Rethinking Crowd-Sourced Evaluation of Neuron Explanations

Tuomas Oikarinen, Ge Yan, Akshay Kulkarni, Tsui-Wei (Lily) Weng - **UCSD**

GitHub: https://github.com/Trustworthy-ML-Lab/efficient_neuron_eval

Paper: https://arxiv.org/abs/2506.07985

ICML International Conference On Machine Learning

UC San Diego

Trustworthy ML Lab @ UCSD
lilywenglab.github.io

Motivation: Existing Crowdsourced studies of neuron explanations only evaluate on highly activating inputs

- Only evaluating highly activating inputs is equivalent to only measuring **Recall**, and ignores whether the concept is present on low activating inputs
- We conduct the first study with a principled metric, Pearson's correlation coefficient
- Evaluating correlation can be very expensive due to need to annotate all inputs and rater noise
  - We propose efficient sampling and error correction strategies to reduce total cost **~60x**

Our Contributions:
1st crowdsourced study measuring correlation coefficient + ~60x cost reduction by efficient sampling and error correction

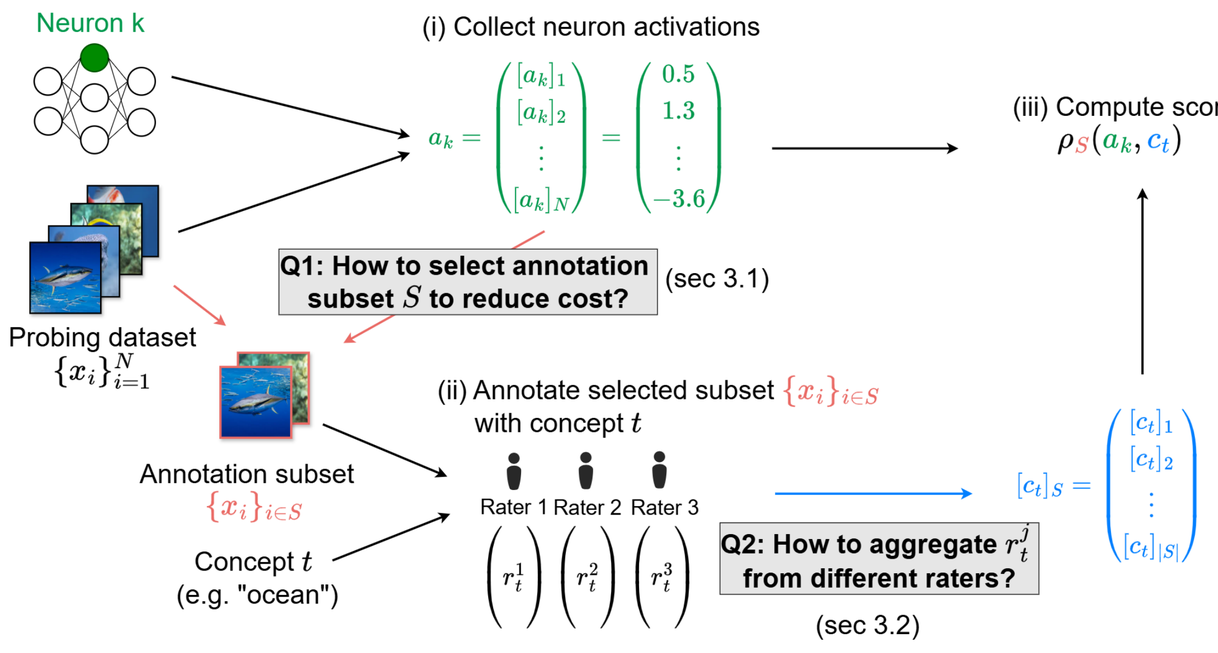

Methods

Contribution 1: Importance Sampling

Rating every input for each concept is not feasible

- Need to sample a subset of inputs to show raters
- We choose samples with Importance sampling (with correction) from distribution q that approximates the theoretical optimum

$$q(x_i) \propto \frac{1}{|\mathcal{D}|} \left| [\bar{a}_k]_i \cdot [\bar{c}_t^{siglip}]_i + \epsilon \right|$$

Contribution 2: Bayes with SigLIP prior

- Crowdsourced ratings are noisy -> Multiple Raters per input
- We show we can get more accurate results by using Bayes rule to estimate P(c | r_1, r_2, …) over typical methods like majority vote

Combined these, we can reduce study cost from $45,000 to $800 with same accuracy!

