

Detection of Human Rights Violations in Images: Can Convolutional Neural Networks Help?

Grigorios Kalliatakis¹, Shoaib Ehsan¹, Maria Fasli¹, Ales Leonardis², Juergen Gall³
and Klaus D. McDonald-Maier¹

¹*School of Computer Science and Electronic Engineering, University of Essex, Colchester, U.K.*

²*School of Computer Science, University of Birmingham, Birmingham, U.K.*

³*Institute of Computer Science III, University of Bonn, Bonn, Germany*
{gkallia, sehsan, mfasli, kdm}@essex.ac.uk, a.leonardis@cs.bham.ac.uk, gall@iai.uni-bonn.de

Keywords: Convolutional Neural Networks, Deep Representation, Human Rights Violations Recognition.

Abstract: After setting the performance benchmarks for image, video, speech and audio processing, deep convolutional networks have been core to the greatest advances in image recognition tasks in recent times. This raises the question of whether there are any benefit in targeting these remarkable deep architectures with the unattempted task of recognising human rights violations through digital images. Under this perspective, we introduce a new, well-sampled human rights-centric dataset called Human Rights Understanding (HRUN). We conduct a rigorous evaluation on a common ground by combining this dataset with different state-of-the-art deep convolutional architectures in order to achieve recognition of human rights violations. Experimental results on the HRUN dataset have shown that the best performing CNN architectures can achieve up to 88.10% mean average precision. Additionally, our experiments demonstrate that increasing the size of the training samples is crucial for achieving an improvement on mean average precision principally when utilising very deep networks.

1 INTRODUCTION

Human rights violations continue to take place in many parts of the world today, while they have been ongoing during the entire human history. These days, organizations concerned with human rights are increasingly using digital images as a mechanism for supporting the exposure of human rights and international humanitarian law violations. However, utilising current advances in technology for studying, prosecuting and possibly preventing such misconduct from occurring have not yet made any progress. From this perspective, supporting human rights is seen as one scientific domain that could be strengthened by the latest developments in computer vision. To support the continued growth of images and videos in human rights and international humanitarian law monitoring campaigns, this study examines how vision based systems can support human rights monitoring efforts by accurately detecting and identifying human rights violations utilising digital images.

This work is made possible by recent progress in Convolutional Neural Networks (CNNs) (LeCun et al., 1989), which has changed the landscape for

well-studied computer vision tasks, such as image classification and object detection (Wang et al., 2010; Huang et al., 2011), by comprehensively outperforming the initial handcrafted approaches (Donahue et al., 2014; Sharif Razavian et al., 2014; Sermanet et al., 2013). These state-of-the-art architectures are now finding their way into a number of vision based applications (Girshick et al., 2014; Oquab et al., 2014; Simonyan and Zisserman, 2014a).

A major contribution of our paper is a new, well-sampled human rights-centric dataset, called the *Human Rights Understanding (HRUN)* dataset, which consists of 4 different categories of human rights violations and 100 diverse images per category. In this paper, we formulate the human rights violation recognition problem as being able to recognise a given input image (from the HRUN dataset) as belonging to one of these 4 categories of human rights violations. See Figure 1 for examples of our data. We use this data for human rights understanding by evaluating different deep representations on this new dataset, while we perform experiments that illustrate the effect of network architecture, image context, and training data size on the accuracy of the system.



Figure 1: Examples from all 4 categories of the Human Rights Understanding (HRUN) Dataset.

In summary, our contribution is two-fold. Firstly we introduce a new human rights-centric dataset HRUN. Secondly, motivated by the great success of deep convolutional networks, we conduct a large set of rigorous experiments for the task of recognising human rights violations. As part of our tests, we delve into the latest, top-performing pre-trained deep convolutional models, allowing a fair, unbiased comparison on a common ground; something that has been largely missing so far in the literature. The remainder of the paper is organised as follows. Section 2 looks into prior works on database construction and image understanding with deep convolutional networks. Section 3 describes the methodology utilised for building our pioneer human rights understanding dataset. Section 4 demonstrates the classification pipeline used for the experiments, while the evaluation results are presented in Section 5, alongside a thorough discussion. Finally, conclusions and future directions are given in Section 6.

