Assessing the Impact of Network Events with User Feedback

Shobha Venkataraman, Jia Wang
{shvenk, jiawang}@research.att.com

August 24, 2018

Motivation

Many users call customer service centers when a network problem occurs

Definitions:

Event: a service-impacting network problem that affects multiple users, typically within a bounded geographic area, e.g., a network outage

Case: Agent-generated summary of user’s problems, troubleshooting steps taken and results.

Question: Can we mine cases to understand user perspective of network events?

AT&T

Problem Statement

Definitions:

- Event-specific Cases: Cases pertaining to a network event
- Normal Cases: Cases that do not pertain to an event

Natural machine learning problem, however:

- Manually labeling data is infeasible at scale: ~~supervised learning~~
- Individual event makes up tiny fraction of total cases: ~~unsupervised learning~~
- Feasible to get a few labeled event cases: semi-supervised learning

Problem Statement: Given an event *e*, a set *U* of unlabeled cases and a set *L* with a few labeled event-specific cases and normal cases, find other cases pertaining to *e* from *U*

Challenges

User-sourced data noise

- Inaccuracy in reporting symptoms
- Symptoms dependent on device model, OS etc.
- Different symptoms at different locations
- Approximation of times observed symptoms
- Location data available at zip code granularity

Agent-sourced data noise

- Differences in writing style and language
- Different details of problem/debugging recorded
- Domain-specific abbreviations and language used
- Differences in understanding of network event

AT&T

Outline

Introduction

Data Characteristics

LOTUS System Design

Experiments

Data Characteristics: Case Definition

Timestamp

Case Attributes

Problem Characteristics

Device Characteristics

Agent Details

Problem Resolution

…

Case Notes

Free Text:

- Problem Symptoms
- Debugging Steps
- Current Status

Zipcode

Symptom information potentially duplicated in attributes and case notes!

Data Characteristics: Case Attribute Statistics

Extremely large number of case attribute combinations!

Table 1

Attribute	Values
Need	90
Subneed	475
Calltype	1367
Task Resolution	3334
Device Model	4168

No. of possible values for most relevant attributes

Many similar values in case attributes!

Table 2

Problem Description	Values
Call connectivity	8
Call quality	4
Text	8
Data connectivity	6
Data speed	5

No. of similar values for Calltypes describing various problems

Data Characteristics: Case Notes

Word	Synonym phrases found in case notes
customer	cx, cust, cus, cu, cxci, cci, ccio, cco, cic, ccii, cciu, ci, ccito, cvci, ccoi, cici, ccin, ccui, cxci, ccvi, cfci, cxcalled, custcalled, ccfi xcci, ccim, cdci, ccit
signal	bars, singal, strength, strenth, reception, strengh, siginal, sig, signa, singnal, recption, bar, receiption, sgnal, serviceat, sigl, signl, sgnl, signals
trouble	trbl, troble, touble, truble, troube, troubl, troule, trble, issue, issuses, issu, issueswith, prob, problem, prblems, promblems, problmes, difficulty, probelms

Observations:

- Many abbreviations and spelling errors for each word
- Often written concisely, as quickly as possible, so may miss key details
- Have unusual vocabulary, sentence structure and idiosyncrasies.

Language models trained on regular English cannot be used!

