Convolutional Dynamic Alignment Networks for Interpretable Classifications

Supplementary Material

Moritz Böhle MPI for Informatics Saarland Informatics Campus Mario Fritz CISPA Helmholtz Center for Information Security Bernt Schiele MPI for Informatics Saarland Informatics Campus

Table of Contents

In this supplement to our work on Convolutional Dynamic Alignment Networks (CoDA-Nets), we provide:

(A)	Additional qualitative results				
	In this section, we show additional <i>qualitative</i> results on the Imagenet subset as well as additional com- parisons between the model-inherent contribution maps and other methods for importance attribution. Further, we show the effect of regularising the linear mappings on the contribution maps for models trained on the CIFAR-10 dataset. Lastly, we show the results of the 'sanity check' by Adebayo et al. [1] as well as the contribution maps of a piece-wise linear model (ResNet-56).				
(B)	Additional quantitative results				
	In this section, we show additional <i>quantitative</i> results. In particular, we show the accuracies of the temperature-regularised models and of models of different sizes (DAU rank ablation). Further, we show interpretability results for models trained with the L2 non-linearity, with explicit regularisation of the linear mapping $W_{0\to L}$, and for a pre-trained ResNet-56 for comparison. Moreover, we show results for the <i>pixel removal metric</i> when removing the most important pixels first.				
(C)	Implementation details				
	In this section, we present architecture and training details for our experiments and describe in detail how the convolutional Dynamic Alignment Units are implemented. Further, we discuss the results of ResNets on the Imagenet subset under the exact same training scheme for comparison.				
(D)	Relation to low-rank matrix approximations				
	In this section, we discuss the relationship between the Dynamic Alignment Units and the problem of low-rank matrix approximation.				
(E)	Relation to capsule networks				
	In this section, we discuss the relationship between the Dynamic Alignment Units and capsules [9]. In particular, we rewrite the standard capsule formulation, which allows us to compare them more easily to our work. Under this new formulation, it becomes clear that the two approaches share similarities, but also that there exist important differences.				



Figure A1: The first 25 most confident classifications decomposed into the contributions from each spatial location, filtered to 1 image per class. Positive (negative) contributions for the ground truth class are shown in red (blue).

A. Additional qualitative results

Additional Imagenet examples. In Figs. A1 and A2 we present additional qualitative examples of the model-inherent contribution maps. In particular, we show the decomposition of the model predictions into input contributions for 50 out of 100 classes; for each class, we show the most confidently classified image and show the classes in sorted order (by confidence). In Figs. A3 and A4, we moreover show the contributions maps for the first 10 of the overall most confidently classified images next to the attribution maps from the post-hoc importance attribution methods for qualitative comparison. We note that GradCam consistently highlights very similar regions to the CoDA-Net contribution maps, but does so at a lower resolution. All contribution maps based on the CoDA-Net use the same linear color scale, which has been set to (-v, v) with v the 99.75th percentile over all absolute values in the contributions maps shown in Figs. A1 and A2. For reproducing the presented contribution maps and more, please visit github.com/moboehle/CoDA-Nets.

Regularising the linear mapping on CIFAR-10. As described in the main paper, the explicit representation of the model computations as a linear mapping, i.e.,

$$\hat{\mathbf{y}}(\mathbf{x}) = \mathbf{W}_{0 \to L}(\mathbf{x})\mathbf{x}$$
,

allows to explicitly regularise the linear mappings and thereby the model-inherent contribution maps. In Fig. A5 we qualitatively show how the regularisation impacts the contribution maps. In particular, we see that without any regularisation the contribution maps are difficult to interpret. As soon as even a small regularisation is applied, however, the maps align well with the discriminative parts of the input image.

Sanity check. In [1], the authors found that many commonly used methods for importance attribution are not *model-faithful*, i.e., they do not reflect the learnt parameters of the model. In order to test this, they proposed to examine how the attributions change if the model parameters are randomised step by step. If the attributions remain stable under model randomisation, they cannot be assumed to explain a specific model, but rather reflect properties of the general architecture and the input data. In Fig. A6, we show how the model-inherent contribution maps behave when re-initialising the CoDA-Net layers one at a time to a random parameter setting, starting from the deepest layer. As can be seen, the contribution maps get significantly perturbed with every layer that is reset to random parameters; thus, the contribution maps pass this 'sanity check' for attribution methods.



Figure A2: The 26th to the 50th most confident classifications decomposed into the contributions from each spatial location, filtered to 1 image per class. Positive (negative) contributions for the ground truth class are shown in red (blue).



Figure A3: Comparison between attribution methods for the 4 most confident classifications. We show positive importance attributions in red, negative attributions in blue; for RISE we use its default visualisation. Note that the model seems to align the weights well with the ornamental eyespots of the peacocks (also see the peacock in Fig. A4) or the heads of the ducks in Fig. A4. While the latter are also highlighted by GCam, the former constitute a structure that is too fine-grained for GCam to resolve properly.



Figure A4: Comparison between attribution methods for the 5th to the 10th most confident classifications. We show positive importance attributions in red, negative attributions in blue; for RISE we use its default visualisation.



