

# Deep Medical Image Analysis with Representation Learning and Neuromorphic Computing

Getty, N.<sup>1,2</sup>, Brettin, T.<sup>1,3</sup>, Stevens. R.<sup>1,3</sup>, Xia, F.<sup>1,3</sup>

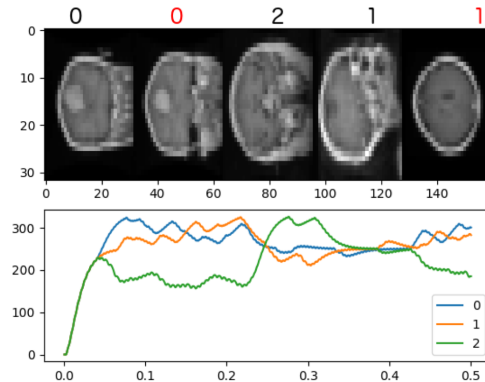
<sup>1</sup>Argonne National Laboratory, Lemont, IL, U.S.A.

<sup>2</sup>Illinois Institute of Technology, Chicago, IL, U.S.A.

<sup>3</sup>University of Chicago, Chicago, IL, U.S.A.

## 1. Overview

Deep learning is increasingly used in medical imaging, improving many steps of the processing chain, from acquisition to segmentation and anomaly detection to outcome prediction. Yet significant challenges remain: (1) Image-based diagnosis depends on the spatial relationships between local patterns, something convolution and pooling often do not capture adequately; (2) data augmentation, the de facto method for learning 3D pose invariance, requires exponentially many points to achieve robust improvement; (3) Labeled medical images are much less abundant than unlabeled ones, especially for heterogenous pathological cases; and (4) Scanning technologies such as magnetic resonance imaging (MRI) can be slow and costly, generally without online learning abilities to focus on regions of clinical interest. To address these challenges, novel algorithmic and hardware approaches are needed for deep learning to reach its full potential in medical imaging.



**Figure 1** Brain tumor image classification with a Spiking Neural Network on Intel's neuromorphic hardware

We explore three representative lines of research and demonstrate the utility of our methods on a classification benchmark of brain cancer MRI data. First, we present a capsule network that explicitly learns a representation robust to rotation and affine transformation. This model requires less training data and outperforms both the original convolution baseline and a previous capsule network implementation. Second, we leverage the latest domain adaptation techniques to achieve a new state-of-the-art accuracy. Our experiments show that non-medical images can be used to improve model performance. Finally, we design a spiking neural network trained on the Intel Loihi neuromorphic chip (Fig. 1 shows an inference snapshot). This model consumes much lower power while achieving reasonable accuracy given model reduction. We posit that more research in this direction combining hardware and learning advancements will power future medical imaging (on-device AI, few-shot prediction, adaptive scanning).

## 2. Results

We tested our methods on a representative medical imaging dataset [1]. This benchmark contains 3,064 MRI slices of 233 patients diagnosed with one of the three brain tumor types: meningioma (23%), glioma (47%), and pituitary tumor (30%). We compare our results with a convolution network baseline and a previous state-of-the-art capsule network implementation.

**Table 1** Comparing our three approaches with two previous models

Classification Model	Test Accuracy	Sample Efficiency (Test accuracy with 10% training data)	Avg. Energy (Measured or estimated*)
ResNet trained from scratch	83.7%	61.9%	329 W
Previous CapsNet [1]	85.6%	N/A	N/A
Our CapsNet	89.3%	<b>75.3%</b>	365 W
Pre-trained ResNet	<b>92.4%</b>	64.7%	329 W
SpikingNet on Intel Loihi chip	70.2%	N/A	2-8* W

Table 1 summarizes model performance differences in terms of test accuracy, sample efficiency, and energy consumption. We stratified the images both by patient and for a balanced tumor classification and used a conservative 30% split for the test set. Our domain adaptation approach, implemented with a pretrained residual network, achieved a new state-of-the-art accuracy of 92.4%. Among models trained from scratch, our capsule network had the highest classification accuracy. To assess sample efficiency, we measured model performance trained on a fraction of the data ranging from 10% to 50%. In these low data scenarios, the capsule network consistently outperformed other models by a large margin (over 10% with 10% training data). Finally, in terms of energy consumption during the inference phase, we measured average energy draw for the models on GPU (Nvidia V100) and estimated it for the spiking network on neuromorphic hardware based on previous experiments. The spiking network unsurprisingly had the lowest consumption with roughly 50 times less energy.

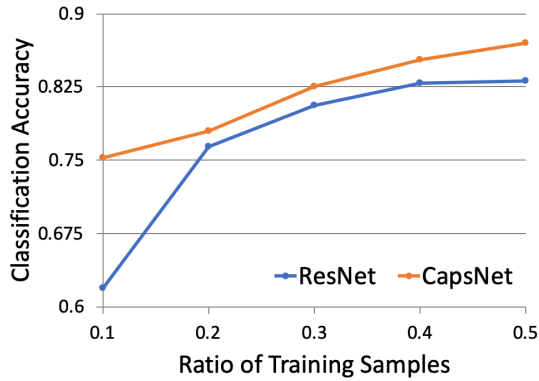
## 3. Methods and Discussion

### Spatial representation learning with Capsule Networks

Capsule networks [2] are a divergence from typical convolution neural networks which rely on pooling to move from simple features detected early in the network to higher level features. What we lose in pooling is spatial resolution. While other techniques such as dilated convolution attempt to assuage this loss, capsules focus on explicitly capturing spatial hierarchies between simple and complex patterns through dynamic routing. The advantage is some level of representational invariance to image transformations including pose, lighting or deformation. This advantage makes capsules a good fit for medical images with diffuse, diverse abnormalities as well as variances from both measuring devices and individual differences.