Outline

Introduction

Data Characteristics

LOTUS System Design

Experiments

LOTUS Overview

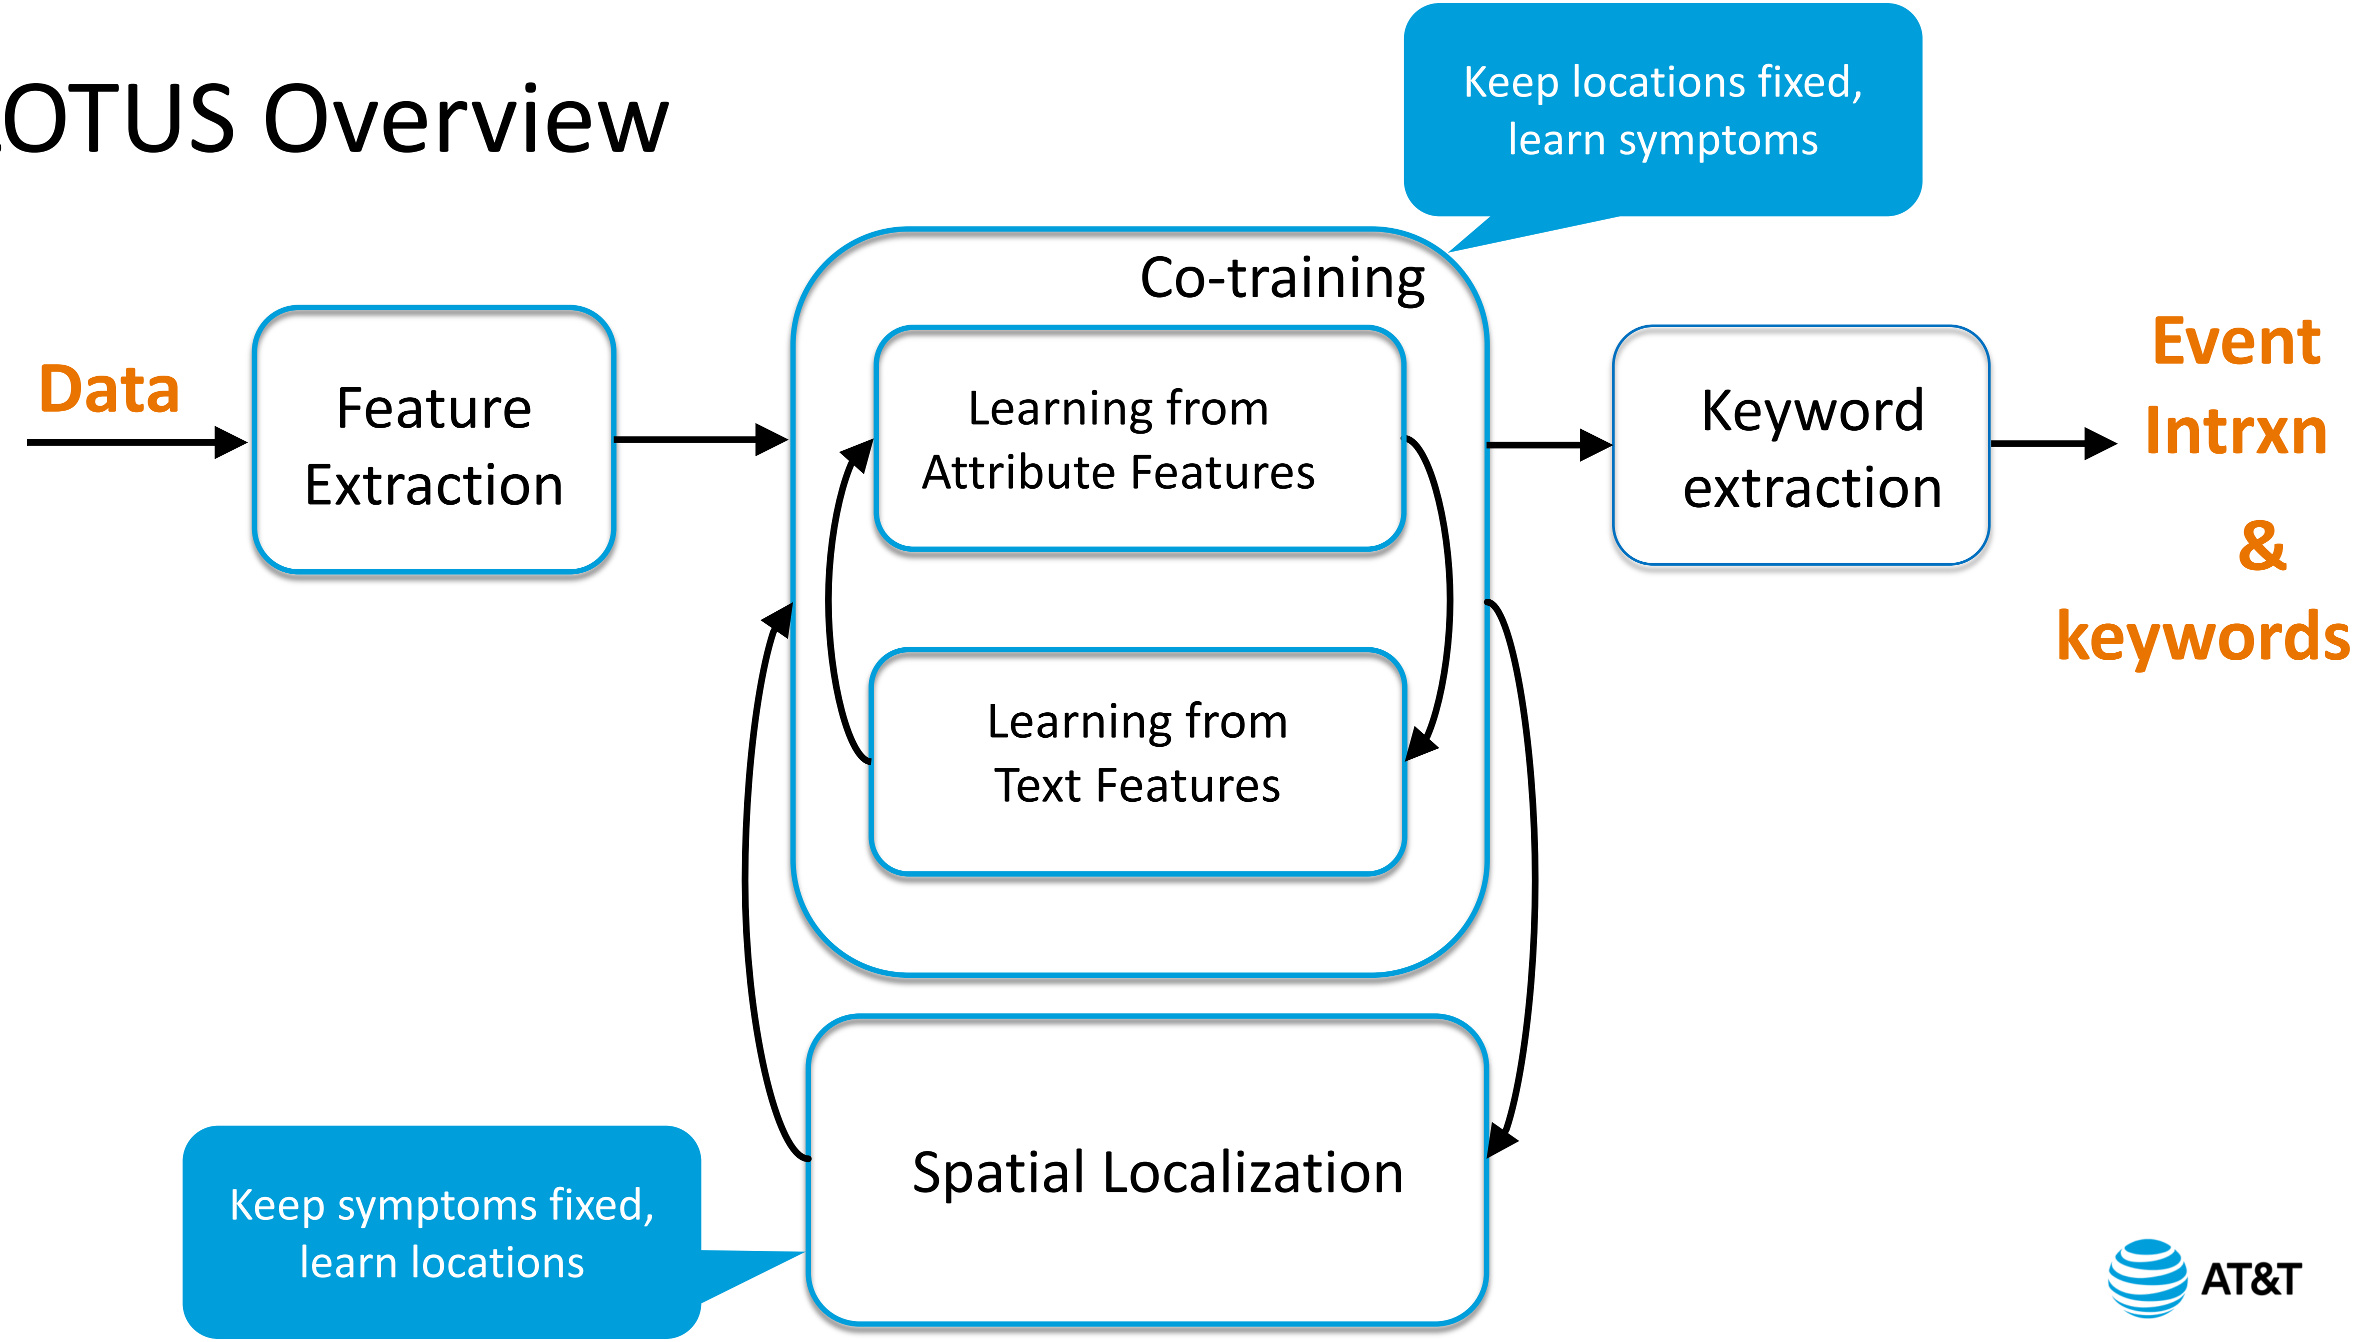

Feature Extraction

Convert each case to an example:

- Location
- Features of case attributes: standard one-hot encoding
- Features of case notes: word vectors [Mikolov+13]

