



A Capsule Network for Hierarchical Multi-label Image Classification

Khondaker Tasrif Noor¹, Antonio Robles-Kelly¹, and Brano Kusy²

¹ School of IT, Deakin University, Waurn Ponds, Geelong, VIC, Australia
knoor@deakin.edu.au

² CSIRO Data 61, Queensland Centre for Advanced Technologies,
Pullenvale, QLD, Australia

Abstract. Image classification is one of the most important areas in computer vision. Hierarchical multi-label classification applies when a multi-class image classification problem is arranged into smaller ones based upon a hierarchy or taxonomy. Thus, hierarchical classification modes generally provide multiple class predictions on each instance, whereby these are expected to reflect the structure of image classes as related to one another. In this paper, we propose a multi-label capsule network (ML-CapsNet) for hierarchical classification. Our ML-CapsNet predicts multiple image classes based on a hierarchical class-label tree structure. To this end, we present a loss function that takes into account the multi-label predictions of the network. As a result, the training approach for our ML-CapsNet uses a coarse to fine paradigm while maintaining consistency with the structure in the classification levels in the label-hierarchy. We also perform experiments using widely available datasets and compare the model with alternatives elsewhere in the literature. In our experiments, our ML-CapsNet yields a margin of improvement with respect to these alternative methods.

Keywords: Hierarchical image classification · Capsule networks · Deep learning

1 Introduction

Image classification is a classical problem in computer vision and machine learning where the aim is to recognise the image features so as to identify a target image class. Image classification has found application in areas such as face recognition [25], medical diagnosis [32], intelligent vehicles [26] and online advertising [12] amongst others. Note that these classification tasks are often aimed at a “flat” class-label set where all the classes are treated equally, devoid of a taxonomical or hierarchical structure between them. This contrasts with hierarchical classification tasks, which require multiple label predictions per instance. Moreover, these should be consistent with the hierarchical structure of the class-label set under consideration.

In this paper, we propose a multi-label image classification model (ML-CapsNet) for hierarchical image classification based on capsule networks [20]. We note that capsule networks (CapsNets) can learn both, the image features and their transformations. This allows for a natural means to recognition by parts. Recall that capsule networks (CapsNets) use a set of neurons to obtain an “activity vector”. These neurons are grouped in capsules, whereby deeper layers account for several capsules at the preceding layer by agreement. Capsules were originally proposed in [10] to learn a feature of instantiation parameters which are robust to variations in position, orientation, scale and lighting. The assertion that CapsNets can overcome viewpoint invariance problems was further explored in [20], where the authors propose a dynamic routing procedure for routing by agreement. Hinton et al. propose in [11] a probabilistic routing approach based upon the EM-algorithm [4] so as to learn part-whole relationships. Building upon the EM routing approach in [11], Bahadori [1] proposes a spectral method to compute the capsule activation and pose.

It is not surprising that, due to their viewpoint invariance, their capacity to address the “Picasso effect” in classifiers and their robustness to input perturbations when compared to other CNNs of similar size [9], CapsNets have been the focus of great interest in the computer vision and machine learning communities. CapsNets have found applications in several computer vision tasks such as text classification [2], 3D data processing [33], target recognition [18] and image classification [28]. They have also been used in architectures such as Siamese networks [16], generative adversarial networks [23] and residual networks (ResNets) [13].

Despite the interest of the community in CapsNets, to our knowledge, they have not been employed for hierarchical multi-label classification (HMC). Since hierarchical multi-label classification can be viewed as a generalisation of multi-class problems with subordinate, not exclusive classes, it is a challenging problem in machine learning and pattern recognition that has attracted considerable attention in the research community. It has been tackled in a number of ways, spanning from kernel methods [19] to decision trees [24] and, more recently, artificial neural networks [27]. Nonetheless this attention and the fact that image HMC has been applied to the annotation of medical images [8], these methods often do not focus on images, but rather on problems such as protein structure prediction [24], data-dependent grouping [22] or text classification [15].