Figure A5: Visualising the qualitative effect of explicitly regularising the linear mapping $\mathbf{W}_{0\to L}$ on CIFAR-10. While the contribution maps without regularisation are noisy, they become sharper as soon as a regularisation is applied.



Figure A6: Sanity check experiment as in [1]. If the importance attributions remain stable under parameter randomisation, they cannot be assumed to faithfully reflect the learnt parameters of the model. Since the contribution maps get significantly perturbed when re-initialising layers from network output (left) to network input (right), the model-inherent contribution maps thus pass this sanity check. Positive contributions are shown in red, negative contributions in blue.



Figure A7: 'Input×Gradient' evaluated on different ResNets. Since ResNets are piece-wise linear functions, s.t. $\mathbf{y}(\mathbf{x}) = \mathbf{M}(\mathbf{x})\mathbf{x} + \mathbf{b}(\mathbf{x})$, this is the ResNet-based equivalent to the CoDA-Net contribution maps.

Contribution maps of a piece-wise linear model. In Fig. A7 we show contribution maps obtained from different pretrained ResNet architectures obtained from https://github.com/akamaster/pytorch_resnet_cifar10. In particular, we visualise the 'Input×Gradient' method. Note that this yields contribution maps, since piece-wise linear models, such as the ResNets, produce input-dependent linear mappings, similar to the CoDA-Nets. These contribution maps, however, are rather noisy and do not reveal particularly relevant features.



Figure B1: Left: Localisation metric results for models trained with different temperatures T, same as in main paper (see Fig. 6, top right). Right: Corresponding accuracies of the models on the CIFAR10 test set. There seems to be a trade-off between the interpretability and the accuracy of the models due to the regularising effect of T.



Figure B2: Effect of scaling the rank of the DAUs in the CoDA-Nets on accuracy. Specifically, all ranks in the S-CoDA-SQ model presented in the main paper were scaled with the same factor, thereby changing the model size. The original model is shown with 8M parameters for comparison.

B. Additional quantitative results

Performance / interpretability trade-off. While the CoDA-Nets were observed to train and perform well over a wide range of choices for the logit temperature T, there seems to be a trade-off between the accuracies of the network and their interpretability—the implicit alignment regularisation comes at a cost. For example, in Fig. **B1**, we contrast the gain in interpretability (left, same figure as in main paper) with the corresponding accuracies (right).

Model size vs. accuracy. Given the quadratic form in the DAUs (cf. eq. (1), $\approx \mathbf{x}^T \mathbf{M} \mathbf{x}$), the number of parameters per DAU scales quadratically with the input dimensions. In order to limit the model size, we decided to explicitly limit the rank of the DAUs by factorising the matrix \mathbf{M} into \mathbf{AB} with $\mathbf{A} \in \mathbb{R}^{d \times r}$ and $\mathbf{B} \in \mathbb{R}^{r \times d}$. While this allows to be more parameter efficient, it, of course, affects the modelling capacity of the DAUs. In Fig. B2, we present how the accuracy changes with the model size; for this, we scaled the ranks of all DAUs per layer with factors of 1/8, 1/4, 1/2, and 4.5 compared to the S-CoDA-SQ model presented in the main paper. This results in models with 1.1M, 2.0M, 4.0M, and 34.6M parameters respectively. For comparison, the original model is also included in Fig. B2 (8M parameters).

Results for L2 non-linearity. In Fig. B3a we show the results of evaluating the different methods for importance attribution on a model with the L2 non-linearity, see eq. (2) in the main paper. As can be seen, the results are very similar to those presented in the main paper in Fig. 6 (center column); in particular, the model-inherent contribution maps outperform the other methods under the localisation metric and are on par with the occlusion methods under the pixel removal metric; note, however, that the occlusion methods are a direct estimate of the behaviour under pixel removal and therefore expected to perform well under this metric.

Regularisation of W_{0 \rightarrow L}. In Fig. B3b, we show the results of the localisation metric and the pixel removal metric for models trained on CIFAR-10 with different regularisation strengths λ , see eq. (10) in the main paper. We note that the localisation metric benefits from an increase in the regularisation strength λ . For the pixel removal metric, we observe that while the models improve at first under this metric, a strong regularisation makes the predictions more brittle.

Removing the most important pixels first. In Fig. B4a, we show the results of the pixel removal metric when removing



(a) The quantitative results for a CoDA-Net trained with the L2 non-linearity are very similar to those of a model trained with the squashing non-linearity, see eq. (2) in the main paper and compare with Fig. 6 (center column) in the main paper. In particular, we observe that the CoDA-Net outperforms the other methods under the localisation metric and achieves similar performance to the occlusion attribution method, which directly estimates the change in model prediction when removing a pixel.



(b) Similar to the effect of temperature scaling, cf. Fig. 6 (right column) in the main paper, here we show quantitative results for different regularisation strengths λ . Similar to increasing the temperature, stronger regularisations also improve the localisation metric (left). While the pixel perturbation metric (right) also improves at first ($\lambda = 0.01$ and $\lambda = 0.05$), the models become less stable for stronger regularisation strengths. In comparison, we found that increasing the temperature also improves the performance of the models under this metric, see Fig. 6 (right column) in the main paper.

