Bayesian Bias Mitigation for Crowdsourcing

Fabian L. Wauthier and Michael I. Jordan

University of California, Berkeley

One-minute Summary:

Bayesian Bias Mitigation for Crowdsourcing (BBMC) unifies three steps of the crowdsourcing pipeline in one Bayesian model. We model the *sources* of bias and can thus account for more complex bias patterns. Active learning is commonly considered infeasible with Gibbs sampling inference. We propose a general approximation strategy for Markov chains and specialize it to active learning with Gibbs sampling in our model.

Overview

Labelers in the crowd may be malicious/unhelpful or tasks ambiguous or hard.

- **Problem: How to learn from biased data?**
- **Currently**: Bias is often addressed separately during *data collection*, *curation*, and *learning*:
  - Data collection: Active learning queries useful labelers.
  - Data curation: Screening or weighting of training data.
  - Learning: Noisy observation model captures bias *effects*
- **Contribution:** Unify all three steps in a Bayesian model.
  - Model the *sources* of bias explicitly. Any labeler not labeling uniformly at random is considered "biased".
  - Labelers are influenced by **accumulated**, **shared** biases.
- **Benefits:**
  - Instead of masking one true labeling by noise, we allow **multiple** "truths" (i.e. biased responses) to coexist.
  - Bayesian inference automatically combines data curation and learning and allows us to drive active learning.
  - Model-based: Don't need to collect multiple labels for each task. Can make predictions for new tasks.
- **Problem: Active learning for Gibbs sampling?**
- **Contribution:** Approximate the stationary distribution of a perturbed Markov chain from an unperturbed chain.
- **Benefits:**
  - General purpose strategy that transfers to other Markov chain problems.
  - Significantly faster scoring of candidate designs than naïve MCMC scoring, while yielding superior results.

Notation

- Tasks i, labelers l, biases b.
- Task covariates $x_i, i = 1, \ldots, n$ in X.
- Labeler labels $y_{i,l} \in \{-1, 0, +1\}, i = 1, \ldots, n; l = 1, \ldots, m$ in Y. Value $y_{i,l} = 0$ indicates no label from labeler l on task i.

Nonparametric Bias Model

- Model *sources* of labeler biases by binary matrix Z. Rows are labelers, columns are biases. If $z_{l,b} = 1$, labeler l expresses bias b.
- Model *effect* of bias b on labels by shared parameter γ_b.
- Parameter β_l predicting l's labels accumulates biases:

$$\beta_l = \sum_{b=1}^{\infty} z_{l,b}\gamma_b$$

- **Likelihood:**

$$p(Y|X, Z, \gamma) = \prod_l \prod_{i:y_{i,l}\neq 0} p(y_{i,l}|\beta_l^\top x_i) = \prod_l \prod_{i:y_{i,l}\neq 0} \Phi(y_{i,l}\beta_l^\top x_i)$$

- **Prior on** Z: Finite approximation to Indian Buffet Process prior. For K bias features Z is $m \times K$. Conditional:

$$p(z_{l,b} = 1|z_{-l,b}) = \frac{\sum_{l'\neq l} z_{l',b} + \frac{\alpha}{K}}{n + \frac{\alpha}{K}}$$

- **Prior on** γ_b: $p(\gamma_b) = \mathcal{N}(0, \sigma^2 I)$ for each b.

Inference

- Suppose r is target labeler who provided some labels $y_{i,r}$.
- **Want:** $p(y_{j,r} = 1|X, Y)$ for unlabeled objects j.
- Since

$$p(y_{j,r} = 1|X, Y) = \int p(y_{j,r} = 1|\beta_r^\top x_j)p(\beta_r|X, Y)d(\beta_r),$$

we need to compute $p(\beta_r|X, Y)$.
- Exact inference is intractable $\Rightarrow$ use a Gibbs sampler and approximate integral by averaging over samples.

Active learning

- Query task-labeler pair (i', l') that maximizes an expected utility of the updated posterior:

$$(i', l') = \text{argmax}_{(i',l')} E_{y_{i',l'}}(U(p(\beta_r|y_{i',l'}, X, Y)))$$

- For us, maximize posterior mean shift:

$$U(p(\beta_r|y_{i',l'}, X, Y)) = \|E_{p(\beta_r|X,Y)}(\beta_r) - E_{p(\beta_r|y_{i',l'},X,Y)}(\beta_r)\|_2$$

- **Problem:** Starting with k task-labeler pairs that could be queried, active learning requires $k \lesssim g \lesssim k^2$ subordinate inferences $p(\beta_r|y_{i',l'}, X, Y)$.
  $\Rightarrow$ Cannot use standard Gibbs sampling for this!
- Suppose we have a Markov chain $p(\beta_r^t|\beta_r^{t-1})$ and a perturbed version $\hat{p}(\hat{\beta}_r^t|\hat{\beta}_r^{t-1})$.
- Let their stationary distributions be $p_\infty(\beta_r)$ and $\hat{p}_\infty(\hat{\beta}_r)$.
- Using samples $\beta_r^s \sim p_\infty(\beta_r)$ $s = 1, \ldots, S$, we approximate

$$\hat{p}_\infty(\hat{\beta}_r) \approx \int \hat{p}(\hat{\beta}_r|\beta_r)p_\infty(\beta_r)d\beta_r \approx \frac{1}{S}\sum_{s=1}^{S} \hat{p}(\hat{\beta}_r|\beta_r^s).$$

