

Gaussian Embeddings: How JEPAs Secretly Learn Your Data Density

Randall Balestriero
Meta-FAIR & Brown University
rbalestr@brown.edu

Nicolas Ballas
Meta-FAIR

Mike Rabbat
Meta-FAIR

Yann LeCun
Meta-FAIR & NYU

Abstract

Joint Embedding Predictive Architectures (JEPAs) learn representations able to solve numerous downstream tasks out-of-the-box. JEPAs combine two objectives: (i) a latent-space prediction term, i.e., the representation of a slightly perturbed sample must be predictable from the original sample’s representation, and (ii) an anti-collapse term, i.e., not all samples should have the same representation. While (ii) is often considered as an obvious remedy to representation collapse, we uncover that JEPAs’ anti-collapse term does much more—it provably estimates the data density. In short, any successfully trained JEPA can be used to get sample probabilities, e.g., for data curation, outlier detection, or simply for density estimation. Our theoretical finding is agnostic of the dataset and architecture used—in any case one can compute the learned probabilities of sample x efficiently and in closed-form using the model’s Jacobian matrix at x . Our findings are empirically validated across datasets (synthetic, controlled, and Imagenet) and across different Self Supervised Learning methods falling under the JEPA family (I-JEPA and DINOv2) and on multimodal models, such as MetaCLIP. We denote the method extracting the JEPA learned density as **JEPA-SCORE**.




Figure 1: Depiction of the 5 least (**left**) and 5 most (**right**) likely samples of class 21 from Imagenet as per **JEPA-SCORE**—JEPAs’ implicit density estimator learned during pretraining. Two striking observations: (i) across all JEPAs (**rows**) the type of samples with low and high probabilities are alike, and (ii) the same samples (amongst 1,000) are found at those extrema. Random samples from that class are provided in fig. 4

1 Introduction

The training procedure of foundation models—Deep Networks (DNs) f_θ able to solve many tasks in zero or few-shot—can take many forms and is at the center of Self Supervised Learning research [2]. Over the years, one principle has emerged as central to all current state-of-the-art methods: encouraging $f_\theta(X)$ to be maximum Entropy given i.i.d pretraining samples X with density p_X [17, 10]. Because the differential Entropy is difficult to estimate in high-dimensional spaces, and $f_\theta(X)$ often contains thousands of dimensions, one solution is to maximize a lower bound by using a decoder and learning to reconstruct X from $f(X)$ [16]. Because this approach comes with known limitations [3], more and more foundation models are trained with Joint-Embedding Predictive Architectures (JEPAs) [12] that directly encourage $f_\theta(X)$ to be Gaussian. In fact, the Gaussian distribution is the one of maximum differential Entropy under covariance constraint, leading to $f_\theta(X)$ producing *Gaussian Embeddings* (GE).

JEPAs can take many forms by employing numerous implicit and explicit regularizers [15, 13]. Today’s JEPAs mostly take three forms: (i) moment-matching objectives (VICReg [4], W-MSE [8]), (ii) non-parametric estimators (SimCLR [6], MoCo [9], CLIP [14]), and (iii) implicit teacher-student methods (DINO [5], I-JEPA [1]). While JEPAs produce state-of-the-art representations, they are currently seen as disconnected from generative models whose goal is to estimate p_X . In fact, the absence of generative modeling is a praised benefit of JEPAs. But one may wonder...

Can the density of $f(X)$ be specified without f learning about p_X ?

That question will be at the core of our study, and the answer is no: producing Gaussian Embeddings can only happen if f_θ learns the underlying data density p_X . But **JEPAs estimate p_X in a highly non standard way, free of input space reconstruction, and free of a parametric model for p_X** . One question remains...

Is there any further benefit of not only specifying a density for $f_\theta(X)$ but using the eponymous Gaussian density?

At it turns out, this choice guarantees that the estimator for p_X implicitly learned during JEPA training can easily be extracted from the final trained model f_θ —an estimator we call the **JEPA-SCORE** (eq. (5)). Our findings not only open new avenues in using **JEPA-SCORE** for outlier detection or data curation, but also shaken the Self Supervised Learning paradigm by showing how **non parametric density estimation in high dimension is now amenable through JEPAs**. Our theory and its corresponding controlled experiments are provided in sections 2.1 and 2.2, and experiments with state-of-the-art models such as DINOv2, MetaCLIP and I-JEPA are provided in section 2.3. **JEPA-SCORE**’s implementation only takes a few lines of code and is provided in listing 1.

