

People, Penguins and Petri Dishes: Adapting Object Counting Models To New Visual Domains And Object Types Without Forgetting

Mark Marsden¹, Kevin McGuinness¹, Suzanne Little¹, Ciara E. Keogh², Noel E. O'Connor¹

¹Insight Centre for Data Analytics
 Dublin City University, Ireland

mark.marsden@insight-centre.org {kevin.mcguinness, suzanne.little, noel.oconnor}@dcu.ie

²Conway Institute, School of Medicine
 University College Dublin, Ireland

ciara.keogh@ucdconnect.ie

Abstract

In this paper we propose a technique to adapt a convolutional neural network (CNN) based object counter to additional visual domains and object types while still preserving the original counting function. Domain-specific normalisation and scaling operators are trained to allow the model to adjust to the statistical distributions of the various visual domains. The developed adaptation technique is used to produce a singular patch-based counting regressor capable of counting various object types including people, vehicles, cell nuclei and wildlife. As part of this study a challenging new cell counting dataset in the context of tissue culture and patient diagnosis is constructed. This new collection, referred to as the Dublin Cell Counting (DCC) dataset, is the first of its kind to be made available to the wider computer vision community. State-of-the-art object counting performance is achieved in both the Shanghaitech (parts A and B) and Penguins datasets while competitive performance is observed on the TRANCOS and Modified Bone Marrow (MBM) datasets, all using a shared counting model.

1. Introduction

Vision-based object counting is an important analysis step in many observational scenarios including wildlife studies, microscopic imaging and CCTV surveillance. An accurate object counting system can provide valuable insights such as the congestion level of a public square (crowd counting), the level of traffic on a motorway (vehicle counting), the inferred migration patterns of a penguin colony (wildlife counting) or the proliferation of cancerous cells in a patient (cell counting). There is a common set of challenges for vision-based object counting which limit counting accuracy in all of these domains. These challenges include object scale



Figure 1: Examples of the challenging images encountered in different object counting domains.

and perspective issues, visual occlusion and poor illumination. To date, these highly related counting tasks have been tackled separately with algorithms engineered specifically to perform object counting in a given domain. Examples of these challenges are shown across several visual domains in Figure 1.

Vision based object counting has been tackled using both object detection [1, 2, 3] and count regression techniques [4, 5, 6, 7, 8] with both types of approach having their own specific weakness (visual occlusion and model overfitting respectively). The utilisation of convolutional neural net-

works (CNN) and hardware accelerated optimisation has led to significant improvements in the accuracy of regression-based counting approaches across several visual domains [9, 10, 11]. This paper looks to build upon this work and investigate if a patch-based object count regressor can be adapted to additional object types and visual domains while maintaining accuracy in the original counting task. The potential benefits of a modular counting architecture include transfer learning and the removal of redundant model parameters. The significant variation in statistical distribution between visual domains is one of the main obstacles to domain adaptation in computer vision [12]. This challenge is addressed in our approach by including a set of domain-specific scaling and normalization layers distributed throughout the network which make up only a small fraction of the overall parameter count [12]. A sequential training procedure inspired by the work of Rebuffi *et al.* [12] allows the network to be extended to new counting tasks over time while preserving all previously learned counting functions. The set of count estimates produced using the proposed patch-based counting model are then refined using an efficient fully convolutional neural network which utilises the wider scene context to mitigate possible count errors. The proposed framework is also extended to perform a visual domain classification in the event of the observed domain being unknown, highlighting the versatility and modular nature of the approach. The core contributions of this paper can be summarised as follows:

- An extendable CNN architecture for patch-based, multi-domain object counting is developed for the first time;
- A fully convolutional neural network is utilised to refine a set of count estimates for an entire image;
- A challenging, representative dataset for cell counting in a tissue culture/patient diagnosis setting is proposed;
- Domain adaptation in object counting is shown to produce more efficient, higher accuracy counting models;
- State-of-the-art counting accuracy is observed on leading benchmarks for several object types including the Shanghaitech [13] and Penguins [14] datasets.

The remainder of this paper is organised as follows: Section 2 presents the related work found in the literature. Section 3 describes the proposed multi-domain object counting approach while Section 4 details the construction of a challenging new cell counting dataset. Finally, Section 5 presents a comprehensive set of experiments highlighting the development of our technique and the benefits associated with multi-domain object counting.

