Towards deep compositional networks

Tabernik, Domen; Kristan, Matej; Wyatt, Jeremy; Leonardis, Ales

DOI:

[10.1109/ICPR.2016.7900171](https://doi.org/10.1109/ICPR.2016.7900171)

License:

Other (please specify with Rights Statement)

Document Version

Peer reviewed version

Citation for published version (Harvard):

Tabernik, D, Kristan, M, Wyatt, J & Leonardis, A 2017, Towards deep compositional networks. in *Proceedings ICPR 2016: 23rd International Conference on Pattern Recognition*. IEEE Computer Society, 23rd International Conference on Pattern Recognition, Cancun, Mexico, 4/12/16. <https://doi.org/10.1109/ICPR.2016.7900171>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

(c) 2016 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other users, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works for resale or redistribution to servers or lists, or reuse of any copyrighted components of this work in other works.

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Towards Deep Compositional Networks

Domen Tabernik^{*†}, Matej Kristan[†], Jeremy L. Wyatt^{*}, Aleš Leonardis^{*†}

^{*}CN-CR Centre, School of Computer Science,
University of Birmingham, Birmingham, UK

[†]Faculty of Computer and Information Science,
University of Ljubljana, Ljubljana, Slovenia

{domen.tabernik,matej.kristan,ales.leonardis}@fri.uni-lj.si
{a.leonardis,jlw}@cs.bham.ac.uk

Abstract—Hierarchical feature learning based on convolutional neural networks (CNN) has recently shown significant potential in various computer vision tasks. While allowing high-quality discriminative feature learning, the downside of CNNs is the lack of explicit structure in features, which often leads to overfitting, absence of reconstruction from partial observations and limited generative abilities. Explicit structure is inherent in hierarchical compositional models, however, these lack the ability to optimize a well-defined cost function. We propose a novel analytic model of a basic unit in a layered hierarchical model with both explicit compositional structure and a well-defined discriminative cost function. Our experiments on two datasets show that the proposed compositional model performs on a par with standard CNNs on discriminative tasks, while, due to explicit modeling of the structure in the feature units, affording a straight-forward visualization of parts and faster inference due to separability of the units.

I. INTRODUCTION

Over the last ten years, computer-vision-based object perception has continuously been moving away from using hand-crafted features like HOG [1] and SIFT [2], and significant efforts have been made on feature learning. An important characteristics of state-of-the-art feature learning methods [3], [4], [5], [6] can be found in the graduated complexity of the features as layers are added to the network, thus forming a hierarchy.

Two paradigms have emerged in the design of hierarchical models, that differ in the definition of a unit learned at each layer. The first paradigm is that of compositional models [7], [8]. We refer to models as being compositional when its features at a layer are modeled as an explicit combination of features in the lower layer. These models allow fast inference via inverted indexing, inherently produce region proposals, offer straight-forward reconstruction from partial observations and visualization of features. However, the learning cost function is weakly defined and is usually performed via co-occurrence learning [7], [8]. While attempts have been made to discriminatively re-interpret the learned reconstructive parts [9], [10], learning discriminative parts directly remains an open issue. On the other hand, the paradigm of convolutional models has recently gained significant momentum in feature learning. These models define the feature units at each layer as filters, which afford learning by back-propagation. In particular, convolutional neural networks (CNNs) [4], [3] have emerged as a highly successful representative of this

Figure 1. Units learned by our deep compositional network (left) and equivalent units learned by the convolutional neural network (right).

class of hierarchical models. A longstanding criticism of CNNs is the lack of precise spatial relationships between high-level parts, a reason advocated to move towards viewpoint-invariant capsule-like systems [11]. The lack of structure in basic units in CNNs also hinders robust handling of occlusions and missing parts [12], [10], and prohibits straight-forward visualization and understanding of networks. Approximate visualization techniques have been developed [13], [14], [15], but further attempts at understanding CNNs uncovered unintuitive behavior when small input perturbations were applied [16]. Recent efforts have also been made toward approximation of the learned filters to reduce the computational complexity of the learned CNNs [17], an issue typically addressed by brute force CPU/GPU power increase.

