Using simulated data to train convolutional neural networks for image registration

Osagie K. Aimiuwu[1,2], Daniel Tward[3]

[1]BIG Summer Program, Institute for Quantitative and Computational Biosciences, UCLA
[2]University of North Carolina, Chapel Hill
[3]Department of Computational Medicine, Department of Neurology, UCLA

Background/Motivation

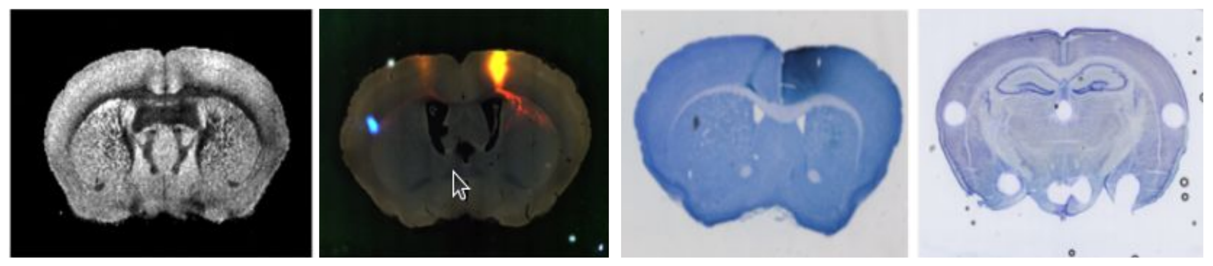

Real Mouse Brain Images. *(Left to right)* Serial two photon tomography, fluorescence with rabies injections, immunohistochemistry, Nissl stain

- Image registration is the process of aligning multiple images to the same coordinate system and is crucial in medical imaging analysis. Registration is complicated by images having varying contrasts or deformations.
- Previous deep learning methods for registration have relied on real data to train the model. However, it is not always possible to acquire large enough datasets to accurately train a model.
- As such, the goal was to develop a convolutional neural network that generated limitless simulated data by randomly coloring and transforming labeled mouse brain images (seen below).

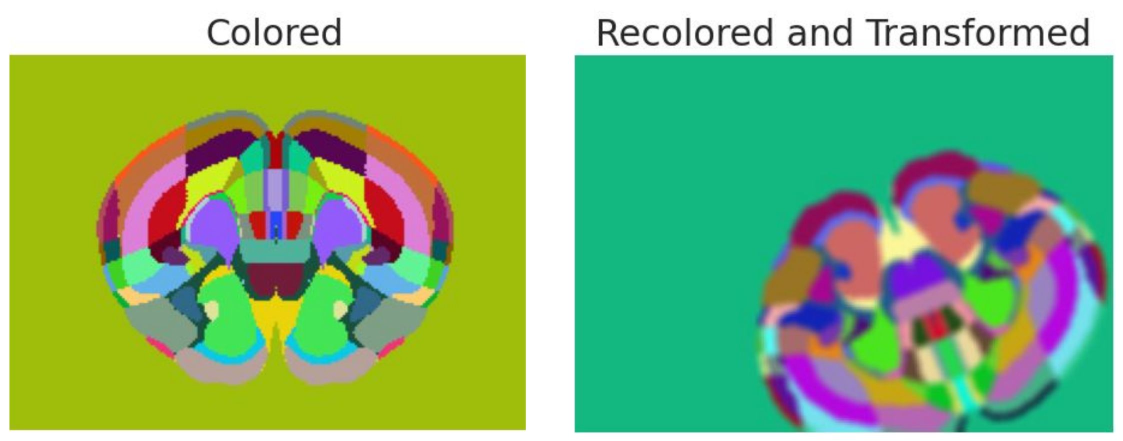

Labeled Mouse Brain Images. *(Left)* Image with different colors for each brain structure. *(Right)* Image recolored and with a random affine transformation applied.