- If $p_\infty(\beta_r) = \hat{p}_\infty(\beta_r)$, the first approximation is exact. For small perturbations, we hope the approximation good.
- Specialize to active learning:
  - Unperturbed chain = Gibbs sampler for $p(\beta_r|X, Y)$.
  - Perturbed chain = Gibbs sampler for the updated posterior $p(\beta_r|y_{i',l'}, X, Y)$ with candidate measurement $y_{i',l'}$.
- **Thus, we approximate samples from** $p(\beta_r|y_{i',l'}, X, Y)$ **with transformed samples from** $p(\beta_r|X, Y)$**.**
- The active learning scores $E_{y_{i',l'}}(U(p(\beta_r|y_{i',l'}, X, Y)))$ are approximated using this strategy for each (i', l') pair.

Results

- **Task:** Is the triangle to the left or above the rectangle

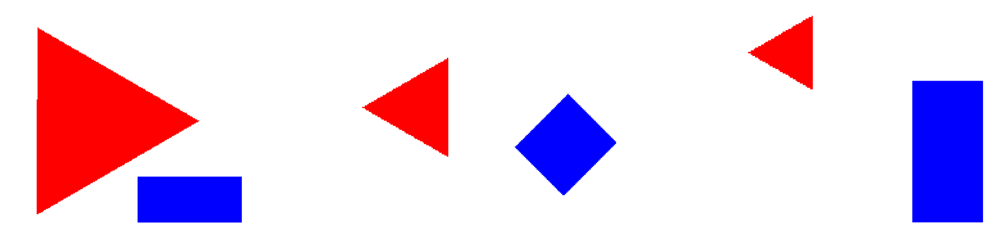

- Ambiguous tasks with many possible interpretations.
- Crowdsourced data: 523 localization tasks, labeled by three distinct labelers each, with disagreement.
- Gold standard compares the centroid positions.
- All methods seeded with 60 gold labels and crowdsourced data. Active learning methods can query gold standard.
- Gibbs sampling with 2000 iterations burnin, and active learning every 200 iterations thereafter. Predictive probabilities computed by averaging over last 200 samples.
- Evaluate by averaged log likelihood and error rate on gold standard labels for a held-out test set of 1101 tasks.

	Algorithm	Final Loglik	Final Error
No Active Learning	GOLD	-3716 ± 1695	0.0547 ± 0.0102
	CONS	-421.1 ± 2.6	0.0935 ± 0.0031
	BBMC	$\mathbf{-219.1 \pm 3.1}$	$\mathbf{0.0309 \pm 0.0033}$
Active Learning	GOLD-ACT	-1957 ± 696	0.0290 ± 0.0037
	CONS-ACT	-396.1 ± 3.6	0.0906 ± 0.0024
	RAND-ACT	-186.0 ± 2.2	0.0292 ± 0.0029
	DIS-ACT	-198.3 ± 5.8	0.0392 ± 0.0052
	MCMC-ACT	-196.1 ± 6.7	0.0492 ± 0.0050
	BBMC-ACT	$\mathbf{-160.8 \pm 3.9}$	$\mathbf{0.0188 \pm 0.0018}$

- BBMC and BBMC-ACT is our model with and without active learning.
- RAND-ACT, DIS-ACT and MCMC-ACT use the BBMC model, but do active learning by random choice, disagreement or naïve MCMC scoring with 10-step subchains.
- GOLD, CONS, GOLD-ACT and CONS-ACT learn logistic regressions on gold or consensus labels, with or without random choice active learning.

References

[1] O. Dekel and O. Shamir. Vox Populi: Collecting High-Quality Labels from a Crowd. In *Proceedings of the 22nd Annual Conference on Learning Theory (COLT)*, Montreal, Quebec, Canada, 2009.

[2] T. L. Griffiths and Z. Ghahramani. Infinite Latent Feature Models and the Indian Buffet Process. Technical report, Gatsby Computational Neuroscience Unit, 2005.

[3] V. C. Raykar, S. Yu, L. H. Zhao, G. H. Valadez, C. Florin, L. Bogoni, and L. Moy. Learning from Crowds. *Journal of Machine Learning Research*, 11:1297–1322, April 2010.

[4] Y. Yan, R. Rosales, G. Fung, and J. G. Dy. Active Learning from Crowds. In L. Getoor and T. Scheffer, editors, *Proceedings of the 28th International Conference on Machine Learning (ICML)*, Bellevue, Washington, 2011.