CPSC 340: Machine Learning and Data Mining

Regularization

Summer 2021

In This Lecture

- Regularization Intro (10 minutes)
- L2-regularization (10 minutes)
- L1-regularization (10 minutes)
- Standardization (10 minutes)

When you don't regularize your model

Coming Up Next

REGULARIZATION INTRO

(This is probably the MOST important topic in this course)

"Feature" Selection vs. "Model" Selection?

- Model selection: "which model should I use?"
  - KNN vs. decision tree, depth of decision tree, degree of polynomial basis.
- Feature selection: "which features should I use?"
  - Using feature 10 or not, using x_i^2 as part of basis.

$$y_i = \underbrace{w_0 + w_1 x_i} + 0 \cdot x_i^2 + 0 \cdot x_i^3 + 0 \cdot x_i^4$$

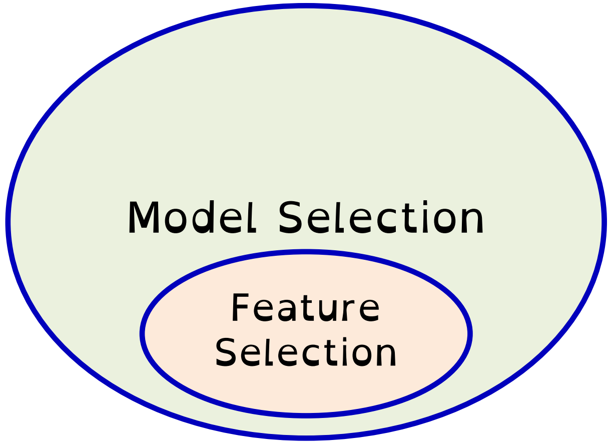

- These two tasks are highly-related:
  - It's a different "model" if we add x_i^2 to linear regression.
  - But the x_i^2 term is just a "feature" that could be "selected" or not.
  - Usually, "feature selection" means choosing from some "original" features.
    - You could say that "feature" selection is a special case of "model" selection.

Is It Good to Throw Features Away?

- (Yes/No), because linear regression can overfit with large 'd'.
  - Even though it's "just" a hyper-plane.
- Consider using d=n, with completely random features.
  - With high probability, you will be able to get a training error of 0.
  - But the features were random, this is completely overfitting.
- You could view "number of features" as a hyper-parameter.
  - Model gets more complex as you add more features.

$w_{d \times 1}$ gets larger

Controlling Complexity

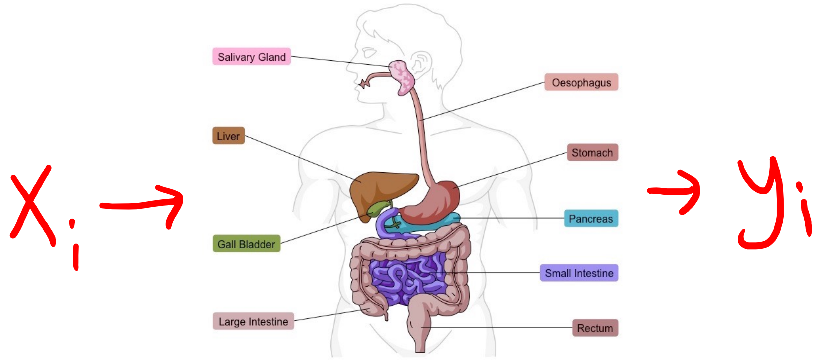

- Usually, "true" mapping from x_i to y_i is complex.
  - Might need high-degree polynomial.
  - Might need to combine many features, and don't know "relevant" ones.
- But complex models can overfit.
- So what do we do???

- Our main tools:
  - Model averaging: average over multiple models to decrease variance.
  - Regularization (today): add a penalty on the complexity of the model.

Would you rather?

- Consider the following dataset and 3 linear regression models:

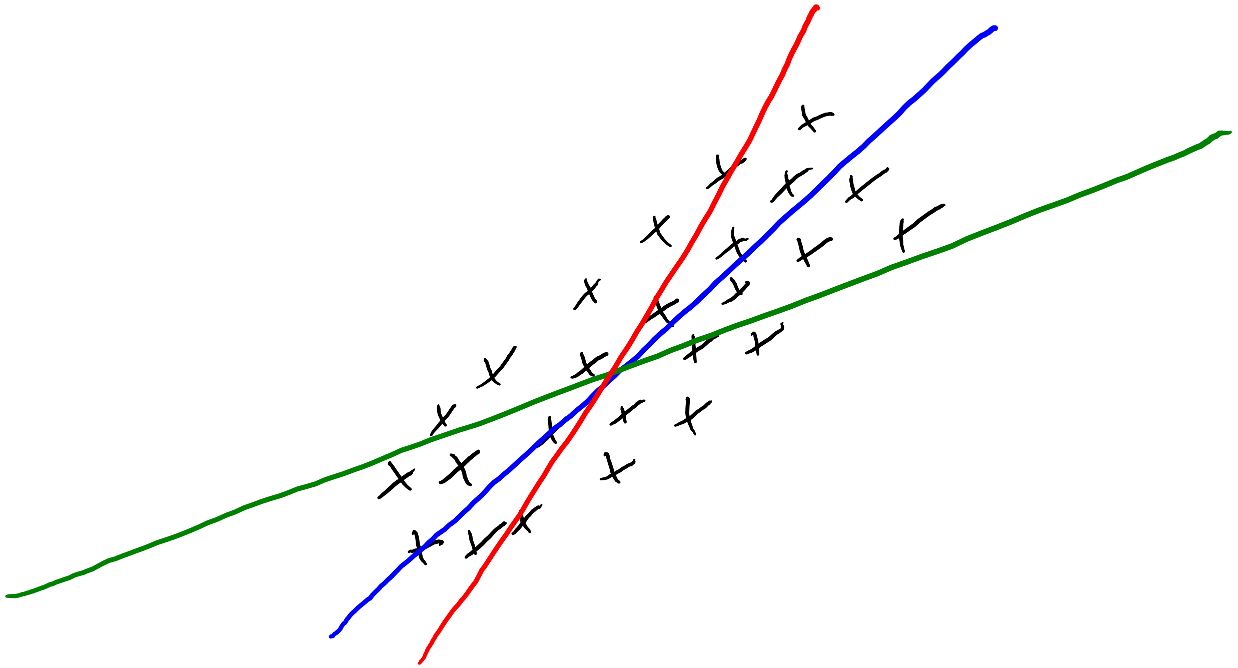

Q: Which one is the "best" model??

Would you rather?

- Consider the following dataset and 3 linear regression models:

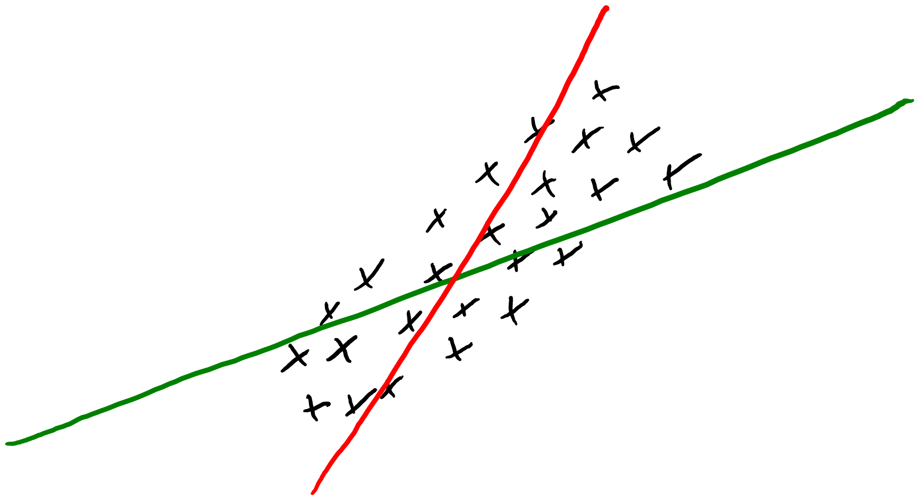

- What if you are forced to choose between red and green?
  - And assume they have the same training error.
- You should pick green.
  - Since slope is smaller, small change in x_i has a __Small change__ in prediction y_i.
    - Green line's predictions are (more/less) sensitive to having 'w' exactly right.
  - Since green 'w' is less sensitive to data, test error might be lower.

"Regularization"

- "Regularization": reducing a property of parameters
  - e.g. L2-norm of w, L1-norm of w, number of non-zeros in w, etc.
  - Optimization must take this term into account when minimizing

sparse reg.

- Assumption: we can express our goal as minimizing some quantity.
  - for linear models, small norm of w => low complexity

- Naive Bayes with Laplace smoothing: an instance of regularization
  - reduce heterogeneity of $p(x_{ij} \mid y_i)$ to control model complexity

$p(x_{ij} | y_i)$

Size of Regression Weights are Overfitting

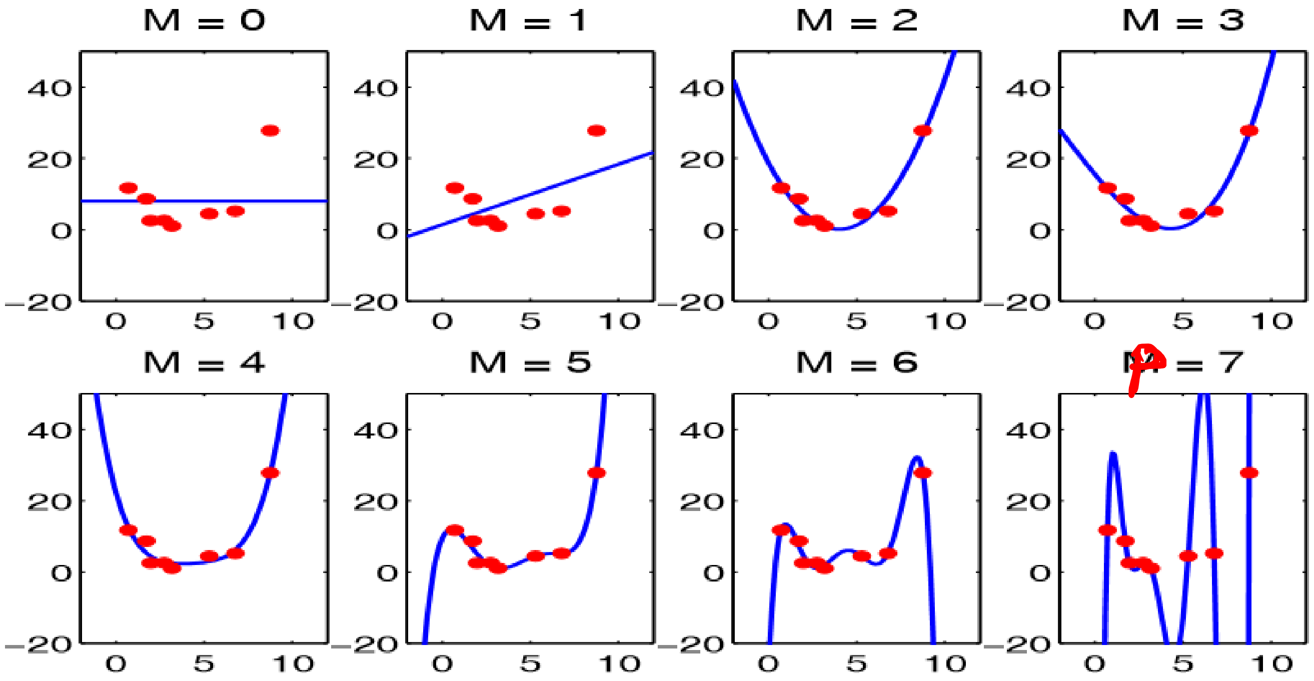

- The regression weights w_j with degree-7 are huge in this example.
- The degree-7 polynomial would be less sensitive to the data, if we "regularized" the w_j so that they are small.

$$\hat{y}_i = 0.0001(x_i)^7 + 0.03(x_i)^3 + 3 \quad \text{vs.} \quad \hat{y}_i = 1000(x_i)^7 - 500(x_i)^6 + 890x_i$$

Should you regularize your model?

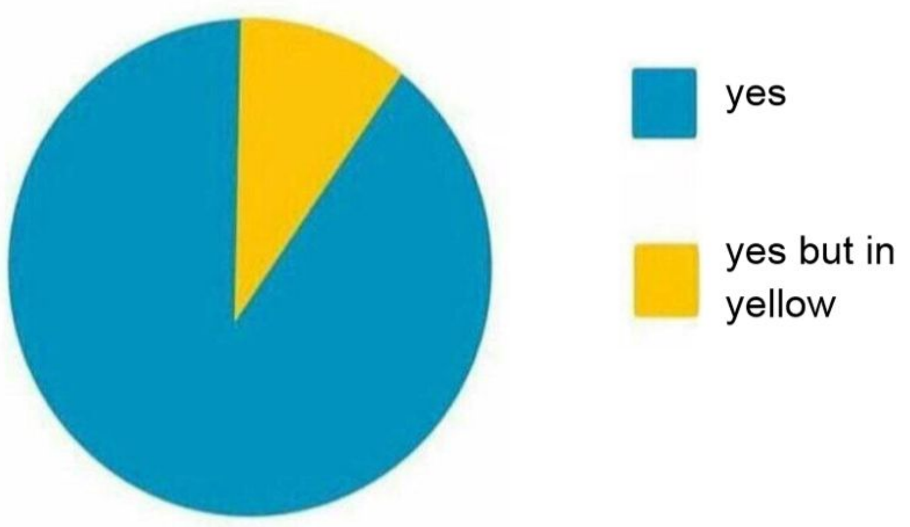

Coming Up Next

L2-REGULARIZATION

L2-Regularization

- Standard regularization strategy is L2-regularization:

$$f(w) = \frac{1}{2}\sum_{i=1}^{n}(w^Tx_i - y_i)^2 + \frac{\lambda}{2}\sum_{j=1}^{d} w_j^2 \quad \text{or} \quad f(w) = \frac{1}{2}\|Xw - y\|^2 + \frac{\lambda}{2}\|w\|^2$$

- Intuition: large slopes w_j tend to lead to overfitting.

