

Dynamic Pattern Detection with Connectivity and Temporal Consistency Constraints


Skyler Speakman & Daniel B. Neill
Event and Pattern Detection Lab
Heinz College, Carnegie Mellon University

This work was partially supported by NSF grants:
IIS-0916345, IIS-0911032, and IIS-0953330



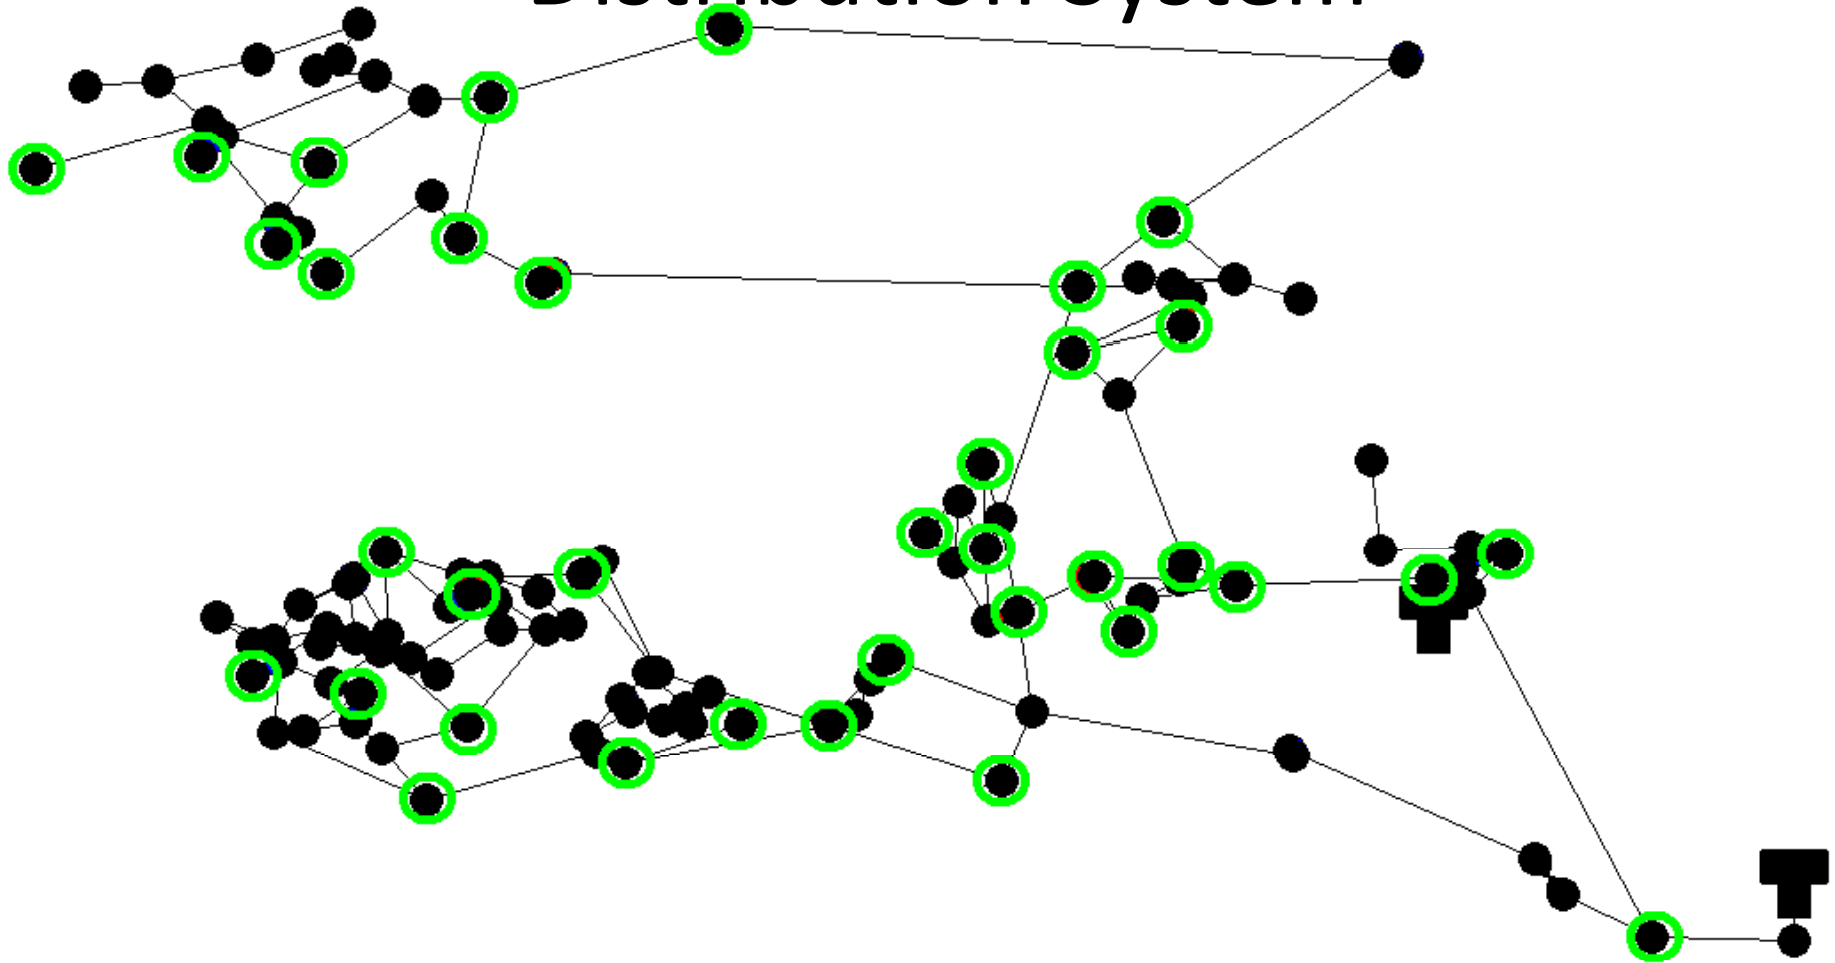
Spreading Contaminants in a Water Distribution System

