

Time-rEversed diffusioN tEnsor Transformer: A new TENET of Few-Shot Object Detection

Shan Zhang^{*,†}, Naila Murray[♣], Lei Wang[♦], and Piotr Koniusz^{*,§,†}

[†]Australian National University [♣]Meta AI

[♦]University of Wollongong [§]Data61/CSIRO

[†]firstname.lastname@anu.edu.au, [♦]leiw@uow.edu.au, [♣]murrayn@fb.com

Abstract. In this paper, we tackle the challenging problem of Few-shot Object Detection. Existing FSOD pipelines (i) use average-pooled representations that result in information loss; and/or (ii) discard position information that can help detect object instances. Consequently, such pipelines are sensitive to large intra-class appearance and geometric variations between support and query images. To address these drawbacks, we propose a Time-rEversed diffusioN tEnsor Transformer (TENET), which i) forms high-order tensor representations that capture multi-way feature occurrences that are highly discriminative, and ii) uses a transformer that dynamically extracts correlations between the query image and the entire support set, instead of a single average-pooled support embedding. We also propose a Transformer Relation Head (TRH), equipped with higher-order representations, which encodes correlations between query regions and the entire support set, while being sensitive to the positional variability of object instances. Our model achieves state-of-the-art results on PASCAL VOC, FSOD, and COCO.

Keywords: few-shot object detection; transformer; multiple order pooling; high order pooling; heat diffusion process;

1 Introduction

Object detectors based on deep learning, usually addressed by supervised models, achieve impressive performance [31,32,33,8,27,13] but they rely on a large number of images with human-annotated class labels/object bounding boxes. Moreover, object detectors cannot be easily extended to new class concepts not seen during training. Such a restriction limits supervised object detectors to predefined scenarios. In contrast, humans excel at rapidly adapting to new scenarios by “*storing knowledge gained while solving one problem and applying it to a different but related problem*” [40], also called as “*transfer of practice*” [41].

Few-shot Object Detection (FSOD) [4,11,12,45,7,50,51] methods mimic this ability, and enable detection of test classes that are disjoint from training classes.

* SZ was mainly in charge of the pipeline/developing the transformer. PK (corresponding author) was mainly in charge of mathematical design of TENET & TSO.

Code: <https://github.com/ZS123-lang/TENET>.

They perform this adaptation using a few “support” images from test classes. Successful FSOD models must (i) find promising candidate regions-of-Interest (RoIs) in query images; and (ii) accurately regress bounding box locations and predict RoI classes, under large intra-class geometric and photometric variations.

To address the first requirement, approaches [45,7,50,51] use the region proposal network [33]. For example, FSOD-ARPN [7], PNSD [50] and KFSOD [51] cross-correlate query feature maps with a class prototype formed from average-pooled (*ie.*, first-order) features, second-order pooled representations and kernel-pooled representations, respectively. These methods use a single class prototype which limits their ability to leverage diverse information from different class samples. Inspired by Transformers [38], approach [23] uses average pooling over support feature maps to generate a vector descriptor per map. Attention mechanism is then used to modulate query image features using such descriptors.

The above methods rely on first- and second-order pooling, while so-called higher-order pooling is more discriminative [17,18,15]. Thus, we propose a non-trivial Time-rEversed diffusioN tEnsor Transformer (TENET). With TENET, higher-order tensors undergo a time-reversed heat diffusion to condense signal on super-diagonals of tensors, after which coefficients of these super-diagonals are passed to a Multi-Head Attention transformer block. TENET performs second-, third- and fourth-order pooling. However, higher-order pooling suffers from several issues, *ie.*, (i) high computational complexity of computing tensors with three/more modes, (ii) non-robust tensor estimates due to the limited number of vectors being aggregated, and (iii) tensor burstiness[17].

To this end, we propose a Tensor Shrinkage Operator (TSO) which generalizes spectral power normalization (SPN) operators [18], such as the Fast Spectral MaxExp operator (MaxExp(F)) [18], to higher-order tensors. As such, it can be used to reduce tensor burstiness. Moreover, by building on the linear algebra of the heat diffusion process (HDP) [34] and recent generalisation of HDP to SPN operators [18], we also argue that such operators can reverse the diffusion of signal in autocorrelation or covariance matrices, and high-order tensors, instead of just reducing the burstiness. Using a parametrization which lets us control the reversal of diffusion, TSO condenses signal captured by a tensor toward its super-diagonal, preserving information along it. This super-diagonal serves as our final representation, reducing the feature size from d'' to d , making our representation computationally tractable. Finally, shrinkage operators are known for their ability to estimate covariances well when only a small number of samples are available [22]. To the best of our knowledge, we are the first to show that MaxExp(F) is a shrinkage operator, and to propose TSO for orders $r \geq 2$.

To address the second requirement, FSOD-ARPN introduces a multi-relation head that captures global, local and patch relations between support and query objects, while PNSD passes second-order autocorrelation matrices to a similarity network. However, FSOD-ARPN and PNSD do not model spatial relations [9]. The QSAM [23] uses attention to highlight the query RoI vectors that are similar to the set of support vectors (obtained using only first-order spatial average pooling). Thus, we introduce a Transformer Relation Head (TRH) to

improve modeling of spatial relations. TRH computes self-attention between spatially-aware features and global spatially invariant first-, second- and higher-order TENET representations of support and/or query RoI features. The second attention mechanism of TRH performs cross-attention between Z support embeddings (for Z -shot if $Z \geq 2$), and a set of global representations of query RoIs. This attention encodes similarities between query RoIs and support samples.

Our FSOD pipeline contains TENET Region Proposal Network (TENET RPN) and the TRH, both equipped with discriminative TENET representations, improving generation of RoI proposals and modeling of query-support relations.

Below are our contributions:

- i. We propose a Time-rEversed diffusiON tEnsor Transformer, called TENET, which captures high-order patterns (multi-way feature cooccurrences) and decorrelates them/reduces tensor burstiness. To this end, we generalize the MaxExp(F) operator [18] for autocorrelation/covariance matrices to higher-order tensors by introducing the so-called Tensor Shrinkage Operator (TSO).
- ii. We propose a Transformer Relation Head (TRH) that is sensitive both to the variability between the Z support samples provided in a Z -shot scenario, and to positional variability between support and query objects.
- iii. In §4.1, we demonstrate that TSO emerges from the MLE-style minimization over the Kullback-Leibler (KL) divergence between the input and output spectrum, with the latter being regularized by the Tsallis entropy [2]. Thus, we show that TSO meets the definition of shrinkage estimator whose target is the identity matrix (tensor).

Our proposed method outperforms the state of the art on novel classes by 4.0%, 4.7% and 6.1% mAP on PASCAL VOC 2007, FSOD, and COCO respectively.

2 Related Works

Below, we review FSOD models and vision transformers, followed by a discussion on feature grouping, tensor descriptors and spectral power normalization.

Few-shot Object Detection. A Low-Shot Transfer Detector (LSTD) [4] leverages rich source domain to construct a target domain detector with few training samples but needs to be fine-tuned to novel categories. Meta-learning-based approach [45] reweights RoI features in the detection head without fine-tuning. Similarly, MPSR [43] deals with scale invariance by ensuring the detector is trained over multiple scales of positive samples. NP-RepMet [46] introduces a negative- and positive-representative learning framework via triplet losses that bootstrap the classifier. FSOD-ARPN [7] is a general FSOD network equipped with a channel-wise attention mechanism and multi-relation detector that scores pair-wise object similarity in both the RPN and the detection head, inspired by Faster R-CNN. PNSD [50], inspired by FSOD-ARPN [7], uses contraction of second-order autocorrelation matrix against query feature maps to produce attention maps. Single-prototype (per class) methods suffer information loss.

Per-sample Prototype FSOD [23] uses the entire support set to form prototypes of a class but it ignores spatial information within regions. Thus, we employ TENET RPN and TRH to capture spatial and high-order patterns, and extract correlations between the query image and the Z -shot support samples for a class.

