CS482/682 Final Project Report Group

MATHIAS: Modeling A Transferable Histopathological Image Analysis System

Hadley VanRenterghem/hvanren1, Matt Figdore/mfigdor1, Max Xu/mxu28, Taha Baig/tbaig3

1 Introduction

Background Interpreting histopathological images is a critical task for diagnosis and treatment guidance. Due to shortages of trained pathologists and expanding applications of histopathology, there is a need for algorithms to accurately diagnose clinically relevant features to increase the efficiency of image interpretation. However, annotating images to train models is expensive and time consuming, so unsupervised methods are needed to leverage all the available data [1].

Related Work Previous studies with histopathology images have focused primarily on supervised learning methods including a variety of CNN architectures, autoencoders, and SVMs. Current research has focused on minimizing overfitting during supervised training with small datasets through active learning [1], as well as unsupervised methods including feature extraction via restricted Boltzmann machines [2] and domain adaptation with adversarial training [3]. Recent work by Facebook AI Research demonstrated that unsupervised representation learning through deep clustering can generate features for image classification with state-of-the-art results on the ImageNet dataset [4].

Proposed Approach We apply a similar representation learning method to histopathological images by first finding a representation through unsupervised deep learning and then training a linear classifier to evaluate the features. For our target task, we have chosen distinguishing between different cell types in colorectal cancer H&E images. We then explore the possibility of applying this representation to other domains.

2 Methods

Dataset Our data is composed of 5,000, 150 x 150 pixel slices of H&E stained colorectal tissue slides. These images are labelled with 8 cell types, including tumor. We split the data into 80% for train, 10% for validation, and 10% for test, with even class distributions. During training, we augmented the data using horizontal and vertical flips.

Deep Clustering Our first method learns an unsupervised representation by alternating between clustering learned features with k-means and predicting the cluster assignments. Specifically, the features are created with the AlexNet CNN architecture, and the clusters are predicted by adding a final classification layer. The k-means clusters are compared with the cluster predictions to calculate the cross-entropy loss. We chose k-means for its simplicity, and AlexNet for a balance between strong image classification and computational efficiency. A variety of different hyperparameters, including batch size, total epochs, and learning rate, was tested but the most impactful was the number of clusters used for the deep clustering, with the optimal amount as 120.

Autoencoder Our second method learns an unsupervised representation with an autoencoder by encoding a latent representation of the image in a lower-dimensional subspace and then reconstructing the original image. The autoencoder is trained with a mean-squared error loss between the input image

and the decoded image. Our architecture uses 4 convolutional and 4 convolutional transpose layers.

Evaluation To evaluate our learned representations, we used them as input features for a logistic regression classifier, similarly to Facebook's Deep Clustering paper [4]. We trained the classifier to predict the cell type using labelled images. As a baseline, we trained a logistic regression classifier using the entire image as the input. Hyperparameters for the logistic regression classifier were kept constant. The best model maximizes the test accuracy of this classifier. To further evaluate the robustness of the features, we froze the best model trained on the colon images and used it to generate features on breast H&E histology images. We then retrained the logistic classifier with these features to detect tumors, and retrained the baseline classifier with the new raw images.

3 Results

Our best deep clustering model obtained 88.26% classification accuracy on the test set, whereas the baseline classifier obtained 58.46% (see confusion matrices). The encoder representation achieved an accuracy of 35.69%.

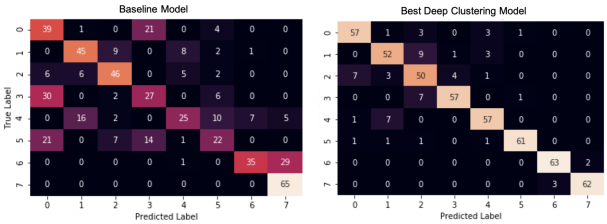

Tumor classification is especially important, and the AUROC for the tumor class with our deep clustering model is 0.97, whereas the baseline's is 0.89.

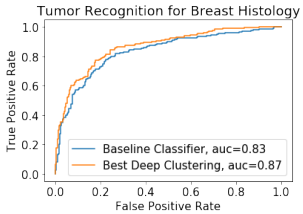

When we attempted domain transfer to breast histology images, our model obtained 81.25% classification accuracy, while the baseline classifier obtained 77.31% and a lower AUC value (see figure).

4 Discussion

Deep Clustering The high test accuracy in our results indicates that our unsupervised learning approach was able to successfully generate an informative representation of the images in the colorectal histology dataset. To construct an even more robust representation, we would integrate a larger set of unlabelled data to train the model. We would also like to experiment with the loss function by jointly optimizing the k-means loss and the cluster prediction loss, which could directly incentivize the model to produce representations that result in more discriminative clusters. Finally, we would like to vary the convolutional network structure by experimenting with other architectures such as VGG16 in addition to AlexNet, since VGG16 typically performs better on image classification tasks [7].

Autoencoder The loss in the autoencoder is designed to be generative and not discriminative, so the latent representation was unable to capture features relevant to classification tasks as reflected by the poor test accuracy. In order to introduce an explicit discriminative component, we would integrate a deep clustering technique to be used in conjunction with an autoencoder. By minimizing the combination of the reconstruction loss and the clustering loss, with a higher bias towards the clustering loss, this model would be able to simultaneously discriminate and preserve local structure [8].

Domain Transfer Upon applying the representation learned from the colorectal cancer H&E images on the new domain of breast H&E histology images, we were able to see an increase in classification accuracy and AUC compared to the baseline model, demonstrating our model's generalizability. To improve further, we would train the model with a wider variety of histology images [4], and integrate a better method of stain normalization [2].

5 References

[1] Raczkowski, Ł., Możejko, M., Zambonelli, J. et al. ARA: accurate, reliable and active histopathological image classification framework with Bayesian deep learning. Sci Rep 9, 14347 (2019) doi:10.1038/$s41598-019-50587-1$

[2] Ren, Jian, et al. "Unsupervised Domain Adaptation for Classification of Histopathology Whole-Slide Images." Frontiers in Bioengineering and Biotechnology, vol. 7, 2019, doi:10.3389/$fbioe$.2019.00102.

[3] Sari, Can Taylan, and Cigdem Gunduz-Demir. "Unsupervised Feature Extraction via Deep Learning for Histopathological Classification of Colon Tissue Images." IEEE Transactions on Medical Imaging, vol. 38, no. 5, 2019, pp. 1139–1149., doi:10.1109/tmi.2018.2879369.

[4] Caron, Mathilde, et al. "Deep Clustering for Unsupervised Learning of Visual Features." Computer Vision – ECCV 2018 Lecture Notes in Computer Science, Mar. 2018, pp. 139–156., doi:10.1007/$978-3-030-01264-9_9$.

[5] Kather, J. N. et al. Multiclass texture analysis in colorectal cancer histology. Scientific Reports (2016).Return to ref 29 in article.

[6] Alom, Md Zahangir, et. al. "A State-of-the-Art Survey on Deep Learning Theory and Architectures." Electronics, vol. 8, no. 3, May 2019, p. 292., doi:10.3390/electronics8030292.

[7] Guo, Xifeng, et al. "Deep Clustering with Convolutional Autoencoders." SpringerLink, Springer, Cham, 14 Nov. 2017,