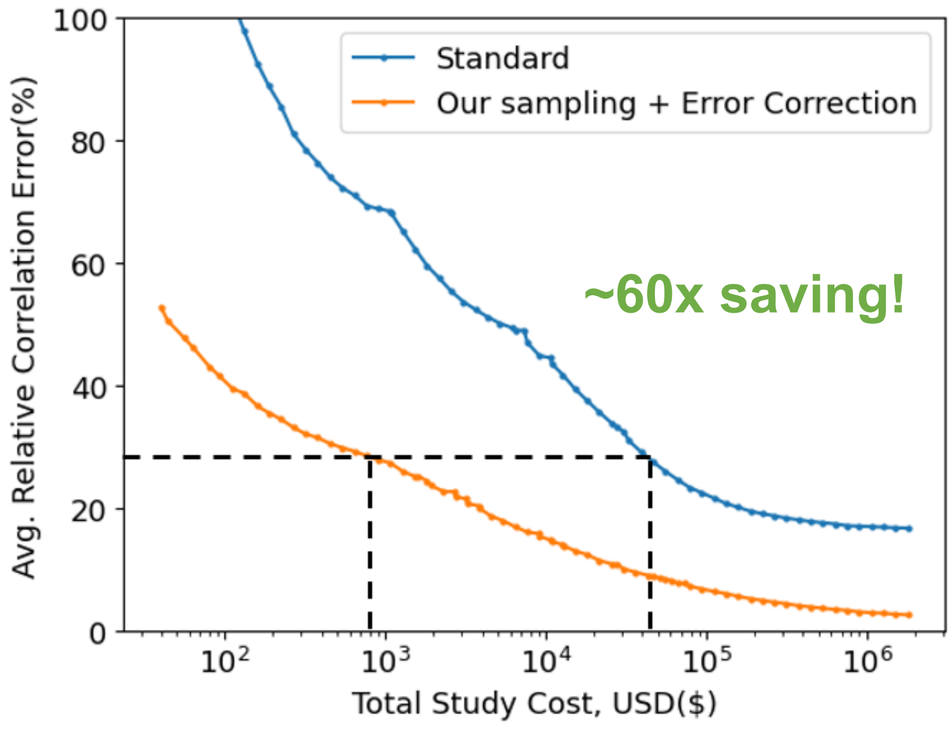

Results

- We evaluated explanations generated by best existing automated interpretability methods for 100 random neurons on two vision different networks
- Linear Explanation LE(SigLIP) performed the best on both Networks studied, even when restricted to produce length 1 explanations
- Notable LE significantly outperformed recent generative model-based methods MAIA and DnD
- Overall correlations relatively low, highlighting the need for more complex explanations or more interpretable architectures

