ROADMAP (Relational Object Analysis Drives Multimodal Attack Prediction) System Faces Big Data Challenges

Rebecca Cathey
BAE Systems
July 12, 2016

Our Partners

DigitalOperatives

StratumPoint
Operationalized Risk Management

University of Maryland

ROADMAP Aims to Predict Cyber-Attacks with High Fidelity

- High fidelity predictions specify attacker, victim, attack type, & timeframe
- Primarily based on publicly available data … which is where the big data challenge arises

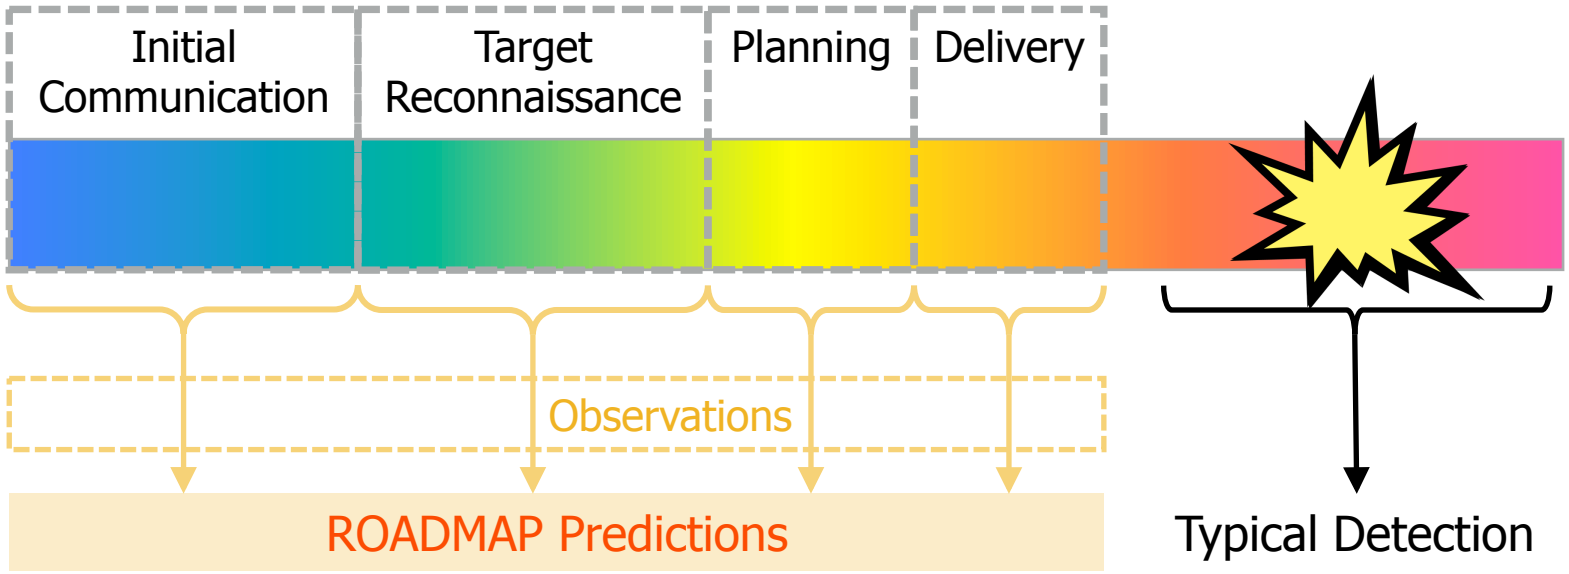

ROADMAP Features a Multi-Stage Big Data Architecture

ROADMAP provides accurate and timely forecasts of cyber-attacks through continuous, real-time, global monitoring of diverse external – publicly available – data

