Learning Dependency Structures for Weak Supervision Models

Fred Sala, Paroma Varma, Ann He, Alex Ratner, Chris Ré

STANFORD DAWN

Snorkel and Weak Supervision

snorkel

A system for rapidly creating training sets with weak supervision

Ratner et al., Snorkel: "Rapid Training Data Creation with Weak Supervision", VLDB 2017.

Frequent use in industry!

Bach et al., "Snorkel DryBell: A Case Study in Deploying Weak Supervision at Industrial Scale", SIGMOD (Industrial) 2019.

The Snorkel/Weak Supervision Pipeline

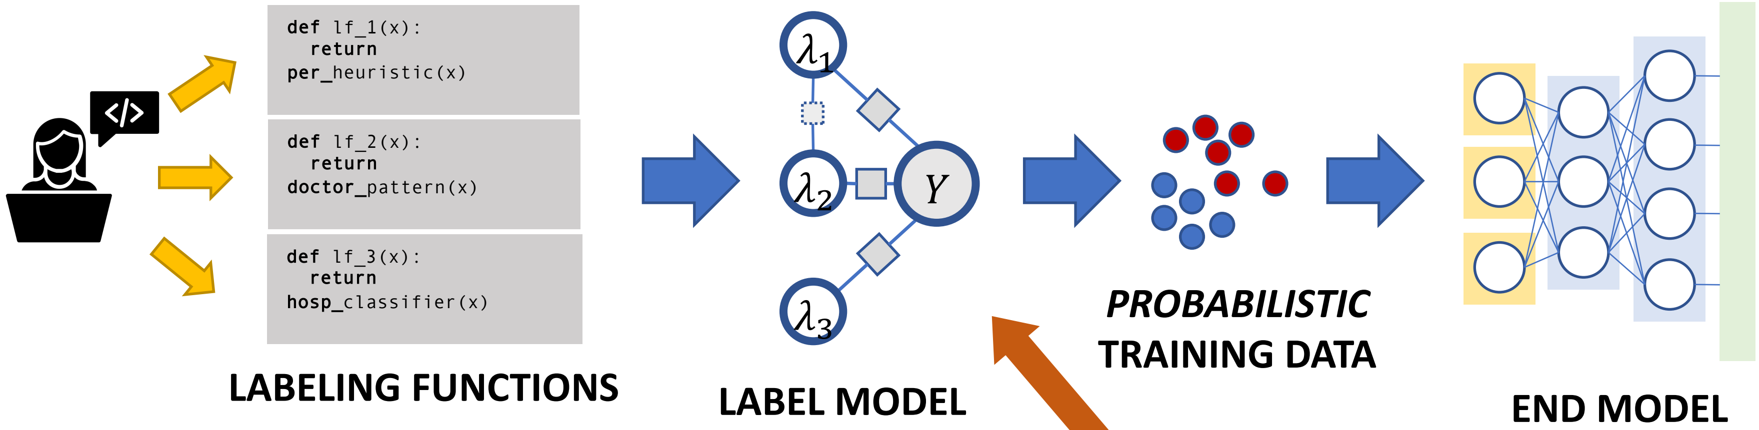

1 Users write *labeling functions* to noisily label data

2 We model the labeling functions' behavior to *de-noise* them

3 We use the *probabilistic* labels to train an *arbitrary end model*

Requires Dependency Structure!

Takeaway: No hand-labeled training data needed!

