Automatic Summarization and Visualization of Incident Reports

Robert Gove

robert.gove@twosixtech.com
@rpgove

Nathan Danneman

nathandanneman@datamachines.io

This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA). The views, opinions, and/or findings expressed are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government. Distribution Statement "A" (Approved for Public Release, Distribution Unlimited)

Big Goals

Visualize complex incident reports.

Summarize incident reports in a way that simplifies and preserves (or even enhances?) accuracy.

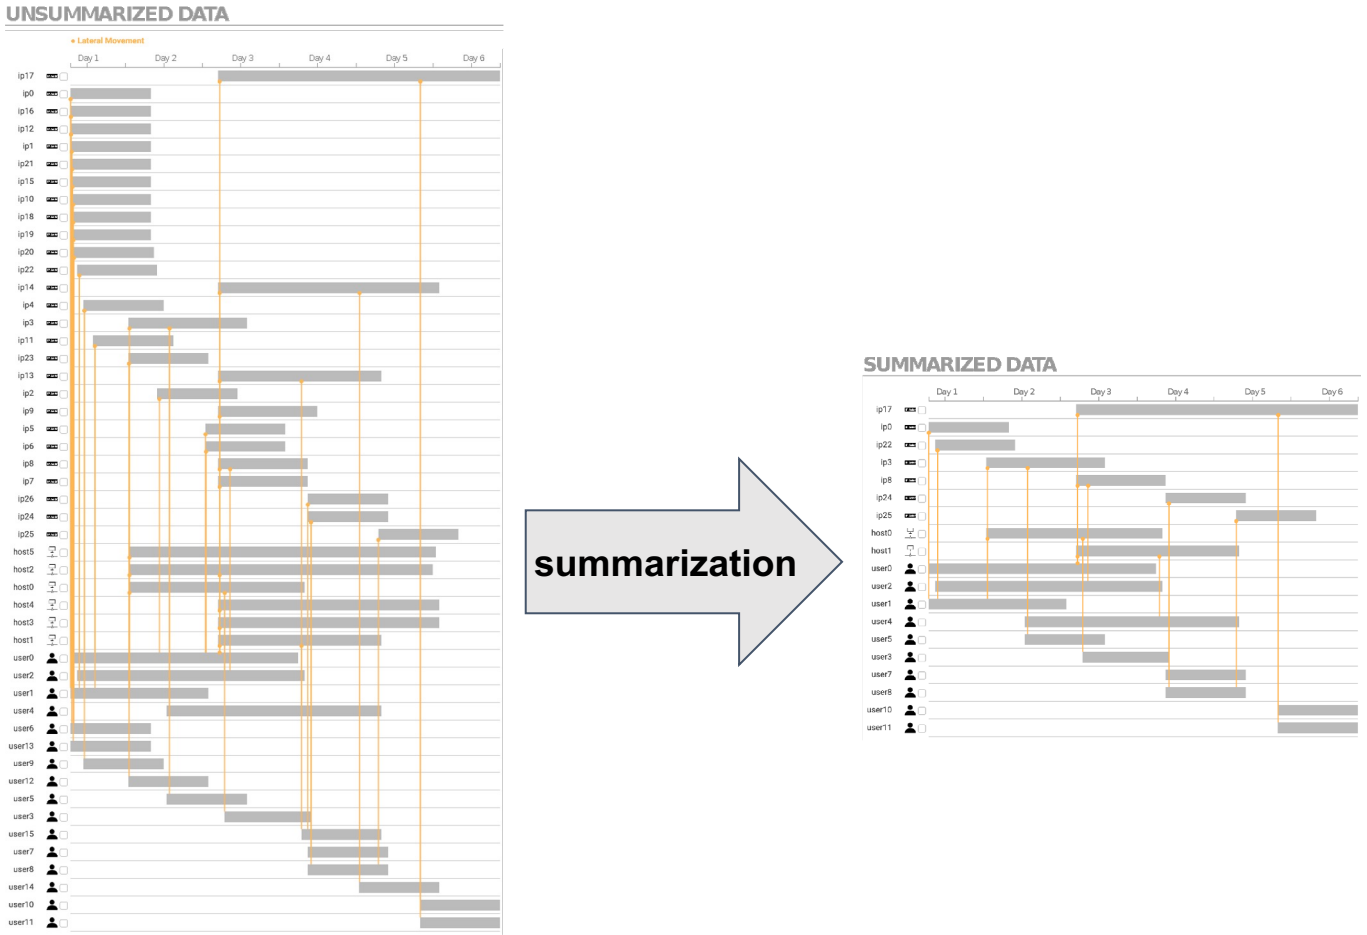

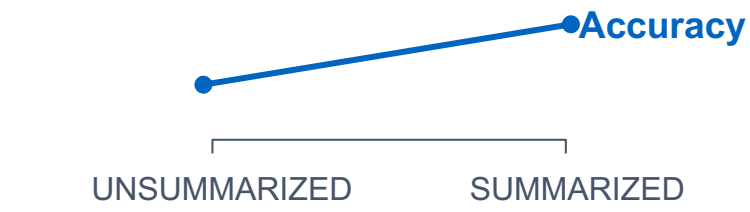

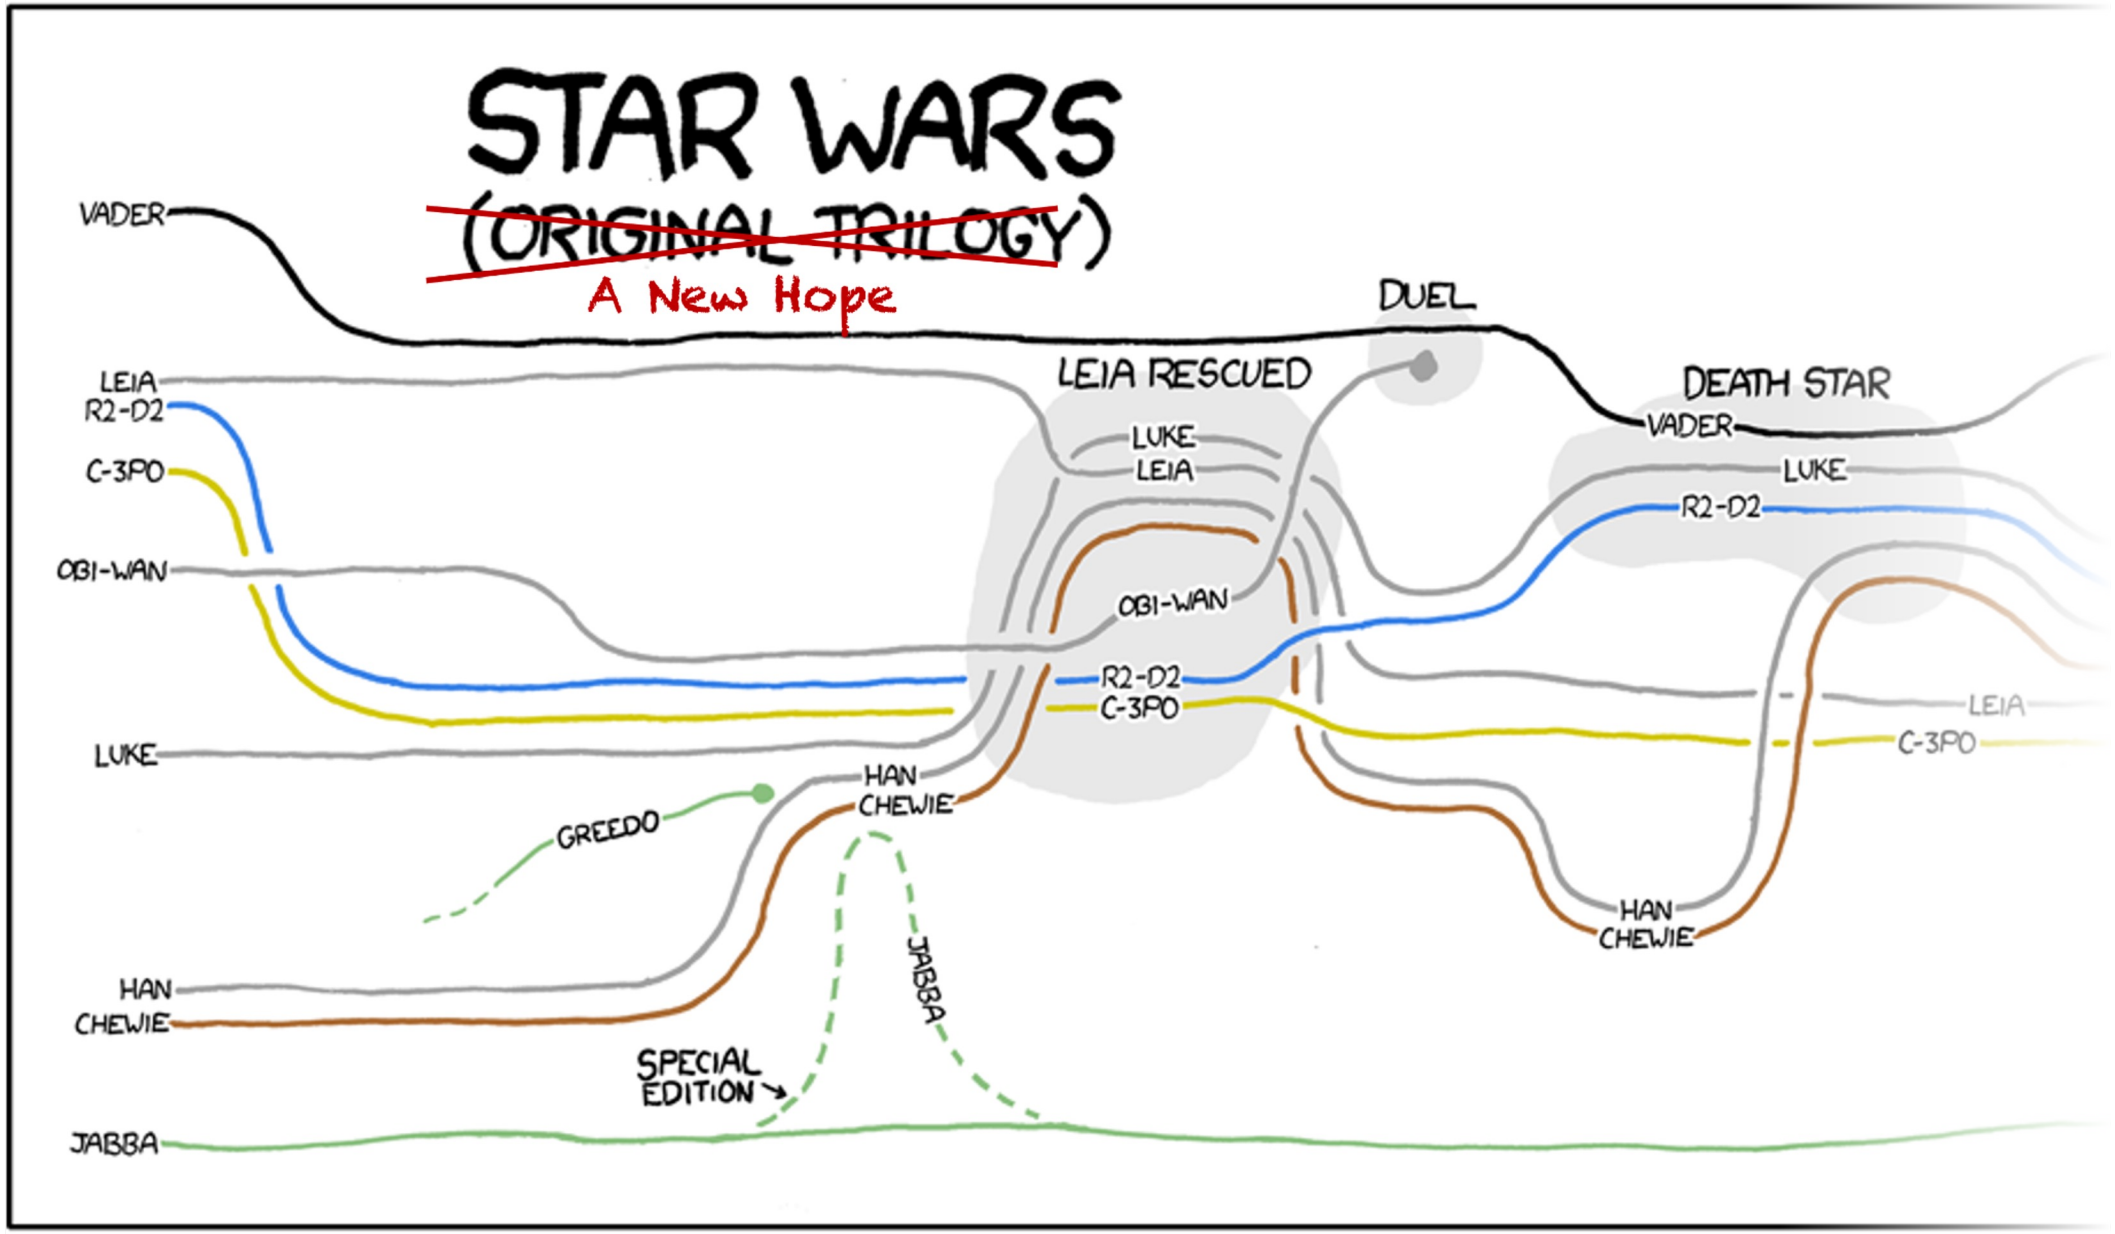

Randall Munroe. Movie Narrative Charts. https://xkcd.com/657/

Summarization Prior Art

Natural language processing

- Summarization vs. simplification
- Extractive vs. abstractive
- Relies on *natural* language
  - We have structured data

Graph theory

- Proxy graphs
  - Derive smaller representative graphs
  - Sampling, filtering, graph filtrations, etc.