Transformers in Vision. Transformers, popular in natural language processing [38], have also become popular in computer vision. Pioneering works such as ViT [9] show that transformers can achieve the state of the art in image recognition. DETR [3] is an end-to-end object detection framework with a transformer encoder-decoder used on top of backbone. Its deformable variant [52] improves the performance/training efficiency. SOFT [29], the SoftMax-free transformer approximates the self-attention kernel by replacing the SoftMax function with Radial Basis Function (RBF), achieving linear complexity. In contrast, our TENET is concerned with reversing the diffusion of signal in high-order tensors via the shrinkage operation, with the goal of modeling spatially invariant high-order statistics of regions. Our attention unit, so-called Spatial-HOP in TRH, also uses RBF to capture correlations of between spatial and high-order descriptors.

Multi-path and Groups of Feature Maps. GoogleNet [36] has shown that multi-path representations (several network branches) lead to classification improvements. ResNeXt [44] adopts group convolution [20] in the ResNet bottleneck block. SK-Net [25], based on SE-Net [10], uses feature map attention across two network branches. However, these approaches do not model feature statistics. ReDRO [30] samples groups of features to apply the matrix square root over submatrices to improve the computational speed. In contrast, TENET forms fixed groups of features to form second-, third- and fourth-order tensors (simply using groups of features to form second-order matrices is not effective).

Second-order Pooling (SOP). Region Covariance Descriptors for texture [37] and object recognition [17] use SOP. Approach [19] uses spectral pooling for fine-grained image classification, whereas SoSN [48] and its variants [47,49] leverage SOP and element-wise Power Normalization (PN) [17] for end-to-end few-shot learning. In contrast, we develop a multi-object few-shot detector. Similarly to SoSN, PNSD [50] uses SOP with PN as representations which are passed to the detection head. So-HoT [14] that uses high-order tensors for domain adaptation is also somewhat related to TENET but So-HoT uses multiple polynomial kernel matrices, whereas we apply TSO to achieve decorrelation and shrinkage. TENET without TSO reduces to polynomial feature maps and performs poorly.

Power Normalization (PN). Burstiness is “the property that a given visual element appears more times in an image than a statistically independent model would predict”. PN [16] limits the burstiness of first- and second-order statistics due to the binomial PMF-based feature detection factoring out feature counts [16,17,19]. Element-wise MaxExp pooling [16] gives likelihood of “at least one particular visual word being present in an image”, whereas SigmE pooling [19] is its practical approximation. Noteworthy are recent Fast Spectral MaxExp operator, MaxExp(F) [18], which reverses the heat diffusion on the underlying loop graph of second-order matrix to some desired past state [34], and Tensor Power-Euclidean (TPE) metric [15]. TPE alas uses the Higher Order Singular Value

Decomposition [21], which makes TPE intractable for millions of region proposals per dataset. Thus, we develop TENET and TSO, which reverses diffusion on high-order tensors by shrinking them towards the super-diagonal of tensor.

3 Background

Below, we detail notations and show how to calculate multiple higher-order statistics and Power Normalization, followed by revisiting the transformer block.

Notations. Let $\mathbf{x} \in \mathbb{R}^d$ be a d -dimensional feature vector. \mathcal{I}_N stands for the index set $\{1, 2, \dots, N\}$. We define a vector of all-ones as $\mathbf{1} = [1, \dots, 1]^T$. Let $\mathcal{X}^{(r)} = \uparrow \otimes_r \mathbf{x}$ denote a tensor of order r generated by the r -th order outer-product of \mathbf{x} , and $\mathcal{X}^{(r)} \in \mathbb{R}^{d \times d \times \dots \times d}$. Typically, capitalised boldface symbols such as Φ denote matrices, lowercase boldface symbols such as ϕ denote vectors and regular case such as $\Phi_{i,j}$, ϕ_i , n or Z denote scalars, *e.g.*, $\Phi_{i,j}$ is the (i, j) -th coefficient of ϕ .

High-order Tensor Descriptors (HoTD). Below we formalize HoTD [14].

Proposition 1. Let $\Phi \equiv [\phi_1, \dots, \phi_N] \in \mathbb{R}^{d \times N}$ and $\Phi' \equiv [\phi'_1, \dots, \phi'_M] \in \mathbb{R}^{d \times M}$ be feature vectors extracted from some two image regions. Let $\mathbf{w} \in \mathbb{R}_+^N$, $\mathbf{w}' \in \mathbb{R}_+^M$ be some non-negative weights and $\mu, \mu' \in \mathbb{R}^d$ be the mean vectors of Φ and Φ' , respectively. A linearization of the sum of polynomial kernels of degree r ,

$$\langle \mathcal{M}^{(r)}(\Phi; \mathbf{w}, \mu), \mathcal{M}^{(r)}(\Phi'; \mathbf{w}', \mu') \rangle = \frac{1}{NM} \sum_{n=1}^N \sum_{m=1}^M w_n^r w_m'^r \langle \phi_n - \mu, \phi'_m - \mu' \rangle^r, \quad (1)$$

yields the tensor feature map

$$\mathcal{M}^{(r)}(\Phi; \mathbf{w}, \mu) = \frac{1}{N} \sum_{n=1}^N w_n^r \uparrow \otimes_r (\phi_n - \mu) \in \mathbb{R}^{d \times d \times \dots \times d}. \quad (2)$$

In our paper, we set $\mathbf{w} = \mathbf{w}' = 1$ and $\mu = \mu' = 0$, whereas orders $r = 2, 3, 4$.

(Eigenvalue) Power Normalization ((E)PN). For second-order matrices, MaxExp(F), a state-of-the-art EPN [18], is defined as

$$g(\lambda; \eta) = 1 - (1 - \lambda)^\eta \quad (3)$$

on the ℓ_1 -norm normalized spectrum from SVD ($\lambda_i := \lambda_i / (\varepsilon + \sum_{i'} \lambda_{i'})$), and on symmetric positive semi-definite matrices as

$$\hat{\mathcal{G}}_{\text{MaxExp}}(\mathbf{M}; \eta) = \mathbf{I} - (\mathbf{I} - \mathbf{M})^\eta, \quad (4)$$

where \mathbf{M} is a trace-normalized matrix, *ie.*, $\mathbf{M} := \mathbf{M} / (\varepsilon + \text{Tr}(\mathbf{M}))$ with $\varepsilon \approx 1e-6$, and $\text{Tr}(\cdot)$ denotes the trace. The time-reversed heat diffusion process is adjusted by integer $\eta \geq 1$. The larger the value of η is, the more prominent the time reversal is. $\hat{\mathcal{G}}_{\text{MaxExp}}$ is followed by the element-wise PN, called SigmE [18]:

$$\mathcal{G}_{\text{SigmE}}(p; \eta') = 2 / (1 + e^{-\eta' p}) - 1, \quad (5)$$

where p takes each output coefficient of Eq. (4), $\eta' \geq 1$ controls detecting feature occurrence *vs.* feature counting trade-off.

Transformers. An architecture based on blocks of attention and MLP layers forms a transformer network [38]. Each attention layer takes as input a set of query, key and value matrices, denoted \mathbf{Q} , \mathbf{K} and \mathbf{V} , respectively. Let $\mathbf{Q} \equiv [\mathbf{q}_1, \dots, \mathbf{q}_N] \in \mathbb{R}^{d \times N}$, $\mathbf{K} \equiv [\mathbf{k}_1, \dots, \mathbf{k}_N] \in \mathbb{R}^{d \times N}$, and $\mathbf{V} \equiv [\mathbf{v}_1, \dots, \mathbf{v}_N] \in \mathbb{R}^{d \times N}$, where N is the number of input feature vectors, also called tokens, and d is the channel dimension. A generic attention layer can then be formulated as:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \alpha(\gamma(\mathbf{Q}, \mathbf{K}))\mathbf{V}^T. \quad (6)$$

Self-attention is composed of $\alpha(\gamma(\cdot, \cdot))$, $\alpha(\cdot)$ is a non-linearity and $\gamma(\cdot, \cdot)$ computes similarity. A popular choice is SoftMax with the scaled dot product [38]:

$$\gamma(\cdot) = \text{SoftMax}(\cdot) \quad \text{and} \quad \gamma(\mathbf{Q}, \mathbf{K}) = \frac{\mathbf{Q}^T \mathbf{K}}{\sqrt{d}}. \quad (7)$$

Note the LayerNorm and residual connections are added at the end of each block.

