CPSC 340:
Machine Learning and Data Mining

Feature Selection

Fall 2018

Admin

- Assignment 3 is due tonight.
  - You can use a "late day" to submit up to 48 hours late.
  - Solutions will be posted Monday.

- Midterm is Thursday of next week.
  - October 18th at 6:30pm (BUCH A102 and A104).
  - 80 minutes.
  - Closed-book.
  - One doubled-sided 'cheat sheet' for midterm.
  - Auditors do not take the midterm.

- There will be two types of questions on the midterm:
  - 'Technical' questions requiring things like pseudo-code or derivations.
    - Similar to assignment questions, and will only be on topics related to those in assignments.
  - 'Conceptual' questions testing understanding of key concepts.
    - All lecture slide material except "bonus slides" is fair game here.

Last Time: Finding the "True" Model

- What if y_i really is a polynomial function of x_i?
  - How can we find the "true" degree 'p' of the polynomial?

- Training error does not work:
  - It goes down as 'p' goes up.

- Cross-validation may also not work:
  - Tends to overestimate 'p'.
  - Due to optimization bias.
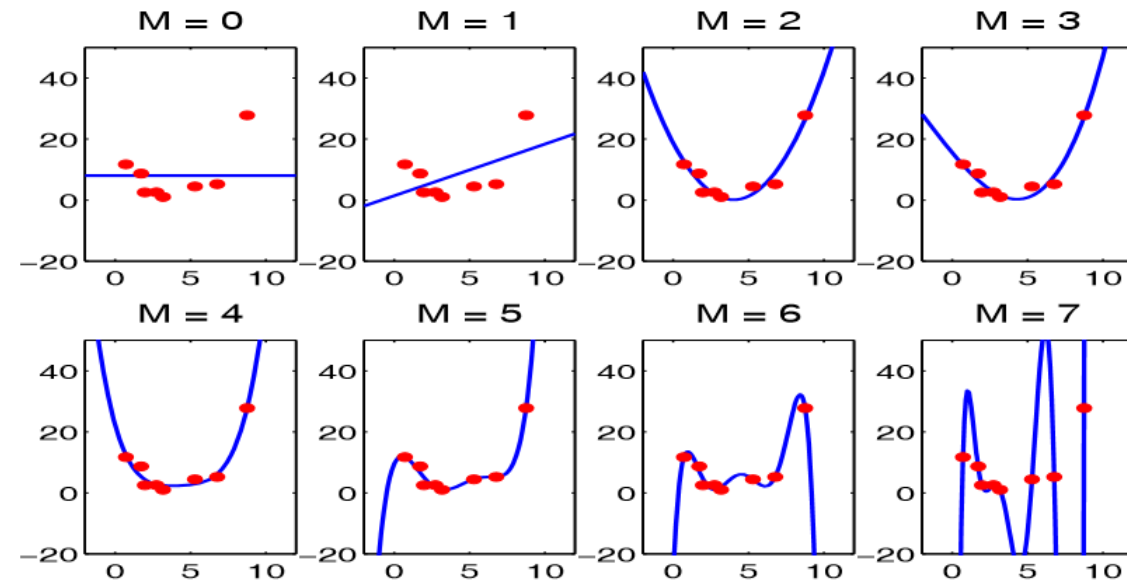
For example, imagine that the true model is $y_i = 2x_i^2 - 5 + (noise)$
We might choose $d=3$ and a model like $\hat{y}_i = 0.001x_i^3 + 2x_i^2 - 5$,
since it might get a slightly smaller validation error.

Last Time: Complexity Penalties

- We discussed putting a penalty on the model complexity.
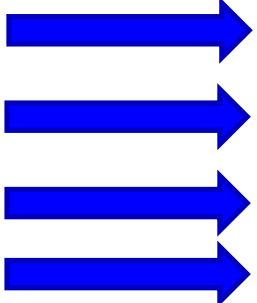 - Want to **fit the data and have a simple model**.

Find 'v' and 'p' minimizing:

$$Score(p) = \frac{1}{2}\|Z_p v - y\|^2 + \lambda K$$

usual error — "strength" of penalty — number of "degrees of Freedom"

  - "To increase the degrees of freedom by one, need to decrease error by λ.
- Prefers smaller degrees of freedom, if errors are similar.
  - Can't optimize this using gradient descent, since it's discontinuous in 'p'.
    - Need to search over values of 'p'.

Motivation: Discovering Food Allergies

- Recall the food allergy example:

Egg	Milk	Fish	Wheat	Shellfish	Peanuts	...		Sick?
0	0.7	0	0.3	0	0			1
0.3	0.7	0	0.6	0	0.01			1
0	0	0	0.8	0	0			0
0.3	0.7	1.2	0	0.10	0.01			1

- Instead of predicting "sick", we want to do feature selection:
  – Which foods are "relevant" for predicting "sick".

Feature Selection

- General feature selection problem:

$$X = \begin{bmatrix} & & \\ & & \\ & & \end{bmatrix} \quad \xrightarrow{\text{feature 'j'}} \quad y = \begin{bmatrix} \\ \\ \end{bmatrix}$$

  - Find the features (columns) of 'X' that are important for predicting 'y'.
    - "What are the relevant factors?"
    - "Which basis functions should I use among these choices?"
    - "What types of new data should I collect?"
    - "How can I speed up computation?"

- One of most important problems in ML/statistics, but very messy.
  - For now, we'll say a feature is "relevant" if it helps predict y_i from x_i.

"Association" Approach

- A simple/common way to do feature selection:
  - For each feature 'j', compute correlation between feature values x^j and 'y'.
    - Say that 'j' is relevant if correlation is above 0.9 or below -0.9.

- Turns feature selection into hypothesis testing for each feature.
    - There are many other measures of "dependence" (Wikipedia).

- Usually gives unsatisfactory results as it ignores variable interactions:
  - Includes irrelevant variables: "Taco Tuesdays".
    - If tacos make you sick, and you often eat tacos on Tuesdays, it will say "Tuesday" is relevant.
  - Excludes relevant variables: "Diet Coke + Mentos Eruption".
    - Diet coke and Mentos don't make you sick on their own, but *together* they make you sick.