Incident Reports → Dynamic Graphs

Incident report

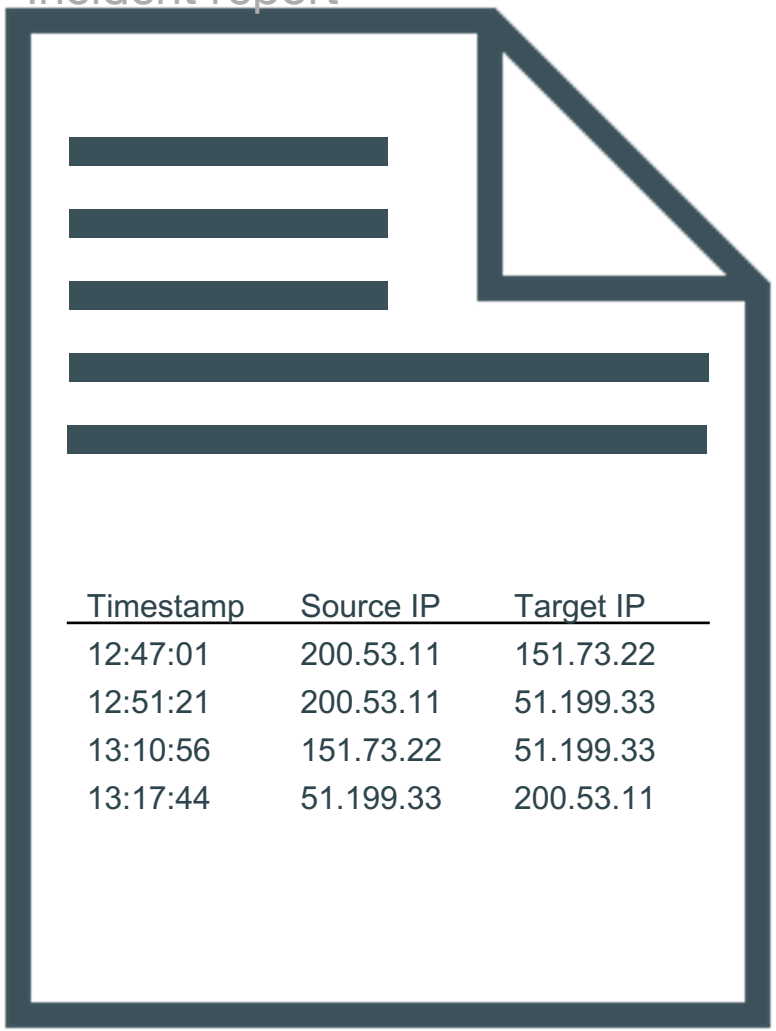

Timestamp	Source IP	Target IP
12:47:01	200.53.11	151.73.22
12:51:21	200.53.11	51.199.33
13:10:56	151.73.22	51.199.33
13:17:44	51.199.33	200.53.11

Incident Reports → Dynamic Graphs

Incident report

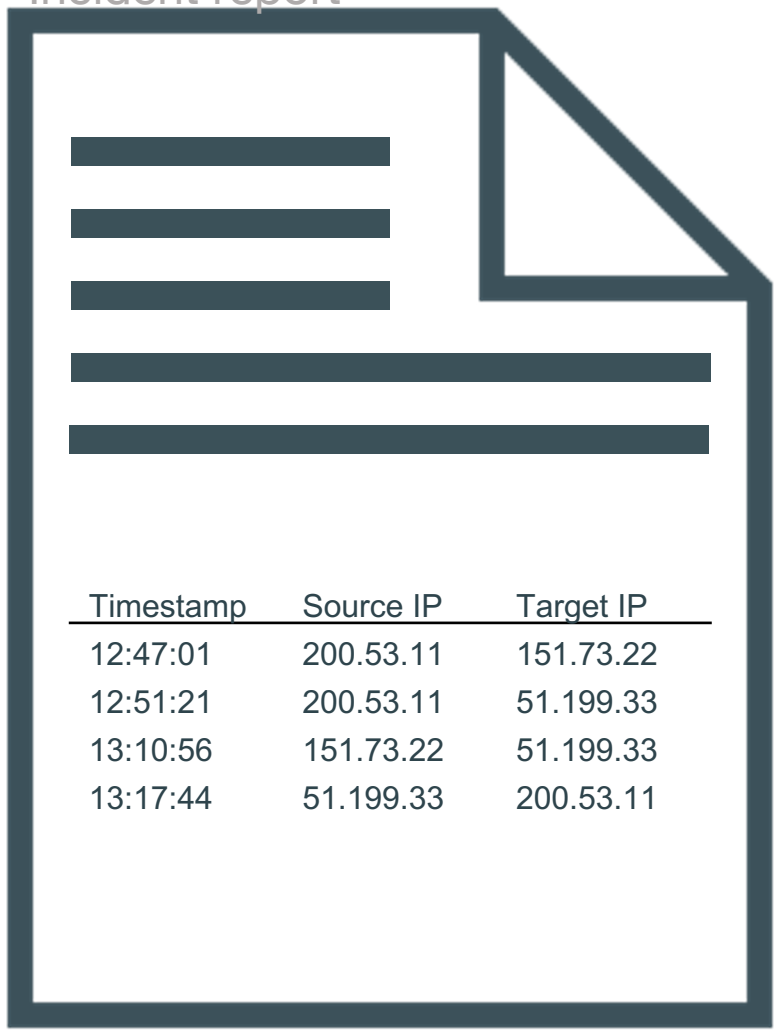

Timestamp	Source IP	Target IP
12:47:01	200.53.11	151.73.22
12:51:21	200.53.11	51.199.33
13:10:56	151.73.22	51.199.33
13:17:44	51.199.33	200.53.11

Tabular log data

Incident Reports → Dynamic Graphs

Incident report

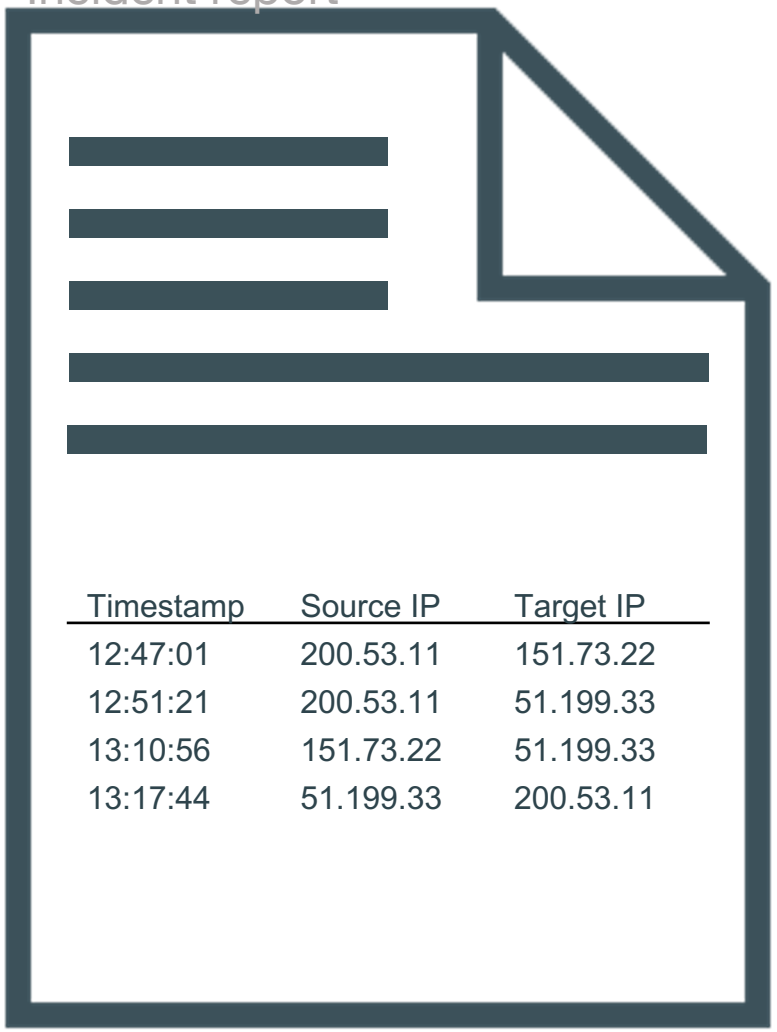

Timestamp	Source IP	Target IP
12:47:01	200.53.11	151.73.22
12:51:21	200.53.11	51.199.33
13:10:56	151.73.22	51.199.33
13:17:44	51.199.33	200.53.11

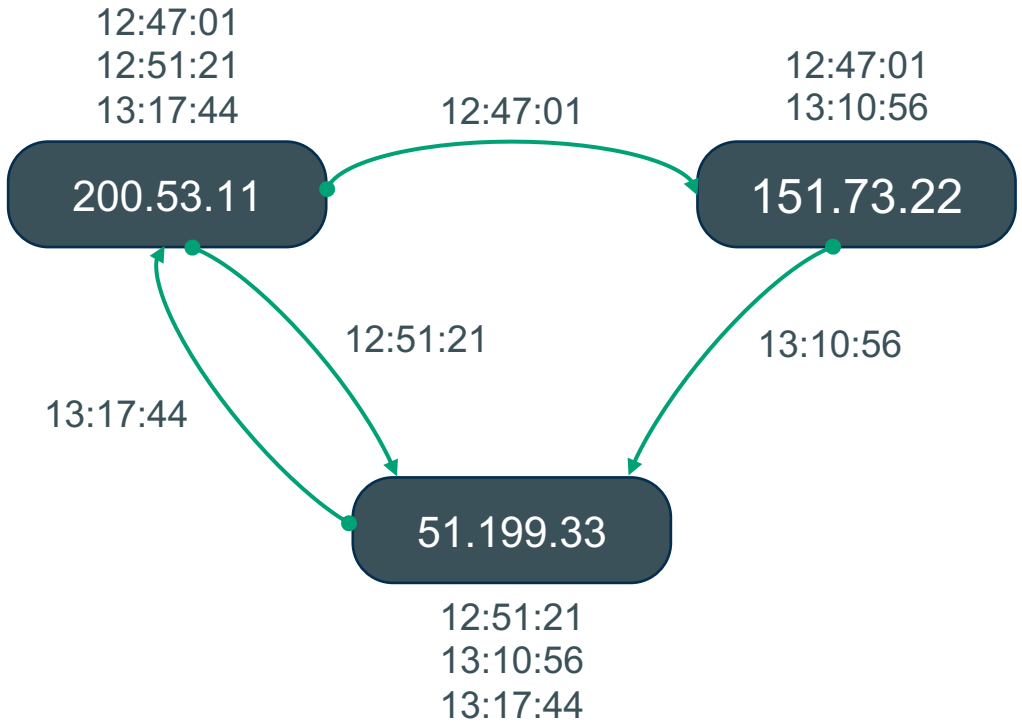

Tabular log data converted to a dynamic graph

Visualization Design

Goals:

- Succinctness
- Consistency
- Activity progression
- Patterns
- Learnability

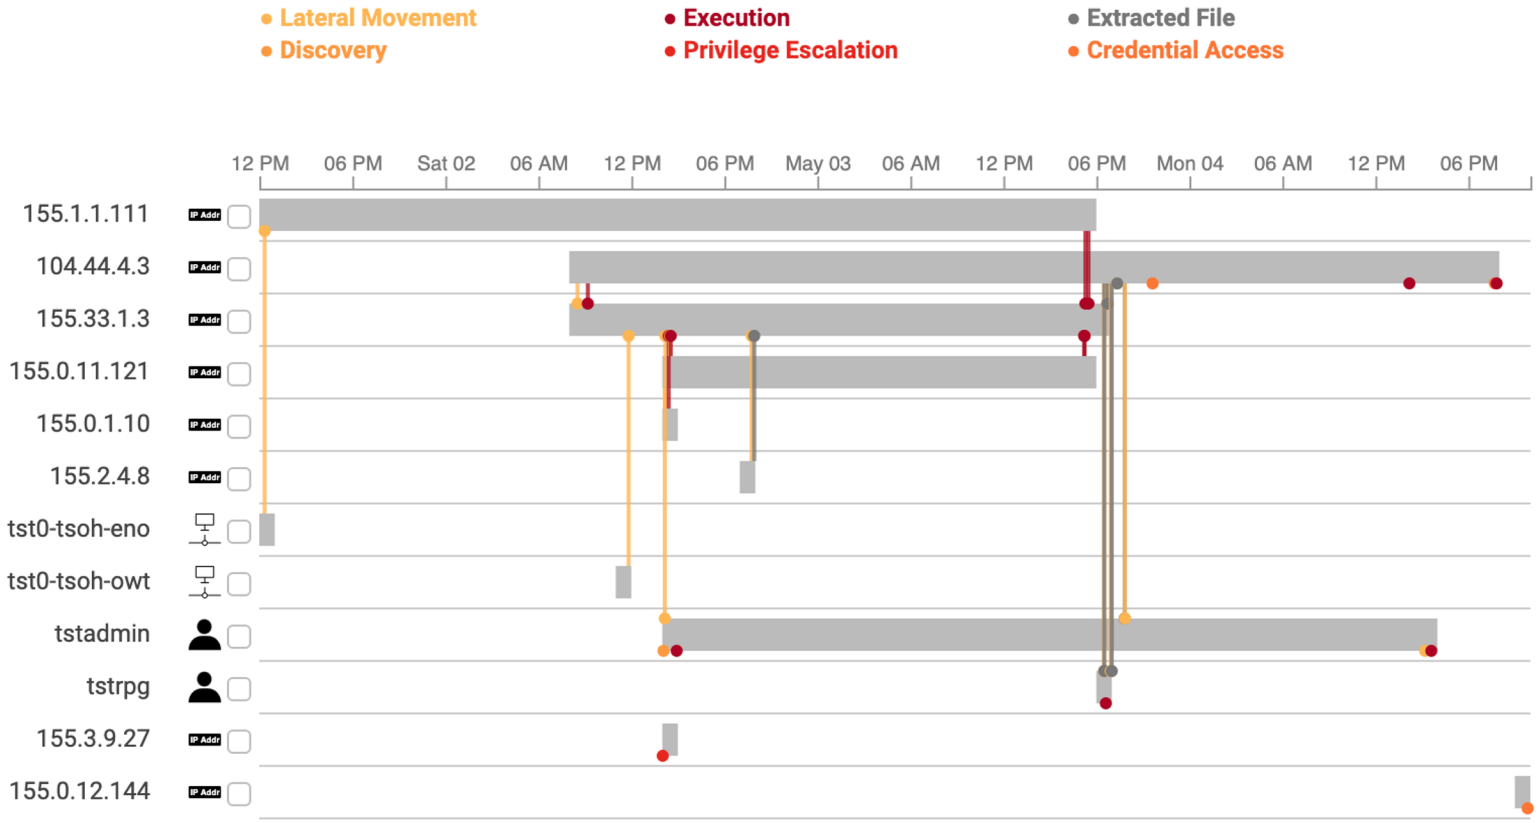

Visualization Design

Goals:

- Succinctness
- Consistency
- Activity progression
- Patterns
- Learnability

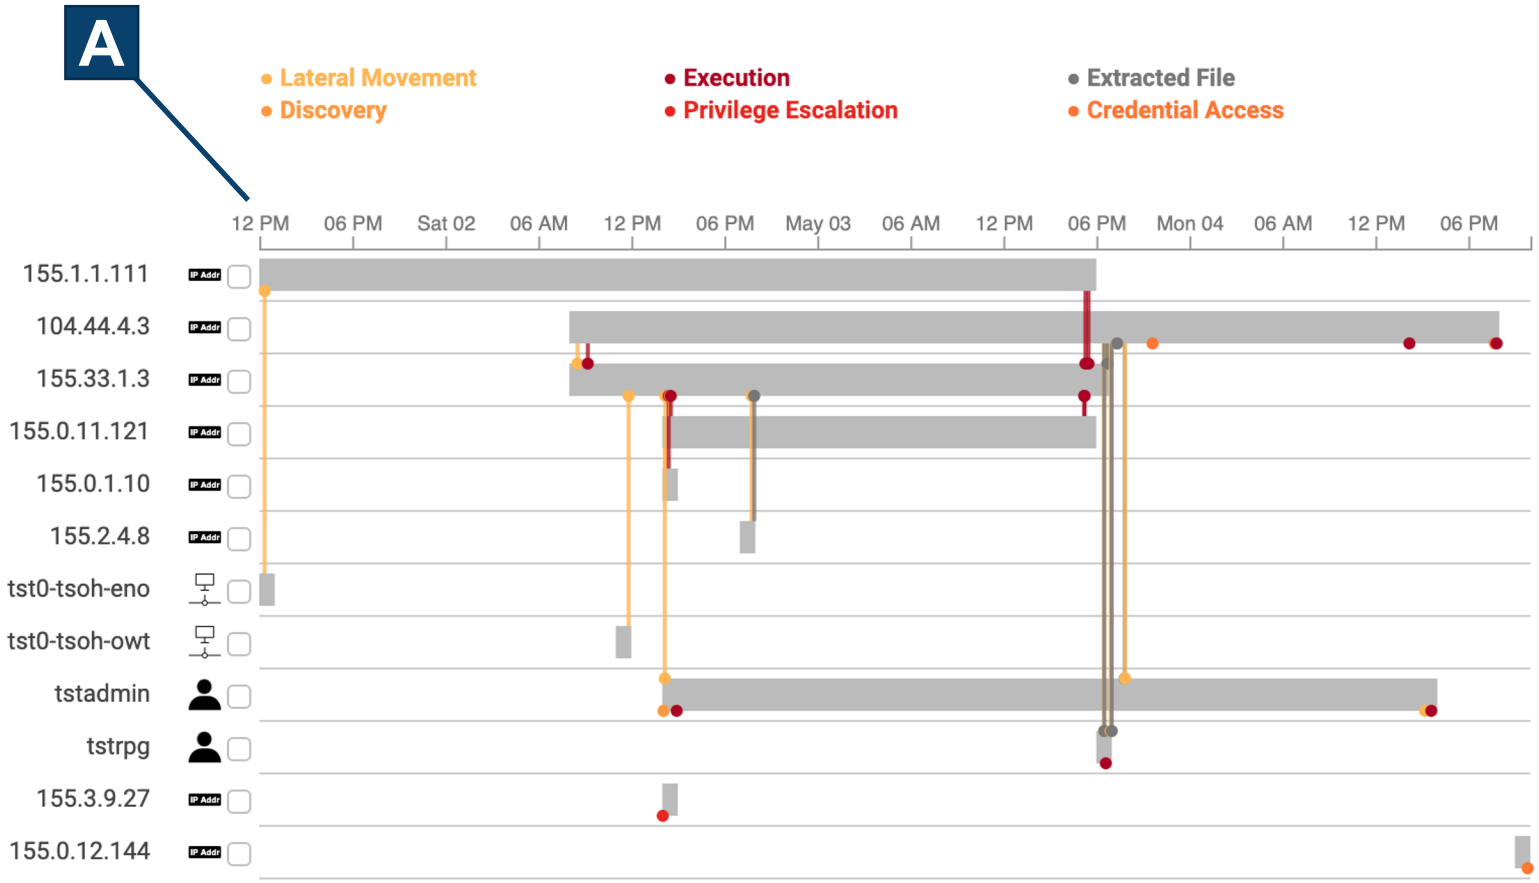

Visualization Design

Goals:

- Succinctness
- Consistency
- Activity progression
- Patterns
- Learnability

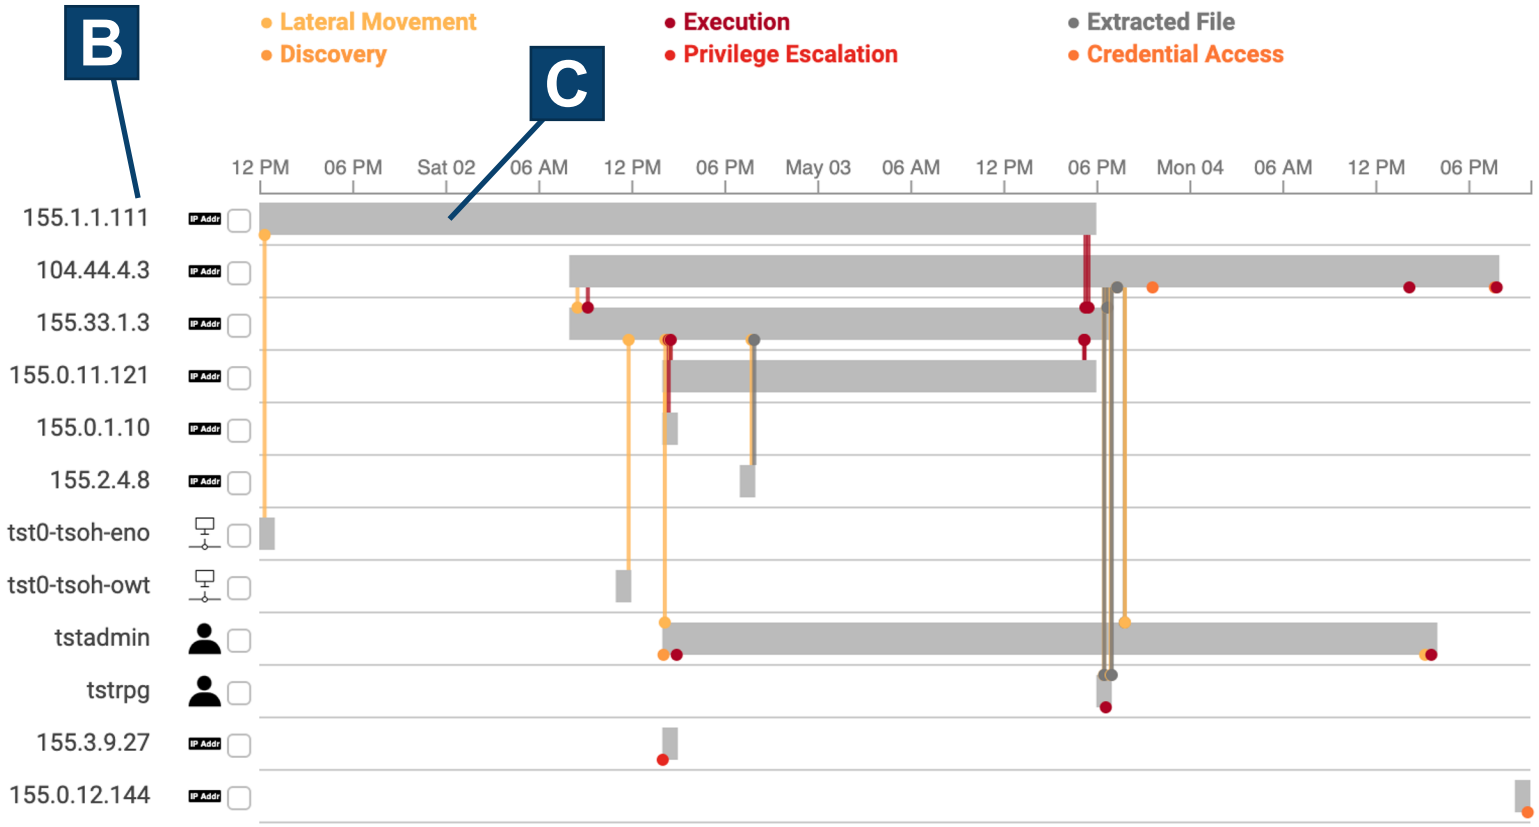

Visualization Design

Goals:

- Succinctness
- Consistency
- Activity progression
- Patterns
- Learnability

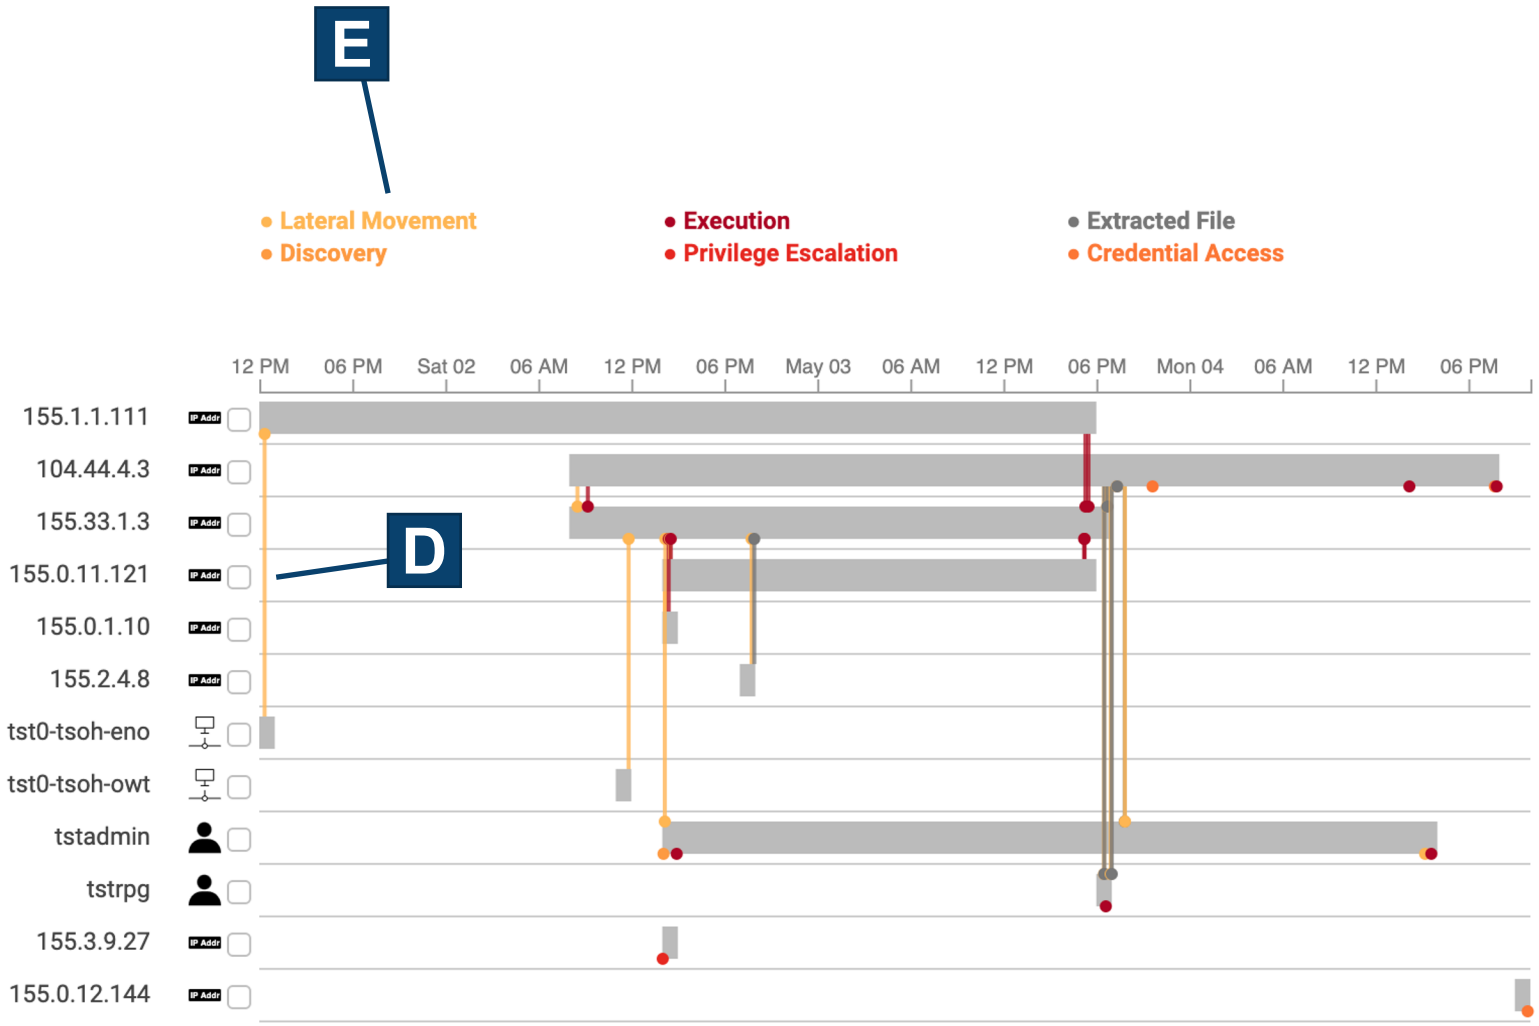

Visualization Design

Goals:

- Succinctness
- Consistency
- Activity progression
- Patterns
- Learnability

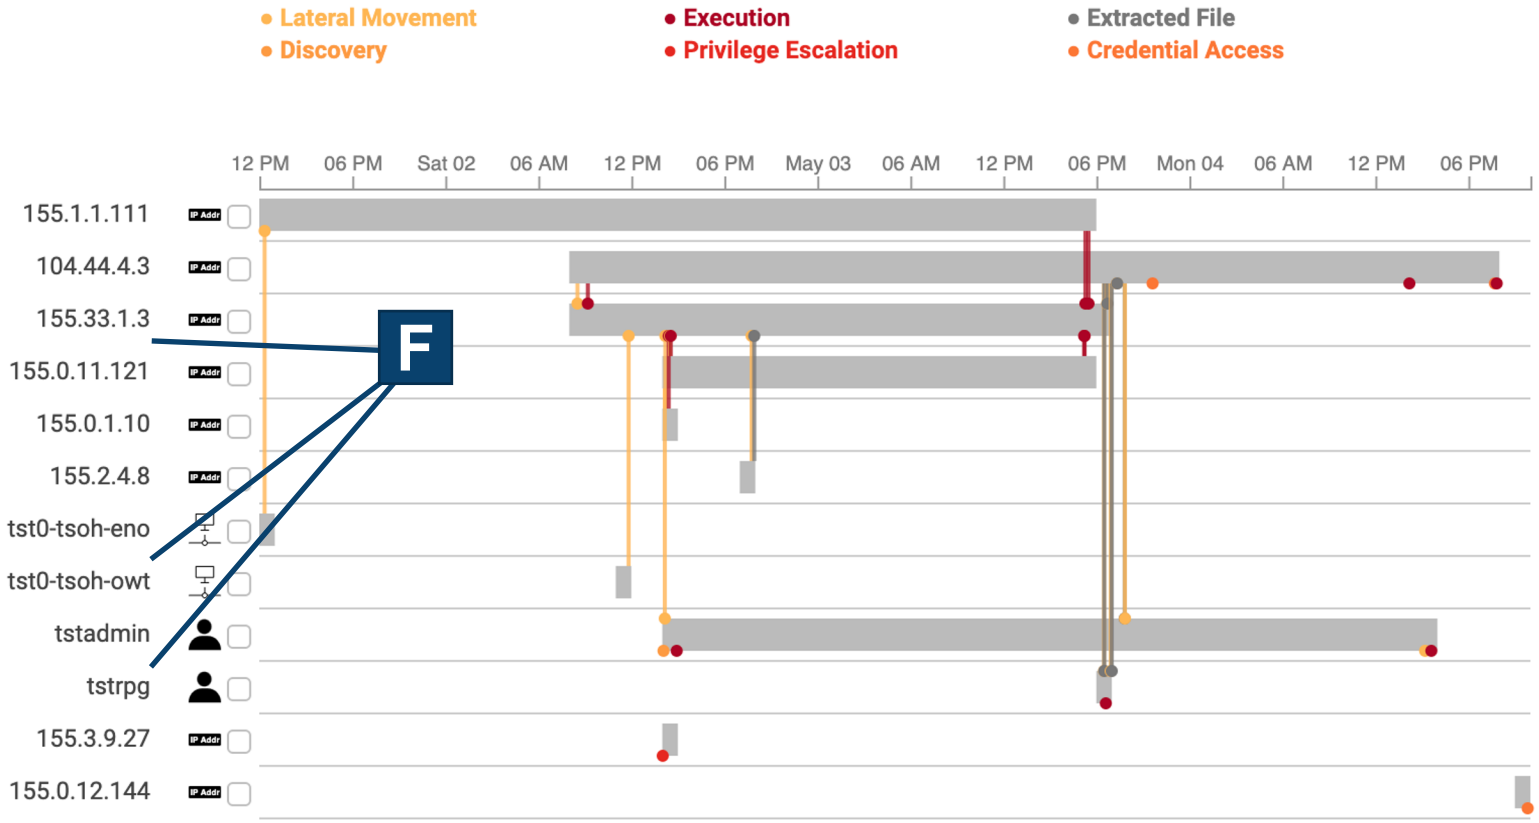

Summarization: What is good?

Evaluation Criteria

- Size: smaller incident reports are better
- Accuracy: don’t drop entities that are true positives
- Useful for analysts

Wait, is that reasonable?