In this paper we propose a novel form of the unit in a deep hierarchical model with an explicit compositional structure (Fig. 1). While stacked weights in deep networks can be considered as compositions, they lack an explicit structure that could further expose and leverage compositional properties. Our proposed unit exposes an explicit structure of compositions as a parametric model over spatial clustering of responses from the lower layer, and can directly be embedded into

the learning framework used in CNNs. This allows learning of hierarchical compositional models with a well-defined, potentially discriminative, cost function similar to that used in convolutional neural networks, and retains the benefits of compositional models, such as a precise encoding of the spatial relationship between parts. We derive the necessary equations for back-propagation and propose a compositional model trained by a discriminative cost function, which is the major contribution of our paper. We experimentally evaluate the proposed model on CIFAR-10 [18] and PaCMan [19] datasets and show that our model achieves comparable performance to a standard CNN, while allowing simple compositional mean-reconstruction of parts. Since our units are separable by design, we demonstrate a significant speedup in inference compared to CNNs.

The paper is structured as follows: Section II describes our model, Section III provides experimental evaluation, in Section IV visualization and speed-up in inference are presented, and conclusions are drawn in Section V.

II. DEEP COMPOSITIONAL NETWORK

We first provide notation for deep convolutional neural networks and then derive our compositional network. A neuron activation in a deep neural network is modeled with a linear function wrapped inside a specific non-linearity: $y_i = f(\sum_s w_s \cdot x_s + b_s)$, with output activation y , input activations x , bias b and weights w determining a linear combination of input activations further modified by a non-linear function $f(\cdot)$. In the image domain the neuron activations are organized in a 3-dimensional matrix $N \times M \times S$, where two dimensions, $N \times M$, represent the image or feature plane, and the third dimension, S represents channels. With the introduction of weight-sharing along the 2D feature plane, the convolutional neural network models the neuron's activation function Y as:

$$Y_i = f(\sum_s W_s * X_s + b_s), \quad (1)$$

where $*$ is a convolution of X_s , the $N \times M$ activation map from the s -th channel, with W_s , the $K_w \times K_h$ weights for the s -th channel. The W_s are basic units in the CNNs that take the form of convolution filters and have to be learned from the data. The convolutions give an $N - K_w \times M - K_h$ output map for each channel. Element-wise summation over the channels represents the final activation map Y_i after non-linearity $f(\cdot)$ is applied. Typically, several activation maps Y_i are created, and the network is organized in layers, such that X_s in the l -th layer is the output activation Y_i from $l - 1$ layer.

We propose a new basic unit that explicitly models the composition of a feature. We define it as a weighted Gaussian component:

$$\tilde{W}_k = \tilde{w}_k G(\theta_k), \quad (2)$$

with weight $\tilde{w}_k$ and Gaussian parameters $\theta_k = [\vec{\mu}_k, \sigma_k]$, containing mean $\vec{\mu}_k$ and variance σ_k^2 . Multiple $\tilde{W}_k$ can be applied to the same input channel s , but we omit this in the notation in the interest of clarity. Units $\tilde{W}_k$ take the form of

convolution filters and have $K_w \times K_h$ elements. We therefore define $G(\cdot)$ as a two dimensional matrix of the same size with 2-dimensional index $\vec{x}$ over $K_w \times K_h$ elements:

$$G(\vec{x}, \theta) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp(-\frac{\|\vec{x} - \vec{\mu}\|^2}{2\sigma^2}). \quad (3)$$

We use two dimensional means but single dimensional variance for simplification. Commonly used unit sizes are 3×3 , 5×5 or 9×9 ; however, such small sizes can lead to significant discretization errors in $G(\vec{x}, \theta_k)$. We avoid this by replacing the normalization factor computed in continuous space with one computed in the discretized space, leading to our final distribution function $G(\vec{x}, \theta)$:

$$G(\vec{x}, \theta) = \frac{1}{N(\theta)} g(\vec{x}, \theta), \quad (4)$$

where $g(\vec{x}, \theta)$ is a non-normalized Gaussian distribution and $N(\theta)$ is a sum over this non-normalized Gaussian distribution computed for a filter of size $K_w \times K_h$:

$$N(\theta) = \sum_{\vec{x}} g(\vec{x}, \theta), \quad g(\vec{x}, \theta) = \exp(-\frac{\|\vec{x} - \vec{\mu}\|^2}{2\sigma^2}). \quad (5)$$

