Bayesian inference and prediction in finite regression models

Carl Edward Rasmussen

October 13th, 2016

Key concepts

Bayesian inference in finite, parametric models

- we contrast maximum likelihood with Bayesian inference
- when both prior and likelihood are Gaussian, all calculations are tractable
  - the posterior on the parameters is Gaussian
  - the predictive distribution is Gaussian
  - the marginal likelihood is tractable
- we observe the contrast
  - in maximum likelihood the data fit gets better with larger models (overfitting)
  - the marginal likelihood prefers an intermediate model size (Occam's Razor)

Maximum likelihood, parametric model

Supervised parametric learning:

- data: $\mathbf{x}, \mathbf{y}$
- model $\mathcal{M}$: $y = f_{\mathbf{w}}(x) + \varepsilon$

Gaussian likelihood:

$$p(\mathbf{y}|\mathbf{x}, \mathbf{w}, \mathcal{M}) \;\propto\; \prod_{n=1}^{N} \exp(-\tfrac{1}{2}(y_n - f_{\mathbf{w}}(x_n))^2/\sigma^2_{\text{noise}}).$$

Maximize the likelihood:

$$\mathbf{w}_{\text{ML}} \;=\; \operatorname*{argmax}_{\mathbf{w}} p(\mathbf{y}|\mathbf{x}, \mathbf{w}, \mathcal{M}).$$

Make predictions, by plugging in the ML estimate:

$$p(y_*|x_*, \mathbf{w}_{\text{ML}}, \mathcal{M})$$

Bayesian inference, parametric model

Posterior parameter distribution by Bayes rule ($p(a|b) = p(a)p(b|a)/p(b)$):

$$p(\mathbf{w}|\mathbf{x}, \mathbf{y}, \mathcal{M}) \;=\; \frac{p(\mathbf{w}|\mathcal{M})p(\mathbf{y}|\mathbf{x}, \mathbf{w}, \mathcal{M})}{p(\mathbf{y}|\mathbf{x}, \mathcal{M})}$$

Making predictions (marginalizing out the parameters):

$$\begin{aligned} p(y_*|x_*, \mathbf{x}, \mathbf{y}, \mathcal{M}) \;&=\; \int p(y_*, \mathbf{w}|\mathbf{x}, \mathbf{y}, x_*, \mathcal{M})d\mathbf{w} \\ &=\; \int p(y_*|\mathbf{w}, x_*, \mathcal{M})p(\mathbf{w}|\mathbf{x}, \mathbf{y}, \mathcal{M})d\mathbf{w}. \end{aligned}$$

Marginal likelihood:

$$p(\mathbf{y}|\mathbf{x}, \mathcal{M}) \;=\; \int p(\mathbf{w}|\mathbf{x}, \mathcal{M})p(\mathbf{y}|\mathbf{x}, \mathbf{w}, \mathcal{M})d\mathbf{w}$$

Posterior and predictive distribution in detail

For a linear-in-the-parameters model with Gaussian priors and Gaussian noise:

- Gaussian *prior* on the weights: $p(\mathbf{w}|\mathcal{M}) = \mathcal{N}(\mathbf{w};\ \mathbf{0},\ \sigma_{\mathbf{w}}^2\,\mathbf{I})$
- Gaussian *likelihood* of the weights: $p(\mathbf{y}|\mathbf{x},\mathbf{w},\mathcal{M}) = \mathcal{N}(\mathbf{y};\ \boldsymbol{\Phi}\,\mathbf{w},\ \sigma_{\text{noise}}^2\,\mathbf{I})$

Posterior parameter distribution by Bayes rule $p(a|b) = p(a)p(b|a)/p(b)$:

$$p(\mathbf{w}|\mathbf{x},\mathbf{y},\mathcal{M}) \;=\; \frac{p(\mathbf{w}|\mathcal{M})p(\mathbf{y}|\mathbf{x},\mathbf{w},\mathcal{M})}{p(\mathbf{y}|\mathbf{x},\mathcal{M})} \;=\; \mathcal{N}(\mathbf{w};\ \boldsymbol{\mu},\ \boldsymbol{\Sigma})$$

$$\boldsymbol{\Sigma} \;=\; \big(\sigma_{\text{noise}}^{-2}\boldsymbol{\Phi}^\top\boldsymbol{\Phi} + \sigma_{\mathbf{w}}^{-2}\,\mathbf{I}\big)^{-1} \quad \text{and} \quad \boldsymbol{\mu} \;=\; \Big(\boldsymbol{\Phi}^\top\boldsymbol{\Phi} + \frac{\sigma_{\text{noise}}^2}{\sigma_{\mathbf{w}}^2}\,\mathbf{I}\Big)^{-1}\boldsymbol{\Phi}^\top\mathbf{y}$$

