Machine Learning for Data Exploration and Generation

Charles Sutton
University of Edinburgh and the Alan Turing Institute
13 June 2017

http://bit.ly/sutton-ml-exploration

Google: “`charles sutton`” `talks` —> find this talk

THE UNIVERSITY *of* EDINBURGH informatics   The Alan Turing Institute   EPSRC Engineering and Physical Sciences Research Council

Prediction: A small part of a big picture

CRISP-DM

[Chapman et al 2000]

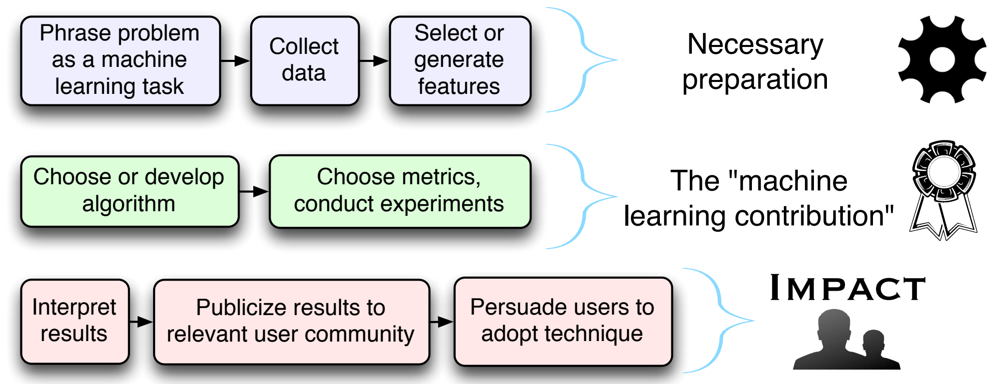

Research process

[Wagstaff, ICML 2012]

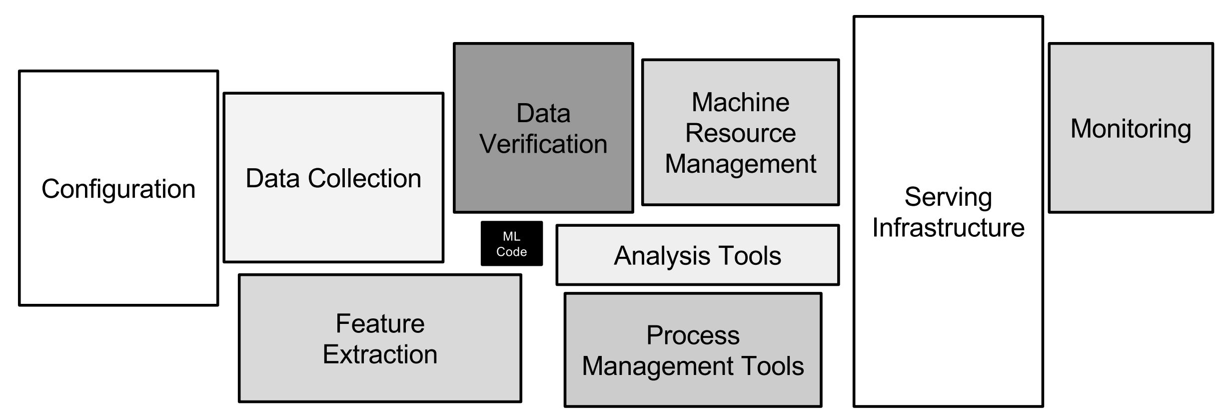

System level

[Sculley et al, NIPS 2015]

“Making data science easier”: an application area for machine learning!

Towards an Artificial Intelligence for Data Science

Data understanding

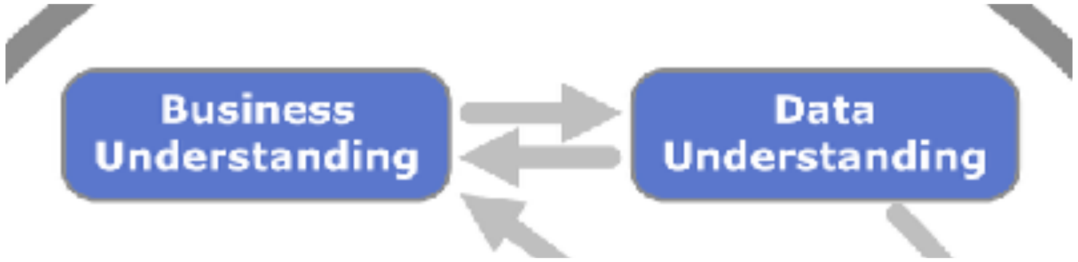

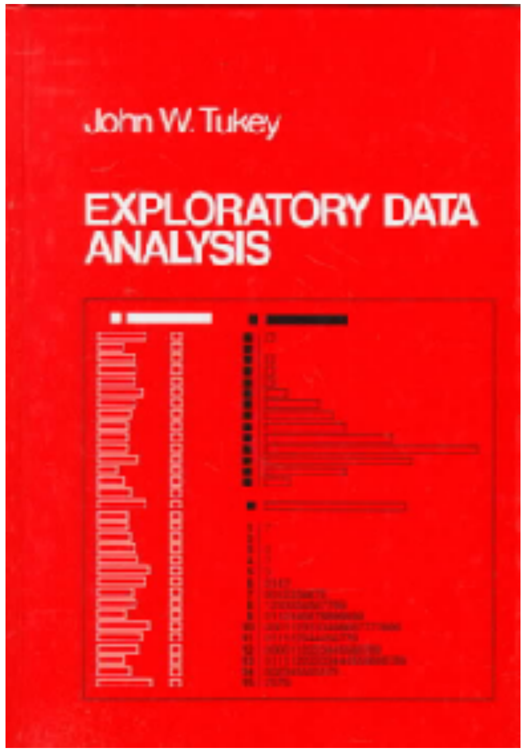

When you get a new data set….

- What's in it?
- What's wrong with it?
- What should I do with it?

Automating exploratory data analysis?

Contradiction in terms?

A task in visual analytics

Scalability a challenge

Our theme

Summarise data with probabilistic ML

Visualize resulting patterns

Patterns are "first class citizens" of model

Exploratory data analysis

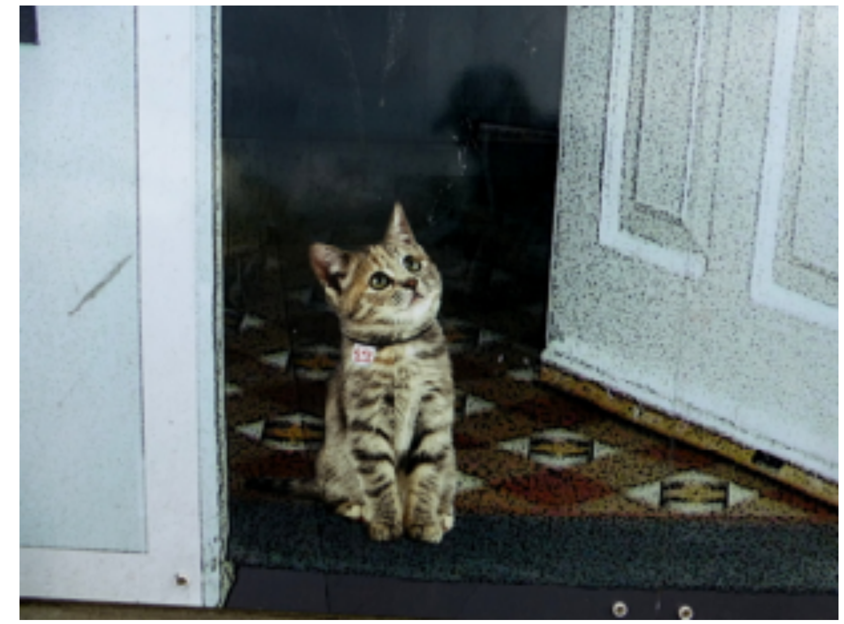

Data analysts are like cats.

1. Want to explore their data
2. Don't know what they want.

Machine learning for analysts

Whose information need is not explicit

Whose domain knowledge is difficult to encode

Explore data via learned patterns

… not just for dummies!

Mining Patterns

[Fowkes & Sutton, KDD 2016]

[Fowkes & Sutton, PKDD 2016]

http://bit.ly/sutton-ml-exploration

Association Rules (def'n)

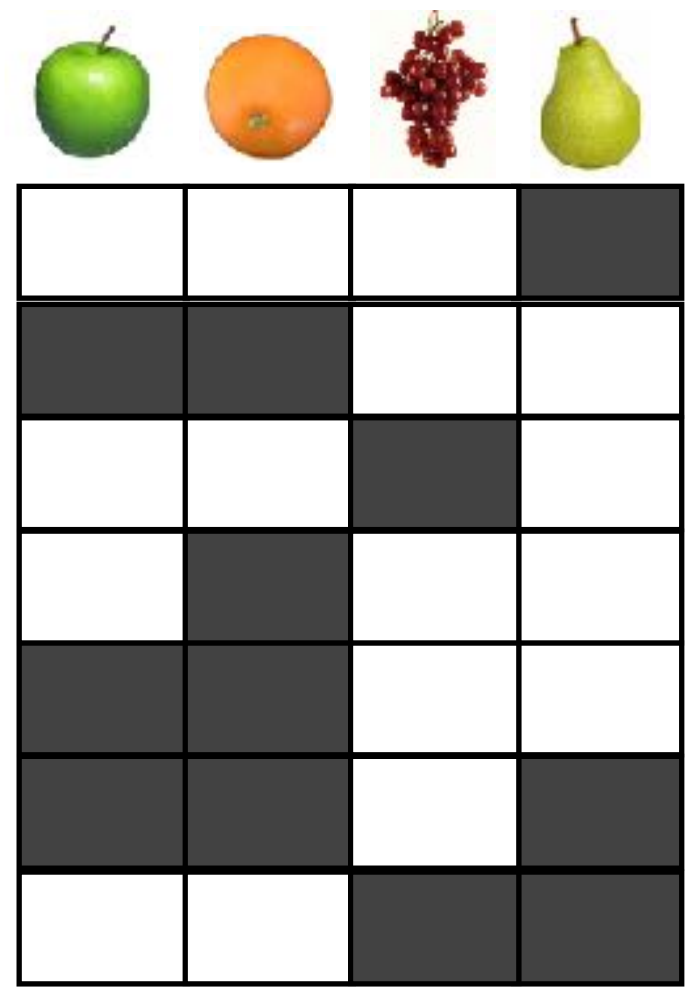

Database of transactions

Association rule mining:

Find set of all rules

that have

Frequent itemsets:

Why? Exploratory data analysis

Association Rules (alg)

1. Identify all frequent item sets

via exhaustive search (APriori, FP-Growth, etc.)

2. For each item set, consider all possible partitions

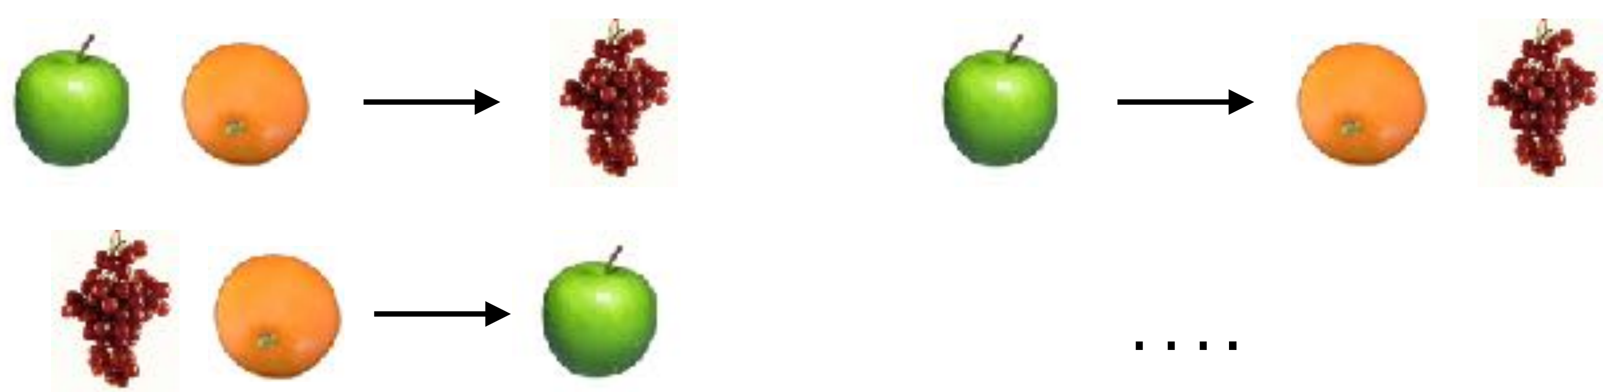

3. Rank the resulting list (e.g., by confidence) and enjoy

Pathologies

List of association rules unwieldy, difficult to understand

Procedure as a whole is statistically incoherent.

Essentially just repeated counting

Redundancy

If { 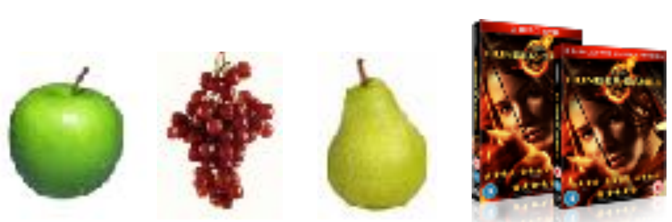 } frequent,

so are all 14 nontrivial subsets.
(Association rules "filter" item sets)

"Free riders"

If both of these

have support >> M and **independent**

{  }

usually still support > M

(Confidence and lift do not fix this!)

