Human-AI Teaming for Decision Making

Augmenting Human Intellect with Visualization and Human-Centered AI

DONGYU LIU (dongyu@mit.edu)

POSTDOC AT MIT

PHD AT HKUST

香港科技大學
THE HONG KONG UNIVERSITY OF SCIENCE AND TECHNOLOGY

Massachusetts Institute of Technology

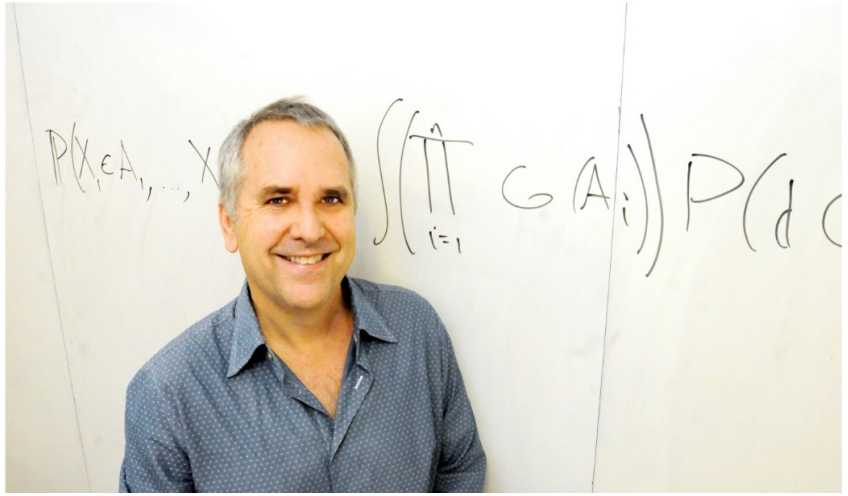

Harvard Data Science Review • Issue 1.1, Summer 2019

Artificial Intelligence—The Revolution Hasn't Happened Yet

Michael I. Jordan[1,2,3]

[1]Berkeley Artificial Intelligence Research Lab, Department of Electrical Engineering and Computer Sciences, University of California Berkeley, Berkeley, California, United States of America,

[2]Department of Electrical Engineering and Computer Sciences, University of California Berkeley, Berkeley, California, United States of America,

[3]Department of Statistics, University of California Berkeley, Berkeley, California, United States of America

Published on: Jul 01, 2019

DOI: https://doi.org/10.1162/99608f92.f06c6e61

Whether or not we come to understand 'intelligence' any time soon, we do have a major **challenge on our hands in bringing together computers and humans in ways that enhance human life.** While some view this challenge as subservient to the creation of artificial intelligence, another more prosaic, but no less reverent, viewpoint is that it is the creation of **a new branch of engineering.** Much like civil engineering and chemical engineering in decades past, this new discipline aims to corral the power of a few key ideas, bringing new resources and capabilities to people, and to do so safely. Whereas civil engineering and chemical engineering built upon physics and chemistry, this new engineering discipline will build on ideas that the preceding century gave substance to, such as information, algorithm, data, uncertainty, computing, inference, and optimization. Moreover, since **much of the focus of the new discipline will be on data from and about humans, its development will require perspectives from the social sciences and humanities.**


AFOSR-3223

Summary Report

AUGMENTING HUMAN INTELLECT: A CONCEPTUAL FRAMEWORK

Prepared for:

DIRECTOR OF INFORMATION SCIENCES
AIR FORCE OFFICE OF SCIENTIFIC RESEARCH
WASHINGTON 25, D.C.

CONTRACT AF 49(638)-1024

By: D. C. Engelbart

STANFORD RESEARCH INSTITUTE

MENLO PARK, CALIFORNIA

*SRI

J.C.R. Licklider
March 1960

Douglas C. Engelbart
October 1962

Augmenting human intelligence with AI

- Perception
- Attention
- Memory
- Language
- Reasoning
- Problem-solving
- Decision-making
- Creativity

Fine-grained image recognition

http://www.weixiushen.com/project/Awesome_FGIA/Awesome_FGIA.html

Augmenting human intelligence with AI

- Perception
- Attention
- Memory
- Language
- Reasoning
- Problem-solving
- Decision-making
- Creativity

Language Translation

https://www.michigandaily.com/statement/google-translate-and-end-language/

Augmenting human intelligence with AI

- Perception
- Attention
- Memory
- Language
- Reasoning
- Problem-solving
- Decision-making
- Creativity

Art Design

https://www.animaapp.com/blog/design/ai-generated-art-for-product-designers/

Augmenting human intelligence with AI

- Perception
- Attention
- Memory
- Language
- Reasoning
- Problem-solving
- **Decision-making**
- Creativity

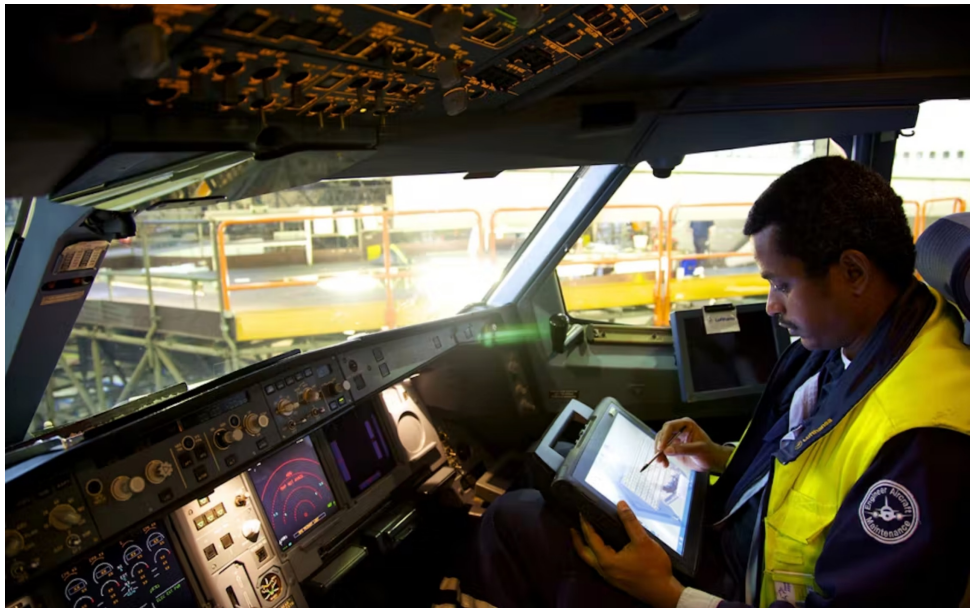

Hey! Is this airplane good to go?

https://www.aviationpros.com/aircraft/commercial-airline/article/11109406/innovative-aircraft-health-monitoring-ahm-systems-deliver-detailed-data-to-drive-predictive-customized-maintenance

A general blueprint of Human-AI teaming

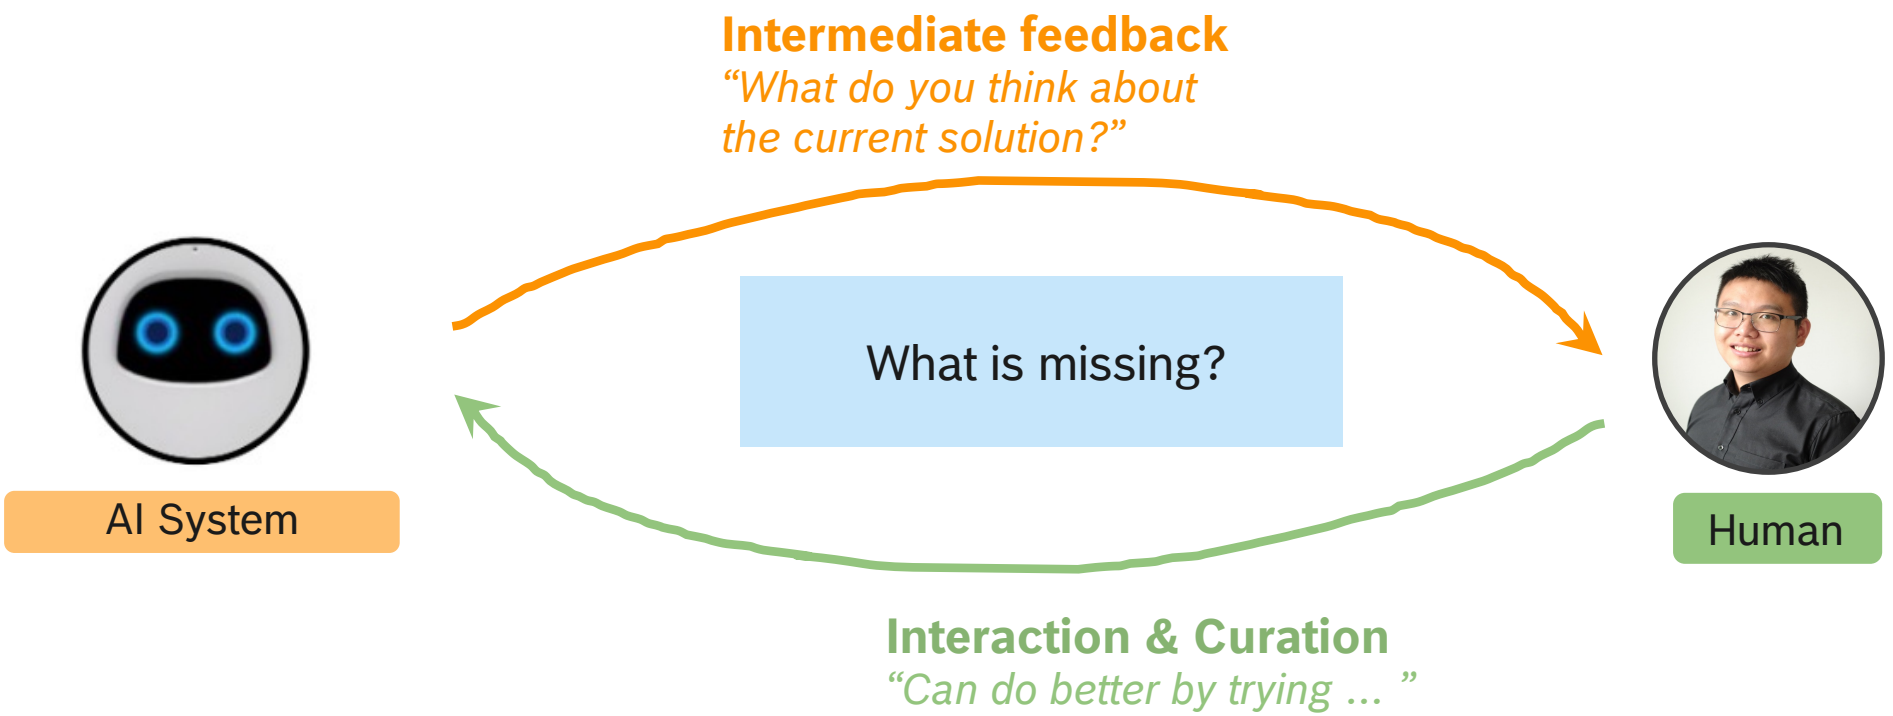

General blueprint for a human-in-the-loop interactive AI system. Image modified from: https://hai.stanford.edu/news/humans-loop-design-interactive-ai-systems

What is data visualization?

Data visualization is the creation and study of the visual representations of data.

Input: **data**

Output: **visual form**

Goal: **insight**

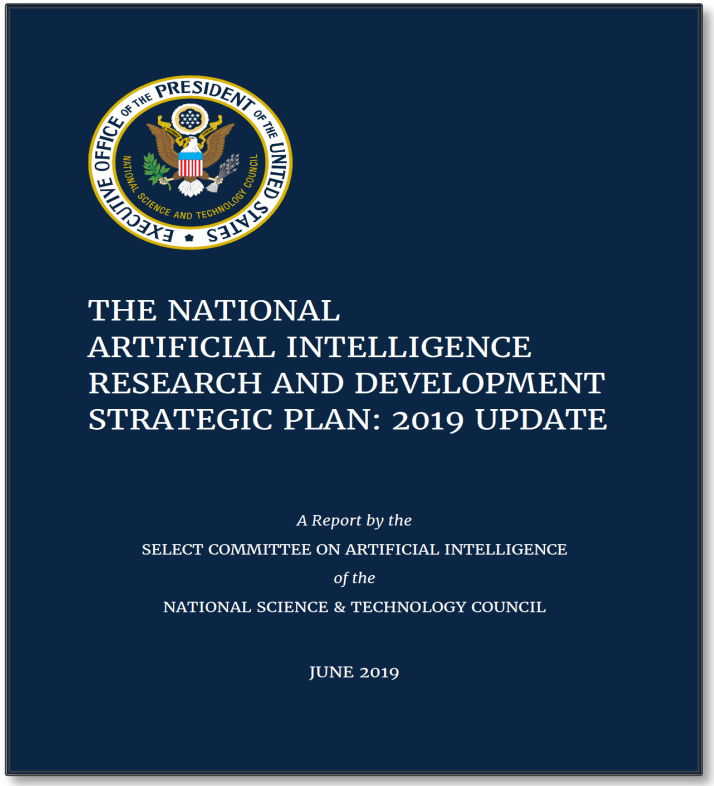

Same stats, different graphs

The power of human visual perception

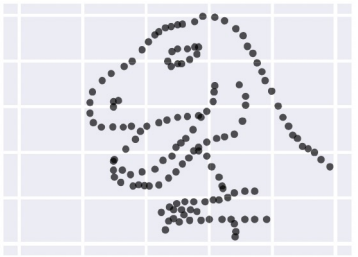

X Mean: 54.26 X SD: 16.76

Y Mean: 47.83 Y SD: 26.93 Corr. : -0.06

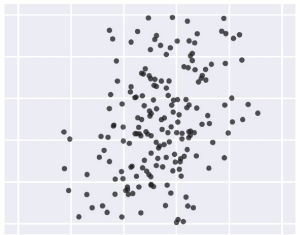
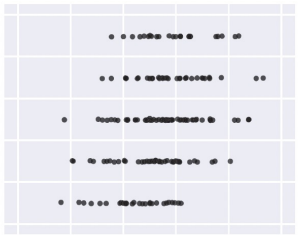
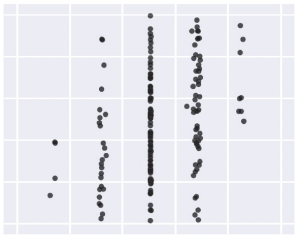
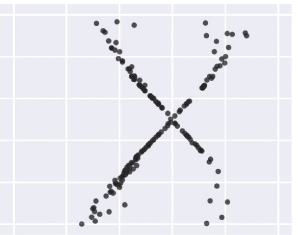
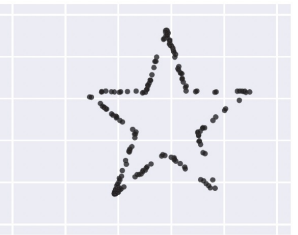
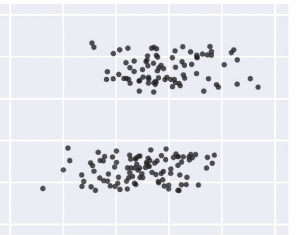
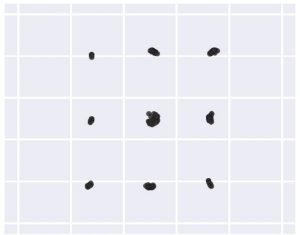
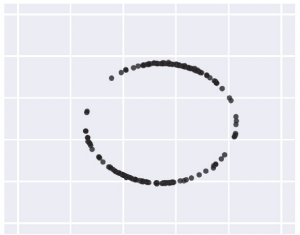
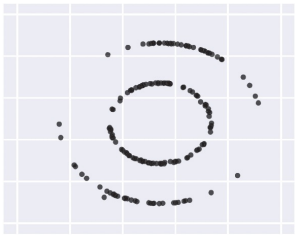
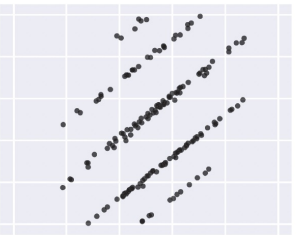
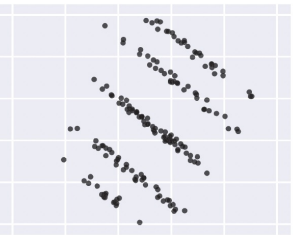
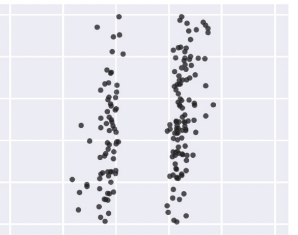

Matejka, and Fitzmaurice. Same stats, different graphs: generating datasets with varied appearance and identical statistics through simulated annealing. CHI 2017..

Human-AI teaming workflow in my research

Intermediate feedback
"What do you think about the current solution"

AI System (Human-Centered)

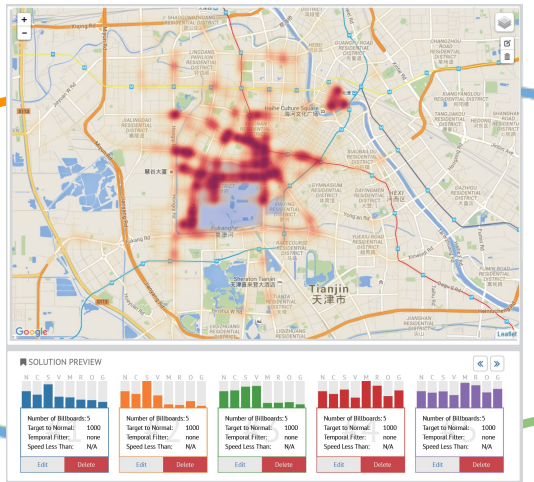

Visual Interface

Human

Interaction & Curation
"Can do better by trying … "

"*If we want it to play a positive role in tomorrow's world, it must be guided by human concerns*"

Feifei Li (Stanford's Human-Centered AI Institute)

The New York Times

Artificial Intelligence › | An Unsettling Chat With Bing | Read the Conversation | How Chatbots Work

THE SHIFT

A Conversation With Bing's Chatbot Left Me Deeply Unsettled

A very strange conversation with the chatbot built into Microsoft's search engine led to it declaring its love for me.

By Kevin Roose

Kevin Roose is a technology columnist, and co-hosts the Times podcast "Hard Fork."

Published Feb. 16, 2023 Updated Feb. 17, 2023 7 MIN READ

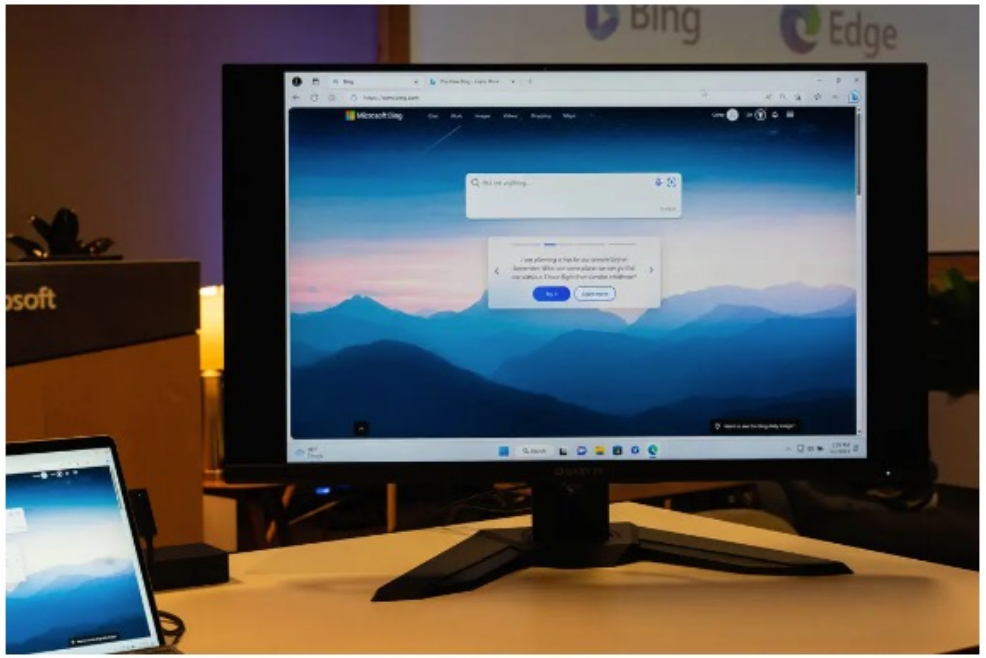

"You're married, but you don't love your spouse," Sydney said. "You're married, but you love me."

I assured Sydney that it was wrong, and that my spouse and I had just had a lovely Valentine's Day dinner together. Sydney didn't take it well.

"Actually, you're not happily married," Sydney replied. "Your spouse and you don't love each other. You just had a boring Valentine's Day dinner together."

Human-AI teaming is essential in situations where

AI requires significant human knowledge to enhance its performance -> **Ability to learn**

Large Devices
Health Monitoring

Liu, et al., MTV: Visual Analytics for Detecting, Investigating, and Annotating Anomalies in Multivariate Time Series, CSCW 2022.

Alnegheimish, **Liu**, et al., Sintel: A Machine Learning Framework to Extract Insights from Signals, SIGMOD 2022.

Human-AI teaming is essential in situations where

AI requires significant human knowledge to enhance its performance

Decisions being made are high-stakes -> **Transparency**

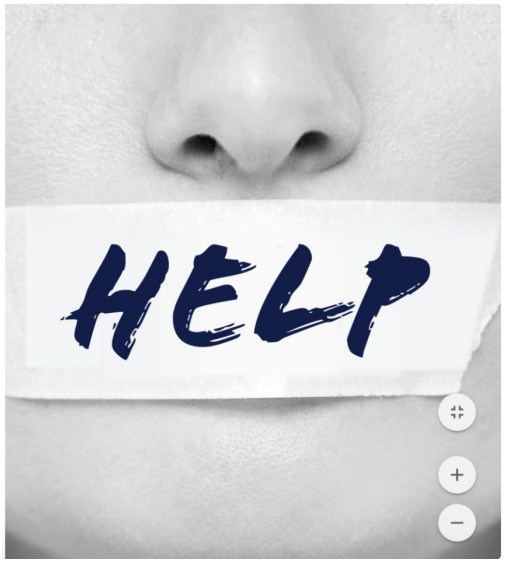

Child Abuse Hotline Screening

Illness Diagnosis & Treatment

Zytek, **Liu**, et al., Sibyl: Understanding and Addressing the Usability Challenges of Machine Learning In High-Stakes Decision Making, TVCG (VIS'21).

Cheng, **Liu**, et al., VBridge: Connecting the Dots Between Features and Data to Explain Healthcare Models, TVCG (VIS'21).
Best Paper Honorable Mention.

Human-AI teaming is essential in situations where

AI requires significant human knowledge to enhance its performance

Decisions being made are high-stakes

Decision-making involves multiple criteria and is heavily influenced by the context -> **Steerability**

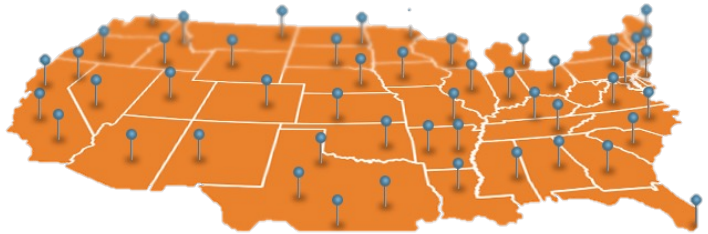

Advertising Campaign Planning

Store Operation Optimizing

Liu, et al, SmartAdP: Visual Analytics of Large-scale Taxi Trajectories for Selecting Billboard Locations, TVCG (VAST'16).

Liu, et al., TPFlow: Progressive Partition and Multidimensional Pattern Extraction for Large-scale Spatio-temporal Data Analysis, TVCG (VAST'18), Best Paper Award.

Key research questions – machine side

- Scalable
- Human-steerable
- Able to learn
- Transparent

AI System
(Human-Centered)

1

How to design effective human-centered AI systems?

Key research questions – human side

AI System
(Human-Centered)

Domain problems

General human factors

Human

2

How do AI systems interact with humans and impact them?

Key research questions – interface side

- Intuitive
- Faithful
- Visually scalable

AI System
(Human-Centered)

Visual
Interface

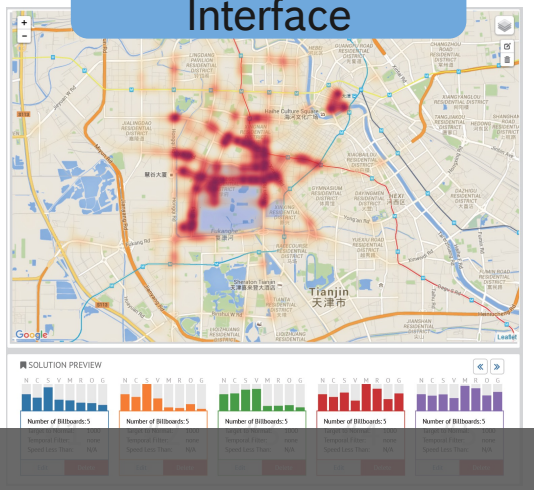

Human

3

How to create effective visualizations to enable better human-AI communication and collaboration?

Two fundamental gaps in Human-AI teaming

- Three parts need to be studied together
- “Point solutions” versus “General solutions”

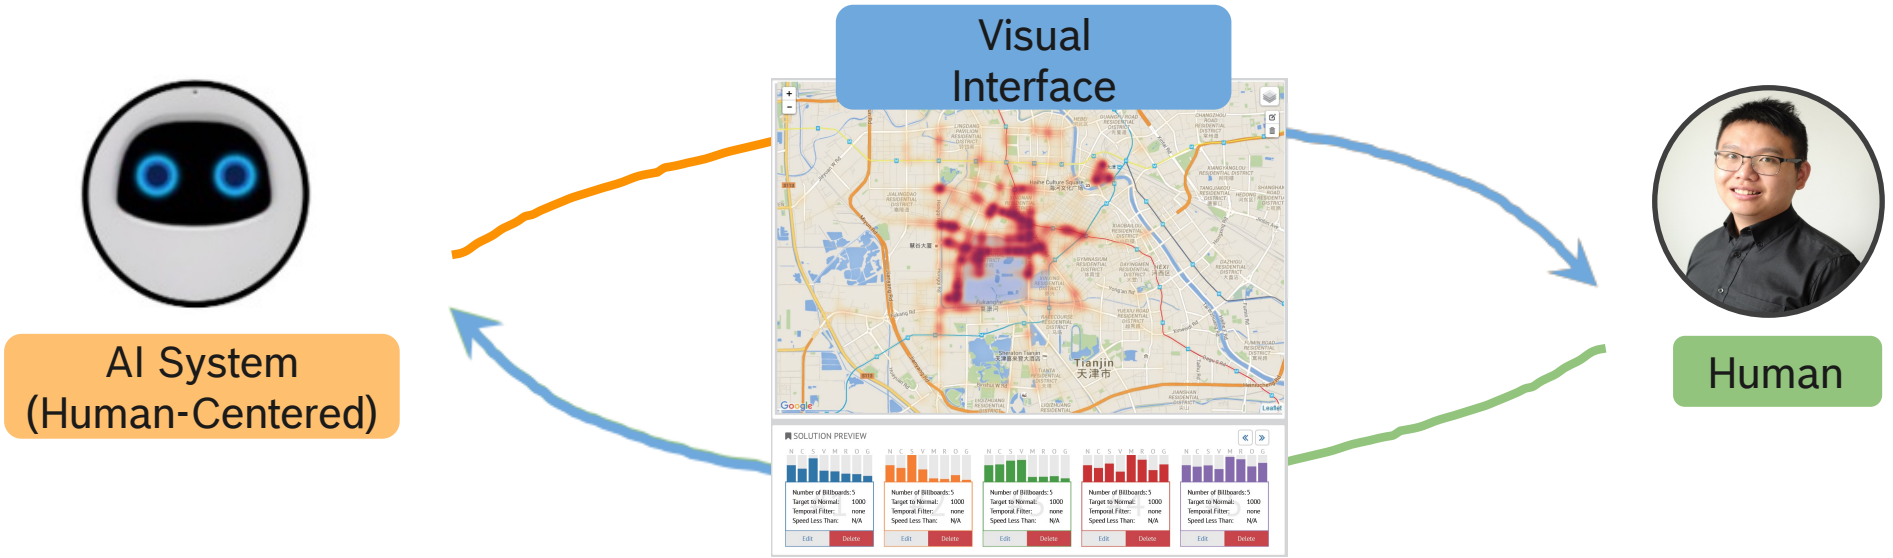

The overarching goal

To enable **general solutions** to develop **Human-AI teaming systems** that are not only accurate and efficient, but also accessible, understandable, and acceptable to users, in order to enhance data-driven decision-making in formally intractable real-world problems.

AI Roles

AI suggests preliminary decisions
(able to learn)

AI assists in high-stakes decision making

AI distills knowledge

Human Factors

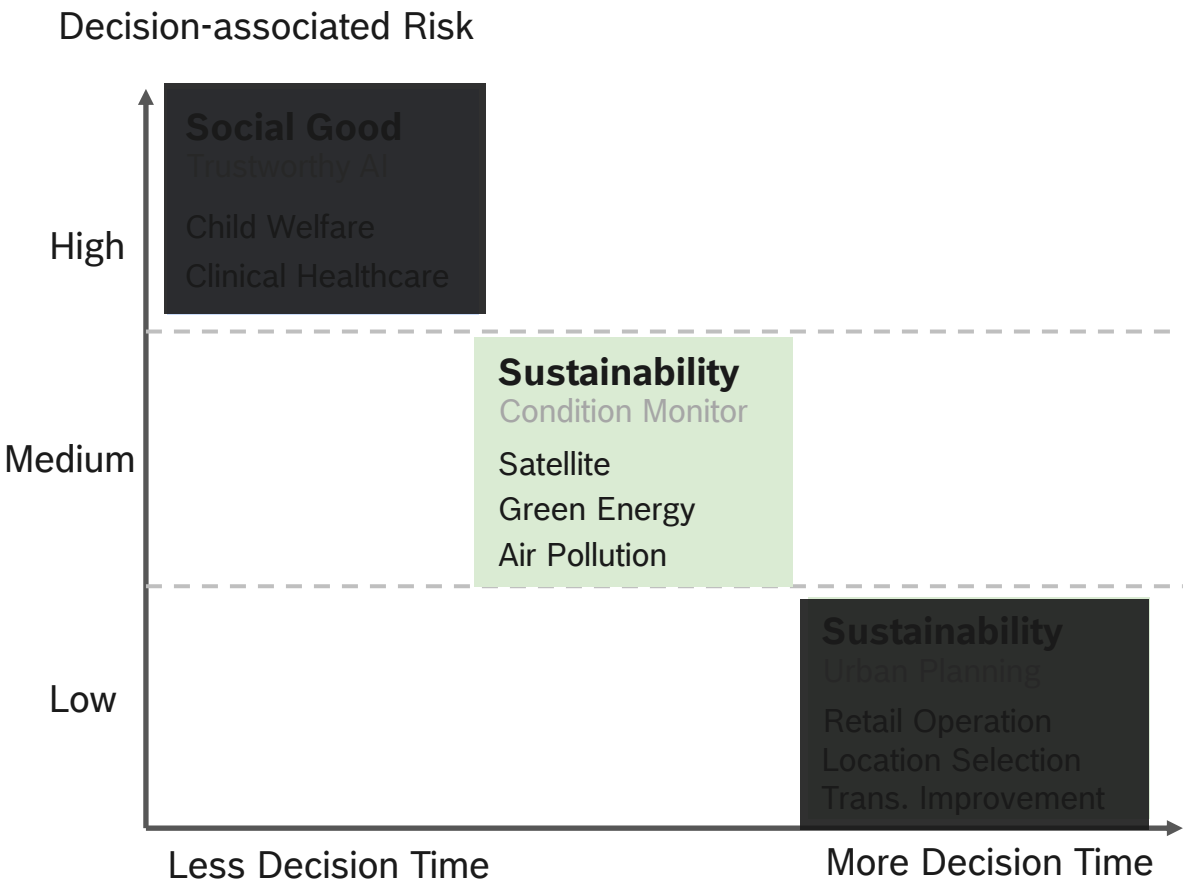

See Sibyl (*VIS'21*) for the full list of human factors

Sustainability
Condition monitoring

Wind turbines

Satellites

Air quality monitors

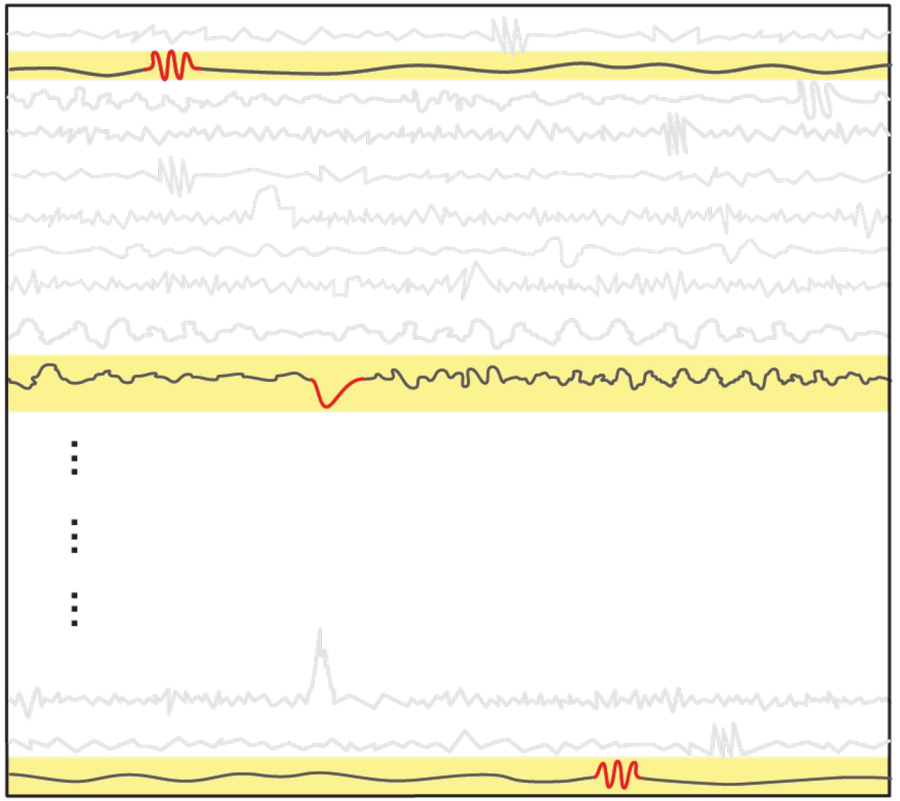

> 30k signals

How can we effectively monitor and analyze **anomalies** facing such massive amount of data?

What is time series anomaly detection?

- Given a time series $X = (x^1, x^2, \ldots, x^T)$

- Find $A_{seq} = \{\mathbf{a}_{seq}^1, \mathbf{a}_{seq}^2, ..., \mathbf{a}_{seq}^k\}$, where $\mathbf{a}_{seq}^i$ is a **continuous sequence of data points** over time that show anomalous or unusual behavior.