- Incident reports are the results of algorithms with access to more information…
- Detector authors have strong incentives to optimize precision and recall…

Featurizing Incident Report Dynamic Graphs

Calculate scores

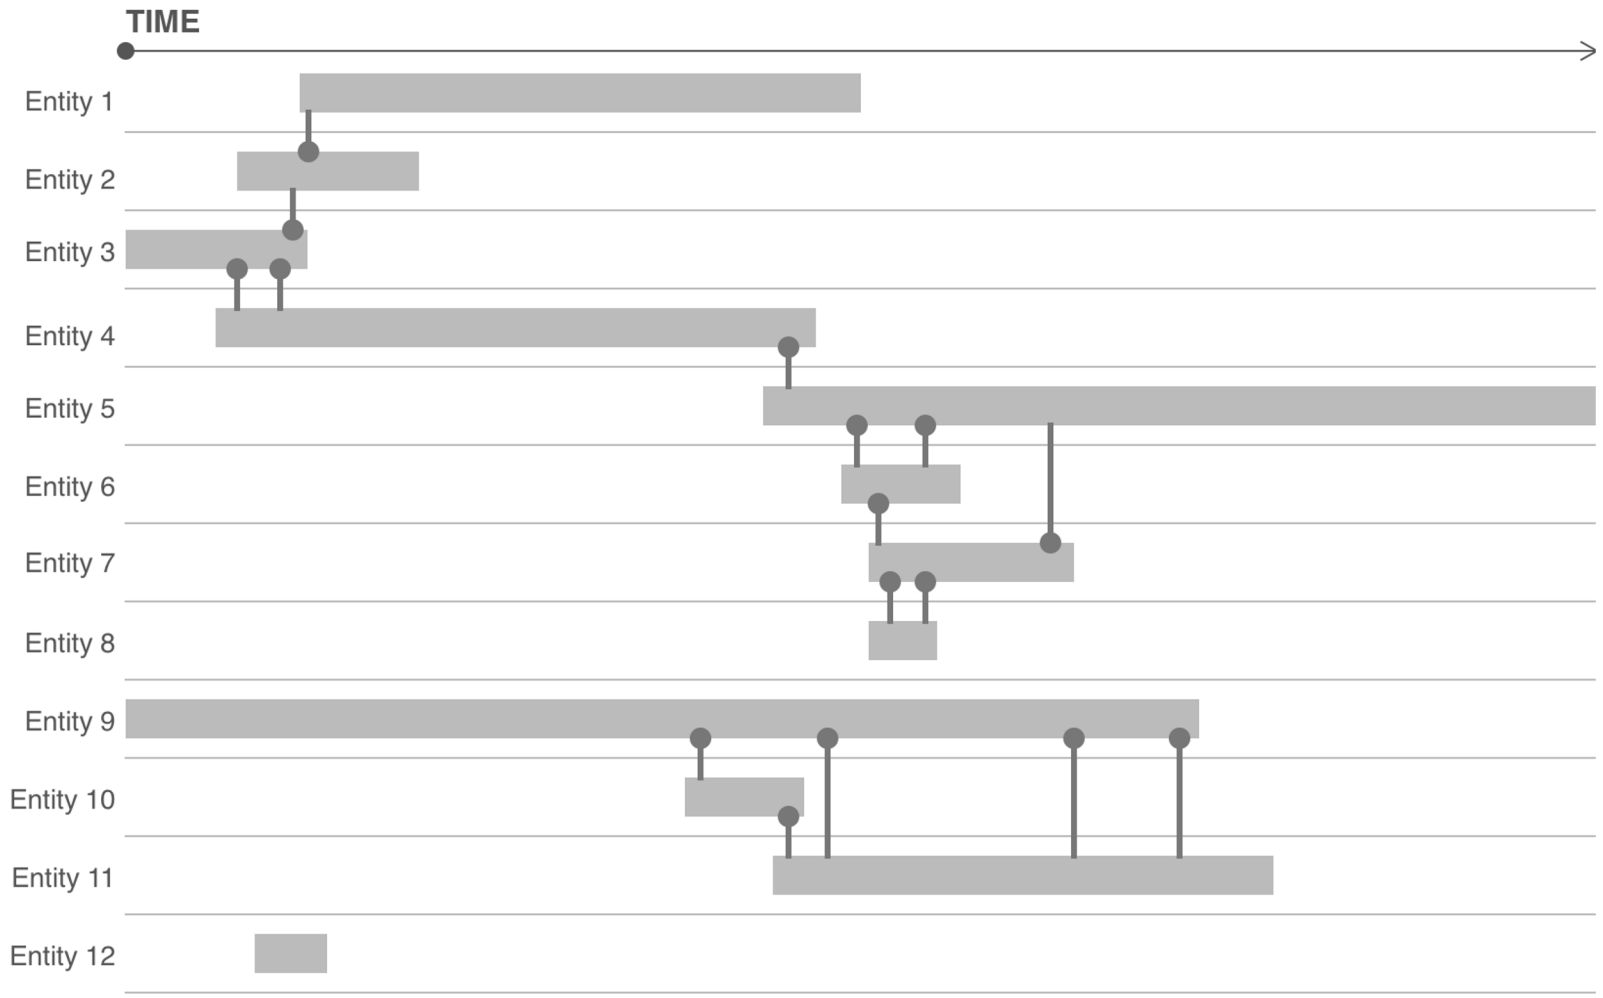

Featurizing Incident Report Dynamic Graphs

Calculate scores

- Component scores

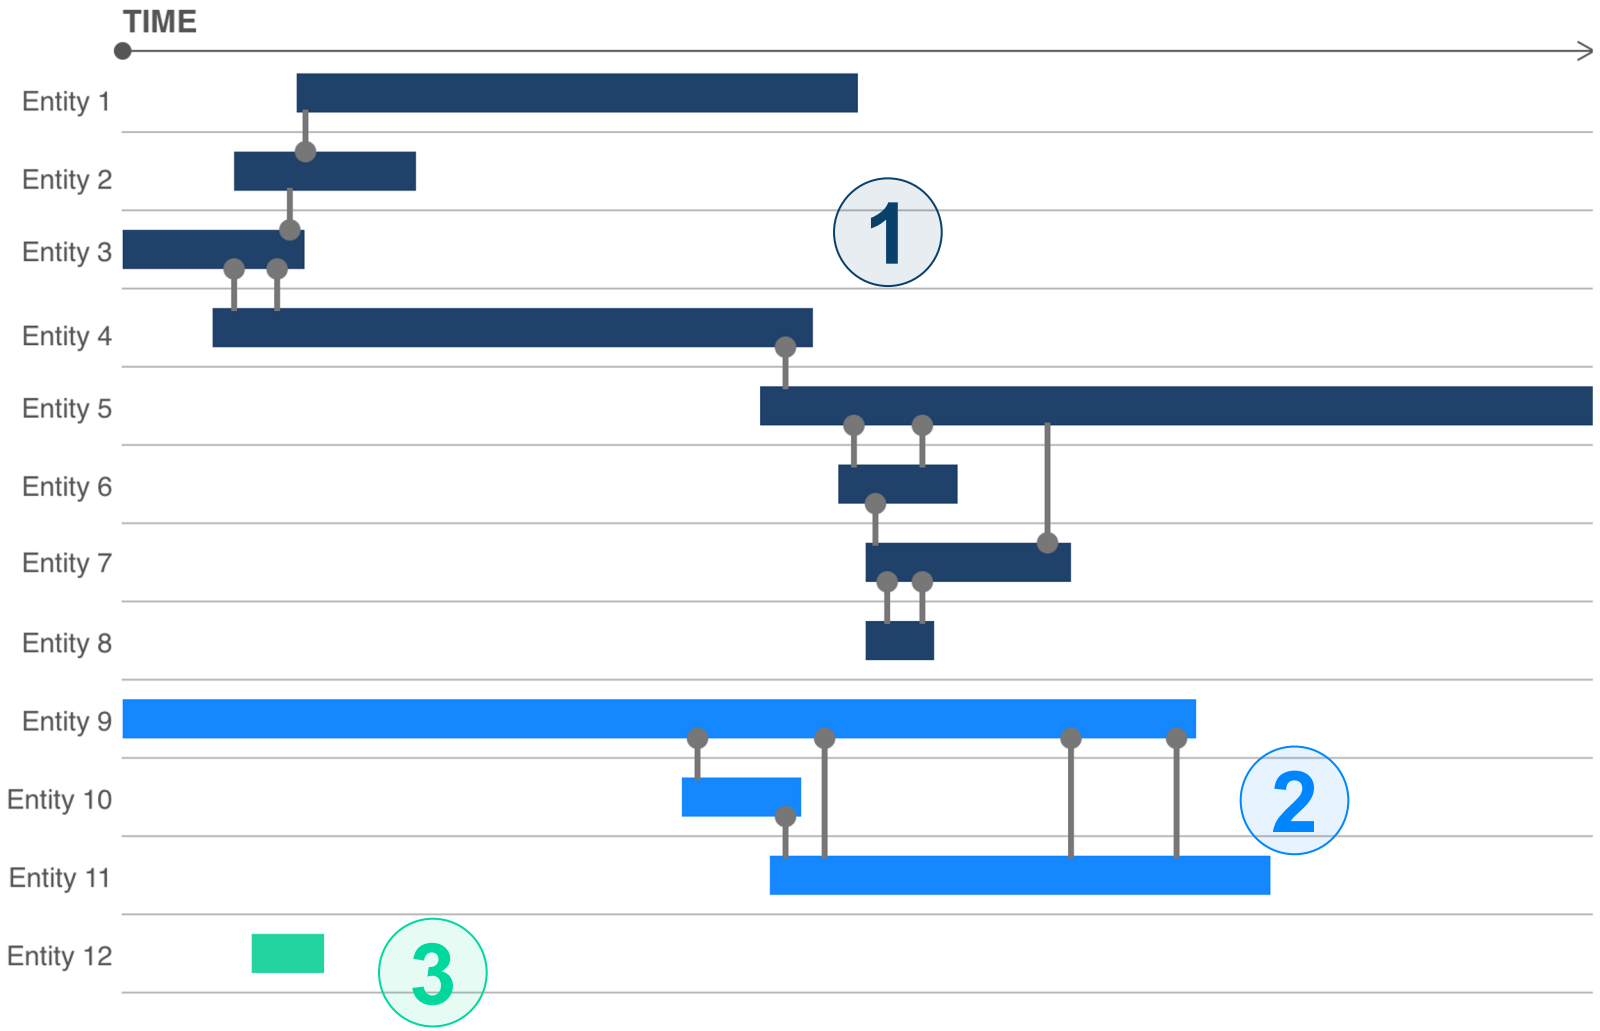

Featurizing Incident Report Dynamic Graphs

Calculate scores

- Component scores
  - Relative duration

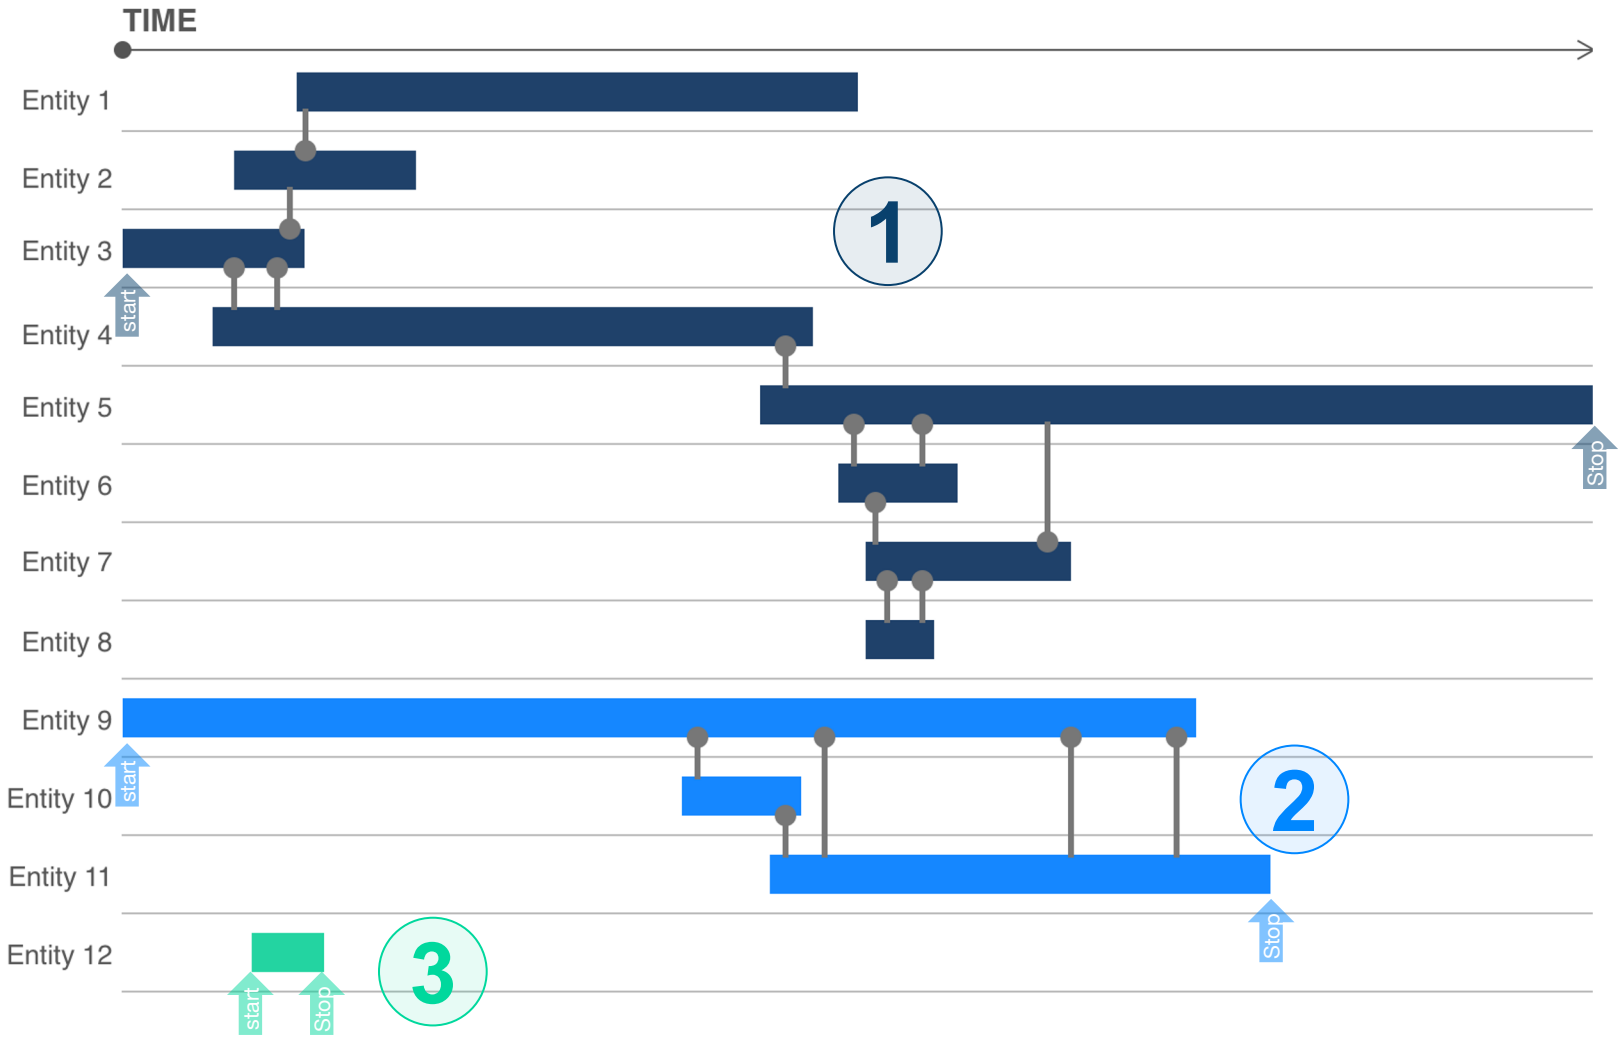

Featurizing Incident Report Dynamic Graphs