A new, compositional unit can be embedded into a CNN by grouping multiple instances applied to the same input channel s and deriving the basic CNN unit from Eq. (1):

$$W_s = \sum_k \tilde{W}_k = \sum_k \tilde{w}_k G(\theta_k). \quad (6)$$

This proposed model for each W_s unit is similar to a standard Gaussian mixture model, but we do not enforce $\sum \tilde{w} = 1$ and component weights can take any value, $\tilde{w} \in [-\infty, \infty]$. Having negative weights is important to approximate edges as differences between neighboring components, whereas positive components can be interpreted as requirements for a presence of a feature and negative components as requirements for an absence of a feature. In principle, this could be satisfied by normalizing the sum of absolute weights, $\sum |\tilde{w}| = 1$, but this would significantly complicate the gradient computation without any performance gain.

Learning an individual unit $\tilde{W}_k$ consists of learning its parameters for the Gaussian distribution, mean $\vec{\mu}_k$ and variance σ_k^2 , together with the weight $\tilde{w}_k$. The number of units, i.e., the number of Gaussian components per input channel, can be considered as a hyper-parameter. Learning can be performed in the same way as in convolutional networks via gradient descent. Parameters are optimized by computing the gradients w.r.t. the cost function $C(y, \vec{y})$, which leads to three different types of gradients. By applying the chain rule we can define the gradient for the component weight $\partial C / \partial \tilde{w}_k$ as a dot-product of back-propagated error and the input feature X_s convolved with the k -th Gaussian component:

$$\frac{\partial C}{\partial \tilde{w}_k} = \sum_{n,m} \frac{\partial C}{\partial Z} \cdot \frac{\partial Z}{\partial \tilde{w}_k} = \sum_{n,m} \frac{\partial C}{\partial Z} \cdot \sum_{\vec{x}} X_s * G(\vec{x}, \theta), \quad (7)$$

where $Z = \sum_s W_s * X_s + b_s$ and $\partial C / \partial Z$ is back-propagated error. Note, that only the s -th channel of input features are

Figure 2. Classification results on CIFAR-10 database.

used since the weight component $\tilde{w}_k$ appears only in W_s . The back-propagated error for layer l is computed the same as in a standard convolutional network:

$$\frac{\partial C}{\partial Z_s^l} = \frac{\partial C}{\partial Z_s^{l+1}} * \text{rot}(W_s), \quad (8)$$

where the back-propagated error from the higher layer $l+1$ is convolved with the 180° rotated unit (a weight filter) $\text{rot}(W_s)$ which can be computed from Eq. 6. We can similarly apply the chain rule to obtain the gradient for the mean and the variance:

$$\frac{\partial C}{\partial \mu_k} = \sum_{n,m} \frac{\partial C}{\partial Z} \cdot \sum_{\vec{x}} X_s * \frac{\partial G(\vec{x}, \theta_k)}{\partial \mu_k}, \quad (9)$$

$$\frac{\partial C}{\partial \sigma_k} = \sum_{n,m} \frac{\partial C}{\partial Z} \cdot \sum_{\vec{x}} X_s * \frac{\partial G(\vec{x}, \theta_k)}{\partial \sigma_k}, \quad (10)$$

where the derivatives of the Gaussian are:

$$\frac{\partial G(\vec{x}, \theta_k)}{\partial \mu_k} = \tilde{w}_k \frac{N(\theta_k) \cdot \frac{g(\vec{x}, \theta_k)}{\partial \mu_k} - g(\vec{x}, \theta_k) \cdot \frac{\partial N(\theta_k)}{\partial \mu_k}}{[N(\theta_k)]^2}, \quad (11)$$

$$\frac{\partial G(\vec{x}, \theta_k)}{\partial \sigma_k} = \tilde{w}_k \frac{N(\theta_k) \cdot \frac{g(\vec{x}, \theta_k)}{\partial \sigma_k} - g(\vec{x}, \theta_k) \cdot \frac{\partial N(\theta_k)}{\partial \sigma_k}}{[N(\theta_k)]^2}. \quad (12)$$

III. CLASSIFICATION PERFORMANCE

This section analyzes the discriminative properties of our deep compositional network. The method is evaluated on two classification datasets, CIFAR-10 [18] and PaCMan [19].