Model as Generative Process

```
def existing_classifier(x):
  return off_shelf_classifier(x)
```

→ "PERSON"

```
def upper_case_existing_classifier(x):
  if all(map(is_upper, x.split())) and \
off_shelf_classifier(x) == 'PERSON':
    return PERSON
```

→ "PERSON"

```
def is_in_hospital_name_DB(x):
  if x in HOSPITAL_NAMES_DB:
    return HOSPITAL
```

→ "HOSPITAL"

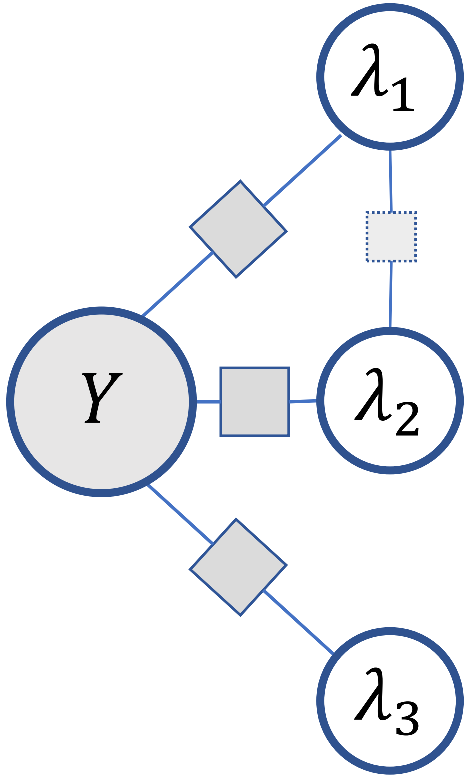

Problem: learn the parameters of this model (accuracies & correlations) without Y?

Solution Sketch: Using the covariance

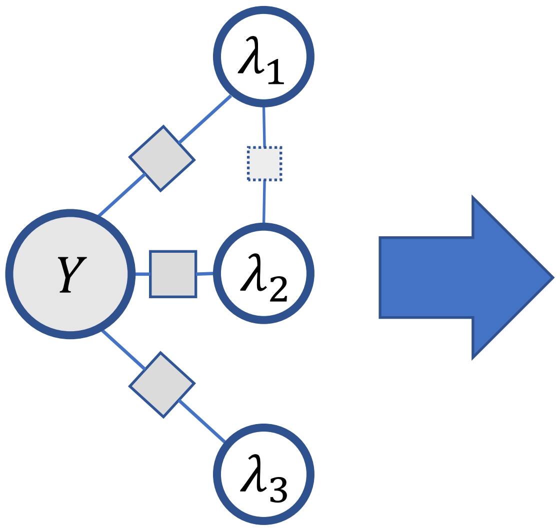

$\Sigma =$

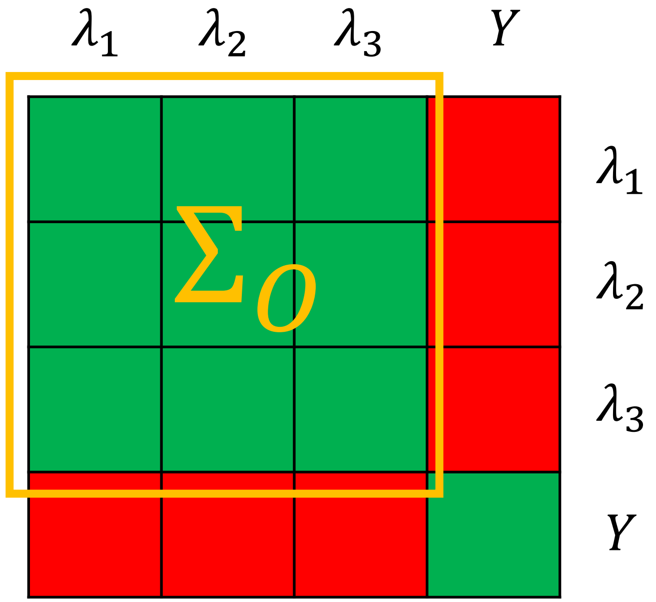

Can only observe **part** of the covariance…

Idea: Use graph-sparsity of the inverse

$$(\Sigma^{-1})_O = \Sigma_O^{-1} + zz^T$$

Observed overlaps

Low-rank parameters to solve for

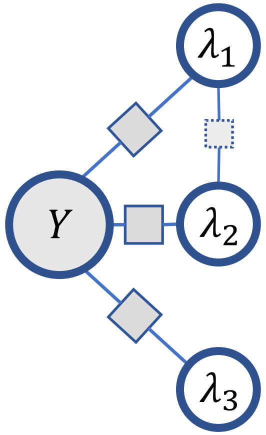

Is zero where corresponding pair of variables has no edge [Loh & Wainwright 2013]

