac{\alpha\beta}{4}\,\mathbb{E}_l\left[\mathrm{sech}^2(\Psi_l)\right]. \tag{E.49}$$

The stationarity condition $\partial_{\bar{\Pi}_l}\mathcal{A}^{\mathrm{RS}}_{\mathrm{DLAM}} = 0$ therefore gives

$$-\frac{\alpha\beta}{4}(1-H_l) + \frac{\alpha\beta}{4}\,\mathbb{E}_l\left[\mathrm{sech}^2(\Psi_l)\right] = 0, \tag{E.50}$$

i.e.

$$H_l = 1 - \mathbb{E}_l\left[\mathrm{sech}^2(\Psi_l)\right] = \mathbb{E}_l\left[\tanh^2(\Psi_l)\right]. \tag{E.51}$$

This is the exact RS self-consistency equation for H_l, and it holds for generic L.

Stationarity with respect to H_l: equation for $\bar{\Pi}_l$ (Generic L, matrix form)

The variable H_l enters the RS pressure only through the LAM streaming term $-\frac{\alpha\beta}{4}\bar{\Pi}_l(1-H_l)$ and the LAM Gaussian one-body term $-\frac{\alpha}{2}\log\det\mathbf{C}(H) + \frac{\alpha}{2}\mathbb{D}_{\mathrm{LAM}}(H)$. At fixed $\bar{\Pi}_l$, the effective local field Ψ_l carries no explicit H_l-dependence in the present RS parametrisation, so the entropic functional contributes nothing to $\partial_{H_l}\mathcal{A}^{\mathrm{RS}}_{\mathrm{DLAM}}$. The stationarity condition therefore reads

$$0 = \frac{\partial \mathcal{A}^{\mathrm{RS}}_{\mathrm{DLAM}}}{\partial H_l} = \frac{\alpha\beta}{4}\bar{\Pi}_l - \frac{\alpha}{2}\frac{\partial}{\partial H_l}\log\det\mathbf{C}(H) + \frac{\alpha}{2}\frac{\partial}{\partial H_l}\mathbb{D}_{\mathrm{LAM}}(H), \tag{E.52}$$

which produces the matrix-form equation

$$\bar{\Pi}_l = \frac{2}{\beta}\left[\frac{\partial}{\partial H_l}\log\det\mathbf{C}(H) - \frac{\partial}{\partial H_l}\mathbb{D}_{\mathrm{LAM}}(H)\right]. \tag{E.53}$$

An equivalent trace form follows from the standard matrix-calculus identity

$$\partial_{H_l}\log\det\mathbf{C}(H) = \mathrm{Tr}\left(\mathbf{C}(H)^{-1}\,\partial_{H_l}\mathbf{C}(H)\right) \tag{E.54}$$

together with

$$\partial_{H_l}\mathbb{D}_{\mathrm{LAM}}(H) = \partial_{H_l}\mathbb{E}_{\boldsymbol{\delta}}[\boldsymbol{b}^{\mathsf{T}}\mathbf{C}(H)^{-1}\boldsymbol{b}], \tag{E.55}$$

which keep (E.53) compact and rigorous even when no explicit scalar formula is available.

Explicit $L = 3$ closure for $\bar{\Pi}_l$

For $L = 3$ the LAM Gaussian sector admits closed-form expressions,

$$\mathbb{D}_{\mathrm{LAM}}\Big|_{L=3} = \beta\,\frac{\mathcal{D}^{\mathrm{num}}_{\mathrm{LAM}}(\mathbf{H})}{\det\mathbf{C}(H)}, \qquad \mathbf{H} := (H_1, H_2, H_3), \tag{E.56}$$

with $\det\mathbf{C}(H)$ and $\mathcal{D}^{\mathrm{num}}_{\mathrm{LAM}}(\mathbf{H})$ given by (D.23)-(D.24) and (D.26). Substituting into (E.53) yields

$$\bar{\Pi}_l = \frac{2}{\beta}\,\frac{\partial_{H_l}\det\mathbf{C}(H)}{\det\mathbf{C}(H)} - 2\,\frac{\det\mathbf{C}(H)\,\partial_{H_l}\mathcal{D}^{\mathrm{num}}_{\mathrm{LAM}} - \mathcal{D}^{\mathrm{num}}_{\mathrm{LAM}}\,\partial_{H_l}\det\mathbf{C}(H)}{(\det\mathbf{C}(H))^2}, \quad l = 1,2,3. \tag{E.57}$$

Equivalently,

$$\bar{\Pi}_l = \frac{2}{\beta}\,\partial_{H_l}\log\det\mathbf{C}(H) - 2\,\partial_{H_l}\left(\frac{\mathcal{D}^{\mathrm{num}}_{\mathrm{LAM}}}{\det\mathbf{C}(H)}\right). \tag{E.58}$$

As an illustration, the matrix identity (D.27) delivers $\partial_{H_l}\log\det\mathbf{C}(H)$ explicitly, while the analogous expression for $\partial_{H_l}\mathcal{D}^{\mathrm{num}}_{\mathrm{LAM}}$ follows from direct differentiation of (D.24). The $l = 1$ instance is then a one-step substitution into (E.57), and the remaining layers $l = 2, 3$ are obtained by cyclic permutation of the indices.

To summarise, the LAM-sector RS saddle-point system reads

$$H_l = \mathbb{E}_l[\tanh^2(\Psi_l)],$$
$$\bar{\Pi}_l = \frac{2}{\beta}\left[\partial_{H_l}\log\det\mathbf{C}(H) - \partial_{H_l}\mathbb{D}_{\mathrm{LAM}}(H)\right] \ \text{(generic } L\text{)}, \tag{E.59}$$

the second equation becoming fully explicit at $L = 3$ through (E.57). Consistency of the RS solution requires the physical branch conditions $0 \le H_l \le 1$, $\bar{\Pi}_l \ge 0$ and $\Re\,\mathbf{C}(H) \succ 0$ (or the corresponding admissibility condition in the $L = 3$ explicit form).

E.8. Compact RS self-consistency system (Generic L, with explicit LAM closure at $L = 3$)

We now collect the replica-symmetric self-consistency equations in compact form. The dreaming/LAM sector closes for arbitrary L; the LAM equation for $\bar{\Pi}_l$ is retained in matrix form at generic L and assumes an explicit scalar form once specialised to $L = 3$. For each layer $l \in \{1, \dots, L\}$ we define

$$u_l := \sum_{m\neq l}\frac{g_{lm}\beta}{\Gamma_{lm}}\hat{M}_m, \qquad \bar{D}_l := 1 + \alpha\frac{t_l}{1+t_l}(\bar{P}_l - \bar{p}_l), \qquad \Delta_l := \frac{1}{\bar{D}_l}, \tag{E.60}$$

and

$$I_l := \sqrt{\frac{\alpha\beta}{2}\bar{\Pi}_l}, \qquad \sigma_l := \frac{\sqrt{\alpha\beta\,\bar{p}_l}}{\bar{D}_l}, \qquad \upsilon_l := \frac{\beta(1+t_l)}{\Gamma_l}\,\frac{\bar{\eta}_l}{\bar{D}_l}. \tag{E.61}$$

Let $\iota, \varphi \sim \mathcal{N}(0,1)$ be independent standard Gaussians and let

$$\hat{X}_l \overset{d}{=} \hat{\xi}^1_{i,l}, \qquad \mathbb{E}[\hat{X}_l^2] = \Gamma_l, \tag{E.62}$$

stand for a representative quenched empirical-archetype component on layer l. The effective RS local field takes the form

$$\Psi_l := I_l\,\iota + \sigma_l\,\varphi + \hat{X}_l\,(u_l + \upsilon_l), \tag{E.63}$$

and we denote by $\mathbb{E}_l[\cdot]$ the expectation over $(\iota, \varphi, \hat{X}_l)$. We also set

$$S_l := 1 - \beta(1+t_l)(\bar{Q}_l - \bar{q}_l). \tag{E.64}$$

For each layer l,

$$\hat{M}_l = \mathbb{E}_l\left[\hat{X}_l\,\tanh(\Psi_l)\right]. \tag{E.65}$$

Moreover,

$$\bar{\eta}_l = \frac{\hat{M}_l}{t_l + \bar{D}_l}. \tag{E.66}$$

The off-diagonal overlap satisfies

$$\bar{q}_l = \mathbb{E}_l\left[\tanh^2(\Psi_l)\right]. \tag{E.67}$$

The diagonal overlap of the k-variables satisfies

$$\bar{Q}_l = \bar{q}_l + \frac{1-\bar{q}_l}{\bar{D}_l^2} - \frac{t_l}{\beta(1+t_l)\bar{D}_l}. \tag{E.68}$$

The conjugate/algebraic variables obey

$$\bar{p}_l = \beta\,\frac{(1+t_l)^2}{S_l^2}\,\bar{q}_l, \qquad \bar{P}_l = \frac{1+t_l}{S_l} + \beta\,\frac{(1+t_l)^2}{S_l^2}\,\bar{q}_l. \tag{E.69}$$

Equivalently,

$$\bar{P}_l - \bar{p}_l = \frac{1+t_l}{S_l}, \qquad \bar{D}_l = 1 + \frac{\alpha t_l}{S_l}. \tag{E.70}$$

The LAM spin overlap satisfies

$$H_l = \mathbb{E}_l\left[\tanh^2(\Psi_l)\right]. \tag{E.71}$$

The corresponding LAM noise order parameter $\bar{\Pi}_l$ is determined, for generic L, by the matrix-calculus stationarity equation

$$\bar{\Pi}_l = \frac{2}{\beta}\left[\partial_{H_l}\log\det\mathbf{C}(H) - \partial_{H_l}\mathbb{D}_{\mathrm{LAM}}(H)\right], \qquad l = 1, \dots, L, \tag{E.72}$$

where $\mathbf{C}(H)$ is the LAM complex-symmetric Gaussian kernel matrix and

$$\mathbb{D}_{\mathrm{LAM}}(H) := \mathbb{E}_{\boldsymbol{\delta}}\left[\boldsymbol{b}^{\mathsf{T}}\mathbf{C}(H)^{-1}\boldsymbol{b}\right]. \tag{E.73}$$

Explicit LAM closure for $L = 3$

For $L = 3$, using the closed-form expressions of Appendix D.5,

$$\mathbb{D}_{\mathrm{LAM}}\Big|_{L=3} = \beta\,\frac{D^{\mathrm{num}}_{\mathrm{LAM}}(\mathbf{H})}{\det \mathbf{C}(H)}, \qquad \mathbf{H} = (H_1, H_2, H_3), \tag{E.74}$$

so that substitution into (E.72) delivers

$$\bar{\Pi}_l = \frac{2}{\beta}\,\partial_{H_l} \log\det \mathbf{C}(H) - 2\,\partial_{H_l}\left(\frac{D^{\mathrm{num}}_{\mathrm{LAM}}(\mathbf{H})}{\det \mathbf{C}(H)}\right), \qquad l = 1, 2, 3. \tag{E.75}$$

Equivalently (quotient rule),

$$\bar{\Pi}_l = \frac{2}{\beta}\,\frac{\partial_{H_l}\det \mathbf{C}(H)}{\det \mathbf{C}(H)} - 2\,\frac{\det \mathbf{C}(H)\,\partial_{H_l} D^{\mathrm{num}}_{\mathrm{LAM}} - D^{\mathrm{num}}_{\mathrm{LAM}}\,\partial_{H_l}\det \mathbf{C}(H)}{(\det \mathbf{C}(H))^2}. \tag{E.76}$$

For generic L, the RS system closes on the $8L$ real unknowns $\{\hat{M}_l, \bar{\eta}_l, \bar{q}_l, \bar{Q}_l, \bar{p}_l, \bar{P}_l, H_l, \bar{\Pi}_l\}_{l=1}^{L}$ via (E.65)–(E.72), with the LAM stationarity treated in matrix form. At $L = 3$ the system becomes fully explicit in scalar form, through (E.75) together with the closed expressions for $\det \mathbf{C}$ and $D^{\mathrm{num}}_{\mathrm{LAM}}$ derived in Appendix D.5.

