FoDA

L13 : Cross Validation

Data (X, y) $\quad X \in \mathbb{R}^{n \times d}$ $(d=1)$ $\quad y \in \mathbb{R}^n$

$\hat{y}_i = M_\alpha(x_i) \rightarrow y_i$

polynomial model

$M_\alpha^{(p)}(x) = \sum_{j=0}^{p} \alpha_j x^p = \alpha_0 + \alpha_1 x + \alpha_2 x^2 + \dots$

residual $\quad \hat{y} - y_i$

$SSE = \sum_i \varepsilon_i^2 (r_i)^2$

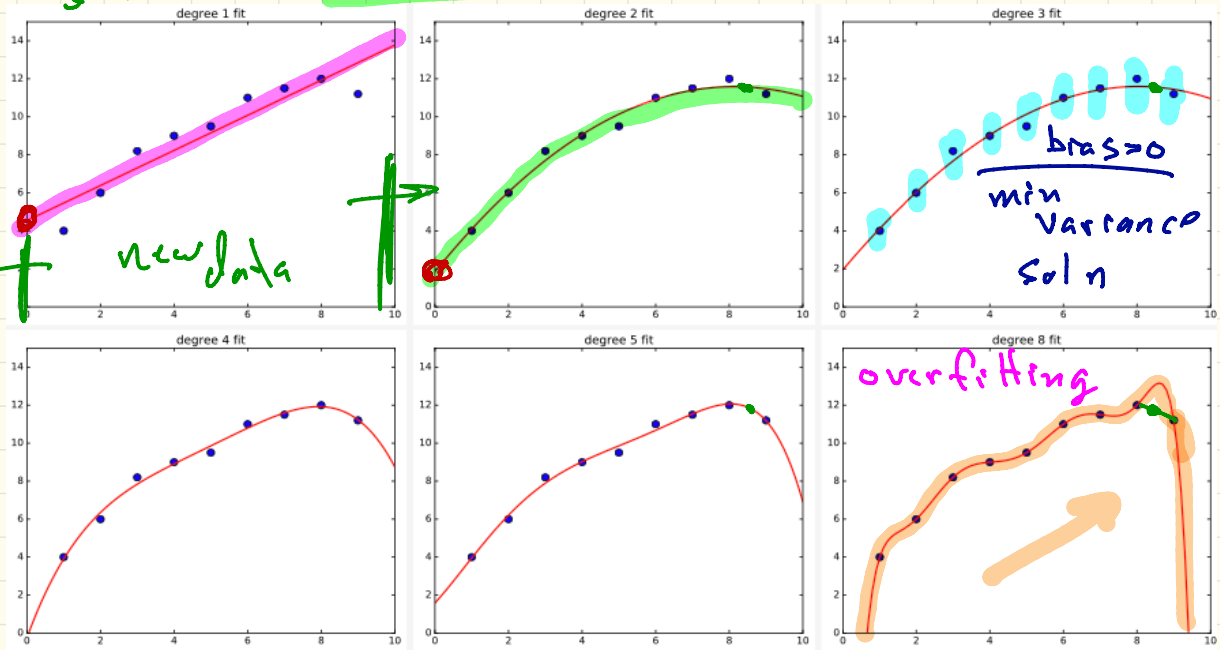

Goal: Make good predictions on new unseen data.

modeling as residual $r_i = |M_k(x_i) - y_i|$

"generalization"