Day 1, 6:00 PM

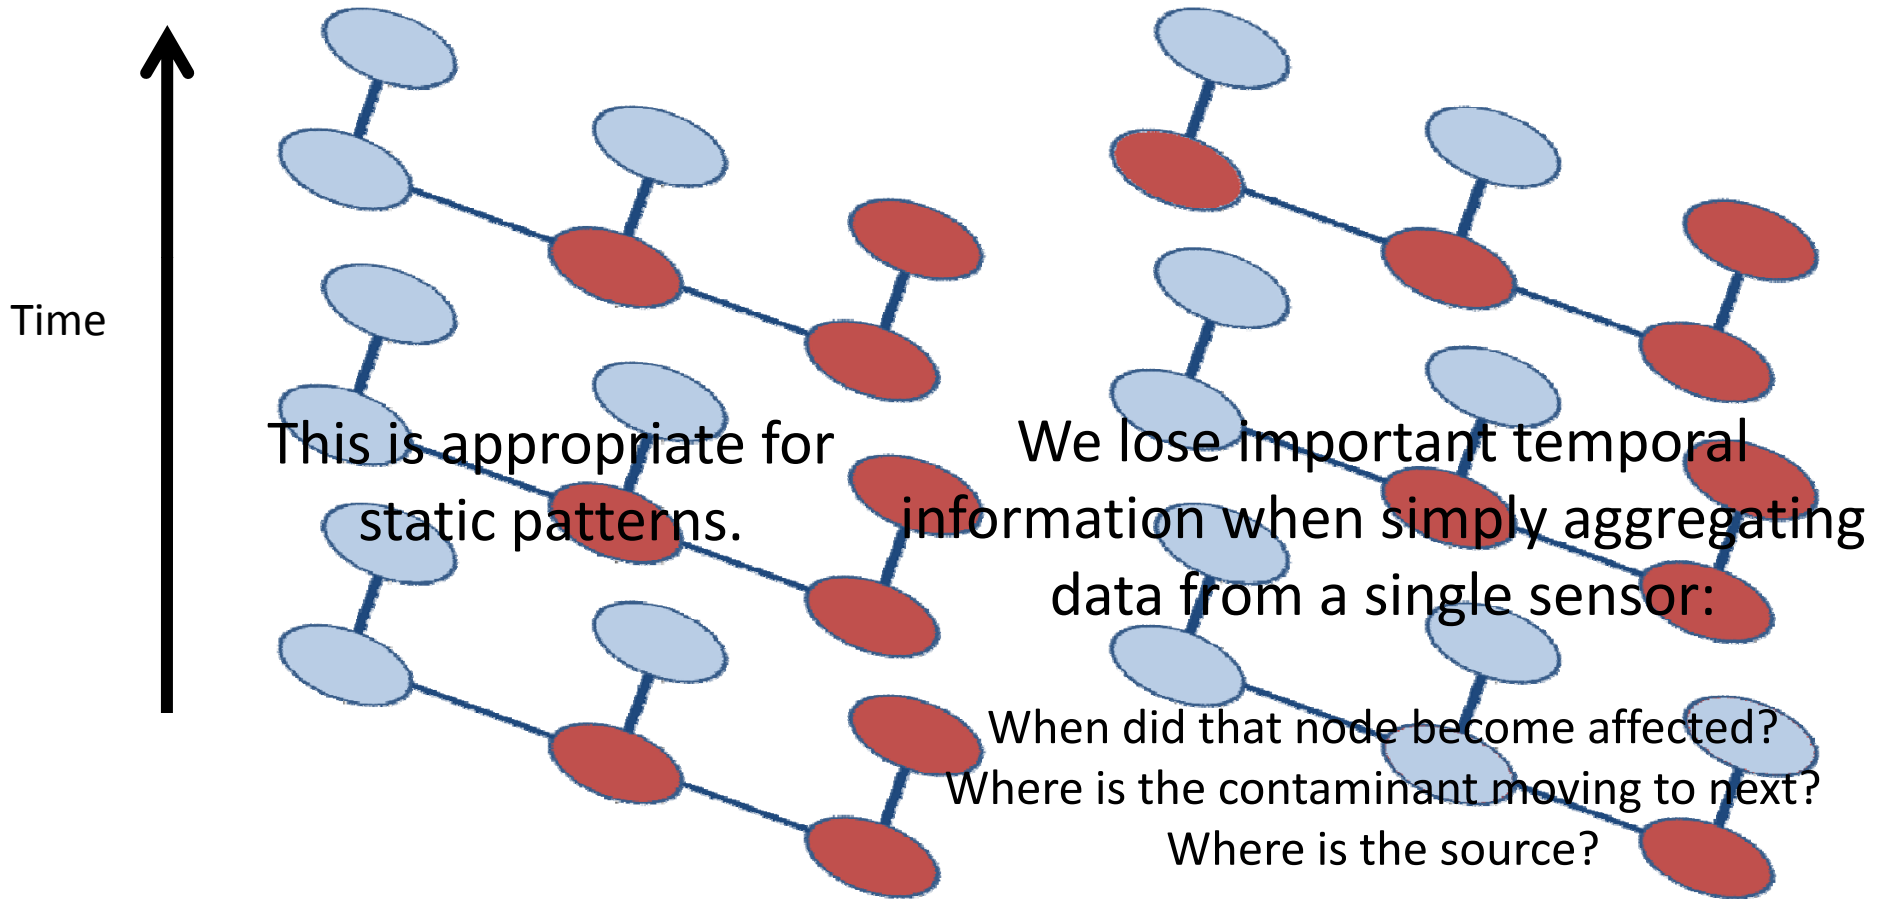


Spreading Contaminants in a Water Distribution System


Day 1, 6:00 PM



Temporal Information: Aggregating



Temporal Consistency



Goals of this Presentation

Provide a framework for incorporating
“soft constraints”
(i.e. temporal consistency)
without violating the properties that allow for the efficient
search over the network

Describe properties of the
expectation-based binomial
scoring function which is appropriate
when dealing with binary sensor data

Show empirical results which
demonstrate the utility of incorporating
soft constraints when detecting
dynamic patterns such as contaminant
plumes in a water distribution system

Pattern Detection as a Search Over Subsets

Pattern Detection can be framed as a search over subsets of the data with the goal of finding the subset which best matches a probabilistically modeled pattern.

This “match” is quantified by a scoring function, typically a **likelihood ratio**.

Computational Problems: Infeasible to perform exhaustive search for more than 30 records

Linear-time Subset Scanning (LTSS)