A physically admissible RS saddle point must additionally satisfy the branch conditions $0 \le \bar{q}_l \le \bar{Q}_l$, $0 \le H_l \le 1$ and $\bar{\Pi}_l \ge 0$ for every l, together with positive-definiteness of the real part of the LAM Gaussian kernel, $\Re\,\mathbf{C}(H) \succ 0$ (equivalently, positivity of every principal minor of $\Re\,\mathbf{C}(H)$ at $L = 3$). The condition is stated on the real part rather than on $\mathbf{C}(H)$ itself, since $\mathbf{C}(H)$ is complex-symmetric rather than Hermitian, and convergence of the underlying Gaussian integral is governed by $\Re\,\mathbf{C}(H)$.

Remark 4 (Functional coincidence of H_l and $\bar{q}_l$ in the RS ansatz)**.** Although H_l (the LAM spin-glass overlap) and $\bar{q}_l$ (the dreaming-sector overlap) are distinct order parameters–with different microscopic definitions and different conjugate variables–within the present RS closure they obey the same functional equation,

$$H_l = \mathbb{E}_l[\tanh^2(\Psi_l)], \qquad \bar{q}_l = \mathbb{E}_l[\tanh^2(\Psi_l)].$$

This coincidence is not an ad hoc identification: both stationarity conditions probe the same effective one-body entropy $\mathbb{E}_l \log\cosh(\Psi_l)$, and Gaussian integration by parts (Stein's lemma) produces the same nonlinear kernel $\tanh^2$. Hence, within the RS ansatz and for the effective field Ψ_l introduced above, one has $H_l = \bar{q}_l$ at the saddle point. We nevertheless retain both symbols throughout, since they belong to distinct sectors of the theory (LAM vs. dreaming) and are closed by different companion equations.

E.9. Low-load regime ($\alpha = 0$): reduced RS stationarity equations

Setting $K = \mathcal{O}(1)$ as $N \to \infty$, i.e. $\alpha = K/N \to 0$, decouples large-N fluctuations from retrieval and reduces the full RS system to a finite-dimensional fixed-point problem for the condensed variables. Since every LAM contribution to the RS pressure carries an explicit α prefactor, the LAM sector switches off entirely at low load, and $H_l, \bar{\Pi}_l$ play no role; likewise, the dreaming-variance block $(\bar{q}_l, \bar{Q}_l, \bar{p}_l, \bar{P}_l)$ becomes flat, as all its stationarity equations reduce to identities once the α-weighted terms vanish. The active low-load variables are therefore $\{\hat{M}_l, \bar{\eta}_l\}_{l=1}^{L}$. From (D.12) one has $\bar{D}_l|_{\alpha=0} = 1$, and both Gaussian amplitudes in Ψ_l vanish,

$$I_l^2 = \frac{\alpha\beta}{2}\bar{\Pi}_l = 0, \qquad \sigma_l = \frac{\sqrt{\alpha\beta\,\bar{p}_l}}{\bar{D}_l} = 0, \tag{E.77}$$

so that, once conditioned on $\hat{X}_l$, the effective RS field reduces to a purely deterministic form,

$$\Psi_l^{(0)} = \hat{X}_l\left(\sum_{m\neq l}\frac{g_{lm}\beta}{\Gamma_{lm}}\hat{M}_m + \frac{\beta(1+t_l)}{\Gamma_l}\bar{\eta}_l\right), \tag{E.78}$$

and the layer-wise expectation $\mathbb{E}_l$ collapses to an average over $\hat{X}_l$ alone. The magnetisation equation accordingly becomes

$$\hat{M}_l = \mathbb{E}_{\hat{X}_l}\left[\hat{X}_l\,\tanh\left(\Psi_l^{(0)}\right)\right], \qquad l = 1, \dots, L, \tag{E.79}$$

while the auxiliary relation (E.18), evaluated at $\bar{D}_l = 1$, gives

$$\bar{\eta}_l = \frac{\hat{M}_l}{t_l + 1}, \qquad l = 1, \dots, L. \tag{E.80}$$

Substituting (E.80) into (E.78), the sleep-time factor cancels and the field collapses to

$$\Psi_l^{(0)} = \hat{X}_l\left(\sum_{m\neq l}\frac{g_{lm}\beta}{\Gamma_{lm}}\hat{M}_m + \frac{\beta}{\Gamma_l}\hat{M}_l\right). \tag{E.81}$$

The low-load fixed-point equations for the Mattis magnetisations are therefore

$$\hat{M}_l = \mathbb{E}_{\hat{X}_l}\left[\hat{X}_l\,\tanh\left(\beta\hat{X}_l\left(\frac{\hat{M}_l}{\Gamma_l} + \sum_{m\neq l}\frac{g_{lm}\hat{M}_m}{\Gamma_{lm}}\right)\right)\right], \qquad l = 1, \dots, L. \tag{E.82}$$

Together with (E.80), Eq. (E.82) forms the complete low-load RS self-consistency system. Introducing the symmetric effective coupling matrix

$$\mathbb{J}^{(0)}_{lm} := \begin{cases}\Gamma_l^{-1}, & l = m,\\ g_{lm}\Gamma_{lm}^{-1}, & l \neq m,\end{cases} \qquad h_l^{(0)} := \beta\sum_{m=1}^{L}\mathbb{J}^{(0)}_{lm}\hat{M}_m, \tag{E.83}$$

the same equation rewrites compactly as

$$\hat{M}_l = \mathbb{E}_{\hat{X}_l}\left[\hat{X}_l\tanh\left(\hat{X}_l h_l^{(0)}\right)\right], \qquad l = 1, \dots, L. \tag{E.84}$$

Although the auxiliary $\bar{\eta}_l$ does depend on t_l through (E.80), the resulting field (E.81) and the magnetisation Eq. (E.82) are themselves independent of t_l: at $\alpha = 0$ the dreaming sector merely rescales the auxiliary representation $\bar{\eta}_l$, while leaving the condensed retrieval fixed points invariant, and the retrieval physics is governed entirely by the effective couplings Γ_l^{-1} and $g_{lm}\Gamma_{lm}^{-1}$. Substituting (E.80) into the $\alpha = 0$ RS pressure eliminates all explicit t_l-dependence, leaving (up to additive constants independent of $\hat{M}$)

$$\mathcal{A}^{\mathrm{RS}}_{\alpha=0} = -\sum_{1\le l<m\le L}\frac{g_{lm}\beta}{\Gamma_{lm}}\hat{M}_l\hat{M}_m + \sum_{l=1}^{L}\mathbb{E}_{\hat{X}_l}\log\cosh\left(\hat{X}_l h_l^{(0)}\right) - \sum_{l=1}^{L}\frac{\beta}{2\Gamma_l}\hat{M}_l^2, \tag{E.85}$$

whose stationarity with respect to $\hat{M}_l$ reproduces (E.82).

At low load, then, the condensed Mattis variables $\hat{M}_l$ carry the full retrieval physics, with $\bar{\eta}_l$ slaved to them by (E.80) and the extensive LAM/dreaming sectors inactive.

E.10. RS consistency reductions: hybrid $t = 0$ limit, LAM-only limit, and decoupled dreaming limit

Three exact reductions of the RS theory connect the DLAM construction to its supervised reference limits and elucidate the distinct roles of the dreaming-time parameter t_l and the genuine on/off switching of the autoassociative sector. All three are stated and proved at the level of the RS pressure and the compact self-consistency equations, starting from the decomposition

$$\mathcal{A}^{\mathrm{RS}}_{\mathrm{DLAM}} = \mathcal{A}^{\mathrm{RS}}_{\mathrm{sig}} + \mathcal{A}^{\mathrm{RS}}_{\mathrm{LAM}} + \sum_{l=1}^{L}\mathcal{A}^{\mathrm{RS}}_{\mathrm{Dream},l}, \tag{E.86}$$

cf. (D.18), together with the compact RS field representation (Appendix E.8)

$$\Psi_l = I_l\,\iota + \sigma_l\,\varphi + \hat{X}_l\,(u_l + v_l), \qquad l = 1, \dots, L, \tag{E.87}$$

where

$$u_l := \sum_{m\neq l}\frac{g_{lm}\beta}{\Gamma_{lm}}\hat{M}_m, \qquad I_l := \sqrt{\frac{\alpha\beta}{2}\bar{\Pi}_l},$$

$$\sigma_l := \frac{\sqrt{\alpha\beta\,\bar{p}_l}}{\bar{D}_l}, \qquad \upsilon_l := \frac{\beta(1+t_l)}{\Gamma_l}\frac{\bar{\eta}_l}{\bar{D}_l}, \tag{E.88}$$

with

$$\bar{D}_l = 1 + \alpha\frac{t_l}{1+t_l}(\bar{P}_l - \bar{p}_l), \qquad \bar{\eta}_l = \frac{\hat{M}_l}{t_l + \bar{D}_l}. \tag{E.89}$$

In the decomposition (E.87), u_l carries the condensed heteroassociative signal, I_l is the LAM Gaussian contribution, and the pair (σ_l, υ_l) belongs to the autoassociative dreaming sector; each reduction below isolates a limit by switching off the corresponding piece of (E.86)-(E.87). The resulting picture can be previewed as follows. Setting $t_l = 0$ across all layers produces a *hybrid* limit–heteroassociation supplemented by a raw Hebbian intra-layer autoassociative term–rather than the supervised LAM closure. Suppressing the *entire* dreaming sector (not merely setting $t_l = 0$) yields the LAM-supervised RS equations. Finally, suppressing every heteroassociative channel decouples the L layers into independent one-layer supervised-dreaming problems.