Genome-Wide Association Studies

- Genome-wide association studies:
  - Measure if there exists a dependency between each individual "single-nucleotide polymorphism" in the genome and a particular disease.
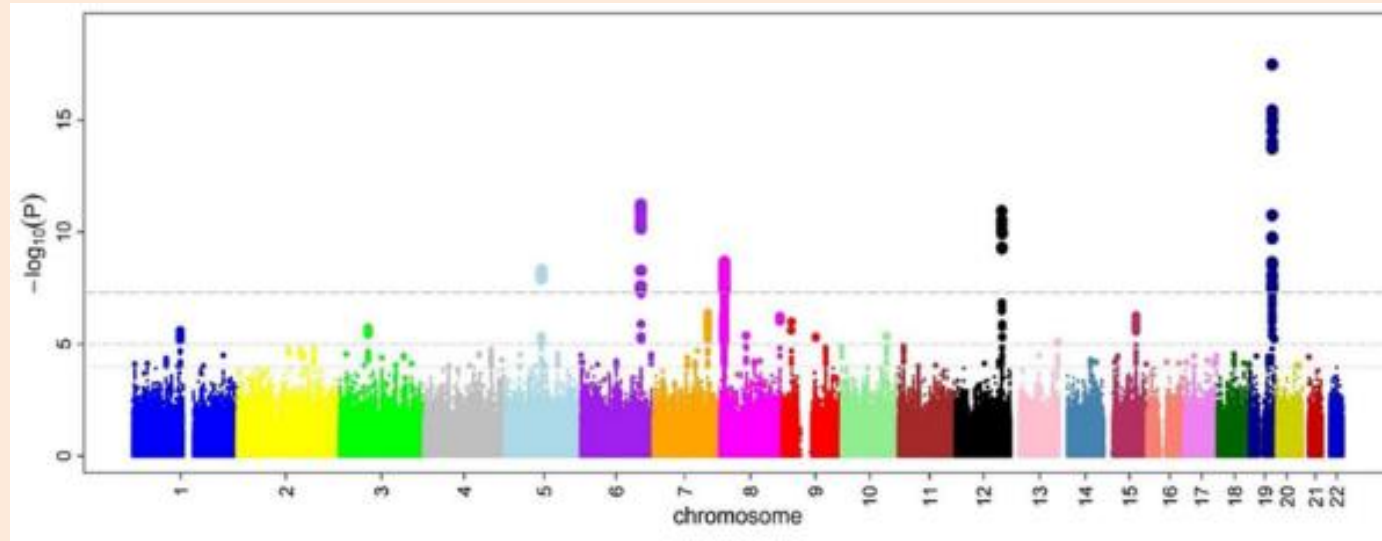
  - Has identified thousands of genes "associated" with diseases.
    - But *by design* this has a huge numbers of false positives (and many false negatives).

"Regression Weight" Approach

- A simple/common approach to feature selection:
  - Fit regression weights 'w' based on **all** features (maybe with least squares).
  - Take all features 'j' where weight $|w_j|$ is greater than a threshold.

- This could recognize that "Tuesday" is irrelevant.
  - If you get enough data, and you sometimes eat tacos on other days.
    (And the relationship is actually linear.)

- This could recognize that "Diet Coke" and "Mentos" are relevant.
  - Assuming this combination occurs enough times in the data.

"Regression Weight" Approach

- A simple/common approach to feature selection:
  - Fit regression weights 'w' based on **all** features (maybe with least squares).
  - Take all features 'j' where weight $|w_j|$ is greater than a threshold.

- Has major problems with collinearity:
  - If the "Tuesday" variable always equals the "taco" variable,
    it could say that Tuesdays are relevant but tacos are not.

$$\hat{y}_i = w_1 * taco + w_2 * Tuesday = 0 * taco + (w_1 + w_2) * Tuesday$$

  - If you have two copies of an irrelevant feature,
    it could take both irrelevant copies.

$$\hat{y}_i = 0 * irrelevant + 0 * irrelevant = 10000 * irrelevant + (-10000) * irrelevant$$

(pause)

Search and Score Methods

- Most common feature selection framework is search and score:
  1. Define score function f(S) that measures quality of a set of features 'S'.
  2. Now search for the variables 'S' with the best score.

- Example with 3 features:
  - Compute "score" of using feature 1.
  - Compute "score" of using feature 2.
  - Compute "score" of using feature 3.
  - Compute "score" of using features {1,2}.
  - Compute "score" of using features {1,3}.
  - Compute "score" of using features {2,3}.
  - Compute "score" of using features {1,2,3}.
  - Compute "score" of using features {}.
  - Return the set of features 'S' with the best "score".

Which Score Function?

- The score can't be the training error.

  - Training error goes down as you add features, so will select all features.

- A more logical score is the validation error.

  - "Find the set of features that gives the lowest validation error."

  - To minimize test error, this is what we want.

- But there are problems due to the large number of sets of variables:

  - If we have 'd' variables, there are 2^d sets of variables.

  - Optimization bias is high: we're optimizing over 2^d models (not 10).

  - Prone to false positives: irrelevant variables will sometimes help by chance.

"Number of Features" Penalties

- To reduce false positives, we can again use complexity penalties:

$$\text{score}(S) = \frac{1}{2} \sum_{i=1}^{n} (w_S^T x_{iS} - y_i)^2 + \text{size}(S)$$

  - E.g., we could use squared error and number of non-zeroes.
  - We're using 'x_{iS}' as the features 'S' of example x_i.

- If two 'S' have similar error, this prefers the smaller set.
  - It prefers having removing feature 3 instead of having $w_3 = 0.00001$.

- Instead of "size(S)", we usually write this using the "L0-norm"...

L0-Norm and "Number of Features We Use"

- In linear models, setting $w_j = 0$ is the same as removing feature 'j':

$$\hat{y}_i = w_1 x_{i1} + w_2 x_{i2} + w_3 x_{i3} + \cdots + w_d x_{id}$$

set $w_2 = 0$

$$\hat{y}_i = w_1 x_{i1} + 0 + w_3 x_{i3} + \cdots + w_d x_{id}$$

ignore x_{i2}

- The L0 "norm" is the number of non-zero values.

If $w = \begin{bmatrix} 1 \\ 0 \\ 2 \\ 0 \\ 3 \end{bmatrix}$ then $\|w\|_0 = 3$ If $w = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$ then $\|w\|_0 = 0$.

  – Not actually a true norm.
  – If 'w' has a small L0-norm, then it doesn't use many features.

L0-penalty: optimization

- L0-norm penalty for feature selection:

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

training error   degrees of freedom 'k'

- Suppose we want to use this to evaluate the features S = {1,2}:
  - First fit the 'w' just using features 1 and 2.
  - Now compute the training error with this 'w' and features 1 and 2.
  - Add λ*2 to the training error to get the score.

- We repeat this with other choices of 'S' to find the "best" features.