Key: we must know the dependency structure

Idea: Use graph-sparsity of the inverse

$$(\Sigma^{-1})_O = \Sigma_O^{-1} + zz^T$$

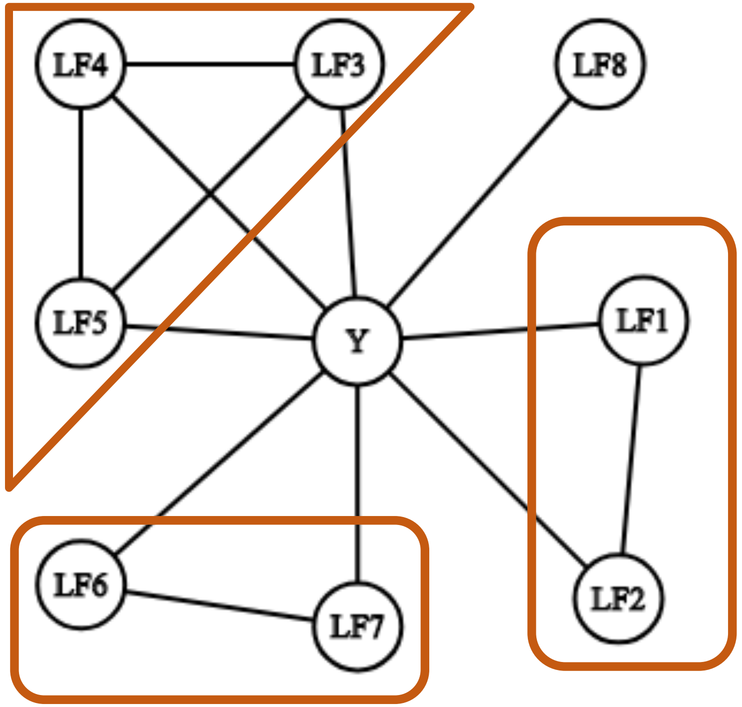

Example: 8 LFs