The $t_l = 0$ limit is hybrid: simple autoassociation + heteroassociation

Fix a layer l and set $t_l = 0$ (and analogously for every layer). This operation removes the dreaming deformation but does not extinguish the autoassociative sector. From (E.89) one immediately obtains the limits of the two key denominators,

$$\bar{D}_l = 1 + \alpha\frac{t_l}{1+t_l}(\bar{P}_l - \bar{p}_l) \quad\Longrightarrow\quad \bar{D}_l\Big|_{t_l=0} = 1, \tag{E.90}$$

$$\bar{\eta}_l = \frac{\hat{M}_l}{t_l + \bar{D}_l} \quad\Longrightarrow\quad \bar{\eta}_l\Big|_{t_l=0} = \hat{M}_l, \tag{E.91}$$

and the dreaming signal contribution υ_l correspondingly reduces to

$$\upsilon_l\Big|_{t_l=0} = \frac{\beta(1+t_l)}{\Gamma_l}\frac{\bar{\eta}_l}{\bar{D}_l}\Big|_{t_l=0} = \frac{\beta}{\Gamma_l}\hat{M}_l. \tag{E.92}$$

The condensed deterministic field thereby becomes

$$\hat{X}_l(u_l + \upsilon_l)\Big|_{t_l=0} = \hat{X}_l\left(\sum_{m\neq l}\frac{g_{lm}\beta}{\Gamma_{lm}}\hat{M}_m + \frac{\beta}{\Gamma_l}\hat{M}_l\right), \tag{E.93}$$

a superposition of the inter-layer heteroassociative signal ($m \neq l$) and a simple intra-layer autoassociative term, the diagonal $\propto \hat{M}_l/\Gamma_l$. In the same limit only $\bar{D}_l$ is modified inside σ_l, while I_l remains unaffected:

$$I_l = \sqrt{\frac{\alpha\beta}{2}\bar{\Pi}_l}\,\iota, \qquad \sigma_l\Big|_{t_l=0} = \sqrt{\alpha\beta\,\bar{p}_l}, \tag{E.94}$$

so that the full RS effective field at $t_l = 0$ reads

$$\Psi_l\Big|_{t_l=0} = I_l\,\iota + \sqrt{\alpha\beta\,\bar{p}_l}\,\varphi + \hat{X}_l\left(\sum_{m\neq l}\frac{g_{lm}\beta}{\Gamma_{lm}}\hat{M}_m + \frac{\beta}{\Gamma_l}\hat{M}_l\right). \tag{E.95}$$

The RS self-consistency equations therefore remain coupled across layers and retain the entire LAM Gaussian block; the dreaming-time deformation has disappeared, replaced by a raw Hebbian intra-layer term.

The layer-wise dreaming contribution $\mathcal{A}^{\mathrm{RS}}_{\mathrm{Dream},l}$ (cf. (D.14)) contains terms carrying explicit factors $1/t_l$, so the limit $t_l \to 0$ must be performed only after algebraic simplification. The potentially singular contribution is

$$T_{1,l} := \frac{\beta}{2}\frac{\bar{D}_l - 1}{\bar{D}_l}\frac{1+t_l}{t_l}, \tag{E.96}$$

which, using the auxiliary identity

$$\bar{D}_l - 1 = \alpha\frac{t_l}{1+t_l}(\bar{P}_l - \bar{p}_l) \tag{E.97}$$

that follows from (E.89), simplifies to

$$T_{1,l} = \frac{\beta}{2}\frac{\alpha\frac{t_l}{1+t_l}(\bar{P}_l - \bar{p}_l)}{\bar{D}_l}\frac{1+t_l}{t_l} = \frac{\alpha\beta}{2}\frac{\bar{P}_l - \bar{p}_l}{\bar{D}_l}, \tag{E.98}$$

which is manifestly regular at $t_l = 0$. The magnetisation-dependent dreaming term is likewise regular: setting

$$T_{2,l} := -\frac{\beta(1+t_l)}{2\Gamma_l}\bar{\eta}_l^2\frac{t_l + \bar{D}_l}{\bar{D}_l} \tag{E.99}$$

and substituting $\bar{\eta}_l = \hat{M}_l/(t_l + \bar{D}_l)$ produces

$$T_{2,l} = -\frac{\beta(1+t_l)}{2\Gamma_l}\frac{\hat{M}_l^2}{(t_l+\bar{D}_l)^2}\frac{t_l+\bar{D}_l}{\bar{D}_l} = -\frac{\beta(1+t_l)}{2\Gamma_l}\frac{\hat{M}_l^2}{(t_l+\bar{D}_l)\bar{D}_l}, \tag{E.100}$$

whence

$$\lim_{t_l\to 0} T_{2,l} = -\frac{\beta}{2\Gamma_l}\hat{M}_l^2. \tag{E.101}$$

Since the term proportional to $t_l/(1+t_l)$ in (D.14) likewise vanishes as $t_l \to 0$, the layer-wise contribution $\mathcal{A}^{\mathrm{RS}}_{\mathrm{Dream},l}$ admits a finite–and in general non-zero–$t_l \to 0$ limit. The $t_l = 0$ limit is therefore *not* the LAM-supervised closure; it yields instead the hybrid RS limit

$$\begin{aligned}\text{DLAM at } t_l = 0 \;&=\; \text{heteroassociative (LAM)}\\ &+\; \text{simple intra-layer autoassociation.}\end{aligned} \tag{E.102}$$

Switching off the autoassociative sector yields the LAM-supervised RS closure

Suppressing the autoassociative sector amounts to removing the entire dreaming contribution from the pressure, together with its order-parameter block from the variational problem,

$$\sum_{l=1}^{L}\mathcal{A}^{\mathrm{RS}}_{\mathrm{Dream},l} \equiv 0, \qquad \{\bar{\eta}_l, \bar{q}_l, \bar{Q}_l, \bar{p}_l, \bar{P}_l\}_{l=1}^{L}$$

removed from the RS saddle-point sector, (E.103)

and constitutes the correct RS implementation of "switching off the autoassociative links"–in particular, it is *not* equivalent to setting $t_l = 0$. From (E.86) the pressure reduces to

$$\mathcal{A}^{\mathrm{RS}}_{\mathrm{DLAM}} \longrightarrow \mathcal{A}^{\mathrm{RS}}_{\mathrm{sig}} + \mathcal{A}^{\mathrm{RS}}_{\mathrm{LAM}}, \tag{E.104}$$

which coincides with the LAM-supervised sector up to the present notation and normalisations. Since both σ_l and υ_l belong to the dreaming block (cf. (E.88)), suppressing the latter forces

$$\sigma_l = 0, \qquad \upsilon_l = 0, \tag{E.105}$$

while the heteroassociative LAM signal u_l and the LAM Gaussian contribution I_l remain active. The effective RS field thereby collapses to

$$\Psi_l^{(\mathrm{LAM})} = I_l\,\iota + \hat{X}_l\,u_l = \sqrt{\frac{\alpha\beta}{2}\bar{\Pi}_l}\,\iota + \hat{X}_l\sum_{m\neq l}\frac{g_{lm}\beta}{\Gamma_{lm}}\hat{M}_m. \tag{E.106}$$

The surviving compact RS equations are then the LAM ones:

$$\hat{M}_l = \mathbb{E}_l\left[\hat{X}_l\tanh\left(\Psi_l^{(\mathrm{LAM})}\right)\right], \tag{E.107}$$

$$H_l = \mathbb{E}_l\left[\tanh^2\left(\Psi_l^{(\mathrm{LAM})}\right)\right], \tag{E.108}$$

$$\bar{\Pi}_l = \frac{2}{\beta}\left[\partial_{H_l}\log\det\mathbf{C}(H) - \partial_{H_l}\mathbb{D}_{\mathrm{LAM}}(H)\right], \tag{E.109}$$

for $l = 1, \ldots, L$–precisely the LAM-supervised closure expressed in the present empirical-archetype notation.

A direct comparison of (E.95) with (E.106) highlights the structural difference between the two reductions: at $t_l = 0$ the term $\upsilon_l = \beta\hat{M}_l/\Gamma_l$ is still present, whereas in the dreaming-off limit it vanishes identically. Hence

$$t_l = 0 \qquad \not\equiv \qquad \text{dreaming-off (autoassociative sector off).} \tag{E.110}$$

Switching off all heteroassociative channels yields decoupled supervised dreaming

In what follows we make the hetero-off reduction explicit layer by layer, mapping each resulting RS saddle-point equation onto its counterpart in the one-layer supervised-dreaming system of Agliari et al.

(2024a) (Appendix S5, Eqs. (E.5) and (E.12)–(E.16) therein). At the level of the RS pressure and self-consistency equations, suppressing the heteroassociative sector in its entirety requires switching off both the condensed cross-layer LAM signal and the LAM Gaussian heteroassociative sector, which we implement by setting

$$\begin{aligned} &g_{lm} = 0 \quad (l \neq m), \qquad \text{hence} \qquad u_l := \sum_{m \neq l} \frac{g_{lm}\beta}{\Gamma_{lm}} \hat{M}_m = 0 \quad \forall l, \\ &g_{lm} = 0 \quad (l \neq m), \qquad I_l = 0 \\ &\quad \text{(equivalently, by } I_l^2 = \frac{\alpha\beta}{2}\bar{\Pi}_l,\ \bar{\Pi}_l = 0). \end{aligned} \tag{E.111}$$

For $L = 3$ this amounts to $g_{12} = g_{13} = g_{23} = 0$ and the choice of branch $\bar{\Pi}_l = I_l = 0$. Substituting (E.111) into the RS decomposition (E.86), the pressure reduces to a sum of layer-wise dreaming contributions and the variational problem factorises,

$$\mathcal{A}^{\mathrm{RS}}_{\mathrm{DLAM}} \longrightarrow \sum_{l=1}^{L} \mathcal{A}^{\mathrm{RS}}_{\mathrm{Dream},l}, \tag{E.112}$$

while the effective field on each layer reduces to

$$\Psi_l^{(\mathrm{Dream})} = \sigma_l\, \varphi + \hat{X}_l\, v_l, \qquad \sigma_l = \frac{\sqrt{\alpha\beta\,\bar{p}_l}}{\bar{D}_l}, \qquad v_l = \frac{\beta(1+t_l)}{\Gamma_l}\frac{\bar{\eta}_l}{\bar{D}_l}. \tag{E.113}$$

Each layer therefore obeys the closed one-layer dreaming RS system

$$\hat{M}_l = \mathbb{E}_l\left[\hat{X}_l \tanh\left(\Psi_l^{(\mathrm{Dream})}\right)\right], \tag{E.114}$$

$$\bar{\eta}_l = \frac{\hat{M}_l}{t_l + \bar{D}_l}, \tag{E.115}$$

$$\bar{q}_l = \mathbb{E}_l\left[\tanh^2\left(\Psi_l^{(\mathrm{Dream})}\right)\right], \tag{E.116}$$

$$\bar{Q}_l = \bar{q}_l + \frac{1-\bar{q}_l}{\bar{D}_l^2} - \frac{t_l}{\beta(1+t_l)\bar{D}_l}, \tag{E.117}$$

$$\begin{aligned} &\bar{p}_l = \beta\frac{(1+t_l)^2}{S_l^2}\bar{q}_l, \qquad \bar{P}_l = \frac{1+t_l}{S_l} + \beta\frac{(1+t_l)^2}{S_l^2}\bar{q}_l, \\ &S_l = 1 - \beta(1+t_l)(\bar{Q}_l - \bar{q}_l), \end{aligned} \tag{E.118}$$

with

$$\bar{D}_l = 1 + \alpha\frac{t_l}{1+t_l}(\bar{P}_l - \bar{p}_l). \tag{E.119}$$

Since no equation for layer l depends on variables associated with $m \neq l$, the RS variational problem splits into L independent one-layer supervised-dreaming saddle-point problems.