Figure B3: Quantitative results for two ablations: (a) using the L2 non-linearity and (b) increasing the regularisation of the linear mapping $\mathbf{W}_{0\to L}$, see eqs. (2) and (10) in the main paper respectively. Compare with Fig. 6 in the main paper.

the most important pixels first. As can be seen, the model-inherent contributions better predict the importance of pixels and outperform the other attribution methods under this metric; in particular, the confidence drops most rapidly when removing pixels according to the ranking given by the model-inherent contribution maps.

Evaluating a pre-trained ResNet. In order to establish a baseline for the performance of the different attribution methods on a classical neural network architecture, we additionally evaluated the attribution methods on a pre-trained ResNet and show the results in Fig. B4b. Specifically, we rely on a publicly available pre-trained ResNet-56 obtained from https://github.com/akamaster/pytorch_resnet_cifar10, which achieves a classification accuracy of 93.39%. The results show that the performance of the CoDA-Net-derived contribution maps is not only strong when comparing them to attribution methods evaluated *on the same model*. Instead, they also perform well in comparison to those methods evaluated *on a different model*. In particular, the model-inherent contribution metric. Further, only the occlusion attributions produce similarly strong pixel importance rankings for the ResNet; note, however, that the occlusion methods are a direct estimate of the behaviour under pixel removal and therefore expected to perform well under this metric.

C. Implementation details

C.1. Training and architecture details

Architectures. The architectures used for the experiments in section 4 are given in Table C1. All activation maps are padded with (k-1)/2 zeros on each side, such that the spatial dimensions are only reduced by the strides; here, k refers to the kernel size. As can be seen, the activation maps thus still have a spatial resolution after the last layer, which we further reduce with a global sum-pooling layer. Note that global sum-pooling is just a linear layer with no trainable parameters and therefore



(a) Results for the pixel removal metric when first removing the most important pixels. As can be seen, the ranking given by the modelinherent contribution maps seems to best reflect the pixel importance, since the confidence most rapidly drops when removing pixels according to this ranking.



(b) Quantitative results for methods for importance attribution evaluated on a pre-trained ResNet-56 on CIFAR-10. Comparing these results to Fig. 6 (center column) in the main paper or Fig. B3a, it can be seen that the model-inherent contribution maps of the CoDA-Net are also strong when compared to importance attributions evaluated on a different model.

Figure B4: In (a) we show the results of the pixel-removal metric when removing those pixels first that are considered the most important ones according to the importance attribution method. Moreover, in (b) we plot the quantitative results for the evaluation metrics of the importance attributions for a pre-trained ResNet-56.

still allows for linear decomposition. The input itself consists of 6 channels, the image and its negative, as explained in section 3.4; hence, the first layer takes an input with 6 channels per pixel.

Training details. We use the pytorch library [6] and optimise all networks with the Adam optimiser [5] with default values. As for the loss function, we use the binary cross entropy loss to optimise class probabilities individually (as 'one-vs-all' classifiers). For all networks, we used a base learning rate of 2.5×10^{-4} ; for the Imagenet experiment, we employed learning rate warm-up and linearly increased the learning rate from 2.5×10^{-4} to 1×10^{-3} over the first 15 epochs. Further, we trained for 200 epochs on CIFAR-10, for 100 epochs on TinyImagenet, and for 60 epochs on the Imagenet subset; we decreased the learning rate by a factor of 2 after every 60/30/20 epochs on CIFAR-10/TinyImagenet/Imagenet. We used a batch size of 16, 128, and 64 for CIFAR-10, TinyImagenet, and Imagenet respectively. For the Imagenet subset, we additionally used RandAugment [3] with parameters n = 2 and m = 9; for this, we relied on the publicly available implementation at https://github.com/ildoonet/pytorch-randaugment and followed their augmentation scheme. The qualitatively evaluated model (see Figs. 5, 8, and A1-A4) for the Imagenet subset was trained with T = 1e5 and achieved a top-1 accuracy of 76.5%. For comparison, we trained several ResNet-50 models (taken from the pytorch library [6]) with the exact same training procedure, i.e., batch size, learning rate, optimiser, augmentation, etc.). The best ResNet-50 out of 4 runs achieved 79.16% top-1 accuracy¹, which outperforms the CoDA-Net but is nevertheless comparable. While it is surely possible to achieve better accuracies for both models, long training times for the CoDA-Nets have thus far prevented us from properly optimising the architectures both on the 100 classes subset, as well as on the full Imagenet dataset. In order to scale the CoDA-Net models to larger datasets, we believe it is important to first improve the model efficiency in future work. Lastly, when regularising the matrix entries of $W_{0 \rightarrow L}$, see eq. (11), we regularised the absolute values for the true class c, $[\mathbf{M}_{0\to L}]_c$, and a randomly sampled incorrect class per image.

¹The best test accuracies per run are given by 79.16%, 79.04%, 78.86%, and 78.7% respectively.