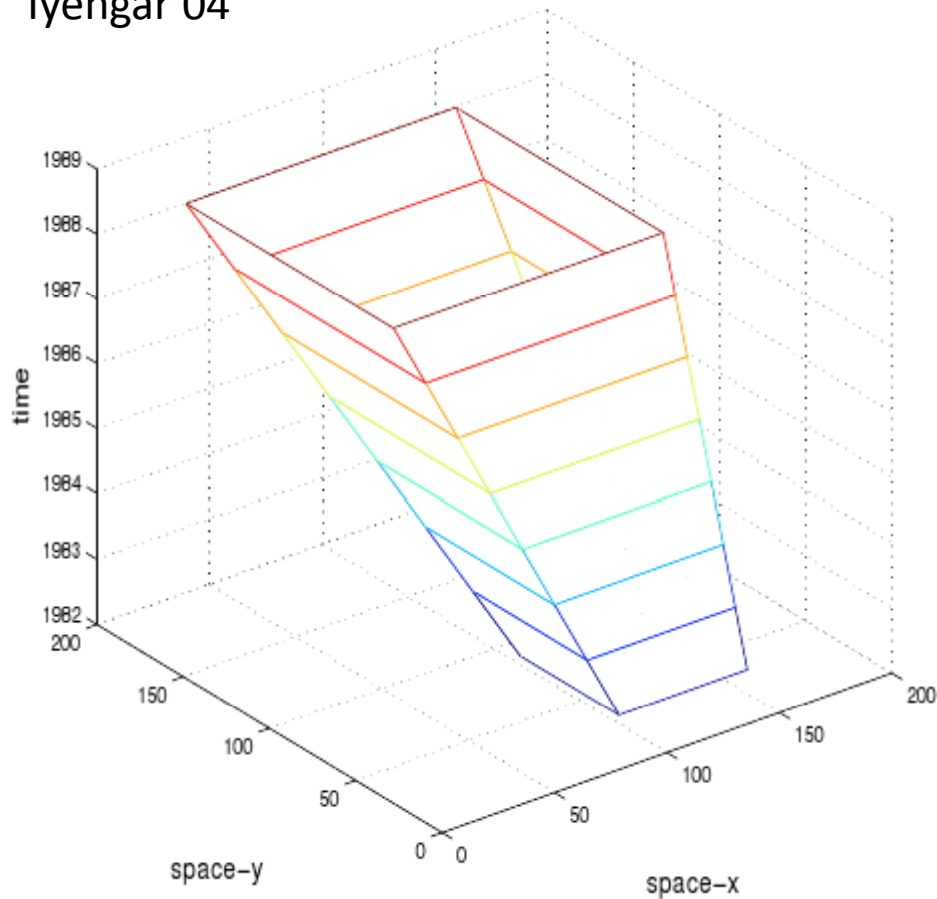
property allows for exact, efficient identification of “highest scoring” subset without an exhaustive search

Neill 2010

GraphScan applied LTSS to only consider **connected** subsets. Increases power to detect patterns that affect a subgraph of a larger network.

Speakman & Neill 2010

Iyengar 04




Use geometric shapes (truncated cones) to aggregate temporal information

Assumes linear growth or movement of the pattern

Potential shapes to search over grow exponentially and therefore relies on a heuristic to approximate best subset

(Neill, 2010)

The highest scoring subset is guaranteed to be one of the following subsets



Decreases the search space from 2^N to N

Naively altering the scoring function to enforce soft constraints violates LTSS

PROBLEM:	We must alter the <i>scoring function</i> instead of restricting the search space.
When applied directly, these constraints <i>violate</i> the <i>LTSS property</i> of the scoring function and make exact, efficient search <i>impossible</i> .	
SOLUTION:	Interpret the scoring function as a <i>sum</i> of <i>contributions</i> from each record in the subset.
<i>Maximizing</i> the scoring function is then equivalent to selecting all records that are making a <i>positive contribution</i> .	
INSIGHT:	When treated as an additive function, <i>further terms</i> may be introduced without interfering with the maximization step.

Enforcing Soft Constraints with LTSS


Scoring Function

$$F(S) = \log \frac{P(\text{Data} \mid H_1(S))}{P(\text{Data} \mid H_0)}$$

$$H_0 : c_i \sim \text{Bernoulli}(p_0)$$


$$H_1 : c_i \sim \text{Bernoulli}(qp_0) \quad 1 < q < \frac{1}{p_0}$$

$$F(S) = \max_q \log \frac{P(\text{Data} \mid H_1(S))}{P(\text{Data} \mid H_0)}$$




Large number
nodes with a
moderate risk

Small number of
nodes with a
high risk



Adding Soft Constraints to the Scoring Function

$$F(S) + \sum_{s_i \in S} \Delta_i$$


SOLUTION:

Interpret the scoring function as a **sum** of **contributions** from each record in the subset.

Maximizing the scoring function is then equivalent to selecting all records that are making a **positive contribution**.

INSIGHT:

When treated as an additive function, **further terms** (i.e., soft constraints) may be introduced without interfering with the maximization step.

$$F(S) = \max_q \sum_{s_i \in S} [F(s_i | q) + \Delta_i]$$



Enforcing Soft Constraints with LTSS

PROBLEM: We must alter the *scoring function* instead of restricting the search space.

When applied directly, these constraints *violate* the *LTSS property* of the scoring function and make exact, efficient search *impossible*.

SOLUTION: Interpret the scoring function as a *sum* of *contributions* from each record in the subset.

Maximizing the scoring function is then equivalent to selecting all records that are making a *positive contribution*.

