

Power Normalizations in Fine-grained Image, Few-shot Image and Graph Classification

Piotr Koniusz and Hongguang Zhang

Abstract—Power Normalizations (*PN*) are useful non-linear operators which tackle feature imbalances in classification problems. We study PNs in the deep learning setup via a novel PN layer pooling feature maps. Our layer combines the feature vectors and their respective spatial locations in the feature maps produced by the last convolutional layer of CNN into a positive definite matrix with second-order statistics to which PN operators are applied, forming so-called Second-order Pooling (*SOP*). As the main goal of this paper is to study Power Normalizations, we investigate the role and meaning of MaxExp and Gamma, two popular PN functions. To this end, we provide probabilistic interpretations of such element-wise operators and discover surrogates with well-behaved derivatives for end-to-end training. Furthermore, we look at the spectral applicability of MaxExp and Gamma by studying Spectral Power Normalizations (*SPN*). We show that SPN on the autocorrelation/covariance matrix and the Heat Diffusion Process (HDP) on a graph Laplacian matrix are closely related, thus sharing their properties. Such a finding leads us to the culmination of our work, a fast spectral MaxExp which is a variant of HDP for covariances/autocorrelation matrices. We evaluate our ideas on fine-grained recognition, scene recognition, and material classification, as well as in few-shot learning and graph classification.

Index Terms—CNN, Second-order Aggregation, Eigenvalue Power Normalization, Bilinear Pooling, Tensor Pooling, Heat Diffusion



1 INTRODUCTION

Second-order statistics of data features are used in object recognition, texture categorization, action representation, and human tracking [1, 2, 3, 4, 5]. For example, the popular region covariance descriptors [1] compute a covariance matrix over multiple features extracted from image regions. Given Bag-of-Words histograms or local descriptors of an image, second-order co-occurrence pooling of such vectors captures correlations between pairs of features, and improves performance of semantic segmentation and visual recognition compared to first-order methods [4, 5, 6]. Extensions to higher-order descriptors [5, 6, 7] improve results further.

However, second- and higher-order statistics require robust aggregation/pooling mechanisms to obtain the best classification results [4, 5, 6]. Once the statistics are captured in the matrix form, they undergo a non-linearity such as Power Normalization [8] whose role is to reduce/boost contributions from frequent/infrequent visual stimuli in an image, respectively. The popular Bag-of-Words provide numerous insights into the role played by pooling during the aggregation step. The theoretical relation between Average and Max-pooling was studied in [9]. A likelihood-inspired analysis of pooling [10] led to a *theoretical expectation of Max-pooling*. Max-pooling was recognized as a lower bound of the likelihood of ‘at least one particular visual word being present in an image’ [11] while Power Normalization was also applied to Fisher Kernels [12]. According to [8], pooling methods are closely related but [8] does not study second-order pooling or end-to-end training. Element-wise Power Normaliza-

tion (PN) and Eigenvalue Power Normalization (EPN) were first applied to autocorrelation/covariance matrices and tensors in [6].

In this paper, we revisit the above pooling methods in end-to-end setting and interpret them in the context of second-order matrices. Firstly, we formulate a kernel which combines feature vectors collected from the last convolutional layer of ResNet-50 together with so-called spatial location vectors [6, 8, 13] which contain Cartesian coordinates (spatial locations) of feature vectors in feature maps. We linearize such a kernel into a second-order matrix to capture correlations of combined feature vectors. Next, we study the role of Power Normalizations in end-to-end setting. We show that PNs have well-founded probabilistic interpretation in the context of second-order statistics. We propose PN surrogates with well-behaved derivatives for end-to-end training. Finally, we study PNs in the spectral domain, so-called Spectral Power Normalizations (SPN). We show that the Heat Diffusion Process (HDP) [14] on a graph Laplacian is closely related to SPNs: HDP and SPN play the same role for graph Laplacian and autocorrelation/covariance matrices, resp. As SPN and the HDP share properties, we propose a fast spectral MaxExp. To summarize:

- i. We aggregate feature vectors extracted from CNNs and their spatial coordinates into a second-order matrix.
- ii. We revisit Power Normalization functions, derive them for second-order representations and show how PNs emerge if we assume that features are drawn from the Bernoulli or Normal distributions. We also suggest PN surrogates with well-behaved derivatives for end-to-end training.
- iii. We show that Spectral Power Normalizations are in fact a time-reversed ($t < 1$) Heat Diffusion Process, an important connection that explains the role of SPN. Thus, we propose a fast spectral MaxExp whose profile closely resembles HDP.
- iv. In addition to our standard fine-grained pipeline, we develop second-order relational representations for few-shot learning and we even consider graph classification.

Figures 1a and 5a show our classification pipeline (we use

- P. Koniusz and H. Zhang are with Data61/CSIRO (formerly known as NICTA) and the Australian National University, Canberra, Australia, ACT2601. E-mail: see <http://claret.wikidot.com>
- Equal contribution: P. Koniusz was mainly concerned with the mathematical analysis/modeling while H. Zhang with the deep learning modeling.

Manuscript submitted Dec-2018. Manuscript accepted by TPAMI on 02-Jul-2020.

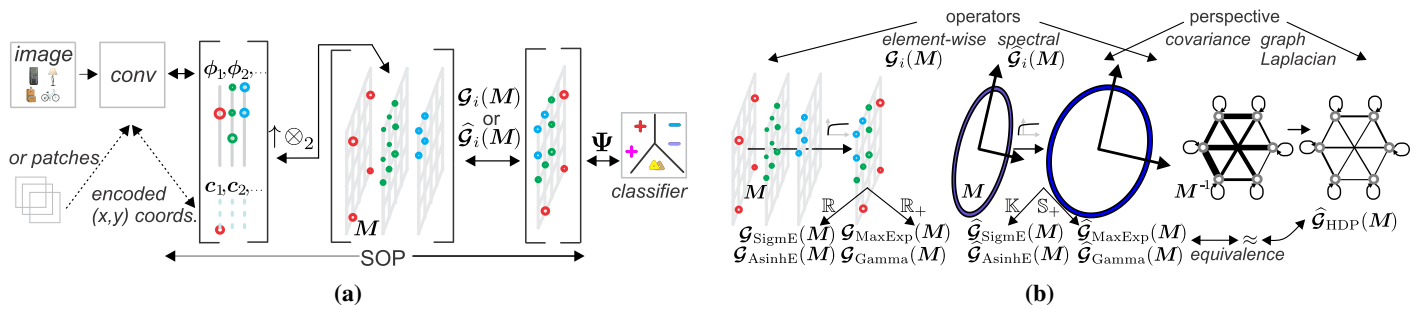


Fig. 1: Our end-to-end classification setting (Fig. 1a). We pass an image (or patches) via CNN and extract feature vectors ϕ from its last conv. layer and augment them by encoded spatial coordinates c . We perform pooling on a second-order matrix M by the Element-wise or Spectral Power Normalization \mathcal{G}_i or $\hat{\mathcal{G}}_i$, resp. Index i indicates the operator *i.e.*, MaxExp or Gamma. Figure 1b shows the taxonomy of pooling. We distinguish element-wise and spectral operators (which operate on autocorrelation/covariance or the graph Laplacian matrices). Element-wise and spectral operators can be further divided into working with non-negative \mathbb{R}_+ and real \mathbb{R} values, and $\mathbb{S}_+/\mathbb{S}_{++}$ and Krein (\mathbb{K}) spaces, resp. We derive so-called element-wise and spectral MaxExp operators, and we show that spectral MaxExp/Gamma (on autocorrelation/covariance) are approx. equivalent to the time-reversed Heat Diffusion Process (on the loopy graph Laplacian). MaxExp makes the underlying multivariate Gaussian closer to isotropic. Equivalently, the strong connections in the graph (thick edges) become weaker (thin lines) and more equalized.

ResNet-50 pre-trained on ImageNet) and our few-shot learning Second-order Similarity Learning Network (*SoSM*). We experiment on ImageNet, Flower102, MIT67, FMD and Food-101 (classification setting), *miniImageNet*, Flower102, Food-101 and Open MIC (few-shot setting), and MUTAG, PTC, PROTEINS, NCI1, COLLAB, REDDIT-BINARY/MULTI-5K (graph classification).

We explore Power Normalizing functions for second-order image [15] and graph classification, and few-shot learning [16].

Motivated by the need to discuss PNs (*i.e.*, see tutorial [17]), we propose a taxonomy in Fig. 1b and analyze (i) element-wise operators (fast but less robust) and (ii) spectral operators (exploit correlations between features but are slower). Second level branches of Fig. 1b consider non-negative \mathbb{R}_+ and real \mathbb{R} inputs to PN, and $\mathbb{S}_+/\mathbb{S}_{++}$ and Krein (\mathbb{K}) input spaces of SPN.

To interpret the statistical meaning of PNs w.r.t. inputs, we show that element-wise operators similar to Gamma [6, 8] emerge from statistical models assuming features being drawn from an i.i.d. Bernoulli or Normal distribution. This encourages us to look at PN as a wider family of functions (*c.f.* Gamma/square root). As the i.i.d. assumption in element-wise operators is limiting, we consider SPNs [6] due to their feature decorrelating properties. We show that SPNs are in fact an equivalent of the Heat Diffusion Process [14], specifically a time-reversed ($t < 1$) HDP. As Fig. 1b shows, SPN and HDP operate on autocorrelation/covariance and the graph Laplacian matrix respectively (different theoretical perspectives). Finally, to tackle the speed bottleneck, we propose a fast SVD-free back-propagation through spectral MaxExp.

Importantly, on few-shot learning, we find that SPNs benefit fine-tuning but not training from scratch. Second-order matrices alone cannot reduce the diffusion of signal between correlated features during fine-tuning, a cause of so-called catastrophic forgetting. However, SPNs limit the diffusion between features of second-order matrices and thus reduce catastrophic forgetting.

Sections 2 and 3 provide the related work, background and notations. Section 4 and 5 present our mathematical analysis. Sections 6, 7 and 8 present the pipeline, experiments and conclusions.

2 RELATED WORK

Below we introduce Region Covariance Descriptors (RCD) which are perhaps the oldest second-order descriptors [1, 18].

Region Covariance Descriptors (RCD). RCDs typically capture co-occurrences of luminance, first- and/or second-order derivatives

of texture patterns, and have been applied to tracking [2], semantic segmentation [4] and object category recognition [5, 6]. As RCDs typically require a Non-Euclidean distance to compare positive (semi-)definite RCD datapoints, we list below popular choices.

Non-Euclidean distances. The distance between two positive definite datapoints is typically measured according to the Riemannian geometry while Power-Euclidean distances [19] extend to positive semi-definite datapoints. In particular, Affine-Invariant Riemannian Metric [20, 21], KL-Divergence Metric (*KLDM*) [22], Jensen-Bregman LogDet Divergence (*JBLD*) [23] and Log-Euclidean (*LogE*) [24] have been used in diffusion imaging, RCD-based methods, dictionary and metric learning [25, 26, 27].

Second-order Pooling in CNNs. There has been a revived interest in co-occurrence patterns in CNN setting. Methods [28, 29, 30] fuse CNN streams via the outer product for the fine-grained image recognition. Approach [31] uses co-occurrences of CNN feature vectors and facial attribute vectors for face recognition.

We note that the Log-Euclidean distance and Power Normalization have recently been implemented in the CNN setting [32, 33, 34, 35] for the purpose of image classification. For instance, approaches [34, 35] build on so-called Eigenvalue Power Normalization as introduced by [5, 6], however, they extend it to end-to-end CNN setting with back-propagation via SVD explored in [32] which is both slow and unstable due to so-called non-simple eigenvalues. To this end, recent advances in [34, 35] propose so-called Newton-Schulz iterations to compute the square root of matrix (a special case of Gamma [6, 8]). Authors of [34, 35] motivate the use of spectral Gamma with the notions of burstiness and whitening, something considered first in the context of Eigenvalue Power Normalization in early works [5, 6]. A visualization approach into bilinear models including α -pooling, a form of Power Normalization, is proposed in [36].

Power Normalizations. Image representations suffer from the so-called burstiness which is ‘the property that a given visual element appears more times in an image than a statistically independent model would predict’ [37]. Power Normalization [12, 37, 38] suppresses the burstiness which has been studied/evaluated in the context of Bag-of-Words [5, 8]. The theoretical study of Average and Max-pooling [9, 10] highlighted their statistical models and a connection to Max-pooling. A relationship between the likelihood of ‘at least one particular visual word being present in an image’ and Max-pooling was studied in [11]. Survey [8] shows that all Power Normalization functions are closely related.

We show that MaxExp for element-wise co-occurrence pooling emerges from the Multinomial modeling while authors of [9, 10] use a Binomial setting for first-order signals. To paraphrase, we show why it is theoretically meaningful to use MaxExp with co-occurrences, and by connecting MaxExp to HDP in the spectral setting, we show why MaxExp can work with the spectrum.

One- and Few-shot Learning, motivated by the human ability to learn from few samples, has been widely studied in both shallow [39, 40, 41] and deep learning scenarios [42, 43, 44, 45].

Matching Network [42], Prototypical Networks [43], Model-Agnostic Meta-Learning (MAML) [44] and Relation Net [45] learn similarity between pairs of images rather than class concepts. Our few-shot learning pipeline is similar to Relation Net [45] which uses first-order representations. However, we build on few-shot learning with second-order representations and Power Normalization as proposed in [16]. Moreover, we consider element-wise and spectral operators, and we show a theoretical analysis that PN is especially beneficial in relation learning. Finally, we note that second-order representations are gaining momentum in few-shot learning [46, 47, 48] which use localization mechanisms while approach [49] uses subspace-based class-wise prototypes.

3 BACKGROUND

Below we review our notations, the background on kernel linearizations and the Power Normalization family.

3.1 Notations

Let $\mathbf{x} \in \mathbb{R}^d$ be a d -dimensional feature vector. \mathcal{I}_N stands for the index set $\{1, 2, \dots, N\}$. The spaces of symmetric positive semi-definite and definite matrices are \mathcal{S}_+^d and \mathcal{S}_{++}^d . $\text{Sym}(\mathbf{X}) = \frac{1}{2}(\mathbf{X} + \mathbf{X}^T)$. A vector with all coefficients equal one is denoted by $\mathbb{1}$, \mathbf{j}_k is a vector of all zeros except for the k -th coefficient which is equal one, and \mathbf{J}_{kl} is a matrix of all zeros with a value of one at the position (k, l) . Moreover, \odot is the Hadamard product, $\text{Vec}(\mathbf{X})$ vectorizes matrix \mathbf{X} in analogy to $\mathbf{X}(:)$ in Matlab and ‘ \dagger ’ is the Moore-Penrose pseudoinverse. We use the Matlab notation $\mathbf{v} = [\text{begin} : \text{step} : \text{end}]$ to generate a vector \mathbf{v} with elements starting as *begin*, ending as *end*, with stepping *step*. Operator ‘ \cdot ’ in $[\mathbf{x}; \mathbf{y}]$ is the concat. of vectors \mathbf{x} and \mathbf{y} (or scalars).

3.2 Autocorrelation matrix

Below we show that autocorrelation (second-order) matrices emerge from a linearization of sum of Polynomial kernels.

Proposition 1. Let $\Phi_A \equiv \{\phi_n\}_{n \in \mathcal{N}_A}$, $\Phi_B \equiv \{\phi_n^*\}_{n \in \mathcal{N}_B}$ be datapoints from two images Π_A and Π_B , and $N = |\mathcal{N}_A|$ and $N^* = |\mathcal{N}_B|$ be the numbers of data vectors e.g., obtained from the last convolutional feature map of CNN for images Π_A and Π_B . Autocorrelation feature maps result from a linearization of the sum of Polynomial kernels of degree 2:

$$\begin{aligned} K(\Phi_A, \Phi_B) &= \frac{1}{NN^*} \sum_{n \in \mathcal{N}_A} \sum_{n' \in \mathcal{N}_B} \langle \phi_n, \phi_{n'}^* \rangle^2 = \langle \text{Vec}(\mathbf{M}(\Phi_A)), \text{Vec}(\mathbf{M}(\Phi_B)) \rangle \\ &= \langle \text{Vec}\left(\frac{1}{N} \sum_{n \in \mathcal{N}_A} \phi_n \phi_n^T\right), \text{Vec}\left(\frac{1}{N^*} \sum_{n' \in \mathcal{N}_B} \phi_{n'}^* \phi_{n'}^{*T}\right) \rangle. \end{aligned} \quad (1)$$

Proof. See [50] for the details of such an expansion. \square

Thus, we obtain the following (kernel) feature map $\mathbf{M}(\Phi)$ on features Φ which coincides with the autocorrelation matrix:

$$\mathbf{M}(\Phi) = \frac{1}{N} \sum_{n \in \mathcal{N}} \phi_n \phi_n^T. \quad (2)$$

For simplicity of notation, we drop $\Phi \equiv \{\phi_n\}_{n \in \mathcal{N}}$ from $\mathbf{M}(\Phi)$ where possible, that is, we often write \mathbf{M} rather than $\mathbf{M}(\Phi)$.

3.3 Power Normalization Family (first-order variants)

Pooling is an aggregation step of feature vectors Φ that produces their signature used in for instance training an SVM. We study pooling of second-order representations to preserve second-order statistics of input vectors. For clarity, we firstly introduce highly-related first-order PNs, that is MaxExp and Gamma operators. Traditionally, PN is a function $g([0, \iota]) \rightarrow [0, \nu]$ such that $g(0) = 0$, $g(\iota) = \nu$ (i.e., $\iota = \nu = 1$), $g(p)$ is monotonically non-decreasing on $0 \leq p \leq \iota$ and the slope of $g(p)$ rises fast/slow for $p \approx 0$ and slow/fast for $p \approx \iota$, resp. Fig. 2d illustrates such properties.

MaxExp (first-order). Drawing features from the Bernoulli distribution under the i.i.d. assumption [10] leads to so-called *Theoretical Expectation of Max-pooling (MaxExp)* operator [8] related to Max-pooling [9] and *Gamma* [8]. The following proposition formalizes this process.

Proposition 2. Assume a vector $\phi \in \{0, 1\}^N$ which stores N outcomes of drawing from Bernoulli distribution under the i.i.d. assumption for which the probability p of an event ($\phi_n = 1$) and $1 - p$ for ($\phi_n = 0$) can be estimated as an expected value e.g., $p = \text{avg}_n \phi_n$. Then the probability of at least one positive event in ϕ from N trials becomes $\psi = 1 - (1 - p)^N$.

Proof. The proof can be found in [15]. \square

A practical implementation of this pooling strategy [8] is given by $\psi_k = 1 - (1 - \text{avg}_n \phi_{kn})^\eta$, where $0 < \eta \approx N$ is an adjustable parameter and ϕ_{kn} is a k -th feature of an n -th feature vector e.g., as defined in Prop. 1, normalized to range 0–1.

Gamma (first-order). It was shown in [8] that Power Normalization (Gamma) given by $\psi_k = (\text{avg}_n \phi_{kn})^\gamma$, where $0 < \gamma \leq 1$ is an adjustable parameter, is in fact an approximation of MaxExp.

4 PROBLEM FORMULATION

The main goal of this paper is a theoretical study of Power Normalizations for second-order representations, their interpretation and theoretical connections. Fig. 1b introduces the taxonomy of operators we follow. We choose existing PN operators Gamma and MaxExp (other operators are typically their modifications). We look at (i) element-wise PN (fast but suboptimal) and (ii) SPN (slow but exploiting feat. correlations). We then consider SPNs on autocorr./covariances and their connection to graph Laplacians.

We show that the Heat Diffusion Process on graphs and SPN are two sides of the same mathematical process. Thus, SPNs reverse the diffusion of signal in autocorrelation/covariance matrices rather than just reduce the burstiness of features. The culmination of our work is MaxExp(F), the fast spectral MaxExp, (and a fast approx. HDP for completeness) whose runtime scales sublinearly w.r.t. its parameter. MaxExp(F) rivals the approx. matrix square root via Netwon-Schulz iter. [34, 35].

Sec. 4.1 proposes a spatially augmented autocorrelation matrix (Fig. 1a) that can be seen as introducing spatially localized nodes into a graph. Sec. 4.2 explains why MaxExp is applicable to co-occurrences. While MaxExp for first-order signals emerges from Binomial modeling of features, for co-occurrences the same operator emerges from Multinomial modeling. Following the taxonomy (Fig. 1b), in Sec. 4.3 we generalize MaxExp/Gamma [5, 8] (work on \mathbb{R}_+) to Logistic a.k.a. Sigmoid (*SigmE*) and the Arcsin hyperbolic (*AsinhE*) functions (work on \mathbb{R}). SigmE emerges from modeling Normal distr. Both functions extend readily to Krein spaces (Table 2). Sec. 5.1 introduces spectral pooling.




Fig. 2: Gamma, AsinhE, MaxExp and SigE are illustrated in Figures 2a and 2b while derivatives of Gamma and AsinhE are shown in Figure 2c. Lastly, Gamma for several values of γ is shown in Figure 2d from which its similarity to MaxExp in range 0–1 is clear.

4.1 Augmented autocorrelation matrix

The autocorrelation matrix defined in Section 3.2 is perfectly applicable to considerations in our paper. We enhance autocorrelation matrices by two steps detailed below for the best performance but are secondary to the main analysis of Power Normalizations.

β -centering. As in Eq. (2), let $\Phi \equiv \{\phi_n\}_{n \in \mathcal{N}}$ but ϕ_n are rectified so that $\phi_n \geq 0$. Subsequently, β -centering w.r.t. data mean $\mu = \text{avg}_{n \in \mathcal{N}} \phi_n$ is obtained as $\phi_n^{(\beta)} = \phi_n - \beta\mu$ for $0 \leq \beta \leq 1$. For brevity, we drop superscript (β) from $\phi_n^{(\beta)}$.

The role of β -centering is to address anti-occurrences *i.e.*, some Bag-of-Word models use so-called negative visual words, the evidence that a given visual stimulus is missing from an image. Authors of [51] define it as ‘*the negative evidence i.e., a visual word that is mutually missing in two descriptions being compared*’. Lack of certain visual stimuli correlates with some visual classes *e.g.*, lack of the sky may imply an indoor scene. Thus, we offset vectors ϕ by $\beta\mu$ (μ computed per-image) so that the positive/negative values yield correlations/anti-correlations.