To compare (E.114)–(E.119) with Eqs. (E.5) and (E.12)–(E.16) of Agliari et al. (2024a), we must account for the layer-wise normalisations of the present multi-layer notation. The structural identifications

$$t \leftrightarrow t_l, \qquad \rho \leftrightarrow \rho_l, \qquad \bar{D} \leftrightarrow \bar{D}_l, \qquad (\bar{p}, \bar{q}, \bar{Q}, \bar{P}) \leftrightarrow (\bar{p}_l, \bar{q}_l, \bar{Q}_l, \bar{P}_l) \tag{E.120}$$

must be supplemented with the LaD-normalised magnetisation-like variables

$$m_{\eta,l} := \frac{\hat{M}_l}{r_l(1+\rho_l)}, \qquad \eta_l^{(\mathrm{S5})} := \frac{\bar{\eta}_l}{r_l(1+\rho_l)}, \tag{E.121}$$

together with the layer-wise normalisation

$$\Gamma_l = r_l^2(1+\rho_l). \tag{E.122}$$

With this dictionary the one-layer equations of Agliari et al. (2024a) follow directly. Eq. (E.5) of Agliari et al. (2024a), $\bar{\eta} = \bar{m}_\eta/(t+\bar{D})$, matches (E.115) once both sides of $\bar{\eta}_l = \hat{M}_l/(t_l+\bar{D}_l)$ are divided by $r_l(1+\rho_l)$ and (E.121) is applied:

$$\eta_l^{(\mathrm{S5})} = \frac{m_{\eta,l}}{t_l + \bar{D}_l}. \tag{E.123}$$

For the $\bar{D}$-Eq. (E.15) of Agliari et al. (2024a), we start from (E.118) and compute directly

$$\bar{P}_l - \bar{p}_l = \frac{1+t_l}{S_l}, \tag{E.124}$$

which, inserted into (E.119), yields

$$\bar{D}_l = 1 + \alpha\frac{t_l}{1+t_l}\cdot\frac{1+t_l}{S_l} = 1 + \frac{\alpha t_l}{S_l}, \tag{E.125}$$

and, after substituting $S_l = 1 - \beta(1+t_l)(\bar{Q}_l - \bar{q}_l)$, recovers the form of Eq. (E.15) of Agliari et al. (2024a),

$$\bar{D}_l = 1 + \frac{\alpha t_l}{1 - \beta(1+t_l)(\bar{Q}_l - \bar{q}_l)}. \tag{E.126}$$

Eq. (E.14) of Agliari et al. (2024a) is already encoded in (E.118): indeed

$$\bar{p}_l = \beta\frac{(1+t_l)^2}{S_l^2}\bar{q}_l, \qquad S_l = 1 - \beta(1+t_l)(\bar{Q}_l - \bar{q}_l), \tag{E.127}$$

matches the S5 form verbatim and requires no further rescaling. For the $\bar{Q} - \bar{q}$ relation (E.12) of Agliari et al. (2024a), starting from (E.117), subtracting $\bar{q}_l$, and multiplying by $\bar{D}_l^2$ gives

$$\begin{aligned} \bar{D}_l^2(\bar{Q}_l - \bar{q}_l) &= \bar{D}_l^2\left[\frac{1-\bar{q}_l}{\bar{D}_l^2} - \frac{t_l}{\beta(1+t_l)\bar{D}_l}\right] \\ &= 1 - \bar{q}_l - \frac{t_l\bar{D}_l}{\beta(1+t_l)}, \end{aligned} \tag{E.128}$$

which is precisely the structure of Eq. (E.12) of Agliari et al. (2024a): the one-layer quantity there denoted by $\hat{q}$ coincides with the right-hand side of the present one-layer $\tanh^2$-expectation, $\hat{q}_l \equiv \bar{q}_l$, whence

$$\bar{D}_l^2(\bar{Q}_l - \bar{q}_l) = 1 - \hat{q}_l - \frac{t_l\bar{D}_l}{\beta(1+t_l)}, \qquad \hat{q}_l \equiv \bar{q}_l \text{ (present compact notation).} \tag{E.129}$$

The matching to Eq. (E.13) of Agliari et al. (2024a)–the one-Gaussian $\hat{q}$ representation–is the only step that is not a one-line symbol substitution: it reduces the two-source random field in (E.113) to the single-Gaussian effective field of Agliari et al. (2024a). Using (E.122) and (E.121), the deterministic coefficient in v_l becomes

$$\begin{aligned} \hat{X}_l\, v_l &= \hat{X}_l\,\frac{\beta(1+t_l)}{\Gamma_l}\frac{\bar{\eta}_l}{\bar{D}_l} \\ &= \hat{X}_l\,\frac{\beta(1+t_l)}{r_l^2(1+\rho_l)}\frac{r_l(1+\rho_l)\,\eta_l^{(\mathrm{S5})}}{\bar{D}_l} \\ &= \hat{X}_l\,\frac{\beta(1+t_l)}{r_l}\frac{\eta_l^{(\mathrm{S5})}}{\bar{D}_l} \\ &= \hat{X}_l\,\frac{\beta(1+t_l)}{r_l}\frac{m_{\eta,l}}{(t_l+\bar{D}_l)\bar{D}_l}, \end{aligned} \tag{E.130}$$

where in the last step we invoked (E.123). Comparison with the one-layer S5 field then requires the one-dimensional decomposition that underlies the S5 derivation–equivalently, the CLT reduction of the empirical-class average along the condensed direction–namely

$$\hat{X}_l \overset{d}{=} \xi_l\, r_l\,(1 + \sqrt{\rho_l}\, z), \qquad \xi_l \in \{\pm 1\}, \quad z \sim \mathcal{N}(0,1), \tag{E.131}$$

with z independent of φ. Inserting (E.131) into (E.113) and using (E.130) produces, in distribution,

$$\begin{aligned} \Psi_l^{(\mathrm{Dream})} &\overset{d}{=} \frac{\sqrt{\alpha\beta\,\bar{p}_l}}{\bar{D}_l}\,\varphi + \xi_l\, r_l(1+\sqrt{\rho_l}\, z)\frac{\beta(1+t_l)}{r_l}\frac{m_{\eta,l}}{(t_l+\bar{D}_l)\bar{D}_l} \\ &= \frac{\sqrt{\alpha\beta\,\bar{p}_l}}{\bar{D}_l}\,\varphi + \xi_l\,\frac{\beta(1+t_l)}{\bar{D}_l+t_l}\frac{m_{\eta,l}}{\bar{D}_l} + \xi_l\,\frac{\beta(1+t_l)}{\bar{D}_l+t_l}\frac{m_{\eta,l}\sqrt{\rho_l}}{\bar{D}_l}\, z. \end{aligned} \tag{E.132}$$

The last expression has the form "deterministic shift plus sum of two independent Gaussians"; the two Gaussian terms can therefore be fused

into a single standard normal $\psi \sim \mathcal{N}(0,1)$, yielding

$$\Psi_l^{(\mathrm{Dream})} \overset{d}{=} \frac{1}{\bar{D}_l}\sqrt{\alpha\beta\,\bar{p}_l + \left(\beta m_{\eta,l}\frac{1+t_l}{\bar{D}_l+t_l}\right)^2}\,\rho_l\,\psi + \xi_l\,\frac{\beta m_{\eta,l}}{\bar{D}_l}\,\frac{1+t_l}{\bar{D}_l+t_l}, \quad \text{(E.133)}$$

so that the one-layer $\tanh^2$ Eq. (E.116) rewrites as

$$\hat{q}_l = \mathbb{E}_\psi\left[\tanh^2\left(\frac{1}{\bar{D}_l}\sqrt{\alpha\beta\,\bar{p}_l + \left(\beta m_{\eta,l}\frac{1+t_l}{\bar{D}_l+t_l}\right)^2}\,\rho_l\,\psi + \frac{\beta m_{\eta,l}}{\bar{D}_l}\,\frac{1+t_l}{\bar{D}_l+t_l}\right)\right], \quad \text{(E.134)}$$

which is precisely the form of Eq. (E.13) of Agliari et al. (2024a): the same layer-wise structure and the same coefficients, modulo the dictionary above.

The match with Eq. (E.16) of Agliari et al. (2024a) is the only remaining step that is not a one-line symbol substitution, because the present compact notation uses $\bar{P}_l$ as a primary saddle variable, whereas Agliari et al. (2024a) recasts the free energy in terms of $\bar{D}$ as an explicit variational variable and imposes stationarity with respect to it. The two formulations are related by the invertible change of variables (valid for $t_l > 0$, and by continuity in the limit $t_l \to 0$)

$$\bar{D}_l = 1 + \alpha\frac{t_l}{1+t_l}(\bar{P}_l - \bar{p}_l) \qquad \Longleftrightarrow \qquad \bar{P}_l = \bar{p}_l + \frac{1+t_l}{\alpha t_l}(\bar{D}_l - 1), \quad \text{(E.135)}$$

so that S5 E.16 is the one-layer saddle condition obtained by restricting to the hetero-off pressure $\mathcal{A}^{\mathrm{RS}}_{\mathrm{Dream},l}$, eliminating $\bar{\eta}_l$ via (E.123), replacing $\bar{P}_l$ by $\bar{D}_l$ through (E.135), and imposing stationarity with respect to $\bar{D}_l$. In the present $\bar{P}_l$-based parametrisation the same information is encoded by the algebraic block (E.116)–(E.119). To render the equivalence transparent at the level of coefficients, consider the magnetisation-dependent term of (D.14), namely

$$-\frac{\beta(1+t_l)}{2\Gamma_l}\bar{\eta}_l^2\,\frac{t_l+\bar{D}_l}{\bar{D}_l}, \quad \text{(E.136)}$$

which, using (E.122) together with $\bar{\eta}_l = r_l(1+\rho_l)\eta_l^{(\mathrm{S5})}$, becomes

$$\begin{aligned}-\frac{\beta(1+t_l)}{2\Gamma_l}\bar{\eta}_l^2\,\frac{t_l+\bar{D}_l}{\bar{D}_l} &= -\frac{\beta(1+t_l)}{2r_l^2(1+\rho_l)}\,r_l^2(1+\rho_l)^2\,(\eta_l^{(\mathrm{S5})})^2\,\frac{t_l+\bar{D}_l}{\bar{D}_l}\\ &= -\frac{\beta(1+t_l)}{2}(1+\rho_l)(\eta_l^{(\mathrm{S5})})^2\,\frac{t_l+\bar{D}_l}{\bar{D}_l}. \quad \text{(E.137)}\end{aligned}$$

Substituting Eq. (E.5) of Agliari et al. (2024a) in the form (E.123)–namely $\eta_l^{(\mathrm{S5})} = m_{\eta,l}/(t_l+\bar{D}_l)$–gives

$$-\frac{\beta(1+t_l)}{2}(1+\rho_l)\,\frac{m_{\eta,l}^2}{\bar{D}_l(t_l+\bar{D}_l)}, \quad \text{(E.138)}$$

which is the magnetisation-sector structure entering the $\bar{D}$-stationarity Eq. (E.16) of Agliari et al. (2024a), confirming that the present layer-wise notation reproduces the S5 coefficients exactly. For each fixed layer l, the hetero-off compact RS system (E.114)–(E.119) therefore coincides with Eqs. (E.5) and (E.12)–(E.16) of Agliari et al. (2024a) under the structural identifications (E.120), the magnetisation normalisation (E.121), and the change of variables (E.135). Thus the hetero-off reduction yields not merely a structurally decoupled system but L independent copies of the one-layer supervised-dreaming RS theory of Agliari et al. (2024a).