Word vectors: Distributed representation of word/phrase using its context

For example:

"Customer wants to know why he cannot dial outgoing numbers successfully"

"Cx unable to make calls"

Why does it work on our data?

- Repetitive sentence structure
- Mixture of styles and vocabulary in sentences

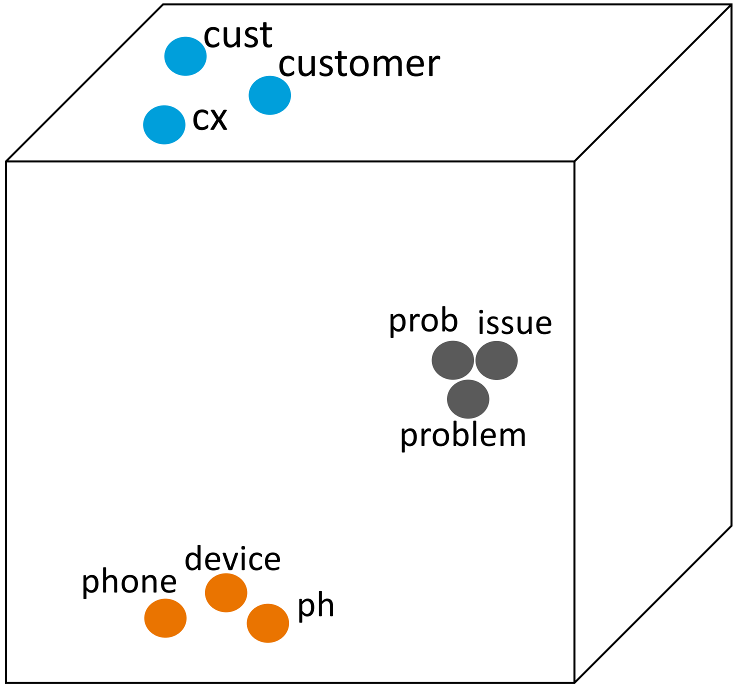

Sample word vectors in unit n-dimensional cube $[0,1]^n$

Distributed Representations of Words and Phrases and their Compositionality. Mikolov, Tomas and Sutskever, Ilya and Chen, Kai and Corrado, Greg S and Dean, Jeff. NIPS 2013

