Active Learning under Label Shift

Eric Zhao, Anqi Liu, Anima Anandkumar, Yisong Yue

RealML@ICML2020

Caltech

The Label Shift Problem

How to efficiently label despite distribution shift?

Label shift: certain classes are over/under-represented.

$P(X \mid Y) = Q(X \mid Y)$ but $P(Y) \neq Q(Y)$

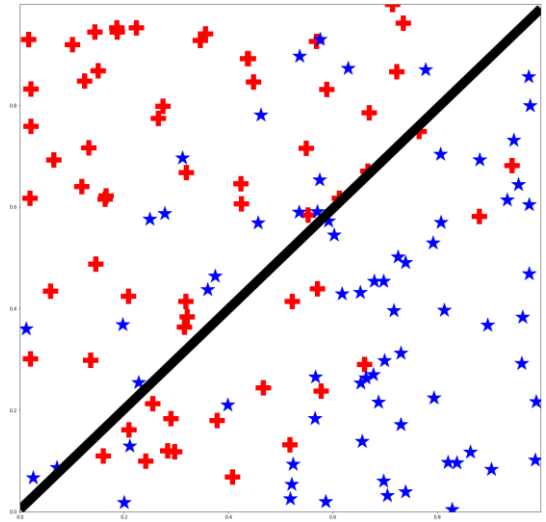

Uniform Distribution

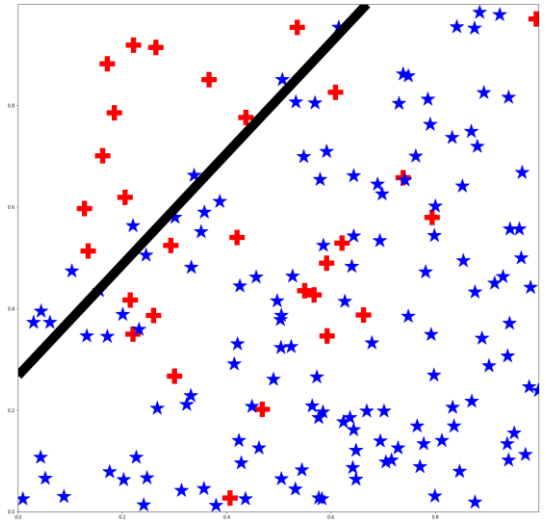

Label Shift

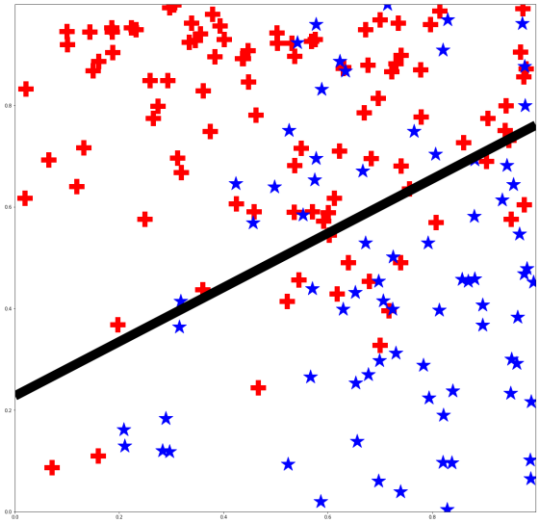

Covariate Shift

The Label Shift Problem

Label shift may arise from:

- Societal bias
- Class imbalance

Distorts uncertainty and disagreement measures.

Gender Classifier	Darker Subjects Accuracy	Lighter Subjects Accuracy	Error Rate Diff.
Microsoft	87.1%	99.3%	12.2%
FACE++	83.5%	95.3%	11.8%
IBM	77.6%	96.8%	19.2%

Source: Gender Shades (Buolamwini, Gebru)

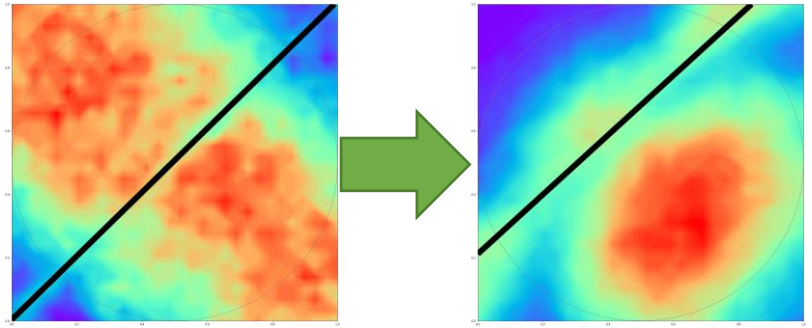

Change in sampling probability (DHM) due to label shift.