1 triangle, 2 pairs, 1 singleton

Inverse Encodes The Structure…

$$(\Sigma^{-1})_O = \Sigma_O^{-1} + zz^T$$

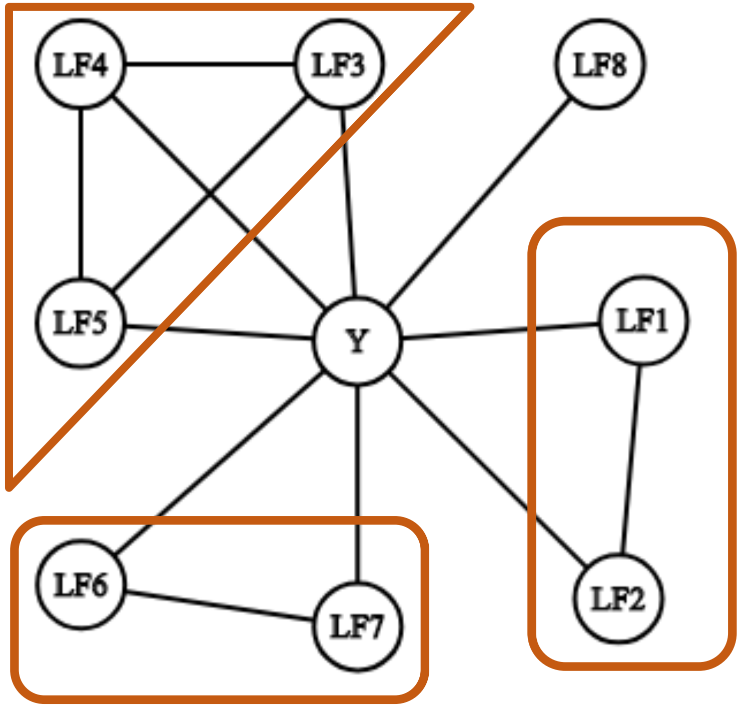

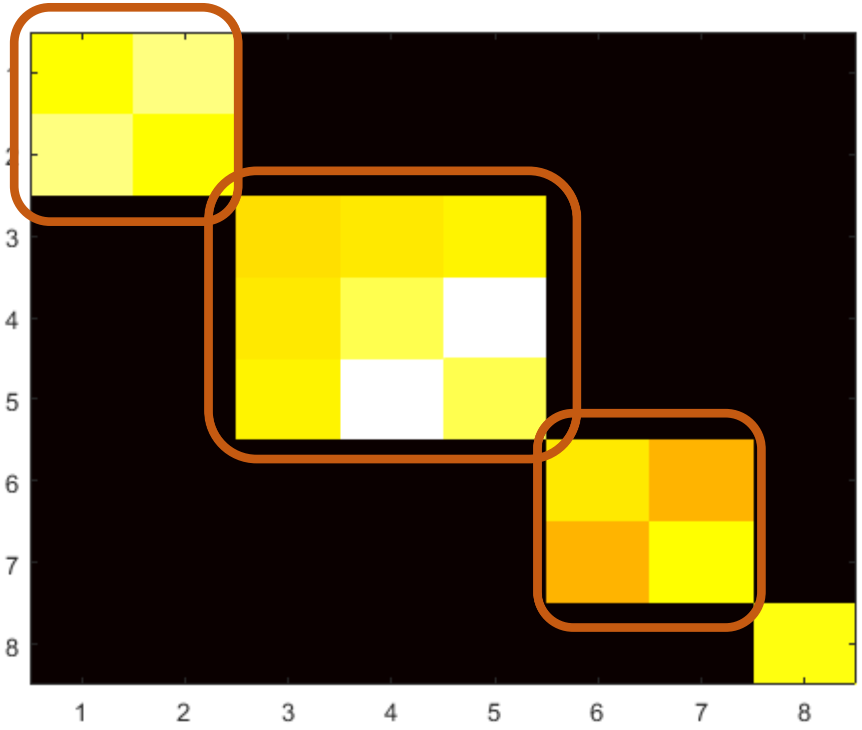

But Observed Matrix Doesn't