Positional embedding. As in Prop. 1, let Φ_A and Φ_B be obtained from the last conv. feature maps of CNN for images Π_A and Π_B . Then, Cartesian (spatial) coordinates [5, 13] at which feature vectors are extracted along the channel mode from conv. feature maps (Φ_A and Φ_B) are normalized to range [0, 1] yielding (x, y) and (x^*, y^*) which are embedded into a non-linear Hilbert space. Firstly, the normalization is performed as $x_n = x'_n / (W - 1)$ and $y_n = y'_n / (H - 1)$ w.r.t. the width W and height H of conv. feature maps, where x'_n and y'_n are Cartesian coordinates. Then, we form the following sum kernel and its linearization:

$$K([x_n; y_n], [x'_n; y'_n]) = \alpha^2 G_\sigma(x_n - x'_n) + \alpha^2 G_\sigma(y_n - y'_n) \approx \langle \alpha\varphi(x_n, \zeta), \alpha\varphi(x'_n, \zeta) \rangle + \langle \alpha\varphi(y_n, \zeta), \alpha\varphi(y'_n, \zeta) \rangle, \quad (3)$$

where $\varphi(x)$ and $\varphi(x^*)$ are feature maps linearizing an RBF kernel $G_\sigma(x - x^*) = \exp(-(x - x^*)^2 / 2\sigma^2) \approx \langle \sqrt{c}\varphi(x), \sqrt{c}\varphi(x^*) \rangle$:

$$\varphi(x) = [G_{\sigma/\sqrt{2}}(x - \zeta_1), \dots, G_{\sigma/\sqrt{2}}(x - \zeta_Z)]^T, \quad (4)$$

where $\sigma > 0$ is the RBF bandwidth, const. $c = 1$. For Z pivots $\zeta = [\zeta_1; \dots; \zeta_Z]$, we use Z in range 3–10 and equally spaced intervals $\zeta = [-0.2 : 1.4 / (Z - 1) : 1.2]$ to encode x_n and y_n . The derivation and the choice of pivots are explained in Appendix L.

The above formulation extends to the aggregation over patches extracted from images as shown in Figure 1a. We form vectors $\bar{\phi}_n = [\phi_n; \mathbf{c}_n]$ augmented by encoded spatial coordinates $\mathbf{c}_n = [\alpha\varphi(x_n, \zeta); \alpha\varphi(y_n, \zeta)]$. Thus, we define the total length of \mathbf{c}_n as $Z' = 2Z$. Then, we define $\bar{\Phi} \equiv \{\bar{\phi}_n\}_{n \in \mathcal{N}}$. Combining augmented vectors with the Prop. 1 and Eq. (2) yields $M(\bar{\Phi})$ (*c.f.* $M(\Phi)$).

We note that the autocorrelation matrix is sometimes called as

co-occurrence matrix in the literature. If $\{\bar{\phi}_n, n \in \mathcal{N}\}$ contains only binary features $\{0, 1\}$ then the autocorrelation matrix is a form of co-occurrence matrix normalized by N that captures counts $M_{kl} = \frac{1}{N} \sum_{n \in \mathcal{N}} \bar{\phi}_{kn} \cdot \bar{\phi}_{ln}$.

Our pipelines apply element-wise or spectral pooling $\Psi = \mathcal{G}_i(M(\bar{\Phi}); \cdot)$ or $\Psi = \hat{\mathcal{G}}_i(M(\bar{\Phi}); \cdot)$, resp., where $M(\bar{\Phi})$ is an autocorrelation/covariance matrix defined above, $\mathcal{G}_i(M(\bar{\Phi}); \cdot)$ and $\hat{\mathcal{G}}_i(M(\bar{\Phi}); \cdot)$ are element-wise and spectral Power Normalizations on $M(\bar{\Phi})$ with some parameter ‘ \cdot ’ and i is replaced by a specific name of PN. Finally, Ψ is the resulting PN feature map. For brevity, we often drop Ψ , i and ‘ \cdot ’, and write M (*c.f.* $M(\bar{\Phi})$).

4.2 Well-motivated Pooling Approaches (second-order element-wise variants on non-negative features: $M_{kl} \geq 0$)

Following the first branch of the taxonomy in Fig. 1b (element-wise operators), we extend PN operators MaxExp and Gamma (first-order variants) introduced in Section 3.3 to their second-order element-wise counterparts (pooling acts on elements of autocorrelation matrix). As such models lack any previous analysis, we demonstrate that MaxExp for co-occurrences emerges naturally from the Multinomial modeling which also gives it a nice interpretation as a co-occurrence detector.

Derivation. Prop. 2 states that $1 - (1 - p)^N$ is the probability of at least one success being detected in the pool of the N i.i.d. trials $\phi \in \{0, 1\}^N$ following the Bernoulli dist. (success prob. p). We extend Prop. 2 to the case of co-occurrences as follows.

Theorem 1. Let two event vectors $\phi, \phi' \in \{0, 1\}^N$ store N trials each, performed according to the Bernoulli distribution under i.i.d. assumption, for which the probability p of an event $((\phi_n = 1) \wedge (\phi'_n = 1))$ denotes a co-occurrence, and $1 - p$ for $((\phi_n = 0) \vee (\phi'_n = 0))$ denotes the lack of it. Let p be estimated as $p = \text{avg}_n \phi_n \phi'_n$. Then the probability of at least one co-occurrence event $((\phi_n = 1) \wedge (\phi'_n = 1))$ in ϕ_n and ϕ'_n simultaneously in N trials becomes:

$$\psi = 1 - (1 - p)^N. \quad (5)$$

Proof. The probability of all N outcomes to be $\{((\phi_1 = 0) \vee (\phi'_1 = 0)), \dots, ((\phi_N = 0) \vee (\phi'_N = 0))\}$ is $(1 - p)^N$. The probability of at least one positive outcome $((\phi_n = 1) \wedge (\phi'_n = 1))$ amounts to the probability of event $((\phi_1 = 1) \wedge (\phi'_1 = 1)) \vee \dots \vee ((\phi_N = 1) \wedge (\phi'_N = 1))$ equal $1 - (1 - p)^N$, where $p = \text{avg}_n \phi_n \phi'_n$.

A stricter proof uses a Multinomial distribution model with four events for (ϕ_n) and (ϕ'_n) which describe all possible outcomes. Let probabilities p, q, s and $1 - p - q - s$ add up to 1 and correspond to events $((\phi_n = 1) \wedge (\phi'_n = 1))$, $((\phi_n = 1) \wedge (\phi'_n = 0))$, $((\phi_n = 0) \wedge (\phi'_n = 1))$ and $((\phi_n = 0) \wedge (\phi'_n = 0))$. The first event is a co-occurrence, the latter two are occurrences only and the last

event is the lack of the first three events. The probability of at least one co-occurrence $((\phi_n = 1) \wedge (\phi'_n = 1))$ in N trials becomes:

$$\sum_{n=1}^N \sum_{n'=0}^{N-n} \sum_{n''=0}^{N-n-n'} \binom{N}{n, n', n'', N-n-n'-n''} p^n q^{n'} s^{n''} (1-p-q-s)^{N-n-n'-n''}. \quad (6)$$

One can verify algebraically/numerically that Eq. (6) and (5) are equivalent w.r.t. p which completes the proof. \square

A proof with a Multinomial distribution shows that MaxExp for co-occurrences (Eq. (5)) has exactly the same form as MaxExp for vectors (Prop. 2), and it acts as a co-occurrence detector. This justifies why Eq. (5) is applicable both to first- and second-order representations in the element-wise pooling regime. Below we explain practical details of this pooling *e.g.*, how to apply it to an autocorrelation/covariance matrix and why this is meaningful.

MaxExp pooling (second-order element-wise). In practice, we have $\psi_{kl} = 1 - (1 - \tilde{p}/\kappa)^\eta$, where $0 < \eta \approx N$ is an adjustable parameter, $\tilde{p} = \text{avg}_n \phi_{kn} \phi_{ln}$, scalars ϕ_{kn} and ϕ_{ln} are k -th and l -th features of an n -th feature vector *e.g.*, as defined in Prop. 1, and $\kappa > 0$ ensures that $0 \leq \tilde{p}/\kappa \leq 1$. As \tilde{p} is an expected value over N channel-wise correlations (product op.) of feature pairs from rectified CNN maps rather than co-occurrences of binary variables, we assume that $\phi_{kn} \phi_{ln}$ is proportional to the confidence that simultaneous detection of stimuli that channels k and l represent is correct, and $\kappa \geq \max_n \phi_{kn} \phi_{ln}$ is a value corresponding to the confidence equal one. We observe that:

$$1 - \prod_{n \in N} (1 - \phi_{kn} \phi_{ln} / \kappa) \geq 1 - (1 - \alpha \tilde{p} / \kappa)^{\eta^*} \approx 1 - (1 - p)^N, \quad (7)$$

where $\eta^* = N$ and $\alpha = 1$. The left-hand side eq. of (7) is the likelihood of at least one co-occurrence if drawing from an unknown distribution under the i.i.d. assumption, thus we may desire to adjust the middle eq. in (7) toward this upper bound. As the proportionality assumption may be violated in practice, $\alpha \neq 1$ helps achieve a good estimate $p \approx \alpha \tilde{p} / \kappa$ and/or adjust the middle eq. in (7) toward the left-hand eq. in (7). Parameters η^* and α may be tied together as we observe that $\log((1 - \alpha \tilde{p} / \kappa)^{\eta^*}) = \eta^* \log(1 - \alpha \tilde{p} / \kappa) \approx -\eta^* \alpha \tilde{p} / \kappa = -\eta \tilde{p} / \kappa \approx \log((1 - \tilde{p} / \kappa)^\eta)$ as $\log(1 - x) \approx -x$ and $\eta = \eta^* \alpha$ (tied parameter). Reversing the logarithm operation yields $1 - (1 - \tilde{p} / \kappa)^\eta \approx 1 - (1 - \alpha \tilde{p} / \kappa)^{\eta^*}$ where $\eta \approx N$ refines $1 - (1 - \tilde{p} / \kappa)^\eta$ towards left- and right-hand side equations in (7). In matrix form, we have:

$$\mathcal{G}_{\text{MaxExp}}(\mathbf{M}; \eta) = 1 - (1 - \mathbf{M} / (\text{Tr}(\mathbf{M}) + \varepsilon))^\eta, \quad (8)$$

where $\kappa = \text{Tr}(\mathbf{M}) + \varepsilon$ to ensure $\kappa \geq \max_{k,l} M_{kl}$ is sufficiently large, $\varepsilon \approx 1e-6$, the global param. η is chosen via cross-validation

Pooling function (operator i)	$g_i(p)$ if $p < 0$	$g'_i(p)$ if $p \rightarrow 0^+ (0^-)$	$g_i(p)$	$g'_i(p)$
Gamma [5]	inv.	$\infty (-\infty \sqrt{-1})$	p^γ	$\gamma p^{\gamma-1}$
MaxExp [5]	inv.	fin.: η	$1 - (1-p)^\eta$	$\eta(1-p)^{\eta-1}$
AsinhE	ok	fin.: $\frac{\gamma'}{\sqrt{1+\gamma'^2}}$	$\text{Asinh}(\gamma' p)$	$\frac{\gamma'}{\sqrt{1+\gamma'^2 p^2}}$
SigmE	ok	fin.: $0.5\eta'$	$\frac{2}{1+e^{-\eta' p}} - 1$	$\frac{2\eta' e^{-\eta' p}}{(1+e^{-\eta' p})^2}$
HDP	inv.	$0 (\infty)$	$e^{-t/p}$	$\frac{te^{-t/p}}{p^2}$

TABLE 1: A collection of Power Normalization functions. Variables $\gamma > 0$, $\gamma' > 0$, $\eta \geq 1$, $\eta' \geq 1$ and $t > 0$ control the level of power normalization. For specific pooling type indicated by i , we indicate properties of $g_i(p)$ such as finite (*fin.*) or infinite (∞) derivative $g'_i(p)$ at $p \rightarrow 0^+ (0^-)$, and invalid (*inv.*) or valid (*ok*) power norm. if $p < 0$. HDP is a spectral operator (Sec. 5.3) but it might act element-wise too.

to compensate for an estimate of κ , violation of the proportionality assumption, variations of distr., and the approx. of logarithm.

Remark 1. $\mathcal{G}_{\text{MaxExp}}^*(\mathbf{M}; \eta) = \mathcal{G}_{\text{MaxExp}}(\mathbf{M}; \eta) (\text{Tr}(\mathbf{M}) + \varepsilon)^\gamma$, $\gamma \approx \frac{1}{2}$, compensates for the trace in (8) which affected the input-output ratio of norms. $\mathcal{G}_{\text{MaxExp}}^\ddagger(\mathbf{M}, \eta) = \mathcal{G}_{\text{MaxExp}}(\mathbf{M}; \eta) + \kappa \mathbf{M}$ prevents vanishing gradients in pooling. Both terms can be combined. We suggest adding a small linear slope by $\mathcal{G}^\ddagger(\kappa > 0)$ if one experiences vanishing gradients. We use \mathcal{G}^* only for fine-tuning on off-the-shelf pre-trained CNN which produces feature vectors with the ℓ_1 norms varying from region to region and/or image to image. As these norms have an impact on the quality of separation between different class concepts (because originally they were not excluded from training), they need to be adapted to the new dataset.

Gamma pooling (second-order element-wise). For completeness, we present Gamma pooling based on the def. in Sec. 3.3:

$$\mathcal{G}_{\text{Gamma}}(\mathbf{M}; \gamma) = (\mathbf{M} + \varepsilon)^\gamma. \quad (9)$$

Rising \mathbf{M} to the power of γ is element-wise, ε is a small reg. constant. Appendix B lists derivatives of Gamma and MaxExp.

4.3 From MaxExp to MaxExp(\pm) to SigmE (motivating second-order element-wise variants for $M_{kl} \in \mathbb{R}$).

Moving one branch down in the taxonomy from Fig. 1b, we note that matrix \mathbf{M} is built from features that may have negative values (β -centering or non-rectified \mathbf{M}). Negative entries of M_{kl} break MaxExp/Gamma. Thus, we extend Prop. 1 to MaxExp(\pm) pooling which works also with negative co-occurrences interpreted by us as two anti-correlating visual words. To interpret such a pooling variant, we show that if trials follow a mixture of two Normal distributions, we obtain SigmE pooling (a zero-centered sigmoid function) which acts as a detector of co-occurrence/negative co-occurrence hypothesis. To establish affinity between MaxExp(\pm) and SigmE, we show that both functions are identical if their parameters $\eta, \eta' \rightarrow \infty$. However, the derivative of MaxExp(\pm) is non-smooth at 0 while SigmE has an almost identical profile to MaxExp(\pm) but its der. is smooth (important in optimization).

Derivation. MaxExp(\pm) and SigmE are derived in Proposition 3 and Theorem 2 while a non-essential Remark 2 proves their affinity.

Proposition 3. Let event vectors $\phi, \phi' \in \{-1, 0, 1\}^N$ store N trials performed according to the Multinomial distr. under i.i.d. assumption, for which we have the probability p of a co-occurrence event $((\phi_n = 1) \wedge (\phi'_n = 1)) \vee ((\phi_n = -1) \wedge (\phi'_n = -1))$, the probability q of a anti-correlating co-occurrence event $((\phi_n = 1) \wedge (\phi'_n = -1)) \vee ((\phi_n = -1) \wedge (\phi'_n = 1))$, and $1-p-q$ for $((\phi_n = 0) \vee (\phi'_n = 0))$ which denotes the lack of the first two events. Then the prob. difference between at least one co-occurrence event and at least one anti-correlating co-occurrence event in N trials, encoded by ϕ, ϕ' becomes $\psi = (1-q)^N - (1-p)^N$.

Proof. See Appendix A. \square

MaxExp(\pm) (second-order element-wise). In practice, we simply estimate p and q as $p = \max(0, p^*)$ and $q = \max(0, -p^*)$, $p^* = \text{avg}_n \phi_n^*$. Note $p^* < 0$ or $p^* > 0$ if the majority of co-occurrences between event vectors ϕ, ϕ' captured as $\phi^* = \phi \odot \phi' \in \{-1, 0, 1\}^N$ are anti-correlating or correlating, respectively.

Below we show that SigmE has a derivation that follows a slightly different statistical interpretation than MaxExp, which explains that SigmE is a likelihood-based detector of co-occurrence vs. negative co-occurrence hypothesis.




Fig. 3: The profile of HDP is shown in Fig. 3a. Note the similarity of HDP to MaxExp. Fig. 3b shows that MaxExp and Gamma given by $g_{\text{MaxExp}}(\lambda)$ and $g_{\text{Gamma}}(\lambda)$ are in fact upper bounds of HDP ($g_{\text{HDP}}(\lambda)$).

Theorem 2. Assume an event vector $\phi^* \in \mathbb{R}^N$ whose coefficients ϕ_n^* represent anti-occurrences or occurrences drawn from $\mathcal{N}(-1, 2)$ or $\mathcal{N}(1, 2)$, resp. Then the probability of an event ϕ_n^* being an anti-occurrence and occurrence is given by $\psi^{(-)}(\phi_n^*) = \frac{G_\sigma(\phi_n^*+1)}{G_\sigma(\phi_n^*-1)+G_\sigma(\phi_n^*+1)}$ and $\psi^{(+)}(\phi_n^*) = \frac{G_\sigma(\phi_n^*-1)}{G_\sigma(\phi_n^*-1)+G_\sigma(\phi_n^*+1)}$, resp. We note that $\psi^{(-)}(0) = \psi^{(+)}(0) = 0.5$ which means that for $\phi_n^*=0$, the determination of event type that the feature represents cannot be made in the statistical sense. As we want to factor out such cases, we simply set $\psi = \psi^{(+)}(\phi_n^*) - \psi^{(-)}(\phi_n^*)$ which reduces to $\frac{2}{1+e^{-\eta'\phi_n^*}} - 1$ for $\eta' = 2/\sigma^2$. Thus, if $p^* = \text{avg}_n \phi_n^*$, SigmE given by $\frac{2}{1+e^{-\eta'p^*}} - 1$ (see Eq. (10)) simply tells, on average, whether events ϕ_n came more likely from $\mathcal{N}(-1, 2)$ or $\mathcal{N}(1, 2)$.

Proof. Appendix J is the proof. \square

Remark 2. One can verify that for $\eta \rightarrow \infty$ and $\eta' \rightarrow \infty$, we have $\lim_{\eta, \eta' \rightarrow \infty} \int_0^1 \left| \frac{2}{1+e^{-\eta'p}} - 1 - (1 - (1-p)^\eta) \right| dp = 0$. In the limit, both formulations are identical on interval $[0, 1]$. For the finite η and η' , minimizing the above integral has no closed form but parametrizations $\eta'(\eta) = \frac{\log(\sqrt{3}+2)}{1 - ((4-2\sqrt{3})/3)^{1/(2\eta)}}$ and $\eta(\eta') = \frac{\log((4-2\sqrt{3})/3)}{2 \log(1 - \log(\sqrt{3}+2)/\eta')}$ yield a low approx. error due to aligning SigmE with MaxExp at a point $p''(\eta') = \log(\sqrt{3}+2)/\eta'$ for which the concavity of SigmE on interval $[0, 1]$ is at its maximum.

Proof. Appendix K is the proof. \square

SigmE pooling (second-order element-wise). In practice, zero-centered Logistic a.k.a. Sigmoid (*SigmE*) functions below may be used in lieu of MaxExp(\pm) and MaxExp (Eq. (8)):

$$\mathcal{G}_{\text{SigmE}}(\mathbf{M}; \eta') = \frac{2}{1+e^{-\eta' M}} - 1 \text{ or } \frac{2}{1+e^{\frac{-\eta' M}{\text{Tr}(\mathbf{M})+\varepsilon}}} - 1. \quad (10)$$

AsinhE pooling (second-order element-wise). For completeness, we present AsinhE, an alternative to Gamma in Eq. (9) as Gamma has an infinite derivative for $M_{kl} \rightarrow 0$ and $\varepsilon \rightarrow 0$, and assumes $M_{kl} \geq 0$. Its reg. $\varepsilon > 0$ may affect results as Gamma magnifies signals close to 0 i.e., $\varepsilon \approx 1e-3$ will mask smaller signals.

AsinhE pooling (Arcsin hyperbolic) by contrast has an almost identical profile with Gamma, it has a finite/smooth derivative and operates on $M_{kl} \in \mathbb{R}$ w/o regularization:

$$\mathcal{G}_{\text{AsinhE}}(\mathbf{M}; \gamma') = \log(\gamma' M + \sqrt{1 + \gamma'^2 M^2}), \quad (11)$$

Table 1 lists properties of Power Normalizations. Figure 2 illustrates MaxExp/Gamma and their extensions SigmE/AsinhE (for $p \in \mathbb{R}$) whose derivatives (Appendix C) are smooth/finite.

Discussion. Pooling functions with similar profiles can be formulated in numerous ways. MaxExp is shown to act as a burstiness-suppressing detector of at least one event in the collection of events

drawn from the Bernoulli distribution. The Gamma operator can be viewed as a function that whitens signal. SigmE simply tells, on average, whether a collection of events ϕ_n is more likely to come from the Normal distribution representing anti-occurrences ($\mu = -1$) or its counterpart representing occurrences ($\mu = 1$). The unifying factor for MaxExp and SigmE is how events are modeled e.g., according to the Bernoulli or Normal distributions. As element-wise operators do not take into account correlations between features, below we introduce and study Spectral Power Normalizations (second main taxonomy branch of Fig. 1b).

5 SPECTRAL POWER NORMALIZATIONS

According to the main branch two in Fig. 1b, Spectral Power Normalizations act on the spectrum of autocorrelation/covariance matrices instead of individual pairs of coefficients in order to respect the data correlation between principal directions of the multivariate Normal Distribution represented by \mathbf{M} . In this section, we provide a generalized recipe on computations of SPNs, Moreover, we show that Gamma/MaxExp are upper bounds of the time-reversed ($t < 1$) Heat Diffusion Process [14] which explains their good performance in classification (fourth main taxonomy branch of Fig. 1b). From these considerations emerges our culminating contribution, a fast spectral MaxExp that rivals recent Newton-Schulz iterations [34, 35].

5.1 Generalized SPNs

Spectral pooling versions of our pooling operators can be obtained via SVD given as $\mathbf{U}\boldsymbol{\lambda}\mathbf{U}^T = \mathbf{M}$. Let us define:

$$\widehat{\mathcal{G}}_i(\mathbf{M}) = \mathbf{U} \text{diag}^\dagger(\mathbf{g}_i(\mathbf{p}(\boldsymbol{\lambda}))) \mathbf{U}^T, \quad (12)$$

where Spectral Power Normalizations are realized by operators $\mathbf{g}_i(\cdot)$ in Table 1, $\mathbf{p}(\boldsymbol{\lambda}) = [p_1, \dots, p_K] = \frac{\text{diag}(\boldsymbol{\lambda})}{\text{Tr}(\boldsymbol{\lambda})+\varepsilon}$, the normalization by $\text{Tr}(\boldsymbol{\lambda})+\varepsilon$ applies to MaxExp and SigmE. A small constant $\varepsilon \approx 1e-6$ prevents the vanishing trace, $\text{diag}(\cdot)$ extracts the diagonal from $\boldsymbol{\lambda}$ into a vector while $\text{diag}^\dagger(\cdot)$ places a vector into the diagonal of a matrix (off-diagonal coeffs. equal zero).