2 JEPA-SCORE: the Data Density Implicitly Learned by JEPAs

We now derive our main result stating that in order to minimize the JEPA objective, a DN must learn the data density. We start with some preliminary results in section 2.1, we formalize our general finding in section 2.2, culminating in the JEPA result of section 2.3 and theorem 1. An efficient implementation is also provided in section 2.3.

2.1 Preliminaries: Gaussian Embeddings Are Uniform on the Hypersphere

Our derivations will rely on a simple observations widely known in high-dimensional statistics: K -dimensional standard Gaussians, appropriately normalized, converge to being Uniform on the hypersphere. Let’s denote the K -dimensional standard Gaussian random variable by Z , and the normalized version by $X \triangleq \frac{Z}{\sqrt{K}}$ with density $f_{\mathcal{N}(0, I/K)}$. Let’s also denote the Uniform distribution on the K -dimensional hypersphere surface by $f_{\mathcal{U}(\mathbb{S}(0, R, K))}$ with radius $R > 0$.

Lemma 1. *As K grows, X quickly concentrates around the hypersphere of radius 1, converging to a Uniform density over the hypersphere surface. (Proof in section A.1.)*

Lemma 1 provides an interesting geometric fact which we turn into a practical result for SSL in the following section 2.2, where we demonstrate how learning to produce Gaussian embeddings equates with learning the Energy function of the training data.




Figure 2: **Top left:** Visual illustration of **JEPA-SCORE**—the DN f_θ must learn p_X for its Jacobian matrix to expand or contract the density in order to produce a Uniform density on the hypersphere surface in its embedding space (lemmas 1 and 2). **Top right:** Pearson correlation between **JEPA-SCORE** and the true $\log p(x)$ on a GMM data model for various input dimensions (rows) and number of samples (columns). In all cases, producing Gaussian embeddings make the backbone f_θ internalize the data density which can be easily extracted using our proposed **JEPA-SCORE**. **Bottom:** as **JEPA-SCORE** is an approximation of the true score function, it is possible to perform Langevin sample to recover the true data distribution as shown here in two dimensions.

2.2 Producing Gaussian Embeddings Equates Learning an Energy Function

This section builds upon lemma 1 to demonstrate how learning to produce Gaussian embeddings implies learning about the data density.

Consider two densities, one on the input domain (p_X) and one on the output domain ($p_{f(X)}$). For $p_{f(X)}$ to have a particular form, e.g., $\mathcal{N}(0, \mathbf{I}/K)$, f must learn something about p_X . To see that, we will have leverage the eponymous change of variable formula expressing the embedding density p_{f_θ} as a function of the data density and the DN’s Jacobian matrix:

$$p_{f(X)}(f(\mathbf{x})) = \int_{\{\mathbf{u} \in \mathbb{R}^D | f(\mathbf{u}) = f(\mathbf{x})\}} \frac{p_X(\mathbf{x})}{\prod_{k=1}^{\text{rank}(J_f(\mathbf{x}))} \sigma_k(J_f(\mathbf{x}))} d\mathcal{H}^r(\mathbf{x}), \quad (1)$$

where $\mathcal{H}^r$ denotes r -dimensional Hausdorff measure, with $r \triangleq \dim(\{\mathbf{u} \in \mathbb{R}^D | f(\mathbf{u}) = f(\mathbf{x})\})$ being the dimension of the level set of f at $\mathbf{x}$. We note that eq. (1) does not require f to be bijective, which will be crucial for our JEPA result in section 2.3; for details see [11, 7]. Combining eq. (1) and lemma 1 leads us to the following result.

Lemma 2. *In order for $f(X)$ to be distributed as $\mathcal{N}(0, \mathbf{I}/K)$ for large K , f must learn the data density p_X up to mean-preserving rescaling within each level-set $\{\mathbf{u} \in \mathbb{R}^D | f(\mathbf{u}) = f(\mathbf{x})\}$. (Proof in section A.2.)*

Empirical validation. Before broadening lemma 2 to JEPAs, we first provide empirical validations that learning to produce Gaussian embeddings implies learning the data density in fig. 2. We show that, in fact, the data density can be recovered with high accuracy, and it is even possible to draw samples from the estimated density through Langevin dynamics.

