Summarizing Event Sequences with Generalized Sequential Patterns

Joscha Cüppers
Jilles Vreeken

Cüppers & Vreeken Below the Surface: Summarizing Event Sequences with Generalized Sequential Patterns.
In: Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), ACM, 2023

CISPA
HELMHOLTZ CENTER FOR INFORMATION SECURITY

Problem

Given only a set of event sequences, D.

Goal, report a set of **generalized events** and a set P of **generalized patterns**, that succinctly summarize D.

Generalized Pattern?

Existing Methods *"Surface Level Patterns"*	**Our Method – FLOCK** *"Generalized Patterns"*
The cat meows The dog barks	The [pet] [makes noise] cat meows dog barks fish blubs

only "strong enough" in the context of the generalized pattern

Generalized Patterns

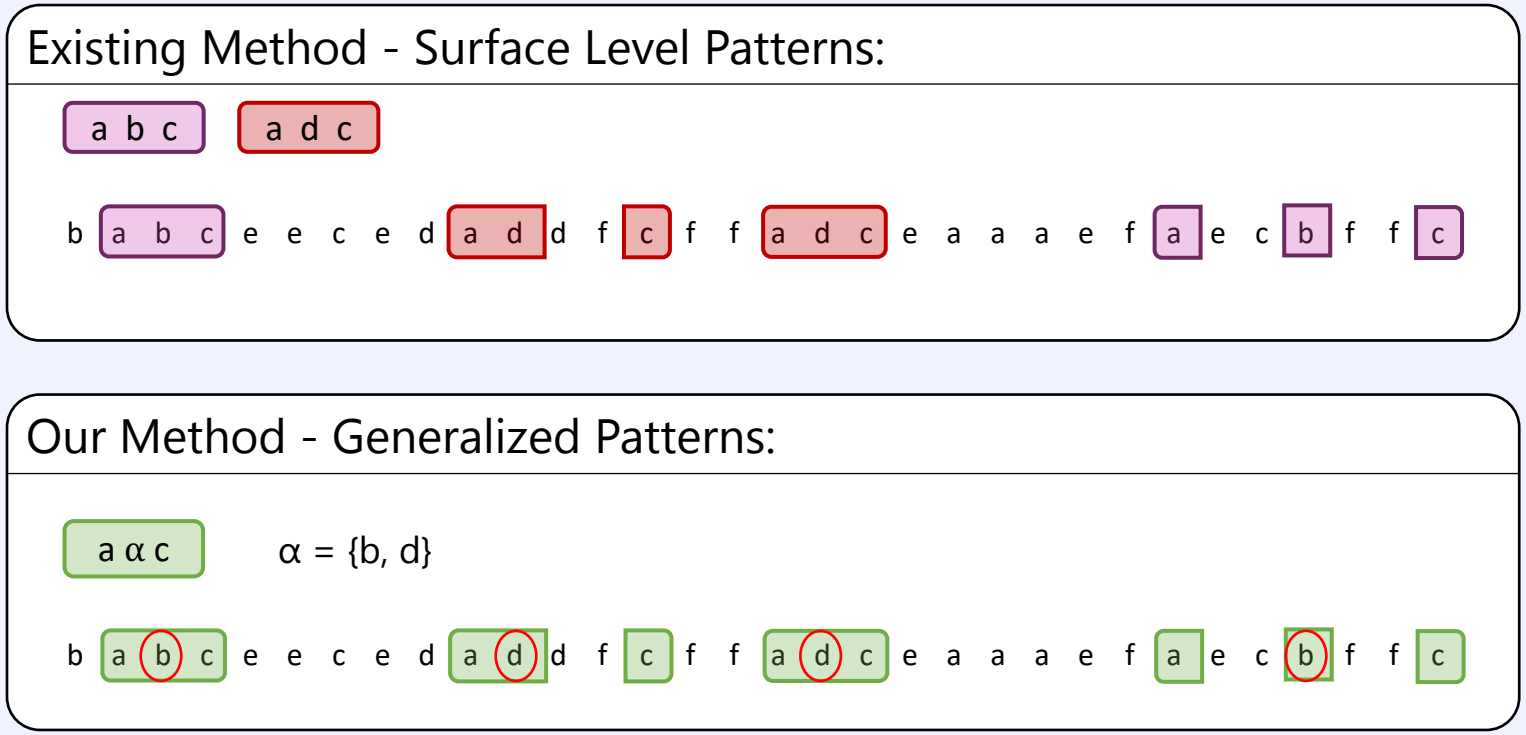

- Set of Observed Events - Ω_o e.g. $\Omega_o = \{a, b, c, \ldots\}$
- Set of Generalized Events - Ω_g e.g. $\Omega_g = \{\alpha\}$
- Alphabet – $\Omega = \Omega_o \cup \Omega_g$

How do we do that?

The Minimum Description Length (MDL) principle:

given a model class $\mathcal{M}$, the best model $M \in \mathcal{M}$
is that M that minimizes

$$L(D, M) = L(M) + L(D \mid M)$$

where:

$L(M)$ is the length of the model, in bits
$L(D \mid M)$ is the length of the data, in bits, when encoded using M

Length of Model

$$L(M) = L(CT) + L(\Omega_g)$$

Code Table – Pattern set and usage of each pattern

$$L(\mathrm{CT}) = \underbrace{L_{\mathbb{N}}(|P'|)}_{\text{how many patterns}} + \underbrace{L_{\mathbb{N}}(usage(P))}_{\text{usage sum over all patterns}} + \underbrace{\log \binom{usage(P)-1}{|P|-1}}_{\text{usage of each pattern}} + \underbrace{\sum_{p \in P'} L(p)}_{\text{encoding of patterns}}$$

Set of Generalized Events $\boldsymbol{\Omega_g}$

$$L(\Omega_g) = \sum_{e \in \Omega_g} L(e)$$

Model M:

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

α : { e f }

p : d α b c
q : α e a

Length of Data

Cover 1: (Singletons):