Further, image hierarchical multi-label classification methods are relatively few elsewhere in the literature. This is even more surprising since it is expected that incorporating hierarchies in the model would improve generalisation, particularly when the training data is limited. Image HMC methods often employ the hierarchical semantic relationships of the target classes so as to improve visual classification results, making use of the hierarchical information as a guide to the classifier. Along these lines, word hierarchies have been applied to provide consistency across multiple datasets [17], deliver a posterior confidence estimate [3] and to optimise the accuracy-specificity trade-off in visual recognition [5]. In [7], order-preserving embeddings are used to model label-to-label and image-to-image hierarchical interactions. In [34], the authors propose a Branch-CNN (B-CNN), a CNN architecture for hierarchical data organised as a tree. The B-CNN




Fig. 1. Architecture of our multi-label capsule network. Our Network has multiple secondary capsules, each of these corresponding to a class-hierarchy. The output of these is used for class prediction and the reconstruction of the input image using the decoder outputs as shown in the figure.

architecture is such that the model outputs multiple predictions ordered from coarse to fine along concatenated convolutional layers. In [30], the authors propose a hierarchical deep CNN (HD-CNN) that employs component-wise pre-training with global fine-tuning making use of a multinomial logistic loss.

Here, we present a CapsNet for image HMC. Our choice of capsules is motivated by their capacity to learn relational information, seeking to profit from the capacity of CapsNets to model the semantic relationships of image features for hierarchical classification. Since CapsNets were not originally proposed for HMC tasks, we propose a modified loss function that considers the consistency of the predicted label with that endowed by the hierarchy under consideration. Further, our reconstruction loss is common to all the secondary capsule decoders, making use of the combined predictions for the hierarchies and further imposing consistency on the HMC classification task.

2 Hierarchical Multi-label Capsules

Our ML-CapsNet uses a capsule architecture to predict multiple classes based on a hierarchical/taxonomical label tree. The overall architecture of ML-CapsNet is

presented in Fig. 1. Note our ML-CapsNet has a feature extraction block aimed at converting pixel intensities into local features that can be used as inputs for the primary capsule network. In our architecture, the feature map is reshaped accordingly to allow for the input features to be passed on to the primary capsule network $\mathcal{P}$, which is shared amongst all the secondary ones. These secondary capsule networks, denoted $\mathcal{S}_n$ in the figure, correspond to the n^{th} hierarchy in the class-label set. Thus, each of the hierarchical levels, from coarse-to-fine, have their own secondary capsule networks and, therefore, there are as many of these as hierarchies. For simplicity, each of these secondary capsules sets contain K_n capsules, where K_n corresponds to the number of classes in the hierarchy indexed n . Note that, as shown in the figure, each of the secondary capsule networks $\mathcal{S}_n$ has their own prediction layer $\mathcal{Y}_n$, which computes the class probabilities using a logit function. Thus, the $\mathcal{Y}_n$ layers provide predictions for the K_n classes corresponding to the n^{th} hierarchy class level.

It is also worth noting that, since CapsNets train making use of the reconstruction loss, here we use an auxiliary decoder network to reconstruct the input image. To this end, we make use of the activity vector delivered by the secondary capsule networks $\mathcal{S}_n$ for each hierarchy. We do this by combining the decoder outputs for each class-hierarchy into a final one via concatenation as shown in Fig. 1. As a result, the overall loss function L_T for our ML-CapsNet is a linear combination of all prediction losses and the reconstruction loss defined as follows

$$L_T = \tau L_R + (1 - \tau) \sum_{n=1}^N \lambda_n L_n \quad (1)$$

where n is the index for hierarchical level in the level tree, N is the total number of levels, λ_n a weight that moderates the contribution of each class hierarchy to the overall loss and τ is a scalar that controls the balance between the classification loss L_n and the reconstruction one L_R . In Eq. 1 the reconstruction loss is given by the L-2 norm between the input instance x and the reconstructed one $\hat{x}$. Thus, the loss becomes

$$L_R = \|x - \hat{x}\|_2^2 \quad (2)$$

For L_n we employ the hinge loss given by

$$L_n = T_n \max(0, m^+ - \|\mathbf{v}_k\|)^2 + \gamma(1 - T_n) \max(0, \|\mathbf{v}_k\| - m^-)^2 \quad (3)$$

where $T_n = 1$ if and only if the class indexed k is present, m^+ , m^- and γ are hyper-parameters and $\mathbf{v}_k$ is the output vector for the capsule for the k^{th} class under consideration.

