APPENDIX: TUNED MODELS OF PEER ASSESSMENT IN MOOCS

CHRIS PIECH, JONATHAN HUANG, ZHENGHAO CHEN,
CHUONG DO, ANDREW NG, DAPHNE KOLLER

In this document, we describe the inference/learning procedures used in our paper.

1. Gibbs sampling for Model $\mathbf{PG}_1$

Model $\mathbf{PG}_1$ is given as follows:

$$\begin{aligned}
\text{(Reliability) } \tau_v &\sim \mathcal{G}(\alpha_0,\, \beta_0) \text{ for every grader } v,\\
\text{(Bias) } b_v &\sim \mathcal{N}(0,\, 1/\eta_0) \text{ for every grader } v,\\
\text{(True score) } s_u &\sim \mathcal{N}(\mu_0,\, 1/\gamma_0) \text{ for every user } u, \text{ and}\\
\text{(Observed score) } z_u^v &\sim \mathcal{N}(s_u + b_v,\, 1/\tau_v),\\
&\quad \text{for every observed peer grade.}
\end{aligned}$$

The joint posterior distribution is:

$$\begin{aligned}
&P(Z|\{s_u\}_{u\in U},\{b_v\}_{v\in G},\{\tau_v\}_{v\in G})\\
&\quad = \prod_u P(s_u|\mu_0,\gamma_0)\cdot \prod_v P(b_v|\eta_0)\cdot P(\tau_v|\alpha_0,\beta_0)\prod_{z_u^v} P(z_u^v|s_u,b_v,\tau_v).
\end{aligned}$$

The pseudocode for Gibbs sampling from Model $\mathbf{PG}_1$ is:

- Generate an initial assignment to all non-observed variables, s_u, τ_v, b_v for all true grades, grader reliabilities and grader biases.
- For $t = 1, \ldots, T$:
  - For each user score s_{u_i}:
    - Sample $s \sim \mathcal{N}\left(s\ ;\ \frac{\gamma_0}{\gamma_0+\sum_{v:v\to u_i}\tau_v}\mu_0 + \frac{\sum_{v:v\to u_i}\tau_v(z_{u_i}^v+b_v)}{\gamma_0+\sum_{v:v\to u_i}\tau_v}\ ,\ \gamma_0+\sum_{v:v\to u_i}\tau_v\right)$
    - $s_{u_i} \leftarrow s$
  - For each grader reliability τ_{v_i}:
    - Sample $\tau \sim \mathcal{G}\left(\tau\ ;\ \alpha_0 + \frac{n_{v_i}}{2},\ \beta_0 + \frac{1}{2}\sum_{u:u\to v_i}(z_u^{v_i} - (s_u+b_{v_i}))^2\right)$
    - $\tau_{v_i} \leftarrow \tau$
  - For each grader bias b_{v_i}:
    - Sample $b \sim \mathcal{N}\left(b\ ;\ \frac{\sum_{u:u\to v_i}\tau_{v_i}(z_u^{v_i}-s_u)}{\eta+n_{v_i}\tau_{v_i}},\ \ \eta + n_{v_i}\tau_{v_i}\right)$
    - $b_{v_i} \leftarrow b$
  - Save sample $\zeta^{(t)} \leftarrow (\{s_u\}_{u\in U}, \{\tau_v\}_{v\in U}, \{b_v\}_{v\in U})$
- Return samples from $\zeta^{(B)}$, $\zeta^{(B+1)}$, …, $\zeta^{(T)}$ for some large enough number B.