C_p: d e b f c d e a

Event Sequence: d e b f c d e a

Cover 2: (Patterns):

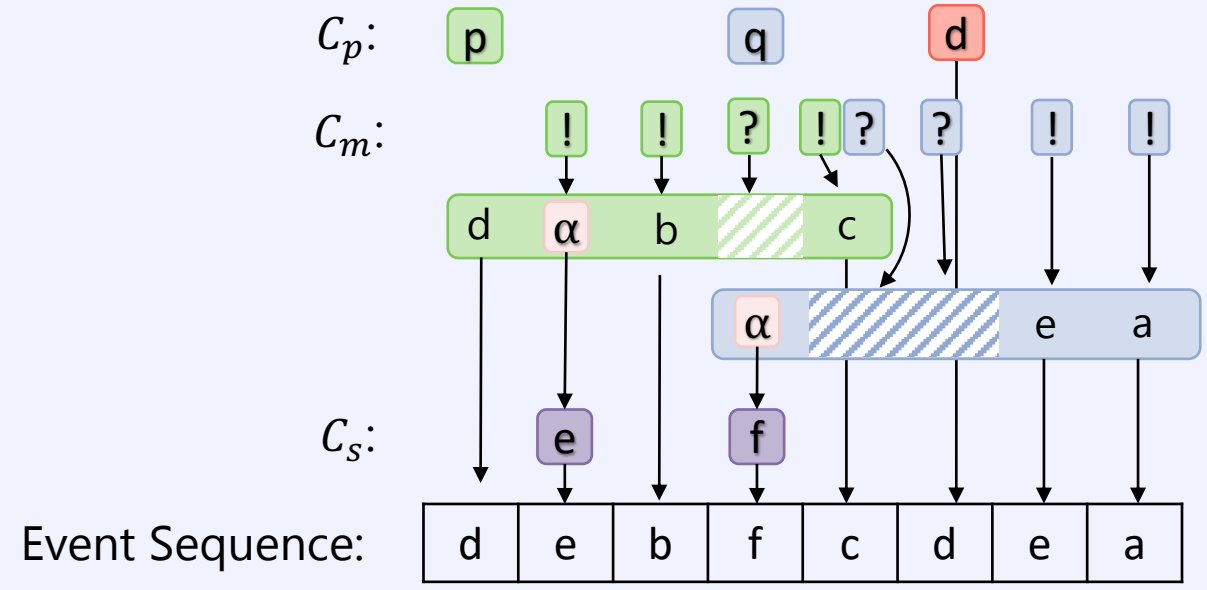

Model M:

a : a

b : b

c : c

d : d

e : e

f : f

α : { e f }

p : d α b c

q : α e a

$$L(D|M) = L(C_p) + L(C_m) + L(C_s)$$

Mining Models

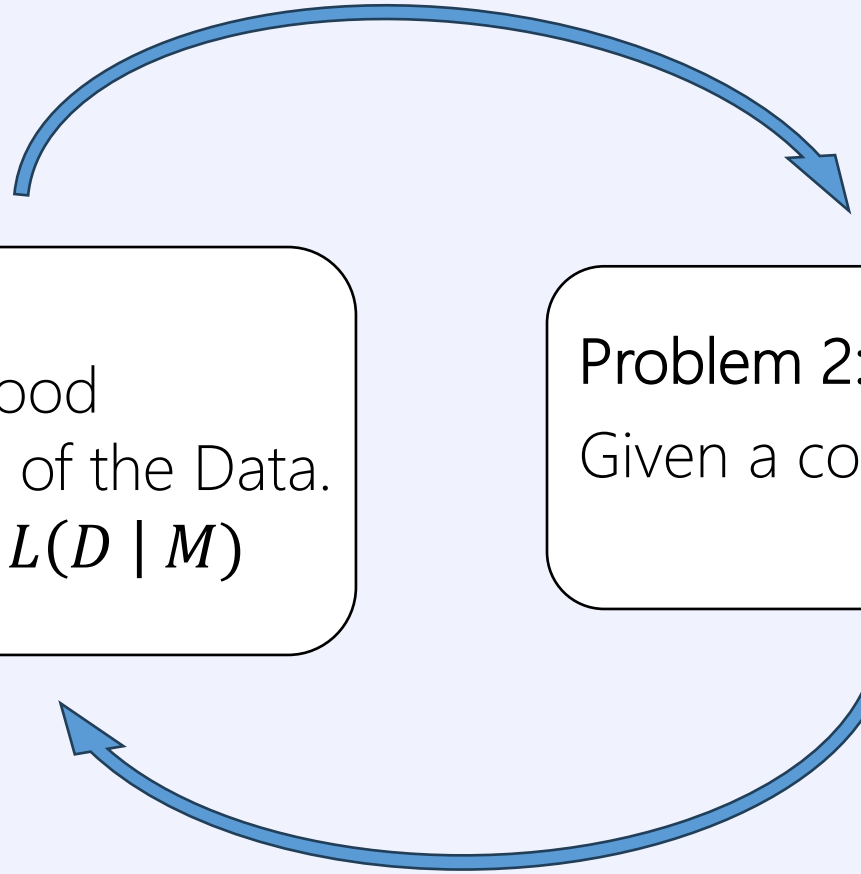

Problem 1:

Given a model M find a good description (i.e. a cover C) of the Data. Objective is to minimize - $L(D \mid M)$

Problem 2:

Given a cover C find a good model M

FLOCK Algorithm – Basic Idea

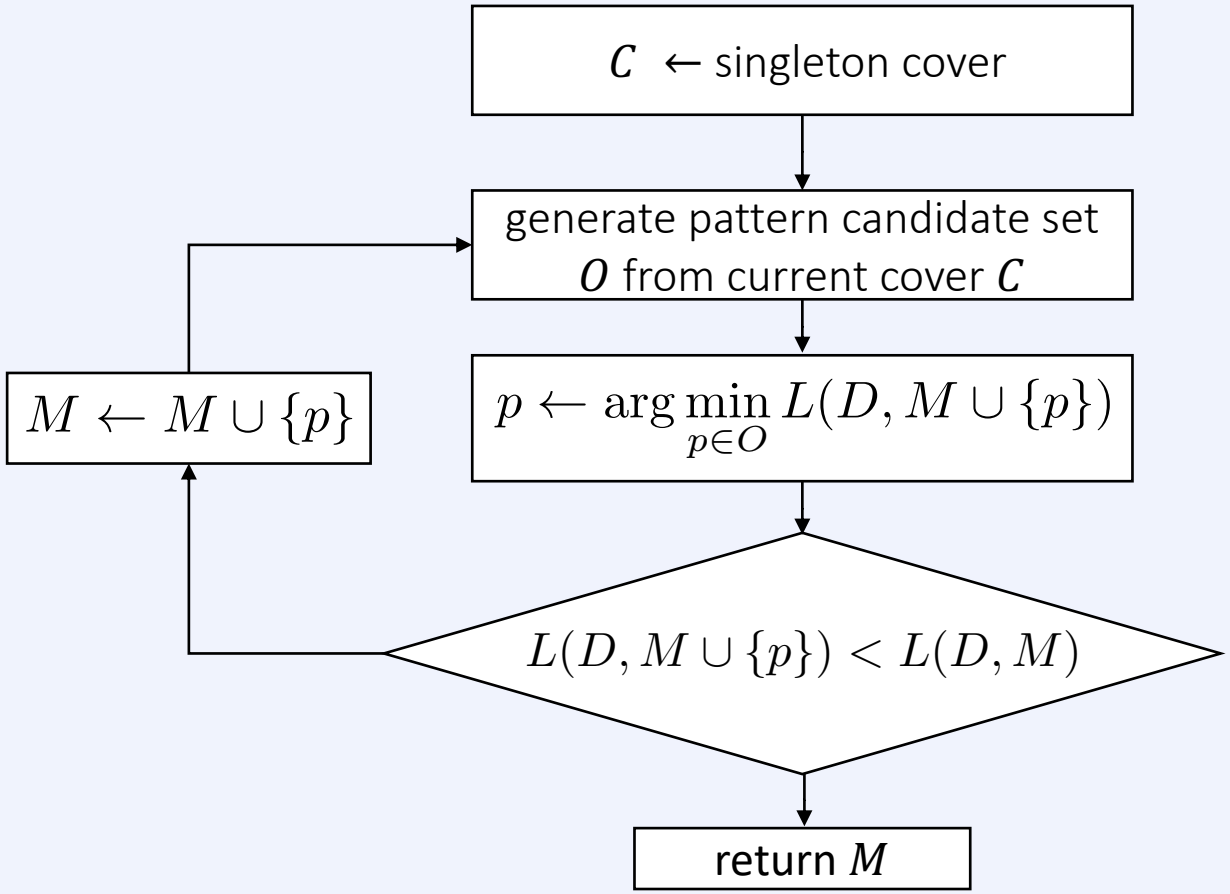

Discovering Generalized Events

Merge

1. mine "surface level" patterns
2. merge patterns

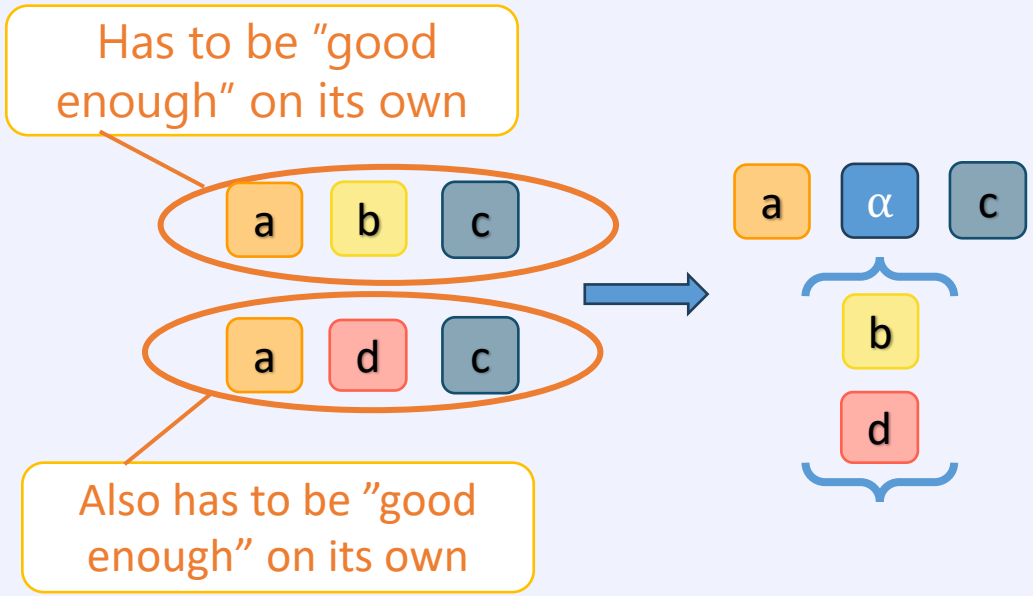

Candidate Generation

Suppose a is often followed by b and d, with similar number of gaps.