3 Experiments

We have performed experiments making use of the MNIST [6], Fashion-MNIST [29], CIFAR-10 and CIFAR-100 [14] datasets. We have also compared our results with those yielded by the B-CNN approach proposed in [34] and used the CapsNet in [20] as a baseline. Note that, nonetheless the CapsNet in [20] is not a

hierarchical classification model, it does share with our approach the routing procedure and the structure of the primary and secondary capsules.

3.1 Implementation Details and Datasets

We have implemented our ML-CapsNet on TensorFlow and, for all our experiments, have used the Adam optimiser with TensorFlow’s default settings. For all the experiments we use exponential decay to adjust the learning rate value after every epoch, where the initial learning rate is set to 0.001 with a decay rate of 0.995.

As mentioned earlier, our network has a common feature extraction block and primary capsule $\mathcal{P}$ shared amongst the secondary capsules $\mathcal{S}_n$. In our implementation, the feature extraction block is comprised by two convolutional layers for the MNIST dataset and 5 convolutional layers for Fashion-MNIST, CIFAR-10 and CIFAR-100 datasets. We do this based on the complexity of the dataset so that the feature extraction block can adapt accordingly. For all the convolutional layers we have used 3×3 filters with zero-padding, *ReLU* activations and gradually increased the number of filters from 32 filters for the first to 512 in the following layers, i.e. 64, 128, 256 and 512. For all our results, we have used an 8-dimensional primary capsule and all the secondary capsules $\mathcal{S}_n$ are 16-dimensional with dynamic routing [20]. In all our experiments we have set m^+ , m^- and γ to 0.9, 0.1 and 0.5, respectively. The value of τ in Eq. 1 has been set to 0.0005.

We have assigned coarse and medium classes to the datasets to construct a hierarchical label tree. In this setting, coarse labels are a superclass of several corresponding medium-level labels for the dataset and these labels in turn are a superclass for the fine labels prescribed for the dataset under consideration. As a result, each instance will have multiple labels in the hierarchy, i.e. one per level. Recall that the MNIST and Fashion-MNIST dataset contains 28×28 grey-scale images. Both datasets have 60,000 training and 10,000 testing images. For the MNIST, which contains images of handwritten digits, we have followed [34] and added five coarse classes for the dataset. The Fashion-MNIST dataset is similar to the MNIST dataset and contains images of fashion products for ten fine classes. We use the coarse to medium hierarchy presented in [21], which

Table 1. Accuracy yielded by our ML-CapsNet, the B-CNN [34] and the baseline [20] on the MNIST and Fashion-MNIST datasets. The absolute best are in bold.

Model	Accuracy (%)				
	MNIST		Fashion-MNIST		
	Coarse	Fine	Coarse	Medium	Fine
CapsNet	—	99.3	—	—	91.2
B-CNN	99.2	99.28	99.57	95.52	92.04
ML-CapsNet	99.65	99.5	99.73	95.93	92.65

Table 2. Accuracy yielded by our ML-CapsNet, the B-CNN [34] and the baseline [20] on the CIFAR-10 and CIFAR-100 datasets. The absolute best are in bold.