To facilitate the design of linear self-attention, approach [29] introduces a SoftMax-free transformer with the dot product replaced by the RBF kernel:

$$\gamma = \text{Exp}(\cdot) \quad \text{and} \quad \gamma(\mathbf{Q}, \mathbf{K}) = -\frac{1}{2\sigma^2} \left[\|\mathbf{q}_i - \mathbf{k}_j\|_2^2 \right]_{i,j \in \mathcal{I}_N}, \quad (8)$$

where $[\cdot]$ stacks distances into a matrix, σ^2 is the kernel variance set by cross-validation, and \mathbf{q}_i and \mathbf{k}_j are ℓ_2 -norm normalized.

The Multi-Head Attention (MHA) layer uses T attention units whose outputs are concatenated. Such an attention splits input matrices \mathbf{Q} , \mathbf{K} , and \mathbf{V} along their channel dimension d into T groups and computes attention on each group:

$$\text{MHA}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = [\mathbf{A}_1, \dots, \mathbf{A}_T], \quad (9)$$

where $[\cdot]$ performs concatenation along the channel mode, the m^{th} head is $\mathbf{A}_m = \text{Attention}(\mathbf{Q}_m, \mathbf{K}_m, \mathbf{V}_m)$, and $\mathbf{Q}_m \in \mathbb{R}^{\frac{d}{T} \times N}$, $\mathbf{K}_m \in \mathbb{R}^{\frac{d}{T} \times N}$, and $\mathbf{V}_m \in \mathbb{R}^{\frac{d}{T} \times N}$.

4 Proposed Approach

Given a set of Z support crops $\{\mathbf{X}_z\}_{z \in \mathcal{I}_Z}$ and a query image \mathbf{X}^* per episode, our approach learns a matching function between representations of query RoIs and support crops. Fig. 1 shows our pipeline comprised of three main modules:

- i. **Encoding Network (EN)** extracts feature map $\Phi \in \mathbb{R}^{d \times N}$ per image (of $N = W \times H$ spatial size) from query and support images via ResNet-50.
- ii. **TENET RPN** extracts RoIs from the query image and computes embeddings for the query RoIs and the support crops. Improved attention maps are obtained by T-Heads Attention (THA) that operates on TENET descriptors.
- iii. **Transformer Relation Head** captures relations between query and support features using self- and cross-attention mechanisms. This head produces representations for the classifier and bounding-box refinement regression loss.

Next, we describe TENET RPN, followed by the Transformer Relation Head.




Fig. 1: Our pipeline. (*top*) We pass ground truth support crops $\{X_z\}_{z \in \mathcal{I}_Z}$ for Z-shot problem and a query image X^* to the Encoding Network (EN). The resulting convolutional feature maps, $\{\Phi_z\}_{z \in \mathcal{I}_Z}$ for support and Φ^* for query, are passed to the TENET-RPN module to produce a set of B descriptors $\{\Phi_b^*\}_{b \in \mathcal{I}_B}$ for B Region Proposals (RP a.k.a. RoIs) of query image, and high-order pooled (HOP) representations for both support crops $\{\psi_z\}_{z \in \mathcal{I}_Z}$ and query image RoIs $\{\psi_b^*\}_{b \in \mathcal{I}_B}$. TENET contains HOP units which compute high-order tensor descriptors and then apply a novel tensor shrinkage operator to them, yielding spatially orderless HOP representations. Sets of features $\{\psi_z\}_{z \in \mathcal{I}_Z}$ and $\{\psi_b^*\}_{b \in \mathcal{I}_B}$ are then passed to the Transformer Relation Head (TRH), along with the convolutional features $\{\Phi_z\}_{z \in \mathcal{I}_Z}$ and $\{\Phi_b^*\}_{b \in \mathcal{I}_B}$. The TRH consists of Z-shot and Spatial-HOP transformer heads for measuring similarities across support regions and query proposals, and refining the localization of target objects, respectively. (*bottom*) Details of HOP, “Orderless FO, Spatial, Orderless HO”, Z-shot T-RH and Spatial-HOP T-RH blocks. See §4.3 for details.

4.1 Extracting representations for support and query RoIs

Fig. 1 (*top*) shows our TENET RPN that produces embeddings and query RoIs. Central to this module is our HOP unit that produces Higher-order Tensor Descriptors (HoTDs). HOP splits features along the channel mode into multiple groups of feature maps, from which second-, third- and fourth- order tensors are aggregated over desired regions. HoTDs use a generalization of the $\text{MaxExp}(F)$ to higher-order tensors, called the Tensor Shrinkage Operator (TSO).

Tensor Shrinkage Operator. Ledoit and Wolf [22] define autocorrelation/covariance matrix estimation as a trade-off between the sample matrix M and a highly structured operator F , using the linear combination $(1-\delta)M + \delta F$. For symmetric positive semi-definite matrices and tensors, one can devise a convex shrinkage operator by minimizing some divergence $d(\lambda, \lambda')$ between the source and target spectra, where λ' is regularized by $\Omega(\lambda')$ with weight $\delta \geq 0$:

$$\lambda^* = \arg \min_{\lambda' \geq 0} d(\lambda, \lambda') + \delta \Omega(\lambda'). \quad (10)$$

Below we derive the TSO as a generalization of MaxExp(F)¹.

With $\widehat{\mathcal{G}}_{\text{TSO}}(\cdot)$ we extract representations $\hat{\psi}_r$ from HoTDs $\mathcal{M}^{(r)}$:

$$\widehat{\mathcal{G}}_{\text{TSO}}(\mathcal{M}^{(r)}; \eta) = \mathcal{I}_r - (\mathcal{I}_r - \mathcal{M}^{(r)})^\eta, \quad (11)$$

$$\hat{\psi}_r = \text{Diag}(\widehat{\mathcal{G}}_{\text{TSO}}(\mathcal{M}^{(r)}; \eta_r)), \quad (12)$$

$$\psi_r = \mathcal{G}_{\text{SigmE}}(\hat{\psi}_r; \eta'_r), \quad (13)$$

where $\text{Diag}(\cdot)$ extracts the super-diagonal of tensor. The identity tensor \mathcal{I}_r of order r is defined such that $\mathcal{I}_{1,\dots,1} = \mathcal{I}_{2,\dots,2} \dots = \mathcal{I}_{d,\dots,d} = 1$ and $\mathcal{I}_{i_1,\dots,i_r} = 0$ if $i_j \neq i_k$ and $j \neq k$, $j, k \in \mathcal{I}_r$.

Theorem 1. Let λ be the ℓ_1 -norm normalized spectrum, as in Eq. (3). Then $g(\lambda; \eta) = 1 - (1 - \lambda)^\eta$ is a shrinkage operator as it is a solution to Eq. (10) if $\delta = \frac{\eta t^{1/\eta}}{s}(1 - \frac{1}{\eta})$, and the Kullback-Leibler divergence and the Tsallis entropy are substituted for the divergence term d and the regularization term Ω , respectively. In the above theorem, $\lambda \in \mathbb{R}_+^d$, $s = d - 1$, $t = d - t'$ and $t' > 0$ is the trace of $g(\lambda; \eta)$.