Appendix F. Signal-to-noise analysis: one-step magnetisations and the learning threshold

The replica-symmetric theory of Appendices A–E provides the exact equilibrium phase structure of the DLAM. It is complemented here by a signal-to-noise (one-step) analysis carried out in the noiseless limit $\beta \to +\infty$ (equivalently $T \to 0^+$), in the spirit of the supervised Hebbian treatment of Alemanno et al. (2023), Alessandrelli et al. (2025b). The aim is twofold: to obtain, by an elementary and fully transparent route, the conditions under which the network locally retrieves an archetype and to recast these conditions in Hinton's one-step Monte Carlo language (Agliari et al., 2019), which is precisely the update rule implemented in the simulations of Section 4. The technique trades the sharp control of statistical mechanics for analytical economy. Nevertheless, it reproduces the equilibrium thresholds and isolates the same signal, dreaming-noise and inter-layer-noise channels that appear in the effective field (28).

F.1. One-step dynamics and the stability criterion

The single-spin-flip dynamics of layer l is governed by the local field

$$h_{i,l}(\boldsymbol{s}) := \frac{\partial\big(-\beta\,\mathcal{H}_{\mathrm{DLAM}}(\boldsymbol{s})\big)}{\partial s_{i,l}}, \qquad -\beta\,\mathcal{H}_{\mathrm{DLAM}}(\boldsymbol{s}) = \frac{1}{2}\sum_{l=1}^{L}\sum_{i=1}^{N} h_{i,l}(\boldsymbol{s})\,s_{i,l}, \quad \text{(F.1)}$$

the second identity holding because $\mathcal{H}_{\mathrm{DLAM}}$ is a quadratic form in the spins. The stochastic update reads $s_{i,l}^{(n+1)} = s_{i,l}^{(n)}\,\mathrm{sign}\big[\tanh\big(h_{i,l}^{(n)}s_{i,l}^{(n)}\big) + \zeta_{i,l}^{(n)}\big]$ with $\zeta_{i,l}^{(n)} \sim \mathcal{U}(-1,1)$; in the fast-noise-free limit $\beta \to +\infty$ it collapses to the deterministic map

$$s_{i,l}^{(n+1)} = \mathrm{sign}\big(h_{i,l}^{(n)}\big). \quad \text{(F.2)}$$

To probe the local stability of the retrieval of archetype $\mu = 1$ we prepare every layer in the corresponding true archetype, $s_{i,l}^{(1)} = \xi_{i,l}^1$ (the network never observes the ξ_l^1; it only sees the dataset through the empirical archetypes $\hat{\xi}_{i,l}^{\mu}$), and evaluate the magnetisation produced by a single synchronous step:

$$m_l^{(2)} := \frac{1}{N}\sum_{i=1}^{N}\xi_{i,l}^1\,s_{i,l}^{(2)} = \frac{1}{N}\sum_{i=1}^{N}\mathrm{sign}\big(h_{i,l}^{(1)}(\boldsymbol{\xi}^1)\,\xi_{i,l}^1\big). \quad \text{(F.3)}$$

For $N \gg 1$ the field $h_{i,l}(\boldsymbol{\xi}^1)$ is a sum of many weakly correlated contributions; by the central limit theorem the product $h_{i,l}\xi_{i,l}^1$ is asymptotically Gaussian, with mean and variance controlled by the first two moments

$$\mu_1^{(l)} := \mathbb{E}\big[h_{i,l}(\boldsymbol{\xi}^1)\,\xi_{i,l}^1\big], \qquad \mu_2^{(l)} := \mathbb{E}\big[h_{i,l}(\boldsymbol{\xi}^1)^2\big], \quad \text{(F.4)}$$

where $\mathbb{E} = \mathbb{E}_\xi\,\mathbb{E}_{\chi|\xi}$ averages over archetypes and example-noise. Consequently

$$\begin{aligned}m_l^{(2)} &\underset{N\gg1}{\simeq} \int\frac{dz}{\sqrt{2\pi}}\,e^{-z^2/2}\,\mathrm{sign}\Big(\mu_1^{(l)} + z\sqrt{\mu_2^{(l)} - (\mu_1^{(l)})^2}\Big)\\ &= \mathrm{erf}\left(\frac{\mu_1^{(l)}}{\sqrt{2\big(\mu_2^{(l)} - (\mu_1^{(l)})^2\big)}}\right). \quad \text{(F.5)}\end{aligned}$$

The archetype is one-step stable on layer l whenever $m_l^{(2)} > \mathrm{erf}(\theta)$, i.e.

$$\mathrm{SNR}_l := \frac{\mu_1^{(l)}}{\sqrt{\mu_2^{(l)} - (\mu_1^{(l)})^2}} > \sqrt{2}\,\theta, \quad \text{(F.6)}$$

where $\theta \in \mathbb{R}^+$ fixes the tolerance ($\mathrm{erf}(\theta)$ lower-bounds the recovered overlap); the canonical choice $\theta = 1/\sqrt{2}$ reproduces the textbook condition $\mu_1^{(l)} > \sqrt{\mathrm{Var}(h_{i,l}\xi_{i,l}^1)}$, namely signal exceeding the fluctuation of the field. Joint retrieval on all layers requires (F.6) for every l, in agreement with the success criterion of Section 2.6.

F.2. The internal field of the DLAM

Differentiating (6)-(4) according to (F.1) splits the field into an intra-layer (dreaming) and an inter-layer (LAM) contribution,

$$h_{i,l} = h_{i,l}^{\mathrm{Dream}} + h_{i,l}^{\mathrm{LAM}}, \quad \text{(F.7)}$$

$$h_{i,l}^{\mathrm{Dream}}(\boldsymbol{s}) = \frac{\beta}{N\Gamma_l}\sum_{\mu,\nu=1}^{K}\hat{\xi}_{i,l}^{\mu}\,[\mathbb{A}_l(t_l)]_{\mu\nu}\sum_{j\neq i}\hat{\xi}_{j,l}^{\nu}\,s_{j,l}, \quad \text{(F.8)}$$

$$h_{i,l}^{\mathrm{LAM}}(s) = \frac{\beta}{N}\sum_{m\neq l}\frac{g_{lm}}{\Gamma_{lm}}\sum_{\mu=1}^{K}\hat{\xi}_{i,l}^{\mu}\sum_{j=1}^{N}\hat{\xi}_{j,m}^{\mu}\, s_{j,m}. \tag{F.9}$$

Here the intra-layer (autoassociative) kernel of (F.8) is normalised in the symmetric Hopfield convention (each unordered neuron pair counted once) so that (F.1) holds and the one-step field is consistent with the effective field (28); equivalently, the bare auto-coupling equals unity and is promoted to g_{auto} under the parametrisation of footnote 5. To treat both sectors on the same footing we write $g_{ll} \equiv 1$ for the dreaming auto-coupling, so that the index $m = l$ in the sums below denotes the intra-layer channel and $m \neq l$ the heteroassociative channels. Throughout we use the dataset statistics

$$\begin{aligned}&\mathbb{E}\big[\hat{\xi}_{i,l}^{1}\,\xi_{i,l}^{1}\big] = r_l, \qquad \mathbb{E}\big[(\hat{\xi}_{i,l}^{\mu})^2\big] = \Gamma_l = r_l^2(1+\rho_l),\\ &\mathbb{E}\big[\hat{\xi}_{i,l}^{\mu}\hat{\xi}_{j,l}^{\mu}\big] = r_l^2\,\xi_{i,l}^{\mu}\xi_{j,l}^{\mu} \quad (i\neq j),\end{aligned} \tag{F.10}$$

with distinct modes/layers independent, all of which follow from $\hat{\xi}_{i,l}^{\mu} = M^{-1}\sum_k \chi_{i,l}^{\mu,k}\,\xi_{i,l}^{\mu}$, $\mathbb{E}[\chi] = r_l$, and the definition (2) of ρ_l.

F.3. First and second moments of the bare model ($t_l = 0$)

At $t_l = 0$ the kernel reduces to the identity, $\mathbb{A}_l(0) = \mathbb{I}$, and (F.8) becomes the supervised Hebbian autoassociative field. We compute the moments (F.4) at the archetype-aligned configuration $s_{j,m} = \xi_{j,m}^{1}$.

Signal (First moment). Only the condensed mode $\mu = 1$ survives the average, since $\mathbb{E}[\hat{\xi}_{i,l}^{\mu}\xi_{i,l}^{1}] = r_l\,\delta_{\mu 1}$. Using (F.10) and $\Gamma_{lm} = r_l r_m\sqrt{(1+\rho_l)(1+\rho_m)}$,

$$\mu_1^{(l)} = \underbrace{\frac{\beta}{N\Gamma_l}(N-1)\,r_l^2}_{\text{dreaming/auto}} + \underbrace{\frac{\beta}{N}\sum_{m\neq l}\frac{g_{lm}}{\Gamma_{lm}}\,N\,r_l r_m}_{\text{LAM/hetero}} \underset{N\gg 1}{=} \beta\,\bar{\mu}_l,$$

$$\bar{\mu}_l := \frac{1}{1+\rho_l} + \sum_{m\neq l}\frac{g_{lm}}{\sqrt{(1+\rho_l)(1+\rho_m)}}. \tag{F.11}$$

The heteroassociative couplings enlarge the signal: each partner layer relays a contribution $g_{lm}/\sqrt{(1+\rho_l)(1+\rho_m)}$, exactly the condensed term $r_l u_l$ of the effective field (29) evaluated at the retrieval value $\hat{M}_m \to r_m$.