Methods/Results

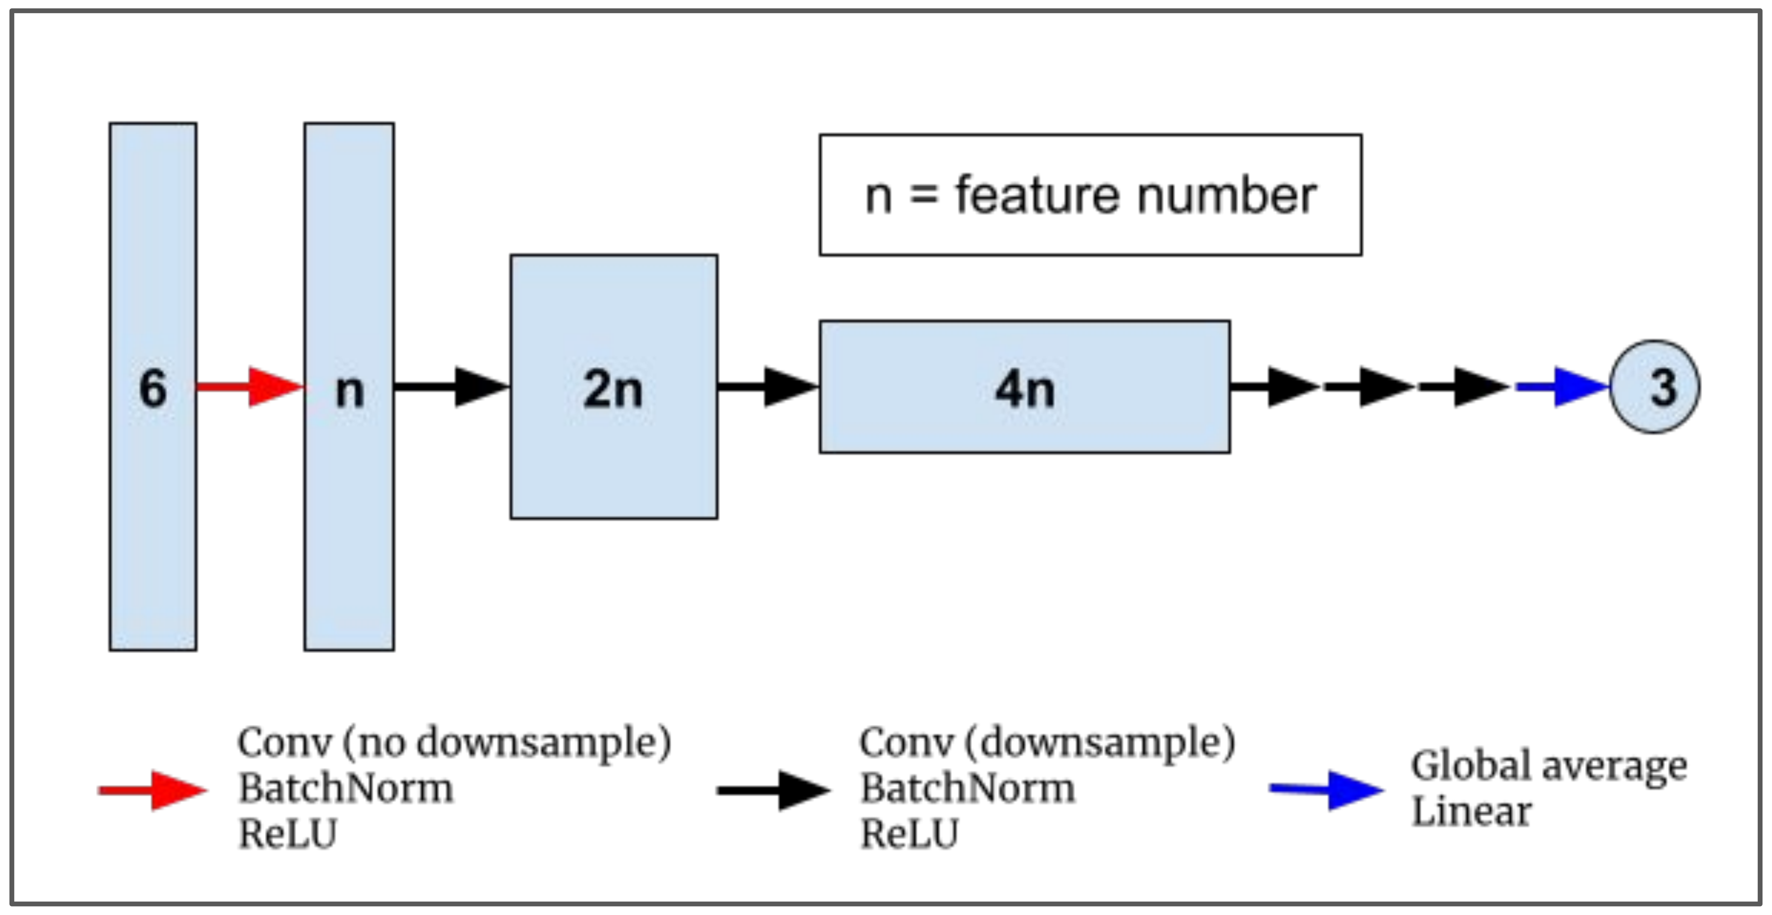

Network Architecture. Network where the number of convolutional layers with downsampling and the number of features (n) could be varied. The height represents the number of pixels.

Eight different networks were created by varying network depth, number of features, and the learning rate during training. We used PyTorch and the Adam optimizer and trained each network for 2,000 epochs. After evaluating accuracy on the validation data, **Network 1** was found to perform the best.

Network	Depth	Features	Learning Rate
1	7	64	0.0010
2	7	64	0.0001
3	7	32	0.0010
4	7	32	0.0001
5	6	64	0.0010
6	6	64	0.0001
7	6	32	0.0010
8	6	32	0.0001

Tested Networks. Networks that were trained and evaluated when finding optimal hyperparameters.