Generate Candidates:

1. a b
2. a d
3. a α
   b
   d

Experiments / Related Work

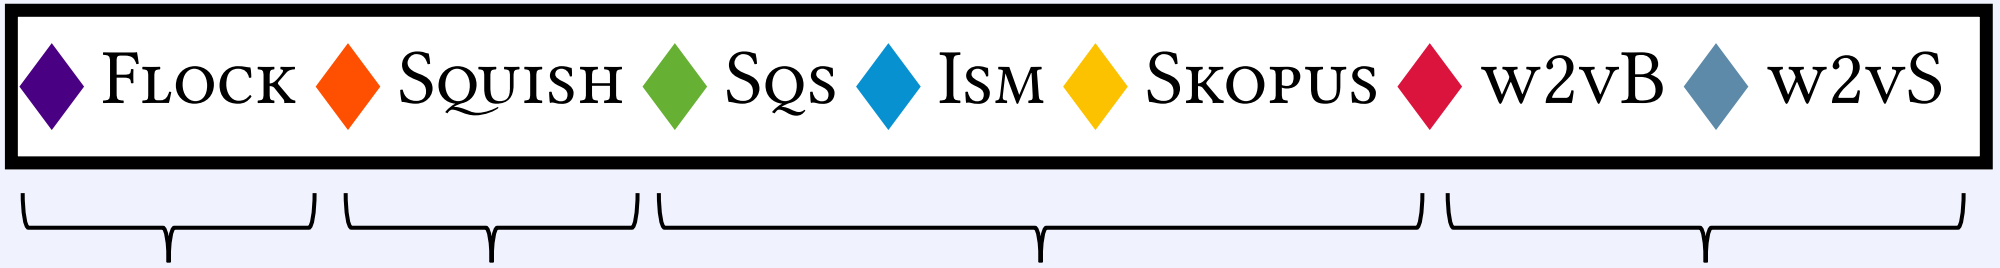

Evaluation – Synthetic Data

Data with known ground truth.