The evaluation on both datasets is performed with a network containing three layers. The first two layers are convolutional/compositional and the third one is fully-connected. We use soft-max with multinomial logistic loss as the cost function. Either three-channel RGB (CIFAR-10) or a single-channel gray-scale image (PaCMan) is used as input data with zero-mean normalization. The data is not normalized to unit variance, since between each layer we use ReLU non-linearity which is less sensitive to data variance. Note that we use slightly bigger filters (basic units) in our network, but use fewer components to approximately match the number of parameters with the standard CNN model. We also restrict

Table I
CONFIGURATION TABLE FOR NEURAL NETWORK USED ON CIFAR-10 DATABASE.

Deep Compositional Network (our)				standard CNN	
	num features	stride	unit/filter size	num components	unit/filter size
conv1	32	1	7×7	2×2	5×5
relu1	/	1	/	/	/
pool1	/	1	3×3		
conv2	32	1	9×9	3×3	5×5
relu2	/	1	/	/	/
pool2	/	2	3×3		
ip1	10	/	15×15	/	15×15

components' positions and standard deviation to ensure derivative of the resulting filters would not have non-zero values outside of the valid window. This also prevents collapsing to a single point and stalling. Positions are restricted to at least 1.5 pixels away from the borders of the valid window, and standard deviations are restricted to $\sigma_s > 0.5$. We apply the AdaDelta [20] with momentum of 0.8 and no weight-decay to achieve proper behavior in gradient descent.

A. Classification on CIFAR-10

The CIFAR-10 [18] dataset consists of 60,000 images split into 50,000 training and 10,000 testing images. We perform training with a mini-batch size of 100 images per iteration and run learning for 5000 iterations. A detailed configuration of the network used is shown in Table I. Comparing performance of both models, shown in Fig. 2, we can see the same accuracy with both achieving slightly less than 70% on testing set. The left-most two graphs reveal similar learning rates with standard CNN learning slightly faster, but they both converge to a similar loss in the end.

B. Classification on PaCMan database

The proposed compositional network was also evaluated on the PaCMan dataset [19]. This dataset contains gray-scale and depth images generated from 3D models of 20 categories of various kitchen objects, with each category containing 20 different instances of objects. Each object is captured at dense, regular viewpoint intervals, but we use only 28 different viewpoints, summing to a total of around 50,000 images. We use only gray-scale images. The dataset was split into

Figure 3. Evaluation results on PaCMan database. Note, only a subset of first-layer and second-layer features are shown. Each first-layer feature consists of one filter, and each second-layer feature consists of 16 filters. Visualized weights trained with Gaussian modeling have component’s means depicted as small circles, and variance as big circle. Note, background color within each feature (i.e. individual patches in first layer and patches within column for the second layer) is considered as zero value. Background colors differ between different features due to different max and min range of weights.

Table II
CONFIGURATION TABLE FOR NEURAL NETWORK USED ON PACMAN DATABASE.

	Deep Compositional Network (our)				standard CNN
	num features	stride	unit/filter size	num components	unit/filter size
conv1	16	1	9×9	2×2	5×5
relu1	/	1	/	/	/
pool1	/	1	3×3	/	/
conv2	16	1	9×9	2×2	5×5
relu2	/	1	/	/	/
pool2	/	2	3×3	/	/
ip1	20	/	15×15	/	15×15

approximately 25,000 samples for training and 25,000 for testing. We ensure that all viewpoints of the same object are in the same split and each category has proportionally the same number of objects in testing and training. The input images are resized to 128×96 to fit the network into the GPU. The network configuration as shown in Table II is slightly modified to accommodate higher resolution images.

Fig. 3 confirms that the proposed compositional network achieves discriminative performance on a par with the CNN – both models attain an accuracy of 64-67%.

