CS 273A: Machine Learning

Winter 2021

Lecture 8: VC Dimension

Roy Fox
Department of Computer Science
Bren School of Information and Computer Sciences
University of California, Irvine

All slides in this course adapted from Alex Ihler & Sameer Singh

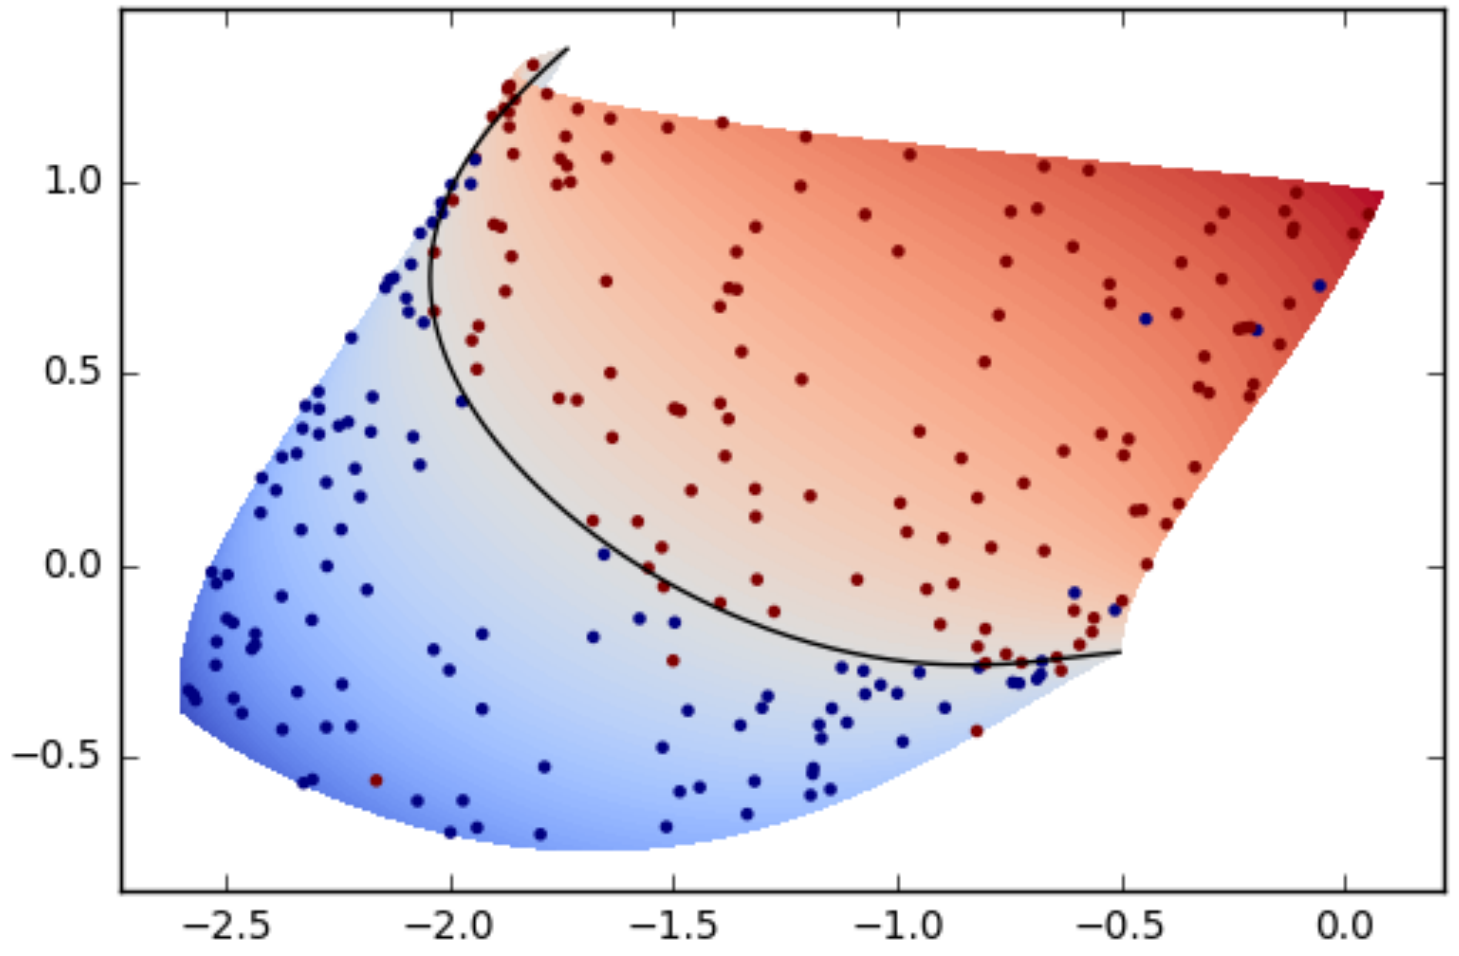

Logistics

project

- Team rosters due Monday, Feb 1 **on Canvas**
- Team-forming spreadsheet posted on piazza

midterm

- Midterm exam on Feb 9, 2–4pm **on Canvas**
- We'll accommodate other timezones — let us know

Today's lecture

Logistic Regression

Multi-class classifiers

VC dimension

Perceptron

- Perceptron = linear classifier
  - Parameters θ = weights (also denoted w)
  - Response = weighted sum of the features $r = \theta^\intercal x$
  - Prediction = thresholded response $\hat{y}(x) = T(r) = T(\theta^\intercal x)$

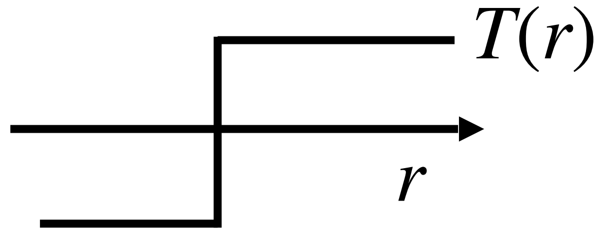

  - Decision function: $\hat{y}(x) = \begin{cases} +1 & \text{if } \theta^\intercal x > 0 \\ -1 & \text{otherwise} \end{cases}$ (for $T(r) = \text{sign}(r)$)
- Update rule: $\theta \leftarrow \theta - \alpha(\underbrace{y - \hat{y}}_{\text{error}})x$

Surrogate loss functions

- Alternative: use differentiable loss function
  - E.g., approximate the step function with a smooth function

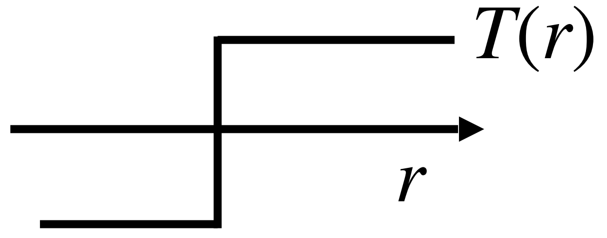

  - Popular choice: logistic / sigmoid function (sigmoid = ”looks like s”)