For a generic back-propagation through Eq. (12), one relies on the back-propagation through SVD and the following chain rule:

$$\frac{\partial \widehat{\mathcal{G}}_i(\mathbf{M})}{\partial M_{kl}} = \frac{\partial \mathbf{U} \text{diag}^\dagger(\mathbf{g}_i(\mathbf{p}(\boldsymbol{\lambda}))) \mathbf{U}^T}{\partial M_{kl}} = 2 \text{Sym} \left(\frac{\partial \mathbf{U}}{\partial M_{kl}} \text{diag}^\dagger(\mathbf{g}_i(\mathbf{p}(\boldsymbol{\lambda}))) \mathbf{U}^T \right) + \mathbf{U} \text{diag}^\dagger \left(\frac{\partial \mathbf{g}_i(\mathbf{p}(\boldsymbol{\lambda}))}{\partial M_{kl}} \right) \mathbf{U}^T. \quad (13)$$

The back-propagation through eigenvectors and eigenvalues of SVD is a well-studied problem [52, 53, 54]:

$$\frac{\partial \lambda_{ii}}{\partial M} = \mathbf{u}_i \mathbf{u}_i^T, \quad \frac{\partial u_{ij}}{\partial M} = u_{ij} (\lambda_{jj} \mathbf{I} - M)^\dagger. \quad (14)$$

Algorithm 1 Generalized Spectral Power Normalization

Input: \mathbf{M} for a given forward pass, ε' : desired spectral gap, d : side dim. of \mathbf{M} , $i \in \{\text{Gamma}, \text{MaxExp}, \text{AsinhE}, \text{SigmE}, \text{HDP}\}$

- 1: $k = 0$ /*prevent non-simple eigenvalues*/
- 2: **while** True:
- 3: $k \leftarrow k + 1$, $\{\xi_i \sim \mathcal{U}(\varepsilon', \varepsilon' + d\varepsilon')\}_{i=1, \dots, d}$
- 4: **try:**

- 5: $\mathbf{U}\boldsymbol{\lambda}\mathbf{U}^T = \mathbf{M} + \text{diag}(k\xi)$
- 6: get $\widehat{\mathcal{G}}_i(\mathbf{M})$ and $\widehat{\mathcal{G}}_i'(\mathbf{M})$ by Eq. (12), (13) and (14)
- 7: **else:**
- 8: break

Output: $\widehat{\mathcal{G}}_i(\mathbf{M})$ and $\widehat{\mathcal{G}}_i'(\mathbf{M})$

As the back-propagation through SVD breaks down in the presence of so-called non-simple eigenvalues, that is $\lambda_i = \lambda_j : i \neq j$, Algorithm 1 ensures that we draw small regularization coefficients ξ_i from the uniform distribution until the spectral gap ε' is ensured.

Combining Eq. (13) and (14) yields the same back-propagation equation as in [32]. Table 2 shows closed-form expressions for Power Normalizations realized via operators in Table 1.

5.2 Fast Spectral MaxExp.

Below we present our culminating contribution: the fast spectral MaxExp. Running SVD is slow and back-prop. via SVD suffers from large errors as the spectral gap between eigenvalues narrows. Thus, a very recent trend in spectral pooling is to use Newton-Schulz iterations [35] to obtain a fast stable approximate matrix square root and its derivative, a special case of spectral Gamma with $\gamma = 0.5$ which is often argued to be a close-to-optimal parameter for Power Normalization. However, contradictory observations come from [5, 55] (see Table 1 in [55]). To this end, we point that our spectral MaxExp has a nice property—its forward pass and its derivative can be computed fast for any $\eta \geq 1$ with matrix-matrix multiplications (no SVD is needed).

Without the loss of generality, if M is trace-normalized then spectral MaxExp $\hat{\mathcal{G}}_{\text{MaxExp}}(M) = \mathbb{I} - (\mathbb{I} - M)^\eta$ (15) has the closed-form derivative defined as:

$$\frac{\partial \hat{\mathcal{G}}_{\text{MaxExp}}(M)}{\partial M_{kl}} = - \sum_{n=0}^{\eta-1} (\mathbb{I} - M)^n \mathbf{J}_{kl} (\mathbb{I} - M)^{\eta-1-n}, \quad (16)$$

whose detailed form given in Appendix N scales linearly (runtime) w.r.t. η . However, evaluation time of MaxExp and its derivative can scale sublinearly w.r.t. η given Algorithm 2.

Forward pass. Given an integer $\eta \geq 1$, computing $(\mathbb{I} - M)^\eta$ has a lightweight complexity $\mathcal{O}(m \log(\eta))$, where m is the cost of the matrix-matrix multiplication in exponentiation by squaring [56] whose cost is $\log(\eta)$. In contrast, the matrix square root via k Newton-Schulz iterations has complexity $\mathcal{O}(mk)$. Head-to-head, our subroutine performed 8 matrix-matrix multiplications for $\eta = 50$ (typically $20 \leq \eta \leq 80$ as in Fig. 9d). Newton-Schulz iter. performed $3k = 60$ matrix-matrix mult. for $k = 20$ set in [35].

Backward pass. Given an integer $\eta \geq 1$, the required powers of $(\mathbb{I} - M)$ come from the forward-pass. Auto-differentiation runs along the recursion path of exponentiation by squaring whose complexity is $\mathcal{O}(m \log(\eta))$. In contrast, der. of the matrix square root via Newton-Schulz iterations has complexity $\mathcal{O}(mk)$ ($k \approx 20$ [35]). Head-to-head, we required 11 matrix-matrix multiplications (3×2 and 5×1 for the derivative in line 4 and 8 respectively of Alg.

¹Application of HDP to fine-grained classification is also our contribution.

oper. i	$\hat{\mathcal{G}}_i(M)$	type of back-propagation (speed)
Gamma	M^γ (e.g., $M^{\frac{1}{2}}$)	Sylv. Eq. (37) of App. D (extremely slow) / Sylv. Eq. (37) via Bartels-Stewart alg. [35] (slow) / SVD back-prop. Eq. (13) (slow) / Newton-Schulz [35] (fast)
MaxExp	$\mathbb{I} - (\mathbb{I} - \frac{M}{\text{Tr}(M)+\varepsilon})^\eta$	SVD Eq. (13) (slow) / Eq. (55) (fast) / Alg. 2 (very fast)
AsinhE	$\text{Log}(\gamma' M + (\mathbb{I} + \gamma'^2 M^2)^{\frac{1}{2}})$	SVD back-prop. Eq. (13) (slow)
SigmE	$2 \left(\mathbb{I} + \text{Exp}(\frac{-\eta' M}{\text{Tr}(M)+\varepsilon}) \right)^{-1} \mathbb{I}$	SVD back-prop. Eq. (13) (slow)
HDP	$\text{Exp}(-t M^\dagger)$	SVD back-prop. Eq. (13) (slow)

TABLE 2: Spectral Power Normalizations. The square, square root, power, log and exp are matrix operations. The type of back-propagation available and speed are indicated. Gamma is the only well-explored Power Norm. yet other SPNs also work well. They are all approx. of the time-reversed Heat Diffusion Process (Sec. 5.3).

2) for $\eta = 50$. Newton-Schulz iter. require $6k = 120$ matrix-matrix mult. for $k = 20$. Memory-wise, MaxExp and Newton-Schulz iter. need to store $\sim 2 \log \eta$ and $2k$ matrices, respectively.

Finally, the complexity of an SVD is $\mathcal{O}(d^\omega)$ with $2 < \omega < 2.376$. Non-spectral PNs are faster by the order of magnitude.

5.3 Spectral MaxExp is a Time-reversed Heat Diffusion.

Our key theoretical contribution below shows that the Heat Diffusion Process¹ (HDP) on a graph Laplacian is closely related to the Spectral Power Normalization (SPN) of the autocorrelation/covariance matrix (see second/fourth main taxonomy branch of Fig. 1b), whose inverse forms a loopy graph Laplacian. To this end, we firstly explain the relation between an autocorrelation/covariance and a graph Laplacian matrices. Subsequently, we establish that for HDP with $t < 1$, MaxExp and Gamma functions are tight upper bounds of HDP for some parametrizations $\eta(t)$ and $\gamma(t)$. Finally, we reconsider the system of Ordinary Differential Equations describing HDP and show that MaxExp and Gamma correspond to modified heat diffusion ODEs. Figure 4 shows that HDP, spectral MaxExp and Gamma play the same role *i.e.*, they boost or dampen the magnitudes of the eigenspectrum to concentrate eigenvalues around a single peak, thus reversing the diffusion of signal in autocorrelation/covariance matrices.

Theorem 3. A structured multivariate Gaussian distribution with a covariance matrix M is associated with a weighted graph such that the precision matrix $Q \equiv M^\dagger$ corresponds to the loopy graph Laplacian [57]. Let M be trace-normalized so that $0 < \sum_i \lambda_i \leq 1$, and $\hat{\mathcal{G}}_{\text{MaxExp}}(M; \eta) = U \text{diag}(1 - (1 - \lambda)^\eta) U^T$ be our spectral MaxExp operator. Let $K(Q; t) = \text{Exp}(-tQ)$ with $t > 0$ be the Heat Diffusion Process on the graph. Then $\forall \eta > 1$, $\hat{\mathcal{G}}_{\text{MaxExp}}(M; \eta)$ can be well approximated by some $K(Q; t)$, $0 < t < 1$. Similarly, $K(Q; t)$ can approximate Gamma.

Proof. Appendix E is the proof. \square

The loopy graph [57] Laplacian (see Fig. 1b) represents each feature (that corresponds to a given CNN filter) by a node while edges quantify the similarity between features. The loopy graph is characterized by: (i) symmetric positive definite matrix (*c.f.* semi-definite Laplacian) due to self-loops of each node, (ii) dense connectivity (each node may connect with other nodes), (iii) the underlying multivariate Gaussian distribution, (iv) structural design (*i.e.*, we augment the autocorrelation matrix with spatial coordinates making some nodes location-specific).

As the time parameter t is in the range $(0, 1)$, we call $K(Q; t)$ a time-reversed Heat Diffusion Process. This means that rather than diffusing the heat between the nodes ($t > 1$), the

Algorithm 2 Exponentiation by Squaring in Fast Spectral MaxExp

Input: M for a given forward pass, $\eta \geq 1$

- 1: $G_1 = \mathbb{I}$, $M_1^* = \mathbb{I} - M$, $n = \text{int}(\eta)$, $t = 1$, $q = 1$
- 2: **while** $n \neq 0$:
- 3: **if** $n \& 1$:
- 4: $G_{t+1} = G_t M_q^*$, $\frac{\partial G_{t+1}}{\partial M} = \frac{\partial G_t}{\partial M} M_q^* + G_t \frac{\partial M_q^*}{\partial M}$
- 5: $n \leftarrow n - 1$, $t \leftarrow t + 1$
- 6: $n \leftarrow \text{int}(n/2)$
- 7: **if** $n > 0$:
- 8: $M_{q+1}^* = M_q^* M_q^*$, $\frac{\partial M_{q+1}^*}{\partial M} = 2 \text{Sym} \left(\frac{\partial M_q^*}{\partial M} M_q^* \right)$
- 9: $q \leftarrow q + 1$

Output: $\hat{\mathcal{G}}_{\text{MaxExp}}(M) = \mathbb{I} - G_t$ and $\hat{\mathcal{G}}'_{\text{MaxExp}}(M) = -\frac{\partial G_t}{\partial M}$

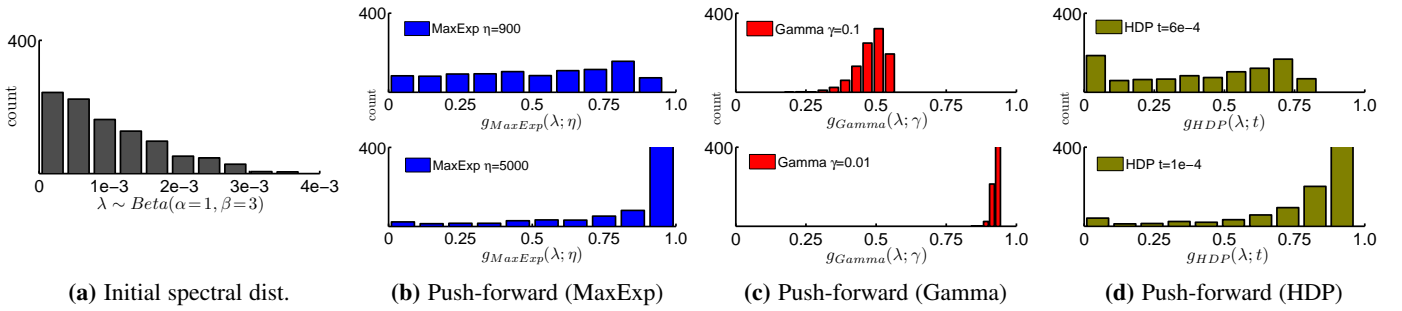


Fig. 4: The intuitive principle of the SPN. Given a discrete eigenspectrum following a Beta distribution in Fig. 4a, the push-forward distributions of MaxExp and HDP in figs. 4b and 4d are very similar. For small γ , Gamma in Fig. 4c is also similar to MaxExp and HDP. Note that all three SPN functions whiten the spectrum (map the majority of values to be ~ 1) thus reversing diffusion (acting as a spectral detector).

model reverses the process in the direction of the identity matrix, that is $\lim_{t \rightarrow 0} \mathbf{K}(\mathbf{Q}; t) \rightarrow \mathbf{I}$, $\lim_{\eta \rightarrow \infty} \hat{\mathcal{G}}_{\text{MaxExp}}(\mathbf{M}; \eta) \rightarrow \mathbf{I}$ and $\lim_{\gamma \rightarrow 0} \mathbf{M}^\gamma \rightarrow \mathbf{I}$. This coincides with so-called eigenspectrum whitening which prevents burstiness [15]. As Theorem 3 does not state any approximation results, we have following theorems.

Theorem 4. $\forall \eta > 1, \exists t(\eta)$ such that MaxExp function is an upper bound of HDP: $1 - (1 - \lambda)^\eta \geq \exp(-t(\eta)/\lambda)$, $\forall \lambda \in [0, 1]$, and gaps ϵ_1 and ϵ_2 between these two functions at $\lambda = t(\eta)$ and $\lambda = \frac{1}{\eta+1}$, where the auxiliary bound $(\frac{1}{e} - 1) \frac{t(\eta)}{\lambda} + 1$ (Appendix F) touches HDP and MaxExp as in Fig. 12b (supp. mat.), resp., satisfy:

$$\frac{e-1}{e} - \left(1 - \frac{e}{e-1} \frac{\eta^\eta}{(\eta+1)^{\eta+1}}\right)^\eta = \epsilon_1 \leq \epsilon_2 = 1 - \left(\frac{\eta}{\eta+1}\right)^\eta - \frac{e}{e-1} \left(\frac{\eta}{\eta+1}\right)^\eta, \quad (17)$$

One possible parametrization $t(\eta)$ satisfying the above condition is:

$$t(\eta) = \frac{e}{e-1} \frac{\eta^\eta}{(\eta+1)^{\eta+1}}, \quad (18)$$

and conversely:

$$\eta(t) \approx 0.5 \sqrt{1 + 4/(t^2(e-1)^2)} - 0.5. \quad (19)$$

Proof. Appendix F is the proof. \square

Theorem 5. Gamma function is an upper bound of HDP, that is $\lambda^{\gamma(t)} \geq \exp(-t/\lambda)$, and there exist a direct point other than at $\lambda=0$ where Gamma and HDP touch. Moreover, the corresponding parametrizations are $\gamma(t) = et$ and $t(\gamma) = e^{-1/\gamma}$.

Proof. Appendix G is the proof. \square

Theorems 4 and 5 give us a combined tighter bound $\min(1 - (1 - \lambda)^{\eta(t)}, \lambda^{\gamma(t)}) \geq \exp(-t/\lambda)$. Theorems 4 and 5 can now be used in the following theorems connecting MaxExp and Gamma to the Heat Diffusion Equation (HDE) [58] which is the system of the Ordinary Differential Equations describing HDP. The HDE is given as:

$$\frac{\partial \mathbf{v}(t)}{\partial t} + \mathbf{L} \mathbf{v}(t) = 0, \quad (20)$$

where vector $\mathbf{v}(t) \in \mathbb{R}^d$ describes some heat quantity of graph nodes at a time t , where $\mathbf{L} \in \mathcal{S}_+^d$ (or \mathcal{S}_{++}^d) is the graph Laplacian (or the loopy graph Laplacian).

Theorem 6. MaxExp can be expressed as a modified Heat Diffusion Equation, where the largest eigenvalue of \mathbf{M} is assumed to be $\lambda_{\text{max}} \leq 1$, and $\mathbf{M} = \mathbf{L}^\dagger$, thus we have:

$$\frac{\partial \mathbf{v}(t)}{\partial t} + \frac{\partial \eta}{\partial t} \text{Log}(\mathbf{I} - \mathbf{M})(1 - \mathbf{v}(t)) = 0. \quad (21)$$

Proof. Appendix H is the proof. \square

Theorem 7. Gamma can be expressed as a modified Heat Diffusion Equation: $\frac{\partial \mathbf{v}(t)}{\partial t} + e \text{Log}(\mathbf{L}) \mathbf{v}(t) = 0$,

where $\text{Log}(\mathbf{L})$ is a Log-Euclidean map of the (loopy) graph Laplacian i.e., $\mathbf{L} = \mathbf{M}^\dagger$. Thus, Gamma is equal to HDP on a Log-Euclidean map of the loopy graph Laplacian, a theoretical connection between Power-Euclidean and Log-Euclidean metrics.

Proof. Appendix I is the proof. \square

Fast approximate HDP. For completeness of our theoretical discussions, we parametrize MaxExp and Gamma w.r.t. time to devise HDP whose runtime scales sublinearly with t . MaxExp and Gamma are upper bounds of HDP and they both can approximately realize the time-reversed ($t \in (0, 1)$) and time-forward ($t \geq 1$) HDP. MaxExp in the time-reversed regime (and its derivative) can be evaluated very fast for integers $\eta \geq 1$ via matrix-matrix multiplications (Sec. 5.2, Alg. 2). Moreover, Gamma \mathbf{M}^γ in the time-forward regime can be evaluated very fast for integers $\gamma \geq 1$ via matrix-matrix multiplications. While the evaluation time of the derivative in Appendix O scales linearly w.r.t. $\gamma \geq 1$, Gamma and its derivative can be computed in the sublinear time by modified Alg. 2 (modify $\mathbf{M}_1^* := \mathbf{M}$ (line 1), $\hat{\mathcal{G}}_{\text{Gamma}}(\mathbf{M}) := \mathbf{G}_t$ and $\hat{\mathcal{G}}'_{\text{Gamma}}(\mathbf{M}) := \frac{\partial \mathbf{G}_t}{\partial \mathbf{M}}$ (output), replace variable η with γ).

FAHDP uses MaxExp/Gamma for the fast approximation of the time-reversed/time-forward HDP, respectively. For the trace-normalized \mathbf{M} , we have²:

$$\hat{\mathcal{G}}_{\text{FAHDP}}(\mathbf{M}; t) = e^{-t} \cdot \begin{cases} \mathbf{I} - (\mathbf{I} - \mathbf{M})^{h_t(\tilde{\eta}(t))} & \text{if } t < 1 \\ \mathbf{M}^{h_t(\tilde{\gamma}(t))} & \text{if } t \geq 1. \end{cases} \quad (23)$$

In the above eq., the scaling factor e^{-t} ensures the Fast Approximate HDP (FAHDP) and HDP have the same magnitude at $\lambda=1$ (trace-normalized \mathbf{M} has eigenvalues $0 \leq \lambda \leq 1$). Note that for $\lambda=1$ and $t \geq \log(10/9) \approx -0.105$, HDP ($e^{-t/\lambda}$) yields $e^{-t} \leq 0.9$ (or less for larger t) while MaxExp ($1 - (1 - \lambda)^\eta$) and Gamma (λ^γ) yield 1 for $\lambda=1$. The scaling factor has no impact on classification results as it is a constant that depends on t fixed throughout an experiment. However, e^{-t} makes FAHDP and HDP visually similar. Thus, we reparametrize η and γ of MaxExp and Gamma by $\tilde{\eta}(t) = 0.5 + \sqrt{0.25 + (1/(t(e-1)) - 1/e)^2}$ and $\tilde{\gamma}(t) = t$. Appendix P contains proofs and further expansions. Finally, $h_t(\cdot)$ may be the ‘round’ function to ensure that MaxExp/Gamma receive an integer parameter (for which forward/backward steps run fast). If $h_t(x) = x$ or $h_t(x) = \lceil x \rceil$ for $t < 1$ and $h_t(x) = \lfloor x \rfloor$ for $t \geq 1$, FAHDP is an upper bound of HDP on $t \in (0, \infty)$.

Discussion. We have shown in Theorems 4 and 5 that MaxExp and Gamma are tight upper bounds of time-reversed HDP thus having

²For the time-reverse regime of FAHDP, one could use $\mathbf{M}^{\tilde{\gamma}(t)}$, $\tilde{\gamma}(t) < 1$, $0 < t < 1$ in place of $\mathbf{I} - (\mathbf{I} - \mathbf{M})^{h_t(\tilde{\eta}(t))}$ e.g., $\mathbf{M}^{0.5}$ reverses \mathbf{M}^2 . However, Gamma and its back-propagation are slow to compute if $0 < t < 1$.

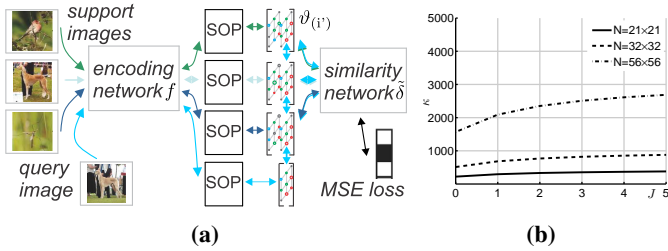


Fig. 5: Figure 5a shows our few-shot pipeline. We use *feature encoding* and *similarity* networks. Power Normalized second-order (SOP) support-query pairs are formed and passed to the similarity net. Support-query pairs of the same class receive positive labels. Fig. 5b is the κ ratio w.r.t. the J -shot value (see Eq. (27)). The curves show that as J grows (0 denotes the regular classification), the similarity learner has to memorize $\kappa \times$ more co-occurrence (k, l) configurations if MaxExp is not used. The curves show for larger N ($N \propto$ image size), not using MaxExp requires even more memorization.

the same role as HDP. We note that MaxExp and Gamma do not require an inversion of spectrum thus they are natural choices for autocorrelation/covariance matrices while HDP is a natural choice for the graph Laplacian matrix. Moreover, in Theorems 6 and 7, we are the first to cast MaxExp and Gamma in the form of modified differential heat equations well-known for HDP. For instance, Gamma is equal to HDP on a Log-Euclidean map of the graph Laplacian, a first concrete result of this kind made in the literature showing how Gamma and HDP relate. The time reversal to a desired state is achieved by setting $t < 1$ of HDP which simply redistributes the heat back to individual graph nodes *e.g.*, all nodes become disconnected in the extreme case. Thus, applying MaxExp and Gamma to autocorrelation matrices has a similar effect, that is, it reduces the level of correlation between features. Thus, some features will remain ‘untouched’ or will be modified to a lesser degree during fine-tuning if they are not related to a new task. We believe this is a very useful property that reduces catastrophic forgetting during fine-tuning. As it reduces the correlation between features, it should also implicitly decorrelate CNN filters.