Proof. Let $d(\lambda, \lambda') = D_{\text{KL}}(\lambda^\circ \| \lambda'^\circ) = -(\sum_{i \in \mathcal{I}_d} \lambda_i^\circ \log \lambda_i'^\circ) + (\sum_{i \in \mathcal{I}_d} \lambda_i^\circ \log \lambda_i^\circ)$, where $\lambda^\circ = \frac{1-\lambda}{s}$ and $\lambda'^\circ = \frac{1-\lambda'}{t}$ are complements of λ and λ' normalized by $s = d - 1$ and $t = d - t'$ to obtain normalized distributions to ensure valid use of the Kullback-Leibler divergence and the Tsallis entropy. Let $\Omega(\lambda'; \alpha) = \frac{1}{\alpha-1}(1 - \sum_{i \in \mathcal{I}_d} \lambda_i'^{\circ/\alpha})$. Define $f(\lambda, \lambda') = -(\sum_{i \in \mathcal{I}_d} \lambda_i^\circ \log \lambda_i'^\circ) + (\sum_{i \in \mathcal{I}_d} \lambda_i^\circ \log \lambda_i^\circ) + \frac{1}{\alpha-1}(1 - \sum_{i \in \mathcal{I}_d} \lambda_i'^{\circ/\alpha})$ (following Eq. (10)) which we minimize w.r.t. λ' by computing $\frac{\partial f}{\partial \lambda'_i} = 0$, that is,

$$\frac{\partial f}{\partial \lambda'_i} = 0 \Rightarrow \lambda'_i = 1 - t \left(\frac{\eta}{\delta s} \right)^\eta (1 - 1/\eta)^\eta (1 - \lambda_i)^\eta, \quad (14)$$

where $\frac{1}{\alpha} = \eta$. Solving $t \left(\frac{\eta}{\delta s} \right)^\eta (1 - \frac{1}{\eta})^\eta = 1$ for δ completes the proof.

Theorem 2. The highly structured operator \mathbf{F} (target of shrinkage) equals \mathbb{I} .

Proof. Notice $\lim_{\eta \rightarrow \infty} 1 - (1 - \lambda_i)^\eta = 1$ if $\lambda \neq 0$ is the ℓ_1 -norm normalized spectrum from SVD, i.e., $\mathbf{U} \lambda \mathbf{U}^T = \mathbf{M} \succ 0$ with $\lambda_i := \lambda_i / (\sum_{i'} \lambda_{i'} + \varepsilon)$, $\varepsilon > 0$. Thus, $\mathbf{U} \mathbf{U}^T = \mathbb{I}$.

Note $\mathcal{I}_r - (\mathcal{I}_r - \mathcal{M})^\eta \rightarrow \mathcal{I}_r$ if $\eta \rightarrow \infty$. Thus, for sufficiently large $1 \ll \eta \ll \infty$, the diffused heat reverses towards the super-diagonal, where the majority of the signal should concentrate. Thus, we limit the number of coefficients of feature representations by extracting the super-diagonals from the TSO-processed $\mathcal{M}^{(r)}$ as in Eq. 12, where r is the order of HoTD \mathcal{M} . In our experiments, $r = 2, 3, 4$.

¹ For $r = 2$, Eq. (12) yields $\text{Diag}(\mathbb{I} - (\mathbb{I} - \mathbf{M})^{\eta_2})$. $\text{Diag}(\text{Sqrtm}(\mathbf{M}))$ is its approximation.

Algorithm 1 Tensor Shrinkage Operator with Exponentiation by Squaring, left part for even orders and right part for odd orders r .

Input: \mathcal{M} , $\eta \geq 1$, $r = 2, 4, \dots$

```

1:  $\mathcal{M}_1^* = \mathcal{I}_r - \mathcal{M}$ ,  $n = \text{int}(\eta)$ ,  $t = 1$ ,  $q = 1$ 
2: while  $n \neq 0$ :
3:   if  $n \& 1$ :
4:     if  $t > 1$ :  $\mathcal{G}_{t+1} = \mathcal{G}_t \times_{1, \dots, r/2} \mathcal{M}_q^*$ ,
5:     else:  $\mathcal{G}_{t+1} = \mathcal{M}_q^*$ 
6:      $n \leftarrow n - 1$ ,  $t \leftarrow t + 1$ 
7:    $n \leftarrow \text{int}(n/2)$ 
8:   if  $n > 0$ :
9:      $\mathcal{M}_{q+1}^* = \mathcal{M}_q^* \times_{1, \dots, r/2} \mathcal{M}_q^*$ 
10:     $q \leftarrow q + 1$ 

```

Output: $\hat{\mathcal{G}}_{\text{TSO}}(\mathcal{M}) = \mathcal{I}_r - \mathcal{G}_t$

Input: \mathcal{M} , $\eta = 3^0, 3^1, 3^2, \dots$, $r = 3, 5, \dots$

```

1:  $\mathcal{M}_1^* = \mathcal{I}_r - \mathcal{M}$ ,  $n = \text{int}(\eta)$ ,  $q = 1$ 
2: while  $n \neq 0$ :
3:    $n \leftarrow \text{int}(n/3)$ 
4:   if  $n > 0$ :
5:      $\mathcal{M}_{q+1}^* = \mathcal{M}_q^* \times_{1, \dots, \lfloor r/2 \rfloor} \mathcal{M}_q^*$ 
6:      $\times_{1, \dots, \lceil r/2 \rceil} \mathcal{M}_q^*$ 
7:      $q \leftarrow q + 1$ 

```

Output: $\hat{\mathcal{G}}_{\text{TSO}}(\mathcal{M}) = \mathcal{I}_r - \mathcal{M}_q^*$

As super-diagonals contain information obtained by multiplying η times the complement $1 - \lambda_i$ (spectral domain), TSO can be seen as $\eta - 1$ aggregation steps along the tensor product mode(s). Thus, we pass $\hat{\psi}_r$ via the element-wise SigmE from Eq. (5) (as in Eq. (13)) to detect the presence of at least one feature being detected in $\hat{\psi}_r$ after such an aggregation. For brevity, we drop subscript r of ψ_r .

Complexity. For integers $\eta \geq 2$ and even orders $r \geq 2$, computing $\eta - 1$ tensor-tensor multiplications $(\mathcal{I}_r - \mathcal{M}^{(r)})^\eta$ has the complexity $\mathcal{O}(d^{\frac{3}{2}r}\eta)$. For odd orders $r \geq 3$, due to alternations between multiplications in $\lfloor \frac{r}{2} \rfloor$ and $\lceil \frac{r}{2} \rceil$ modes, the complexity is $\mathcal{O}(d^{\lfloor \frac{r}{2} \rfloor} d^{2\lceil \frac{r}{2} \rceil} \eta) \approx \mathcal{O}(d^{\frac{3}{2}r}\eta)$. Thus, the complexity of Eq. (11) w.r.t. integer $\eta \geq 2$ scales linearly. However, for even orders r , one can readily replace $(\mathcal{I}_r - \mathcal{M}^{(r)})^\eta$ with exponentiation by squaring [1], whose cost is $\log(\eta)$. This readily yields the sublinear complexity $\mathcal{O}(d^{\frac{3}{2}r} \log(\eta))$ w.r.t. η .

Implementation of TSO. Algorithm 1 shows fast TSO for even/odd orders r . We restrict the odd variant to $r = 3^0, 3^1, 3^2, \dots$ for brevity. Finally, we note that matrix-matrix and tensor multiplications with cuBLAS are highly parallelizable so the $d^{\frac{3}{2}r}$ part of complexity can be reduced in theory even to $\log(d)$.

TENET-RPN. To extract query RoIs, we firstly generate a set \mathcal{Z} of TSO representations $\Psi \equiv \{\psi_z\}_{z \in \mathcal{I}_Z}$ from the support images, with $Z = |\mathcal{Z}|$. We then perform cross-attention between $\Psi \in \mathbb{R}^{d \times Z}$ and the N feature vectors $\Phi^* \in \mathbb{R}^{d \times N}$ extracted from the query image. The attention input Q is then generated from Φ^* , while K and V are both generated from Ψ . The output of the transformer block in Eq. 6 is fed into an RPN layer to output a set \mathcal{B} of $B = |\mathcal{B}|$ query RoIs.

Query RoIs are represented by spatially ordered features $\{\Phi_b^*\}_{b \in \mathcal{I}_B}$ and TSO orderless features $\{\psi_b^*\}_{b \in \mathcal{I}_B}$, passed to the Transformer Relation Head.