0 Generalized Patterns ⇨ 1 Generalized Patterns ⇨ ...

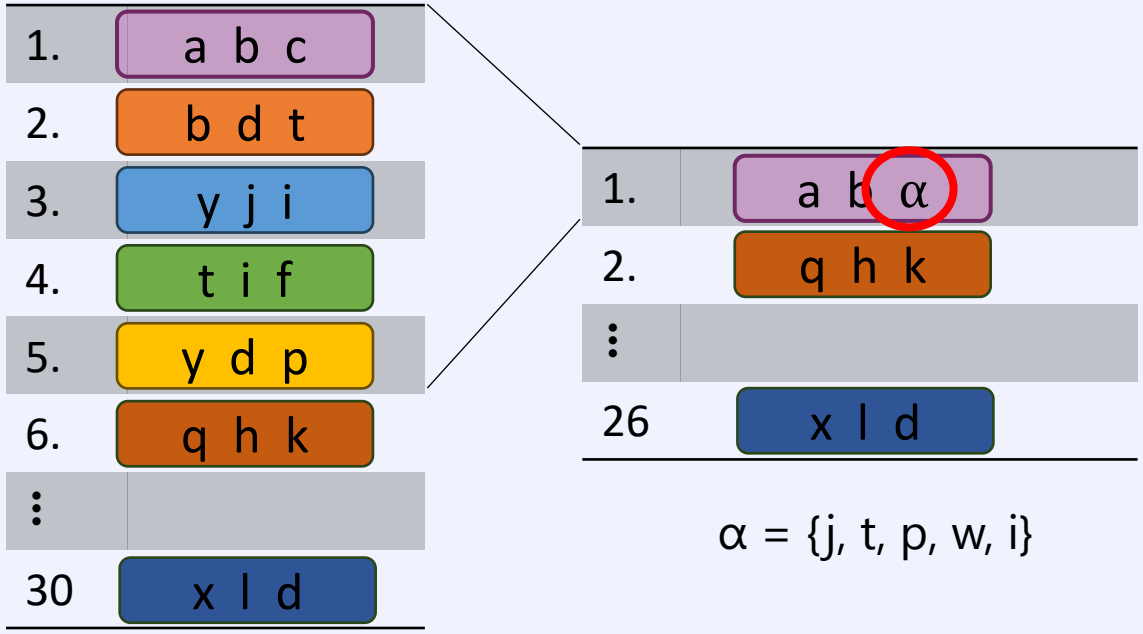

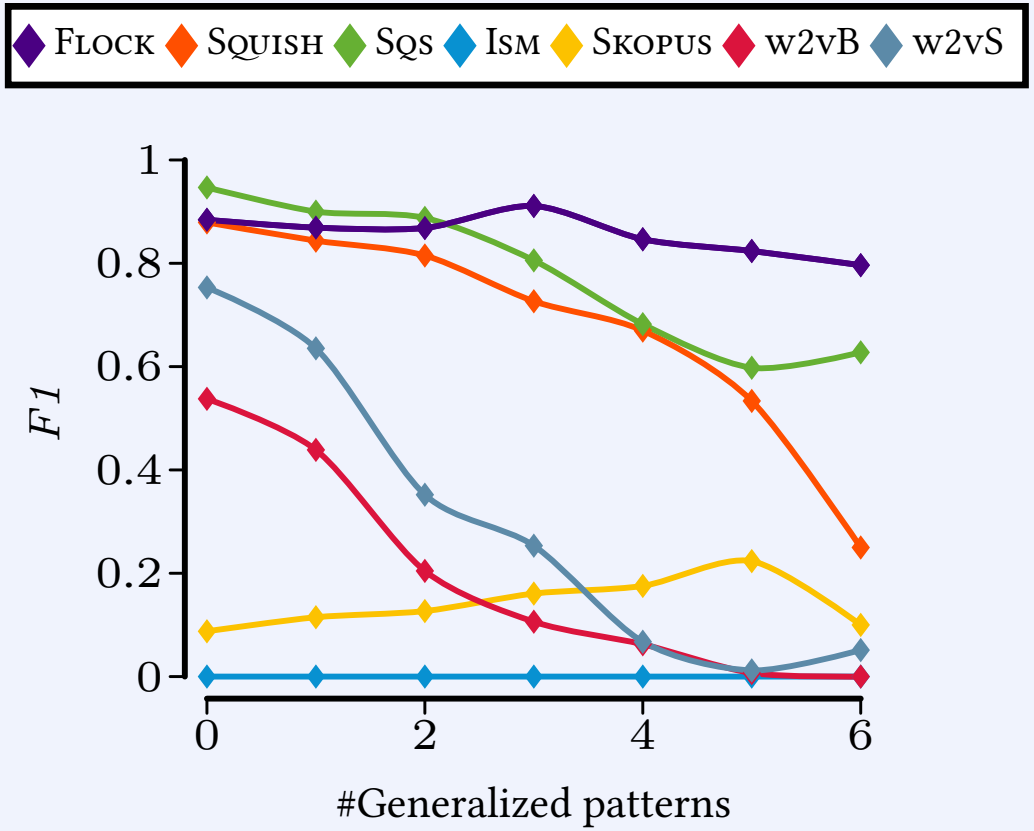

Evaluation – Synthetic Data

- 5 patterns
- 5 generalizations in total
- 2 generalizations per pattern

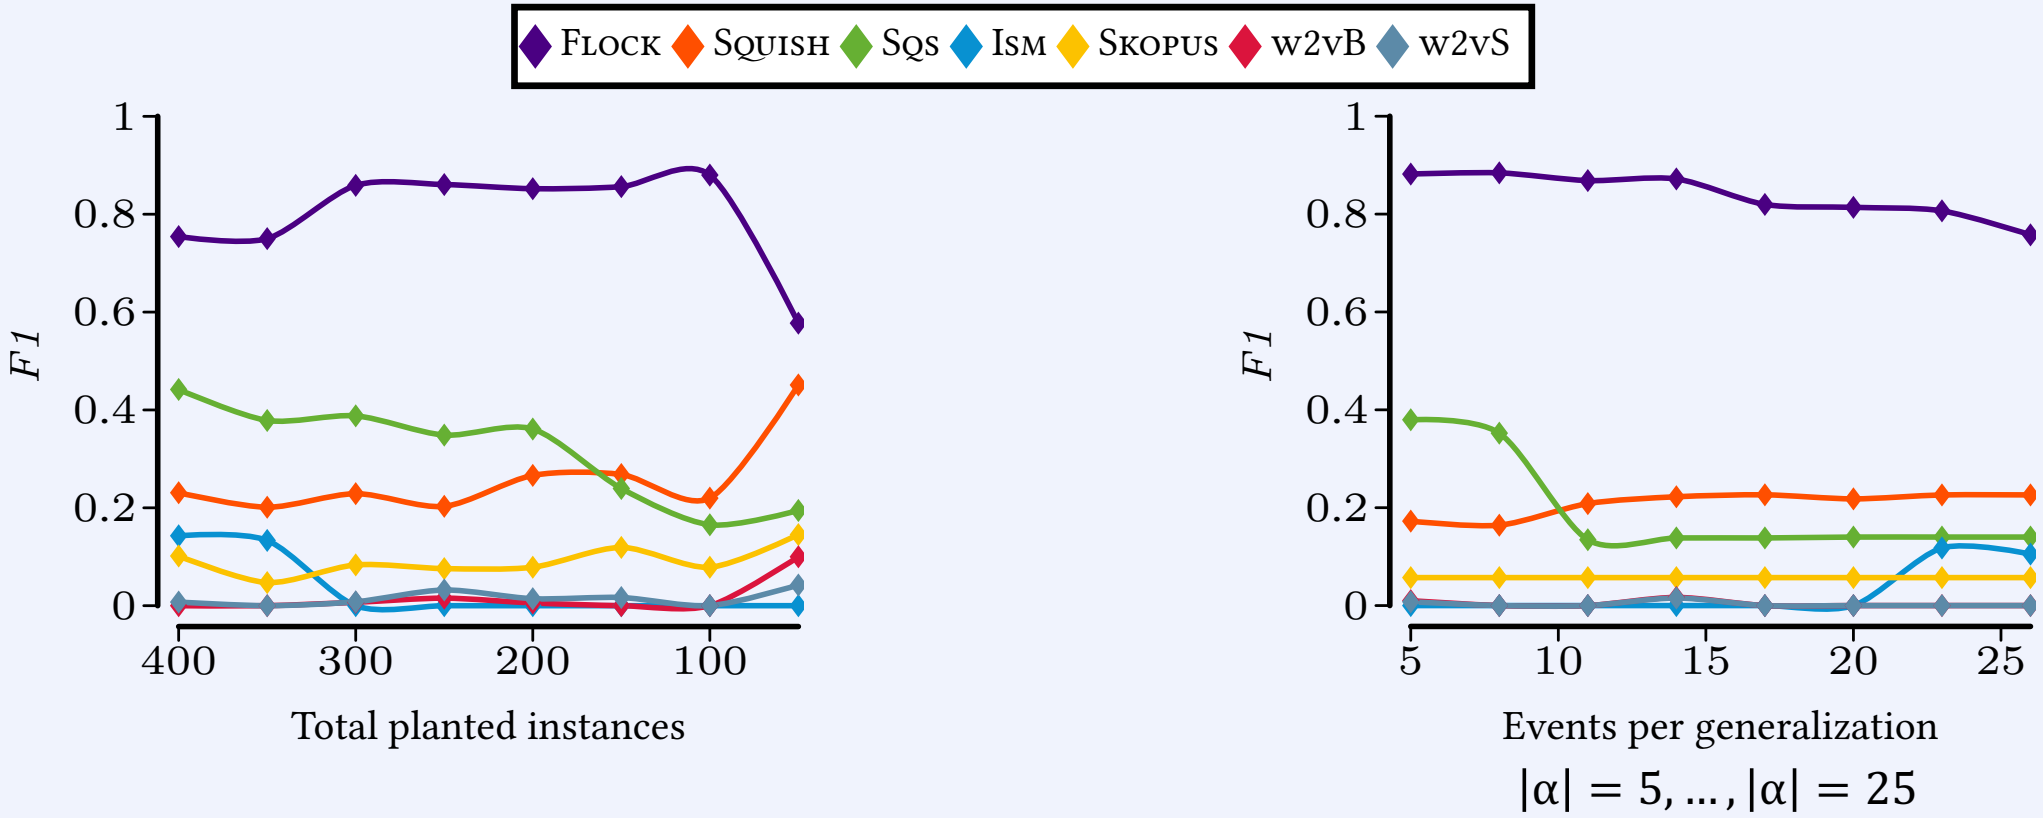

Real World Pattern Example

- Data: Production Log of Steel Rolling Mill

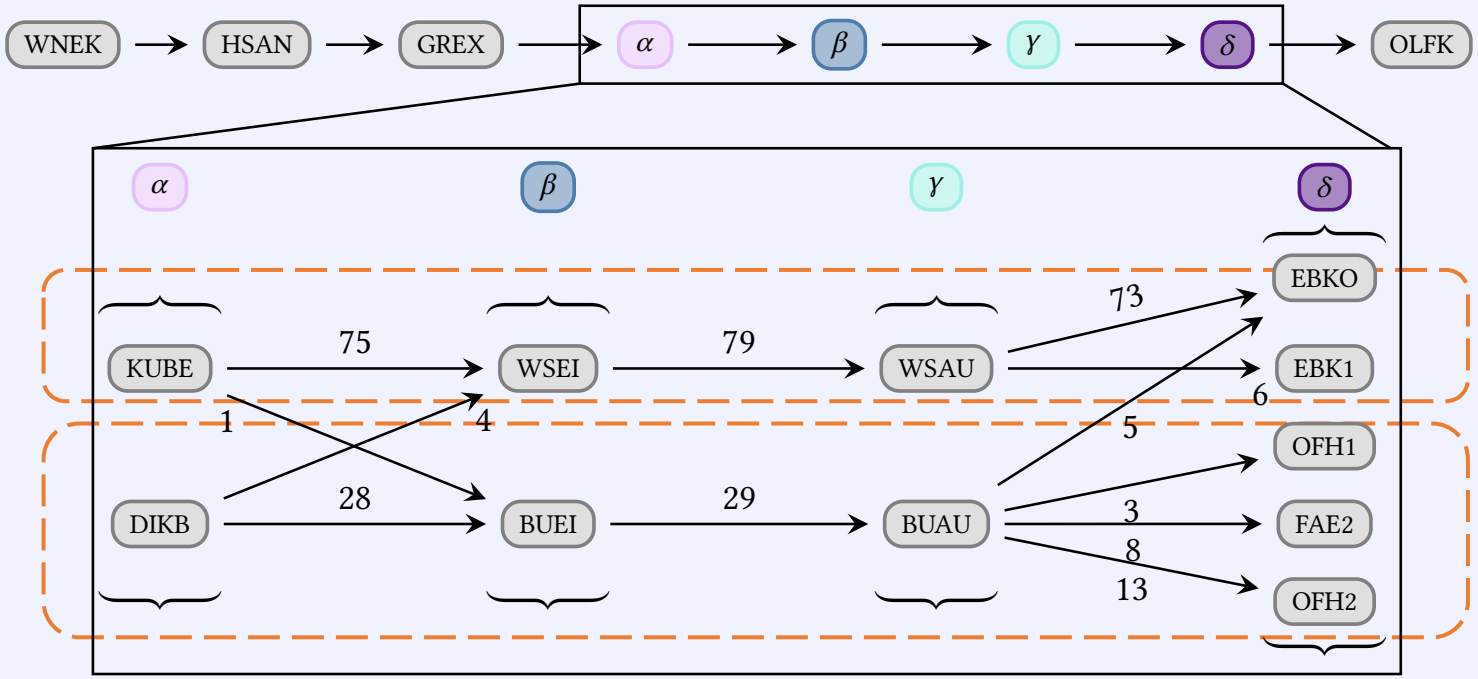

line for wide steel

line for thin steel

Image: https://c1.wallpaperflare.com/preview/82/931/800/bloom-rolling-mill-rail-hot-product.jpg

Something Else

Patterns with Predictable Inter-Event Delays

- Explicit modeling of delays between events
- Ability to model and **discover** patterns with long inter-event delays

Conclusion

We consider the problem of finding a succinct set of generalized patterns that describes the data

- Generalized pattern describe general "behavior"– not instances
- Capture infrequent instances of general patterns

Formalized the problem with the Minimum Description Length (MDL) principle

- Define model class and encoding of model
- Encoding of Data given a Model

Present greedy algorithm to mine patterns and generalized events

Evaluation shows that we can discover generalized patterns

- Recover ground truth well on synthetic data

Thank you!

We consider the problem of finding a succinct set of generalized patterns that describes the data

- Generalized pattern describe general "behavior"– not instances
- Capture infrequent instances of general patterns

Formalized the problem with the Minimum Description Length (MDL) principle

- Define model class and encoding of model
- Encoding of Data given a Model

Present greedy algorithm to mine patterns and generalized events

Evaluation shows that we can discover generalized patterns

- Recover ground truth well on synthetic data

Transition probabilities / frequencies

Pattern:

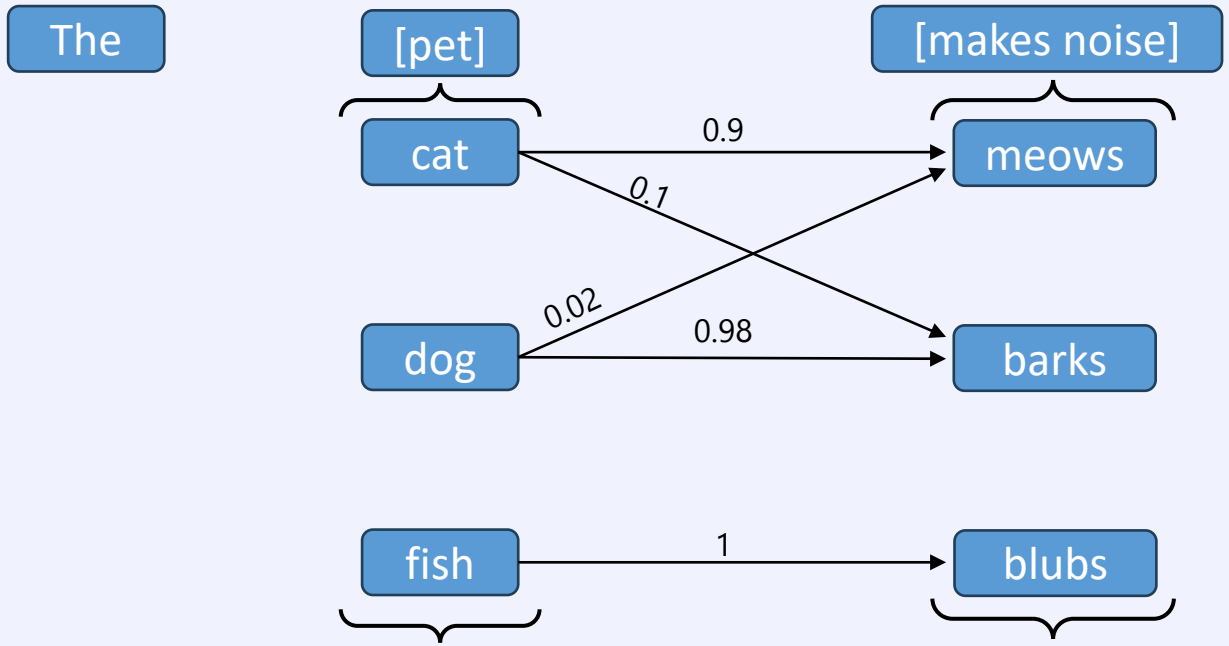

Evaluation – Synthetic Data

- Data with known ground truth.

