Qeios

Peer Review

Review of: "Neural Networks Use Distance Metrics"

Marcos Eduardo Valle[1]

1. Universidade Estadual de Campinas, Brazil

This paper explores the impact of perturbations on a neuron's inductive field (activation strength) prior to the application of the activation function. The author specifically examines how the inductive field of a hidden neuron in a single-hidden-layer multilayer perceptron network influences a classification task. Through computational experiments, the observations suggest that the network's performance tends to decline more rapidly with perturbations affecting the offset compared to those impacting scale or clip values.

In reviewing the paper, I find the topic to be quite intriguing; however, there seems to be a need for greater mathematical rigor. The author acknowledges in the conclusion that “the lack of a precise mathematical definition for intensity metrics makes it difficult to definitively rule out intensity-based interpretations.” Additionally, I believe there is a similar need for clarity regarding the definition of the distance metric, which raises some questions about whether the reported results align with the conjecture that neural networks utilize distance metrics. Accordingly, the author notes, “A distance metric quantifies how far an input is from a learned statistical property of the data. While an intensity metric reflects a confidence level — larger values indicate higher certainty that the input belongs to the node’s feature set.” However, the relationship between these concepts and the perturbations on the offset, scale, and clip values could benefit from further clarification.

Beyond this primary concern, there are a few other aspects that might warrant further elaboration and consistency. Below are some observations I noted in the manuscript:

1. In Section 4, equation (1) indicates that the MLP network contains 784 hidden neurons; however, lines 31 and 53 of models.py refer to networks with 128 hidden neurons.

2. The computational experiments outline varying configurations, yet it is unclear if the scale is defined as scale=(1-p)r or scale=p.

 doi.org/10.32388/UYB556

3. The relative magnitudes of perturbations on the offset, scale, and clip values also appear ambiguous, particularly since the value of r used in experiments could not be found.

4. Section 4.2 states, “Intensity and cutoff range over [1%...1000%],” while the plot in Figure 1 extends from 0% to 120%.

5. The methodology for creating the plots in Figure 2 could benefit from further explanation.

6. Lastly, while the paper claims that “neural networks use distance metrics,” it primarily focuses on an MLP network with a single hidden layer. Drawing broader conclusions about modern deep learning architectures based on this framework may not be entirely straightforward.

Overall, I believe that addressing these points could enhance the clarity and rigor of the paper.

Declarations

Potential competing interests: No potential competing interests to declare.