Noise (Second moment). Write $h_{i,l} = h_{i,l}^{(1)} + h_{i,l}^{(\geq 2)}$, separating the condensed mode $\mu = 1$ from the non-condensed modes $\mu \geq 2$. The two pieces are uncorrelated at leading order (distinct modes are independent and the $\mu \geq 2$ block has zero mean), so $\mu_2^{(l)} = \mathbb{E}[(h_{i,l}^{(1)})^2] + \mathbb{E}[(h_{i,l}^{(\geq 2)})^2]$.

a. Condensed mode. Both sectors carry the common site factor $\hat{\xi}_{i,l}^{1}$, hence $h_{i,l}^{(1)} = \hat{\xi}_{i,l}^{1}\,\Phi_{i,l}$ with $\Phi_{i,l}$ depending only on sites $j \neq i$ (and on the partner layers). Since $\hat{\xi}_{i,l}^{1}$ is independent of $\Phi_{i,l}$, and $\Phi_{i,l}$ concentrates with $\mathbb{E}[\Phi_{i,l}] = \mu_1^{(l)}/r_l$,

$$\mathbb{E}\big[(h_{i,l}^{(1)})^2\big] = \mathbb{E}\big[(\hat{\xi}_{i,l}^{1})^2\big]\,\mathbb{E}\big[\Phi_{i,l}^2\big] \underset{N\gg 1}{=} \Gamma_l\cdot\frac{(\mu_1^{(l)})^2}{r_l^2} = (1+\rho_l)\,(\mu_1^{(l)})^2. \tag{F.12}$$

The amplification factor $\Gamma_l/r_l^2 = 1+\rho_l$ is the finite-sample fluctuation of the empirical archetype: it is the sole signature of the dataset entropy in the second moment, and it is shared coherently by the dreaming and LAM channels because both read out neuron i through the same $\hat{\xi}_{i,l}^{1}$.

b. Non-condensed modes. For $\mu \geq 2$ the cross-channel and cross-mode terms vanish on average (orthogonality of distinct archetypes and independence of distinct layers), leaving only the diagonal interference,

$$\begin{aligned}\mathbb{E}\big[(h_{i,l}^{(\geq 2)})^2\big] &= \frac{\beta^2}{N^2}\sum_{m=1}^{L}\frac{g_{lm}^2}{\Gamma_{lm}^2}\sum_{\mu\geq 2}\mathbb{E}\big[(\hat{\xi}_{i,l}^{\mu})^2\big]\,\mathbb{E}\Big[\big(\textstyle\sum_j \hat{\xi}_{j,m}^{\mu}\xi_{j,m}^{1}\big)^2\Big]\\ &\underset{N\gg 1}{=} \alpha\beta^2\sum_{m=1}^{L} g_{lm}^2,\end{aligned} \tag{F.13}$$

having used $\Gamma_{ll} = \Gamma_l$, $\mathbb{E}[(\hat{\xi}_{i,l}^{\mu})^2] = \Gamma_l$, $\mathbb{E}[(\sum_j \hat{\xi}_{j,m}^{\mu}\xi_{j,m}^{1})^2] = N\Gamma_m$, and $K - 1 \simeq \alpha N$. Combining (F.12)-(F.13) the variance of the field is

$$\mu_2^{(l)} - (\mu_1^{(l)})^2 = \underbrace{\rho_l\,(\mu_1^{(l)})^2}_{\text{example interference}} + \underbrace{\alpha\beta^2\sum_{m=1}^{L} g_{lm}^2}_{\text{inter-pattern crosstalk}}, \qquad g_{ll} = 1. \tag{F.14}$$

The two noise sources are the one-step counterparts of the dreaming-noise channel $\sigma_l\varphi$ ($m = l$ term) and the LAM-noise channel I_{ll} ($m \neq l$ terms) of the effective field (28): the inter-pattern crosstalk $\alpha\beta^2 g_{ll}^2$ matches σ_l^2 at $t_l = 0$, while $\alpha\beta^2\sum_{m\neq l} g_{lm}^2$ matches $I_l^2 = \frac{\alpha\beta}{2}\bar{\Pi}_l$. Both estimates (F.11) and (F.14) agree with direct Monte Carlo evaluation of the one-step field to within finite-size corrections.

F.4. One-step magnetisation and the threshold for learning

Inserting (F.11) and (F.14) into (F.5), and noting that β cancels in the ratio, the one-step magnetisation of layer l in the bare DLAM is

$$\begin{aligned}m_l^{(2)}\Big|_{t=0} &= \mathrm{erf}\left(\frac{\bar{\mu}_l}{\sqrt{2\Big[\rho_l\,\bar{\mu}_l^2 + \alpha\sum_{m=1}^{L} g_{lm}^2\Big]}}\right),\\ \bar{\mu}_l &= \frac{1}{1+\rho_l} + \sum_{m\neq l}\frac{g_{lm}}{\sqrt{(1+\rho_l)(1+\rho_m)}}.\end{aligned} \tag{F.15}$$

With the tolerance $\theta = 1/\sqrt{2}$, the stability criterion (F.6) becomes $\bar{\mu}_l^2 > \rho_l\bar{\mu}_l^2 + \alpha\sum_m g_{lm}^2$, i.e. the *learning threshold*

$$\rho_l + \frac{\alpha\sum_{m=1}^{L} g_{lm}^2}{\bar{\mu}_l^2} < 1 \qquad\Longleftrightarrow\qquad \alpha\sum_{m=1}^{L} g_{lm}^2 < (1-\rho_l)\,\bar{\mu}_l^2, \qquad g_{ll} = 1, \tag{F.16}$$

to hold simultaneously on every layer. Condition (F.16) makes the cooperative trade-off explicit: a hetero-coupling g_{lm} adds crosstalk αg_{lm}^2 to the numerator but enlarges the signal $\bar{\mu}_l$ in the denominator through the cross term $g_{lm}/\sqrt{(1+\rho_l)(1+\rho_m)}$; whenever the latter gain dominates -in particular when a poorly-sampled layer (large ρ_l, small $\bar{\mu}_l$) borrows signal from a better-informed partner (small ρ_m)- the inter-layer channel widens the retrieval region, which is the mechanism behind the disentanglement and data-saving results of Section 4.

Reductions. *(i)* For a single decoupled layer ($L = 1$, $g_{11} = 1$) one has $\bar{\mu}_1 = 1/(1+\rho_1)$ and (F.16) collapses to

$$\alpha\,(1+\rho)^2 + \rho < 1, \tag{F.17}$$

the supervised single-layer threshold of Alemanno et al. (2023); setting in addition $\rho = 0$ (noiseless patterns) returns the Amit–Gutfreund–Sompolinsky bound $\alpha < 1$ for one-step stability.

(ii) Switching off dreaming-induced effects ($t_l = 0$) while retaining the couplings, (F.16) reproduces the multidirectional supervised result of Alessandrelli et al. (2025b): in the fully symmetric case $\rho_l \equiv \rho$, $g_{lm} \equiv g$ ($m \neq l$), one finds $\bar{\mu} = \big[1 + (L-1)g\big]/(1+\rho)$ and the boundary

$$\alpha\,\frac{1+(L-1)g^2}{\big[1+(L-1)g\big]^2}\,(1+\rho)^2 + \rho < 1. \tag{F.18}$$

Since $1 + (L-1)g^2 \leq \big[1+(L-1)g\big]^2$ for $g \in [0,1]$, the prefactor of α is at most $(1+\rho)^2$, with equality only at $g = 0$: heteroassociation can only relax the single-layer load constraint, never tighten it.

F.5. The dreaming kernel at finite sleeping time: exact one-step moments

For $t_l > 0$ the kernel $\mathbb{A}_l(t_l) = (1+t_l)(\mathbb{I} + t_l\Sigma_l)^{-1}$ must be kept in full. We now compute the dreaming-sector moments exactly, in closed form and for arbitrary t_l, by transferring the problem from the K-dimensional

pattern space to the N-dimensional neuron space, where the dreaming coupling becomes a resolvent amenable to Marchenko–Pastur (MP) analysis.

With the rescaled patterns $J^{\mu}_{i,l} = \hat{\xi}^{\mu}_{i,l}/\sqrt{\Gamma_l}$, the $N \times K$ matrix $J_l = (J^{\mu}_{i,l})$, the pattern-space correlation $\Sigma_l = \frac{1}{N} J_l^{\top} J_l$ and the neuron-space overlap matrix

$$G_l := \tfrac{1}{N}\, J_l J_l^{\top} \in \mathbb{R}^{N\times N}, \tag{F.19}$$

the intra-layer coupling $\mathcal{J}^{(l)} := \frac{1}{N\Gamma_l}\hat{\Xi}_l\, \mathbb{A}_l\, \hat{\Xi}_l^{\top} = \frac{1}{N} J_l\, \mathbb{A}_l\, J_l^{\top}$ collapses, via the push-through identity $J(\mathbb{I} + \frac{t}{N} J^{\top}J)^{-1} = (\mathbb{I} + \frac{t}{N} JJ^{\top})^{-1}J$, to a function of G_l alone:

$$\mathcal{J}^{(l)} = (1+t_l)\, G_l(\mathbb{I} + t_l G_l)^{-1} = \frac{1+t_l}{t_l}\big(\mathbb{I} - \mathbb{M}_l\big), \qquad \mathbb{M}_l := (\mathbb{I} + t_l G_l)^{-1}. \tag{F.20}$$

Since $h^{\mathrm{Dream}}_{i,l}(\boldsymbol{\xi}^1) = \beta \sum_{j\neq i} \mathcal{J}^{(l)}_{ij}\, \xi^1_{j,l}$, the aligned field reduces to a quadratic form in the resolvent $\mathbb{M}_l$:

$$h^{\mathrm{Dream}}_{i,l}\, \xi^1_{i,l} = \beta\, \frac{1+t_l}{t_l}\Big[(\mathbb{M}_l)_{ii} - \xi^1_{i,l}(\mathbb{M}_l \boldsymbol{\xi}^1_l)_i\Big] = -\beta\, \frac{1+t_l}{t_l} \sum_{j\neq i} (\mathbb{M}_l)_{ij}\, \xi^1_{i,l}\xi^1_{j,l}. \tag{F.21}$$

All one-step observables are therefore controlled by three self-averaging spectral functionals of $\mathbb{M}_l$,

$$m_l := \tfrac{1}{N}\mathrm{Tr}\, \mathbb{M}_l, \qquad q_l := \tfrac{1}{N}\, \boldsymbol{\xi}^{1\top}_l \mathbb{M}_l\, \boldsymbol{\xi}^1_l, \qquad q^{(2)}_l := \tfrac{1}{N}\, \boldsymbol{\xi}^{1\top}_l \mathbb{M}_l^2\, \boldsymbol{\xi}^1_l. \tag{F.22}$$

Averaging (F.21) and using the concentration of $(\mathbb{M}_l)_{ii} \to m_l$,

$$\mu^{\mathrm{Dream}}_1(t_l) = \beta\, \frac{1+t_l}{t_l}\, \big(m_l - q_l\big). \tag{F.23}$$