We also visualize CNN filters and units in our network for features on the first and the second layer in Fig. 3. The units in our network applied to the same input channel are visualized in a single filter based on Eq. (6), while individual component’s means are plotted as small circles, and variances as large circles. Filters in the standard CNN have a certain

structure but they are still noisy and incoherent. It would also be difficult to capture this structure without human interpretation. On the other hand, our unit with Gaussian models explicitly captures the spatial structure as can be seen in the first-layer filters. Many components converge to the same location and the configuration of components directly points to different edge or blob detectors. On the second layer only a small set of components have high weights, indicating that most are irrelevant for the final classification, offering further simplifications of the network. Compared to the filters from the standard CNN they are more compact. With Gaussian modeling a spatial position of a sub-feature is much clearer and easily determined, and, as we show in the next section, can be further utilized.

IV. UTILIZING COMPOSITIONAL REPRESENTATION

In this section we demonstrate two advantages of having a rich, compositional representation. We propose a novel visualization of deep networks using mean reconstruction of compositions and demonstrate the inference speed-up due to inherent kernel separability in compositions.

A. Visualization by mean reconstruction

Feature visualization in standard deep networks is difficult due to lack of structure. Visualization is thus usually performed indirectly by deconvolution [13] or optimization of input pixels to maximize the output activations [14], [15]. Such visualization techniques are applicable to our model as well, but having explicit compositions enables exploration of more straight-forward visualization techniques.

Figure 4. Mean reconstruction of 8 features from the second layer of PaCMan dataset. The first and the third rows depict compositions and their variances back-projected to pixel values. They can be considered as detectors for specific blob configurations, where yellow values request high positive values, and dark blue, low negative values. 3D visualization of the last feature is also on the left. The second and the fourth rows then show corresponding edge configurations between positive (yellow) and negative (blue) components.

A feature can be visualized by finding an image that produces a maximum output response for its unit. Based on Eq. (1) a maximum response is obtained when all sub-features match to specific patterns defined in weights W_s . In our model each W_s consists of individual compositions $\tilde{W}_k$, thus the premise can further be extended to having a maximum response when individual sub-features are present at specific position, i.e., at a mean of a Gaussian in our case. We propose to visualize a single feature by recursively projecting compositions top-down to image pixels by following (indexing) the corresponding means in network units. We term this process as mean reconstruction, similar to visualization techniques for hierarchical compositions [21].

During each step of back-projection, several properties of the compositions need to be correctly accounted for: (a) expected uncertainty of the position of a sub-feature (variance), (b) the importance of a sub-feature (weight) and (c) requested presence or absence of a sub-feature (weight sign). The uncertainty is defined by a Gaussian variance and grows with each step of back-projection. We account for this by summing the variances along each step to arrive at the final uncertainty at the pixel level. The importance of a sub-feature is accounted for by multiplying the magnitude of weight at each back-projected step. We consider only sub-feature components with a positive weight, i.e., ones that request the presence of a sub-feature. Sub-feature components with negative weights are ignored. This is applied to all layers, except the first one, since negative values of features are truncated by the ReLU layer on all layers except for the first one. We therefore consider the first layer as a special case and use compositions with negative weights as well.

Positive and negative weights in the first layer in most cases define edges. But edges are defined indirectly since means in Gaussians define positions of blobs, i.e., regions with low intensity or high intensity values. Edges can be inferred from neighboring blobs with opposing signs and are free to occur anywhere between blobs, either as a smooth transition or a sharp edge. Consequently, after all compositions are back-projected from the top to the bottom we are left with two

Figure 5. Forward-pass in second (left) and speed-up factor separable implementation (right). Shown for pruned second layer PaCMan model. Note, the values are shown in the logarithmic scale.

sets of Gaussian distributions, ones with the positive sign and ones with the negative sign. They are visualized in the second and the fourth row of Fig. 4. This visualization is performed by summing over all positive and negative distributions for each pixel. We refer to a map obtained this way as a reconstructed blobs or distribution map.