Derivation of updates. We examine the problems of sampling s_u and τ_v separately. Consider now a fixed user u_i. We derive the sampling step for s_u as follows:

$$
\begin{aligned}
s &\sim P(s_{u_i}|MB(s_{u_i})),\\
&\propto P(s_{u_i}|\mu_0,\gamma_0)\cdot \prod_{v:v\to u_i} P(z^v_{u_i}|s_u,b_v,\tau_v),\\
&\propto \exp\left(-\frac{1}{2}\gamma_0(s_{u_i-\mu_0})^2+\sum_{v:v\to u_i}\left(-\frac{1}{2}\tau_v\left(z^v_{u_i}-(s_{u_i}+b_v)\right)^2\right)\right),\\
&\propto \exp\left(-\frac{1}{2}\left[\gamma_0(s_{u_i}-\mu_0)^2+\sum_{v:v\to u_i}\tau_v\left(z^v_{u_i}-(s_{u_i}+b_v)\right)^2\right]\right).
\end{aligned}
$$

The expression inside the exponent is quadratic — we thus complete the square, obtaining:

$$
\begin{aligned}
&\gamma_0(s_{u_i}-\mu_0)^2+\sum_{v:v\to u_i}\tau_v\left(z^v_{u_i}-(s_{u_i}+b_v)\right)^2\\
&\qquad= \text{const.}+\gamma_0(s^2_{u_i}-2\mu_0 s_{u_i})+\sum_{v:v\to u_i}\tau_v\left((s_{u_i}+b_v)^2-2z^v_{u_i}(s_{u_i}+b_v)\right),\\
&\qquad= \text{const.}+\left(\gamma_0+\sum_{v:v\to u_i}\tau_v\right)s^2_{u_i}-2\left(\gamma_0\mu_0+\sum_{v:v\to u_i}\tau_v(z^v_{u_i}-b_v)\right)s_{u_i},\\
&\qquad= \text{const.}+R\left(s_{u_i}-\frac{1}{R}\left(\gamma_0\mu_0+\sum_{v:v\to u_i}\tau_v(z^v_{u_i}-b_v)\right)\right)^2,\\
&\qquad\qquad(\text{where } R=\gamma_0+\textstyle\sum_{v:v\to u_i}\tau_v).
\end{aligned}
$$

Therefore the sampling distribution is Gaussian:

$$
s\sim\mathcal{N}\left(s\ ;\ \frac{\gamma_0}{\gamma_0\sum_{v:v\to u_i}\tau_v}\mu_0+\frac{\sum_{v:v\to u_i}\tau_v(z^v_{u_i}-b_v)}{\gamma_0+\sum_{v:v\to u_i}\tau_v},\gamma_0+\sum_{v:v\to u_i}\tau_v\right)
$$

Now consider a fixed user v_i. We derive the sampling step for grader reliability τ_v as follows:

$$
\begin{aligned}
\tau &\sim P(\tau_{v_i} \,|\, \mathrm{MB}(\tau_{v_i})), \\
&\propto P(\tau_{v_i}|\alpha_0, \beta_0) \cdot \prod_{u:u\to v_i} P(z_u^{v_i}|s_u, \tau_{v_i}, b_{v_i}), \\
&\propto \tau_{v_i}^{\alpha_0 - 1} \exp\left(-\beta_0 \tau_{v_i} + \sum_{u:u\to v_i} \frac{1}{2} \left(\log \tau_{v_i} - \log 2\pi - \tau_{v_i} \left(z_u^{v_i} - (s_u + b_{v_i}) \right)^2 \right) \right), \\
&\propto \tau_{v_i}^{\alpha_0 + \frac{n_{v_i}}{2} - 1} \exp\left(- \left[\beta_0 + \frac{1}{2} \sum_{u:u\to v_i} \left(z_u^{v_i} - (s_u + b_{v_i}) \right)^2 \right] \tau_{v_i} \right).
\end{aligned}
$$

From this, we can recognize the sampling distribution to be Gamma with:

$$
\tau \sim \mathcal{G}\left(\tau \,;\, \alpha_0 + \frac{n_{v_i}}{2} \,,\, \beta_0 + \frac{1}{2} \sum_{u:u\to v_i} \left(z_u^{v_i} - (s_u + b_{v_i}) \right)^2 \right).
$$