For the variance, differentiating the resolvent gives the exact operator identity $\mathbb{M}_l^2 = \mathbb{M}_l + t_l\, d\mathbb{M}_l/dt_l$ (from $\frac{d}{dt_l}\mathbb{M}_l = -\mathbb{M}_l G_l \mathbb{M}_l$ and $t_l G_l \mathbb{M}_l = \mathbb{I} - \mathbb{M}_l$), whence $q^{(2)}_l = \frac{d}{dt_l}(t_l q_l)$; carrying the variance of (F.21) through (again with $(\mathbb{M}_l)_{ii} \to m_l$) the diagonal terms cancel and one is left with

$$\mu^{\mathrm{Dream}}_2 - (\mu^{\mathrm{Dream}}_1)^2 = \beta^2\, \frac{(1+t_l)^2}{t_l^2}\, \big(q^{(2)}_l - q_l^2\big). \tag{F.24}$$

Split the overlap matrix into its non-condensed (Wishart) block and the rank-one signal contribution,

$$G_l = G_0 + u\, u^{\top}, \qquad G_0 := \tfrac{1}{N} \sum_{\mu\geq 2} J^{\mu}_l (J^{\mu}_l)^{\top}, \qquad u := \tfrac{1}{\sqrt{N}}\, J^1_l. \tag{F.25}$$

Under the Gaussian equivalence (A.32) the nonzero spectrum of G_0 obeys the MP law of ratio α; its resolvent density $m_l = (1-\alpha) + \alpha \int \frac{d\mu^{(\alpha)}_{\mathrm{MP}}(g)}{1+t_l g}$ admits the closed form fixed by the MP self-consistency

$$t_l\, m_l\, (m_l + \alpha - 1) = 1 - m_l \iff m_l = \frac{\Delta_l - 1 + (1-\alpha)t_l}{2t_l},$$
$$\Delta_l := \sqrt{1 + 2(1+\alpha)t_l + (1-\alpha)^2 t_l^2}\,. \tag{F.26}$$

The condensed column carries the signal, $J^1_{i,l} = \sqrt{1-\delta_l}\, \xi^1_{i,l} + \sqrt{\delta_l}\, g_{i,l}$, with sampling-noise fraction

$$\delta_l := 1 - \frac{r_l^2}{\Gamma_l} = \frac{\rho_l}{1+\rho_l}, \qquad g_{i,l} \sim \mathcal{N}(0,1) \text{ independent of } G_0. \tag{F.27}$$

Treating $uu^{\top}$ as a rank-one perturbation, the Sherman–Morrison formula gives, exactly in the thermodynamic limit,

$$q_l = m_l\, \frac{1 + t_l\, \delta_l\, m_l}{1 + t_l\, m_l}, \qquad q^{(2)}_l = \frac{d}{dt_l}\big(t_l\, q_l\big). \tag{F.28}$$

Eqs. (F.26)–(F.28) are the exact closed-form moments at arbitrary sleeping time; we have verified them, and the resulting $\mu^{\mathrm{Dream}}_1(t_l)$ and variance (F.24), against direct Monte Carlo evaluation of (F.21) over the whole range $t_l \in [0, +\infty)$ (agreement within finite-size corrections).

Using (F.28) one finds $m_l - q_l = \dfrac{t_l\, m_l^2(1-\delta_l)}{1 + t_l m_l}$, so the signal (F.23) takes the compact form

$$\mu^{\mathrm{Dream}}_1(t_l) = \frac{\beta}{1+\rho_l}\, \frac{(1+t_l)\, m_l^2}{1 + t_l m_l} =: \frac{\beta}{1+\rho_l}\, c_S(t_l), \tag{F.29}$$

with signal form factor $c_S(0) = 1$ (since $m_l \to 1$) and $c_S(\infty) = 1 - \alpha$. The prefactor $\beta(1+t_l)/t_l$ cancels in the ratio, so the one-step magnetisation and signal-to-noise ratio are temperature-independent and read

$$m^{(2)}_l(t_l) = \mathrm{erf}\left(\frac{m_l - q_l}{\sqrt{2\,(q^{(2)}_l - q_l^2)}}\right), \qquad \mathrm{SNR}^2_l(t_l) = \frac{(m_l - q_l)^2}{q^{(2)}_l - q_l^2}. \tag{F.30}$$

The learning threshold $\mathrm{SNR}^2_l > 1$ (tolerance $\theta = 1/\sqrt{2}$) interpolates between two closed-form endpoints, obtained by inserting $m_l = 1$ and $m_l = 1 - \alpha$ into (F.28)–(F.30):

$$t_l = 0: \quad \mathrm{SNR}^2_l = \frac{1}{\alpha(1+\rho_l)^2 + \rho_l} \implies \alpha\,(1+\rho_l)^2 + \rho_l < 1, \tag{F.31}$$

$$t_l \to +\infty: \quad \mathrm{SNR}^2_l = \frac{1-\alpha}{\rho_l\,(1+\alpha\rho_l)} \implies \alpha\,(1+\rho_l^2) + \rho_l < 1. \tag{F.32}$$

The Hebbian boundary (F.31) reproduces (F.17); the $t_l \to +\infty$ limit is the Kanter–Sompolinsky projector ($\mathbb{A}_l \to \Sigma_l^{-1}$, $\mathcal{J}^{(l)} \to P_l = \hat{\Xi}_l(\hat{\Xi}_l^{\top}\hat{\Xi}_l)^{-1}\hat{\Xi}_l^{\top}$) reached by Agliari et al. (2019), Fachechi et al. (2019). The two boundaries differ *only* in the load penalty multiplying α: dreaming replaces the Hebbian factor $(1+\rho_l)^2 = 1 + 2\rho_l + \rho_l^2$ by the projector factor $1 + \rho_l^2$, i.e. it removes the linear term $2\rho_l$: precisely the portion of the inter-pattern crosstalk that the finite-sample fluctuation of the empirical archetypes amplifies. Consolidation thus buys back, in the load budget, an amount $2\alpha\rho_l$ of tolerable noise, the signal-to-noise face of the data–computation trade-off of Section 3.4; at $\rho_l = 0$ both endpoints collapse to the one-step capacity $\alpha < 1$.

The interpolation between the two endpoints is, however, not monotone whenever the data are imperfect ($\rho_l > 0$). Differentiating (F.30) shows that $\mathrm{SNR}^2_l(t_l)$ rises from its Hebbian value, attains a maximum at a finite optimal sleeping time $t^{\star}_l(\alpha, \rho_l)$, and then decreases back towards the projector value (F.32), with an overshoot $\mathrm{SNR}^2_l(t^{\star}_l) \geq \mathrm{SNR}^2_l(\infty)$ that grows with ρ_l. The mechanism is transparent in (F.28): increasing t_l suppresses the inter-pattern crosstalk (the term $\propto C$), but, because the kernel acts on the empirical correlation Σ_l, the projector limit $\mathbb{A}_l \to \Sigma_l^{-1}$ inverts, hence amplifies, the finite-sample fluctuation δ_l of the archetypes, re-injecting noise into the retrieval direction. The balance between crosstalk removal and sampling-noise amplification is optimised at $t^{\star}_l$; over-consolidation ($t_l > t^{\star}_l$) is the signal-to-noise counterpart of the *overfitting* regime identified, for the regularised/early-stopped dreaming kernel, in Agliari et al. (2024b). Only in the noiseless limit $\rho_l \to 0$ does $t^{\star}_l \to +\infty$ and the dependence becomes monotone, recovering the projector optimum. Consequently the maximal tolerable entropy $\rho_c(t_l)$ likewise rises to a peak at $t^{\star}_l$ and relaxes to $\rho_c(\infty)$, root of $\alpha(1+\rho^2) + \rho = 1$; the widest retrieval region – and the largest data saving – is reached at finite $t^{\star}_l$, not at $t_l \to +\infty$.

The exact moments (F.24)–(F.29) are the one-step (retrieval-initialised, $\beta \to +\infty$) shadow of the effective field (28): identifying m_l with the resolvent functional that fixes $\bar{D}_l$ and $\bar{p}_l$, the signal μ^{Dream}_1 matches $r_l v_l$ with $v_l = \beta(1+t_l)\bar{p}_l/(\Gamma_l \bar{D}_l)$, while $\beta^2(1+t_l)^2(q^{(2)}_l - q_l^2)/t_l^2$ matches the dreaming-noise variance $\sigma^2_l = \alpha\beta\bar{p}_l/\bar{D}_l^2$. As t_l grows the crosstalk noise contracts while the signal is largely preserved, so (F.30) increases up to the optimal sleeping time $t^{\star}_l$ before the sampling-noise amplification of the projector limit sets in; the resulting non-monotone retrieval boundary is in quantitative agreement with the self-consistency Eqs. (30)–(34).

F.6. The full DLAM at finite sleeping time

Combining the kernel-free inter-layer sector of Appendix F.3 with the finite-t_l dreaming sector above yields the complete one-step description

of the DLAM. The signal is additive (the LAM term carries no kernel),

$$\mu_1^{(l)}(t_l) = \beta\,\bar{\mu}_l(t_l), \qquad \bar{\mu}_l(t_l) := \frac{c_S(t_l)}{1+\rho_l} + \sum_{m\neq l} \frac{g_{lm}}{\sqrt{(1+\rho_l)(1+\rho_m)}}, \tag{F.33}$$

with $c_S(t_l)$ the dreaming signal form factor of (F.29). Because the dreaming and LAM sectors read neuron i through the common factor $\hat{\xi}^1_{i,l}$, their condensed-mode (example-interference) fluctuations add coherently on the total signal, whereas the two interference channels – the kernel-filtered intra-layer crosstalk and the inter-layer crosstalk – are statistically independent. Writing the dreaming variance (F.24) as the sum of a coherent part $\rho_l(\mu_1^{\mathrm{Dream}})^2$ and a residual crosstalk $\alpha\beta^2\,C(t_l)$, with intra-layer crosstalk form factor

$$C(t_l) := \frac{1}{\alpha\beta^2}\Big[\mu_2^{\mathrm{Dream}} - (\mu_1^{\mathrm{Dream}})^2 - \rho_l(\mu_1^{\mathrm{Dream}})^2\Big],$$
$$C(0) = 1, \quad C(\infty) = (1-\alpha)\,\delta_l = (1-\alpha)\frac{\rho_l}{1+\rho_l}, \tag{F.34}$$

a quantity decreasing monotonically in t_l from unity (it is the dreaming-filtered analogue of the inter-pattern interference $\alpha\beta^2$ of (F.13)), the total field variance is

$$\mu_2^{(l)} - (\mu_1^{(l)})^2 = \rho_l\,(\mu_1^{(l)})^2 + \alpha\beta^2\Big[C(t_l) + \sum_{m\neq l} g_{lm}^2\Big], \tag{F.35}$$

which we have confirmed against Monte Carlo simulation of the full three-layer field. The one-step magnetisation and the joint learning threshold then read

$$m_l^{(2)}(t_l) = \mathrm{erf}\left(\frac{\bar{\mu}_l(t_l)}{\sqrt{2\big[\rho_l\,\bar{\mu}_l(t_l)^2 + \alpha\big(C(t_l) + \sum_{m\neq l} g_{lm}^2\big)\big]}}\right), \tag{F.36}$$

$$\rho_l + \frac{\alpha\big[C(t_l) + \sum_{m\neq l} g_{lm}^2\big]}{\bar{\mu}_l(t_l)^2} < 1 \quad \text{for all } l. \tag{F.37}$$

At $t_l = 0$ one has $c_S = 1$, $C = 1$, and (F.37) collapses exactly to the bare-model threshold (F.16). As t_l increases, $C(t_l)$ decreases (intra-layer crosstalk suppressed) while $\bar{\mu}_l(t_l)$ stays $\mathcal{O}(1)$ (signal preserved), so the left-hand side decreases and the retrieval region expands with the sleeping time – up to the layer-wise optimal $t_l^\star$ beyond which over-consolidation slightly re-narrows it (cf. the over-consolidation discussion above) – on every layer simultaneously, including the uninformed ones in the disentanglement protocol, for which $C(t_l)$ controls how clean the intra-layer basins must be before the relayed inter-layer signal $\sum_{m\neq l} g_{lm}/\sqrt{(1+\rho_l)(1+\rho_m)}$ can drive recovery. A further consequence of the cooperative signal is a saturation in the number of layers: since the example-interference term scales coherently as $\rho_l(\mu_1^{(l)})^2$ while the inter-layer crosstalk grows only linearly in L, the signal-to-noise ratio increases with the number of coupled layers but is bounded above by $1/\rho_l$, so adding modalities never degrades retrieval and its benefit saturates once the cooperative signal dominates the residual sampling noise.