Calculate scores

- Component scores
  - Relative duration
  - Relative number of entities

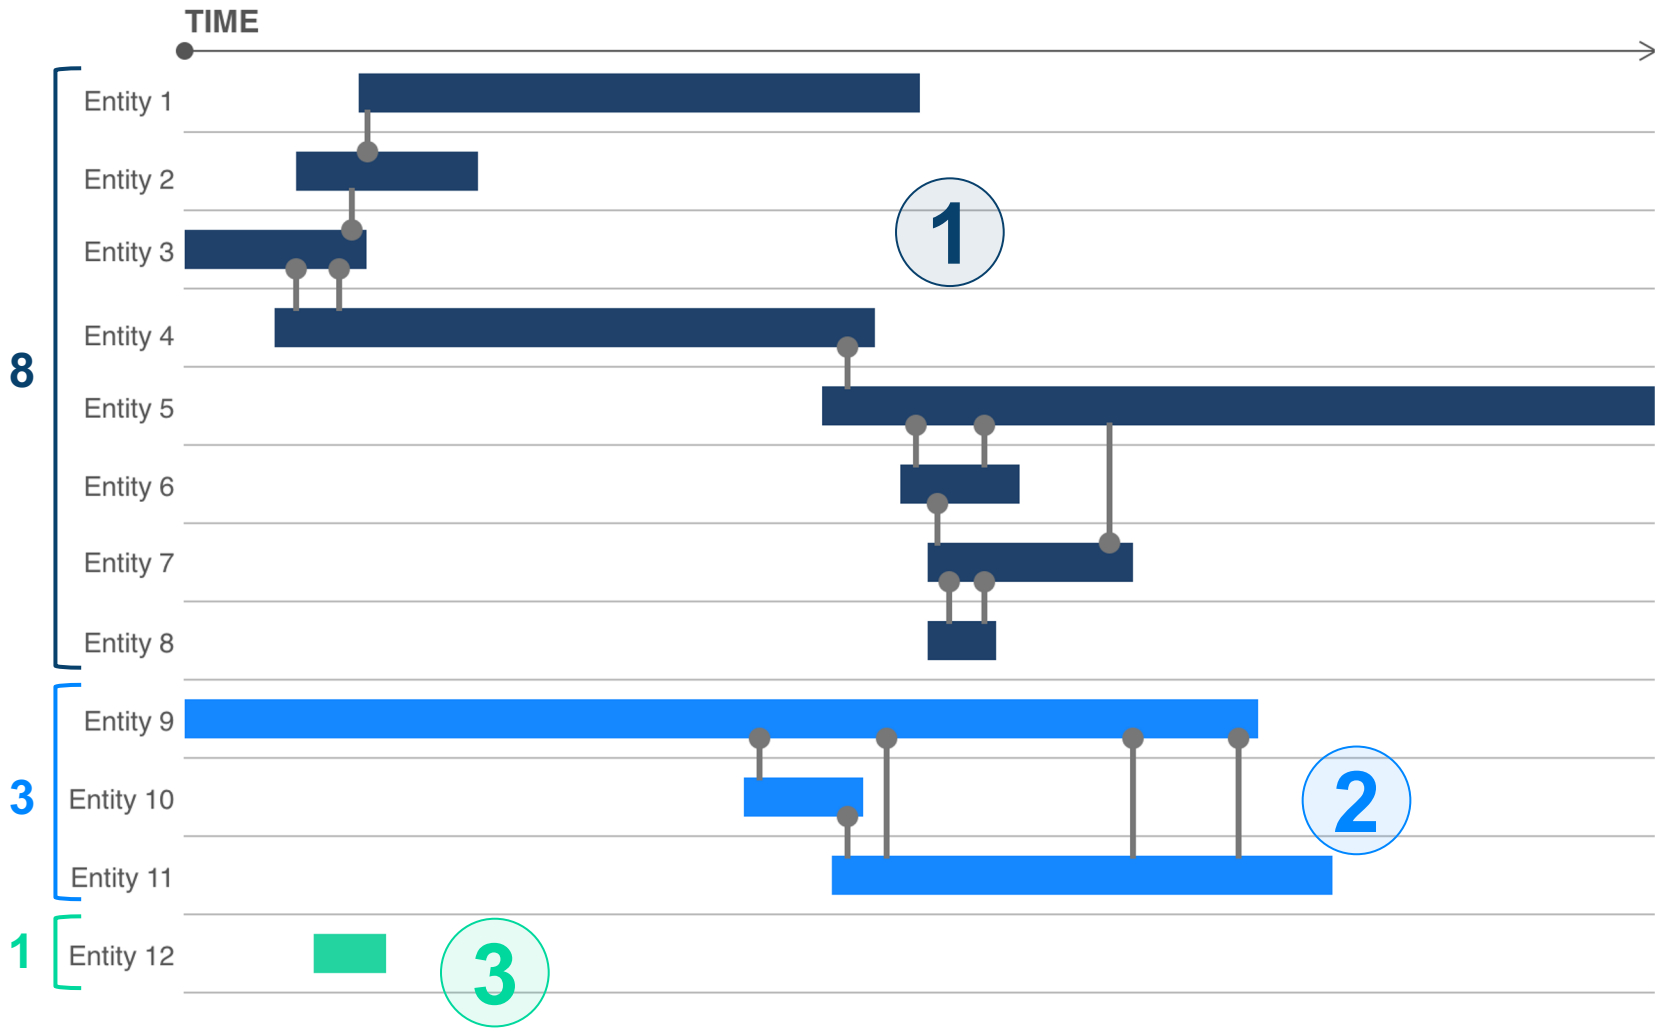

Featurizing Incident Report Dynamic Graphs

Calculate scores

- Component scores
  - Relative duration
  - Relative number of entities
  - Relative number of relationships

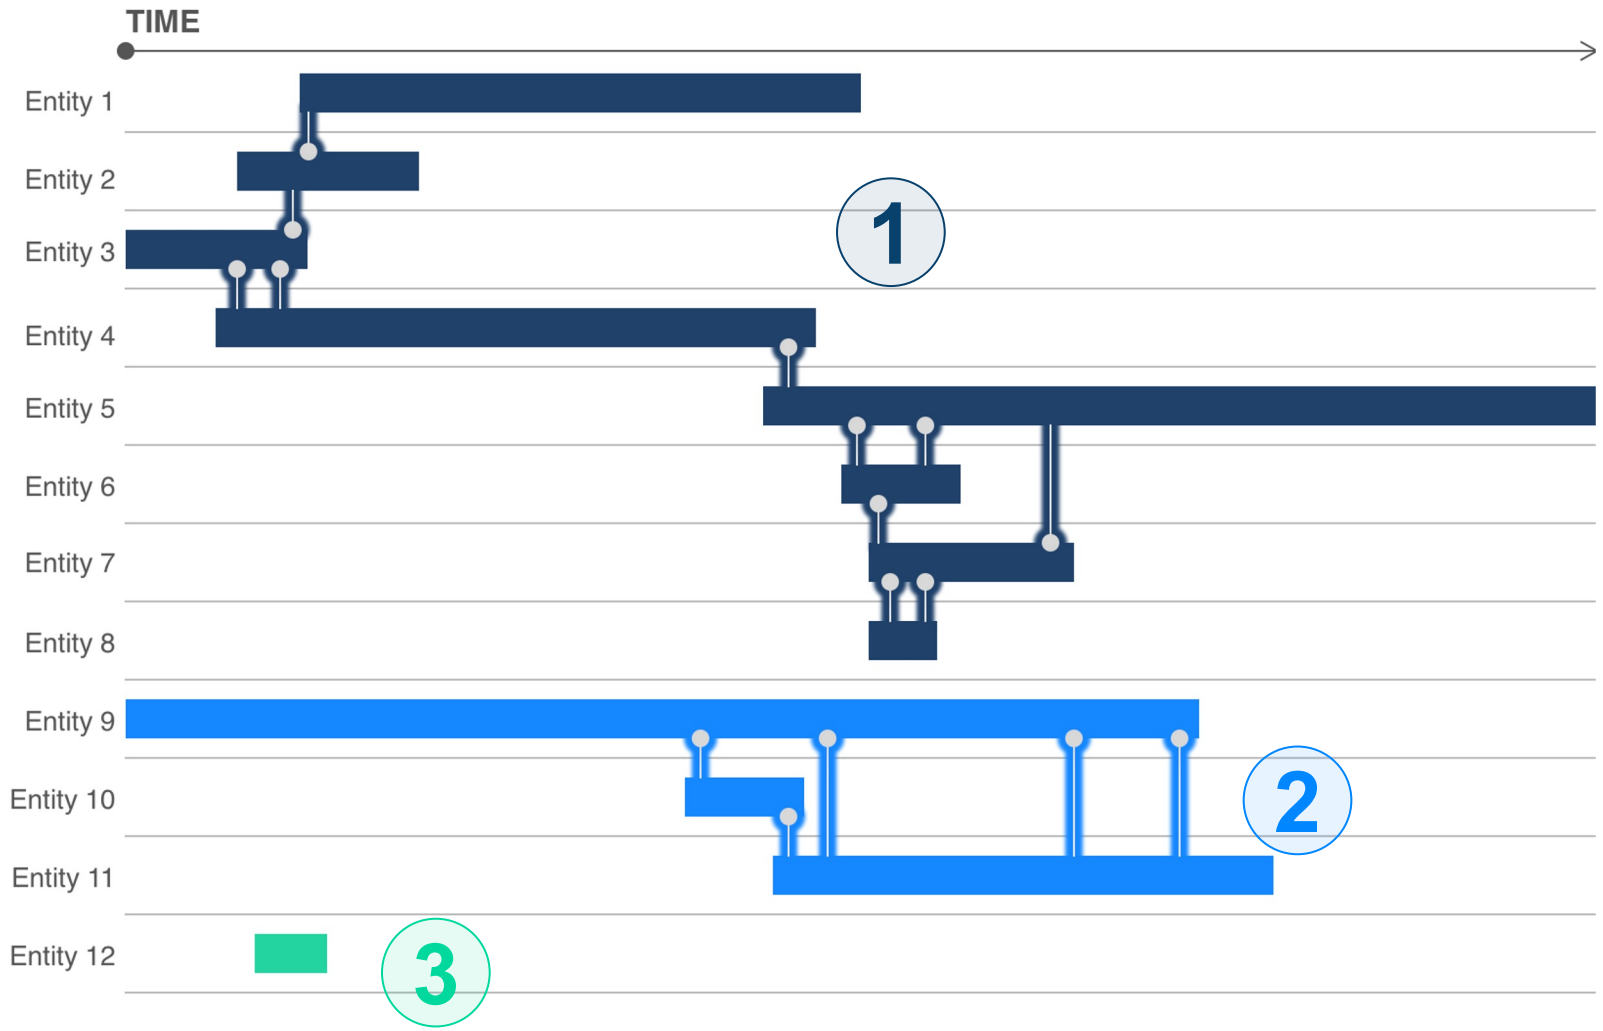

Featurizing Incident Report Dynamic Graphs

Calculate scores

- Component scores
  - Relative duration
  - Relative number of entities
  - Relative number of relationships
  - Relative number of timestamps