6 PIPELINES

Classification pipeline from Figure 1a has a straightforward implementation. We pass each image via ImageNet pre-trained ResNet-50, we extract feature vectors Φ from the last conv. layer. Formally, we have the feature encoding network $f : (\mathbb{R}^{3 \times W \times H}; \mathbb{R}^{\mathcal{F}}) \rightarrow \mathbb{R}^{K \times N}$, where W and H are the width and height of an input image, K is the length of feature vectors (number of filters), $N = N_W \cdot N_H$ is the total number of spatial locations in the last convolutional feature map. For brevity, we denote an image descriptor by $\Phi \in \mathbb{R}^{K \times N}$, where $\Phi = f(\mathbf{X}; \mathcal{F})$ for an image $\mathbf{X} \in \mathbb{R}^{3 \times W \times H}$ and \mathcal{F} are the parameters-to-learn of the encoding network. Moreover, where stated, assume that spatial locations discussed in Section 4.1 are concatenated with Φ to obtain $\tilde{\Phi}$. Subsequently, we form autocorrelation matrix M per image according to details of Section 4.1 which is then passed via pooling $\mathcal{G}_i(M)$ or $\tilde{\mathcal{G}}_i(M)$ to the classifier, in end-to-end setting.

Few-shot learning pipeline, called Second-order Similarity Network (SoSN), is shown in Figure 5a. It is inspired by the end-to-end relationship-learning network [45] and consists of two major parts which are (i) feature encoding network and (ii) similarity network. The role of the feature encoding network is to generate convolutional feature vectors which are then used as image descriptors. The role of the similarity network is to learn the relation and compare so-called support and query image embeddings. Our

work is different to the Relation Net [45] in that we apply second-order representations built from image descriptors followed by Power Normalizing functions. For instance, we construct the support and query second-order feature matrices followed by a non-linear Power Normalization unit. In SoSN, the feature encoding network remains the same as Relation Net [45], however, the similarity network learns to compare from second- rather than the first-order statistics. SoSN is illustrated in Figures 1a and 6.

We use the feature encoding network f illustrated in Figure 6 (top) for which spatial locations may be concatenated with representations Φ (see Section 4.1) used by the similarity network.

Figure 6 (bottom) shows the similarity network, which compares two datapoints encoded as $2 \times K \times K$ dim. second-order representations, is denoted by $\tilde{\delta} : (\mathbb{R}^{2 \times K \times K}; \mathbb{R}^{\mathcal{S}}) \rightarrow \mathbb{R}$, where \mathcal{S} are the parameters-to-learn of the similarity network.

Next, let an operator $\vartheta_{(i')} : (\{\mathbb{R}^{K \times N}\}^J; \mathbb{R}^{K \times N}) \rightarrow 2 \times K \times K$ encode a relationship between the descriptors built from the J -shot support images and a query image. This relationship is encoded via computing second-order statistics followed by Power Normalization and applying concatenation (inner-product, sum, subtraction, *etc.* are among other possible choices) to capture a relationship between features of two images. Finally, $i' \in \{\otimes, \otimes+L\}$ takes on one of specific operator variants defined below.

For the L -way J -shot problem, assume some J support images $\{\mathbf{X}_n\}_{n \in \mathcal{J}}$ from some set \mathcal{J} and their corresponding image descriptors $\{\Phi_n\}_{n \in \mathcal{J}}$ which can be considered as a J -shot descriptor if stacked along the third mode. Moreover, we assume one query image \mathbf{X}^* with its image descriptor Φ^* . In general, we use ** to indicate query-related variables. Both the J -shot and the query descriptors belong to one of L classes in the subset $\mathcal{C}^\ddagger \equiv \{c_1, \dots, c_L\} \subset \mathcal{I}_{\mathcal{C}} \equiv \mathcal{C}$ chosen randomly per episode. Similarly to approach [45], we employ the Mean Square Error (MSE) objective in our end-to-end SoSN model. Then, we perform the L -way J -shot learning by:

$$\arg \min_{\mathcal{F}, \mathcal{S}} \sum_{c, c' \in \mathcal{C}^\ddagger} \left(\tilde{\delta}(\vartheta_{(i')}(\{\Phi_n\}_{n \in \mathcal{J}_c}, \Phi_{q \in \mathcal{Q}: \ell(q)=c'}^*); \mathcal{S}) - \delta(c - c') \right)^2,$$

$$\text{where } \Phi_n = f(\mathbf{X}_n; \mathcal{F}) \text{ and } \Phi_q^* = f(\mathbf{X}_q^*; \mathcal{F}). \quad (24)$$

\mathcal{J}_c is a randomly chosen set of J support image descriptors of class $c \in \mathcal{C}^\ddagger$, \mathcal{Q} is a randomly chosen set of L query image descriptors so that its consecutive elements belong to the consecutive classes in $\mathcal{C}^\ddagger \equiv \{c_1, \dots, c_L\}$. Lastly, $\ell(q)$ corresponds to the label of $q \in \mathcal{Q}$ while $\delta(x) = 1$ if $x = 0$, otherwise $\delta(x) = 0$ (note that $c - c' = 0$ if class labels c and c' are the same).

Relationship Descriptor (operator ϑ). We consider two choices for the operator $\vartheta_{(i')}(\{\Phi_n\}_{n \in \mathcal{J}}, \Phi^*) \in \mathbb{R}^{2 \times K \times K}$ whose role is to capture/summarize the information held in support/query image representations to pass it to the similarity network $\tilde{\delta}$ for learning similarity. Below we detail two operators ϑ used by us.

Relationship Descriptor (\otimes) averages J feature maps of support images per class followed by the outer product on the mean support and query vectors, Power Normalization and concatenation. This strategy, beneficial for 84×84 images, is defined as:

$$\vartheta_{(\otimes)}(\{\Phi_n\}_{n \in \mathcal{J}}, \Phi^*) = \left[\mathcal{G}_i \left(\frac{1}{N} \tilde{\Phi} \tilde{\Phi}^T \right);_1 \mathcal{G}_i \left(\frac{1}{N} \Phi^* \Phi^{*T} \right) \right], \quad \tilde{\Phi} = \frac{1}{J} \sum_{n \in \mathcal{J}} \Phi_n, \quad (25)$$

where ‘;1’ is the concatenation along the channel mode, that is, $[\mathbf{X};_1 \mathbf{Y}] \equiv \text{cat}(1, \mathbf{X}, \mathbf{Y})$ in the Matlab notation, $N = WH$, and $i \in \{\text{Gamma}, \text{MaxExp}, \text{AsinhE}, \text{SigmE}, \text{HDP}\}$.

Relationship Descriptor $\otimes+L$ denotes the outer product of feature vectors per support image followed by Power Normalization




Fig. 6: The network architecture used in our SoSN model.

of each matrix and then the average of J such obtained support matrices and concatenation with the query matrix. This strategy, beneficial for large resolution images, is defined as:

$$\vartheta_{(\otimes+L)}(\{\Phi_n\}_{n \in \mathcal{J}}, \Phi^*) = \left[\frac{1}{J} \sum_{n \in \mathcal{J}} \mathcal{G}_i \left(\frac{1}{N} \Phi_n \Phi_n^T \right); \mathcal{G}_i \left(\frac{1}{N} \Phi^* \Phi^{*T} \right) \right]. \quad (26)$$

MaxExp in few-shot learning. Our final analysis shows that MaxExp (second-order element-wise) reduces the burstiness which is otherwise exacerbated in few-shot learning (compared to the regular classification) due to the concatenation operation in relation descriptors as explained below. Theorem 1 states that MaxExp (Power Normalizations in general) performs a co-occurrence detection rather than counting. For classification problems, assume a probability mass function $p_{X_{kl}}(x) = 1/(N+1)$ if $x = 0, \dots, N$, $p_{X_{kl}}(x) = 0$ otherwise, that tells the probability that co-occurrence between ϕ_{kn} and ϕ_{ln} given an image happened $x = 0, \dots, N$ times. Note that classification often depends on detecting a co-occurrence (e.g., is there a flower co-occurring with a pot?) rather than counts (e.g., how many flowers and pots co-occur?). Using second-order pooling without MaxExp requires a classifier to observe $N+1$ tr. samples of *flower and pot* co-occurring in quantities $0, \dots, N$ to memorize all possible co-occurrence configurations. For 1-shot learning, our ϑ stacks pairs of samples to compare, thus a similarity learner now has to deal with a probability mass function of $R_{kl} = X_{kl} + Y_{kl}$ capturing co-occurrence configurations of *flowers and pots* whose $\text{support}(p_{R_{kl}}) = 2N + 1 > \text{support}(p_{X_{kl}}) = N + 1$ as random variable $X = Y$ (same class). The same is reflected by variances *i.e.*, $\text{var}(p_{R_{kl}}) > \text{var}(p_{X_{kl}})$. For J -shot learning, $R'_{kl} = X_{kl}^{(1)} + \dots + X_{kl}^{(J)} + Y_{kl}$, $X^{(j)} = Y, \forall j \in \mathcal{J}$, we have $\text{support}(p_{R'_{kl}}) = (J+1)N + 1$ and the variance grows further indicating that the similarity learner has to memorize more configurations of co-occurrence (k, l) as J grows. However, this situation is alleviated by MaxExp whose probability mass function yields $p_{X_{kl}^{\text{MaxExp}}}(x) = 1/2$ if $x = \{0, 1\}$, $p_{X_{kl}^{\text{MaxExp}}}(x) = 0$ otherwise, as MaxExp detects a co-occurrence (or its lack). For J -shot learning, $\text{support}(p_{R'_{kl}^{\text{MaxExp}}}) = J+2 \ll \text{support}(p_{R'_{kl}}) = (J+1)N+1$.

The ratio

$$\kappa = \frac{\text{support}(p_{R'_{kl}})}{\text{support}(p_{R'_{kl}^{\text{MaxExp}}})} = \frac{(J+1)N+1}{J+2} \quad (27)$$

shows that the similarity learner has to memorize more configurations for (k, l) if no pooling is used relative to configurations if MaxExp is used. As J and/or N increase ($N = WH$ of encoder feature maps), this effect becomes more prominent.

Figure 5b shows how κ varies w.r.t. J and N . Our modeling assumptions are very basic *e.g.*, we use mass functions with uniform probabilities and their set support rather than variances to model the variability of co-occurrences (k, l) . More sophisticated choices *i.e.*, Binomial PMF and variance-based modeling in Appendix M lead to the same theoretical conclusions that: (i) MaxExp (and PN) benefits few-shot learning ($J \geq 1$) even more than it benefits the regular classification ($J = 0$) in terms of reducing possible configurations of (k, l) to memorize, and (ii) for



Fig. 7: Each column shows examples of images from the Flower102, MIT67, FMD and Food-101 dataset, respectively.



Fig. 8: Each column shows examples of fine-grained objects from the Open MIC dataset which look similar but belong to different classes.

larger images (large N), MaxExp (and PN) must reduce a larger number of configurations of (k, l) than for smaller images (smaller N). While classifiers and similarity learners do not memorize all configurations of (k, l) thanks to their generalization ability, they learn quicker if the number of configurations of (k, l) is reduced.

7 EXPERIMENTS

Below we demonstrate experimentally merits of our second-order pooling via Power Normalization functions.

Datasets. For the standard classification setting, we use five publicly available datasets and report the mean top-1 accuracy on them. The Flower102 dataset [59] is a fine-grained category recognition dataset that contains 102 categories of various flowers. Each class consists of between 40 and 258 images. The MIT67 dataset [60] contains a total of 15620 images belonging to 67 indoor scene classes. We follow the standard evaluation protocol, which uses a train and test split of 80% and 20% of images per class. The FMD dataset contains in total 100 images per category belonging to 10 categories of materials (*e.g.*, glass, plastic, leather) collected from the Flickr website. The Food-101 dataset [61], a fine-grained collection of food images from 101 classes, has 101000 images in total and 1000 images per category. Finally, we report top-1 and -5 error on the ImageNet 2012 dataset [62] with 1000 object categories. The dataset contains 1.28M images for training, 50K images for validation and 100K images for testing. As testing labels are withheld, we follow the common practice [34, 63] and report the results on the validation set.

For few-shot classification problems, we also use four publicly available datasets and report the mean top-1 accuracy for so-called L -way J -shot problems [45]. The *mini*ImageNet dataset [42] is a standard benchmark for evaluating few-shot learning approaches. It consists of 60000 RGB images from 100 classes. We follow [42] and use 64 classes for training, 16 classes for validation, remaining 20 classes for testing, and we use images of size 84×84 . We also investigate larger sizes, *e.g.* 224×224 , as our few-shot learning SoSN model can use richer spatial information from larger images

Method	top-1 accuracy	
<i>Second-order Bag-of-Words</i> [5]	90.2	
<i>Factors of Transferability</i> [65]	91.3	
<i>Reversal-inv. Image Repr.</i> [66]	94.0	
<i>Optimal two-stream fusion</i> [67]	94.5	
<i>Neural act. constellations</i> [68]	95.3	
Method	AlexNet	ResNet-50
<i>Baseline</i>	82.00	94.06
<i>FOP</i>	85.40	94.08
<i>FOP+AsinhE</i>	85.64	94.60
<i>SOP</i>	87.20	94.70
<i>SOP+AsinhE</i>	88.40	95.12
<i>SOP+SC+AsinhE</i>	90.70	95.74
<i>SOP+SC+SigmE</i>	91.71	96.78
<i>SOP+SC+Spec. Gamma</i>	-	96.88
<i>SOP+SC+Spec. HDP</i>	-	97.05
<i>SOP+SC+Spec. MaxExp</i>	-	97.28
<i>SOP+SC+Spec. MaxExp(F)</i>	-	97.62

TABLE 3: The Flower102 dataset. The bottom part shows our results for AlexNet and ResNet-50. The horizontal lines separate first- and second-order element-wise pooling, and the spectral pooling. The top part of the table lists state-of-the-art results from the literature.

Method	acc.	Method	acc.
<i>Baseline</i>	81.9	<i>SOP+SC+SigmE</i>	87.5
<i>SOP</i>	83.0	<i>SOP+SC+Spec. MaxExp</i>	87.8
<i>Kern. Pool.</i> [69]	85.5	<i>SOP+SC+Spec. MaxExp(F)</i>	88.4

TABLE 4: The Food-101 dataset. Our (right) vs. other methods (left).

to obtain high-rank autocorrelation matrices without a need to modify the similarity network to larger feature inputs. Moreover, we investigate the fine-grained Flower102 and Food-101 datasets in the few-shot learning scenario. We take the first 80 categories of each dataset for training/validation and the remaining 21 and 22 for testing, respectively. Lastly, we introduce few-shot learning protocols on a fine-grained Open MIC dataset [64] detailed next.

Open MIC. The Open Museum Identification Challenge [64] contains photos of various exhibits *e.g.*, paintings, timepieces, sculptures, glassware, relics, science exhibits, natural history pieces, ceramics, pottery, tools and indigenous crafts, captured from 10 museum exhibition spaces according to which this dataset is divided into 10 sub-problems. In total, it has 1–20 images per class and 866 diverse classes, many of which are fine-grained *e.g.*, fossils, jewelery, cultural relics, as shown in Figure 8. The within-class images undergo various geometric and photometric distortions as the data was captured by wearable cameras. Thus, Open MIC challenges one-shot learning algorithms. We combine (*shn+hon+clv*), (*clk+gls+scl*), (*sci+nat*) and (*shx+rlc*) into sub-problems $p1, \dots, p4$. We form 12 possible pairs in which sub-problem x is used for training and y for testing ($x \rightarrow y$). Our first protocol aims at the generalization from one task to another task, thus we use the target part of Open MIC with the 12 above sub-problems. Our second protocol aims at the generalization from one domain to another domain, thus we use the source and target parts of Open MIC for training/testing on 10 original sub-problems.

Graph datasets. We use seven popular graph benchmarks MUTAG, PTC, PROTEINS, NCI1, COLLAB, REDDIT-BINARY and REDDIT-MULTI-5K [70]. MUTAG contains mutable molecules, 188 chemical compounds, and 7 node labels. PTC includes a num-

ber of carcinogenicity tasks for toxicology prediction, it contains 417 compounds from four species, and 18 node labels. PROTEINS are sets of proteins from the BRENDA database [71] with 3 node labels. NCI is a collection of datasets for anticancer activity prediction with 37 node labels. Finally, COLLAB, REDDIT-BINARY and REDDIT-MULTI-5K represent social networks.

Experimental setup (classification setting). For Flower102 [59], we extract 12 cropped 224×224 patches per image and use mini-batch of size 5 to fine-tune the ResNet-50 model [63] pre-trained on ImageNet 2012 [62]. We obtain 2048 dim. $12 \times 7 \times 7$ conv. feature vectors from the last conv. layer for our second-order pooling layer. For MIT67 [60], we resize original images to 336×336 and use mini-batch of size 32, then fine-tune it on the ResNet-50 model [63] pre-trained on the Places-205 dataset [72]. With 336×336 image size, we obtain 2048 dim. 11×11 conv. feature vectors from the last conv. layer for our second-order pooling layer. For FMD [73] and Food-101 [61], we resize images to 448×448 , use mini-batch of size 32 and fine-tune ResNet-50 [63] pre-trained on ImageNet 2012 [62]. We use the 2048 dim. 14×14 conv. feature vectors from the last conv. layer. For ImageNet 2012 [62], we crop 224×224 patches and allow left-right flip. We obtain 2048 dim. $12 \times 7 \times 7$ conv. feature vectors. For ResNet-50, we fine-tune all layers for ~ 20 epochs with learning rates $1e-6$ – $1e-4$. We use RMSprop [74] with the moving average 0.99. Where stated, we use AlexNet [75] with fine-tuned last two conv. layers. We use 256 dim. 6×6 conv. feature vectors from the last conv. layer.

Experimental setup (few-shot setting). For *miniImageNet*, we use standard 5-way 1-shot and 5-way 5-shot protocols. For every training/testing episode, we randomly select 5/3 query samples per class. We average over 600 episodes to obtain results. We use the initial learning rate $1e-3$ and train the model with 200K episodes. For Flower102 and Food-101, we follow the same setting and train models with 40K and 200K episodes. For Open MIC, we mean-center images per sub-problem. As some classes have less than 5 images, we use the 5- to 90-way 1-shot learning protocol. During training, to form an episode, we select 1 image for the support set and another 2 images for the query set per class. During testing, we use the same number of support/query samples in every episode and compute the accuracy over 1000 episodes. We use the initial learning rate $1e-4$ and train over 15K episodes. For all datasets, we resize images to 84×84 or 224×224 where stated.

Experimental setup (graph classification). We use the Graph Isomorphism Network (*GIN0*) from package [70]. We remove the classifier to produce covariances. We tune the neighborhood size, hidden units and the number of layers between 10–50, 16–128 and 2–5. We use the Adam optimizer with learning rate $1e-2$.

Our methods. We evaluate the generalizations of MaxExp and Gamma, that is Sigmoid (*SigmE*) and Arcsin hyperbolic (*AsinhE*) pooling functions. We focus mainly on our second-order representation (*SOP*) but we also occasionally report results for the first-order approach (*FOP*). For the baseline, we use the classifier on top of the *fc* layer (*Baseline*). The hyperparameters are selected via cross-validation. The use of spatial coordinates and spectral operators is indicated by (*SC*) and (*Spec.*), resp. For few-shot setting, we evaluate second-order similarity network variants (*SoSN*(\otimes)) and (*SoSN*(\otimes +L)) defined in Eq. (25) and (26).

7.1 Evaluations

Below we investigate element-wise and spectral pooling in fine-grained classification, few-shot learning and graph classification.

Method	top-1 accuracy
<i>CNNs with Deep Supervision</i> [76]	76.1
<i>Places-205</i> [77]	80.9
<i>Deep Filter Banks</i> [78]	81.0
<i>Spectral Features</i> [79]	84.3
<i>Baseline</i>	84.0
<i>SOP+AsinhE</i>	85.3
<i>SOP+SigME</i>	85.6
<i>SOP+SC+AsinhE</i>	85.9
<i>SOP+SC+SigME</i>	86.3
<i>SOP+SC+Spec. Gamma</i>	86.4
<i>SOP+SC+Spec. HDP</i>	86.3
<i>SOP+SC+Spec. MaxExp</i>	86.5
<i>SOP+SC+Spec. MaxExp(F)</i>	86.8

TABLE 5: The MIT67 dataset. The bottom part shows our results for ResNet-50 pre-trained on the Places-205 dataset. The horizontal lines separate first- and second-order element-wise pooling, and the spectral pooling. The top part are state-of-the-art results from the literature.

Fine-grained datasets. The Flower102 dataset is evaluated in Table 3 which shows that AlexNet performs worse than ResNet-50, which is consistent with the literature. For the standard ResNet-50 fine-tuned on Flower102, we obtain 94.06% accuracy. The first-order Average and AsinhE pooling (*FOP*) and (*FOP+AsinhE*) score 94.08 and 94.6% accuracy. The second-order pooling (*SOP+AsinhE*) outperforms (*FOP+AsinhE*). For element-wise operators, we obtain the best result of **96.78%** for the second-order representation combined with spatial coordinates and SigME (*SOP+SC+SigME*), which is $\sim 2.7\%$ higher than our baseline. In contrast, a recent more complex method [68] obtained 95.3% accuracy. Our scores highlight that capturing co-occurrences of visual features and passing them via a well-defined Power Normalization function such as SigME works well for our fine-grained problem. We attribute the good performance of SigME to its ability to act as a detector of co-occurrences. The role of the Hyperbolic Tangent non-linearity (popular in deep learning) may be explained by its similarity to SigME.

Furthermore, we note that the spectral approaches (*Spec.*) outperform element-wise second-order pooling. However, the differences are not drastically large. Spectral Gamma and MaxExp, both with spatial coordinates, denoted as (*SOP+SC+Spec. Gamma*) and (*SOP+SC+Spec. MaxExp*), perform similarly to the time-reversed Heat Diffusion Process (HDP) which experimentally validates their similarity to HDP. Finally, our fast spectral MaxExp (*SOP+SC+Spec. MaxExp(F)*) yields **97.62%** accuracy. We expect that the back-prop. through the fast spectral MaxExp is more stable




Fig. 9: Performance w.r.t. hyperparameters. Figures 9a and 9b: β -centering on Flower102 and α for spatial coordinate encoding on FMD. Figures 9c and 9d: the accuracy w.r.t. the η' and η parameters given SigME and the spectral MaxExp.