C.2. Convolutional Dynamic Alignment Units

In Algorithm 1, we present the implementation of the convolutional DAUs (CoDAUs). As can be seen, a Convolutional Dynamic Alignment Layer applies dynamic (input-dependent) filters to each of the patches extracted at different spatial locations. In detail, the Dynamic Alignment Units are implemented as two consecutive convolutions (lines 11 and 15), which are equivalent to first applying matrix **B** (line 24) and then **A** to each patch and adding a bias term **b** (line 29). After applying the non-linearity (line 33), we obtain the dynamic weight vectors for CoDAUs as described in eq. (1) in the main paper. In particular, for every patch \mathbf{p}_{hw} extracted at the spatial positions hw in the input, we obtain the dynamic weight $\mathbf{w}_j(\mathbf{p}_{hw})$ for the *j*-th DAU as

$$\mathbf{w}_{i}(\mathbf{p}_{hw}) = g(\mathbf{A}_{i}\mathbf{B}\mathbf{p}_{hw} + \mathbf{b}_{i}); \qquad (C.1)$$

note that the projection matrices **B** are thus shared between the DAUs. These weights are then applied to the respective locations (line 41) to yield the outputs of the DAUs per spatial location. As becomes apparent in line 41, the outputs are linear transformations (weighted sums) of the input and can be written as

$$\mathbf{a}_{l+1}(\mathbf{a}_l) = \mathbf{W}(\mathbf{a}_l)\mathbf{a}_l \quad , \tag{C.2}$$

with $\mathbf{a}_l \in \mathbb{R}^d$ the vectorised input to layer l and $\mathbf{W} \in \mathbb{R}^{f \times d}$ and f the number of filters (DAUs). The rows in matrix \mathbf{W} correspond to exactly one filter (DAU) applied to exactly one patch \mathbf{p}_{hw} and are non-zero only at those positions that correspond to this specific patch in the input.

C.3. Attribution methods

In section 4.2, we compare the model-inherent contribution maps of the CoDA-Nets to those of the following methods for importance attribution: the gradient of the class logits with respect to the input image [2] (Grad), 'Input×Gradient' (IxG, cf. [1]), GradCam [10] (GCam), Integrated Gradients [13] (IntG), DeepLIFT [11], several occlusion sensitivities (Occ-K, with K the size of the occlusion patch) [14], RISE [7], and LIME [8].

For RISE and LIME, we relied on the official implementations available at https://github.com/eclique/RISE and https://github.com/marcotcr/lime respectively. For RISE, we generated 6000 masks with parameters s = 6 and $p_1 = 0.1$. For LIME, we evaluated on 256 samples per image and used the top 3 features for the attribution; for the segmentation, we also used the default parameters, namely 'quickshift' with $max_dist = 200$, ratio = 0.2, and a kernel size of 4.

For Grad, GCam, IxG, IntG, DeepLIFT, and the occlusion sensitivities, we relied on the publicly available pytorch library 'captum' (https://github.com/pytorch/captum). GCam was used on the last activation map before global sumpooling. The occlusion sensitivities were used with strides of 2 on CIFAR-10 and strides of 4 for TinyImagenet. Finally, for IntG we used 20 steps for the integral approximation.

C.4. Evaluation metrics

In section 4.2, we evaluated the attribution methods against 2 quantitative metrics: (1) the adapted *pointing game* [15] and (2) the prediction stability under removing the *least important pixels* as in [12]. In section **B**, we further show results for removing the *most important pixels* first.

For (1), we constructed 500 (250) 3×3 multi-images for CIFAR-10 (TinyImagenet); for an example with 2×2 , see Fig. 7 in the main paper. In each of these multi-images, every class occurred at most once. As stated in section 4.2, we measured the fraction of positive contributions falling inside the correct mini-image. Further, the images were sorted according to their confidence for each of the classes. For every multi-image, a random set of classes was sampled. For each of the classes, we included the most confidently classified class image in the multi-image that had not been used yet in previous multi-images.

For (2), we followed [12] and successively replaced one pixel at a time by [0, 0, 0, 0, 0], until up to 25% of the image were removed. The pixels were removed in order, sorted by their assigned importance.

