Recitation 11

Gradient Boosting

Colin

Spring 2022

Apr 13

Motivation

- We started with linear models (regression, classification)
- Introduced feature transformation to incorporate non-linearity
- Learned interpretable non-parametric models that are "very non-linear"
- Combining small models to tackle complex problems

Additive Models

1. *Additive models* over a *base hypothesis space* $\mathcal{H}$ take the form

$$\mathcal{F} = \left\{ f(x) = \sum_{m=1}^{M} \nu_m h_m(x) \mid h_m \in \mathcal{H}, \nu_m \in \mathbb{R} \right\}.$$

2. Since we are taking linear combinations, we assume the h_m functions take values in $\mathbb{R}$ or some other vector space.
3. Empirical risk minimization over $\mathcal{F}$ tries to find

$$\arg\min_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^{n} \ell(y_i, f(x_i)).$$

4. This in general is a difficult task, as the number of base hypotheses M is unknown, and each base hypothesis h_m ranges over all of $\mathcal{H}$.

Forward Stagewise Additive Modeling (FSAM)

The FSAM method fits additive models using the following (greedy) algorithmic structure:

1. Initialize $f_0 \equiv 0$.
2. For stage $m = 1, \ldots, M$:
   1. Choose $h_m \in \mathcal{H}$ and $\nu_m \in \mathbb{R}$ so that
   $$f_m = f_{m-1} + \nu_m h_m$$
   has the minimum empirical risk.
   2. The function f_m has the form
   $$f_m = \nu_1 h_1 + \cdots + \nu_m h_m.$$

- When choosing h_m, ν_m during stage m, we must solve the minimization
$$(\nu_m, h_m) = \operatorname*{arg\,min}_{\nu \in \mathbb{R}, h \in \mathcal{H}} \sum_{i=1}^{n} \ell(y_i, f_{m-1}(x_i) + \nu h(x_i)).$$

Gradient Boosting

1. Can we simplify the following minimization problem:

$$(\nu_m, h_m) = \underset{\nu\in\mathbb{R}, h\in\mathcal{H}}{\arg\min} \sum_{i=1}^{n} \ell(y_i, f_{m-1}(x_i) + \nu h(x_i)).$$

2. How about just take a step along the steepest descent direction?
3. Issue 1: h is a function instead of a vector
4. Issue 2: h must lies in $\mathcal{H}$, the base hypothesis space,
5. Solution 1: Treat h as a vector of the size of the training set $(h(x_1), \ldots, h(x_n))$ rather than a function.
6. Solution 2: Compute unconstrained steepest descent direction, and then find the closest choices in $\mathcal{H}$.

Gradient Boosting Machine

1. Initialize $f_0 \equiv 0$.
2. For stage $m = 1, \ldots, M$:
   1. Compute the steepest descent direction (also called *pseudoresiduals*):
   $$r_m = -\left(\frac{\partial}{\partial f_{m-1}(x_1)}\ell(y_1, f_{m-1}(x_1)), \ldots, \frac{\partial}{\partial f_{m-1}(x_n)}\ell(y_n, f_{m-1}(x_n))\right).$$
   2. Find the closest base hypothesis (using Euclidean distance):
   $$h_m = \arg\min_{h \in \mathcal{H}} \sum_{i=1}^{n} ((r_m)_i - h(x_i))^2.$$
   3. Choose fixed step size $\nu_m \in (0, 1]$ or line search:
   $$\nu_m = \arg\min_{\nu \geq 0} \sum_{i=1}^{n} \ell(y_i, f_{m-1}(x_i) + \nu h_m(x_i)).$$
   4. Take the step:
   $$f_m(x) = f_{m-1}(x) + \nu_m h_m(x).$$

Gradient Boosting Machine

1. Each stage we need to solve the following step:

$$h_m = \arg\min_{h \in \mathcal{H}} \sum_{i=1}^{n} ((r_m)_i - h(x_i))^2.$$

   How do we do this?

2. This is a standard least squares regression task on the "mock" dataset

$$\mathcal{D}^{(m)} = \{(x_1, (r_m)_1), \ldots, (x_n, (r_m)_n)\}.$$

3. We assume that we have a learner that (approximately) solves least squares regression over $\mathcal{H}$.

Gradient Boosting Comments

1. The algorithm above is sometimes called AnyBoost or Functional Gradient Descent.
2. The most commonly used base hypothesis space is small regression trees (between 4 and 8 leaves).

Practice With Different Loss Functions

Question

Explain how to perform gradient boosting with the following loss functions:

1. Square loss: $\ell(y, a) = (y - a)^2/2$.
2. Absolute loss: $\ell(y, a) = |y - a|$.
3. Exponential margin loss: $\ell(y, a) = e^{-ya}$.

Solution: Square loss

Using $\ell(y, a) = (y - a)^2/2$
To compute an arbitrary pseudoresidual we first note that

$$\frac{\partial \ell}{\partial a} = -(y - a)$$

giving

$$-\frac{\partial \ell}{\partial f_{m-1}(x_i)} = (y_i - f_{m-1}(x_i)).$$

In words, for the square loss, the pseudoresiduals are simply the residuals from the previous stage's fit. Thus, in stage m our step direction h_m is given by solving

$$h_m := \arg\min_{h \in \mathcal{H}} \sum_{i=1}^{n} ((y_i - f_{m-1}(x_i)) - h(x_i))^2.$$

Solution: Absolute Loss

Using $\ell(y, a) = |y - a|$
Note that

$$\frac{\partial \ell}{\partial a} = -\operatorname{sgn}(y - a)$$

giving

$$-\frac{\partial \ell}{\partial f_{m-1}(x_i)} = \operatorname{sgn}(y_i - f_{m-1}(x_i)).$$

The absolute loss only cares about the sign of the residual from the previous stage's fit. Thus, in stage m our step direction h_m is given by solving

$$h_m := \operatorname*{arg\,min}_{h \in \mathcal{H}} \sum_{i=1}^{n} (\operatorname{sgn}(y_i - f_{m-1}(x_i)) - h(x_i))^2.$$

Solution: Exponential Loss

Using $\ell(y, a) = e^{-ya}$
Note that

$$\frac{\partial \ell}{\partial a} = -ye^{-ya}$$

giving

$$-\frac{\partial \ell}{\partial f_{m-1}(x_i)} = y_i e^{-y_i f_{m-1}(x_i)}.$$

Thus, in stage m our step direction h_m is given by solving

$$h_m := \underset{h \in \mathcal{H}}{\arg\min} \sum_{i=1}^{n} (y_i e^{-y_i f_{m-1}(x_i)} - h(x_i))^2.$$

Review

- Notice the minimization step does not require any differentiability.
- One can optimize/train base on hypothesis space.
- This almost is as if we are taking grading with respect to a function which opens a lot of possibilities, expanding our choice of functions.