0 Generalized Patterns ⇨ 1 Generalized Patterns ⇨ ...

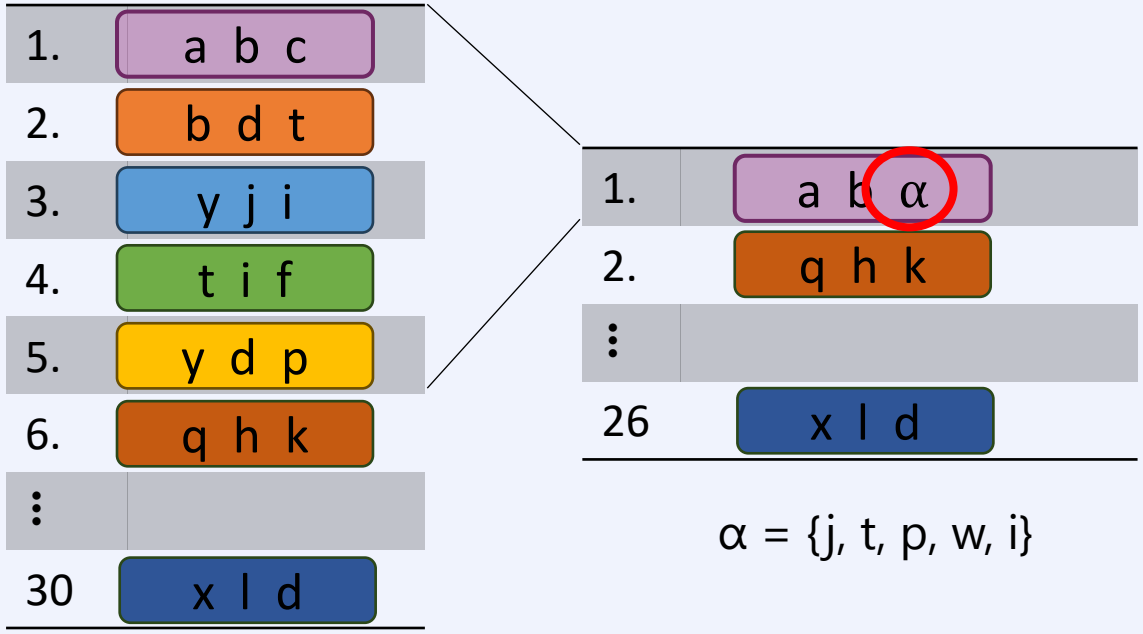

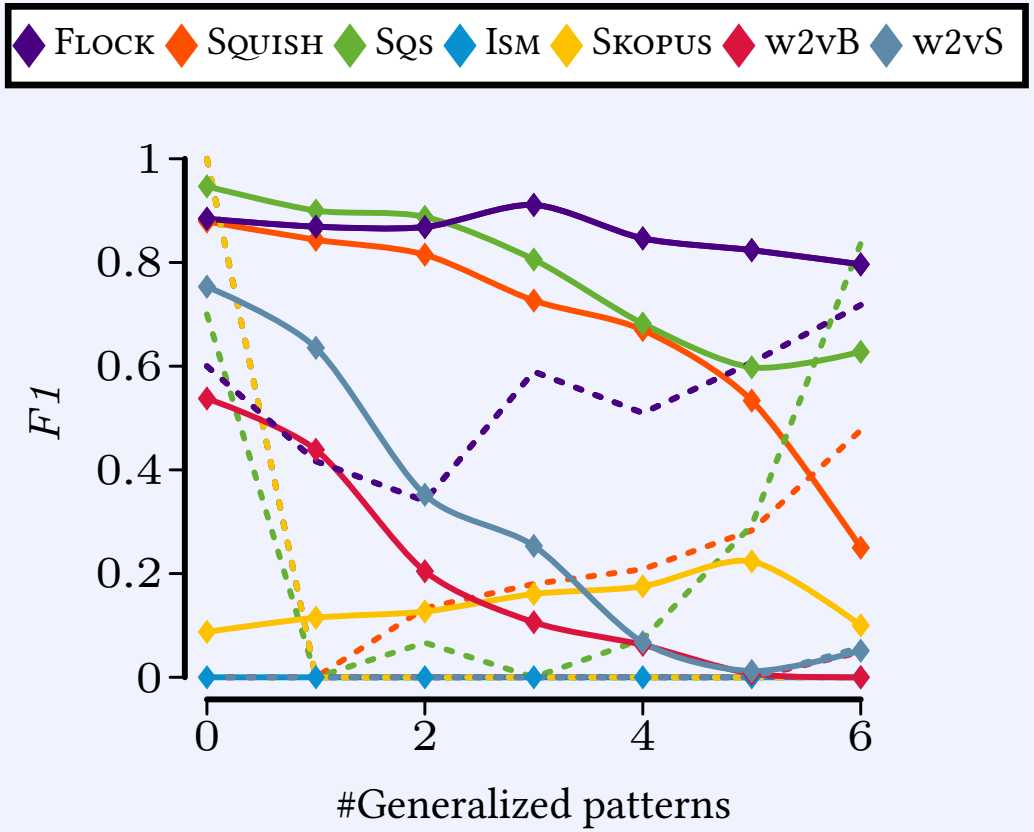

Evaluation – Synthetic Data

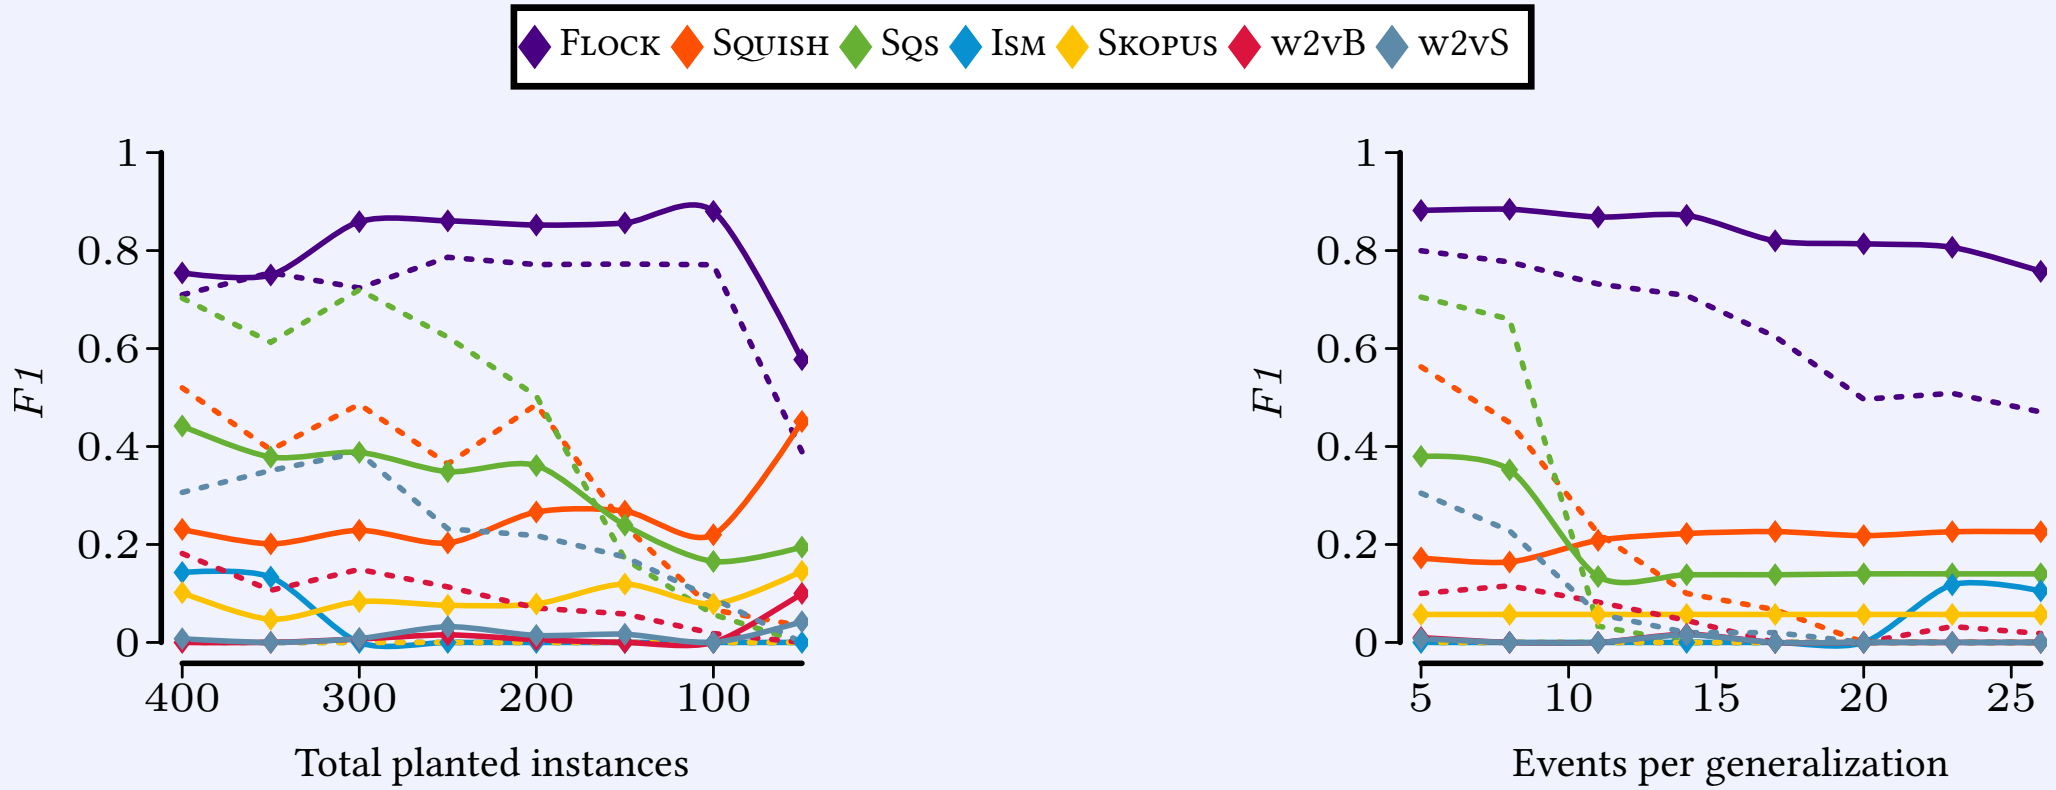

$|\alpha| = 5, \ldots, |\alpha| = 25$