2. Related work

Crowd Counting. Crowd counting has been approached using a wide variety of techniques, from HOG-based head detectors [15] to CNN-based regressors [9]. Heatmap-based crowd counting using a fully convolutional neural network was firstly investigated by Zhang *et al.* [13] and subsequently by Marsden *et al.* [16], with notable performance gains observed. A novel model switching technique for crowd counting was proposed by Sam *et al.* [17] which firstly classifies the crowd density level of an image region before performing heatmap-based counting using a network which has been optimised for the detected crowd density level. The standard datasets for evaluating crowd counting techniques include the UCF_CC_50 [15] and Shanghaitech [13] datasets.

Cell counting. Cell nuclei counting techniques have evolved from SIFT-based heatmap estimation [6] to fully convolutional neural networks [18, 19], with significant improvements in accuracy observed. Significant variation in cell morphology and visual occlusion limit the accuracy of these techniques. The VGG Cell dataset [20] (made up entirely of synthetic images) is the main public benchmark used to compare cell counting techniques. Recently the Modified Bone Marrow (MBM) dataset [19] was introduced, containing images of real cells observed in histological slides (a 2D cross section of a 3D tissue structure). However, a cell counting dataset in the context of tissue culture and patient diagnosis is still needed. Cell counting in tissue culture is often done by hand on a daily basis, especially in cases where highly expensive automated systems are not available. A vision-based cell counting system for tissue culture can lead to a more affordable solution and faster diagnoses overall.

Vehicle counting. Vehicle counting has been tackled using a variety of techniques including SIFT-based regression [21], CNN regression [10], density heatmap generation [22] as well as an LSTM (Long Short-Term Memory) based approach from Zhang *et al.* [23] which adds a temporal dimension to vehicle counting. Vehicle counting techniques are trained and evaluated using the WebCamT [23] and TRANCOS [21] datasets.

Wildlife counting. Counting of wildlife in ecological studies has not received significant attention from the computer vision community. The core work in this area is that of Arteta *et al.* [14] who produced a large-scale dataset for counting penguin colonies.

Domain adaptation and shared learning models. While this paper focuses primarily on domain adaptation it also touches on concepts including transfer learning, feature extraction and learning without forgetting (LwF). The ability to train a machine learning model to perform additional tasks over time while maintaining accuracy in previous tasks has generated significant interest in the computer vision community. This notion is inspired largely by the human visual system, which learns a universal representation for vision in

the early life and uses this representation for a variety of problems [12]. Fine-tuning is a common process used to adapt a given neural network to a new task, the main downside being that the original function is often lost during optimisation. Multi-Task learning (MTL) approaches attempt to train a model to simultaneously perform several tasks, often within a *specific* visual domain. MTL approaches typically involve including sets of task specific layers at end of a given neural network and jointly training to perform all tasks. This type of approach is cumbersome to extend to new tasks and visual domains as all tasks must be retrained.

While intra-domain, multi-objective learning has been utilized heavily to extend the functionality of neural networks, training a model to perform tasks in various *distinct* visual domains (CCTV scenes, medical imaging) has proven to be more challenging due to the observed variation in statistical distributions between domains. These issues have been addressed by Bilen *et al.* [24] who trained a CNN model to learn a universal vision representation which can jointly perform non-related tasks from distinct visual domains by including domain-specific scaling and normalisation layers throughout the network. This work was extended by Rebuffi *et al.* [12] who added domain-specific convolutions and proposed a sequential training procedure for learning new tasks over time without discarding the previously learned functions.

3. Our approach

The proposed object counting technique consists of a patch-based CNN regressor with significant inter-domain parameter sharing that can be quickly switched between a learned set of visual domains by interchanging a subset of domain specific parameters. The estimated object count for a given image is calculated by summing the object count for all patches. A patch based approach to counting enables a more robust regressor to be learned as objects within a smaller image region are observed to be approximately uniform in size. A lightweight fully convolutional neural network is then used to refine a set of patch estimates for a given image by including the context of adjacent patch estimates to mitigate possible errors.

3.1. Base object counting regressor

The base object counting regressor (shown in Figure 2) consists of two distinct steps. First a set of high-level features are extracted from each image patch using a pre-trained image classification network, the parameters of which are frozen during model training. The N feature maps generated by the given image classification network’s final convolutional layer are average pooled globally to produce an N -dimensional feature representation. This N -dimensional feature representation is then mapped to an object count value using a fully connected neural network. This fully connected

network consists of 5 layers with the following configuration of neurons 256-128-64-64-1. Rectified linear unit (ReLU) activations are applied after each fully connected layer. The trainable portion of the network consists of just 330,000 parameters when the pre-trained base network produces a 1024-D descriptor. A variety of pre-trained object classification networks and image patch sizes are investigated in Section 5.