L0-penalty: interpretation

- L0-norm penalty for feature selection:

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

- Balances between training error and number of features we use.
  - With λ=0, we get least squares with all features.
  - With λ=∞, we must set w=0 and not use any features.

  - With other λ, balances between training error and number of non-zeroes.
    - Larger λ puts more emphasis on having zeroes in 'w' (more feature selection).
    - Different values give AIC, BIC, and so on.

Forward Selection (Greedy Search Heuristic)

- In search and score, it's also just hard to search for the best 'S'.
  - There are 2^d possible sets.

- A common greedy search procedure is forward selection:

1. Compute score if we use no features.
2. Try adding "taco", "milk", "egg", and so on (computing score of each)
3. Add "milk" because it got the best score.
4. Try {milk, taco}, {milk, egg}, and so on (computing score of each variable with milk)
5. Continue until no single-variable addition improves the score.

Forward Selection (Greedy Search Heuristic)

- Forward selection algorithm for variable selection:
  1. Start with an empty set of features, S = [ ].
  2. For each possible feature 'j':
     - Compute scores of features in 'S' combined with feature 'j'.
  3. If no 'j' improves the score, stop.
  4. Otherwise, add the 'j' that improves the score the most to 'S'.
     - Then go back to Step 2.

- Not guaranteed to find the best set, but reduces many problems:
  – Considers $O(d^2)$ models: cheaper, ovefits less, has fewer false positives.

Backward Selection and RFE

- Forward selection often works better than naïve methods.

- A related method is backward selection:
  – Start with all features, remove the one that most improves the score.

- If you consider adding or removing features, it's called stagewise.

- Stochastic local search is a class of fancier methods.
  – Simulated annealing, genetic algorithms, ant colony optimization, etc.

- Recursive feature elimination is another related method:
  – Fit parameters of a regression model.
  – Prune features with small regression weights.
  – Repeat.

(pause)

Is "Relevance" Clearly Defined?

- Consider a supervised classification task:

gender	mom	dad
F	1	0
M	0	1
F	0	0
F	1	1

SNP
1
0
0
1

- Predict whether someone has particular genetic variation (SNP).
  - Location of mutation is in "mitochondrial" DNA.
    - "You almost always have the same value as your mom".

 - For simplicity we'll assume 1950s-style gender and parentage.

Is "Relevance" Clearly Defined?

- Consider a supervised classification task:

gender	mom	dad
F	1	0
M	0	1
F	0	0
F	1	1

SNP
1
0
0
1

Mitochondrial

Unaffected father Affected mother   Affected father Unaffected mother

Affected children   Unaffected children

■ Unaffected
□ Affected

U.S. National Library of Medicine

- True model:
  - (SNP = mom) with very high probability.
  - (SNP != mom) with some very low probability.
- What are the "relevant" features for this problem?
  - Mom is relevant and {gender, dad} are not relevant.
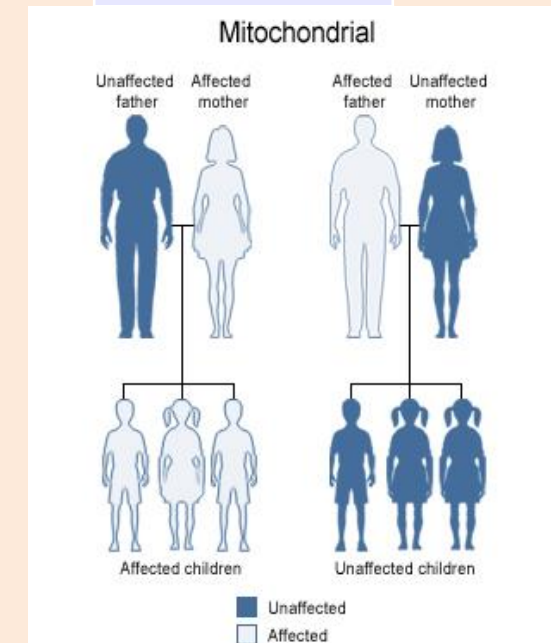
Is "Relevance" Clearly Defined?

- What if "mom" feature is repeated?

gender	mom	dad	mom2
F	1	0	1
M	0	1	0
F	0	0	0
F	1	1	1

SNP
1
0
0
1

- Are "mom" and "mom2" relevant?
  - Should we pick them both?
  - Should we pick one because it predicts the other?

- General problem ("dependence", "collinearity" for linear models):
  - If features can be predicted from features, don't know one(s) to pick.

Neither of these is "correct", but not picking either is incorrect.

Is "Relevance" Clearly Defined?

- What if we add (maternal) "grandma"?

gender	mom	dad	grandma
F	1	0	1
M	0	1	0
F	0	0	0
F	1	1	1

SNP
1
0
0
1

- Is "grandma" relevant?
  - You can predict SNP very accurately from "grandma" alone.
  - But "grandma" is irrelevant if I know "mom".
- General problem (conditional independence):
  - "Relevant" features may be irrelevant given other features.

Is "Relevance" Clearly Defined?

- What if we don't know "mom"?

gender	grandma	dad
F	1	0
M	0	1
F	0	0
F	1	1

SNP
1
0
0
1

- Now is "grandma" is relevant?
  - Without "mom" variable, using "grandma" is the best you can do.
- General problem ("taco Tuesday"):
  - Features can be relevant due to missing information.

Is "Relevance" Clearly Defined?

- What if we don't know "mom" or "grandma"?

gender	dad
F	0
M	1
F	0
F	1

SNP
1
0
0
1

- Now there are no relevant variables, right?
  - But "dad" and "mom" must have some common maternal ancestor.
  - "Mitochondrial Eve" estimated to be ~200,000 years ago.
- General problem (effect size):
  - "Relevant" features may have small effects.

Is "Relevance" Clearly Defined?

- What if we don't know "mom" or "grandma"?

gender	dad
F	0
M	1
F	0
F	1

SNP
1
0
0
1

- Now there are no relevant variables, right?
  - What if "mom" likes "dad" because he has the same SNP as her?
- General problem (confounding):
  - Hidden effects can make "irrelevant" variables "relevant".

Is "Relevance" Clearly Defined?

- What if we add "sibling"?

gender	dad	sibling
F	0	1
M	1	0
F	0	0
F	1	1

SNP
1
0
0
1

- Sibling is "relevant" for predicting SNP, but it's not the cause.
- General problem (non-causality or reverse causality):
  - A "relevant" feature may not be causal, or may be an effect of label.

Is "Relevance" Clearly Defined?