$$\sigma(r) = \frac{1}{1 + \exp(-r)}$$

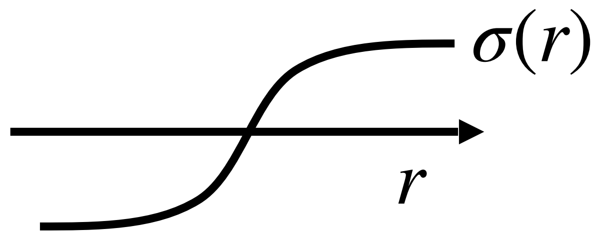

- MSE loss: $\mathscr{L}_\theta(x, y) = (y - \sigma(r(x)))^2$

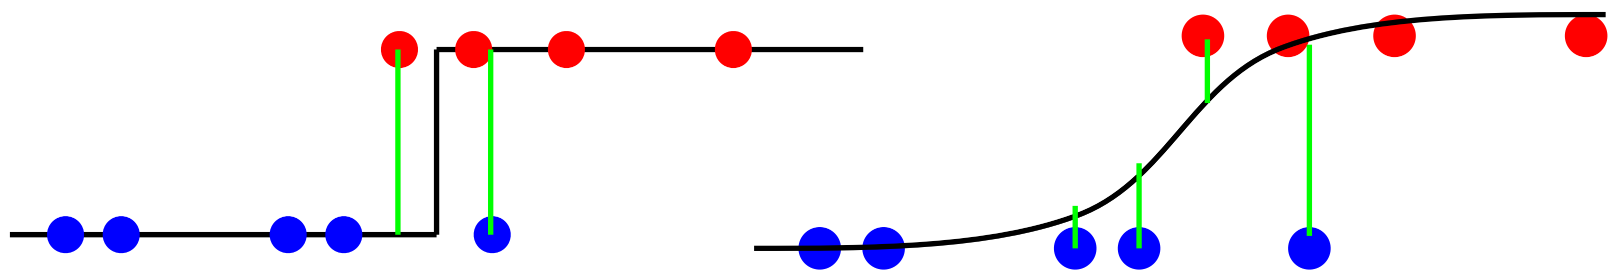

  - Far from the boundary: $\sigma \approx 0$ or 1, loss approximates 0–1 loss
  - Near the boundary: $\sigma \approx \frac{1}{2}$, loss near $\frac{1}{4}$, but clear improvement direction

Learning smooth linear classifiers

- Use gradient-based optimizer on the loss $\mathscr{L}_\theta(x, y) = (y - \sigma(\theta^\intercal x))^2$

$$-\nabla_\theta \mathscr{L}_\theta(x, y) = 2\underbrace{(y - \sigma(\theta^\intercal x))}_{\text{error}}\underbrace{\sigma'(\theta^\intercal x)}_{\text{sensitivity}}\, x$$

- Logistic / sigmoid function: $\sigma(r) = \dfrac{1}{1 + \exp(-r)}$

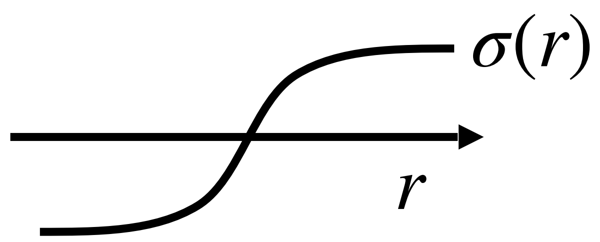

- It's derivative: $\sigma'(r) = \sigma(r)(1 - \sigma(r))$
  - Saturates for both $r \to \infty$, $r \to -\infty$
- Confidently correct prediction: $\sigma(r) \approx y \in \{0,1\} \implies \nabla_\theta \mathscr{L}_\theta \approx 0$ ← good
- Confidently incorrect prediction: $\sigma(r) \approx 1 - y \implies \nabla_\theta \mathscr{L}_\theta \approx 0$ ← bad

Maximum likelihood

- What if we had a probabilistic predictor $p_\theta(y \mid x)$?
- The better the parameter θ, the more likely the training data:

$$p_\theta(y^{(1)}, \dots, y^{(m)} \mid x^{(1)}, \dots, x^{(m)}) = \prod_j p_\theta(y^{(j)} \mid x^{(j)})$$

- Bayesian interpretation?

  except, often there's no uniform distribution over parameter space

  - MAP: $\arg\max_\theta p(\theta \mid \mathscr{D}) = \arg\max_\theta p(\theta) p(\mathscr{D}_x) p_\theta(\mathscr{D}_y \mid \mathscr{D}_x) = \arg\max_\theta p_\theta(\mathscr{D}_y \mid \mathscr{D}_x)$

  average over training dataset

- Maximum log-likelihood: $\max_\theta \frac{1}{m} \sum_j \log p_\theta(y^{(j)} \mid x^{(j)})$

Logistic Regression

- Can we turn a linear response into a probability? Sigmoid! $\sigma : \mathbb{R} \to [0,1]$
- Think of $\sigma(\theta^\intercal x) = p_\theta(y = 1 \,|\, x)$
- Negative Log-Likelihood (NLL) loss:

$$\mathscr{L}_\theta(x, y) = -\log p_\theta(y \,|\, x) = -\underbrace{y \log \sigma(\theta^\intercal x)}_{\textbf{for } y=1} - \underbrace{(1-y)\log(1-\sigma(\theta^\intercal x))}_{\textbf{for } y=0}$$

