

Fermi-Bose Machine

Mingshan Xie¹, Yuchen Wang¹, and Haiping Huang^{1,2*}

¹PMI Lab, School of Physics, Sun Yat-sen University,
Guangzhou 510275, People's Republic of China and

²Guangdong Provincial Key Laboratory of Magnetoelectric Physics and Devices,
Sun Yat-sen University, Guangzhou 510275, People's Republic of China

(Dated: April 23, 2024)

Distinct from human cognitive processing, deep neural networks trained by backpropagation can be easily fooled by adversarial examples. To design a semantically meaningful representation learning, we discard backpropagation, and instead, propose a local contrastive learning, where the representation for the inputs bearing the same label shrink (akin to boson) in hidden layers, while those of different labels repel (akin to fermion). This layer-wise learning is local in nature, being biological plausible. A statistical mechanics analysis shows that the target fermion-pair-distance is a key parameter. Moreover, the application of this local contrastive learning to MNIST benchmark dataset demonstrates that the adversarial vulnerability of standard perceptron can be greatly mitigated by tuning the target distance, i.e., controlling the geometric separation of prototype manifolds.

Introduction.— Deep neural networks (DNNs) are widely used in current artificial intelligence studies, including Chat GPT and generative diffusion models, producing remarkable revolution in many fields of science and technology [1–3]. These networks are commonly trained by backpropagation algorithms, which are criticized to be far from biological intelligence [4], mainly because a global cost-function should be defined and each weight update requires propagating errors. In addition, a serious caveat of DNNs is their adversarial vulnerability [5, 6], i.e., a human-imperceptible perturbation to an input can completely change the classification of that input. There do not yet exist effective ways to control this vulnerability. One standard strategy in practice is the adversarial training which takes adversarial examples as a part of training data [7]. But a trade-off between standard accuracy and robustness might be inevitable [8, 9].

We resolve these two challenges by a unified principle formulated as a statistical mechanics problem. Semantically meaningful representations must emerge during a good representation learning [10], as commonly observed in information disentangling process in the ventral stream of visual cortex [11, 12]. We argue that this semantic hierarchy of concepts can be learned by designing a contrastive Hamiltonian. The data pair belonging to the same class forms a boson pair, as expected from physics that the Euclidean distance between the pair must be minimized, while the pair belonging to different classes forms a fermion pair, and the distance must be enlarged (a physics intuition). Both types of distance construct the Hamiltonian, which can be trained layer by layer and thus backpropagating a global error is not needed. This principle, namely Fermi-Bose machine (FBM), realizes geometry separation of internal representations in the latent space of DNNs, bearing similarity with recent empirical observation of hierarchical nucleation in DNNs [13], latent space clustering [14] and emergence of compact latent space [15, 16].

Remarkably, *without* adversarial training, our FBM demonstrates the ability to mitigate the adversarial vulnerability by tuning only a target fermion-pair distance, aligning with a recent hypothesis of relationship between data concentration and adversarial robust classifier [17]. Our work shows that the geometrically separated representations encode semantic hierarchy that facilitates *both* discrimination and robustness against class-preserving perturbations, and moreover, this can be realized by a local learning, amenable to a statistical mechanics analysis.

Model.— We consider a classification task implemented by a deep network with L layers. N_ℓ ($\ell = 1, \dots, L$) denotes the layer size. $L - 2$ hidden layers are individually learned by our FBM. We first consider a single layer perceptron of $N-1$ structure ($N_1 = N$ and $N_2 = 1$) for statistical mechanical analysis, which can be straightforward to extend to more layers and multiple labels. In the toy theoretical model, the data of Gaussian mixture are considered. We then test FBM in a three layer network for classifying real benchmark dataset. In the toy model, we analyze only the representation learning, while in learning real dataset, we add a final readout layer to leverage well-separated representations for classification.

The network parameters are specified by all weight matrices between layers. The weight value from neuron j at the upstream layer $\ell - 1$ to neuron i at the downstream layer ℓ is denoted by w_{ij}^ℓ . The activation of neuron i in the layer ℓ is given by $h_i^\ell = \phi(z_i^\ell)$, where $z_i^\ell = \sum_{j=1}^{N_{\ell-1}} w_{ij}^\ell h_j^{\ell-1}$, and $\phi(\cdot)$ is the nonlinear function $\tanh(\cdot)$. In the toy model setting, the hidden layer has a single neuron, and thus the subscript of z in the following analysis indicates which input of the fermion (or boson) pair.

To learn a meaningful representation at each layer, a layer-wise training is carried out. The inputs are decomposed into boson pair and fermion pair (this is only a physics metaphor but not talking about real




FIG. 1: Schematic illustration of FBM. A network with structure 3-4-3-1 is used to learn geometrically separated representations, which become increasingly compact per class, such that the final linear separation is achievable. A layer-wise training is carried out, and thus the learning is local.

physical particles). The dataset is defined by $\mathcal{D} \equiv \{(\mathbf{x}_1^\mu, \mathbf{x}_2^\mu, \sigma^\mu)\}_{\mu \in [P]}$, where $\mathbf{x}_1^\mu, \mathbf{x}_2^\mu \in \mathbb{R}^N$ and $\sigma^\mu \in \{0, 1\}$. $\sigma^\mu = 1$ if $\mathbf{x}_1^\mu, \mathbf{x}_2^\mu$ belong to the same category (boson pair), while $\sigma^\mu = 0$ if $\mathbf{x}_1^\mu, \mathbf{x}_2^\mu$ belong to different categories (fermion pair). The loss function can thus be constructed below:

$$\mathcal{L} = \sum_{\mu=1}^P \frac{1}{2} [\sigma^\mu D_\mu^2 + (1 - \sigma^\mu) \varphi(d_F - D_\mu^2)] + \frac{\lambda_w}{2} \|\mathbf{w}\|_2^2, \quad (1)$$

where λ_w is a weight regularization parameter, $\sigma^\mu = \frac{1+y_1^\mu y_2^\mu}{2}$, y_1^μ and y_2^μ are the labels of \mathbf{x}_1^μ and \mathbf{x}_2^μ , respectively, and $D_\mu^2 = [\phi(z_1^\mu) - \phi(z_2^\mu)]^2$ is the Euclidean distance. $\varphi(\cdot)$ denotes a relu-like function. The first term encourages contraction of boson pairs, while the second term repels fermion pairs, and d_F sets the target distance of fermion pairs. An illustration of this idea is given in Fig. 1. One can also add the activity regularization into the above loss function for more biological sparse activation. The layer-wise Hamiltonian in Eq. (1) can also be adapted to unsupervised learning where positive pairs (the transformation of one input and the input itself) and negative pairs (two different inputs from the training batch) are used [18].

For the sake of statistical mechanics analysis, we use the Gaussian mixture data, which is defined by $P(\mathbf{x}_1, \mathbf{x}_2, y_1, y_2) = P(\mathbf{x}_1|y_1)P(\mathbf{x}_2|y_2)P(y_1, y_2)$, where $P(y_1, y_2) = \frac{1}{2}\rho\delta(y_1 - y_2) + \frac{1}{2}(1 - \rho)\delta(y_1 + y_2)$. ρ specifies the fraction of boson pairs. $P(\mathbf{x}|y) = \mathcal{N}\left(\frac{my}{N}\mathbf{1}, \frac{\Delta^2}{N}\mathbf{I}\right)$, where $\mathbf{1}$ is an all-one vector, \mathbf{I} is an identity matrix, and $y \in \{\pm 1\}$ for two classes. m and Δ determine the mean and variance of each Gaussian distribution.

Statistical mechanics analysis.— We focus on the limit of large number of training pairs and input dimension, while keeping the ratio $\alpha = \frac{P}{N}$ fixed. The partition function of the model reads,

$$Z(\beta, \mathcal{D}, \lambda_w, d_F) = \int d\mathbf{w} e^{-\beta \mathcal{L}(\mathcal{D}, \mathbf{w}, \lambda_w, d_F)}, \quad (2)$$

where β denotes the inverse temperature, the partition function is data-dependent, and thus the free energy must

be averaged over the data ($\langle \cdot \rangle$ below), which is highly non-trivial and we have to use the replica method [19, 20]:

$$-\beta f = \langle \ln Z \rangle = \lim_{n \rightarrow 0, N \rightarrow \infty} \frac{\ln \langle Z^n \rangle}{nN}, \quad (3)$$

where f is the free energy density, and we have assumed the limit of zero replica number and thermodynamic limit are exchangeable [19].

A lengthy calculation [21] shows that the free energy density depends on physically relevant order parameters $\{M, q, Q\}$ together with the conjugate counterparts $\{\hat{M}, \hat{q}, \hat{Q}\}$ as follows:

$$-\beta f = \text{Extr}_{M, q, Q, \hat{M}, \hat{q}, \hat{Q}} \hat{M}M + \hat{q}q - \frac{\hat{Q}Q}{2} + \Psi_S(\hat{Q}, \hat{q}, \hat{M}) + \Psi_E(Q, q, M), \quad (4)$$

where the entropy term $\Psi_S(\hat{Q}, \hat{q}, \hat{M}) = \frac{1}{2} \ln \frac{2\pi}{2\hat{q} + 2\beta\lambda_w - \hat{Q}} + \frac{\hat{M}^2 - \hat{Q}}{2(2\hat{q} + 2\beta\lambda_w - \hat{Q})}$, and the energy term $\Psi_E(Q, q, M) = \alpha \mathbb{E}_{\mathbf{y}} \int D\mathbf{v} \ln \int D\mathbf{u} e^{-\beta \mathcal{H}(y_1, y_2, z_1(y_1, u_1, v_1), z_2(y_2, u_2, v_2))}$, where $\mathbf{y} = (y_1, y_2)$, $\mathbf{v} = (v_1, v_2)$, $\mathbf{u} = (u_1, u_2)$, Dv is a Gaussian measure, and the effective Hamiltonian $\mathcal{H} = \frac{1}{2} \left[\frac{1+y_1 y_2}{2} D^2 + \frac{1-y_1 y_2}{2} \varphi(d_F - D^2) \right]$, where $D^2 = [\phi(z_1) - \phi(z_2)]^2$, and the pre-activation z_1 and z_2 can be reparametrized by the order parameters [21]. To get Eq. (4), the replica symmetry (RS) ansatz is used [19, 20]. M characterizes the mean of learned weight values; q denotes the variance of weight values; Q denotes the overlap between two typical weight values sampled from different replicas (equivalent to states in physics [19]). This minimal set of order parameters obeys the following iterative equations [21]:

$$\begin{aligned} \hat{q} &= \frac{\alpha\beta\Delta^2}{2} \mathbb{E}_{\mathbf{y}, \mathbf{v}} \langle \langle \mathcal{H}_{11} + \mathcal{H}_{22} - \beta(\mathcal{H}_1^2 + \mathcal{H}_2^2) \rangle \rangle, \\ \hat{Q} &= -\alpha\beta^2\Delta^2 \mathbb{E}_{\mathbf{y}, \mathbf{v}} [\langle \langle \mathcal{H}_1 \rangle \rangle^2 + \langle \langle \mathcal{H}_2 \rangle \rangle^2], \\ \hat{M} &= \alpha\beta M \mathbb{E}_{\mathbf{y}, \mathbf{v}} [y_1 \langle \langle \mathcal{H}_1 \rangle \rangle + y_2 \langle \langle \mathcal{H}_2 \rangle \rangle], \\ q &= \left(2\hat{q} - \hat{Q} + \beta\lambda_w \right)^{-1} + \frac{\hat{M}^2 - \hat{Q}}{\left(2\hat{q} - \hat{Q} + \beta\lambda_w \right)^2}, \\ Q &= \frac{\hat{M}^2 - \hat{Q}}{\left(2\hat{q} - \hat{Q} + \beta\lambda_w \right)^2}, \\ M &= -\frac{\hat{M}}{\left(2\hat{q} - \hat{Q} + \beta\lambda_w \right)}, \end{aligned} \quad (5)$$

where $\langle \langle \bullet \rangle \rangle = \frac{\int D\mathbf{u} e^{-\beta \mathcal{H}} \bullet}{\int D\mathbf{u} e^{-\beta \mathcal{H}}}$ denotes the thermal average under the effective Hamiltonian, $\mathcal{H}_i = \frac{\partial \mathcal{H}}{\partial z_i}$, and \mathcal{H}_{11} and \mathcal{H}_{22} are corresponding second derivatives. Focusing on the ground state, the zero temperature limit can be taken [21].

