Sequential crowdsourced labeling as an epsilon-greedy exploration in a Markov Decision Process
Supplementary material

Vikas C. Raykar
IBM Research, Bangalore, India

Priyanka Agrawal
IBM Research, Bangalore, India

A Variational Bayes updates for the single coin model

The joint likelihood of the observed labels $\boldsymbol{y}_i$ and the true label z_i for the i^{th} instance given the annotator parameters $\boldsymbol{\theta}$ can be factored as

$$p(\boldsymbol{y}_i, z_i|\boldsymbol{\theta}) \quad = \quad p(z_i)p(\boldsymbol{y}_i|z_i,\boldsymbol{\theta}) = p(z_i)\prod_{j\in\mathcal{M}_i} p(y_i^j|z_i,\theta^j) = p(z_i)\prod_{j\in\mathcal{M}_i}(\theta^j)^{\delta(y_i^j,z_i)}(1-\theta^j)^{1-\delta(y_i^j,z_i)},$$

where $\delta(y_i^j, z_i) = 1$ if $y_i^j = z_i$ and 0 otherwise. We have made an assumption that the labels provided by the different annotators for a given instance are independent conditional on the true label, which is a typical assumption made in most crowdsourcing algorithms. Hence

$$\ln p(\boldsymbol{y}_i, z_i|\boldsymbol{\theta}) = \ln p(z_i) + \sum_{j\in\mathcal{M}_i}\left[\delta(y_i^j,z_i)\ln\theta^j + (1-\delta(y_i^j,z_i))\ln(1-\theta^j)\right]. \tag{1}$$

VBE-step: Assuming the n instances are independent the updates for $\boldsymbol{q}_{\boldsymbol{z}}^{(t+1)}(\boldsymbol{z})$ can be broken down across the n instances as $\boldsymbol{q}_{\boldsymbol{z}}^{(t+1)}(\boldsymbol{z}) = \prod_{i=1}^n \boldsymbol{q}_{z_i}^{(t+1)}(z_i)$ where

$$\boldsymbol{q}_{z_i}^{(t+1)}(z_i) \propto \exp\left[\mathbb{E}_{q_{\boldsymbol{\theta}}^{(t)}(\boldsymbol{\theta})}\left[\ln p(\boldsymbol{y}_i, z_i|\boldsymbol{\theta})\right]\right].$$

Taking the expectation of (1) with respect to $q_{\boldsymbol{\theta}}^{(t)}(\boldsymbol{\theta})$ we have

$$\mathbb{E}_{q_{\boldsymbol{\theta}}^{(t)}(\boldsymbol{\theta})}\left[\ln p(\boldsymbol{y}_i, z_i|\boldsymbol{\theta})\right] = \ln p(z_i) + \sum_{j\in\mathcal{M}_i}\left[\delta(y_i^j,z_i)\ln A^j + (1-\delta(y_i^j,z_i))\ln B^j\right],$$

where $A^j := \exp\left(\mathbb{E}_{q_{\theta^j}^{(t)}(\theta^j)}\left[\ln\theta^j\right]\right)$ and $B^j := \exp\left(\mathbb{E}_{q_{\theta^j}^{(t)}(\theta^j)}\left[\ln(1-\theta^j)\right]\right)$. If $x \sim \text{Beta}(a,b)$ then $\mathbb{E}[\ln x] = \text{Digamma}(a) - \text{Digamma}(a+b)$ and $\mathbb{E}[\ln(1-x)] = \text{Digamma}(b) - \text{Digamma}(a+b)$. Hence we have the following updates for the hidden variable z_i

$$q_{z_i}^{(t+1)}(z_i) \propto p(z_i)\prod_{j\in\mathcal{M}_i}(A^j)^{\delta(y_i^j,z_i)}(B^j)^{1-\delta(y_i^j,z_i)}, \quad i = 1,\ldots,n. \tag{2}$$

VBM-step: Similarly the updates for $q_{\boldsymbol{\theta}}^{(t+1)}(\boldsymbol{\theta})$ can be broken down across the m annotators as $q_{\boldsymbol{\theta}}^{(t+1)}(\boldsymbol{\theta}) = \prod_{j=1}^m q_{\theta_j}^{(t+1)}(\theta_j)$ where

$$q_{\theta^j}^{(t+1)}(\theta^j) \propto p(\theta^j)\cdot\exp\left[\mathbb{E}_{q_{\boldsymbol{z}}^{(t+1)}(\boldsymbol{z})}\left[\ln p(\boldsymbol{y},\boldsymbol{z}|\theta^j)\right]\right].$$

Assuming the instances are independent and taking the expectation of (1) with respect to $q_{\boldsymbol{z}}^{(t+1)}(\boldsymbol{z})$

$$\mathbb{E}_{q_{\boldsymbol{z}}^{(t+1)}(\boldsymbol{z})}\left[\ln p(\boldsymbol{y}, z|\theta^j)\right] = \sum_{i\in\mathcal{N}^j}\ln p(z_i) + q_{z_i}^{(t+1)}(y_i^j)\ln\theta^j + (1-q_{z_i}^{(t+1)}(y_i^j))\ln(1-\theta^j),$$

since $\mathbb{E}_{q_{z_i}^{(t+1)}(z_i)}\left[\delta(y_i^j,z_i)\right] = q_{z_i}^{(t+1)}(y_i^j)$. Hence we have the following updates for the annotator accuracy θ^j

$$q_{\theta^j}^{(t+1)}(\theta^j) \propto p(\theta^j)\prod_{i\in\mathcal{N}_j}(\theta^j)^{q_{z_i}^{(t+1)}(y_i^j)}(1-\theta^j)^{1-q_{z_i}^{(t+1)}(y_i^j)}, \quad j = 1,\ldots,m. \tag{3}$$

As a consequence of using a beta prior for θ^j the posterior is again a beta distribution

$$q_{\theta^j}^{(t+1)}(\theta^j) = \text{Beta}\left(\theta^j|a^j + \sum_{i\in\mathcal{N}_j} q_{z_i}^{(t+1)}(y_i^j), b^j + \sum_{i\in\mathcal{N}_j} 1 - q_{z_i}^{(t+1)}(y_i^j)\right).$$