Method	acc.	Method	acc.
<i>IFV+DeCAF</i> [80]	65.5	<i>Baseline</i>	83.4
<i>FV+FC+CNN</i> [78]	82.2	<i>SOP+SC+AsinhE</i>	85.0
<i>SMO Task</i> [81]	82.3	<i>SOP+SC+SigME</i>	85.5
		<i>SOP+SC+Spec. MaxExp(F)</i>	86.4

TABLE 6: The FMD dataset. Our (right) vs. other methods (left).

Method	top-1 err	top-5 err
<i>ResNet-50</i> [63]	24.7	7.8
<i>MPN-COV</i> [34]	22.73	6.54
<i>Newton-Schulz</i> [34, 35]	22.14	6.22
<i>Baseline</i>	25.0	8.1
<i>SOP+SC+Spec. Gamma</i>	22.51	6.85
<i>SOP+SC+Spec. MaxExp(F)</i>	22.05	6.04

TABLE 7: The ImageNet 2012 dataset. The top part of the table lists state-of-the-art results from the literature.

than the back-prop. via SVD which we use for other methods.

For Food-101, we apply our second-order representations (*SOP+SC+SigME*) and (*SOP+SC+Spec. MaxExp*) and obtain **87.5%** and **87.8%** accuracy. Furthermore, the fast spectral MaxExp (*SOP+SC+Spec. MaxExp(F)*) yields **88.4%** accuracy. In contrast, a recent more involved kernel pooling [69] reports 85.5% accuracy while the baseline approach scores only 81.9% in the same testbed. This shows the strength of our approach on fine-grained problems. **Scene recognition.** Next, we validate our approach on MIT67—a larger dataset for scene recognition. Table 5 shows that all second-order approaches (*SOP*) outperform the standard ResNet-50 network (*Baseline*) pre-trained on the Places-205 dataset and fine-tuned on MIT67. Moreover, (*SigME*) yields marginally better results than (*AsinhE*). Using spatial coordinates (*SC*) also results in additional gain in the classification performance. The second-order representation combined with spatial coordinates and SigME pooling (*SOP+SC+SigME*) yields **86.3%** accuracy and outperforms our baseline and [79] by 2.3% and 2%, respectively.

For spectral operators, we observe a similar trend to results on Flower-102. Spectral Gamma with spatial coordinates (*SOP+SC+Spec. Gamma*) marginally outperforms HDP (*SOP+SC+Spec. HDP*). We expect this is due to the inversion of eigenvalues in HDP which makes it unstable for rank-deficient autocorrelation matrices. We note that the fast spectral MaxExp (*SOP+SC+Spec. MaxExp*) outperforms other spectral operators due to a more stable back-propagation it enjoys.

Material classification. For the FMD dataset (material/texture recognition), Table 6 demonstrates that our second-order representation (*SOP+SC+SigME*) scores **85.5%** accuracy and outperforms our baseline approach by 2.1%. Moreover, using the fast spectral MaxExp (*SOP+SC+Spec. MaxExp*) yields **86.4%** accuracy. We note that our approach and the baseline use the same testbed, that is, the only difference is the addition of our second-order representations, spatial coordinates and Power Normalization.

ImageNet 2012. Our fast spectral MaxExp (*SOP+SC+Spec. MaxExp(F)*) is shown to outperform (*SOP+SC+Spec. Gamma*) in Table 7. This is expected due to instabilities in back-propagation through SVD. Our method is also comparable with the recent approaches while enjoying strong theoretical connections to HDP.

Performance w.r.t. hyperparameters. Figure 9a demonstrates that β -centering has a positive impact on image classification with ResNet-50. This strategy, detailed in Section 4.1, is trivial to combine with our pooling. Figure 9b shows that setting non-zero α , which lets encode spatial coordinates according to Eq.

Model	Image Res.	Fine Tune	1-shot	5-shot
<i>Meta-Learn LSTM</i> [82]	84×84	N	43.44±0.7	60.60±0.7
<i>Prototypical Net</i> [43]		N	49.42±0.7	68.20±0.6
<i>MAML</i> [44]		Y	48.70±1.8	63.11±0.9
<i>Relation Net</i> [45]		N	50.44±0.8	65.32±0.7
<i>SoSN</i> (⊗) (no Power Norm.)	84×84	N	50.88±0.8	66.71±0.6
<i>SoSN</i> (⊗)+ <i>AsinhE</i>			52.10±0.8	67.79±0.6
<i>SoSN</i> (⊗)+ <i>SigmE</i>			52.96±0.8	68.63±0.6
<i>SoSN</i> (⊗)+ <i>L</i> + <i>AsinhE</i>			54.36±0.7	70.80±0.6
<i>SoSN</i> (⊗)+ <i>L</i> + <i>SC</i> + <i>AsinhE</i>			55.01±0.6	71.23±0.6
<i>SoSN</i> (⊗)+ <i>L</i> + <i>SigmE</i>			54.19±0.6	70.94±0.6
<i>SoSN</i> (⊗)+ <i>L</i> + <i>SC</i> + <i>SigmE</i>			55.36±0.7	71.23±0.6
<i>SoSN</i> (⊗)+ <i>SigmE</i>			60.35±0.7	74.01±0.6
<i>SoSN</i> (⊗)+ <i>L</i> + <i>AsinhE</i>			60.32±0.6	75.10±0.5
<i>SoSN</i> (⊗)+ <i>L</i> + <i>SC</i> + <i>AsinhE</i>			60.42±0.6	75.89±0.6
<i>SoSN</i> (⊗)+ <i>L</i> + <i>SigmE</i>	224×224	N	60.35±0.7	75.02±0.5
<i>SoSN</i> (⊗)+ <i>L</i> + <i>SC</i> + <i>SigmE</i>			60.49±0.6	75.54±0.5
<i>SoSN</i> (⊗)+ <i>L</i> + <i>SC</i> + <i>Spec. Gamma</i>			60.38±0.5	75.70±0.5
<i>SoSN</i> (⊗)+ <i>L</i> + <i>SC</i> + <i>Spec. MaxExp(F)</i>			60.40±0.6	75.75±0.5
<i>SoSN</i> (⊗)+ <i>L</i> + <i>Pretr.</i> + <i>SC</i> + <i>SigmE</i>			60.51±0.5	75.56±0.6
<i>SoSN</i> (⊗)+ <i>L</i> + <i>Pretr.</i> + <i>SC</i> + <i>Spec. Gamma</i>			60.95±0.6	76.20±0.5
<i>SoSN</i> (⊗)+ <i>L</i> + <i>Pretr.</i> + <i>SC</i> + <i>Spec. MaxExp(F)</i>			61.32±0.6	76.45±0.5

TABLE 8: Evaluations on the *miniImageNet* dataset (5-way acc. given). Refer to [45] for references to baselines.

(3), brings additional gain in accuracy at no extra cost. Figure 9c demonstrates that over 1% accuracy can be gained by tuning our SigmE pooling. Moreover, Figure 9d shows that the spectral MaxExp can yield further gains over element-wise SigmE and MaxExp for carefully chosen η . Lastly, we observed that our spectral and element-wise MaxExp converged in 3–12 and 15–25 iterations, respectively. This shows that both spectral and element-wise pooling have their strong and weak points.

Timing and variance in SPN. Below we present timing experiments performed on a TitanX GPU with the use of autograd profiler of PyTorch. To evaluate the forward runtime t_{fw+SPN} of Fast Spectral MaxExp, Newton-Schulz iter. (the approximate matrix square root) and the Generalized Spectral Power Normalization, we applied the *record_function()* subroutine of the profiler. To time the autograd-based back-propagation runtime through each of these operators, we firstly recorded the total GPU time $t_{tot+SPN}$ per method before removing these operators from the code and recording the total GPU time t_{tot} . Thus, we obtain the backward runtime $t_{bw+SPN} = t_{tot+SPN} - t_{tot} - t_{fw+SPN}$. We normalize results by the number of batches and datapoints per mini-batch.

Figure 11a shows the speed of forward/backward passes of Fast Spectral MaxExp from Alg. 2 w.r.t. η . The plot shows that the runtime grows sublinearly w.r.t. η which is the major advantage over the Generalized Spectral Power Normalization (GSPN) that uses SVD (runtime scales with d^ω where $2 < \omega < 2.376$) and the Newton-Schulz iterations which realize only the approximate matrix square root ($\gamma = 0.5$) whose quality depends on the number of iterations k (runtime scales linearly w.r.t. k).

Figures 11b and 11c compare the forward and backward speeds which show that our Fast Spectral MaxExp is faster than the Newton-Schulz iter. and GSPN. Notably, the backward pass appears $\sim 2\times$ more costly than the forward pass in all cases. We suspect this is due to the autograd recomputing intermediate variables for the forward pass in the backward pass. Thus, the runtime of optimized backward pass can be halved. Another downside of the Newton-Schulz iter. and GSPN compared to the Fast Spectral MaxExp was their larger memory footprint.

Finally, we investigate if SPN reduces the correlation between features that M represents. It is known that an isotropic Gaussian can be thought of as being constructed from uncorrelated features. Thus, we measure the variance of $j = 5$ leading eigenvalues of M passed via MaxExp. Figure 11d shows that as η increases,

Model	1-shot	5-shot
<i>Relation Net</i>	68.52 ± 0.94%	81.11 ± 0.66%
<i>SoSN</i> + <i>SigmE</i>	77.62 ± 0.88%	88.60 ± 0.53%
<i>SoSN</i> + <i>SC</i> + <i>SigmE</i>	78.50 ± 0.75%	89.95 ± 0.62%
<i>SoSN</i> + <i>SC</i> + <i>Spec. Gamma</i>	78.45 ± 0.72%	89.75 ± 0.55%
<i>SoSN</i> + <i>SC</i> + <i>Spec. MaxExp(F)</i>	78.51 ± 0.68%	89.70 ± 0.58%
<i>SoSN</i> + <i>Pretr.</i> + <i>SC</i> + <i>SigmE</i>	78.55 ± 0.70%	89.99 ± 0.80%
<i>SoSN</i> + <i>Pretr.</i> + <i>SC</i> + <i>Spec. Gamma</i>	79.35 ± 0.72%	90.65 ± 0.84%
<i>SoSN</i> + <i>Pretr.</i> + <i>SC</i> + <i>Spec. MaxExp(F)</i>	80.01 ± 0.74%	91.50 ± 0.65%

TABLE 9: Evaluations on the Flower102 dataset (5-way acc. given). For SoSN, we evaluate only our ((⊗+L)+SigmE) aggregator (equiv. to ((⊗)+SigmE) for 1-shot problems) given images of res. 224×224.

Model	1-shot	5-shot
<i>Relation Net</i>	36.89 ± 0.72%	49.07 ± 0.65%
<i>SoSN</i> + <i>SigmE</i>	42.44 ± 0.75%	60.70 ± 0.65%
<i>SoSN</i> + <i>SC</i> + <i>SigmE</i>	42.80 ± 0.72%	60.95 ± 0.68%
<i>SoSN</i> + <i>SC</i> + <i>Spec. Gamma</i>	42.30 ± 0.72%	60.71 ± 0.60%
<i>SoSN</i> + <i>SC</i> + <i>Spec. MaxExp(F)</i>	42.40 ± 0.75%	60.70 ± 0.59%
<i>SoSN</i> + <i>Pretr.</i> + <i>SC</i> + <i>SigmE</i>	42.60 ± 0.69%	60.85 ± 0.70%
<i>SoSN</i> + <i>Pretr.</i> + <i>SC</i> + <i>Spec. Gamma</i>	43.72 ± 0.52%	62.02 ± 0.68%
<i>SoSN</i> + <i>Pretr.</i> + <i>SC</i> + <i>Spec. MaxExp(F)</i>	45.21 ± 0.62%	64.50 ± 0.61%

TABLE 10: Evaluations on the Food-101 dataset (5-way acc. given). For SoSN, we evaluate only our ((⊗+L)+SigmE) aggregator (equiv. to ((⊗)+SigmE) for 1-shot problems) given images of res. 224×224.

the variance of the leading eigenvalues decreases. However, while leading eigenvalues for smaller η become equalized and pulled towards the value of one, non-leading eigenvalues may remain unaffected which increases the variance. This behavior is desired as only a certain j leading eigenvalues correspond to the signal and the remaining non-leading eigenvalues represent the noise (by analogy to the Principal Component Analysis).

Few-shot learning. Below we evaluate our SoSN model and compare it against state-of-the-art models *e.g.*, Relation Net [45].

For *miniImageNet*, Table 8 shows that our method outperforms others on 5-way 1- and 5-shot learning. For experiments with image size of 84×84, our SoSN model (*SoSN*(⊗)+*SigmE*) achieved $\sim 2.5\%$ and $\sim 3.3\%$ higher accuracy than Relation Net [45]. Our SoSN models also outperformed Prototypical Net by 0.43–0.75% accuracy on the 5-way 5-shot protocol. For 224×224 images on (*SoSN*(⊗)+*L*)+*SigmE*, the accuracies on both protocols increase by $\sim 5.5\%$ and $\sim 6\%$ over 84×84 counterpart, which shows that SoSN benefits from larger image sizes as second-order matrices used in Eq. (26) become full-rank (the similarity network needs no modifications). Finally, Table 8 shows gains on spectral methods (*Spec.*) pre-trained on Food-101 (*Pretr.*) *i.e.*, (*Pretr.*+*SC*+*Spec. Gamma*) and (*Pretr.*+*SC*+*Spec. MaxExp(F)*) outperforms non-spectral (*SC*+*SigmE*). However, spectral methods without pre-training (*SC*+*Spec. Gamma*) and (*SC*+*Spec. MaxExp(F)*) and even pre-trained non-spectral (*Pretr.*+*SC*+*SigmE*) fail to bring any benefits over non-spectral (*SC*+*SigmE*).

We believe this is an important finding—we benefit from pre-training only when we use spectral Power Norms. As spectral Power Norms act as time-reversed HDP, the level of correlation between co-occurring features is reduced. This reduces so-called catastrophic forgetting on the dataset used for pre-training.

Our best results for (*Pretr.*+*SC*+*Spec. MaxExp(F)*) gave **61.32%** and **76.45%** accuracy on 5-way 1- and 5-shot learning.

Fine-grained few-shot learning. For Open MIC, Figure 10a introduces results for the protocol that tests the generalization from one task to another task (*Protocol I*). As only 1-shot protocol can

Method	MUTAG	PTC	PROTEINS	NCII
S^2GC [83]	85.1±7.4	-	75.5±4.1	-
DGCNN [84]	85.8±1.7	58.6±2.5	75.5±0.9	74.4±0.5
GCAPS-CNN [85]	-	66.0±5.9	76.4±4.2	82.7±2.4
BC+CAPS [86]	88.9±5.5	69.0±5.0	74.1±3.2	65.9±1.1
GINO [70]	86.1±5.8	56.5±6.8	72.2±4.9	77.9±2.5
SOP	86.2±5.2	57.5±10.1	71.2±4.9	78.3±3.0
SOP+AsinhE	86.2±6.2	58.2±5.8	72.0±3.8	79.5±2.0
SOP+SigME	87.8±6.1	58.4±5.5	71.8±3.6	79.6±1.9
SOP+Newton-Schulz [34, 35]	86.2±6.1	59.3±5.8	75.3±2.8	79.9±2.3
SOP+Spec. Gamma	86.5±6.1	61.5±3.8	75.7±4.0	80.0±2.1
SOP+Spec. HDP	86.2±7.9	61.2±6.4	75.5±2.8	79.6±2.0
SOP+Spec. MaxExp	86.8±6.6	61.9±2.4	76.8±2.9	79.8±2.4
SOP+Spec. MaxExp(F)	88.9±5.8	68.3±9.3	76.2±2.8	80.3±2.4
Method	MUTAG	COLLAB	REDDIT-B	REDDIT-5K
S^2GC [83]	85.1±7.4	80.2±1.2	-	-
DGCNN [84]	85.8±1.7	73.8±0.5	76.0±1.7	48.7±4.5
GCAPS-CNN [85]	-	77.7±2.5	87.6±2.5	50.1±1.7
AWE [87]	-	71.0±1.5	83.0±2.7	54.7±2.9
GINO [70]	89.2±4.8	79.9±1.7	92.1±1.9	55.5±2.1
SOP	89.9±9.5	80.6±1.2	91.8±2.1	55.6±3.0
SOP+AsinhE	92.0±4.9	81.2±1.6	92.5±1.7	56.8±2.1
SOP+SigME	91.5±5.6	81.4±1.8	92.4±1.9	57.1±1.8
SOP+Newton-Schulz [34, 35]	93.7±8.5	81.2±8.9	92.3±1.8	57.1±1.9
SOP+Spec. MaxExp(F)	94.7±5.0	81.7±1.7	92.6±1.6	57.1±1.9

TABLE 11: Classification with GINO and various pooling methods on (top) MUTAG, PTC, PROTEINS and NCII (the validation split was used only for validation) and (bottom) MUTAG, COLLAB, REDDIT-BINARY and REDDIT-MULTI-5K (after obtaining hyperparameters, the train and validation splits were combined for retraining). State-of-the-art results from the literature are listed directly below ‘Method’.

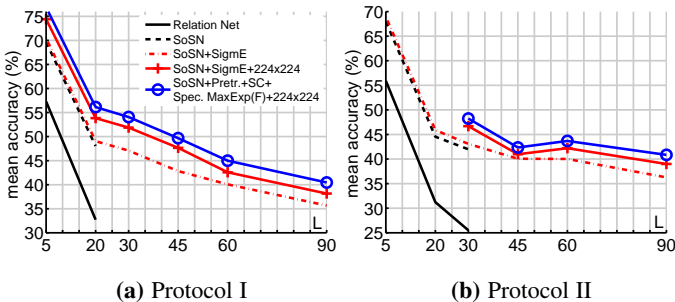


Fig. 10: Evaluations on the Open MIC dataset. Fig. 10a shows 1-shot mean accuracy on Protocol I. Each point is averaged over 12 possible testing results from train-test pairs $x \rightarrow y$ where we use p1: shn+hon+clv, p2: clk+gls+scl, p3: sci+nat, p4: shx+rlc. Fig. 10b shows 1-shot mean accuracy on Protocol II. Training is performed on source images and testing on target images for every exhibition. Then every point in the plot is average over results on 10 exhibitions.

be applied to this dataset, we use ($SoSN(\otimes)$) which is equivalent to ($SoSN(\otimes+L)$) for 1-shot problems and we denote it by ($SoSN$). All our methods outperform the Relation Net [45]. For 5- and 20-way, Relation Net scores 55.45 and 31.58%. In contrast, our ($SoSN$) scores 68.23 and 45.31%, resp. Our ($SoSN+SigME$) scores 69.06 and 46.53%, resp. Increasing resolution to 224×224 on ($SoSN+SigME$) yields **74.43%** and **53.84%** accuracy.

Moreover, our ($SoSN+Pretr.+SC+Spec. MaxExp(F)$) with 224×224 res. pre-trained on *miniImageNet* yielded a $\sim 2.0\%$ increase over non-spectral ($SoSN+SigME$). A similar trend is observed on 5- to 90-way protocols which test the stability of our idea across matching testing queries each with support images from 90 distinct classes. A realistic setting with a large ‘way’ number is frequently avoided in few-shot learning community.

Figure 10b introduces results for the protocol that tests the generalization from one domain to another domain (*Protocol II*)

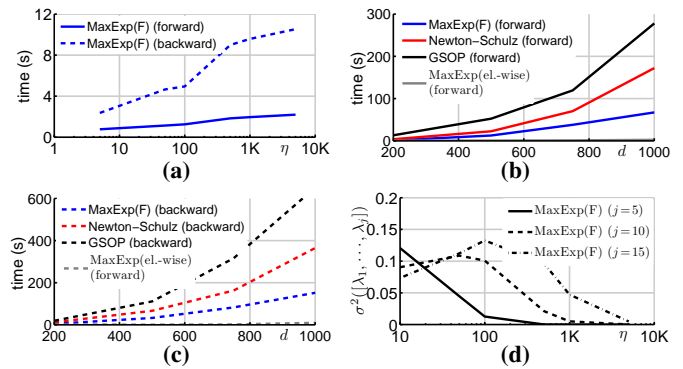


Fig. 11: Timing and variance of operators. Figure 11a evaluates the speed (per 1000 images) of forward and backward passes our Fast Spectral MaxExp (Alg. 2) as a function of η given the side size $d=200$ of M . Figures 11b and 11c compare the speed of forward and backward passes of the proposed Fast Spectral MaxExp ($MaxExp(F)$), the Newton-Schulz iterations algorithm, the Generalized Spectral Power Normalization ($GSPN$) that uses SVD and the Element-wise MaxExp ($MaxExp(el-wise)$). We vary the side size $200 \leq d \leq 1000$ of M . Figure 11d shows how the variance over $5 \leq j \leq 15$ leading eigenvalues passed via MaxExp(F) varies as a function of η . As η increases the variance of the leading eigenvalues decreases.

which is also often avoided in the few-shot learning community but helps ascertain how well the algorithm generalizes between different domains. Our ($SoSN+SigME$) outperforms Relation Net by up to $\sim 35\%$. As this protocol measures the generalization ability of few-shot learning methods under a large domain shift, the average scores are $\sim 20\%$ below scores from Figure 10a. However, our second-order relationship descriptor with SigME pooling is beneficial for similarity learning. Moreover, our ($SoSN+Pretr.+SC+Spec. MaxExp(F)$) with 224×224 res. pre-trained on *miniImageNet* obtained the best performance.

To conclude our experiments on fine-grained few-shot learning, Tables 9 and 10 introduce results on 5-way 1-shot and 5-way 5-shot evaluation protocols on Flower102 and Food-101. Both tables demonstrate that our ($SoSN+SigME$) model outperforms Relation Net by $\sim 5\%$ to $\sim 11\%$ accuracy.

Spectral operators on $SoSN$ pre-trained with *miniImageNet* performed well. On 1-/5-shot protocols, our ($SoSN+Pretr.+SC+Spec. MaxExp(F)$) scored 1.5%/1.55% and 2.4%/3.5% accuracy over ($SoSN+SC+SigME$) on Flower102 and Food-101, resp. In contrast, pre-trained non-spectral ($SoSN+Pretr.+SC+SigME$) and spectral ($SoSN+SC+Spec. Gamma$) and ($SoSN+SC+Spec. MaxExp(F)$) without pre-training fail to bring further benefits over ($SoSN+SC+SigME$) without pre-training. This supports our hypothesis about the connection of spectral operators to time-reversed HDP which reduces correlation between co-occurrences thus limiting catastrophic forgetting.