2 PRIOR WORK

2.1 Database Construction

Challenging databases are important for many areas of research, while large-scale datasets combined with CNNs have been key to recent advances in computer vision and machine learning applications. While the field of computer vision has developed several databases to organize knowledge about object categories (Deng et al., 2009; Griffin et al., 2007; Torralba et al., 2008; Fei-Fei et al., 2007), scenes (Zhou et al., 2014; Xiao et al., 2010) or materials (Liu et al., 2010; Sharan et al., 2009; Bell et al., 2015) a well-inspected dataset of images depicting human rights violations does not currently exist. The first reference point in standardized dataset of images and annotations was the VOC2010 dataset (Everingham et al., 2010), which was constructed by utilizing images collected by non-vision/machine learning researchers, by querying Flickr with a number of related keywords, including the class name, synonyms and scenes or situations where the class is likely to appear. Similarly, an extensive scene understanding (SUN) database was introduced by (Xiao et al., 2010), containing 899

environments and 130,519 images. The primary objectives of this work were to build the most complete dataset of scene image categories. Microsoft's work in regard to detection and segmentation of objects taking place in their natural context, was marked with the introduction of common objects in context (Lin et al., 2014) (MS COCO) dataset including 328,000 images of complex everyday scenes consisted of 91 different object categories and 2.5 million labelled instances. More recently (Yu et al., 2015) presented their first version of a scene-centric database (LSUN) with millions of label images in each category alongside an integrated framework which makes use of deep learning techniques in order to achieve large-scale image annotation. To our knowledge, this particular work is the first attempt to construct a well-sampled image database in the domain of human rights understanding.

2.2 Deep Convolutional Networks

For decades, traditional machine learning systems demanded accurate engineering and significant domain expertise in order to design a feature extractor capable of converting raw data (such as the pixel values of an image) into a convenient internal representation or feature vector from which a classifier could classify or detect patterns in the input. Today, representation learning methods and principally CNNs (LeCun et al., 1989) are driving advances at a dramatic pace in the computer vision field after enjoying a great success in large-scale image recognition and object detection tasks (Krizhevsky et al., 2012; Sermanet et al., 2013; Simonyan and Zisserman, 2014a; Tompson et al., 2015; Taigman et al., 2014; LeCun et al., 2015). The key aspect of deep learning representations is that the layers of features are not manually hand-crafted, but are learned from data using a generic-purpose learning scheme. The architecture of a typical deep-learning system can be considered as a multilayer stack of simple modules, each one transforming its input to increase both the selectivity and the invariance of the representation as stated in (LeCun et al., 2015). In the last few years vision tasks became feasible due to high-performance computing systems such as GPUs, extensive public image repositories (Deng et al., 2009), a new regularisation tech-

nique called *dropout* (Srivastava et al., 2014) which prevents deep learning systems from overfitting, rectified linear units (ReLU) (Nair and Hinton, 2010), softmax layer and techniques able to generate more training examples by deforming the existing ones.

Since (Krizhevsky et al., 2012) first used an eight layer CNN (also known as AlexNet) trained on ImageNet to perform 1000-way object classification, a number of other works have used deep convolutional networks (ConvNets) to elevate image classification further (Simonyan and Zisserman, 2014b; He et al., 2015a; Szegedy et al., 2015; He et al., 2015b; Chatfield et al., 2014). (Simonyan and Zisserman, 2014b) use a very deep CNN (also known as VGGNet) with up to 19 weight layers for large-scale image classification. They demonstrated that a substantially increased depth of a conventional ConvNet (LeCun et al., 1989; Krizhevsky et al., 2012) can result in state-of-the-art performance on the ImageNet challenge dataset (Deng et al., 2009). They also perform localization for the same challenge by training a very deep ConvNet to predict the bounding box location instead of the class scores at the last fully connected layer. Another deep network architecture that has been recently used to great success is the GoogLeNet model of (Szegedy et al., 2015) where an inception layer is composed of a shortcut branch and a few deeper branches in order to improve utilization of the computing resources inside the network. The two main ideas of that architecture are: (i) to create a multi-scale architecture capable of mirroring correlation structure in images and (ii) dimensional reduction and projections to keep their representation sparse along each spatial scale. Most recently (He et al., 2015a) announced the even deeper residual network (also known as ResNet), featured 152 layers, which has considerably improved the state-of-the-art performance of ImageNet (Deng et al., 2009) classification and object detection on PASCAL (Everingham et al., 2010). Residual networks are inspired by the observation that neural networks lean towards gaining higher training errors as the depth of the network increases to very large values. The authors argue that although the network gains more parameters by increasing its depth, the network becomes inferior at function approximation because of the gradients and training signals loss when they are propagated through numerous layers. Therefore, they give convincing theoretical and practical evidence that residual connections (re-formulated layers for learning residual functions with reference to the layer input) are inherently necessary for training very deep convolutional models.