- What if don't have "mom" but we have "baby"?

gender	dad	baby
F	0	1
M	1	1
F	0	0
F	1	1

SNP
1
0
0
1

- "Baby" is relevant when (gender == F).
  - "Baby" is relevant (though causality is reversed).
  - Is "gender" relevant?
    - If we want to find relevant causal factors, "gender" is not relevant.
    - If we want to predict SNP, "gender" is relevant.
- General problem (context-specific relevance):
  - Adding a feature can make an "irrelevant" feature "relevant".

Is "Relevance" Clearly Defined?

- Warnings about feature selection:
  - A feature is only "relevant" in the context of available features.
    - Adding/removing features can make features relevant/irrelevant.

  - Confounding factors can make "irrelevant" variables the most "relevant".

  - If features can be predicted from features, you can't know which to pick.
    - Collinearity is a special case of "dependence" (which may be non-linear).

  - A "relevant" feature may have a tiny effect.

  - "Relevance" for prediction does not imply a causal relationship.

Is this hopeless?

- We often want to do feature selection we so have to try!

- Different methods are affected by problems in different ways.
  - We'll ignore causality and confounding issues (bonus slides).

- These "problems" don't have right answers but have wrong answers:
  - Variable dependence ("mom" and "mom2" have same information).
    - But should take at least one.
  - Conditional independence ("grandma" is irrelevant given "mom").

- These "problems" have application-specific answers:
  - Tiny effects.
  - Context-specific relevance (is "gender" relevant if given "baby"?).

Rough Guide to Feature Selection

Method\Issue	Dependence	Conditional Independence	Tiny effects	Context-Specific Relevance
Association (e.g., measure correlation between features 'j' and 'y')	Ok (takes "mom" and "mom2")	Bad (takes "grandma", "great-grandma", etc.)	Ignores	Bad (misses features that must interact, "gender" irrelevant given "baby")

Rough Guide to Feature Selection

Method\Issue	Dependence	Conditional Independence	Tiny effects	Context-Specific Relevance
Association (e.g., measure correlation between features 'j' and 'y')	Ok (takes "mom" and "mom2")	Bad (takes "grandma", "great-grandma", etc.)	Ignores	Bad (misses features that must interact, "gender" irrelevant given "baby")
Regression Weight (fit least squares, take biggest $|w_j|$)	Bad (can take irrelevant but collinear, can take none of "mom1-3")	Ok (takes "mom" not "grandma", if linear and 'n' large.	Ignores (unless collinear)	Ok (if linear, "gender" relevant give "baby")

Rough Guide to Feature Selection

Method\Issue	Dependence	Conditional Independence	Tiny effects	Context-Specific Relevance
Association (e.g., measure correlation between features 'j' and 'y')	Ok (takes "mom" and "mom2")	Bad (takes "grandma", "great-grandma", etc.)	Ignores	Bad (misses features that must interact, "gender" irrelevant given "baby")
Regression Weight (fit least squares, take biggest $|w_j|$)	Bad (can take irrelevant but collinear, can take none of "mom1-3")	Ok (takes "mom" not "grandma", if linear and 'n' large.	Ignores (unless collinear)	Ok (if linear, "gender" relevant give "baby")
Search and Score w/ Validation Error	Ok (takes at least one of "mom" and "mom2")	Bad (takes "grandma", "great-grandma", etc.)	Allows	Ok ("gender" relevant given "baby")

Rough Guide to Feature Selection

Method\Issue	Dependence	Conditional Independence	Tiny effects	Context-Specific Relevance
Association (e.g., measure correlation between features 'j' and 'y')	Ok (takes "mom" and "mom2")	Bad (takes "grandma", "great-grandma", etc.)	Ignores	Bad (misses features that must interact, "gender" irrelevant given "baby")
Regression Weight (fit least squares, take biggest $\lvert w_j \rvert$)	Bad (can take irrelevant but collinear, can take none of "mom1-3")	Ok (takes "mom" not "grandma", if linear and 'n' large.	Ignores (unless collinear)	Ok (if linear, "gender" relevant give "baby")
Search and Score w/ Validation Error	Ok (takes at least one of "mom" and "mom2")	Bad (takes "grandma", "great-grandma", etc.)	Allows (many false positives)	Ok ("gender" relevant given "baby")
Search and Score w/ L0-norm	Ok (takes exactly one of "mom" and "mom2")	Ok (takes "mom" not grandma if linear-ish).	Ignores (even if collinear)	Ok ("gender" relevant given "baby")

Summary

- Feature selection is task of choosing the "relevant" features.
  - Obvious simple approaches have obvious simple problems.
- Search and score: find features that optimize some score.
  - L0-norm penalties are the most common scores.
  - Forward selection is a heuristic to search over a smaller set of features.
- "Relevance" is really hard to define.

- Next time: getting a good test error even with irrelevant features.

Mallow's Cp

- Older than AIC and BIC is Mallow's Cp:

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

least squares weights if we used all features

- Minimizing this score is equivalent to L0-regularization:

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

$$\text{with } \lambda = \frac{\|X\hat{w} - y\|^2}{n}$$

- So again, viewing λ as hyper-parameter, this score is special case.

Adjusted R²

- Older than AIC and BIC and Mallow's Cp is adjusted R²:

$$f(w) = 1 - (1 - R^2) \frac{n-1}{n - \|w\|_0 - 1} \quad \text{where} \quad R^2 = 1 - \frac{\|Xw - y\|^2}{\|X\hat{w} - y\|^2}$$

- Maximizing this score is equivalent to L0-regularization:

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

$$\text{with} \quad \lambda = \frac{\|X\hat{w} - y\|^2}{2(n-1)}$$

- So again, viewing λ as hyper-parameter, this score is special case.

ANOVA

- Some people also like to compute this "ANOVA" quantity:

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

mean of y_i values repeated 'n' times

- This is based on the decomposition of "total squared error" as:

$$\|y - \bar{y}\|^2 = \|Xw - \bar{y}\|^2 + \|Xw - y\|^2$$

"total" error   "explained" error   "residual" (usual) error.

- Notice that "explained error" goes up as our usual ("residual") error goes down.
- Trying to find the 'k' features that maximize 'f' ("explain the most variance") is equivalent to L0-regularization with a particular λ (so another special case).

Information Criteria with Noise Variance

- We defined AIC/BIC for feature selection in least squares as:

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

- The first term comes from assuming $y_i = w^T x_i + \varepsilon$,
  where ε comes from a normal distribution with a variance of 1.
  - We'll discuss why when we discuss MLE and MAP estimation.
- If you treat variance as a parameter, then after some manipulation:

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

- However, this is again equivalent to just changing λ.
- If you aren't doing least squares, replace first term by "log-likelihood".