Alg	Algorithm 1: Implementation of a Convolutional Dynamic Alignment Layer						
1 f 1	om torch import nn						
2 ir	nport torch.nn.functional as F						
3							
4 cl	ass DAUConv2d(nn.Module):						
5							
6	definit (self, in_channels, out_channels, rank, kernel_size, stride, padding, act_func):						
7	<i># act_func: non-linearity for scaling the weights. E.g., L2 or SQ.</i>						
8	<i># out_channels: Number of convolutional DAUs for this layer.</i>						
9	# rank: Rank of the matrix AB.						
10	# 'dim_reduction' applies matrix B .						
11	self .dim_reduction = nn.Conv2d(<i>in_channels, rank, kernel_size, stride, padding</i> , bias= False)						
12	# Total dimensionality of a single patch						
13	<pre>self.patch_dim = in_channels * kernel_size * kernel_size</pre>						
14	# 'weightings' applies matrix \mathbf{A} and adds bias \mathbf{b} .						
15	<pre>self.weightings = nn.Conv2d(rank, out_channels * self.patch_dim , kernel_size=1, bias=True)</pre>						
16	<pre>self.act_func = act_func</pre>						
17	self .out_channels = out_channels						
18	self .kernel_size = kernel_size						
19	self.stride = stride						
20	self .padding = padding						
21							
22	<pre>def forward (self, in_tensor) :</pre>						
23	# Project to lower dimensional representation, i.e., apply matrix B . This yields Bp for every patch p .						
24	reduced = self .dim_reduction(in_tensor)						
25	# Get new spatial size height h and width w						
26	h, w = reduced.shape $[-2:]$						
27	batch_size = in_tensor.shape[0]						
28	# Apply matrix A and add bias b , yielding $ABp + b$ for every patch p .						
29	weights = self.weightings(reduced)						
30	# Reshape for every location to size patch_dim×out_channels						
31	weights = weights.view(batch_size, self.patch_dim, out_channels, h, w)						
32	# Apply non-linearity to the weight vectors, yielding $\mathbf{w}(\mathbf{p}) = g(\mathbf{ABp} + \mathbf{b})$ as in eq. (1) for every patch \mathbf{p}						
33	weights = self.act_func(weights, dim=1)						
34	# Extract patches from the input to apply dynamic weights to patches.						
35	patches = F.unfold(in_tensor, self.kernel_size, padding=self.padding, stride=self.stride)						
36	# Reshape for applying weights.						
37	patches = patches.view(batch_size, self .patch_dim, 1, h, w)						
38	# Apply the weights to the patches.						
39	# As can be seen, the output is just a weighted combination of the input, i.e., a linear transformation.						
40	# The output can thus be written as $\mathbf{o} = \mathbf{W}(\mathbf{x})\mathbf{x}$.						
41	return (patches * weights).sum(1)						

Network	Input dimensions	Layer	Number of DAUs	Rank of AB	Kernel size	Stride
		1	16	32	3	1
		2	16	32	3	1
		3	32	64	3	2
Ψ		4	32	64	3	1
QD	$6 \times 32 \times 32$	5	32	64	3	1
S-O		6	64	64	3	2
		7	64	64	3	1
		8	64	64	3	1
		9	10	64	1	1
		1	16	64	3	1
		2	16	64	3	1
		3	32	128	3	2
AC		4	32	128	3	1
Col	$6 \times 32 \times 32$	5	32	128	3	1
0-M		6	64	256	3	2
-		7	64	256	3	1
		8	64	256	3	1
		9	10	256	1	1
		1	16	64	7	3
		2	32	64	3	1
		3	32	64	3	1
AC		4	64	128	3	2
CoI	$6 \times 240 \times 240$	5	64	128	3	1
Γ		6	64	128	3	1
		7	64	256	3	2
		8	64	256	3	1
		9	100	256	3	1
		1	16	64	5	1
		2	32	64	3	1
		3	32	128	3	2
DA		4	64	128	3	1
ç	$6\times 64\times 64$	5	64	128	3	1
Ļ.		6	64	256	3	2
\sim		7	64	256	3	1
		8	64	256	3	1
		9	200	256	3	2

Table C1: Architecture details for the experiments described in section 4.

D. Low-rank matrix approximations

In the following, we first introduce the standard formulation of the low-rank matrix approximation problem and present the well-known solution via singular value decomposition. We then introduce an additional constraint to this standard formulation and show that at their optimum (maximal average output), the DAUs solve this constrained approximation problem.

Low-rank approximation problem. Given a data matrix² $\mathbf{M} \in \mathbb{R}^{m \times n}$, the low-rank approximation problem is typically formulated as $\min_{\mathbf{M}'}(||\mathbf{M} - \mathbf{M}'||_F)$ with $rank(\mathbf{M}') \leq r$ and $||\cdot||_F$ the Frobenius norm. The Frobenius norm can be written as a sum over the column differences $||\mathbf{\Delta}_i||_F^2 = ||\mathbf{m}_i||_2^2 + ||\mathbf{m}'_i||_2^2 - 2\mathbf{m}_i^T\mathbf{m}'_i$ with $\mathbf{m}_i, \mathbf{m}'_i \in \mathbb{R}^m$ the *i*-th columns of the matrices \mathbf{M} and \mathbf{M}' respectively. Note that since the \mathbf{m}_i are fixed (they just correspond to the fixed input matrix \mathbf{M}), the optimisation objective is equivalent to

$$\min \sum_{i} ||||\mathbf{m}_{i}'||_{2}^{2} - 2\mathbf{m}_{i}^{T}\mathbf{m}_{i}' =$$
(D.1)

$$\min \sum_{i} ||||\mathbf{m}'_{i}||_{2}^{2} - 2||\mathbf{m}_{i}||||\mathbf{m}'_{i}||\cos(\alpha_{i}) \quad , \tag{D.2}$$

²In the context of our work, we can think of **M** as storing *n* data samples of dimensionality *m*; e.g., when considering images, each column might correspond to one vectorised image from a dataset of *n* images. The low-rank approximation problem aims to approximate this dataset with *r* independent variables (dimensionality reduction).