Finally we derive the sampling set for grader bias b_v as follows:

$$
\begin{aligned}
b &\sim P(b_{v_i}|\mathrm{MB}(b_{v_i})), \\
&\propto P(b_{v_i}|\eta_0) \cdot \prod_{u:u\to v_i} P(z_u^{v_i}|s_u, \tau_{v_i}, b_{v_i}), \\
&\propto \exp\left(-\frac{1}{2} \eta_0 b_{v_i}^2 - \frac{1}{2} \sum_{u:u\to v_i} \tau_{v_i} \left(z_u^{v_i} - (s_u + b_{v_i}) \right)^2 \right), \\
&\propto \exp\left(-\frac{1}{2} \left[\eta_0 b_{v_i}^2 + \sum_{u:u\to v_i} \tau_{v_i} \left((s_u + b_{v_i})^2 - 2 z_u^{v_i} (s_u + b_{v_i}) \right) \right] \right).
\end{aligned}
$$

The expression inside square brackets is quadratic, again allowing us to complete-the-square as follows:

$$
\begin{aligned}
\eta b_{v_i}^2 + &\sum_{u:u\to v_i} \tau_{v_i} \left((s_u + b_{v_i})^2 - 2 z_u^{v_i} (s_u + b_{v_i}) \right) \\
&= \text{const.} + (\eta_0 + \sum_{u:u\to v_i} \tau_{v_i}) b_{v_i}^2 - 2 \left(\sum_{u:u\to v_i} \tau_{v_i} (z_u^{v_i} - s_u) \right) b_{v_i}, \\
&= \text{const.} + R \left(b_{v_i} - \frac{1}{R} \left(\sum_{u:u\to v_i} \tau_{v_i} (z_u^{v_i} - s_u) \right) \right)^2,
\end{aligned}
$$

where $R = \eta_0 + \sum_{u:u\to v_i} \tau_{v_i} = \eta_0 + n_{v_i} \tau_{v_i}$. The sampling distribution for b is thus Gaussian with:

$$
b \sim \mathcal{N}\left(b \,;\, \frac{\sum_{u:u\to v_i} \tau_{v_i} (z_u^{v_i} - s_u)}{\eta_0 + n_{v_i} \tau_{v_i}},\ \eta + n_{v_i} \tau_{v_i} \right).
$$

2. Handling multiple assignments (Model **PG**$_2$)

Model **PG**$_2$ looks almost identical to **PG**$_1$ with the exception of the fact that a grader's bias depends on her bias at the last homework assignment.

$$\tau_v^{(T)} \sim \mathcal{G}(\alpha_0, \beta_0) \text{ for every grader } v,$$
$$b_v^{(T)} \sim \mathcal{N}(b_v^{(T-1)}, 1/\omega_0) \text{ for every grader } v,$$
$$s_u^{(T)} \sim \mathcal{N}(\mu_0, 1/\gamma_0) \text{ for every user } u, \text{ and}$$
$$z_u^{v,(T)} \sim \mathcal{N}(s_u^{(T)} + b_v^{(T)}, 1/\tau_v^{(T)}),$$
$$\text{for every observed peer grade.}$$

Since we handle assignments in an online fashion, we do not consider the possibility of using grades from Assignment T to retroactively go back and modify earlier grades. Due to the Markov nature of the model for bias in Model **PG**$_2$, inference at each timeslice (i.e. each homework assignment) is the same as that of Model **PG**$_1$ with the exception that instead of using the same bias for all graders, each grader now has his own prior over bias.