Rare itemsets

Strongly associated but rare:
Not a frequent itemset

{ 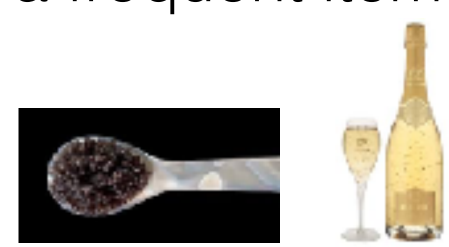 }

e.g., champagne, caviar

[Hastie et al., 2009]

Alternative: *Interesting* Itemsets

Optimise the collection of itemsets *as a whole,* rather than each in isolation

transaction X

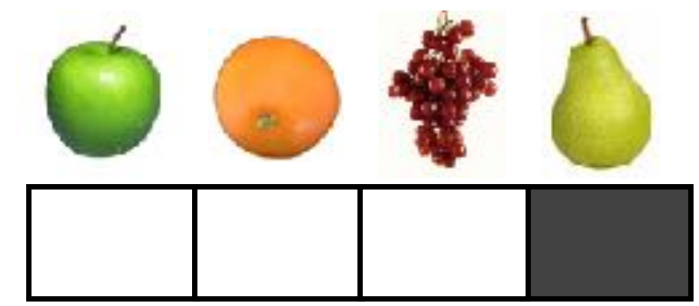

define probability model

$$p(X|\mathcal{I})$$

e.g., $\mathcal{I} = \{\{$  $\},\{$  $\}\}$

choose $\mathcal{I}$ to best fit data

$\mathcal{I}$ are the *interesting* itemsets

(unlike frequent itemsets, these are suitable for data analysis)

Interesting Sequence Mining

define a goodness measure on a *set* of patterns

Minimum description length

[Vreeken et al, 2011; Tatti and Vreeken, 2012; Lam et al 2014]

Use patterns to define a compression algorithm for database

Search for patterns that best compress

Probabilistic methods

[Fowkes and Sutton, KDD 2016, PKDD 2016]

Use patterns to define a probability distribution over database

Search for patterns that maximise database probability

(actually isomorphic; see MacKay, 2003)

Sequences more meaningful, less redundant

Also, see tiling: *[Geerts, Goethals, and Mielikäinen, 2004]*

Model

To sample a transaction,

1. For each itemset, sample

$$z_S \sim \textbf{Bernoulli}(\pi_S).$$

2. Deterministically set

$$X = \bigcup_{z_s = 1} S.$$

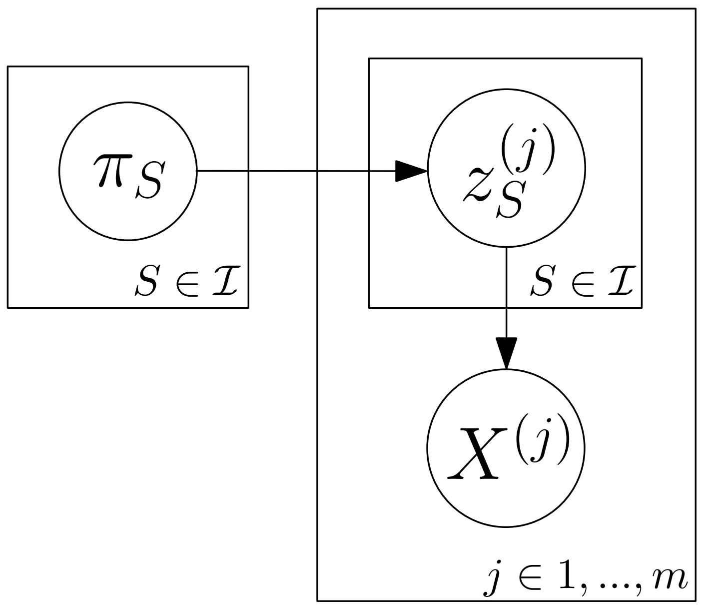

Parameters:

$\mathcal{I}$ Collection of "interesting" itemsets

$\pi_S \in [0, 1]$ for each $S \in \mathcal{I}$

probability of occurrence

Inference / Learning

Infer z from X

$$\max_{\mathbf{z}} \sum_{S \in \mathcal{I}} z_S \ln\left(\frac{\pi_S}{1-\pi_S}\right) + \ln(1-\pi_S)$$
$$\text{s.t.} \sum_{S | i \in S} z_S \geq 1 \qquad \forall\, i \in X$$
$$z_S \in \{0,1\} \qquad \forall\, S \in \mathcal{I}$$

NP-hard but submodular
(weighted set cover)
use greedy algorithm

Infer $\mathcal{I}$

Structural EM

Propose new itemset S

Add S to model

Re-infer $\mathcal{Z}$

Check if cost improved

“Implicit regularization”

Redundancy

Average distance between itemsets in one ranked list
(symmetric distance, higher is better)

	Plants	Mammals	ICDM	Uganda
Interesting Itemsets	**3.50**	**5.30**	**3.66**	**3.72**
KRIMP	**1.53**	**2.02**	**2.22**	**2.24**
CHARM	**1.53**	**1.52**	**1.47**	**1.45**

Facebook posts

IIM	MTV	KRIMP
soul, rest, peace	heal, jesus, amen	whi, ?
chris, brown	god, amen	?, !
bebe, cool	2, 4	2, 4
airtel, red	whi, ?	wat, ?
everi, thing	god, heal	time, !
time, wast	2, !	soul, rest, peace

[courtesy John Quinn, UN Global Pulse]

Trending

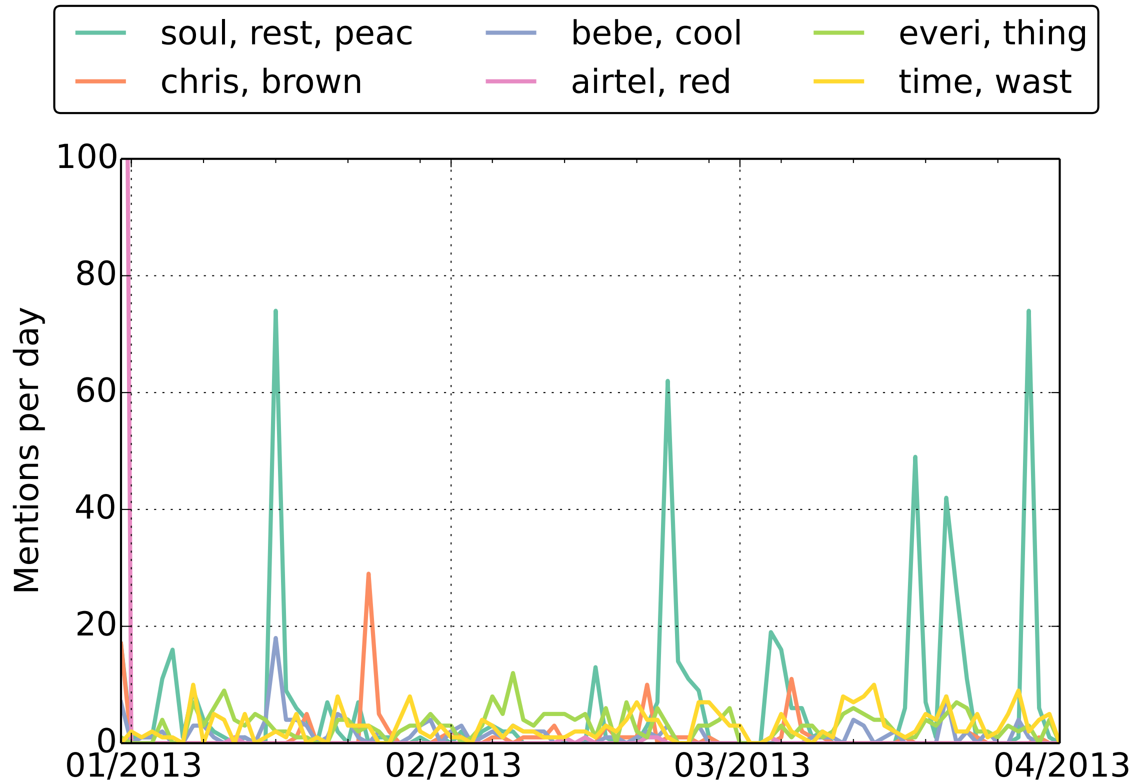

IIM

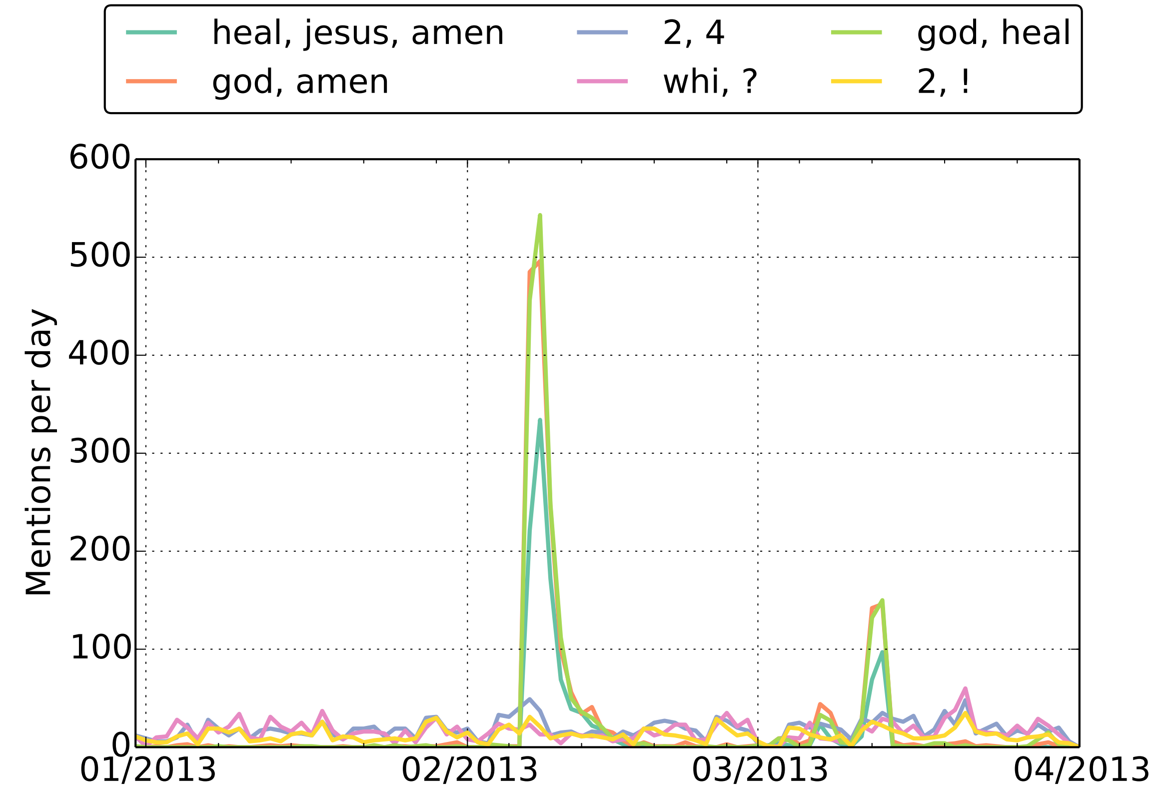

MTV

Facebook posts from public Ugandan pages

[courtesy John Quinn, UN Global Pulse]