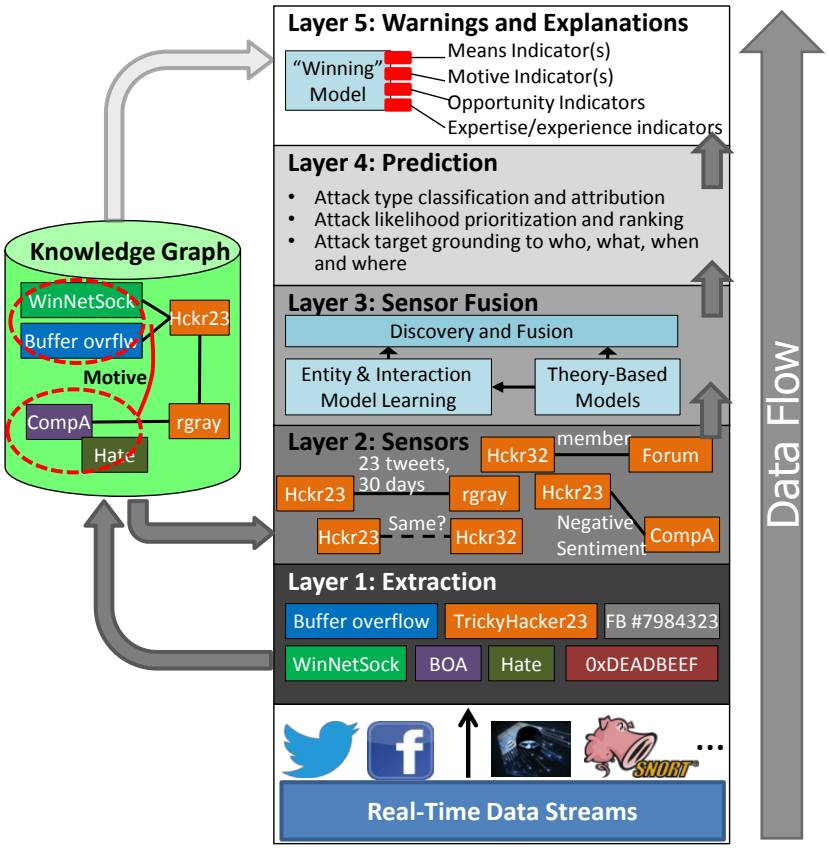

5. **Warnings & Explanations**: Generates warnings with source, victim, attack type, & date traceability
4. **Prediction**: Evaluates intermediate hypotheses; Generates globally consistent attack prediction
3. **Sensor Fusion**: Fuses sensor inputs via attack models; Generates attack prediction hypotheses
2. **Sensors**: Applies environmental and behavioral theories to generate indicators of attacks
1. **Extraction**: Extracts entities & relations from multi-modal streaming data to populate knowledge graph

ROADMAP Faces Multiple Big Data Challenges

Layered architecture filters data to progressively refine information into reliable warnings

