0.1 Partial Specification for Linear Regression

So far we have made no assumptions about the relation between Y and X, and consequently have only been able to handle small shifts in the distribution $p(x)$ (i.e. $D_{\chi^2}(\tilde{p}(x) \| p^*(x))$ must be small). We will next make very strong assumptions about how Y relates to X, and handle much larger shifts. Then we will try to relax those assumptions using an idea called *partial specification*. We will move away from the causal inference setting above, and instead consider linear regression: we wish to predict $Y \in \mathbb{R}$ from $X \in \mathbb{R}^d$ using some linear predictor $\langle \beta, X \rangle$. As before we make the covariate shift assumption that $\tilde{p}(y \mid x) = p^*(y \mid x)$, but $\tilde{p}(x)$ and $p^*(x)$ could differ. Our goal, rather than to construct a good predictor on p^*, is to estimate the error of the ordinary least squares estimator.

Starting point: linear response with Gaussian errors. In the simplest setting, suppose that we completely believe our model:

$$Y = \langle \beta, X \rangle + Z, \text{ where } Z \sim \mathcal{N}(0, \sigma^2 I). \tag{1}$$

We observe samples $(x_1, y_1), \dots, (x_n, y_n) \sim \tilde{p}$, and samples $\bar{x}_1, \dots, \bar{x}_m \sim p^*(x)$. Suppose that we estimate β using the ordinary least squares estimator:

$$\hat{\beta} = \arg\min_{\beta} \frac{1}{n} \sum_{i=1}^{n} (y_i - \langle \beta, x_i \rangle)^2 = (\sum_{i=1}^{n} x_i x_i^\top)^{-1} \sum_{i=1}^{n} x_i y_i. \tag{2}$$

Define $\tilde{\Omega} = \frac{1}{n} \sum_{i=1}^{n} x_i x_i^\top$. Then since $y_i = x_i^\top \beta + z_i$, we can further write

$$\begin{aligned}
\hat{\beta} &= (\sum_{i=1}^{n} x_i x_i^\top)^{-1} (\sum_{i=1}^{n} x_i x_i^\top \beta + x_i z_i) && (3)\\
&= (n\tilde{\Omega})^{-1} (n\tilde{\Omega}\beta + \sum_{i=1}^{n} x_i z_i) && (4)\\
&= \beta + \frac{1}{n} \tilde{\Omega}^{-1} \sum_{i=1}^{n} x_i z_i. && (5)
\end{aligned}$$

From this we see that, conditional on the x_i, $\hat{\beta} - \beta$ is a zero-mean Gaussian distribution. Its covariance matrix is given by

$$\frac{1}{n^2} \tilde{\Omega}^{-1} \sum_{i=1}^{n} \mathbb{E}[z_i^2 \mid x_i] x_i x_i^\top \tilde{\Omega}^{-1} = \frac{\sigma^2}{n} \tilde{\Omega}^{-1}. \tag{6}$$

Now suppose that we wish to estimate the error on the samples $\bar{x}_{1:m}$. The expected error on sample $\bar{x}_i$ is $\sigma^2 + \langle \hat{\beta} - \beta, \bar{x}_i \rangle^2$. If we let $\Omega^* = \frac{1}{m} \sum_{i=1}^{m} \bar{x}_i \bar{x}_i^\top$, then the overall average expected error (conditional on $x_{1:n}, \bar{x}_{1:m}$) is

$$\begin{aligned}
\sigma^2 + \mathbb{E}_Z[\frac{1}{m} \sum_{i=1}^{m} (\bar{x}_i^\top (\beta - \hat{\beta}))^2] &= \sigma^2 + \langle \frac{1}{m} \sum_{i=1}^{m} \bar{x}_i \bar{x}_i^\top, \mathbb{E}_Z[(\beta - \hat{\beta})(\beta - \hat{\beta})^\top] \rangle && (7)\\
&= \sigma^2 + \langle \Omega^*, \frac{\sigma^2}{n} \tilde{\Omega}^{-1} \rangle && (8)\\
&+ \sigma^2 \Big(1 + \frac{1}{n} \langle \Omega^*, \tilde{\Omega}^{-1} \rangle \Big). && (9)
\end{aligned}$$

This shows that the error depends on the divergence between the second moment matrices of $\tilde{p}(x)$ and $p^*(x)$:

- When $\tilde{p}(x) = p^*(x)$, then $\langle \Omega^*, \tilde{\Omega}^{-1} \rangle = \mathrm{tr}(\Omega^* \tilde{\Omega}^{-1}) \approx \mathrm{tr}(I) = d$, so the error decays as $\frac{d}{n}$.