- Objective balances getting low error vs. having small slopes 'w_j'.
  - "You can increase the training error if it makes 'w' much smaller."
  - Nearly-always reduces overfitting.

  - Regularization parameter $\lambda > 0$ controls strength of regularization.
    - Large λ puts large penalty on slopes.

Small $\lambda \rightarrow$ completely prioritize error.

L2-Regularization

- Standard regularization strategy is L2-regularization:

$$f(w) = \frac{1}{2}\sum_{i=1}^{n}(w^Tx_i - y_i)^2 + \frac{\lambda}{2}\sum_{j=1}^{d} w_j^2 \quad \text{or} \quad f(w) = \frac{1}{2}\|Xw - y\|^2 + \frac{\lambda}{2}\|w\|^2$$

$\in \mathbb{R}$

- In terms of fundamental trade-off:
  - Regularization (increases/decreases) training error.
  - Regularization (increases/decreases) approximation error.
- How should you choose λ?
  - Theory: as 'n' grows λ should be in the range O(1) to (√n).
  - Practice: optimize validation error or cross-validation error.
    - This almost always decreases the test error.

Q: Does this mean optimization bias is not a problem anymore?

NO

L2-Regularization "Shrinking" Example

- Solution to a "least squares with L2-regularization" for different λ:

λ	w_1	w_2	w_3	w_4	w_5	$\|\|Xw - y\|\|^2$	$\|\|w\|\|^2$
0	-1.88	1.29	-2.63	1.78	-0.63	285.64	15.68
1	-1.88	1.28	-2.62	1.78	-0.64	285.64	15.62
4	-1.87	1.28	-2.59	1.77	-0.66	285.64	15.43
16	-1.84	1.27	-2.50	1.73	-0.73	285.71	14.76
64	-1.74	1.23	-2.22	1.59	-0.90	286.47	12.77
256	-1.43	1.08	-1.70	1.18	-1.05	292.60	8.60
1024	-0.87	0.73	-1.03	0.57	-0.81	321.29	3.33
4096	-0.35	0.31	-0.42	0.18	-0.36	374.27	0.56

- We get least squares with $\lambda = 0$.
  - But we can achieve similar training error with smaller $||w||$.
- $||Xw - y||$ increases with λ, and $||w||$ decreases with λ.
  - Though individual w_j can increase or decrease with lambda.
  - Because we use the L2-norm, the large ones decrease the most.

Regularization Path

- Regularization path is a plot of the optimal weights 'w_j' as 'λ' varies:

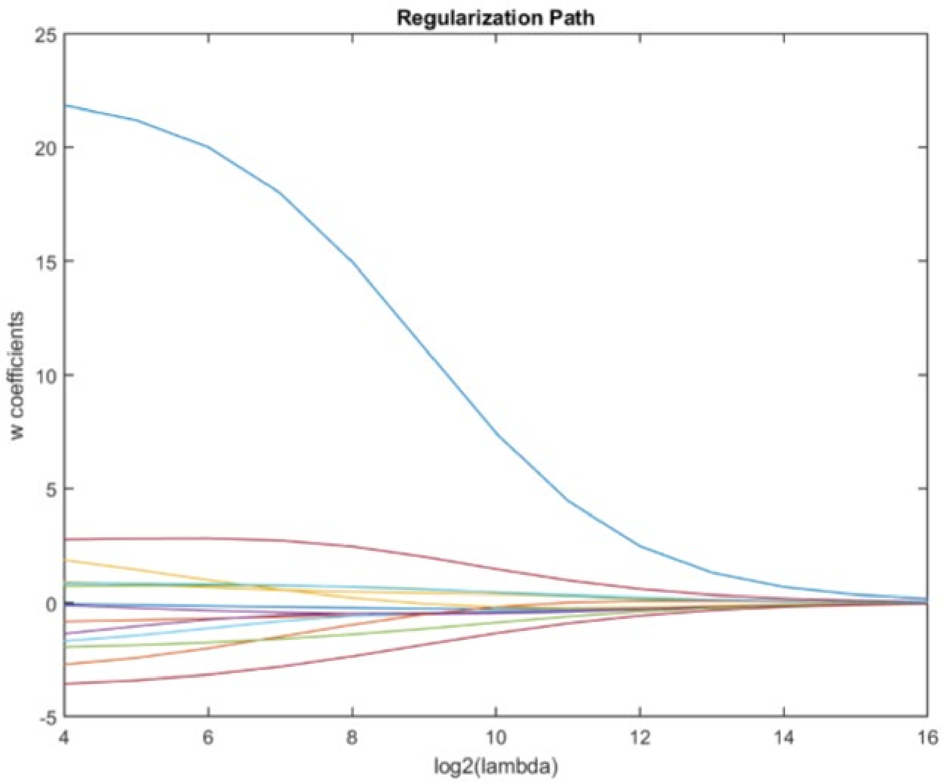

- Starts with least squares with $\lambda = 0$, and w_j converge to 0 as λ grows.

L2-regularized Least Squares Normal Equations

- When using L2-regularized squared error, we can solve for ∇ f(w) = 0.
- Loss before: $f(w) = \frac{1}{2}\|Xw - y\|^2$
- Loss after: $f(w) = \frac{1}{2}\|Xw - y\|^2 + \frac{\lambda}{2}\|w\|^2$
- Gradient before: $\nabla f(w) = X^T X w - X^T y$
- Gradient after: $\nabla f(w) = X^T X w - X^T y + \lambda w$ $= (X^T X + \lambda I)w - X^T y$ ($\lambda I \cdot w$, $d \times d$)
- Linear system before: $X^T X w = X^T y$
- Linear system after: $(X^T X + \lambda I)w = X^T y$ (PD)
- But unlike $X^T X$, the matrix $(X^T X + \lambda I)$ is always invertible:
  - Multiply by its inverse for unique solution: $w = (X^T X + \lambda I)^{-1}(X^T y)$

Gradient Descent for L2-Regularized Least Squares

- The L2-regularized least squares objective and gradient:

$$f(w) = \frac{1}{2}\|Xw - y\|^2 + \frac{\lambda}{2}\|w\|^2 \qquad \nabla f(w) = X^T(Xw - y) + \lambda w$$

- Gradient descent iterations for L2-regularized least squares:

$$w^{t+1} = w^t - \alpha^t \underbrace{\left[X^T(Xw^t - y) + \lambda w^t\right]}_{\nabla f(w^t)}$$

- Cost of gradient descent iteration is still O(nd).
  - Can show number of iterations decrease as λ increases (not obvious).

Why use L2-Regularization?

- It's a weird thing to do, but Mark says "always use regularization".
  - "Almost always decreases test error" should already convince you.

- But here are 6 more reasons:
  1. Solution 'w' is unique.
  2. X^TX does not need to be invertible (no collinearity issues).
  3. Less sensitive to changes in X or y.
  4. Gradient descent converges faster (bigger λ means fewer iterations).
  5. Stein's paradox: if $d \geq 3$, 'shrinking' moves us closer to 'true' w.
  6. Worst case: just set λ small and get the same performance.

Regularizing the y-Intercept?

- Should we regularize the y-intercept? $y_i = w_0 + w_1 x_i$
- No! Why encourage it to be closer to zero? (It could be anywhere.)
  - You should be allowed to shift function up/down globally.
- Yes! It makes the solution unique and it easier to compute 'w'.
- Compromise: regularize by a smaller amount than other features.

$$f(w, w_0) = \frac{1}{2}\|Xw + w_0 - y\|^2 + \frac{\lambda}{2}\|w\|^2 + \frac{\lambda_0}{2} w_0^2$$

Xw

Coming Up Next

L1-REGULARIZATION

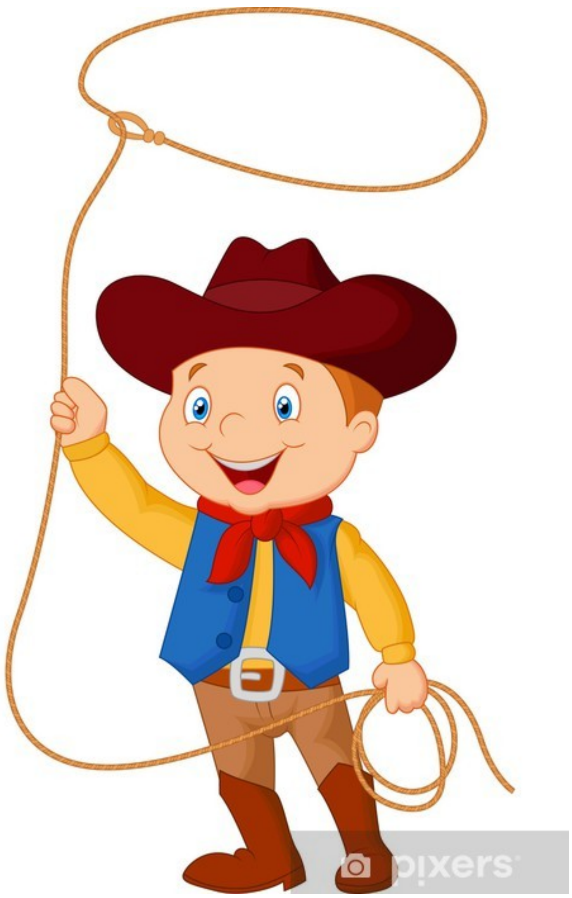

Lasso, pronounced "la sue"

Previously: Search and Score

- We talked about search and score for feature selection:
  - Define a "score" and "search" for features with the best score.
- Usual scores count the number of non-zeroes ("L0-norm"):

$$f(w) = \frac{1}{2}\|Xw - y\|^2 + \lambda \|w\|_0$$

number of non-zeroes in 'w'

- But it's hard to find the 'w' minimizing this objective.
- We discussed forward selection, but requires fitting O(d^2) models.

$$d + (d-1) + (d-2) + (d-3) \cdots = \binom{d}{2} = O(d^2)$$

Previously: Search and Score

- What if we want to pick among millions or billions of features?

- If 'd' is large, forward selection is too slow (A4):
  - For least squares, need to fit $O(d^2)$ models at cost of $O(nd^2 + d^3)$.
  - Total cost $O(nd^4 + d^5)$.

- The situation is worse if we aren't using basic least squares:
  - For robust regression, need to run gradient descent $O(d^2)$ times.
  - With regularization, need to search for lambda $O(d^2)$ times.

L1-Regularization

- Instead of L0- or L2-norm, consider regularizing by the L1-norm:

$$\|w\|_1 = \sum_{j=1}^{d} |w_j|$$

$$f(w) = \frac{1}{2}\|Xw - y\|^2 + \lambda \|w\|_1$$

- Like L2-norm, it's convex and improves our test error.
- Like L0-norm, it encourages elements of 'w' to be exactly zero.

- L1-regularization simultaneously regularizes and selects features.
  - Very fast alternative to search and score.
  - Sometimes called "LASSO" regularization.
    - least absolute shrinkage and selection operator

L2-Regularization vs. L1-Regularization

- Regularization path of w_j values as 'λ' varies:

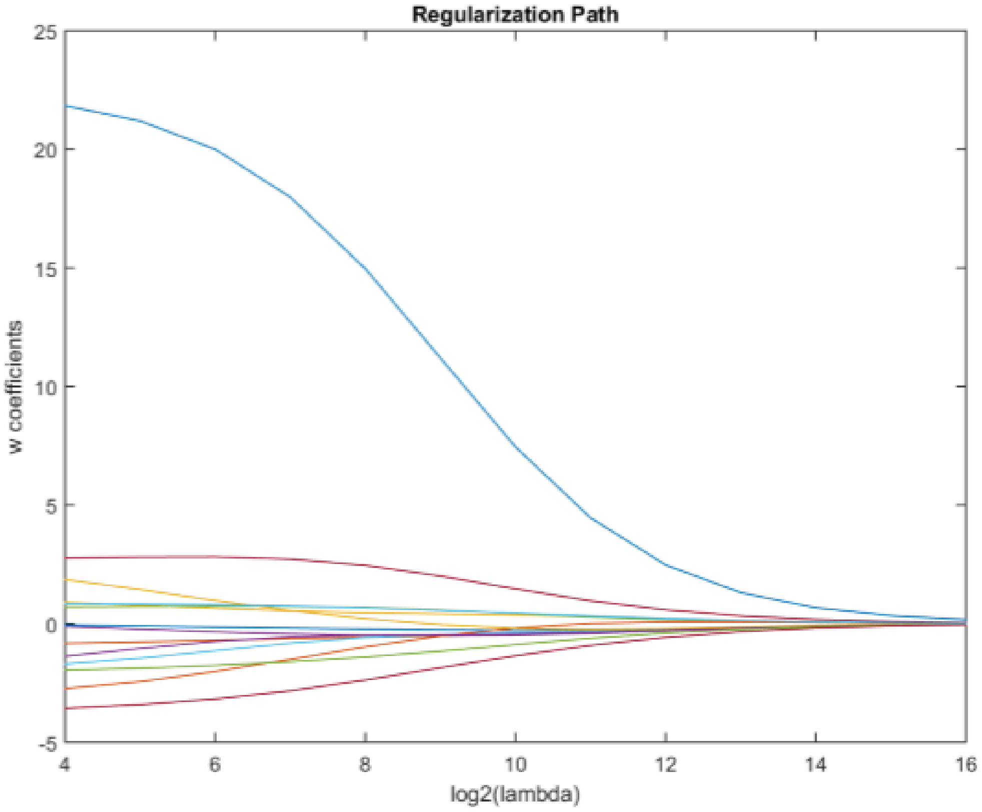

L2-regularization

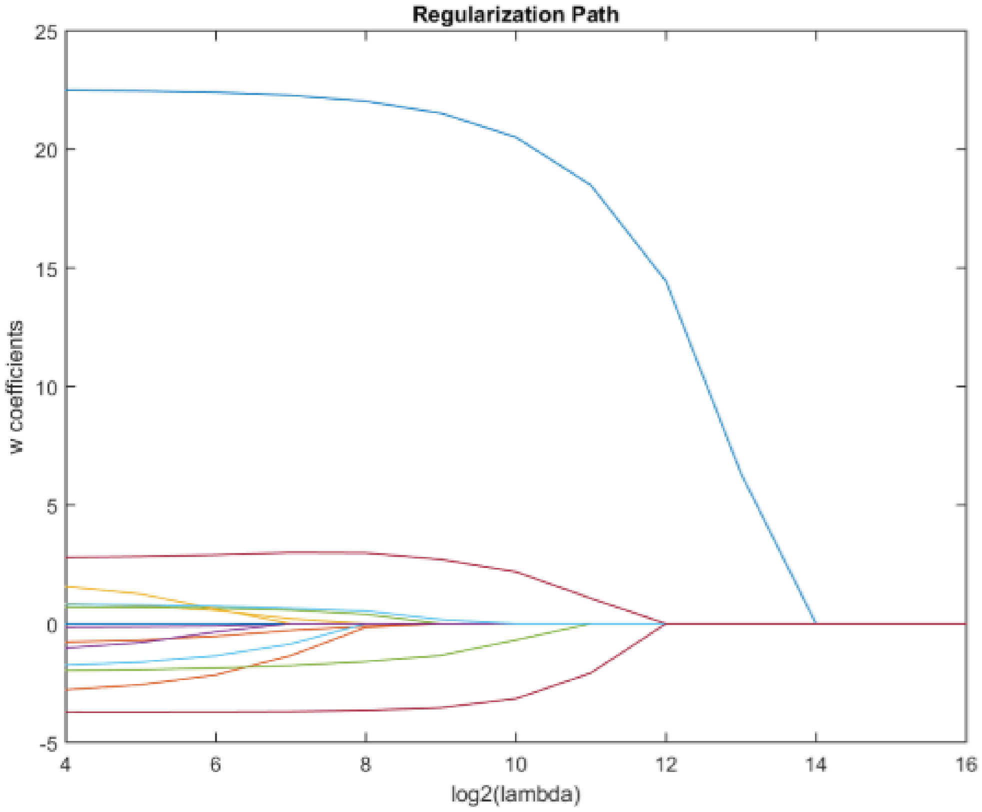

L1-regularization

- L1-Regularization sets values to exactly 0 (WHY?)

Regularizers and Sparsity

- L1-regularization gives sparsity but L2-regularization doesn't.
  - But don't they both shrink features towards zero?

- What is the penalty for setting $w_j = 0.00001$?

- L0-regularization: penalty of λ.
  - A constant penalty for any non-zero value.
  - Encourages you to set w_j exactly to zero, but otherwise doesn't care if w_j is small or not.

- L2-regularization: penalty of $(\lambda/2)(0.00001)^2 = 0.0000000005\lambda$.
  - The penalty gets smaller as you get closer to zero.
  - The penalty asymptotically vanishes as w_j approaches 0 (no incentive for "exact" zeroes).

- L1-regularization: penalty of $\lambda|0.00001| = 0.00001\lambda$.
  - The penalty stays is proportional to how far away w_j is from zero.
  - There is still something to be gained from making a tiny value exactly equal to 0.

L2-Regularization vs. L1-Regularization

- L2-Regularization:
  - Insensitive to changes in data.
  - Decreased variance:
    - Lower test error.
  - Closed-form solution.
  - Solution is unique.
  - All 'w_j' tend to be non-zero.
  - Can learn with *linear* number of irrelevant features.
    - E.g., only O(d) relevant features.

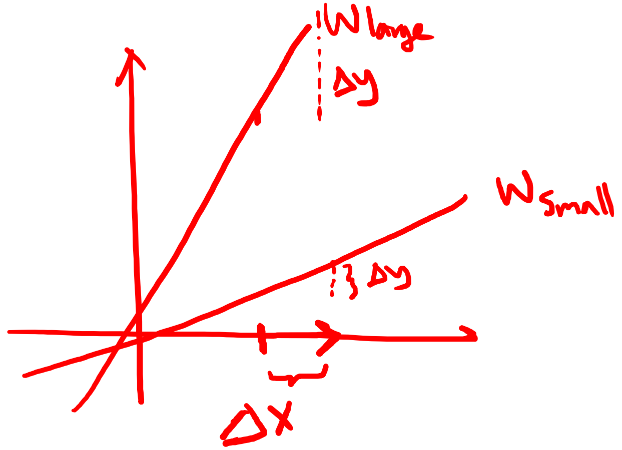

- L1-Regularization:
  - Insensitive to changes in data.
  - Decreased variance:
    - Lower test error.
  - Requires iterative solver.
  - Solution is not unique.
  - Many 'w_j' tend to be zero.
  - Can learn with **exponential** number of irrelevant features.
    - E.g., only O(log(d)) relevant features.

Paper on this result by Andrew Ng

L1-Regularization Applications

- Used to give super-resolution in imaging black holes.
  - Sparsity arises in a particular basis.

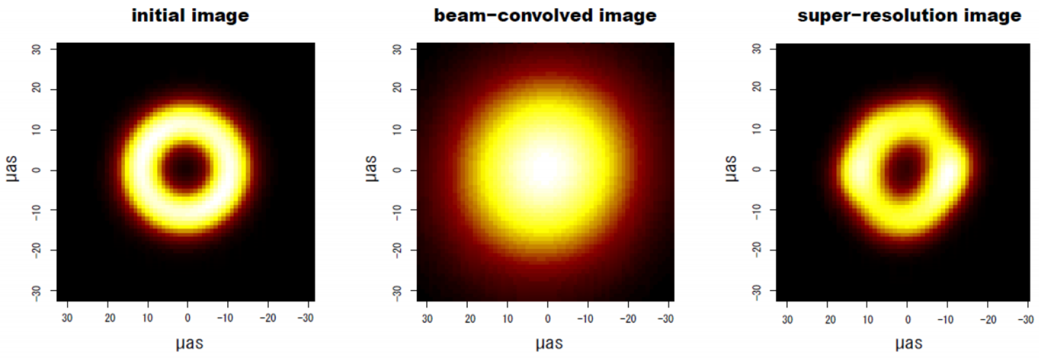

Figure 2. Simulated images of M87. From left to right, the initial model, the image with 0-filling, and the image with LASSO. Improvement of resolution in the LASSO image is significant.

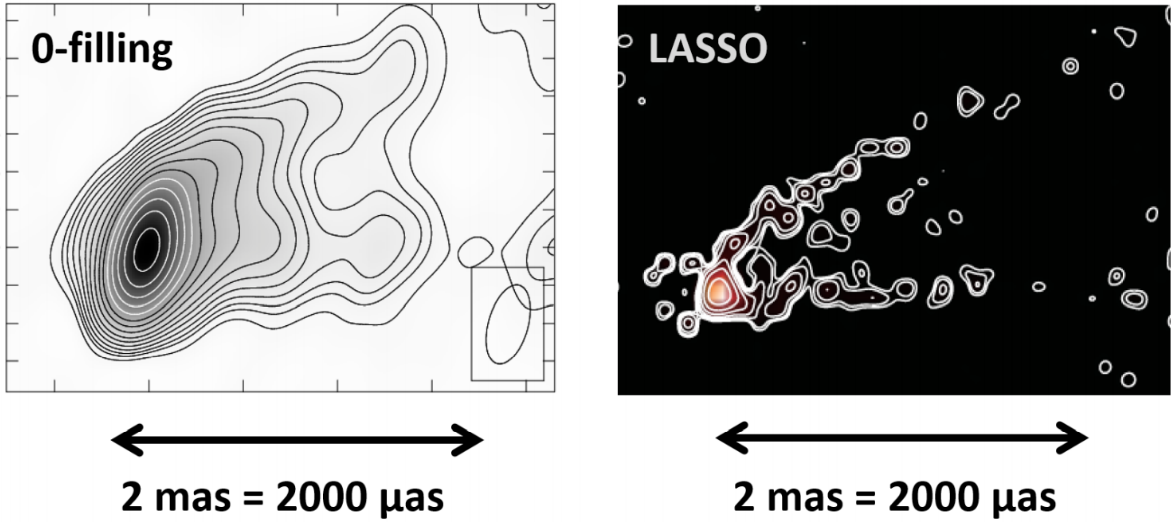

Figure 3. Standard and LASSO images of M87 observed with VLBA at a wavelength of 7 mm. In the two plots, exactly the same data are used. The angular resolution is better in the LASSO image, and the detailed structure of the M87 jet can be traced in more detail.

https://iopscience.iop.org/article/10.1088/1742-6596/699/1/012006/pdf

L1-loss vs. L1-regularization

- Don't confuse the L1 loss with L1-regularization!
  - L1-loss is robust to outlier data points.
    - You can use this instead of removing outliers.
  - L1-regularization is robust to irrelevant features.
    - You can use this instead of removing features.
- And note that you can be robust to outliers and irrelevant features:

$$f(w) = \underbrace{\|Xw - y\|_1}_{L_1\text{-loss}} + \underbrace{\lambda \|w\|_1}_{L_1\text{-regularizer}}$$

- Can we smooth and use "Huber regularization"?
  - Huber regularizer is still robust to irrelevant features.
  - But it's the non-smoothness that sets weights to exactly 0.

L*-Regularization

- L0-regularization (AIC, BIC, Mallow's Cp, Adjusted R^2, ANOVA):
  - Adds penalty on the number of non-zeros to select features.

$$f(w) = \|Xw - y\|^2 + \lambda \|w\|_0$$

- L2-regularization (ridge regression):
  - Adding penalty on the L2-norm of 'w' to decrease overfitting:

$$f(w) = \|Xw - y\|^2 + \frac{\lambda}{2} \|w\|^2$$

- L1-regularization (LASSO):
  - Adding penalty on the L1-norm decreases overfitting and selects features:

$$f(w) = \|Xw - y\|^2 + \lambda \|w\|_1$$

L0- vs. L1- vs. L2-Regularization

	Sparse 'w' (Selects Features)	Speed	Unique 'w'	Coding Effort	Irrelevant Features
L0-Regularization	Yes	Slow	No	Few lines	Not Sensitive
L1-Regularization	Yes*	Fast*	No	1 line*	Not Sensitive
L2-Regularization	No	Fast	Yes	1 line	A bit sensitive

- L1-Regularization isn't as sparse as L0-regularization.
  - L1-regularization tends to give more false positives (selects too many).
  - And it's only "fast" and "1 line" with specialized solvers (optimizers.py).
- Cost of L2-regularized least squares is $O(nd^2 + d^3)$.
  - Changes to O(ndt) for 't' iterations of gradient descent (same for L1).
- "Elastic net" (L1- and L2-regularization) is sparse, fast, and unique.
- Using L0+L2 does not give a unique solution.

Ensemble Feature Selection

- We can also use ensemble methods for feature selection.
  - Usually designed to reduce false positives or reduce false negatives.
    - False positive: irrelevant feature is selected
    - False negative: relevant feature is excluded

- In this case of L1-regularization, we want to reduce false positives.
  - Unlike L0-regularization, continuous tension between performance/selection
    - "Irrelevant" features can be included before "relevant" w_j reach best value.

- A bootstrap approach to reducing false positives:
  - Apply the method to bootstrap samples of the training data.
  - Only take the features selected in all bootstrap samples.