with α_i the angle between \mathbf{m}_i and \mathbf{m}'_i . Hence, optimising the Frobenius norm of the difference matrix finds a trade-off between aligning the directions (angles) of the columns of \mathbf{M} and \mathbf{M}' and approximating the correct norms of the original columns in \mathbf{M} . The optimal solution for this problem is given by singular value decomposition (SVD) [4] of \mathbf{M} , and the optimal \mathbf{M}' can be written as

$$\mathbf{M}' = \mathbf{U}_r \mathbf{U}_r^T \mathbf{M} = \mathbf{U} \boldsymbol{\Sigma}_r \mathbf{V}^T \quad , \tag{D.3}$$

where \mathbf{U}_r is the matrix of left singular vectors of \mathbf{M} up to the *r*-th vector³; the remaining vectors are replaced by zero-vectors. Note that the first formulation of the solution $(\mathbf{U}_r \mathbf{U}_r^T \mathbf{M})$ emphasises a property that we will encounter again for the DAUs. In particular, this formulation highlights the fact that the optimal matrix \mathbf{M}' is given by reconstructing the individual data points from the *r*-dimensional eigenspace spanned by the eigenvectors of the data points; i.e., each column \mathbf{m}'_i can be calculated as $\mathbf{U}_r \mathbf{U}_r^T \mathbf{m}_i$. The right hand side of the equation shows the conventional SVD-notation of the low-rank approximation solution, with $\mathbf{U} \in \mathbb{R}^{m \times n}$, $\boldsymbol{\Sigma}_r \in \mathbb{R}^{m \times n}$ $\mathbf{V} \in \mathbb{R}^{n \times n}$ and $\boldsymbol{\Sigma}_r$ being a rectangular diagonal matrix with the first *r* singular values as non-zero entries.

Constrained low-rank approximations Now we diverge from the standard low-rank approximation formulation and impose an additional constraint on \mathbf{M}' (red box):

$$\min_{\mathbf{M}'}(||\mathbf{M} - \mathbf{M}'||_F), \qquad (D.4)$$

s.t.
$$rank(\mathbf{M}') \le r \wedge ||\mathbf{m}'_i||_2^2 = 1 \ \forall i$$
 (D.5)

When combining eqs. (D.2) and (D.5), it becomes clear that the optimisation problem can now be formulated as

$$\max \sum_{i} \mathbf{m}_{i}^{T} \mathbf{m}_{i}^{\prime} = \max \sum_{i} ||\mathbf{m}_{i}||_{2} \cos\left(\angle(\mathbf{m}_{i}, \mathbf{m}_{i}^{\prime})\right) , \qquad (D.6)$$

since the norm of \mathbf{m}'_i is fixed and the addition of or multiplication with a fixed constant does not change the location of the optima.

As a corollary, we note that *any* matrix \mathbf{M}' with a maximum rank r that maximises eq. (D.6) gives a solution to the constrained approximation problem, *independent* of the norm of its columns: the columns can be normalised *post-hoc* to fulfill the new constraint, since rescaling does not change neither the rank of the matrix, nor the score of the objective function in eq. (D.6).

Further, we note that the unnormalised optimal matrix $\widetilde{\mathbf{M}}'$ can be factorised as

$$\mathbf{M}' = \mathbf{ABM} , \tag{D.7}$$

s.t.
$$\mathbf{M}' = G(\mathbf{M}') = G(\mathbf{ABM})$$
 (D.8)

with $\mathbf{A} \in \mathbb{R}^{m \times r}$ and $\mathbf{B} \in \mathbb{R}^{r \times n}$, similar to the solution to the unconstrained matrix approximation problem in eq. (D.3). Here, G is a function that normalises the columns of a matrix to unit norm (as discussed in the previous paragraph, in order to fulfill the norm constraint, we can normalise the columns *post-hoc*).

Finally, this allows us to rewrite the objective function in eq. (D.6) as

$$\max \sum_{i} g(\mathbf{ABm}_{i})\mathbf{m}_{i} = \max \sum_{i} \mathbf{w}(\mathbf{m}_{i})\mathbf{m}_{i}$$
(D.9)

with $g(\mathbf{v})$ normalising its input vector to unit norm. Comparing this with eq. (1) in the main paper, we see that this is equivalent to maximising a DAU without bias term for maximal output over the *n* columns in M and therefore, at their optimum, the DAUs solve the constrained low-rank matrix approximation problem in eq. (D.5). In particular, the DAUs achieve this by encoding common inputs in the eigenvectors of **AB**, which allows for an optimal angular reconstruction of the inputs, similar to reconstructing from the first PCA components in a PCA decomposition. The main difference to PCA is that PCA yields optimal reconstructions under the L2-norm, whereas the DAUs yield optimal angular reconstructions.

³A short proof of the equality to the right can be found at the end of this section.