To support the replica symmetry ansatz, we also derive belief propagation equation to infer the weights by using cavity method [19, 20]. We first define the cavity probability $P_{i \rightarrow \mu}(w_i)$ in the absence of one data pair μ , while $\hat{P}_{\nu \rightarrow i}(w_i)$ denotes the probability contribution from neighboring weights of pair ν except i . These two probabilities obey the following self-consistent equations under the cavity approximation in disordered system theory [19, 20]:

$$\begin{aligned} P_{i \rightarrow \mu}(w_i) &\propto e^{-\frac{\beta \lambda_w w_i^2}{2}} \prod_{\nu \in \partial i \setminus \mu} \hat{P}_{\nu \rightarrow i}(w_i), \\ \hat{P}_{\nu \rightarrow i}(w_i) &= \int \prod_{j \in \partial \nu \setminus i} [dw_j P_{j \rightarrow \nu}(w_j)] P_\nu(\mathbf{y}^\nu, \mathbf{z}^\nu), \end{aligned} \quad (6)$$

where ∂i ($\partial \mu$) indicates neighbors of node i (μ), $\mathbf{z}^\nu = (z_1^\nu, z_2^\nu)$ are the pre-activations triggered by $(\mathbf{x}_1^\nu, \mathbf{x}_2^\nu)$, and $P_\nu = e^{-\beta \mathcal{L}_\nu}$ (\mathcal{L}_ν is the ν -th pair related loss). Making a Gaussian assumption for $P_{i \rightarrow \mu} \sim \mathcal{N}(m_{i \rightarrow \mu}, v_{i \rightarrow \mu})$ and taking a large N limit (see also the previous work [22] for committee machine), we get a simplified belief propagation (BP) that is used in our experiments.

$$\begin{aligned} \omega_{\mu \rightarrow i} &= \sum_{j \in \partial \mu \setminus i} m_{j \rightarrow \mu} \mathbf{x}_j^\mu, \\ \mathbf{V}_{\mu \rightarrow i} &= \sum_{j \in \partial \mu \setminus i} v_{j \rightarrow \mu} \mathbf{x}_j^\mu (\mathbf{x}_j^\mu)^\top, \\ B_{\nu \rightarrow i} &= (\mathbf{x}_i^\nu)^\top \mathbf{f}_\nu, \\ A_{\nu \rightarrow i} &= -(\mathbf{x}_i^\nu)^\top \frac{\partial \mathbf{f}_\nu}{\partial \omega_{\nu \rightarrow i}} \mathbf{x}_i^\nu, \\ m_{i \rightarrow \mu} &= \frac{\sum_{\nu \neq \mu} B_{\nu \rightarrow i}}{\beta \lambda_w + \sum_{\nu \neq \mu} A_{\nu \rightarrow i}}, \\ v_{i \rightarrow \mu} &= \frac{1}{\beta \lambda_w + \sum_{\nu \neq \mu} A_{\nu \rightarrow i}}, \end{aligned} \quad (7)$$

where $\mathbf{x}_i^\mu = [x_{1,i}^\mu, x_{2,i}^\mu]^\top$, ω is similarly defined, and a modified measure $\tilde{P}_\mu = \frac{1}{\tilde{Z}_\mu} e^{-\frac{1}{2}(\mathbf{z}^\mu - \omega_{\mu \rightarrow i})^\top \mathbf{V}_{\mu \rightarrow i}^{-1}(\mathbf{z}^\mu - \omega_{\mu \rightarrow i})} P_\mu$, from which we derive $\mathbf{f}_\mu \equiv \frac{\partial \ln \tilde{Z}_\mu}{\partial \omega_{\mu \rightarrow i}}$. Detailed derivations are given in SM [21].

We next show that our theory confirms computational power of FBM (Fig. 2). As the training data size increases, the test representation loss [Eq. (1)] decreases first rapidly and then slowly, indicating that the output representation space is well separated given sufficient data size. The separation is well supported by the simulation results in the insets of Fig. 2 (b). In particular, the theory is in an excellent agreement with algorithmic results of BP running on single instances of finite sizes. As expected, the increase of Gaussian variance (Δ) makes the clustering hard, and thus the test generalization loss grows sharply [Fig. 2 (b)]. In this case, multiple layers must be added and learned.




FIG. 2: Generalization loss ϵ_g of FBM as a function of α (a) and Δ (b). $m = 1, \rho = 0.5, d_F = 1$, and $\lambda_w = 0.05$. Theory matches well experiments ($N = 200$). Error bars are computed from five independent runs. $\Delta = 0.5$ in (a). $\alpha = 2.5$ and $\beta = 50$ in (b). The insets from bottom to top describe the distribution of 10 pairs of data (fermion and boson pairs) in the output representation space learned by BP with $\Delta = 0.2, 0.7$, and 1.5 , respectively. Blue points indicate the class $y = -1$, and red points indicate $y = 1$. (c) Test accuracy of FBM vs d_F . (d) Fermion pair distance vs d_F (dashed line indicates equality). $\alpha = 2.5, \Delta = 0.5$ in (c,d).

The fermion-pair distance is predicted to increase non-linearly with d_F , as also holds for the classification accuracy ($\langle \Theta(-yz(\mathbf{w}, \mathbf{x})) \rangle_{\mathbf{w}, P(\mathbf{x}, y)}$, see details at SM [21]) [Fig. 2 (c,d)]. This prediction qualitatively matches the following experimental results on more complex networks and datasets. We next apply FBM to real data learning and show how this framework resolves the conflict between generalization and robustness. In particular, by tuning d_F , the FBM can outperform backpropagation when the data size is not sufficient, and even achieve better adversarial robustness.

Accuracy and adversarial robustness in practical learning.—We compare FBM with multilayer perceptron (MLP) trained using backpropagation in this section. The FBM is trained layer by layer with the contrastive cost (Eq. (1), but the preactivation of each neuron becomes a vector because of the input data-pair [21]), and only the readout layer is trained by the cross entropy of the classification. In contrast, MLP is trained jointly by a single cost of cross entropy.

Compared with MLP, FBM with local learning has better generalization accuracy [Fig. 3 (a)], while the geometry of hidden representation space is well controlled [Fig. 3 (b)]. Interestingly, by applying principal component analysis (PCA) to the hidden representation space, we find that the neural activity is approximately distributed on the surface of a sphere [a section displayed in Fig. 3 (c)] in the case of FBM, while the MLP training

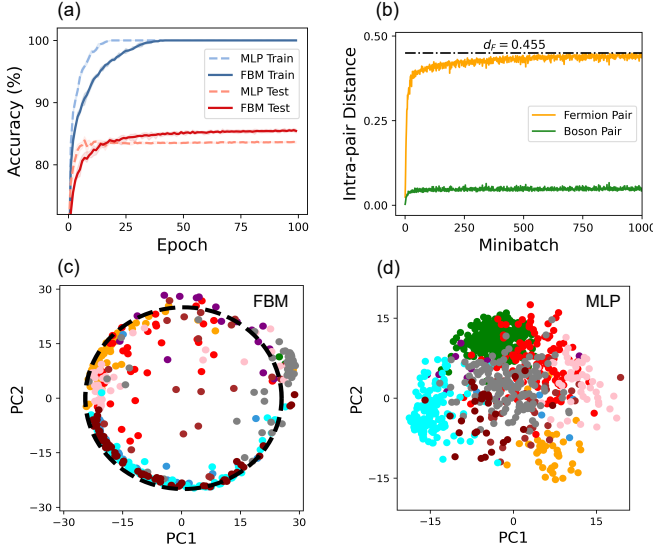


FIG. 3: Performance comparison between FBM and MLP. (a) Training and test trajectories on the MNIST dataset [23]. The fluctuation is estimated from five independent runs. The network structure is specified by 784-1000-10, trained by 1000 handwritten digits and tested by 4000 digits. $d_F = 0.465$. Stochastic gradient descent is used with mini-batch size of 50. (b) The distance of fermion and boson pairs during training. Other parameters are the same as (a). (c,d) The two-dimensional section of three-dimensional PCA projection of the hidden representation learned by FBM and MLP. The first three eigenvectors of PCA are considered with explained variance ratio r . (c) FBM ($r = 74.6\%$). (d) MLP ($r = 54.5\%$). Different colors encode different classes. $d_F = 1.2$.

has no such intriguing geometric properties [Fig. 3(d)]. This observation demonstrates that the local contrastive learning leads to low-dimensional disentangled representations, resembling a similar process in visual cortical hierarchy [11, 12]. FBM thus bears the computational benefit of learning semantically meaningful clusters, and each cluster acts as a prototype of noisy sensory inputs.

The target distance d_F plays a key role, as predicted by the theory. By varying d_F , we find that the generalization accuracy displays double ascent phenomenon [Fig. 4 (a)], and we also find that the fermion-pair distance first increases approximately linearly with d_F but gets saturated when $d_F > 0.465$ [Fig. 4 (b)]. This critical threshold may vary with dataset and architecture details, but this phenomenon is qualitatively the same across different settings (e.g., changing width of hidden layer and training data size), which is also consistent with the theory [Fig. 2]. The profile of double peaks may depend on real structured dataset. The double ascent phenomenon implies that there exists an optimal representation geometry, which is intimately related to adversarial robustness.

The standard classifier can be easily fooled by an im-




FIG. 4: Generalization accuracy tuned by d_F in FBM. Error bars are the fluctuation with five independent runs. The dash line marks the performance of MLP. The point S_1 marks the distance where the accuracy of FBM first matches that of MLP; The point P_1 marks the first peak; The point V marks the local minimum; The point P_2 marks the second peak (the best performance); The point S_2 marks the place where the accuracy of FBM drops down to that of MLP; The point B_1 marks the stationary part. The bottom figure shows the fermion-pair distance as a function of d_F , where the grey dash line indicates an identity function.

perceptible perturbation (adversarial attack). The standard training learns highly predictive yet non-robust features by focusing on spurious correlations between input and output [24–27]. We argue that a semantically meaningful feature representation should be learned before the classification of last layer. This representation can be formed by our local contrastive learning, resulting in geometrically separated prototype clusters, such that decision boundaries are low density areas. We consider fast gradient sign method (FGSM) and white noise attacks to the trained neural network [6, 7], i.e., the perturbation to the input $\delta\mathbf{x} = \epsilon \text{sign}(\nabla_{\mathbf{x}}\mathcal{C})$, and $\epsilon\xi$, respectively. \mathcal{C} is the cross entropy cost, ϵ is the attack strength, and independent $\xi_i \sim \mathcal{N}(0, 1)$. We find that d_F is a key parameter to control the adversarial robustness, as predicted by the theory [inset of Fig. 5 (c)]. In a narrow range of d_F , the robustness measured by the area under the accuracy- ϵ curve (Fig. 5) increases sharply for both types of attack. For FGSM, the robustness gets saturated while the Gaussian noise attack exhibits a peak. This demonstrates that FBM without adversarial training is more adversarial robust than MLP, providing a principled route toward alignment with human cognition processing using disentangled invariant representations [11, 12, 28, 29].