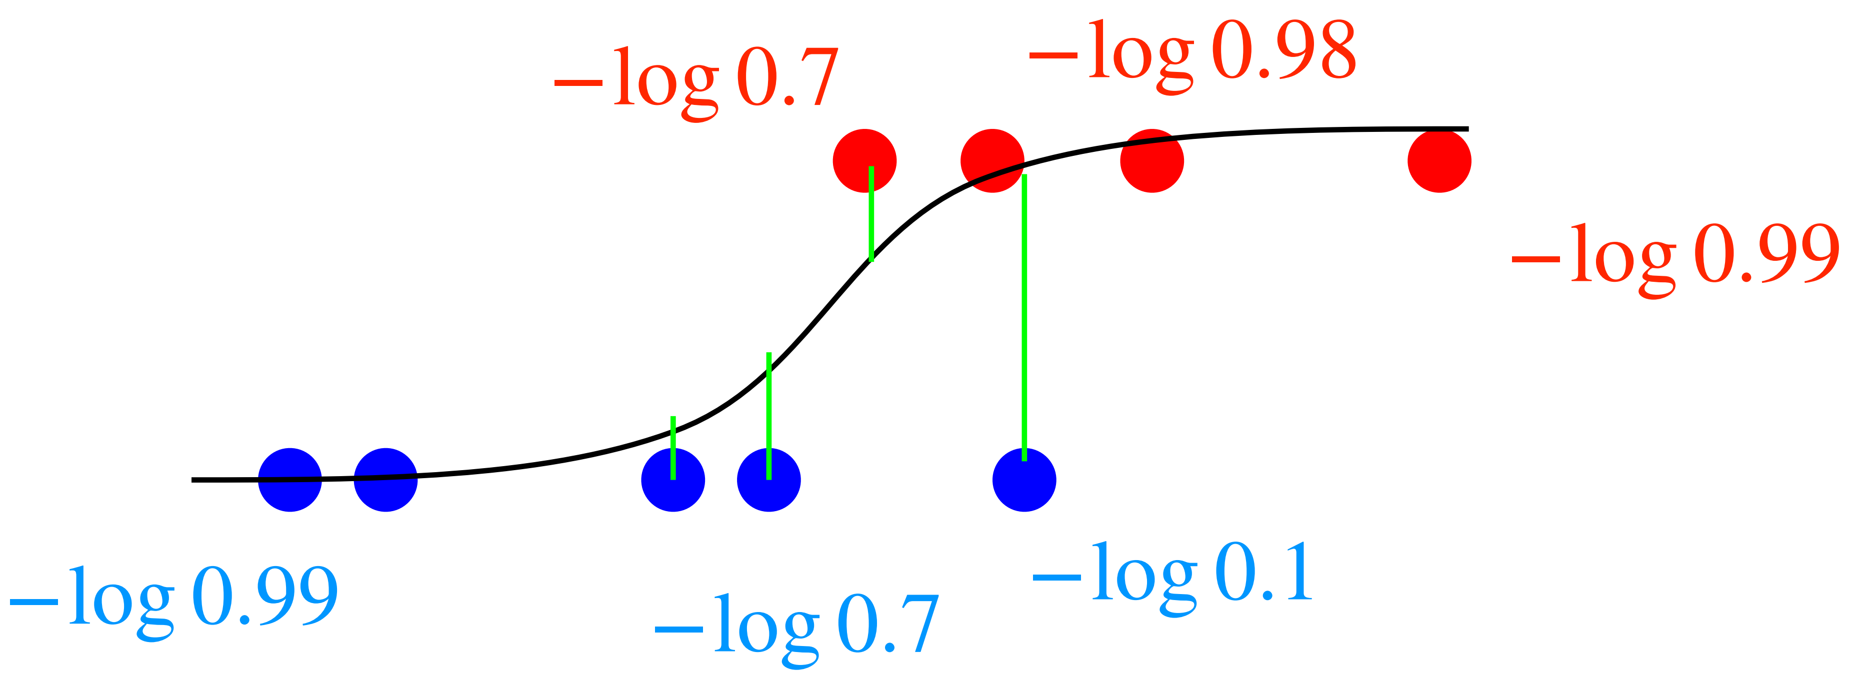

Logistic Regression: gradient

- Logistic NLL loss: $\mathscr{L}_\theta(x, y) = -y \log \sigma(\theta^\intercal x) - (1-y)\log(1-\sigma(\theta^\intercal x))$

- Gradient:

$$-\nabla_\theta \mathscr{L}_\theta(x,y) = \left(y\frac{\sigma'(\theta^\intercal x)}{\sigma(\theta^\intercal x)} - (1-y)\frac{\sigma'(\theta^\intercal x)}{1-\sigma(\theta^\intercal x)} \right) x$$

$$= (\underbrace{y\,(1-\sigma(\theta^\intercal x))}_{\text{error for } y=1} - (1-y)\underbrace{\sigma(\theta^\intercal x)}_{\text{error for } y=0})x$$

but update toward $-x$

$$= (y - p_\theta(y=1\,|\,x))x$$

- Compare:

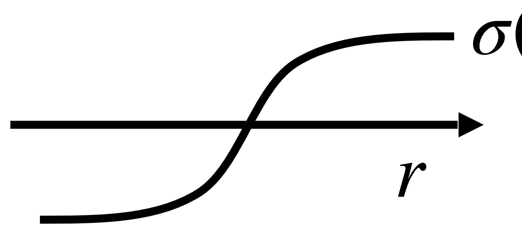

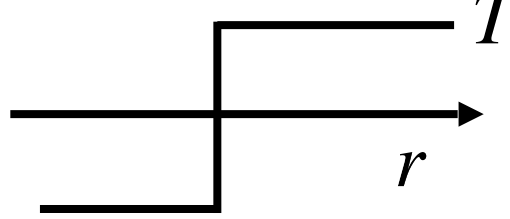

  - Perceptron: $(y - \hat{y})x$ ⟵ constant error (± 2), insensitive to margin
  - Logistic MSE: $-\nabla_\theta \mathscr{L}_\theta(x,y) = 2(y - \sigma(\theta^\intercal x))\sigma'(\theta^\intercal x)x$ ⟵ 0 gradient for bad mistakes

Recap

- Linear classifiers:
  - Perceptron
  - Logistic classifier
- Measuring decision quality:
  - Error rate / 0–1 loss
  - MSE loss
  - Negative log-likelihood (Logistic Regression)
- Learning the weights
  - Perceptron algorithm — not quite gradient-based (or gradient of weird loss)
  - Gradient-based optimization of surrogate loss (MSE / NLL)

Today's lecture

Logistic Regression

Multi-class classifiers

VC dimension

Multi-class linear models

- How to predict multiple classes?
- Idea: have a linear response per class $r_c = \theta_c^\intercal x$
  - Choose class with largest response: $f_\theta(x) = \arg\max_c \theta_c^\intercal x$
- Linear boundary between classes c_1, c_2:
  - $\theta_{c_1}^\intercal x \lessgtr \theta_{c_2}^\intercal x \iff (\theta_{c_1} - \theta_{c_2})^\intercal x \lessgtr 0$

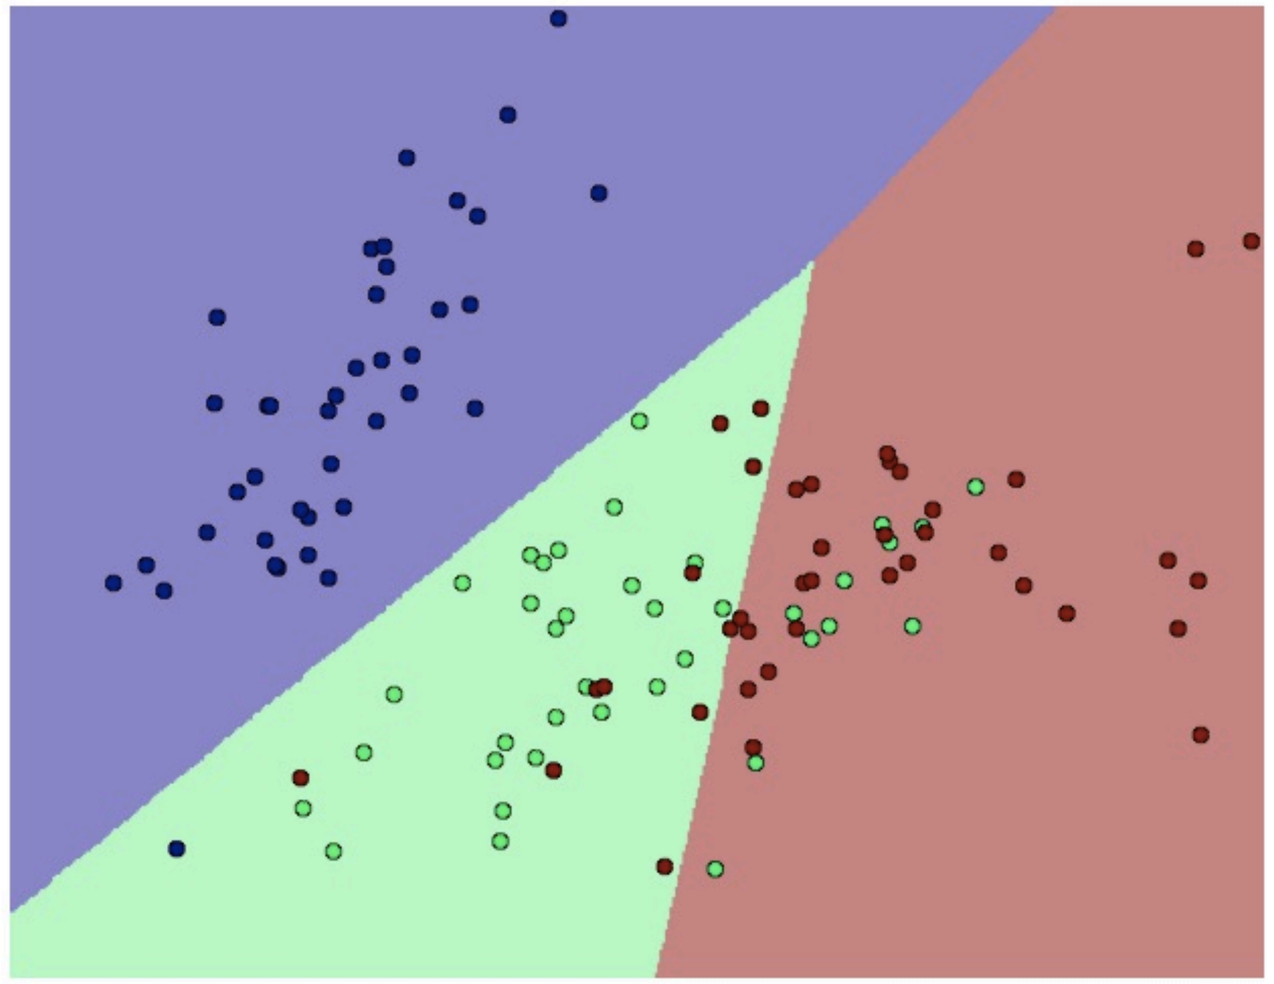

Multi-class linear models

- More generally: add features — can even depend on y!