$$(\Sigma^{-1})_O = \Sigma_O^{-1} + zz^T$$

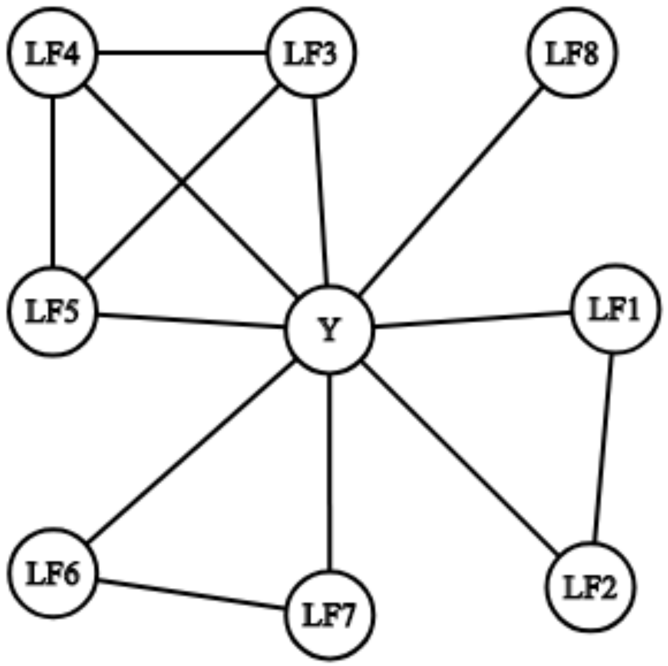

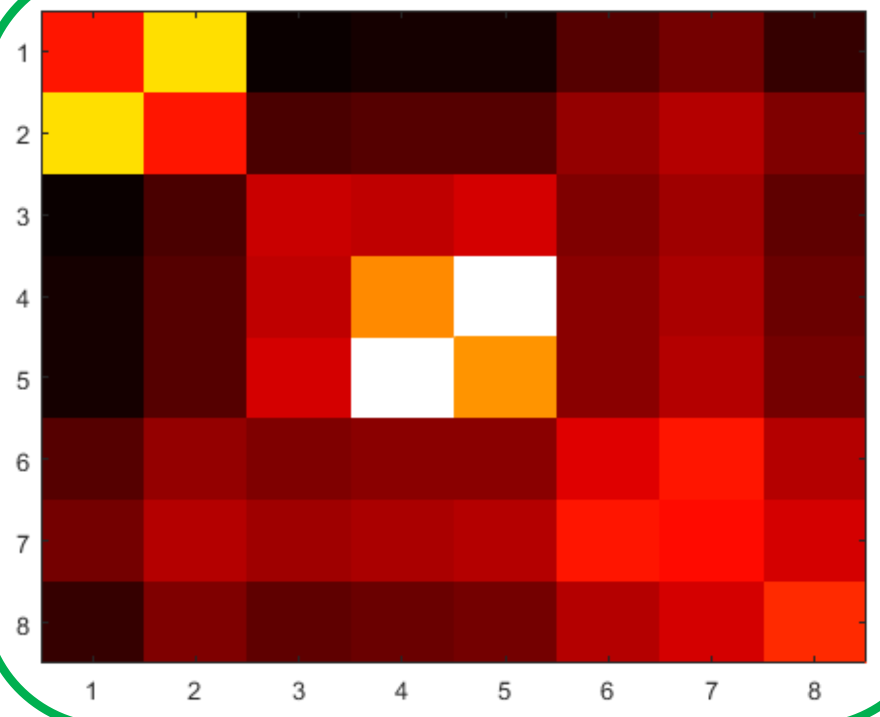

Need the Sparse Component…

Can we extract the sparse part?