- Collect new data
  ↳ try on this new data
- "Save" some data for testing

Data

assumption each (x_i, y_i) iid f

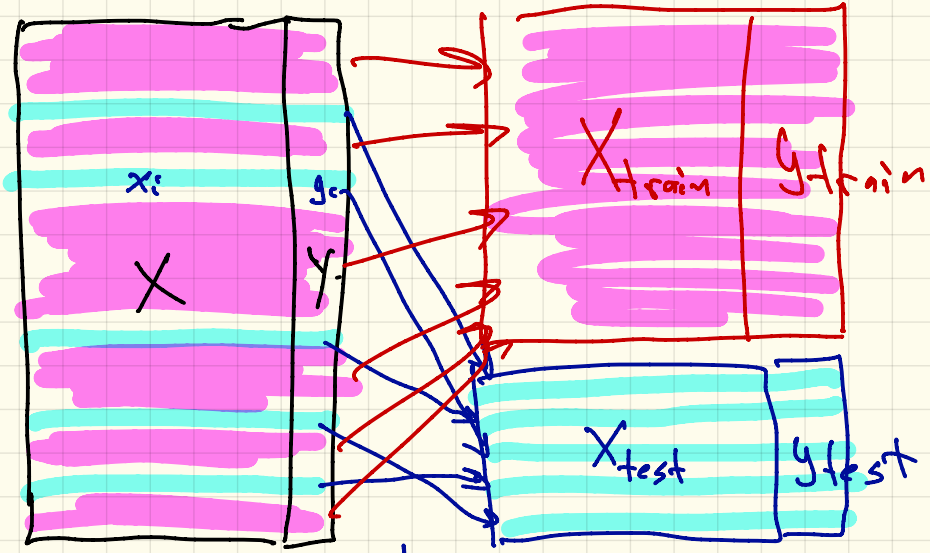

random hold out

$$\alpha = (X_{train}^T X_{train})^{-1} X_{train}^T y_{train}$$

↓

f_{train} eval

$$SSE\left((X_{test}, y_{test}), \alpha_{train}\right) = \sum_{(x_i, y_i) \in (X_{test}, y_{test})} \left(M_{\alpha_{train}}(x_i) - y_i\right)^2$$

How well will work on new data?

$$\frac{1}{|X_{test}|} SSE\left((X_{test}, y_{test}), \alpha_{train}\right)$$

RMSE

How large should the test set be?

Common test size

- 10%
- 33%

Evaluate expected value of error by averaging n_{test} observations.

↳ CLT

→ unbiased

→ variance $\frac{Var}{n_{test}}$

more data
less test percentage

more complex model
more test size

What is cross-validation used for?

- See how model generalizes to new data.

or

- to select a parameter in model (ex. p in $M_\alpha^{(p)}$)

$$p^* = \underset{p \in [1 \ldots 8]}{\arg\min} \; SSE\left((X_{test}, y_{test}), M_{\alpha_{train}}^{(p)}\right)$$

choose best

train & model

not both

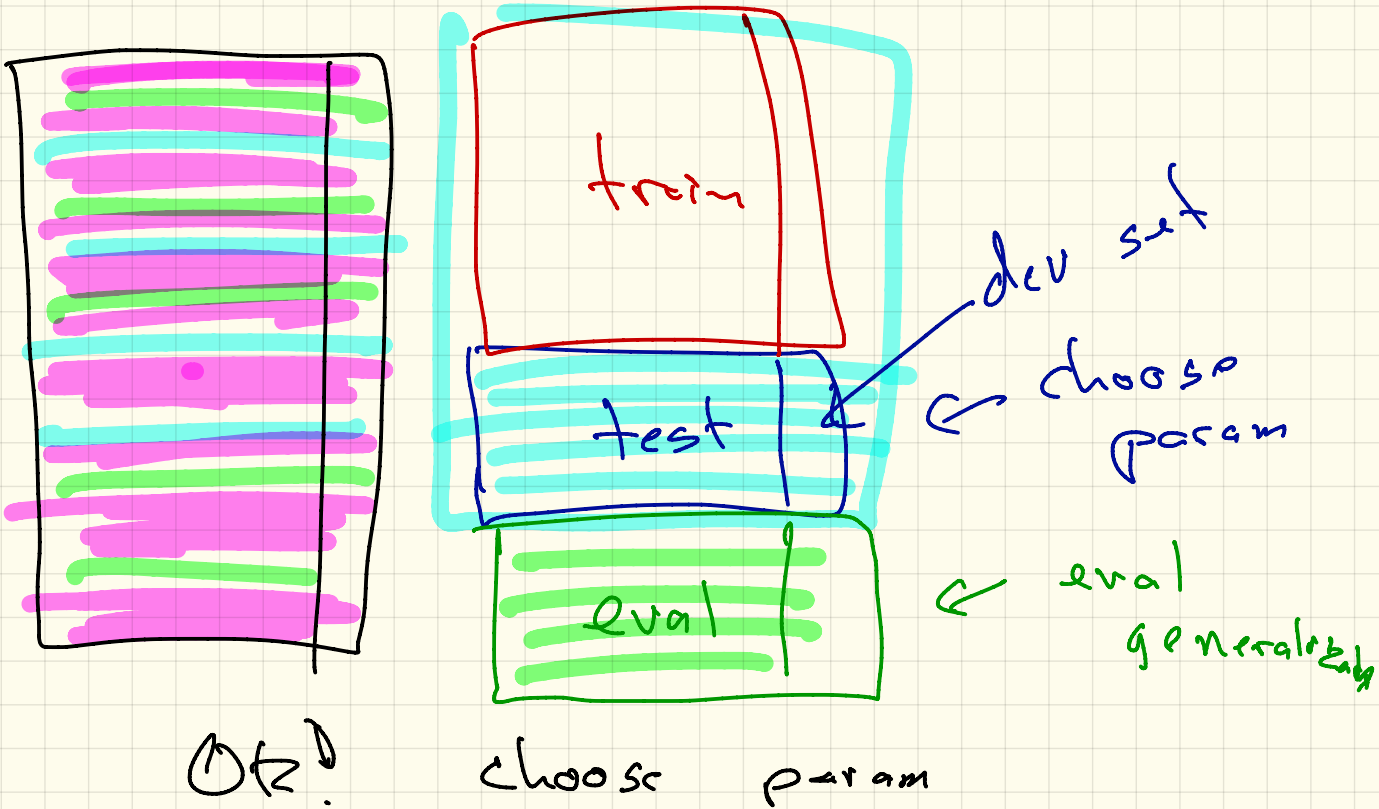
train
dev set
test
choose param
eval
generalization
Otz!
choose param
+
evaluate generalization