Model	Accuracy (%)					
	CIFAR-10			CIFAR-100		
	Coarse	Medium	Fine	Coarse	Medium	Fine
CapsNet	—	—	70.42	—	—	34.93
B-CNN	95.63	86.95	84.95	71.04	61.94	55.52
ML-CapsNet	97.52	89.27	85.72	78.73	70.15	60.18

manually adds a coarse and medium level for the dataset. In [34], the coarse level has two classes, and the medium level has six classes. The CIFAR-10 and CIFAR-100 dataset consist 32×32 colour images and have 50,000 training images and 10,000 testing images in 10 and 100 fine classes, respectively. Here, we use the class-hierarchy used in [34], which adds an additional coarse and medium levels for the dataset whereby the coarse level contains two classes and the medium contains seven classes. The CIFAR-100 dataset is similar to CIFAR-10, except it has 100 fine and 20 medium classes. We have used the hierarchy in [34], which groups the medium and fine classes into eight coarse ones.

3.2 Experimental Setup

As mentioned earlier, in order to compare our ML-CapsNet with alternatives elsewhere in the literature, we have used the CapsNet as originally proposed in [20] as a baseline and compared our results with those yielded by the B-CNN [34]. In all our experiments, we normalise the training and testing data by subtracting the mean and dividing by the standard deviation. For all the datasets, the B-CNN follows the exact model architecture and training parameters used by the authors in [34]. For training the B-CNN model on the Fashion-MNIST dataset, we followed the architecture for the MNIST dataset and added an additional branch for the fine level.

When training our ML-CapsNet on the CIFAR-10 and CIFAR-100 datasets, we apply MixUp data augmentation [31] with $\alpha = 0.2$. Also, recall that the λ_n values in Eq. 1 govern the contribution of each hierarchy level to the overall loss function. Thus, for our ML-CapsNet, we adjust the λ_n values as the training progresses so as to shift the importance of the class-level hierarchy from coarse-to-fine. As a result, on the MNIST dataset we set the initial λ_n to 0.90, 0.10 for the coarse and fine levels, respectively. The value of λ_n is then set to 0.10, 0.90 after 5 epochs and 0.02, 0.98 after 10. For the Fashion-MNIST dataset the initial loss weight values for λ_n are 0.98, 0.01, 0.01 for the coarse, medium and fine hierarchies. The value of λ_n after 5 epochs becomes 0.10, 0.70, 0.20, after 10 epochs shifts to 0.07, 0.10, 0.83, at 15 epochs is set to 0.05, 0.05, 0.90 and, finally, at 25 epochs assumes the value of 0.01, 0.01, 0.98. For the CIFAR-10 dataset, the initial λ_n values are 0.90, 0.05 and 0.05 for the coarse, medium




Fig. 2. Accuracy as a function of training epoch for all the models under consideration. The left-hand panel shows the plots for the MNIST dataset whereas the right-hand panel corresponds to the Fashion-MNIST dataset.

and fine class-hierarchy levels. The λ_n then changes to 0.10, 0.70 and 0.20 after Epoch 5 and, at Epoch 11 is set to 0.07, 0.20 and 0.73. At Epoch 17 is set to 0.05, 0.15 and 0.80, taking its final value of 0.05, 0.10 and 0.85 at Epoch 24. For the CIFAR-100 dataset the initial values are 0.90, 0.08 and 0.02, which then change to 0.20, 0.70 and 0.10 after 7 epochs. After Epoch 15 we set these to 0.15, 0.30 and 0.55. At Epoch 22 we modify them to be 0.10, 0.15 and 0.75 and, at Epoch 33, they take their final value of 0.05, 0.15 and 0.80.

3.3 Results

We now turn our attention to the results yielded by our network, the CapsNet as originally proposed in [20] and the B-CNN in [34] when applied to the four datasets under consideration. In Table 1 we show the performance yielded by these when applied to the MNIST and Fashion-MNIST dataset. Similarly, in Table 2, we show the accuracy for our network and the alternatives when applied to the CIFAR-10 and CIFAR-100 datasets. In the tables, we show the accuracy for all the levels of the applicable class-hierarchies for both, our network and the B-CNN. We also show the performance on the CapsNet in [20] for the fine class-hierarchy. We do this since the baseline is not a hierarchical classification one, rather the capsule network as originally proposed in [20] and, therefore, the medium and coarse label hierarchies do not apply.

