

Peer Review

Review of: "TabularGRPO: Modern Mixture-Of-Experts Transformer with Group Relative Policy Optimization GRPO for Tabular Data Learning"

Russell D. Wolfinger¹

1. Independent researcher

Research on novel tabular architectures is always welcome, and the authors appear to have some good ideas. I only have time for a brief review and would like to suggest the following improvements for the examples and comparisons with other common approaches:

1. Add Cross Entropy Loss as a metric, since this is what is actually being optimized for binary targets.
2. Add LightGBM and even another model or two, like TabNet and FTTransformer, given their popularity.
3. Use leak-free, repeated k-fold cross-validation for all examples and report standard errors on performance metrics to provide a much better form of comparison.
4. Include hyperparameter tuning for all models, performed in a nested, leak-free manner within each fold, e.g., scikit-learn RandomizedSearchCV.
5. Add one more example with a continuous target.

Declarations

Potential competing interests: No potential competing interests to declare.