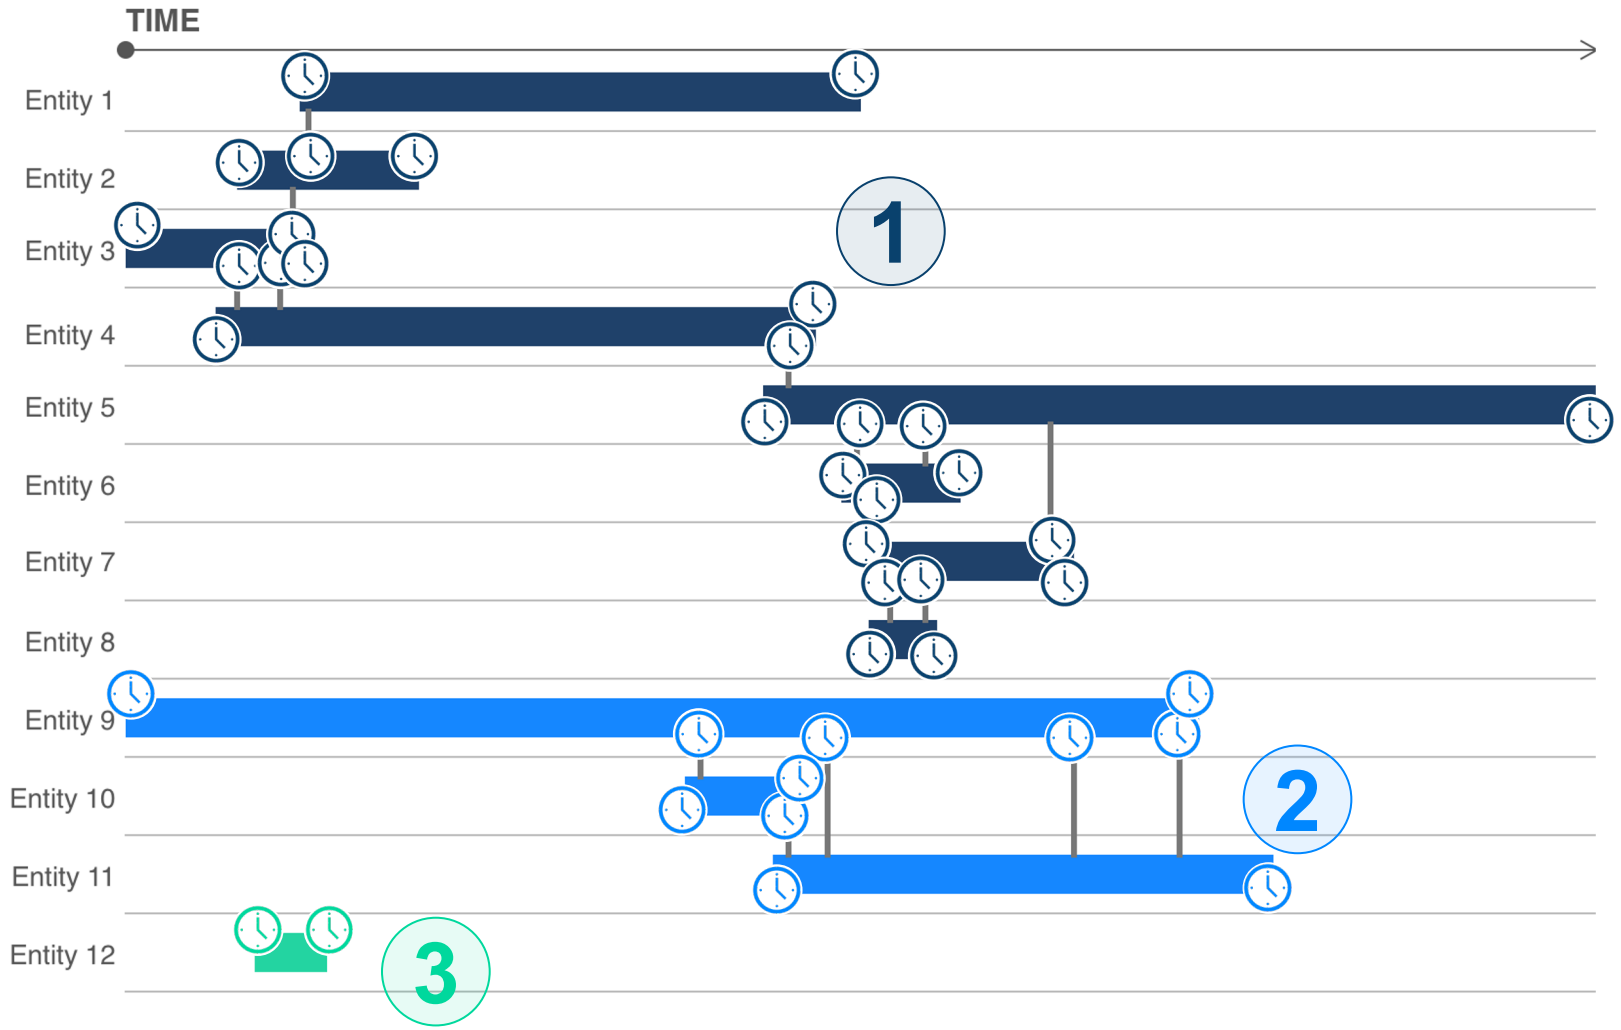

Featurizing Incident Report Dynamic Graphs

Calculate scores

- Component scores
- Branch scores

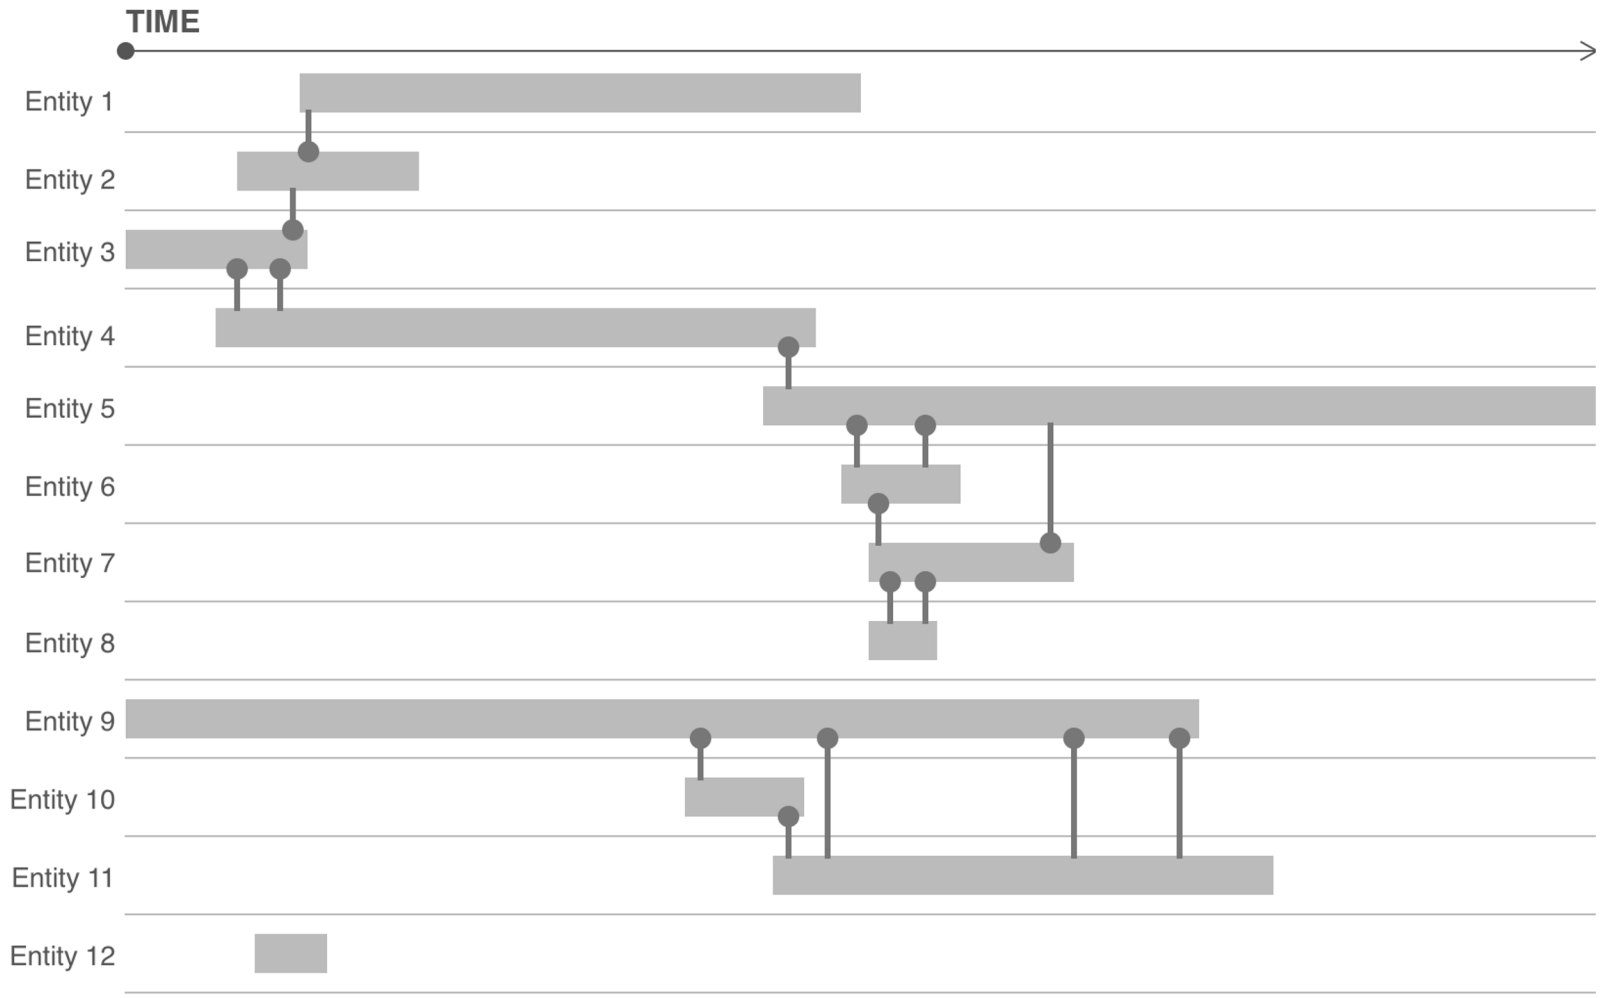

Featurizing Incident Report Dynamic Graphs

Calculate scores

- Component scores
- Branch scores
  - Core sequence of events

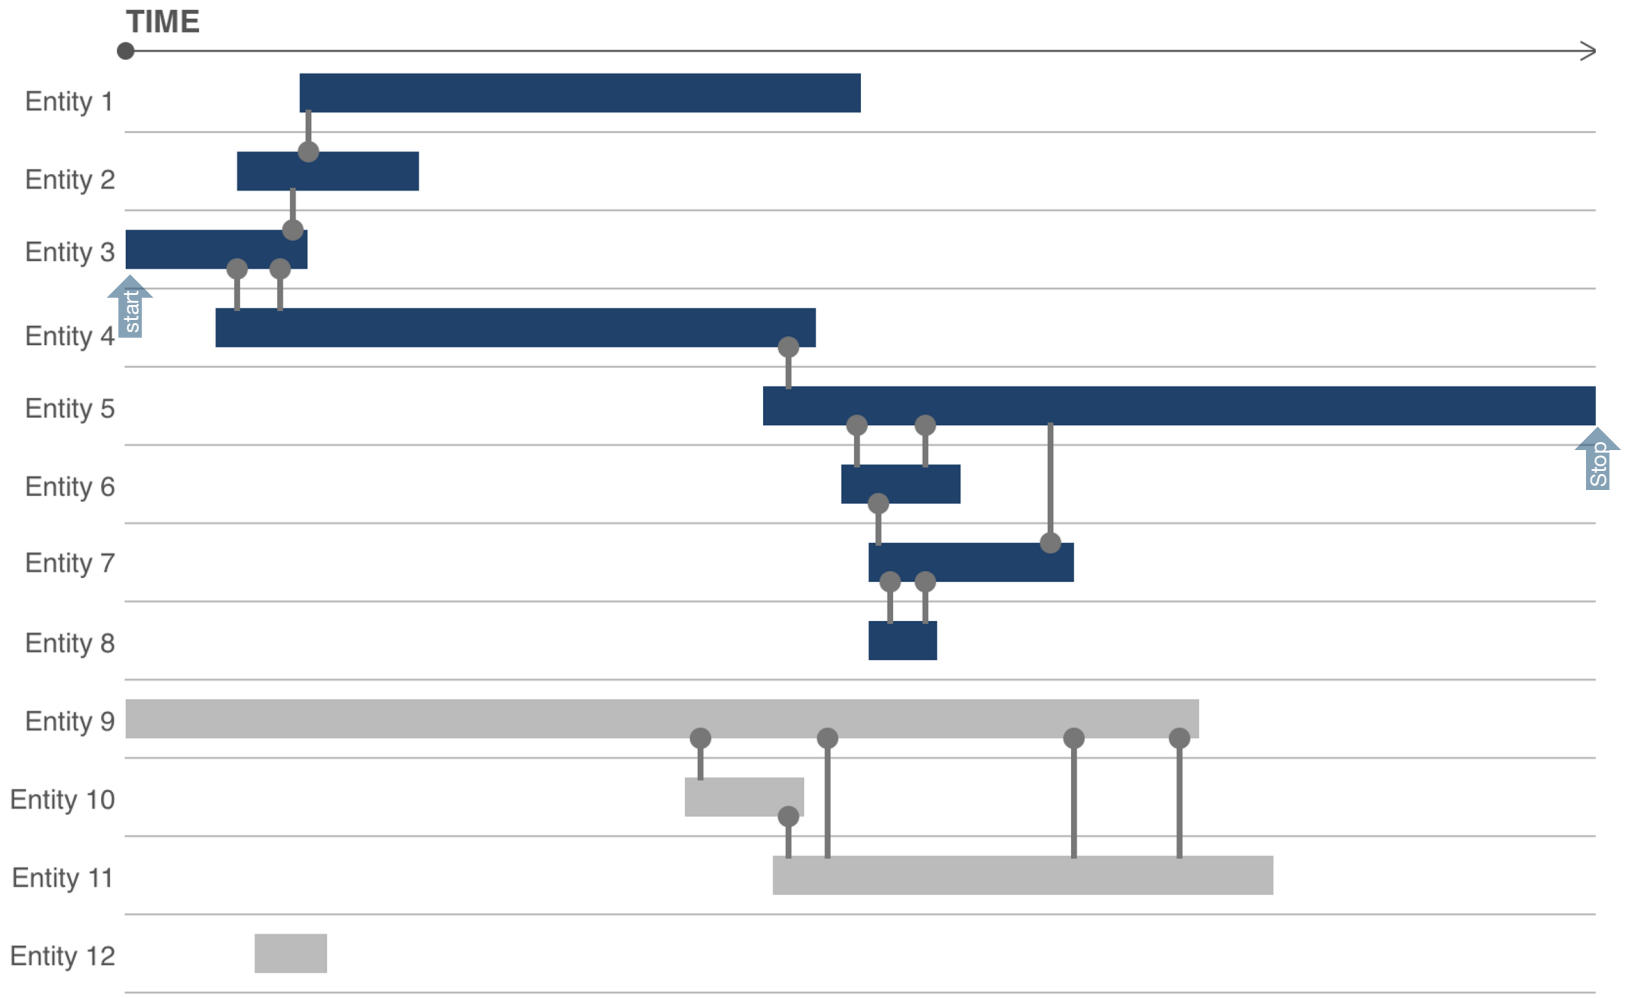

Featurizing Incident Report Dynamic Graphs

Calculate scores

- Component scores
- Branch scores
  - Core sequence of events

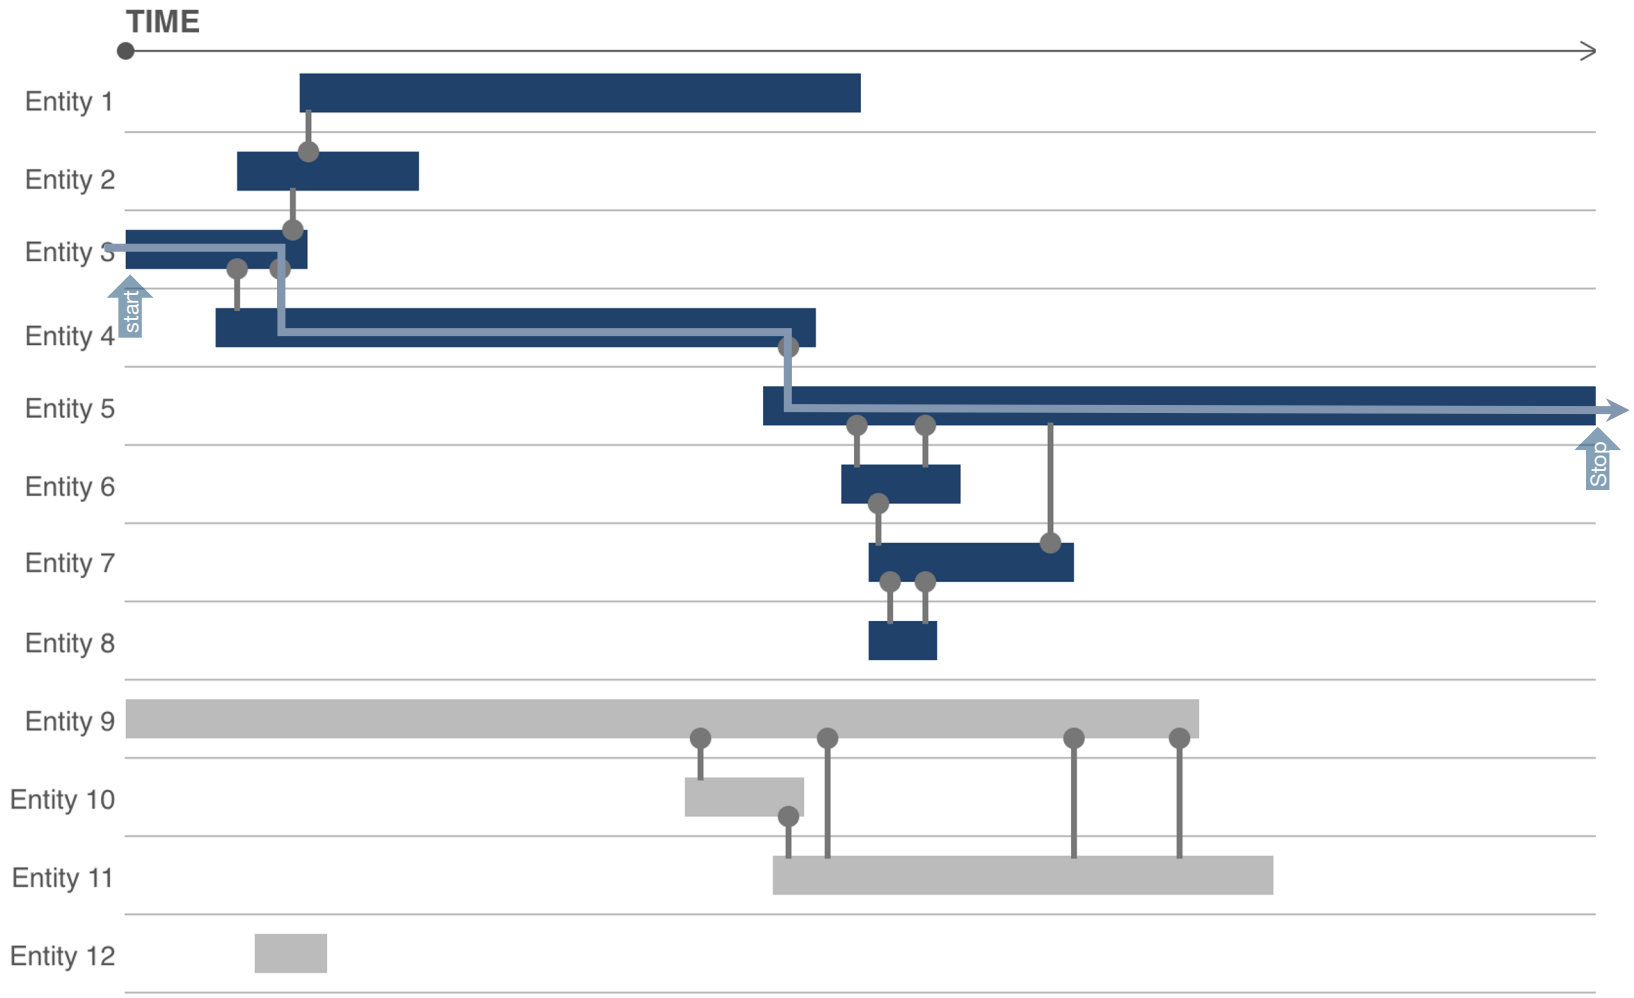

Featurizing Incident Report Dynamic Graphs

Calculate scores

- Component scores
- Branch scores
  - Core sequence of events
  - Earliness of branch
  - Relative branch duration
  - Relative number of timestamps
  - Relative number of entities
  - Relative number of relationships
  - MITRE ATT&CK severity