3.2. Domain-specific layers

To enable the proposed counting regressor to adapt to various visual domains (e.g. people, vehicles, wildlife, cell nuclei) a set of domain-specific modules are included before each fully connected layer and then after the final fully connected layer. These domain-specific modules are interchanged during training and inference depending on the chosen visual domain (this switching concept is highlighted in Figure 3). Each domain-specific module contains the residual adapter of Rebuffi *et al.* [12] which has been adjusted to work the fully connected layers of the proposed network. An input vector input of length N is firstly fed through a 1-D batch normalisation layer, then a 1-D convolutional layer with a single $N \times 1$ convolutional kernel and finally another 1-D batch normalisation layer, with a residual skip connection also included to improve model convergence. The structure of the proposed adapter module is shown in figure 4. Each adapter module requires just $9N + 1$ parameters where N is the input vector length. These residual adapter modules allows for the network to adapt to the distinct statistical distributions of the various visual domains through domain-specific normalisation and scaling [25].

3.3. Sequential training

A sequential training procedure inspired by the work of Rebuffi *et al.* [12] is used to allow the proposed counting regressor to be extended to new domains over time while still performing the original set of counting tasks. First the network is primed by training to convergence on a given visual domain (e.g. crowd counting), after which the domain-agnostic parameters (i.e. the fully connected layers) are frozen and only the subset of domain-specific parameters are trained to convergence for each of the remaining visual domains, one at a time. Freezing the domain-agnostic parameters enables the network to retain the original counting regression function it has learned. Euclidean distance, given in equation 1 is optimised during training. Θ corresponds to the set of network parameters to optimise, N is the batch size, X_i is the i^{th} batch image while F_i is the corresponding ground truth count value. $F(X_i; \Theta)$ is the estimated count value for a given batch image X_i .

$$L_{l2}(\Theta) = \frac{1}{2N} \sum_{i=1}^N \|F(X_i; \Theta) - F_i\|_2^2. \quad (1)$$

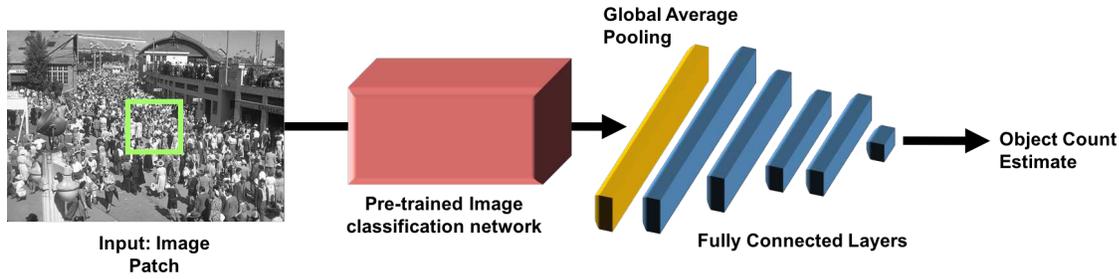


Figure 2: Network architecture for the proposed patch-based object counting regressor

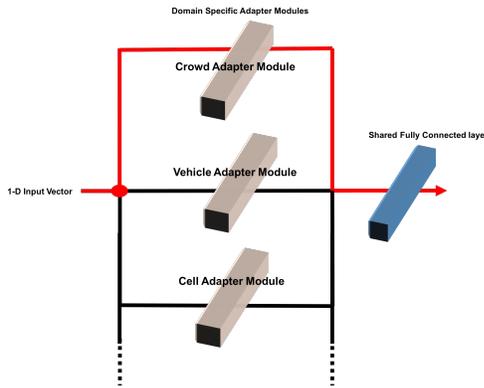


Figure 3: Domain specific modules are interchanged during training and inference depending on the chosen counting task (red path).

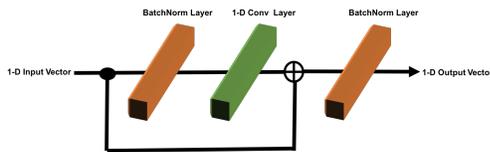


Figure 4: The residual adapter module of Rebuffi *et al.* [12].

The AdaGrad optimiser [26] is used to avoid learning rate selection issues with the initial learning rate set to 1×10^{-1} . L_2 weight regularisation (i.e. weight decay) is also included during model training with λ set to 1×10^{-3} . Model weights are initialised using the uniform initialiser of Glorot and Bengio [27] while the bias terms are initialised to zero. Training is carried out for 10,000 iterations for each domain. The choice of visual domain used to prime the network is investigated in section 5.

3.4. Fully convolutional refinement network

A patch-based counting regressor does not include the wider scene context when processing a given image as patches are analysed on an individual basis. To address this, a fully convolutional neural network is trained to refine

a grid of patch estimates for a given image. The proposed fully convolutional network, shown in Figure 5, consists of 4 convolutional layers with the following number of kernels per layer: 16,16,16,1. Rectified linear unit (ReLU) activations are applied after each layer. All convolutional kernels are 3×3 resulting in a total parameter count of 4950.