Complexity Penalties for Other Models

- Scores like AIC and BIC can also be used in other contexts:
  - When fitting a decision tree, only split a node if it improves BIC.
  - This makes sense if we're looking for the "true tree", or maybe just a simple/interpretable tree that performs well.

- In these cases we replace "L0-norm" with "degrees of freedom".
  - In linear models fit with least squares, degrees of freedom is number of non-zeroes.
  - Unfortunately, it is not always easy to measure "degrees of freedom".

Alternative to Search and Score: good old p-values

- Hypothesis testing ("constraint-based") approach:
  - Generalization of the "association" approach to feature selection.
  - Performs a sequence of conditional independence tests.

$$x_{ij} \perp y_i \mid x_{is}$$

feature 'j' — label — other features 's'
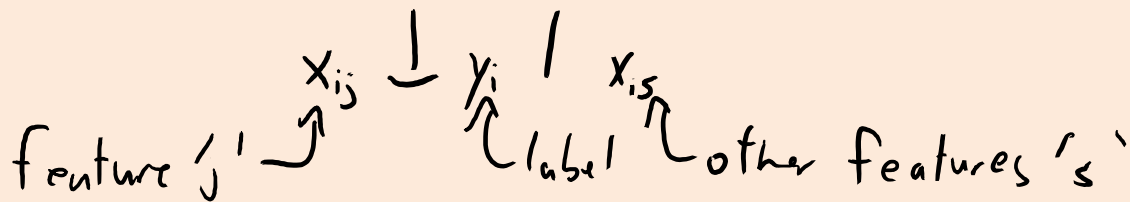
"If I know features in 's' does feature 'j' tell me anything about label?"

  - If they are independent (like "p < .05"), say that 'j' is "irrelevant".
- Common way to do the tests:
  - "Partial" correlation (numerical data).
  - "Conditional" mutual information (discrete data).

Testing-Based Feature Selection

- Hypothesis testing ("constraint-based") approach:

- Two many possible tests, "greedy" method is for each 'j' do:

First test if $x_{ij} \perp y_i$

If still dependent test $x_{ij} \perp y_i \mid x_{is}$ where 's' has <u>one</u> feature $\Big\}$ Often choose features to minimize dependence.

If still dependent test $x_{ij} \perp y_i \mid x_{is}$ where 's' now has two features

⋮

If still dependent when 's' includes all other features, declare 'j' relevant.

- "Association approach" is the greedy method where you only do the first test (subsequent tests remove a lot of false positives).

Hypothesis-Based Feature Selection

- Advantages:
  - Deals with conditional independence.
  - Algorithm can explain why it thinks 'j' is irrelevant.
  - Doesn't necessarily need linearity.

- Disadvantages:
  - Deals badly with exact dependence: doesn't select "mom" or "mom2" if both present.
  - Usual warning about testing multiple hypotheses:
    - If you test $p < 0.05$ more than 20 times, you're going to make errors.
  - Greedy approach may be sub-optimal.

- Neither good nor bad:
  - Allows tiny effects.
  - Says "gender" is irrelevant when you know "baby".
  - This approach is sometimes better for finding relevant factors, not to select features for learning.

Causality

- None of these approaches address causality or confounding:
  - "Mom" is the only relevant direct causal factor.
  - "Dad" is really irrelevant.
  - "Grandma" is causal but is irrelevant if we know "mom".

- Other factors can help prediction but aren't causal:
  - "Sibling" is predictive due to confounding of effect of same "mom".
  - "Baby" is predictive due to reverse causality.
  - "Gender" is predictive due to common effect on "baby".
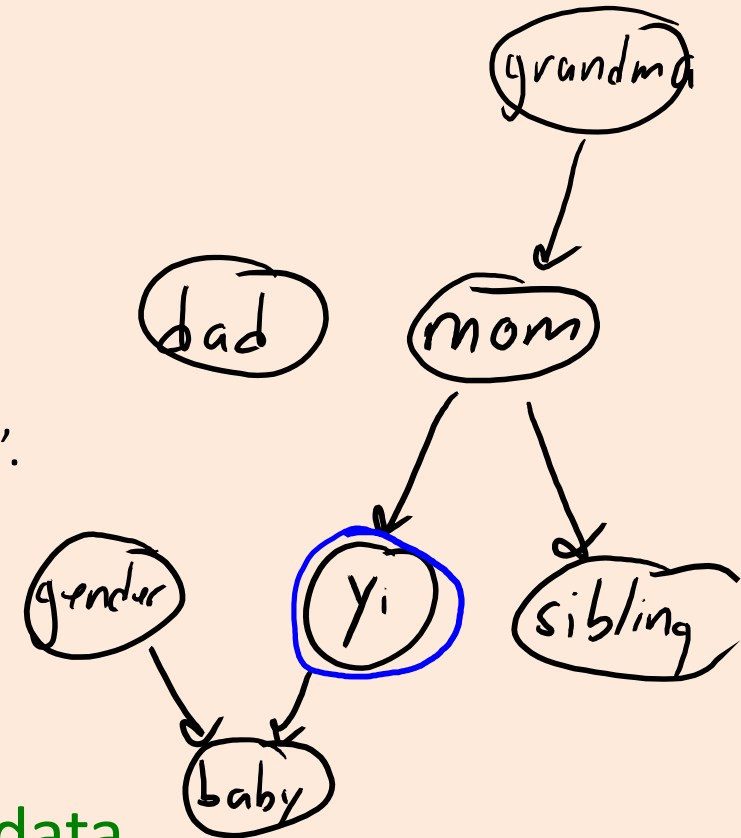
- We can sometimes address this using interventional data…

Interventional Data

- The difference between observational and interventional data:
  - If I see that my watch says 10:45, class is almost over (observational).
  - If I set my watch to say 10:45, it doesn't help (interventional).

- The intervention can help discover causal effects:
  - "Watch" is only predictive of "time" in observational setting (so not causal).

- General idea for identifying causal effects:
  - "Force" the variable to take a certain value, then measure the effect.
    - If the dependency remains, there is a causal effect.
    - We "break" connections from reverse causality, common effects, or confounding.

Causality and Dataset Collection

- This has to do with the way you collect data:
  - You can't "look" for variables taking the value "after the fact".
  - You need to manipulate the value of the variable, then watch for changes.

- This is the basis for randomized control trial in medicine:
  - Randomly assigning pills "forces" value of "treatment" variable.
  - Include a "control" as a value to prevent placebo effect as confounding.