4.2 Transformer Relation Head

TRH, in Fig. 1, models relations between support crops and query RoIs. TSO representations are derived from features of layer 4 of ResNet-50, leading to a

channel dimension $d = 1024$. Spatially ordered representations of support features $\{\bar{\Phi}_z \in \mathbb{R}^{2d \times N}\}_{z \in \mathcal{I}_Z}$ and query RoIs $\{\bar{\Phi}_b^* \in \mathbb{R}^{2d \times N}\}_{b \in \mathcal{I}_B}$ are both extracted from layer 5 of ResNet-50, leading to a channel dimension $2d = 2048$.

They are then fed into 2 different transformers: (i) a **Z-shot transformer head**, which performs cross-attention between globally-pooled representations of the query images and support images; and (ii) a **Spatial-HOP transformer head**, which performs self-attention between spatially ordered representations and spatially orderless high-order representations for a given image.

Z-shot transformer head consists of a cross-attention layer formed with:

$$\begin{pmatrix} \mathbf{Q} \\ \mathbf{K} \\ \mathbf{V} \end{pmatrix} = \begin{pmatrix} [\mathbf{q}_1, \dots, \mathbf{q}_B] \\ [\mathbf{k}_1, \dots, \mathbf{k}_Z] \\ [\mathbf{v}_1, \dots, \mathbf{v}_Z] \end{pmatrix} \quad \text{where} \quad \begin{pmatrix} \mathbf{q}_b \\ \mathbf{k}_z \\ \mathbf{v}_z \end{pmatrix} = \begin{pmatrix} \mathbf{W}^{(q)}(\bar{\Phi}_b^* + \mathbf{W}^{(p)}\psi_b^*) \\ \mathbf{W}^{(k)}(\bar{\Phi}_z + \mathbf{W}^{(p)}\psi_z) \\ \mathbf{W}^{(v)}(\bar{\Phi}_z + \mathbf{W}^{(p)}\psi_z) \end{pmatrix}. \quad (15)$$

Moreover, $\bar{\Phi}_b^*$ and $\bar{\Phi}_z$ are average-pooled features $\frac{1}{N}\bar{\Phi}_b^*\mathbf{1}$ and $\frac{1}{N}\bar{\Phi}_z\mathbf{1}$, respectively. The matrices $\mathbf{W}^{(q)} \in \mathbb{R}^{2d \times 2d}$, $\mathbf{W}^{(k)} \in \mathbb{R}^{2d \times 2d}$, $\mathbf{W}^{(v)} \in \mathbb{R}^{2d \times 2d}$, and $\mathbf{W}^{(p)} \in \mathbb{R}^{2d \times d}$ is a linear projection mixing spatially orderless TSO representations with spatially ordered representations. Thus, each attention query vector \mathbf{q}_b combines the extracted spatially orderless TSO representations with spatially ordered representations for a given query RoI. Similarly, each key vector \mathbf{k}_z and value vector \mathbf{v}_z combine such two types of representations for a support crop.

Spatial-HOP transformer head consists of a layer that performs self-attention on spatially orderless TSO representations and spatially ordered representations, extracted either from Z support crops, or B query RoIs. Below we take one support crop as an example. For the set \mathcal{Z} , we compute ‘‘Spatial’’, a spatially ordered Z -averaged representation $\Phi^\dagger = \frac{1}{Z} \sum_{z \in \mathcal{I}_Z} \Phi_z \in \mathbb{R}^{2d \times N}$. We also compute HO, a spatially orderless High-Order Z -pooled representation $\psi^\dagger = \frac{1}{Z} \sum_{z \in \mathcal{I}_Z} \psi_z \in \mathbb{R}^d$. We split Φ^\dagger along the channel mode of dimension $2d$ to create two new matrices $\Phi^{\dagger u} \in \mathbb{R}^{d \times N}$ and $\Phi^{\dagger l} \in \mathbb{R}^{d \times N}$. We set $\Phi^{\dagger l} = [\phi_1^{\dagger l}, \dots, \phi_N^{\dagger l}] \in \mathbb{R}^d$. Also, we form FO (from $\Phi^{\dagger u}$), a spatially orderless First-Order Z -pooled representation. Self-attention is then performed over the matrix of token vectors:

$$\mathcal{T} = [\phi_1^{\dagger l}, \dots, \phi_N^{\dagger l}, \bar{\Phi}^{\dagger u}, \mathbf{W}^{(g)}\psi^\dagger], \quad (16)$$

where $\mathbf{W}^{(g)} \in \mathbb{R}^{d \times d}$ a linear projection for HO. \mathcal{T} is projected onto the query, key and value linear projections. Spatial-HOP attention captures relations among spatially-aware first-order and spatially orderless high-order representations. The outputs of Z -shot and Spatial-HOP transformer heads are combined and fed into the classifier/bounding-box regressor. See §B of **Suppl. Material** for details.

4.3 Pipeline Details (Figure 1 (bottom))

HOP unit uses \odot to split the channel mode into groups (*e.g.*, 2:1:1 split means two parts of the channel dimension are used to form $\mathcal{M}^{(2)}$, one part to form $\mathcal{M}^{(3)}$, and one part to form $\mathcal{M}^{(4)}$). TSOs with parameters $\eta_2 = \dots = \eta_r = \eta$ are applied for orders $r = 2, \dots, r$, and diagonal entries are extracted from each tensor and concatenated by \odot . Element-wise SigmE with $\eta'_2 = \dots = \eta'_r = \eta'$ is applied.

“Orderless FO, Spatial, Orderless HO” block combines the First-Order (FO), spatial and High-Order (HO) representations. Operator “Z-avg” performs average pooling along Z -way mode, operator “Sp-avg” performs average pooling along the spatial modes of feature maps, operator \odot simply splits the channel mode into two equally sized groups (each is half of the channel dimension), and operator \otimes performs concatenation of FO, spatial and HO representations along the spatial mode of feature maps, *e.g.*, we obtain $N+2$ fibers times 1024 channels.

Z-shot T-RH is a transformer which performs attention on individual Z -shots. The spatial representation Φ is average pooled along spatial dimensions by “Sp-avg” and combined with high-order Ψ (passed by a FC layer) via addition \oplus . Another FC layer follows and subsequently the value, key and query matrices are computed, an RBF attention formed. Operator \bullet multiplies the value matrix with the RBF matrix. Head is repeated T times, outputs concatenated by \odot .

Spatial-HOP T-RH takes inputs from the “Orderless FO, Spatial, Orderless HO” block, and computes the value, key and query matrices. The attention matrix (RBF kernel) has $(N+2) \times (N+2)$ size, being composed of spatial, FO-spatial and HO-spatial attention. After multiplying the attention matrix with the value matrix, we extract the spatial, FO and HO representations. Support and query first-order representations (FO) (and high-order representations (HO)) are element-wisely multiplied by \bullet (multiplicative relationship). Finally, support and query spatial representations use the subtraction operator \ominus . After the concatenation of FO and HO relational representations by \odot , passing via an MLP (FC+ReLU+ FC), and concatenation with the spatial relational representations, we get one attention block, repeated T times.

5 Experiments

Datasets and Settings . For PASCAL VOC 2007/12 [6], we adopt the 15/5 base/novel category split setting and use training/validation sets from PASCAL VOC 2007 and 2012 for training, and the testing set from PASCAL VOC 2007 for testing, following [11,7,50,23]. For MS COCO [28], we follow [45], and adopt the 20 categories that overlap with PASCAL VOC as the novel categories for testing, whereas the remaining 60 categories are used for training. For the FSOD dataset [7], we split its 1000 categories into 800/200 for training/testing.

Implementation Details. TENET uses ResNet-50 pre-trained on ImageNet [5] and MS COCO [28]. We fine-tune the network with a learning rate of 0.002 for the first 56000 iterations and 0.0002 for another 4000 iterations. Images are resized to 600 pixels (shorter edge) and the longer edge is capped at 1000 pixels. Each support image is cropped based on ground-truth boxes, bilinearly interpolated and padded to 320×320 pixels. We set, via cross-validation) SigmE parameter $\eta' = 200$ and TSO parameters $\eta_2 = \eta_3 = \eta_4 = 7$. We report standard metrics for FSOD, namely mAP , AP , AP_{50} and AP_{75} .