Co-training

Step overview: Location fixed, want to learn symptoms

Available initially: A few labeled cases and a lot of unlabeled cases

Learning model:

- Each case has both attribute features and text features
- Both attributes and text individually describe symptom and/or resolution

Thus, attributes and text each individually sufficient to judge whether case is event-specific

Co-training idea [Blum+98]:

- Simultaneously learn separate functions over attribute features and over text features
- Use learning over attribute features to bootstrap better learning over text features and vice versa

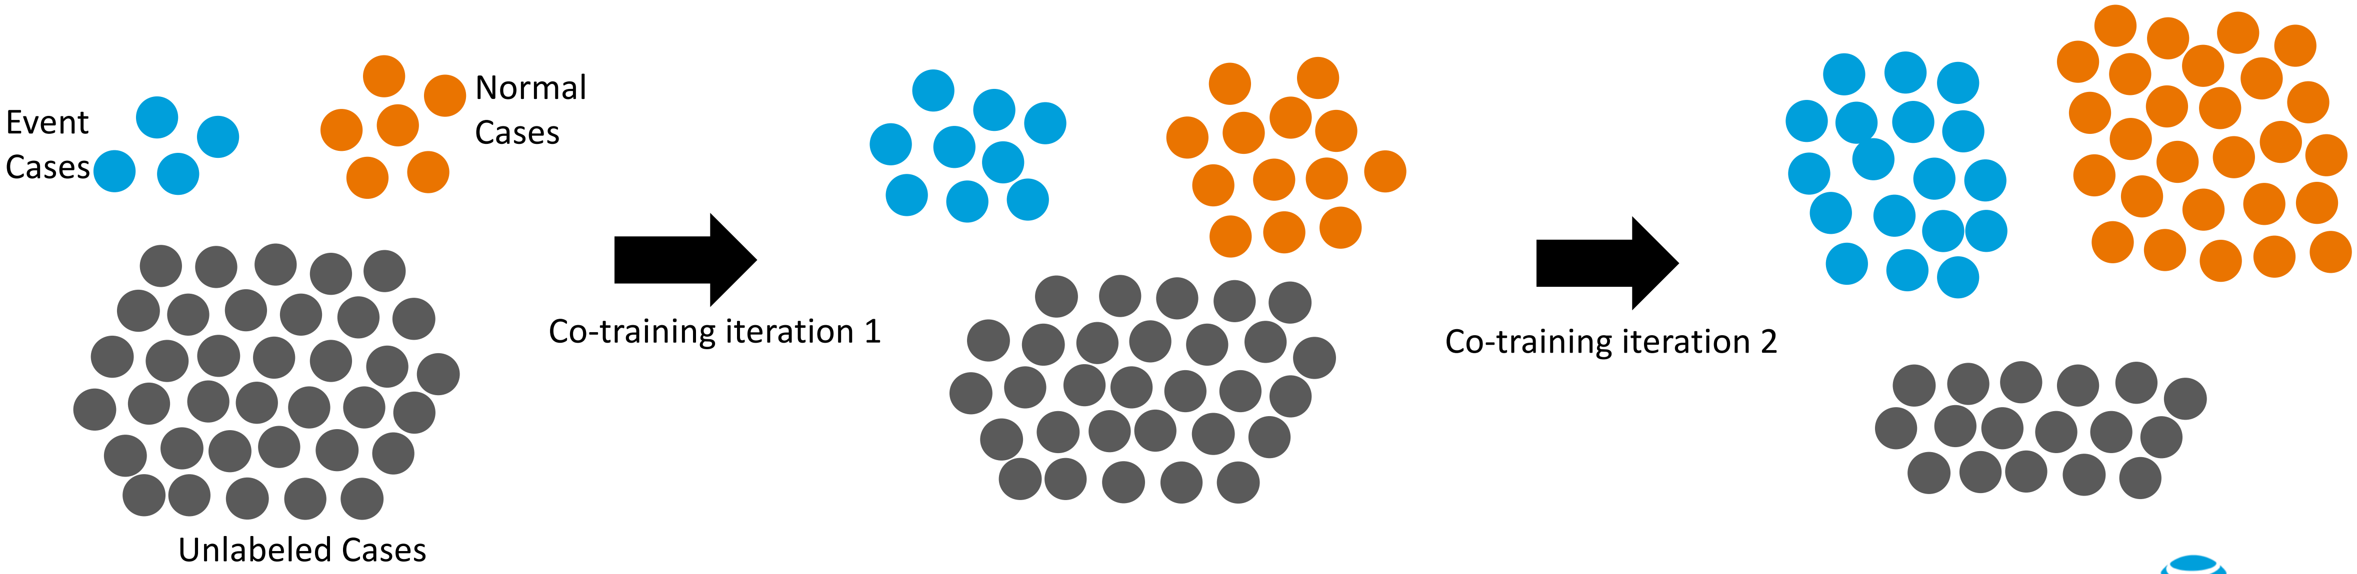

A. Blum and T. Mitchell. Combining labeled and unlabeled data with co-training. In COLT 1998.

Spatial Localization

Step overview: Symptoms fixed, want to learn location

Satscan [Kuldorff97]: Algorithm and software for computing spatial scan statistics to identify clusters of spatially correlated cases

- Uses likelihood ratio tests for common statistical models
- Normalizes for inhomogeneous population sizes