A refinement model can be trained for the count regressor of a given domain by firstly producing a grid of patch estimates for each training image used as well as a corresponding ground truth grid. Training is then carried out using these grid pairs for 10,000 iterations with Euclidean distance again minimised. Once trained, a refinement model can be applied to a set of count estimates of any width and height due to the fully convolutional nature of the model. During inference this step results in a near negligible increase in processing time.

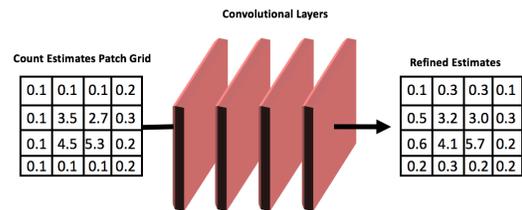


Figure 5: Fully convolutional network used for count estimation refinement.

3.5. Domain classification

If the visual domain observed during inference is unknown this can be predicted by extending the core network to also perform domain classification. To accomplish this the final fully connected layer is interchanged with a K -neuron fully connected layer, where K is the number of visual domains to distinguish between. Following this final layer a softmax activation is applied and the subsequent domain adapter module is not included. Training can then be performed by freezing the common model parameters (the initial 4 fully connected layers) and creating a fresh set of adapter modules. Categorical cross-entropy loss, defined as:

$$L_{\text{CCE}}(\Theta) = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^K S_{ij} \log(\hat{S}_{ij}), \quad (2)$$

is then minimized. Θ corresponds to the set of trainable network parameters, N and K are the batch size and number of visual domains to be distinguished between while \hat{S}_{ij} is the predicted probability score for concept j in the i^{th} batch image and S_{ij} is the corresponding ground truth value.

4. Dublin cell counting dataset

When considering the application of computer vision to tissue culture and patient diagnosis there is a clear lack of publicly available and fully annotated cell microscopy datasets. The main dataset used to evaluate techniques for this task consists entirely of synthetic images [18]. To address this, the Dublin Cell Counting (DCC) dataset was constructed.

This dataset consists of 177 images containing a wide array of tissues and species. Amongst these are examples of stem cells derived from embryonic mice, isolated human lung adenocarcinoma and examples of primary human monocytes isolated from a healthy human volunteer. Several factors were varied during image capture to provide a more representative set of images. First, the density of cells loaded onto the slide naturally varies as cell lines proliferate at different rates. Second, the morphology and size of the cells for each cell line can vary significantly. Furthermore, the objective lens used during imaging was varied as was the diameter of the diaphragm which controls the amount of light hitting the sample. Finally, the haemocytometer grid size was varied to produce a representative set of non-cellular image artifacts. Cell images were obtained via a camera mounted on an Olympus CKX41 microscope using both $4\times$ and $10\times$ objectives. The high levels of variation in this collection allow for a more robust cell counting function to be learned. After the full set of image were acquired, a dot annotation process was performed by a domain expert with a background in molecular biology.

The mean cell count across these images is 34.1 with a standard deviation of 21.8, showing the significant variation in cell density. 100 images are used for training and validation while the remaining 77 form an unseen test set. Sample images from this collection are shown in Figure 6.

5. Experiments

The proposed multi-domain object counting technique is evaluated using a challenging and representative dataset for each visual domain including the proposed DCC dataset. These collections are detailed in full in Table 1. Horizontal flips are used for training set augmentation in all cases. A 30% subset of the provided training data is set aside as

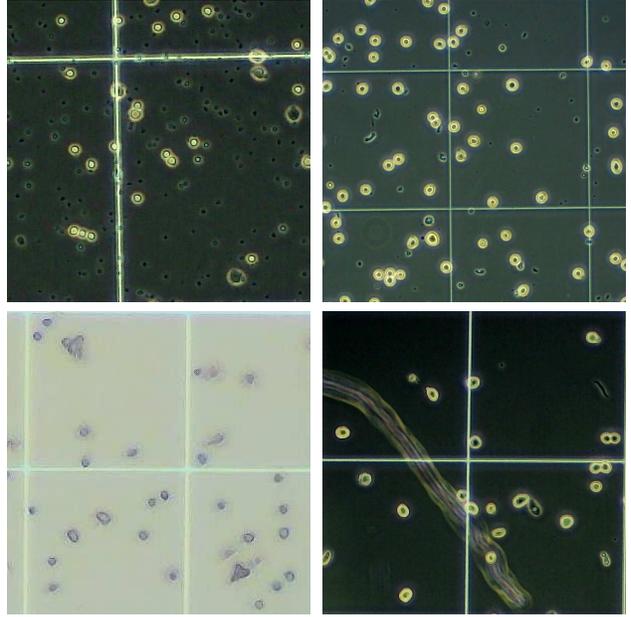


Figure 6: DCC dataset examples showing the significant variation within this collection.