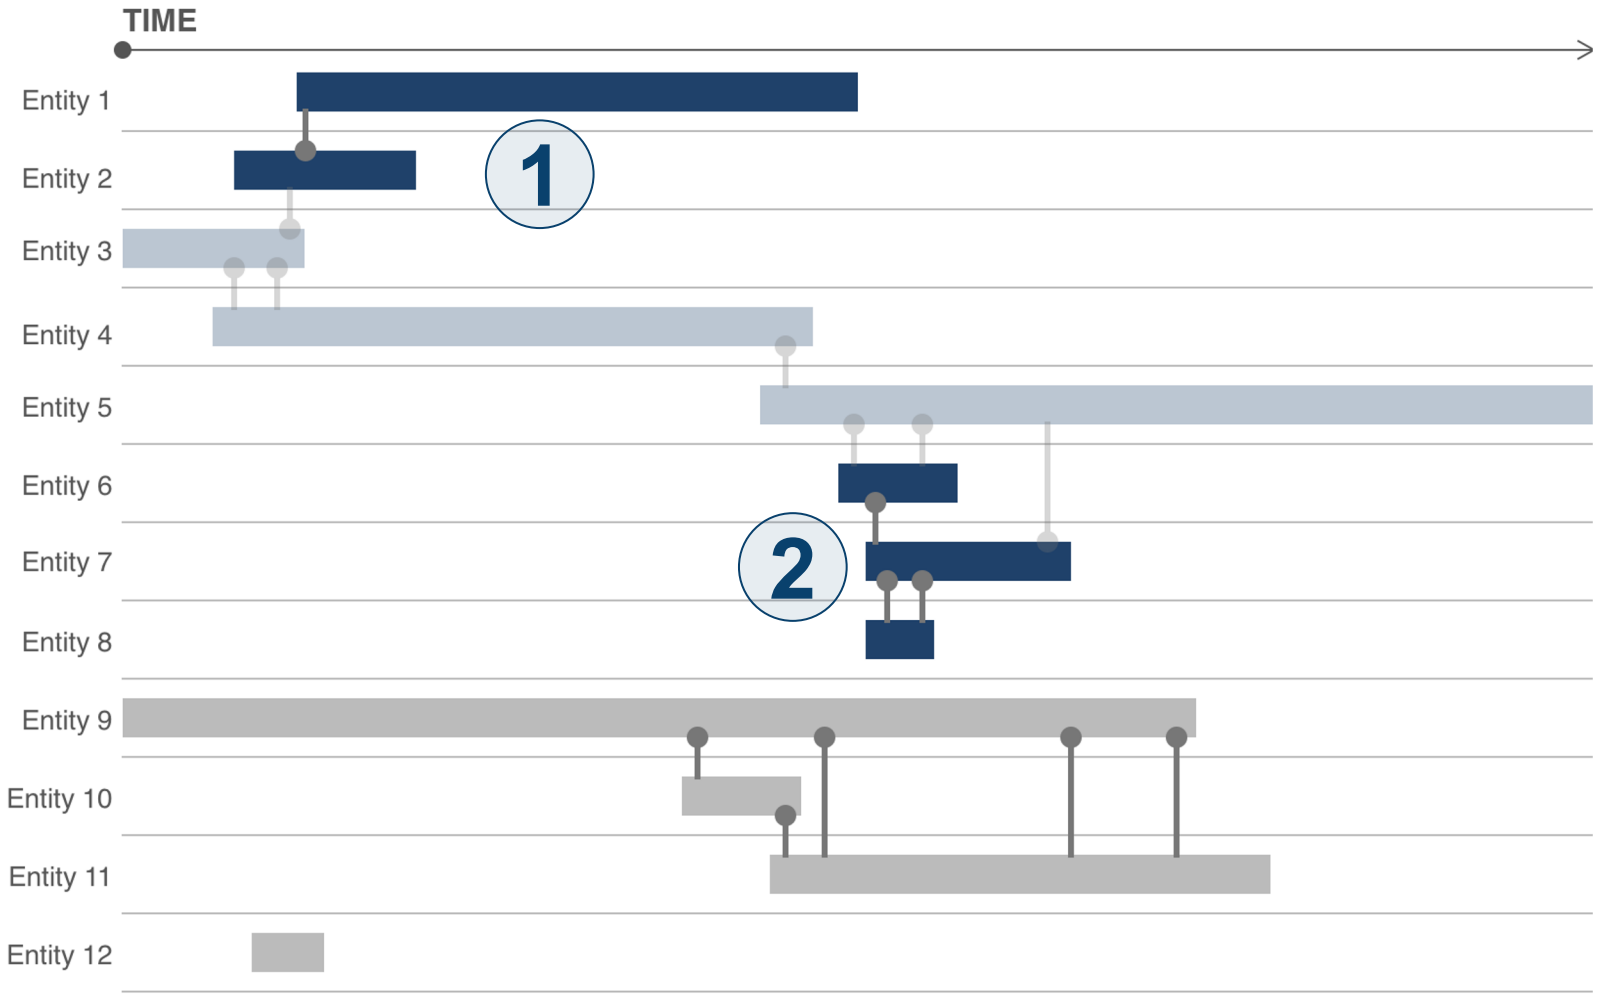

Featurizing Incident Report Dynamic Graphs

Calculate scores

- Component scores
- Branch scores
- Entity score

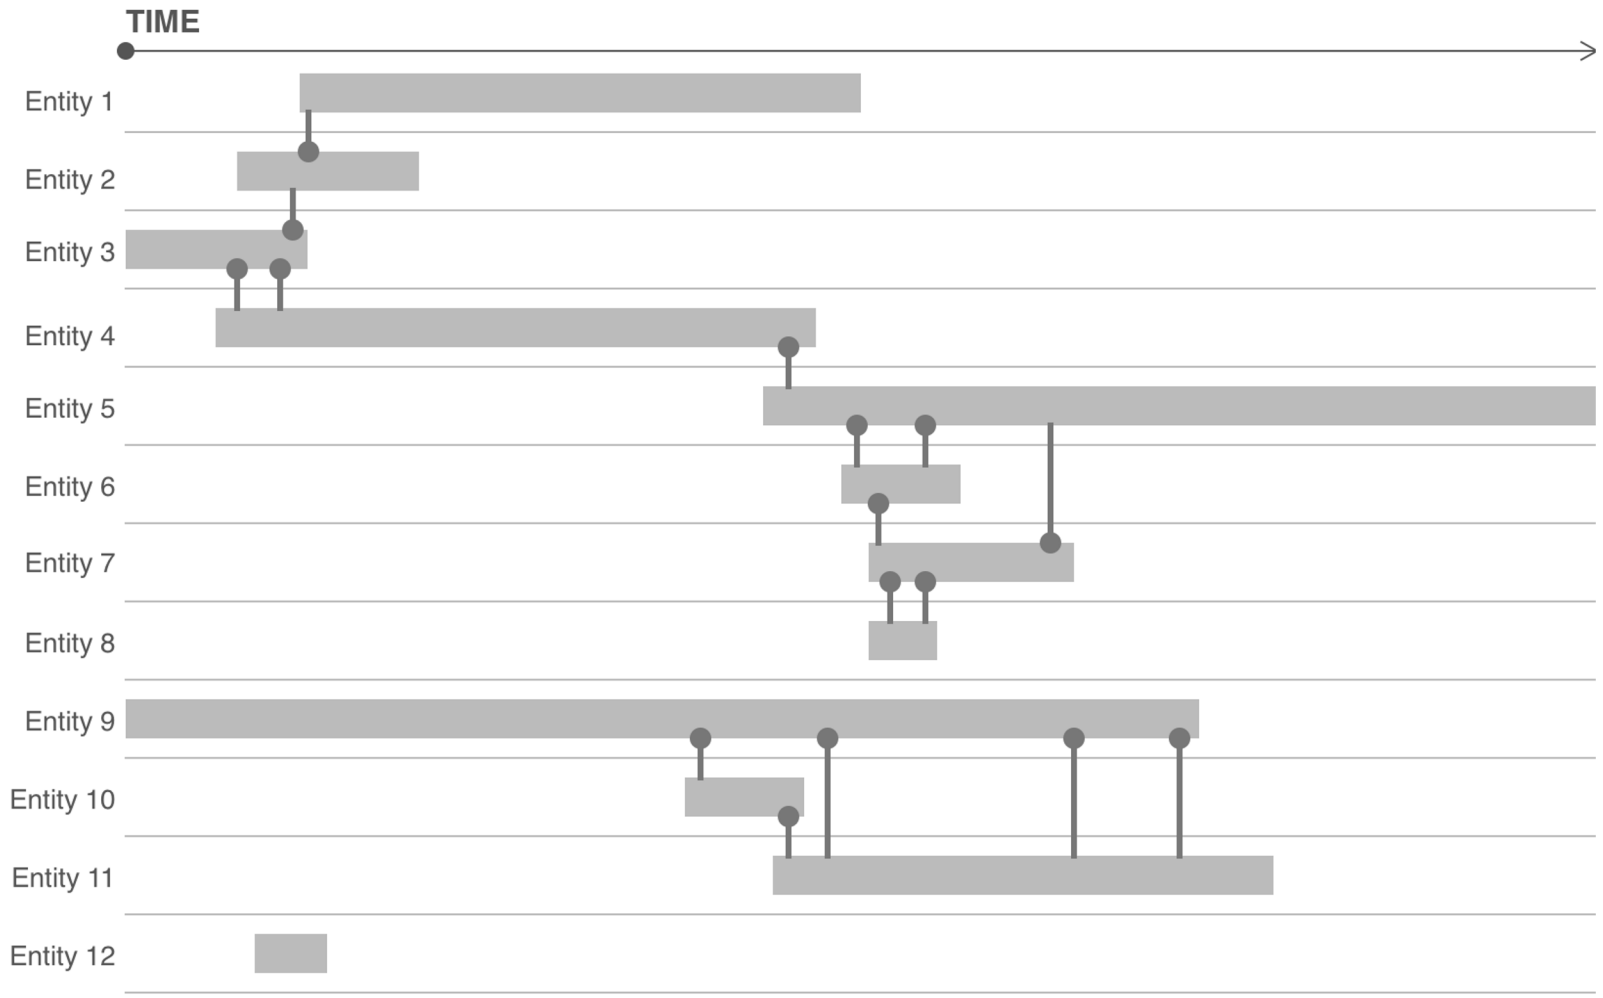

Featurizing Incident Report Dynamic Graphs

Calculate scores

- Component scores
- Branch scores
- Entity score
  - Severity from a cyber security analytic

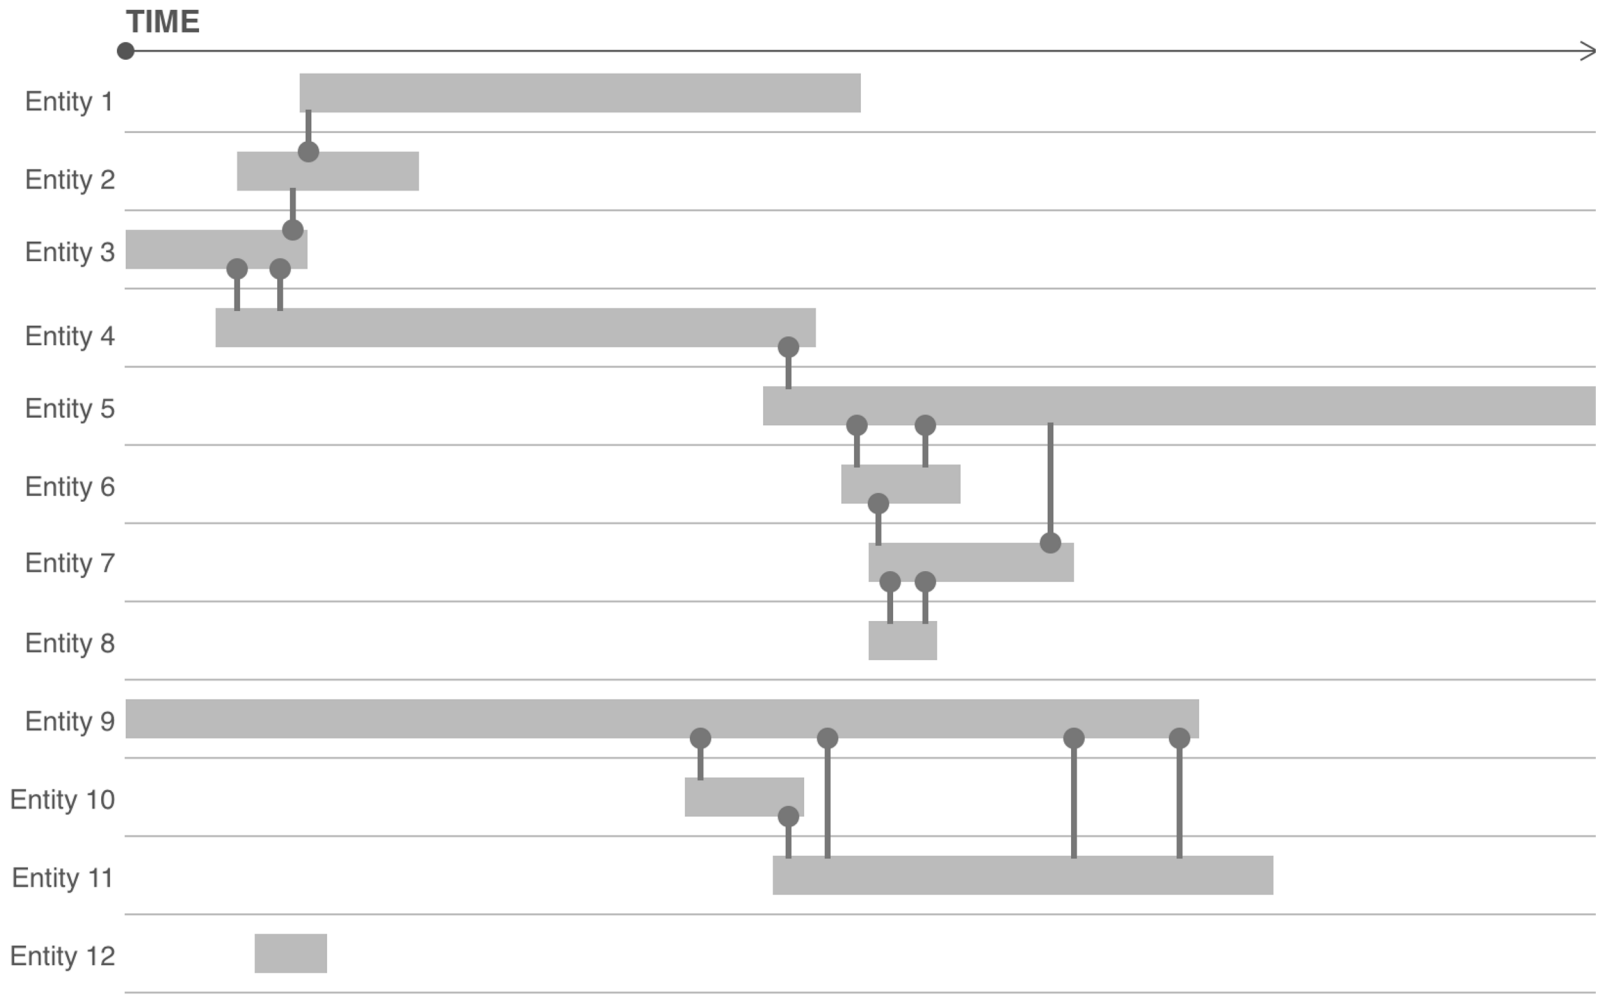

Summarizing Incident Reports: Naive Approach

Average and filter

- Remove an entity if:

  For a summarization threshold *t:*
  mean(*component scores*) < *t*, or
  mean(*branch & entity scores*) < *t*