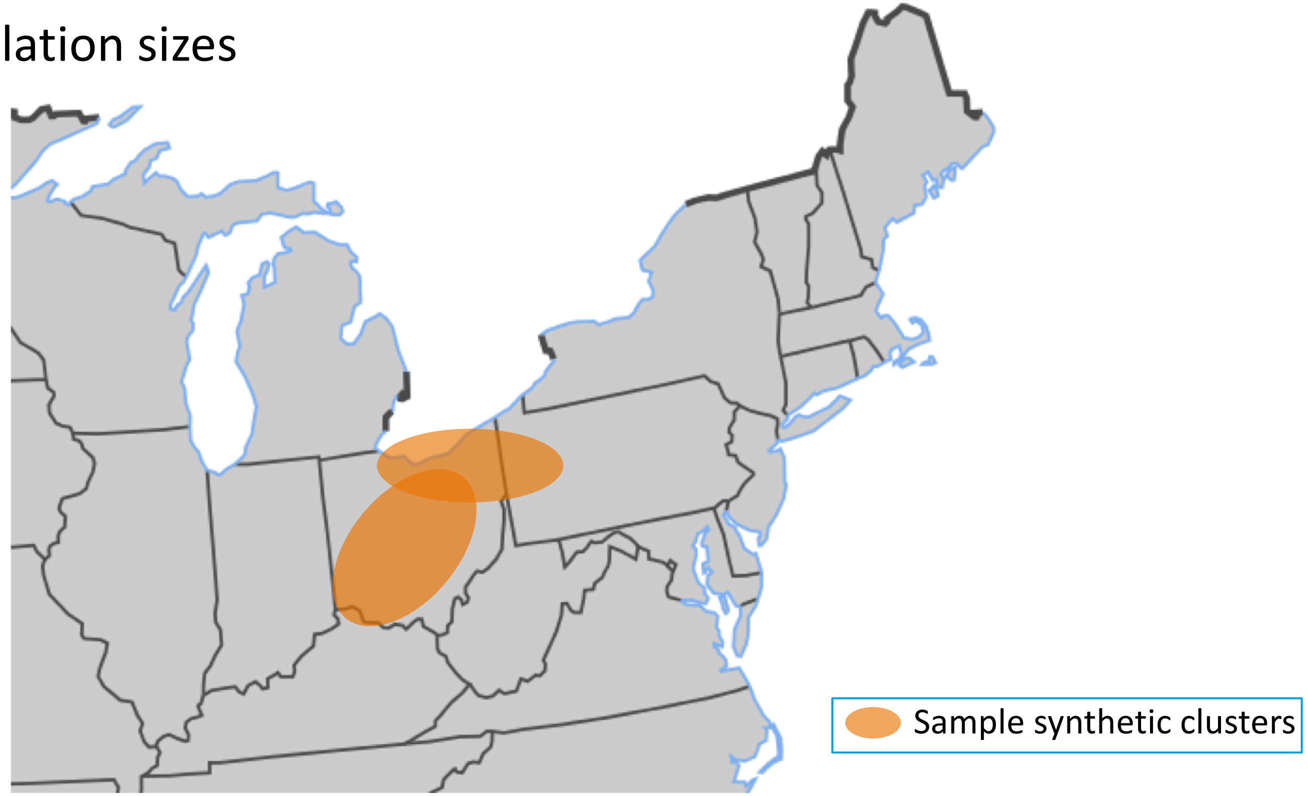

Kulldorff M., A spatial scan statistic. Communications in Statistics: Theory and Methods, 1997; 26:1481-1496. https://www.satscan.org/

Outline

Introduction

Data Characteristics

LOTUS System Design

Experiments

Experimental Setup

Synthetic Events to Measure Accuracy:

Procedure for creation of synthetic data

- Choose from a fixed possible combination of attributes
- Choose text that matches attributes from real event cases
- Choose zipcode based on population of affected area
- Background data for synthetic events: 350k normal cases

Baseline Algorithm:

- Supervised learning algorithm based on case attributes and case notes features
- Ensemble classifier to allow ease of fitting complex hypotheses

Run-time:

- About 10-20 minutes for typical run: around 1-2 days, 2-4 states
- Over 80% of time is spatial scan analysis (SatScan)

LOTUS Accuracy on Synthetic Events

Experimented on three types of synthetic events:

- Synthetic Event 1: Single fixed location, multiple fixed symptoms
- Synthetic Event 2: Multiple locations, changing symptoms
- Synthetic Event 3: Multiple locations, changing symptoms x 2