Conclusion.—A fundamental question in machine

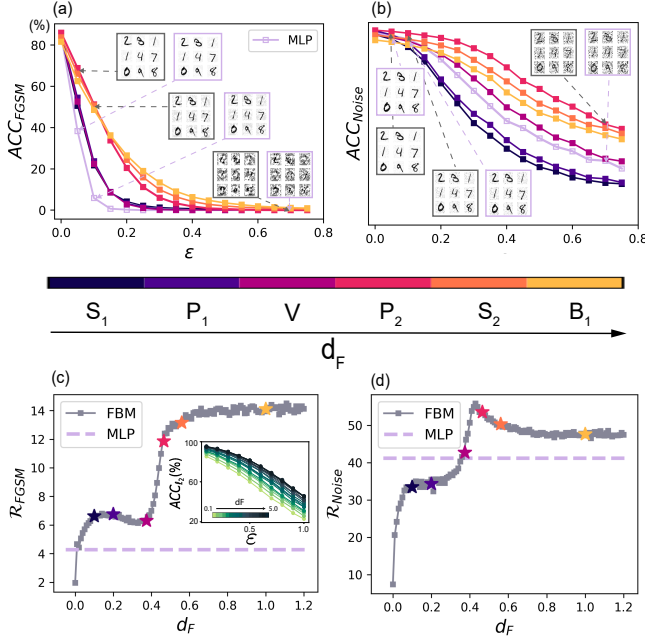


FIG. 5: Robustness properties of FBM and MLP. The line color indicates increasing value of d_F corresponding to marked points in Fig. 4. (a) FGSM attack. (b) Gaussian white noise attack. The insets show adversarially perturbed inputs. (c,d) Adversarial robustness measured by the area under the accuracy- ϵ curves. The inset in (c) shows a theoretical prediction of ℓ_2 norm attack [21].

learning and neuroscience is what is a good representation learning that is biological plausible and moreover adversarial robust. We propose a local contrastive learning, namely Fermi-Bose machine to resolve this challenging question. Our theoretical toy model based on Gaussian mixture data captures main characteristics of FBM, explaining effects of data density, entanglement intensity and target fermion-pair distance. When applied to MNIST data learning, FBM exhibits intriguing geometric separation, double ascent of accuracy and a great improvement of adversarial robustness over joint training of MLP. Many future directions arise, such as connection to supervised Hebbian learning [30], precise mathematical mechanism of robustness versus representation geometry, generalization to unsupervised learning and so on.

ACKNOWLEDGMENTS

H.H. benefited from the international workshop of machine learning physics held in Kyoto University (Nov 16-18, 2023). This research was supported by the National Natural Science Foundation of China for Grant number 12122515, and Guangdong Provincial Key Laboratory of Magnetoelectric Physics and De-

vices (No. 2022B1212010008), and Guangdong Basic and Applied Basic Research Foundation (Grant No. 2023B1515040023).

* Electronic address: huanghp7@mail.sysu.edu.cn

- [1] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, 2016).
- [2] S. Bubeck, V. Chandrasekaran, R. Eldan, J. A. Gehrke, E. Horvitz, E. Kamar, P. Lee, Y. T. Lee, Y.-F. Li, S. M. Lundberg, et al., arXiv:2303.12712 (2023).
- [3] T. Brooks, B. Peebles, C. Holmes, W. DePue, Y. Guo, L. Jing, D. Schnurr, J. Taylor, T. Luhman, E. Luhman, et al. (2024), URL <https://openai.com/research/video-generation-models-as-world-simulators>.
- [4] T. P. Lillicrap, A. Santoro, L. Marris, C. J. Akerman, and G. Hinton, *Nature Reviews Neuroscience* **21**, 335 (2020).
- [5] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, in *ICLR 2014 : International Conference on Learning Representations (ICLR) 2014* (2014).
- [6] I. J. Goodfellow, J. Shlens, and C. Szegedy, in *ICLR 2015 : International Conference on Learning Representations 2015* (2015).
- [7] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, in *International Conference on Learning Representations* (2018).
- [8] D. Tsipras, S. Santurkar, L. Engstrom, A. Turner, and A. Madry, arXiv:1805.12152 (2018).
- [9] A. Javanmard and M. Soltanolkotabi, arXiv:2010.11213 (2020).
- [10] H. Huang, arXiv:2306.11232 (2023).
- [11] J. J. DiCarlo and D. D. Cox, *Trends in Cognitive Sciences* **11**, 333 (2007).
- [12] J. J. DiCarlo, D. Zoccolan, and N. C. Rust, *Neuron* **73**, 415 (2012).
- [13] D. Doimo, A. Glielmo, A. Ansuini, and A. Laio, in *Advances in Neural Information Processing Systems*, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin (Curran Associates, Inc., 2020), vol. 33, pp. 7526–7536.
- [14] K. Kamnitsas, D. Castro, L. L. Folgoc, I. Walker, R. Tanno, D. Rueckert, B. Glocker, A. Criminisi, and A. Nori, in *Proceedings of the 35th International Conference on Machine Learning*, edited by J. Dy and A. Krause (PMLR, 2018), vol. 80 of *Proceedings of Machine Learning Research*, pp. 2459–2468.
- [15] Z. Jiang, J. Zhou, and H. Huang, *Chinese Phys. B* **30**, 048702 (2021).
- [16] C. Li and H. Huang, *Phys. Rev. Res.* **5**, L022011 (2023).
- [17] A. Pal, J. Sulam, and R. Vidal, arXiv:2309.16096 (2023).
- [18] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, in *Proceedings of the 37th International Conference on Machine Learning*, edited by H. D. III and A. Singh (PMLR, 2020), vol. 119 of *Proceedings of Machine Learning Research*, pp. 1597–1607.
- [19] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).
- [20] H. Huang, *Statistical Mechanics of Neural Networks* (Springer, Singapore, 2022).
- [21] See the supplemental material at <http://...> for experi-

- mental details and theoretical derivations.
- [22] B. Aubin, A. Maillard, J. Barbier, F. Krzakala, N. Macris, and L. Zdeborová, *Journal of Statistical Mechanics: Theory and Experiment* **2019**, 124023 (2019).
 - [23] Y. LeCun, The MNIST database of handwritten digits, retrieved from <http://yann.lecun.com/exdb/mnist>.
 - [24] A. Shafahi, W. R. Huang, C. Studer, S. Feizi, and T. Goldstein, arXiv:1809.02104 (2018).
 - [25] A. Fawzi, H. Fawzi, and O. Fawzi, *Advances in neural information processing systems* **31**, 1186–1195 (2018).
 - [26] A. Ilyas, S. Santurkar, D. Tsipras, L. Engstrom, B. Tran, and A. Madry, arXiv:1905.02175 (2019).
 - [27] R. Geirhos, J.-H. Jacobsen, C. Michaelis, R. Zemel, W. Brendel, M. Bethge, and F. A. Wichmann, *Nature Machine Intelligence* **2**, 665 (2020).
 - [28] H. Huang, *Phys. Rev. E* **98**, 062313 (2018).
 - [29] J. Zhou and H. Huang, *Phys. Rev. E* **103**, 012315 (2021).
 - [30] F. Alemanno, M. Aquaro, I. Kanter, A. Barra, and E. Agliari, *Europhysics Letters* **141**, 11001 (2023).
 - [31] D. P. Kingma and J. Ba, arXiv:1412.6980 (2014).
 - [32] M. Xie, Y. Wang, and H. Huang (2024), <https://...>
 - [33] J. T. Barron, arXiv:2112.11687 (2021).

Algorithmic details of FBM training

The deep network we consider consists of L layers, with $L-2$ hidden layers. Because the network is trained layer by layer, we consider training the weights between ℓ^{th} and $(\ell+1)^{th}$ layers as an example to detail the training method. The ℓ^{th} layer of the network has N_ℓ neurons. The connection weight between the neuron i in the $(\ell+1)^{th}$ layer and the neuron j in the ℓ^{th} layer is denoted by $w_{ij}^{\ell+1}$.

Firstly, we introduce how to construct fermion or boson pairs for training. The pairing process is exhaustive and does not miss any possible pairs. For instance, we have a total of 10 images in a minibatch, consisting of two images of handwritten digit zero and eight images of handwritten digit one. We thus have sixteen fermion pairs and twenty-nine boson pairs. Subsequently, we assign pair-labels to boson pairs as $(+1)$ and to fermion pairs as (-1) . This label setting is consistent with that introduced in the main text. Note that this label is not the original label of digit images, but used to define the cost function below.

For the sake of compactness, we define the pre-activation in the form of vector:

$$z_j^\ell = \begin{pmatrix} z_{L,j}^\ell \\ z_{R,j}^\ell \end{pmatrix}, \quad (8)$$

where we use L and R to distinguish which element of the data pair triggers the network output. In the standard training of MLP, z_j^ℓ is a scalar quantity, but in FBM, it is a vector as defined above. The pre-activations in the next layer can be estimated according to the following affine transformation:

$$z_j^{\ell+1} = \sum_j w_{ij}^{\ell+1} \begin{pmatrix} h_{L,j}^\ell \\ h_{R,j}^\ell \end{pmatrix}, \quad (9)$$

where the activation \mathbf{h}_j^ℓ is also a vector whose components $h_{X,j}^\ell = \phi(z_{X,j}^\ell)$ ($X = L, R$), and in practice we use $\phi(x) = \frac{1}{2}[\tanh(x) + 1]$, but other choices are also possible, and we would not discuss the effects of transfer function here.

The loss function to minimize for each data pair is expressed below:

$$\mathcal{L} = \frac{1+\sigma}{4}D^2 + \frac{1-\sigma}{4}\varphi(d_F - D^2) + \lambda_w\|\mathbf{w}\|^2, \quad (10)$$

where the pair label $\sigma \in \{-1, 1\}$, and $D^2 = \|\mathbf{h}_L^{\ell+1} - \mathbf{h}_R^{\ell+1}\|_2^2$. The first term of boson pair distance should be minimized, and the second term is related to the fermion pair distance, which must reach the target separation d_F . $\varphi(\cdot)$ is some relu-like function. The loss function is layerwisely minimized by stochastic gradient descent (update rules of weights are detailed below). The full training data is first divided into mini-batches each of which contains a few randomly sampled images. These images inside one mini-batches are paired according the aforementioned rule. After a layer is trained, the weights before this layer are all frozen, and then the next layer is trained following a similar procedure. All layers except the last readout layer are trained by minimizing the loss function in in Eq. (10). The readout layer is trained by minimizing a cross-entropy cost function for classification we consider in this paper. The test accuracy is estimated from a test data set which is fed into the trained network one by one. Hyper-parameters for the network and algorithm are summarized in the Table I.