- See also Simpson's Paradox:
  - https://www.youtube.com/watch?v=ebEkn-BiW5k

Structure Learning: Unsupervised Feature Selection

- "News" data: presence of 100 words in 16k newsgroup posts:

car	drive	files	hockey	mac	league	pc	win
0	0	1	0	1	0	1	0
0	0	0	1	0	1	0	1
1	1	0	0	0	0	0	0
0	1	1	0	1	0	0	0
0	0	1	0	0	0	1	1

- Which words are related to each other?

- Problem of structure learning: unsupervised feature selection.

Structure Learning: Unsupervised Feature Selection

- Optimal tree structure:
  (ignore arrows)
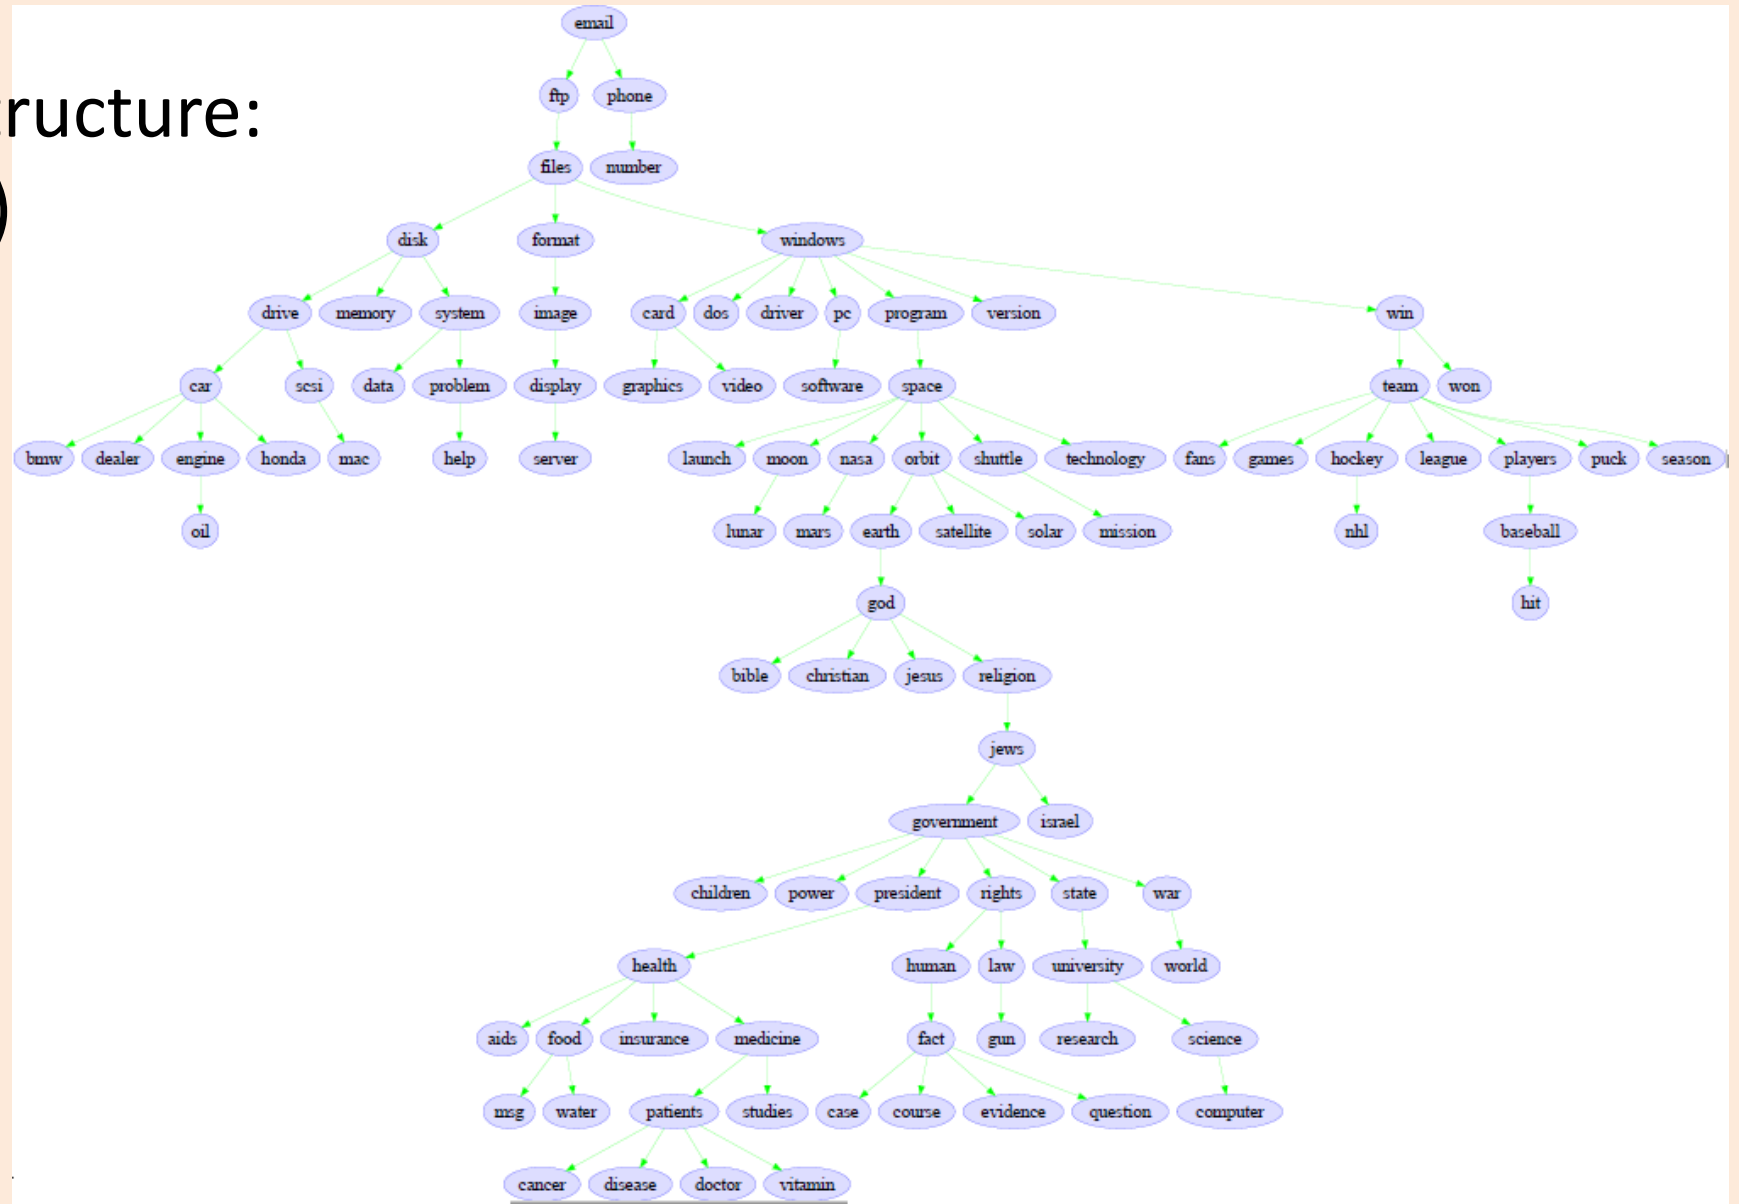
Naïve Approach: Association Networks

- A naïve approach to structure learning ("association networks"):
  - For each pair of variables, compute a measure of similarity or dependence.