Proof of $\mathbf{U}_r \mathbf{U}_r^T \mathbf{M} = \mathbf{U} \boldsymbol{\Sigma}_r \mathbf{V}^T$ **.**

In order to see that this equality holds, we first write matrix \mathbf{M} as

$$\mathbf{M} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^T \quad (\text{SVD-form}) \tag{D.10}$$

Now, we multiply both sides with $\mathbf{U}_r \mathbf{U}_r^T$ from the left:

$$\mathbf{U}_r \mathbf{U}_r^T \mathbf{M} = \mathbf{U}_r \mathbf{U}_r^T \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^T \tag{D.11}$$

$$\xrightarrow{\mathbf{U}_r^T \mathbf{U} = \mathbf{I}_r} \mathbf{U}_r \mathbf{U}_r^T \mathbf{M} = \mathbf{U}_r \mathbf{I}_r \mathbf{\Sigma} \mathbf{V}^T$$
(D.12)

$$\Rightarrow \mathbf{U}_r \mathbf{U}_r^T \mathbf{M} = \mathbf{U} \boldsymbol{\Sigma}_r \mathbf{V}^T \qquad \Box.$$
(D.13)

Here, we made use of the fact that U is an orthogonal matrix and therefore $\mathbf{U}^T \mathbf{U} = \mathbf{I}$; since we only use the first r eigenvectors, we obtain the truncated identity matrix \mathbf{I}_r when multiplying $\mathbf{U}_r^T \mathbf{U}$. Further, in the last line, we used that $\mathbf{U} \boldsymbol{\Sigma}_r = \mathbf{U}_r \boldsymbol{\Sigma}_r$.

E. Comparison to capsule networks

In order to discuss the relationship to capsule networks, in section E.1 we will first rewrite the classical capsule formulation in 'Dynamic Routing Between Capsules' by Sabour et al. [9] to mitigate the notational differences between their work and our work. In section E.2 we then show that while capsules and Dynamic Alignment Units share some computations, there are several important differences that we summarise in Table E1.

	Classical capsules	Dynamic Alignment Units
Non-Lin. g	SQ	SQ, L2,
Activations	$g\left(\mathbf{Vx} ight)$	$g\left(\mathbf{ABx}+\mathbf{b} ight)$
Routing	yes	no
Low-rank	no	yes
Output	$CAP(\mathbf{x})$	$\operatorname{CAP}(\mathbf{x})^T \mathbf{x}$

Table E1: Comparison between capsules and Dynamic Alignment Units (DAUs). Importantly, DAUs produce a linear transformation of the input by multiplying the 'capsule activations' with the input (see 'Output') and allow for constraining the rank of the transformation. The dynamic weights in the DAUs can be seen as the activations of a capsule, s.t. $\mathbf{w}(\mathbf{x}) = CAP(\mathbf{x})$, see 'Output' in the table.

E.1. Reformulating capsules

In this subsection, we will show that the classical capsule formulation (eq. (E.1)), in which input capsules 'vote' for the activations of an output capsule, can be written as a simple linear transformation $\mathbf{s} = \mathbf{V}\mathbf{x}$ if just one iteration of the dynamic routing algorithm is applied; here, s is a vector containing the activations of the output capsule, \mathbf{V} stores the 'votes' of the input capsules to an output capsule, and x is a vector containing the activations of all input capsules. In the following, we will start from how capsules are formulated in [9] and rewrite this formulation step by step.

In [9] eq. (2), the authors calculate the activations s of a capsule⁴ before any routing or non-linearity as

$$\mathbf{s} = \sum_{i} c_i \hat{\mathbf{u}}_i , \quad \hat{\mathbf{u}}_i = \mathbf{W}_i \mathbf{u}_i \quad .$$
(E.1)

Here, \mathbf{u}_i is the capsule vector of capsule *i* from the incoming layer and \mathbf{W}_i a matrix which transforms the capsule activations to generate the votes $\hat{\mathbf{u}}_i$, i.e., the vote of the *i*-th incoming capsule to the output capsule in the current layer. Note that c_i

⁴As we only discuss a single output capsule, we omitted the subscript j for the j-th output capsule for simplicity.

are the coefficients for the dynamic routing algorithm. If no routing is applied, they can be combined with \mathbf{W}_i to yield $\widetilde{\mathbf{W}}_i = c_i \mathbf{W}_i$, which simplifies eq. (E.1) to

$$\mathbf{s} = \sum_{i} \hat{\mathbf{u}}_{i} , \quad \hat{\mathbf{u}}_{i} = \widetilde{\mathbf{W}}_{i} \mathbf{u}_{i} .$$
(E.2)

Further, we note that the linear transformation of \mathbf{u}_i can be represented as votes of the individual entries t of \mathbf{u}_i :

$$\hat{\mathbf{u}}_{i} = \widetilde{\mathbf{W}}_{i} \mathbf{u}_{i} = \sum_{t} \left[\mathbf{u}_{i}\right]_{t} \left[\widetilde{\mathbf{W}}_{i}^{T}\right]_{t} \quad .$$
(E.3)

Here, $[\mathbf{W}]_t$ denotes the *t*-th row in a matrix \mathbf{W} (equivalently for a vector). Inserting this formulation of $\hat{\mathbf{u}}_i$ in eq. (E.2) yields

$$\mathbf{s} = \sum_{i} \sum_{t} \left[\mathbf{u}_{i} \right]_{t} \left[\widetilde{\mathbf{W}}_{i}^{T} \right]_{t} \quad .$$
(E.4)