Table 1: Evaluations (mAP %) on three splits of the VOC 2007 testing set.

Method/Shot		Split 1				Split 2				Split 3				Mean			
		1	3	5	10	1	3	5	10	1	3	5	10	1	3	5	10
FRCN	ICCV12	11.9	29.0	36.9	36.9	5.9	23.4	29.1	28.8	5.0	18.1	30.8	43.4	7.6	23.5	32.3	36.4
FR	ICCV19	14.8	26.7	33.9	47.2	15.7	22.7	30.1	39.2	19.2	25.7	40.6	41.3	16.6	25.0	34.9	42.6
Meta	ICCV19	19.9	35.0	45.7	51.5	10.4	29.6	34.8	45.4	14.3	27.5	41.2	48.1	14.9	30.7	40.6	48.3
FSOD	CVPR20	29.8	36.3	48.4	53.6	22.2	25.2	31.2	39.7	24.3	34.4	47.1	50.4	25.4	32.0	42.2	47.9
NP-RepMet	NeurIPS20	37.8	41.7	47.3	49.4	41.6	43.4	47.4	49.1	33.3	39.8	41.5	44.8	37.6	41.6	45.4	47.8
PNSD	ACCV20	32.4	39.6	50.2	55.1	30.2	30.3	36.4	42.3	30.8	38.6	46.9	52.4	31.3	36.2	44.5	49.9
MPSR	ECCV20	41.7	51.4	55.2	61.8	24.4	39.2	39.9	47.8	35.6	42.3	48.0	49.7	33.9	44.3	47.7	53.1
TFA	ICML20	39.8	44.7	55.7	56.0	23.5	34.1	35.1	39.1	30.8	42.8	49.5	49.8	31.4	40.5	46.8	48.3
FSCE	CVPR21	44.2	51.4	61.9	63.4	27.3	43.5	44.2	50.2	22.6	39.5	47.3	54.0	31.4	44.8	51.1	55.9
CGDP+FRCN	CVPR21	40.7	46.5	57.4	62.4	27.3	40.8	42.7	46.3	31.2	43.7	50.1	55.6	33.1	43.7	50.0	54.8
TIP	CVPR21	27.7	43.3	50.2	56.6	22.7	33.8	40.9	46.9	21.7	38.1	44.5	50.9	24.0	38.4	45.2	52.5
FSOD ^{up}	ICCV21	43.8	50.3	55.4	61.7	31.2	41.2	44.2	48.3	35.5	43.9	50.6	53.5	36.8	45.1	50.1	54.5
QSAM	WACV22	31.1	39.2	50.7	59.4	22.9	32.1	35.4	42.7	24.3	35.0	50.0	53.6	26.1	35.4	45.4	51.9
TENET	(Ours)	46.7	55.4	62.3	66.9	40.3	44.7	49.3	52.1	35.5	46.0	54.4	54.6	40.8	48.7	55.3	57.9

Table 2: Evaluations on the MS COCO minival set (2a) and FSOD testset (2b).

Shot	Method		AP	AP ₅₀	AP ₇₅	Shot	Method		AP ₅₀	AP ₇₅
10	LSTD	AAAI18	3.2	8.1	2.1	5	LSTD (FRN)	AAAI18	23.0	12.9
	FR	ICCV12	5.6	12.3	4.6		LSTD	AAAI18	24.2	13.5
	Meta	ICCV19	8.7	19.18	6.6		FSOD	CVPR20	27.5	19.4
	MPSR	ECCV20	9.8	17.9	9.7		PNSD	ACCV20	29.8	22.6
	FSOD	CVPR20	11.1	20.4	10.6		QSAM	WACV22	30.7	25.9
	PNSD	ACCV20	12.3	21.7	11.7					
	TFA	ICML20	9.6	10.0	9.3					
	FSCE	CVPR21	10.7	11.9	10.5					
	CGDP+FRCN	CVPR21	11.3	20.3	11.5					
	FSOD ^{up}	ICCV21	11.6	23.9	9.8					
QSAM	WACV22	13.0	24.7	12.1						
TENET		(Ours)	19.1	27.4	19.6	TENET		(Ours)	35.4	31.6

(a)

(b)

5.1 Comparisons with the State of the Art

PASCAL VOC 2007/12. We compare our method to QSAM[23], FSOD^{up} [42], CGDP+FRCN [26], TIP [24], FSCE [35], TFA [39], Feature Reweighting (FR) [11], LSTD [4], FRCN [33], NP-RepMet [46], MPSR[43], PSND [50] and FSOD [7]. Table 1 shows that our TENET outperforms FSOD by a 7.1–15.4% margin. For the 1- and 10-shot regime, we outperform QSAM [23] by ~14.7%.

MS COCO. Table 2a compares TENET with QSAM[23], FSOD^{up} [42], CGDP+FRCN [26], TIP [24], FSCE [35], TFA [39], FR [11], Meta R-CNN [45], FSOD [7] and PNSD[50] on the MS COCO minival set (20 novel categories, 10-shot protocol). Although MS COCO is more challenging in terms of complexity and the dataset size, TENET boosts results to 19.1%, 27.4% and 19.6%, surpassing the SOTA method QSAM by 6.1%, 2.7% and 7.5% on AP , AP_{50} and AP_{75} .

FSOD. Table 2b compares TENET (5-shot) with PNSD [50], FSOD [7], LSTD [4] and LSTD (FRN [33]). We re-implement BD&TK, modules of LSTD, based on Faster-RCNN for fairness. TENET yields SOTA 35.4% AP_{50} and 31.6% AP_{75} .

Table 3: Results on VOC2007 testset for applying TENET in RPN or TRH (3a, top panel of 3b). TRH ablation shown in bottom panel of 3b.

r		dim. split	Shot(Novel)		Shot(Base)		Speed
2	3	4	5	10	5	10	(img/ms)
✓			56.5	64.2	71.7	75.5	32
	✓		55.7	63.2	67.0	72.1	69
		✓	51.4	58.9	68.7	74.8	78
✓	✓	3:1	58.3	63.2	69.3	75.1	42
✓	✓	3:1	56.1	62.4	70.8	75.4	68
✓	✓	2:2	51.8	61.7	68.1	73.6	71
		6:1:1	53.6	62.7	69.4	72.8	
		5:2:1	62.3	66.9	73.8	77.9	59
		5:1:2	53.9	63.1	69.7	73.3	
✓	✓	4:2:2	61.4	65.0	70.4	74.9	
		4:3:1	59.1	63.6	71.8	75.2	
		4:1:3	61.0	64.1	68.9	72.5	

(a)

	RPN	TRH	Shot(Novel)		Shot(Base)	
	r		5	10	5	10
a	1	1	53.4	61.8	64.9	72.1
b	2,3,4	1	57.2	63.7	68.8	76.6
c	2,3,4	2,3,4	61.0	65.4	71.3	77.3
d	2,3,4	1,2,3,4	62.3	66.9	73.8	78.2

TRH		Shot(Novel)		Shot(Base)	
Z-shot	Spatial-HOP	5	10	5	10
✓		58.5	63.2	69.3	75.1
	✓	61.0	65.8	71.7	76.5
✓	✓	62.3	66.9	73.8	78.2

(b)

5.2 Hyper-parameter and ablation analysis

TENET. Table 3a shows that among orders $r=2$, $r=3$ and $r=4$, variant $r=2$ is the best. We next consider pairs of orders, and the triplet $r=2, 3, 4$. As the number of tensor coefficients grows quickly w.r.t. r , we split the 1024 channels into groups, *e.g.*, $r=2, 3$. A 3:1 split means that second- and third-order tensors are built from 768 and 256 channels ($768+256=1024$). We report only the best splits. For pairs of orders, variant $r=2, 3$ was the best. Triplet $r=2, 3, 4$, the best performer, outperforms $r=2$ by 5.8% and 2.7% in novel classes (5- and 10-shot), and 2.1% and 2.4% in base classes. As all representations are 1024-dimensional, we conclude that multi-order variants are the most informative.