Next, we detail how the stochastic gradient descent works, which is divided into two coupled steps. The first forward

Parameter	Value	Description
N_0	784	input dimension
N_1	1000	hidden layer width
N_3	10	the number of output units
M	1000	training dataset size
\mathcal{B}	50	mini-batch size
λ_w	0.01	weight decay strength

TABLE I: Hyperparameters used in training FBM to learn MNIST.

propagation step is summarized as follows:

$$\begin{aligned} \mathbf{z}_j^\ell &= \begin{pmatrix} z_{L,j}^\ell \\ z_{R,j}^\ell \end{pmatrix} = \sum_k w_{jk}^\ell \begin{pmatrix} h_{L,k}^{\ell-1} \\ h_{R,k}^{\ell-1} \end{pmatrix}, \\ \mathbf{h}_j^\ell &= \phi \begin{pmatrix} z_{L,j}^\ell \\ z_{R,j}^\ell \end{pmatrix}. \end{aligned} \quad (11)$$

The second backward propagation proceeds as follows:

$$\begin{aligned} \Delta_j^\ell &\equiv \frac{\partial \mathcal{L}}{\partial \mathbf{z}_j^\ell} = \begin{pmatrix} \frac{\partial \mathcal{L}}{\partial z_{L,j}^\ell} \\ \frac{\partial \mathcal{L}}{\partial z_{R,j}^\ell} \end{pmatrix} = \begin{pmatrix} \frac{\partial \mathcal{L}}{\partial h_{L,j}^\ell} \\ \frac{\partial \mathcal{L}}{\partial h_{R,j}^\ell} \end{pmatrix} \odot \begin{pmatrix} \frac{\partial h_{L,j}^\ell}{\partial z_{L,j}^\ell} \\ \frac{\partial h_{R,j}^\ell}{\partial z_{R,j}^\ell} \end{pmatrix} = \begin{pmatrix} \frac{\partial \mathcal{L}}{\partial h_{L,j}^\ell} \\ \frac{\partial \mathcal{L}}{\partial h_{R,j}^\ell} \end{pmatrix} \odot \begin{pmatrix} \phi'(z_{L,j}^\ell) \\ \phi'(z_{R,j}^\ell) \end{pmatrix}, \\ \delta_{jk}^\ell &\equiv \frac{\partial \mathcal{L}}{\partial w_{jk}^\ell} = \frac{\partial \mathcal{L}}{\partial \mathbf{z}_j^\ell} \frac{\partial \mathbf{z}_j^\ell}{\partial w_{jk}^\ell} = (\mathbf{h}_k^{\ell-1})^\top \Delta_j^\ell, \end{aligned} \quad (12)$$

where \odot represents the element-wise product. We remark here that this is just a gradient descent without propagating a global error from the top layer, thereby unlike backpropagation in standard deep learning. The derivative $\frac{\partial \mathcal{L}}{\partial \mathbf{h}_j^\ell}$ in Eq. (12) can be directly computed using the definition of the contrastive loss \mathcal{L} . The weight is then updated as follows:

$$w_{jk}^\ell \leftarrow w_{jk}^\ell - \eta \delta_{jk}^\ell, \quad (13)$$

where η is the learning rate, set by default values of the Adam algorithm [31]. Codes realizing the FBM training are available in our GitHub page [32] (to be released upon formal publication of the paper), which also includes two kinds of adversarial attacks.

Theoretical details of replica calculation

Finite temperature

To get a quantitative description of local contrastive learning, we semi-rigorous analyze FBM in a simple setting of N -1 structure, as already detailed in the main text. This statistical mechanics analysis can be carried out by application of replica method, a seminal method in disordered system theory [19]. The disorder in our model comes from the design of Gaussian mixture data, which is detailed as follows.

The training data is given by $\{(\mathbf{x}_1^\mu, y_1^\mu), (\mathbf{x}_2^\mu, y_2^\mu)\}_{\mu \in [P]}$, where $\mathbf{x}_1^\mu, \mathbf{x}_2^\mu \in \mathbb{R}^N$, and y_1^μ or y_2^μ is a scalar. Each μ is an input data-pair, which can be a boson or fermion pair. The statistics of each pair $P(\mathbf{x}_1^\mu, y_1^\mu, \mathbf{x}_2^\mu, y_2^\mu) = P(y_1^\mu, y_2^\mu) P(\mathbf{x}_1^\mu | y_1^\mu) P(\mathbf{x}_2^\mu | y_2^\mu)$, where

$$\begin{aligned} P(\mathbf{x}_1^\mu | y_1^\mu) &= \mathcal{N}\left(y_1^\mu \mathbf{1}_N \frac{m}{N}, \frac{\Delta^2}{N} \mathbf{I}_N\right), \\ P(\mathbf{x}_2^\mu | y_2^\mu) &= \mathcal{N}\left(y_2^\mu \mathbf{1}_N \frac{m}{N}, \frac{\Delta^2}{N} \mathbf{I}_N\right), \\ P(y_2^\mu, y_1^\mu) &= P(y_1^\mu) P(y_2^\mu | y_1^\mu) = \frac{1}{2} \rho \delta(y_1^\mu - y_2^\mu) + \frac{1}{2} (1 - \rho) \delta(y_1^\mu + y_2^\mu), \end{aligned} \quad (14)$$

where $\delta(\cdot)$ is a Dirac delta function, $\mathbf{1}_N$ is an all-one vector, and $\mathbf{I}_N \in \mathbb{R}^{N \times N}$ is an identity matrix.

The single-node output is given by $h^\mu = \phi(z^\mu) = \phi(\sum_i w_i x_i^\mu)$. The loss is given below:

$$\mathcal{L} = \sum_{\mu=1}^P \frac{1}{2} [\sigma^\mu D_\mu^2 + (1 - \sigma^\mu) \varphi(d_F - D_\mu^2)] + \frac{\lambda_w}{2} \|\mathbf{w}\|_2^2, \quad (15)$$

where λ_w is a weight regularization parameter commonly used in deep learning, $\sigma^\mu = \frac{1+y_1^\mu y_2^\mu}{2}$, y_1^μ and y_2^μ are the labels of \mathbf{x}_1^μ and \mathbf{x}_2^μ , respectively, and $D_\mu^2 = [\phi(z_1^\mu) - \phi(z_2^\mu)]^2$ is the Euclidean distance. $z_1(z_2)$ is the preactivation for the first (second) element of the data pair input. $\varphi(\cdot)$ denotes a relu-like function, and we choose $\varphi(x) = \frac{x + \sqrt{x^2 + a}}{2}$ ($a = 0.01$) for function smoothness at the origin point [33].

The partition function for our model reads,

$$Z = \int \prod_{i=1}^N dw_i P_0(w_i) \exp \left[-\beta \sum_{\mu=1}^P \mathcal{L}^\mu(\mathbf{w}) \right], \quad (16)$$

where $\mathcal{L}^\mu \equiv \frac{1}{2} [\sigma^\mu D_\mu^2 + (1 - \sigma^\mu) \varphi(d_F - D_\mu^2)]$, $\mathbf{w} \in \mathbb{R}^N$, $P_0(w_i) = e^{-\frac{\beta \lambda_w}{2} w_i^2}$ ($\forall i$) corresponding to weight regularization commonly used in deep learning. In our case, the partition function is a random variable, resulting in a random free energy in physics. We must do an average of the free energy over Gaussian mixture data. We have thus to use replica method as follows:

$$-\beta f = \frac{\mathbb{E}_{\mathbf{X}_1, \mathbf{y}_1, \mathbf{X}_2, \mathbf{y}_2} \ln Z}{N} = \lim_{n \rightarrow 0} \frac{\ln \left(\mathbb{E}_{\mathbf{X}_1, \mathbf{y}_1, \mathbf{X}_2, \mathbf{y}_2} Z^n \right)}{Nn}, \quad (17)$$

where f is free energy per synapse (density), n is a replica number first taken to an integer and finally sent to zero (a bit weird but exact in some cases [20]), $\mathbf{X}_1 \in \mathbb{R}^{N \times P}$ whose μ -th column $\mathbf{X}_1^\mu = \mathbf{x}_1^\mu$ (\mathbf{X}_2 has a similar meaning), and $\mathbf{y} \in \mathbb{R}^P$ (the subscript 1 (or 2) indicates which element of an input data-pair). In the following, we use \mathbb{E} to represent the average over the data disorder.

Now, we explicitly write down the integer power of the random partition function averaged over the data pairs.

$$\begin{aligned} \mathbb{E} Z^n &= \sum_{\{\mathbf{y}_1, \mathbf{y}_2\}} \int \prod_{\mu} [d\mathbf{x}_1^\mu d\mathbf{x}_2^\mu] \prod_{\mu} \left[P(y_1^\mu, y_2^\mu) \prod_i [P((\mathbf{x}_1^\mu)_i | y_1^\mu) P((\mathbf{x}_2^\mu)_i | y_2^\mu)] \right] \int \prod_{a=1}^n d\mathbf{w}^a P_0(\mathbf{w}^a) \\ &\times \prod_{a=1}^n \prod_{\mu} \exp \left\{ -\beta \left[\frac{\sigma^\mu}{2} (D_\mu^a)^2 + \frac{1 - \sigma^\mu}{2} \varphi(d_F - (D_\mu^a)^2) \right] \right\}, \end{aligned} \quad (18)$$

where $D_\mu^a = \sqrt{[\phi(\mathbf{w}^a \mathbf{x}_1^\mu) - \phi(\mathbf{w}^a \mathbf{x}_2^\mu)]^2}$, and a is the index of replicas ($a = 1, \dots, n$).