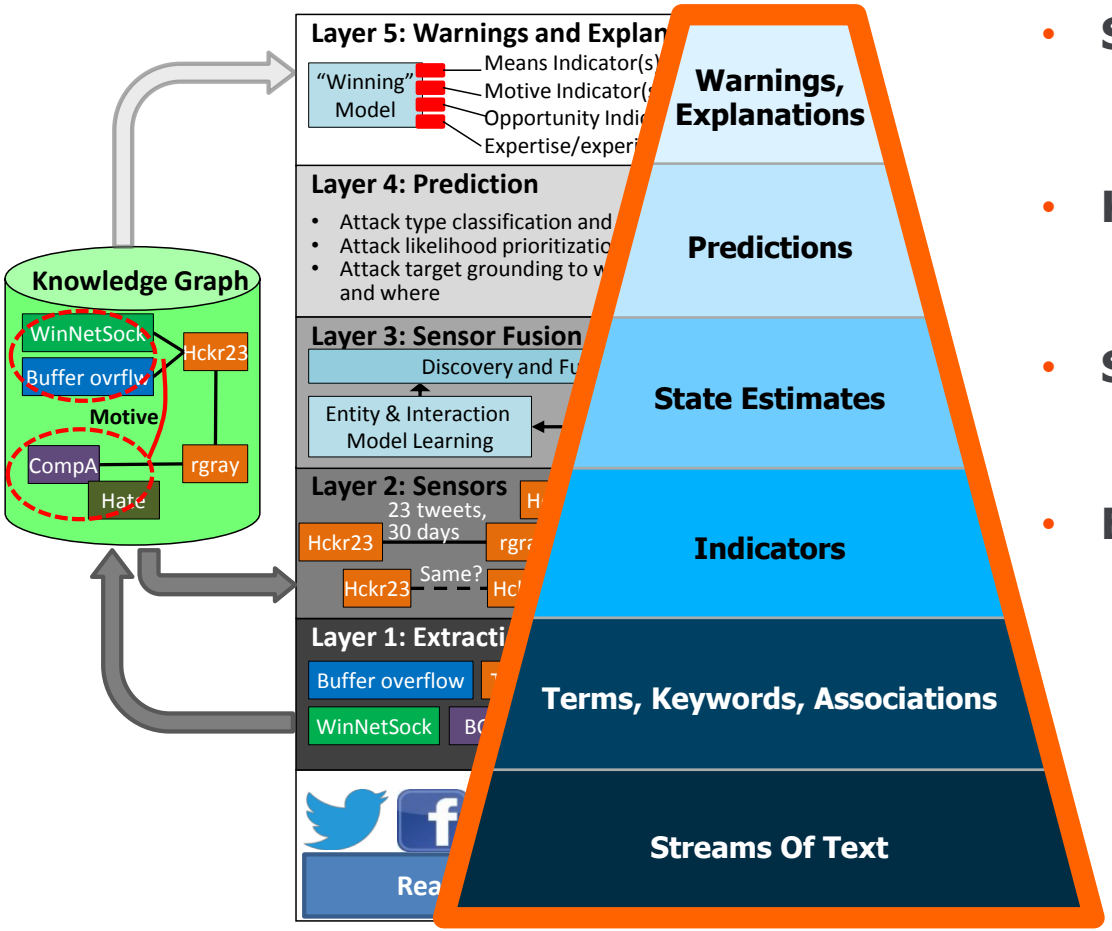

- **Sensor Fusion & Prediction layers**
  - Distributed learning & fusion
  - Real-time prediction generation
- **Knowledge Graph**
  - Large number of ambiguous entities & relationships
- **Sensor layer**
  - High volume of associations and observations
- **Extraction layer**
  - High volume & velocity of streaming data
  - Multiple, heterogeneous sources
  - Out-of-order data & pedigree

throughout ROADMAP

Extraction Layer Populates Knowledge Graph with Entities & Relations

- Extracts names of organizations, people, locations, & abstract concepts (e.g., events & technical concepts)
- Identifies direct associations between extracted data items (representative)
  - Per Tweet = 1+(2*#terms)+#followers +#followees
  - Per Post = (#users*avg#keywords) +#answers
- Large expected data volume (representative)
  - Twitter: ~6k / second; 500M / day
  - Facebook: >55M status updates / day
  - Stack Exchange: ~40k questions / week; ~60k answers / week