Graph classification. Table 11 shows results for spectral and element-wise PN operators on covariance matrices employed on top of the Graph Isomorphism Network ($GINO$) [70]. Element-wise PN ($SOP+SigME$) outperforms (SOP) and the first-order average pooling ($GINO$), one of the strongest baselines. Moreover, the Fast Spectral MaxExp ($SOP+MaxExp(F)$) typically outperforms the Newton-Schulz inter. (approx. matrix square root) and element-wise operators. As package [70] uses the validation split only for validation (in contrast to other packages), we retrain on train+validation splits (see the bottom of Table 11).

8 CONCLUSIONS

We have studied Power Normalizations in the context of element-wise co-occurrence representations and demonstrated their theoretical role which is to ‘detect’ co-occurrences. We have shown that different assumptions on distributions from which features are drawn result in similar non-linearities *e.g.*, MaxExp vs. SigmE. Thus, we have proposed surrogate functions SigmE and AsinhE with well-behaved derivatives for end-to-end training which can handle so-called negative evidence. Moreover, we have proposed a fast spectral MaxExp which can be computed as faster than the matrix square root via iterative Newton-Schulz while enjoying the adjustable parameter. Finally, we have shown that Spectral Power Normalizations fulfill a similar role to the time-reversed Heat Diffusion Process well-known from the graph theory, thus paving a strong theoretical foundation for further studies of SPNs.

Acknowledgements. We thank Dr. Ke Sun for brainstorming, Ondrej Hlinka/Garry Swan for help with HPC, Hao Zhu for checks of some SOP codes, and Lei Wang for quick checks of text.


REFERENCES

- [1] O. Tuzel, F. Porikli, and P. Meer, “Region covariance: A fast descriptor for detection and classification,” *ECCV*, 2006. 1, 2
- [2] F. Porikli and O. Tuzel, “Covariance tracker,” *CVPR*, 2006. 1, 2
- [3] K. Guo, P. Ishwar, and J. Konrad, “Action recognition from video using feature covariance matrices,” *TIP*, vol. 22, no. 6, pp. 2479–2494, 2013. 1
- [4] J. Carreira, R. Caseiro, J. Batista, and C. Sminchisescu, “Semantic Segmentation with Second-Order Pooling,” *ECCV*, 2012. 1, 2
- [5] P. Koniusz, F. Yan, P.-H. Gosselin, and K. Mikolajczyk, “Higher-order occurrence pooling for bags-of-words: Visual concept detection,” *TPAMI*, vol. 39, no. 2, pp. 313–326, 2017. 1, 2, 3, 4, 5, 7, 11
- [6] P. Koniusz, F. Yan, P. Gosselin, and K. Mikolajczyk, “Higher-order Occurrence Pooling on Mid- and Low-level Features: Visual Concept Detection,” INRIA, Tech. Rep. hal-00922524, 2013. [Online]. Available: <https://hal.inria.fr/hal-00922524/> 1, 2
- [7] P. Koniusz and A. Cherian, “Sparse coding for third-order super-symmetric tensor descriptors with application to texture recognition,” *CVPR*, 2016. 1
- [8] P. Koniusz, F. Yan, and K. Mikolajczyk, “Comparison of Mid-Level Feature Coding Approaches And Pooling Strategies in Visual Concept Detection,” *CVIU*, 2012. 1, 2, 3
- [9] Y. Boureau, F. Bach, Y. LeCun, and J. Ponce, “Learning Mid-Level Features for Recognition,” *CVPR*, 2010. 1, 2, 3
- [10] Y. Boureau, J. Ponce, and Y. LeCun, “A Theoretical Analysis of Feature Pooling in Vision Algorithms,” *ICML*, 2010. 1, 2, 3
- [11] L. Lingqiao, L. Wang, and X. Liu, “In Defence of Soft-assignment Coding,” *ICCV*, 2011. 1, 2
- [12] F. Perronnin, J. Sánchez, and T. Mensink, “Improving the Fisher Kernel for Large-Scale Image Classification,” *ECCV*, 2010. 1, 2
- [13] P. Koniusz and K. Mikolajczyk, “Spatial coordinate coding to reduce histogram representations, dominant angle and colour pyramid match,” *ICIP*, 2011. 1, 4
- [14] A. J. Smola and R. Kondor, “Kernels and regularization on graphs,” *Learning Theory and Kernel Machines*, pp. 144–158, 2003. 1, 2, 6
- [15] P. Koniusz, H. Zhang, and F. Porikli, “A deeper look at power normalizations,” *CVPR*, pp. 5774–5783, 2018. 2, 3, 8
- [16] H. Zhang and P. Koniusz, “Power normalizing second-order similarity network for few-shot learning,” *WACV*, 2019. 2, 3
- [17] P. Koniusz, M. Harandi, L. Wang, and R. Wang, “Second- and higher-order representations in computer vision,” ICCV Tutorial, <http://users.cecs.anu.edu.au/~koniusz/secondcv-iccv19>, accessed: 02-11-2019. 2
- [18] O. Tuzel, F. Porikli, and P. Meer, “Pedestrian detection via classification on riemannian manifolds,” *TPAMI*, vol. 30, no. 10, pp. 1713–1727, 2008. 2
- [19] I. L. Dryden, A. Koloydenko, and D. Zhou, “Non-euclidean statistics for covariance matrices, with applications to diffusion tensor imaging,” *The Annals of Applied Statistics*, vol. 3, no. 3, pp. 1102–1123, 2009. 2
- [20] X. Pennec, P. Fillard, and N. Ayache, “A Riemannian Framework for Tensor Computing,” *IJCV*, vol. 66, no. 1, pp. 41–66, 2006. 2
- [21] R. Bhatia, *Positive definite matrices*. Princeton University Press, 2007. 2
- [22] Z. Wang and B. C. Vemuri, “An affine invariant tensor dissimilarity measure and its applications to tensor-valued image segmentation,” *CVPR*, 2004. 2
- [23] A. Cherian, S. Sra, A. Banerjee, and N. Papanikolopoulos, “Jensen-Bregman LogDet Divergence with Application to Efficient Similarity Search for Covariance Matrices,” *TPAMI*, vol. 35, no. 9, pp. 2161–2174, 2013. 2
- [24] V. Arsigny, P. Fillard, X. Pennec, and N. Ayache, “Log-euclidean metrics for fast and simple calculus on diffusion tensors,” *Magnetic resonance in medicine*, vol. 56, no. 2, pp. 411–421, 2006. 2
- [25] M. Harandi, R. Hartley, C. Shen, B. Lovell, and C. Sanderson, “Extrinsic methods for coding and dictionary learning on grassmann manifolds,” *IJCV*, 2015. 2
- [26] S. Kumar Roy, Z. Mhammedi, and M. Harandi, “Geometry aware constrained optimization techniques for deep learning,” *CVPR*, 2018. 2
- [27] M. Harandi, M. Salzmann, and R. Hartley, “Joint dimensionality reduction and metric learning: A geometric take,” *ICML*, p. 1404–1413, 2017. 2
- [28] T.-Y. Lin, A. R. Chowdhury, and S. Maji, “Bilinear cnn models for fine-grained visual recognition,” *ICCV*, 2017. 2
- [29] —, “Bilinear convolutional neural networks for fine-grained visual recognition,” *TPAMI*, 2017. 2
- [30] Y.-F. Shih, Y.-M. Yeh, Y.-Y. Lin, M.-F. Weng, Y.-C. Lu, and Y.-Y. Chuang, “Deep co-occurrence feature learning for visual object recognition,” *CVPR*, 2017. 2
- [31] G. Hu, Y. Hua, Y. Yuan, Z. Zhang, Z. Lu, S. S. Mukherjee, T. M. Hospedales, N. M. Robertson, and Y. Yang, “Attribute-enhanced face recognition with neural tensor fusion networks,” *ICCV*, 2017. 2
- [32] C. Ionescu, O. Vantzos, and C. Sminchisescu, “Matrix backpropagation for deep networks with structured layers,” *ICCV*, 2015. 2, 7
- [33] Z. Huang and L. V. Gool, “A riemannian network for spd matrix learning,” *AAAI*, pp. 2036–2042, 2017. 2
- [34] P. Li, J. Xie, Q. Wang, and Z. Gao, “Towards faster training of global covariance pooling networks by iterativematrix square root normalization,” *CVPR*, 2018. 2, 3, 6, 10, 12, 14
- [35] T.-Y. Lin and S. Maji, “Improved Bilinear Pooling with CNNs,” *BMVC*, 2017. 2, 3, 6, 7, 12, 14
- [36] M. Simon, E. Rodner, T. Darrell, and J. Denzler, “The whole is more than its parts? from explicit to implicit pose normalization,” *TPAMI*, 2018. 2
- [37] H. Jégou, M. Douze, and C. Schmid, “On the Burstiness of Visual Elements,” *CVPR*, pp. 1169–1176, 2009. 2
- [38] S. Boughorbel, J.-P. Tarel, and N. Boujemaa, “Generalized Histogram Intersection Kernel for Image Recognition,” *ICIP*, 2005. 2
- [39] E. G. Miller, N. E. Matsakis, and P. A. Viola, “Learning from one example through shared densities on transforms,” *CVPR*, vol. 1, pp. 464–471, 2000. 3
- [40] F. F. Li, R. VanRullen, C. Koch, and P. Perona, “Rapid natural scene categorization in the near absence of attention,” *Proceedings of the National Academy of Sciences*, vol. 99, no. 14, pp. 9596–9601, 2002. 3
- [41] M. Fink, “Object classification from a single example utilizing class relevance metrics,” *NIPS*, pp. 449–456, 2005. 3

- [42] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra *et al.*, “Matching networks for one shot learning,” *NIPS*, 2016. 3, 10
- [43] J. Snell, K. Swersky, and R. Zemel, “Prototypical networks for few-shot learning,” *NIPS*, pp. 4077–4087, 2017. 3, 13
- [44] C. Finn, P. Abbeel, and S. Levine, “Model-agnostic meta-learning for fast adaptation of deep networks,” *ICML*, pp. 1126–1135, 2017. 3, 13
- [45] F. Sung, Y. Yang, L. Zhang, T. Xiang, P. H. Torr, and T. M. Hospedales, “Learning to compare: Relation network for few-shot learning,” *CVPR*, 2018. 3, 9, 10, 13, 14
- [46] D. Wertheimer and B. Hariharan, “Few-shot learning with localization in realistic settings,” *CVPR*, pp. 6558–6567, 2019. 3
- [47] H. Zhang, J. Zhang, and P. Koniusz, “Few-shot learning via saliency-guided hallucination of samples,” *CVPR*, 2019. 3
- [48] S. Zhang, D. Luo, L. Wang, and P. Koniusz, “Few-shot object detection by second-order pooling,” *ACCV*, 2020. 3
- [49] C. Simon, P. Koniusz, R. Nock, and M. Harandi, “Adaptive subspaces for few-shot learning,” *CVPR*, 2020. 3
- [50] P. Koniusz, Y. Tas, and F. Porikli, “Domain adaptation by mixture of alignments of second-or higher-order scatter tensors,” *CVPR*, vol. 2, 2017. 3
- [51] H. Jegou and O. Chum, “Negative evidences and co-occurrences in image retrieval: the benefit of pca and whitening,” *ECCV*, 2012. 4
- [52] L. C. Rogers, “Derivatives of eigenvalues and eigenvectors,” *AIAA Journal*, vol. 8, no. 5, pp. 943–944, 1970. 6
- [53] C. S. Rudisill, “Derivatives of eigenvalues and eigenvectors for a general matrix,” *AIAA Journal*, vol. 12, no. 5, pp. 721–722, 1974. 6
- [54] J. R. Magnus, “On differentiating eigenvalues and eigenvectors,” *Econometric Theory*, 1985. 6
- [55] T.-Y. Lin, S. Maji, and P. Koniusz, “Second-order democratic aggregation,” *ECCV*, 2018. 7
- [56] “Exponentiation by squaring,” Wikipedia, https://en.wikipedia.org/wiki/Exponentiation_by_squaring, accessed: 27-11-2019. 7
- [57] C. Zhang, D. Florencio, and P. A. Chou, “Graph signal processing - a probabilistic framework,” Microsoft, Tech. Rep. MSR-TR-2015-31, 2015. 7, 17
- [58] L. E. Zhukov, “Diffusion and random walks on graphs,” www.leonidzhukov.net/hse/2015/networks/lectures/lecture11.pdf, 2015. 8, 18
- [59] M.-E. Nilsback and A. Zisserman, “Automated Flower Classification over a Large Number of Classes,” *ICVGIP*, 2008. 10, 11
- [60] A. Quattoni and A. Torralba, “Recognizing indoor scenes,” *CVPR*, 2009. 10, 11
- [61] L. Bossard, M. Guillaumin, and L. J. V. Gool, “Food-101 - mining discriminative components with random forests,” *ECCV*, pp. 446–461, 2014. 10, 11
- [62] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, “ImageNet large scale visual recognition challenge,” *IJCV*, vol. 115, no. 3, pp. 211–252, 2015. 10, 11
- [63] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” *CVPR*, 2016. 10, 11, 12
- [64] P. Koniusz, Y. Tas, H. Zhang, M. Harandi, F. Porikli, and R. Zhang, “Museum exhibit identification challenge for the supervised domain adaptation and beyond,” *ECCV*, 2018. 11
- [65] H. Azizpour, A. S. Razavian, J. Sullivan, A. Maki, and S. Carlsson, “Factors of transferability for a generic convnet representation,” *CoRR*, vol. abs/1406.5774, 2015. 11
- [66] L. Xie, J. Wang, W. Lin, B. Zhang, and Q. Tian, “Towards reversal-invariant image representation,” *IJCV*, vol. 123, no. 2, pp. 226–250, 2017. 11
- [67] J. Liu, C. Gao, D. Meng, and W. Zuo, “Two-stream contextualized cnn for fine-grained image classification,” *AAAI*, 2016. 11
- [68] M. Simon and E. Rodner, “Neural activation constellations: Un-supervised part model discovery with convolutional networks,” *ICCV*, pp. 1143–1151, 2015. 11, 12
- [69] Y. Cui, F. Zhou, J. Wang, X. Liu, Y. Lin, and S. Belongie, “Kernel pooling for convolutional neural networks,” *CVPR*, 2017. 11, 12
- [70] M. Fey and J. E. Lenssen, “Fast graph representation learning with pytorch geometric,” *ICLR*, 2019. 11, 14
- [71] I. Schomburg, A. Chang, C. Ebeling, M. Gremse, C. Heldt, G. Huhn, and D. Schomburg, “Brenda, the enzyme database: updates and major new developments,” *Nucleic Acids Research*, vol. 32, pp. D431–D433, 2004. 11
- [72] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva, “Learning deep features for scene recognition using places database,” *NIPS*, 2014. 11
- [73] L. Sharan, R. Rosenholtz, and E. H. Adelson, “Material perception: What can you see in a brief glance?” *Journal of Vision*, vol. 14, no. 9, 2014. 11
- [74] G. Hinton, “Neural Networks for Machine Learning–Lecture 6e,” www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf, accessed: 10-11-2017. 11
- [75] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “ImageNet classification with deep convolutional neural networks,” *NIPS*, pp. 1106–1114, 2012. 11
- [76] L. Wang, C.-Y. Lee, Z. Tu, and S. Lazebnik, “Training deeper convolutional networks with deep supervision,” *CoRR*, vol. abs/1505.02496, 2015. 12
- [77] L. Wang, S. Guo, W. Huang, and Y. Qiao, “Places205-vggnet models for scene recognition,” *CoRR*, vol. abs/1508.01667, 2015. 12
- [78] M. Cimpoi, S. Maji, and A. Vedaldi, “Deep filter banks for texture recognition and segmentation,” *CVPR*, 2015. 12
- [79] S. H. Khan, M. Hayat, and F. Porikli, “Scene categorization with spectral features,” *ICCV*, pp. 5638–5648, 2017. 12
- [80] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, and A. Vedaldi, “Describing textures in the wild,” *CVPR*, 2014. 12
- [81] Y. Zhang, M. Ozay, X. Liu, and T. Okatani, “Integrating deep features for material recognition,” *ICPR*, 2016. 12
- [82] S. Ravi and H. Larochelle, “Optimization as a model for fewshot learning,” *ICLR*, 2017. 13
- [83] H. Zhu and P. Koniusz, “Simple spectral graph convolution,” *ICLR*, 2021. 14
- [84] M. Zhang, Z. Cui, M. Neumann, and Y. Chen, “An end-to-end deep learning architecture for graph classification,” *AAAI*, 2018. 14
- [85] S. Verma and Z.-L. Zhang, “Graph capsule convolutional neural networks,” *CoRR*, vol. abs/1805.08090, 2018. 14
- [86] M. D. G. Mallea, P. Meltzer, and P. J. Bentley, “Capsule neural networks for graph classification using explicit tensorial graph representations,” *CoRR*, vol. abs/1902.08399, 2019. 14
- [87] S. Ivanov and E. Burnaev, “Anonymous walk embeddings,” *ICML*, pp. 2186–2195, 2018. 14



Piotr Koniusz. A Senior Researcher in Machine Learning Research Group at Data61/CSIRO (NICTA), and a Senior Honorary Lecturer at the Australian National University (ANU). He was a postdoctoral researcher in the team LEAR, INRIA, France. He received his BSc in Telecommunications and Software Engineering in 2004 from the Warsaw University of Technology, Poland, and completed his PhD in Computer Vision in 2013 at CVSSP, University of Surrey, UK.



Hongguang Zhang. From 2016, Hongguang Zhang is a PhD student in computer vision and machine learning at the Australian National University and Data61/CSIRO, Canberra, Australia. He received the BSc degree in electrical engineering and automation from Shanghai Jiao Tong University, Shanghai, China in 2014. He received his MSc degree in electronics science and technology from National University of Defense Technology, Changsha, China in 2016. His interests include fine-grained image classification, zero-shot learning, few-shot learning and deep learning methods.

SUPPLEMENTARY MATERIAL

Below we provide proofs and derivations.

A. Proof of Eq. (28)

Let the following difference of probabilities:

$$\psi = (1-q)^N - (1-p)^N. \quad (28)$$

Proof. One can derive Eq. (28) by directly applying the Multinomial calculus as follows:

$$\sum_{n=1}^N \sum_{n'=0}^{N-n} \binom{N}{n, n', N-n-n'} (p^n q^{n'} - p^{n'} q^n) (1-p-q)^{N-n-n'}. \quad (29)$$

One can verify algebraically/numerically that Eq. (29) and (28) are equivalent. \square

B. Deriv. of Average, Gamma and MaxExp

Let $\Phi = [\phi_1, \dots, \phi_N] \in \mathbb{R}^{d \times N}$, $\mathbf{C} = [\mathbf{c}_1, \dots, \mathbf{c}_N] \in \mathbb{R}^{Z' \times N}$, $\bar{\phi}_n = [\phi_n; \mathbf{c}_n]$, $n \in \mathcal{I}_N$. Let some class. loss $\ell(\Psi, \mathbf{W})$, $\Psi \in \mathcal{S}_{++}^{d+Z'}$ (or \mathcal{S}_{++}) and \mathbf{W} are our descriptor and a hyperplane. Observe that:

$$\frac{\partial \sum_n \bar{\phi}_n \bar{\phi}_n^T}{\partial \phi_{kl}} = \begin{bmatrix} \mathbf{j}_k \phi_l^T + \phi_l \mathbf{j}_k^T & \mathbf{j}_k \mathbf{c}_l^T \\ \mathbf{c}_l \mathbf{j}_k^T & [0]_{Z' \times Z'} \end{bmatrix}, \quad (30)$$

where $[0]_{Z' \times Z'}$ denotes an array of size $Z' \times Z'$ filled with zeros.

Average pooling is set by $\Psi = \mathcal{G}_{\text{Avg}}(\mathbf{M}) = \mathbf{M}$ and $\mathbf{D} = \mathbb{1}\mathbb{1}^T$ so that $\Psi = \frac{1}{N} \sum_n \bar{\phi}_n \bar{\phi}_n^T$. Thus, the full derivative becomes:

$$\sum_{k,l} \frac{\partial \ell(\Psi, \mathbf{W})}{\partial \Psi_{kl}} \frac{\partial \Psi_{kl}}{\partial \Phi} = \frac{2}{N} \text{Sym} \left(\frac{\partial \ell(\Psi, \mathbf{W})}{\partial \Psi} \odot \mathbf{D} \right)_{(1:d,:)} \begin{bmatrix} \Phi \\ \mathbf{C} \end{bmatrix}. \quad (31)$$

$\mathbf{X}_{(1:d,:)}$ returns 1, ..., d rows/all col. as Matlab oper. $\mathbf{X}(1:d, :)$.

Gamma pooling is set by $\Psi = \mathcal{G}_{\text{Gamma}}(\mathbf{M}; \gamma) = (\mathbf{M} + \varepsilon)^\gamma$, where rising \mathbf{M} to the power of γ is element-wise and ε is a reg. constant. Thus, we obtain:

$$\frac{\partial \Psi}{\partial \phi_{kl}} = \frac{1}{N} \gamma (\mathbf{M} + \varepsilon)^{\gamma-1} \odot \frac{\partial \sum_n \bar{\phi}_n \bar{\phi}_n^T}{\partial \phi_{kl}}. \quad (32)$$

The derivative is given by Eq. (31) if $\mathbf{D} = \gamma (\mathbf{M} + \varepsilon)^{\gamma-1}$.

MaxExp pooling $\Psi = \mathcal{G}_{\text{MaxExp}}(\mathbf{M}; \eta) = 1 - (1 - \mathbf{M} / (\text{Tr}(\mathbf{M}) + \varepsilon))^\eta$ has the derivative given by Eq. (31) with the following \mathbf{D} :

$$\mathbf{D} = \eta \left(1 - \frac{\mathbf{M}}{\text{Tr}(\mathbf{M}) + \varepsilon} \right)^{\eta-1} \odot \mathbf{T} \text{ and } \mathbf{T} = \left(\frac{1}{\text{Tr}(\mathbf{M}) + \varepsilon} - \frac{\mathbf{M} \odot \mathbf{I}}{(\text{Tr}(\mathbf{M}) + \varepsilon)^2} \right), \quad (33)$$

where multiplication \odot , division, rising to the power *etc.* are all element-wise operations.

C. Derivatives of SigmE and AsinhE pooling

SigmE pooling is set by $\Psi = \mathcal{G}_{\text{SigmE}}(\mathbf{M}; \eta) = \frac{2}{1+e^{-\eta/\mathbf{M}}} - 1$ or trace-normalized $\frac{2}{1+e^{-\eta/\mathbf{M}}} - 1$. The first expression yields:

$$\frac{\partial \Psi}{\partial \phi_{kl}} = \frac{1}{N} \frac{2\eta' e^{-\eta/\mathbf{M}}}{(1+e^{-\eta/\mathbf{M}})^2} \odot (\mathbf{j}_k \phi_l^T + \phi_l \mathbf{j}_k^T), \quad (34)$$

where multiplication \odot , division, and exponentiation are all element-wise operations.