2.3 JEPA-SCORE: The Data Density Learned by JEPAs

Most Joint-Embedding Predictive Architecture (JEPA) methods aim to achieve two goals: (i) predictive invariance, and (ii) representation diversity, as seen in the following loss function

$$\begin{aligned} \mathcal{L} \triangleq & \sum_{n=1}^N \mathbb{E}_{(\mathbf{x}_n^{(1)}, \mathbf{x}_n^{(2)}) \sim \mathcal{G}(\mathbf{x}_n)} \left[\text{dist} \left(\text{Pred} \left(\text{Enc} \left(\mathbf{x}_n^{(1)} \right) \right), \text{Enc} \left(\mathbf{x}_n^{(2)} \right) \right) \right] \quad (\text{predictive invariance}) \\ & + \text{diversity} \left((\text{Enc}(\mathbf{x}_n))_{n \in [N]} \right), \quad (\text{anti-collapse}), \end{aligned} \quad (2)$$




Figure 3: Depiction of **JEPA-SCORE** for 5,000 samples from different datasets (imagenet-1k/a/r, MNIST and Galaxy). We clearly observe that as the pretraining dataset size increases (all models against IJEPA-1k) as MNIST and Galaxy images are seen as lower probability samples, i.e., *those images are less and less represented within the overall pretraining dataset*. While our score does not rely on singular vectors, we provide some examples in fig. 7. This can be used to assess if a model is ready or not to handle particular data domains at test time for zero-shot tasks.

where $\mathbf{x}_n^{(1)}, \mathbf{x}_n^{(2)}$ are two generated “views” from the original sample through the stochastic operator $\mathcal{G}$, and dist is a distance function (e.g., L2). For images, $\mathcal{G}$ typically involves two different data-augmentations. At this point, lemma 2 only takes into account the non-collapse term of JEPAs. But an interesting observation from lemma 2 is that the integration occurs over the level set of the function f_θ which coincides with the JEPA’s invariance term when Pred is near identity.

Data assumption. To make our result as precise as possible, we focus on the case where the views are generated from stochastic transformations, with density p_T . We also denote the density of *generators* as p_μ , from which the data density p_X is defined as

$$p_X \triangleq p_\mu \otimes p_T. \quad (3)$$

In other words, each training sample is seen as some transformation of some generator. We do not impose further restrictions, e.g., that there is only one generator per class. *In practical setting, the generators (p_μ) are the original training samples prior to applying any augmentation—hence estimating p_μ will amount to estimating the data density.*

JEPA-SCORE. Combining eqs. (2) and (3) and lemma 2 leads to the following result proved in section A.3.

Theorem 1. *At optimality, JEPA embeddings estimate the data density as per*

$$p_\mu(\mu) \propto \mathbb{E}_{p_T} \left[\frac{1}{\prod_{k=1}^{\text{rank}(J_f(\mathbf{x}))} \sigma_k(J_f(\mu, T))} \right]^{-1}. \quad (4)$$

We define our **JEPA-SCORE** for input $\mathbf{x}$ as the Monte Carlo estimator of eq. (4), for a single-sample estimate we have (in log-scale)

$$\text{JEPA-SCORE}(\mathbf{x}) \triangleq \sum_{k=1}^{\text{rank}(J_f(\mathbf{x}))} \log(\sigma_k(J_f(\mathbf{x}))), \quad (5)$$

which exactly recovers p_μ as long as the JEPA loss is minimized. We use a logarithmic scale to align with the definition of a score function. A visual illustration is provided at the top of fig. 2. We empirically validate eq. (5) by using pretrained JEPA models and visualizing a few Imagenet-1k classes after ordering the images based on their JEPA-SCORE in figs. 1, 5 and 6. We obtain that for bird classes, high probability samples depict flying birds while low probability ones are seated. We also conduct an additional experiment where we compute JEPA-SCORE for 5,000 samples of different dataset (Imagenet, MNIST and Galaxy) and depict the distribution of their JEPA-SCORE in fig. 3. We clearly see that datasets that weren’t seen during pretraining, e.g., Galaxy, have much lower JEPA-SCORE than Imagenet samples.

Conclusion. We provided a novel connection between JEPAs and score-based methods, two families of methods thought to be unrelated until now. Controlled experiments on synthetic data confirm the validity of our theory and qualitative experiments with state-of-the-art large scale JEPAs also qualitatively validate our findings. Although this is only a first step, we hope that JEPA-SCORE will open new avenues for outlier detection and model assessment for downstream tasks.