$$\Sigma_O^{-1} = (\Sigma^{-1})_O - zz^T$$

Observed **Sparse** **Low-Rank**

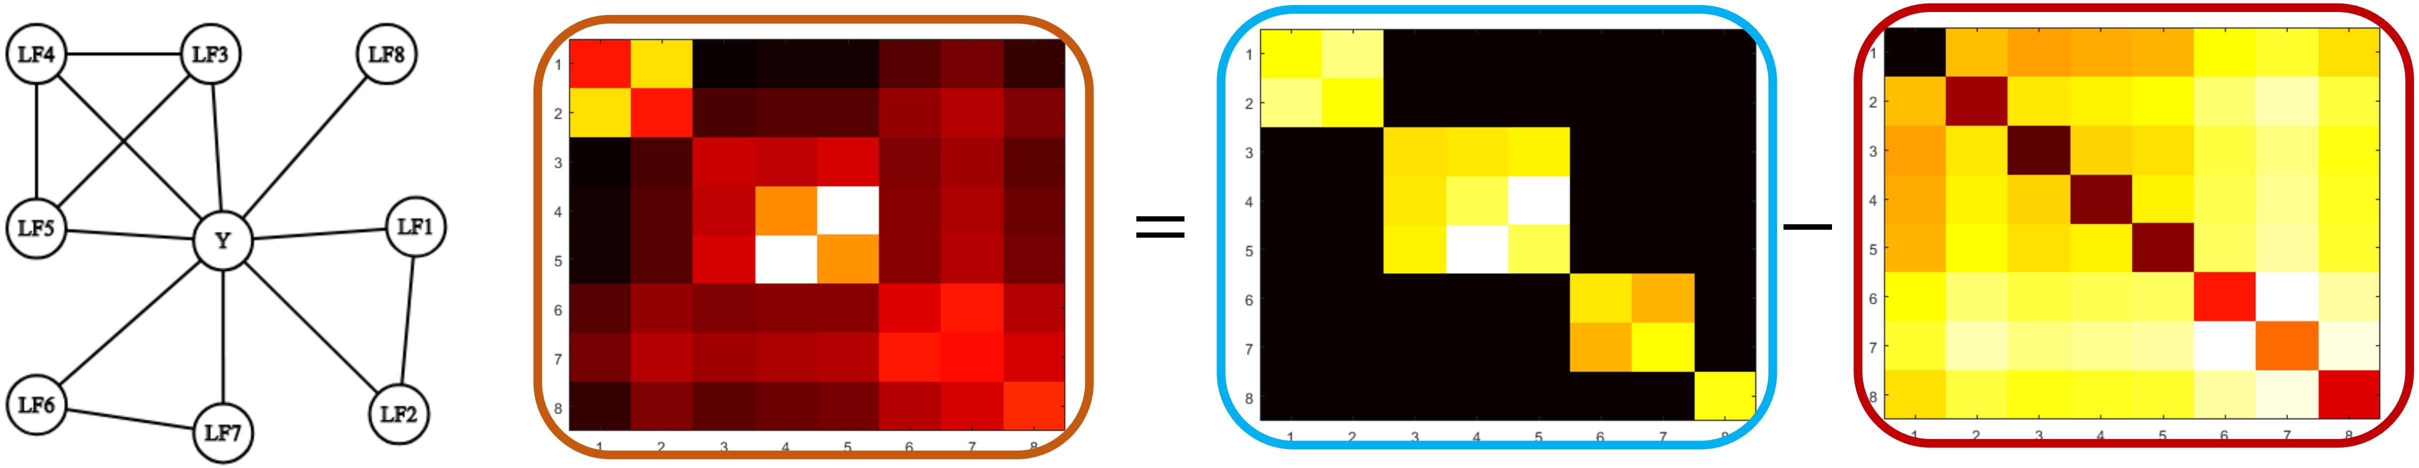

… & Robust PCA Recovers It!

Need to decompose:

$$\Sigma_O^{-1} = (\Sigma^{-1})_O - zz^T$$

Observed Sparse Low-Rank

Robust PCA : Decompose a matrix into **sparse** and **low-rank** components; **sparse** part contains graph structure

Convex optimization: $\text{argmin}_{S,L} ||\Sigma_O^{-1} - (S + L)||_2 + \lambda||S||_1 + \mu||L||_*$

Candes et al., "Robust Principal Components Analysis?", Chandrasekaran et al., "Rank-Sparsity Incoherence for Matrix Decomposition"

Theory Results: Sample Complexity

m is # of LFs, d is largest degree for a dependency

- **Prior work**: samples to recover WS dependency structure w. h. p.

S. Bach, B. He, A. Ratner, C. Ré, "Learning the structure of generative models without labeled data", ICML 2017.

$$\Omega(m \log m)$$

Doesn't exploit *d:* sparsity of the graph structure

- Recent application of RPCA for general latent-variable structure learning

C. Wu, H. Zhao, H. Fang, M. Deng, "Graphical model selection with latent variables", EJS 2017.

$$\Omega(d^2 m)$$

Linear in *m.*

Theory Results: Sample Complexity

m is # of LFs, d is largest degree for a dependency

Ours: for $\tau < 1$, an eigenvalue decay factor in blocks of LFs

$$\Omega(d^2 m^\tau)$$

Ours: When there is a dominant block of correlated LFs

$$\Omega(d^2 \log m)$$

Idea: exploit sharp concentration inequalities on sample covariance matrix Σ_O via the *effective rank* [Vershynin '12]