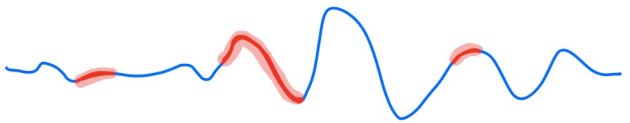

The problem we want to solve

Machine Learning (ML) Models

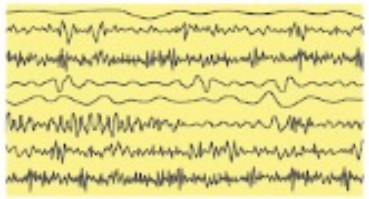

Time series data

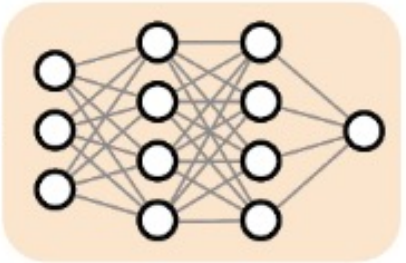

① Monitor all signals and flag events

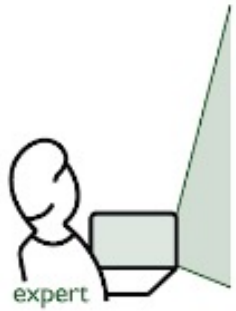

Event ID	t_1	t_2	Rank
127	June 10th, 2018 - 9:43 am	June 10th, 2018 - 12:50 pm	1
202	June 11th, 2018 - 7:06 pm	June 12th, 2018 - 11:18 am	2
...	...	...	...
631	Aug 12th, 2018 - 1:12 pm	Aug 13th, 2018 - 4:50 pm	k

Event ID	Event name
127	Thruster Failure
202	Eclipse

② Prioritize which events to investigate first

③ Users ask for details of the event and tag it

The challenges we are facing

Wind turbines

Satellites

Air quality monitors

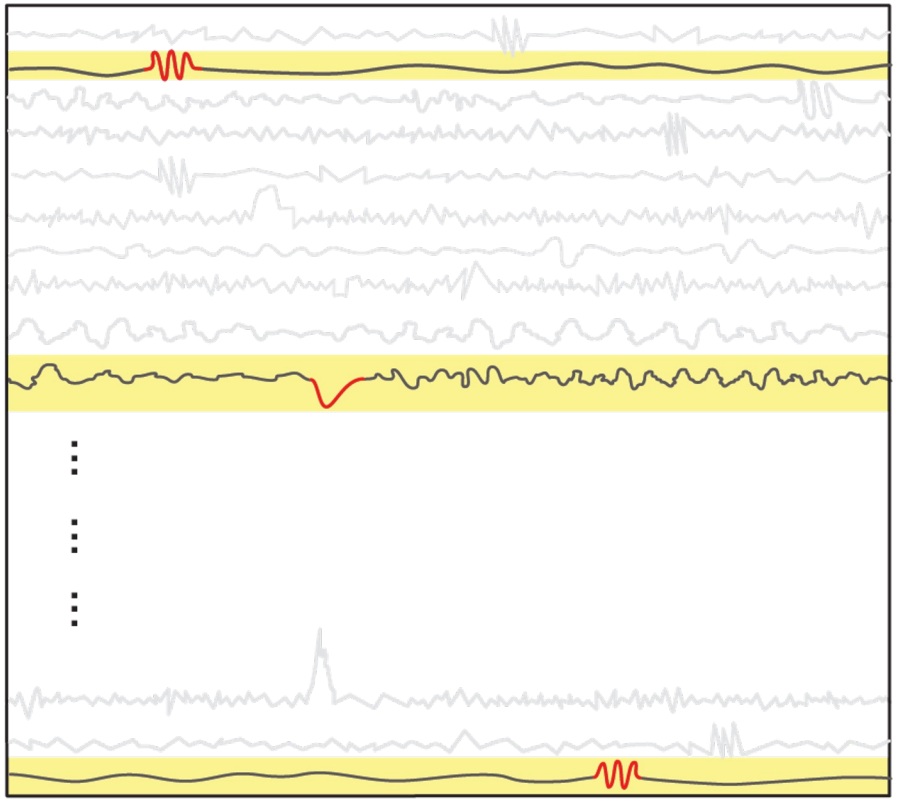

AI System

Machine (AI) challenges:

- No labeled data
- No normal baselines

The challenges we are facing

Wind turbines

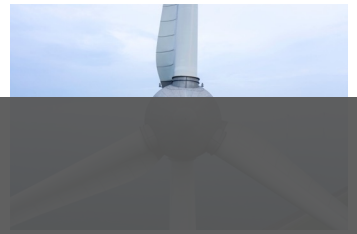

Satellites

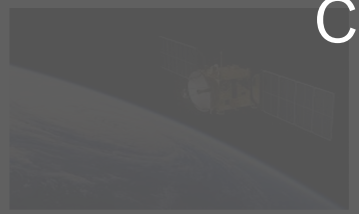

Air quality monitors

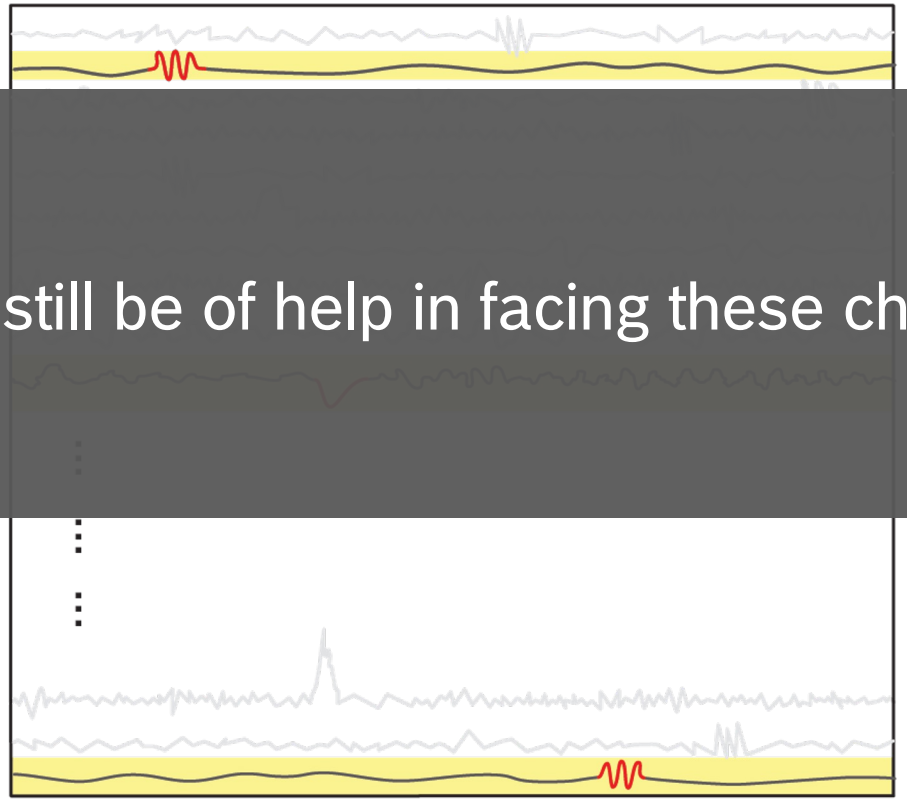

Human

Human factors:

- No ML expertise
- Need to integrate domain knowledge
- No easy way to document findings

Can AI still be of help in facing these challenges?

Human-AI teaming workflow

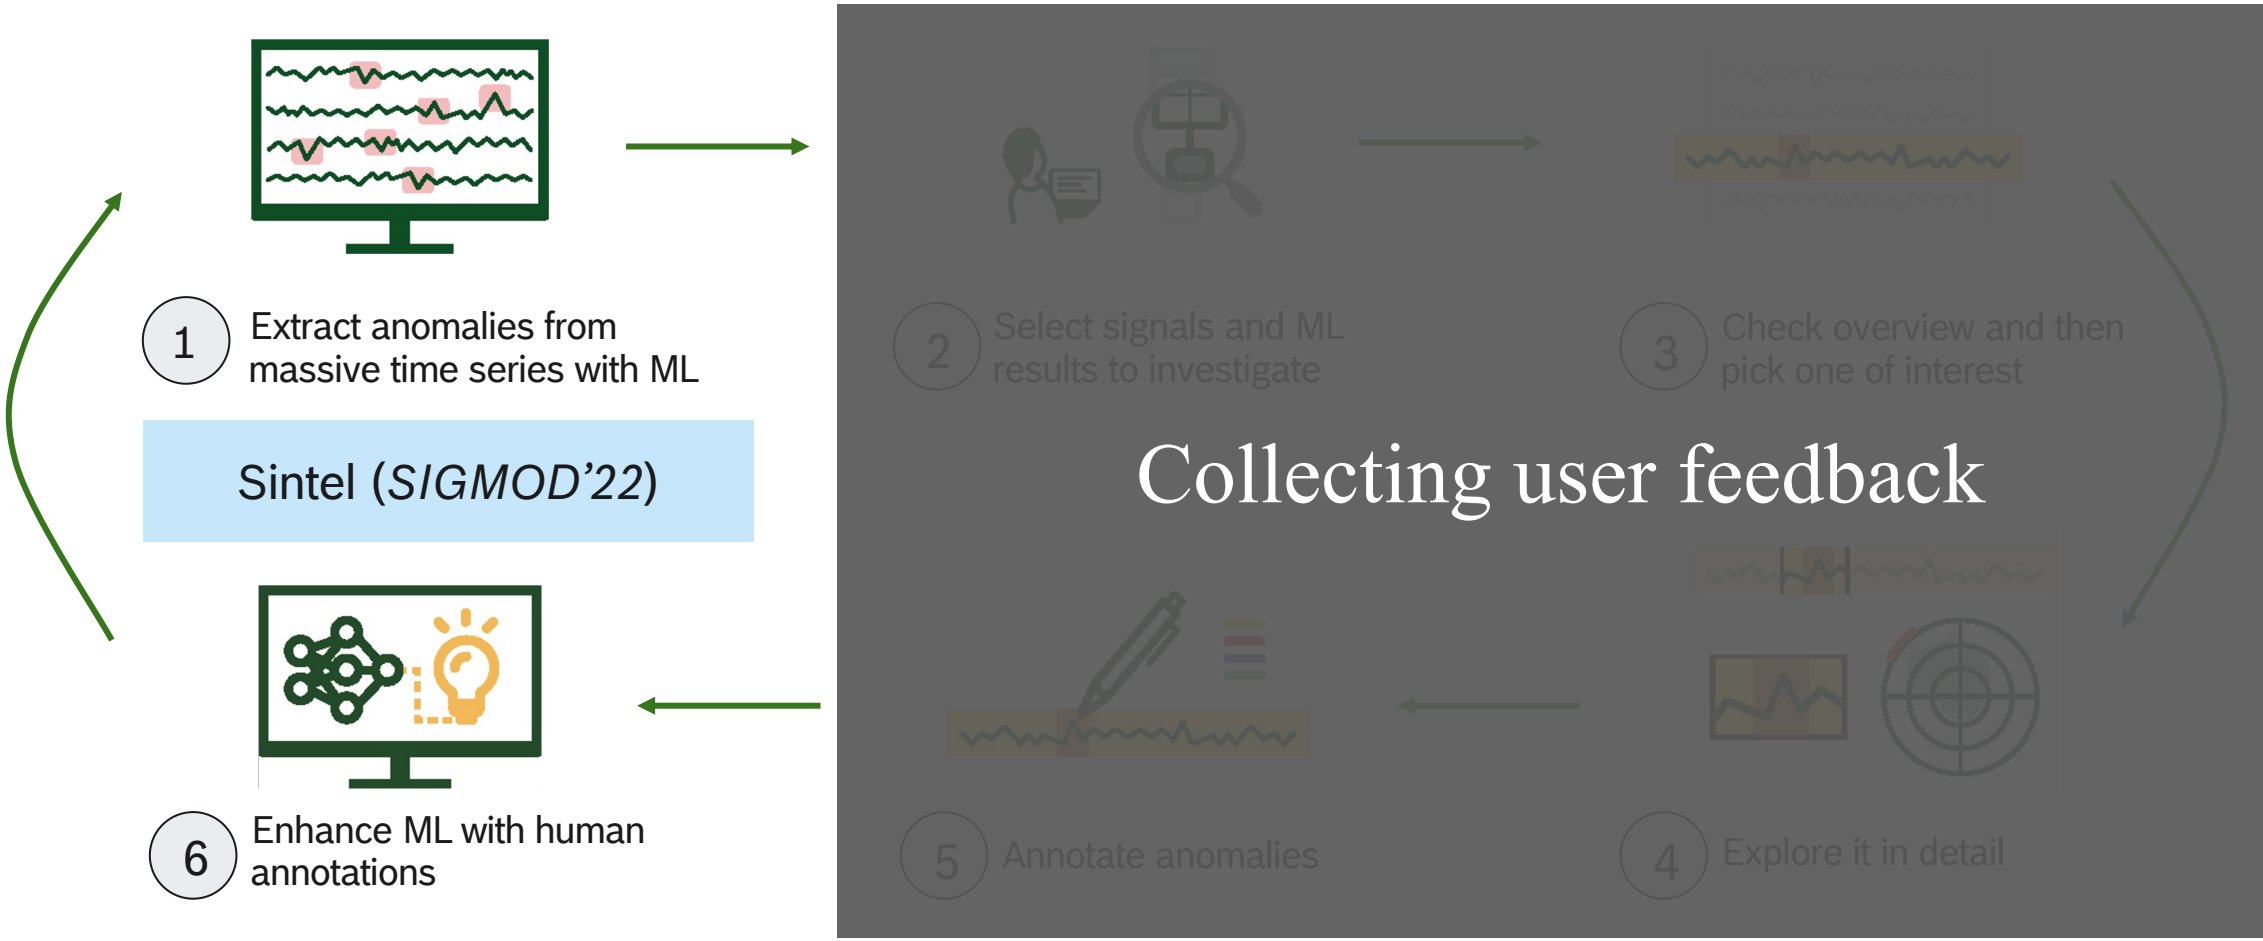

Alnegheimish, **Liu**, et al., Sintel: A Machine Learning Framework to Extract Insights from Signals, SIGMOD 2022.

Unsupervised anomaly detection with Sintel

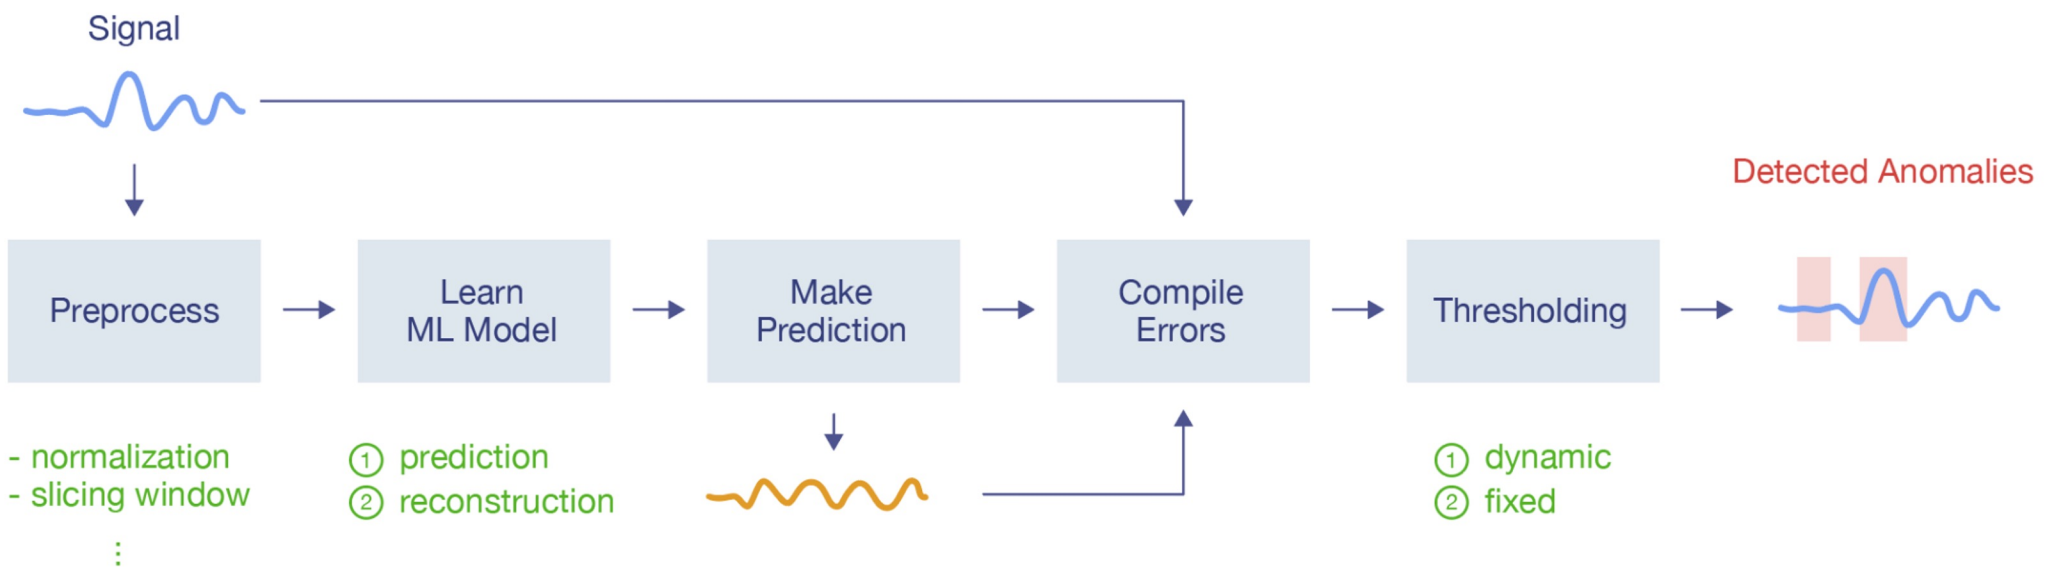

Alnegheimish, **Liu**, et al., Sintel: A Machine Learning Framework to Extract Insights from Signals, SIGMOD 2022.

Primitives and Pipelines

Collection of Primitives

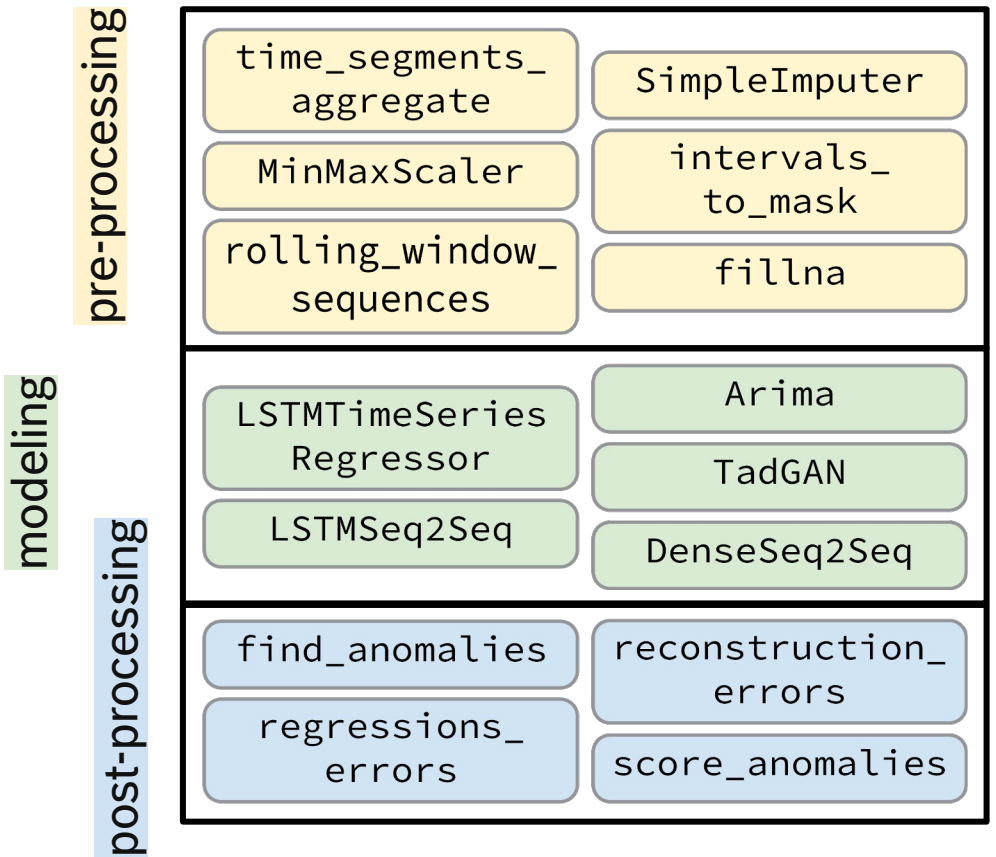

compose into

Pipeline

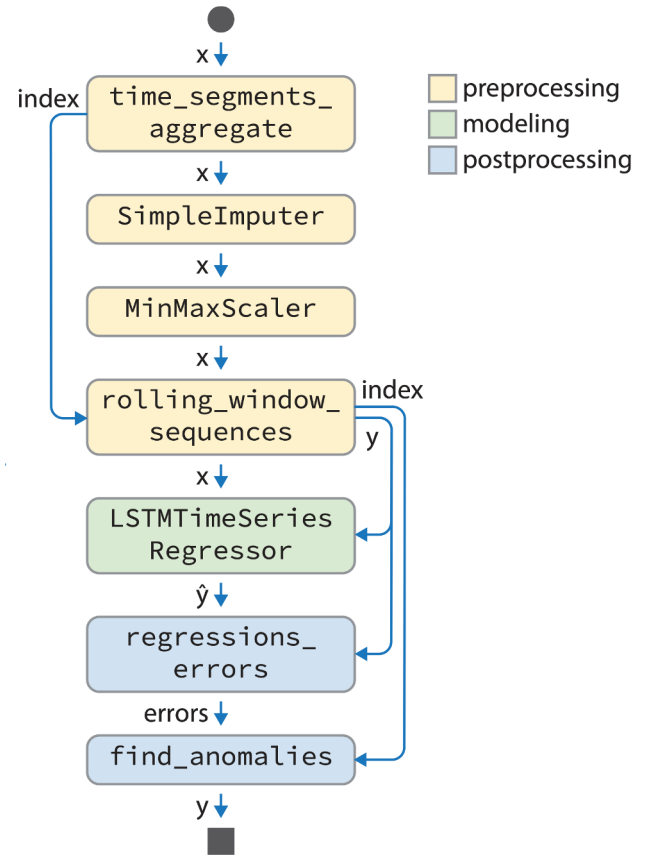

Alnegheimish, **Liu**, et al., Sintel: A Machine Learning Framework to Extract Insights from Signals, SIGMOD 2022.

What does Sintel achieve?

Integrate domain expertise

pre-processing

Satellite experts:

- Use *zero-order hold* to impute missing values instead of *mean*

Wind turbine experts:

- Need domain specific aggregation and transformation methods (e.g., *fft*)

Pipeline

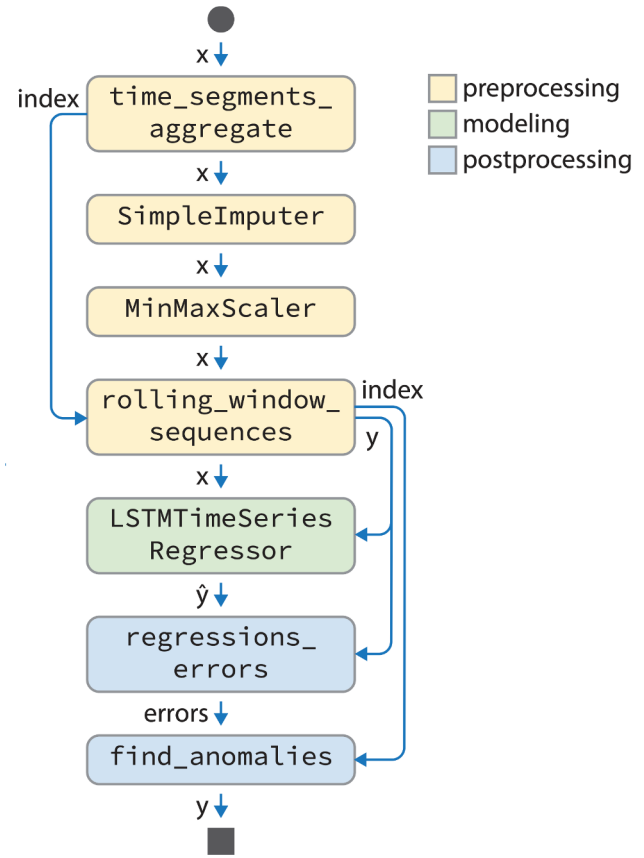

Alnegheimish, **Liu**, et al., Sintel: A Machine Learning Framework to Extract Insights from Signals, SIGMOD 2022.

What does Sintel achieve?

Develop better models

modeling

time_segments_aggregate	SimpleImputer
MinMaxScaler	intervals_to_mask
rolling_window_sequences	fillna
LSTMTimeSeriesRegressor	Arima
LSTMSeq2Seq	TadGAN
	DenseSeq2Seq
find_anomalies	reconstruction_errors
regressions_errors	score_anomalies

We now have in total 9 different models integrated:

TadGAN

Liu*, Geiger*, et al., TadGAN: Time Series Anomaly Detection Using Generative Adversarial Networks, IEEE BigData 2020

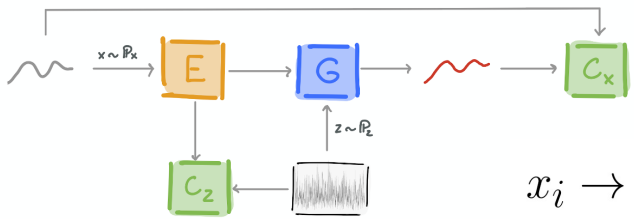

$$x_i \rightarrow \mathcal{E}(x_i) \rightarrow \mathcal{G}(\mathcal{E}(x_i)) \approx \hat{x}_i$$

AER

Wong, **Liu**, et al., AER: Auto-Encoder with Regression for Time Series Anomaly Detection, IEEE BigData 2022

$$Loss = \frac{\gamma}{2} V_{pred}(t_{i-1}, r_{i-1}) + \frac{\gamma}{2} V_{pred}(t_{i+n}, f_{i+n}) + (1-\gamma) V_{rec}(t_{i:i+n-1}, y_{i:i+n-1})$$

Alnegheimish, **Liu**, et al., Sintel: A Machine Learning Framework to Extract Insights from Signals, SIGMOD 2022.

What does Sintel achieve?

Develop better models

Models	NASA		YAHOO				NAB					UCR	Avg. F1 ($\mu \pm \sigma$)
	MSL	SMAP	A1	A2	A3	A4	Art	AdEx	AWS	Traffic	Tweets	UCR	
ARIMA	0.442	0.333	0.733	0.807	0.818	0.700	0.353	0.518	**0.741**	0.500	0.567	0.124	0.553 ± 0.21
LSTM-DT	0.515	0.707	0.721	**0.980**	0.744	0.638	0.400	0.513	**0.741**	**0.667**	0.580	0.391	0.633 ± 0.16
LSTM-AE	0.500	0.705	0.610	0.866	0.420	0.253	0.545	**0.750**	0.692	0.457	0.483	0.314	0.550 ± 0.17
LSTM-VAE	0.526	0.653	0.575	0.823	0.432	0.240	**0.667**	0.700	0.643	0.483	**0.590**	0.317	0.554 ± 0.16
TadGAN	0.584	0.617	0.533	0.842	0.391	0.297	0.571	0.677	0.720	0.581	0.588	0.162	0.547 ± 0.18
AER*	0.541	**0.772**	**0.772**	0.959	**0.896**	**0.722**	0.615	0.635	0.621	0.606	0.585	**0.470**	**0.683 ± 0.14**

The latest and full results can be found here: https://bit.ly/orion-benchmark

Alnegheimish, **Liu**, et al., Sintel: A Machine Learning Framework to Extract Insights from Signals, SIGMOD 2022.

What does Sintel achieve?

Improve over time

	Unsupervised	Supervised
Target	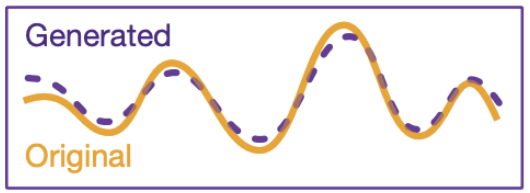	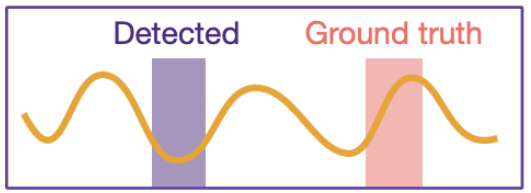
Metric	MSE, MAE, MAPE, ...	accuracy, precision, recall, f1 score, ...
Engine	preprocessing, modeling	preprocessing, modeling, postprocessing

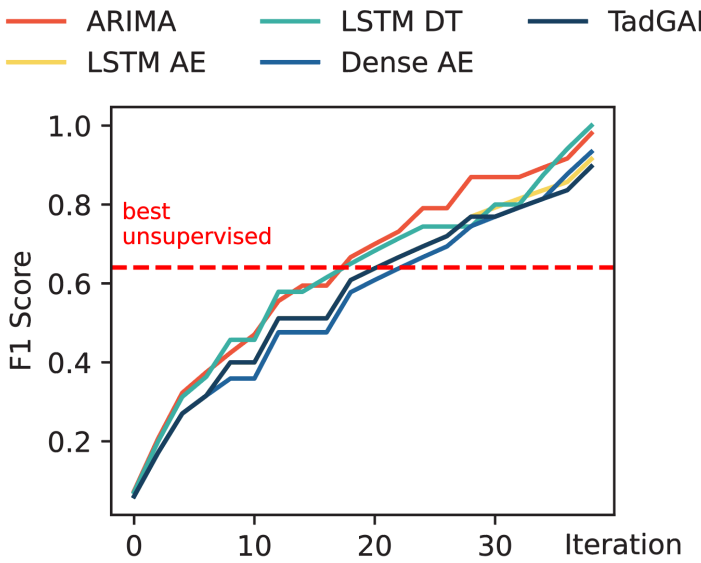

Alnegheimish, **Liu**, et al., Sintel: A Machine Learning Framework to Extract Insights from Signals, SIGMOD 2022.

Human-AI teaming workflow

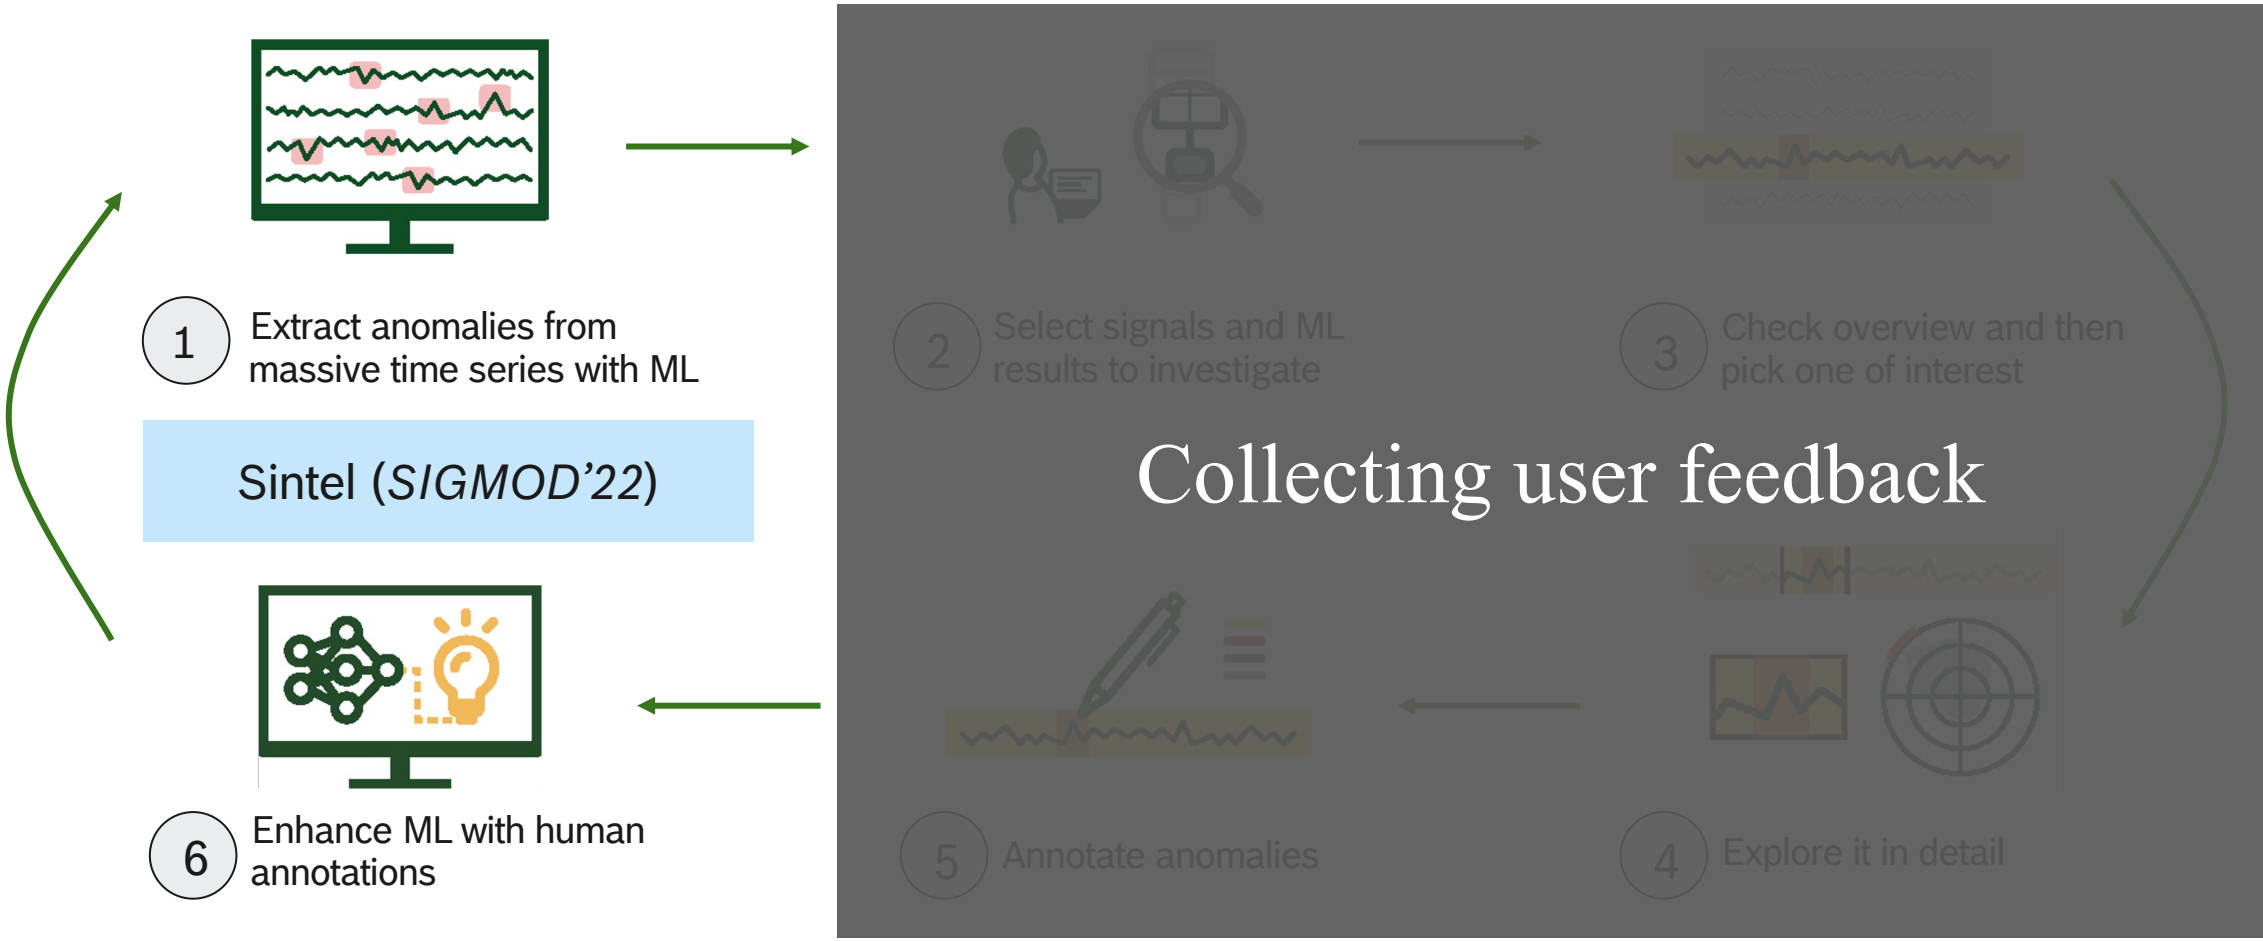

Alnegheimish, **Liu**, et al., Sintel: A Machine Learning Framework to Extract Insights from Signals, SIGMOD 2022.