- If $\tilde{\Omega}$ is low-rank and is missing any directions that appear in Ω, then the error is infinite. This makes sense, as we have no way of estimating β along the missing directions, and we need to be able to estimate β in those directions to get good error under p^*. We can get non-infinite bounds if we further assume some norm bound on β^*; e.g. if $\|\beta^*\|_2$ is bounded then the missing directions only contribute some finite error.
- On the other hand, if $\tilde{\Omega}$ is full-rank but Ω^* is low-rank, then we still achieve finite error. For instance, suppose that $\tilde{\Omega} = I$ is the identity, and $\Omega^* = \frac{d}{k}P$ is a projection matrix onto a k-dimensional subspace, scaled to have trace d. Then we get a sample complexity of $\frac{d}{n}$, although if we had observed samples with second moment matrix Ω^* at training time, we would have gotten a better sample complexity of $\frac{k}{n}$.
- In general it is always better for $\tilde{\Omega}$ to be bigger. This is partially an artefact of the noise σ^2 being the same for all X, so we would always rather have X be as far out as possible since it pins down β more effectively. If the noise was proportional to $\|X\|$ (for instance) then the answer would be different.

Of course, this all so far rests on the assumption of Gaussian error. Can we do better?

Calculation from moment assumptions. It turns out that our calculation above relied only on conditional moments of the errors, rather than Gaussianity. We will show this explicitly by doing the calculations more carefully. Re-using steps above, we have that

$$\hat{\beta} - \beta = \frac{1}{n}\tilde{\Omega}^{-1}\sum_{i=1}^{n} x_i z_i. \tag{10}$$

In particular, assuming that the (x_i, y_i) are i.i.d., we have

$$\mathbb{E}[\hat{\beta} - \beta \mid x_1, \ldots, x_n] = \frac{1}{n}\tilde{\Omega}^{-1}\sum_{i=1}^{n} x_i \mathbb{E}[z_i \mid x_i] = \tilde{\Omega}^{-1}\tilde{b}, \tag{11}$$

where $\tilde{b} \stackrel{\text{def}}{=} \frac{1}{n}\sum_{i=1}^{n} x_i \mathbb{E}[z_i \mid x_i]$.

In particular, as long as $\mathbb{E}[Z \mid X] = 0$ for all X, $\hat{\beta}$ is an unbiased estimator for X. In fact, since this only needs to hold on average, as long as $\mathbb{E}[ZX] = 0$ (covariates uncorrelated with noise) then $\mathbb{E}[\hat{\beta} - \beta] = 0$, and $\mathbb{E}[\hat{\beta} - \beta \mid x_{1:n}]$ converges to zero as $n \to \infty$. This yields an insight that is important more generally:

> *Orindary least squares yields an unbiased estimate of β whenever the covariates X and noise Z are uncorrelated.*

This partly explains the success of OLS compared to other alternatives (e.g. penalizing the absolute error or fourth power of the error). While OLS might initially look like the maximum likelihood estimator under Gaussian errors, it yields consistent estimates of β under much weaker assumptions. Minimizing the fourth power of the error requires stronger assumptions for consistency, while minimizing the absolute error would yield a different condition in terms of medians rather than expectations.

Next we turn to the covariance of $\hat{\beta}$. Assuming that (x_i, y_i) are independent like before, we have

$$\begin{aligned}
\mathsf{Cov}[\hat{\beta} \mid x_{1:n}] &= \mathsf{Cov}[\frac{1}{n}\tilde{\Omega}^{-1}\sum_{i=1}^{n} x_i z_i \mid x_{1:n}] & (12)\\
&= \frac{1}{n^2}\tilde{\Omega}^{-1}\sum_{i,j=1}^{n} x_i \mathsf{Cov}[z_i, z_j \mid x_i, x_j] x_j^\top \tilde{\Omega}^{-1} & (13)\\
&= \frac{1}{n^2}\tilde{\Omega}^{-1}\sum_{i=1}^{n} x_i \mathsf{Var}[z_i \mid x_i] x_i^\top \tilde{\Omega}^{-1}, & (14)
\end{aligned}$$

where the final line is because z_i, z_j are independent for $i \neq j$. If we define $\tilde{M} = \frac{1}{n}\sum_{i=1}^{n} x_i \mathsf{Var}[z_i \mid x_i] x_i^\top$, then the final term becomes $\frac{1}{n}\tilde{\Omega}^{-1}\tilde{M}\tilde{\Omega}$.