Note that, for all our experiments, our network outperforms the alternatives. It is also worth noting that, as compared with the baseline, it fairs much better for the more complex datasets of CIFAR-10 and CIFAR-100. This is also consistent with the plots in Figs. 2 and 3, which show the model accuracy on the dataset under consideration for every one of their corresponding class-hierarchies. Note that in several cases, our network not only outperforms B-CNN, but also




Fig. 3. Accuracy as a function of training epoch for all the models under consideration. The left-hand panel shows the plots for the CIFAR-10 dataset whereas the right-hand panel corresponds to the CIFAR-100 dataset.

converges faster. It is also important to note that, as compared to the baseline, our ML-CapsNet employs the hierarchical structure of the class-set to impose consistency through the loss. This appears to introduce a structural constraint on the prediction that helps improve the results.

This is consistent with the reconstruction results shown in Fig. 4. In the figure, we show the reconstruction results for the classes in the Fashion-MNIST dataset for both, our ML-CapsNet and the capsule network as originally proposed in [20]. In the figure, from top-to-bottom, the rows show the input image, the image reconstructed by our network and that yielded by that in [20]. Note the images reconstructed using our network are much sharper, showing better detail and being much less blurred. This is somewhat expected, since the reconstruction




Fig. 4. Sample reconstructed images for the Fashion-MNIST dataset. From top to bottom we show the input image, the image reconstructed by our ML-Capsnet and that reconstructed using the capsule network in [20].

error is also used by the loss and, hence, better reconstruction should yield higher accuracy and vice versa.

4 Conclusions

In this paper, we have presented a capsule network for image classification, which uses capsules to predict multiple hierarchical classes. The network presented here, which we name ML-CapsNet, employs a shared primary capsule, making use of a secondary one for each class-label set. To enforce the multi-label structure into the classification task, we employ a loss which balances the contribution of each of the class-sets. The loss proposed here not only enforces consistency with the label structure, but incorporates the reconstruction loss making use of a common encoder. We have shown results on four separate widely available datasets. In our experiments, our ML-CapsNet outperforms the B-CNN [34] and the classical capsule network in [20]. As expected, it also delivers better reconstructed images than those yielded by the network in [20].

References

1. Bahadori, M.T.: Spectral capsule networks. In: International Conference on Learning Representations Workshops (2018)
2. Chen, B., Huang, X., Xiao, L., Jing, L.: Hyperbolic capsule networks for multi-label classification. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 3115–3124 (2020)
3. Davis, J., Liang, T., Enouen, J., Ilin, R.: Hierarchical classification with confidence using generalized logits. In: 2020 25th International Conference on Pattern Recognition (ICPR), pp. 1874–1881 (2021)
4. Dempster, A., Laird, N., Rubin, D.: Maximum-likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. Ser. B Methodol.* **39**, 1–38 (1977)
5. Deng, J., Krause, J., Berg, A.C., Fei-Fei, L.: Hedging your bets: optimizing accuracy-specificity trade-offs in large scale visual recognition. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition, pp. 3450–3457 (2012)
6. Deng, L.: The MNIST database of handwritten digit images for machine learning research [best of the web]. *IEEE Signal Process. Mag.* **29**(6), 141–142 (2012)
7. Dhall, A., Makarova, A., Ganea, O., Pavlo, D., Greeff, M., Krause, A.: Hierarchical image classification using entailment cone embeddings. In: Computer Vision and Pattern Recognition Workshops, pp. 836–837 (2020)
8. Dimitrovski, I., Koccev, D., Loskovska, S., Dzeroski, S.: Hierarchical annotation of medical images. *Pattern Recogn.* **44**(10), 2436–2449 (2011)
9. Hahn, T., Pyeon, M., Kim, G.: Self-routing capsule networks. In: Advances in Neural Information Processing Systems, vol. 32 (2019)
10. Hinton, G.E., Krizhevsky, A., Wang, S.D.: Transforming auto-encoders. In: International Conference on Artificial Neural Networks, pp. 44–51 (2011)
11. Hinton, G.E., Sabour, S., Frosst, N.: Matrix capsules with EM routing. In: International Conference on Learning Representations (2018)
12. Hussain, Z., et al.: Automatic understanding of image and video advertisements. In: Computer Vision and Pattern Recognition, pp. 1705–1715 (2017)