To further simplify the calculation, we define the following conditional probability distributions:

$$\begin{aligned} P(z_1^{a\mu} | y_1^\mu) &= \langle \delta(z_1^{a\mu} - \mathbf{w}^a \mathbf{x}_1^\mu) \rangle_{\mathbf{x}_1^\mu | y_1^\mu}, \\ P(z_2^{a\mu} | y_2^\mu) &= \langle \delta(z_2^{a\mu} - \mathbf{w}^a \mathbf{x}_2^\mu) \rangle_{\mathbf{x}_2^\mu | y_2^\mu}. \end{aligned} \quad (19)$$

After a straightforward computation, we get Gaussian distributions for the preactivations as follows:

$$\begin{aligned} P(\mathbf{z}_1 | \mathbf{y}_1) &= \prod_{\mu} P(\mathbf{z}_1^\mu | y_1^\mu) \\ &= \frac{1}{\sqrt{(2\pi)^{nP} (\det \Delta^2 \mathbf{Q})^P}} \prod_{\mu} \exp \left[(\mathbf{z}_1^\mu - m y_1^\mu \mathbf{M})^T (2\Delta^2 \mathbf{Q})^{-1} (\mathbf{z}_1^\mu - m y_1^\mu \mathbf{M}) \right], \end{aligned} \quad (20)$$

and

$$\begin{aligned} P(\mathbf{z}_2 | \mathbf{y}_2) &= \prod_{\mu} P(\mathbf{z}_2^\mu | y_2^\mu) \\ &= \frac{1}{\sqrt{(2\pi)^{nP} (\det \Delta^2 \mathbf{Q})^P}} \prod_{\mu} \exp \left[(\mathbf{z}_2^\mu - m y_2^\mu \mathbf{M})^T (2\Delta^2 \mathbf{Q})^{-1} (\mathbf{z}_2^\mu - m y_2^\mu \mathbf{M}) \right], \end{aligned} \quad (21)$$

where the preactivations $\mathbf{z}_1^\mu, \mathbf{z}_2^\mu \in \mathbb{R}^n$, and $\mathbf{M} \in \mathbb{R}^n, \mathbf{Q} \in \mathbb{R}^{n \times n}$ are mean and correlation of weights in the replica space. Their more precise expressions are given below.

$$\mathbf{M} = \begin{pmatrix} \frac{\sum_i w_i^a}{N} \\ \vdots \\ \frac{\sum_i w_i^a}{N} \\ \vdots \\ \frac{\sum_i w_i^n}{N} \end{pmatrix}, \quad \mathbf{Q} = \begin{pmatrix} \frac{\sum_i w_i^1 w_i^1}{N} & \dots & \frac{\sum_i w_i^1 w_i^n}{N} \\ \vdots & \ddots & \vdots \\ \frac{\sum_i w_i^n w_i^1}{N} & \dots & \frac{\sum_i w_i^n w_i^n}{N} \end{pmatrix}. \quad (22)$$

One can verify \mathbf{M} and \mathbf{Q} are related to the mean and variance of the above Gaussian distribution of preactivations. This is demonstrated below.

$$\begin{aligned} \langle \mathbf{z}_1^\mu \rangle_{\mathbf{x}_1^\mu | y_1^\mu} &= m y_1^\mu \mathbf{M}, \\ \langle \mathbf{z}_2^\mu \rangle_{\mathbf{x}_2^\mu | y_2^\mu} &= m y_2^\mu \mathbf{M}, \\ \langle \mathbf{z}_1^\mu (\mathbf{z}_1^\nu)^\top \rangle_{\mathbf{x}_1^\mu | y_1^\mu} - \langle \mathbf{z}_1^\mu \rangle_{\mathbf{x}_1^\mu | y_1^\mu} \langle \mathbf{z}_1^\nu \rangle_{\mathbf{x}_1^\mu | y_1^\mu}^\top &= \Delta^2 \mathbf{Q} \delta_{\mu\nu}, \\ \langle \mathbf{z}_2^\mu (\mathbf{z}_2^\nu)^\top \rangle_{\mathbf{x}_2^\mu | y_2^\mu} - \langle \mathbf{z}_2^\mu \rangle_{\mathbf{x}_2^\mu | y_2^\mu} \langle \mathbf{z}_2^\nu \rangle_{\mathbf{x}_2^\mu | y_2^\mu}^\top &= \Delta^2 \mathbf{Q} \delta_{\mu\nu}, \\ \langle \mathbf{z}_2^\mu (\mathbf{z}_1^\nu)^\top \rangle_{\mathbf{x}_2^\mu | y_2^\mu, \mathbf{x}_1^\mu | y_1^\mu} - \langle \mathbf{z}_2^\mu \rangle_{\mathbf{x}_2^\mu | y_2^\mu} \langle \mathbf{z}_1^\nu \rangle_{\mathbf{x}_1^\mu | y_1^\mu}^\top &= 0, \end{aligned} \quad (23)$$

where the last equality is derived due to the fact that given the label y , the data sample \mathbf{x} is independently generated. \mathbf{M} and \mathbf{Q} can then be used to represent the Gaussian distribution, which would greatly simplify the replica calculation. Hence, we can insert Dirac delta functions to specify \mathbf{M} and \mathbf{Q} , and then use their Fourier representations:

$$\begin{aligned} \prod_a \delta \left(N M^a - \sum_i w_i^a \right) &= \int_{-i\infty}^{+i\infty} \prod_a \frac{d\hat{M}^a}{(2\pi i)} \exp \left[\hat{M}^a \left(N M^a - \sum_i w_i^a \right) \right], \\ \prod_a \delta \left(N Q^{aa} - \sum_i w_i^a w_i^a \right) &= \int_{-i\infty}^{+i\infty} \prod_a \frac{d\hat{Q}^{aa}}{(2\pi i)} \exp \left[\hat{Q}^{aa} \left(N Q^{aa} - \sum_i w_i^a w_i^a \right) \right], \\ \prod_{a < b} \delta \left(N Q^{ab} - \sum_i w_i^a w_i^b \right) &= \int_{-i\infty}^{+i\infty} \prod_{a < b} \frac{d\hat{Q}^{ab}}{(2\pi i)} \exp \left[\hat{Q}^{ab} \left(N Q^{ab} - \sum_i w_i^a w_i^b \right) \right]. \end{aligned} \quad (24)$$

Therefore,

$$\begin{aligned} \mathbb{E} Z^n &= \int_{-i\infty}^{+i\infty} \left(\prod_{a \leq b} \frac{d\hat{Q}^{ab} dQ^{ab}}{(2\pi i)} \right) \left(\prod_a \frac{d\hat{M}^a dM^a}{(2\pi i)} \right) \\ &\times \int \prod_a^n d\mathbf{w}^a \prod_a^n P_0(\mathbf{w}^a) \exp \left(N \sum_{a < b} (\hat{Q}^{ab} Q^{ab}) - \sum_{a < b} \hat{Q}^{ab} \sum_i w_i^a w_i^b \right) \\ &\times \exp \left(N \sum_a \hat{M}^a M^a - \sum_a \hat{M}^a \sum_i w_i^a \right) \exp \left(N \sum_a \hat{Q}^{aa} Q^{aa} - \sum_a \hat{Q}^{aa} \sum_i w_i^a w_i^a \right) \\ &\times \prod_{\mu}^P \sum_{\{y_1^\mu, y_2^\mu\}} \int d\mathbf{z}_1^\mu d\mathbf{z}_2^\mu P(y_1^\mu, y_2^\mu) P(\mathbf{z}_2^\mu | y_2^\mu) P(\mathbf{z}_1^\mu | y_1^\mu) \\ &\times \prod_{\mu, a} \exp \left[-\beta \left(\frac{\sigma^\mu}{2} (D_\mu^a)^2 + \frac{1 - \sigma^\mu}{2} \varphi \left[d_F - (D_\mu^a)^2 \right] \right) \right]. \end{aligned} \quad (25)$$

Finally, the integer power of the partition function averaged over data can be recast into a mathematically compact form:

$$\mathbb{E} Z^n = \int \left(\prod_{a \leq b} \frac{d\hat{Q}^{ab} dQ^{ab}}{2\pi i} \right) \left(\prod_a \frac{d\hat{M}^a dM^a}{2\pi i} \right) e^{-N\beta f}, \quad (26)$$

where $-\beta f$ is an action in physics and can be expressed as the sum of the following three terms:

$$\begin{aligned}\Psi_0 &= \sum_a^n \hat{M}^a M^a + \sum_a^n \hat{Q}^{aa} Q^{aa} + \sum_{a<b} \hat{Q}^{ab} Q^{ab}, \\ \Psi_S &= \frac{1}{N} \ln \int \prod_a^n d\mathbf{w}^a \prod_a^n P_0(\mathbf{w}^a) \exp \left(- \sum_a^n \hat{Q}^{aa} \sum_i w_i^a w_i^a - \sum_{a<b}^{n(n-1)/2} \hat{Q}^{ab} \sum_i w_i^a w_i^b - \sum_a^n \hat{M}^a w_i^a \right), \\ \Psi_E &= \alpha \ln \sum_{\{y_1, y_2\}} \int d\mathbf{z}_1 d\mathbf{z}_2 P(y_1, y_2) P(\mathbf{z}_2 | y_2) P(\mathbf{z}_1 | y_1) \prod_a^n \exp \left[-\beta \left(\frac{\sigma}{2} (D^a)^2 + \frac{1-\sigma}{2} \varphi [d_F - (D^a)^2] \right) \right],\end{aligned}\quad (27)$$

where $(D^a)^2 = [\phi(z_1^a) - \phi(z_2^a)]^2$, and $\sigma = \frac{1+y_1 y_2}{2}$.

To proceed, we have to assume the replica symmetry (RS) ansatz specified as follows:

$$\begin{aligned}\hat{M}^a &= \hat{M}, \quad M^a = M \quad \forall a, \\ \hat{Q}^{aa} &= \hat{q}, \quad Q^{aa} = q \quad \forall a, \\ \hat{Q}^{ab} &= \hat{Q}, \quad Q^{ab} = Q \quad \forall a \neq b.\end{aligned}\quad (28)$$

In essence, RS assumes permutation symmetry of the replica overlap matrix, which is the first level of approximation that should be cross-checked by numerical experimental results. If this ansatz is not correct, the resulting saddle-point equation (SDE, see below) does not converge during iterations. Then an advanced level of approximation is required [20]. Under the RS ansatz, the replica overlap matrices can be expressed as follows:

$$\mathbf{M} = \begin{pmatrix} M \\ \vdots \\ M \\ \vdots \\ M \end{pmatrix}, \quad \mathbf{Q} = \begin{pmatrix} q & Q & \cdots & Q \\ Q & q & \cdots & Q \\ \vdots & \vdots & \ddots & \vdots \\ Q & \cdots & q \end{pmatrix}.\quad (29)$$

Applying the RS ansatz to the entropy term Ψ_S , we obtain

$$\begin{aligned}\Psi_S &= \frac{1}{N} \ln \int \prod_a^n d\mathbf{w}^a \prod_a^n P_0(\mathbf{w}^a) \exp \left(- \sum_a^n \hat{q} \sum_i w_i^a w_i^a - \sum_{a<b}^N \hat{Q} \sum_i w_i^a w_i^b - \sum_a^n \hat{M} \sum_i w_i^a \right) \\ &= \frac{1}{N} \ln \int \prod_a^n d\mathbf{w}^a \prod_a^n P_0(\mathbf{w}^a) \exp \left(- \frac{\sum_a^n (2\hat{q} - \hat{Q}) \sum_i w_i^a w_i^a}{2} - \frac{\hat{Q}}{2} \sum_i \left(\sum_a^n w_i^a \right)^2 - \sum_a^n \hat{M} \sum_i w_i^a \right) \\ &= \frac{1}{N} \ln \int \prod_i D\xi_i \int \prod_{a,i}^n d w_i^a \prod_{a,i} P_0(w_i^a) \exp \left(- \frac{\sum_i^n (2\hat{q} - \hat{Q}) \sum_i^N w_i^a w_i^a}{2} + \sum_i^N \left(\xi_i \sqrt{-\hat{Q}} \sum_a^n w_i^a \right) - \sum_a^n \hat{M} \sum_i^N w_i^a \right) \\ &= \ln \int D\xi \int \prod_a^n d w^a \exp \left(- \frac{\sum_a^n (2\hat{q} - \hat{Q} + \beta \lambda_w) (w^a)^2}{2} + \left(\xi \sqrt{-\hat{Q}} \sum_a^n w^a \right) - \sum_a^n \hat{M} w^a \right),\end{aligned}\quad (30)$$

where $D\xi \equiv \prod_i D\xi_i$ is a Gaussian measure, and the integral identity $\int D\xi e^{b\xi} = e^{b^2/2}$ is used.

