SpacSec 2026

Formalization of the semantic gap in projecting cyber threat models onto computational forensics tasks

Evgeny Abramov and Margarita Knyazeva

Southern Federal University

Motivation of the paper

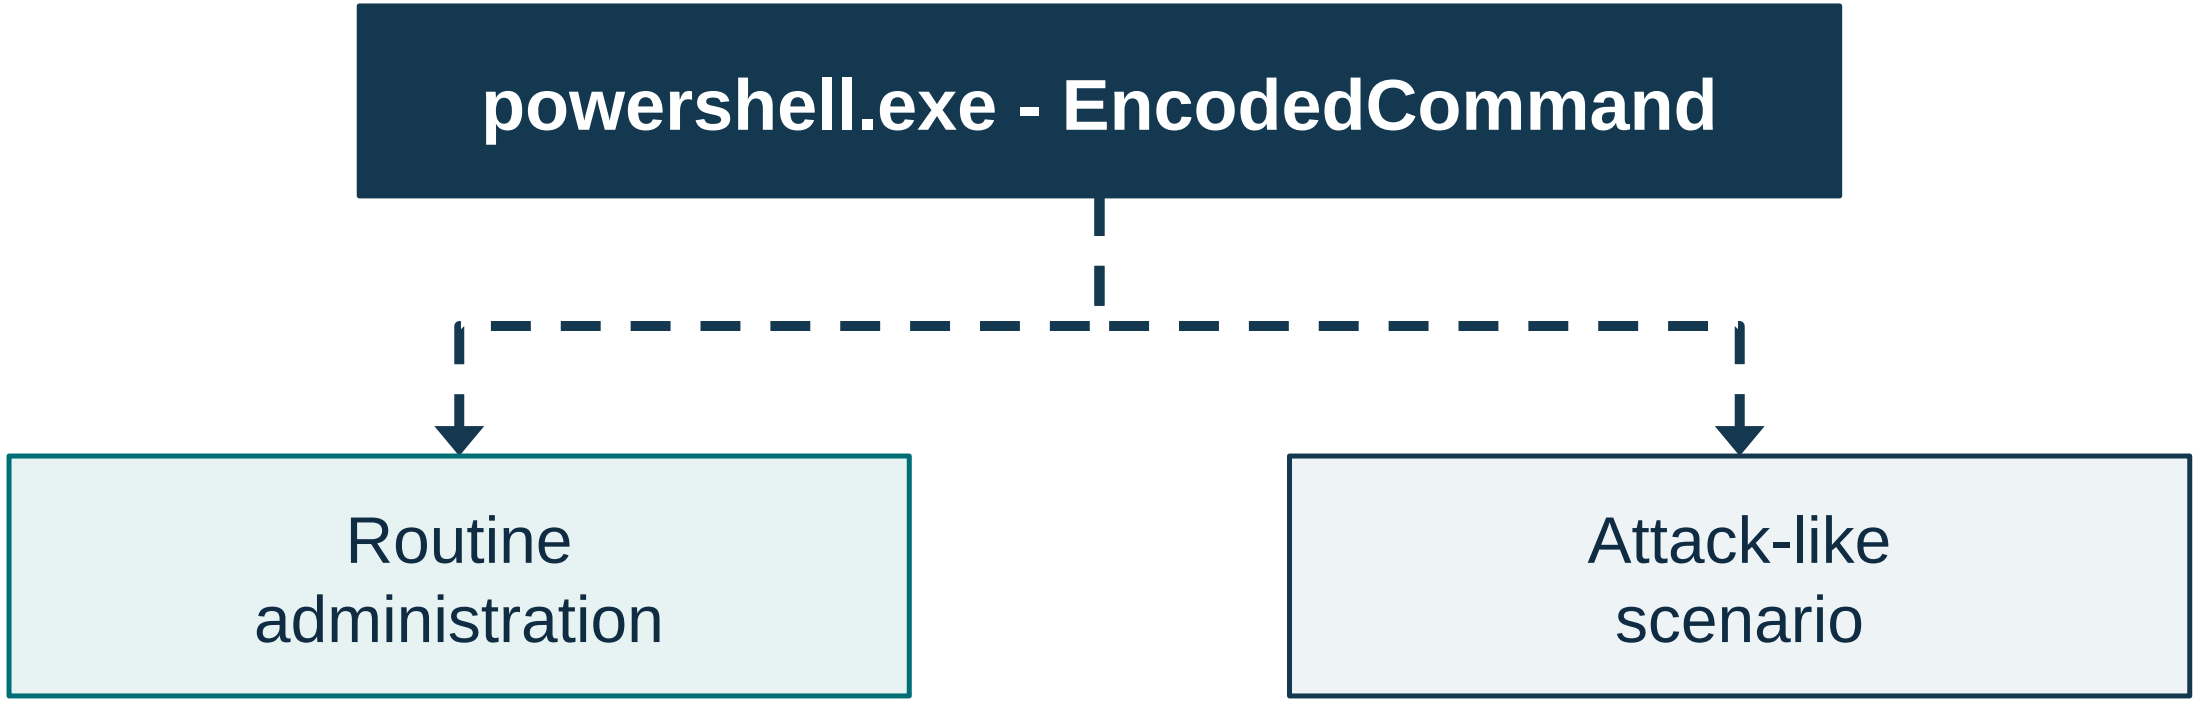
1. Show how can a CTI description of possible behavior become a forensic technical claim whose grounds are checkable against acquired digital traces.