We can further visualize Gaussian compositions by their boundaries that separate positive and negative components. We achieve this by translating the problem to a graph-cut problem, where pixels are represented as a graphical model connected to neighbors to either a sink or a source. We use a sum over all positive distributions for one pixel as a cost for sink and a sum over all negative distributions as a cost for source. The cost for the two neighboring pixels is considered as a squared difference between the distribution maps for that pixel. The optimal edge between them is obtained by finding a minimal cut that maximizes the flow in a graph. We highlight edges with strong borders between positive and negative distributions by multiplying with a difference of neighboring pixels in the distribution maps. The resulting image is finally re-sized and is visualized in the first and the third row of Fig. 4.

The proposed visualization is applied to the second layer features trained on PaCMan database. Most features are still representing edges at different orientations, but some (e.g., 4th and 8th feature) are compositions of edges in form of corners as is evident from the reconstructed boundaries. The features are visualized on a pruned model, i.e, we merged any overlapping compositions in each sub-feature to remove duplicate components, and compositions with weights below a 2% of a maximal weight are discarded since they do not contribute to the final score. With pruning we removed or merged approximately 400 out of 1024 Gaussians at the second layer, while reducing the score by less than 1%. This pruning process is another benefit of our explicit compositions, which can lead to a network with significantly reduced complexity.

B. Inference speed-up with separable filters

Another advantage of our deep compositional network is the ability to decompose the units into separate filters by leveraging the separability of Gaussians to speed-up the inference. In this case, the complexity for the forward pass at a single layer is reduced from $O(S \cdot F \cdot W \cdot H \cdot K_w \cdot K_h)$ for the standard CNN to $O(S \cdot F \cdot G \cdot (W \cdot H \cdot (K_w + K_h) + c))$ for implementation with separable filters, with S sub-features, F features, G number

of Gaussian components per sub-feature, $W \times H$ feature map size, $K_w \times K_h$ kernel filter size and c as an additional overhead in a separable filter implementation. An important factor in this separable implementation is the number of Gaussian components G , but we can reduce this number with the same pruning process described in the previous section, where we reduced the number of components by almost a half.

Based on time complexities a separable implementation should gain significant speed-up for $G \cdot (K_w + K_h) < K_w \cdot K_h$. We evaluate this separable implementation on a pruned PaC-Man model considering different kernel sizes. Results are depicted in Fig. 5. We achieve a slight speed-up at kernel sizes of 10×10 while a significant speed-up of 3-fold or more is achieved with kernel sizes of 15×15 pixels or bigger.

We use the Caffe implementation with convolution as matrix multiplication using the CBLAS library and implement separable convolution as multiple calls to AXPY methods using the same CBLAS library. The demonstrated speed-up factors are fairly conservative considering the Caffe implementation performs 2D convolution with a fully optimized single call while our implementation adds some overhead with multiple calls. We evaluated only a CPU implementation and enforced a single-core process. Since both implementations have a similar level of parallelism it is fair to assume that speed-up can be maintained in multicore CPU or GPU implementations as well.

V. CONCLUSION

A new deep compositional network is introduced in this paper. The new network is based on a novel form of an element unit (a filter) that applies a parametric model. We demonstrated that parametrization with Gaussian distributions retains the spatial structure of compositions of features and affords learning by optimizing a well-defined cost function. We derived the necessary equations for back-propagation and embedded our model into a deep neural network framework to evaluate discriminative learning of compositional parts on CIFAR-10 and PaCMan datasets. We showed that having a compositional representation is advantageous for deep networks by presenting a novel visualization and an inference speed-up. We performed visualization of deep network features using a mean reconstruction of parts. Other visualization techniques for deep networks typically rely on approximation with de-convolution [13] or a complex optimization [14], and need to process additional data. In contrast, the compositional representation allowed us to generate a representation of a feature with a fairly simple technique that uses only the model itself and no additional data. A simpler visualization and the 3-fold speed-up of inference speak of the advantages of using the new parametric units in deep networks with a convolutional layered architecture.

In the future we plan to explore other venues opened by combining compositional and convolutional hierarchies, such as, pre-training with co-occurrence statistics, performing generative and discriminative learning concurrently, and further

leveraging inverted indexing of parts for inference.