Outside of the aforementioned top-performing networks, other works worth mentioning are: (Chat-

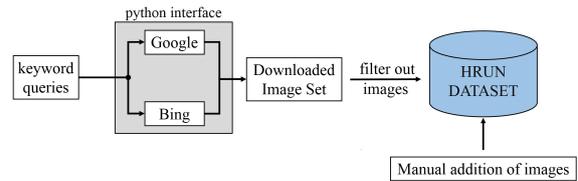


Figure 2: Constructing HRUN Dataset.

field et al., 2014) where a rigorous evaluation study on different CNN architectures for the task of object recognition was conducted and (Zhou et al., 2014) where a brand-new scene-centric database called Places was introduced and established state-of-the-art results on different scene recognition tasks, by learning deep representations from their extensive database. Despite these impressive results, human rights advocacy is one of the high profile domains which remain broadly missing from the curated list of problems which were benefited from the continuing growth of deep convolutional networks. We build on this body of work in deep learning to solve the untrodden problem of recognising human rights violations utilising digital images.

3 HRUN DATASET

Recent achievements in computer vision can be mainly ascribed to the ever growing size of visual knowledge in terms of labelled instances of objects, scenes, actions, attributes, and the dependent relationships between them. Therefore, obtaining effective high-level representations has become increasingly important, while a key question arises in the context of human rights understanding: *how will we gather this structured visual knowledge?*

This section describes the image collection procedure utilised for the formulation of the HRUN dataset, as captured by Figure 2.

Initially, the keywords, with a view to formulate the query terms, were collected in collaboration with specialists in the human rights domain. This happens in order to include more than one query term for every ‘targeted class’. For instance, for the class police violence the queries ‘police violence’, ‘police brutality’ and ‘police abuse of force’ were all used for retrieving results. Work commenced with the Flickr photo-sharing website, but in a short time, it

Table 1: Image Collection Analysis from Search Engines.

Query Term	Retrieved Images		Relevant Images			HRUN
	Google	Bing	Google	Bing	Manually	
Child labour	99	137	18	5	77	100
Child Soldiers	176	159	31	13	56	100
Police Violence	149	232	10	16	74	100
Refugees	111	140	10	39	51	100



Figure 3: Side by side examples of irrelevant images with their respective query term which were eliminated during the filtering process.

became apparent that its limitations resulted in a huge number of irrelevant results returned for the given queries. This happens because Flickr users are authorised to tag their uploaded images without restriction. Subsequently there have been situations where the given keyword was ‘armed conflict’ and the majority of the returned images had to do with military parades. Another similar example was with the given keyword ‘genocide’ where the returned results included protesting campaigns against genocide, something that may be consider close to the keyword, but it can not serve our purpose by any means. Another shortcoming was the case when people massively tagged an image deliberately incorrectly in order to acquire an increased number of hits on the photo-sharing website. Consequently, Google and Bing search engines were chosen as a better alternative. Images were downloaded for each class using a python interface to the Google and Bing application programming interfaces (APIs), with the maximum number of images permitted by their respective API for each query term. All exact duplicate images were eliminated from the downloaded image set, alongside images regarded as inappropriate during the filtering step as illustrated by Figure 3. Nonetheless, the number of filtered images generated was still insufficient as shown in Table 1.