2. Introduce three elements of semantic gap: temporal (attack stages vs existing traces), attributive (strategic attribution vs technical entities and observed data) and contextual (behavioral label vs execution context).

3. Show the gap between data volatility and observed data.

Encoded PowerShell

Candidate technique ATT&CK T1059.001

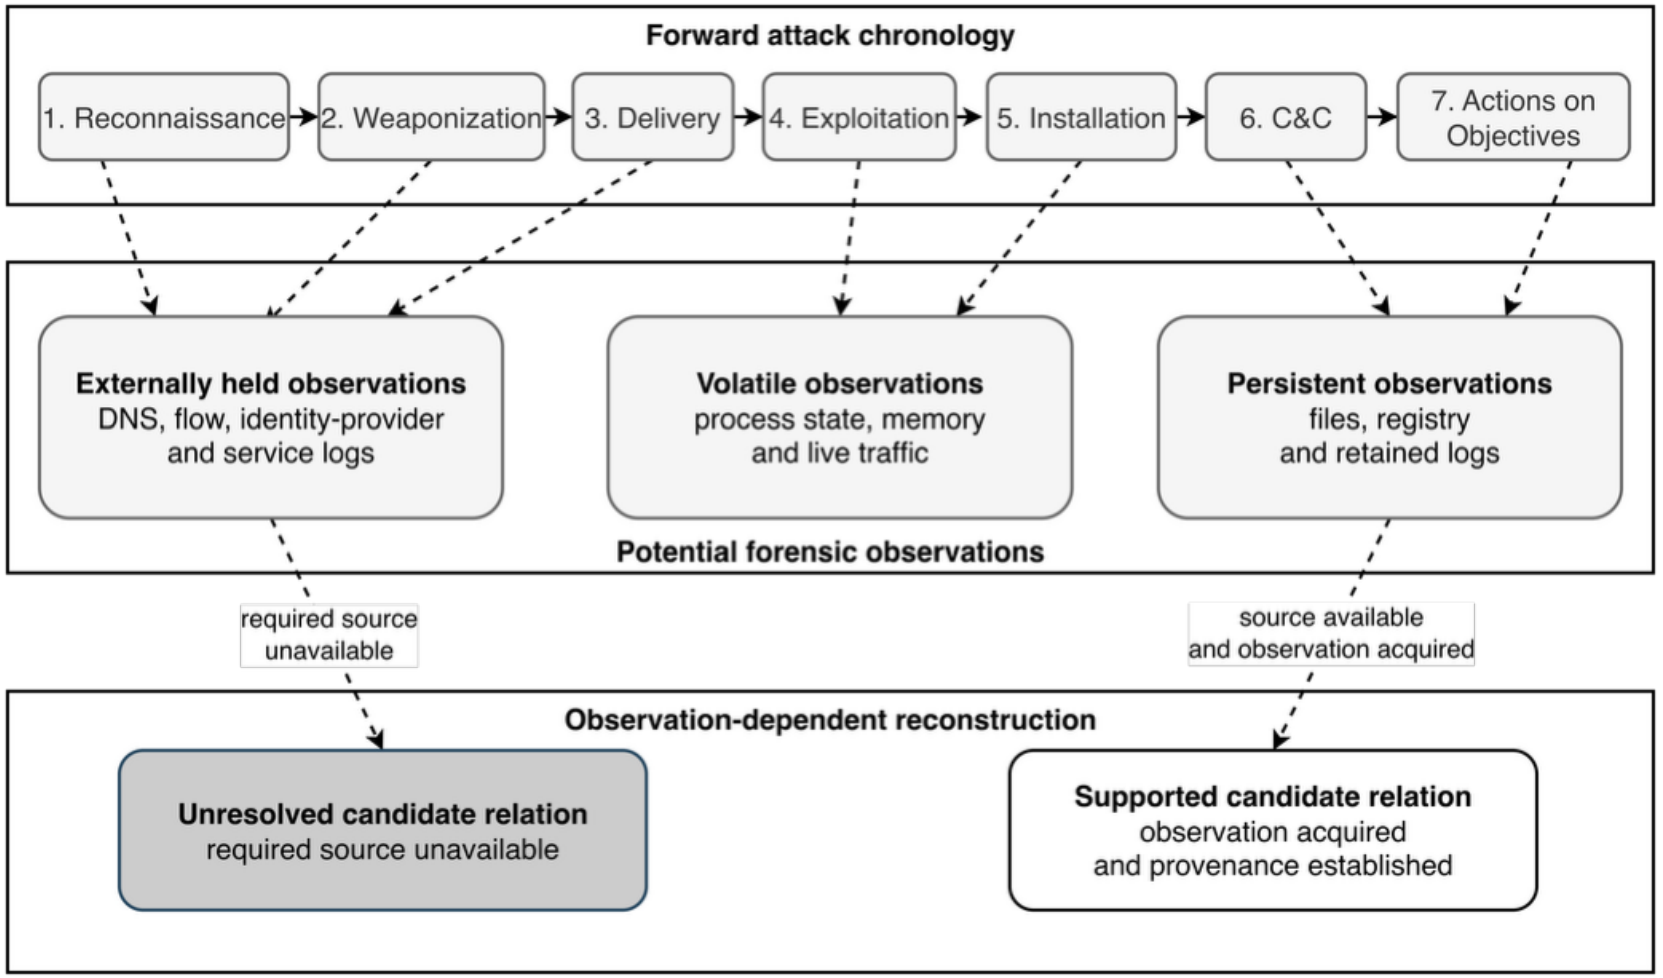