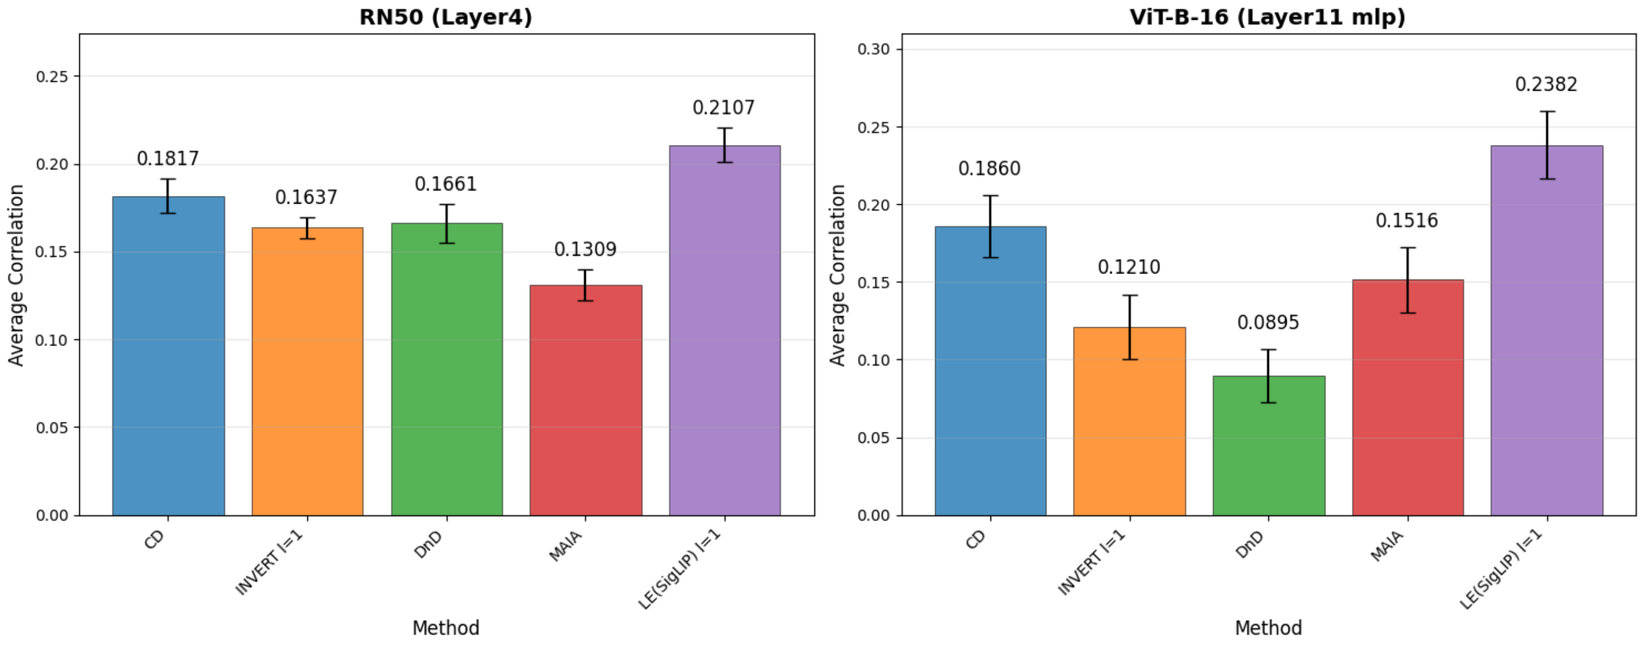

User Interface

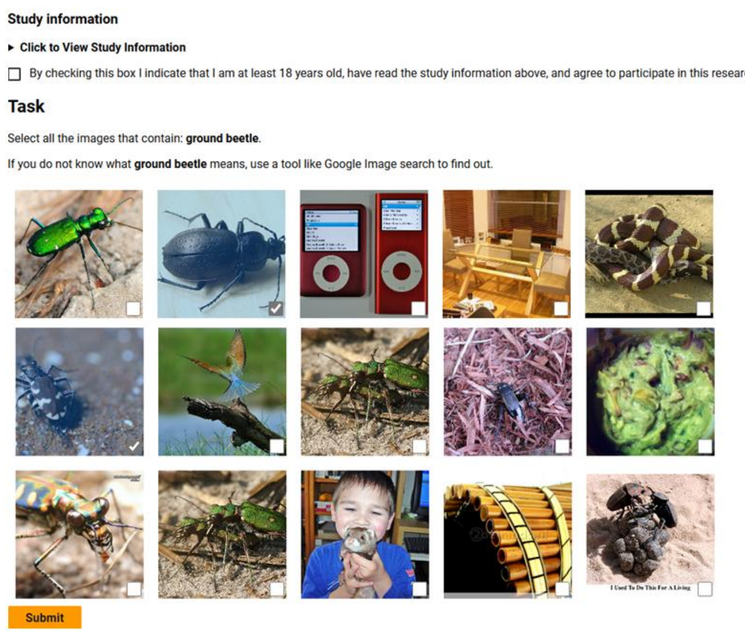

Paper

Code