Application: Bone Tumor Task

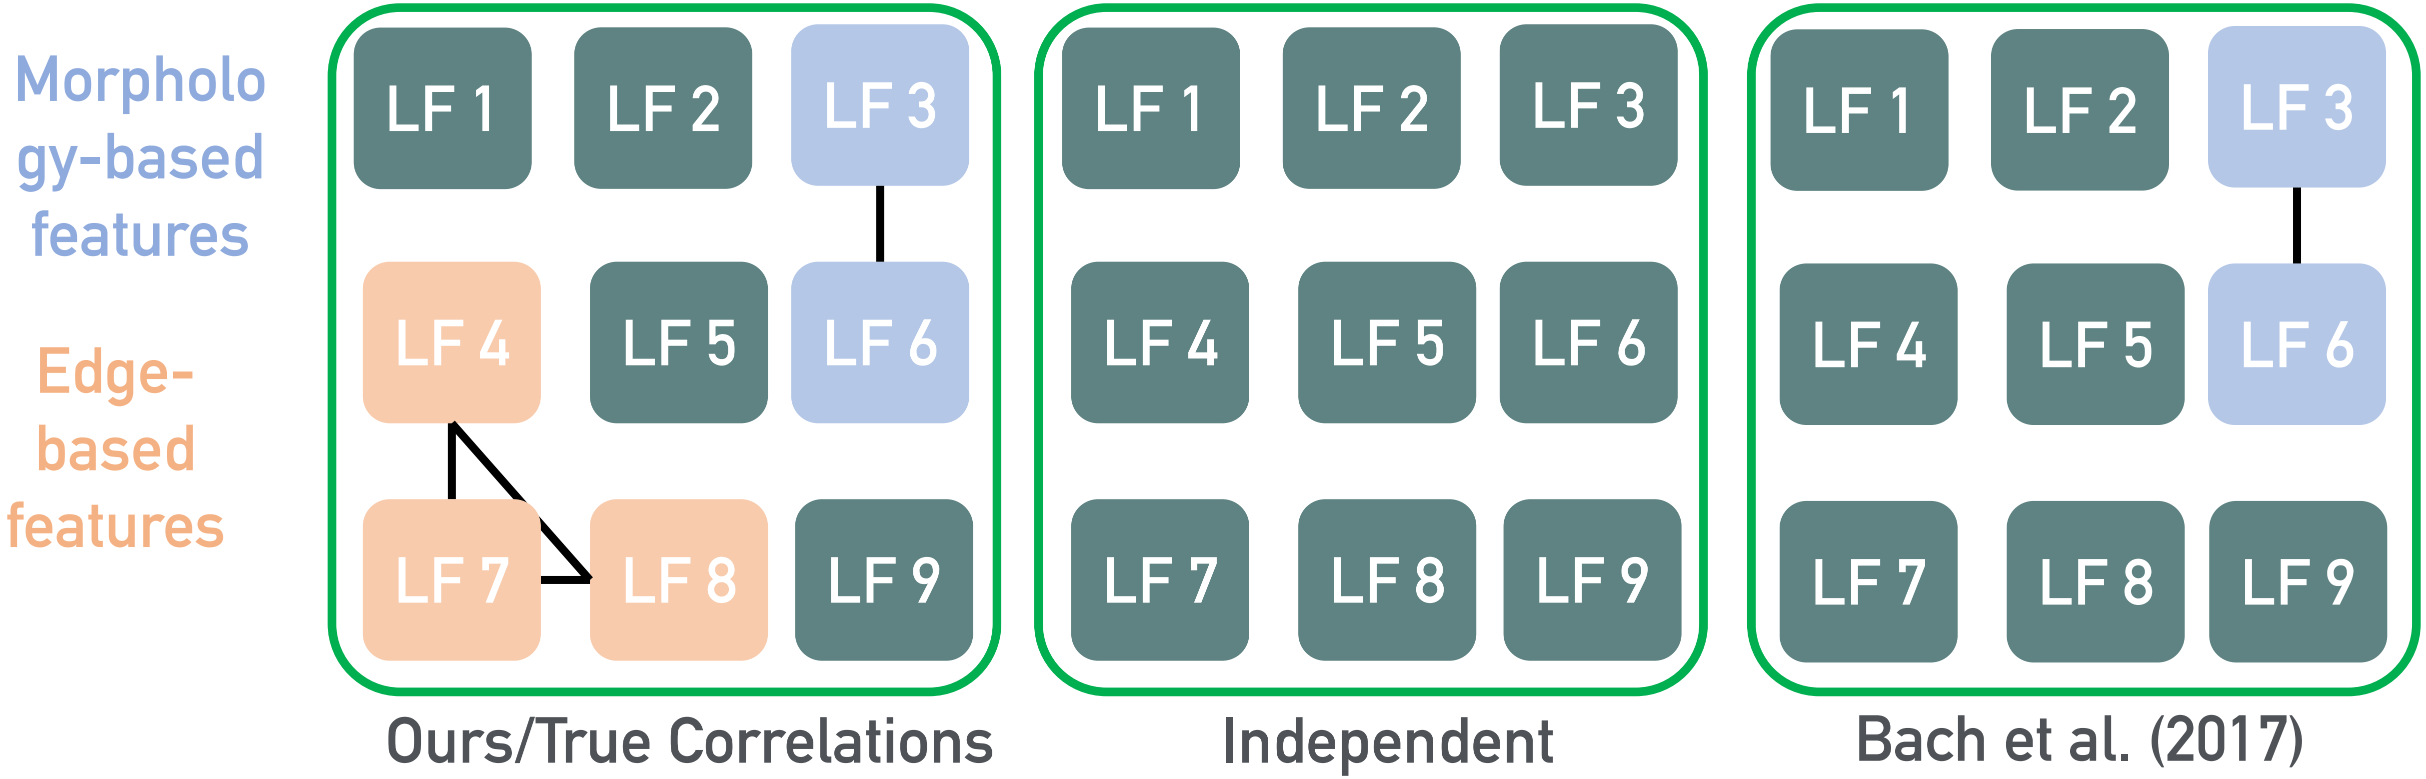

We pick up all the edges---

+4.64 F1 points, over indep., + 4.13 over Bach et al.

More Resources

- **Blog Post:** Intro to weak supervision
  https://dawn.cs.stanford.edu/2017/12/01/snorkel-programming/