- Using these n^2 similarity values either:
  - Select all pairs whose similarity is above a threshold.
  - Select the "top k" most similar features to each feature 'j'.

- Main problems:
  - Usually, most variables are dependent (too many edges).
    - "Sick" is getting connected to "Tuesdays" even if "tacos" are a variable.
  - "True" neighbours may not have the highest dependence.
    - "Sick" might get connected to "Tuesdays" before it gets connected to "milk".

- (Variation: best tree can be found as minimum spanning tree problem.)

Example: Vancouver Rain Data

- Consider modeling the "Vancouver rain" dataset.

	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7	Day 8	Day 9	...
Month 1	0	0	0	1	1	0	0	1	1	
Month 2	1	0	0	0	0	0	1	0	0	
Month 3	1	1	1	1	1	1	1	1	1	
Month 4	1	1	1	1	0	0	1	1	1	
Month 5	0	0	0	0	1	1	0	0	0	
Month 6	0	1	1	0	0	0	0	1	1	

- The strongest signal in the data is the simple relationship:
  - If it rained yesterday, it's likely to rain today (> 50% chance that $x^{t-1} = x^t$).
  - But an "association network" might connect all days (all dependent).
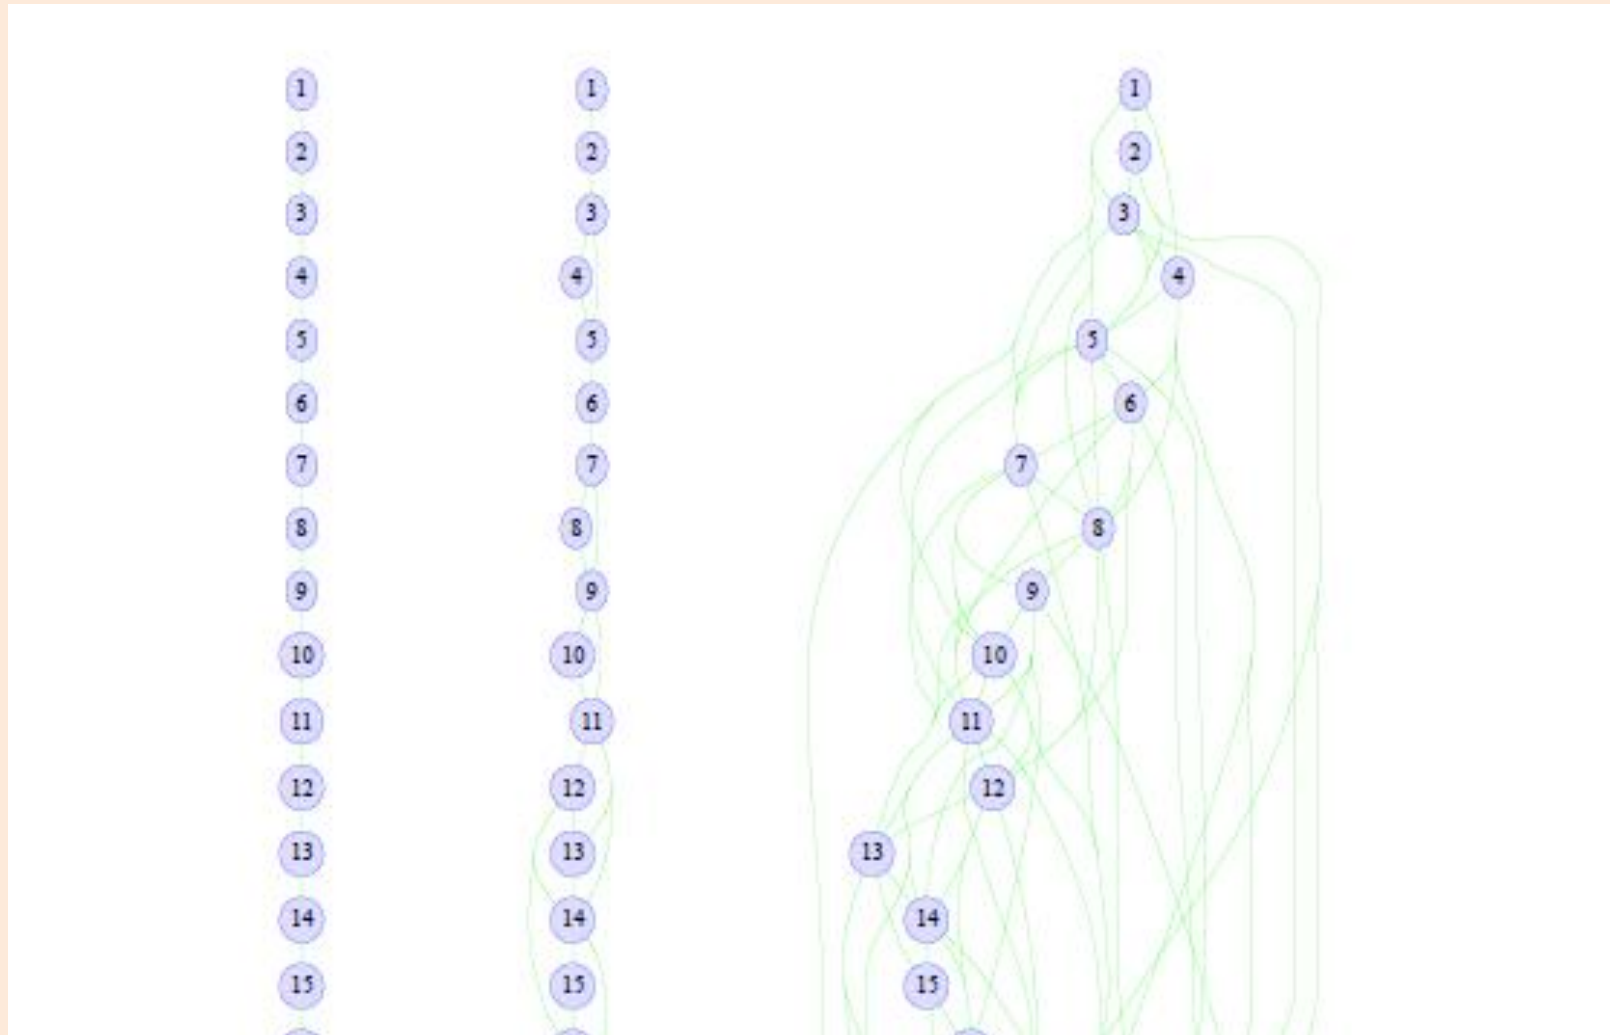
Dependency Networks