AsinhE pooling is set by $\Psi = \mathcal{G}_{\text{AsinhE}}(\mathbf{M}; \gamma) = \text{arcsinh}(\gamma/\mathbf{M}) = \log(\gamma/\mathbf{M} + \sqrt{1 + \gamma^2/\mathbf{M}^2})$ which yields the following:

$$\frac{\partial \Psi}{\partial \phi_{kl}} = \frac{1}{N} \frac{\gamma'}{\sqrt{\gamma^2/\mathbf{M}^2 + 1}} \odot (\mathbf{j}_k \phi_l^T + \phi_l \mathbf{j}_k^T), \quad (35)$$

where multiplication \odot , division, square root and the square are all element-wise operations.

For SigmE, trace-normalized SigmE and AsinhE pooling, final derivatives are given by Eq. (31) with the following \mathbf{D} , resp.:

$$\mathbf{D} = \frac{2\eta' e^{-\eta/\mathbf{M}}}{(1+e^{-\eta/\mathbf{M}})^2} \text{ or } \mathbf{D} = \frac{2\eta' e^{-\frac{\eta/\mathbf{M}}{\text{Tr}(\mathbf{M})+\varepsilon}}}{(1+e^{-\frac{\eta/\mathbf{M}}{\text{Tr}(\mathbf{M})+\varepsilon})^2}} \odot \mathbf{T} \text{ and } \mathbf{D} = \frac{\gamma'}{\sqrt{\gamma'^2/\mathbf{M}^2 + 1}}. \quad (36)$$

Moreover, for SigmE and AsinhE we allow β -centering so its derivative has to be included in the chain rule.

D. Derivative of Spectral Gamma

Gamma pooling has derivative which can be solved by the SVD back-propagation or the Sylvester equation if $\gamma = 0.5$:

$$2 \text{Res} \left(\text{Vec} \left(\text{Sym} \left(\frac{\partial \ell(\Psi, \mathbf{W})}{\partial \Psi} \right) \right)^T \mathbf{M}^* \right)_{d+Z' \times d+Z'} \text{ and } \mathbf{M}^* = (\mathbf{I} \otimes \mathbf{M}^{\frac{1}{2}} + \mathbf{M}^{\frac{1}{2}} \otimes \mathbf{I})^\dagger, \quad (37)$$

where \otimes and \dagger are the Kronecker product and the pseudo-inverse. Matrix reshaping to the size $m \times n$ is by $\text{Res}(\mathbf{X})_{m \times n}$.

E. Outline Proof of Theorem 3

Proof. The connection of GMRF and \mathcal{E} follows [57]. For the reminder of the proof, a simple visual inspection of profiles $g_{\text{MaxExp}}(\lambda) = 1 - (1 - \lambda)^\eta$ and $g_{\text{HDP}}(\lambda) = \exp(-t/\lambda)$ shows that $g_{\text{MaxExp}}(\lambda) \approx g_{\text{HDP}}(\lambda)$, or even $0 \leq g_{\text{MaxExp}}(\lambda) - g_{\text{HDP}}(\lambda) < \epsilon$ for some sufficiently small $\epsilon > 0$, which shows that $g_{\text{MaxExp}}(\lambda)$ is an upper bound of $g_{\text{HDP}}(\lambda)$ on the interval $\lambda \in [0, 1]$. Figure 3b shows $g_{\text{MaxExp}}(\lambda)$, $g_{\text{Gamma}}(\lambda)$ and $g_{\text{HDP}}(\lambda)$ for MaxExp, Gamma and HDP. The plot also shows that $g_{\text{Gamma}}(\lambda) \approx g_{\text{HDP}}(\lambda)$. \square

F. Proof of Theorem 4

Working with MaxExp and HDP according to their original parametrization as used in Figure 12b is difficult/intractable. Thus, for this bound, we start by a parametrization $y = \frac{t}{\lambda}$ and we note that t and η can be tied together, that is $t\eta = \alpha$. We obtain:

$$e^{-y} \leq \left(\frac{1}{e} - 1 \right) y + 1 \leq 1 - \left(1 - \frac{t}{y} \right)^\eta = 1 - \left(1 - \frac{\alpha}{\eta y} \right)^\eta, \quad (38)$$

where $\left(\frac{1}{e} - 1 \right) y + 1$ is an upper bound of e^{-y} on $y \in (0, 1)$ and a lower bound of $1 - \left(1 - \frac{\alpha}{\eta y} \right)^\eta$. Moreover, the latter equation can be tightened (as in 'lowered down') to touch $\left(\frac{1}{e} - 1 \right) y + 1$ on $y \in (0, 1)$. This process is illustrated in Figure 12a. To this end, we need to solve for the system of equations to obtain (y, α) :

$$\begin{cases} \left(\frac{1}{e} - 1 \right) y + 1 = 1 - \left(1 - \frac{\alpha}{\eta y} \right)^\eta \\ \frac{\partial \left(\frac{1}{e} - 1 \right) y + 1}{\partial y} = - \frac{\partial \left(1 - \frac{\alpha}{\eta y} \right)^\eta}{\partial y} \Rightarrow - \frac{\alpha \eta \left(1 - \frac{\alpha}{\eta y} \right)^\eta}{y(\alpha - \eta y)} = \left(\frac{1}{e} - 1 \right), \end{cases} \quad (39)$$

which simply says that we search for (y, α) for which both functions on the left- and right-hand side touch and their slopes/tangents (thus derivatives w.r.t. y) are equal. We could try directly 'lower down' $1 - \left(1 - \frac{\alpha}{\eta y} \right)^\eta$ onto e^{-y} but such an approach yields an intractable system of equations requiring numerical approximations and the use of the LambertW function. By solving Eq. (39) we get:

$$\begin{cases} \alpha(\eta) = \frac{e}{e-1} \left(\frac{\eta}{\eta+1} \right)^{\eta+1} \\ y(\eta) = \frac{e}{e-1} \left(\frac{\eta}{\eta+1} \right)^\eta. \end{cases} \quad (40)$$

Recall that we have assumed parametrization $t(\eta) = \frac{\alpha(\eta)}{\eta}$, thus $t(\eta) = \frac{e}{e-1} \frac{\eta^\eta}{(\eta+1)^{\eta+1}}$ (Eq. (18)). Furthermore, we have to check if $y \in (0, 1)$ for $1 \leq \eta \leq \infty$ in order for the bound to hold as $\left(\frac{1}{e} - 1 \right) y + 1$ is an upper bound of e^{-y} only for $y \in (0, 1)$. To this end, we firstly notice that $y(\eta)$ is monotonically decreasing

on $0 \leq \eta \leq \infty$ as shown in Figure 12c. Therefore, it suffices to check extremes of η for $1 \leq \eta \leq \infty$, that is $y(1) = \frac{1}{2} \frac{e}{e-1}$ and $\lim_{\eta \rightarrow \infty} y(\eta) = \frac{1}{e-1}$ which verifies that $y(1 \leq \eta \leq \infty) \subset (0, 1)$. This completes the first part of the proof.

The next part of the proof requires solving: $\epsilon_2(\eta) = 1 - \left(1 - \frac{\alpha(\eta)}{\eta y}\right)^\eta - e^{-y}$ which can be solved by plugging (y, α) from Eq. (40) into it. After a few of algebraic manipulations we have $\epsilon_2(\eta) = 1 - \left(\frac{\eta}{\eta+1}\right)^\eta - e^{-\frac{e}{e-1} \left(\frac{\eta}{\eta+1}\right)^\eta}$ (the right part of Eq. (17)).

We also note that, by design, $\epsilon_1(\eta) = 1 - \left(1 - \frac{\alpha}{\eta y}\right)^\eta - \left(\frac{1}{e} - 1\right) y + 1$ where $\left(\frac{1}{e} - 1\right) y + 1$ touches e^{-y} at $y=1$. Thus, we readily obtain $\epsilon_1(\eta) = \frac{e-1}{e} - \left(1 - \frac{e}{e-1} \frac{\eta^\eta}{(\eta+1)^{\eta+1}}\right)^\eta$ (the left part of Eq. (17)).

Finally, obtaining $\eta(t)$ (an inverse of $t(\eta)$) follows from simple algebraic manipulations based on the Stirling approximation.

G. Proof of Theorem 5

As in F, let parametrization $y = \frac{t}{\lambda}$ which yields a set of following equations:
$$\begin{cases} e^{-y} = \left(\frac{t}{y}\right)^\gamma \\ \frac{\partial e^{-y}}{\partial y} = \frac{\partial \left(\frac{t}{y}\right)^\gamma}{\partial y} \Rightarrow -e^{-y} = -\frac{\gamma}{\lambda} \left(\frac{t}{y}\right)^\gamma. \end{cases} \quad (41)$$

We again seek a parametrization $t(\gamma)$ for which the Gamma and HDP functions touch. However, this time the strict bound can be achieved analytically (e.g., the analytical solution to (41) exists) and thus an intermediate bounding function is not needed. After a few of algebraic manipulations, we obtain a candidate solution $y = \gamma$, which, if combined with the intermediate equation $e^y \left(\frac{t}{y}\right)^\gamma = 1$, readily yields $\gamma(t) = et$, which completes the proof.

H. Proof of Theorem 6

To obtain the proof, we note that the system of ODE from Eq. (20) and thus also Eq. (21) can be written in the span of eigenvectors of graph Laplacian [58]. Thus, we write the standard Heat Diffusion Equation as:

$$\frac{\partial \mathbf{v}'(t)}{\partial t} + \lambda \mathbf{v}'(t) = 0, \quad (42)$$

where $\mathbf{v}'(t)$ is now expressed in the span of new bases ($\mathbf{v}'(t)$ is not a derivative).

Now simply write MaxExp as $1 - (1 - \lambda^{-1})^{\eta(t)}$. In this parametrization, we use λ^{-1} as we start from the eigenvectors of the graph Laplacian rather than an autocorrelation/covariance

matrix, and we use the parametrization $\eta(t)$ derived earlier. Thus $\frac{\partial \mathbf{v}'(t)}{\partial t} = -\log(1 - \lambda^{-1}) (1 - \lambda^{-1})^{\eta(t)} \frac{\partial \eta(t)}{\partial t}$. Plugging this result into Eq. (42), we obtain:

$$-\log(1 - \lambda^{-1}) \frac{\partial \eta(t)}{\partial t} (1 - \lambda^{-1})^{\eta(t)} + f(\lambda, t) \lambda \mathbf{v}'(t) - h(\lambda, t) = 0. \quad (43)$$

After simple algebraic manipulations we find $f(\lambda, t) = -\lambda^{-1} \log(1 - \lambda^{-1}) \frac{\partial \eta(t)}{\partial t}$ and $h(\lambda, t) = -\log(1 - \lambda^{-1}) \frac{\partial \eta(t)}{\partial t}$ such that Eq. (43) holds. Putting these results together we obtain the set of ODE given as:

$$\frac{\partial \mathbf{v}'(t)}{\partial t} + \frac{\partial \eta(t)}{\partial t} \log(1 - \lambda^*) (1 - \mathbf{v}'(t)) = 0. \quad (44)$$

which is equivalent to Eq. (21) as, in the above equation, eigenvalues $\lambda^* = \lambda^{-1}$ correspond to the autocorrelation/covariance and graph Laplacian, respectively, which completes the proof.

I. Proof of Theorem 7

As above, we note that the system of ODE from Eq. (20) and thus also Eq. (22) can be rewritten in the span of eigenvectors of the graph Laplacian [58], that is Eq. (42).

Now simply write Gamma as $(\lambda^{-1})^{et}$. In this parametrization, we use λ^{-1} as we start from the eigenvectors of the graph Laplacian rather than an autocorrelation/covariance matrix, and we use the previous result stating that $\gamma = et$. Thus $\frac{\partial \mathbf{v}'(t)}{\partial t} = e (\lambda^{-1})^{et} \log(\lambda^{-1})$. Plugging this result into Eq. (42), we obtain:

$$-e \log(\lambda) (\lambda^{-1})^{et} + f(\lambda) \lambda \mathbf{v}'(t) = 0, \quad (45)$$

where $f(\lambda)$ must be equal $e \lambda^{-1} \log(\lambda)$ for Eq. (45) to hold. Putting together these results we obtain:

$$-e \log(\lambda) (\lambda^{-1})^{et} + e \log(\lambda) \lambda^{-1} \lambda \mathbf{v}'(t) = 0, \quad (46)$$

which simply tells us that the desired result is yielded by the set of ODE:

$$\frac{\partial \mathbf{v}'(t)}{\partial t} + e \log(\lambda) \mathbf{v}'(t) = 0. \quad (47)$$

which is equivalent to Eq. (22) which completes the proof.

J. Proof of Theorem 2

Probabilities $\psi^{(-)}$ and $\psi^{(+)}$ follow a simple calculus for the probability of selecting a component in a mixture model given a sample. The rest follows simple algebraic manipulations.

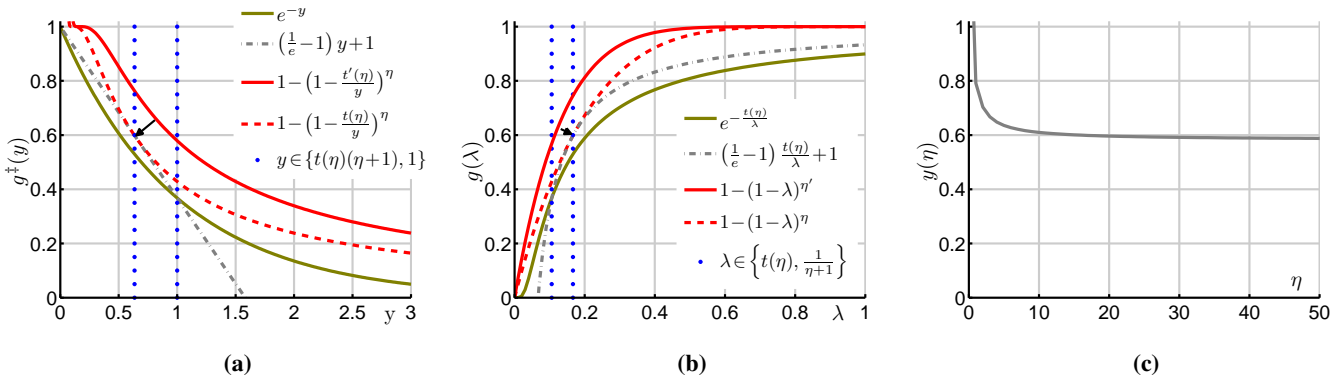


Fig. 12: In Fig. 12a, we show pooling functions reparametrized according to $y = \frac{t}{\lambda}$. Specifically, we have HDP given by e^{-y} , its upper bound $(\frac{1}{e} - 1) y + 1$ for $y \in (0, 1)$, and MaxExp given by $1 - \left(1 - \frac{t(\eta)}{y}\right)^\eta$ which we ‘lower down’ onto $(\frac{1}{e} - 1) y + 1$ as illustrated by the black arrow. As we tighten the bound, some initial MaxExp with $t'(\eta)$ becomes MaxExp with $t(\eta)$. Blue vertical lines indicate y at which we measure ϵ_1 and ϵ_2 . Fig. 12b illustrates the same pooling operations as in Fig. 12a but without the reparametrization, that is, we show $g(\lambda)$ rather than $g^+(y)$. Note the corresponding y and λ ranges in both figures indicated by the blue dashed lines. Fig. 12c shows that $y(\eta)$ is monotonically decreasing on $\eta \in (0, \infty)$.

K. Proof of Remark 2

We note that the concavity of SigmE on interval $[0, 1]$ is at its maximum for a point $p''(\eta') = \log(\sqrt{3}+2)/\eta'$ which we obtain as a solution to:

$$\frac{\partial^3}{\partial p^3} \frac{2}{1+e^{-\eta'p}} - 1 = 0. \quad (48)$$

Subsequently, we formulate the square loss between SigmE and MaxExp at p'' , and take its derivative to solve it for η (or η'):

$$\frac{\partial}{\partial \eta} \left(\frac{2}{1+e^{-\eta' \frac{\log(\sqrt{3}+2)}{\eta'}}} - 1 - \left(1 - \left(1 - \frac{\log(\sqrt{3}+2)}{\eta'} \right)^\eta \right) \right)^2 = 0, \quad (49)$$

Regarding the maximum error between SigmE and MaxExp, one should technically find the maximum square difference between SigmE and MaxExp parametrized by $\eta'(\eta)$. However, such an equation has no closed form. Thus, an easier approximate measure is to consider the difference between SigmE and MaxExp at $2p''(\eta')$ or $p'''(\eta') = \log(\sqrt{26\sqrt{105}+270}/2 + \sqrt{105}/2 + 13/2)/\eta'$, which is the solution to the fifth derivative of MaxExp:

$$\frac{\partial^5}{\partial p^5} \frac{2}{1+e^{-\eta'p}} - 1 = 0. \quad (50)$$

L. Derivation of the kernel linearization in Eq. (4)

Let $G_\sigma(\mathbf{x}-\mathbf{y}) = \exp(-\|\mathbf{x}-\mathbf{y}\|_2^2/2\sigma^2)$ be a Gaussian RBF kernel with a bandwidth σ . Kernel linearization refers to rewriting G_σ as an inner-product of two (in)finite-dimensional feature maps. Specifically, we employ the inner product of d' -dimensional isotropic Gaussians centered at $\mathbf{x}, \mathbf{y} \in \mathbb{R}^{d'}$:

$$G_\sigma(\mathbf{x}-\mathbf{y}) = \left(\frac{2}{\pi\sigma^2} \right)^{\frac{d'}{2}} \int_{\zeta \in \mathbb{R}^{d'}} G_{\sigma/\sqrt{2}}(\mathbf{x}-\zeta) G_{\sigma/\sqrt{2}}(\mathbf{y}-\zeta) d\zeta, \quad (51)$$

Eq. (51) can be thought of as a convolution of function $G_{\sigma/\sqrt{2}}(\mathbf{x}-\zeta)$ with $G_{\sigma/\sqrt{2}}(\mathbf{y}-\zeta)$ centered at \mathbf{x} and \mathbf{y} , respectively. Both functions are isometric multivariate Normal distributions if normalized by $\left(\frac{2}{\pi\sigma^2}\right)^{\frac{d'}{2}}$. To prove that Eq. (51) holds, we consider the sum (denoted as S) of two i.i.d. random variables distributed according to two Normal distributions (variance $\sigma/\sqrt{2}$). We note that S is then also distributed according to the Normal distribution (variance σ).

Eq. (51) is approximated by replacing the integral with the so-called Riemann sum over Z pivots ζ_1, \dots, ζ_Z which represent centers of so-called approximating rectangles in the sum:

$$G_\sigma(\mathbf{x}-\mathbf{y}) \approx c \sum_{i \in \mathcal{I}_Z} \varphi_i(\mathbf{x}; \{\zeta_i\}_{i \in \mathcal{I}_Z}) \cdot \varphi_i(\mathbf{y}; \{\zeta_i\}_{i \in \mathcal{I}_Z}), \quad (52)$$

where $\varphi(\mathbf{x}; \{\zeta_i\}_{i \in \mathcal{I}_Z}) = \left[G_{\sigma/\sqrt{2}}(\mathbf{x}-\zeta_1), \dots, G_{\sigma/\sqrt{2}}(\mathbf{x}-\zeta_Z) \right]^T$, and c is a normalization constant related to the normalization in Eq. (51) and the width of approximating rectangles.

Finally, one dim. input features (Cartesian coordinates normalized in range 0–1) from which we form feature maps correspond to spatial locations in the conv. layer, and so are distributed uniformly. Thus, to cover the entire support set ($\{s: f(s) > 0\}$) of random variable $s \sim S$ and obtain roughly a uniform approximation quality across the support (size of support set), we select Z pivots at equally spaced intervals, that is $[\zeta_1; \dots; \zeta_Z] = [-0.2 : 1.4/(Z-1) : 1.2]$. The range exceeds $[0, 1]$ as the arms ($2 \times$ standard deviation) of Gaussians ($\sigma/\sqrt{2}$) at extreme locations 0 and 1 require the support of S to be roughly $[-2\sigma/\sqrt{2}; 1+2\sigma/\sqrt{2}]$ to cover ~ 95.0 of the support set. In practice, we found that

the support $[-0.2; 1.2]$ and Z in range 3–10 are sufficient. As $0 < c < \infty$ is a constant, its exact value does not influence the information captured by the maps, thus we set $c=1$.

M. MaxExp in few-shot learning.

Below we present a motivation similar to one presented in Section 6. However, we employ the Binomial PMF and the variance-based modeling in contrast to the uniform PMF and its set support modeling from Section 6. Theorem 1 states that MaxExp performs a co-occurrence detection rather than counting. For classification problems, let a probability mass function $p_{X_{kl}}(x) = \text{Binom}(x; N, p)$ for $x = 0, \dots, N$ and some p tell the probability that co-occurrence (k, l) between ϕ_{kn} and ϕ_{ln} happened $x = 0, \dots, N$ times given an image, where N is the number of feature vectors to aggregate. Using second-order pooling without MaxExp requires a classifier to observe $N+1$ training samples of two features co-occurring in quantities $0, \dots, N$ to memorize their possible co-occurrence counts. For similarity learning, our ϑ stacks pairs of samples to compare, thus a similarity learner has to deal with a probability mass function $R_{kl} = X_{kl} + Y_{kl}$ describing configurations of two features co-occurring whose $\text{var}(p_{R_{kl}}) = 2Np(p-1) > \text{var}(p_{X_{kl}}) = Np(p-1)$ as random variable $X=Y$ (same class). For J -shot learning which stacks one query and J support matrices (per class) in relation descriptor, $R'_{kl} = X_{kl}^{(1)} + \dots + X_{kl}^{(J)} + Y_{kl}$, $X^{(j)} = Y, \forall j \in \mathcal{I}_J$ and we have $\text{var}(p_{R'_{kl}}) = (J+1)Np(p-1)$ indicating that the similarity learner has to memorize more configurations as J and/or N grow.

However, this situation is alleviated by MaxExp whose probability mass function can be modeled as $p_{X_{kl}^{\text{MaxExp}}}(x) = p$ if $x=1$, $p_{X_{kl}^{\text{MaxExp}}}(x) = 1-p$ if $x=0$, as MaxExp detects a co-occurrence (or its lack). Thus, for J -shot learning, $\text{var}(p_{R'_{kl}^{\text{MaxExp}}}) = (J+1)p(p-1) \ll \text{var}(p_{R'_{kl}}) = (J+1)Np(p-1)$. The ratio of variances of MaxExp normalized to non-normalized few-shot learning equals $\kappa' = 1/N$ which shows that if MaxExp is used, the similarity learner has to memorize representations which have $N \times$ less variance compared to the case without MaxExp.