INSIGHT: When treated as an additive function, *further terms* (i.e. soft constraints) may be introduced without interfering with the maximization step.

$$F(S) = \max_q \sum_{s_i \in S} F(s_i | q)$$



Enforcing Soft Constraints with LTSS

PROBLEM: We must alter the *scoring function* instead of restricting the search space.

When applied directly, these constraints *violate* the *LTSS property* of the scoring function and make exact, efficient search *impossible*.

SOLUTION: Interpret the scoring function as a *sum* of *contributions* from each record in the subset.

Maximizing the scoring function is then equivalent to selecting all records that are making a *positive contribution*.

INSIGHT: When treated as an additive function, *further terms* (i.e. soft constraints) may be introduced without interfering with the maximization step.

$$F(S) = \max_q \sum_{s_i \in S} [F(s_i | q) + \Delta_i]$$



Enforcing Soft Constraints with LTSS

PROBLEM: We must alter the *scoring function* instead of restricting the search space.

When applied directly, these constraints *violate* the *LTSS property* of the scoring function and make exact, efficient search *impossible*.

SOLUTION: Interpret the scoring function as a *sum* of *contributions* from each record in the subset.

Maximizing the scoring function is then equivalent to selecting all records that are making a *positive contribution*.

INSIGHT: When treated as an additive function, *further terms* (i.e. soft constraints) may be introduced without interfering with the maximization step.

Demonstration with Expectation-based Binomial

$$F(S) = \max \log \frac{P(\text{Data} | H_1(S))}{P(\text{Data} | H_0)} \quad \begin{array}{l} H_0 : c_i \sim \text{Bernoulli}(p_0) \\ H_1 : c_i \sim \text{Bernoulli}(qp_0) \quad 1 < q < \frac{1}{p_0} \end{array}$$

$$F(S) = \max_{1 < q < 1/p_0} \log \prod_{s_i \in S} \frac{(qp_0)^{c_i} (1 - qp_0)^{N_i - c_i}}{p_0^{c_i} (1 - p_0)^{N_i - c_i}}$$

$$F(S) = \max_{1 < q < 1/p_0} \sum_{s_i \in S} c_i \log q + \sum_{s_i \in S} (N_i - c_i) \log \left(\frac{1 - qp_0}{1 - p_0} \right)$$

$$F(S) = \underbrace{C}_{\text{\#Triggers}} \log \left(\frac{C}{Np_0} \right) + (N - C) \log \left(\frac{1 - C/N}{1 - p_0} \right) \text{ for } C > Np_0, 0 \text{ otherwise}$$

\#Trials
Baseline rate

Demonstration with Expectation-based Binomial

$$F(S) = \max_{1 < q < 1/p_0} \sum_{s_i \in S} c_i \log q + \sum_{s_i \in S} (N_i - c_i) \log \left(\frac{1 - qp_0}{1 - p_0} \right)$$

$$F(S) = \max_{1 < q < 1/p_0} \sum_{s_i \in S} \left[c_i \log q + (N_i - c_i) \log \left(\frac{1 - qp_0}{1 - p_0} \right) \right]$$

$$F(S | q) = \sum_{s_i \in S} \left[c_i \log q + (N_i - c_i) \log \left(\frac{1 - qp_0}{1 - p_0} \right) \right]$$

Contribution from each sensor in subset

$$F(S | q) = \sum_{s_i \in S} \left[c_i \log q + (N_i - c_i) \log \left(\frac{1 - qp_0}{1 - p_0} \right) + \triangle_i \right]$$

Log-likelihood $F(s_i | q)$

Reward /Penalty from constraints

Demonstration with Expectation-based Binomial

$$F(S) = \max \log \frac{P(\text{Data} | H_1(S))}{P(\text{Data} | H_0)}$$

$$F(S) = \max_{1 < q < 1/p_0} \log \prod_{s_i \in S} \frac{(qp_0)^{c_i} (1-qp_0)^{N_i-c_i}}{p_0^{c_i} (1-p_0)^{N_i-c_i}}$$