a validation set for all model selection experiments, apart from where an explicit validation set is provided. No count estimate refinement is applied until subsection 5.5. Counting performance is evaluated for all visual domains using Mean Absolute Error (MAE) and Root Mean Squared Error (MSE), which are defined as follows:

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |z_i - \tilde{z}_i|, \quad (3)$$

$$\text{MSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (z_i - \tilde{z}_i)^2}, \quad (4)$$

All network optimisation and testing is performed using an NVIDIA GeForce GTX 970 GPU with a batch size of 256 and implemented using the Keras API [28] with a Tensorflow backend [29].

5.1. Patch size selection

Several image patch sizes (50×50 , 100×100 , 200×200) are compared by evaluating crowd counting validation performance on the Shanghaitech dataset (part A). Feature extraction for the base counting regressor is performed using the VGG16 network of Simonyan and Zisserman [30] for all runs. A given patch size is used during both dataset construction and inference. Both the domain-agnostic and domain-specific network parameters are optimised for each run in this experiment as the network has not been primed on any other dataset. Table 2 highlights the performance

Visual Domain	Dataset	No. of Images	Count Mean	Count STD	Count Range
Crowd	Shanghaitech (Part A) [13]	482	501.4	456.4	33-3139
Vehicles	TRANCOS [21]	1641	36.54	14.9	9-95
Wildlife	Penguins [14]	80095	7.18	5.71	0-67
Cells	Dublin Cell Counting	177	34.1	21.8	0-101

Table 1: Details of the datasets used to evaluate the proposed multi-domain object counting technique.

of the various patch sizes. 100×100 patches result in the best overall performance and will be used in all subsequent experiments. This size likely strikes the optimal balance between uniform object size and the inclusion of wider scene context needed to learn a robust function. On the other hand, the use of 200×200 patches results in significantly inferior performance due to the non-uniform object sizes observed in larger image regions.

Patch Size	MAE	MSE
50×50	100.8	152.5
100×100	97.5	145.4
200×200	158.5	247.3

Table 2: Crowd counting validation performance of various image patch sizes on a validation set taken from the Shanghaitech (Part A) dataset.

5.2. Feature extractor selection

Several pre-trained image classification networks are evaluated as feature extractors for object counting. These networks include the ResNet50 network of He *et al.* [31] the VGG16 network of Simonyan and Zisserman [30] as well as the MobileNet architecture of Howard *et al.* [32]. Crowd counting validation performance is again evaluated on the Shanghaitech (Part A) dataset with a patch size of 100×100 used. Table 3 details the various pre-trained networks and presents the performance achieved by each. The number of parameters associated with each network does not include the original set of fully connected layers (which have been replaced entirely). The MobileNet feature extractor achieves the best overall performance despite having significantly fewer parameters and is used for all subsequent experiments. The use of this lightweight architecture also potentially makes the developed method more suitable for edge processing on low power devices.

5.3. Priming the network

In the work of Rebuffi *et al.* [12] a multi-domain image classification network is primed on the ImageNet dataset [33] due to its size and variety before being adapted to other domains. However, for the proposed multi-domain object counting model the choice of which domain with which to

Net	Parameters	MAE	MSE
VGG16 [30]	14M	101.2	151.4
Resnet50 [31]	24M	97.5	145.4
MobileNet ($\alpha = 1.0$) [32]	3.5M	94.1	138.9

Table 3: Validation performance on the Shanghaitech dataset (part A) for various image classification networks being used for feature extraction.

prime the network is not clear. Therefore all 4 domains (crowds, vehicles, wildlife, cells) are evaluated for this purpose by priming the network from scratch and then adapting to the remaining 3 domains one by one using the sequential training process discussed previously. Table 4 presents the MAE score observed for each pairing. Each row in this table corresponds to the domain used to prime the network while each column corresponds to the MAE performance achieved when adapting the model. The diagonal entries correspond to the performance achieved when training the network from scratch for each domain. A concurrent training baseline is also included, where the model is trained to perform all 4 counting tasks, switching between them in round robin fashion.