Airline Tweet Example

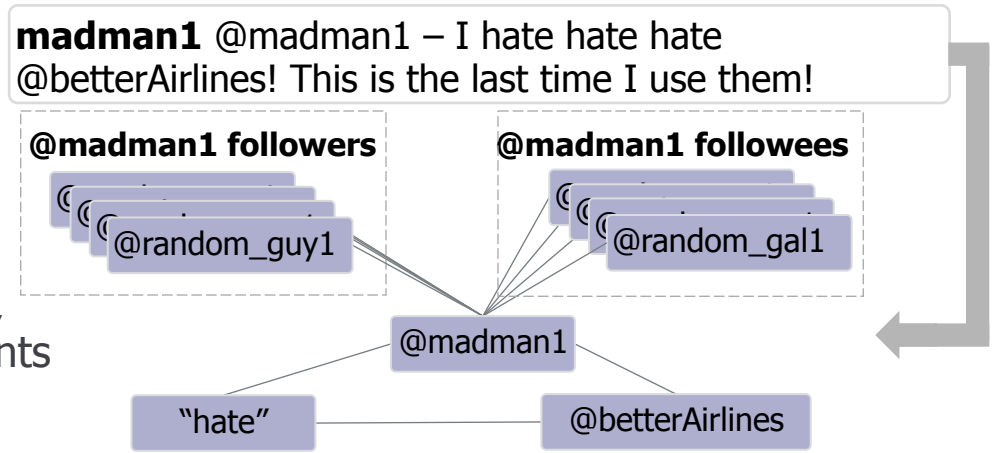

Forum Post Example

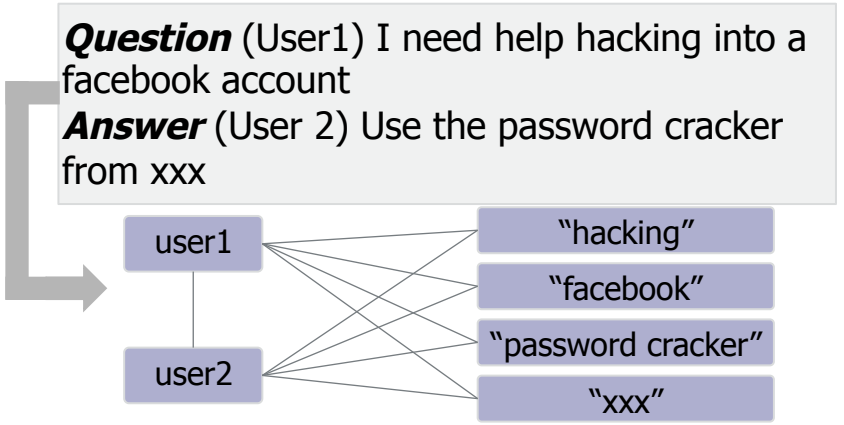

Extraction Layer Processes Multiple Free Text Data Sources

- Extract entities from full text
  - Automatic identification of jargon terms relating to specific types of attack
  - Train using historical documents based on attack types & multiple topics / domains
- Extract relationships from metadata
  - Find relationships between individuals & work items, snippets of code & malware
  - Find individuals who serve specific roles within the cyber-attack lifecycle
    - Researcher, reconnaissance operative, developer, or attacker
    - Detect alias detection, relationship correlation