$$f_\theta(x) = \arg\max_y \theta^\intercal \Phi(x, y)$$

- Example: $y = \pm 1$
  - $\Phi(x, y) = xy$

$$\implies f_\theta(x) = \arg\max_y y\theta^\intercal x = \begin{cases} +1 & +\theta^\intercal x > -\theta^\intercal x \\ -1 & +\theta^\intercal x < -\theta^\intercal x \end{cases}$$

$$= \operatorname{sign}(\theta^\intercal x) \longleftarrow \textbf{perceptron!}$$

Multi-class linear models

- More generally: add features — can even depend on y!

$$f_\theta(x) = \arg\max_y \theta^\intercal \Phi(x, y)$$

- Example: $y \in \{1,2,\ldots,C\}$
  - $\Phi(x, y) = [0 \;\; 0 \;\; \cdots \;\; x \;\; \cdots \;\; 0] =$ one-hot$(y) \otimes x$
  - $\theta = [\theta_1 \;\; \cdots \;\; \theta_C]$

$$\implies f_\theta(x) = \arg\max_c \theta_c^\intercal x \longleftarrow \textbf{largest linear response}$$

Multi-class perceptron algorithm

- While not done:
  - For each data point $(x, y) \in \mathscr{D}$:
    - Predict: $\hat{y} = \arg\max_c \theta_c^\intercal x$
    - Increase response for true class: $\theta_y \leftarrow \theta_y + \alpha x$
    - Decrease response for predicted class: $\theta_{\hat{y}} \leftarrow \theta_{\hat{y}} - \alpha x$
- More generally:
  - Predict: $\hat{y} = \arg\max_y \theta^\intercal \Phi(x, y)$
  - Update: $\theta \leftarrow \theta + \alpha(\Phi(x, y) - \Phi(x, \hat{y}))$

Multilogit Regression

- Define multi-class probabilities: $p_\theta(y \mid x) = \dfrac{\exp(\theta_y^\intercal x)}{\sum_c \exp(\theta_c^\intercal x)} = \operatorname{soft\,max}_c \theta_c^\intercal x \Big|_y$
  - For binary y:
    $$p_\theta(y = 1 \mid x) = \frac{\exp(\theta_1^\intercal x)}{\exp(\theta_1^\intercal x) + \exp(\theta_2^\intercal x)}$$
    $$= \frac{1}{1 + \exp((\theta_2 - \theta_1)^\intercal x)} = \sigma((\theta_1 - \theta_2)^\intercal x)$$

Logistic Regression with $\theta = \theta_1 - \theta_2$

- Benefits:
  - Probabilistic predictions: knows its confidence
  - Linear decision boundary: $\arg\max_y \exp(\theta_y^\intercal x) = \arg\max_y \theta_y^\intercal x$
  - NLL is convex

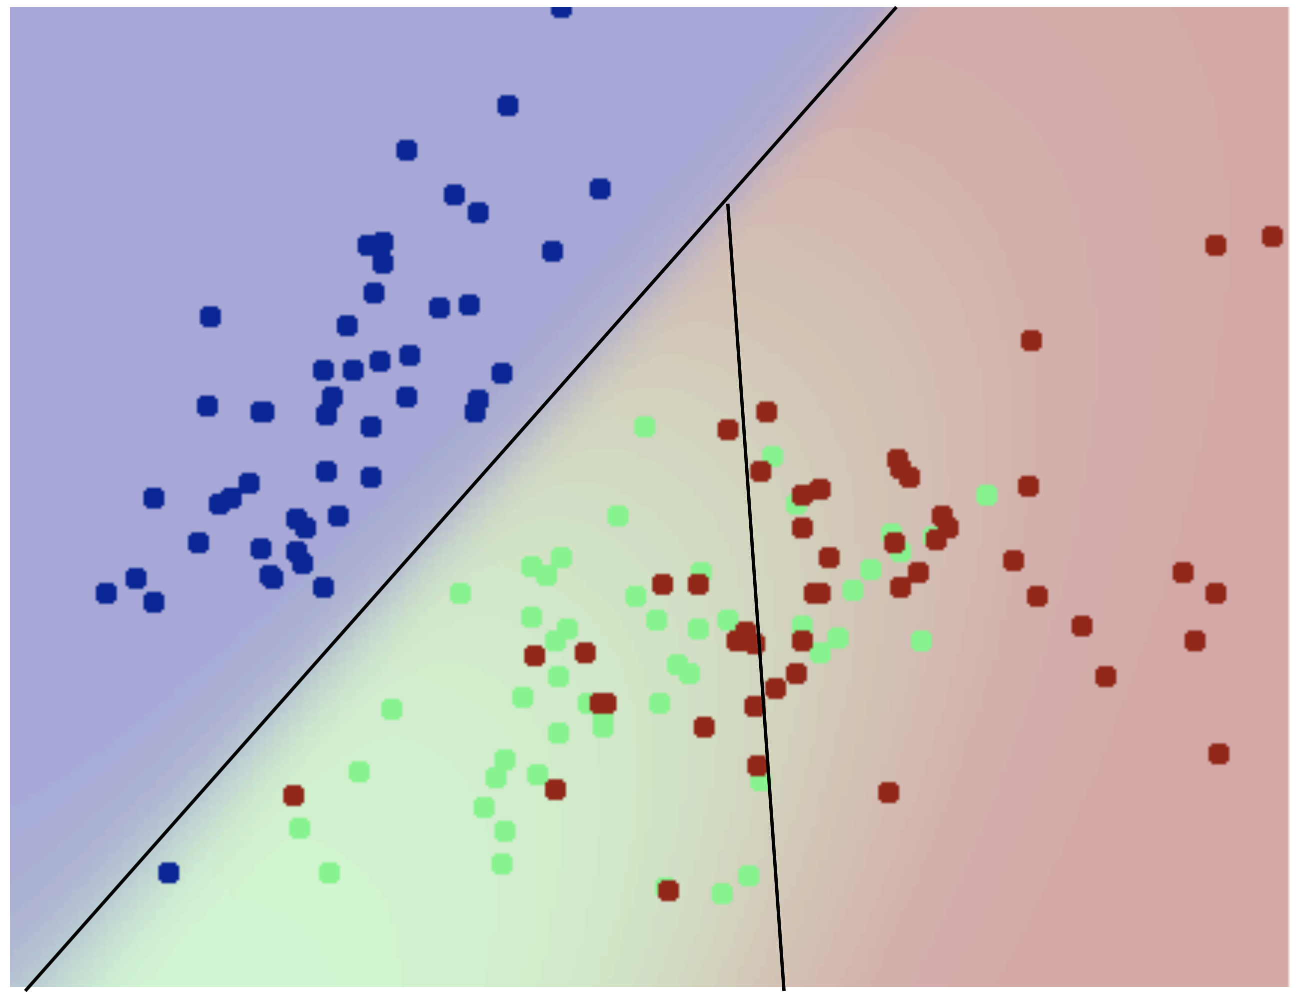

Multilogit Regression: gradient

- NLL loss: $\mathscr{L}_\theta(x,y) = -\log p_\theta(y\,|\,x) = -\theta_y^\intercal x + \log \sum_c \exp(\theta_c^\intercal x)$