The predictive distribution is given by:

$$\begin{aligned} p(y_*|x_*,\mathbf{x},\mathbf{y},\mathcal{M}) \;&=\; \int p(y_*|\mathbf{w},x_*,\mathcal{M})p(\mathbf{w}|\mathbf{x},\mathbf{y},\mathcal{M})d\mathbf{w} \\ &=\; \mathcal{N}(y_*;\ \boldsymbol{\phi}(x_*)^\top\boldsymbol{\mu},\ \boldsymbol{\phi}(x_*)^\top\boldsymbol{\Sigma}\boldsymbol{\phi}(x_*) + \sigma_{\text{noise}}^2). \end{aligned}$$

Multiple explanations of the data

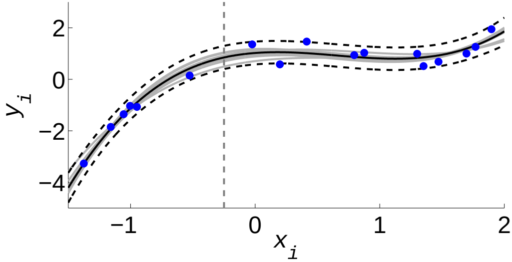

Remember that a finite linear model $f(x_n) = \boldsymbol{\phi}(x_n)^\top \mathbf{w}$ with prior on the weights $p(\mathbf{w}) = \mathcal{N}(\mathbf{w};\ \mathbf{0}, \sigma_\mathbf{w}^2 \mathbf{I})$ has a posterior distribution

$$p(\mathbf{w}|\mathbf{x}, \mathbf{y}, \mathcal{M}) \ = \ \mathcal{N}(\mathbf{w};\ \boldsymbol{\mu},\ \boldsymbol{\Sigma}) \quad \text{with} \quad \begin{aligned} \boldsymbol{\Sigma} \ &= \ \big(\sigma_{\text{noise}}^{-2} \boldsymbol{\Phi}^\top \boldsymbol{\Phi} + \sigma_\mathbf{w}^{-2}\big)^{-1} \\ \boldsymbol{\mu} \ &= \ \Big(\boldsymbol{\Phi}^\top \boldsymbol{\Phi} + \frac{\sigma_{\text{noise}}^2}{\sigma_\mathbf{w}^2}\, \mathrm{I}\Big)^{-1} \boldsymbol{\Phi}^\top \mathbf{y} \end{aligned}$$

and predictive distribution

$$p(y_*|x_*, \mathbf{x}, \mathbf{y}, \mathcal{M}) \ = \ \mathcal{N}(y_*;\ \boldsymbol{\phi}(x_*)^\top \boldsymbol{\mu},\ \boldsymbol{\phi}(x_*)^\top \boldsymbol{\Sigma} \boldsymbol{\phi}(x_*) + \sigma_{\text{noise}}^2\, \mathbf{I})$$

Marginal likelihood (Evidence) of our polynomials

Marginal likelihood, or "evidence" of a finite linear model:

$$\begin{aligned} p(\mathbf{y}|\mathbf{x},\mathcal{M}) &= \int p(\mathbf{w}|\mathbf{x},\mathcal{M})p(\mathbf{y}|\mathbf{x},\mathbf{w},\mathcal{M})d\mathbf{w} \\ &= \mathcal{N}(\mathbf{y};\ \mathbf{0}, \sigma_{\mathbf{w}}^2\,\mathbf{\Phi}\,\mathbf{\Phi}^\top + \sigma_{\text{noise}}^2\,\mathbf{I}). \end{aligned}$$

Luckily for Gaussian noise there is a closed-form analytical solution!

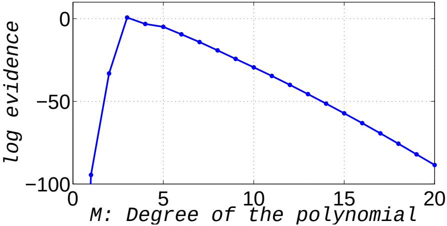

- The evidence prefers $M = 3$, not simpler, not more complex.
- Too simple models consistently miss most data.
- Too complex models frequently miss some data.