Plants

KRIMP

Plants

IIM

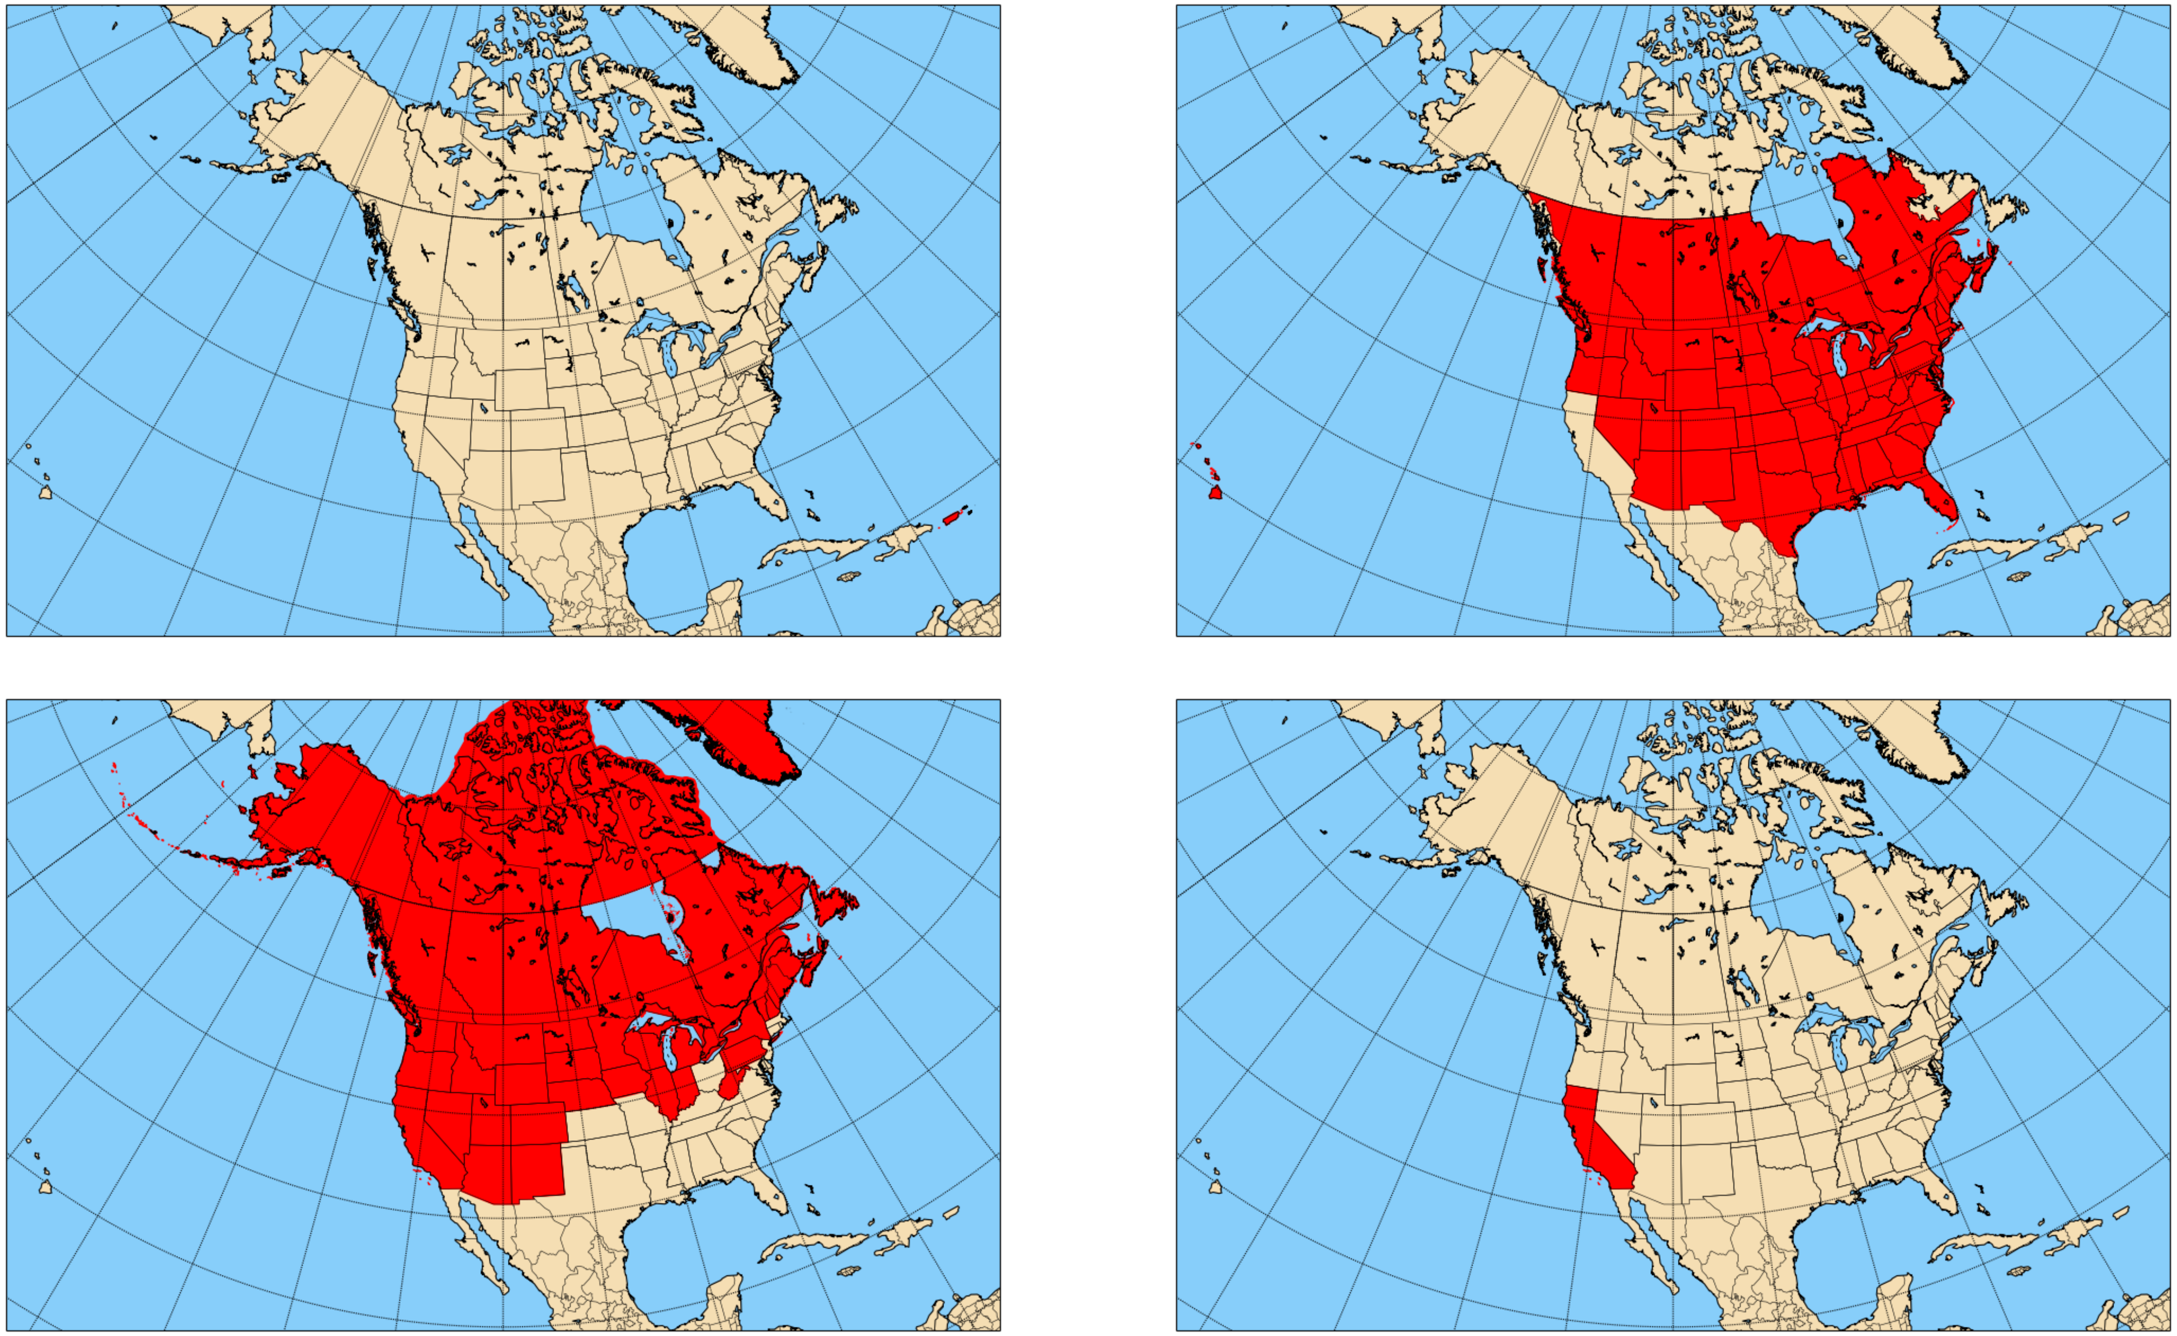

Frequent Sequence Mining

Return all patterns with >= given support

[Agrawal and Srikant, 1995; Wang and Han, 2004]

Support of pattern: Number of database sequences that contain it

b d b a f e c
b c e a
e d a f c
a e f b
b d a e f c

Database of sequences

→

d a f c
b a f c
a e
b e
e c
...

Sequence patterns

(e.g. minimum support = 3*)*

Problem: Frequent can be trivial!

Fundamental Pathologies

Truncation

d a f c → a c

Real pattern

Could be returned (more frequent!)

Spurious correlation

Support(**a**) = 90%

Support(**d**) = 90%

… but independent …

→ d a

Pattern at 81% min_support

Effect: Redundant list of patterns

Freerider

a f c real pattern

Support(**d**) = 90%

… but independent …

→ a d f c

for high enough min_support

Probabilistic Sequence Mining

[Fowkes and Sutton, KDD 2016]

Define a distribution P(database | patterns)

$\mathcal{I}$

[b c e]	: 0.1,	**0.6**
[d f]	: 0.7,	**0.3**
[d f]	: 0.8,	**0.2**
[e f]	: **0.8**,	0.1

Sequence patterns (with probabilities)

Sample

Inclusion variables:

z_1: **1**

z_2: **1**

z_3: **1**

z_4: **0**

Interleave randomly

X b d c e d f f

Sampled database sequence

$P(X, z|\mathcal{I})$

probability of generating X, z from this process

Probabilistic Sequence Mining

[Fowkes and Sutton, KDD 2016]

Model:

$$p(X, \mathbf{z}|\mathbf{\Pi}) = \frac{1}{|\mathcal{P}|} \prod_{S \in \mathcal{I}} \prod_{m=0}^{|\boldsymbol{\pi}_S|-1} \pi_{S_m}^{[z_S = m]}$$

Inference: Determine $z|X, \mathcal{I}$

X b d c e d f f

cover these?

or these?

$\mathcal{I}$

[b c e] : 0.1, 0.6
[d f] : 0.7, 0.3
[d f] : 0.8, 0.2
[e f] : 0.8, 0.1

Use greedy algorithm to

$$\max_z \log p(z|X, \mathcal{I})$$

(extension of weighted set cover)

Probabilistic Sequence Mining

[Fowkes and Sutton, KDD 2016]

Output of inference

	z [b c e]	[d f]	[d f]	[e f]
b d c e d f f	1	1	1	0
e e d f f f	0	1	0	1
d f d d f f	0	1	1	1

Learning step: Infer $\mathcal{I}$

- Update probabilities
  - (average of z)
- Propose new patterns
  - Add to model
  - See if probability increases

$\mathcal{I}$

[b c e]	: 0.3, 0.7
[d f]	: 0.0, 1.0
[d f]	: 0.7, 0.3
[e f]	: 0.3, 0.7

Formally: Structural Expectation Maximization

Application to Software Engineering

[Fowkes & Sutton, FSE 2016]

http://bit.ly/sutton-ml-exploration

Modern development is layers of libraries

Average Java file on Github:

Imports from **2.1** packages outside project

45% of files import an external package

(Not counting `java.*` `javax.*` `sun.*`)

Github Java corpus (Allamanis and Sutton, 2013)

13000+ projects with at least one fork, 2M+ Java files

http://groups.inf.ed.ac.uk/cup/javaGithub/

(heuristic analysis)

API Mining

[Zhong et al, 2009; Dang et al 2013]

Library

e.g. TWITTER4J

Coding

```
private FinchTwitterFactory(Context context) {
    mContext = context;

    installHttpResponseCache();

    ConfigurationBuilder configurationBuilder = new ConfigurationBuilder();
    configurationBuilder.setOAuthConsumerKey(ConsumerKey.CONSUMER_KEY);
    configurationBuilder.setOAuthConsumerSecret(ConsumerKey.CONSUMER_SECRET);
    configurationBuilder.setUseSSL(true);
    Configuration configuration = configurationBuilder.build();
    mTwitter = new TwitterFactory(configuration).getInstance();
}
```

brk3 / finch

```
public Twitter getTwitterInstance() {
    ConfigurationBuilder cb = new ConfigurationBuilder();
    cb.setOAuthConsumerKey(Keys.consumerKey);
    cb.setOAuthConsumerSecret(Keys.consumerSecret);
    cb.setOAuthAccessToken(mSettings.getString("accessToken", null));
    cb.setOAuthAccessTokenSecret(mSettings.getString("accessSecret", null));
    TwitterFactory tf = new TwitterFactory(cb.build());
    return tf.getInstance();
}
```

jrupac/CleanTwitter

```
private void startOAuth() {
    ConfigurationBuilder configurationBuilder = new ConfigurationBuilder();
    configurationBuilder.setOAuthConsumerKey(Const.CONSUMER_KEY);
    configurationBuilder.setOAuthConsumerSecret(Const.CONSUMER_SECRET);
    twitter = new TwitterFactory(configurationBuilder.build()).getInstance();

    try {
        requestToken = twitter.getOAuthRequestToken(Const.CALLBACK_URL);
        Toast.makeText(this, "Please authorize this app!", Toast.LENGTH_LONG).show();
        this.startActivity(new Intent(Intent.ACTION_VIEW,
            Uri.parse(requestToken.getAuthenticationURL() + "&force_login=true")));
    } catch (TwitterException e) {
        e.printStackTrace();
    }
}
```

katahirado/tsubunomi

Corpus of client code

API Mining

```
ConfigurationBuilder.<init>
ConfigurationBuilder.setOAuthConsumerKey
ConfigurationBuilder.setOAuthConsumerSecret
ConfigurationBuilder.build
TwitterFactory.<init>
TwitterFactory.getInstance
```

```
TwitterFactory.<init>
TwitterFactory.getInstance
Twitter.setOAuthConsumer
Twitter.setOAuthAccessToken
```

API patterns

Documentation
Suggestion

Frequent Sequence Mining

Each transaction: client method

Each element: a method call to an API method

b d b a f e c

b c e a

e d a f c

a e f b

b d a e f c

Database of sequences

d a f c

b a f c

a e

b e

e c

...

Sequence patterns

(e.g. minimum support = 3*)*

For API Mining…

Top 10 API patterns from pure sequence mining (BIDE)

```
TwitterFactory.<init>
TwitterFactory.getInstance
```

```
TwitterFactory.<init>
Twitter.setOAuthConsumer
```

```
Status.getUser
Status.getText
```

```
auth.AccessToken.<init>
Twitter.setOAuthAccessToken
```

```
TwitterFactory.<init>
TwitterFactory.getInstance
Twitter.setOAuthConsumer
Twitter.setOAuthAccessToken
```

```
TwitterFactory.getInstance
Twitter.setOAuthConsumer
```

```
TwitterFactory.<init>
TwitterFactory.getInstance
Twitter.setOAuthConsumer
```

```
TwitterFactory.<init>
Twitter.setOAuthAccessToken
```

```
TwitterFactory.<init>
TwitterFactory.getInstance
Twitter.setOAuthAccessToken
```

```
TwitterFactory.getInstance
Twitter.setOAuthAccessToken
```

```
TwitterFactory.<init>
Twitter.setOAuthConsumer
Twitter.setOAuthAccessToken
```

Previous Approach: Cluster before/after

[Zhong et al, 2009; Dang et al 2013]

Probabilistic API Miner (PAM)

Interesting sequence mining for API mining

```
ConfigurationBuilder.<init>
ConfigurationBuilder.setOAuthConsumerKey
ConfigurationBuilder.setOAuthConsumerSecret
ConfigurationBuilder.setUseSSL
ConfigurationBuilder.build
TwitterFactory.<init>
TwitterFactory.getInstance
```

```
ConfigurationBuilder.<init>
ConfigurationBuilder.setOAuthConsumerKey
ConfigurationBuilder.setOAuthConsumerSecret
ConfigurationBuilder.setOAuthAccessToken
ConfigurationBuilder.setOAuthAccessTokenSecret
ConfigurationBuilder.build
TwitterFactory.<init>
TwitterFactory.getInstance
```

```
ConfigurationBuilder.<init>
ConfigurationBuilder.setOAuthConsumerKey
ConfigurationBuilder.setOAuthConsumerSecret
ConfigurationBuilder.build
TwitterFactory.<init>
TwitterFactory.getInstance
TwitterFactory.getOAuthRequestToken
RequestToken.getAuthenticationURL
```

Sequence database

```
private FinchTwitterFactory(Context context) {
    mContext = context;

    installHttpResponseCache();

    ConfigurationBuilder configurationBuilder = new ConfigurationBuilder();
    configurationBuilder.setOAuthConsumerKey(ConsumerKey.CONSUMER_KEY);
    configurationBuilder.setOAuthConsumerSecret(ConsumerKey.CONSUMER_SECRET);
    configurationBuilder.setUseSSL(true);
    Configuration configuration = configurationBuilder.build();
    mTwitter = new TwitterFactory(configuration).getInstance();
}
```

brk3 / finch

```
public Twitter getTwitterInstance() {
    ConfigurationBuilder cb = new ConfigurationBuilder();
    cb.setOAuthConsumerKey(Keys.consumerKey);
    cb.setOAuthConsumerSecret(Keys.consumerSecret);
    cb.setOAuthAccessToken(mSettings.getString("accessToken", null));
    cb.setOAuthAccessTokenSecret(mSettings.getString("accessSecret", null));
    TwitterFactory tf = new TwitterFactory(cb.build());
    return tf.getInstance();
}
```

jrupac/CleanTwitter

```
private void startOAuth() {
    ConfigurationBuilder configurationBuilder = new ConfigurationBuilder();
    configurationBuilder.setOAuthConsumerKey(Const.CONSUMER_KEY);
    configurationBuilder.setOAuthConsumerSecret(Const.CONSUMER_SECRET);
    twitter = new TwitterFactory(configurationBuilder.build()).getInstance();

    try {
        requestToken = twitter.getOAuthRequestToken(Const.CALLBACK_URL);
        Toast.makeText(this, "Please authorize this app!", Toast.LENGTH_LONG).show();
        this.startActivity(new Intent(Intent.ACTION_VIEW,
            Uri.parse(requestToken.getAuthenticationURL() + "&force_login=true")));
    } catch (TwitterException e) {
        e.printStackTrace();
    }
}
```

katahirado/tsubunomi

Corpus

Probabilistic sequence mining

Data

Target projects: 17 Java libraries, all that:

- Library source on Github
- Library in top 1000 Github projects
- Called by >50 other methods on Github
- At least 10k lines of `example/` code
- Total: Over 300k lines of example code

Client methods: all that called any targets

- 967 client projects
- Total: Over 4M lines of client code

Experimental Questions

- Quality
  - Match to “held-out” client code
  - Match to examples from library developers
  - Measure: sequence overlap, precision, recall
- Redundancy
  - Why? Ease of use, diversity
  - Measure: number of containing sequences
- All results averaged over the 17 libraries

Handwritten Examples

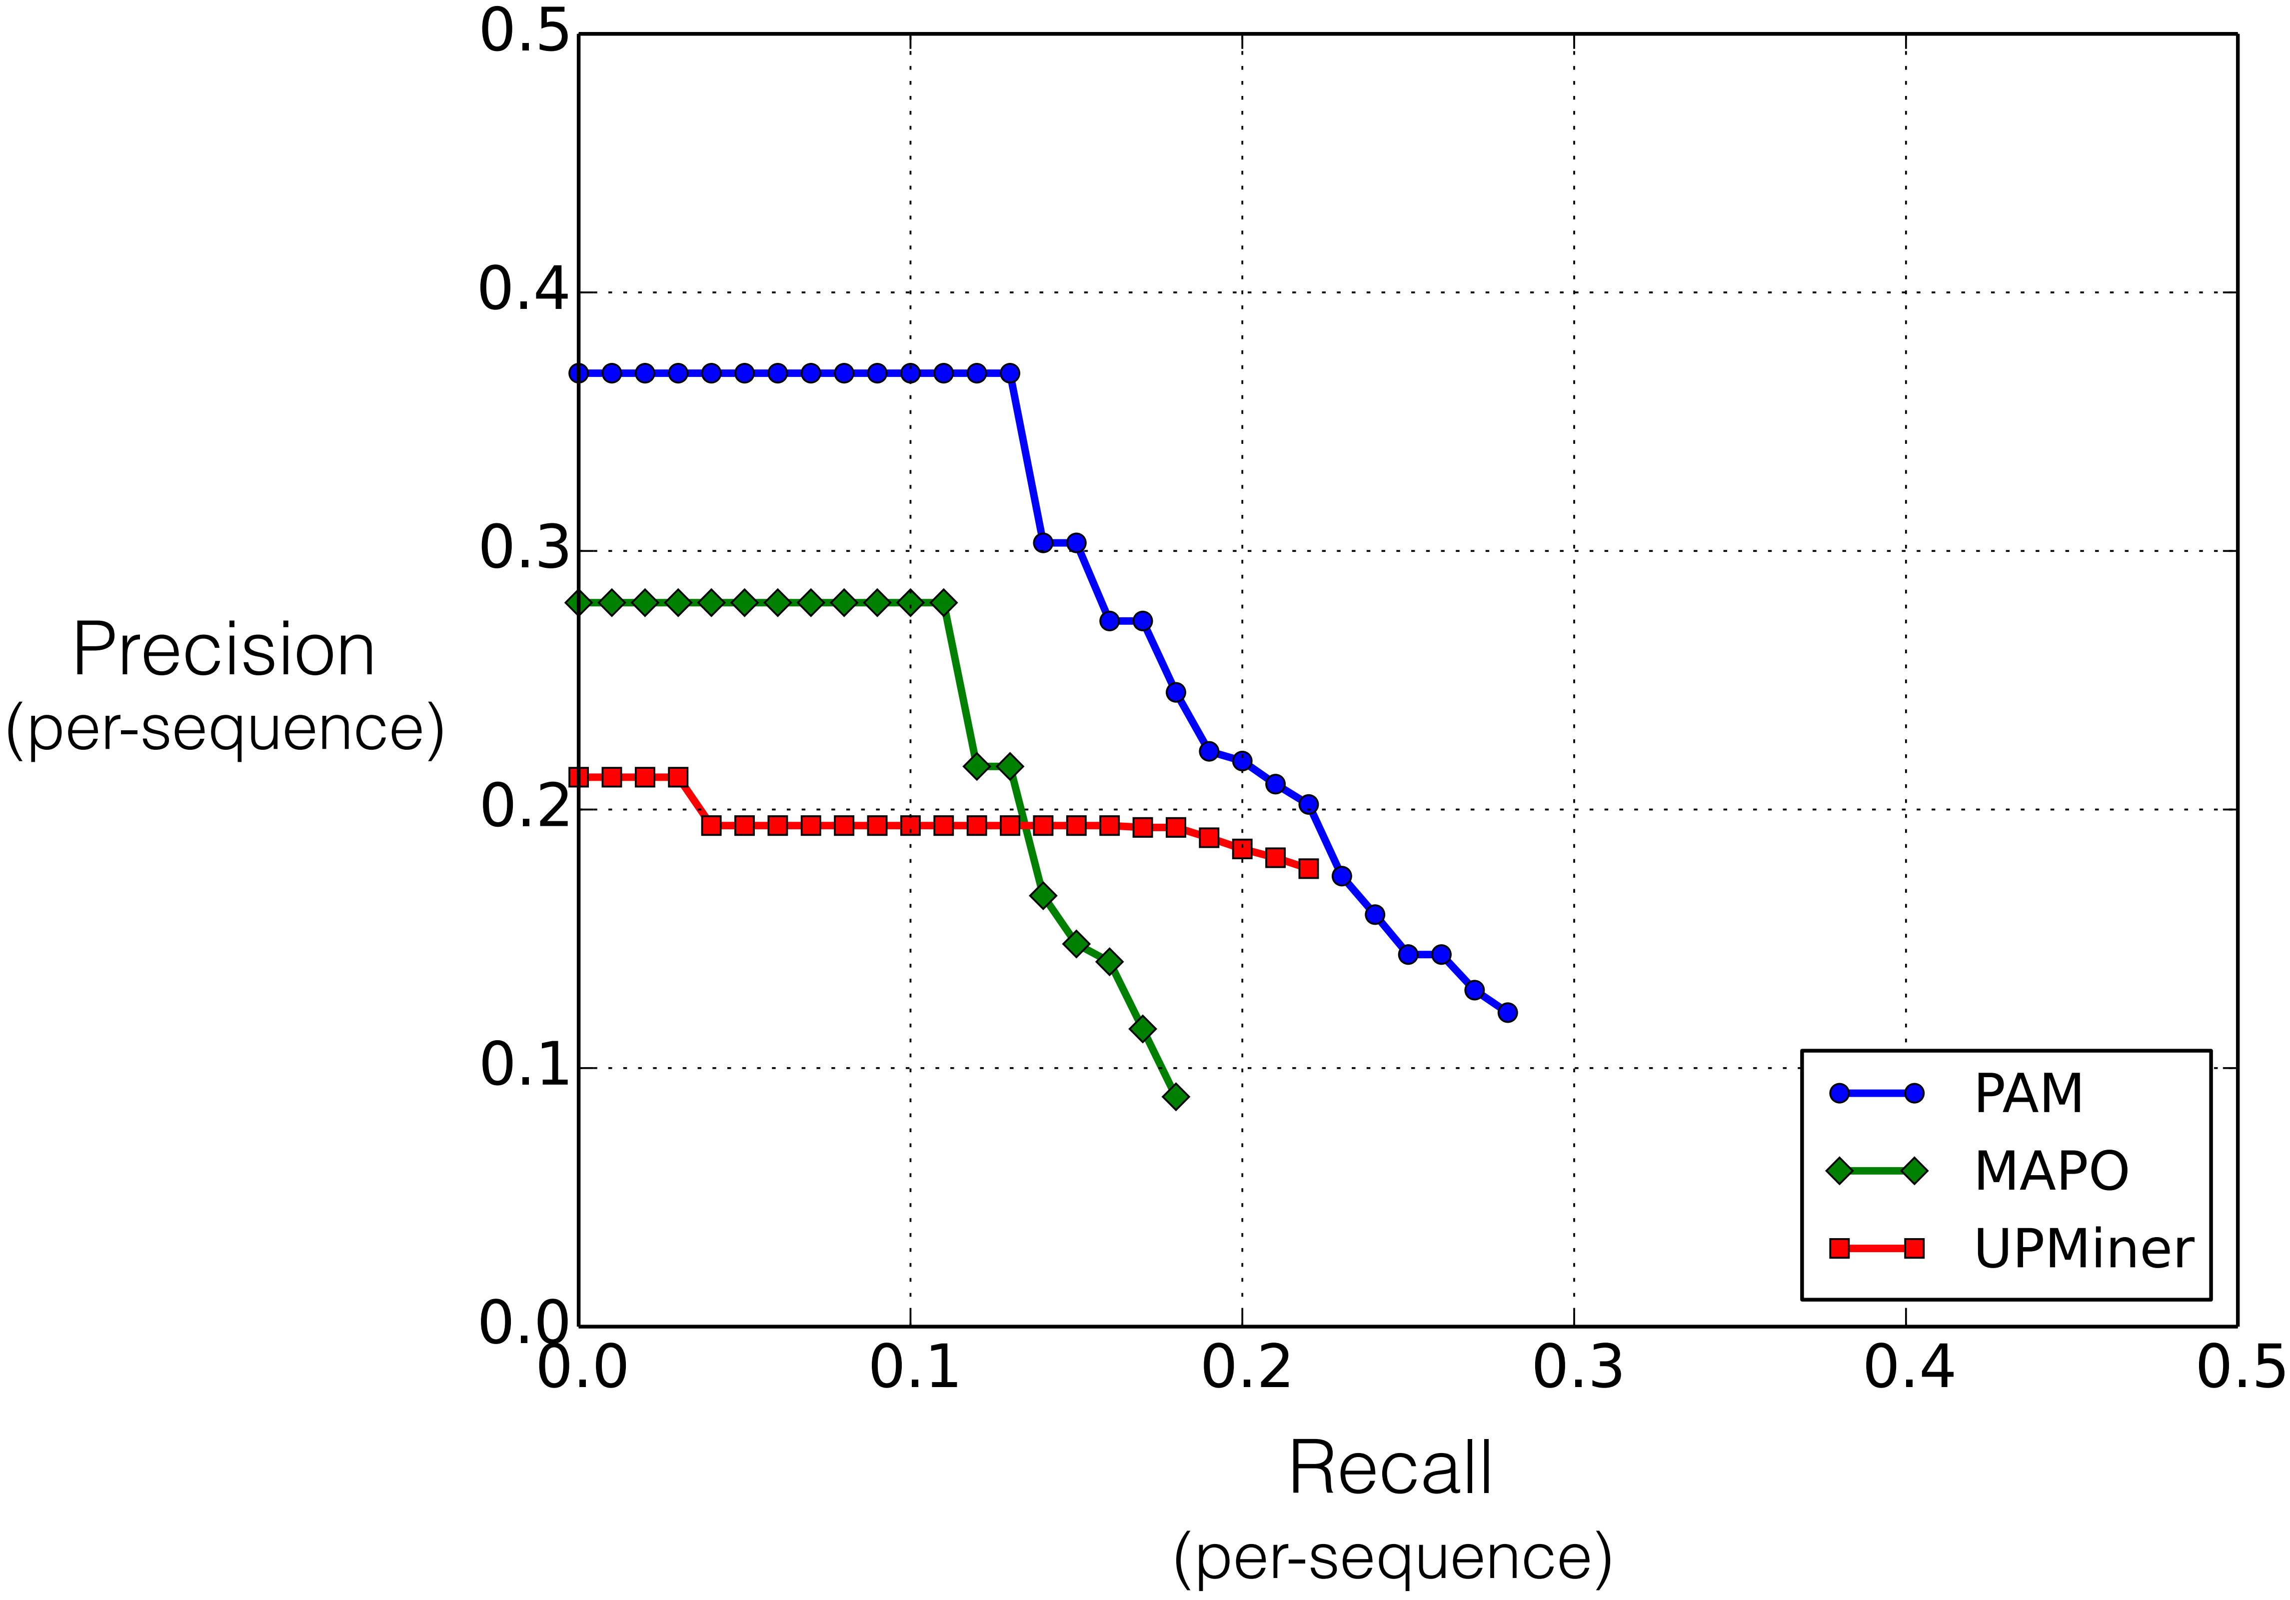

Why Low Recall?