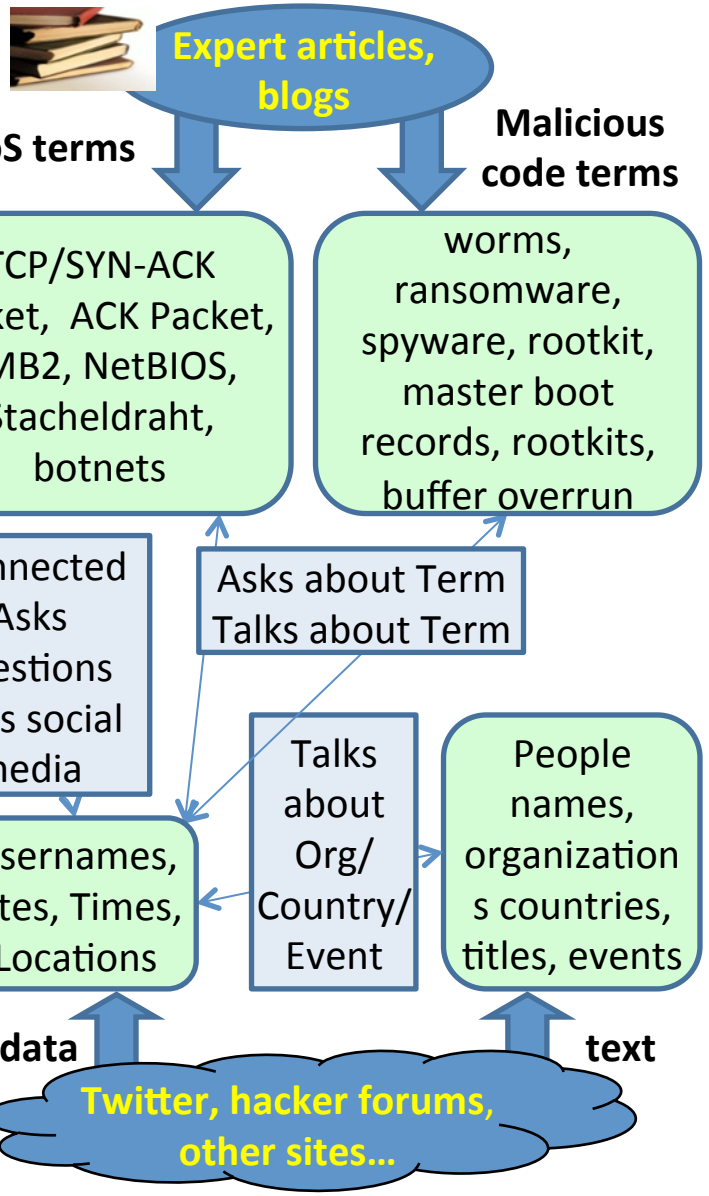

Sensor Layer Finds Attack Indicator Patterns in Knowledge Graph

- Multiple sensors reason over Knowledge Graph to identify indicators of predictive events across multiple actor & cyber-attack types
- Ingested data include many examples with potential to indicate a pending attack
  - 5,160 tweets about airlines / day
  - 229 discussion threads / day about DDoS attacks & DDoS-related activities on hacker forums
  - 15k-30k observations of trending public perception
- Sensors measure characteristics of entities such as organizations, technical concepts, & snippets of code
- Sensors measure attacker means, motivation, opportunity, experience, & expertise

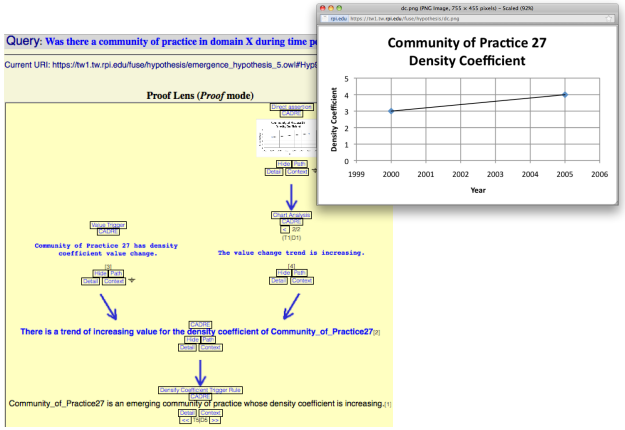

Identify patterns in Knowledge Graph

Data source	Indicator Processing
Web of Science	317 ind./sec
LNUSP	168 ind./sec
LNCN	65 ind./sec
CNKI	229 ind./sec

LNUSP = English patents
LNCN = Chinese patents
CNKI = Chinese Science Articles

Sensors Search for Predictive Signals of Pending Attacks

- Sensors identify signals indicative of attacks as patterns in Knowledge Graph
  - Wagon wheel connections, indirect associations (len.<=2)
  - Sentiment sensor: connections to negative terms
  - Vulnerability sensor: wheels with exploits at center
  - Apply templates to analyze communities

Example Sensor	Indicators	Measures
Social Networking	• Connections, contacts between people • Identification of frequent users, experts, friends	Experience, Means
Sentiment	• Sentiment-related characteristics of individuals or groups • General trends including growth of negative sentiment toward target organization	Motivation
Vulnerability	• Vulnerabilities / exploits being built • Which users learn about which exploits • Experience level of each user with each exploit	Expertise, Means
(ROADMAP includes many other sensors)		

Sensor Fusion Provides Attack State Estimates to Prediction for Decision to Predict Attack