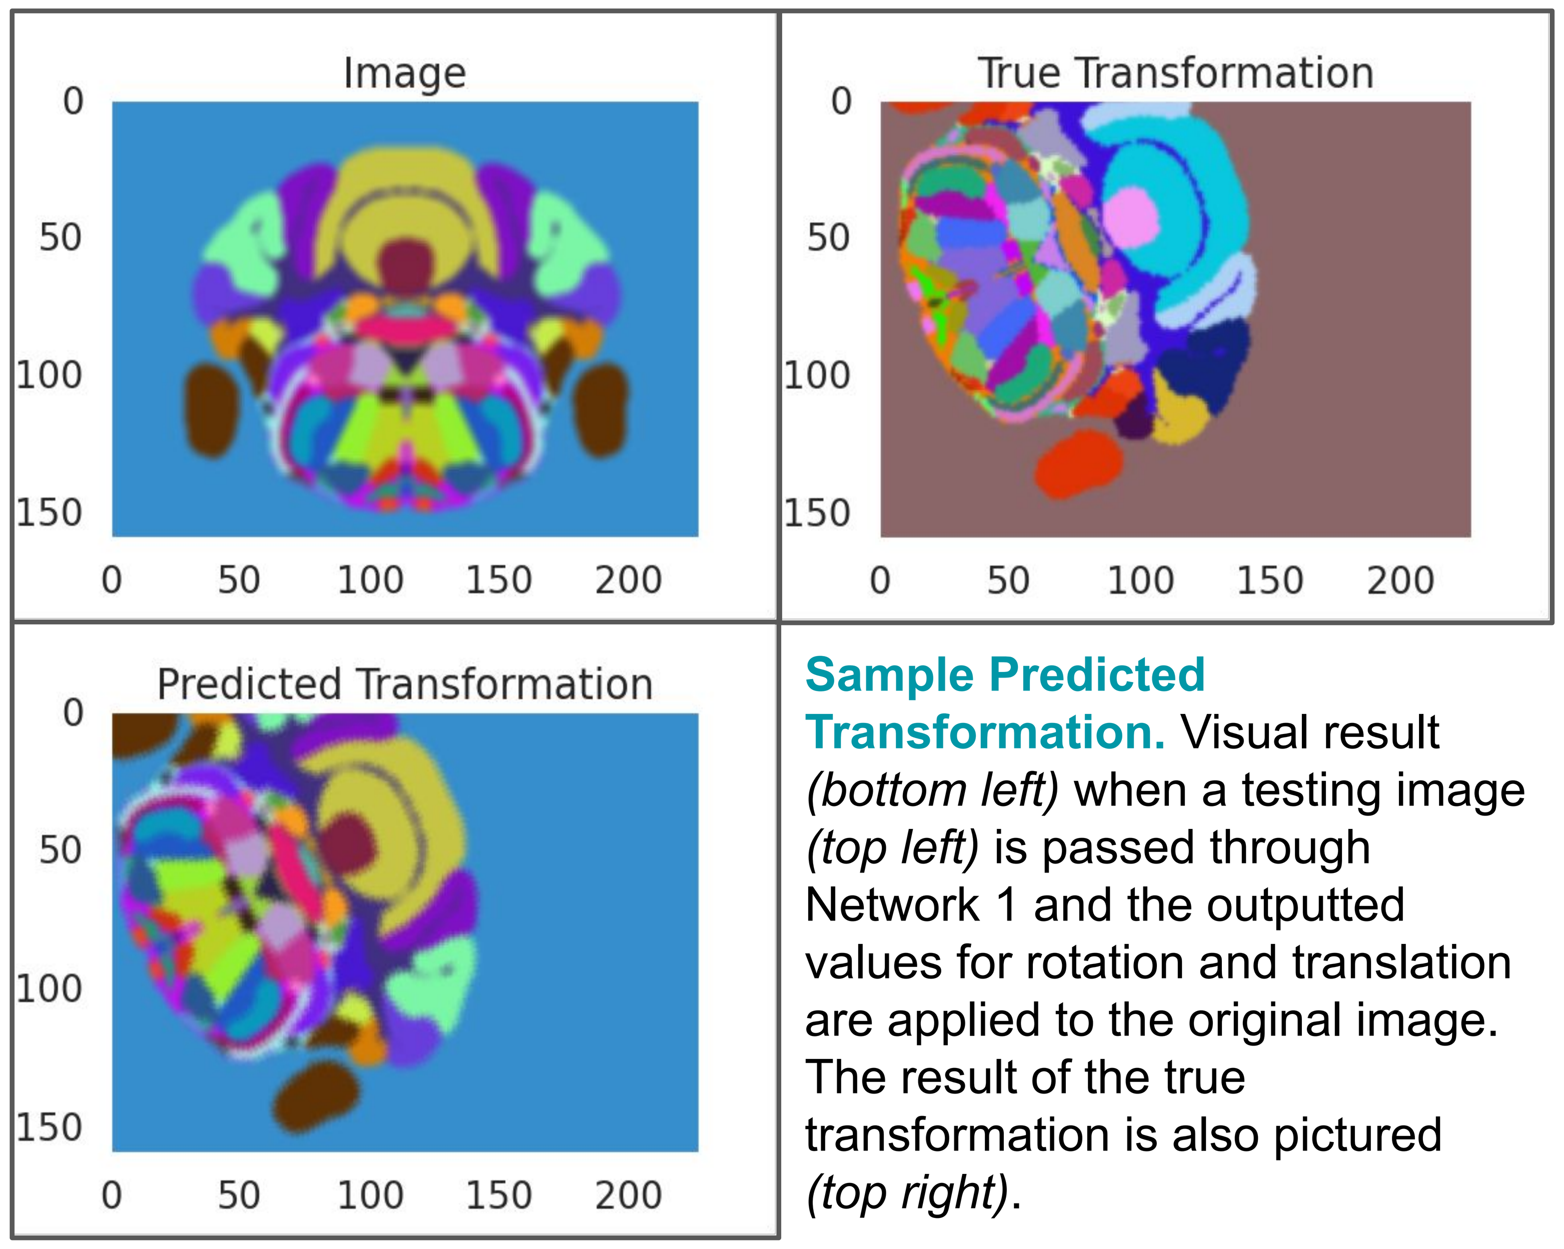

Sample Predicted Transformation. Visual result *(bottom left)* when a testing image *(top left)* is passed through Network 1 and the outputted values for rotation and translation are applied to the original image. The result of the true transformation is also pictured *(top right)*.

Network 1 was then applied to the test data. The results was a root mean squared error (RMSE) of 11.246 degrees for rotation and 1.378 pixels for translation.