References

- [1] Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15619–15629, 2023.
- [2] Randall Balestriero, Mark Ibrahim, Vlad Sobal, Ari Morcos, Shashank Shekhar, Tom Goldstein, Florian Bordes, Adrien Bardes, Gregoire Mialon, Yuandong Tian, et al. A cookbook of self-supervised learning. *arXiv preprint arXiv:2304.12210*, 2023.
- [3] Randall Balestriero and Yann LeCun. Learning by reconstruction produces uninformative features for perception. *arXiv preprint arXiv:2402.11337*, 2024.
- [4] Adrien Bardes, Jean Ponce, and Yann LeCun. Vicreg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021.
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021.
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PmlR, 2020.
- [7] Milan Cvitkovic and Günther Koliander. Minimal achievable sufficient statistic learning. In *International Conference on Machine Learning*, pages 1465–1474. PMLR, 2019.
- [8] Aleksandr Ermolov, Aliaksandr Siarohin, Enver Sangineto, and Nicu Sebe. Whitening for self-supervised representation learning. In *International conference on machine learning*, pages 3015–3024. PMLR, 2021.
- [9] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020.
- [10] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. *arXiv preprint arXiv:1808.06670*, 2018.
- [11] Steven Krantz and Harold Parks. *Geometric integration theory*. Springer, 2008.
- [12] Yann LeCun. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. *Open Review*, 62(1):1–62, 2022.
- [13] Etai Littwin, Omid Saremi, Madhu Advani, Vimal Thilak, Preetum Nakkiran, Chen Huang, and Joshua Susskind. How japa avoids noisy features: The implicit bias of deep linear self distillation networks. *Advances in Neural Information Processing Systems*, 37:91300–91336, 2024.
- [14] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmlR, 2021.
- [15] Manu Srinath Halvagal, Axel Laborieux, and Friedemann Zenke. Implicit variance regularization in non-contrastive ssl. *Advances in Neural Information Processing Systems*, 36:63409–63436, 2023.
- [16] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pages 1096–1103, 2008.
- [17] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International conference on machine learning*, pages 9929–9939. PMLR, 2020.

A Proofs

A.1 Proof of lemma 1

Proof. **Proof 1:** The proof consists in expressing both densities in spherical coordinates, and studying their convergence as K increases. Let's first express the Uniform distribution in spherical coordinates:

$$f_{\mathcal{U}(\mathbb{S}(0,R,K))}(\mathbf{x}) = \delta(\|\mathbf{x}\|_2 - R) \frac{\Gamma(K/2)}{2\pi^{K/2} R^{K-1}}, \quad (\text{Cartesian coordinates})$$

$$f_{\mathcal{U}(\mathbb{S}(0,R,K))}(r, \boldsymbol{\theta}) = \delta(r - R) \frac{\Gamma(K/2)}{2\pi^{K/2}} \prod_{i=1}^{K-1} \sin(\theta_i)^{K-i-1}, \quad (\text{spherical coordinates}),$$

and let's now express the rescaled standard Gaussian density $\frac{Z}{\sqrt{K}}$ in spherical coordinates:

$$f_{\mathcal{N}(0, \mathbf{I}/K)}(\mathbf{x}) = \left(\frac{K}{2\pi}\right)^{K/2} e^{-\frac{K}{2}\|\mathbf{x}\|_2^2}, \quad (\text{Cartesian coordinates})$$

$$\begin{aligned} f_{\mathcal{N}(0, \mathbf{I}/K)}(r, \boldsymbol{\theta}) &= \left(\frac{K}{2\pi}\right)^{K/2} e^{-\frac{K}{2}r^2} r^{K-1} \prod_{i=1}^{K-1} \sin(\theta_i)^{K-i-1} \\ &= \frac{K^{\frac{K}{2}}}{2^{K/2-1}\Gamma(K/2)} e^{-\frac{Kr^2}{2}} \frac{\Gamma(K/2)}{2\pi^{K/2}} r^{K-1} \prod_{i=1}^{K-1} \sin(\theta_i)^{K-i-1} \\ &= \underbrace{\frac{K^{\frac{K}{2}}}{2^{K/2-1}\Gamma(K/2)} r^{K-1} e^{-\frac{Kr^2}{2}}}_{\text{scaled Chi-distribution} \xrightarrow{K \rightarrow \infty} \delta(r-1)} f_{\mathcal{U}(\mathbb{S}(0,R,K))}(r, \boldsymbol{\theta}), \quad (\text{spherical coordinates}). \end{aligned}$$

As K increases, as the scaled Chi-distribution converges to a Dirac function at 1, leading to our desired result.