- API mining bad, or examples incomplete?
  - Match test set to examples: is test set covered?
  - 73% of client API sequences not covered
  - 36% of examples used in client code
- Manual error analysis
  - 3 random projects, top 5 unmatching patterns
  - 7 referred to API method not in examples
  - 3 referred to API class not in examples

Prevalence in Client Code

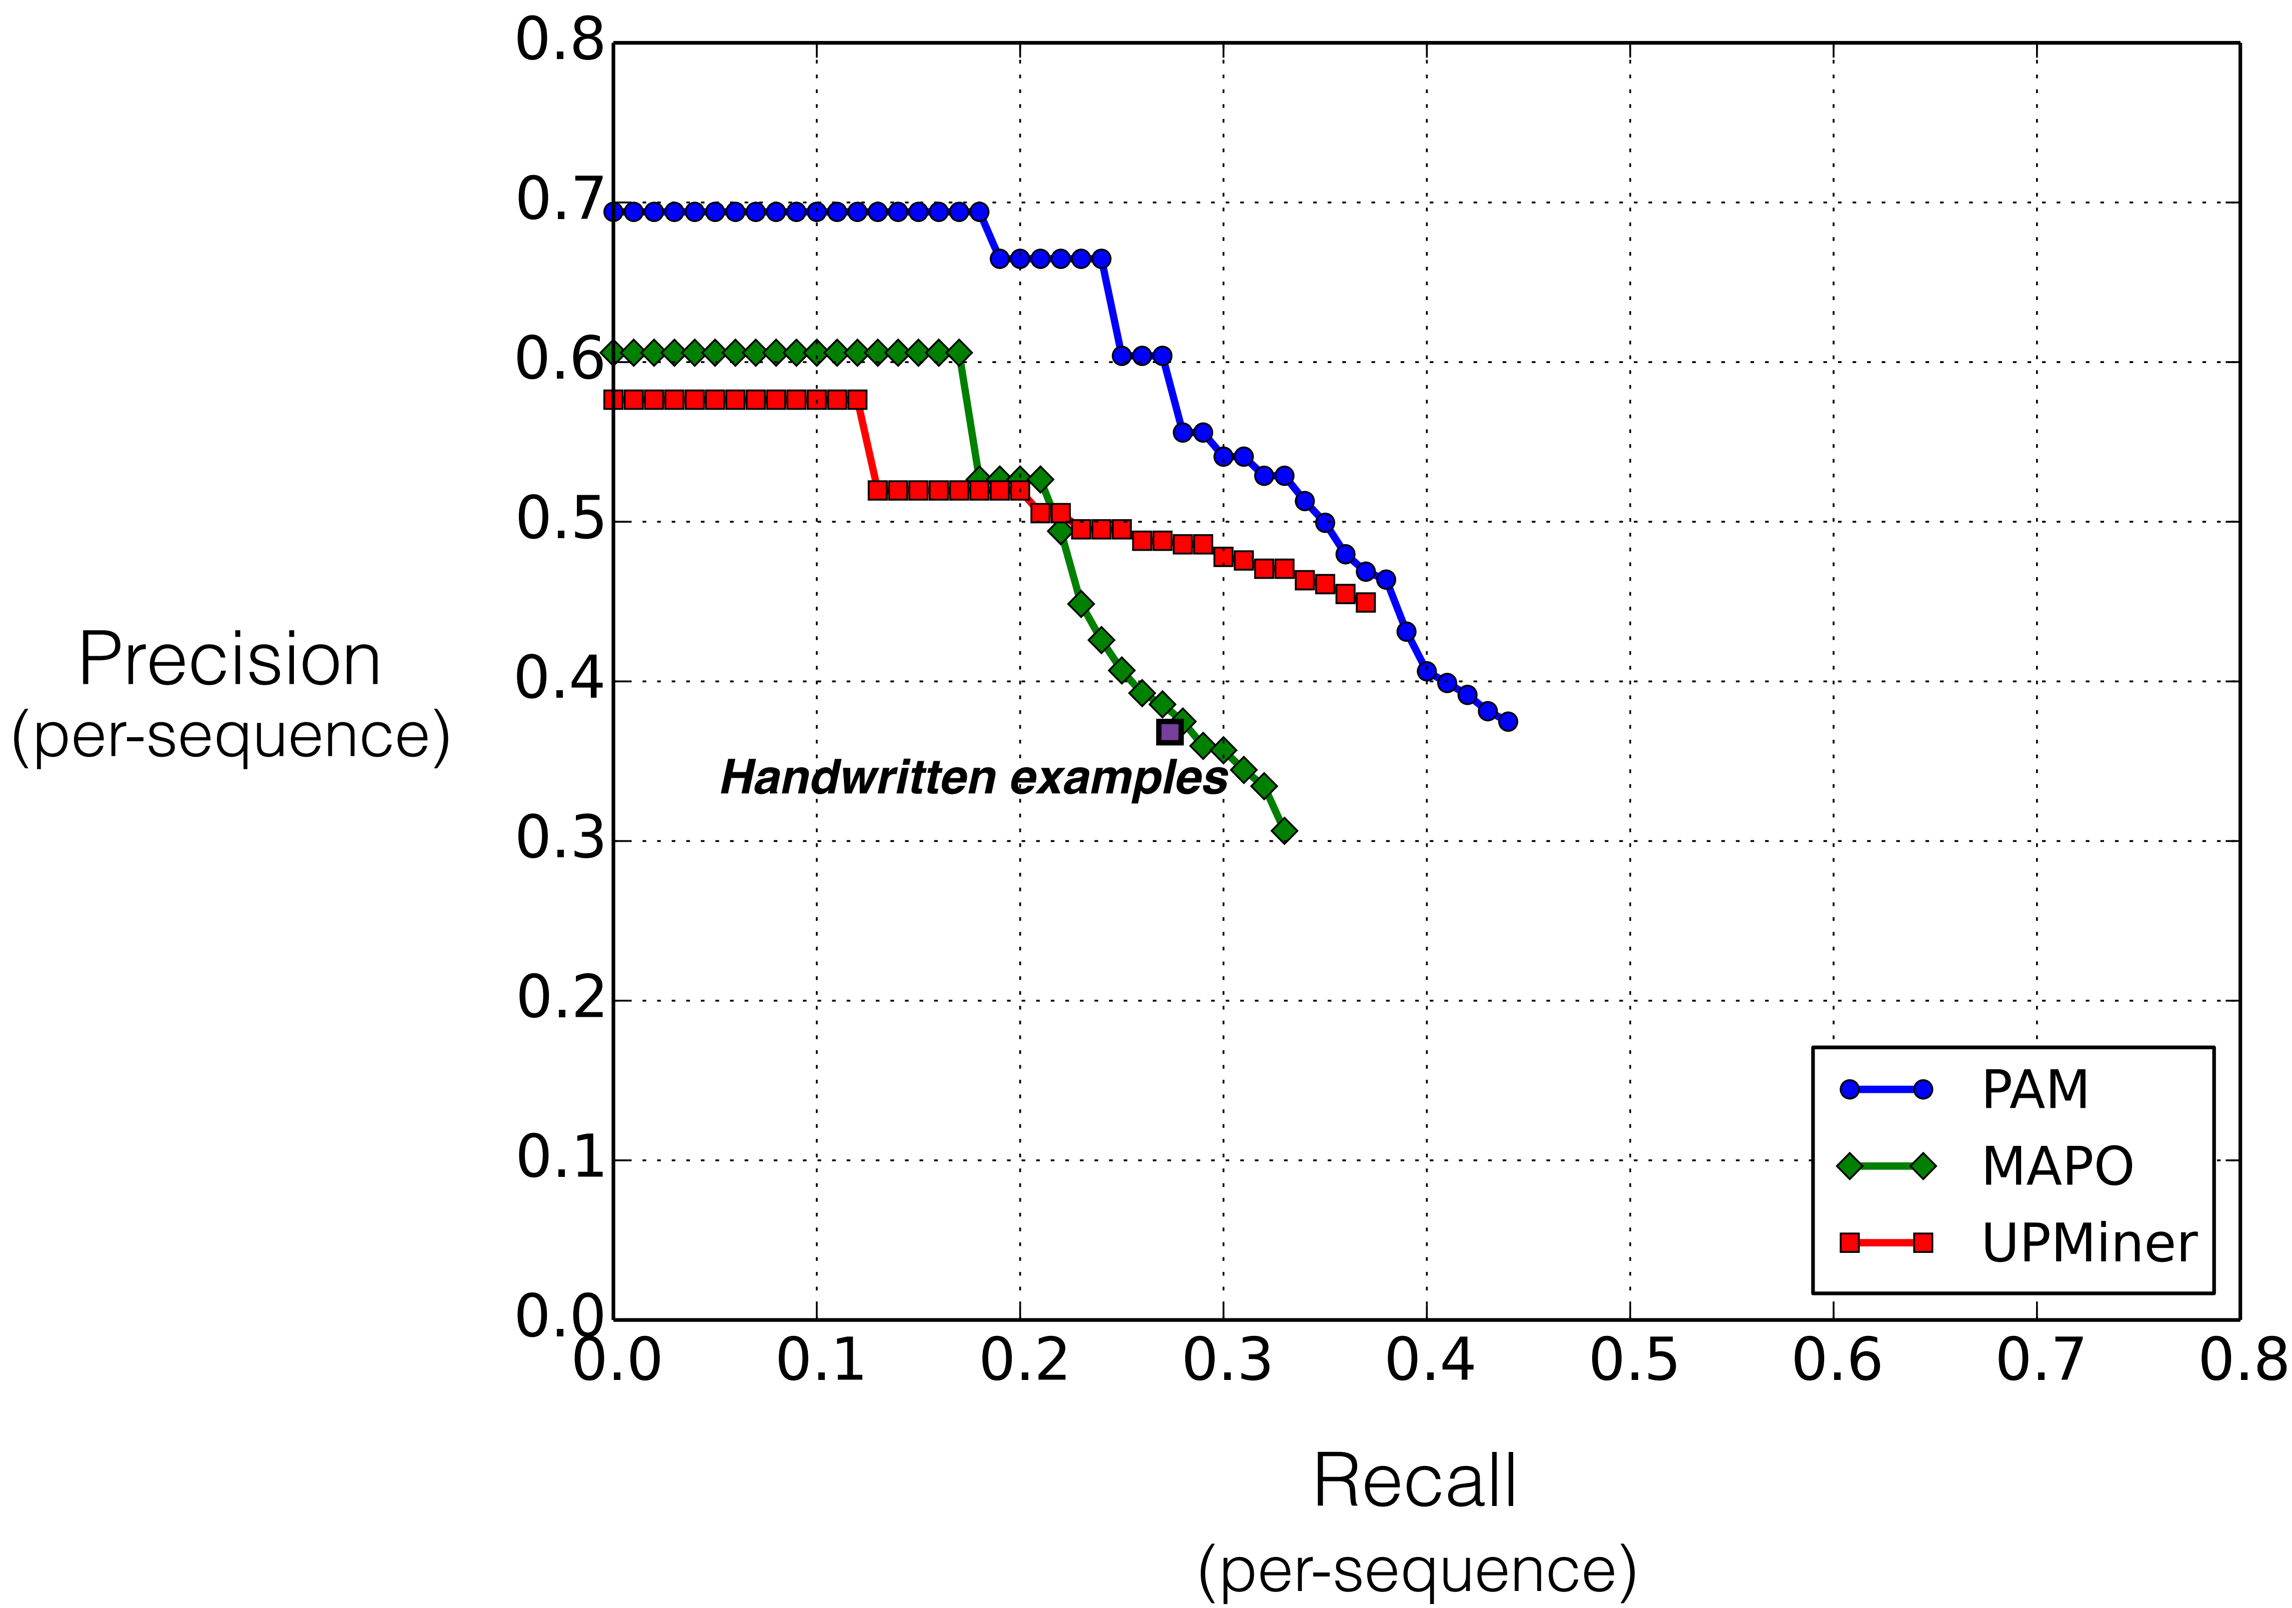

Example: twitter4j

PAM	MAPO *[Zhong et al, '09]*	UPMiner *[Wang et al, '13]*
TwitterFactory.<init> TwitterFactory.getInstance	TwitterFactory.<init> TwitterFactory.getInstance	TwitterFactory.<init> TwitterFactory.getInstance
TwitterFactory.<init> TwitterFactory.getInstance Twitter.setOAuthConsumer Twitter.setOAuthAccessToken	Status.getUser Status.getText	TwitterFactory.getInstance Twitter.setOAuthConsumer
Status.getUser Status.getText	ConfigurationBuilder.<init> ConfigurationBuilder.build	TwitterFactory.<init> TwitterFactory.getInstance Twitter.setOAuthConsumer
AccessToken.getToken AccessToken.getTokenSecret	ConfigurationBuilder.<init> TwitterFactory.<init>	Status.getUserStatus.getText
ConfigurationBuilder.<init> ConfigurationBuilder.build TwitterFactory.<init> TwitterFactory.getInstance	ConfigurationBuilder.<init> ConfigurationBuilder.setOAuthConsu merKey	Twitter.setOAuthConsumer Twitter.setOAuthAccessToken

: two main types of twitter initialization call

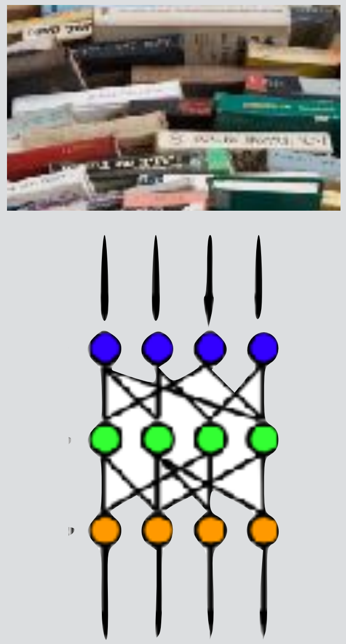

Super-Fast Neural Network Topic Modelling

[Srivastava and Sutton, ICLR 2017]

http://bit.ly/sutton-ml-exploration

Topic Models

motherboard meg printer quadra hd windows processor vga mhz connector
armenian genocide turks turkish muslim massacre turkey armenians armenia greek
voltage nec outlet circuit cable wiring wire panel motor install
season nhl team hockey playoff puck league flyers defensive player
israel israeli lebanese arab lebanon arabs civilian territory palestinian militia

Topic Models: The Industry

Latent dirichlet allocation
DM Blei, AY Ng, MI Jordan - Journal of machine Learning research, 2003 - jmlr.org
Abstract We describe latent Dirichlet allocation (LDA), a generative probabilistic model for collections of discrete data such as text corpora. LDA is a three-level hierarchical Bayesian model, in which each item of a collection is modeled as a finite mixture over an underlying
Cited by 17215 Related articles All 123 versions Cite Save