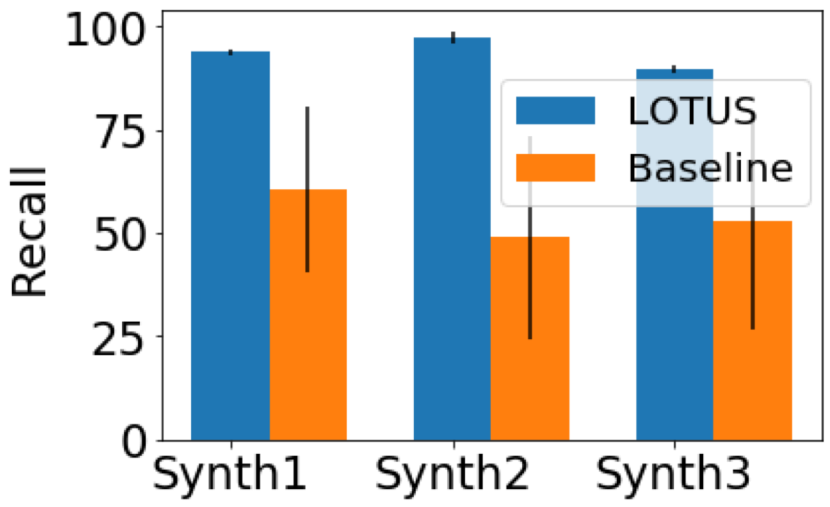

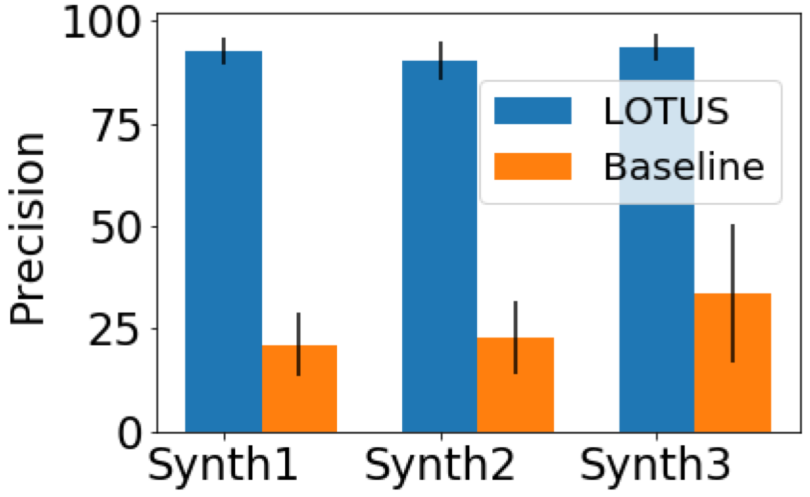

LOTUS precision & recall significantly exceed baseline in all cases!

LOTUS Results on Live Data

Two sample case studies:

Case Study 1

- Symptoms: all services affected, mainly voice + data complaints
- Locations: multiple metro areas
- Networks affected: 3G and LTE

Case Study 2

- Symptoms: only data issues
- Locations: multiple metro areas
- Networks affected: 3G and LTE

Observations:

- Over 75% of known event cases detected in both studies
- 97% of detected cases validated as event cases
- Keywords reflect event and symptom parameters