For this reason, there were manually added other suitable images in order to reach the final structure of the HRUN dataset. We finally ended up with a total of four different categories, each one containing 100 distinct images of human rights violations captured in real world situations and surroundings. With this first attempt, our main intention was to produce a high quality dataset for the task in hand. For that reason, the number of categories was kept to a certain degree for the time being. Expanding the dataset both in categories and number of images has already been included in our actual future plans and many other online repositories that might be related to human rights

violations are being checked into thoroughly.

4 LEARNING DEEP REPRESENTATIONS FOR HUMAN RIGHTS VIOLATIONS RECOGNITION

4.1 Transfer Learning

Our goal is to train a system that recognises different human rights violations from a given input image of the HRUN dataset. One high-priority research issue in our work is how to find a good representation for instances in such a unique domain. More than that, having a dataset of sufficient size is problematic for this task as described in the previous section. For those two reasons, a conventional alternative to training a deep ConvNet from the very beginning, is to use a pre-trained model and then use the ConvNet as a fixed feature extractor for the task of interest. This method, referred to as *transfer learning* (Donahue et al., 2014; Zeiler and Fergus, 2014), is implemented by taking a pre-trained CNN, replacing the fully-connected layers (and potentially the last convolutional layer), and consider the rest of the ConvNet as a fixed feature extractor for the relevant dataset. By freezing the weights of the convolutional layers, the deep ConvNet can still extract general image features such as edges, while the fully connected layers can take this information and use it to classify the data in a way that is applicable to the problem.

4.2 Pipeline for Human Rights Violations Recognition

The entire pipeline used for the experiments is depicted in Figure 4, and detailed further below. In this pipeline, every block is fixed except the feature extractor as different deep convolutional networks are plugged in, one at a time, to compare their performance utilizing the mean average precision (mAP) metric.

Given a training dataset T_r consisting of m human rights violation categories, a test dataset T_s comprising unseen images of the categories given in T_r , and a set of n pre-trained CNN architectures (C_1, \dots, C_n), the pipeline operates as follows: The training dataset T_r is used as input to the first CNN architecture C_1 . The output of C_1 , as described above, is then utilized to train m SVM classifiers. Once trained, the test dataset T_s is employed to assess the performance

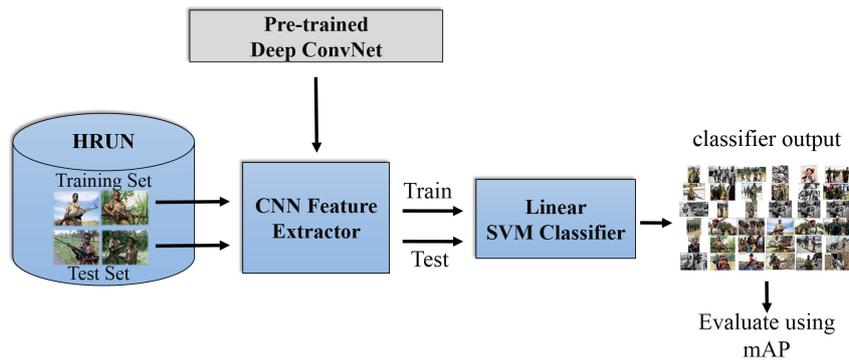


Figure 4: An overview of the human rights violations recognition pipeline used here. Different deep convolutional models are plugged into the pipeline one at a time, while the training and test samples taken from the HRUN dataset remain fixed. Mean average precision(mAP) metric is used for evaluating the results.