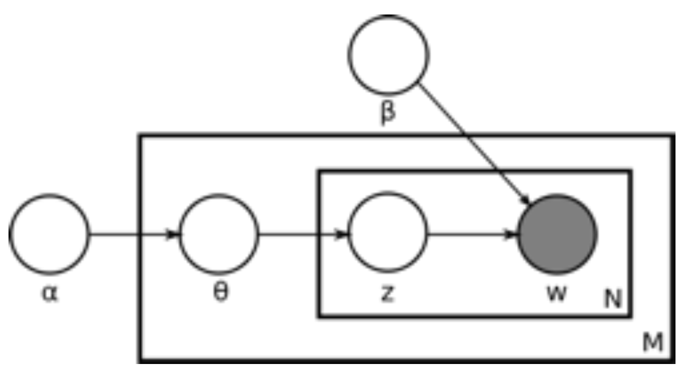

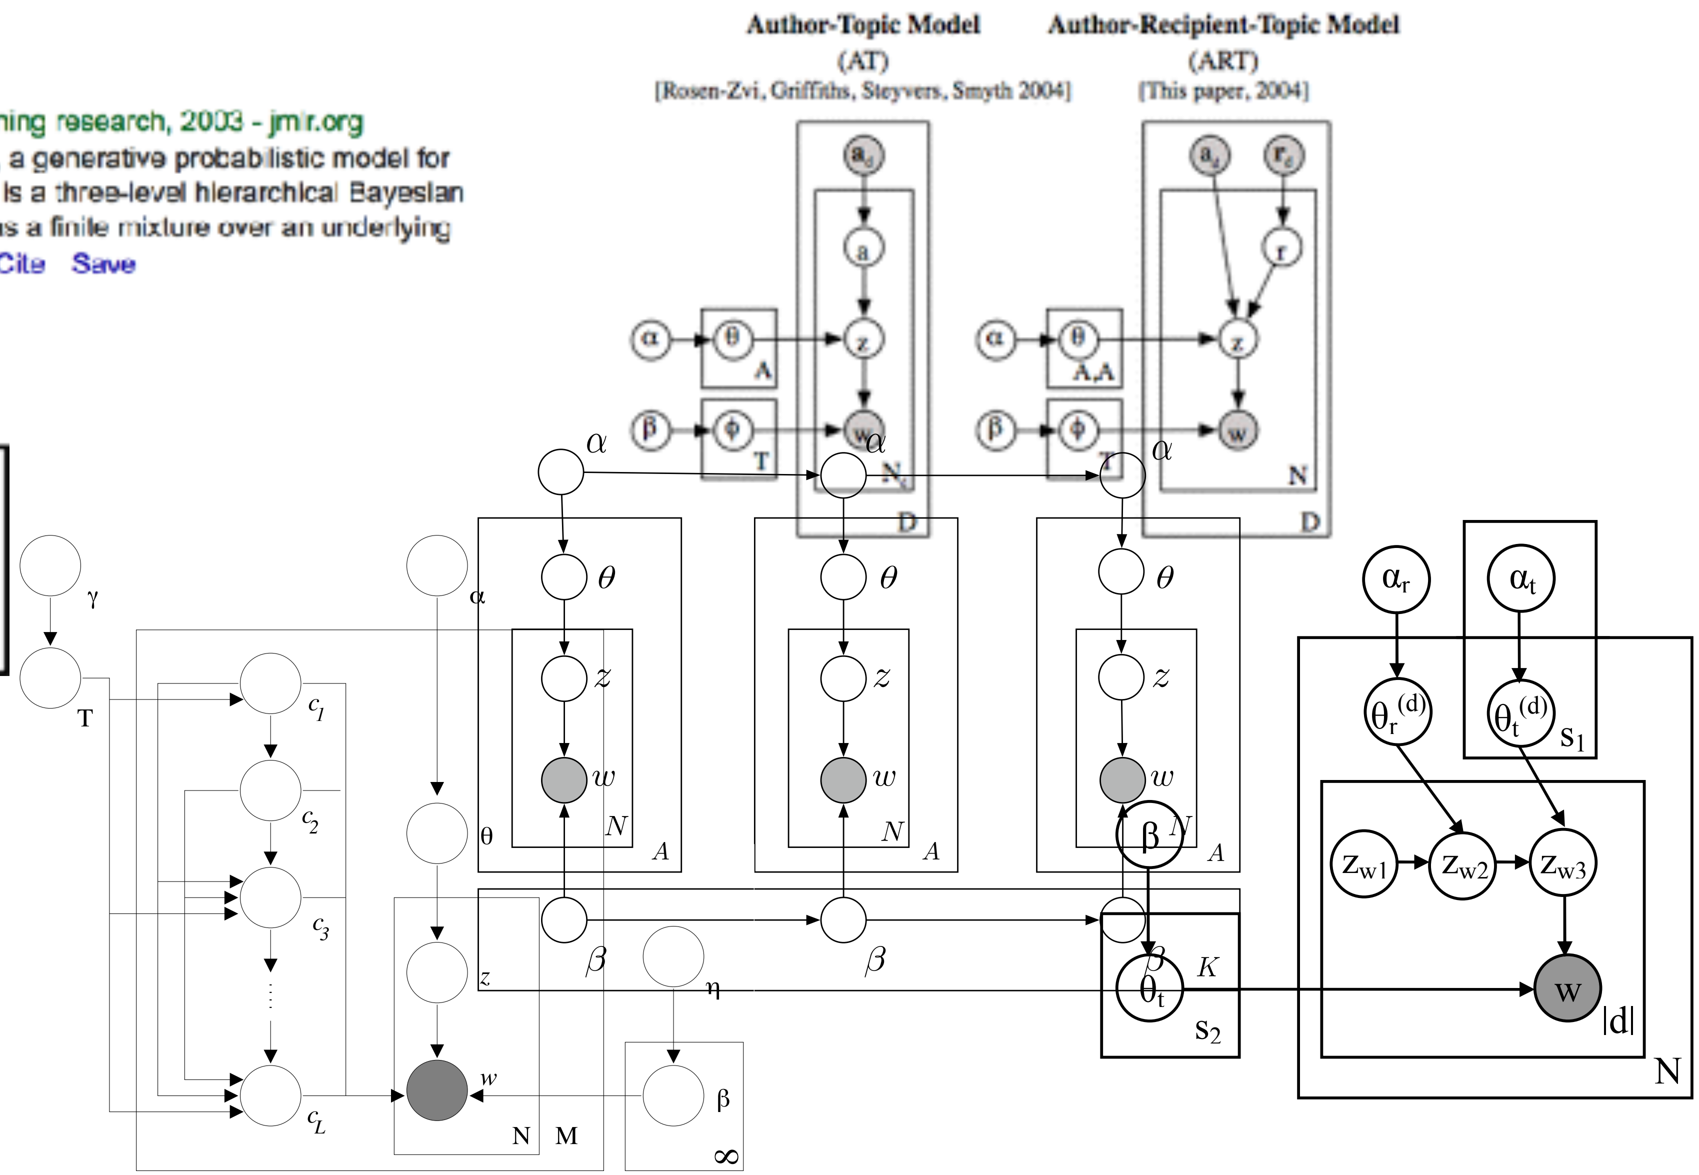

But…

- New topic model means new inference algorithm
- MCMC general but slow; variational fast but not as general, and lower quality topics

Recognition Networks

[Hinton et al, 1995]

Generator

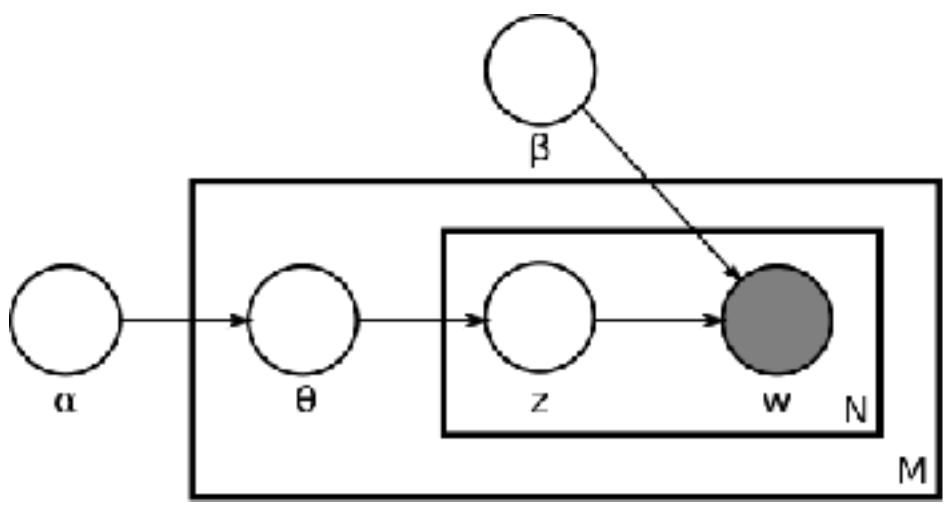

$$p(\mathbf{w}, \theta, z|\alpha)$$

LDA

Inference network

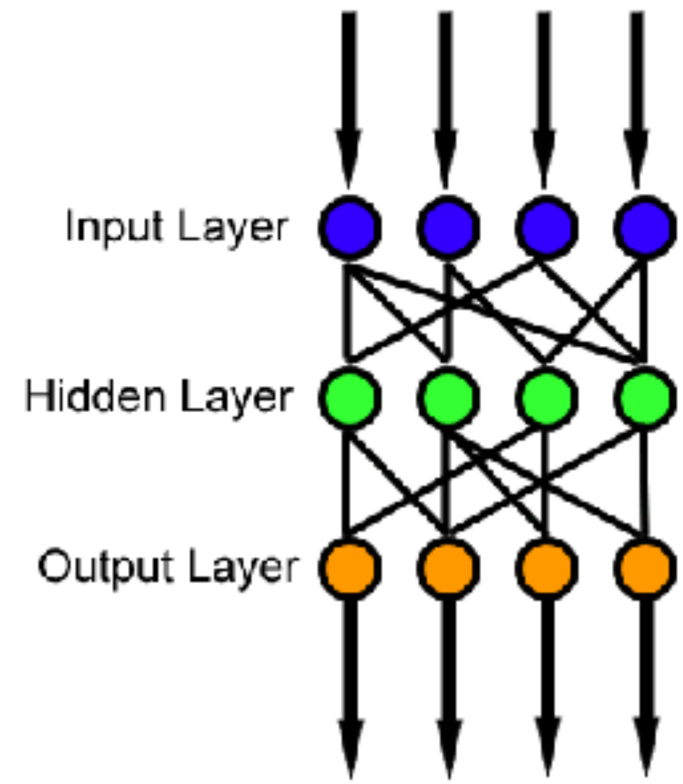

$$\theta$$

$$q(\theta|\gamma, \phi)$$

Appx posterior

Variational autoencoder

[Kingma and Welling, 2013]

$$L(\gamma, \phi \mid \alpha, \beta) = -D_{KL}\left[q(\theta, z|\gamma, \phi)||p(\theta, z|\alpha)\right] + \mathbb{E}_{q(\theta, z|\gamma, \phi)}[\log p(\mathbf{w}|z, \theta, \alpha, \beta)]$$

How to make it work

- Discrete variables
  - Marginalize them out
- Dirichlet difficult to reparameterize
  - Use Laplace approximation *[MacKay, 1998]*
- Topic collapsing (a bad local minimum)
  - High momentum in ADAM *[Kingma and Ba, 2014]*
  - Batch normalization *[Ioffe and Szegeti, 2015]*

Deriving new models

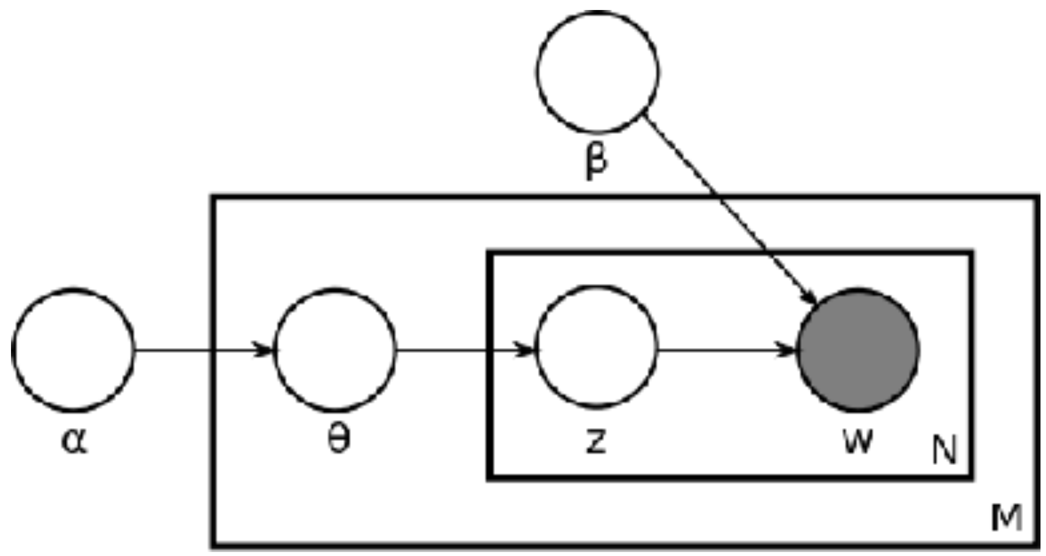

LDA

$$p(w_n|\theta, \beta) = \sum_k \theta_k p(w_n|z_n = k, \beta)$$

ProdLDA

$$p(w_n|\theta, \beta) \propto \prod_k p(w_n|z_n = k, \beta)^{\theta_k}$$

Very simple change to LDA, but hasn't been seen before.

Why?

- Before, I would have derived a variational inference algorithm
- Now, change one line of code

Evaluation

# topics	ProdLDA	LDA NVLDA	LDA DMFVI	LDA Collapsed Gibbs	NVDM
50	**0.24**	0.20	0.11	0.17	0.08
200	**0.19**	0.12	0.06	0.14	0.06

Topic coherence (20 Newsgroups)

# topics	ProdLDA	LDA NVLDA	LDA DMFVI	LDA Collapsed Gibbs	NVDM
50	**0.14**	0.07	-	0.04	0.07
200	**0.12**	0.05	-	0.06	0.05

Topic coherence (RCV1)

Look Ma, no inference!

# Topics	Inference Network Only	Inference Network + Optimization
50	1172	1162
200	1168	1151

Test set perplexity (20 Newsgroups)

Accurate topic inference on new topic

with one pass of a feedforward neural network

VEEGAN: Reducing Mode Collapse in Generative Adversarial Learning

http://bit.ly/sutton-ml-exploration

Generative Adversarial Networks

Classical probabilistic modelling

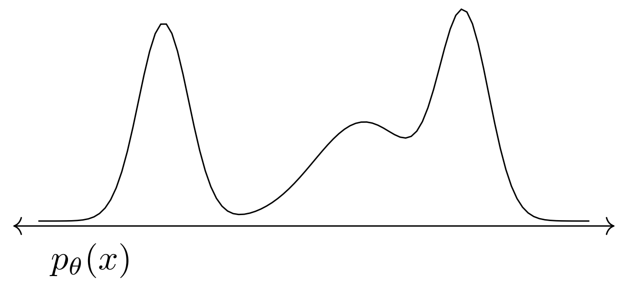

0.00003

Input

Explicit model

Density value

Implicit probabilistic modelling

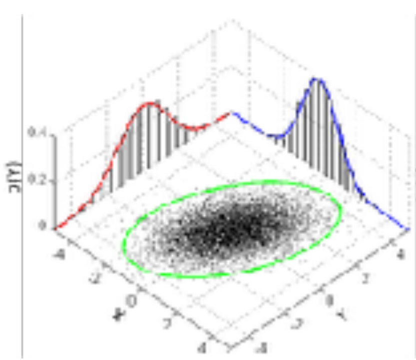

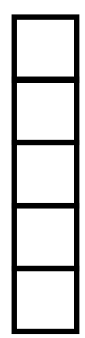

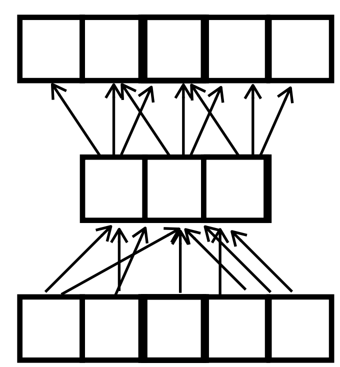

Gaussian
$p(z)$

Representation
z

Generator network
$G_\gamma(z)$

Image
x

Sampling procedure for $p_\theta(x)$

How to train?