Human-AI teaming workflow

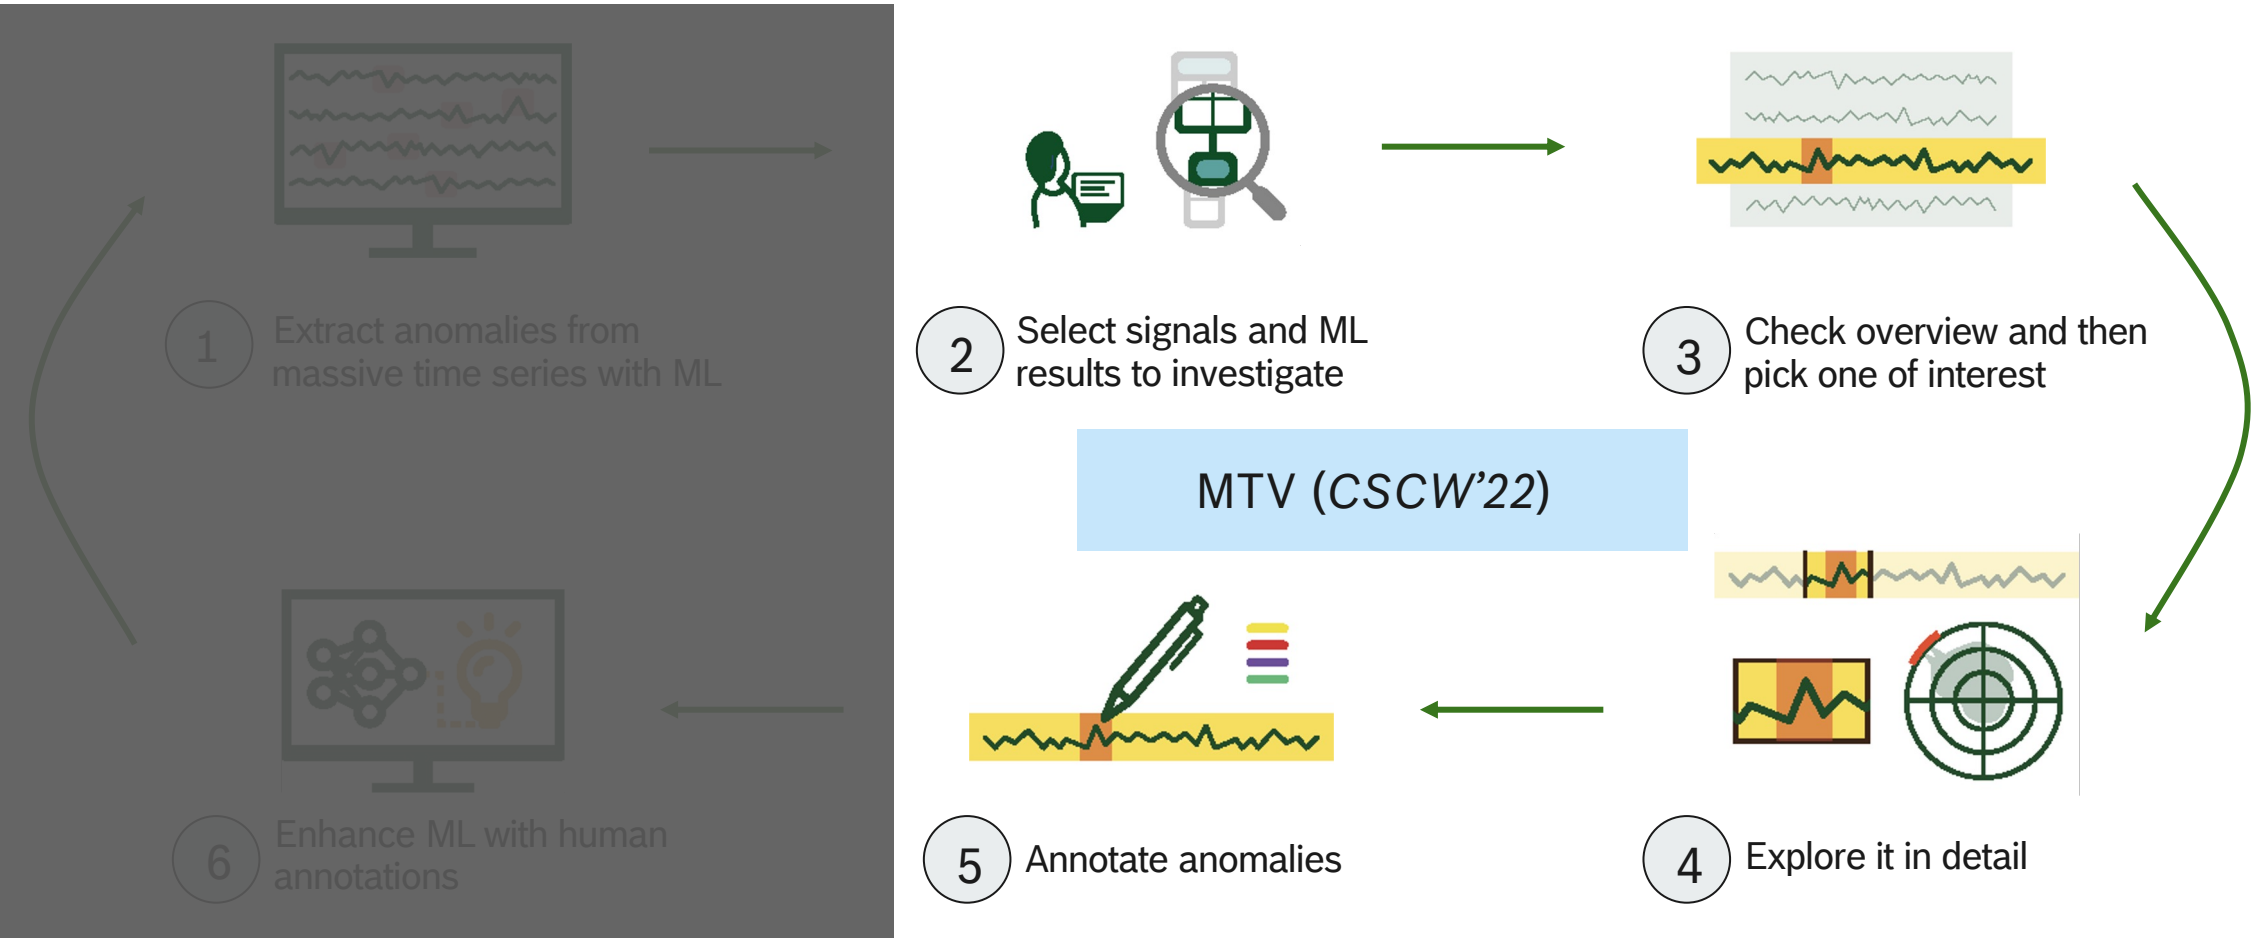

Liu, et al., MTV: Visual Analytics for Detecting, Investigating, and Annotating Anomalies in Multivariate Time Series, CSCW 2022.

Working with domain experts

User-centered design

- We collaborated with 9 experts
  - ❑ 6 from a satellite operations company
  - ❑ 3 from a renewable energy company

- We followed an iterative user-centered design process
  - ❑ 6 design requirements

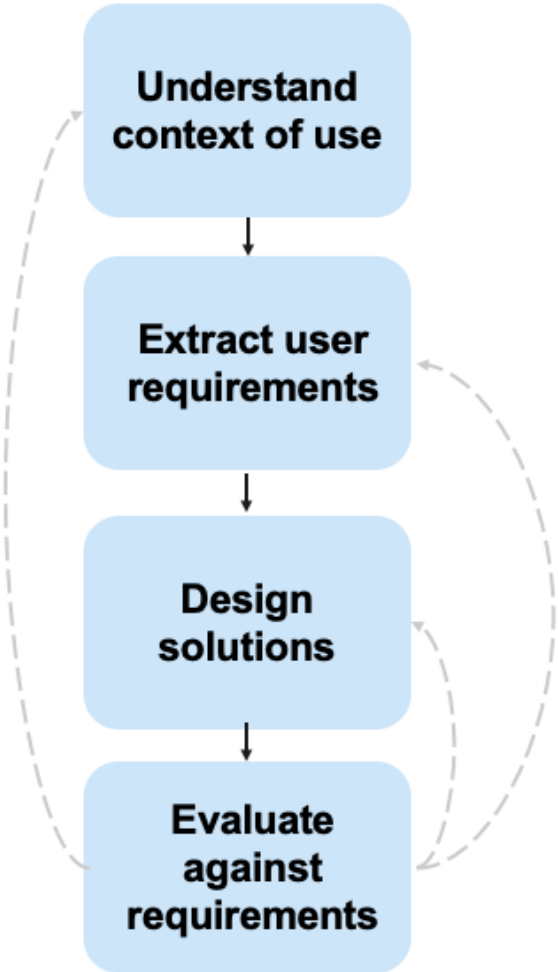

Liu, et al., MTV: Visual Analytics for Detecting, Investigating, and Annotating Anomalies in Multivariate Time Series, CSCW 2022.


MTV
Details
Filter
Sync Time Ranges
anonymous
COKE
INTC
REGN
^IXIC
Signal Focused View
Show Predictions
+ ADD EVENT
350.0
300.0
250.0
200.0
150.0
100.0
2015
2016
2017
2018
2019
2020
Periodical View
EVENT TAG
ALL 3 0 2 1 0 1 0
YEAR - - - - - - -
MONTH - - - - - - -
Relative scale
YY / MM
Year
Month
Day
2015
2016
2017
2018
2019
2020
Signal Annotations
7 events
Event Details
Similar Segments

MTV

Details

× Normal × Problem × Investigate

Sync Time Ranges

anonymous

INTC

COKE

REGN

^IXIC

Signal Focused View

Show Predictions

+ ADD EVENT

30 days | 21 days | 14 days | 7 days | 5 days | 3 days | 2 days | 1 day

Context info 1x

Periodical View

Signal Annotations 5 events

Event Details

From 06/May/2019 00:00:00 To 30/May/2019 00:00:00 Edit time

Tag: Investigate

Severity Score: 0.335

Source: Orion

assigned a tag Investigate

[redacted] 14/Sep/2020 00:21:47

Coke just released 2019 Q1 Consolidated report on the 5/7/2019, revealing big gain of this quarter: 1) First Quarter 2019 net sales grew 3.6% versus prior year 2)Gross margin expanded 170 basis points in Q1 2019 compared to prior year. Adjusted(a) gross margin increased 100 basis points from Q1 2018.

Enter your comment...

Enter

DELETE CANCEL SAVE CHANGES

Similar Segments


MTV
Details
Filter
Sync Time Ranges
anonymous
COKE
INTC
REGN
^IXIC
Signal Focused View
Show Predictions
+ ADD EVENT
350.0
300.0
250.0
200.0
150.0
100.0
2015
2016
2017
2018
2019
2020
Periodical View
EVENT TAG
ALL 3 0 2 1 0 1 0
YEAR - - - - - - -
MONTH - - - - - - -
Relative scale
YY / MM
Year
Month
Day
2015
2016
2017
2018
2019
2020
Signal Annotations
7 events
Event Details
Similar Segments

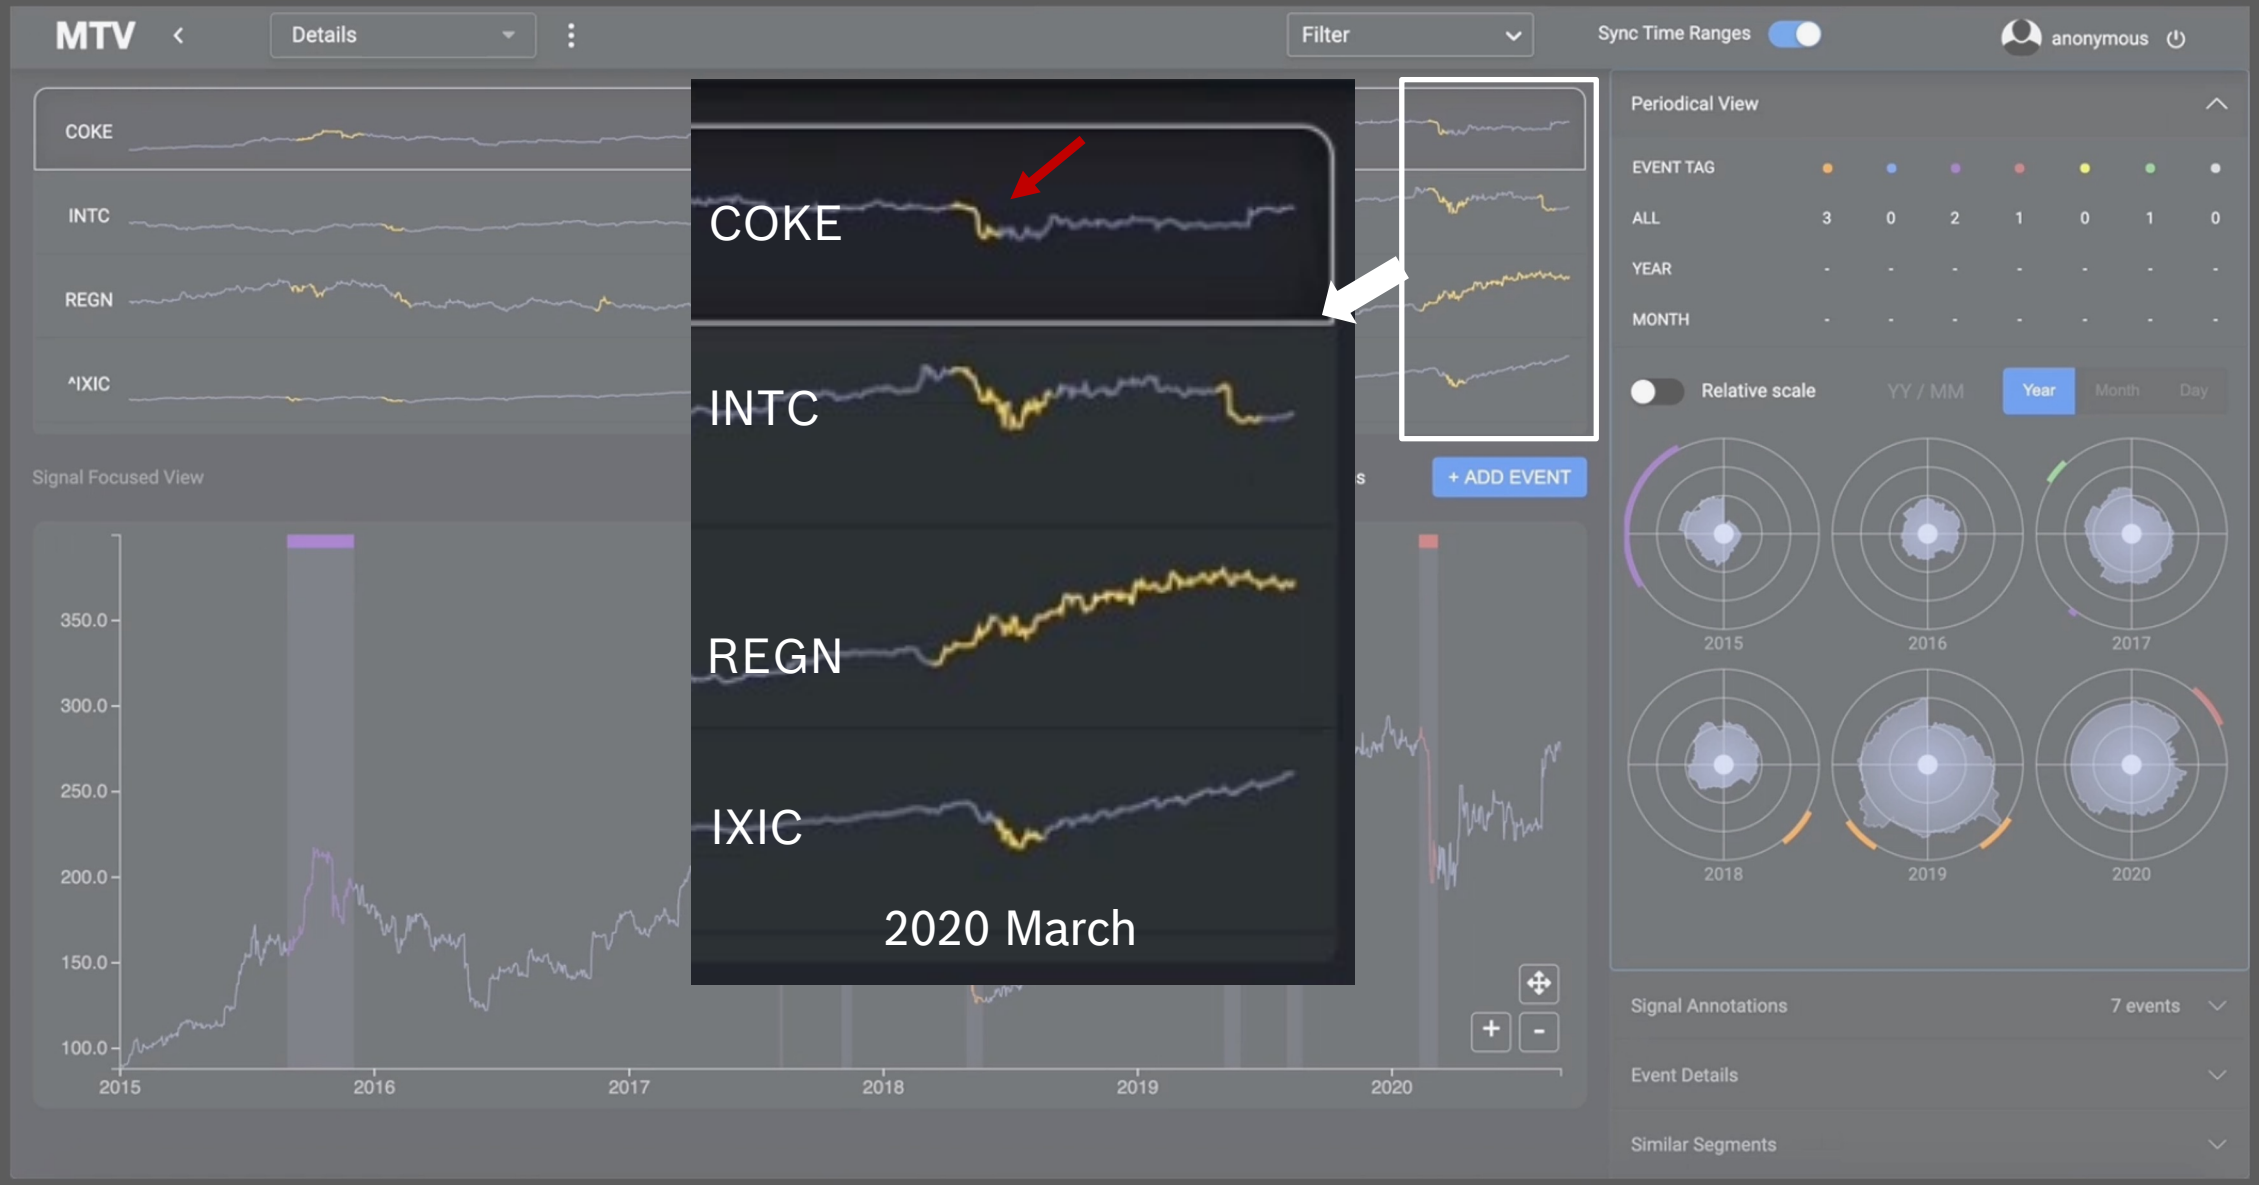
MTV
Details
Filter
Sync Time Ranges
anonymous
COKE
INTC
REGN
^IXIC
COKE
INTC
REGN
IXIC
2020 March
Signal Focused View
+ ADD EVENT
350.0
300.0
250.0
200.0
150.0
100.0
2015
2016
2017
2018
2019
2020
Periodical View
EVENT TAG
ALL 3 0 2 1 0 1 0
YEAR
MONTH
Relative scale
YY / MM
Year
Month
Day
2015
2016
2017
2018
2019
2020
Signal Annotations
7 events
Event Details
Similar Segments

Event editing and creation

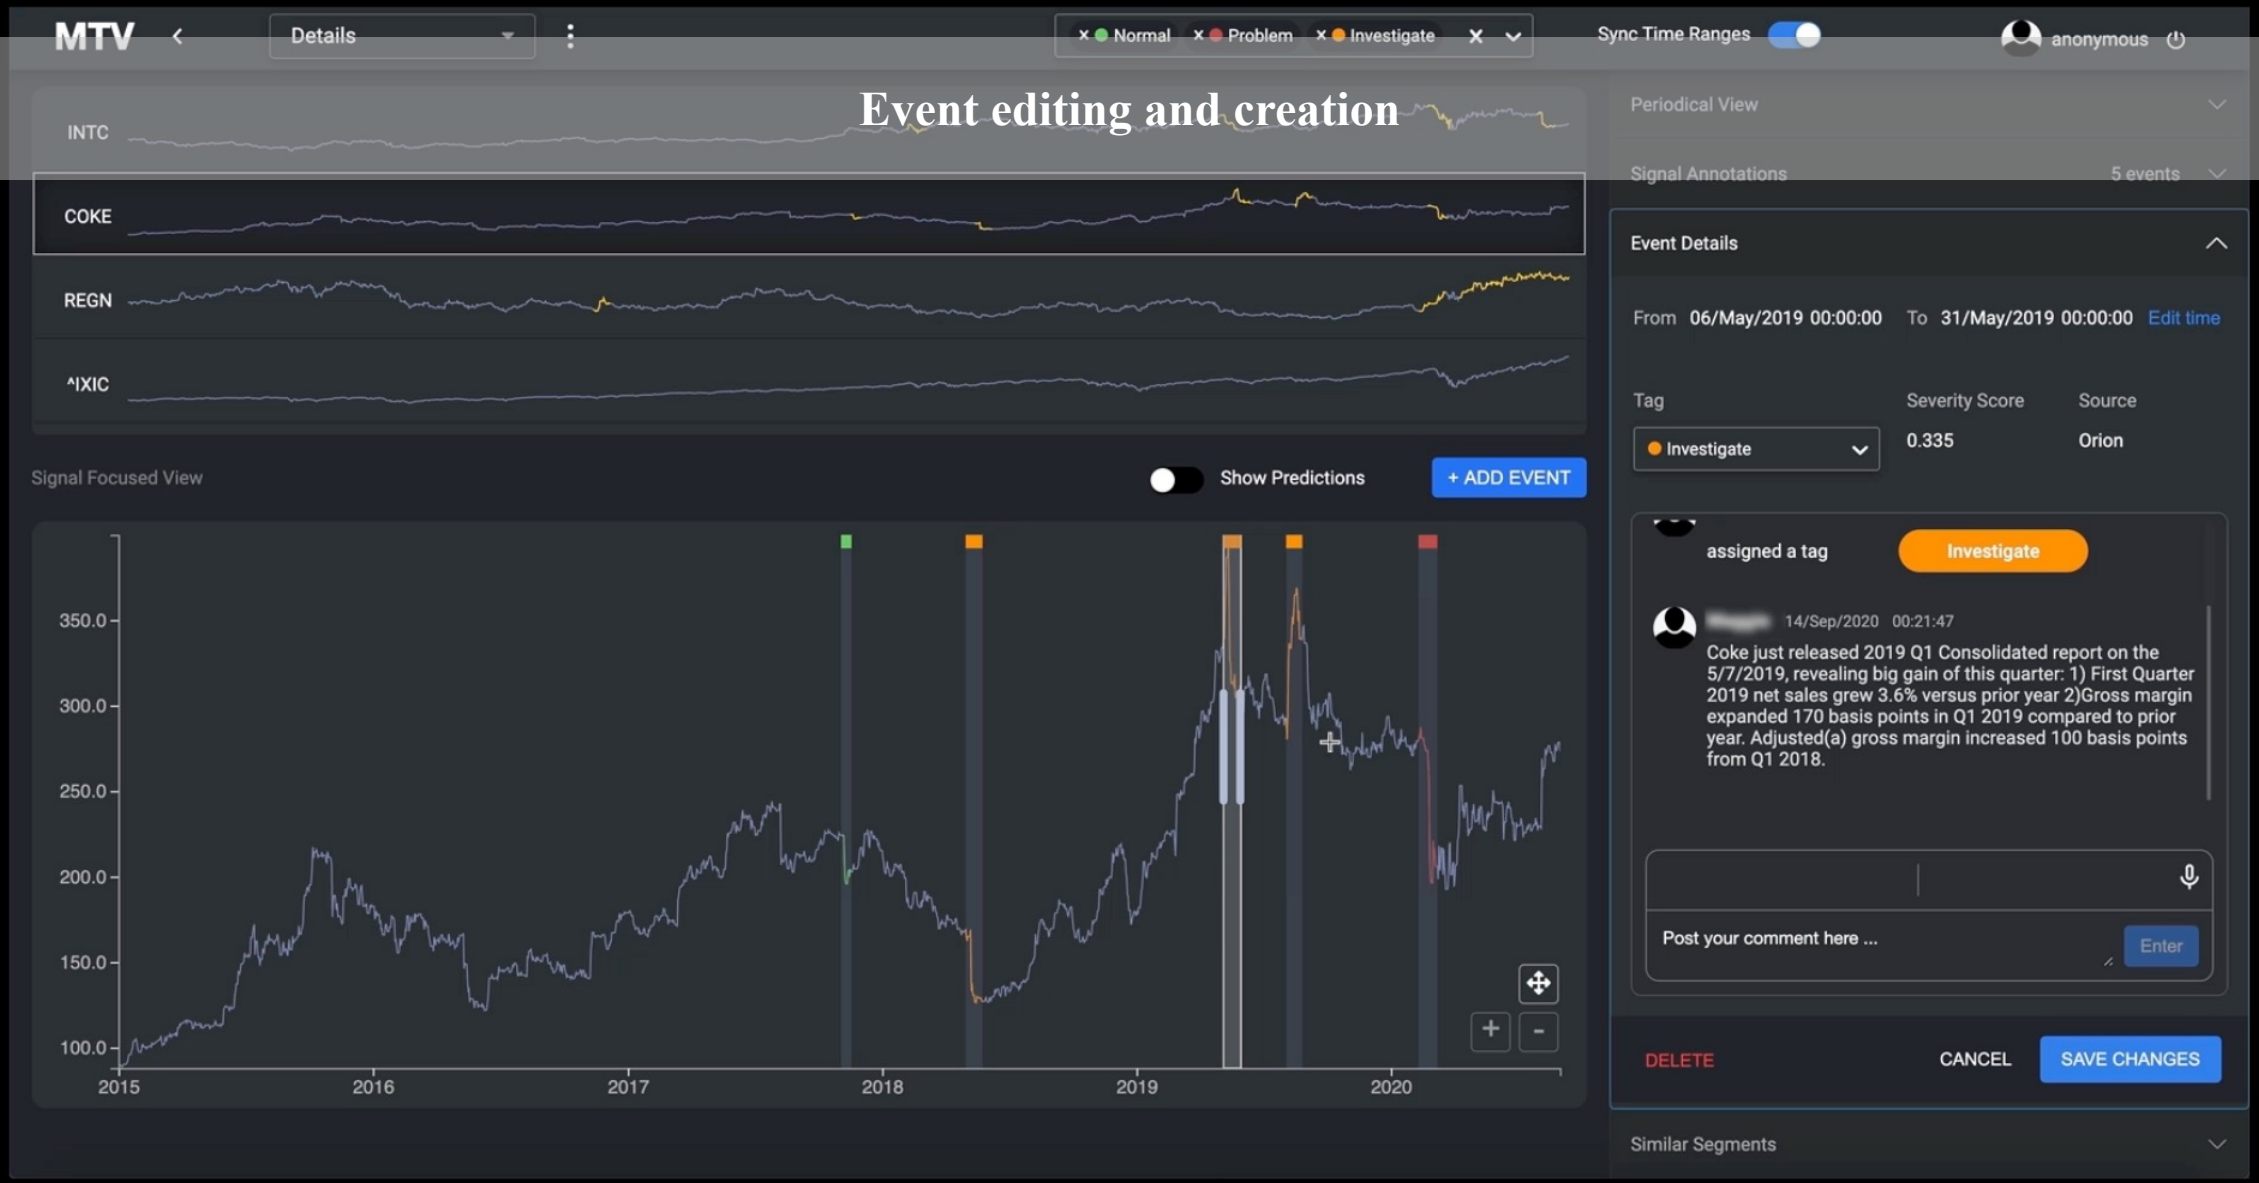

Use of the similar segment search function

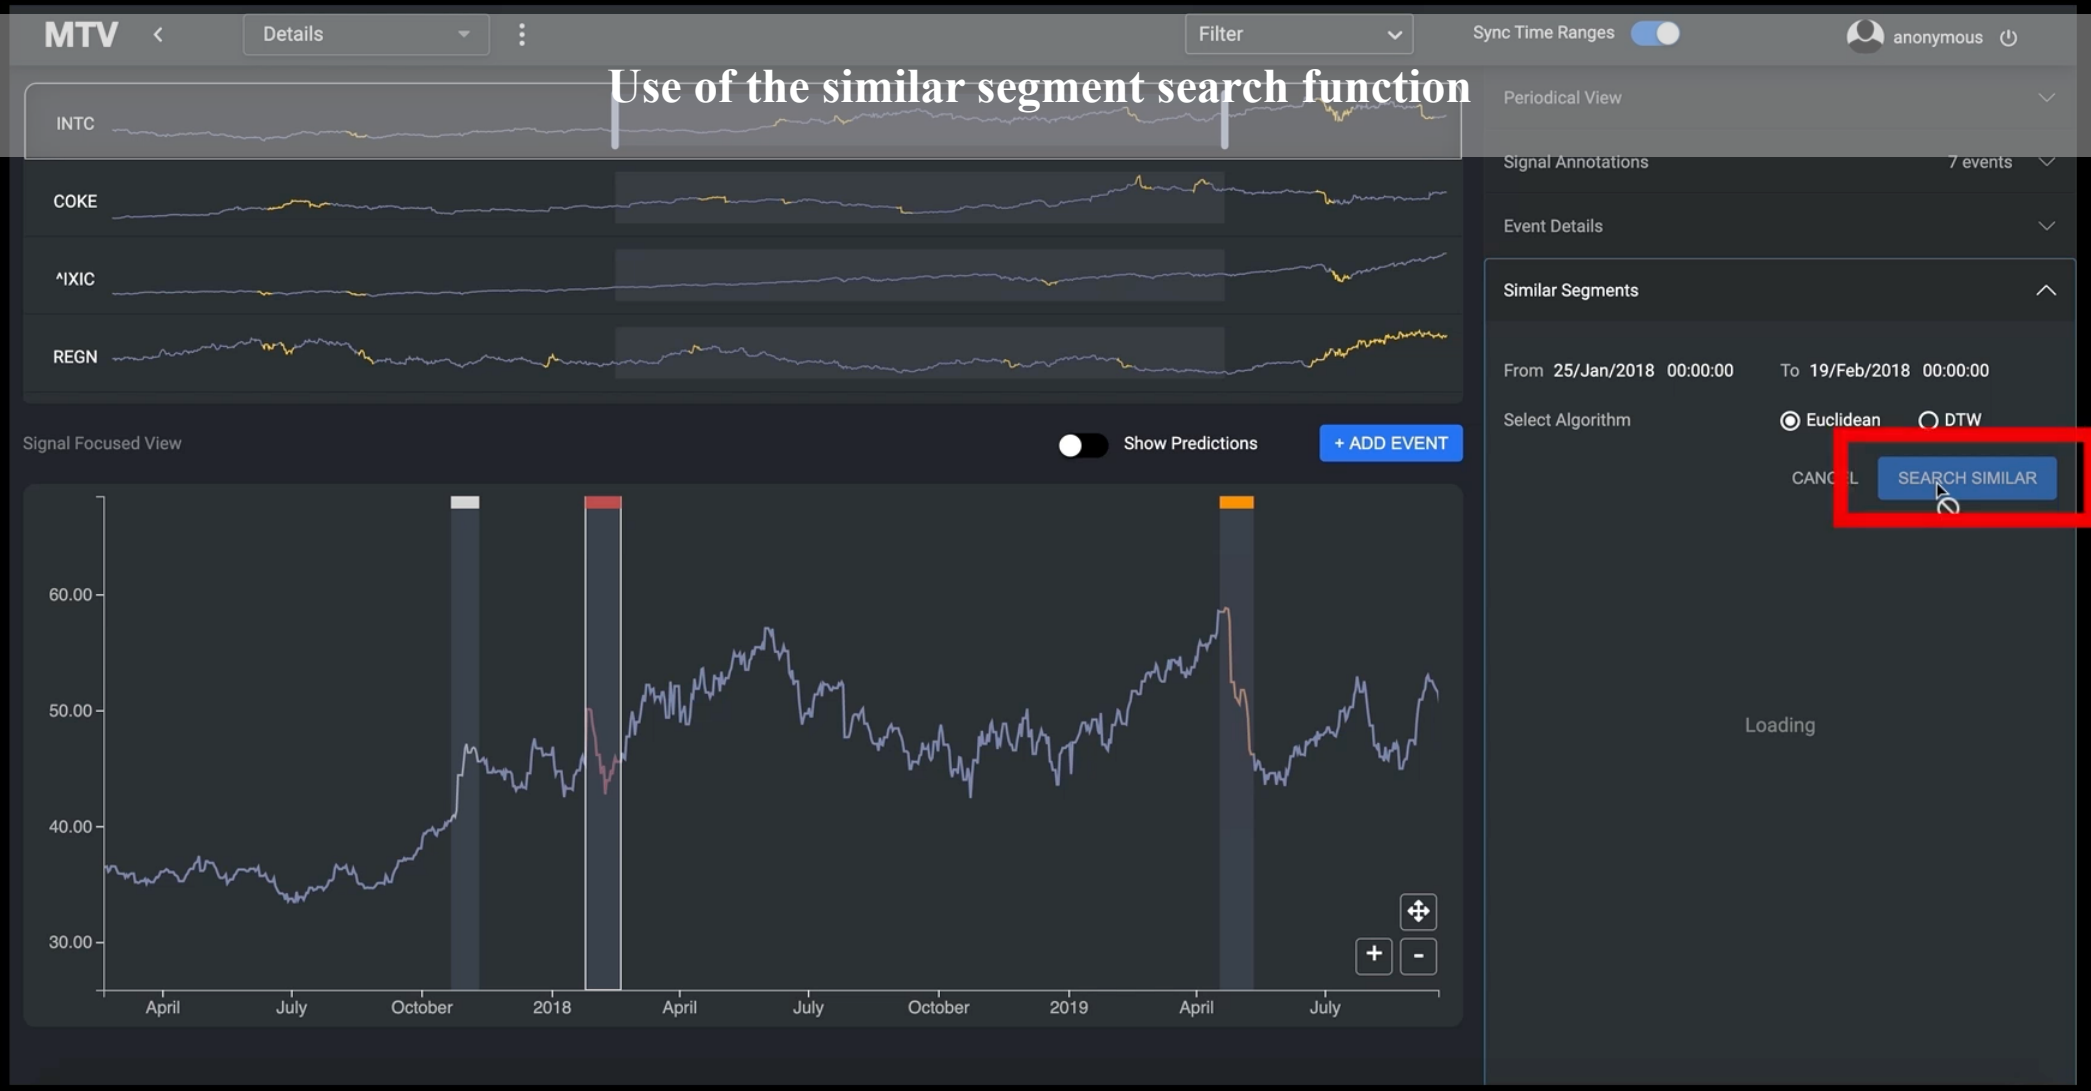

User study

- 6 experts from a satellite operations company

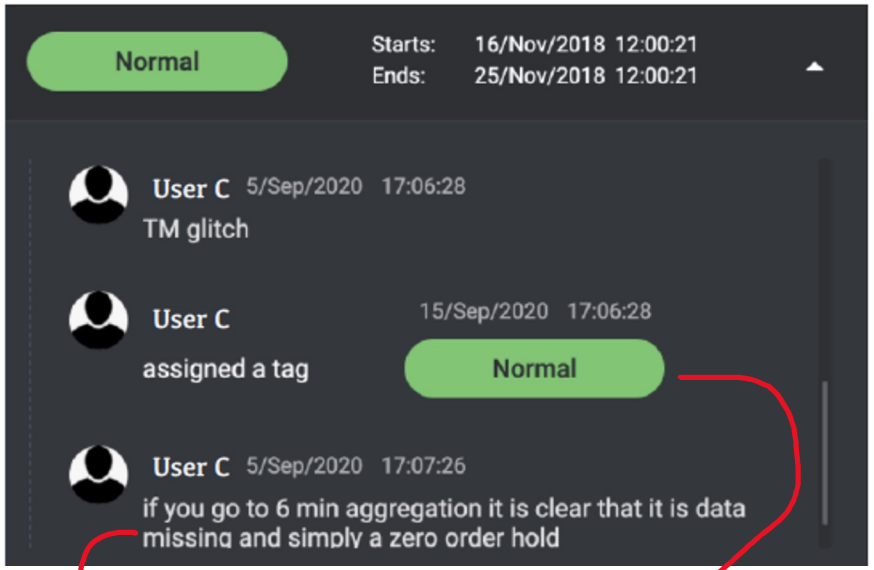

Case	ML Event	User-created Event	Comment (avg.)	Tag (avg.)
1	38	15	162 (3.1)	58 (1.1)
2	45	12	87 (1.5)	60 (1.1)
3	40	8	96 (2.0)	48 (1.0)
4	23	10	66 (2.0)	40 (1.2)

Liu, et al., MTV: Visual Analytics for Detecting, Investigating, and Annotating Anomalies in Multivariate Time Series, CSCW 2022.