- Gradient:

$$-\nabla_{\theta_c}\mathscr{L}_\theta(x,y) = \delta(y=c)x - \frac{\nabla_{\theta_c}\sum_{c'}\exp(\theta_{c'}^\intercal x)}{\sum_{c'}\exp(\theta_{c'}^\intercal x)}$$

$$= \left(\delta(y=c) - \frac{\exp(\theta_c^\intercal x)}{\sum_{c'}\exp(\theta_{c'}^\intercal x)}\right)x$$

$$= (\delta(y=c) - p_\theta(c\,|\,x))x$$

make true class more likely

make all other classes less likely

- Compare to multi-class perceptron: $(\delta(y=c) - \delta(\hat{y}=c))x$

Today's lecture

Logistic Regression

Multi-class classifiers

VC dimension

Complexity measures

- What are we looking for in a measure of model class complexity?
  - Tell us something about generalization error $\mathscr{L}_{\text{test}} - \mathscr{L}_{\text{training}}$
  - Tell us how it depends on amount of data m

    also called: risk – empirical risk
  - Be easy to find for a given model class — haha jk not gonna happen (more later)
- Ideally: a way to select model complexity (other than validation)
  - Akaike Information Criterion (AIC) — roughly: loss + #parameters
  - Bayesian Information Criterion (BIC) — roughly: loss + #parameters $\cdot \log m$
    - But what's the #parameters, effectively? Should $f_{\theta_1,\theta_2} = g_{\theta=h(\theta_1,\theta_2)}$ change the complexity?

Model expressiveness

- Model complexity also measures expressiveness / representational power
- Tradeoff:
  - More expressive $\Longrightarrow$ can reduce error, but may also overfit to training data
  - Less expressive $\Longrightarrow$ may not be able to represent true pattern / trend
- Example: $\mathrm{sign}(\theta_0 + \theta_1 x_1 + \theta_2 x_2)$

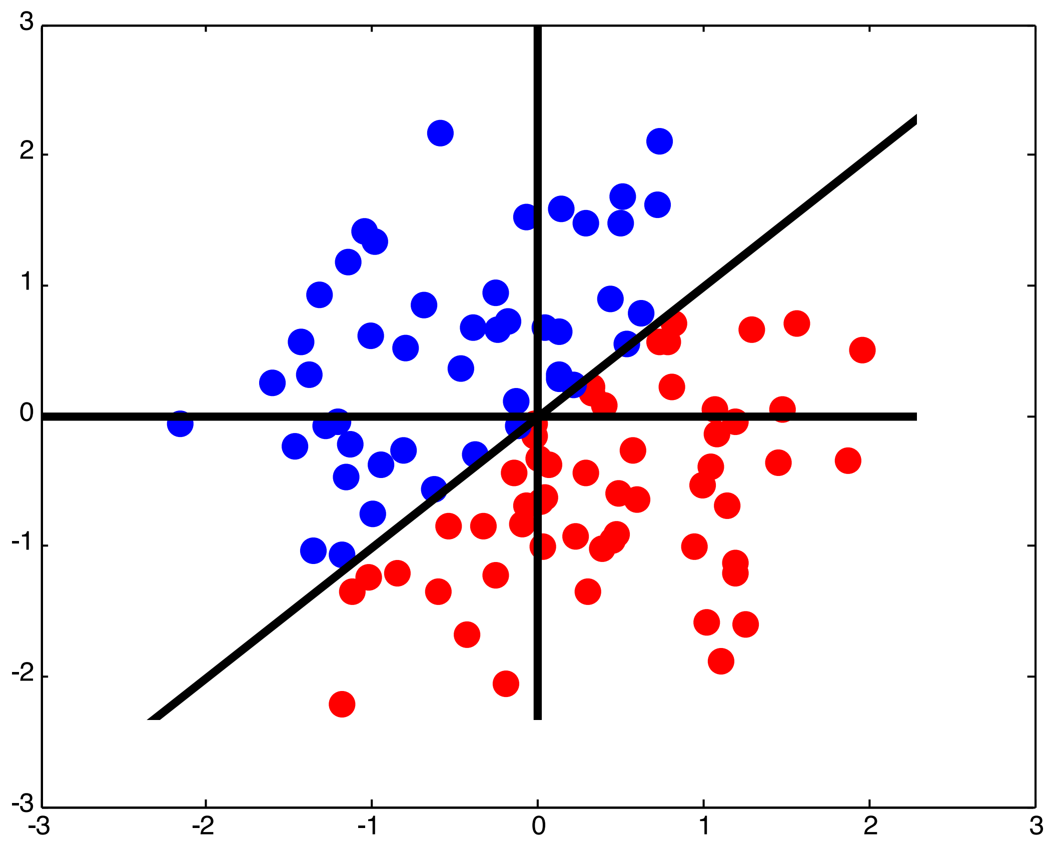

Model expressiveness

- Model complexity also measures expressiveness / representational power
- Tradeoff:
  - More expressive $\implies$ can reduce error, but may also overfit to training data
  - Less expressive $\implies$ may not be able to represent true pattern / trend
- Example: $\text{sign}(x_1^2 + x_2^2 - \theta)$

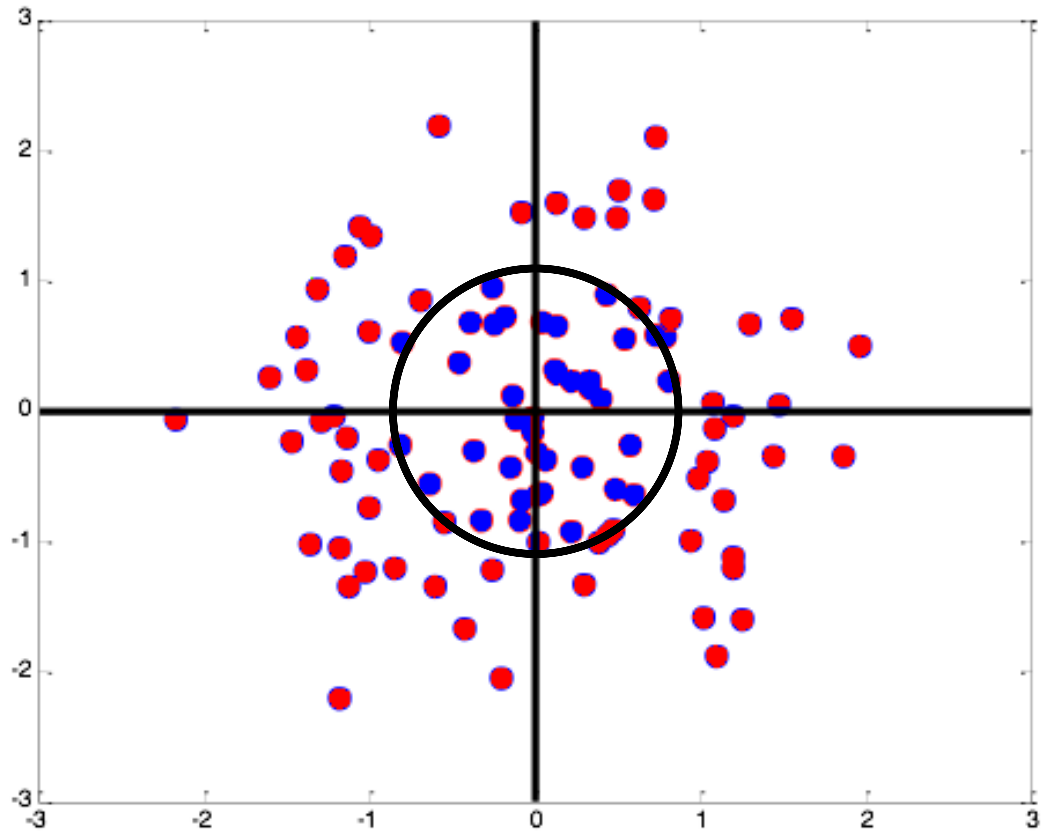

Shattering

- Separability / realizability: there's a model that classifies all points correctly
- Shattering: the points are separable regardless of their labels
  - Our model class can shatter points $x^{(1)}, \ldots, x^{(h)}$

    if for any labeling $y^{(1)}, \ldots, y^{(h)}$

    there exists a model that classifies all of them correctly
    - The model class must have at least as many models as labelings C^h

Shattering

- Separability / realizability: there's a model that classifies all points correctly
- Shattering: the points are separable <u>regardless of their labels</u>
    - Our model class can shatter points $x^{(1)}, \ldots, x^{(h)}$

      if for <u>any</u> labeling $y^{(1)}, \ldots, y^{(h)}$

      there <u>exists</u> a model that classifies all of them correctly
- Example: can $f_\theta(x) = \text{sign}(\theta_0 + \theta_1 x_1 + \theta_2 x_2)$ shatter these points?