TSO. Based on the best channel-wise splits in Table 3a, we study the impact of η_r (shrinkage/decorrelation) of TSO to verify its effectiveness. Figure 2a shows mAP w.r.t. the individual η_2, η_3 and η_4 for $r=2$, $r=3$ and $r=4$. We then investigate the impact of η_r on pairwise representations, where we set the same η_r for pairwise variants, *e.g.*, $\eta_2 = \eta_3$. Again, the same value of η_r is used for triplet $r=2, 3, 4$. Note that for $\eta_r=1$, TSO is switched off and all representations reduce to the polynomial feature maps in So-HoT [14]. As shown in Figure 2a, TSO is very beneficial ($\sim 5\%$ gain for triplet $r=2, 3, 4$ over not using TSO).

TRH. Below we investigate the impact of Z-shot and Spatial-HOP T-RHs on results. Table 3b (bottom) shows that both heads are highly complementary.

Other hyperparameters. We start by varying σ of the RBF kernel from 0.3 to 3. Fig 2b shows that $\sigma=0.5$ gives the best result. We now fix σ and investigate the impact of varying the number of heads used in T-Heads Attention (TA). Table 4a shows best performance with $TA=4$. Lastly, we vary the number of TENET blocks (TB). Table 4b shows that results are stable especially if $TB \geq 2$.

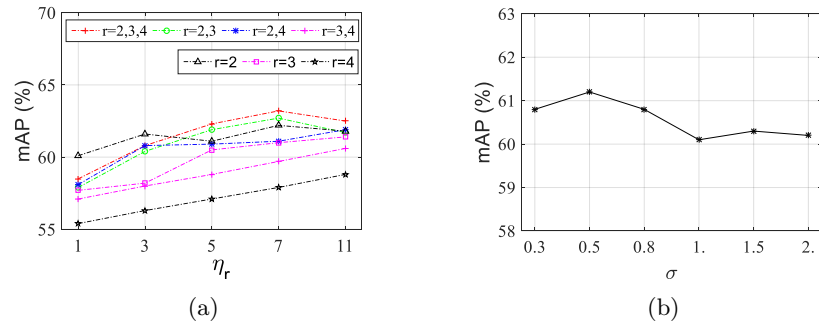


Fig. 2: mAP (VOC2007 dataset, novel classes, 10-shot) w.r.t. varying η_r in TSO (Fig. 2a) and the σ of RBF kernel in self-attention (Fig. 2b).

Table 4: Effect of varying (a) group within MHA in Tab. 4a and (b) TENET block in Tab. 4b on PASCAL VOC 2007 (5/10-shot, novel classes). When varying TENET block, group number is fixed to 4 (best value).

TA	1	2	4	8	16	32	64	
$\text{Shot}_{(\text{Novel})}$	5	58.3	59.1	61.8	60.5	58.4	58.5	56.0
	10	61.2	62.8	65.8	64.2	61.2	61.7	60.4

(a)

TB	1	2	3	4	5	
$\text{Shot}_{(\text{Novel})}$	5	61.8	62.3	62.1	61.8	61.9
	10	65.8	66.9	66.4	66.1	66.4

(b)

Unless otherwise noted, $TA = 2$ and $TB = 4$, respectively, on VOC dataset. See §C of **Suppl. Material** for more results on FSOD and MS COCO.

Impact of TENET on RPN and TRH. Table 3b shows ablations w.r.t. TENET variants in: 1) either RPN or TRH, or 2) both RPN and TRH. Comparing results for settings a,b, and c confirms that using second-, third- and fourth-orders simultaneously benefits both RPN and TRH, achieving 3.8%/1.9% as well as 3.8%/1.8% improvement on novel classes over the first-order-only variant. Results for settings c and d show that TRH encodes better the information carried within regions if leveraging both first- and higher-order representations.

6 Conclusions

We have proposed TENET, which uses higher-order tensor descriptors, in combination with a novel Tensor Shrinkage Operator, to generate highly-discriminative representations with tractable dimensionality. We use these representations in our proposed Transformer Relation Head to dynamically extract correlations between query image regions and support crops. TENET has heightened robustness to large intra-class variations, leading to SOTA performance on all benchmarks.

References

1. Exponentiation by squaring. Wikipedia, https://en.wikipedia.org/wiki/Exponentiation_by_squaring, accessed: 12-03-2021 **9**
2. Tsallis entropy. Wikipedia, https://en.wikipedia.org/wiki/Tsallis_entropy, accessed: 12-03-2021 **3**
3. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: ECCV 2020. Lecture Notes in Computer Science, vol. 12346, pp. 213–229. Springer (2020) **4**
4. Chen, H., Wang, Y., Wang, G., Qiao, Y.: LSTD: A low-shot transfer detector for object detection. In: McIlraith, S.A., Weinberger, K.Q. (eds.) Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018. pp. 2836–2843. AAAI Press (2018) **1, 3, 12**
5. Deng, J., Dong, W., Socher, R., Li, L., Li, K., Li, F.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20-25 June 2009, Miami, Florida, USA. pp. 248–255. IEEE Computer Society (2009). <https://doi.org/10.1109/CVPR.2009.5206848> **11, 19**
6. Everingham, M., Gool, L.V., Williams, C.K.I., Winn, J.M., Zisserman, A.: The pascal visual object classes (VOC) challenge. *Int. J. Comput. Vis.* **88**(2), 303–338 (2010). <https://doi.org/10.1007/s11263-009-0275-4> **11**
7. Fan, Q., Zhuo, W., Tai, Y.: Few-shot object detection with attention-rpn and multi-relation detector. *CoRR* **abs/1908.01998** (2019) **1, 2, 3, 11, 12**
8. Girshick, R.B.: Fast R-CNN. In: 2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015. pp. 1440–1448. IEEE Computer Society (2015). <https://doi.org/10.1109/ICCV.2015.169> **1**
9. Hu, H., Zhang, Z., Xie, Z., Lin, S.: Local relation networks for image recognition. In: ICCV 2019. pp. 3463–3472. IEEE (2019). <https://doi.org/10.1109/ICCV.2019.00356> **2, 4**
10. Hu, J., Shen, L., Albanie, S., Sun, G., Wu, E.: Squeeze-and-excitation networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **42**(8), 2011–2023 (2020). <https://doi.org/10.1109/TPAMI.2019.2913372> **4**
11. Kang, B., Liu, Z., Wang, X., Yu, F., Feng, J., Darrell, T.: Few-shot object detection via feature reweighting. In: 2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019. pp. 8419–8428. IEEE (2019). <https://doi.org/10.1109/ICCV.2019.00851> **1, 11, 12**
12. Karlinsky, L., Shtok, J., Harary, S., Schwartz, E., Aides, A., Feris, R.S., Giryes, R., Bronstein, A.M.: Repmet: Representative-based metric learning for classification and few-shot object detection. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019. pp. 5197–5206. Computer Vision Foundation / IEEE (2019). <https://doi.org/10.1109/CVPR.2019.00534> **1**
13. Kong, T., Yao, A., Chen, Y., Sun, F.: Hypernet: Towards accurate region proposal generation and joint object detection. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016. pp. 845–853. IEEE Computer Society (2016). <https://doi.org/10.1109/CVPR.2016.98> **1**

