

Statement of Research Interests

Branton DeMoss

My research interests involve the nature of emergence, which I investigate using tools from algorithmic information theory, machine learning, and (soon) fractal geometry.

When can we say emergence has occurred? An approach I take in the paper attached to this application is to re-frame the question around compression, using the Minimum Description Length (MDL) principle. If we have a model M that we use to explain some dataset D , a naïve model selection criterion is to choose the model which achieves the best predictive performance on D . How can we understand whether a “simpler” model is a better choice compared to a more complex one which performs better? To give a poetic example, consider a puddle of water: the underlying dynamics of water are certainly governed by quantum mechanics. However, if we have limited memory and computational ability, we prefer to have access to coarse-grained hydrodynamic information, rather than information about the quantum state of the puddle. Though the quantum description is more complete and can in principle be used to predict the behavior of the puddle with greater accuracy, it requires considerably more information and computational resources to achieve this. With only a few bits of information about macroscopic hydrodynamic variables, we can achieve good prediction of the puddle’s dynamics, if we allow for some inaccuracy.

The MDL principle makes this trade-off sharp, and formalizes Occam’s Razor by asserting that we should minimize the following sum:

$$\text{Total Description Length of } D \text{ Given } M = H(D | M) + C(M) \quad (1)$$

Where $H(D | M)$ is the entropy of the data under the model, and $C(M)$ is the model complexity. That is, the best model for the data is the one which minimizes the *sum* of the entropy of the data and the model complexity. Hence, we can identify emergence by the transition of *our preference* for one model to another, as determined by the MDL. This also makes it clear that emergence is relative to our computational resources: Laplace’s demon has no need for emergent descriptions

because its computational resources are unbounded. In the attached paper, I extended this idea to formalize the notion of “coarse-graining” a model, then used the resultant machinery to explain the grokking phenomenon in neural networks, where networks exhibit a sudden phase transition from memorization of their training data to perfect generalization. Intriguingly, we observed a characteristic rise and fall of complexity in the networks during the transition from memorization to generalization. This same rise and fall of complexity was reported by Scott Aaronson and SFI affiliate Sean Carroll in their work which studied the complexity dynamics of a cup of coffee. My next work will explain the double descent phenomenon in terms of network complexity dynamics.

Machine learning researchers tend to use the tools of classical information theory, probability, and statistics. Although successful, these tools limit analysis to stateless processes, and cannot provide a full account of the algorithmic nature of many phenomena in the natural world. I envision a research program which moves beyond the statistical paradigm of random variables and stochastic processes to an *algorithmic* information paradigm, whose mathematical foundations were laid out by Kolmogorov, Chaitin, and Solomonoff. We have good reason to believe that much of the apparent complexity we observe is in fact due to repeated iteration of simple rules. It may be possible to bridge the algorithmic complexity of computation with the apparent complexity of the natural world using ideas from fractal geometry. It is already known that the fractal (effective Hausdorff) dimension d_{eff} of a set X can be understood in terms of its Kolmogorov complexity K , via:

$$d_{\text{eff}} = \liminf_n \frac{K(X | n)}{n} \quad (2)$$

That is, the effective dimension of objects is deeply linked with their computational complexity. This is another sign of emergence: *effective* dimensions which arise from irreducibly complex description. Machine learning must evolve to account for regularities beyond statistical description. The renormalization group and effective field theories gave us our best yet analytical understanding of the nature of emergence. Neural networks, Kolmogorov complexity, and fractal geometry will do the same for analytically intractable complex systems.