Next we take the limit of $n \rightarrow 0$, and obtain the limit of $\Psi_S(n)/n$ as the new Ψ_S which reads,

$$\begin{aligned}
\Psi_S &= \lim_{n \rightarrow 0} \frac{1}{n} \ln \int D\xi \left[\int dw \exp \left(-\frac{(2\hat{q} - \hat{Q} + \beta\lambda_w) w^2}{2} + \left(\xi\sqrt{-\hat{Q}} - \hat{M} \right) w \right) \right]^n \\
&= \lim_{n \rightarrow 0} \frac{1}{n} \ln \int D\xi \left[\sqrt{\frac{2\pi}{2\hat{q} - \hat{Q} + \beta\lambda_w}} \exp \left[\left(\xi\sqrt{-\hat{Q}} - \hat{M} \right)^2 \left(2(2\hat{q} - \hat{Q} + \beta\lambda_w) \right)^{-1} \right] \right]^n \\
&= \int D\xi \ln \left[\sqrt{\frac{2\pi}{2\hat{q} - \hat{Q} + \beta\lambda_w}} \exp \left[\frac{\left(\xi\sqrt{-\hat{Q}} - \hat{M} \right)^2}{2(2\hat{q} - \hat{Q} + \beta\lambda_w)} \right] \right] \\
&= \frac{1}{2} \ln \frac{2\pi}{(2\hat{q} - \hat{Q} + \beta\lambda_w)} + \frac{\hat{M}^2 - \hat{Q}}{2(2\hat{q} - \hat{Q} + \beta\lambda_w)}.
\end{aligned} \tag{31}$$

Before taking the limit of Ψ_E , we parameterize the random pre-activation z_1^a and z_2^a as follows:

$$\begin{aligned}
z_1^a &= my_1 M + \Delta\sqrt{Q}v_1 + \Delta\sqrt{q - Q}u_1^a, \\
z_2^a &= my_2 M + \Delta\sqrt{Q}v_2 + \Delta\sqrt{q - Q}u_2^a,
\end{aligned} \tag{32}$$

where v_1, v_2, u_1^a, u_2^a are standard Gaussian random variables. The above parameterization obeys the statistics of (z_1^a, z_2^a) [see Eq. (23)]. Therefore, the energy term can be recast into the following form:

$$\begin{aligned}
\Psi_E &= \alpha \ln \sum_{\{y_1, y_2\}} P(y_1, y_2) \int Dv_1 Dv_2 \int \prod_a^n Du_1^a Du_2^a \\
&\quad \prod_a^n \exp \left[-\beta \left(\frac{\sigma}{2} (D^a)^2 + \frac{1-\sigma}{2} \varphi \left[d_F - (D^a)^2 \right] \right) \right],
\end{aligned} \tag{33}$$

where $D^a = \sqrt{[\phi(z_1^a) - \phi(z_2^a)]^2}$. After taking the limit $\lim_{n \rightarrow 0} \frac{\Psi_E(n)}{n}$, we get

$$\Psi_E(Q, q, M) = \alpha \sum_{\{y_1, y_2\}} P(y_1, y_2) \int Dv_1 Dv_2 \ln \int Du_1 Du_2 \exp[-\beta \mathcal{H}(Q, q, M)], \tag{34}$$

where the effective single variable Hamiltonian reads

$$\mathcal{H}(Q, q, M) = \frac{\sigma}{2} D^2 + \frac{1-\sigma}{2} \varphi(d_F - D^2), \tag{35}$$

where $\sigma = \frac{1+y_1 y_2}{2}$, $D^2 = [\phi(z_1) - \phi(z_2)]^2$, $z_x = my_x M + \Delta\sqrt{Q}v_x + \Delta\sqrt{q - Q}u_x$ ($x = 1, 2$).

Collecting all three contributions to the replica symmetric free energy, we get

$$-\beta f = \text{Extr}_{M, q, Q, \hat{M}, \hat{q}, \hat{Q}} \hat{M}M + \hat{q}q - \frac{\hat{Q}Q}{2} + \Psi_S(\hat{Q}, \hat{q}, \hat{M}) + \Psi_E(Q, q, M), \tag{36}$$

which has been shown in the main text.

The self-consistent equations the order parameters must obey are derived by taking the derivative $\frac{\partial[-\beta f]}{\partial \mathcal{O}}$ to be zero, where \mathcal{O} represents all order parameters and their conjugate counterparts. We skip the technical details here, but we remind the readers that to simplify the results, one has to use an integral identity $\int Dz z F(z) = \int Dz F'(z)$ (also named the Stein Lemma) which can be proved using integral by parts, and $F'(z)$ is an any differentiable function.

The resultant equations are called saddle-point equations (SDEs), which is summarized as follows.

$$\begin{aligned}
\hat{q} &= \frac{\alpha\beta\Delta^2}{2} \mathbb{E}_{\mathbf{y}, \mathbf{v}} \langle \langle \mathcal{H}_{11} + \mathcal{H}_{22} - \beta(\mathcal{H}_1^2 + \mathcal{H}_2^2) \rangle \rangle, \\
\hat{Q} &= -\alpha\beta^2\Delta^2 \mathbb{E}_{\mathbf{y}, \mathbf{v}} [\langle \langle \mathcal{H}_1 \rangle \rangle^2 + \langle \langle \mathcal{H}_2 \rangle \rangle^2], \\
\hat{M} &= \alpha\beta M \mathbb{E}_{\mathbf{y}, \mathbf{v}} [y_1 \langle \langle \mathcal{H}_1 \rangle \rangle + y_2 \langle \langle \mathcal{H}_2 \rangle \rangle], \\
q &= \left(2\hat{q} - \hat{Q} + \beta\lambda_w \right)^{-1} + \frac{\hat{M}^2 - \hat{Q}}{\left(2\hat{q} - \hat{Q} + \beta\lambda_w \right)^2}, \\
Q &= \frac{\hat{M}^2 - \hat{Q}}{\left(2\hat{q} - \hat{Q} + \beta\lambda_w \right)^2}, \\
M &= -\frac{\hat{M}}{\left(2\hat{q} - \hat{Q} + \beta\lambda_w \right)},
\end{aligned} \tag{37}$$

where $\mathbf{y} = (y_1, y_2)$, $\mathbf{v} = (v_1, v_2)$, $\mathbf{u} = (u_1, u_2)$, $\mathcal{H}_i = \frac{\partial \mathcal{H}}{\partial z_i}$, and \mathcal{H}_{11} and \mathcal{H}_{22} are corresponding second derivatives; $\langle \langle \bullet \rangle \rangle$ denotes the thermal average under the measure of effective Hamiltonian expressed as $\langle \langle \bullet \rangle \rangle = \frac{\int D\mathbf{u} e^{-\beta \mathcal{H}} \bullet}{\int D\mathbf{u} e^{-\beta \mathcal{H}}}$. The Hamiltonian derivatives are specified by $\mathcal{H}_1 = \frac{\partial \mathcal{H}}{\partial z_1}$, and $\mathcal{H}_2 = \frac{\partial \mathcal{H}}{\partial z_2}$, $\mathcal{H}_{11} = \frac{\partial^2 \mathcal{H}}{\partial z_1 \partial z_1}$ and $\mathcal{H}_{22} = \frac{\partial^2 \mathcal{H}}{\partial z_2 \partial z_2}$. Their explicit expressions are given below:

$$\begin{aligned}
\mathcal{H}_1 &= \sigma D\phi'_1 - D(1 - \sigma)\varphi'\phi'_1, \\
\mathcal{H}_2 &= -\sigma D\phi'_2 + (1 - \sigma)\varphi'D\phi'_2, \\
\mathcal{H}_{11} &= \sigma [(\phi'_1)^2 + D\phi''_1] - (1 - \sigma) \{ -2D^2\varphi''(\phi'_1)^2 + \varphi'[(\phi'_1)^2 + D\phi''_1] \}, \\
\mathcal{H}_{22} &= \sigma [(\phi'_2)^2 - D\phi''_2] + (1 - \sigma) \{ 2\varphi''D^2(\phi'_2)^2 - \varphi'[(\phi'_2)^2 - D\phi''_2] \},
\end{aligned} \tag{38}$$

where $\phi_x \equiv \phi(z_x)$ ($x = 1, 2$), $D = \phi_1 - \phi_2$, $\varphi'(x) = \frac{1}{2} \left(1 + \frac{x}{\sqrt{x^2 + a}} \right)$, $\varphi''(x) = \frac{a}{2(x^2 + a)^{\frac{3}{2}}}$, and the argument of φ' (or φ'') is $d_F - D^2$.

Generalization contrastive loss

The order parameters of M, q determine the following generalization contrastive loss:

$$\epsilon_g = \mathbb{E}_{\mathcal{D}} [\langle \mathcal{L}^\mu(\mathbf{w}^*) \rangle], \tag{39}$$

where $\langle \cdot \rangle$ means the average over the Boltzmann measure of weights, $\mathcal{L}^\mu = \frac{1}{2} [\sigma^\mu D_\mu^2 + (1 - \sigma^\mu)\varphi(d_F - D_\mu^2)]$, and $\mathbf{w}^* \in \mathbb{R}^N$ is the optimal weight configuration following the Boltzmann measure. We can then use the fact that given the label \mathbf{y} , the pre-activations of z_1 and z_2 follow a joint Gaussian distribution, and then the average over \mathbf{x} can be integrated out, leading to the following result:

$$\epsilon_g = \mathbb{E}_{y_1, y_2} \int Du_1 \int Du_2 \left[\frac{\sigma}{2} D^2(\phi_1, \phi_2) + \frac{1 - \sigma}{2} \varphi[d_F - D^2(\phi_1, \phi_2)] \right], \tag{40}$$

where $D^2(\phi_1, \phi_2) = [\phi_1(z_1) - \phi_2(z_2)]^2$, and $z_x \equiv my_x M + \Delta\sqrt{q}u_x$ ($x = 1, 2$).

Distance of fermion pair

As in the derivation of generalization contrastive loss, the distance of fermion pair can also be estimated in an analytic form:

$$D_F^2 = \int Du_1 Du_2 [\phi(mM + \Delta\sqrt{q}u_1) - \phi(-mM + \Delta\sqrt{q}u_2)]^2, \tag{41}$$

where we have assumed a fermion pair $(y_1, y_2) = (+1, -1)$ is considered, and the result is the same in the case of $(y_1, y_2) = (-1, +1)$.




FIG. 6: Generalization accuracy vs. the fraction of Boson pair. Other parameters are the same as in Fig. 2 in the main text.

Generalization accuracy

We use a sign function readout to measure how good the contrastive representation learning is. If the hidden representation (in our toy model this is one dimension latent space) is well separated, then the accuracy for a test data point can be defined as follows,

$$\epsilon_g^C = \langle \Theta[-yz(\mathbf{x}, \mathbf{w})] \rangle_{\mathbf{w}, \mathcal{D}}, \quad (42)$$

where $\langle \cdot \rangle$ means the average over the test data and the optimal weight configuration following the Boltzmann measure. Note that the pre-activation z is a Gaussian random variable $\mathcal{N}(z; Mmy, q\Delta^2)$. We also assume that the two clusters corresponding to two classes are correctly separated with corresponding labels ($y \rightarrow -y$ otherwise). We have then the following derivation:

$$\begin{aligned} \epsilon_g^C &= \mathbb{E}_y \int Du \Theta(-y(Mmy + \Delta\sqrt{q}u)) \\ &= \int Du \Theta(-Mm + \Delta\sqrt{q}u) \\ &= \frac{1}{2} \left[1 - \operatorname{erf} \left(\frac{Mm}{\Delta\sqrt{2q}} \right) \right], \end{aligned} \quad (43)$$

where $m = 1$ in our Gaussian mixture data setting, and C indicates classification. We can thus get the generalization accuracy as $1 - \epsilon_g^C$, which is plotted as a function of ρ in Fig. 6. The accuracy deteriorates sharply when the fraction of Boson pairs in the training dataset goes above a threshold.