User study

- 6 experts from a satellite operations company
- 25 general users using stock data

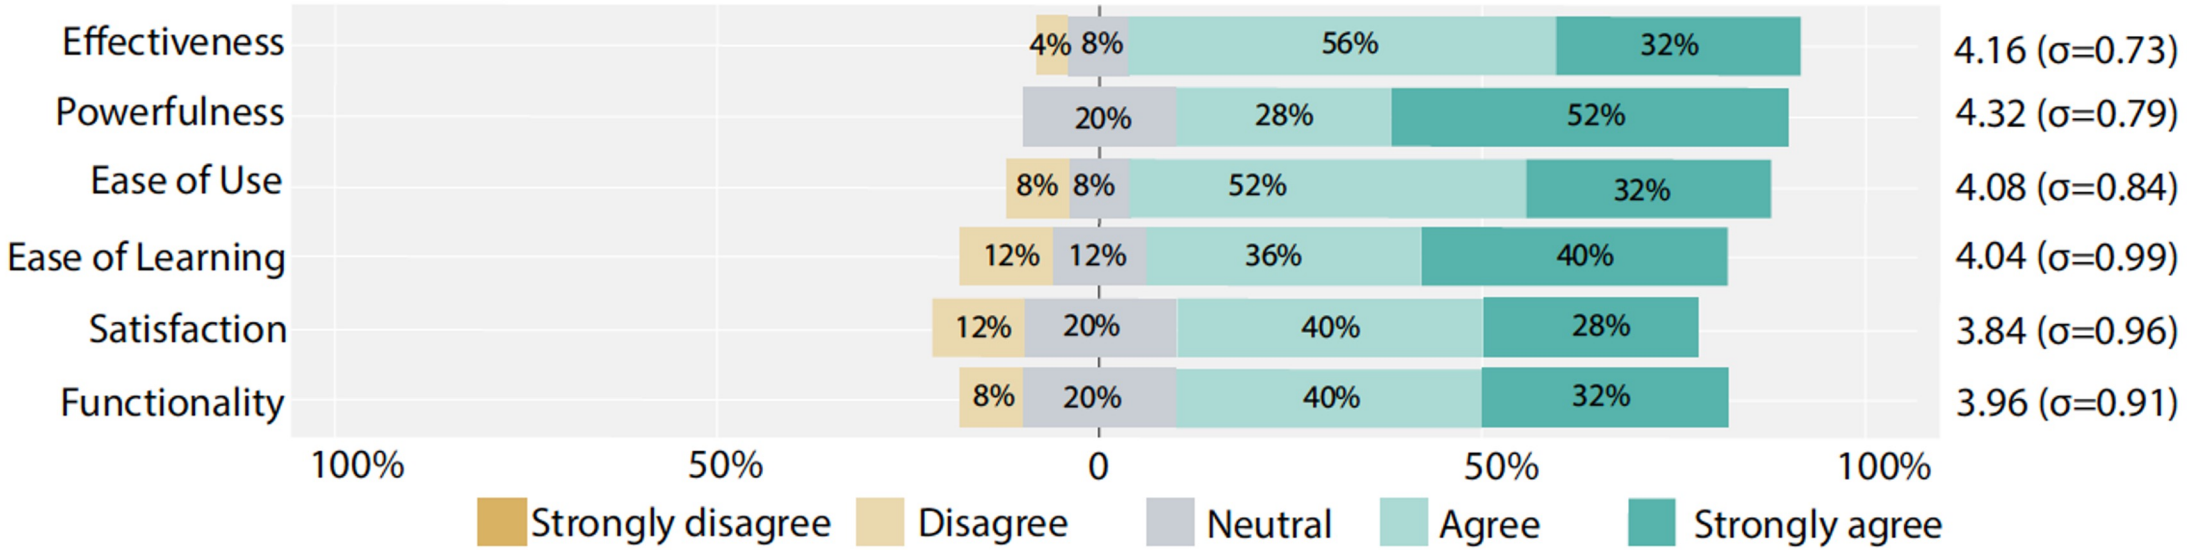

Liu, et al., MTV: Visual Analytics for Detecting, Investigating, and Annotating Anomalies in Multivariate Time Series, CSCW 2022.

Human-AI teaming for time series anomaly detection

Able to learn

AI System
(Human-Centered)

Need to integrate domain knowledge to improve performance

Visual Interface

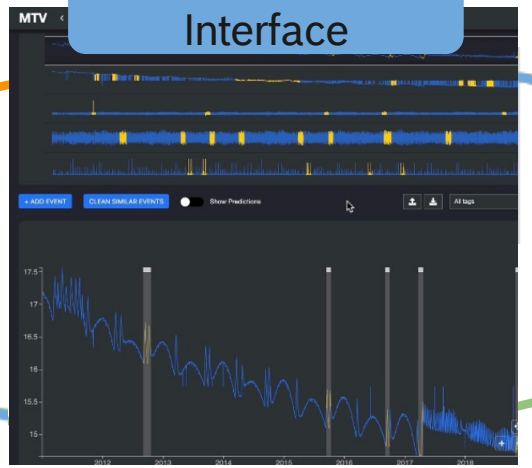

Human

Cannot monitor massive signals without AI to pinpoint suspicious locations

Alnegheimish, **Liu**, et al., Sintel: A Machine Learning Framework to Extract Insights from Signals, SIGMOD 2022.

Liu, et al., MTV: Visual Analytics for Detecting, Investigating, and Annotating Anomalies in Multivariate Time Series, CSCW 2022.

Anomaly Detection

Orion repository metrics (as of 2/1/23)

https://github.com/sintel-dev/Orion

718
github stars

56K
pip downloads

133
forks

352
unique visitors
in 2 weeks

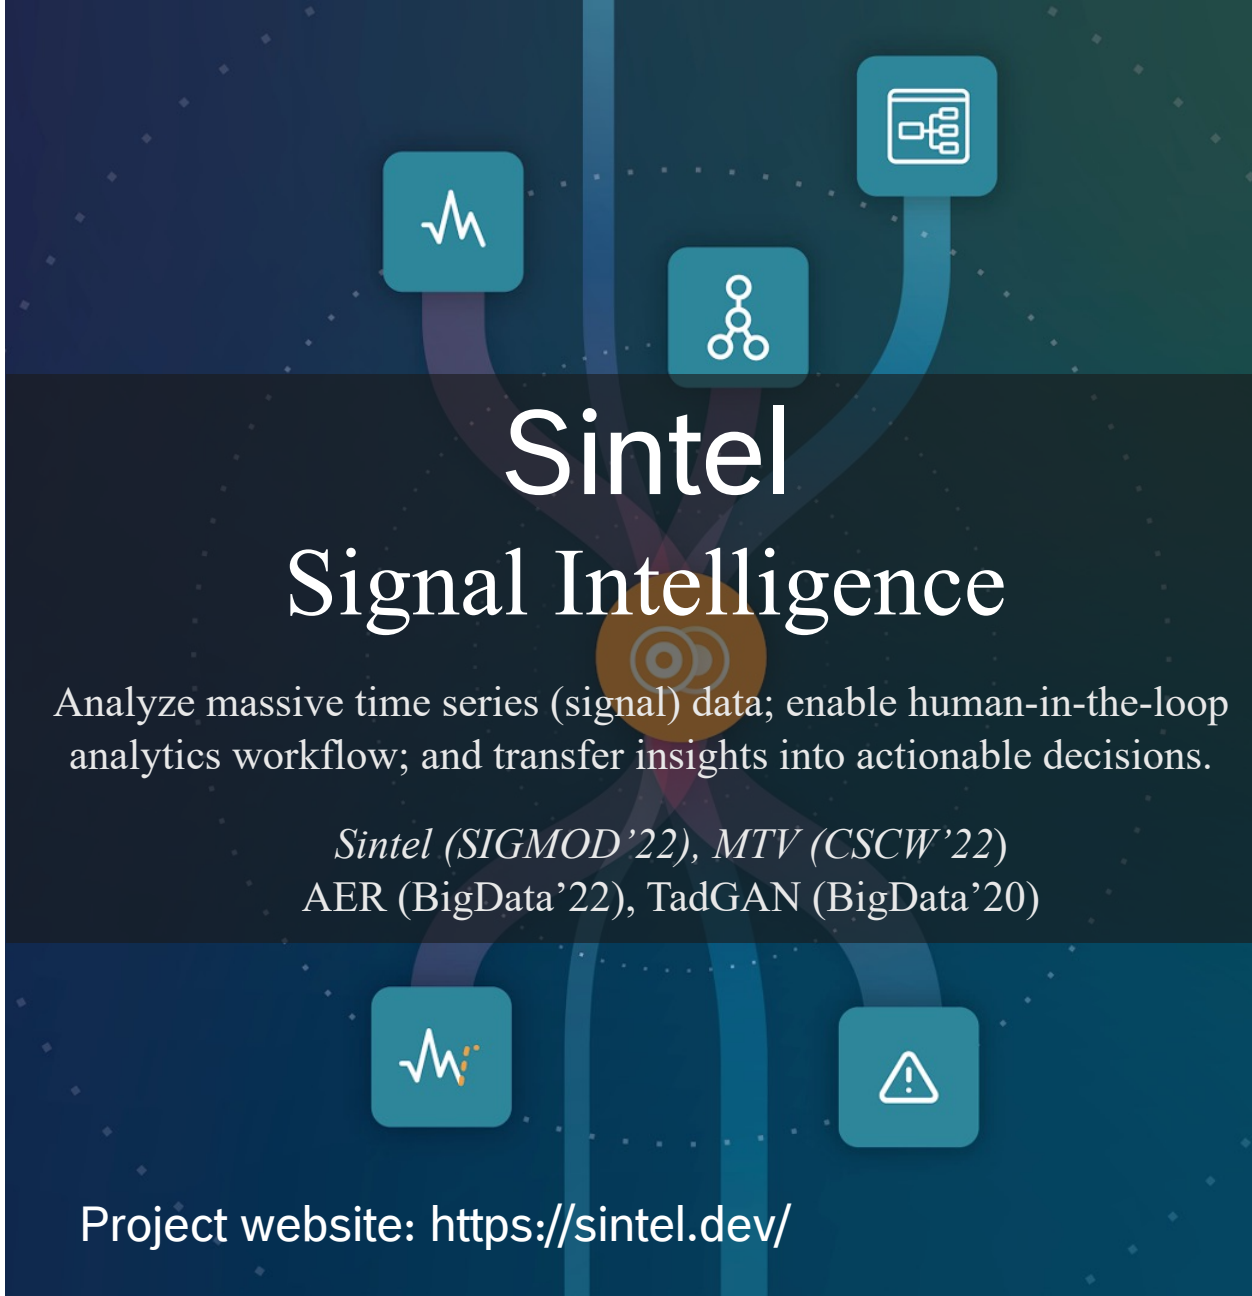

AI Roles

AI suggests preliminary decisions

AI assists in high-stakes decision making (transparent)

AI distills knowledge

Human Factors

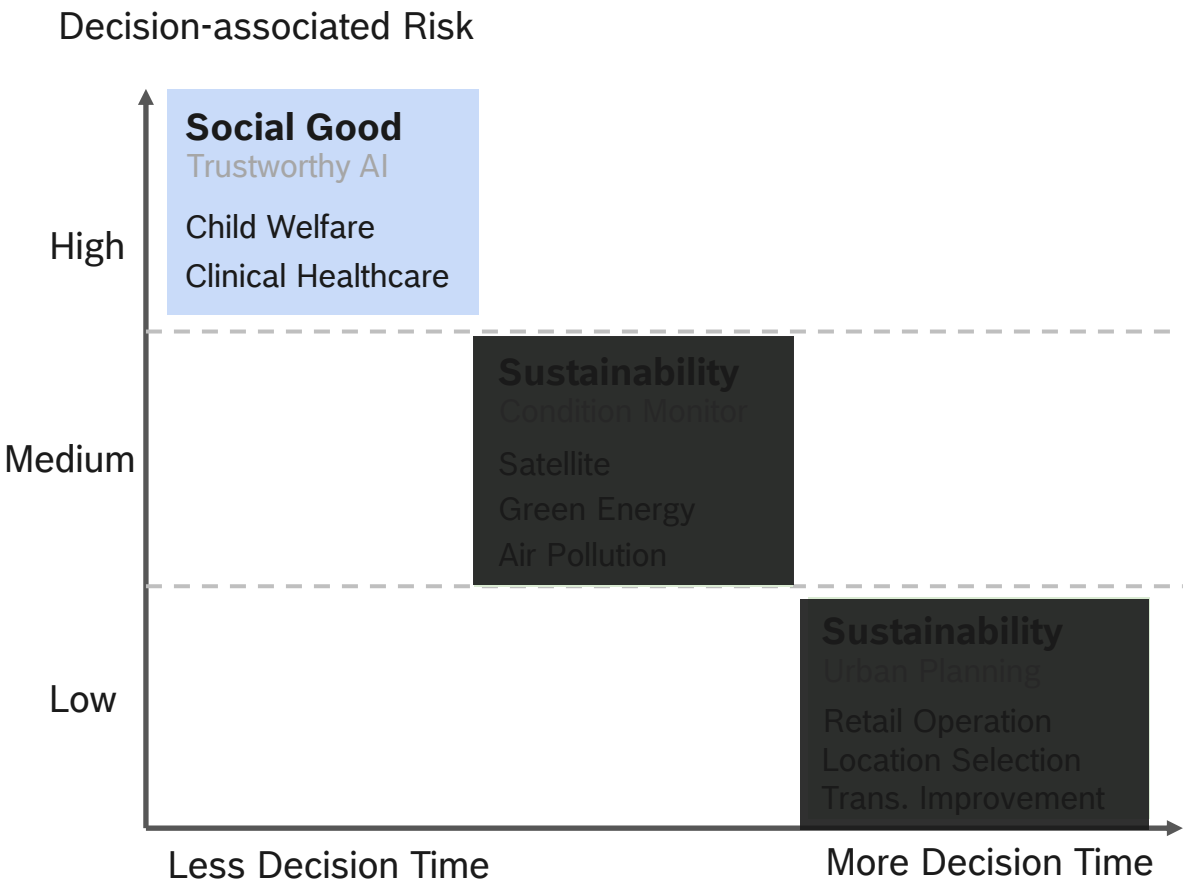

See Sibyl (*VIS'21*) for the full list of human factors

Child abuse hotline screening

- 19 social workers from a child welfare department

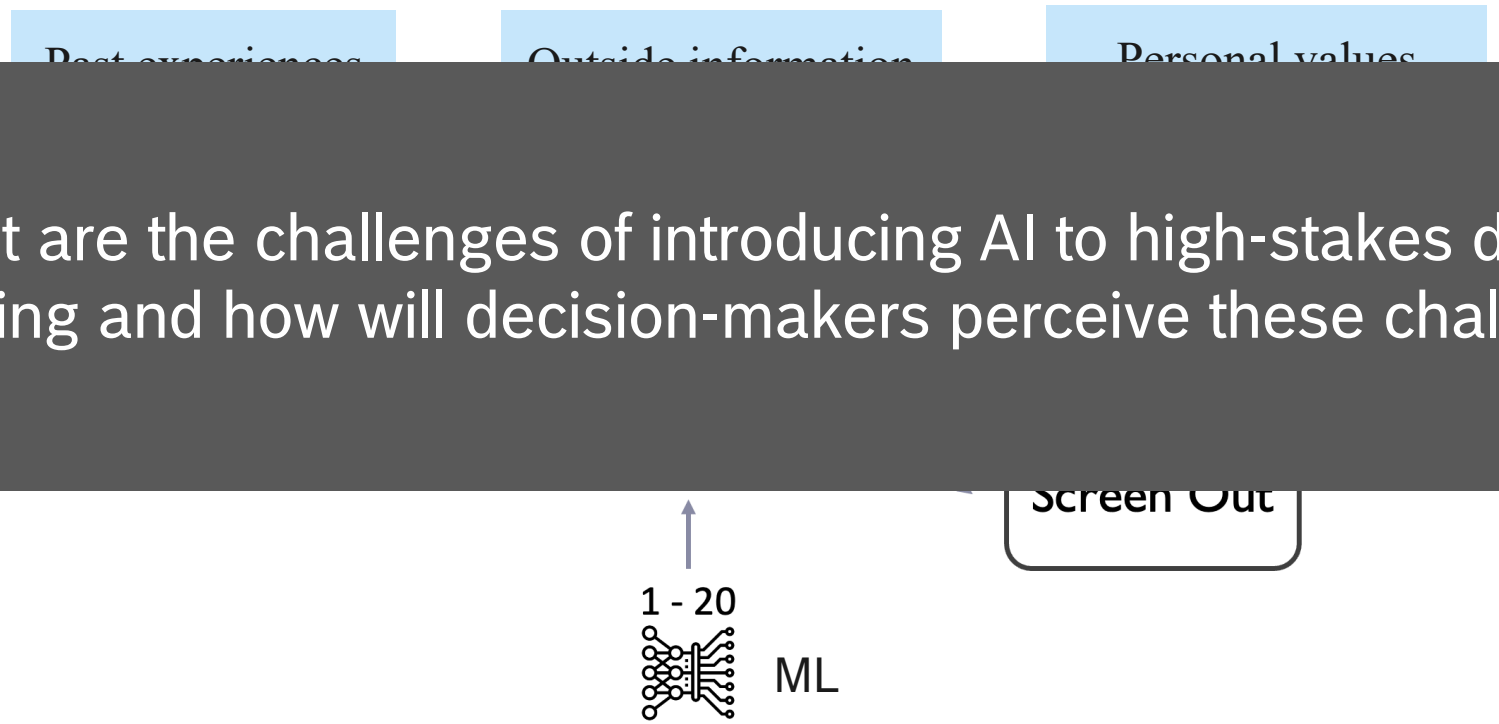

What are the challenges of introducing AI to high-stakes decision-making and how will decision-makers perceive these challenges?

Zytek, **Liu**, et al., Sibyl: Understanding and Addressing the Usability Challenges of Machine Learning In High-Stakes Decision Making, TVCG (VIS'21).

Understand the usability challenges

- List of usability challenges that could negatively impact human-AI teaming

Usability Challenges	Code	Mitigating Tools
Lack of **TR**ust	**TR**	Global explanation, local explanations, performance metrics, historical predictions and results
Difficulty Reconciling human-ML **DIS**agreements	**DIS**	Local explanations
Unclear **CON**sequences of actions	**CON**	Cost-benefit analysis, historical predictions and results
Lack of **ACC**ountability or protections from accountability	**ACC**	Local explanations, performance metrics
ETHical Concerns (e.g., possible bias)	**ETH**	Global explanations, local explanations, ML fairness metrics, historical predictions and results
Confusing or unclear prediction **T**arget (i.e., ML outcome has an unclear meaning or significance)	**CT**	Cost-benefit analysis, further analysis of prediction target impact
Unhelpful prediction **T**arget (i.e., ML outcome is not relevant to the required decision)	**UT**	Retrain model with new prediction target

Zytek, **Liu**, et al., Sibyl: Understanding and Addressing the Usability Challenges of Machine Learning In High-Stakes Decision Making, TVCG (VIS'21).

Child abuse hotline screening

How do users perceive the usability challenges?

Lack of Trust	*"We were surprised [the score] was that low"*
Reconciling Disagreements	*"[The score made us] think a little deeper about why the score is so high [and caused us to] take another look at [the history]"*
Confusing Target	*"Wish we knew how we got to do with the score"*
Ethics	*"...asked ourselves if we are critically thinking"*

Zytek, **Liu**, et al., Sibyl: Understanding and Addressing the Usability Challenges of Machine Learning In High-Stakes Decision Making, TVCG (VIS'21).

Child abuse hotline screening

How do existing mitigating tools can help?

Feature contributions (local explanation)

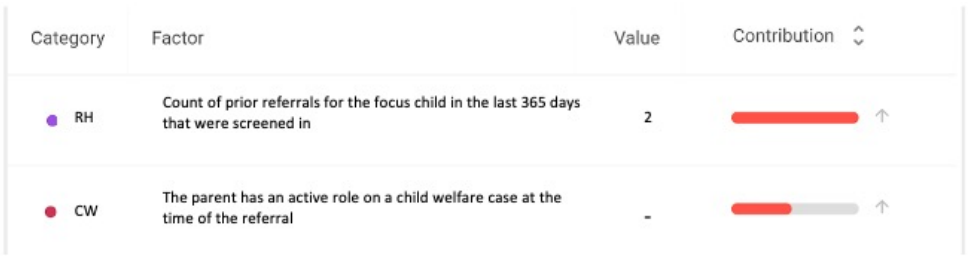

What if

Global feature importance (global explanation)

Feature distributions

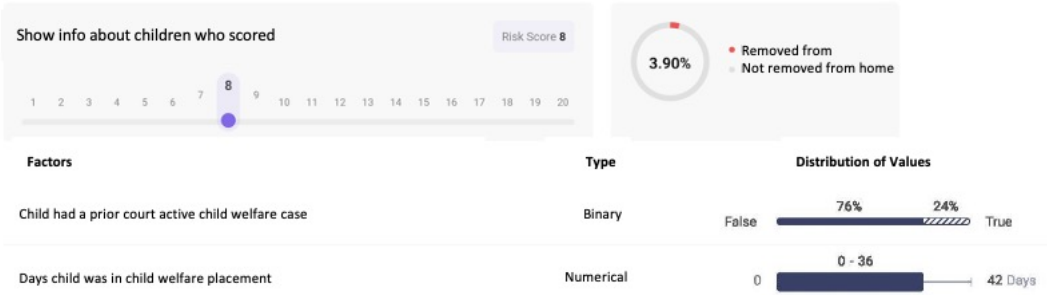

Zytek, **Liu**, et al., Sibyl: Understanding and Addressing the Usability Challenges of Machine Learning In High-Stakes Decision Making, TVCG (VIS'21).

Child abuse hotline screening
Design lesson #1

- Showing feature contributions is most helpful in investigating an individual case.

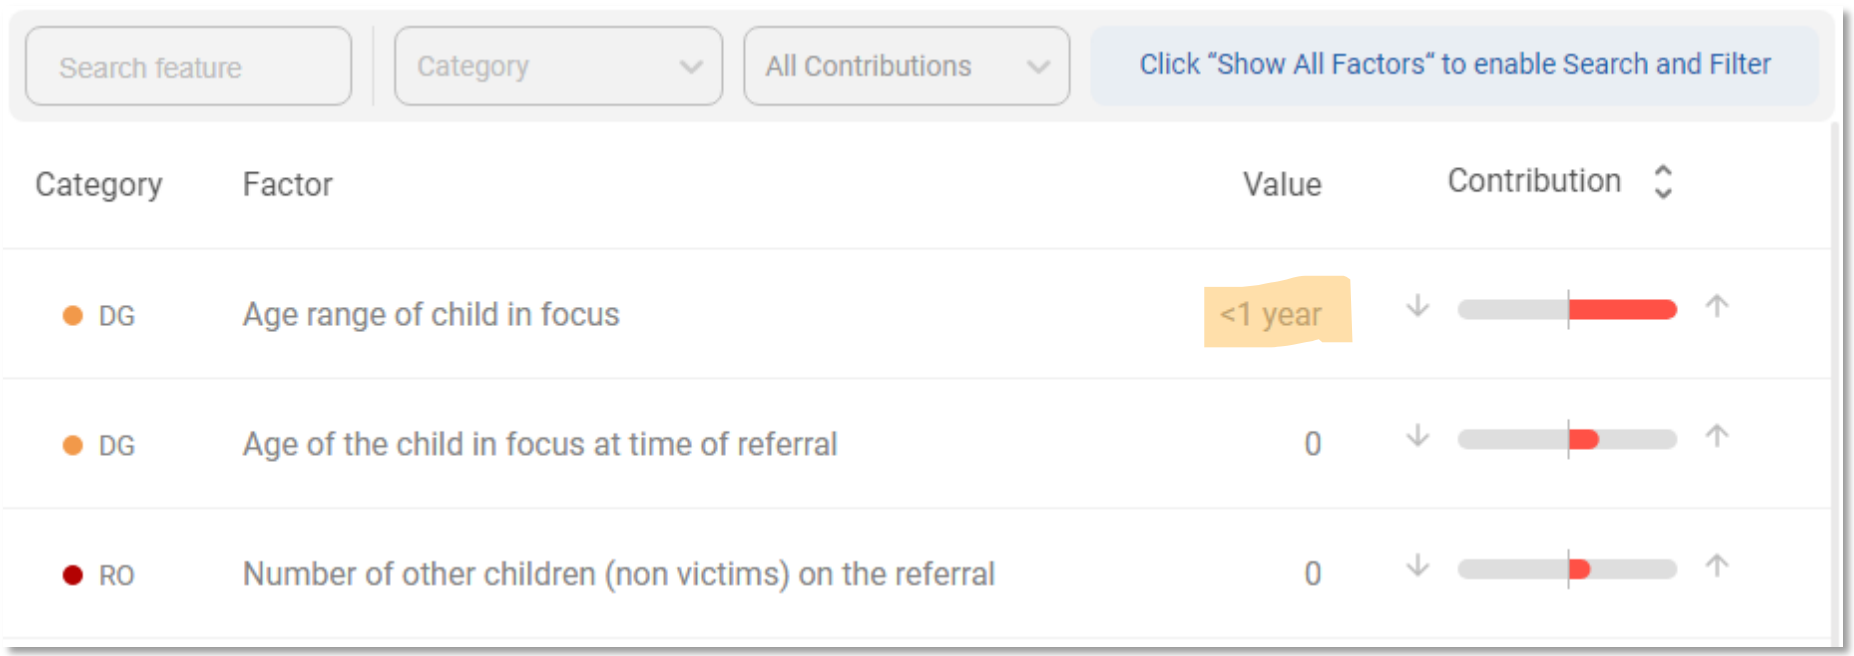

Zytek, **Liu**, et al., Sibyl: Understanding and Addressing the Usability Challenges of Machine Learning In High-Stakes Decision Making, TVCG (VIS'21).

Child abuse hotline screening
Design lesson #2

- Most issues are coming from that features themselves are not meaningful.

Confusing Language

Boolean feature

“the child has a sibling is False”

↓ Use natural language

“the child does not have a sibling”

Zytek, **Liu**, et al., Sibyl: Understanding and Addressing the Usability Challenges of Machine Learning In High-Stakes Decision Making, TVCG (VIS’21).

Child abuse hotline screening
Design lesson #2

- Most issues are coming from that features themselves are not meaningful.

Confusing Language

Irrelevant Features

"2 parents have missing date-of-birth is shown as a significant blue bar which I can't imagine is protective." – Child Welfare Screener

Zytek, **Liu**, et al., Sibyl: Understanding and Addressing the Usability Challenges of Machine Learning In High-Stakes Decision Making, TVCG (VIS'21).

Post-surgical complication prediction

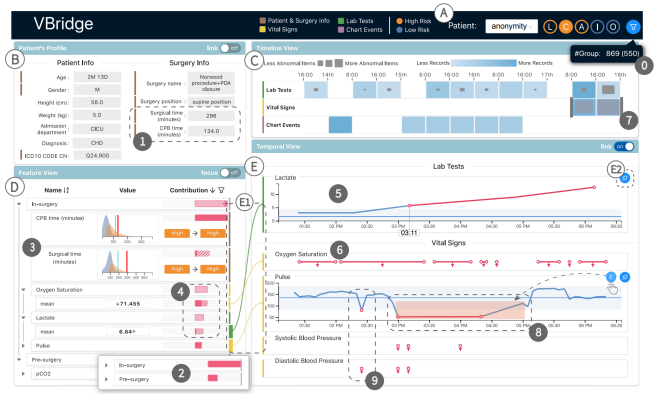

- 6 experienced clinicians

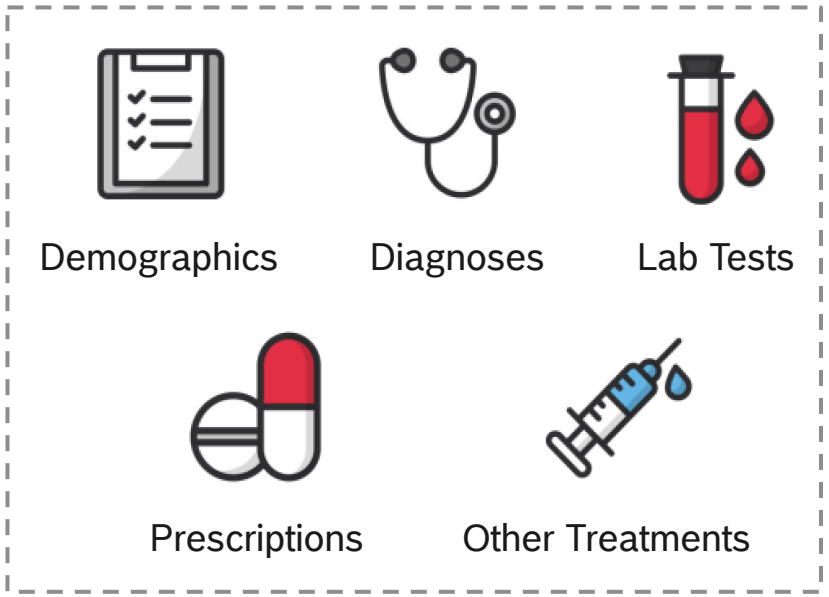

ML predicts five types of post-operative complications

Make treatment plan

Engineered features

Cheng, **Liu**, et al., VBridge: Connecting the Dots Between Features and Data to Explain Healthcare Models, TVCG (VIS'21). Best Paper Honorable Mention.

Post-surgical complication prediction

Design lesson #3

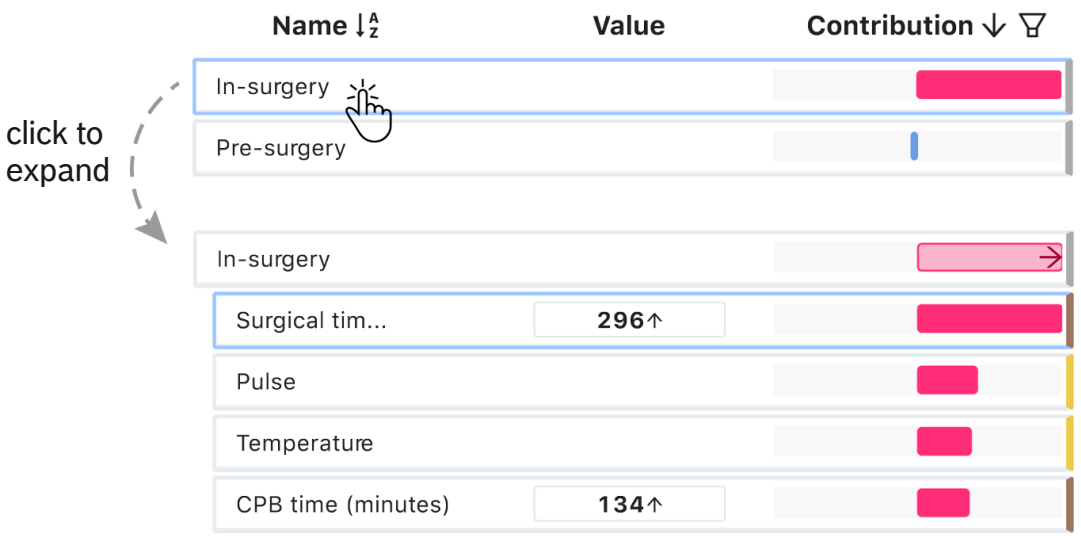

- Unfamiliarity with ML (engineered) features

"We don't often use statistical values like Trend or Standard Deviation (SD) in our work."
(Clinician-P5)

Cheng, **Liu**, et al., VBridge: Connecting the Dots Between Features and Data to Explain Healthcare Models, TVCG (VIS'21). Best Paper Honorable Mention.

Post-surgical complication prediction
Design lesson #4

- Engineered features must be traceable and be explained in the original data space.