In our work, we considered an intuitive relation descriptor:

$$\vartheta_{(\otimes+P)}(\{\Phi_n\}_{n \in \mathcal{J}}, \Phi^*) = \left[\mathfrak{g}_i \left(\frac{1}{J} \sum_{n \in \mathcal{J}} \frac{1}{N} \Phi_n \Phi_n^T \right); \mathfrak{g}_1 \left(\frac{1}{N} \Phi^* \Phi^{*T} \right) \right]. \quad (53)$$

Despite its intuitive nature, $\vartheta_{(\otimes+P)}$ in Eq. (53) performed $\sim 1-2\%$ worse than $\vartheta_{(\otimes+L)}$ in Eq. (26) for $J > 1$, and thus was deemed not fit for presentation. Analyzing variance, one can notice that query and support parts of Eq. (53) are described by $p_{X_{kl}^{\text{MaxExp}}}(x)$ each, thus $\text{var}(p_{X_{kl}^{\text{MaxExp}} + Y_{kl}^{\text{MaxExp}}}) = 2p(p-1)$, $Y=X$. Due to a very low variance, the representational power of this relation descriptor is simply insufficient if $J > 1$. Detecting if two co-occurring features (k, l) between ϕ_{kn} and ϕ_{ln} are detected at least once in JN trials can be thought of as creating one simplified image representing such detections which deprives the similarity learner the individual per-image co-occurrence statistics.

N. Deriving derivative of fast spectral MaxExp.

Let $\ell(\Psi, \mathbf{W})$ be some classification loss (or any layer with param. \mathbf{W}) where $\Psi = \hat{\mathcal{G}}_{\text{MaxExp}}(\mathbf{M}) \in \mathcal{S}_{++}^d$ (or \mathcal{S}_{++}) are our feature maps. Then, we obtain a versatile equation from Eq. (16):

$$\sum_{k,l} \frac{\partial \ell(\Psi, \mathbf{W})}{\partial \Psi_{kl}} \frac{\partial \Psi_{kl}}{\partial \mathbf{M}} = - \sum_{n=0}^{\eta-1} (\mathbb{I} - \mathbf{M})^n \frac{\partial \ell(\Psi, \mathbf{W})}{\partial \Psi} (\mathbb{I} - \mathbf{M})^{\eta-1-n}, \quad (54)$$

which simplifies to:

$$\sum_{k,l} \frac{\partial \ell(\Psi, \mathbf{W})}{\partial \Psi_{kl}} \frac{\partial \Psi_{kl}}{\partial \mathbf{M}} = -2 \text{Sym} \left(\sum_{n=0}^{\lfloor \frac{\gamma}{2} \rfloor - 1} (\mathbb{I} - \mathbf{M})^n \frac{\partial \ell(\Psi, \mathbf{W})}{\partial \Psi} (\mathbb{I} - \mathbf{M})^{\eta-1-n} \right) - \begin{cases} (\mathbb{I} - \mathbf{M})^{\lfloor \frac{\gamma}{2} \rfloor} \frac{\partial \ell(\Psi, \mathbf{W})}{\partial \Psi} (\mathbb{I} - \mathbf{M})^{\lfloor \frac{\gamma}{2} \rfloor} & \text{if } \eta \text{ is odd} \\ 0 & \text{otherwise.} \end{cases} \quad (55)$$

For the second half of indexes $0, \dots, \eta-1$ (even η), Eq. (55) uses transposed summation terms corresponding to the first half of indexes instead of recomputing them as in Eq. (54). For odd η , only the term corresponding to index $\lfloor \frac{\eta}{2} \rfloor$ is not aggregated twice.

O. Derivative of fast spectral Gamma for integers $\gamma \geq 1$.

Let $\ell(\Psi, \mathbf{W})$ be some classification loss/layer with param. \mathbf{W} (as in Eq. (55)) and $\Psi = \hat{\mathcal{G}}_{\text{Gamma}}(\mathbf{M}) \in \mathcal{S}_+^d$ (or \mathcal{S}_{++}) be our feature maps. Then, the derivative of Gamma for integers $\gamma(t) \geq 1$ is:

$$\sum_{k,l} \frac{\partial \ell(\Psi, \mathbf{W})}{\partial \Psi_{kl}} \frac{\partial \Psi_{kl}}{\partial \mathbf{M}} = \sum_{n=0}^{\gamma-1} \mathbf{M}^n \frac{\partial \ell(\Psi, \mathbf{W})}{\partial \Psi} \mathbf{M}^{\gamma-1-n} = \quad (56)$$

$$2 \text{Sym} \left(\sum_{n=0}^{\lfloor \frac{\gamma}{2} \rfloor - 1} \mathbf{M}^n \frac{\partial \ell(\Psi, \mathbf{W})}{\partial \Psi} \mathbf{M}^{\gamma-1-n} \right) + \begin{cases} \mathbf{M}^{\lfloor \frac{\gamma}{2} \rfloor} \frac{\partial \ell(\Psi, \mathbf{W})}{\partial \Psi} \mathbf{M}^{\lfloor \frac{\gamma}{2} \rfloor} & \text{if } \gamma \text{ is odd} \\ 0 & \text{otherwise.} \end{cases}$$

P. Derivations of Fast Approximate HDP (FAHDP).

Let $0 \leq \lambda \leq 1$. Note that for $\lambda = 1$ and $t \geq \log(10/9) \approx -0.105$, HDP given by $e^{-t/\lambda}$ yields $e^{-t} \leq 0.9$ (it drops more for larger t) while MaxExp given by $1 - (1-\lambda)^\eta$ and Gamma given by λ^γ yield 1 for $\lambda = 1$ and $t \geq 0$. For this reason, we reparametrize MaxExp and Gamma as $e^{-t}(1 - (1-\lambda)^{\tilde{\eta}(t)})$ and $e^{-t}\lambda^{\tilde{\gamma}(t)}$, respectively, where e^{-t} is the scaling factor ensuring that MaxExp/Gamma (and thus FAHDP) and HDP have the same magnitude at $\lambda = 1$.

Time-reversed HDP. To obtain a good approximation of time-reversed HDP ($t < 1$) by a scaled MaxExp, we take steps as those in Appendix F, that is we use the substitution $y = \frac{t}{\lambda}$, we tie together t and $\tilde{\eta}$ as $t\tilde{\eta} = \alpha$. Subsequently, we seek to ‘lowered down’ MaxExp given now as $e^{-t}(1 - (1 - \frac{t}{y})^{\tilde{\eta}})$ onto $(\frac{1}{e} - 1)y + 1$, an upper bound of reparametrized HDP given as e^{-y} . We obtain:






Fig. 13: In Fig. 13a, we show pooling functions reparametrized according to $y = \frac{t}{\lambda}$ and scaled where appropriate to match the value of HDP at $\lambda = 1$. Specifically, we have HDP given by e^{-y} , its upper bound $(\frac{1}{e} - 1)y + 1$ for $y \in (0, 1)$, and a scaled by $1 - \frac{e-1}{e}t$ MaxExp given as $(1 - \frac{e-1}{e}t)(1 - (1 - \frac{t}{y})^{\tilde{\eta}})$ which we ‘lower down’ onto $(\frac{1}{e} - 1)y + 1$ as indicated by the black arrow. As we tighten the bound, some initial MaxExp with $\tilde{\eta}(t)$ becomes MaxExp with $\tilde{\eta}(t)$. Blue vertical lines are $y_0 = t$ and $y_1 = t\tilde{\eta}(t)$ at which the scaled MaxExp touches $(\frac{1}{e} - 1)y + 1$. The green lines indicate the range of $y_1 \in [\frac{1}{e-1}, \frac{1}{2} \frac{e}{e-1}]$ given by the image of $y(1 \leq \tilde{\eta} \leq \infty)$. Fig. 13b that the scaled MaxExp given as $(1 - \frac{e-1}{e}t)(1 - (1 - \frac{t}{y})^{\tilde{\eta}})$ is lower-bounded by a tightly-scaled MaxExp $e^{-t}(1 - (1 - \frac{t}{y})^{\tilde{\eta}})$ which is lower-bounded by a linear function $a(y(t, \tilde{\eta}(t), \tilde{\eta}(t)) + b(y(t, \tilde{\eta}(t), \tilde{\eta}(t)))$ (the blue lines are where the latter two functions touch). Again, the green lines indicate the range of y_1 for $1 \leq \tilde{\eta} \leq \infty$ and the gray lines indicate $y_2 = y_0 = t$ and y_3 at which $ay + b$ touches r^{-y} . Finally, 13c verifies that $y_2(t) \leq y_1(t) \leq y_3(t)$ for $t \in [0, 1]$ which is necessary to prove that the tightly-scaled MaxExp is an upper bound of HDP.

$$\begin{cases} (\frac{1}{e} - 1)y + 1 = e^{-t}(1 - (1 - \frac{t}{y})^{\tilde{\eta}}) \\ \frac{\partial (\frac{1}{e} - 1)y + 1}{\partial y} = -\frac{\partial e^{-t}(1 - \frac{t}{y})^{\tilde{\eta}}}{\partial y}. \end{cases} \quad (57)$$

We notice that (57) may not have a closed form solution.

Scaled MaxExp (Fig. 13a). We notice that $1 - \frac{e-1}{e}t$ is an upper bound of e^t on interval $t \in (0, 1)$, thus we use scaling $1 - \frac{e-1}{e}t$ and we solve the following set of equations instead of Eq. (57):

$$\begin{cases} (\frac{1}{e} - 1)y + 1 = (1 - \frac{e-1}{e}t)(1 - (1 - \frac{t}{y})^{\tilde{\eta}}) \\ \frac{\partial (\frac{1}{e} - 1)y + 1}{\partial y} = -(1 - \frac{e-1}{e}t) \frac{\partial (1 - \frac{t}{y})^{\tilde{\eta}}}{\partial y}. \end{cases} \quad (58)$$

By solving Eq. (58), we get $y = \alpha = t\tilde{\eta} \vee y = t = \frac{\alpha}{\tilde{\eta}}$. We denote the first result as $y_1 = \alpha = t\tilde{\eta}$ and the later one as $y_0 = t = \frac{\alpha}{\tilde{\eta}}$. We notice that $y_0 \leq y_1$ for $\tilde{\eta} \geq 1$ and that y_0 corresponds to the solution where $\lambda = \frac{t}{y_0} = \frac{t}{t} = 1$ which is one of the two points for which the set of equations in (58) is fulfilled. The other point for which the above set of eq. is fulfilled is $y = \alpha$. Thus, we have:

$$\alpha(\tilde{\eta}) = y(\tilde{\eta}) = t\tilde{\eta} \quad \text{and} \quad t(\tilde{\eta}) = \frac{\frac{e}{e-1} \left(\frac{\tilde{\eta}-1}{\tilde{\eta}} \right)^{\tilde{\eta}}}{\left(\frac{\tilde{\eta}-1}{\tilde{\eta}} \right)^{\tilde{\eta}} + \tilde{\eta} - 1}, \quad (59)$$

where $\tilde{\eta} > 1$. For $\tilde{\eta} = 1$, $t(\tilde{\eta})$ and $\alpha(\tilde{\eta})$ are undefined but $\lim_{\tilde{\eta} \rightarrow 1} t(\tilde{\eta}) = \lim_{\tilde{\eta} \rightarrow 1} \alpha(\tilde{\eta}) = \frac{e}{2(e-1)}$ from the L’Hospital’s rule.

To obtain $\tilde{\eta}(t)$, the inverse of $t(\tilde{\eta})$, we firstly apply a substitution $\tilde{\eta} = \hat{\eta} + 1$ and evaluate $t(\hat{\eta} + 1)$:

$$t(\hat{\eta} + 1) = \frac{e}{e-1} \frac{\hat{\eta}^{\hat{\eta}}}{\hat{\eta}^{\hat{\eta}} + (\hat{\eta} + 1)^{\hat{\eta} + 1}}. \quad (60)$$

Next, we employ the Stirling approximation $\hat{\eta}^{\hat{\eta}} \approx \hat{\eta}! / \sqrt{2\pi\hat{\eta}} = \tilde{\eta}^{\tilde{\eta}}$. We note that $\hat{\eta}^{\hat{\eta}} \leq \tilde{\eta}^{\tilde{\eta}}, \forall \hat{\eta}, \tilde{\eta} \geq 1$ but the gap decreases monotonically as $\hat{\eta}, \tilde{\eta} \rightarrow \infty$. Subsequently, we obtain:

$$t(\hat{\eta} + 1) \approx \frac{\frac{e-1}{e} \sqrt{\tilde{\eta} + 1}}{\sqrt{\tilde{\eta} + 1} + e(\hat{\eta} + 1)\sqrt{\tilde{\eta}}} \Rightarrow t(\tilde{\eta}) \approx \frac{\frac{e-1}{e} \sqrt{\tilde{\eta}}}{\sqrt{\tilde{\eta}} + e\tilde{\eta}\sqrt{\tilde{\eta} - 1}} = \tilde{t}(\tilde{\eta}). \quad (61)$$

From Eq. (61), it is straightforward to recover $\tilde{\eta}(t)$ which is:

$$\tilde{\eta}(t) \approx 0.5 + \sqrt{0.25 + (1/(t(e-1)) - 1/e)^2} = \tilde{\eta}(t). \quad (62)$$

For Eq. (58), we have to check if $y \in (0, 1)$ for $1 \leq \tilde{\eta} \leq \infty$ in

order for the bound to hold as $(\frac{1}{e}-1)y+1$ is an upper bound of e^{-y} only for $y \in (0, 1)$. To this end, we firstly notice that $y(\bar{\eta})$ is monotonically decreasing on $1 \leq \bar{\eta} \leq \infty$ as Figure 14a shows. Therefore, it suffices to check extremes of $\bar{\eta}$ for $1 \leq \bar{\eta} \leq \infty$, that is $y(1) = \frac{1}{2} \frac{e}{e-1}$ and $\lim_{\bar{\eta} \rightarrow \infty} y(\bar{\eta}) = \frac{1}{e-1}$ which verifies that $y(1 \leq \bar{\eta} \leq \infty) \subset (0, 1)$. Thus, for $y \in (0, 1)$ we have:

$$e^{-y} \leq \left(\frac{1}{e}-1\right)y+1 \leq \left(1-\frac{e-1}{e}t\right)\left(1-\left(1-\frac{t}{y}\right)\bar{\eta}\right). \quad (63)$$

Tightly-scaled MaxExp (Fig. 13b). We tighten the bound by replacing $1 - \frac{e-1}{e}t$ with e^{-t} as $e^{-t} \leq 1 - \frac{e-1}{e}t$ for $t \in [0, 1]$. Thus, we seek $ay+b$ in:

$$e^{-y} \leq ay+b \leq e^{-t}\left(1-\left(1-\frac{t}{y}\right)\bar{\eta}\right). \quad (64)$$

Firstly, we notice that $\tilde{\eta}(t) \geq \bar{\eta}(t)$ whose gap decreases rapidly and monotonically as $t \rightarrow 0$. We can recover $\epsilon_{\tilde{\eta}}(\bar{\eta}) = \tilde{\eta} - \bar{\eta} = \tilde{\eta}(t(\bar{\eta})) - \bar{\eta}$ and $\epsilon_{\tilde{t}}(t) = t - \tilde{\eta}(t)$. This upper bound $\tilde{\eta}(t) \geq \bar{\eta}(t)$ assures that $1 - \left(1 - \frac{t}{y}\right)\bar{\eta} \geq 1 - \left(1 - \frac{t}{y}\right)\tilde{\eta}$ which turns out to be a sufficient upper bound to fulfill (64). To this end, we note that as $\left(\frac{1}{e}-1\right)y+1 = \left(1-\frac{e-1}{e}t\right)\left(1-\left(1-\frac{t}{y}\right)\bar{\eta}\right)$ for y_0 and y_1 (defined under Eq. (58)), that is the nonlinear function touches the linear one only at y_0 and y_1 , then $ay+b = e^{-t}\left(1-\left(1-\frac{t}{y}\right)\bar{\eta}\right)$ at y_0 and y_1 because $\left(\frac{1}{e}-1\right)y+1$ and e^{-t} perform just some scaling of $1 - \left(1 - \frac{t}{y}\right)\bar{\eta}$. Thus, we define $\bar{\eta}(t) = \tilde{\eta}(2t - t(\tilde{\eta}(t)))$ and we readily obtain:

$$a(t, \bar{\eta}, \tilde{\eta}) = \frac{e^{-t} \left(\frac{\bar{\eta}(t)-1}{\bar{\eta}(t)}\right)^{\bar{\eta}}}{t \frac{1-\bar{\eta}(t)}{\bar{\eta}(t)}} \quad \text{and} \quad b(t, \bar{\eta}, \tilde{\eta}) = e^{-t} \left(1 - \frac{\left(\frac{\bar{\eta}(t)-1}{\bar{\eta}(t)}\right)^{\bar{\eta}}}{1-\bar{\eta}(t)}\right). \quad (65)$$

Eq. $a(t, \bar{\eta}, \tilde{\eta})y + b(t, \bar{\eta}, \tilde{\eta})$ (a and b are defined in (65)) is a lower bound of $e^{-t}\left(1-\left(1-\frac{t}{y}\right)\bar{\eta}\right)$ on interval $y \in (t, \infty)$ by design. Now, it suffices to check if $a(t, \bar{\eta}(t), \tilde{\eta}(t))y + b(t, \bar{\eta}(t), \tilde{\eta}(t))$ is also an upper bound of e^{-y} on interval $y \in (y_0, y_1) \equiv (t, t\bar{\eta}(t))$. Thus, we substitute a and b from Eq. (65) into following:

$$e^{-y^*} \leq a(t, \bar{\eta}(t), \tilde{\eta}(t))y^* + b(t, \bar{\eta}(t), \tilde{\eta}(t)). \quad (66)$$

After several manipulations we arrive at the solution $y^* \in [y_2, y_3]$:


$$y_2 = \text{LambertW}_{-1}\left(\frac{e^{\frac{b}{a}}}{a}\right) - \frac{b}{a} \quad \text{and} \quad y_3 = \text{LambertW}_0\left(\frac{e^{\frac{b}{a}}}{a}\right) - \frac{b}{a}, \quad (67)$$

where the LambertW solves $W(z)e^{W(z)} = z$. We observe that $y_2 = y_0$ and $y_2 \leq y_1 \leq y_3$ (Fig. 13c), and thus we have:

$$[t, t\bar{\eta}(t)] \subset \left[t, \text{LambertW}_0\left(\frac{e^{\frac{b}{a}}}{a}\right) - \frac{b}{a}\right], \quad (68)$$

where $a(t, \bar{\eta}(t), \tilde{\eta}(t))$ and $b(t, \bar{\eta}(t), \tilde{\eta}(t))$ depend explicitly on time t . Eq. (68) verifies the tighter bound achieved by Eq. (64).

Finally, bounds ϵ_3 and ϵ_4 at y_1 and y_3 may be evaluated by



(a)

plugging them respectively into $e^{-t}\left(1-\left(1-\frac{t}{y}\right)\bar{\eta}\right) - y^{-y}$:

$$e^{-t}e^{-t}\left(\frac{\bar{\eta}(t)-1}{\bar{\eta}(t)}\right)^{\bar{\eta}(t)} - e^{-t\bar{\eta}(t)} = \epsilon_3 \leq \epsilon_4 = e^{-t}e^{-t}\left(\frac{y_3(t)-t}{y_3(t)}\right)^{\bar{\eta}(t)} - e^{-y_3(t)}, \quad (69)$$

Time-forward HDP. To obtain a good approximation of time-forward HDP ($t \geq 1$) by Gamma, we reparametrize and scale Gamma given as $e^{-t}\lambda^{\bar{\gamma}}$. We solve the following set of equations:

$$\begin{cases} e^{-t/\lambda} = e^{-t}\lambda^{\bar{\gamma}} \\ \frac{\partial e^{-t/\lambda}}{\partial \lambda} = \frac{\partial e^{-t}\lambda^{\bar{\gamma}}}{\partial \lambda} \Rightarrow -\frac{t}{\lambda^2}e^{-t/\lambda} = \bar{\gamma}e^{-t}\lambda^{\bar{\gamma}-1}, \end{cases} \quad (70)$$

which yields the following set of equations:

$$\begin{cases} \bar{\gamma}e^{t/\lambda}\lambda^{\bar{\gamma}+1} = te^t \\ e^{t/\lambda}\lambda^{\bar{\gamma}} = e^t. \end{cases} \quad (71)$$

After some manipulations, the above set of equations further reduces to:

$$\bar{\gamma} - \bar{\gamma} \log(\bar{\gamma}) = t - \bar{\gamma} \log(t). \quad (72)$$

To solve this equation, it suffices to use an approximation of logarithm $\log(x) \approx rx^{1/r} - r$ for order $r = 2$ which is also an upper bound of $\log(x)$, that is $\log(x) \leq rx^{1/r} - r$. Solving Eq. (72) yields $\bar{\gamma}(t) = t$.

Fast Approximate HDP (FAHDP). We combine the right-hand side of Eq. (64) given as $e^{-t}\left(1-\left(1-\lambda\right)\bar{\eta}\right)$ with $\bar{\eta}$ in Eq. (62) for $t \in (0, 1)$ and $e^{-t}\lambda^{\bar{\gamma}}$ defined just above Eq. (70) with $\bar{\gamma}(t) = t$ for $t \geq 0$. We assume that we operate on an SVD of an autocorrelation matrix \mathbf{M} that is normalized by its trace. Thus, we have p_i which correspond to trace-normalized λ_i , and we define:


$$g_{\text{FAHDP}}(p; t) = e^{-t} \cdot \begin{cases} 1 - (1-p)^{h_t(\bar{\eta}(t))} & \text{if } t < 1 \\ p^{h_t(\bar{\gamma}(t))} & \text{if } t \geq 1, \end{cases} \quad (73)$$

where $h_t(x) = x$ or $h_t(\cdot)$ is defined in the same way as for Eq. (23). Eq. (73) operates on the trace-normalized eigenvalues $0 \leq p \leq 1$ and it can be combined with Algorithm 1. Moreover, Eq. (73) has its closed form given in Eq. (23) which enjoys a fast back-propagation. FAHDP and HDP are illustrated together in Fig. 14b.

Q. Additional results on GIN0 with SOP+PN.

Method	IMDB-BIN	IMDB-MULTI	PTC	NCI1
<i>GIN0</i>	75.5±4.0	51.3±3.4	66.1±6.7	82.1±1.8
<i>SOP+Newton-Schulz</i>	77.3±4.4	52.2±3.2	68.8±6.4	82.3±2.2
<i>SOP+Spec. MaxExp(F)</i>	77.8±3.6	53.5±2.1	72.2±7.0	82.5±2.5

TABLE 12: Classification with GIN0 and various pooling methods on IMDB-BIN, IMDB-MULTI, PTC and NCI1 (after obtaining hyperparameters, the train and validation splits were combined for retraining).



(b)

Fig. 14: In Fig. 14a, we show that $y(\bar{\eta})$ decreases monotonically as $\bar{\eta} \rightarrow \infty$. Fig. 14b illustrates that the FAHDP (combination of the tightly-scaled MaxExp and the scaled Gamma) is a tight upper bound of HDP.