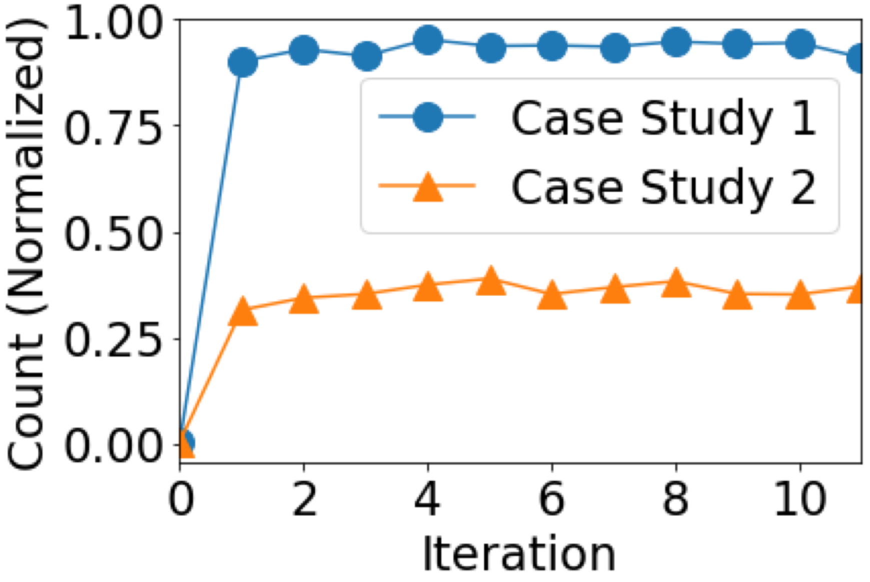

Event	Keyword Phrases
1	call, outage, message, sms, signal, disconnect, LTE, 3G, voicemail, degraded_tower, long_distance, microcell, outbound, video, hd, bar, poor, RAN
2	sms, message, safari, LTE, 3G, internet, RAN, drop, voicemail, website, block, slow, email, power_cycle, imessage_send, text, connect, hotspot, error, outage

Conclusions

Presented LOTUS system to assess impact of network events from user feedback

Novel algorithmic composition of semi-supervised learning and spatial scan statistics

Analyzes typical network events in 10-20 minutes with high accuracy

Questions?

AT&T