14. Koniusz, P., Tas, Y., Porikli, F.: Domain adaptation by mixture of alignments of second-or higher-order scatter tensors. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017. pp. 7139–7148. IEEE Computer Society (2017). <https://doi.org/10.1109/CVPR.2017.755> 4, 5, 13
15. Koniusz, P., Wang, L., Cherian, A.: Tensor representations for action recognition. TPAMI (2020) 2, 4
16. Koniusz, P., Yan, F., Gosselin, P.H., Mikolajczyk, K.: Higher-order occurrence pooling on mid-and low-level features: Visual concept detection. Tech. Report (2013) 4
17. Koniusz, P., Yan, F., Gosselin, P., Mikolajczyk, K.: Higher-order occurrence pooling for bags-of-words: Visual concept detection. IEEE Trans. Pattern Anal. Mach. Intell. **39**(2), 313–326 (2017). <https://doi.org/10.1109/TPAMI.2016.2545667> 2, 4
18. Koniusz, P., Zhang, H.: Power normalizations in fine-grained image, few-shot image and graph classification. TPAMI (2020) 2, 3, 4, 5
19. Koniusz, P., Zhang, H., Porikli, F.: A deeper look at power normalizations. In: 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018. pp. 5774–5783. IEEE Computer Society (2018). <https://doi.org/10.1109/CVPR.2018.00605> 4
20. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. Commun. ACM **60**(6), 84–90 (2017). <https://doi.org/10.1145/3065386> 4
21. Lathauwer, L.D., Moor, B.D., Vandewalle, J.: A multilinear singular value decomposition. SIAM J. Matrix Analysis and Applications **21**, 1253–1278 (2000) 5
22. Ledoit, O., Wolf, M.: Honey, i shrunk the sample covariance matrix. The Journal of Portfolio Management **30**(4), 110–119 (2004). <https://doi.org/10.3905/jpm.2004.110> 2, 7
23. Lee, H., Lee, M., Kwak, N.: Few-shot object detection by attending to per-sample-prototype. In: WACV, 2022, Waikoloa, HI, USA, January 3-8, 2022. pp. 1101–1110. IEEE (2022). <https://doi.org/10.1109/WACV51458.2022.00117> 2, 4, 11, 12
24. Li, A., Li, Z.: Transformation invariant few-shot object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3094–3102 (2021) 12
25. Li, X., Wang, W., Hu, X., Yang, J.: Selective kernel networks. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019. pp. 510–519. Computer Vision Foundation / IEEE (2019). <https://doi.org/10.1109/CVPR.2019.00060> 4
26. Li, Y., Zhu, H., Cheng, Y., Wang, W., Teo, C.S., Xiang, C., Vadakkepat, P., Lee, T.H.: Few-shot object detection via classification refinement and distractor re-treatment. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15395–15403 (2021) 12
27. Lin, T., Dollár, P., Girshick, R.B., He, K., Hariharan, B., Belongie, S.J.: Feature pyramid networks for object detection. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017. pp. 936–944. IEEE Computer Society (2017). <https://doi.org/10.1109/CVPR.2017.106> 1
28. Lin, T., Maire, M., Belongie, S.J., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: common objects in context. In: Fleet, D.J., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part

- V. Lecture Notes in Computer Science, vol. 8693, pp. 740–755. Springer (2014). https://doi.org/10.1007/978-3-319-10602-1_48 11
29. Lu, J., Yao, J., Zhang, J., Zhu, X., Xu, H., Gao, W., Xu, C., Xiang, T., Zhang, L.: SOFT: softmax-free transformer with linear complexity. CoRR **abs/2110.11945** (2021) 4, 6
30. Rahman, S., Wang, L., Sun, C., Zhou, L.: Redro: Efficiently learning large-sized spd visual representation. In: European Conference on Computer Vision (2020) 4
31. Redmon, J., Farhadi, A.: YOLO9000: better, faster, stronger. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21–26, 2017. pp. 6517–6525. IEEE Computer Society (2017). <https://doi.org/10.1109/CVPR.2017.690> 1
32. Redmon, J., Farhadi, A.: Yolov3: An incremental improvement. CoRR **abs/1804.02767** (2018) 1
33. Ren, S., He, K., Girshick, R.B., Sun, J.: Faster R-CNN: towards real-time object detection with region proposal networks. In: Cortes, C., Lawrence, N.D., Lee, D.D., Sugiyama, M., Garnett, R. (eds.) Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7–12, 2015, Montreal, Quebec, Canada. pp. 91–99 (2015) 1, 2, 12
34. Smola, A.J., Kondor, R.: Kernels and regularization on graphs. Learning Theory and Kernel Machines pp. 144–158 (2003) 2, 4
35. Sun, B., Li, B., Cai, S., Yuan, Y., Zhang, C.: FSCE: few-shot object detection via contrastive proposal encoding. CoRR **abs/2103.05950** (2021) 12
36. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S.E., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7–12, 2015. pp. 1–9. IEEE Computer Society (2015). <https://doi.org/10.1109/CVPR.2015.7298594> 4
37. Tuzel, O., Porikli, F., Meer, P.: Region covariance: A fast descriptor for detection and classification. In: Leonardis, A., Bischof, H., Pinz, A. (eds.) Computer Vision - ECCV 2006, 9th European Conference on Computer Vision, Graz, Austria, May 7–13, 2006, Proceedings, Part II. Lecture Notes in Computer Science, vol. 3952, pp. 589–600. Springer (2006). https://doi.org/10.1007/11744047_45 4
38. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Advances in Neural Information Processing Systems 30, 2017. pp. 5998–6008 (2017) 2, 4, 6
39. Wang, X., Huang, T.E., Gonzalez, J., Darrell, T., Yu, F.: Frustratingly simple few-shot object detection. In: ICML 2020. Proceedings of Machine Learning Research, vol. 119, pp. 9919–9928. PMLR (2020) 12
40. West, J., Venture, D., Warnick, S.: Spring research presentation: A theoretical foundation for inductive transfer. Brigham Young University, College of Physical and Mathematical Sciences (2007) 1
41. Woodworth, R.S., Thorndike, E.L.: The influence of improvement in one mental function upon the efficiency of other functions. Psychological Review (I) **8**(3), 247–261 (1901). <https://doi.org/10.1037/h0074898> 1
42. Wu, A., Han, Y., Zhu, L., Yang, Y.: Universal-prototype enhancing for few-shot object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 9567–9576 (October 2021) 12
43. Wu, J., Liu, S., Huang, D., Wang, Y.: Multi-scale positive sample refinement for few-shot object detection. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J. (eds.) Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August

- 23–28, 2020, Proceedings, Part XVI. Lecture Notes in Computer Science, vol. 12361, pp. 456–472. Springer (2020) [3](#), [12](#)
44. Xie, S., Girshick, R.B., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21–26, 2017. pp. 5987–5995. IEEE Computer Society (2017). <https://doi.org/10.1109/CVPR.2017.634> [4](#)
 45. Yan, X., Chen, Z., Xu, A., Wang, X., Liang, X., Lin, L.: Meta R-CNN: towards general solver for instance-level low-shot learning. In: 2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019. pp. 9576–9585. IEEE (2019). <https://doi.org/10.1109/ICCV.2019.00967> [1](#), [2](#), [3](#), [11](#), [12](#)
 46. Yang, Y., Wei, F., Shi, M., Li, G.: Restoring negative information in few-shot object detection. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (eds.) Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6–12, 2020, virtual (2020) [3](#), [12](#)
 47. Zhang, H., Zhang, L., Qi, X., Li, H., Torr, P., Koniusz, P.: Few-shot action recognition with permutation-invariant attention. In: Proceedings of the European Conference on Computer Vision (ECCV 2020). vol. 12350. Springer (2020) [4](#)
 48. Zhang, H., Koniusz, P.: Power normalizing second-order similarity network for few-shot learning. In: IEEE Winter Conference on Applications of Computer Vision, WACV 2019, Waikoloa Village, HI, USA, January 7–11, 2019. pp. 1185–1193. IEEE (2019). <https://doi.org/10.1109/WACV.2019.00131> [4](#)
 49. Zhang, H., Koniusz, P., Jian, S., Li, H., Torr, P.H.S.: Rethinking class relations: Absolute-relative supervised and unsupervised few-shot learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9432–9441 (June 2021) [4](#)
 50. Zhang, S., Luo, D., Wang, L., Koniusz, P.: Few-shot object detection by second-order pooling. In: Proceedings of the Asian Conference on Computer Vision (2020) [1](#), [2](#), [3](#), [4](#), [11](#), [12](#)
 51. Zhang, S., Wang, L., Murray, N., Koniusz, P.: Kernelized few-shot object detection with efficient integral aggregation. In: IEEE Conference on Computer Vision and Pattern Recognition (2022) [1](#), [2](#)
 52. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable DETR: deformable transformers for end-to-end object detection. In: ICLR 2021. OpenReview.net (2021) [4](#)