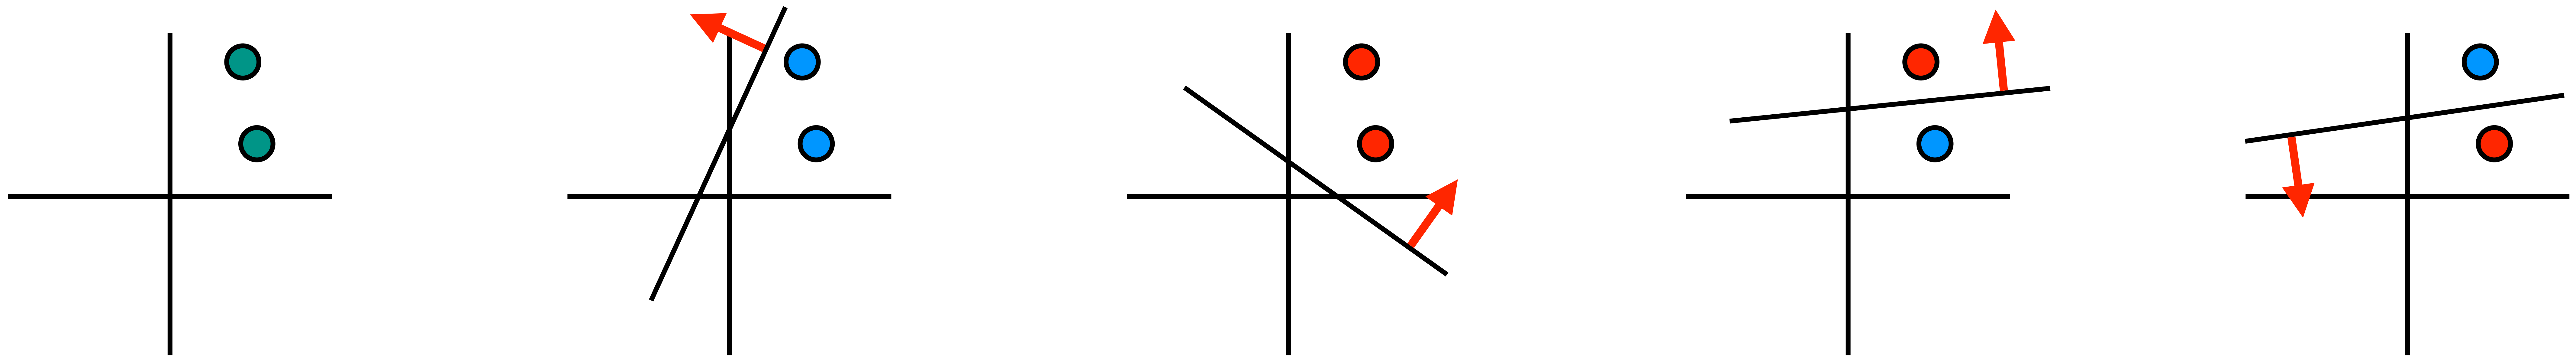

Shattering

- Separability / realizability: there's a model that classifies all points correctly
- Shattering: the points are separable regardless of their labels
  - Our model class can shatter points $x^{(1)}, \ldots, x^{(h)}$

    if for any labeling $y^{(1)}, \ldots, y^{(h)}$

    there exists a model that classifies all of them correctly
- Example: can $f_\theta(x) = \text{sign}(x_1^2 + x_2^2 - \theta)$ shatter these points?

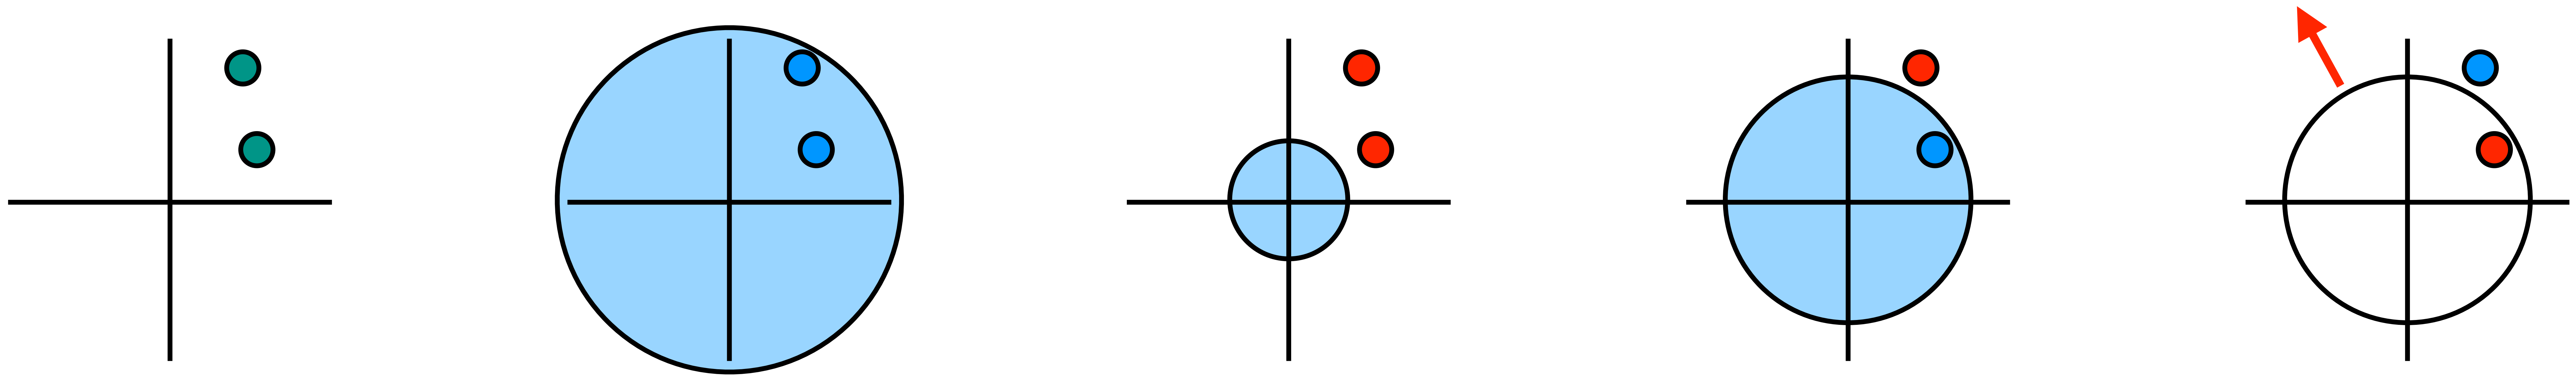

Vapnik–Chervonenkis (VC) dimension

- VC dimension: maximum number H of points that can be shattered by a class
- A game:
  - Fix a model class $f_\theta : x \to y \quad \theta \in \Theta$
  - Player 1: choose h points $x^{(1)}, \ldots, x^{(h)}$
  - Player 2: choose labels $y^{(1)}, \ldots, y^{(h)}$
  - Player 1: choose model θ
  - Are all $y^{(j)} = f_\theta(x^{(j)})$? $\implies$ Player 1 wins $\quad \exists x^{(1)}, \ldots, x^{(h)} : \ \forall y^{(1)}, \ldots, y^{(h)} : \ \exists \theta : \ \forall j : \ y^{(j)} = f_\theta(x^{(j)})$
- $h \leq H \implies$ Player 1 can win, otherwise cannot win