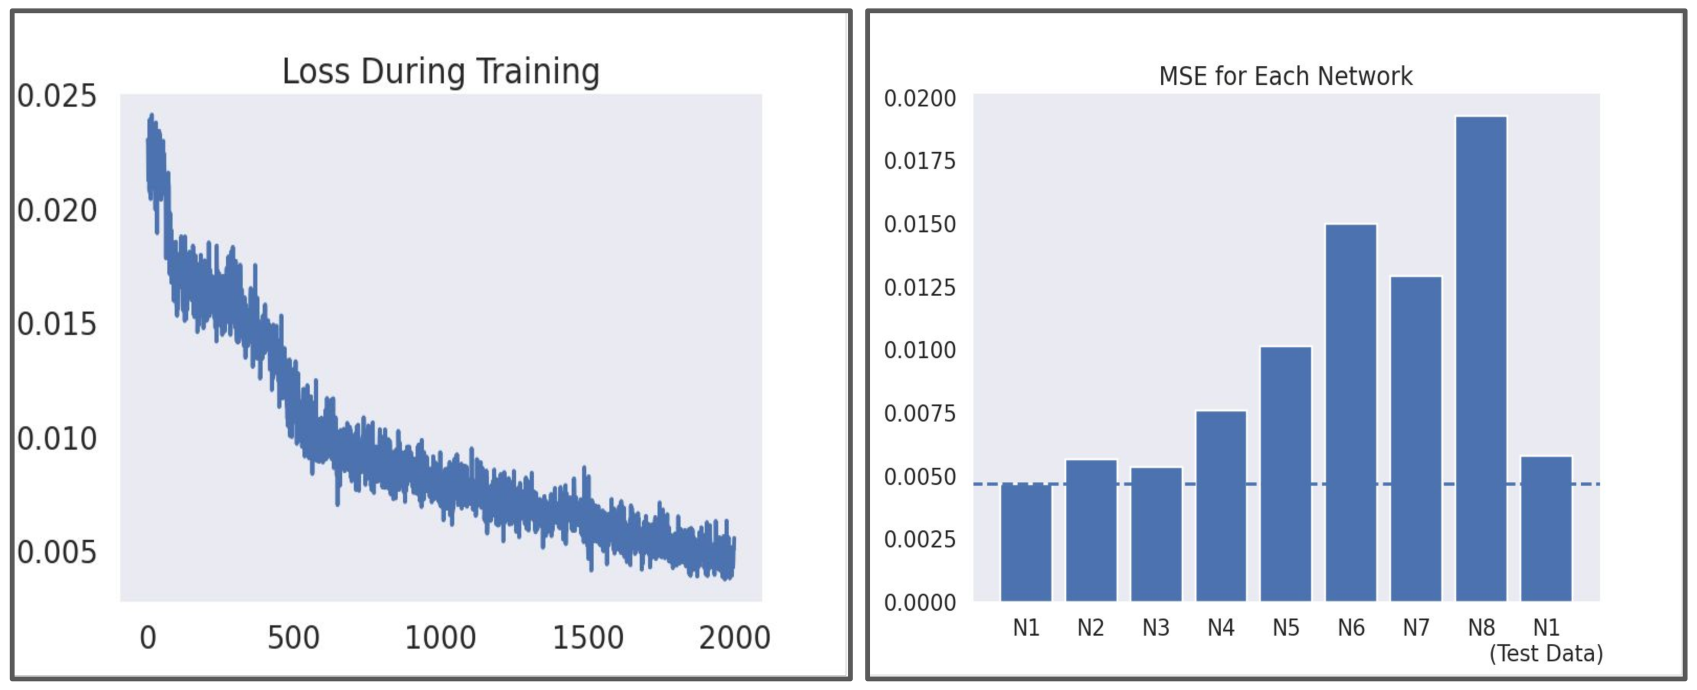

Loss Function. *(Left)* Mean squared error (MSE) for Network 1 during training. *(Right)* MSE for each network during evaluation and for Network 1 when applied to the test data.

Conclusions

Network 1 can be used to quickly align images of different modalities or stains, and was developed using a novel simulated data framework that generalizes to other applications. Future directions include training it to predict non-affine transformations.

Our loss function, MSE, accounted for rotation by using the fact that, for example, +179 and -179 degrees are very close to one another; however, the network architecture sees them as being different by 358 degrees. Our desired function is therefore not continuous and is not well modeled by a continuous neural network. More work is required to model maps to these types of spaces (sets that are not vector spaces).

Acknowledgements

National Institute of Health: U19MH11482, R01NS074980, RF1MH126732, RF1MH128875, RF1MH128888, RF1MH128861, R01 NS121761, U01AG073804
The Karen Toffler Charitable Trust
BIG Summer Program