Ensemble Feature Selection

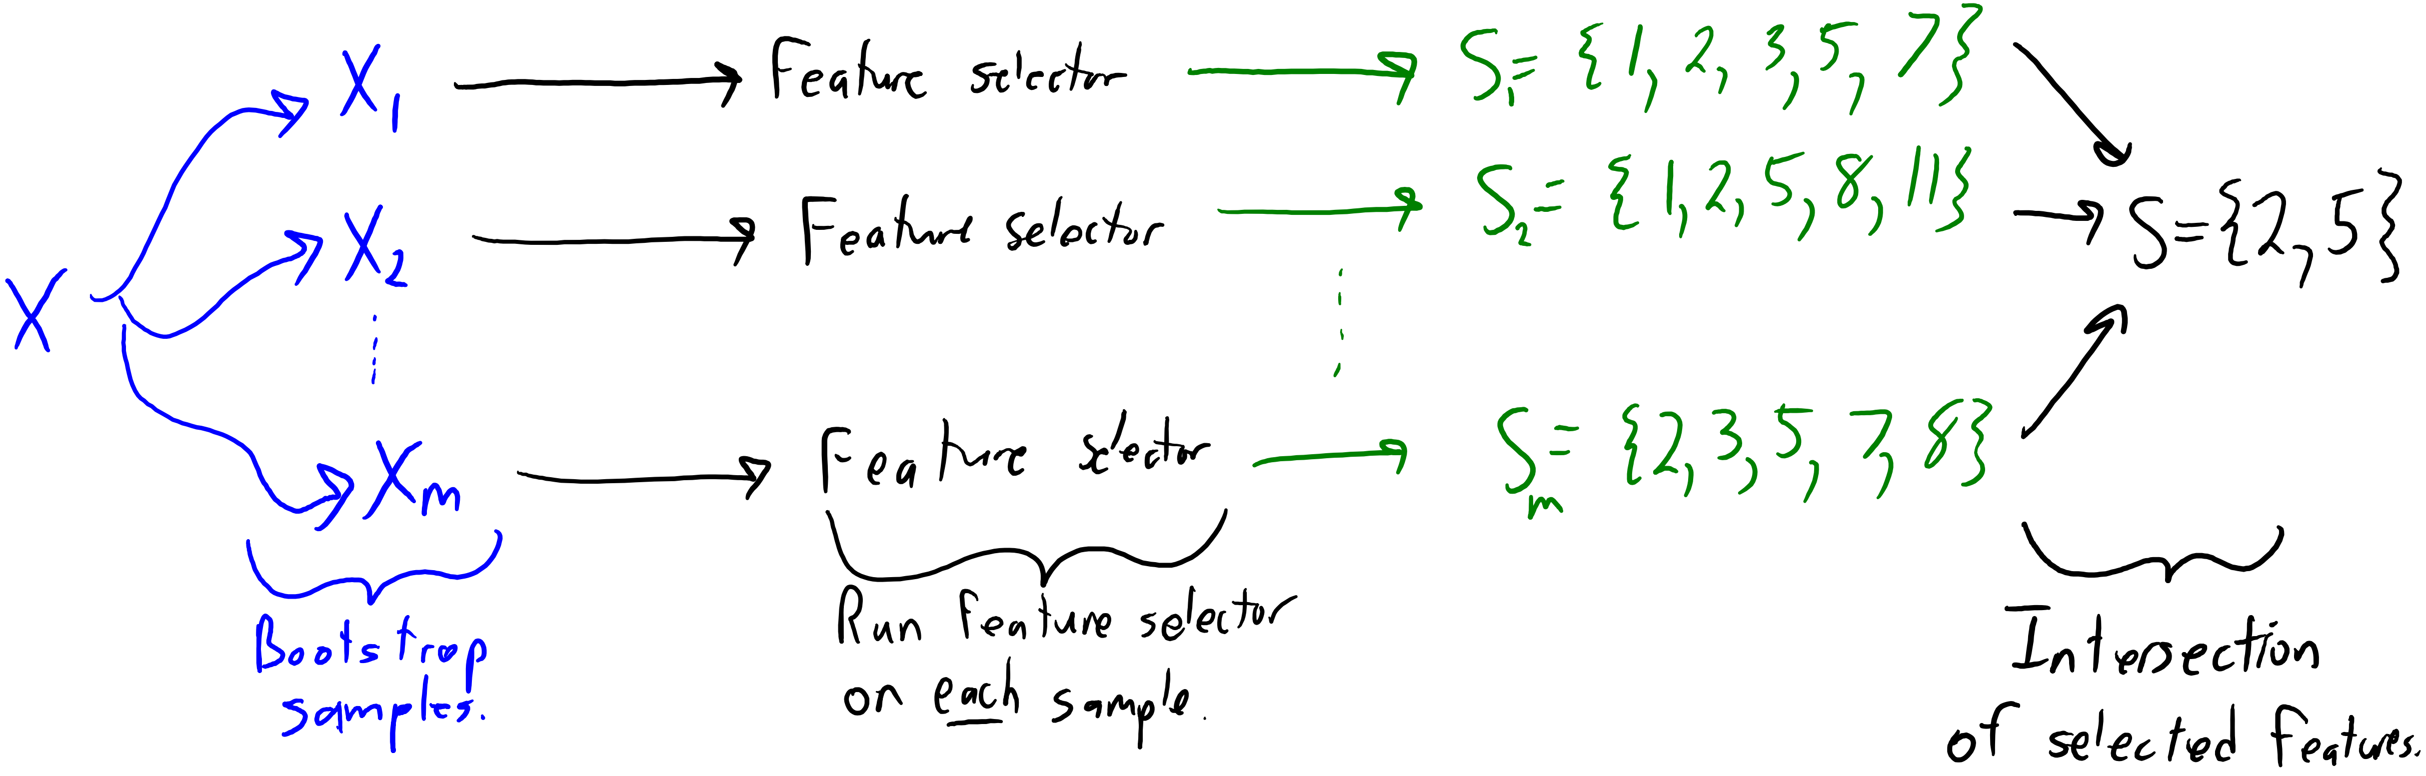

- Example: bootstrapping plus L1-regularization ("BoLASSO").
  - Reduces false positives.
  - It's possible to show it recovers "correct" features with weaker conditions.
    - Can replace "intersection" with "selected frequency" if has false negatives too.

Summary

- Regularization:
  - Adding a penalty on model complexity.
- L2-regularization: penalty on L2-norm of regression weights ‘w’.
  - Almost always improves test error.
- L1-regularization: penalty on L1-norm of regression weights ‘w’.
  - Simultaneous regularization and feature selection.
  - Robust to having lots of irrelevant features.
- ~~Feature standardization:~~
  - ~~Change the unit of every feature into “z-score”~~

- Next time: non-parametric feature transform and linear classifiers

Review Questions

- ~~Q1: In what ways can standardizing the features help reduce a linear model's complexity?~~
- Q2: Why is L1-regularization able to perform feature selection while L2-regularization cannot?
- Q3: Why are we allowed to use $(X^{\top}X + \lambda I)^{-1}$ in the solution to L2-regularized least squares?
- Q4: What happens to colinear features when L1-regularization is used?
- ~~Q5: What parameters are we "learning" for standardization?~~

L2-Regularization

- Standard regularization strategy is L2-regularization:

$$f(w) = \frac{1}{2}\sum_{i=1}^{n}(w^Tx_i - y_i)^2 + \frac{\lambda}{2}\sum_{j=1}^{d} w_j^2 \quad \text{or} \quad f(w) = \frac{1}{2}\|Xw - y\|^2 + \frac{\lambda}{2}\|w\|^2$$

- Equivalent to minimizing squared error but keeping L2-norm small.

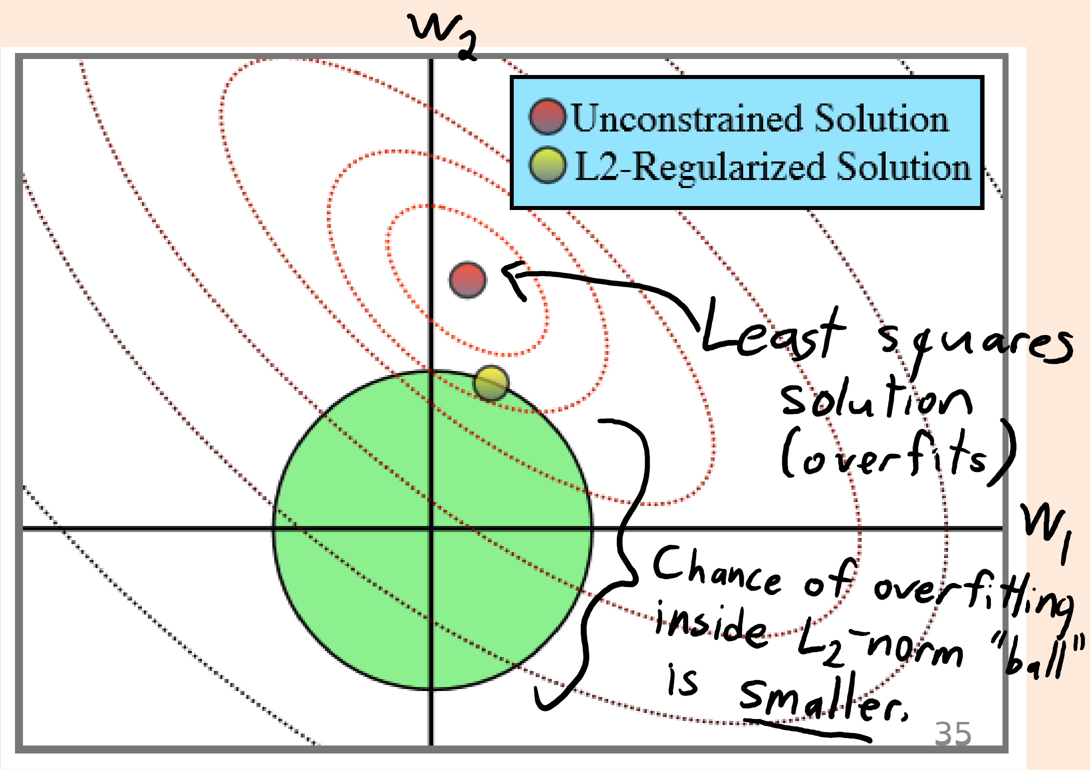

Regularization/Shrinking Paradox

- We throw darts at a target:
  - Assume we don't always hit the exact center.
  - Assume the darts follow a symmetric pattern around center.

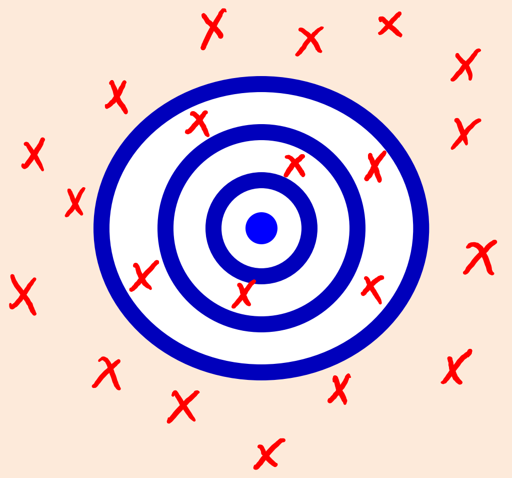

Regularization/Shrinking Paradox

- We throw darts at a target:
  - Assume we don't always hit the exact center.
  - Assume the darts follow a symmetric pattern around center.
- Shrinkage of the darts :
  1. Choose some arbitrary location '0'.
  2. Measure distances from darts to '0'.

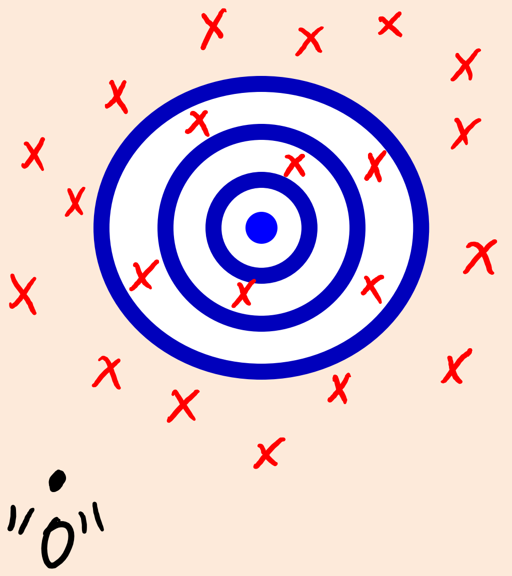

Regularization/Shrinking Paradox

- We throw darts at a target:
  - Assume we don't always hit the exact center.
  - Assume the darts follow a symmetric pattern around center.
- Shrinkage of the darts :
  1. Choose some arbitrary location '0'.
  2. Measure distances from darts to '0'.
  3. Move misses towards '0', by *small* amount proportional to distance from 0.
- If small enough, darts will be closer to center on average.

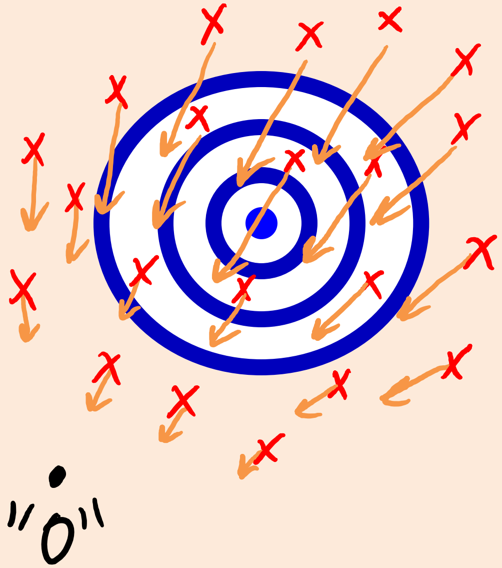

Regularization/Shrinking Paradox

- We throw darts at a target:
  - Assume we don't always hit the exact center.
  - Assume the darts follow a symmetric pattern around center.
- Shrinkage of the darts :
  1. Choose some arbitrary location '0'.
  2. Measure distances from darts to '0'.
  3. Move misses towards '0', by *small* amount proportional to distance from 0.
- If small enough, darts will be closer to center on average.