of the pipeline using mAP. The training and testing procedures are then repeated after replacing C_1 with the second CNN architecture C_2 to evaluate the performance of the human rights violation recognition pipeline. For a set of n pre-trained CNN architectures, the training and testing processes are repeated n times. Since the entire pipeline is fixed (including the training and test datasets, learning procedure and evaluation protocol) for all n CNN architectures, the differences in the performance of the classification pipeline can be attributed to the specific CNN architectures used. For comparison, 10 different deep CNN architectures were identified, grouped by the common paper which they were first made public: a) 50-layer ResNet, 101-layer ResNet and 152-layer ResNet presented in (He et al., 2015a); b) 22-layer GoogLeNet (Szegedy et al., 2015); c) 16-layer VGG-Net and 19-layer VGG-Net introduced in (Simonyan and Zisserman, 2014b); d) 8-layer VGG-S, 8-layer VGG-M and 8-layer VGG-F displayed in (Chatfield et al., 2014); and e) 8-layer Places (Zhou et al., 2014), as they represent the state-of-the-art for image classification tasks. For further design and implementation details for these models, please refer to their respective papers. To ensure a fair comparison, all the standardised CNN models used in our experiments are based on the opensource Caffe framework (Jia et al., 2014) and are pre-trained on 1000 ImageNet (Deng et al., 2009) classes with the exception of Places CNN (Zhou et al., 2014) which was trained on 205 scenes categories of Places database. For the majority of the networks, the dimensionality of the last hidden layer (FC7) leads to a 4096×1 dimensional image representation. Since the GoogLeNet (Szegedy et al., 2015) and the ResNet (He et al., 2015a) architectures do not utilise fully connected layers at the end of their networks, the last hidden layers before average pooling at the top of the ConvNet are exploited with $1024 \times 7 \times 7$

and $2048 \times 7 \times 7$ feature maps respectively, to counter-balance the behaviour of the pool layers, which provide downsampling regarding the spatial dimensions of the input.

5 EXPERIMENTS AND RESULTS

5.1 Evaluation Details

The evaluation process is divided into two different sets of scenarios, each one making use of an explicit split of images between the training and testing samples of the pipeline. For the first scenario, a split of 70/30 was utilised, while for the second scenario the split was adjusted to 50/50 for training and testing images respectively. Additionally, three distinct series of tests were conducted for each scenario, each and every one assembled with a completely arbitrary shift of the entire image set for every category of the HRUN dataset. This approach ensures an unbiased comparison with a rather limited dataset like HRUN at present. The compound results of all three tests are given in Table 2 and Table 3 and analysed below.

5.2 Results and Discussion

It is evident from Table 2 and Table 3 that the *Slow* CNN architecture performs the best for the child labour category for both scenarios. *VGG* with 16 layers performs the best in the case of child soldiers with scenario 1, while on the other hand, scenario's 2 best performing architecture is *Places* with *VGG-16* coming genuinely close. *Places* was also the best performing architecture for the category of police violence for the two scenarios. Lastly, regarding refugees category, the *Slow* version of *VGG* was the dominant architecture for both scenarios. Since our work is the

Table 2: **Human rights violations classification results with a 70/30 split for training and testing images.** Mean average precision (mAP) accuracy for different CNNs. Bold font highlights the leading mAP result for every experiment.

Model	Dimensional Representation	mAP	Child Labour	Child Soldiers	Police Violence	Refugees
ResNet 50	100K	42.59	41.12	43.69	43.81	41.73
ResNet 101	100K	42.07	40.48	44.78	42.56	40.48
ResNet 152	100K	45.80	44.27	44.11	48.08	46.73
GoogLeNet	50K	48.62	42.72	40.71	61.91	49.16
VGG 16	4K	77.46	70.79	77.71	83.46	77.87
VGG 19	4K	47.01	31.69	50.98	73.79	31.57
VGG - M	4K	67.93	59.52	62.96	81.45	67.80
VGG - S	4K	78.19	80.17	64.46	87.46	80.68
VGG - F	4K	64.15	45.42	63.20	84.78	63.21
Places	4K	68.59	55.67	65.60	93.17	59.92

Table 3: **Human rights violations classification results with a 50/50 split for training and testing images.** Mean average precision (mAP) accuracy for different CNNs. Bold font highlights the leading mAP result for every experiment.

Model	Dimensional Representation	mAP	Child Labour	Child Soldiers	Police Violence	Refugees
ResNet 50	100K	70.94	73.15	68.07	70.44	72.09
ResNet 101	100K	68.46	69.50	66.90	68.34	69.09
ResNet 152	100K	76.20	80.60	73.07	72.00	79.12
GoogLeNet	50K	55.92	41.48	60.21	55.52	66.48
VGG 16	4K	84.79	79.15	87.94	89.47	82.59
VGG 19	4K	60.39	35.72	72.67	83.10	50.08
VGG - M	4K	78.94	68.71	82.32	89.99	74.74
VGG - S	4K	88.10	84.84	88.14	91.92	87.50
VGG - F	4K	73.46	53.57	78.78	90.41	71.08
Places	4K	81.40	62.04	89.97	95.70	77.90

first effort in the literature to recognise human rights violations, we are not able to compare our experimental results with other works.