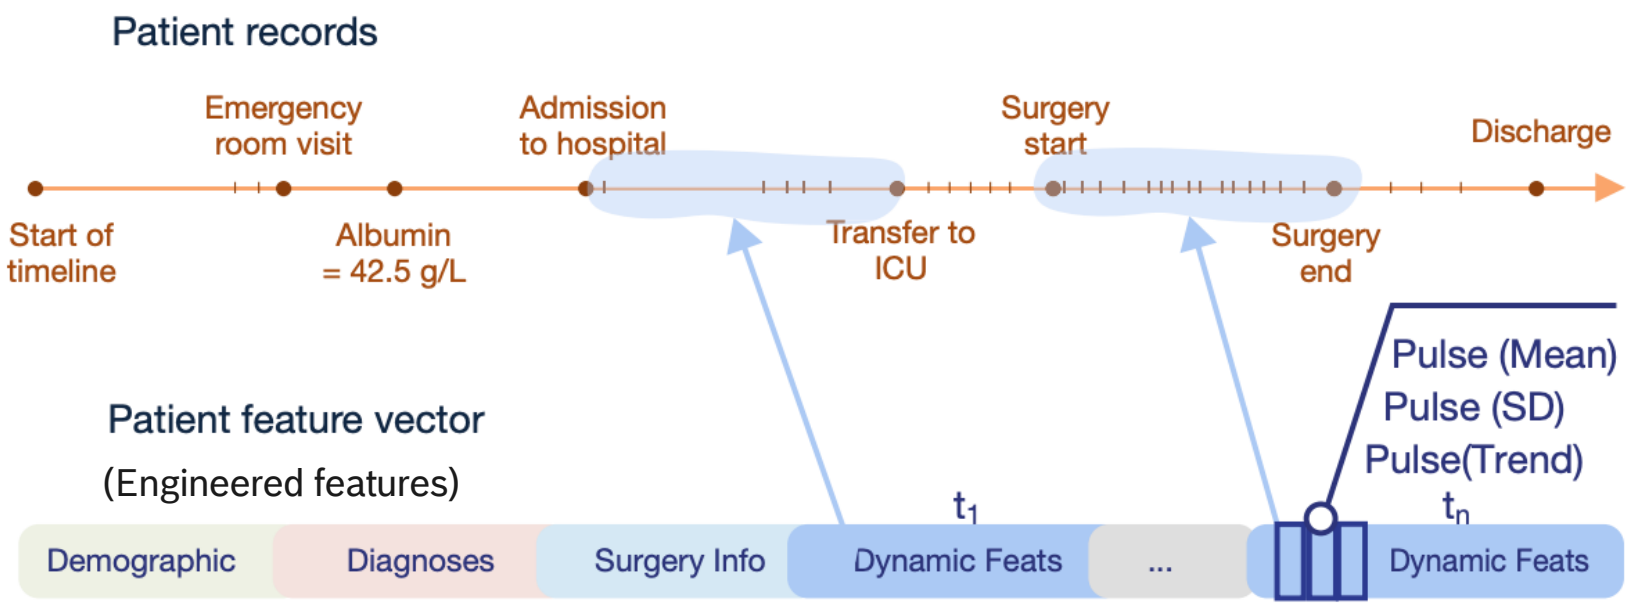

Cheng, **Liu**, et al., VBridge: Connecting the Dots Between Features and Data to Explain Healthcare Models, TVCG (VIS'21). Best Paper Honorable Mention.

Post-surgical complication prediction

Design lesson #4

- Engineered features must be traceable and be explained in the original data space.

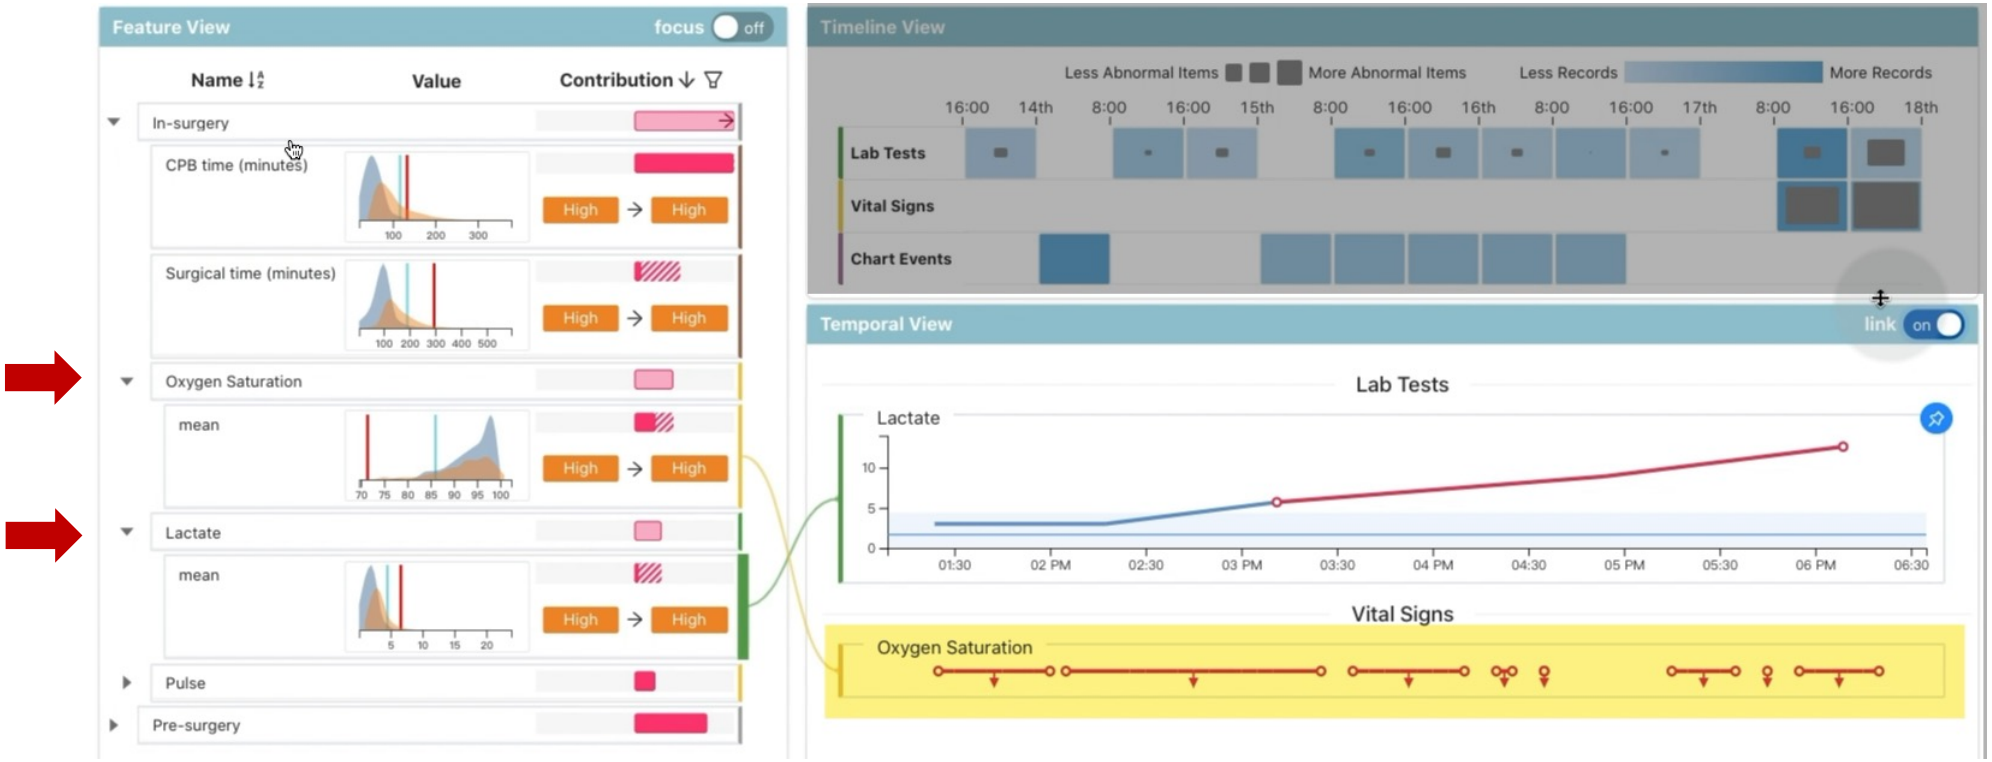

Cheng, **Liu**, et al., VBridge: Connecting the Dots Between Features and Data to Explain Healthcare Models, TVCG (VIS'21). Best Paper Honorable Mention.

What are the features that are most useful and meaningful to users?

Interpretable Features

Zytek, Arnaldo, **Liu**, et al., The Need for Interpretable Features: Motivation and Taxonomy, SIGKDD Explorations Newsletter 2022.

Example: Housing Price Prediction

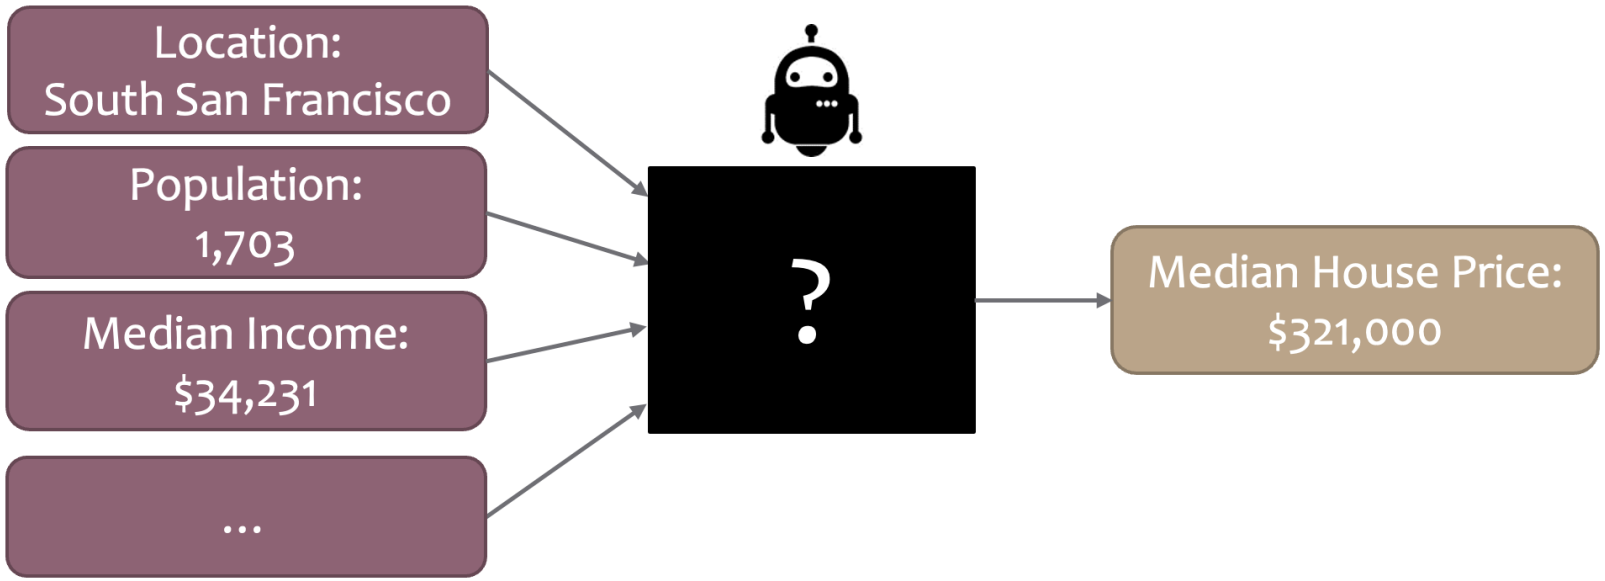

Zytek, Arnaldo, **Liu**, et al., The Need for Interpretable Features: Motivation and Taxonomy, SIGKDD Explorations Newsletter 2022.

Taxonomy of interpretable features

	Area Quality (numeric)	Average House Size (numeric)	Common House Color (categorical)	Normalized Median Income (numeric)	x12 (numeric)
Readable	✓	✓	✓	✓	
Understandable	✓	✓	✓		
Relevant	✓	✓			

… 5 more …

Zytek, Arnaldo, **Liu**, et al., The Need for Interpretable Features: Motivation and Taxonomy, SIGKDD Explorations Newsletter 2022.

How we generate interpretable features in real-world scenarios?

Zytek, Arnaldo, **Liu**, et al., The Need for Interpretable Features: Motivation and Taxonomy, SIGKDD Explorations Newsletter 2022.

Generate interpretable features

Pyreal: a system for interpretable transforms

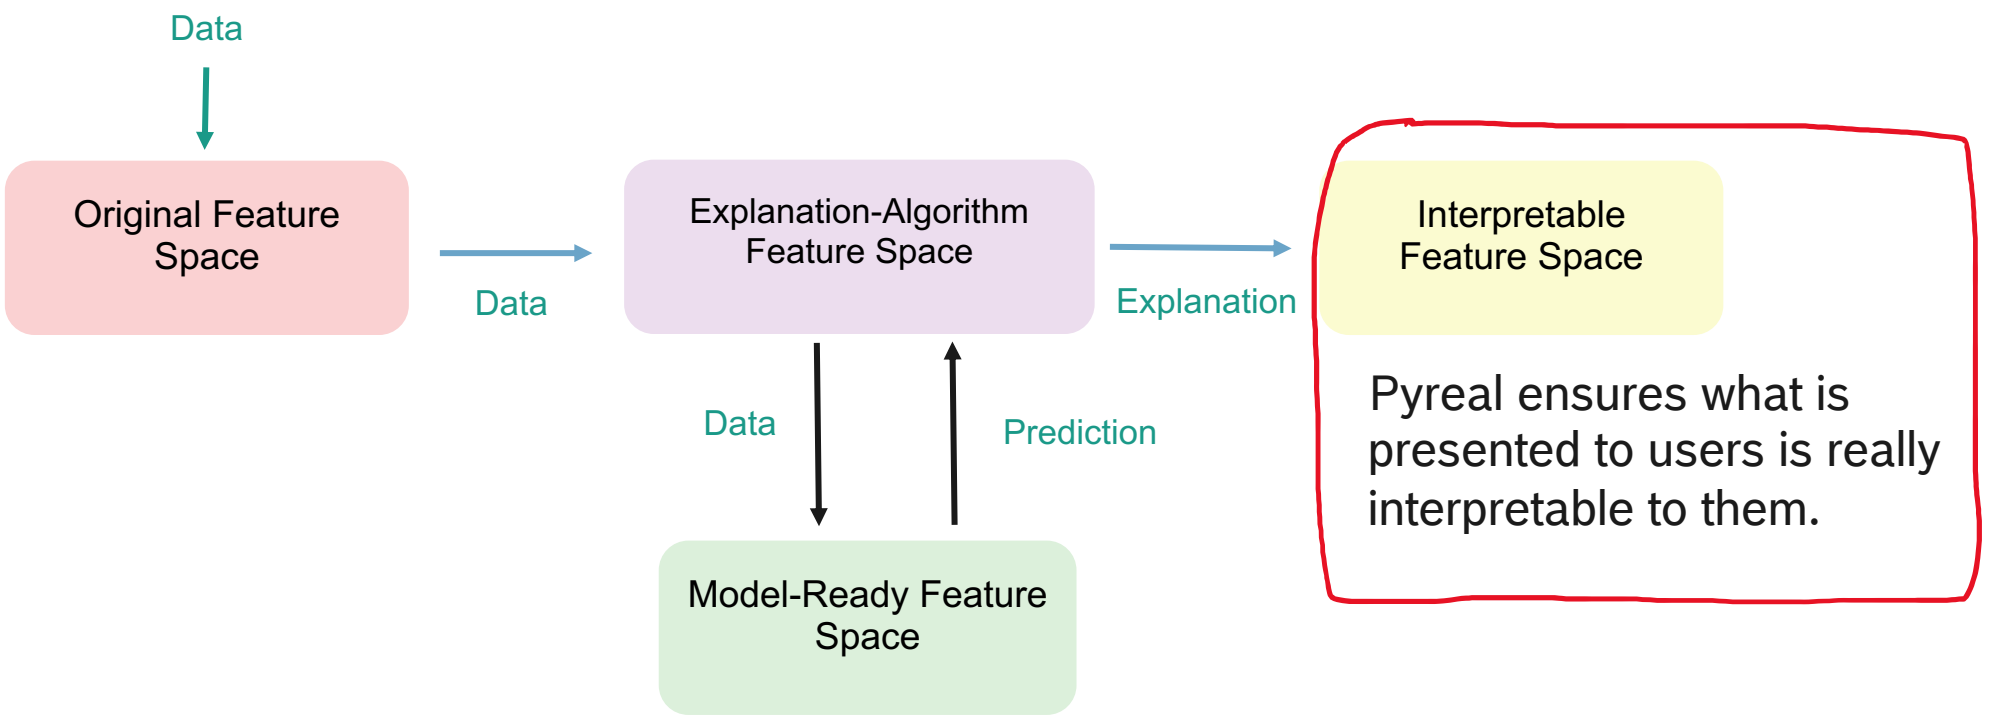

Zytek, **Liu**, et al., An Interpretable ML Explanation Framework, In Submission.

Original outputs of XAI algorithms

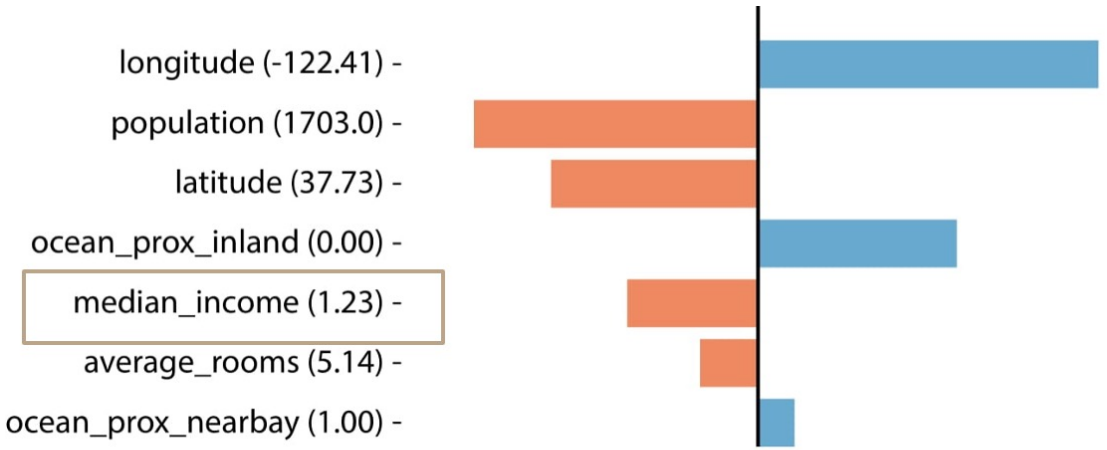

pyreal

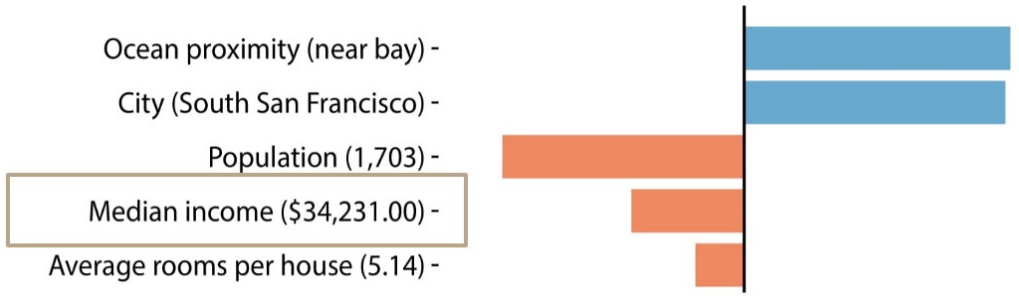

Housing price prediction

Zytek, **Liu**, et al., An Interpretable ML Explanation Framework, In Submission.

Original outputs of XAI algorithms

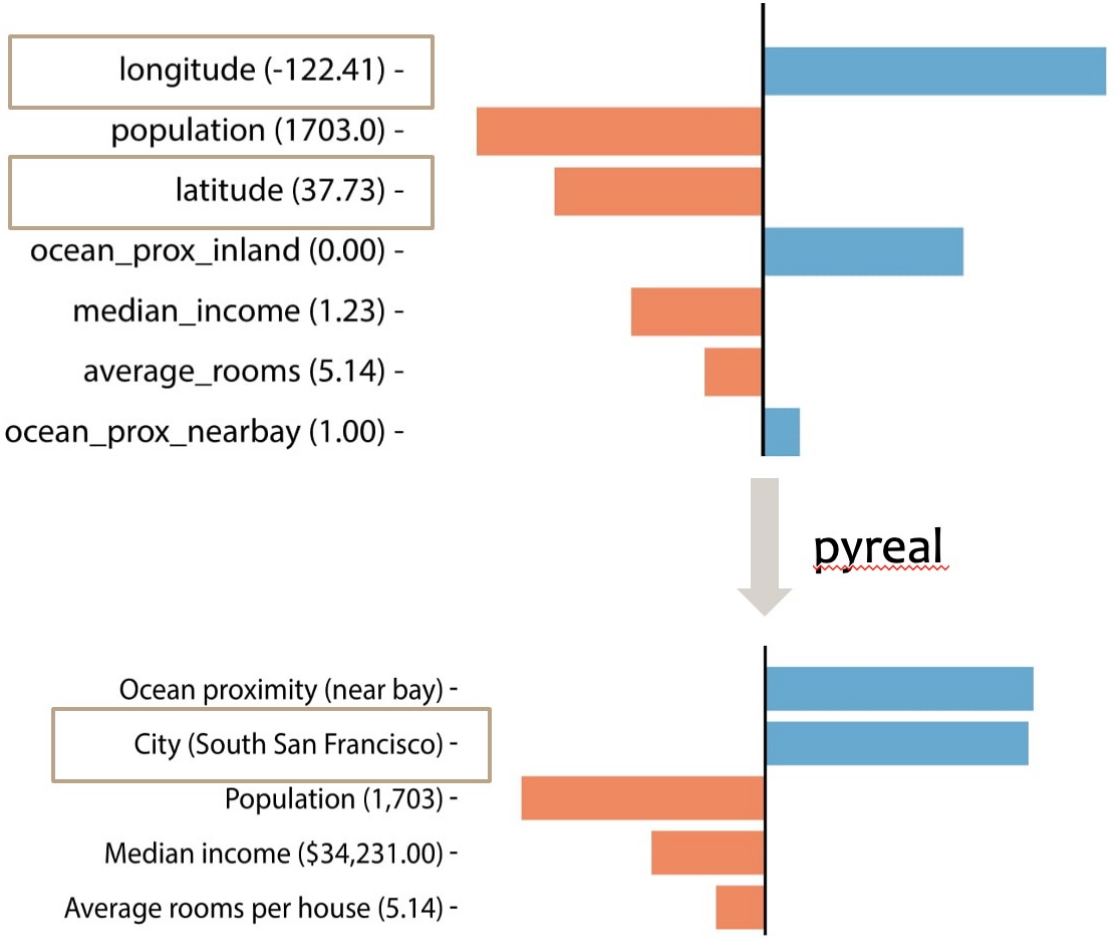

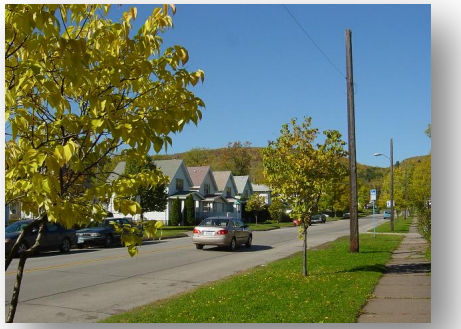

Housing price prediction

Zytek, **Liu**, et al., An Interpretable ML Explanation Framework, In Submission.

Human-AI teaming for high-stakes decision-making

Transparency

AI System
(Human-Centered)

Intransparency raises trust issues

Visual
Interface

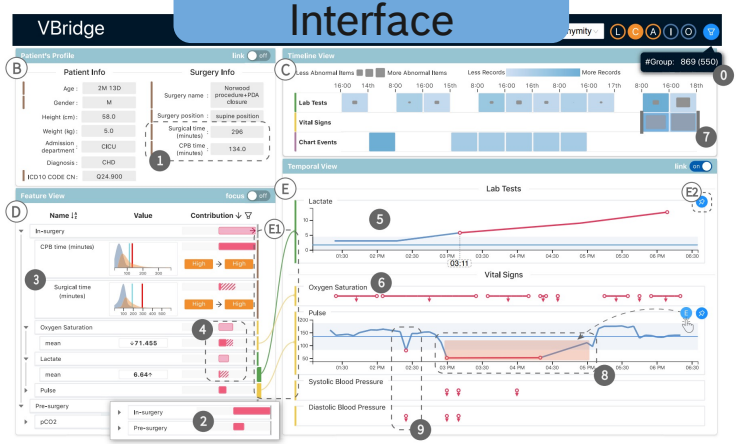

Human

Not as efficient as when
AI provides augmented information

Zytek, **Liu**, et al., Sibyl: Understanding and Addressing the Usability Challenges of Machine Learning In High-Stakes Decision Making, TVCG (VIS'21).
Cheng, **Liu**, et al., VBridge: Connecting the Dots Between Features and Data to Explain Healthcare Models, TVCG (VIS'21). Best Paper Honorable Mention.
Zytek, Arnaldo, **Liu**, et al., The Need for Interpretable Features: Motivation and Taxonomy, SIGKDD Explorations Newsletter 2022.
Zytek, **Liu**, et al., An Interpretable ML Explanation Framework, In Submission.

AI Roles

AI suggests preliminary decisions

AI assists in high-stakes decision making

AI distills knowledge (steerable)

Human Factors

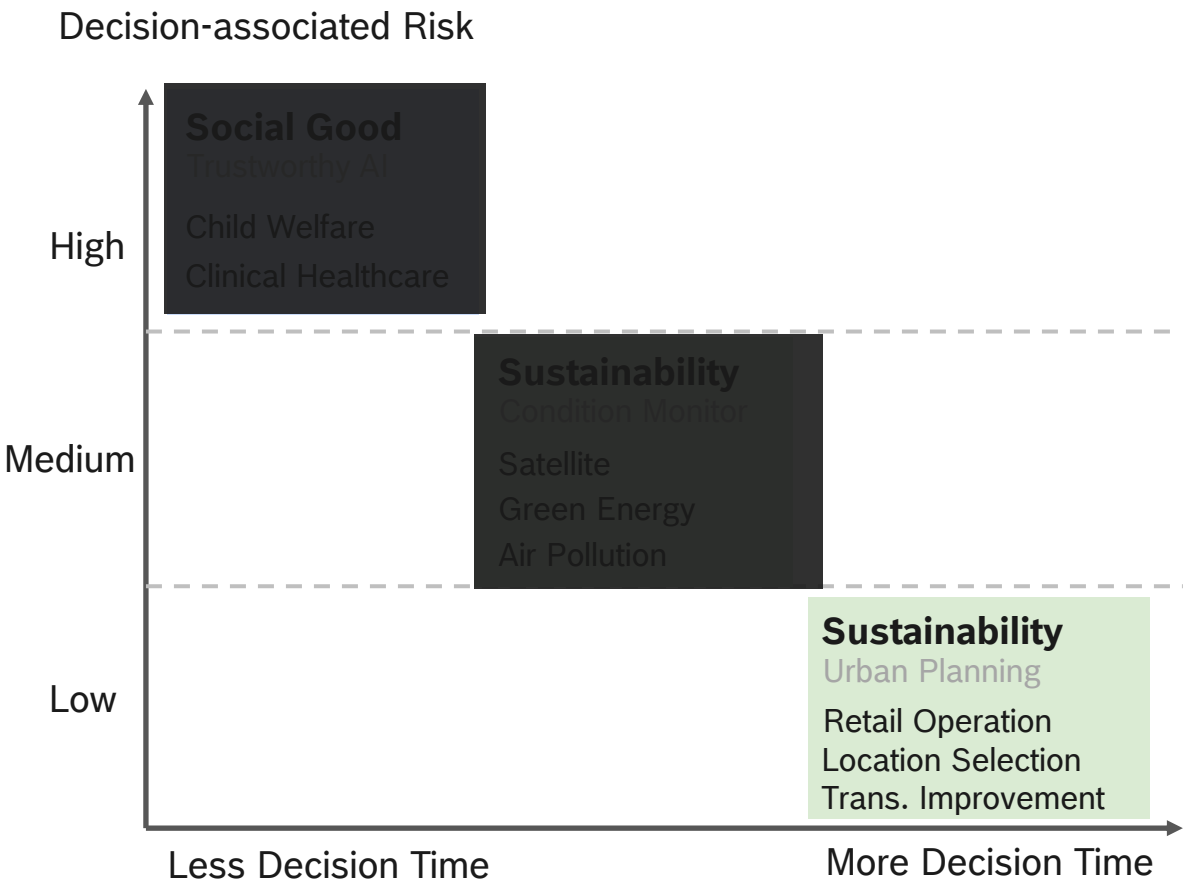

See Sibyl (*VIS'21*) for the full list of human factors

Sustainability
Urban planning

How can we help users extract the most relevant knowledge from data to guide complex decision making?

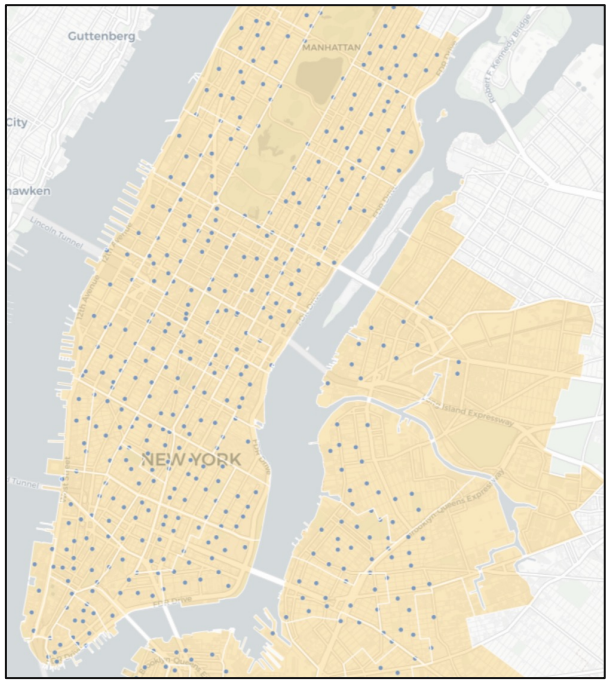

Spatio-temporal data

(an observed event with both time and location information, e.g., vehicle GPS data, mobile phone data)