This quantity is bounded under much weaker assumptions than Gaussianity. If we, for instance, merely assume that $\mathsf{Var}[z_i \mid x_i] \leq \sigma^2$ for all i, then we have that $\tilde{M} \preceq \sigma^2\tilde{\Omega}$ and hence $\mathsf{Cov}[\hat{\beta} \mid x_{1:n}] \preceq \frac{\sigma^2}{n}\tilde{\Omega}^{-1}$.

We can put this together to estimate the squared error. Letting $\bar{z}_i$ be the noise for $\bar{x}_i$, the squared error is then $\frac{1}{m}\sum_{j=1}^m(\langle\beta - \hat{\beta}, \bar{x}_i\rangle + \bar{z}_i)^2$, and computing the expectation given $x_{1:n}, \bar{x}_{1:m}$ yields

$$\mathbb{E}[\frac{1}{m}\sum_{j=1}^m(\langle\beta - \hat{\beta}, \bar{x}_i\rangle + \bar{z}_i)^2 \mid x_{1:n}, \bar{x}_{1:m}] \tag{15}$$

$$= \frac{1}{m}\sum_{i=1}^m \bar{x}_i^\top \mathbb{E}[(\beta - \hat{\beta})(\beta - \hat{\beta})^\top \mid x_{1:n}]\bar{x}_i + 2\bar{x}_i^\top\mathbb{E}[\beta - \hat{\beta} \mid x_{1:n}]\mathbb{E}[\bar{z}_i \mid x_i] + \mathbb{E}[\bar{z}_i^2 \mid \bar{x}_i] \tag{16}$$

$$= \Big\langle \Omega^*, \tilde{\Omega}^{-1}\Big(\frac{1}{n}\tilde{M} + \tilde{b}\tilde{b}^\top\Big)\tilde{\Omega}^{-1}\Big\rangle + 2\Big\langle b^*, \tilde{\Omega}^{-1}\tilde{b}\Big\rangle + \frac{1}{m}\sum_{j=1}^m \mathbb{E}[\bar{z}_i^2 \mid \bar{x}_i]. \tag{17}$$

To interpret this expression, first assume that the true model is "actually linear", meaning that $\tilde{b} = b^* = 0$. Then the expression reduces to $\frac{1}{n}\langle\Omega^*, \tilde{\Omega}^{-1}\tilde{M}\tilde{\Omega}^{-1}\rangle + \frac{1}{m}\sum_{j=1}^m \mathbb{E}[\bar{z}_i^2 \mid x_i]$. The second term is the intrinsic variance in the data, while the first term is similar to the $\frac{1}{n}\langle\Omega^*, \tilde{\Omega}^{-1}\rangle$ term from before, but accounts for correlation between X and the variation in Z. The $\tilde{M}$ term is also reminiscent of our earlier condition for robust linear regression.

If the model is not actually linear, then we need to decide how to define β (since the optimal β is then no longer independent of the distribution). In that case a natural choice is to let β be the minimizer under the training distribution, in which case $\tilde{b} \to 0$ as $n \to \infty$ and thus the $\langle b^*, \tilde{\Omega}^{-1}\tilde{b}\rangle$ term conveniently becomes asymptotically negligible. The twist is that $\mathbb{E}[\bar{z}_i^2 \mid \bar{x}_i]$ now measures not just the intrinsic variance but also the departure from linearity, and could be quite large if the linear extrapolation away from the training points ends up being poor.

Partial specification. In general, we see that we can actually form good estimates of the mean-squared error on p^* making only certain moment assumptions (e.g. $\tilde{b} = b^* = 0$) rather than needing to assume the Gaussian model is correct. This idea is called *partial specification*, where rather than making assumptions that are stringent enough to specify a parametric family, we make weaker assumptions that are typically insufficient to even yield a likelihood, but show that our estimates are still valid under those weaker assumptions. The weaker the assumptions, the more happy we are. Of course $\tilde{b} = b^* = 0$ is still fairly strong, but much better than Gaussianity. The goal of partial specification aligns with our earlier desire to design estimators for the entire family of resilient distributions, rather than specific parametric classes. We will study other variants of partial specification in the next section.