Contribution from each sensor, for a fixed q

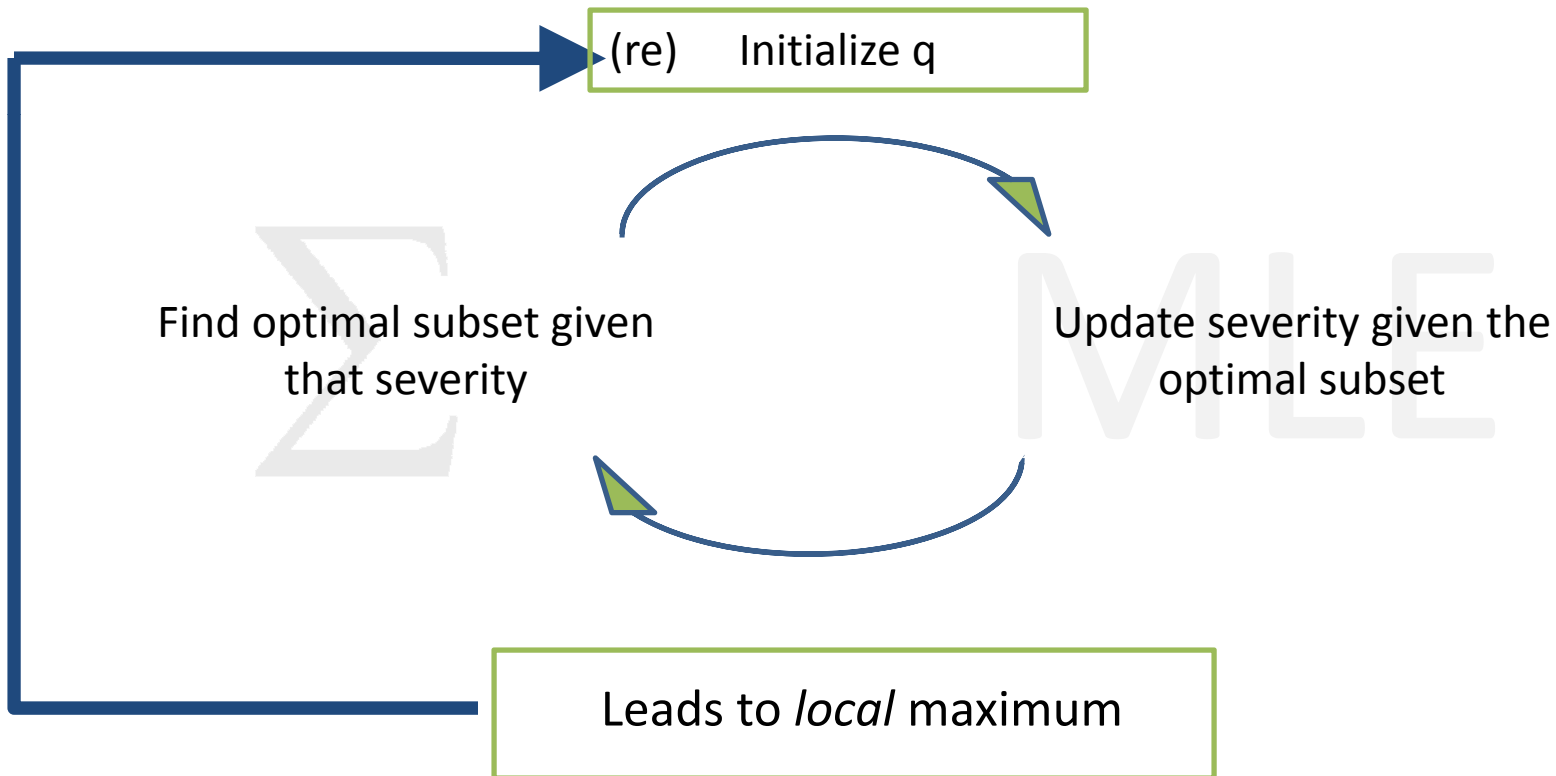
$$F(S | q) = \sum_{s_i \in S} \left[c_i \log q + (N_i - c_i) \log \left(\frac{1-qp_0}{1-p_0} \right) + \Delta_i \right]$$

Log-likelihood $F(s_i | q)$


Reward / Penalty from constraints

From *Fixed* q to *All* q


Our goal is to maximize $F(S)$ *over all* q



Additive GraphScan



Additive GraphScan




Additive GraphScan is 10x faster than regular, even though operating with multiple iterations


We represent groups of subsets as a string of 0's, 1's, and ?'s

Priority Ranking	1	2	3	4	5	6
Bit String	1	0	0	1	?	?

The above bit string represents 4 possible subsets:
 {1,4} {1,4,5} {1,4,6} {1,4,5,6}

A Naïve approach would search all 2^N subsets and is computationally infeasible






Seed nodes have higher priority than all of their neighbors

We can rule out bit strings whose highest priority node is not a seed node

	1	2	3	4	5	6
s_1	1	?	?	?	?	?
s_2	0	1	?	0	0	0
s_3	0	0	1	?	?	?
s_4	0	0	0	1	?	0
s_5	0	0	0	0	1	?
s_6	0	0	0	0	0	1

If we rule out a high priority node, we can also rule out all of its lower priority neighbors...