It can be seen that priming the network on the cell domain using the DCC dataset results in the best overall performance, achieving optimal MAE on 3 of the 4 domains. Adapting from the cell domain achieves performance superior to training from scratch on both the crowd and wildlife domains, which is noteworthy given the small number of domain-specific parameters trained. The high performance observed across domains when adapting from a cell counting model is likely due to the significant morphological variation observed between cell objects in the DCC dataset, resulting in a broader set of learned features. In all subsequent experiments the counting network is primed on the cell counting task.

5.4. Comparison with feature extraction based network adaptation

Feature extraction based network adaptation involves freezing the majority of a given neural network and retraining the final few layers for a new task (with the task-specific parameters then interchanged as required). This approach

Visual Domain	Crowd (Tested)	Vehicles (Tested)	Wildlife (Tested)	Cells (Tested)	Best Overall
Crowd (Primed)	94.1	10.3	6.8	12.9	0/4
Vehicles (Primed)	96.2	9.9	6.1	11.5	1/4
Wildlife (Primed)	95.3	10.3	6.05	10.2	0/4
Cells (Primed)	90.2	10.2	5.7	9.5	3/4
Concurrent Training	92.6	9.95	5.93	10.1	0/4

Table 4: The MAE validation performance achieved when priming the network using each of the 4 visual domains.

retains the original function but often requires a large number of new parameters to be trained. In this experiment we compare the Rebuffi domain adapter [12] with feature extraction as a domain adaptation strategy. The approaches are evaluated in terms of MAE performance in the target domain and the quantity of new model parameters introduced. The number of retrained fully connected layers is varied to investigate the difference in performance. All adapter modules placed before trainable layers are re-trained while those before frozen layers are also frozen. A given cell counting model is adapted to perform crowd counting on the Shanghaitech dataset (part A) in each case. Figure 5 highlights the superior performance of the Rebuffi domain adapter when adapting to crowd counting, despite using significantly fewer parameters.

Approach	Extra Parameters	MAE
From Scratch	330K	94.1
Final 4 layers re-trained	50K	128.3
Final 3 layers re-trained	15K	232.2
Final 2 layers re-trained	5K	245.5
Rebuffi <i>et al.</i> [12]	16K	90.2

Table 5: MAE validation performance on the Shanghaitech dataset (part A) for various domain adaptation strategies.

5.5. Prediction refinement

A fully convolutional refinement model is trained for each of the 4 visual domains and compared to the validation performance of the base regressor model in each case. Results of this experiment are shown in table 6. This efficient post-processing step has a near negligible impact on inference time but results in significant performance boosts in the crowd and cell counting benchmarks. Examples of the final count regressor in action across all 4 domains are shown in figure 7.

5.6. Domain classification performance

A domain classifier is trained by firstly producing a representative dataset taken from all 4 visual domains used in this study. 300 full scene images are taken from each collection and a 6:4 training/test split is then applied, ensuring an even number of images are taken from each domain. Horizontal

Visual Domain	Base MAE	Refined-MAE
Crowds	90.2	86.5
Vehicles	10.2	10.1
Wildlife	5.7	5.6
Cells	9.5	8.4

Table 6: The MAE observed on the validation sets of all 4 visual domains for the base regressor and after the proposed refinement step has been applied.

image flips are applied to produce a 300 image set from the smaller DCC dataset. Table 7 compares classifier accuracy when training from scratch, when adapting from a cell counting network using a fresh set of adapter modules and when adapting from a cell counting network and just training the final fully connected layer.

Near perfect domain classification accuracy is observed when training the network from scratch and also when adapting from the a cell counting network using a new set of adapter modules. However, when we only train the final fully connected layer to perform classification we observe significant performance degradation, showing the importance of the included domain adapter modules despite their low parameter count.

Primer	Training Approach	Accuracy
None	Entire Network Trained	99.7%
Cells	Adapter Modules + Final Layer	98.2%
Cells	Final Layer Only	56.3%

Table 7: Domain classification accuracy as the training approach is varied and domain adapter modules are included.

5.7. Comparison with the state-of-the-art

The developed object counting technique is compared to the leading techniques across the 4 visual domains used in this study. Evaluation is performed in all cases on a previously unseen test set. The network is primed on cell counting using the DCC dataset. The proposed refinement step is applied for each domain. Table 8 compare crowd counting performance on the Shanghaitech dataset (parts A and B), with state-of-the-art performance achieved on