Theory of adversarial attack on FBM

Because the FGSM attack leads to an order parameter $\Lambda = \sum_i |w_i|/N \leq \sqrt{q}$, we would not analyze the FGSM attack using the replica theory of FBM. Instead, we analyze the ℓ_2 norm attack defined as follows:

$$\delta \mathbf{x} = \epsilon \frac{\nabla_{\mathbf{x}} \mathcal{L}}{\sqrt{\sum_i (\nabla_{x_i} \mathcal{L})^2}}, \quad (44)$$

where ϵ is the attack strength. Note that FGSM is an ℓ_∞ norm attack [6, 7]. We conjecture that the theory of ℓ_2 norm attack would qualitatively predict behavior of a broad range of ℓ_p norm bounded attacks including FGSM. The attack is calculated from a mean-squared loss specified as follows:

$$\mathcal{L} = \frac{1}{2} \left[y - \phi \left(\sum_i w_i x_i \right) \right]^2. \quad (45)$$

Therefore, one can compute the increment of the preactivation caused by the attack, i.e., $\mathbf{w}^T \delta \mathbf{x}$, which reads

$$\begin{aligned}\mathbf{w}^T \delta \mathbf{x} &= \epsilon \mathbf{w}^T \frac{-[y - \phi(z)] \phi'(z) \mathbf{w}}{\sqrt{[-[y - \phi(z)] \phi'(z)]^2 \mathbf{w}^T \mathbf{w}}} \\ &= \epsilon \frac{\mathcal{K}_y(z) q}{\sqrt{\mathcal{K}_y^2(z) q}} \\ &= \epsilon \sqrt{q} \text{sgn}[\mathcal{K}_y(z)],\end{aligned}\tag{46}$$

where we have rescaled the attack strength to make the increment of the order one, keeping the same order with $z = \mathbf{w}^T \mathbf{x}$, and $\mathcal{K}_y(z) = -[y - \phi(z)] \phi'(z)$. $q \equiv \frac{\mathbf{w}^T \mathbf{w}}{N}$, which can be treated as the order parameter in our replica theory. The averaged adversarial accuracy as a function of attack strength is thus given by

$$\begin{aligned}\text{ACC}_{adv} &= 1 - \langle \Theta[-y(z + \mathbf{w}^T \delta \mathbf{x})] \rangle_{\mathbf{w}, x, y} \\ &= 1 - \langle \Theta(-yz - y\epsilon\sqrt{q}\text{sgn}[\mathcal{K}_y(z)]) \rangle_{z, y} \\ &= 1 - \frac{1}{2} \int Du \Theta(-mM - \Delta\sqrt{q}u - \epsilon\sqrt{q}\text{sgn}[\mathcal{K}_1(mM + \Delta\sqrt{q}u)]) \\ &\quad - \frac{1}{2} \int Du \Theta(-mM + \Delta\sqrt{q}u + \epsilon\sqrt{q}\text{sgn}[\mathcal{K}_{-1}(-mM + \Delta\sqrt{q}u)]) \\ &= \frac{1}{2} \int Du \Theta(mM + \Delta\sqrt{q}u + \epsilon\sqrt{q}\text{sgn}[\mathcal{K}_1(mM + \Delta\sqrt{q}u)]) \\ &\quad + \frac{1}{2} \int Du \Theta(mM - \Delta\sqrt{q}u - \epsilon\sqrt{q}\text{sgn}[\mathcal{K}_{-1}(-mM + \Delta\sqrt{q}u)]) \\ &= \frac{1}{2} \int Du \Theta(mM + \Delta\sqrt{q}u + \epsilon\sqrt{q}\text{sgn}[\mathcal{K}_1(mM + \Delta\sqrt{q}u)]) \\ &\quad + \frac{1}{2} \int Du \Theta(mM + \Delta\sqrt{q}u - \epsilon\sqrt{q}\text{sgn}[\mathcal{K}_{-1}(-mM - \Delta\sqrt{q}u)]) \\ &= \int Du \Theta(mM + \Delta\sqrt{q}u + \epsilon\sqrt{q}\text{sgn}[\mathcal{K}_1(mM + \Delta\sqrt{q}u)]),\end{aligned}\tag{47}$$

where we have used $z \sim \mathcal{N}(mMy, \Delta^2 q)$, $P(y = \pm 1) = \frac{1}{2}$, $\Theta(x) = 1 - \Theta(-x)$, $\mathcal{K}_{-1}(-z) = -\mathcal{K}_1(z)$, and u is symmetric for the standard Gaussian distribution. Note that $\phi(\cdot) = \tanh(\cdot)$ in this paper for the toy model. We remark that the averaged adversarial accuracy in the absence of attack reduces to the standard accuracy derived in Eq. (43), i.e., $\int Du \Theta(mM + \sqrt{q}\Delta u)$.

Zero temperature

The zero temperature limit would make the Boltzmann measure focus on the ground state of lowest energy. In this limit, the order parameters including their conjugate counterparts will diverge or vanish with β as $\beta \rightarrow \infty$. Therefore, the (conjugated) order parameters need to be properly scaled with β . We have then to check the finite temperature free energy first and determine the following scaling behavior of order parameters:

$$\begin{aligned}(q - Q) &\rightarrow \frac{\chi}{\beta}, \\ (2\hat{q} - \hat{Q}) &\rightarrow \beta\hat{\chi}, \\ \hat{M} &\rightarrow \beta\hat{M}, \\ \hat{Q} &\rightarrow \beta^2\hat{Q}.\end{aligned}\tag{48}$$

Therefore, we estimate the limit $\lim_{\beta \rightarrow \infty} \frac{\Psi_0 + \Psi_S + \Psi_E}{\beta}$. The first limit of Ψ_0 is obtained as

$$\Psi_0^\infty = \frac{1}{2} (q\hat{\chi} - \chi\hat{Q}) + M\hat{M},\tag{49}$$

and the second limit of Ψ_S is obtained as

$$\Psi_S^\infty = \frac{\hat{M}^2 - \hat{Q}}{2(\hat{\chi} + \lambda_w)}, \quad (50)$$

and the final limit of Ψ_E reads

$$\begin{aligned} \Psi_E^\infty &= -\alpha \mathbb{E}_{\mathbf{y}} \int D\mathbf{v} F_{\mathbf{y}}(\mathbf{v}), \\ F_{\mathbf{y}}(\mathbf{v}) &= \min_{u_1, u_2} \left[\frac{u_1^2}{2} + \frac{u_2^2}{2} + \mathcal{H}(y_1, y_2, z_1, z_2) \right], \\ z_x &= m y_x M + \sqrt{\Delta^2 q} v_x + \sqrt{\Delta^2 \chi} u_x, \end{aligned} \quad (51)$$

where $x = 1, 2$. Collecting all three contributions in the zero temperature limit, we obtain the ground state free energy as follows:

$$-f(T=0) = \Psi_0^\infty + \Psi_S^\infty + \Psi_E^\infty. \quad (52)$$

Now, the introduced new set of order parameters, $\{q, \hat{Q}, M, \hat{M}, \chi, \hat{\chi}\}$ must obey the saddle-point equations which can be derived in an analogous way to the finite temperature version shown in this section. The SDEs now become

$$\begin{aligned} M &= -\frac{\hat{M}}{\hat{\chi} + \lambda_w}, \\ q &= \frac{\hat{M}^2 - \hat{Q}}{(\hat{\chi} + \lambda_w)^2}, \\ \chi &= -\frac{1}{\hat{\chi} + \lambda_w}, \\ \hat{M} &= \alpha \mathbb{E}_{\mathbf{y}} \int D\mathbf{v} (\mathcal{H}_1^* m y_1 + \mathcal{H}_2^* m y_2), \\ \hat{Q} &= -\frac{\alpha \Delta^2}{\sqrt{\Delta^2 \chi}} \mathbb{E}_{\mathbf{y}} \int D\mathbf{v} (\mathcal{H}_1^* u_1^* + \mathcal{H}_2^* u_2^*), \\ \hat{\chi} &= \alpha \Delta^2 \mathbb{E}_{\mathbf{y}} \int D\mathbf{v} (\mathcal{H}_{11}^* + \mathcal{H}_{22}^*), \end{aligned} \quad (53)$$

where the superscript $*$ means that the optimal values of u_1 and u_2 obtained by solving the min problem in Eq. (51) are used to estimate the effective Hamiltonian, and to arrive at the last equation, we use the Stein Lemma.

Cavity method for FBM

In this section, we derive the message passing equations applied to single instances of learning, using the cavity method developed in spin glass theory [19, 20]. These equations are used to verify our theoretical results of replica calculation ($N \rightarrow \infty$) on finite-sized systems. The two cavity probabilities defined in the main text obey the following self-consistent equations under the seminal cavity approximation in disordered system theory [19, 20]:

$$\begin{aligned} P_{i \rightarrow \nu}(w_i) &\propto e^{-\frac{\beta \lambda_w w_i^2}{2}} \prod_{\mu \in \partial i \setminus \nu} \hat{P}_{\mu \rightarrow i}(w_i), \\ \hat{P}_{\mu \rightarrow i}(w_i) &= \int \prod_{j \in \partial \mu \setminus i} dw_j P_{j \rightarrow \mu}(w_j) P_\mu(\mathbf{y}^\mu, \mathbf{z}^\mu), \end{aligned} \quad (54)$$

where $\mathbf{z}^\mu = (z_1^\mu, z_2^\mu)$ are the pre-activations triggered by $\mathbf{x}_1^\mu, \mathbf{x}_2^\mu$, respectively, and $P_\mu = e^{-\beta \mathcal{L}_\mu}$ (\mathcal{L}_μ is the μ -th pair related loss). The explicit expression of \mathcal{L}^μ is given below.

$$\begin{aligned} \mathcal{L}^\mu &= \frac{\sigma^\mu}{2} (D^\mu)^2 + \frac{1 - \sigma^\mu}{2} \varphi \left[\lambda_F - (D^\mu)^2 \right], \\ D^\mu &= \sqrt{[\phi(z_1^\mu) - \phi(z_2^\mu)]^2}, \\ \sigma^\mu &= \frac{1 + y_1^\mu y_2^\mu}{2}. \end{aligned} \quad (55)$$

The prior distribution of weight $P_0(\mathbf{w}) = \prod_i P_0(w_i)$ is a Gaussian distribution (i.e., the standard ℓ_2 norm weight decay), as shown below.

$$P_0(w_i) \propto \exp\left(-\frac{\beta\lambda_w w_i^2}{2}\right), \quad (56)$$

where λ_w controls the strength of weight decay.