Proof 2: The above proof provides granular details into the convergence to the Uniform distribution on the hypersphere by studying the scaled Chi-distribution. For completeness, we also provide a more straightforward argument, sufficient to study the limiting case. First, it is known that $\frac{Z}{\sqrt{K}}$ being isotropic Gaussian, the distribution of norms, $\|Z\|_2^2/K$, is a Chi-squared distribution with mean 1 and variance $2/K$. That is, as K increases as the norms distribution converges to a Dirac at 1. Lastly, because $\frac{Z}{\sqrt{K}}$ is isotropic, it will be uniformly distribution on the hypersphere after normalization. But as K increases, as the samples are already normalized, hence leading to our result. $\square$

A.2 Proof of lemma 2

Proof. First and foremost, recall that the density of the random variable $f(X)$ is given by eq. (1). Relying on lemma 1 which stated that for large K , our assumption on the output density reads $f(\mathbf{x}) \sim \mathcal{U}(0, 1)$, we obtain that $\int_{\{\mathbf{u} \in \mathbb{R}^D \mid f(\mathbf{u}) = f(\mathbf{x})\}} \frac{p_X(\mathbf{u})}{\prod_{k=1}^{\text{rank}(J_f(\mathbf{u}))} \sigma_k(J_f(\mathbf{u}))} d\mathcal{H}^r(\mathbf{u}) = \text{cst}$. Now if f

is bijective between $\text{supp}(p_X)$ and $\mathbb{R}^K$, then it is direct to see that $p_X(x) \propto \prod_{k=1}^{\text{rank}(J_f(x))} \sigma_k(J_f(x))$. Now if f is surjective there is no longer a one-to-one mapping between f and p_X . Instead, there is ambiguity over each level set of f . To see that, recall that we only need to maintain a constant value over the integration on the level set. Hence, f is free to scale up one subset of that level set, and scale down another subset, proportionally to p_X to preserve the integration to a constant. $\square$

A.3 Proof of theorem 1

Proof. The role of the predictor in JEPA training is to allow for additional computation to predict one view's embedding from the other view's embedding. While this provides numerous empirical benefits, e.g., in terms of optimization landscape, it actually does not impact the level-set of the encoder—which is what is needed in eq. (1).

To understand the above argument, consider the case where the views are obtained from applying a transformation such as masking. We denote by $\mathcal{M}$ the masking random and by $\text{mask}(\mathbf{x})$ the



Figure 4: Random samples from Imagenet-1k training dataset for class 21.

application of one realization of $\mathcal{M}$ onto the input $\mathbf{x}$. We thus have for the invariance term of sample $\mathbf{x}_n$

$$\mathbb{E}_{(\text{mask}^{(1)}, \text{mask}^{(2)}) \sim (\mathcal{M}, \mathcal{M})} \left[\text{dist} \left(\text{Pred} \left(\text{Enc} \left(\text{mask}^{(1)}(\mathbf{x}_n) \right) \right), \text{Enc} \left(\text{mask}^{(2)}(\mathbf{x}) \right) \right) \right].$$

Because the predictor is only applied on one of the two embeddings, it is clear that for the JEPA loss to be minimized, it must also be true that

$$\mathbb{E}_{\text{mask}^{(2)} \sim \mathcal{M}} \left[\text{dist} \left(\text{Pred} \left(\text{Enc} \left(\text{mask}^{(1)}(\mathbf{x}_n) \right) \right), \text{Enc} \left(\text{mask}^{(2)}(\mathbf{x}) \right) \right) \right] = 0,$$

for any realization of $\text{mask}^{(1)}$. In other word, the encoder’s invariance is over the support of $\mathcal{M}$ no matter if the predictor is identity or nonlinear. Therefore our result directly follows from the above combined with eqs. (1) and (3). $\square$

B Implementation Details

```

1  from torch.autograd.functional import jacobian
2
3  # model returns a tensor of shape (num_samples, features_dim)
4  J = jacobian(lambda x: model(x).sum(0), inputs=images)
5  with torch.inference_mode():
6      J = J.flatten(2).permute(1, 0, 2)
7      svdvals = torch.linalg.svdvals(J)
8      jepa_score = svdvals.clip_(eps).log_().sum(1)
9

```

Listing 1: JEPA-SCORE implementation in PyTorch. Our empirical ablations demonstrate that JEPA-SCORE is not sensitive to the choice of eps (we pick $1e-6$)

C Additional Figures



Figure 5: **Top:** least and most likely samples based on the proposed JEPA-SCORE for a fe different pretrained backbones on class 141. **Bottom:**Random samples from Imagenet-1k training dataset for class 141.




Figure 6: **Top:** least and most likely samples based on the proposed JEPA-SCORE for a fe different pretrained backbones on class 141. **Bottom:**Random samples from Imagenet-1k training dataset for class 147.



Figure 7: Depiction of the top singular vectors of the Jacobian matrix (columns) for a given input image (left column). The corresponding singular value is provided in the title of each subplot.