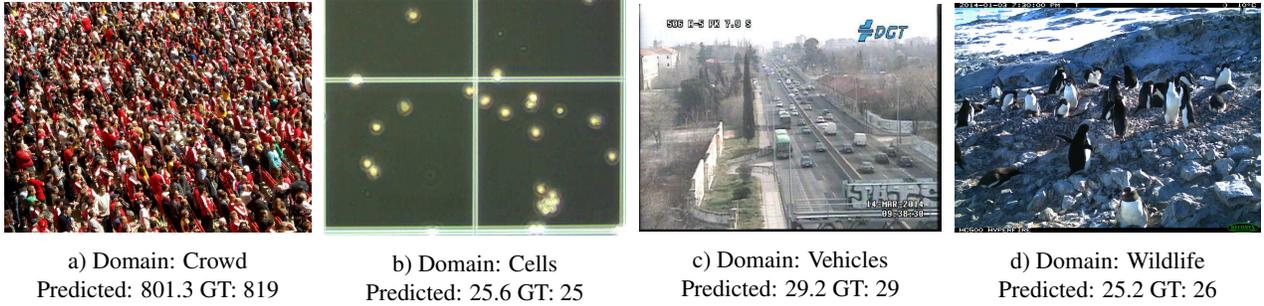


Figure 7: The robust performance of the proposed counting model is highlighted on images taken from each of the 4 visual domains used in this study.

Method	Part A		Part B	
	MAE	MSE	MAE	MSE
Zhang <i>et al.</i> [13]	110.2	173.2	26.4	41.3
Marsden <i>et al.</i> [16]	126.5	173.5	23.76	33.12
Switch-CNN [17]	90.4	135.0	21.6	33.12
Our-Approach	85.7	131.1	17.7	28.6

Table 8: Comparing the performance of various crowd counting approaches on the Shanghaitech dataset.

Method	MAE
Hydra-CNN [10]	10.99
Our Approach	9.7
Zhang <i>et al.</i> [23]	4.2

Table 9: Comparing performance of various vehicle counting approaches on the TRANCOS dataset.

Method	N=5	N=10	N=15
[18]	28.9 ± 22.6	22.2 ± 11.6	21.3 ± 9.4
Ours	23.6 ± 4.6	21.5 ± 4.2	20.5 ± 3.5
[11]	12.6 ± 3.0	10.7 ± 2.5	8.8 ± 2.3

Table 10: Cell counting MAE performance on the MBM dataset. Out of the 44 images in this collection, N are used for training, N for validation and an unseen 14 images for testing. At least 10 runs using random dataset splits are performed for the each N value.

both. The commonly used UCF_CC_50 dataset [15] has been deemed inappropriate for benchmarking as it contains crowds too dense for the human eye to count [34] and is therefore not used to evaluate the proposed technique. Table 9 compares the proposed technique with the leading vehicle counting techniques on the TRANCOS dataset while table 10 highlights cell counting performance on the Modified Bone Marrow dataset [11]. Finally, Table 11 shows the superior performance of our technique on the Penguins dataset [14].

Method	MAE
Arteta <i>et al.</i> [14]	8.11
Our Approach	5.8

Table 11: Comparing performance of various counting techniques on the Penguins dataset test set. MAE is computed w.r.t the max count on each image (as there are multiple annotators). The separate site dataset split is used and no depth information is utilised.

6. Conclusion

In this paper we propose a new multi-domain object counting technique that employs significant parameter sharing and achieves state-of-the-art benchmarking performance for several visual domains. This model can be extended to new counting tasks over time while still maintaining identical performance in all prior tasks. The benefits of this singular approach to object counting include the removal of redundant model parameters as well as noticeable increases in counting accuracy over single-domain baseline runs. The Dublin Cell Counting (DCC) dataset was also introduced as part of this study. This new collection is the first of its kind and contains a challenging and highly-varied set of cellular images. The developed single-model approach achieves state-of-the-art object counting performance in the Shanghaitech dataset (parts A and B) as well as the Penguins dataset. Future work in this area will look to extend the proposed counting model to perform additional regression and classification tasks.

Acknowledgements

This publication has emanated from research conducted with the financial support of the Irish Research Council and Science Foundation Ireland (SFI) under grant numbers SFI/12/RC/2289 and 15/SIRG/3283. A special thanks is also given to the Cummin’s group in the School of Medicine, University College Dublin for making the DCC dataset images available to the wider research community.