Hence, s can be represented as the result of votes from all neurons u contained in any of the incoming capsules (note that we sum over all entries in all capsules). If we represent the activations of these neurons in a single vector x, denote their activations by x_u , and their respective votes for the output capsule by v_u we can write eq. (E.4) as:

$$\mathbf{s} = \sum_{i} \sum_{t} \overbrace{[\mathbf{u}_{i}]_{t}}^{\widehat{=}\mathbf{v}_{u}} \overbrace{[\widetilde{\mathbf{W}_{i}}^{T}]_{t}}^{\widehat{=}\mathbf{v}_{u}} = \sum_{u} x_{u} \mathbf{v}_{u} \quad .$$
(E.5)

The sum on the right hand side in eq. (E.5), in turn, can be expressed as a simple matrix-vector multiplication, such that

$$\mathbf{s} = \sum_{u} x_u \mathbf{v}_u = \mathbf{V} \mathbf{x} \tag{E.6}$$

with v_u as the columns of V. The result is, of course, trivial and only shows that the capsule activations (a weighted sum of the inputs, eq. (E.1)) are obtained as a linear transformation of the input if no dynamic routing is applied.

Finally, we note that subsequent to this linear transformation, the authors in [9] apply the squashing function SQ (see eq. (2) in our main paper) to the output vector \mathbf{s} , which yields the final capsule output CAP(\mathbf{x}):

$$CAP(\mathbf{x}) = SQ(\mathbf{s}(\mathbf{x})) = SQ(\mathbf{V}\mathbf{x})$$
 (E.7)

E.2. Differences to capsules

In the previous section we showed that the activation of a single capsule is computed as a linear transformation of all input activations, which is subsequently squashed (see eq. (E.7)). Comparing this with the computation of the DAU outputs (eq. (1) in our main paper), we see that this is equivalent to the dynamic weight computation if V = AB and b = 0. As such, Dynamic Alignment Units and capsules are related. However, there are important differences, which we discuss in the following and summarise in Table E1.

First, in [9], the squashed capsule activations CAP(x) are used as input to the next layer (after potentially applying the dynamic routing algorithm first). Instead of forwarding the squashed activations directly, in our DAUs they are used to linearly transform the input. Second, by factorising the matrix V into two matrices AB, we are able to control the rank of the linear transformation. Third, when computing the weights in the DAUs we allow for an additional bias term, which in the context of capsules can be considered a 'default vote'. Fourth, we generalise the non-linearity in the DAUs to any non-linearity that only changes the norm. Lastly, we do not apply dynamic routing in the DAUs.

References

- Julius Adebayo, Justin Gilmer, Michael Muelly, Ian J. Goodfellow, Moritz Hardt, and Been Kim. Sanity Checks for Saliency Maps. In Advances in Neural Information Processing Systems (NeurIPS), 2018. 1, 2, 5, 9
- [2] David Baehrens, Timon Schroeter, Stefan Harmeling, Motoaki Kawanabe, Katja Hansen, and Klaus-Robert Müller. How to explain individual classification decisions. *The Journal of Machine Learning Research*, 2010. 9
- [3] Ekin D. Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V. Le. RandAugment: Practical data augmentation with no separate search. *CoRR*, abs/1909.13719, 2019. 8
- [4] Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1936. 12
- [5] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In International Conference on Learning Representations (ICLR), Conference Track Proceedings, 2015. 8
- [6] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In Advances in Neural Information Processing Systems (NeurIPS). Curran Associates, Inc., 2019. 8
- [7] Vitali Petsiuk, Abir Das, and Kate Saenko. RISE: Randomized Input Sampling for Explanation of Black-box Models. In British Machine Vision Conference (BMVC), 2018.
- [8] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. "Why Should I Trust You?": Explaining the predictions of any classifier. In *International Conference on Knowledge Discovery and Data Mining (SIGKDD)*, 2016. 9
- [9] Sara Sabour, Nicholas Frosst, and Geoffrey E. Hinton. Dynamic Routing Between Capsules. In Advances in Neural Information Processing Systems (NeurIPS), 2017. 1, 13, 14
- [10] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. In *International Conference* on Computer Vision (ICCV), 2017. 9
- [11] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning Important Features Through Propagating Activation Differences. In *International Conference on Machine Learning (ICML)*, 2017. 9
- [12] Suraj Srinivas and François Fleuret. Full-Gradient Representation for Neural Network Visualization. In Advances in Neural Information Processing Systems (NeurIPS), 2019. 9
- [13] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic Attribution for Deep Networks. In Doina Precup and Yee Whye Teh, editors, *International Conference on Machine Learning (ICML)*, 2017. 9
- [14] Matthew D. Zeiler and Rob Fergus. Visualizing and Understanding Convolutional Networks. In European Conference on Computer Vision (ECCV), 2014. 9
- [15] Jianming Zhang, Sarah Adel Bargal, Zhe Lin, Jonathan Brandt, Xiaohui Shen, and Stan Sclaroff. Top-Down Neural Attention by Excitation Backprop. Int. J. Comput. Vis., 2018. 9