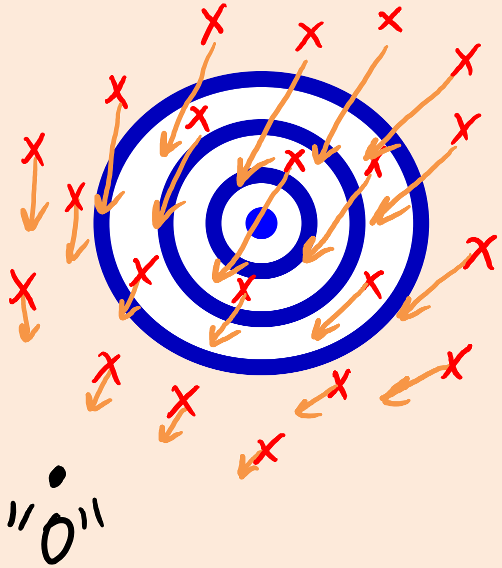

Visualization of the related higher-dimensional paradox that the mean of data coming from a Gaussian is not the best estimate of the mean of the Gaussian in 3-dimensions or higher: https://www.naftaliharris.com/blog/steinviz

Regularizers and Sparsity

- L1-regularization gives sparsity but L2-regularization doesn't.
  - But don't they both shrink variables to zero?
- Consider problem where 3 vectors can get minimum training error:

$$w^1 = \begin{bmatrix} 100 \\ 0.02 \end{bmatrix} \qquad w^2 = \begin{bmatrix} 100 \\ 0 \end{bmatrix} \qquad w^3 = \begin{bmatrix} 99.99 \\ 0.02 \end{bmatrix}$$

- Without regularization, we could choose any of these 3.
  - They all have same error, so regularization will "break tie".
- With L0-regularization, we would choose w^2:

$$\|w^1\|_0 = 2 \qquad \|w^2\|_0 = 1 \qquad \|w^3\|_0 = 2$$

Regularizers and Sparsity

- L1-regularization gives sparsity but L2-regularization doesn't.
  - But don't they both shrink variables to zero?
- Consider problem where 3 vectors can get minimum training error:

- With L2-regularization, we would choose w^3:

$$w^1 = \begin{bmatrix} 100 \\ 0.02 \end{bmatrix} \quad w^2 = \begin{bmatrix} 100 \\ 0 \end{bmatrix} \quad w^3 = \begin{bmatrix} 99.99 \\ 0.02 \end{bmatrix}$$

- L2-regularization focuses on decreasing largest (makes w_j similar).

$$\|w^1\|^2 = 100^2 + 0.02^2 = 10000.0004 \quad \|w^2\|^2 = 100^2 + 0^2 = 10000 \quad \|w^3\|^2 = 99.99^2 + 0.02^2 = 9998.0005$$

Regularizers and Sparsity

- L1-regularization gives sparsity but L2-regularization doesn't.
  - But don't they both shrink variables to zero?
- Consider problem where 3 vectors can get minimum training error:

- With L1-regularization, we would choose w^2:

$$w^1 = \begin{bmatrix} 100 \\ 0.02 \end{bmatrix} \quad w^2 = \begin{bmatrix} 100 \\ 0 \end{bmatrix} \quad w^3 = \begin{bmatrix} 99.99 \\ 0.02 \end{bmatrix}$$

- L1-regularization focuses on decreasing all w_j until they are 0.

$$\|w^1\|_1 = 100 + 0.02 = 100.02 \quad \|w^2\|_1 = 100 + 0 = 100 \quad \|w^3\|_1 = 99.99 + 0.02 = 100.01$$

Sparsity and Least Squares

- Consider 1D least squares objective:

$$f(w) = \frac{1}{2}\sum_{i=1}^{n}(w x_i - y_i)^2$$

- This is a convex 1D quadratic function of ‘w’ (i.e., a parabola):

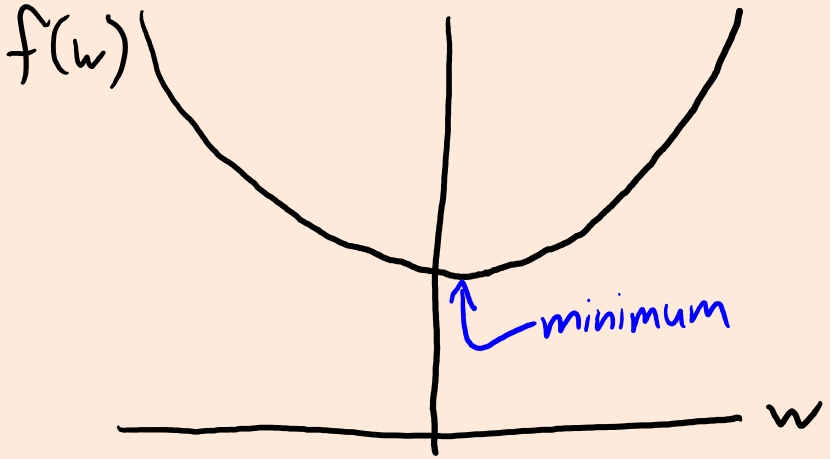

- This variable does not look relevant (minimum is close to 0).
  - But for finite ‘n’ the minimum is unlikely to be exactly zero.

$f'(0) = 0$ only happens if $\sum_{i=1}^{n} y_i x_i = 0$. (bonus)

Sparsity and L0-Regularization

- Consider 1D L0-regularized least squares objective:

$$f(w) = \frac{1}{2}\sum_{i=1}^{n}(w x_i - y_i)^2 + \lambda \|w\|_0$$

$$\lambda \|w\|_0 = \begin{cases} \lambda & \text{if } w \neq 0 \\ 0 & \text{if } w = 0 \end{cases}$$

- This is a convex 1D quadratic function but with a discontinuity at 0:

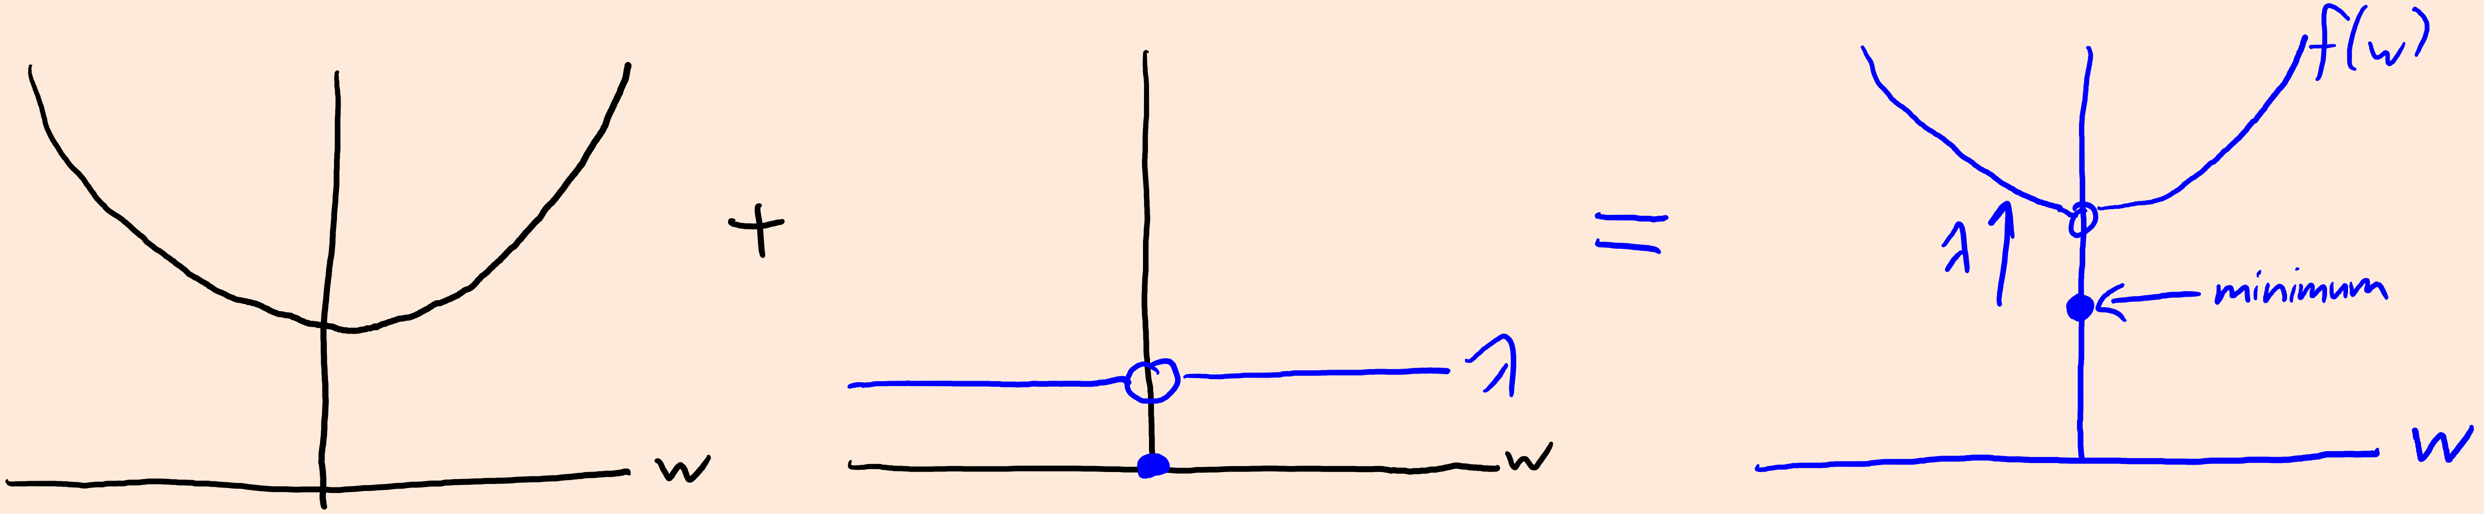

- L0-regularized minimum is often exactly at the ‘discontinuity’ at 0:
  - Sets the feature to exactly 0 (does feature selection), but is non-convex.

Sparsity and L2-Regularization

- Consider 1D L2-regularized least squares objective:

$$f(w) = \frac{1}{2}\sum_{i=1}^{n}(w x_i - y_i)^2 + \frac{\lambda}{2}w^2$$

- This is a convex 1D quadratic function of 'w' (i.e., a parabola):

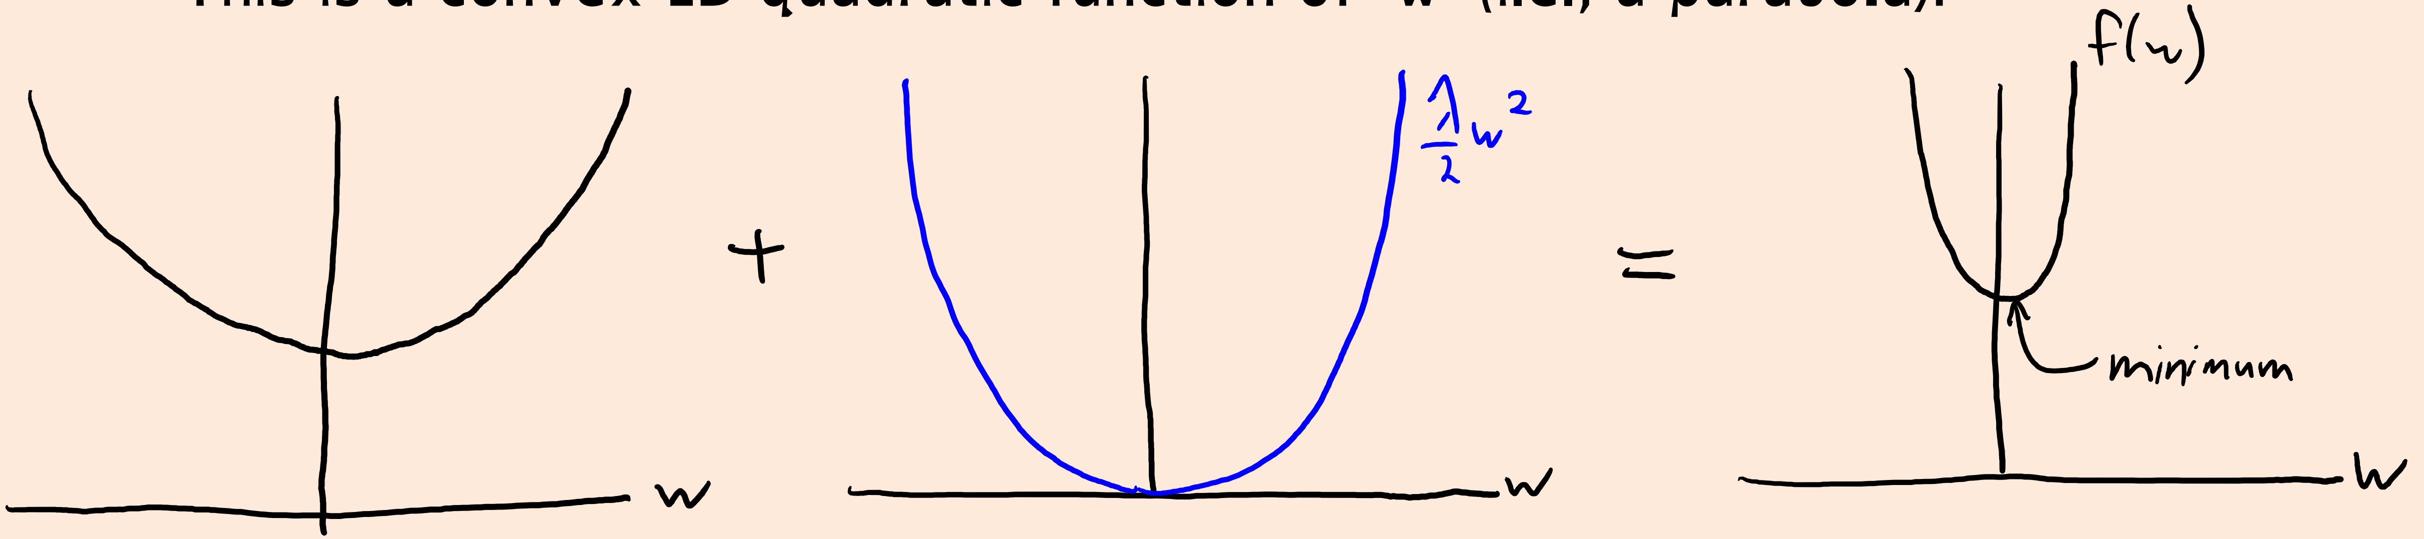

- L2-regularization moves it closer to zero, but not all the way to zero.
  - It doesn't do feature selection ("penalty goes to 0 as slope goes to 0").