We replicated the capsule network architecture [1] by Guo [3] and optimized the preprocessing and model parameters for the raw MRI tumor images rather than the derived tumor

segmentation masks. Furthermore, we utilized learning rate decay, larger batch sizes and longer training time. Ultimately our trained model marks a significant performance improvement on the raw MRI images. While the accuracy did not surpass the pertained ResNet model, it outperformed the ResNet model trained from scratch. The capsule model is also data efficient, with performance degrading less substantially as the amount of labeled data decreases (Fig. 2).



**Figure 2** Low data learning: classification accuracy as a function of the fraction of training examples

Another advantage of capsule networks is explainability. By manipulating the vector output of the capsule layer and reconstructing the image with the jointly learned decoder, one can visualize what each capsule is learning. This is especially vital in clinical settings where transparency may improve both doctor and patient trust and has the potential to uncover previously unremarkable image-based biomarkers that are relevant for diagnosis and prognosis.

### Transfer learning

Our best model in terms of classification accuracy is a 50-layer convolutional residual network, improving on the previous state-of-the-art by 6.8%. The training process, implemented using the FastAI framework, is as follows: We started with the pretrained weights for ImgeNet, and first trained only the top fully-connected layer while keeping the residual blocks frozen. At this stage, the model reached 87.1% test accuracy. We then unfroze all layers and assigned them differential starting learning rates, giving more freedom to the higher layers. A second round of training with cyclical learning rates boosted the final test accuracy to 92.4%.

In an ablation experiment to quantify the value of pretrained weights, we trained the same model architecture from scratch. This baseline model achieved 83.7% test accuracy. This comparison points to an 8% boost from domain adaptation. Although ImageNet contains no medical images, the learned lower-level features were able to generalize to the brain scans. Recent efforts in building large-scale medical imaging databases will likely provide an even better base for transfer learning.

### Spiking network on neuromorphic device

Despite the tremendous success of AI, we understand very little about the human intelligence algorithm. Unlike the point neuron model used in deep learning, real neurons are connected to thousands of excitatory synapses hypothesized to recognize multiple independent patterns.

Neuromorphic computing aims to get one step closer to how brain works with spiking signals and local learning rules. Rather than employing a non-linear activation function, spiking neurons must reach an activation potential before generating an output spike to forward connections and resetting. Learning rules are assigned to synapses, defining how the weight, delay and tag are computed during learning as a function of the pre and post synaptic traces. Neuromorphic hardware is an asynchronous design architecture specifically for implementing spiking based neural networks.

We have partnered with Intel to experiment with their recent research chip, Loihi [4], for potential medical imaging application. What currently sets Loihi apart from other neuromorphic platforms is the online learning capacity and system scalability. We implemented three spiking networks on and off chip using Nengo [5] and the Loihi SDK: (1) A single layer image classifier, trained on chip, which encodes every pixel's intensity via a random poisson spike generation process. (2) A converted convolutional neural network which successfully ran on the chip simulator but resulted in the actual chip timing up. (3) A spiking fully connected network (Table 1) which dramatically reduces the total number of neurons and processing time by first preprocessing images using dimensionality decomposition. All three models are energy efficient. Loihi is still at early development, so we expect the model performance to improve with reduced software constraints.

#### **4. Conclusions**

We have demonstrated three different approaches to improve medical imaging analysis tackling data efficiency, transfer learning and on-device learning. We posit the application of neuromorphic computing coupled with new learning algorithms will stimulate advancements in imaging technologies. Today's MRI scans are slow and often require patients to be still for extended periods. Acceleration and enhancement may be realized at several steps of the process: (1) Real-time adaptive scanning implemented to increase spatial and temporal resolution at automatically detected regions of interest. (2) Accelerated data acquisition by intelligent compressive sampling and reconstructive techniques to reduce number of slices and required TR/TE times per slice. (3) Faster algorithm implementation of the final image reconstruction. (4) Real-time, few-shot learning for image enhancement techniques including motion correction, anomaly detection and anatomical segmentation.

#### **5. References**

- [1] Afshar, P., Mohammadi, A. and Plataniotis, K.N., 2018, Brain tumor type classification via capsule networks. In *2018 25th IEEE International Conference on Image Processing (ICIP)*.
- [2] Sabour, S., Frosst, N. and Hinton, G.E., 2017. Dynamic routing between capsules. In *Advances in neural information processing systems* (pp. 3856-3866).
- [3] Guo, X. "A Keras implementation of CapsNet." <https://github.com/XifengGuo/CapsNet-Keras>, 2018.
- [4] Davies, M., Srinivasa, N., ... and Liao, Y., 2018. Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro*, 38(1), pp.82-99.
- [5] Bekolay, T., Bergstra, ... and Eliasmith, C., 2014. Nengo: a Python tool for building large-scale functional brain models. *Frontiers in neuroinformatics*, 7, p.48.