Active Learning under Label Shift

Option 1: subsample by class

Option 2: importance weighting

ALLS: do a little of both

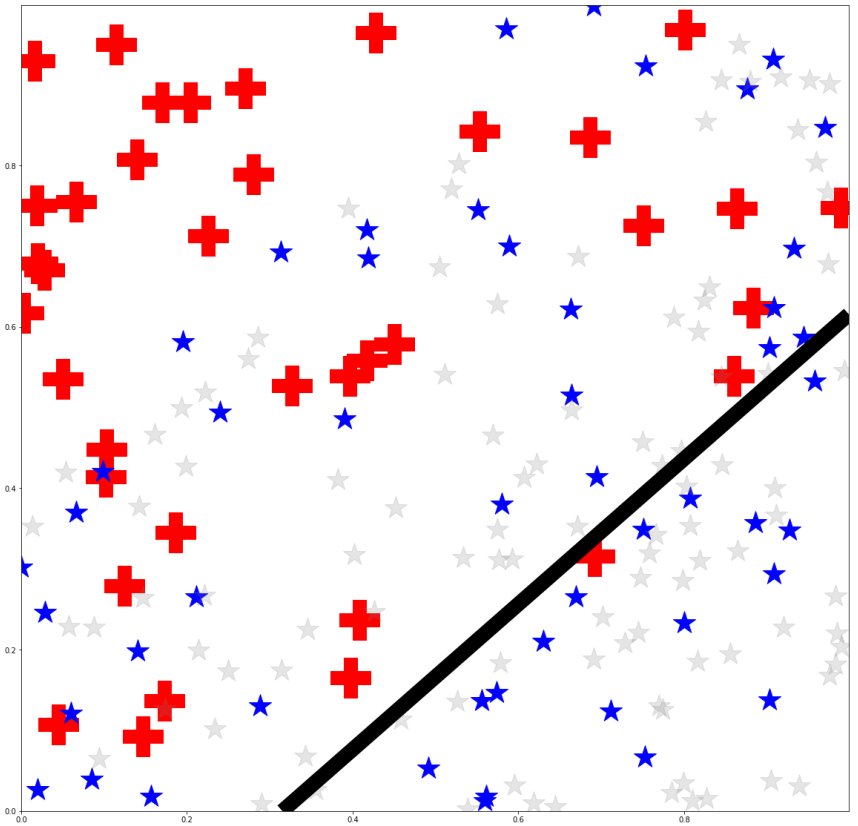

ALLS boosting red-class representation

Main Take-Aways

Preserves **consistency guarantee** with IW.

Reveals **bias-variance trade-off** between IW and subsampling.

Outperforms/matches with all explored settings/datasets.

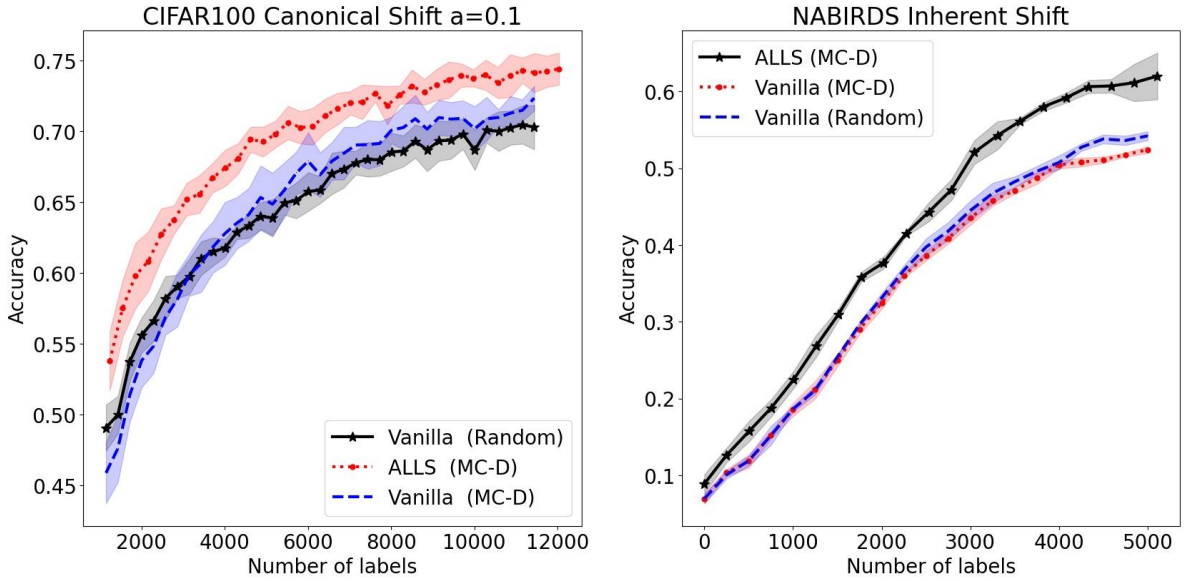

...and **scales** to deep learning models and large output spaces