- Sensor Fusion models entities & their interactions evolving over time searching for attack precursor behavior
- Anticipated modeling and fusion combinatorics & behavior complexity are daunting
  - Millions of entities – only a fraction involved in attack-related activity
  - Billions of interactions – most not attack-related, but may become so
  - Hundreds of attack types & variants – increasing all the time
- Prediction assesses whether evidence of attack imminence is strong enough to make formal prediction

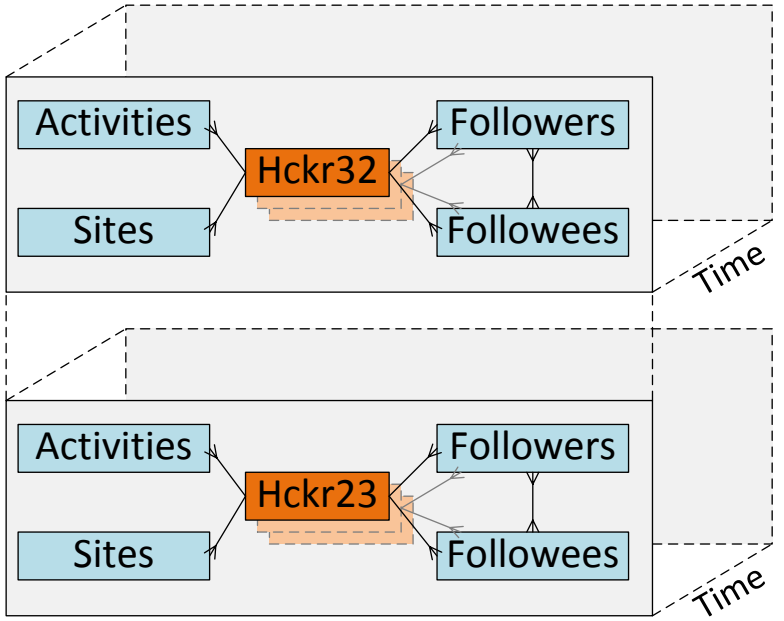

Sensor Fusion Addresses Multiple Attack Prediction Challenges

- Theory-based models developed from academic theories or operational expertise assist in two ways
  - Produce measure of threat-related activity based on expert opinion
  - Generate synthetic training data to overcome scarcity of historical data for certain attack types
- Entity / interaction modeling extracts relevant hidden relationships between entities & monitors their evolution over time
  - Temporal & spatial analysis enables informative warnings & associated causes of attack estimates
  - Provides continually evolving attack state estimates