...and any additional nodes that are disconnected when these nodes are ruled out



	1	2	3	4	5	6
S_1	1	?	?	?	?	?
S_2	0	1	?	0	0	0
S_3	0	0	1	?	?	?
S_4	0	0	0	1	?	0
S_5	0	0	0	0	1	?
S_6	0	0	0	0	0	1

Propagation of bit strings:

Pull off S_1 and consider the two cases of including or excluding node 2

Including node 2 implies including nodes 3 and 6

S_{1a} : 1 1 1 ? ? 1

Excluding node 2 implies excluding nodes 3, 4, 5, and 6


S_{1b} : 1 0 0 0 0 0

Data: Battle of the Water Sensor Networks

Day 1, 3:00 PM

Plumes of contaminants are simulated in a water distribution system

We assume the system is equipped with *imperfect* sensors



Competing Methods

Upper Level Sets:

A heuristic that is not guaranteed to find the most anomalous subgraph

Patil & Taillie, 2004

ULS

GraphScan:

Determines the most anomalous subgraph *without further constraints*


Speakman & Neill, 2010

GS

Competing Methods

GraphScan with basic temporal consistency


$$F(S) = \max_q \sum_{s_i \in S} (F(s_i | q) + \Delta_i) \quad \Delta_i = \begin{cases} +\Delta & \text{if } s_i \in \Omega \\ 0 & \text{otherwise} \end{cases} \quad \text{ADD-GS}$$



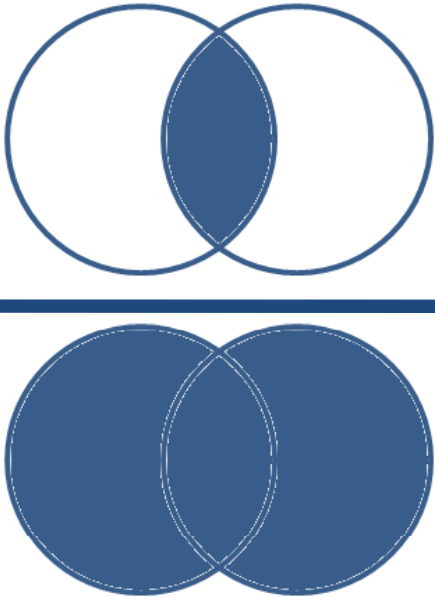
Competing Methods

GraphScan with temporal consistency: Overlap and Neighbors

$$F_{pen}(\mathcal{S}) = \max_q \sum_{s_i \in \mathcal{S}} (F(s_i | q) - \lambda + \Delta_i) \quad \Delta_i = \begin{cases} +\Delta & \text{if } s_i \in \Omega_{neighbor} \\ 0 & \text{otherwise} \end{cases} \quad \text{GS-SON}$$




Results: Spatial Overlap

$$\textit{Overlap} = \frac{A \cap B}{A \cup B} = \frac{\text{Diagram 1}}{\text{Diagram 2}}$$