Can’t use maximum likelihood. There is no likelihood!

Instead define a game

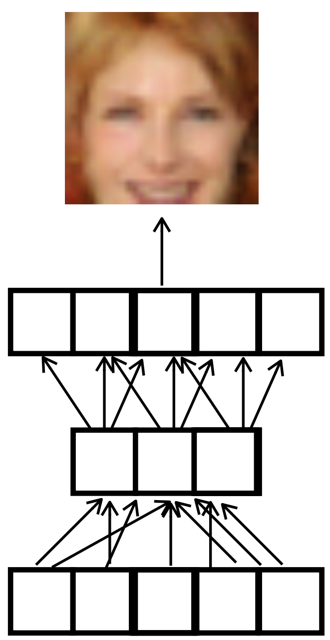

Generator network

$G_\gamma(z)$

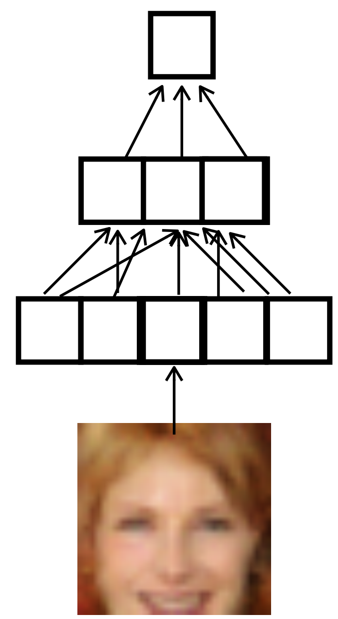

Discriminator network

$D_\omega(x)$

1 if x came from generator
0 if x came from data

Optimize

$$\max_{\omega}\min_{\gamma}\mathcal{O}_{\text{GAN}}(\omega,\gamma) := E_z\left[\log D_\omega(G_\gamma(z))\right] + E_x\left[\log\left(1 - D_\omega(x)\right)\right]$$

Mode Collapse

Example from 2D mixture of Gaussians

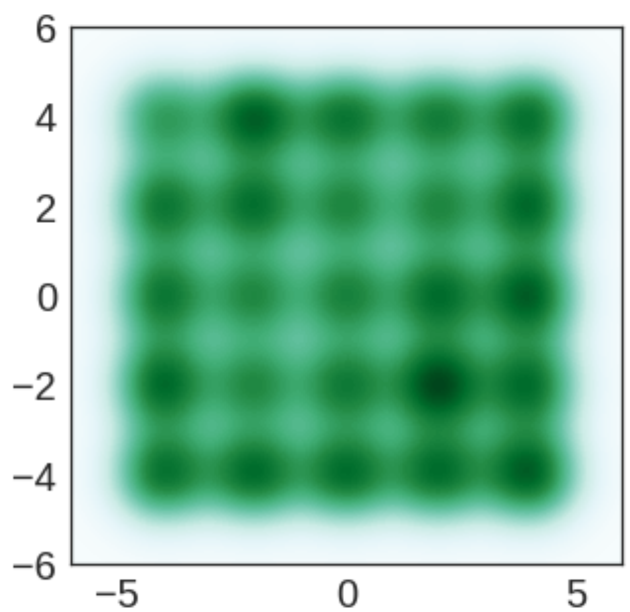

True data

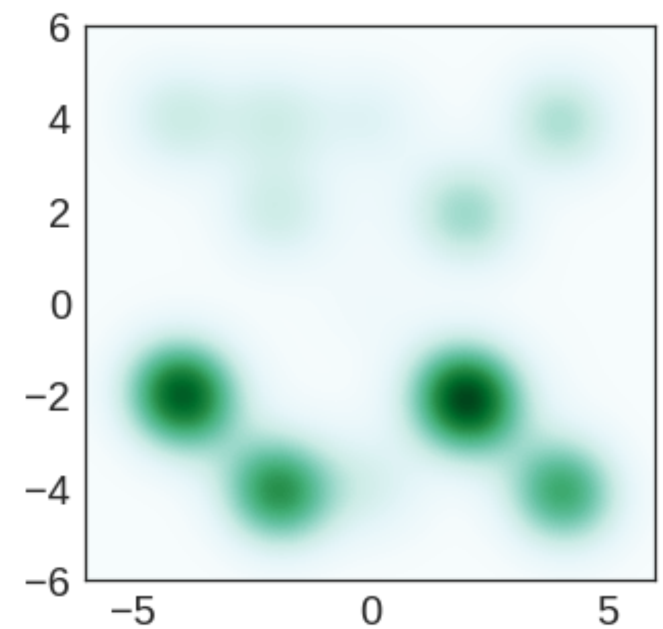

Samples from GAN

VEEGAN: Adding an Autoencoder

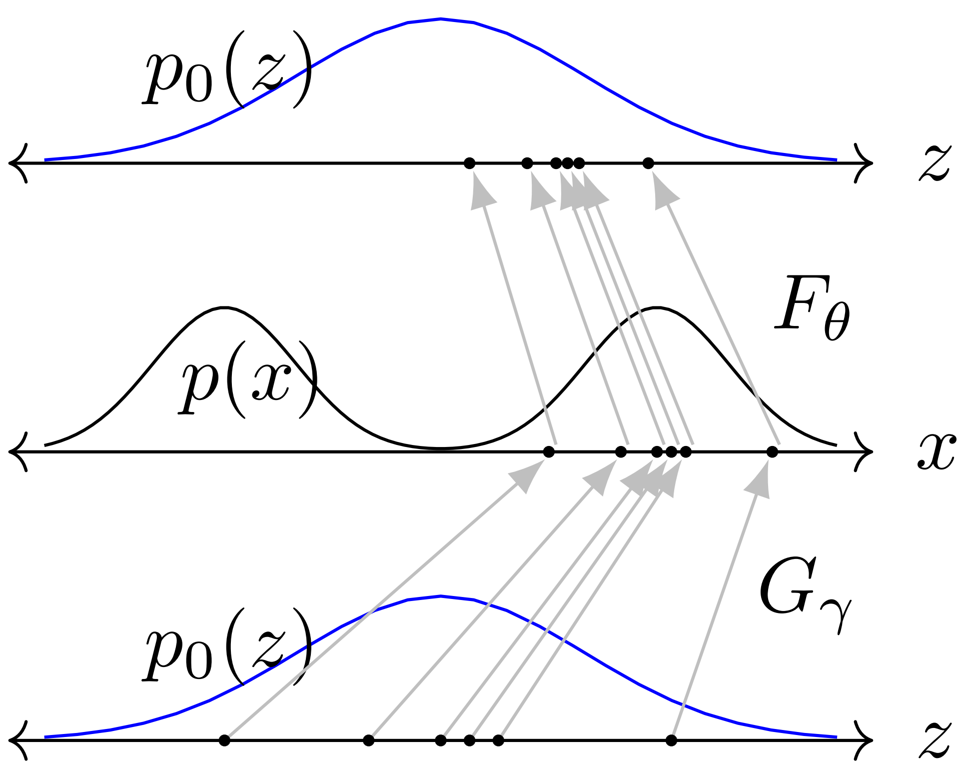

Detecting mode collapse

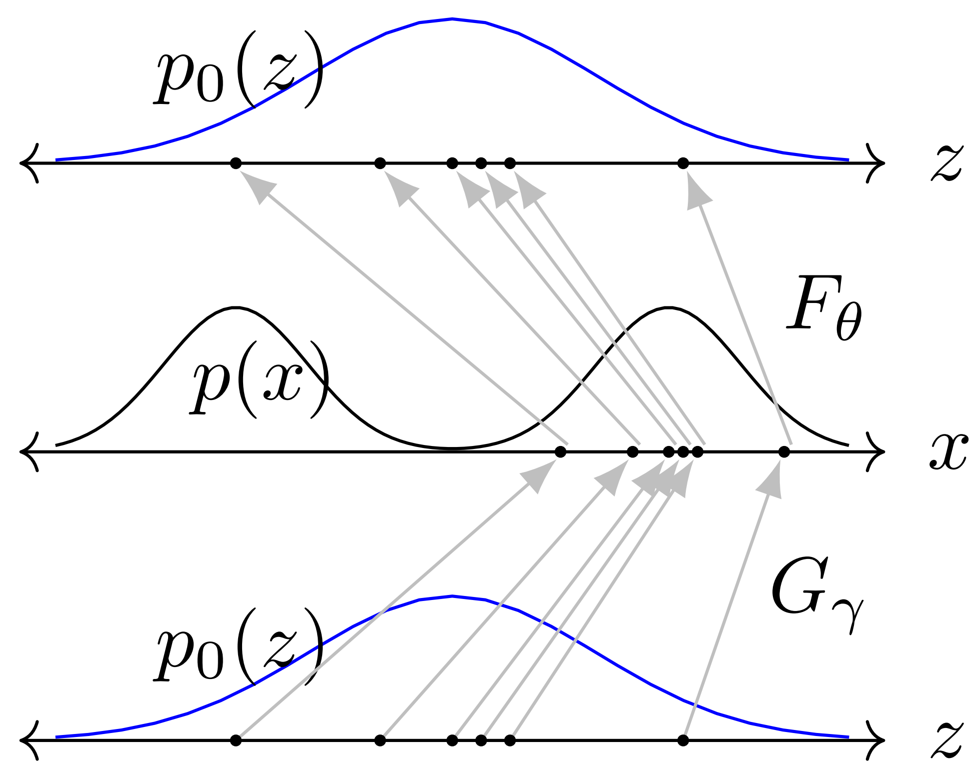

Not detecting mode collapse

Problem: How to train F_θ ?

VEEGAN objectives

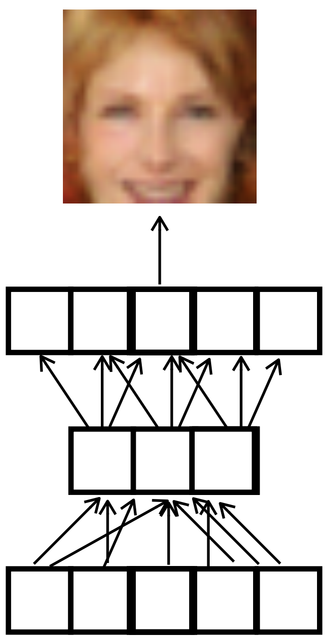

Generator network

$G_\gamma(z)$

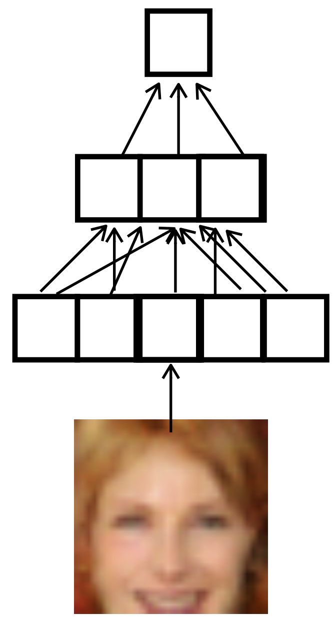

Discriminator network

$D_\omega(x)$

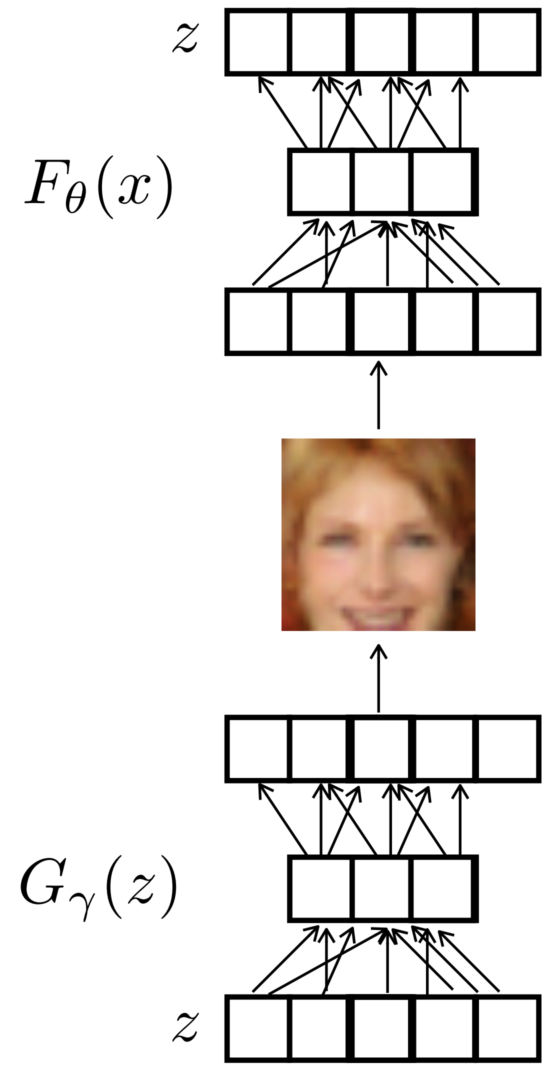

Autoencoder

VEEGAN algorithm

Algorithm 1 VEEGAN training

while not converged **do**
  for $i \in \{1 \dots N\}$ **do**
    Sample $z^i \sim p_0(z)$
    Sample $x_g^i \sim q_\gamma(x|z_i)$
    Sample $x^i \sim p(x)$
    Sample $z_g^i \sim p_\theta(z_g|x_i)$
  $g_\omega \leftarrow -\nabla_\omega \frac{1}{N} \sum_i \log \sigma \left(D_\omega(z^i, x_g^i) \right) + \log \left(1 - \sigma \left(D_\omega(z_g^i, x^i) \right) \right)$ ▷ Compute $\nabla_\omega \hat{\mathcal{O}}_{\text{LR}}$

  $g_\theta \leftarrow \nabla_\theta \frac{1}{N} \sum_i d(z^i, x_g^i)$ ▷ Compute $\nabla_\theta \hat{\mathcal{O}}$

  $g_\gamma \leftarrow \nabla_\gamma \frac{1}{N} \sum_i D_\omega(z^i, x_g^i) + \frac{1}{N} \sum_i d(z^i, x_g^i)$ ▷ Compute $\nabla_\gamma \hat{\mathcal{O}}$

  $\omega \leftarrow \omega - \eta g_\omega$; $\theta \leftarrow \theta - \eta g_\theta$; $\gamma \leftarrow \gamma - \eta g_\gamma$ ▷ Perform SGD updates for ω, θ and γ