$f'(0) = 0$ only if $\sum_{i=1}^{n} y_i x_i = 0$

Sparsity and L1-Regularization

- Consider 1D L1-regularized least squares objective:

$$f(w) = \frac{1}{2}\sum_{i=1}^{n}(w x_i - y_i)^2 + \lambda|w|$$

- This is a convex piecewise-quadratic function of 'w' with 'kink' at 0:

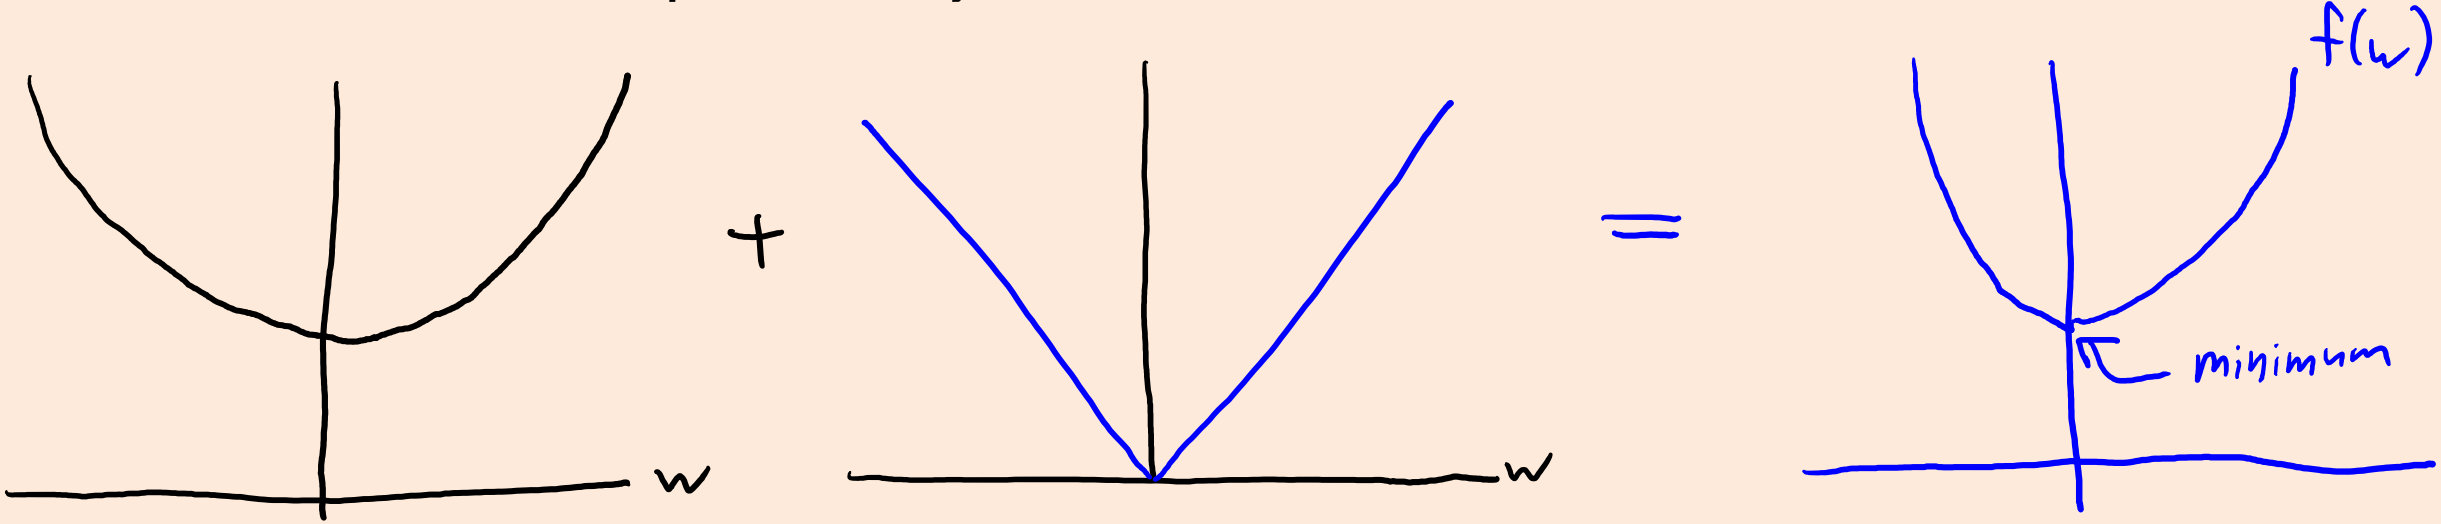

- L1-regularization tends to set variables to exactly 0 (feature selection).
  - Penalty on slope is λ even if you are close to zero.
  - Big λ selects few features, small λ allows many features.

→ Happens when $|\sum_{i=1}^{n} x_i y_i| \leq \lambda$ (bonus)

Sparsity and Regularization (with d=1)

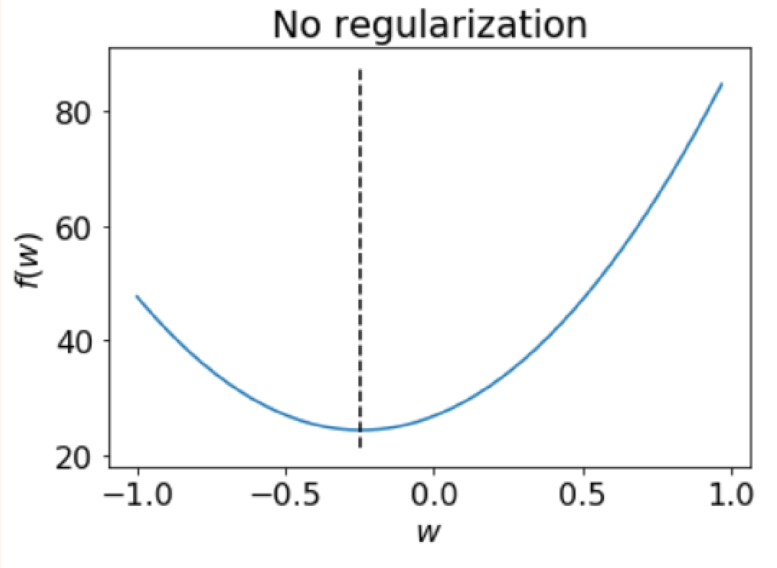

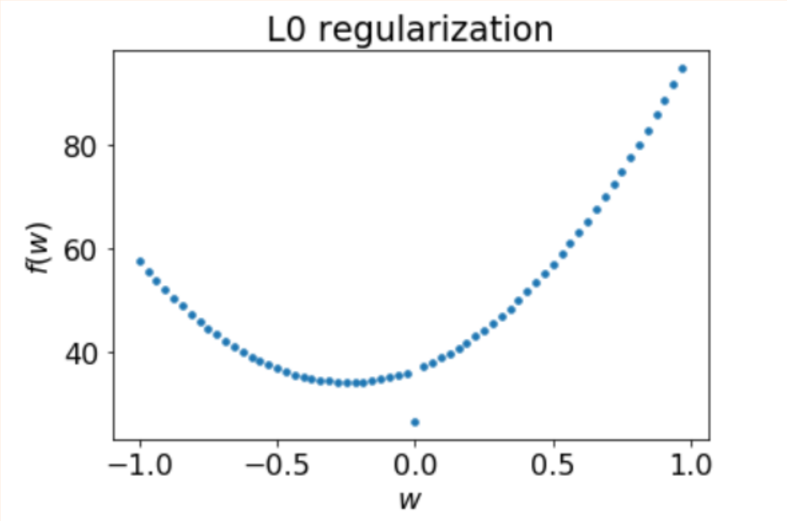

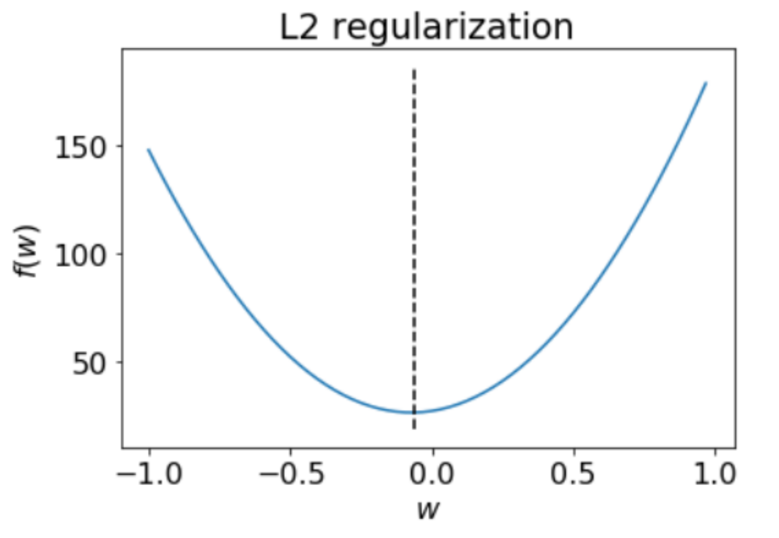

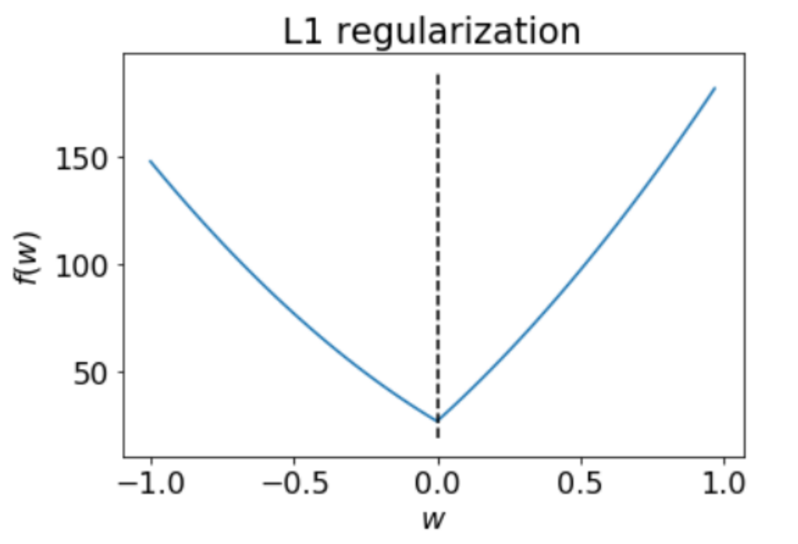

Why doesn't L2-Regularization set variables to 0?

- Consider an L2-regularized least squares problem with 1 feature:

$$f(w) = \frac{1}{2}\sum_{i=1}^{n}(wx_i - y_i)^2 + \frac{\lambda}{2}w^2$$

- Let's solve for the optimal 'w':

$$f'(w) = \sum_{i=1}^{n} x_i(wx_i - y_i) + \lambda w$$

Set equal to 0:

$$\sum_{i=1}^{n} x_i^2 w - \sum_{i=1}^{n} x_i y_i + \lambda w = 0$$

re-arrange

$$w\Big(\underbrace{\sum_{i=1}^{n} x_i^2}_{\|x\|^2} + \lambda\Big) = \underbrace{\sum_{i=1}^{n} x_i y_i}_{y^Tx}$$

or

$$w = \frac{y^Tx}{\|x\|^2 + \lambda}$$

- So as λ gets bigger, 'w' converges to 0.
- However, for all finite λ 'w' will be non-zero unless $y^Tx = 0$ exactly.
  - But it's very unlikely that y^Tx will be exactly zero.

Why doesn't L2-Regularization set variables to 0?