VC dimension: example (1)

- VC dimension: maximum number H of points that can be shattered by a class
- To find H, think like the winning player: 1 for $h \leq H$, 2 for $h > H$
- Example: $f_\theta(x) = \mathrm{sign}(x_1^2 + x_2^2 - \theta)$
  - We can place one point and ”shatter” it
  - We can prevent shattering any two points: make the distant one blue
  - $H = 1$

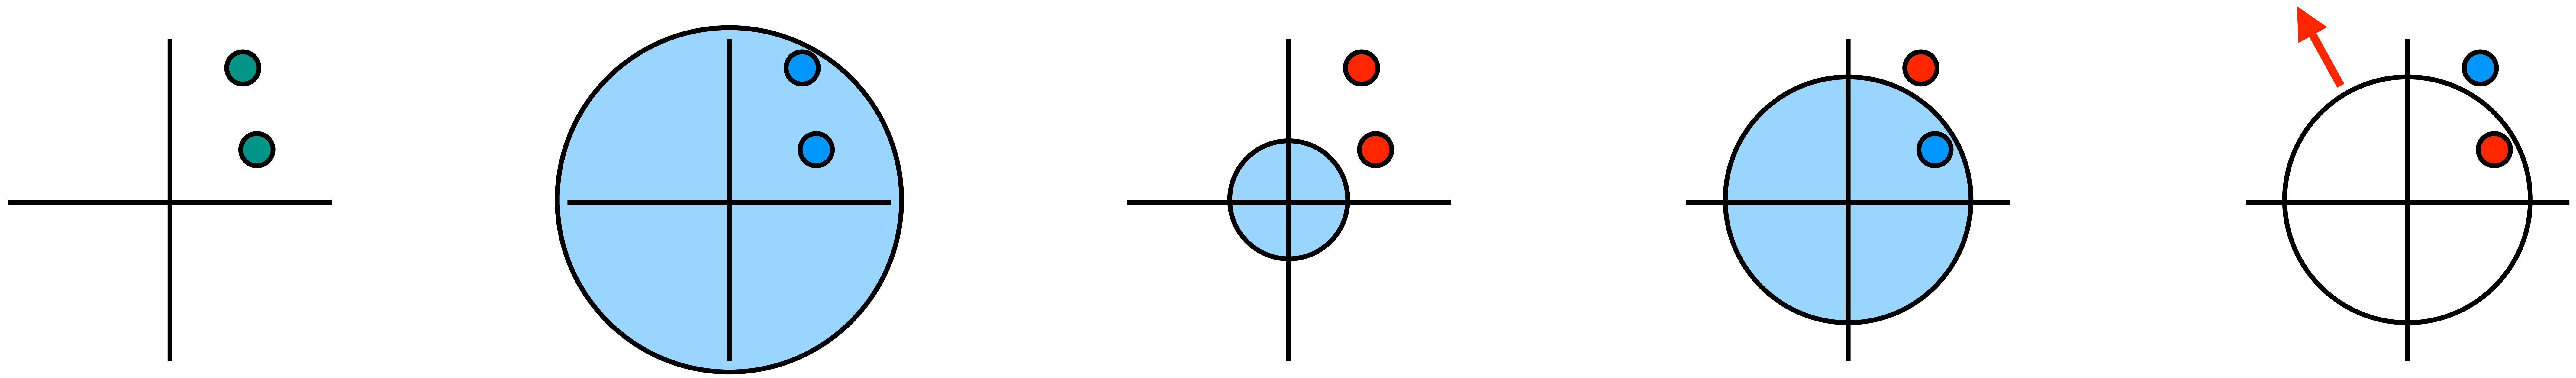

VC dimension: example (2)

- Example: $f_\theta(x) = \text{sign}(\theta_0 + \theta_1 x_1 + \theta_2 x_2)$
  - We can place 3 points and shatter them

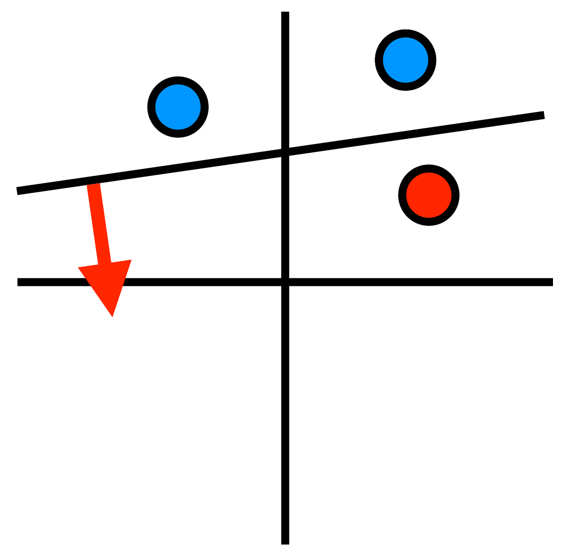

  - We can prevent shattering any 4 points:
    - If they form a convex shape, alternate labels
    - Otherwise, label differently the point in the triangle

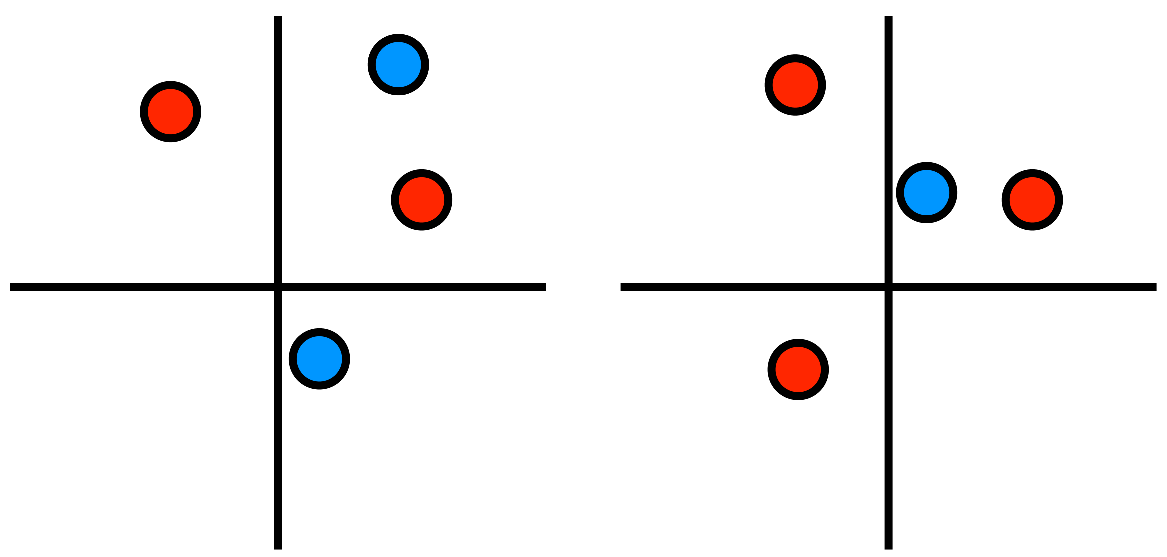

  - $H = 3$
- Linear classifiers (perceptrons) of d features have VC-dim $d + 1$
  - But VC-dim is generally not #parameters