Examples of generated images

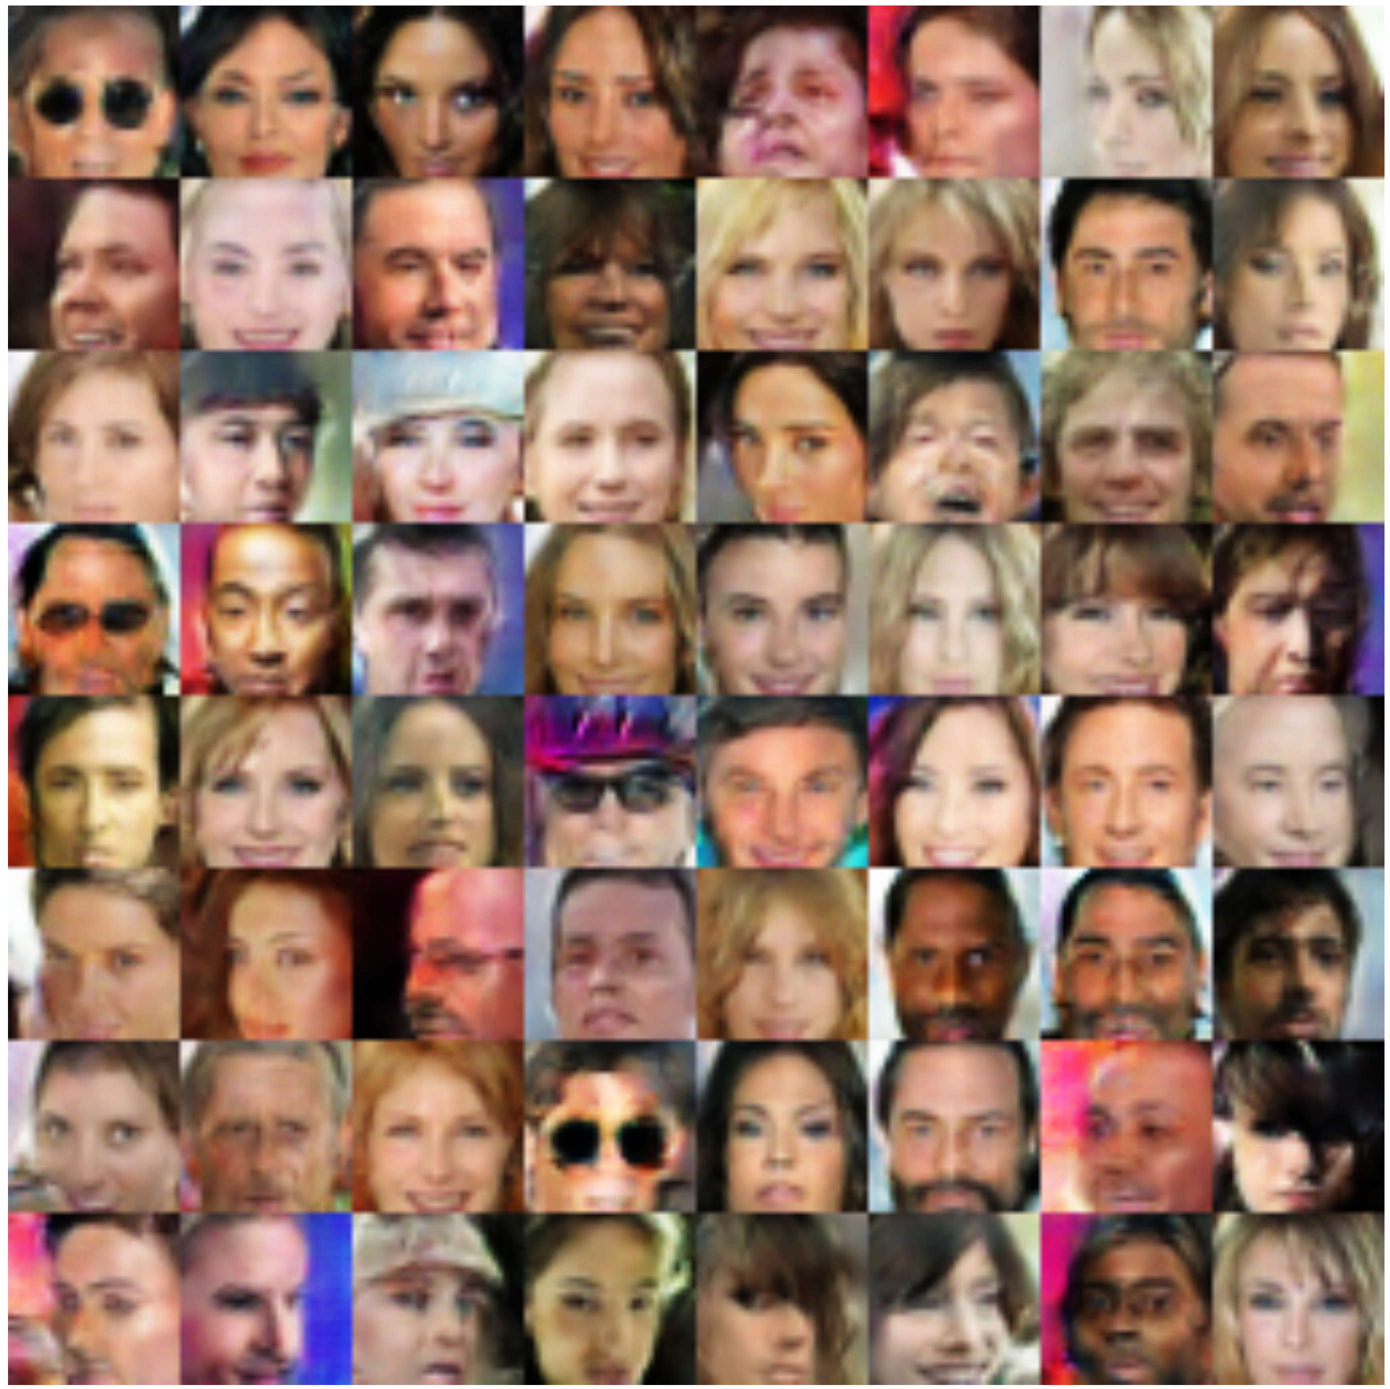

Celebrity faces

Examples of generated images

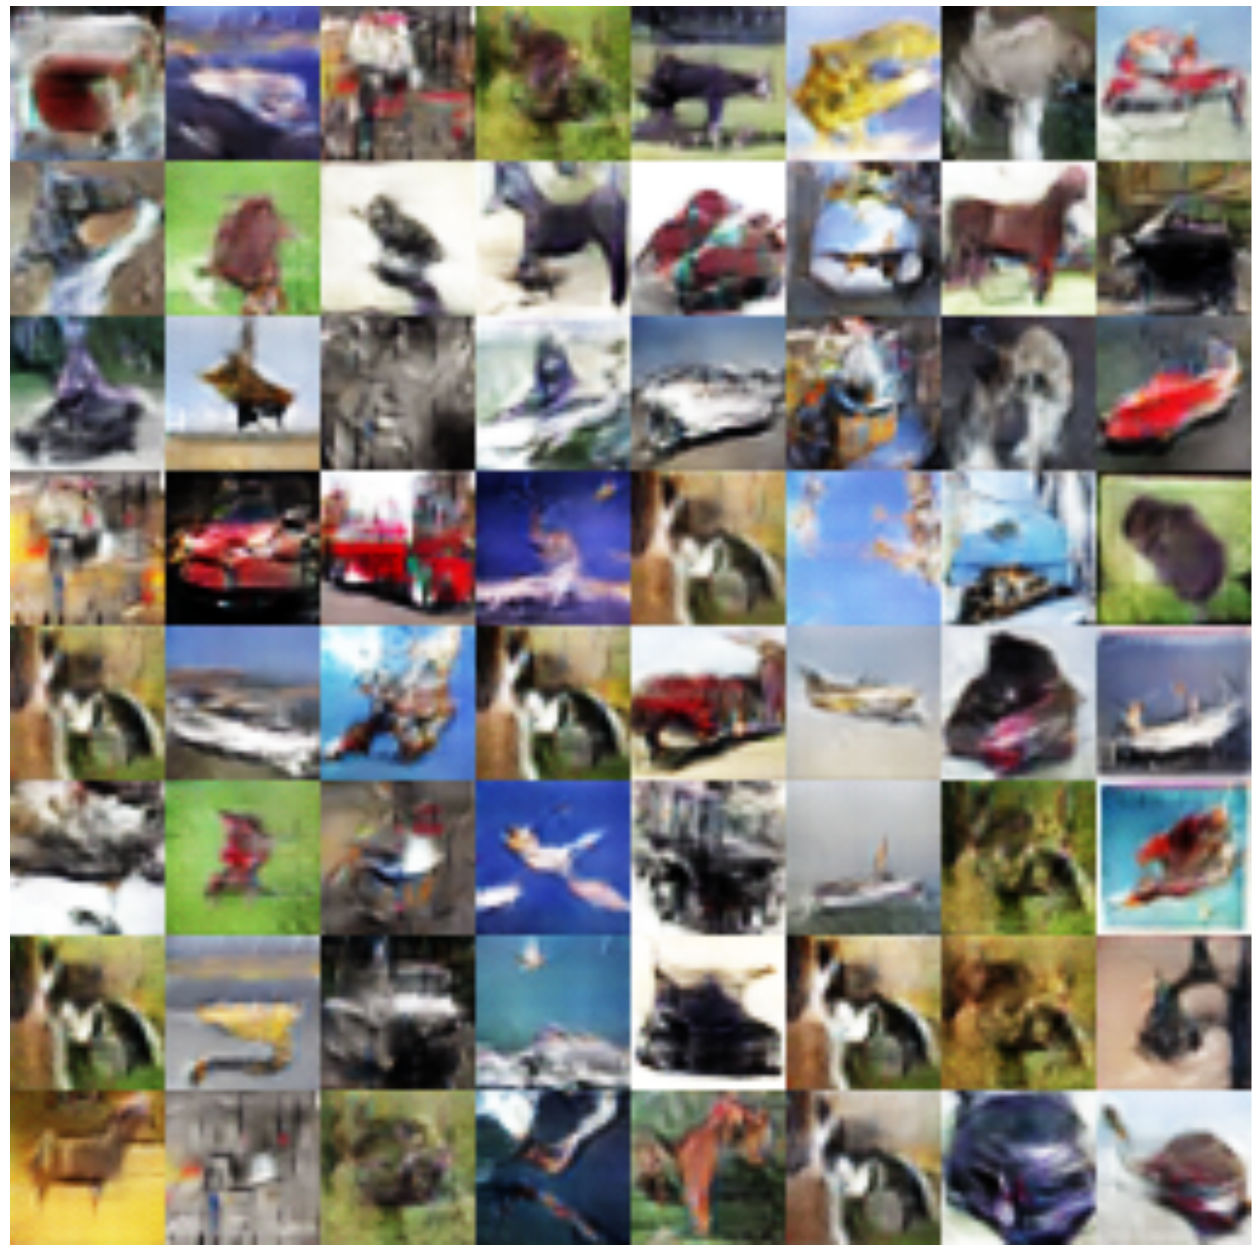

CIFAR-10 natural images

Machine Learning for Data Exploration

Charles Sutton, University of Edinburgh

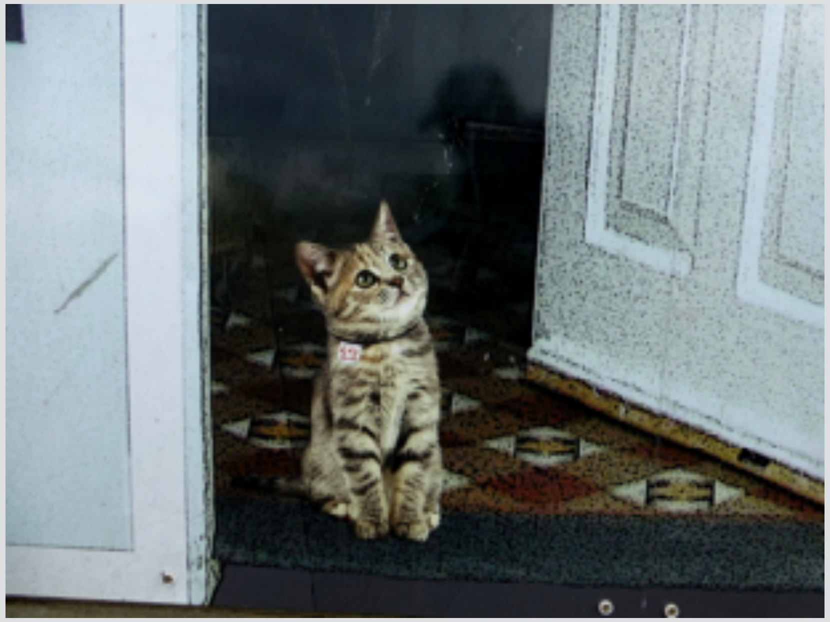

Help cats explore!

Pattern mining

API mining

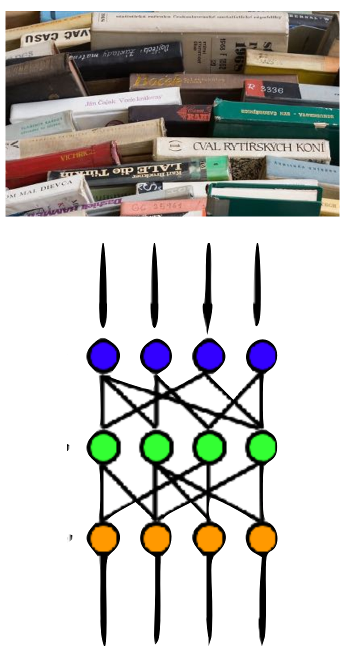

Neural topic modelling

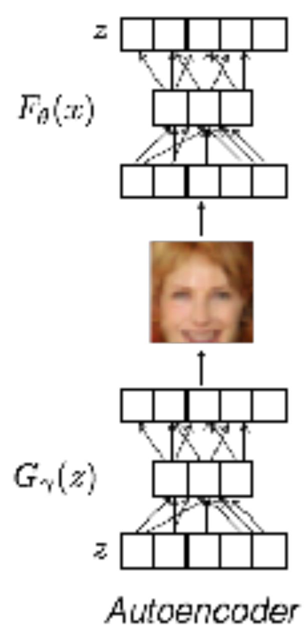

VEEGANs

- Jaroslav Fowkes
- Akash Srivastava
- Lazar Valkov
- Chris Russell
- Michael Gutmann

Thanks!

EPSRC
Engineering and Physical Sciences Research Council

http://bit.ly/sutton-ml-exploration