However the results are unquestionably promising and reveal that the best performing CNN architectures can achieve up to 88.10% mean average precision when recognising human rights violations. On the other hand, some of the regularly top performing deep ConvNets, such as *GoogLeNet* and *ResNet*, fell short for this particular task compared to the others. Such weaker performance occurs primarily because of the limited dataset size, whereby learning millions of parameters of those very deep convolutional networks is usually impractical and may lead to over-fitting. Another interpretation could be due to the inadequate structure of the image representation deducted from the last hidden layer before average pooling compared to the FC7 layer of the others. Furthermore, it is clear that by utilising the 50/50 split of images in the course of scenario 2, there is a considerable boost in performance of the human rights violations recognition pipeline as compared to the first scenario when a split of 70/30 was employed for training and testing images respectively. Figure 5 depicts the effect of two varying training data sizes (scenario 1 vs sce-

nario 2) on the performance of different deep convolutional networks. Remarkably with scenario 2, where the half and half split was applied, accomplishes a notable improvement on mean average precision which spans from 4.03% up to 36.33% across all four HRUN categories which were tested. Only on two occasions scenario 2 was outperformed by scenario 1, both of them while *GoogLeNet* was selected for the categories of ‘child labour’ and ‘police violence’. This observation strengthens the point of view discussed above relative to the last hidden layer of this model. Nonetheless, in all instances a mean average precision greatly above 40% was achieved, which can be regarded as an impressive outcome given the unconventional nature of the problem, the limited dataset which was adopted for learning deep representations and the transfer learning approach that was employed here.

6 CONCLUSIONS

Recognising human rights violations through digital images is a new and challenging problem in the area of computer vision. We introduce a new, open hu-



Figure 5: Comparison of deep convolutional networks performance, with reference to mAP, for the two diverse scenarios appearing in our experiments. The number on the left side of the slash denotes the training proportion of images, while the name on the right implies the testing percentage.

man rights understanding dataset, HRUN, designed to represent human rights and international humanitarian law violations found in the real world. Using this innovative dataset we conduct an evaluation of recent deep learning architectures for human rights violations recognition and achieve results that are comparable to prior attempts on other long-standing hallmark tasks of computer vision in the hope that it would provide a scaffold for future evaluations, and good benchmark for human rights advocacy research. The following conclusions have derived: Digital images that can be rated as appropriate for human rights monitoring purposes are rare and characterising them requires great effort, expertise and vast time. Utilising transfer learning for the task of recognising human rights violations can provide very strong results by employing a straightforward combination of deep representations and a linear SVM. Deep convolutional neural networks are constructed to benefit and learn from massive amounts of data. For this reason and in order to obtain even higher quality recognition results, training a deep convolutional network from scratch on an expanded version of the HRUN dataset is likely to further improve results. Inspired by the high-standard characteristics of legal evidence, in the future we would like to have the means to clarify three different questions set by every human rights monitoring mechanism: *what*, *who* and *how*, and expand our dataset to a wider range of categories in order to include them. We also presume that further analysis of joint object recognition and scene understanding will be beneficial and lead to improvements in both tasks for human rights violations understanding.

ACKNOWLEDGEMENTS

We acknowledge MoD/Dstl and EPSRC for providing the grant to support the UK academics (Ales Leonardis) involvement in a Department of Defense funded MURI project. This work was also supported in part by EU H2020 RoMaNS 645582, EP-

SRC EP/M026477/1 and ES/M010236/1.