VC Generalization bound

- VC-dim of a model class can be used to bound generalization loss:
    - With probability at least $1-\eta$, we will get a "good" dataset, for which
- $\underbrace{\text{test loss} - \text{training loss}}_{\text{generalization loss}} \leq \sqrt{\frac{H\log(2m/H) + H - \log(\eta/4)}{m}}$
- We need larger training size m:
    - The better generalization we need
    - The more complex (higher VC-dim) our model class
    - The more likely we want to get a good training sample

Model selection with VC-dim

- Using validation / cross-validation:
  - Estimate loss on held out set
  - Use validation loss to select model

- Using VC dimension:
  - Use generalization bound to select model
  - Structural Risk Minimization (SRM)
  - Bound not tight, must too conservative

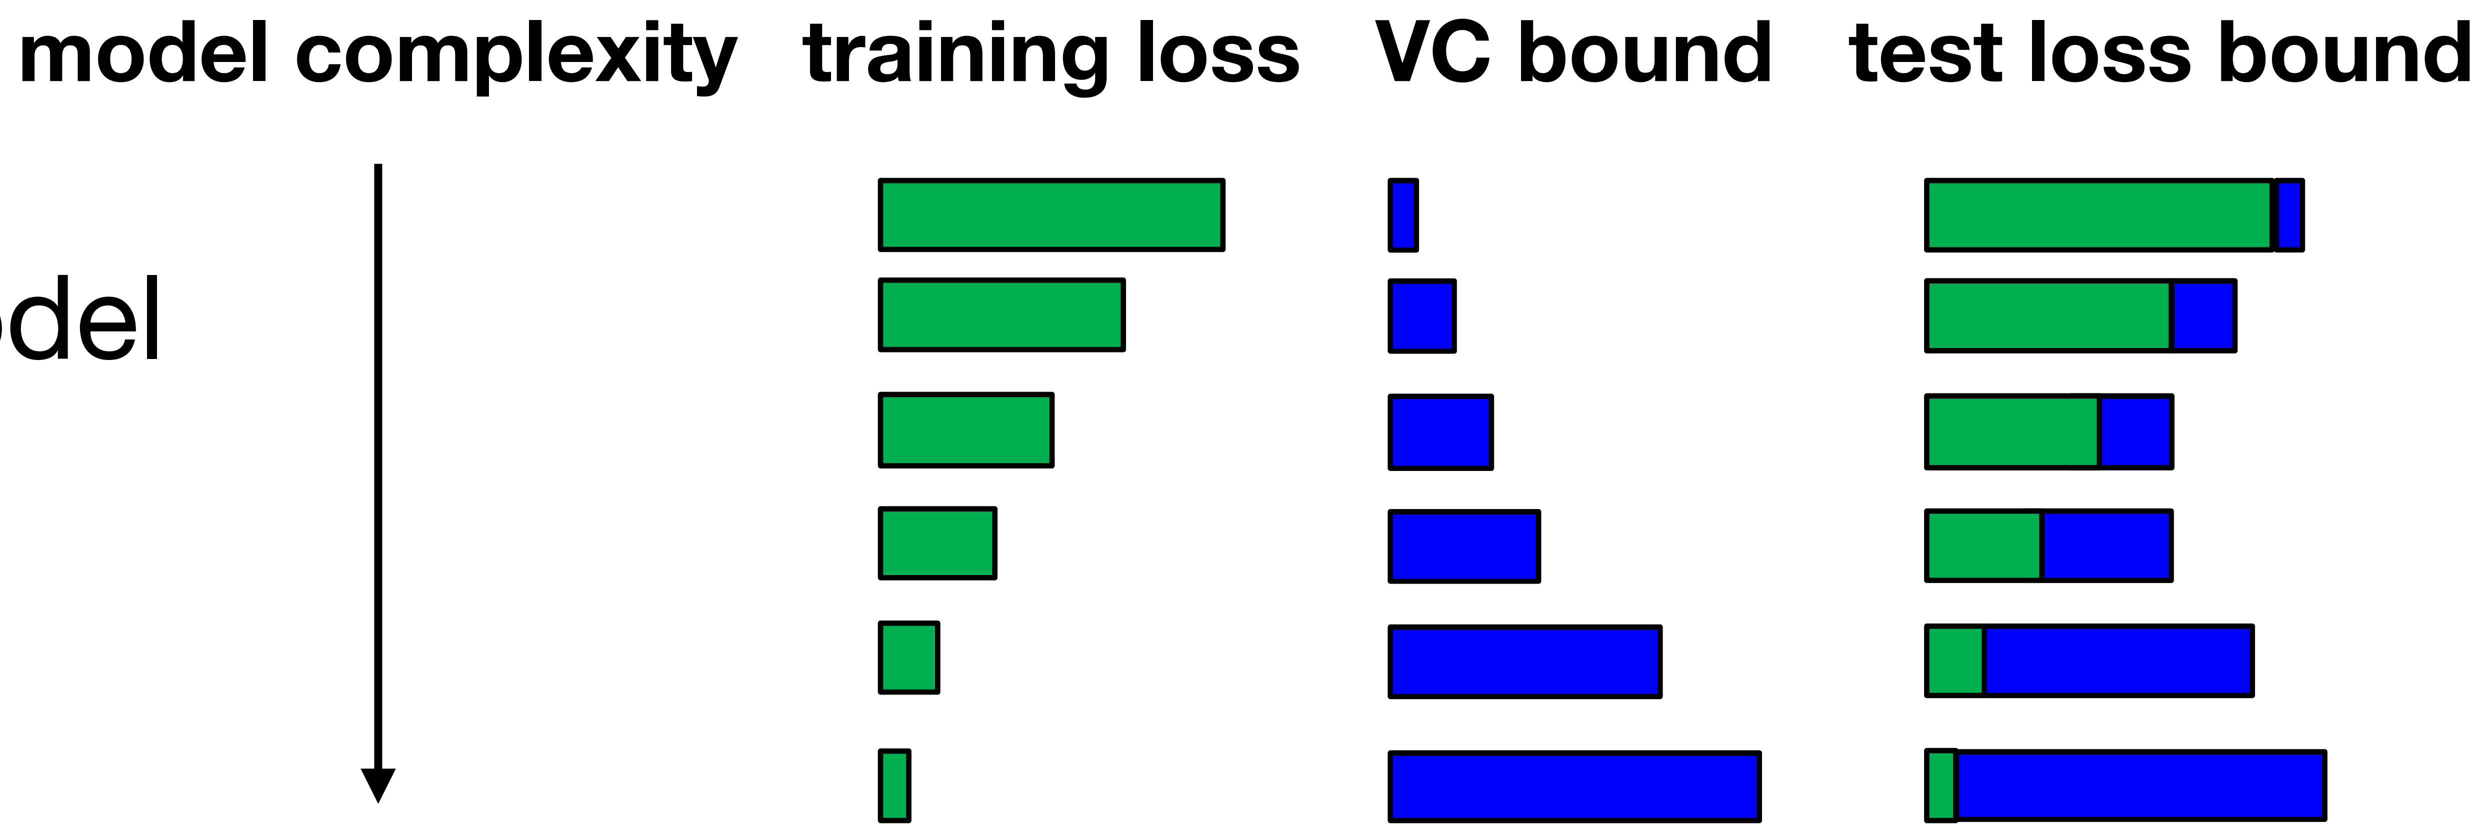

Logistics

project

- Team rosters due Monday, Feb 1 **on Canvas**
- Team-forming spreadsheet posted on piazza

midterm

- Midterm exam on Feb 9, 2–4pm **on Canvas**
- We'll accommodate other timezones — let us know