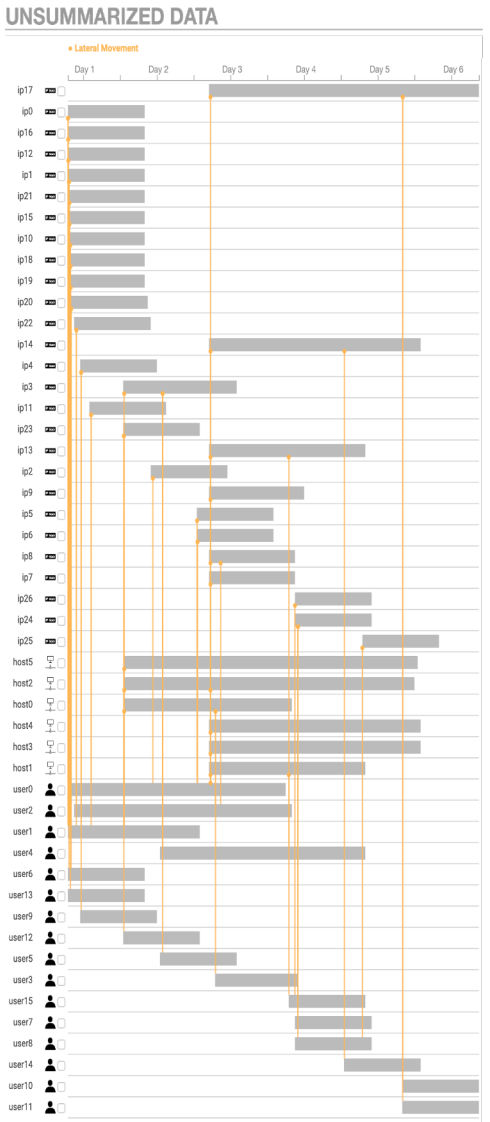

SUMMARIZE

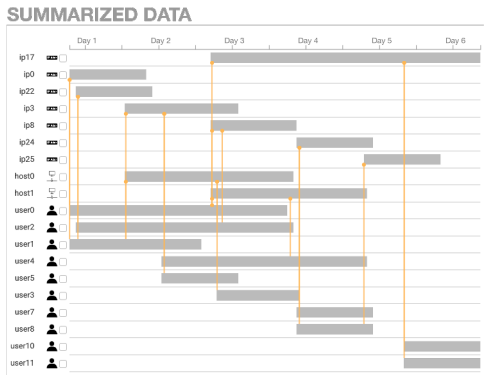

Summarizing Incident Reports: Hierarchical Approach

Data and Challenges

- Data from 2 Red Team events against monitored network with known detectors (ground truth!)
- Small data
  - 460 observations
- Lots of *structure*
  - 2 RT events, 2 detectors, 15 reports
  - **entities** within **branches** within **components** within **reports**
- Heterogeneous covariate availability
  - correlated with detector

Summarizing Incident Reports: Hierarchical Approach

Data and Challenges

- Data from 2 Red Team events against monitored network with known detectors (ground truth!)
- Small data
  - 460 observations
- Lots of *structure*
  - 2 RT events, 2 detectors, 15 reports
  - **entities** within **branches** within **components** within **reports**
- Heterogeneous covariate availability
  - correlated with detector

Approach: Bayesian Hierarchical Model

- Small data: priors and structure instead of "just throw it in a NN"
- Structure: covariates at the entity, branch, component, detector, and RT event levels
  - In practice omitted RT event effects; too few to matter
- Detector-specific data modeled with...detector-specific (entity-level) models

Summarizing Incident Reports: Model Details

- Varying-intercept, fixed slopes model
- entity model: f(intercept, type, MITRE location)
  - detector-specific data modeled via interactions
- branch model: f(core sequence, duration, connections…)
- component model: f(duration, entities, relationships, timestamps)
- logistic link
- scaled inverse-Wishart distribution for priors over related within-level coefficients

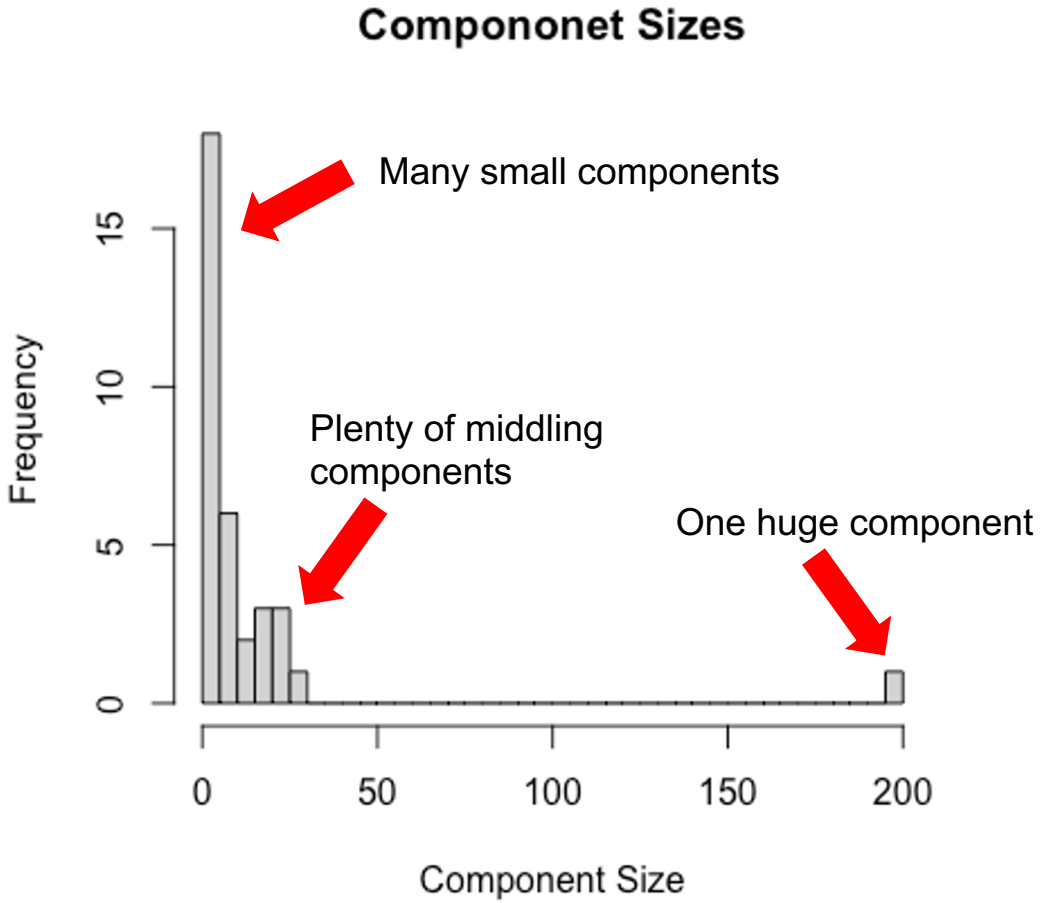

Summarizing Incident Reports: Alternatives?

- A flat model with fixed effects?
  - Low bias, high variance
  - Zero degrees of freedom in many components/branches
- Feed forward neural network?
  - Small data
  - Unclear how to leverage structure
  - Finger-cross strategy for missing data

Summarizing Incident Reports: Alternatives?

- A flat model with fixed effects?
  - Low bias, high variance
  - Zero degrees of freedom in many components/branches
- Feed forward neural network?
  - Small data
  - Unclear how to leverage structure
  - Finger-cross strategy for missing data

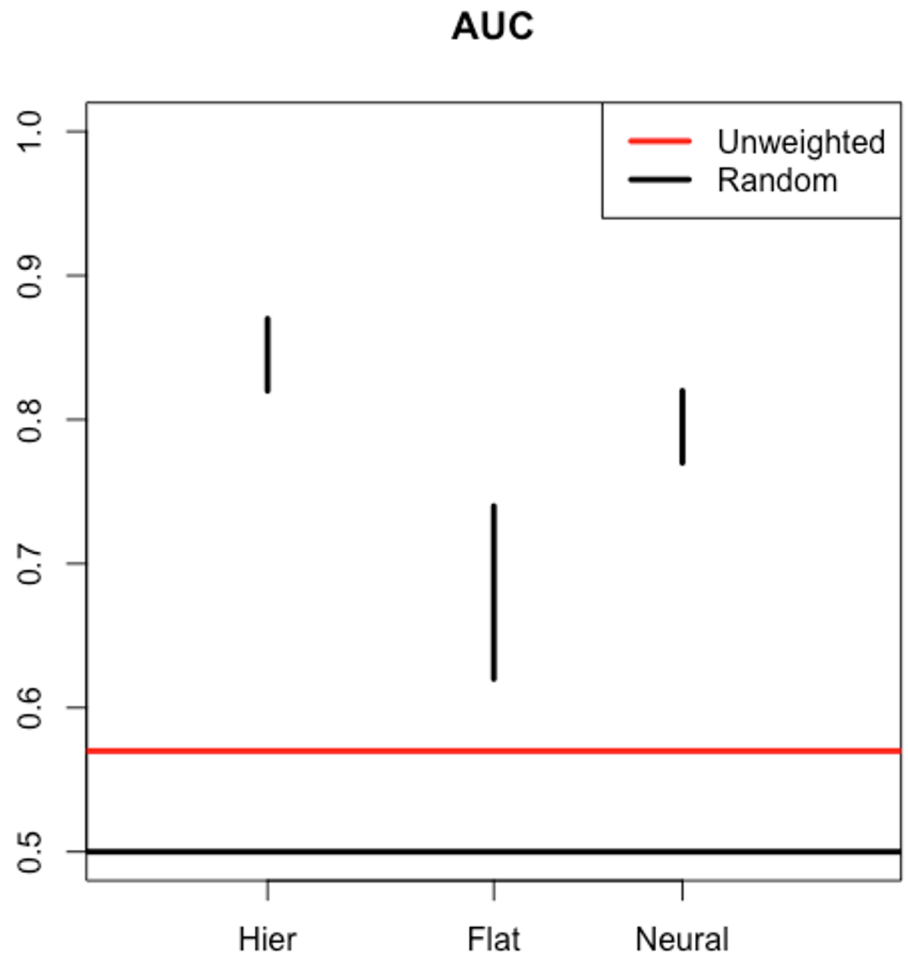

Evaluating Summarization Performance

Goals:

- Size: enable the dynamic creation of smaller incident reports
- Accuracy: don't drop entities that are true positives

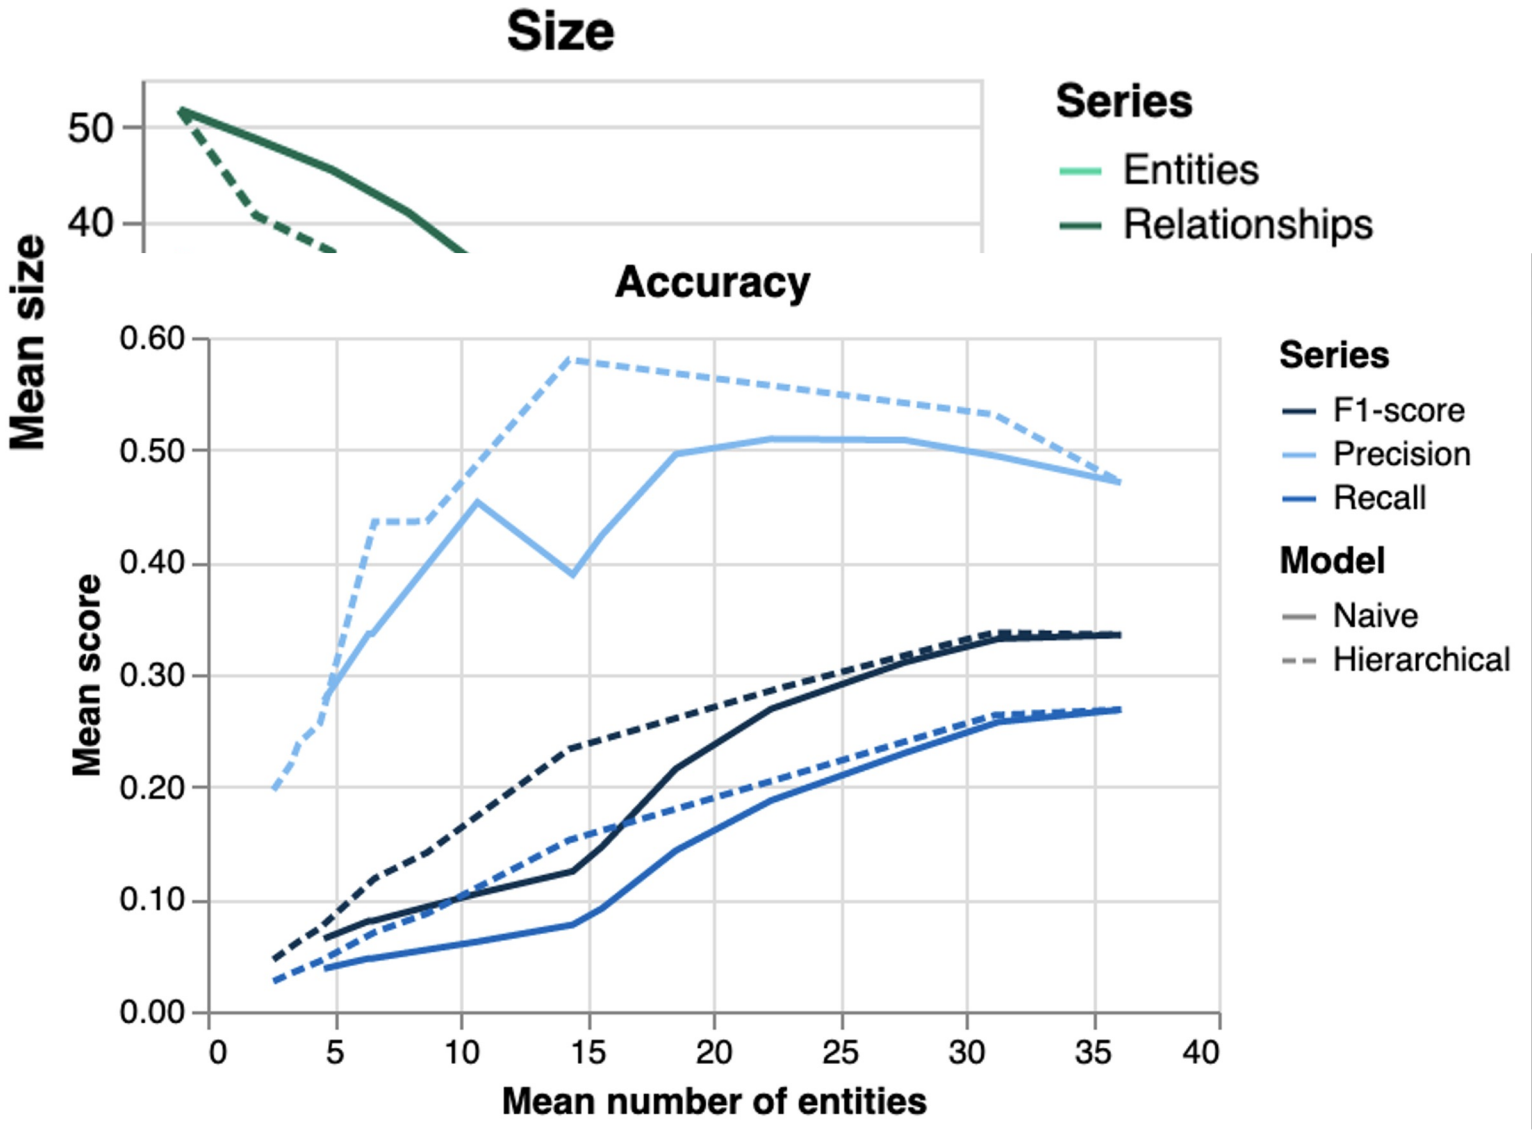

Qualitative Feedback

In live testing in a real environment, a SOC lead gave the following feedback:

He liked the visualization design, saying **“I feel like I can look at this and get an understanding of the key parts faster”** compared to looking at the tables of data contained in typical incident reports. Regarding the summarizations, he commented **“you’re going to save me a bunch of time”** compared to analyzing unsummarized incident reports.

Future Work

- Generalizability
- Cross-tool amalgamation
- Package/deployment

Visualizing incident reports is useful.

Moderate ML effort allows you to accurately summarize incident reports as well.

Robert Gove

robert.gove@twosixtech.com

@rpgove

Nathan Danneman

nathandanneman@datamachines.io

twosix TECHNOLOGIES

datamachines