References

- [1] Sheng-Fuu Lin, Jaw-Yeh Chen, and Hung-Xin Chao. Estimation of number of people in crowded scenes using perspective transformation. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 31(6):645–654, 2001.
- [2] Bo Wu and Ramakant Nevatia. Detection of multiple, partially occluded humans in a single image by bayesian combination of edgelet part detectors. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 1, pages 90–97. IEEE, 2005.
- [3] Weina Ge and Robert T Collins. Marked point processes for crowd counting. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 2913–2920. IEEE, 2009.
- [4] Chen Change Loy, Shaogang Gong, and Tao Xiang. From semi-supervised to transfer counting of crowds. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2256–2263, 2013.
- [5] Ke Chen, Chen Change Loy, Shaogang Gong, and Tony Xiang. Feature mining for localised crowd counting. In *BMVC*, volume 1, page 3, 2012.
- [6] Victor Lempitsky and Andrew Zisserman. Learning to count objects in images. In *Advances in Neural Information Processing Systems*, pages 1324–1332, 2010.
- [7] Tongliang Liu and Dacheng Tao. On the robustness and generalization of cauchy regression. In *2014 4th IEEE International Conference on Information Science and Technology*, pages 100–105. IEEE, 2014.
- [8] Antoni B Chan and Nuno Vasconcelos. Counting people with low-level features and bayesian regression. *IEEE Transactions on Image Processing*, 21(4):2160–2177, 2012.
- [9] Cong Zhang, Hongsheng Li, Xiaogang Wang, and Xiaokang Yang. Cross-scene crowd counting via deep convolutional neural networks. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 07-12-June, pages 833–841, 2015.
- [10] Daniel Onoro-Rubio and Roberto J López-Sastre. Towards perspective-free object counting with deep learning. In *European Conference on Computer Vision*, pages 615–629. Springer, 2016.
- [11] Joseph Paul Cohen, Henry Z Lo, and Yoshua Bengio. Count-ception: Counting by fully convolutional redundant counting. *arXiv preprint arXiv:1703.08710*, 2017.
- [12] Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Learning multiple visual domains with residual adapters. *2017 Conference on Neural Information Processing Systems (NIPS)*, 2017.
- [13] Yingying Zhang, Desen Zhou, Siqin Chen, Shenghua Gao, and Yi Ma. Single-image crowd counting via multi-column convolutional neural network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 589–597, 2016.
- [14] Carlos Arteta, Victor Lempitsky, and Andrew Zisserman. Counting in the wild. In *European Conference on Computer Vision*, pages 483–498. Springer, 2016.
- [15] Haroon Idrees, Imran Saleemi, Cody Seibert, and Mubarak Shah. Multi-source multi-scale counting in extremely dense crowd images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2547–2554, 2013.
- [16] Mark Marsden, Kevin McGuinness, Suzanne Little, and Noel E O’Connor. Fully Convolutional Crowd Counting On Highly Congested Scenes. In *The International Conference on Computer Vision Theory and Applications*, 2017.
- [17] Deepak Babu Sam, Shiv Surya, and R Venkatesh Babu. Switching convolutional neural network for crowd counting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 1, page 6, 2017.
- [18] Weidi Xie, J Alison Noble, and Andrew Zisserman. Microscopy cell counting and detection with fully convolutional regression networks. *Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization*, pages 1–10, 2016.
- [19] Joseph Paul Cohen, Genevieve Boucher, Craig A Glastonbury, Henry Z Lo, and Yoshua Bengio. Count-ception: Counting by fully convolutional redundant counting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 18–26, 2017.
- [20] Weidi Xie, J. Alison Noble, and Andrew Zisserman. Microscopy cell counting and detection with fully convolutional regression networks. *Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization*, pages 1–10, 2016.
- [21] Ricardo Guerrero-Gómez-Olmedo, Beatriz Torre-Jiménez, Roberto López-Sastre, Saturnino Maldonado-Bascón, and Daniel Oñoro-Rubio. Extremely overlapping vehicle counting. In *Iberian Conference on Pattern Recognition and Image Analysis*, pages 423–431. Springer, 2015.
- [22] Di Kang, Zheng Ma, and Antoni B Chan. Beyond counting: Comparisons of density maps for crowd analysis tasks-counting, detection, and tracking. *arXiv preprint arXiv:1705.10118*, 2017.
- [23] Shanghang Zhang, Guanhang Wu, João P Costeira, and José MF Moura. Fcn-rlstm: Deep spatio-temporal neural networks for vehicle counting in city cameras. *arXiv preprint arXiv:1707.09476*, 2017.
- [24] Hakan Bilen and Andrea Vedaldi. Universal representations: The missing link between faces, text, planktons, and cat breeds. *arXiv preprint arXiv:1701.07275*, 2017.
- [25] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. 2015.
- [26] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159, 2011.

- [27] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Aistats*, volume 9, pages 249–256, 2010.
- [28] François Chollet. Keras. <https://github.com/fchollet/keras>, 2015.
- [29] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. 2015.
- [30] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [31] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [32] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- [33] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- [34] Yaocong Hu, Huan Chang, Fudong Nian, Yan Wang, and Teng Li. Dense crowd counting from still images with convolutional neural networks. *Journal of Visual Communication and Image Representation*, 38:530–539, 2016.