Overlap = 1 Perfect Match


Overlap = 0 Completely Disjoint

Overlap Coefficient for "Easy" Case $p_0 = 0.1$ and $p_1 = 0.9$




Hours until Detection	% Detected
7.66	100%
9.65	97.5%
15.4	92.4%

Overlap Coefficient for "Easy" Case $p_0 = 0.1$ and $p_1 = 0.9$



Hours	% Detected
7.66	100%
7.82	100%
10.27	97.6%
9.65	97.5%
15.4	92.4%


Overlap Coefficient for "Easy" Case $p_0 = 0.1$ and $p_1 = 0.9$



Hours	% Detected
7.66	100%
8.2	100%
4.6	100%


Overlap Coefficient for "Medium" Case

$p_0 = 0.2$ and $p_1 = 0.8$




Overlap Coefficient for "Medium" Case

$p_0 = 0.2$ and $p_1 = 0.8$




Overlap Coefficient for "Hard" Case $p_0 = 0.4$ and $p_1 = 0.6$




	Hours Until Detection	% Detected
ADD-GS	27.4	52.5%
GS	31.0	26.8%
ULS	33.5	15.8%

Overlap Coefficient for "Hard" Case $p_0 = 0.4$ and $p_1 = 0.6$




Hours Until Detection	% Detected
27.4	52.5%
31.0	26.8%
33.5	15.8%

Overlap Coefficient for "Hard" Case $p_0 = 0.4$ and $p_1 = 0.6$



Hours	% Detected
27.4	52.5%
30.6	26.8%
28.1	48.0%
22.6	76.8%
33.5	15.8%

Overlap Coefficient for "Hard" Case $p_0 = 0.4$ and $p_1 = 0.6$



Hours	% Detected
27.4	52.5%
30.6	26.8%
28.1	48.0%
22.6	76.8%
33.5	15.8%

Conclusions

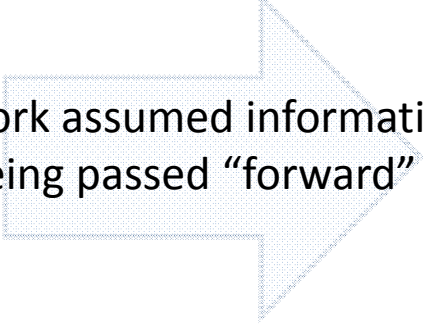
Provided a framework that allows soft constraints to influence the scoring function and give preference to subsets of desired temporal consistency while still allowing an efficient search for the highest scoring connected sub-graph

Applied the EBB scoring function with **Additive GraphScan** to the task of detecting contaminants in a water distribution system

Empirical results showed temporal consistency constraints **reduced the time to detect** the contaminants and **increased spatial accuracy** of the methods

Future work: Forward & Backward Consistency

This work assumed information was only being passed “forward” in time

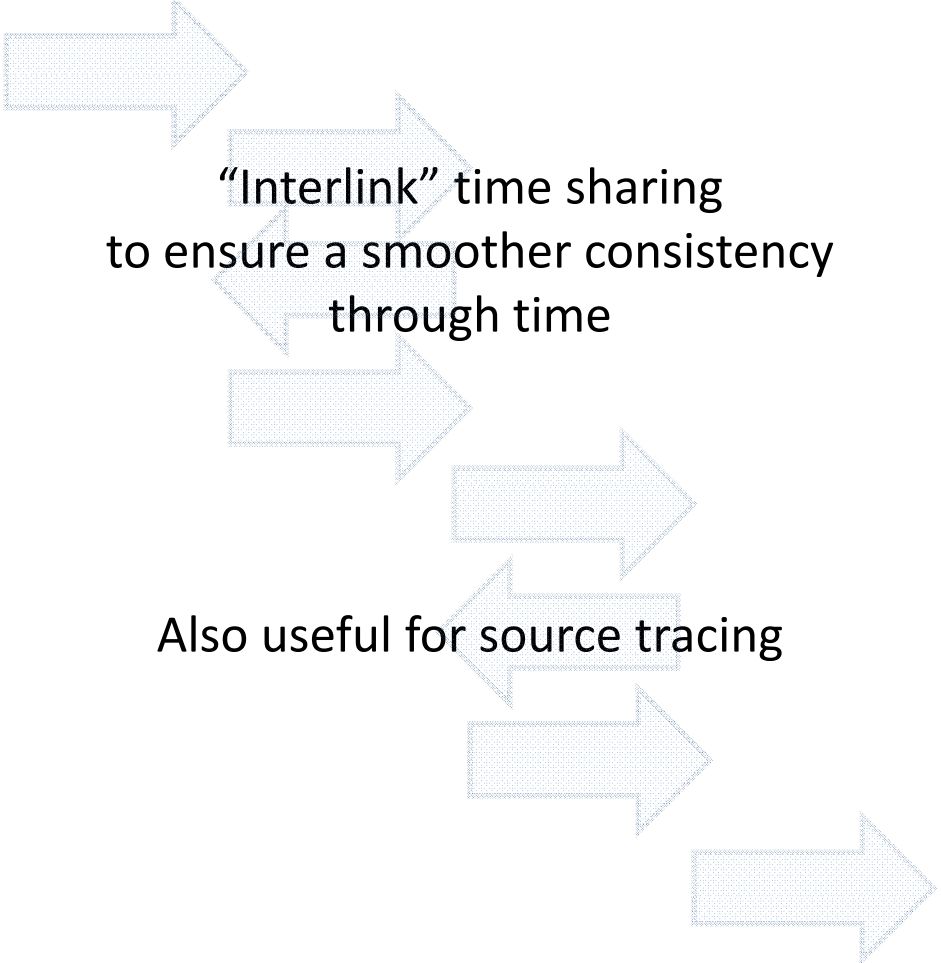


We can also share current information with the past

Alter past subsets to be consistent with current information



Use the altered subsets from the past to make more informed searches in present



“Interlink” time sharing
to ensure a smoother consistency
through time

Also useful for source tracing



Thank you!

speakman@cmu.edu

neill@cs.cmu.edu