If your data is small.

don't want to waste data on test.

Leave-one-out CV.

test set size 1

but try all

Double Descent

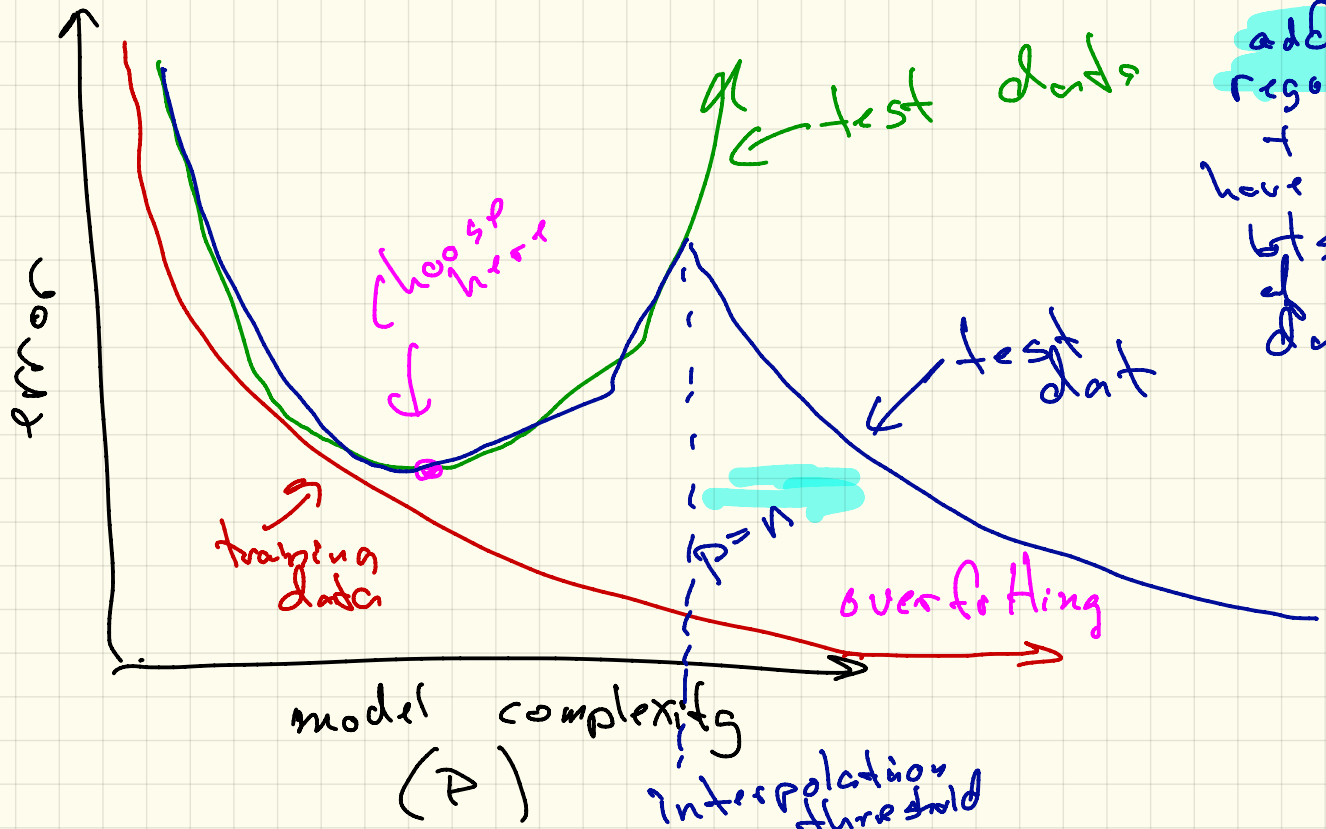

if you add regularization + have lots of data