REFERENCES

- [1] N. Dalal and B. Triggs, "Histograms of Oriented Gradients for Human Detection," in *CVPR*, 2005, pp. 886–893.
- [2] D. Lowe, "Object recognition from local scale-invariant features," in *Proceedings of the Seventh IEEE ICCV*, 1999, pp. 1150–1157.
- [3] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," in *Advances In Neural Information Processing Systems* 25, 2012, pp. 1097–1105.
- [4] Y. Lecun, S. Chopra, R. Hadsell, M. A. Ranzato, and F. J. Huang, "A Tutorial on Energy-Based Learning," pp. 1–59, 2006.
- [5] A. Leonardis and S. Fidler, "Learning hierarchical representations of object categories for robot vision," *Robotics Research*, vol. 66, pp. 99–110, 2011.
- [6] L. L. Zhu, Y. Chen, and A. Yuille, "Learning a hierarchical deformable template for rapid deformable object parsing," *IEEE transactions on pattern analysis and machine intelligence*, vol. 32, no. 6, pp. 1029–43, 6 2010.
- [7] S. Fidler and A. Leonardis, "Towards Scalable Representations of Object Categories: Learning a Hierarchy of Parts," in *Computer Vision and Pattern Recognition*. IEEE Computer Society, 2007, pp. 1–8.
- [8] Z. Si and S.-c. Zhu, "Learning AND-OR Templates for Object Recognition and Detection," *IEEE Transactions on Pattern analysis and machine intelligence*, vol. 35, no. 9, pp. 2189–2205, 2013.
- [9] M. Kristan, M. Boben, D. Tabernik, and A. Leonardis, "Adding discriminative power to hierarchical compositional models for object class detection," in *18th Scandinavian Conference on Image Analysis*, 2013, pp. 1–12.
- [10] D. Tabernik, A. Leonardis, M. Boben, D. Skočaj, and M. Kristan, "Adding discriminative power to a generative hierarchical compositional model using histograms of compositions," *Computer Vision and Image Understanding*, vol. 138, pp. 102–113, 9 2015.
- [11] G. E. Hinton, A. Krizhevsky, and S. D. Wang, "Transforming auto-encoders," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 6791, no. 1, pp. 44–51, 2011.
- [12] A. Yuille and R. Mottaghi, "Complexity of Representation and Inference in Compositional Models with Part Sharing," *International Conference on Learning Representations*, vol. 17, pp. 1–13, 2016.
- [13] M. D. Zeiler and R. Fergus, "Visualizing and Understanding Convolutional Networks," in *European Conference on Computer Vision*, 11 2013, pp. 818–833. [Online]. Available: <http://arxiv.org/abs/1311.2901>
- [14] A. Mahendran and A. Vedaldi, "Understanding Deep Image Representations by Inverting Them," in *Computer Vision and Pattern Recognition*, 2015, pp. 1–9.
- [15] —, "Visualizing Deep Convolutional Neural Networks Using Natural Pre-Images," 2016. [Online]. Available: <http://arxiv.org/abs/1512.02017>
- [16] C. Szegedy, W. Zaremba, and I. Sutskever, "Intriguing properties of neural networks," in *International Conference on Learning Representations*, 2014, pp. 1–9.
- [17] M. Jaderberg, A. Vedaldi, and A. Zisserman, "Speeding up Convolutional Neural Networks with Low Rank Expansions," in *British Machine Vision Conference*, 2014, p. 7.
- [18] A. Krizhevsky, "Learning Multiple Layers of Features from Tiny Images," *Science Department, University of Toronto, TechReport*, pp. 1–60, 2009.
- [19] "Synthetic RGB-D images of 400 objects from 20 classes. Generated from 3D mesh models." [Online]. Available: <http://www.pacman-project.eu/datasets/>
- [20] M. D. Zeiler, "ADADELTA: An Adaptive Learning Rate Method," 2012. [Online]. Available: <http://arxiv.org/abs/1212.5701>
- [21] S. Fidler, M. Boben, and A. Leonardis, "Optimization framework for learning a hierarchical shape vocabulary for object class detection," in *Proceedings of the British Machine Vision Conference 2009*, 2009, pp. 1–93.

Acknowledgments. This work was supported in part by EU IST-600918 project PaCMan and ARRS research project L2-6765. We also acknowledge MoD/Dstl and EPSRC for providing the grant to support the UK academics (Aleš Leonardis) involvement in a Department of Defense funded MURI project.