From Neuroscience to Artificially Intelligent Systems (NAISys), CSHL, April 5th - 9th 2022

WHY DO BRAINS HAVE 2 HEMISPHERES AND WOULD IT IMPROVE DEEP LEARNING? — A STUDY WITH LEFT AND RIGHT NEURAL NETWORKS

Gideon Kowadlo (gideon@cerenaut.ai), Chandramouli Rajagopalan
Cerenaut (cerenaut.ai)

Left and Right Brains

The brains of all bilaterally symmetric animals on Earth are Bicameral - they are divided into left and right hemispheres. It's a *remarkably* conserved feature across species, indicating its importance for intelligence.

The **anatomy** and **functionality** of the hemispheres have a large degree of overlap, but they specialize to possess different attributes. One likely explanation for the emergent specialization is small differences in parameterisation of the substrate e.g. synaptic plasticity, resource allocation (fig. 2) and connectivity patterns within and across layers [1, 2, 3, 4], see fig. 1.

Despite the importance of hemispheric specialization, it is still poorly understood and has not been exploited in AI.

Mimicking Left and Right

The hypothesis:

1. Specialization can be achieved by mimicking substrate differences with architectural and functional features.
2. An ANN with differentially specialized sub-networks can outperform a single comparable network on classification spanning both general and specific classes. I.e. Each image belongs to both a specific and a general class

General and Specific classes, an example:
General = *sea creature*, Specific = *penguin, seal, shark etc.*

Method: We tested the hypothesis with two phases (see fig. 3):

1. We investigated general-specific specialization with supervised training of single hemispheres on *just* Specific or *just* General classes (fig. 3a). The single hemispheres are standard Resnet9. We then also introduced architectural and functional variation.The configurations were thus:
   A. Simple Specialization: Only supervision is used to elicit specialization.
   B. Resource Allocation: Pyramid and Inverted Pyramid architecture represents resource allocation in the brain, fig. 2. The Inverted Pyramid with standard Resnet9, and Pyramid, a Resnet9 with modified layer widths.
   C. Sparsity: Other studies show that the level of sparsity can influence the spatial 'breadth' of features [5], which is important for the General/Specific distinction. Implemented with top-*k* 'winner take all' competitive learning [5].
   D. Combination of B and C above.
2. After demonstrating specialization, we constructed a bicameral network with pre-trained differentially specialized hemispheres, by adding a single fully connected (FC ANN) 'combiner' layer (fig. 3b) and tested it on multilabel classification, implemented with dual heads, one each for Specific or General.

The task: Image classification using the CIFAR100 dataset, which has hierarchical labels to indicate General and Specific classes.

Baselines: compared to comparative networks without specialization, both with dual heads for each class type:

- Bicameral, no specialization, 2x Resnet9 hemispheres (similar to fig. 3b)
- Resnet18, which has approximately the same number of trainable parameters as our bicameral network of 2x Resnet9.

Experiment details:

- We conducted hyperparameter searches for each experiment. The best results are included, the others are omitted for brevity.

Acknowledgements

This project came about from Elkhonon Goldberg's research and ideas. Thank you Elkhonon for the thought-provoking discussions. In an earlier Masters project, co-supervised by Levin Kuhlmann, Amir Maayan explored similar ideas, which formed a valuable background to planning this project. Finally, thanks to Punarjay Chakravarty and David Rawlinson for discussions and write-up reviews.

Specialization
Specific class definitions and familiarity/routine

Attributes
Pyramid: More neurons near input layers (fig. 2)

Short local pathways, modular network organisation.



Specialization
General classes and novelty

Attributes
Inverted pyramid: More neurons near output layers (fig.2)

Long inter-regional pathways, small world network i.e. shallow network.

*Figure 1: **Hemispheric differences:** specialization and attributes.*

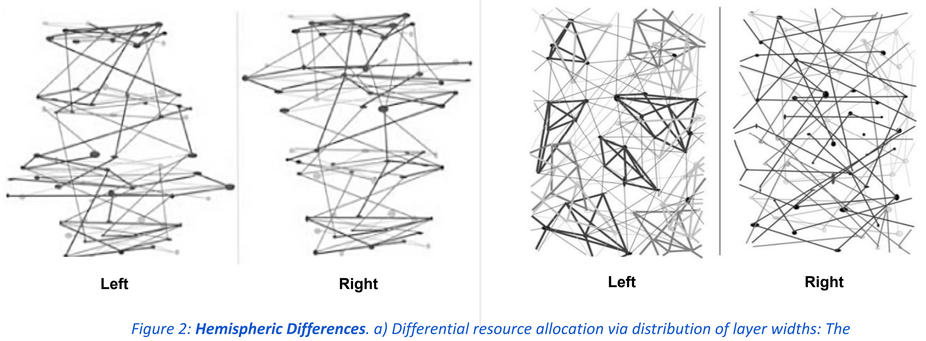


*Figure 2: **Hemispheric Differences**. a) Differential resource allocation via distribution of layer widths: The left has a pyramid architecture, the right an inverted pyramid, b) short local pathways vs long inter-regional pathways. Images reproduced from [2].*