- A better approach is dependency networks:
  - For each variable 'j', make it the target in a supervised learning problem.

$$ X = \begin{bmatrix} | & | & | & | & | \\ x^1 & x^2 & x^3 & x^4 & x^5 \\ | & | & | & | & | \end{bmatrix} \implies \bar{X} = \begin{bmatrix} | & | & | & | \\ x^1 & x^2 & x^3 & x^5 \\ | & | & | & | \end{bmatrix} \quad y = \begin{bmatrix} | \\ x^4 \\ | \end{bmatrix} $$

  - Now we can use any feature selection method to choose j's "neighbours".
    - Forward selection, L1-regularization, ensemble methods, etc.

- Can capture conditional independence:
  - Might connect "sick" to "tacos", and "tacos" to "Tuesdays" (w/o sick-tacos).
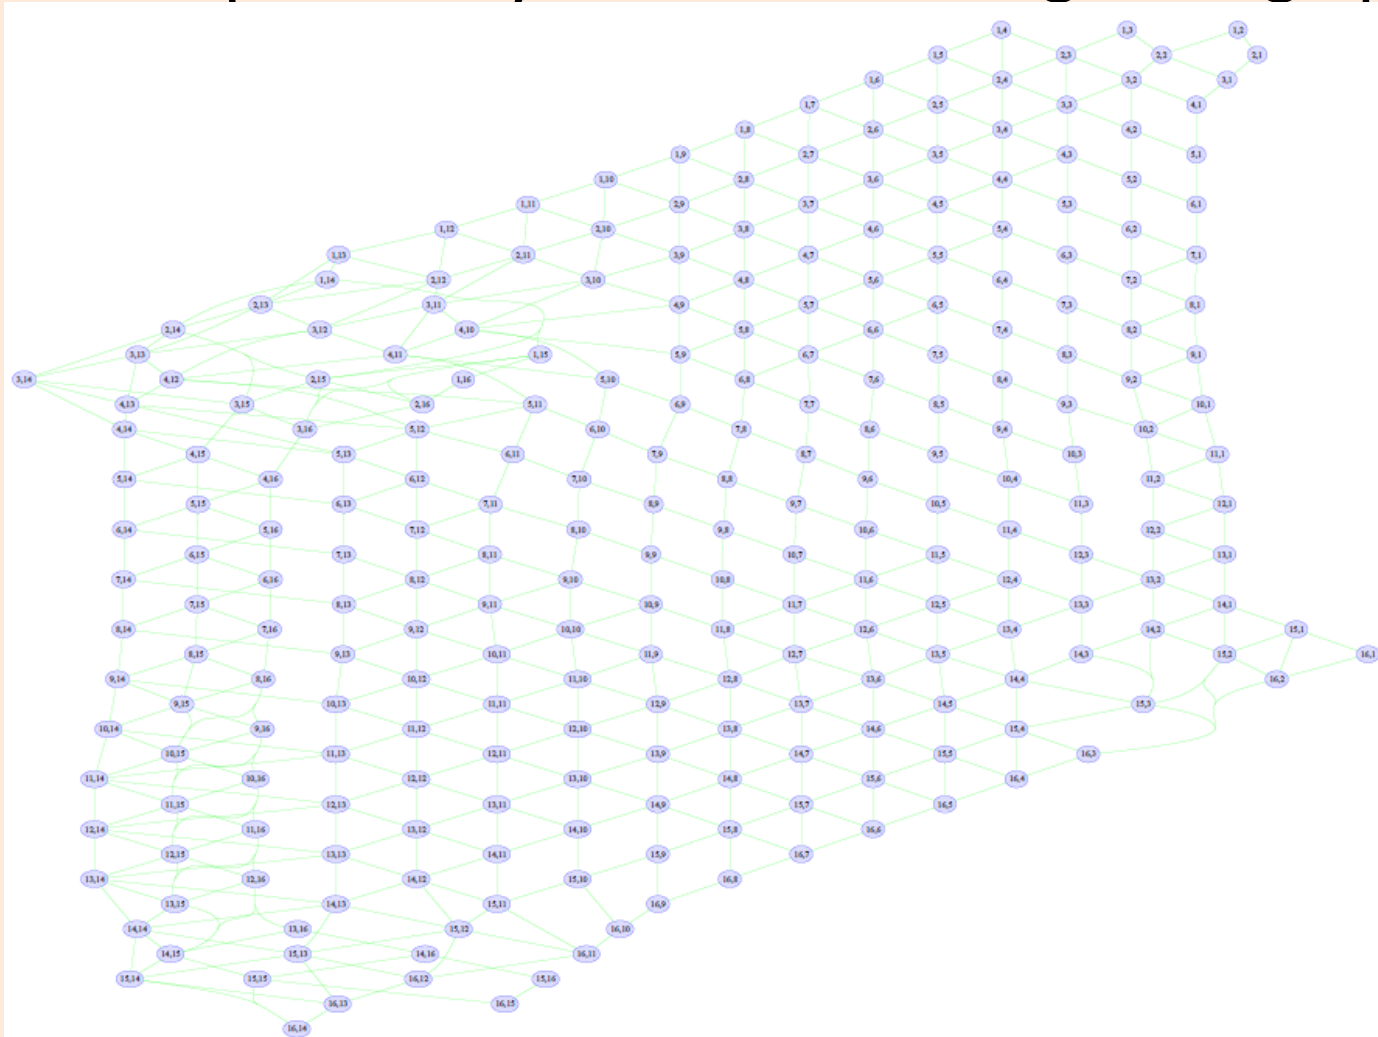
Dependency Networks

- Dependency network fit to Vancouver rain data (different λ values):

Dependency Networks

- Variation on dependency networks on digit image pixels:

Another popular structure learning method is the "PC" algorithm.