- **Blog Post:** Gentle Introduction to Structure Learning
  https://dawn.cs.stanford.edu/2018/06/13/structure

- **Software**: https://github.com/HazyResearch/metal

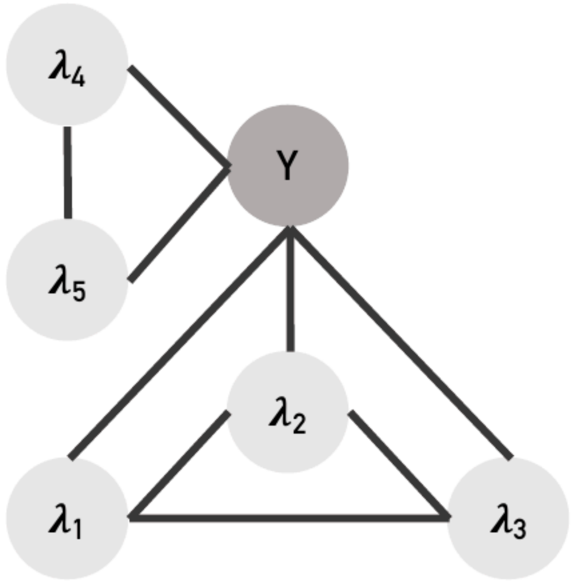

Fred Sala: https://stanford.edu/~fredsala