F.7. One-step reading of the disentanglement cue

The same machinery clarifies why disentanglement is a heteroassociative task. Under the simplex cue (13) a layer l that receives negligible direct evidence (vanishing cue weight $\alpha_l \to 0$ in (13), not to be confused with the load α) starts essentially uninformed, $\widehat{M}_l \simeq 0$, so its autoassociative signal v_l is absent and the first moment (F.11) is carried entirely by the relayed term $\beta\sum_{m\neq l} g_{lm}/\sqrt{(1+\rho_l)(1+\rho_m)}$. This term can exceed the fluctuation (F.35) only if the intra-layer basins are clean enough, that is, only after the dreaming kernel has suppressed the intra-layer crosstalk $\alpha\,C(t_l)$ of (F.34). The one-step criterion (F.37) thus predicts that the joint-success region expands from the simplex barycentre toward its vertices as t_l grows, with a collection of decoupled dreaming layers ($g_{lm} = 0$, $m \neq l$) yielding $\bar{\mu}_l = 0$ and hence zero success on the uninformed layers – precisely the behaviour reported in Fig. 2 and the rationale for benchmarking dreaming+LAM against LAM-only.

References

Agliari, E., Alemanno, F., Aquaro, M., Barra, A., Durante, F., & Kanter, I. (2024a). Hebbian dreaming for small datasets. *Neural Networks*, *173*, 106174.

Agliari, E., Alemanno, F., Aquaro, M., & Fachechi, A. (2024b). Regularization, early-stopping and dreaming: A Hopfield-like setup to address generalization and overfitting. *Neural Networks*, *177*, 106389.

Agliari, E., Alemanno, F., Barra, A., & Fachechi, A. (2019). Dreaming neural networks: Rigorous results. *Journal of Statistical Mechanics: Theory and Experiment*, *2019*(8), 083503.

Agliari, E., Alessandrelli, A., Barra, A., Centonze, M. S., & Ricci-Tersenghi, F. (2025a). Generalized hetero-associative neural networks. *Journal of Statistical Mechanics: Theory and Experiment*, *2025*(1), 013302.

Agliari, E., Alessandrelli, A., Barra, A., Centonze, M. S., & Ricci-Tersenghi, F. (2025b). Networks of neural networks: More is different. arXiv:2501.16789.

Agliari, E., Barra, A., Longo, C., & Tantari, D. (2017). Neural networks retrieving Boolean patterns in a sea of Gaussian ones. *Journal of Statistical Physics*, *168*(5), 1085–1104.

Agliari, E., Fachechi, A., & Marullo, C. (2022). Nonlinear PDEs approach to statistical mechanics of dense associative memories. *Journal of Mathematical Physics*, *63*(10), 103304.

Agliari, E., Mourao, P., Fachechi, A., & Vivo, P. (2026). Spectral criteria for generalization in unsupervised Hebbian nets. arXiv preprint, 2606.04651.

Alemanno, F., Aquaro, M., Kanter, I., Barra, A., & Agliari, E. (2023). Supervised Hebbian learning. *Europhysics Letters*, *141*(1), 11001.

Alessandrelli, A., Barra, A., Ladiana, A., Lepre, A., & Ricci-Tersenghi, F. (2025a). Beyond disorder: Unveiling cooperativeness in multidirectional associative memories. arXiv:2503.04454. Preprint.

Alessandrelli, A., Barra, A., Ladiana, A., Lepre, A., & Ricci-Tersenghi, F. (2025b). Supervised and unsupervised protocols for hetero-associative neural networks. *Physica A: Statistical Mechanics and its Applications*, *676*, 130871.

Alhulayil, M., & López-Benítez, M. (2019). Methods for the allocation of almost blank subframes with fixed duty cycle for improved LTE-U/Wi-Fi coexistence. In *2019 International Conference on Wireless Networks and Mobile Communications (WINCOM)* (pp. 1–6). IEEE.

Amit, D. J., Gutfreund, H., & Sompolinsky, H. (1985). Storing infinite numbers of patterns in a spin-glass model of neural networks. *Physical Review Letters*, *55*(14), 1530.

Bardes, A., Ponce, J., & LeCun, Y. (2022). VICReg: Variance-invariance-covariance regularization for self-supervised learning. In *ICLR 2022-International conference on learning representations*.

Barker, G. R. I., & Warburton, E. C. (2020). Multi-level analyses of associative recognition memory: The whole is greater than the sum of its parts. *Current Opinion in Behavioral Sciences*, *32*, 80–87.

Barra, A., Bernacchia, A., Santucci, E., & Contucci, P. (2012). On the equivalence of Hopfield networks and Boltzmann machines. *Neural Networks*, *34*, 1–9.

Benedetti, M., Carillo, L., Marinari, E., & Mézard, M. (2024). Eigenvector dreaming. *Journal of Statistical Mechanics: Theory and Experiment*, *2024*(1), 013302.

Crick, F., & Mitchison, G. (1983). The function of dream sleep. *Nature*, *304*(5922), 111–114.

Denis, D., Mylonas, D., Poskanzer, C., Bursal, V., Payne, J. D., & Stickgold, R. (2021). Sleep spindles preferentially consolidate weakly encoded memories. *The Journal of Neuroscience*, *41*, 4088–4099. https://doi.org/10.1523/jneurosci.0818-20.2021

Diekelmann, S., & Born, J. (2010). The memory function of sleep. *Nature Reviews Neuroscience*, *11*(2), 114–126.

Fachechi, A., Agliari, E., & Barra, A. (2019). Dreaming neural networks: Forgetting spurious memories and reinforcing pure ones. *Neural Networks*, *112*, 24–40.

Frankland, P. W., & Bontempi, B. (2005). The organization of recent and remote memories. *Nature Reviews Neuroscience*, *6*(2), 119–130.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, *79*(8), 2554–2558.

Hopfield, J. J., Feinstein, D. I., & Palmer, R. G. (1983). 'Unlearning' has a stabilizing effect in collective memories. *Nature*, *304*(5922), 158–159.

Ji, D., & Wilson, M. A. (2007). Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nature Neuroscience*, *10*(1), 100–107.

Kamsma, T. M., Kim, J., Kim, K., Boon, W. Q., Spitoni, C., Park, J., & van Roij, R. (2024). Brain-inspired computing with fluidic iontronic nanochannels. *Proceedings of the National Academy of Sciences*, *121*, e2320242121.

Kanter, I., & Sompolinsky, H. (1987). Associative recall of memory without errors. *Physical Review A*, *35*(1), 380. https://doi.org/10.1103/PhysRevA.35.380

Klinzing, J. G., Niethard, N., & Born, J. (2019). Mechanisms of systems memory consolidation during sleep. *Nature Neuroscience*, *22*(10), 1598–1610.

Kobayashi, M., Hattori, M., & Yamazaki, H. (2002). Multidirectional associative memory with a hidden layer. *Systems and Computers in Japan*, *33*(6), 1–9.

Krotov, D., & Hopfield, J. J. (2016). Dense associative memory for pattern recognition. *Advances in neural information processing systems*, *29*.

Krotov, D., & Hopfield, J. J. (2021). Large associative memory problem in neurobiology and machine learning. In *International Conference on Learning Representations (ICLR)*. https://arxiv.org/abs/2008.06996.

Leonelli, F. E., Agliari, E., Albanese, L., & Barra, A. (2021). On the effective initialisation for restricted Boltzmann machines via duality with Hopfield model. *Neural Networks*, *143*, 314–326. https://doi.org/10.1016/j.neunet.2021.06.017

Marinari, E., Rossi, S., & Zamponi, F. (2026). Dreaming improves memorization in a Hopfield model with bounded synaptic strength. arXiv:2603.09384.

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*(3), 419–457.

Ramsauer, H., Schäfl, B., Lehner, J., Seidl, P., Widrich, M., Adler, T., Gruber, L., Holzleitner, M., Pavlović, M., Sandve, G. K., Greiff, V., Kreil, D., Kopp, M., Klambauer, G., Brandstetter, J., & Hochreiter, S. (2020). Hopfield networks is all you need. arXiv:2008.02217. Preprint; presented at ICLR 2021.

Rasch, B., & Born, J. (2013). About sleep's role in memory. *Physiological Reviews*, *93*(2), 681–766.

Rissman, J., & Wagner, A. D. (2012). Distributed representations in memory: Insights from functional brain imaging. *Annual Review of Psychology*, *63*, 101–128.

Rolls, E. T. (2013). The mechanisms for pattern completion and pattern separation in the hippocampus. *Frontiers in Systems Neuroscience*, *7*, 74.

Serricchio, L., Bocchi, D., Chilin, C., Marino, R., Negri, M., Cammarota, C., & Ricci-Tersenghi, F. (2025). Daydreaming Hopfield networks and their surprising effectiveness on correlated data. *Neural Networks*, *186*, 107216.

Tononi, G., & Cirelli, C. (2006). Sleep function and synaptic homeostasis. *Sleep Medicine Reviews*, *10*(1), 49–62.

Xiang, J., Zhang, J., Zheng, R., Li, X., & Li, M. (2021). NIDM: Network impulsive dynamics on multiplex biological network for disease-gene prediction. *Briefings in Bioinformatics*, *22*, bbab080.

Yassa, M. A., & Stark, C. E. L. (2011). Pattern separation in the hippocampus. *Trends in Neurosciences*, *34*(10), 515–525.

Yu, F., Wang, X., Guo, R., Ying, Z., Cai, S., & Jin, J. (2025). Dynamical analysis, hardware implementation, and image encryption application of new 4D discrete hyperchaotic maps based on parallel and cascade memristors. *Integration*, *102*, 102475.