13. Jampour, M., Abbaasi, S., Javidi, M.: Capsnet regularization and its conjugation with resnet for signature identification. *Pattern Recogn.* **120**, 107851 (2021)
14. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
15. Meng, Y., Shen, J., Zhang, C., Han, J.: Weakly-supervised hierarchical text classification. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, pp. 6826–6833 (2019)
16. Neill, J.O.: Siamese capsule networks. *arXiv E-preprints* (2018)
17. Redmon, J., Farhadi, A.: Yolo9000: better, faster, stronger. In: *Computer Vision and Pattern Recognition*, pp. 7263–7271 (2017)
18. Ren, H., Yu, X., Zou, L., Zhou, Y., Wang, X., Bruzzone, L.: Extended convolutional capsule network with application on SAR automatic target recognition. *Signal Process.* **183**, 108021 (2021)
19. Rousu, J., Saunders, C., Szedmak, S., Shawe-Taylor, J.: Kernel-based learning of hierarchical multilabel classification models. *J. Mach. Learn. Res.* **7**, 1601–1626 (2006)
20. Sabour, S., Frosst, N., Hinton, G.E.: Dynamic routing between capsules. In: *Advances in Neural Information Processing Systems*, vol. 30 (2017)
21. Seo, Y., Shin, K.S.: Hierarchical convolutional neural networks for fashion image classification. *Expert Syst. Appl.* **116**, 328–339 (2019)
22. Ubaru, S., Dash, S., Mazumdar, A., Günlük, O.: Multilabel classification by hierarchical partitioning and data-dependent grouping. In: *Advances in Neural Information Processing Systems* (2020)
23. Upadhyay, Y., Schrater, P.: Generative adversarial network architectures for image synthesis using capsule networks. *arXiv E-preprint* (2018)
24. Vens, C., Struyf, J., Schietgat, L., Džeroski, S., Blockeel, H.: Decision trees for hierarchical multi-label classification. *Mach. Learn.* **73**(2), 185–214 (2008)
25. Wang, M., Deng, W.: Deep face recognition: a survey. *Neurocomputing* **429**, 215–244 (2021)
26. Wang, Z., Zhan, J., Duan, C., Guan, X., Lu, P., Yang, K.: A review of vehicle detection techniques for intelligent vehicles. *IEEE Trans. Neural Netw. Learn. Syst.* (2022)
27. Wehrmann, J., Cerri, R., Barros, R.: Hierarchical multi-label classification networks. In: *International Conference on Machine Learning*, pp. 5075–5084 (2018)
28. Xiang, C., Zhang, L., Tang, Y., Zou, W., Xu, C.: MS-CapsNet: a novel multi-scale capsule network. *IEEE Signal Process. Lett.* **25**(12), 1850–1854 (2018)
29. Xiao, H., Rasul, K., Vollgraf, R.: Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint [arXiv:1708.07747](https://arxiv.org/abs/1708.07747)* (2017)
30. Yan, Z., et al.: HD-CNN: hierarchical deep convolutional neural networks for large scale visual recognition. In: *International Conference on Computer Vision*, pp. 2740–2748 (2015)
31. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: beyond empirical risk minimization. In: *International Conference on Learning Representations* (2018)
32. Zhang, Z., Xie, Y., Xing, F., McGough, M., Yang, L.: MDNet: a semantically and visually interpretable medical image diagnosis network. In: *Computer Vision and Pattern Recognition*, pp. 6428–6436 (2017)
33. Zhao, Y., Birdal, T., Deng, H., Tombari, F.: 3D point capsule networks. In: *Computer Vision and Pattern Recognition* (2019)
34. Zhu, X., Bain, M.: B-CNN: branch convolutional neural network for hierarchical classification. *arXiv E-prints pp. arXiv-1709* (2017)