- Small λ — Big λ

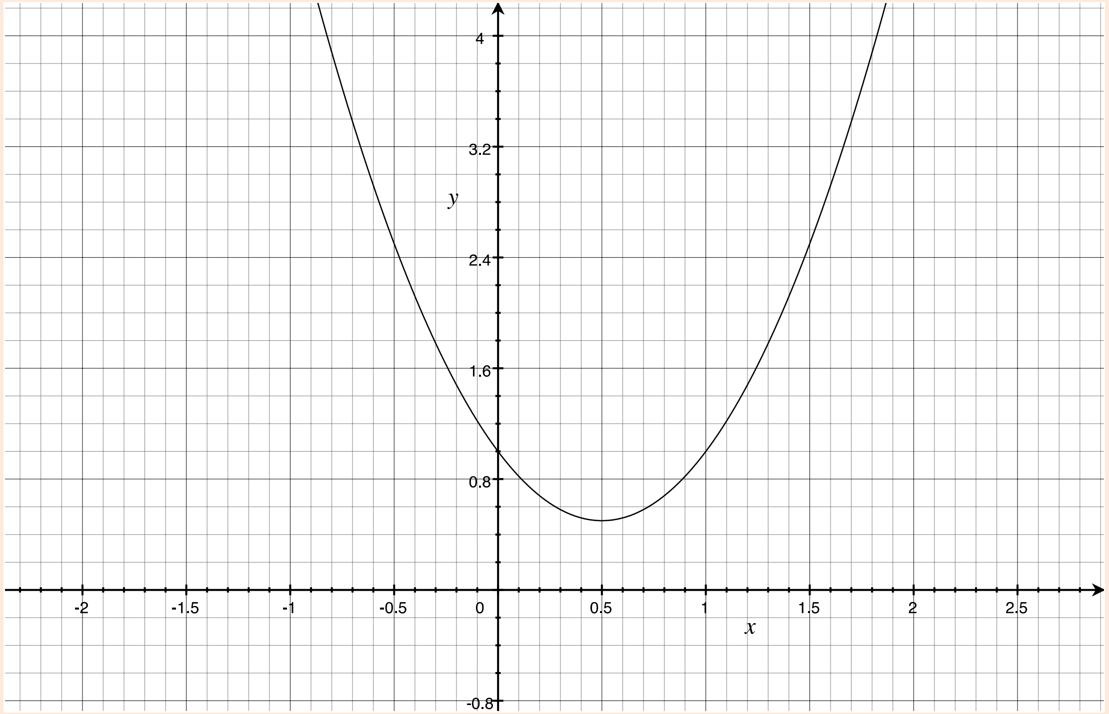

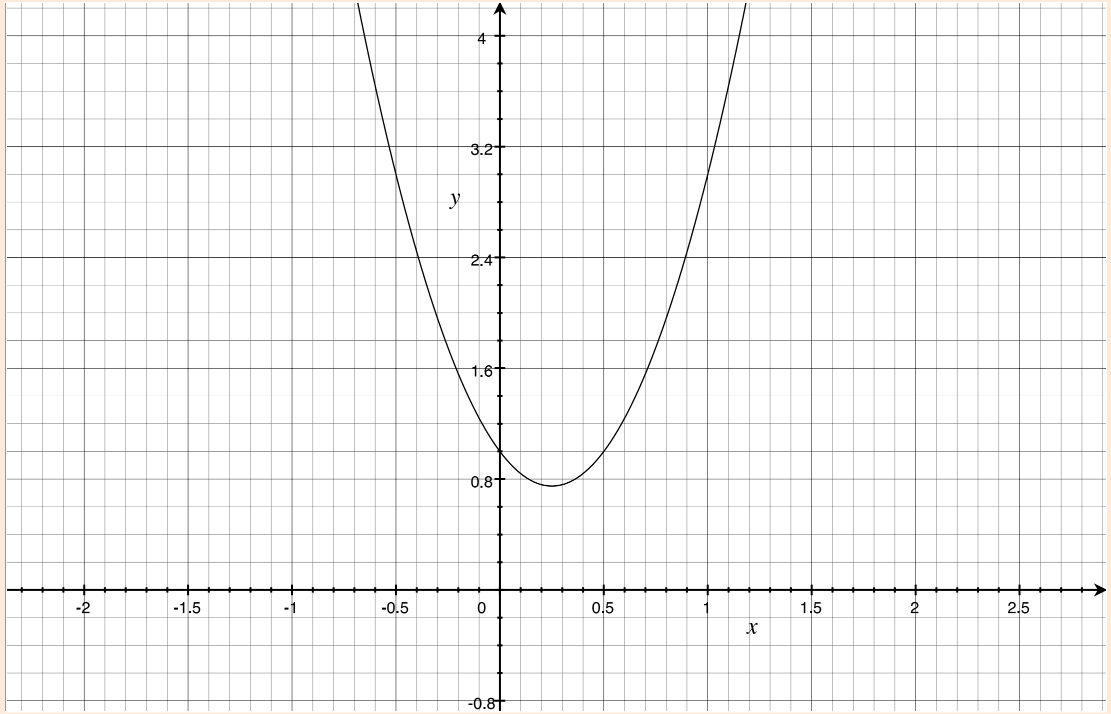

- Solution further from zero — Solution closer to zero (but not exactly 0)

Why does L1-Regularization set things to 0?

- Consider an L1-regularized least squares problem with 1 feature:

$$f(w) = \frac{1}{2}\sum_{i=1}^{n}(wx_i - y_i)^2 + \lambda|w|$$

- If (w = 0), then "left" limit and "right" limit are given by:

$$f^-(0) = \sum_{i=1}^{n} x_i(0x_i - y_i) - \lambda = \sum_{i=1}^{n} x_i y_i - \lambda$$

$$f^+(0) = \sum_{i=1}^{n} x_i(0x_i - y_i) + \lambda = \sum_{i=1}^{n} x_i y_i + \lambda$$

- So which direction should "gradient descent" go in?

$$-f^-(0) = -y^Tx + \lambda$$
$$-f^+(0) = -y^Tx - \lambda$$

If these are positive ($-y^Tx > \lambda$), we can improve by increasing 'w'.

If these are negative ($y^Tx > \lambda$), we can improve by decreasing 'w'.

But if left and right "gradient descent" directions point in opposite directions ($|y^Tx| \leq \lambda$), minimum is 0.

Why does L1-Regularization set things to 0?

- Small λ Big λ

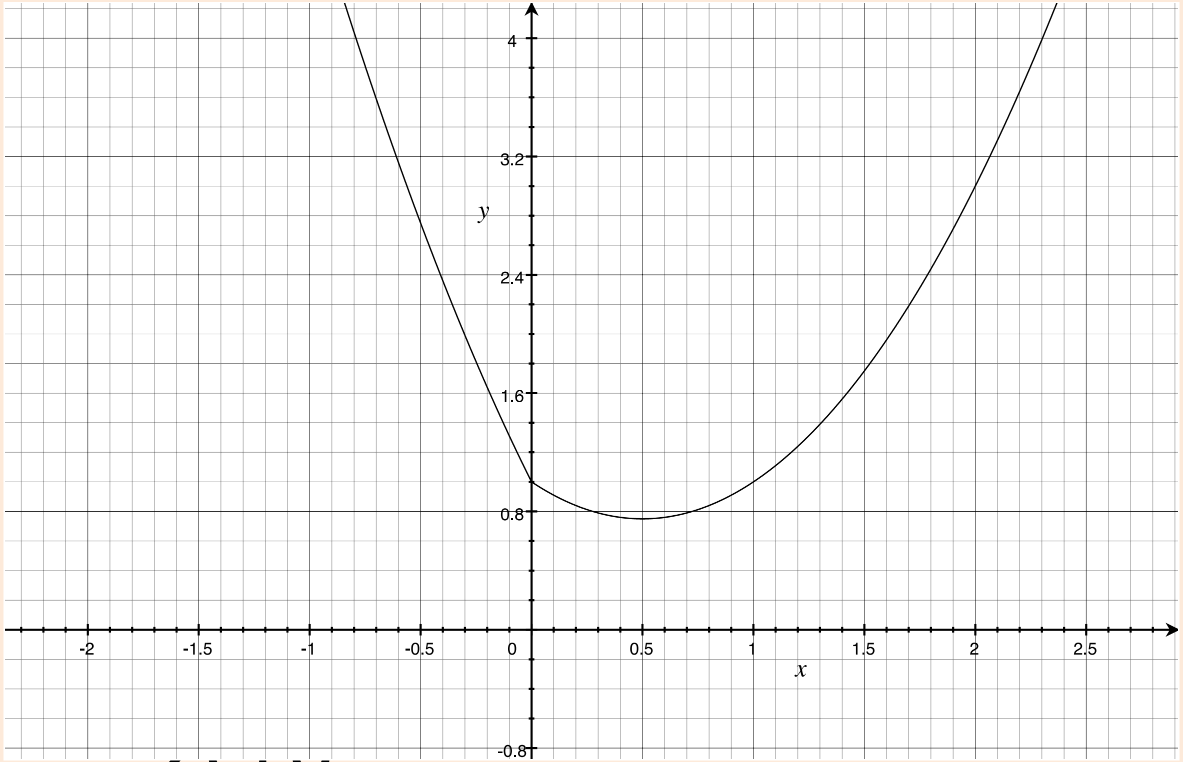

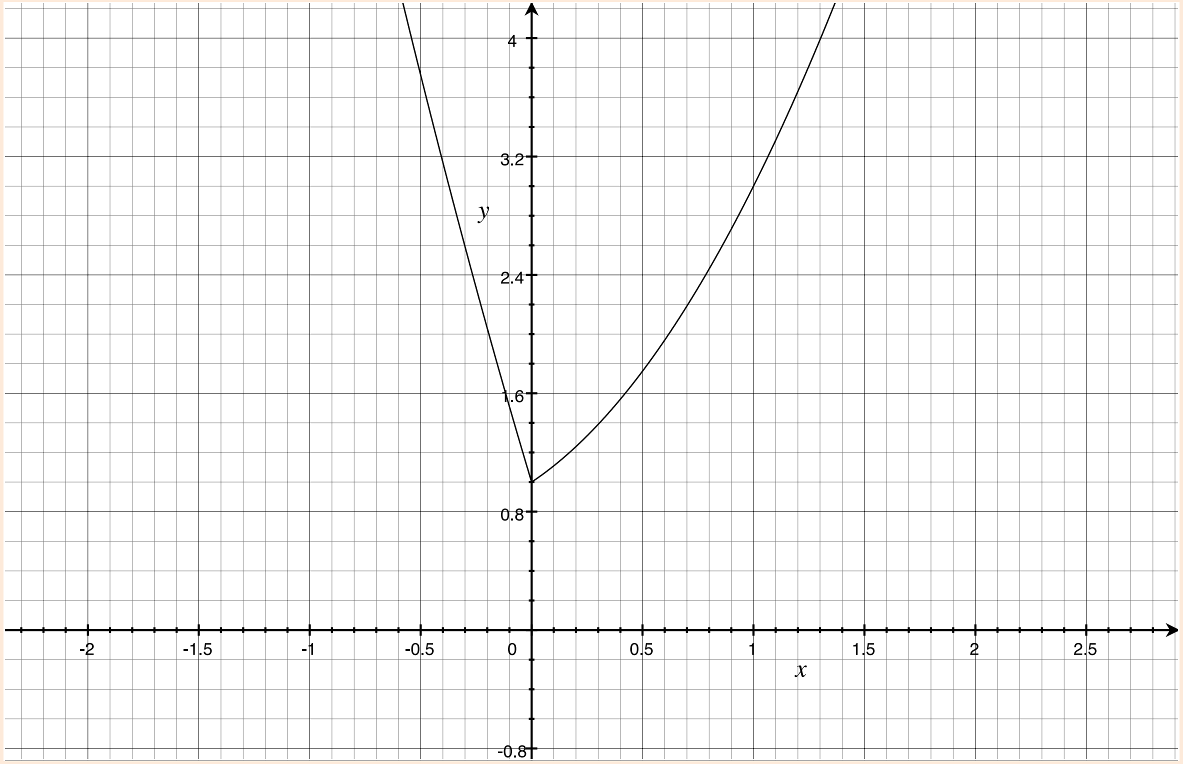

zero

(minimum of left parabola is past origin, but right parabola is not) (minimum of both parabola are past the origin)

L2-regularization vs. L1-regularization

- So with 1 feature:
  - L2-regularization only sets 'w' to 0 if $y^T x = 0$.
    - There is a only a single possible $y^T x$ value where the variable gets set to zero.
    - And λ has nothing to do with the sparsity.
  - L1-regularization sets 'w' to 0 if $|y^T x| \leq \lambda$.
    - There is a range of possible $y^T x$ values where the variable gets set to zero.
    - And increasing λ increases the sparsity since the range of $y^T x$ grows.
- Note that it's important that the function is non-differentiable:
  - Differentiable regularizers penalizing size would need $y^T x = 0$ for sparsity.

L1-Loss vs. Huber Loss

- The same reasoning tells us the difference between the L1 *loss* and the Huber loss. They are very similar in that they both grow linearly far away from 0. So both are both robust but…
  - With the L1 loss the model often passes exactly through some points.
  - With Huber the model doesn't necessarily pass through any points.

- Why? With L1-regularization we were causing the elements of 'w' to be exactly 0. Analogously, with the L1-loss we cause the elements of 'r' (the residual) to be exactly zero. But zero residual for an example means you pass through that example exactly.

Non-Uniqueness of L1-Regularized Solution

- How can L1-regularized least squares solution not be unique?
  - Isn't it convex?
- Convexity implies that minimum value of f(w) is unique (if exists), but there may be multiple 'w' values that achieve the minimum.

- Consider L1-regularized least squares with d=2, where feature 2 is a copy of a feature 1. For a solution (w_1,w_2) we have:

- So we can get the same squared error with different w_1 and w_2 values that have the same sum. Further, if neither w_1 or w_2 changes sign, then $|w_1| + |w_2|$ will be the same so the new w_1 and w_2 will be a solution.

$$\hat{y}_i = w_1 x_{i1} + w_2 x_{i2} = w_1 x_{i1} + w_2 x_{i1} = (w_1 + w_2) x_{i1}$$

Splines in 1D

- For 1D interpolation, alternative to polynomials/RBFs are splines:
  - Use a polynomial in the region between each data point.
  - Constrain some derivatives of the polynomials to yield a unique solution.
- Most common example is cubic spline:
  - Use a degree-3 polynomial between each pair of points.
  - Enforce that f'(x) and f''(x) of polynomials agree at all point.
  - "Natural" spline also enforces f''(x) = 0 for smallest and largest x.
- Non-trivial fact: natural cubic splines are sum of:
  - Y-intercept.
  - Linear basis.
  - RBFs with $g(\varepsilon) = \varepsilon^3$.
    - Different than Gaussian RBF because it *increases with distance*.