REFERENCES

- Bell, S., Upchurch, P., Snavely, N., and Bala, K. (2015). Material recognition in the wild with the materials in context database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3479–3487.
- Chatfield, K., Simonyan, K., Vedaldi, A., and Zisserman, A. (2014). Return of the devil in the details: Delving deep into convolutional nets. In *British Machine Vision Conference, BMVC 2014, Nottingham, UK, September 1-5, 2014*.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 248–255. IEEE.
- Donahue, J., Jia, Y., Vinyals, O., Hoffman, J., Zhang, N., Tzeng, E., and Darrell, T. (2014). Decaf: A deep convolutional activation feature for generic visual recognition. In *ICML*, pages 647–655.
- Everingham, M., Van Gool, L., Williams, C. K., Winn, J., and Zisserman, A. (2010). The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338.
- Fei-Fei, L., Fergus, R., and Perona, P. (2007). Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. *Computer Vision and Image Understanding*, 106(1):59–70.
- Girshick, R., Donahue, J., Darrell, T., and Malik, J. (2014). Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587.
- Griffin, G., Holub, A., and Perona, P. (2007). Caltech-256 object category dataset.
- He, K., Zhang, X., Ren, S., and Sun, J. (2015a). Deep residual learning for image recognition. *CoRR*, abs/1512.03385.
- He, K., Zhang, X., Ren, S., and Sun, J. (2015b). Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of*

- the *IEEE International Conference on Computer Vision*, pages 1026–1034.
- Huang, Y., Huang, K., Yu, Y., and Tan, T. (2011). Salient coding for image classification. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1753–1760. IEEE.
- Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., and Darrell, T. (2014). Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the 22nd ACM international conference on Multimedia*, pages 675–678. ACM.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105.
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature*, 521(7553):436–444.
- LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., and Jackel, L. D. (1989). Back-propagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. (2014). Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, pages 740–755. Springer.
- Liu, C., Sharan, L., Adelson, E. H., and Rosenholtz, R. (2010). Exploring features in a bayesian framework for material recognition. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 239–246. IEEE.
- Nair, V. and Hinton, G. E. (2010). Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pages 807–814.
- Oquab, M., Bottou, L., Laptev, I., and Sivic, J. (2014). Learning and transferring mid-level image representations using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1717–1724.
- Sermanet, P., Eigen, D., Zhang, X., Mathieu, M., Fergus, R., and LeCun, Y. (2013). Overfeat: Integrated recognition, localization and detection using convolutional networks. *CoRR*, abs/1312.6229.
- Sharan, L., Rosenholtz, R., and Adelson, E. (2009). Material perception: What can you see in a brief glance? *Journal of Vision*, 9(8):784–784.
- Sharif Razavian, A., Azizpour, H., Sullivan, J., and Carlsson, S. (2014). Cnn features off-the-shelf: an astounding baseline for recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 806–813.
- Simonyan, K. and Zisserman, A. (2014a). Two-stream convolutional networks for action recognition in videos. In *Advances in Neural Information Processing Systems*, pages 568–576.
- Simonyan, K. and Zisserman, A. (2014b). Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556.
- Srivastava, N., Hinton, G. E., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. (2015). Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9.
- Taigman, Y., Yang, M., Ranzato, M., and Wolf, L. (2014). Deepface: Closing the gap to human-level performance in face verification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1701–1708.
- Tompson, J., Goroshin, R., Jain, A., LeCun, Y., and Bregler, C. (2015). Efficient object localization using convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 648–656.
- Torralba, A., Fergus, R., and Freeman, W. T. (2008). 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 30(11):1958–1970.
- Wang, J., Yang, J., Yu, K., Lv, F., Huang, T., and Gong, Y. (2010). Locality-constrained linear coding for image classification. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 3360–3367. IEEE.
- Xiao, J., Hays, J., Ehinger, K. A., Oliva, A., and Torralba, A. (2010). Sun database: Large-scale scene recognition from abbey to zoo. In *Computer vision and pattern recognition (CVPR), 2010 IEEE conference on*, pages 3485–3492. IEEE.
- Yu, F., Zhang, Y., Song, S., Seff, A., and Xiao, J. (2015). LSUN: construction of a large-scale image dataset using deep learning with humans in the loop. *CoRR*, abs/1506.03365.
- Zeiler, M. D. and Fergus, R. (2014). Visualizing and understanding convolutional networks. In *European Conference on Computer Vision*, pages 818–833. Springer.
- Zhou, B., Lapedriza, A., Xiao, J., Torralba, A., and Oliva, A. (2014). Learning deep features for scene recognition using places database. In *Advances in neural information processing systems*, pages 487–495.