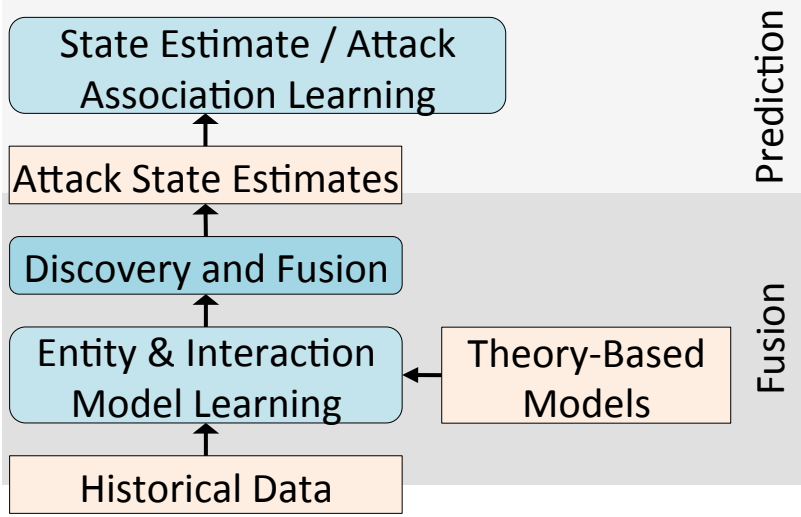

ROADMAP Model Learning

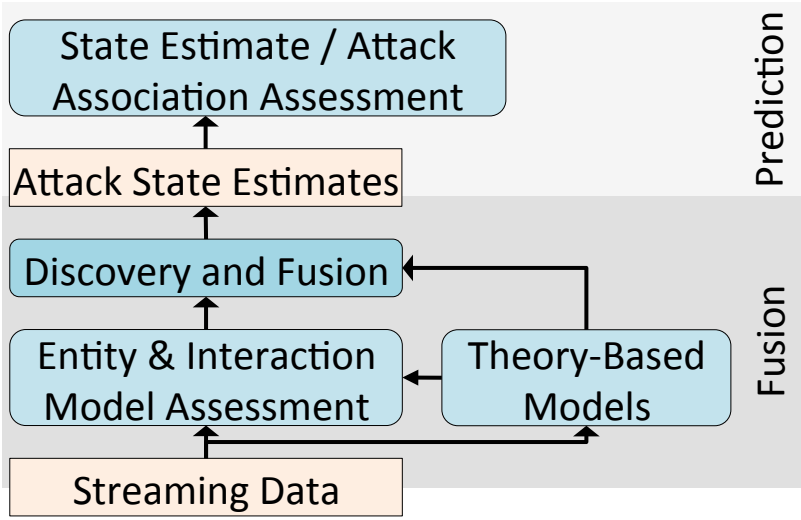

ROADMAP Model Exploitation

Prediction Determines When to Formally Make a Prediction, Which Generates a Warning

- Prediction generates formal attack predictions
  - Learns temporal dependencies between Fusion attack state estimates & attacks
  - Projects forward in time to estimate future attack probability & timing
  - Issues prediction when future likelihood exceeds a threshold parameter
- Warning & Explanation generates warnings from formal attack predictions
  - Provides attack source, victim, attack type, & date
  - Includes direct association to attacker means, motive, opportunity, expertise, & experience

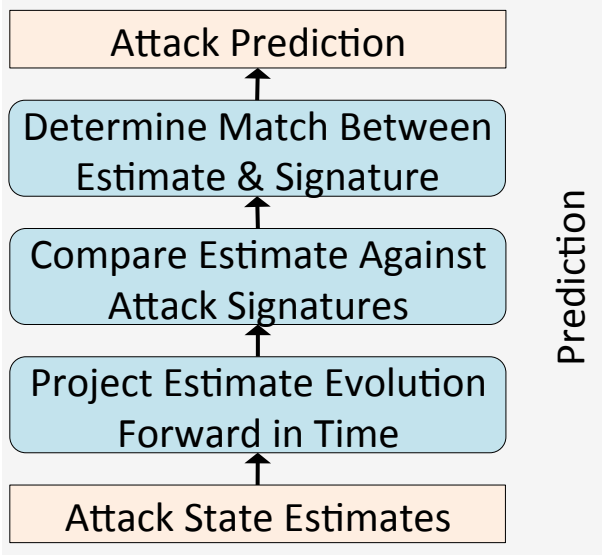

Layer 5: Warnings and Explanations

Supra-Threshold Prediction
- Means Indicator(s)
- Motive Indicator(s)
- Opportunity Indicators
- Expertise /Experience Inds.

Layer 4: Prediction

- Attack type classification & attribution
- Attack likelihood prioritization & ranking
- Attack target grounding to who, what, when, & where

Conclusion: ROADMAP Addresses Myriad Big Data Challenges with Layered, Modular Approach

- ROADMAP handles large numbers of entities and associations through a rich, uniformly represented knowledge graph
  - Simplifies data extraction and sensor detection of indicator patterns
  - Supports prediction activities without sacrificing timeliness or accuracy
- ROADMAP captures patterns and detects attacks via learned spatial and temporal patterns, allowing detailed analysis of causes of attacks
  - Incorporates expert knowledge and accommodates attacks with low historical data via theory-based models
  - Provides detailed warnings with low little leading signal available for sophisticated attackers
- ROADMAP weeds out false positives by analyzing patterns over time and looking forward to predict the likelihood of attacks

Questions?