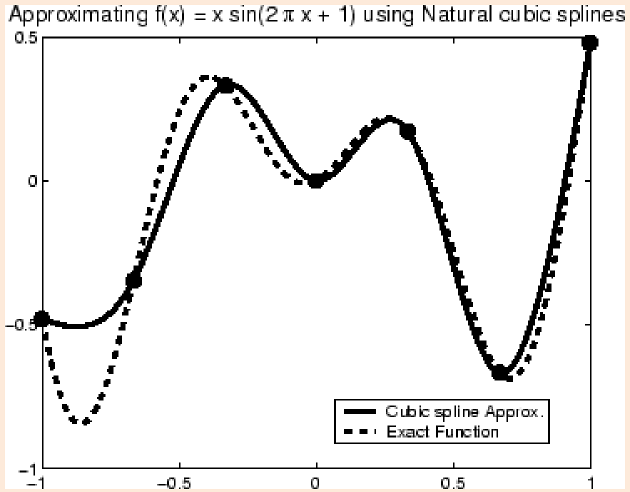

http://www.physics.arizona.edu/~restrepo/475A/Notes/sourcea-/node35.html

Splines in Higher Dimensions

- Splines generalize to higher dimensions if data lies on a grid.
  - Many methods exist for grid-structured data (linear, cubic, splines, etc.).
  - For more general ("scattered") data, there isn't a natural generalization.
- Common 2D "scattered" data interpolation is thin-plate splines:
  - Based on curve made when bending sheets of metal.
  - Corresponds to RBFs with $g(\varepsilon) = \varepsilon^2 \log(\varepsilon)$.
- Natural splines and thin-plate splines: special cases of "polyharmonic" splines:
  - Less sensitive to parameters than Gaussian RBF.

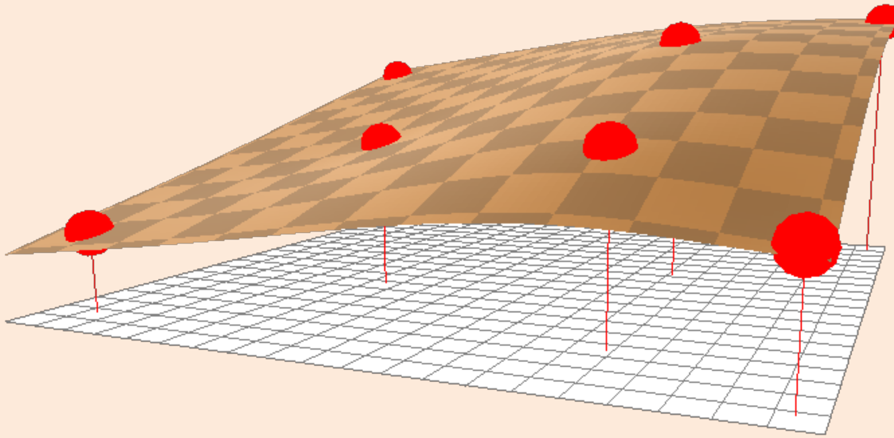

http://step.polymtl.ca/~rv101/thinplates/

L2-Regularization vs. L1-Regularization

- L2-regularization conceptually restricts 'w' to a ball.

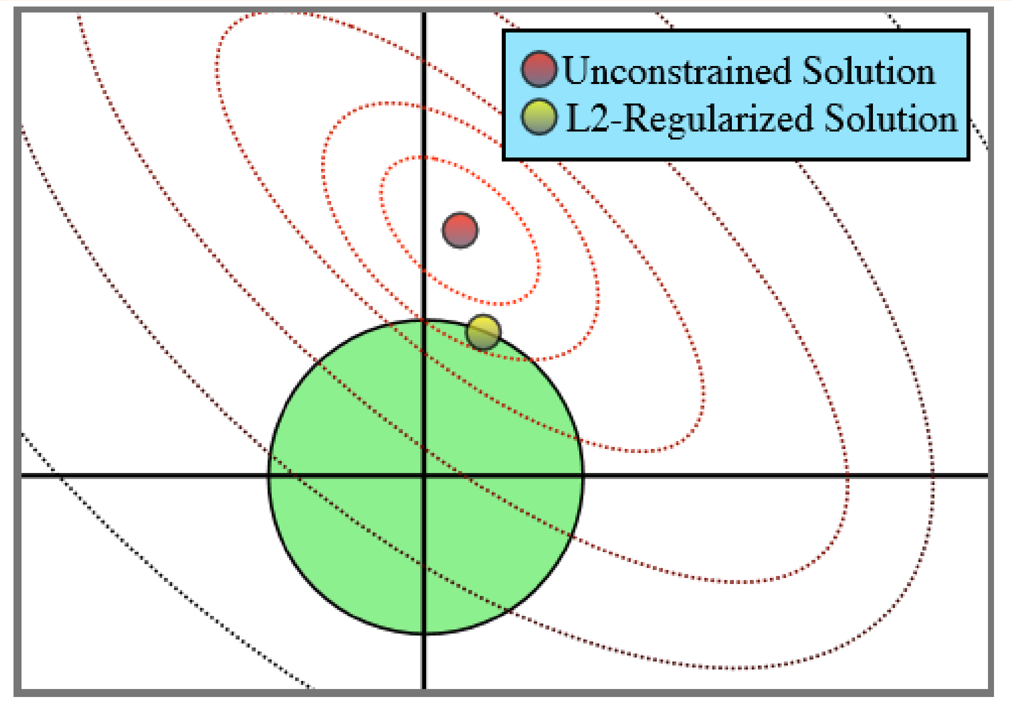

Minimizing $\frac{1}{2}\|Xw - y\|^2 + \frac{\lambda}{2}\|w\|^2$

is equivalent to minimizing

$\frac{1}{2}\|Xw - y\|^2$ subject to

the constraint that $\|w\| \leq \tau$

for some value 'τ'

L2-Regularization vs. L1-Regularization

- L2-regularization conceptually restricts 'w' to a ball.

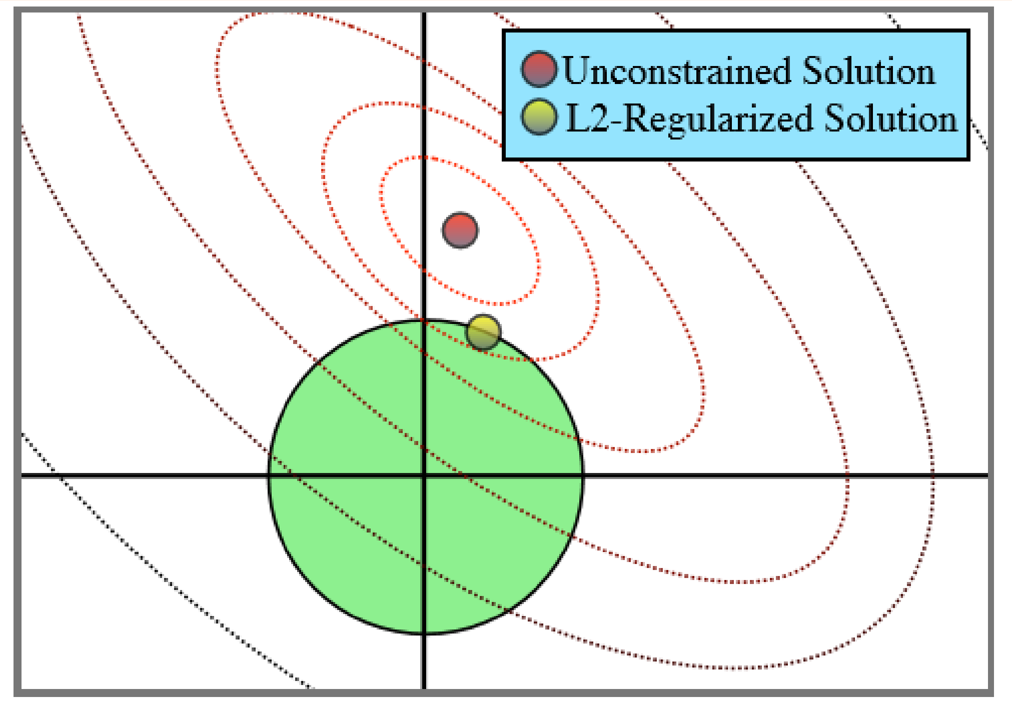

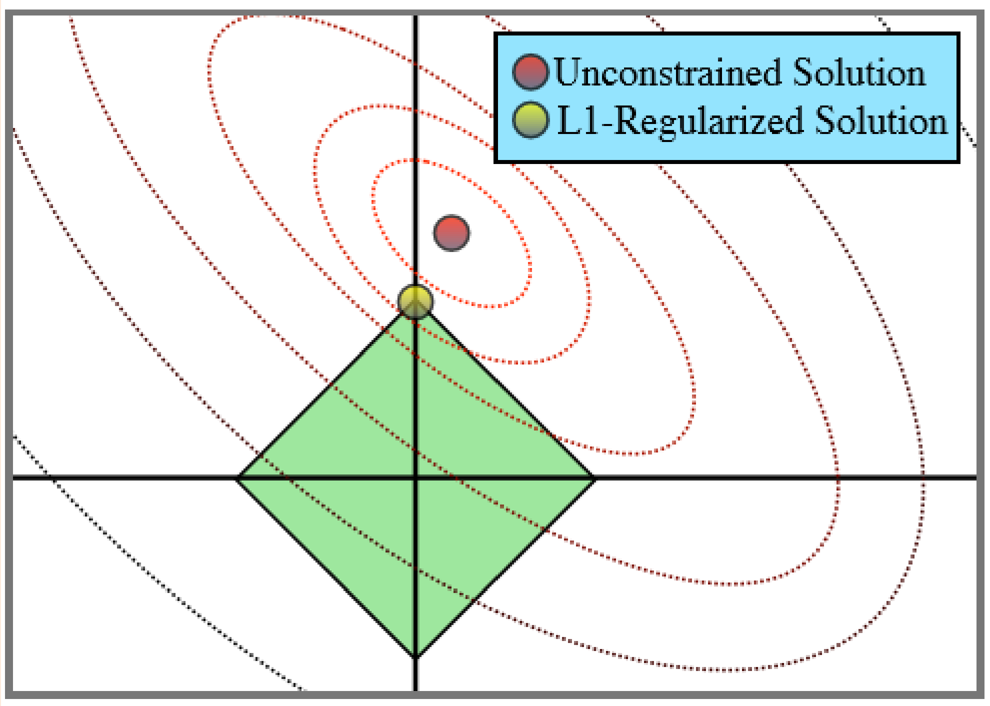

- L1-regularization restricts to the L1 "ball":
  - Solutions tend to be at corners where w_j are zero.

Related Infinite Series

L1-Regularization as a Feature Selection Method

- Advantages:
  - Deals with conditional independence (if linear).
  - Sort of deals with collinearity:
    - Picks at least one of “mom” and “mom2”.
  - Very fast with specialized algorithms.
- Disadvantages:
  - Tends to give false positives (selects too many variables).
- Neither good nor bad:
  - Does not take small effects.
  - Says “gender” is relevant if we know “baby”.
  - Good for prediction if we want fast training and don’t care about having some irrelevant variables included.

“Elastic Net”: L2- and L1-Regularization

- To address non-uniqueness, some authors use L2- and L1-:

$$f(w) = \frac{1}{2}\|Xw - y\|^2 + \frac{\lambda_2}{2}\|w\|^2 + \lambda_1\|w\|_1$$

- Called “elastic net” regularization.
  - Solution is sparse and unique.
  - Slightly better with feature dependence:
    - Selects both “mom” and “mom2”.

- Optimization is easier though still non-differentiable.

L1-Regularization Debiasing and Filtering

- To remove false positives, some authors add a debiasing step:
  - Fit 'w' using L1-regularization.
  - Grab the non-zero values of 'w' as the "relevant" variables.
  - Re-fit relevant 'w' using least squares or L2-regularized least squares.

- A related use of L1-regularization is as a filtering method:
  - Fit 'w' using L1-regularization.
  - Grab the non-zero values of 'w' as the "relevant" variables.
  - Run standard (slow) variable selection restricted to relevant variables.
    - Forward selection, exhaustive search, stochastic local search, etc.

Non-Convex Regularizers

- Regularizing $|w_j|^2$ selects all features.
- Regularizing $|w_j|$ selects fewer, but still has many false positives.
- What if we regularize $|w_j|^{1/2}$ instead?

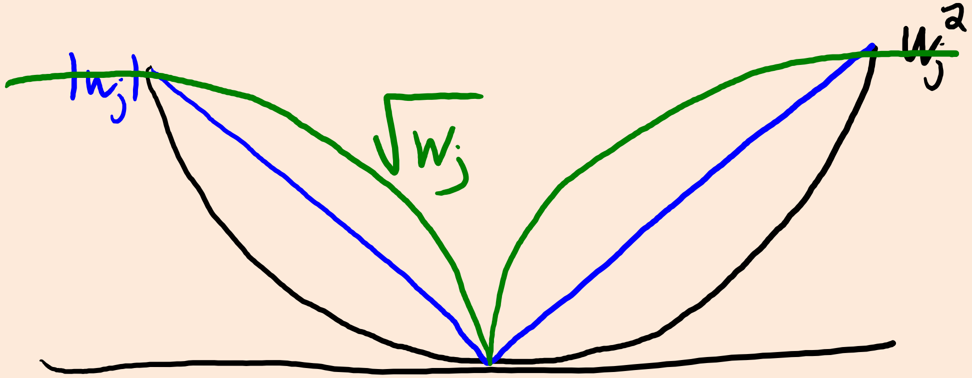

- Minimizing this objective would lead to fewer false positives.
  - Less need for debiasing, but it's not convex and hard to minimize.
- There are many non-convex regularizers with similar properties.
  - L1-regularization is (basically) the "most sparse" convex regularizer.