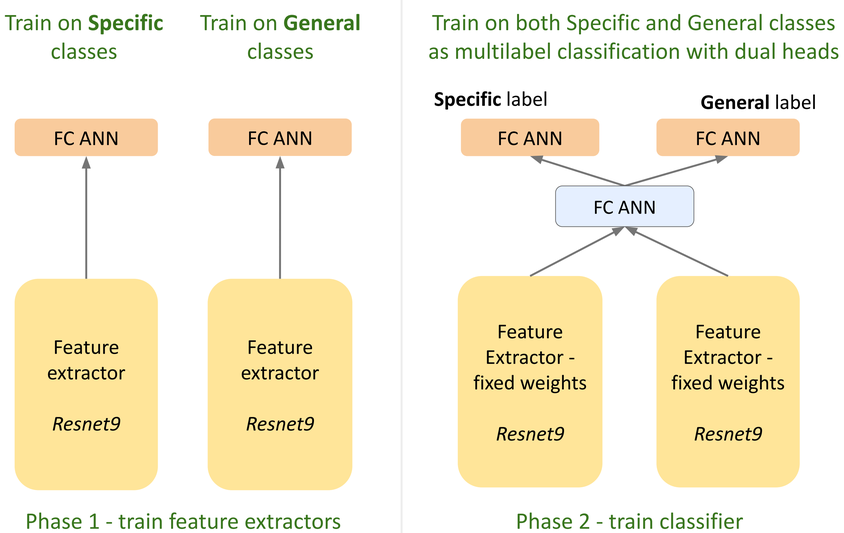


*Figure 3: **Method overview**: First feature extractors are trained, then fixed and used in a bicameral architecture.*

Related Work

Beaulieu et. al (2020) created a dual-stream architecture called Neuromodulator Meta-learner, where one network learns to modulate the other, to enhance continual learning. We did not find examples of bicameral architectures that exploit specialization inspired by biological hemispheric differences.

Results

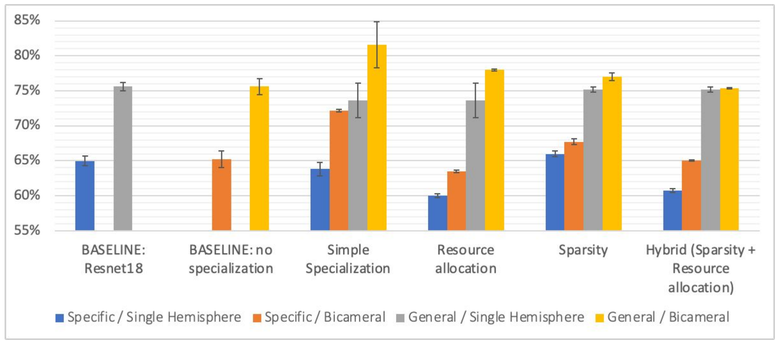


*Figure 4: **Results**: The tests are grouped by architecture.
Refer to how the score changes from single to two hemispheres, for both Specific and General classes, compared for different architectures.
Error bars show 1 std dev, from 5 seeds.*

Specialization:

- Simple Specialization (supervision alone) had the most significant effect
- Sparsity a slight effect
- Resource allocation increased differential specialization, although decreased absolute performance on both classe types
- Sparsity + Resource allocation, the combination was surprisingly deleterious

Bicameral network:

- The simple Bicameral architecture had the best outcome, boosting performance compared to all baselines by ~7% on Specific, and ~6% General.
- Resource Allocation and Sparsity performed better on General, but not Specific Classes.
- *Combining specialized hemispheres always conferred an advantage over a single hemisphere, even when it didn't outperform baseline.*

Put another way, the ability of the classifier on General classes, is better off with representations that are specialized for both General and Specific, than just General alone, and vice versa. Since each image belongs to a General and Specific class, they are informative about the other. The inductive bias of enforcing specialized representations has benefits above what it learns from the explicit objective of multiclass classification alone.

What does it mean?

The results demonstrate that small changes to network parameterisation can achieve specialization, and that specialization can be complementary and beneficial for certain tasks. It is important because it provides evidence that biological hemispheric differences arise this way and it provides insight into which parameters are likely to be most important. Secondly, it shows that neuroscience principles can provide inductive biases for novel deep-learning AI architectures.

Future work will explore more complex biologically inspired interactions between left and right throughout the hierarchy (not just with a combined head).

References

1. E. Goldberg et al., "Hemispheric asymmetries of cortical volume in the human brain", Cortex, 2013, 49(1), pp. 200-10.
2. E. Goldberg, "The New Executive Brain", Oxford University Press, 2009
3. E. Goldberg and L. Costa. Hemisphere differences in the acquisition and use of descriptive systems, Brain and Language, 1981, 14, pp. 144-173
4. E. Goldberg et al., "Cognitive bias, functional cortical geometry, and the frontal lobes: laterality, sex, and handedness", J Cogn Neurosci. 1994, 6(3)
5. A. Makhzani and B. Frey, "k-Sparse Autoencoders", arXiv:1312.5663v2, 2013
6. S. Beaulieu, L. Frati, T. Miconi, J. Lehman, K. O. Stanley, J. Clune and N. Cheney, "Learning to Continually Learn," in 24th European Conference on Artificial Intelligence - ECAI 2020