Because of the continuous nature of weights, we first assume a Gaussian distribution $\mathcal{N}(w_i; m_{i \rightarrow \mu}, v_{i \rightarrow \mu})$ for $P_{i \rightarrow \mu}(w_i)$ whose mean and variance are given below.

$$\begin{aligned} m_{j \rightarrow \mu} &= \int dw_j w_j P_{j \rightarrow \mu}(w_j), \\ v_{j \rightarrow \mu} &= \int dw_j w_j^2 P_{j \rightarrow \mu}(w_j) - m_{j \rightarrow \mu}^2. \end{aligned} \quad (57)$$

We use this approximation to simplify the conjugate cavity probability $\hat{P}_{\mu \rightarrow i}(w_i)$ as follows.

$$\begin{aligned} \hat{P}_{\mu \rightarrow i}(w_i) &= \int \prod_{j \in \partial\mu \setminus i} dw_j P_{j \rightarrow \mu}(w_j) P_\mu(\mathbf{y}^\mu, \mathbf{z}^\mu), \\ &= \sqrt{\frac{4\pi^2}{\det(\mathbf{V}_{\mu \rightarrow i})}} \int d\mathbf{z}_1^\mu d\mathbf{z}_2^\mu P_\mu(\mathbf{y}^\mu, z_1^\mu, z_2^\mu) e^{-\frac{1}{2}(\mathbf{z}^\mu - w_i \mathbf{x}_i^\mu - \boldsymbol{\omega}_{\mu \rightarrow i})^T \mathbf{V}_{\mu \rightarrow i}^{-1} (\mathbf{z}^\mu - w_i \mathbf{x}_i^\mu - \boldsymbol{\omega}_{\mu \rightarrow i})}, \end{aligned} \quad (58)$$

where we have defined $\mathbf{z}^\mu = (z_1^\mu, z_2^\mu)$, $\mathbf{x}_i^\mu = [x_{1,i}^\mu, x_{2,i}^\mu]^T$, $\boldsymbol{\omega}_{\mu \rightarrow i} = \sum_{j \in \partial\mu \setminus i} m_{j \rightarrow \mu} \mathbf{x}_j^\mu$, and $\mathbf{V}_{\mu \rightarrow i} = \sum_{j \in \partial\mu \setminus i} v_{j \rightarrow \mu} \mathbf{x}_j^\mu (\mathbf{x}_j^\mu)^T$.

Next, we approximate the integral in Eq. (58) (defined as I_μ) as follows,

$$\begin{aligned} I_\mu &= \int d\mathbf{z}_1^\mu d\mathbf{z}_2^\mu P_\mu(\mathbf{y}^\mu, z_1^\mu, z_2^\mu) \exp\left[-\frac{1}{2}(\mathbf{z}^\mu - w_i \mathbf{x}_i^\mu - \boldsymbol{\omega}_{\mu \rightarrow i})^T \mathbf{V}_{\mu \rightarrow i}^{-1} (\mathbf{z}^\mu - w_i \mathbf{x}_i^\mu - \boldsymbol{\omega}_{\mu \rightarrow i})\right] \\ &\simeq \int d\mathbf{z}_1^\mu d\mathbf{z}_2^\mu P_\mu(\mathbf{y}^\mu, z_1^\mu, z_2^\mu) \exp\left[-\frac{1}{2}(\mathbf{z}^\mu - \boldsymbol{\omega}_{\mu \rightarrow i})^T \mathbf{V}_{\mu \rightarrow i}^{-1} (\mathbf{z}^\mu - \boldsymbol{\omega}_{\mu \rightarrow i})\right] \\ &\times \left[1 + w_i (\mathbf{z}^\mu - \boldsymbol{\omega}_{\mu \rightarrow i})^T \mathbf{V}_{\mu \rightarrow i}^{-1} \mathbf{x}_i^\mu - \frac{1}{2} (w_i)^2 (\mathbf{x}_i^\mu)^T \mathbf{V}_{\mu \rightarrow i}^{-1} \mathbf{x}_i^\mu + \frac{1}{2} w_i^2 (\mathbf{x}_i^\mu)^T \mathbf{V}_{\mu \rightarrow i}^{-1} (\mathbf{z}^\mu - \boldsymbol{\omega}_{\mu \rightarrow i}) (\mathbf{z}^\mu - \boldsymbol{\omega}_{\mu \rightarrow i})^T \mathbf{V}_{\mu \rightarrow i}^{-1} \mathbf{x}_i^\mu\right], \end{aligned} \quad (59)$$

where the last approximation is based on the large N limit. More precisely, in the large N limit, we assume the entries of matrices \mathbf{A}, \mathbf{C} are of the order one, but entries of \mathbf{B} are of the order $\mathcal{O}(1/N)$. Then we have the following mathematical result:

$$\exp\left[-\frac{1}{2}(\mathbf{A} - \mathbf{B})^T \mathbf{C} (\mathbf{A} - \mathbf{B})\right] \simeq \exp\left[-\frac{1}{2}\mathbf{A}^T \mathbf{C} \mathbf{A}\right] \left[1 + \mathbf{B}^T \mathbf{C} \mathbf{A} - \frac{1}{2}\mathbf{B}^T \mathbf{C} \mathbf{B} + \frac{1}{2}\mathbf{B}^T \mathbf{C} \mathbf{A} \mathbf{A}^T \mathbf{C}^T \mathbf{B}\right], \quad (60)$$

where the order higher than $1/N^2$ has been dropped off in the above expansion, and \mathbf{C} is a symmetric matrix. Equation (60) is used to derive the last equality in Eq. (59).

From Eq. (59), we extract a probability measure $\tilde{P}_\mu = \frac{1}{\tilde{Z}_\mu} e^{-\frac{1}{2}(\mathbf{z}^\mu - \boldsymbol{\omega}_{\mu \rightarrow i})^T \mathbf{V}_{\mu \rightarrow i}^{-1} (\mathbf{z}^\mu - \boldsymbol{\omega}_{\mu \rightarrow i})} P_\mu$, from which we derive $\mathbf{f}_\mu \equiv \frac{\partial \ln \tilde{Z}_\mu}{\partial \boldsymbol{\omega}_{\mu \rightarrow i}}$. More precisely, \mathbf{f}_μ and its derivative with respect to $\boldsymbol{\omega}_{\mu \rightarrow i}$ can be derived below.

$$\begin{aligned} \mathbf{f}_\mu &= \frac{\int d\mathbf{z}^\mu \mathbf{V}_{\mu \rightarrow i}^{-1} (\mathbf{z}^\mu - \boldsymbol{\omega}_{\mu \rightarrow i}) \exp\left[-\frac{1}{2}(\mathbf{z}^\mu - \boldsymbol{\omega}_{\mu \rightarrow i})^T \mathbf{V}_{\mu \rightarrow i}^{-1} (\mathbf{z}^\mu - \boldsymbol{\omega}_{\mu \rightarrow i})\right] P_\mu(\mathbf{y}^\mu, \mathbf{z}^\mu)}{\int d\mathbf{z}^\mu \exp\left[-\frac{1}{2}(\mathbf{z}^\mu - \boldsymbol{\omega}_{\mu \rightarrow i})^T \mathbf{V}_{\mu \rightarrow i}^{-1} (\mathbf{z}^\mu - \boldsymbol{\omega}_{\mu \rightarrow i})\right] P_\mu(\mathbf{y}^\mu, \mathbf{z}^\mu)}, \\ \frac{\partial \mathbf{f}_\mu}{\partial \boldsymbol{\omega}_{\mu \rightarrow i}} &= -\mathbf{V}_{\mu \rightarrow i}^{-1} - \mathbf{f}_\mu (\mathbf{f}_\mu)^T \\ &+ \mathbf{V}_{\mu \rightarrow i}^{-1} \frac{\int d\mathbf{z}^\mu (\mathbf{z}^\mu - \boldsymbol{\omega}_{\mu \rightarrow i}) (\mathbf{z}^\mu - \boldsymbol{\omega}_{\mu \rightarrow i})^T \exp\left[-\frac{1}{2}(\mathbf{z}^\mu - \boldsymbol{\omega}_{\mu \rightarrow i})^T \mathbf{V}_{\mu \rightarrow i}^{-1} (\mathbf{z}^\mu - \boldsymbol{\omega}_{\mu \rightarrow i})\right] P_\mu(\mathbf{y}^\mu, \mathbf{z}^\mu)}{\int d\mathbf{z}^\mu \exp\left[-\frac{1}{2}(\mathbf{z}^\mu - \boldsymbol{\omega}_{\mu \rightarrow i})^T \mathbf{V}_{\mu \rightarrow i}^{-1} (\mathbf{z}^\mu - \boldsymbol{\omega}_{\mu \rightarrow i})\right] P_\mu(\mathbf{y}^\mu, \mathbf{z}^\mu)} \mathbf{V}_{\mu \rightarrow i}^{-1}. \end{aligned} \quad (61)$$

Based on the above two auxiliary quantities, the conjugate cavity probability $\hat{P}_{\mu \rightarrow i}(w_i)$ can be recast as

$$\begin{aligned}\hat{P}_{\mu \rightarrow i}(w_i) &\propto 1 + w_i (\mathbf{x}_i^\mu)^\top \mathbf{f}_\mu + \frac{1}{2} w_i^2 (\mathbf{x}_i^\mu)^\top (\partial_{\boldsymbol{\omega}_{\mu \rightarrow i}} \mathbf{f}_\mu + \mathbf{f}_\mu \mathbf{f}_\mu^\top) \mathbf{x}_i^\mu \\ &= 1 + w_i B_{\mu \rightarrow i} - \frac{1}{2} w_i^2 A_{\mu \rightarrow i} + \frac{1}{2} w_i^2 B_{\mu \rightarrow i}^2 \\ &\simeq \exp \left[w_i B_{\mu \rightarrow i} - \frac{1}{2} w_i^2 A_{\mu \rightarrow i} \right],\end{aligned}\tag{62}$$

where $B_{\mu \rightarrow i} = (\mathbf{x}_i^\mu)^\top \mathbf{f}_\mu$, and $A_{\mu \rightarrow i} = -(\mathbf{x}_i^\mu)^\top \frac{\partial \mathbf{f}_\mu}{\partial \boldsymbol{\omega}_{\mu \rightarrow i}} \mathbf{x}_i^\mu$. Inserting Eq. (62) into the first equation of the message passing equations, and noting that $P_{i \rightarrow \mu} \sim \mathcal{N}(m_{i \rightarrow \mu}, v_{i \rightarrow \mu})$, one can derive self-consistently the following closed equations for the cavity mean and variance.

$$\begin{aligned}m_{i \rightarrow \mu} &= \frac{\sum_{\nu \neq \mu} B_{\nu \rightarrow i}}{\beta \lambda_w + \sum_{\nu \neq \mu} A_{\nu \rightarrow i}}, \\ v_{i \rightarrow \mu} &= \frac{1}{\beta \lambda_w + \sum_{\nu \neq \mu} A_{\nu \rightarrow i}}.\end{aligned}\tag{63}$$

A full version of the mean and variance is derived by restoring the μ -th data pair as follows,

$$\begin{aligned}m_i &= \frac{\sum_{\nu \in \partial i} B_{\nu \rightarrow i}}{\beta \lambda_w + \sum_{\nu \in \partial i} A_{\nu \rightarrow i}}, \\ v_i &= \frac{1}{\beta \lambda_w + \sum_{\nu \in \partial i} A_{\nu \rightarrow i}}.\end{aligned}\tag{64}$$

Finally, $P_i(w_i) = \mathcal{N}(w_i; m_i, v_i)$ from which we can estimate the generalization contrastive loss (see Fig. 2 in the main text).