3. Gibbs sampling for Model **PG**$_3$

Model **PG**$_3$ is given as follows:

$$b_v \sim \mathcal{N}(0,\, 1/\eta_0) \text{ for every grader } v,$$
$$s_u \sim \mathcal{N}(\mu_0,\, 1/\gamma_0) \text{ for every user } u, \text{ and}$$
$$z_u^v \sim \mathcal{N}\left(s_u + b_v,\, \frac{1}{f_\theta(s_v)}\right),$$
$$\text{for every observed peer grade,}$$

where $f_\theta(s) \equiv \theta_1 \cdot s + \theta_0$. **PG**$_3$ is the only model that we cannot Gibbs sample in closed form. The joint probability distribution is written as:

$$P(Z|\{s_u\}_{u\in U}, \{b_v\}_{v\in G}, \{\tau_v\}_{v\in G}) = \prod_u P(s_u|\mu_0, \gamma_0) \cdot \prod_v P(b_v|\eta_0) \prod_{z_u^v} P(z_u^v|s_u, s_v, b_v).$$

Derivation of updates. Again we look at the cases of sampling s_u and b_v separately. Consider now a fixed user u_i. We derive the sampling step for s_u as follows:

$$s \sim P(s_{u_i}|MB(s_{u_i})),$$
$$\propto P(s_{u_i}|\mu_0, \gamma_0) \cdot \prod_{v:v\to u} P(z_u^v|s_u, s_v, b_v) \cdot \prod_{w:u\to w} P(z_u^v|s_w, s_u, b_v),$$
$$\propto \exp\left(-\frac{1}{2}\gamma_0 (s_{u-\mu_0})^2\right)$$

$$
\begin{aligned}
&\cdot \prod_{v:v\to u} \exp\left(-\frac{1}{2} f_\theta(s_v)(z_u^v - (s_u + b_v))^2\right) \\
&\cdot \prod_{w:u\to w} \sqrt{f_\theta(s_u)} \exp\left(-\frac{1}{2} f_\theta(s_v)[z_w^u - (s_w + b_u)]^2\right), \\
\propto\ & \sqrt{f_\theta(s_u)}^{k_u} \cdot \exp\Bigg(-\frac{1}{2}\Big[\gamma_0(s_u - \mu_0)^2 \\
&+ \sum_{v:v\to u} f_\theta(s_v)(z_u^v - (s_u + b_v))^2 \\
&+ \sum_{w:u\to w} f_\theta(s_u)(z_w^u - (s_w + b_u))^2\Big]\Bigg), \\
&\qquad \text{(where } k_u \text{ is the number of people graded by } u\text{)} \\
\propto\ & f_\theta(s_u)^{k_u/2} \cdot \exp\left(-\frac{1}{2}\left[R\left(s_u - \frac{y}{R}\right)^2\right]\right),
\end{aligned}
$$

where:

$$
\begin{aligned}
R &= \gamma_0 + \sum_{v:v\to u} f_\theta(s_u), \text{ and} \\
y &= \mu_0\gamma_0 + \sum_{v:v\to u} f_\theta(s_v)(z_u^v - b_v) + \sum_{w:u\to w} \theta_1(z_w^v - (s_w + b_v))^2.
\end{aligned}
$$

Note that unlike its analog from Model $\mathbf{PG}_1$, the sampling step for s_u in Model $\mathbf{PG}_3$ cannot be performed in closed form. In our experiments, we sample from a discretized approximation of the posterior distribution instead. We expect that a Laplace approximation would also be effective (and fast) for this problem as the posterior distributions typically "look" nearly Gaussian in practice.

We now turn to sampling the bias variables b_v. Note that there are no reliability variables to sample in Model $\mathbf{PG}_3$.

$$
\begin{aligned}
b &\sim P(b_v | MB(b_v)), \\
&\propto P(b_v|\eta_0) \cdot \prod_{u:v\to u} P(z_u^v | s_u, s_v, b_v), \\
&\propto \exp\left(-\frac{1}{2}\left[\eta_0 b_v^2 + \sum_{u:v\to u} f_\theta(s_v)(z_u^v - (s_u + b_v))^2\right]\right), \\
&\propto \exp\left(-\frac{1}{2}\left[R(b_v - \frac{y}{R})^2\right]\right),
\end{aligned}
$$

where:

$$R = \eta_0 + \sum_{u:v \to u} f_\theta(s_v), \text{ and}$$

$$y = \sum_{u:v \to u} f_\theta(s_u)(z_u^v - s_u).$$