Augmenting internal data with population-level statistics from external sources could improve model robustness to data-shift

Learning Robust Models from Limited External Statistics

PRESENTER: **Tal** El Hay

INTRO

- Model robustness is usually assessed by external validation
- In a previous work, we developed a method that estimates model performance on external data sources from their limited statistical characteristics

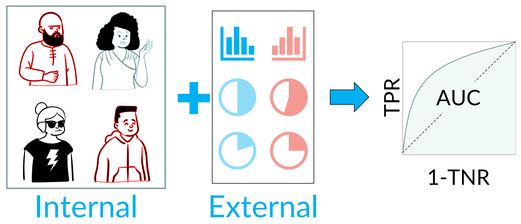

- Can we adopt a similar approach to train robust models, alleviating privacy concerns and communication costs?

METHODS

1. Search for weights that reproduce external statistics; generate a weighted copy of internal data with external characteristics.

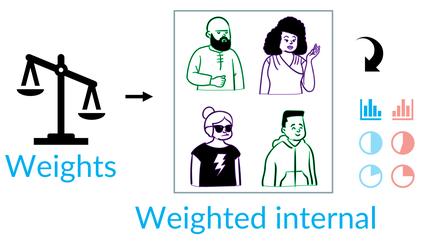

2. Train model on internal and weighted sets

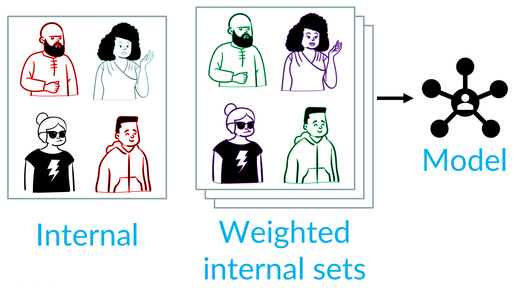

AIM

Train robust predictive models using:

- various machine learning algorithms
- patient-level internal data + population-level statistics from external sources
- a single (or very few) communication round

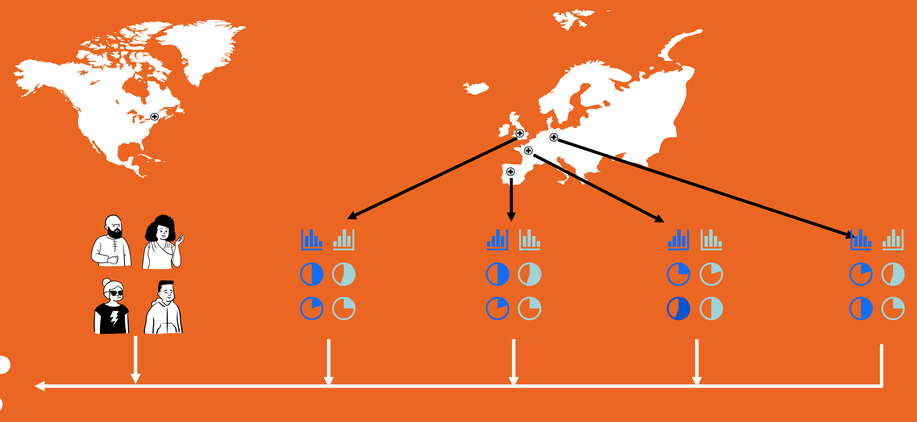

External performance (AUC) of a model trained on internal data degrades faster than for models trained on pooled data or using external statistics and reweighting

Scan for details about the **reweighting algorithm**

Estimating Model Performance on External Samples from Their Limited Statistical Characteristics, *Conference on Health, Inference, and Learning (CHIL) 2022*

Covariate shift, logistic regression

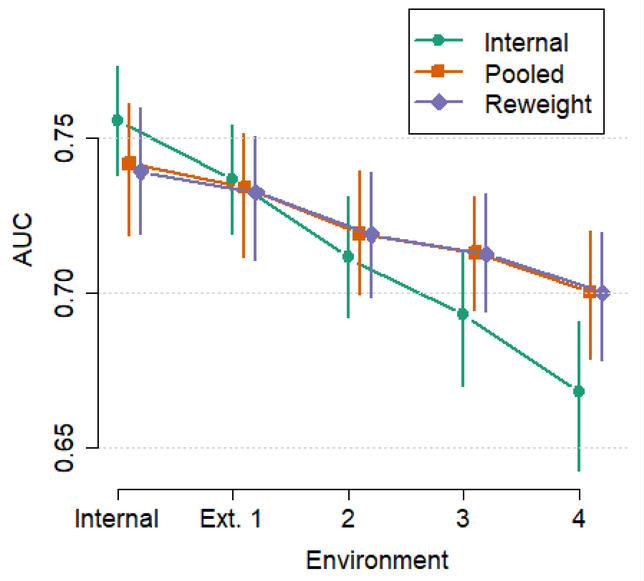

ADDITIONAL RESULTS

1. **Model averaging may fail with linear models** that combine environment-dependent and invariant predictors using Lasso.

2. XGBoost: reweight model only slightly improves over internal one

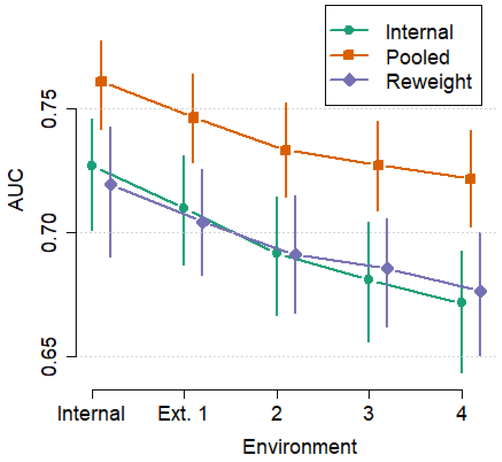

3. Neural network: reweight model outperforms internal but is not as good as pooled one.

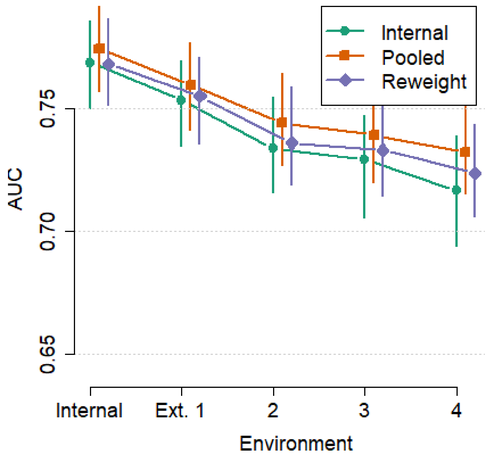

DISCUSSION

Strengths. requires only limited statistics (can use info from characterization studies); a single communication round

Limitations. may fail if insufficient statistics are used; suboptimal in comparison to pooled training.

Future directions. Adapt the method to non-linear models; optimize the choice of stats; introduce a distributionally robust objective.

Tal El Hay and Chen Yanover

The Israeli Institute for Applied Research in Computational Health

OHDSI