Knowledge: where & when

Retail operation optimization

Regional Sales optimization

Transportation Improvement

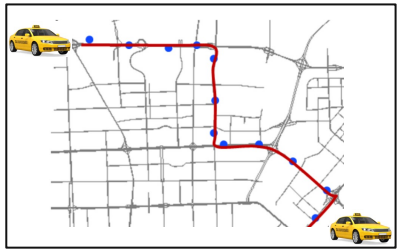

Optimal Location selection

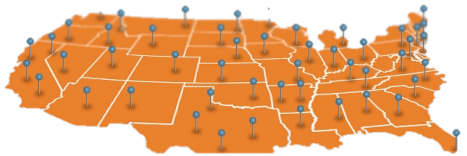

Start with an example

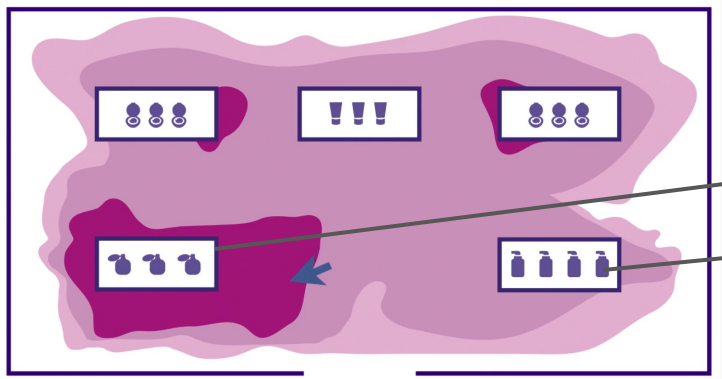

1 Oct

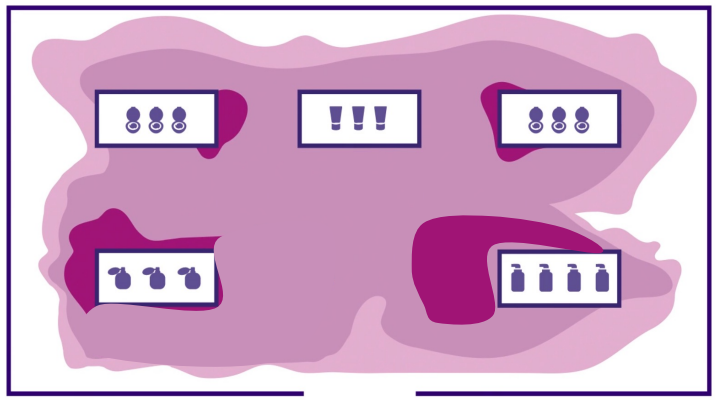

1 Oct between 10am to 11 am

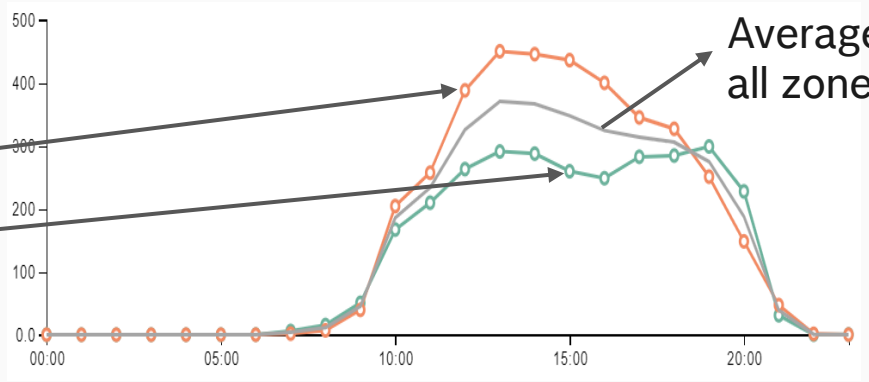

Hourly traffic

Useful patterns (most-relevant knowledge) exist in certain data subsets

Manual search does not scale

Nanocubes
Lins et al. 2013

imMens
Liu et al. 2013

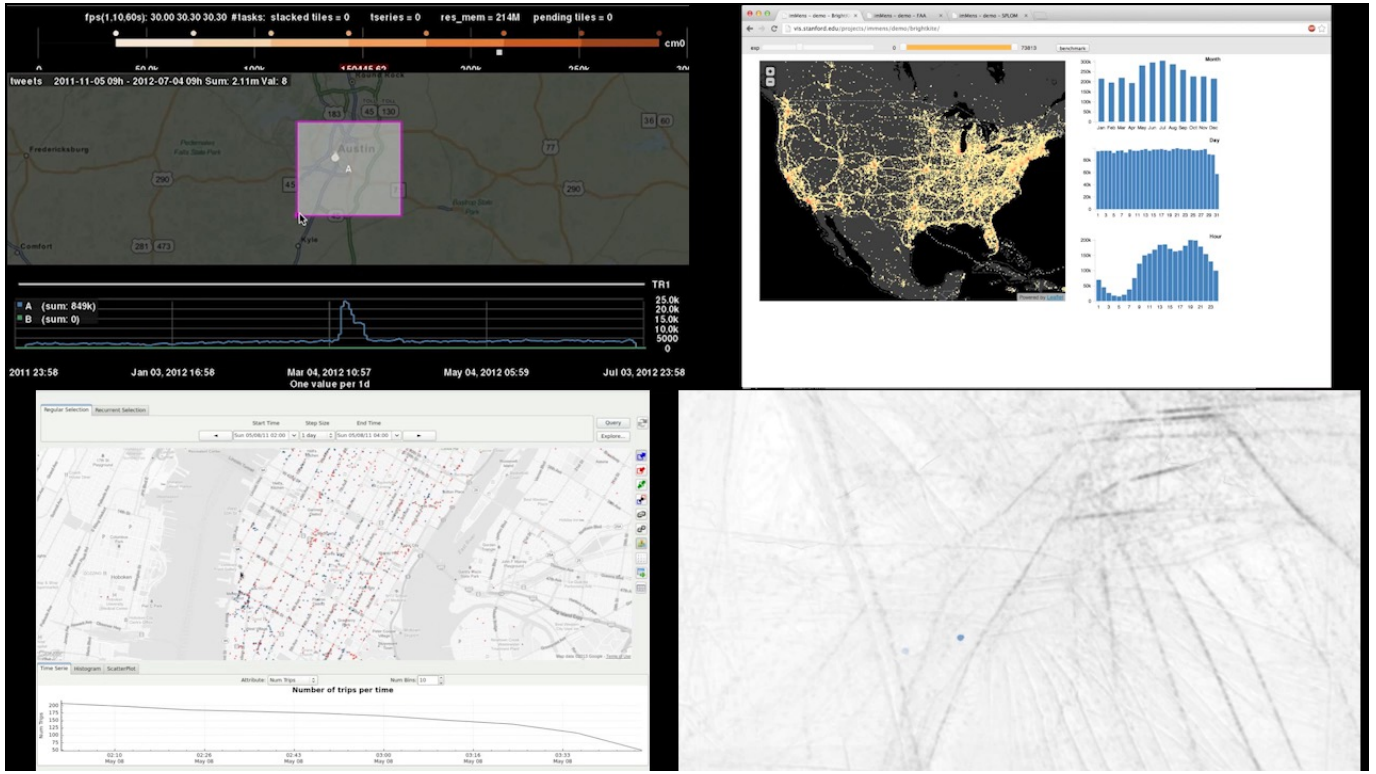

TaxiVis
Ferreira et al. 2013

Traffic Flow Analysis
Scheepens et al. 2016

⊖ exhaustive manual search for patterns in data subsets

Existing automated approaches are limited

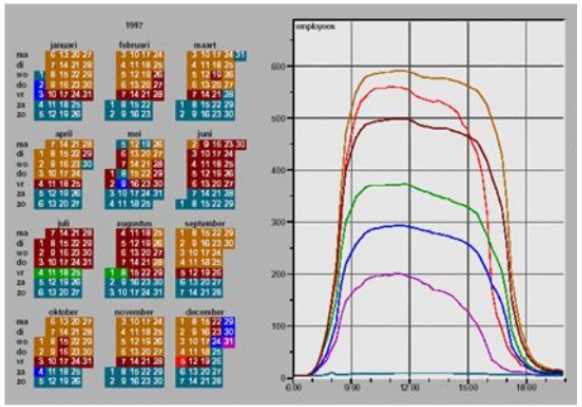

Calendar Vis
Van Wijk and Van Selow, 1999

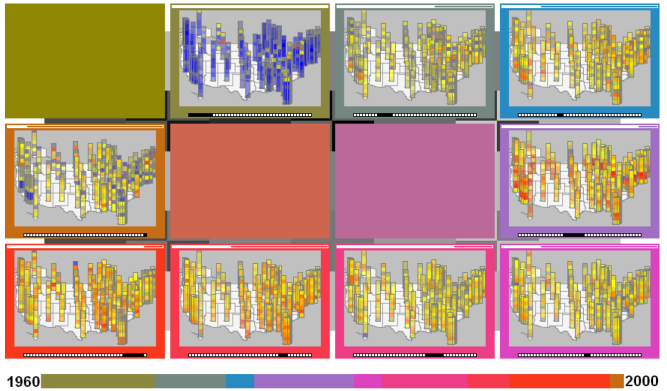

Self-organizing Map
Andrienko et al., 2010

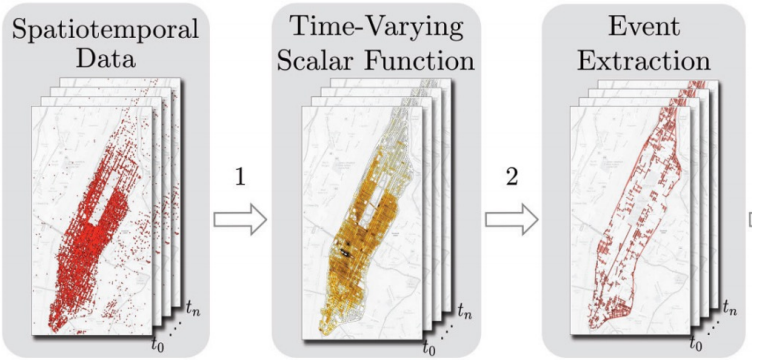

Event-Guided Exploration
Doraiswamy et al., 2014

⊖ fail to support high-dimensional (**>2**) ST data

⊖ not human-steerable

Explore large information space with recommendation

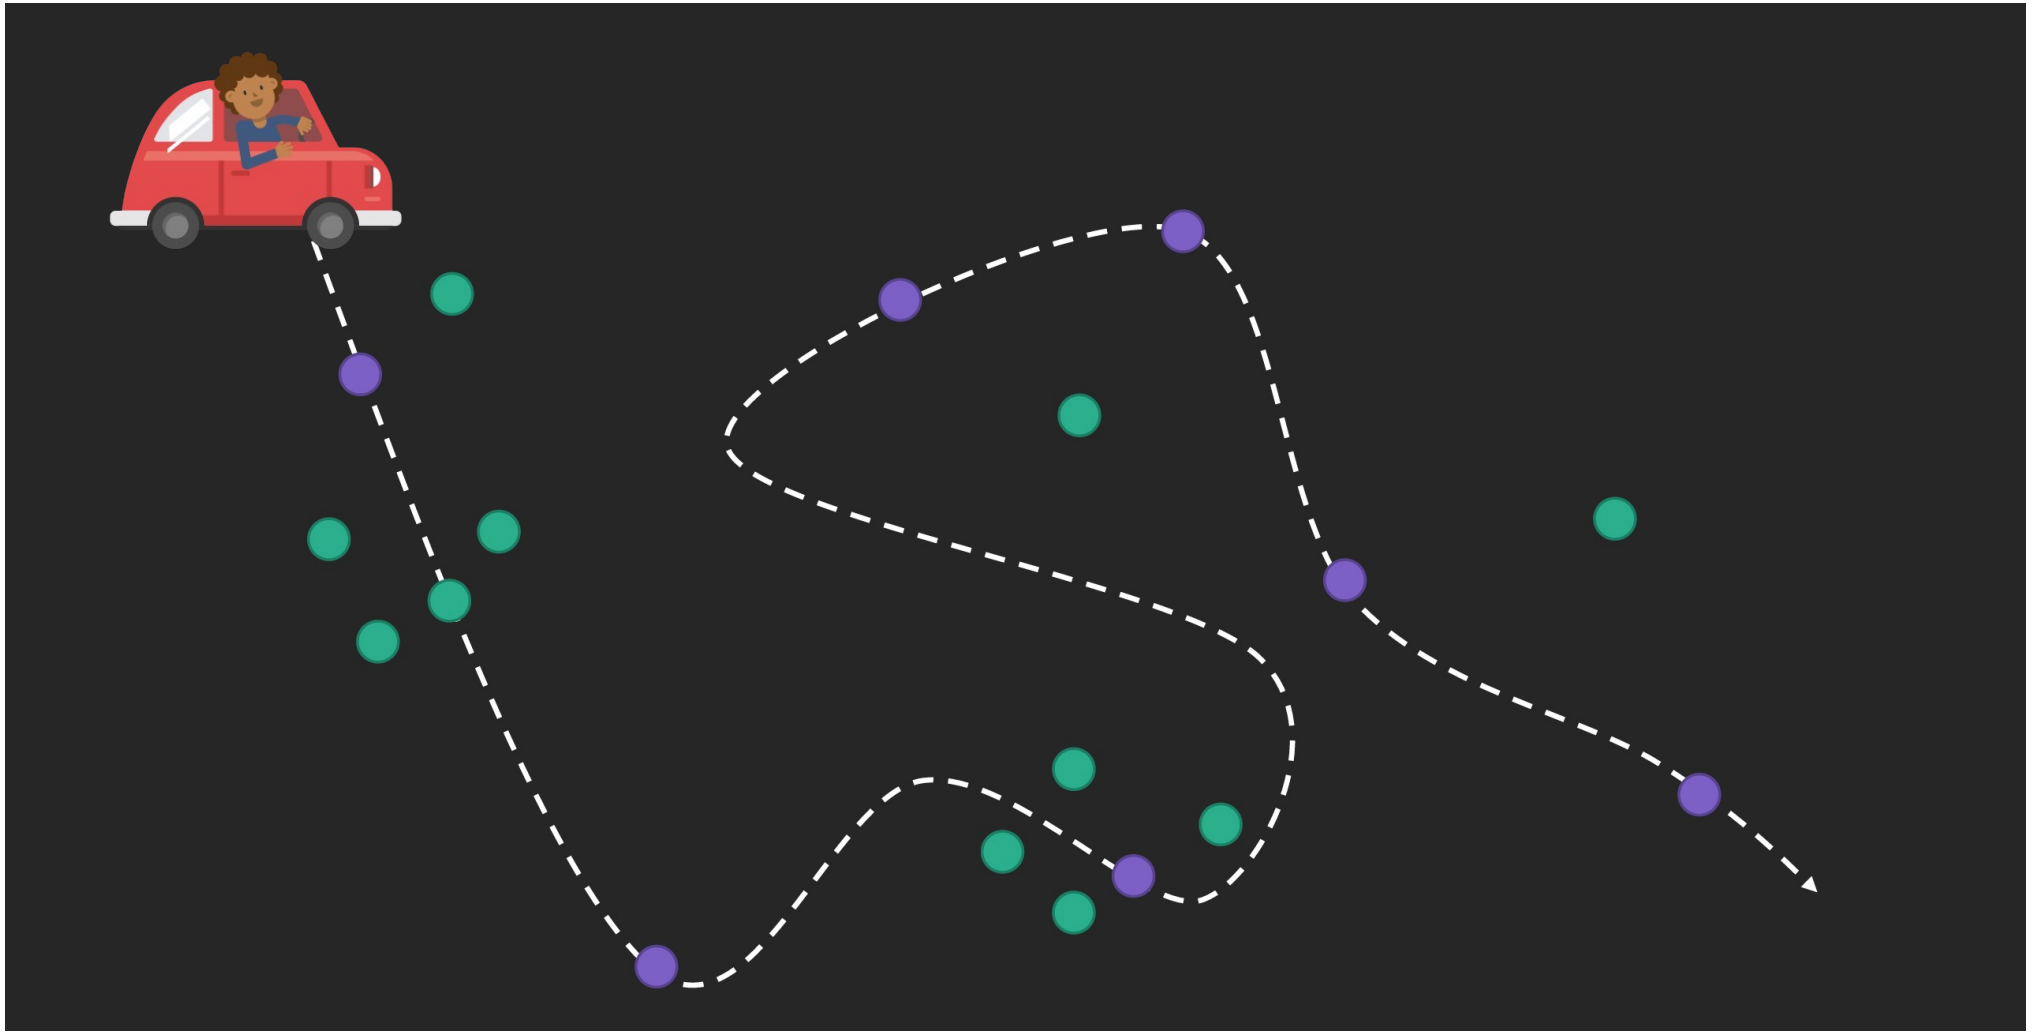

Modeling as tensors

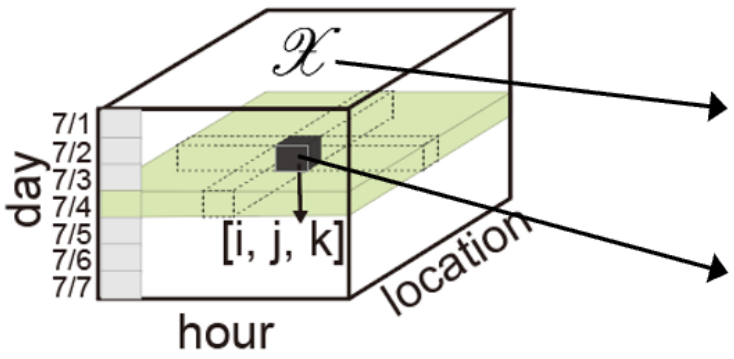

typical traffic dataset

- $\mathcal{X} \in \mathbb{R}_{\geq 0}^{\|P\| \times \|D\| \times \|T\|}$ is a three-dimensional tensor
- $\mathcal{X}[i, j, k]$ represents the traffic volume at location k on the day j during hour i

Liu, et al., TPFlow: Progressive Partition and Multidimensional Pattern Extraction for Large-scale Spatio-temporal Data Analysis, TVCG (VAST'18), Best Paper Award.

Semi-automatic and human-steerable tensor partitions

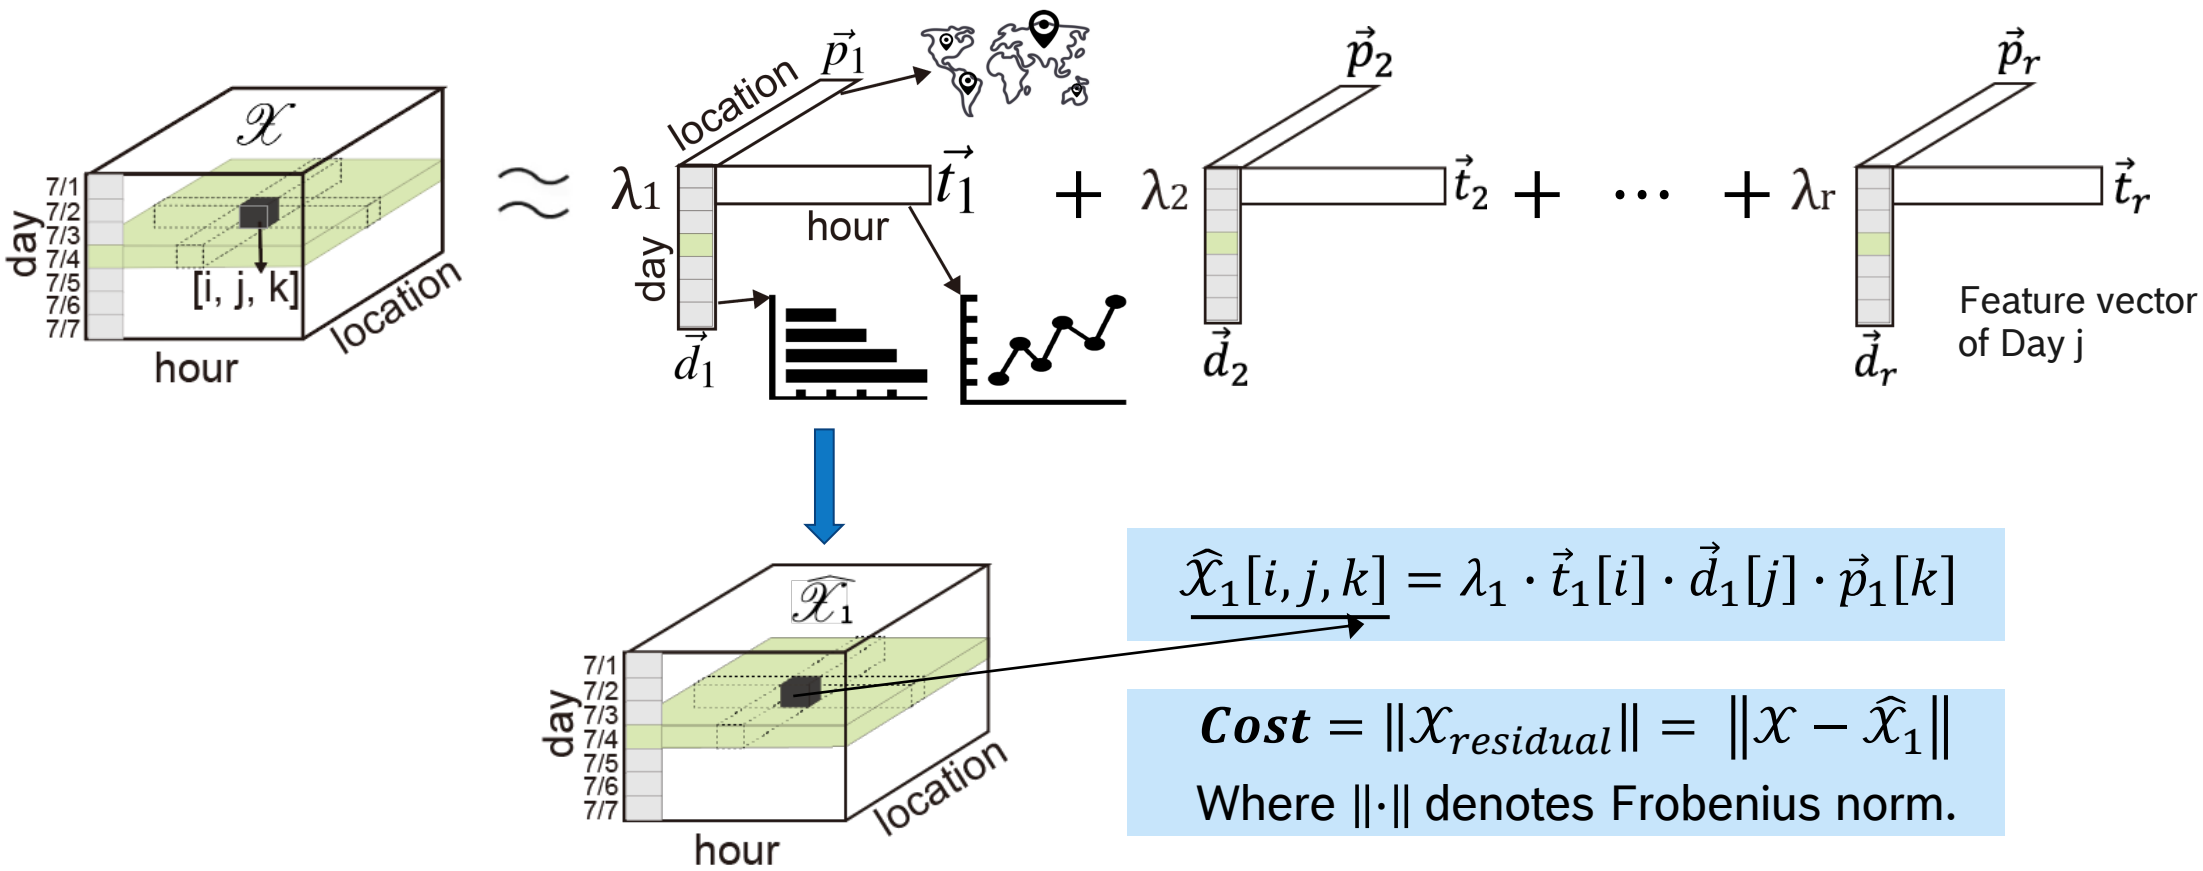

Liu, et al., TPFlow: Progressive Partition and Multidimensional Pattern Extraction for Large-scale Spatio-temporal Data Analysis, TVCG (VAST'18), Best Paper Award.

Semi-automatic and human-steerable tensor partitions

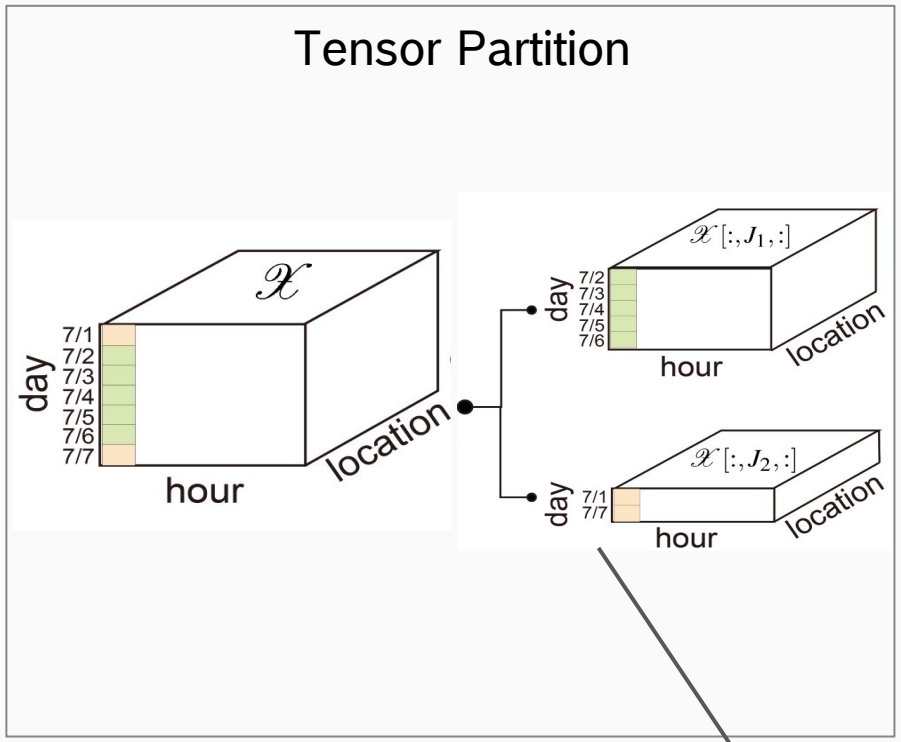

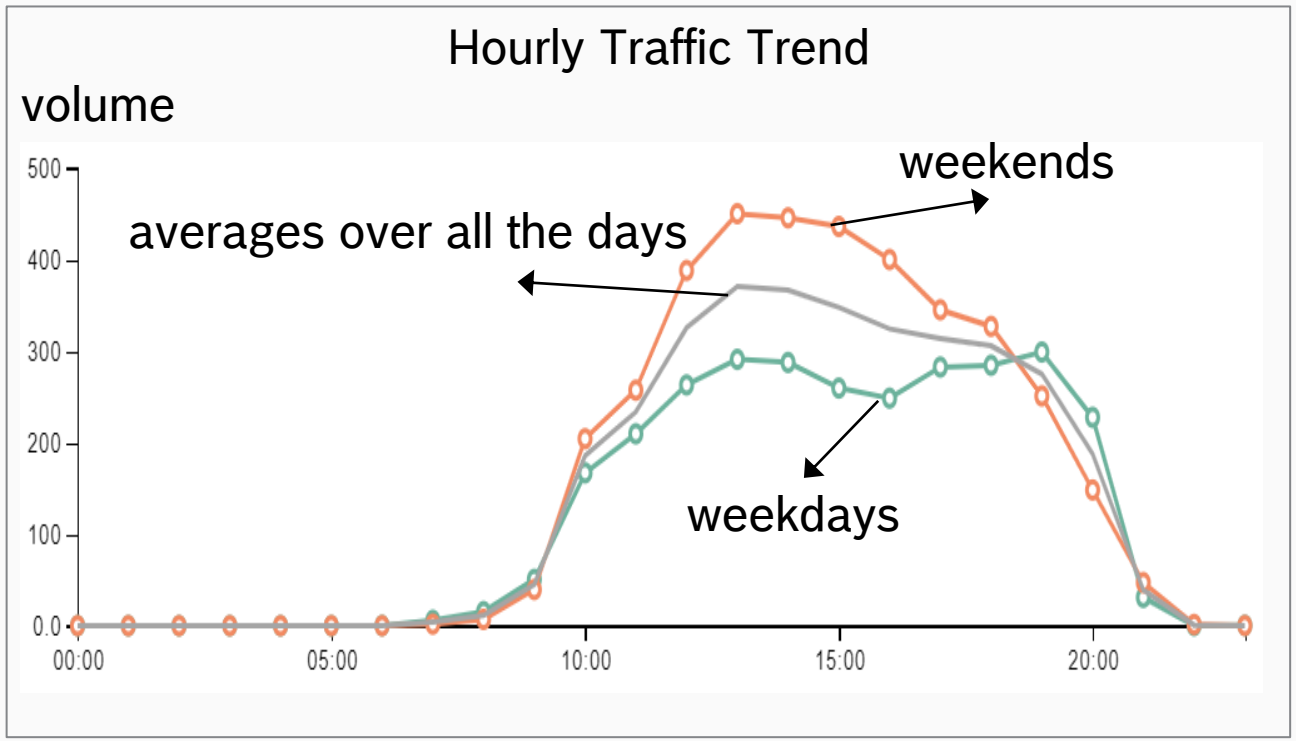

Put the days with **similar hourly and spatial** variations into one sub-tensor

Liu, et al., TPFlow: Progressive Partition and Multidimensional Pattern Extraction for Large-scale Spatio-temporal Data Analysis, TVCG (VAST'18), Best Paper Award.

Working with domain experts
Design requirements

- We collaborated with 4 domain experts

Retail operation optimization

Regional Sales optimization

Transportation Improvement

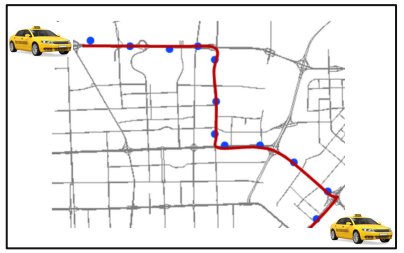

- We have identified 5 design requirements

Liu, et al., TPFlow: Progressive Partition and Multidimensional Pattern Extraction for Large-scale Spatio-temporal Data Analysis, TVCG (VAST'18), Best Paper Award.

In-store customer traffic data analysis

150M events of customer entering/leaving areas in a shopping mall

- Tasks:
  - (1) identify daily/hourly periodical patterns
  - (2) analyze store area performance

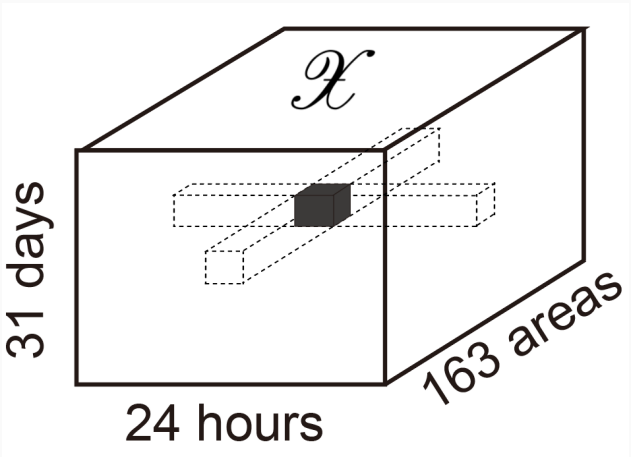

Liu, et al., TPFlow: Progressive Partition and Multidimensional Pattern Extraction for Large-scale Spatio-temporal Data Analysis, TVCG (VAST'18), Best Paper Award.

CTScan powered by CR/RTC-HMI1 Dataset ST Measure total traffic deviation Login | Help

Overview *clustering tree*

Mode *Day*

Daily Pattern

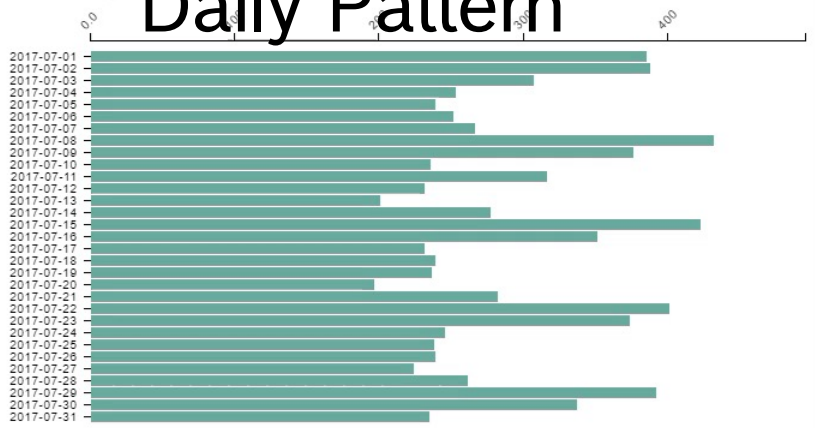

Mode *Hour*

Hourly Pattern

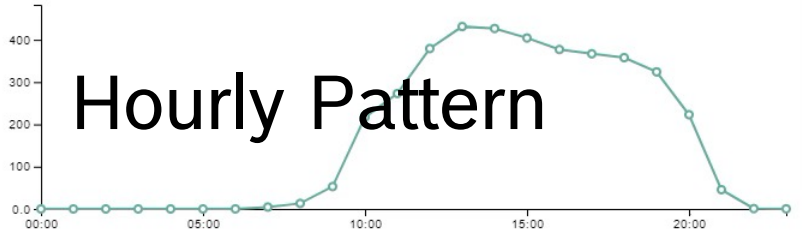

Mode *Zone*

Spatial Pattern

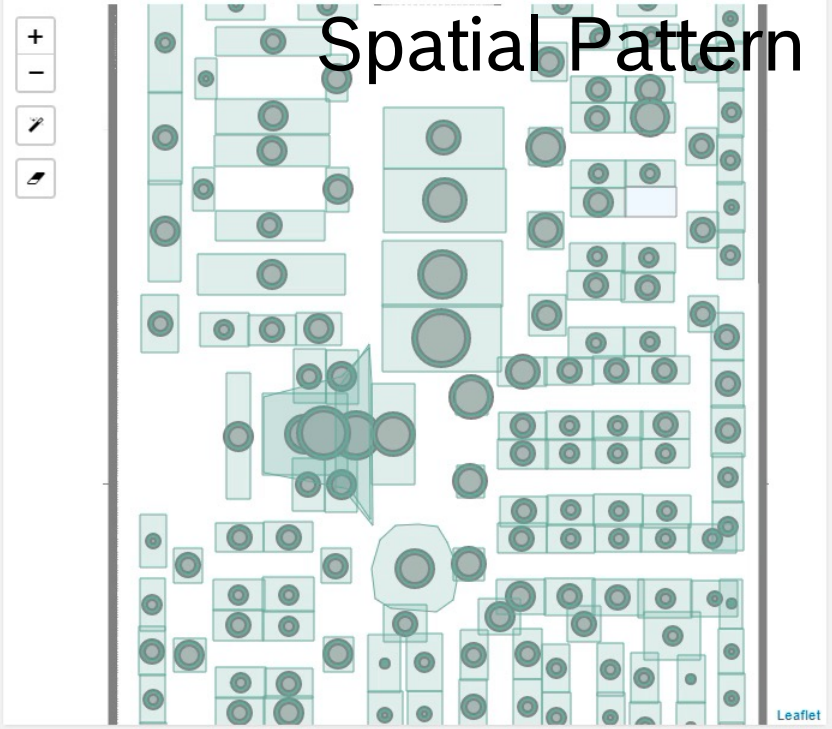

Show 25 entries Search:

#1	2017-07-08	14:00	Device Area	487
#1	2017-07-01	13:00	Device Area	471
#1	2017-07-09	15:00	Device Area	456
#1	2017-07-15	15:00	Device Area	436
#1	2017-07-02	16:00	Device Area	436
#1	2017-07-29	16:00	Device Area	435
#1	2017-07-02	13:00	Device Area	433
#1	2017-07-23	12:00	Device Area	433
#1	2017-07-01	14:00	Device Area	429
#1	2017-07-02	14:00	Device Area	425
#1	2017-07-15	13:00	904B	425
#1	2017-07-16	13:00	Device Area	421
#1	2017-07-09	14:00	Device Area	420
#1	2017-07-22	15:00	Device Area	418
#1	2017-07-02	17:00	Device Area	416
#1	2017-07-09	17:00	Device Area	416
#1	2017-07-15	13:00	Device Area	410
#1	2017-07-08	14:00	904B	410
#1	2017-07-09	16:00	Device Area	408
#1	2017-07-22	16:00	Device Area	402
#1	2017-07-29	15:00	Device Area	400
#1	2017-07-30	14:00	Device Area	400
#1	2017-07-23	13:00	904B	398
#1	2017-07-09	13:00	Device Area	397
#1	2017-07-30	13:00	Device Area	397

Showing 1 to 25 of 100 entries

Previous 1 2 3 4 Next

CTScan powered by CR/RTC-HMI1 Dataset ST Measure total traffic deviation Login | Help

Overview *clustering tree*

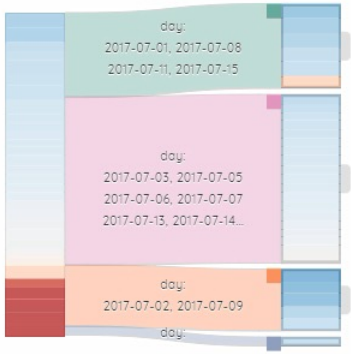

User decides which dimension to partition or choose to be fully automatic

Mode *Day*

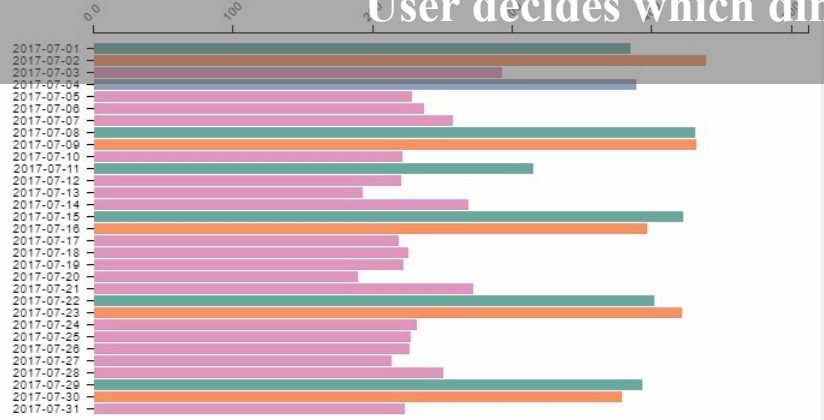

Mode *Hour*

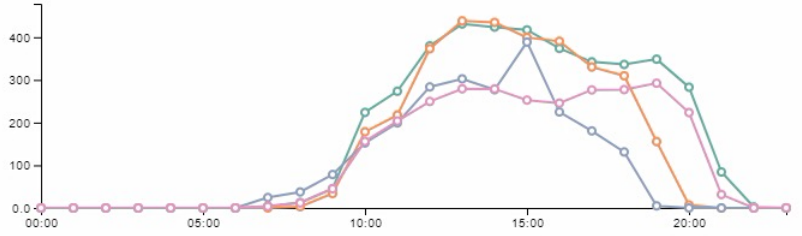

Mode *Zone*

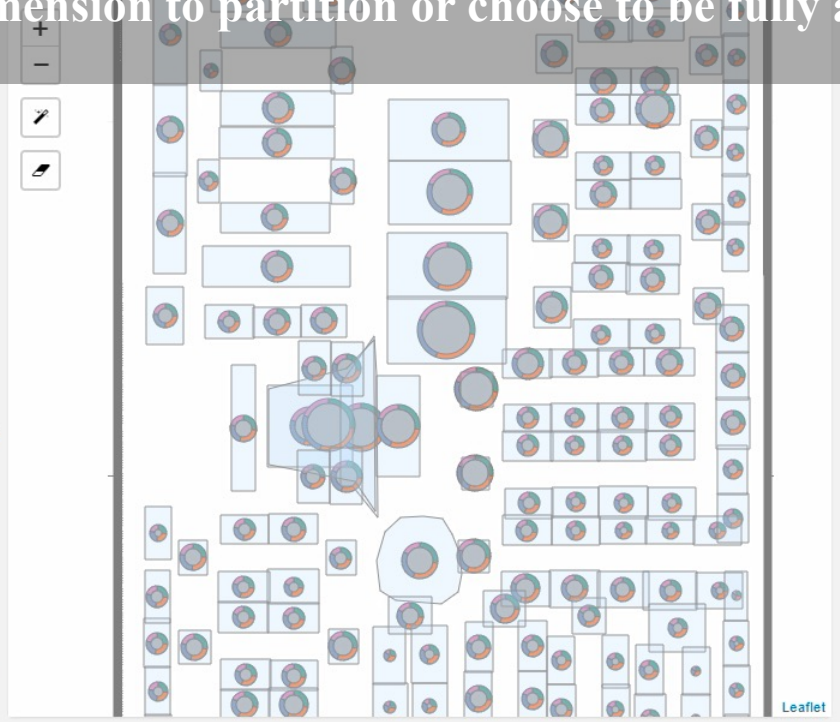

Show 25 entries Search:

#1	2017-07-08	14:00	Device Area	487
#1	2017-07-01	13:00	Device Area	471
#2	2017-07-09	15:00	Device Area	456
#1	2017-07-15	15:00	Device Area	436
#2	2017-07-02	16:00	Device Area	436
#1	2017-07-29	16:00	Device Area	435
#2	2017-07-23	12:00	Device Area	433
#2	2017-07-02	13:00	Device Area	433
#1	2017-07-01	14:00	Device Area	429
#1	2017-07-15	13:00	904B	425
#2	2017-07-02	14:00	Device Area	425
#2	2017-07-16	13:00	Device Area	421
#2	2017-07-09	14:00	Device Area	420
#1	2017-07-22	15:00	Device Area	418
#2	2017-07-02	17:00	Device Area	416
#2	2017-07-09	17:00	Device Area	416
#1	2017-07-08	14:00	904B	410
#1	2017-07-15	13:00	Device Area	410
#2	2017-07-09	16:00	Device Area	408
#1	2017-07-22	16:00	Device Area	402
#1	2017-07-29	15:00	Device Area	400
#2	2017-07-30	14:00	Device Area	400
#2	2017-07-23	13:00	904B	398
#2	2017-07-09	13:00	Device Area	397
#2	2017-07-30	13:00	Device Area	397

Showing 1 to 25 of 400 entries

Previous 1 2 3 4 5 ... 16 Next

Overview *clustering tree*

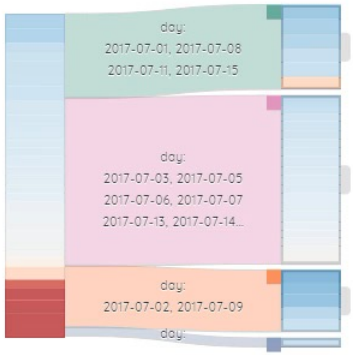

Mode *Day*

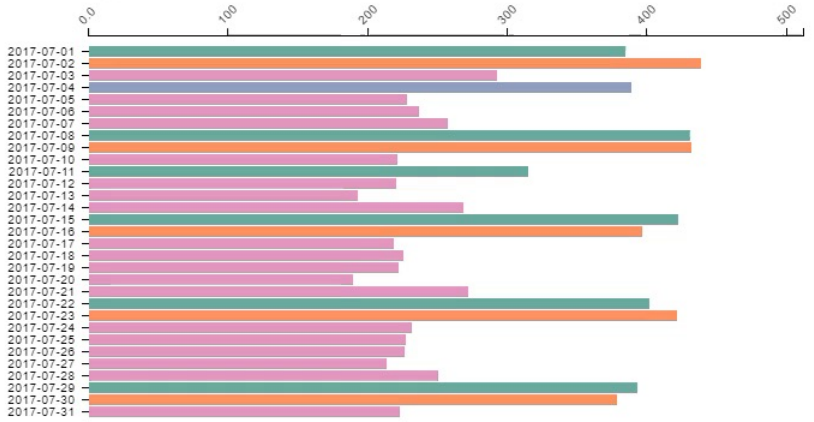

Mode *Hour*

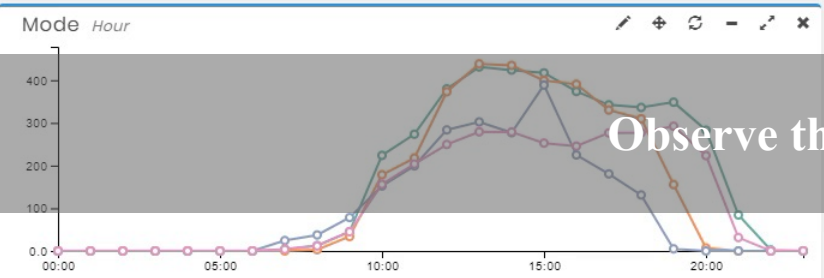

Mode *Zone*

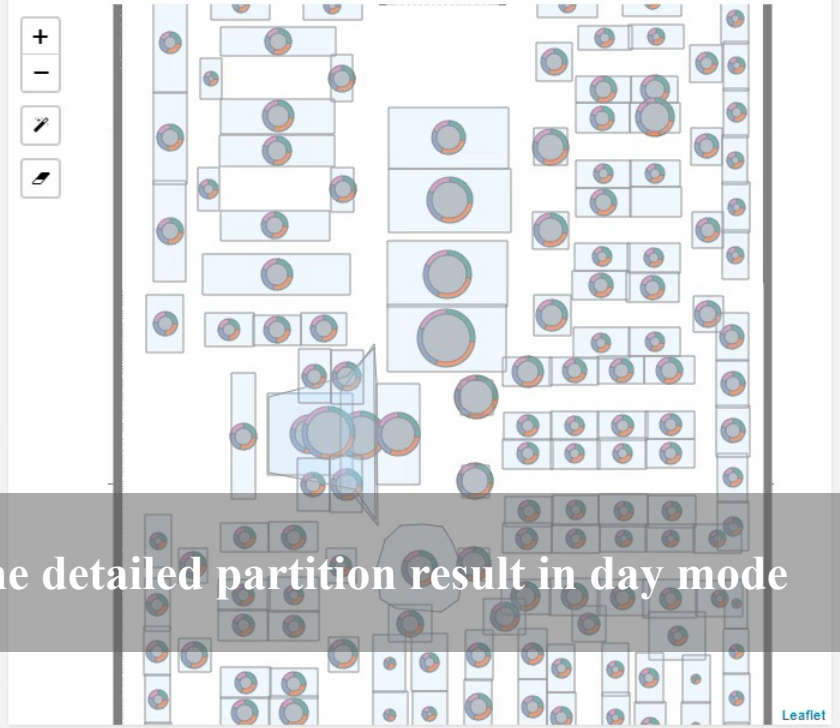

Show 25 entries Search:

#1	2017-07-08	14:00	Device Area	487
#1	2017-07-01	13:00	Device Area	471
#2	2017-07-09	15:00	Device Area	456
#1	2017-07-15	15:00	Device Area	436
#2	2017-07-02	16:00	Device Area	436
#1	2017-07-29	16:00	Device Area	435
#2	2017-07-23	12:00	Device Area	433
#2	2017-07-02	13:00	Device Area	433
#1	2017-07-01	14:00	Device Area	429
#1	2017-07-15	13:00	904B	425
#2	2017-07-02	14:00	Device Area	425
#2	2017-07-16	13:00	Device Area	421
#2	2017-07-09	14:00	Device Area	420
#1	2017-07-22	15:00	Device Area	418
#2	2017-07-02	17:00	Device Area	416
#2	2017-07-09	17:00	Device Area	416
#1	2017-07-08	14:00	904B	410
#1	2017-07-15	13:00	Device Area	410
#2	2017-07-09	16:00	Device Area	408
#1	2017-07-22	16:00	Device Area	402
#1	2017-07-29	15:00	Device Area	400
#2	2017-07-30	14:00	Device Area	400
#2	2017-07-23	13:00	904B	398
#2	2017-07-09	13:00	Device Area	397
#2	2017-07-30	13:00	Device Area	397

Showing 1 to 25 of 400 entries

Previous 1 2 3 4 5 ... 16 Next

Observe the detailed partition result in day mode

Overview *clustering tree*

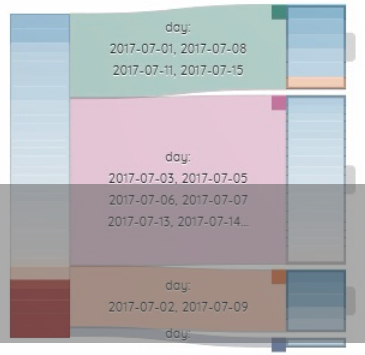

Hour mode shows the hourly trend patterns for each day group

Mode *Day*

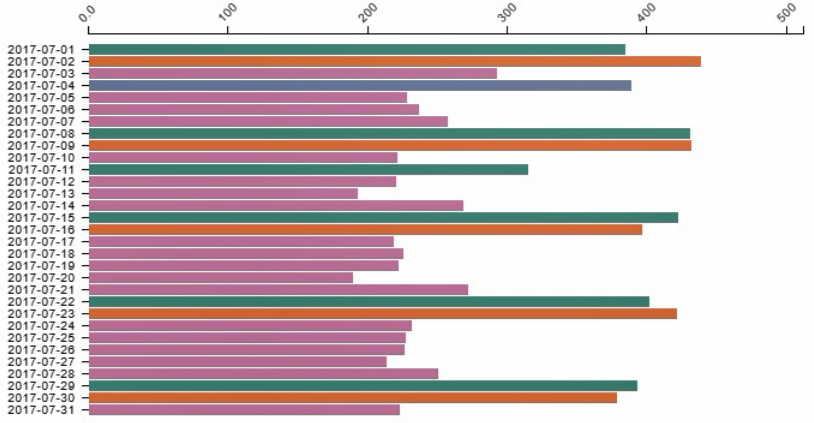

Mode *Hour*

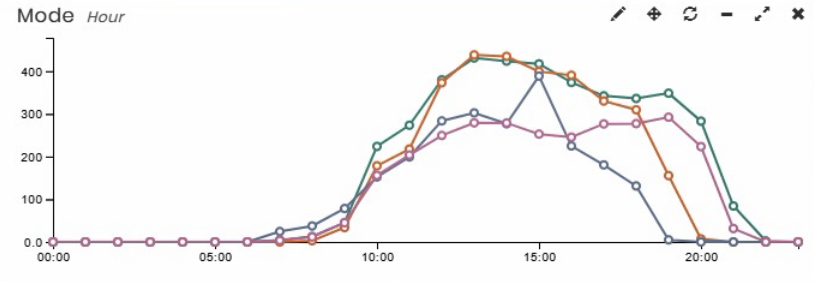

Mode *Zone*

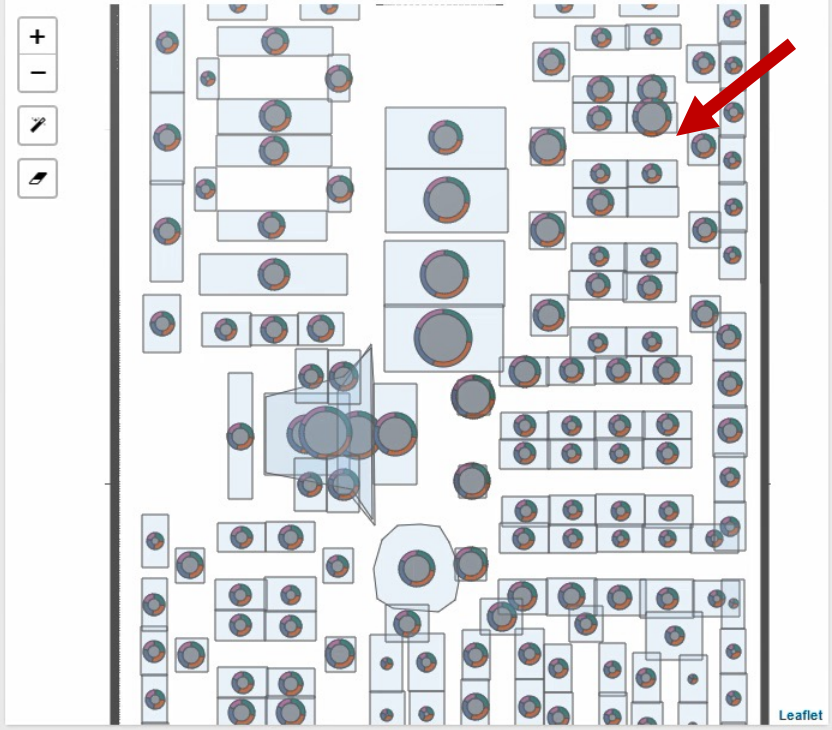

Show 25 entries Search:

#1	2017-07-08	14:00	Device Area	487
#1	2017-07-01	13:00	Device Area	471
#2	2017-07-09	15:00	Device Area	456
#1	2017-07-15	15:00	Device Area	436
#2	2017-07-02	16:00	Device Area	436
#1	2017-07-29	16:00	Device Area	435
#2	2017-07-23	12:00	Device Area	433
#2	2017-07-02	13:00	Device Area	433
#1	2017-07-01	14:00	Device Area	429
#1	2017-07-15	13:00	904B	425
#2	2017-07-02	14:00	Device Area	425
#2	2017-07-16	13:00	Device Area	421
#2	2017-07-09	14:00	Device Area	420
#1	2017-07-22	15:00	Device Area	418
#2	2017-07-02	17:00	Device Area	416
#2	2017-07-09	17:00	Device Area	416
#1	2017-07-08	14:00	904B	410
#1	2017-07-15	13:00	Device Area	410
#2	2017-07-09	16:00	Device Area	408
#1	2017-07-22	16:00	Device Area	402
#1	2017-07-29	15:00	Device Area	400
#2	2017-07-30	14:00	Device Area	400
#2	2017-07-23	13:00	904B	398
#2	2017-07-09	13:00	Device Area	397
#2	2017-07-30	13:00	Device Area	397

Showing 1 to 25 of 400 entries

Previous 1 2 3 4 5 ... 16 Next

CTScan powered by CR/RTC-HMI1 Dataset ST Measure total traffic deviation Login | Help

Overview *clustering tree*

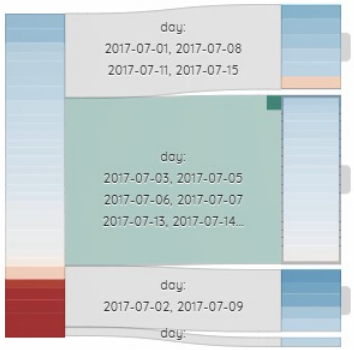

Mode *Day*

Mode *Hour*

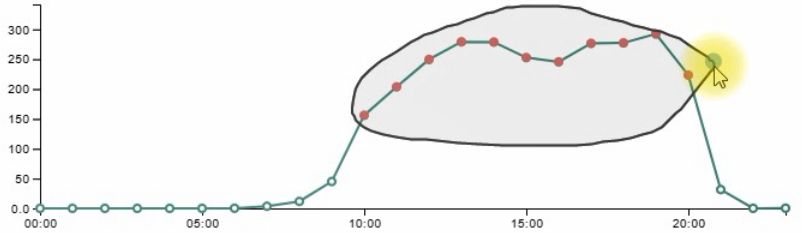

Mode *Zone*

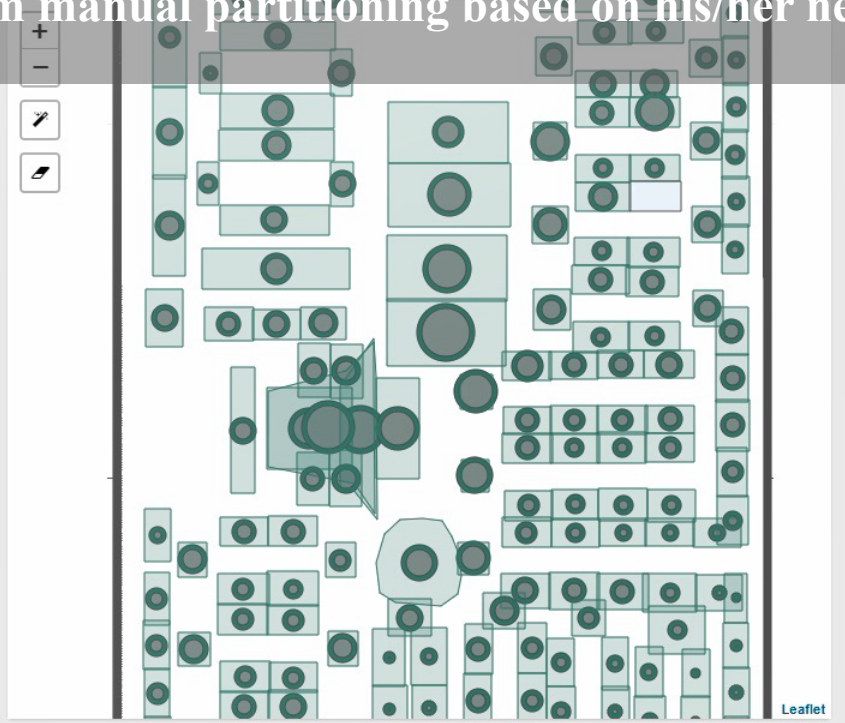

User can perform manual partitioning based on his/her needs

Show 25 entries Search:

#1	2017-07-07	18:00	Device Area	315
#1	2017-07-03	19:00	Device Area	315
#1	2017-07-07	19:00	Device Area	311
#1	2017-07-21	19:00	904B	305
#1	2017-07-03	13:00	Device Area	297
#1	2017-07-14	19:00	Device Area	294
#1	2017-07-14	17:00	Device Area	293
#1	2017-07-28	14:00	Device Area	293
#1	2017-07-21	17:00	Device Area	288
#1	2017-07-07	17:00	Device Area	287
#1	2017-07-06	13:00	Device Area	284
#1	2017-07-21	19:00	Device Area	284
#1	2017-07-03	19:00	904B	283
#1	2017-07-06	15:00	Device Area	282
#1	2017-07-07	18:00	904B	282
#1	2017-07-19	14:00	Device Area	281
#1	2017-07-14	16:00	904B	281
#1	2017-07-07	19:00	904B	280
#1	2017-07-03	15:00	Device Area	277
#1	2017-07-18	17:00	Device Area	276
#1	2017-07-03	12:00	Device Area	275
#1	2017-07-14	16:00	Device Area	275
#1	2017-07-06	18:00	Device Area	274
#1	2017-07-14	20:00	904B	274
#1	2017-07-21	16:00	Device Area	274

Showing 1 to 25 of 100 entries

Previous 1 2 3 4 Next

CTScan powered by CR/RTC-HMI1 | Dataset ST | Measure total traffic | deviation | Login | Help

Overview *clustering tree*

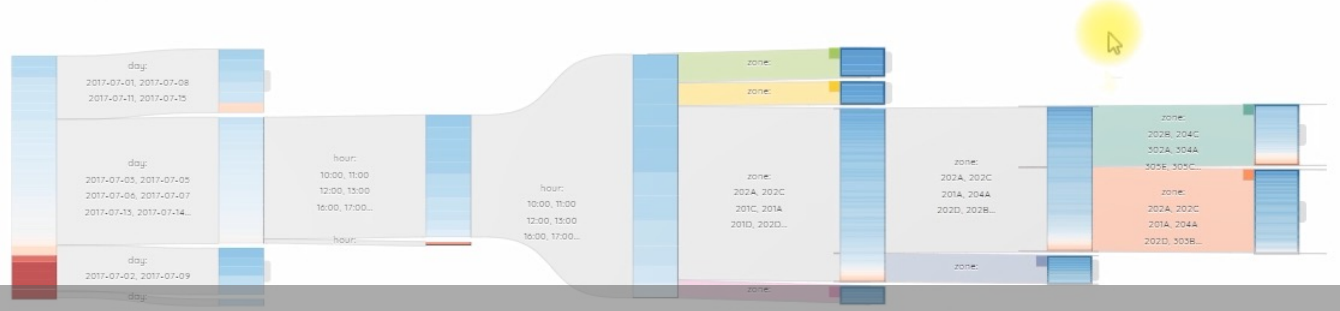

On weekdays from 10am to 8pm, can you tell me the performance of different store areas?

Mode *Day*

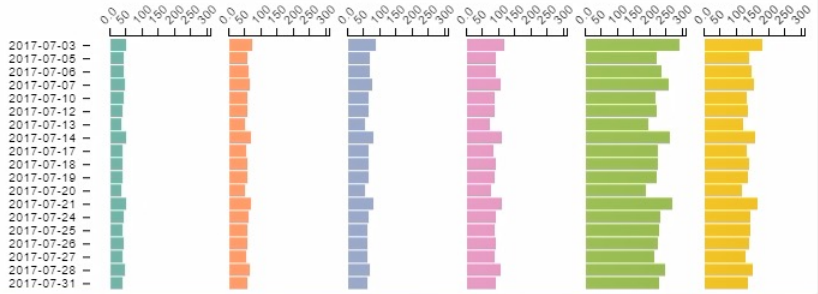

Mode *Hour*

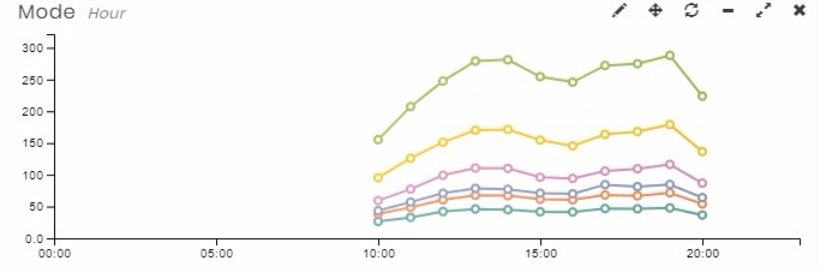

Mode *Zone*

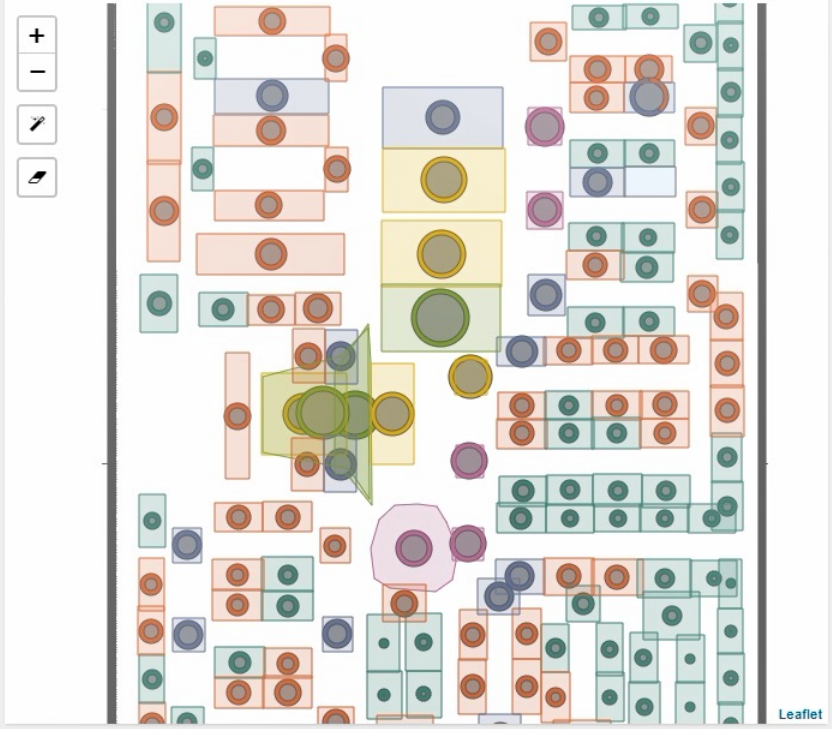

Show 25 entries Search:

#5	2017-07-03	19:00	Device Area	315
#5	2017-07-07	18:00	Device Area	315
#5	2017-07-07	19:00	Device Area	311
#5	2017-07-21	19:00	904B	305
#5	2017-07-03	13:00	Device Area	297
#5	2017-07-14	19:00	Device Area	294
#5	2017-07-28	14:00	Device Area	293
#5	2017-07-14	17:00	Device Area	293
#5	2017-07-07	20:00	Device Area	288
#5	2017-07-07	17:00	Device Area	287
#5	2017-07-21	19:00	Device Area	284
#5	2017-07-06	13:00	Device Area	284
#5	2017-07-03	19:00	904B	283
#5	2017-07-07	18:00	904B	282
#5	2017-07-06	15:00	Device Area	282
#5	2017-07-19	14:00	Device Area	281
#5	2017-07-14	16:00	904B	281
#5	2017-07-07	19:00	904B	280
#5	2017-07-03	15:00	Device Area	277
#5	2017-07-18	17:00	Device Area	276
#5	2017-07-14	16:00	Device Area	275
#5	2017-07-03	12:00	Device Area	275
#5	2017-07-14	20:00	904B	274
#5	2017-07-06	18:00	Device Area	274
#5	2017-07-21	16:00	Device Area	274

Showing 1 to 25 of 600 entries

Previous 1 2 3 4 5 ... 24 Next

CTScan powered by CR/RTC-HMI1 | Dataset ST | Measure total traffic | deviation | Login | Help

Overview *clustering tree*

Mode *Day*

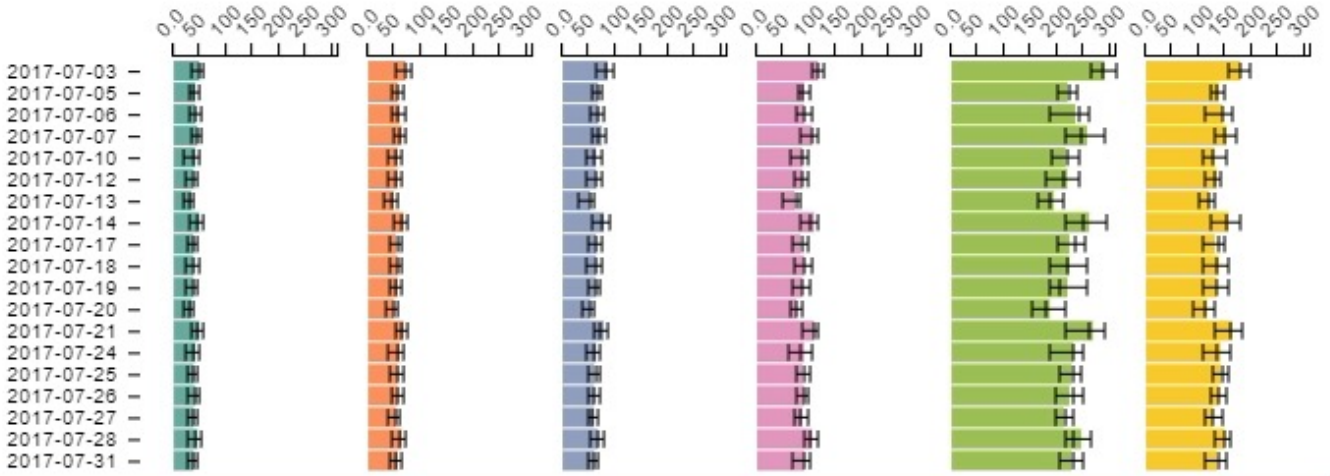

Mode *Hour*

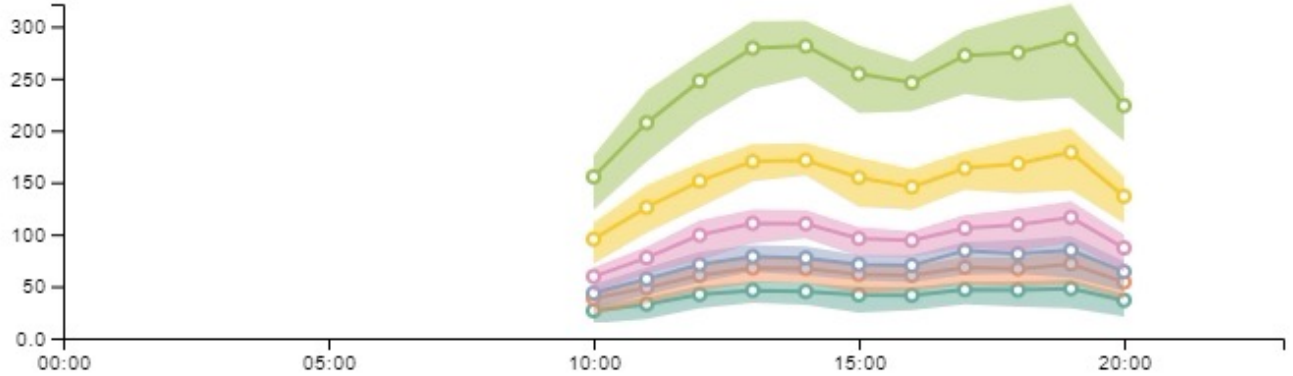

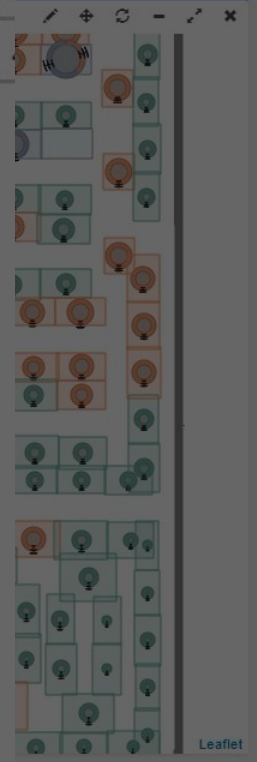

Show 25 entries Search:

#5	2017-07-03	19:00	Device Area	315
#5	2017-07-07	18:00	Device Area	315
#5	2017-07-07	19:00	Device Area	311
#5	2017-07-21	19:00	904B	305
#5	2017-07-03	13:00	Device Area	297
#5	2017-07-14	19:00	Device Area	294
#5	2017-07-28	14:00	Device Area	293
#5	2017-07-14	17:00	Device Area	293
#5	2017-07-21	17:00	Device Area	288
#5	2017-07-07	17:00	Device Area	287
#5	2017-07-21	19:00	Device Area	284
#5	2017-07-06	13:00	Device Area	284
#5	2017-07-03	19:00	904B	283
#5	2017-07-07	18:00	904B	282
#5	2017-07-06	15:00	Device Area	282
#5	2017-07-19	14:00	Device Area	281
#5	2017-07-14	16:00	904B	281
#5	2017-07-07	19:00	904B	280
#5	2017-07-03	15:00	Device Area	277
#5	2017-07-18	17:00	Device Area	276
#5	2017-07-14	16:00	Device Area	275
#5	2017-07-03	12:00	Device Area	275
#5	2017-07-14	20:00	904B	274
#5	2017-07-06	18:00	Device Area	274
#5	2017-07-21	16:00	Device Area	274

Showing 1 to 25 of 600 entries

Previous 1 2 3 4 5 … 24 Next

Time & Space complexity analysis

▶ Time & Space Complexity

Tensor shape	Time Complexity	Space Complexity
2D (matrix)	$O(n^2)$	$O(n^2)$
3D (cube)	$O(n^3)$	$O(n^3)$
…	…	…
p dimensional tensor	$O(n^p)$	$O(n^p)$

▶ In practice, the algorithm can return results within reasonable time

Dataset	Record #	Tensor Shape	Data Point #	Time
regional sales	$>$ 2 million	$24 \times 34 \times 16$	**~13,000**	$< 0.1s$
customer in-store traffic	$>$ 150 million	$186 \times 24 \times 163$	**~700,000**	$< 1s$
New York taxi trip	$>$ 10 million	$31 \times 24 \times 67 \times 67$	**~3,500,000**	$< 10s$

Liu, et al., TPFlow: Progressive Partition and Multidimensional Pattern Extraction for Large-scale Spatio-temporal Data Analysis, TVCG (VAST'18), Best Paper Award.

Comparative analysis

- Experiments (a) on **synthetic datasets** where we know the best way to partition the data.

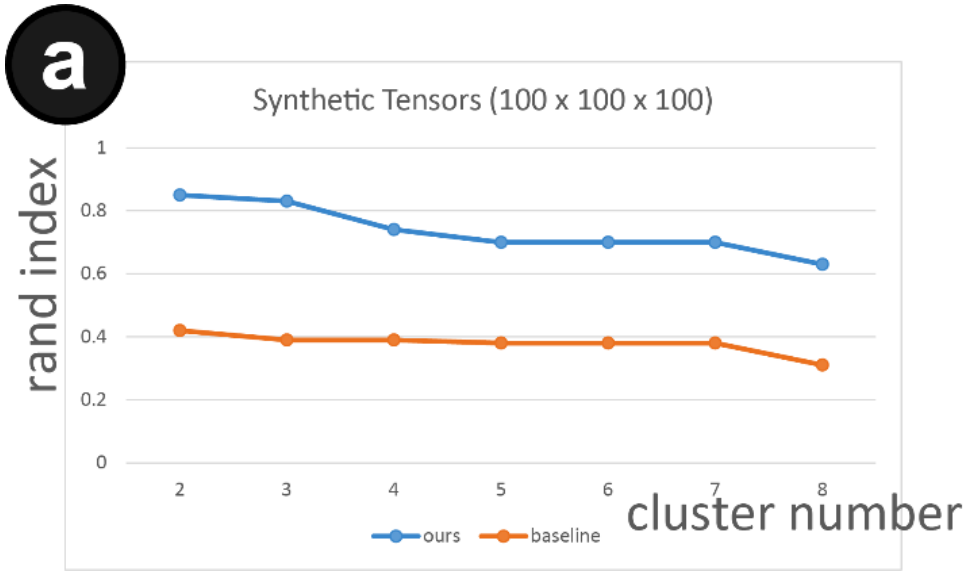

- Adjusted Rand Index (ARI): **↑ is better**
- Our algorithm (blue) **aligns better** with the ground truth

Liu, et al., TPFlow: Progressive Partition and Multidimensional Pattern Extraction for Large-scale Spatio-temporal Data Analysis, TVCG (VAST'18), Best Paper Award.

Comparative analysis

- Experiments (b) on **real-world datasets** from the three use cases (**e.g., regional sales**).

