A Probabilistic Approach to the Positive and Unlabeled Learning Problem

Kristen Jaskie, Andreas Spanias
SenSIP Center, School of ECEE, Arizona State University.

MOTIVATION

Traditional binary classification requires well-labeled data.

- **Both positive and negative labels:**

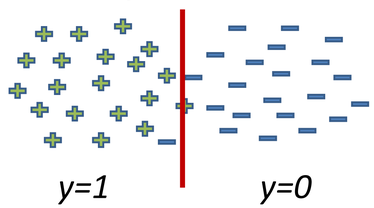

Negative data is EXPENSIVE in many interesting problems.

- **Ex: Cancer Detection**
  - **Known positive set:** People who have cancer
  - **Unlabeled set:** Everyone else
  - Finding true negatives – people who ABSOLUTELY do not have cancer – is expensive or impossible.
- **This leaves us with some positive and no negative labels.**
- **Other examples:**
  - Fraud detection
  - Terrorist detection
  - Threat detection

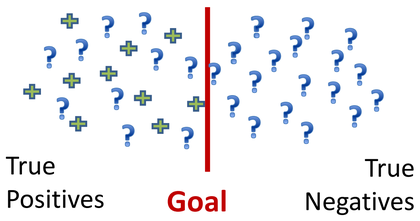

SCAR ASSUMPTION

- **We assume that labeled positives are "Selected Completely At Random" from the set of all positive samples.**
- **Means labeled and unlabeled sets are completely non-separable.**
- **Means that there is a constant probability c that a positive sample is labeled.**

GOAL

- **Given data samples x and data labels y**
- **We want to learn a probabilistic classifier $p(y=1|x)$**

PREVIOUS SOLUTION [1]

- Include a new random variable s:
  If sample is labelled, $s=1$, if unlabelled, $s=0$
- It can be shown that

$$p(y=1|x)=\frac{p(s=1|x)}{c}$$

- Used Standard Logistic Regression (SLR) to learn non-traditional classifier $p(s=1|x)$
- Constructed estimators for c using a validation data.

Found to be INEFFECTIVE in practice.

OUR SOLUTION

- Created a Modified Logistic Regression (MLR) to learn non-traditional classifier $p(s=1|x)$
- Introduced learned variable b into SLR equation:

$$MLR=\frac{1}{1+b^2+e^{-\bar{w}\bar{x}}}$$

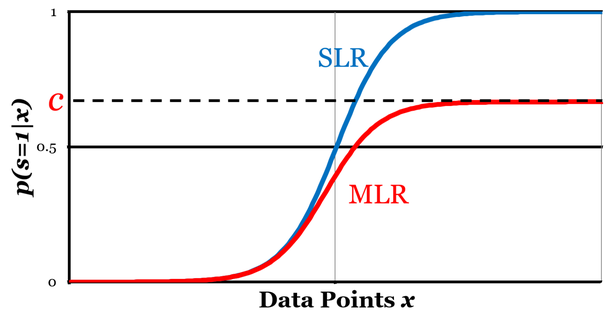

RESULTS

- **We compared SLR estimators in [1] with estimators using our MLR over different data sets.**
- **Metric: mean accuracy of c value estimate.**

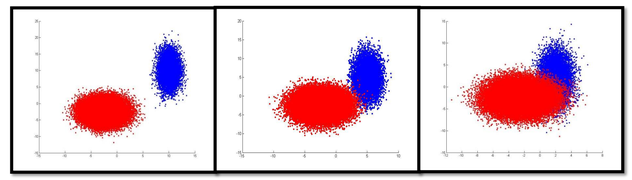

	Accuracy	
	MLR	SLR
Well Separated Data	98.32%	86.53%
Mostly Separable Data	95.38%	65.78%
Poorly Separable Data	90.58%	43.51%

SECURITY APPLICATION USING SENSORS

- **Threat detection on military bases or public venues**

Given sensor input such as audio, video, satellite

Known positive set: Previous attacks

Unlabeled set: Everything else

Just because an attack didn't occur, doesn't mean that a threat wasn't present – perhaps the attack was cancelled at the last minute.

REFERENCES

[1] C. Elkan and K. Noto. Learning classifiers from only positive and unlabeled data. In Proceedings of the Fourteenth International Conference on Knowledge Discovery and Data Mining (KDD '08).