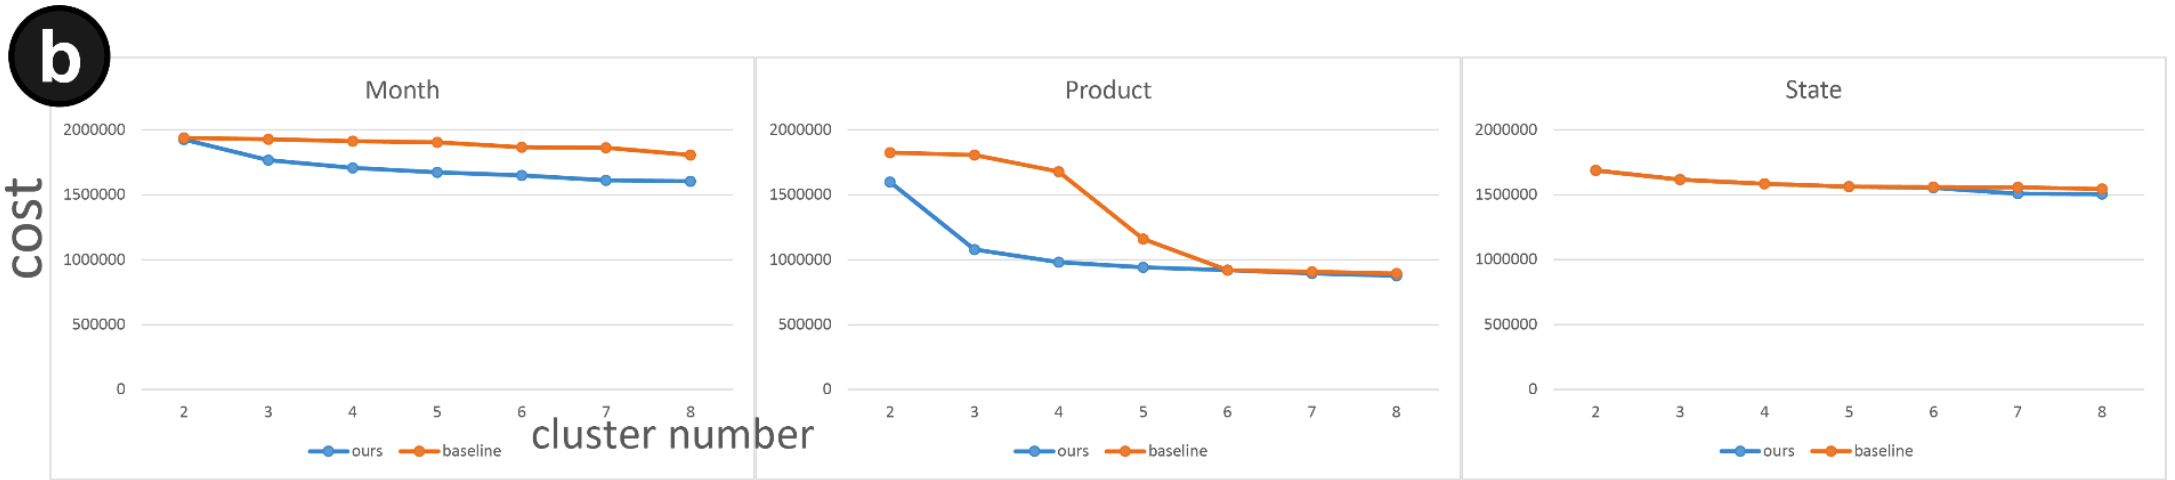

Tensor: 24 months×34 products×16 regions

- Cost function: ↓ **is better**
- Our algorithm (blue) **better optimize** the cost function

Liu, et al., TPFlow: Progressive Partition and Multidimensional Pattern Extraction for Large-scale Spatio-temporal Data Analysis, TVCG (VAST'18), Best Paper Award.

Human-AI teaming for spatio-temporal pattern extraction

Steerable

AI System
(Human-Centered)

May not align with
users' analytical needs

Visual
Interface

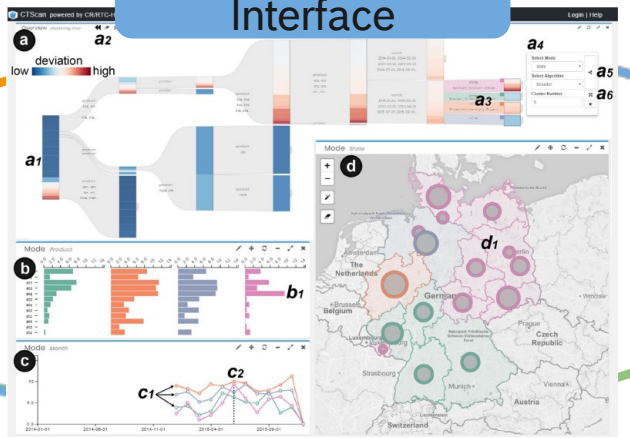

Human

Cannot exhaust every data subsets
for meaningful pattern seeking

Liu, et al., TPFlow: Progressive Partition and Multidimensional Pattern Extraction for Large-scale Spatio-temporal Data Analysis, TVCG (VAST'18), Best Paper Award.

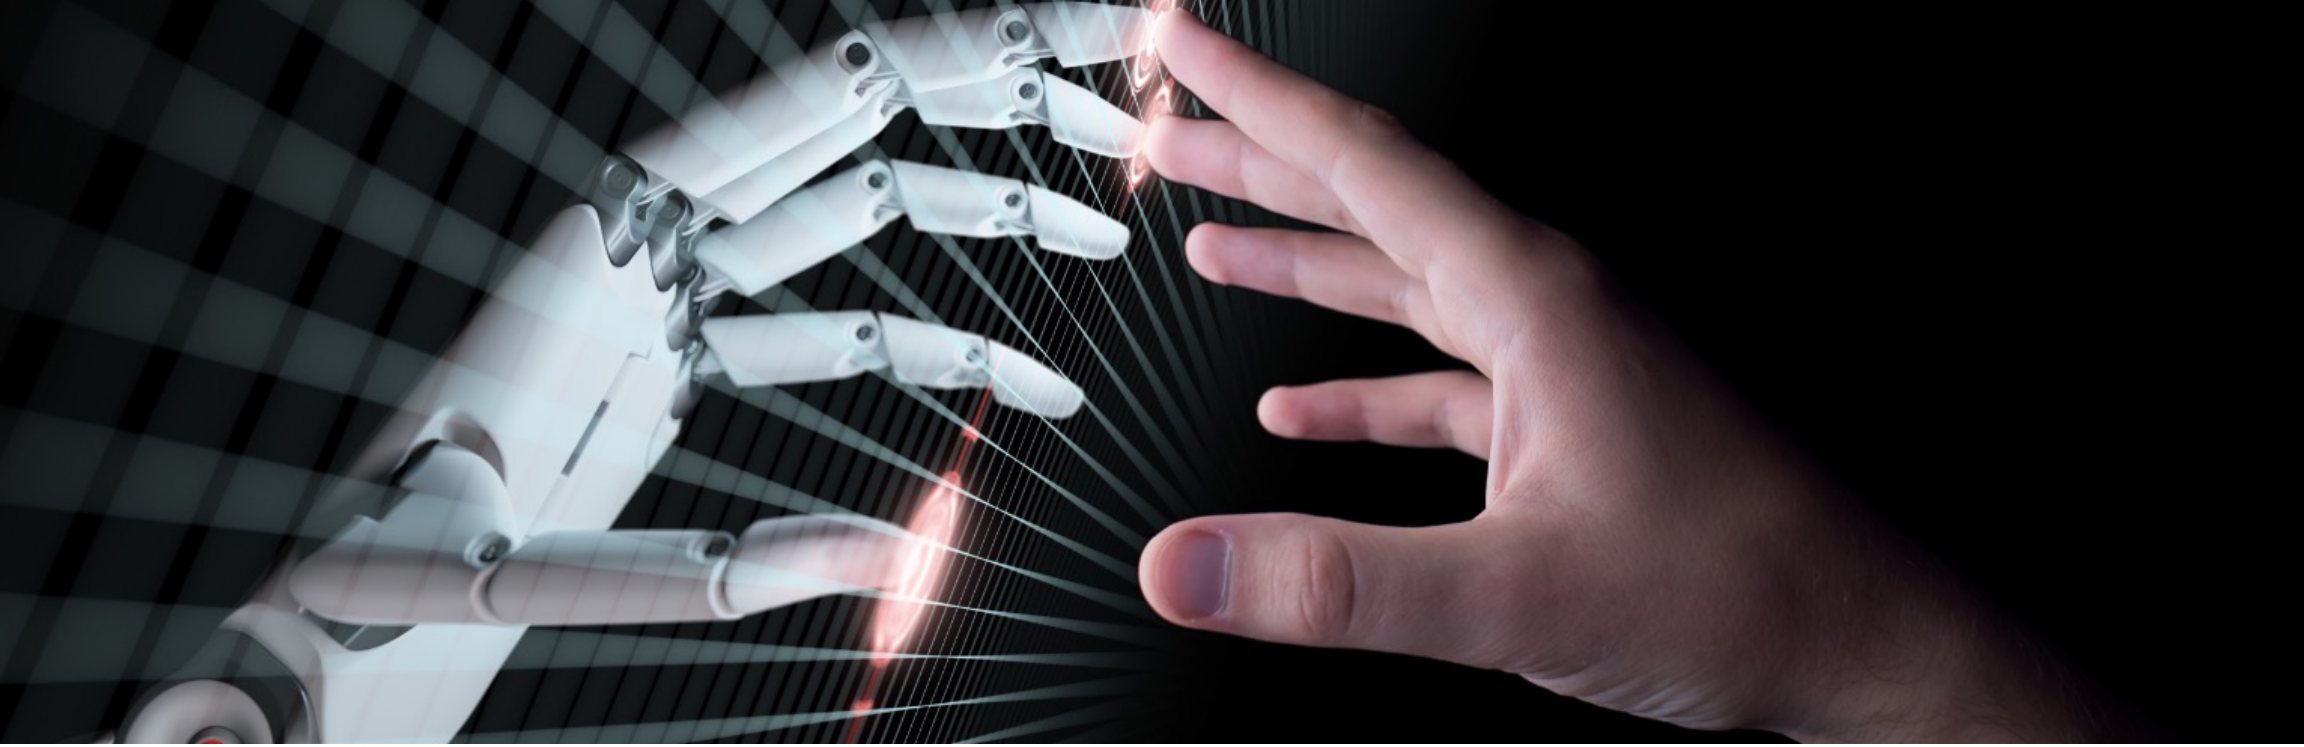

To enable **general solutions** to develop **Human-AI teaming systems** that are not only accurate and efficient, but also accessible, understandable, and acceptable to users, in order to enhance data-driven decision-making in formally intractable real-world problems.

Future directions

- Enhance human-AI teaming experience utilizing more data from and about humans

To enable general solutions to develop **Human-AI teaming systems** that are not only accurate and efficient, but also accessible, understandable, and acceptable to users, in order to enhance data-driven decision-making in formally intractable real-world problems.

Enhance human-AI teaming experience

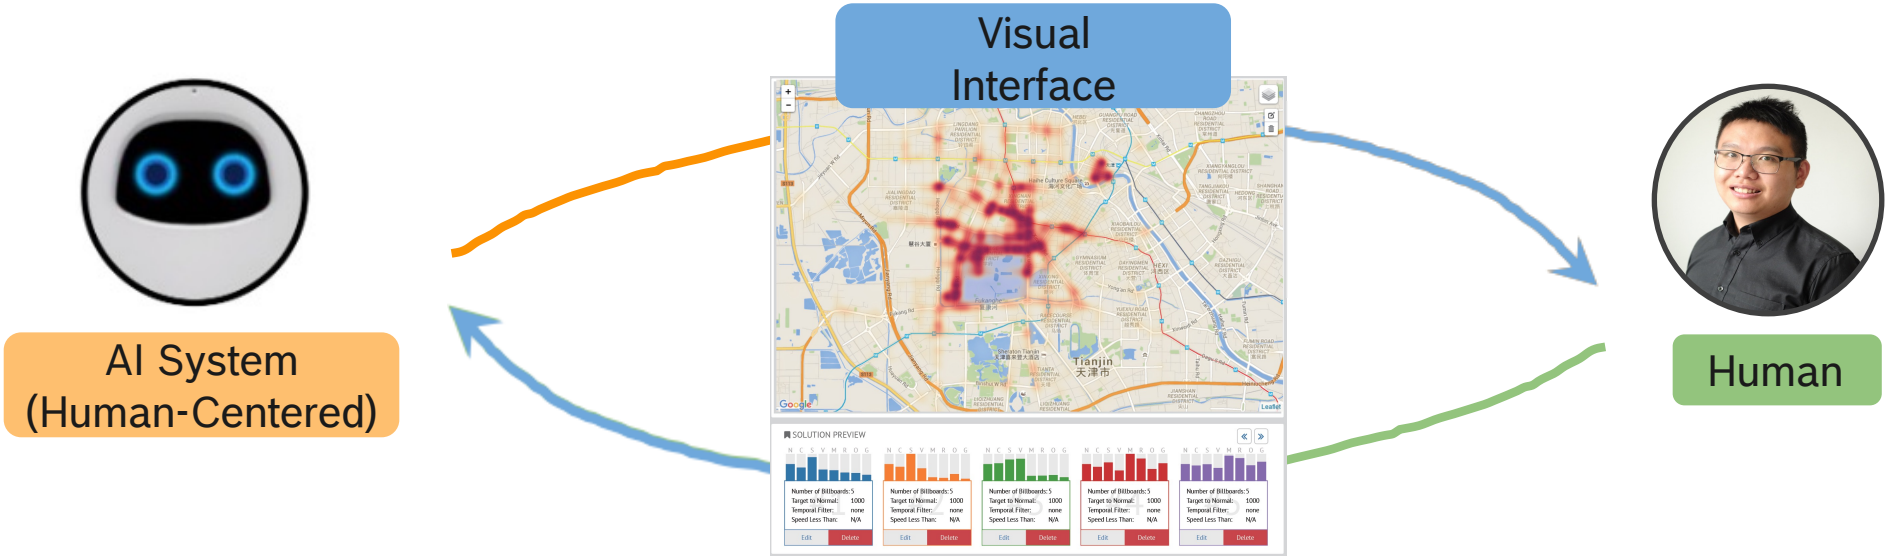

Enhance human-AI teaming experience

AI System
(Human-Centered)

- Understand human intentions/interactions and adapt to context changes

emotion data

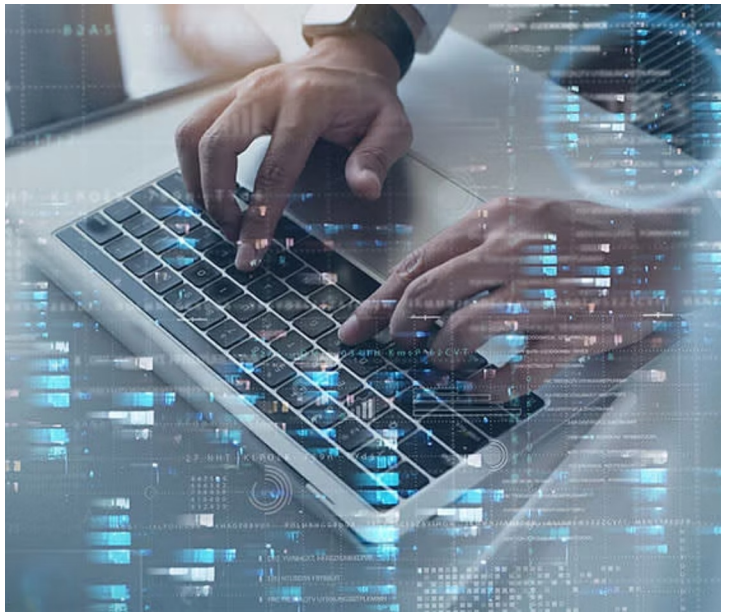

interaction logs

Enhance human-AI teaming experience

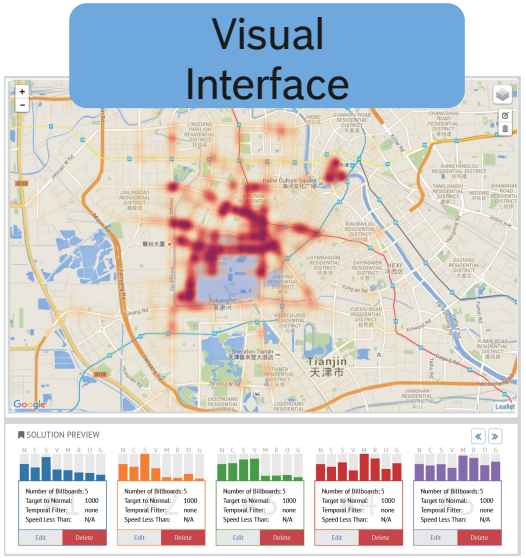

- Intelligent visualization interaction

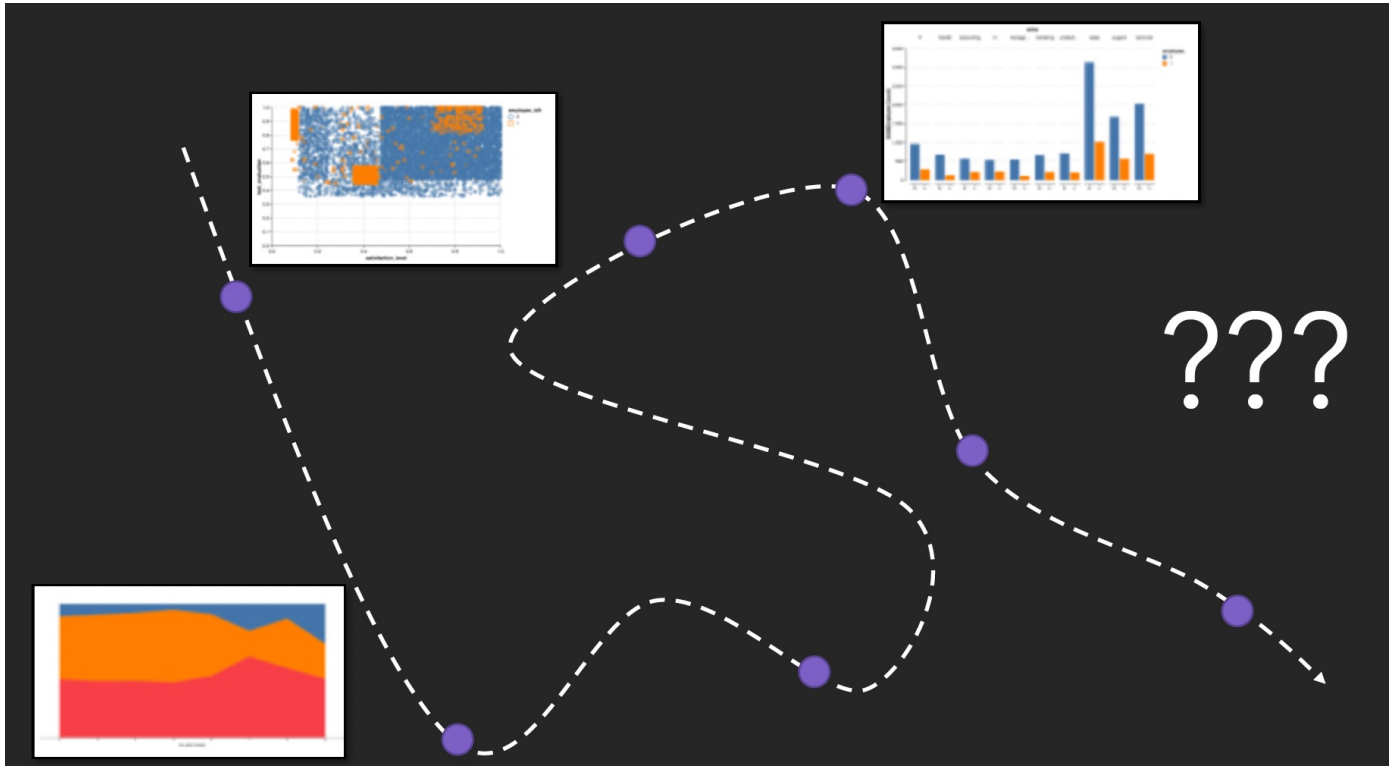

Enhance human-AI teaming experience

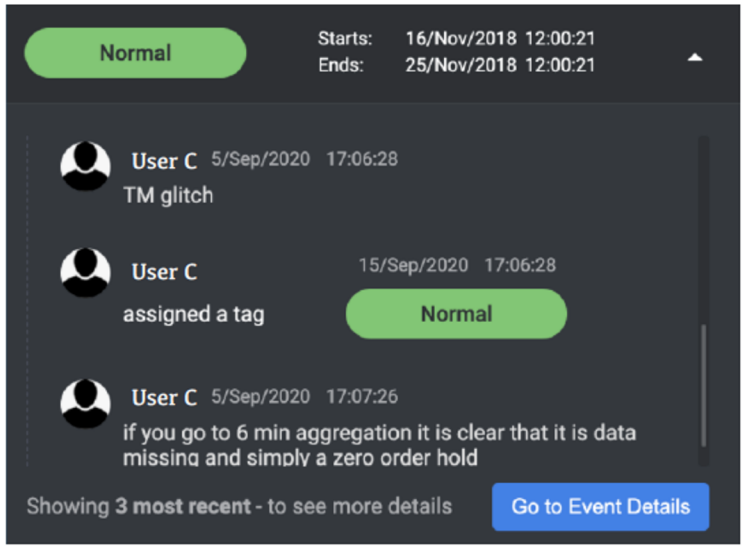

Human

- Optimize human cognitive load
- Actionable decisions

Enhance human-AI teaming experience

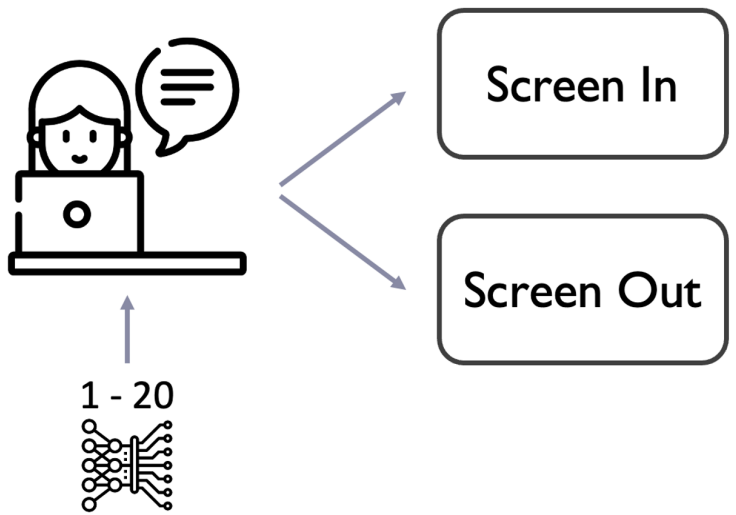

Human

Confusing Target

"Wish we knew how we got to do with the score"

- Optimize human cognitive load
- Actionable decisions

Future directions

- Enhance human-AI teaming experience with more data from and about humans
- Go beyond “point solutions” using open source as the substrate

To enable **general solutions** to develop Human-AI teaming systems that are not only accurate and efficient, but also accessible, understandable, and acceptable to users, in order to enhance data-driven decision-making in formally intractable real-world problems.

Go beyond “point solutions”

Orion

A human-centered machine learning library for detecting anomalies in signals.

Python | pypi v0.4.1 | downloads 58k | launch binder

Cardea

An open source automl library for using machine learning in healthcare.

Python | pypi v0.1.2 | downloads 33k | launch binder

Pyreal

A Human-Centered ML Explanation Framework.

Python | downloads 8k

Zephyr

AutoML for renewable energy industries.

machine-learning | wind-energy | automl

Python | pypi v0.0.0 | downloads 787

Draco

A machine learning library for time series classfication.

Python | downloads 8k | launch binder

Sintel

Sintel (Signal Intelligence) provides Restful APIs to perform massive signal data analysis with human-in-the-loop.

Python | pypi v0.1.0.dev0 | downloads 2k

VBridge

Visualization for Explainable Healthcare Models

Python | TypeScript | React

MTV

A visual analytics system for detecting, investigating, and annotating anomalies in multivariate time series

visualization | anomaly-detection | data-analysis

TypeScript | React | Redux

Go beyond “point solutions”

Improve HITL anomaly detection

https://github.com/sintel-dev/Orion

- How can I better use annotations?

Go beyond “point solutions”

Extend human-centered XAI framework

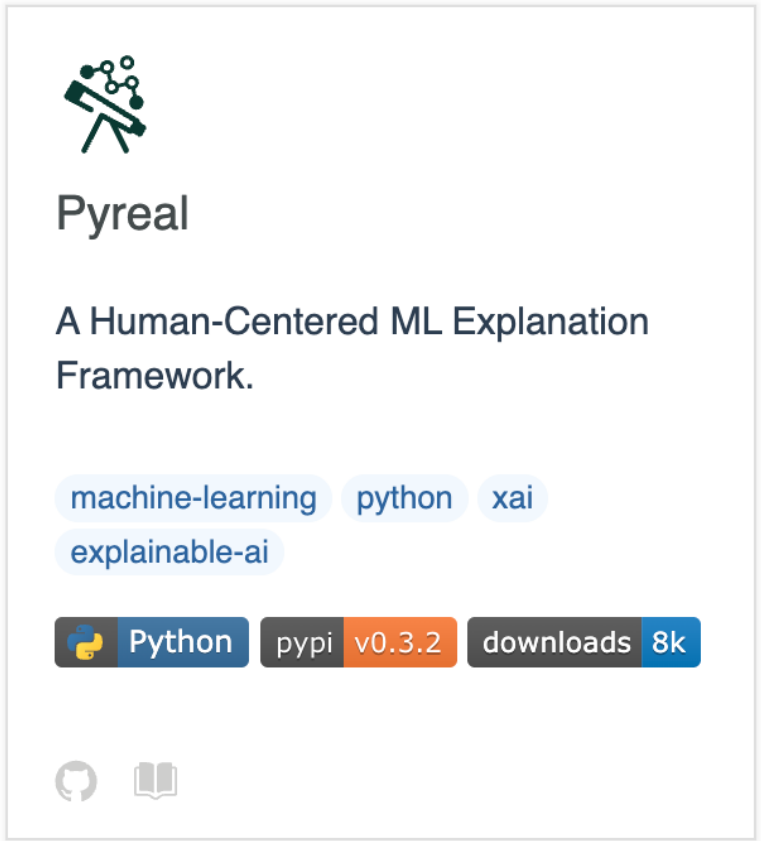

- How can we extend it to support time series data?

https://github.com/sibyl-dev/pyreal

Future directions

- Enhance human-AI teaming experience with more data from and about humans
- Go beyond “point solutions” using open source as the substrate
- Catalyze solutions to critical domains by making use of multimodal data

---

To enable general solutions to develop Human-AI teaming systems that are not only accurate and efficient, but also accessible, understandable, and acceptable to users, in order to enhance data-driven decision-making in formally intractable real-world problems.

Catalyze solutions to critical domains

Social Good (healthcare, child welfare)

- Fragile families challenges
- Early detection of Cognitive Impairment (Dementia)

Sustainability

- Smart curtailment of wind turbines for bird collision mitigation
- Understand and mitigate climate change impacts on migratory birds

Catalyze solutions to critical domains

Mitigating bird collisions with smart curtailment

Catalyze solutions to critical domains

Mitigating bird collisions with smart curtailment

Collect data at time t

Wind 27 km/h
- speed
- direction

Turbines
- layout
- size
- power curve
- RPM

Birds
- species
- flock size
- current location
- prior trajectory

Predict trajectory in next Δt

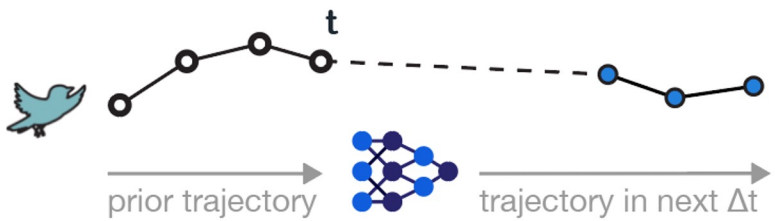

Optimize curtailment strategy

- Wind
- Turbines
- Birds + Pred
- Requirements

mortality, supply need, etc.

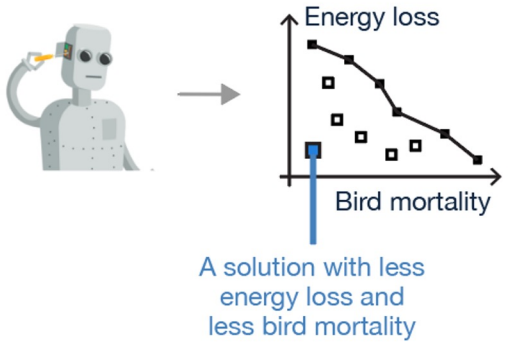

Return solution

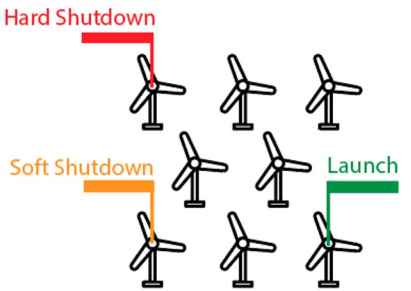

Retrospective review

Turbine ID	Action	Est. Mortality	Est. Loss
18	SSD	0	125 kWh
27	HSD	3	300 kWh
...	...	...	...
89	LAU	0	0

Catalyze solutions to critical domains

Mitigating climate change impacts on migratory birds

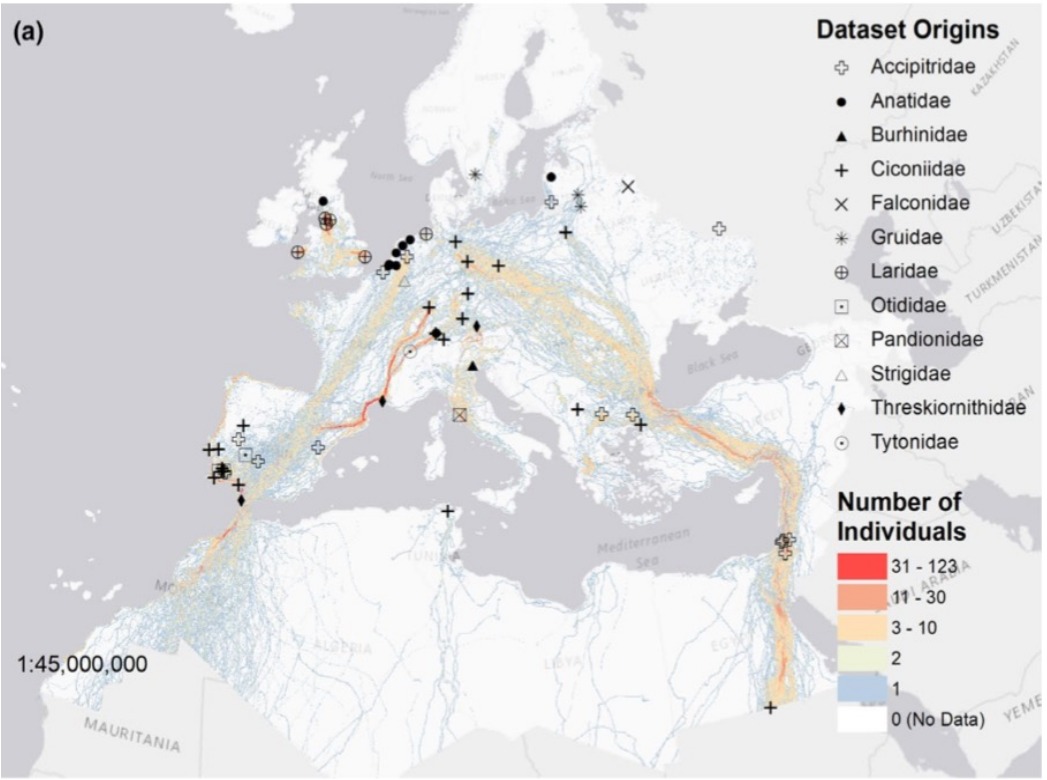

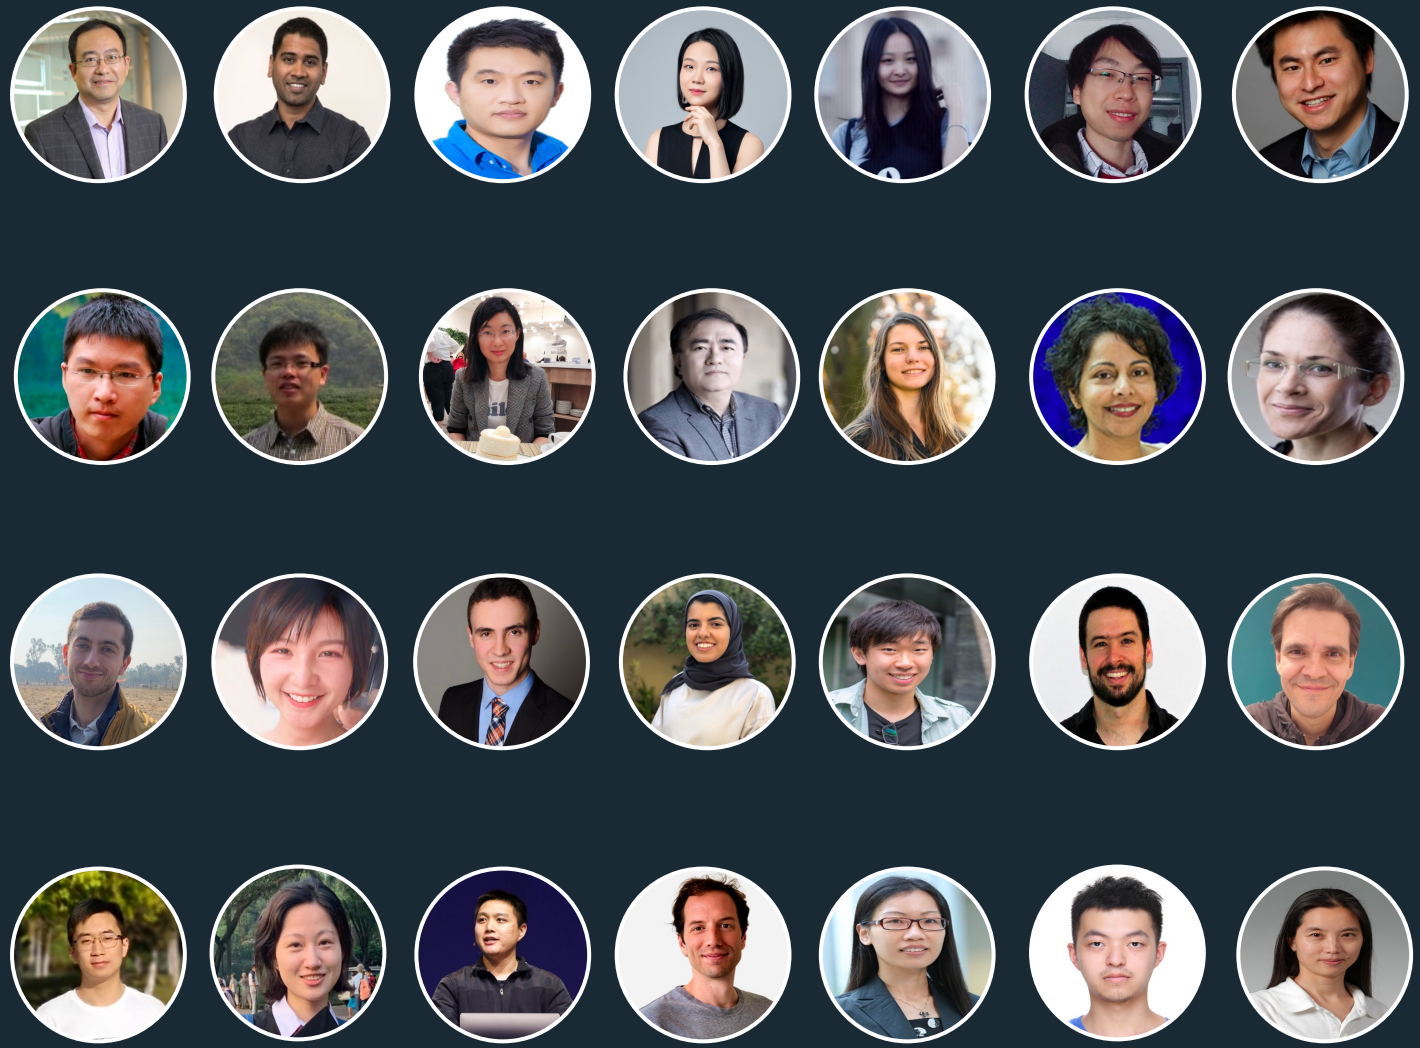

Human concerns and values must be prioritized

AI should aim to augment humans, not replace humans

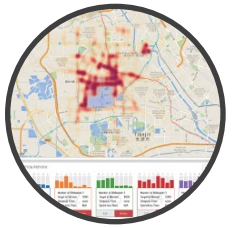

Visualization and user interfaces need much greater development