Which observations distinguish the explanations?

Three axes of the semantic gap

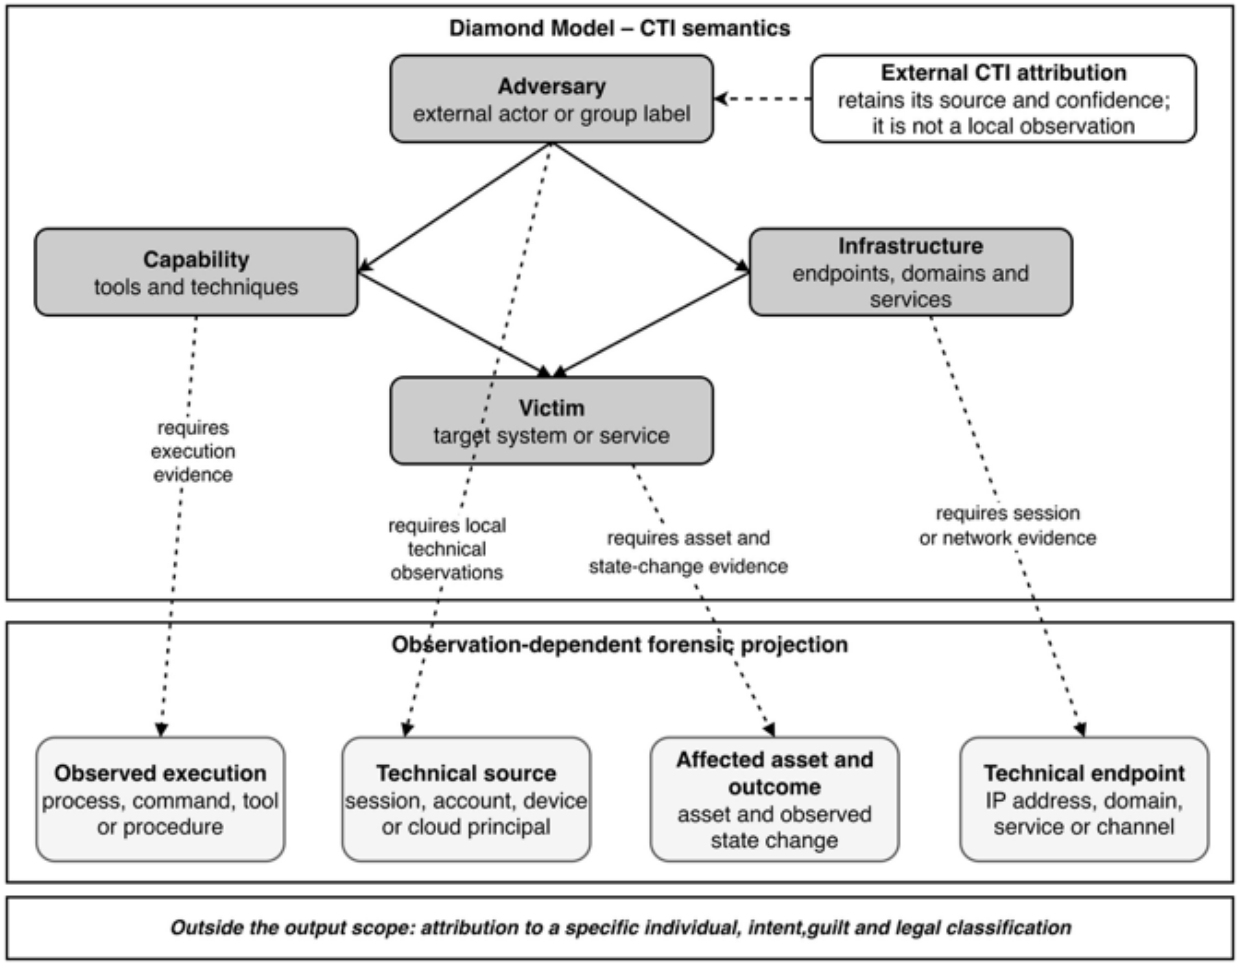
	CTI		Forensic task
Cyber Kill Chain	Attack stages	≠	Surviving traces
Diamond Model	Strategic attribution	≠	Technical entities, observed data
MITRE ATT&CK	Behavioral label	≠	Execution context

Chronology and observability



Examination starts from acquired traces

Attribution levels

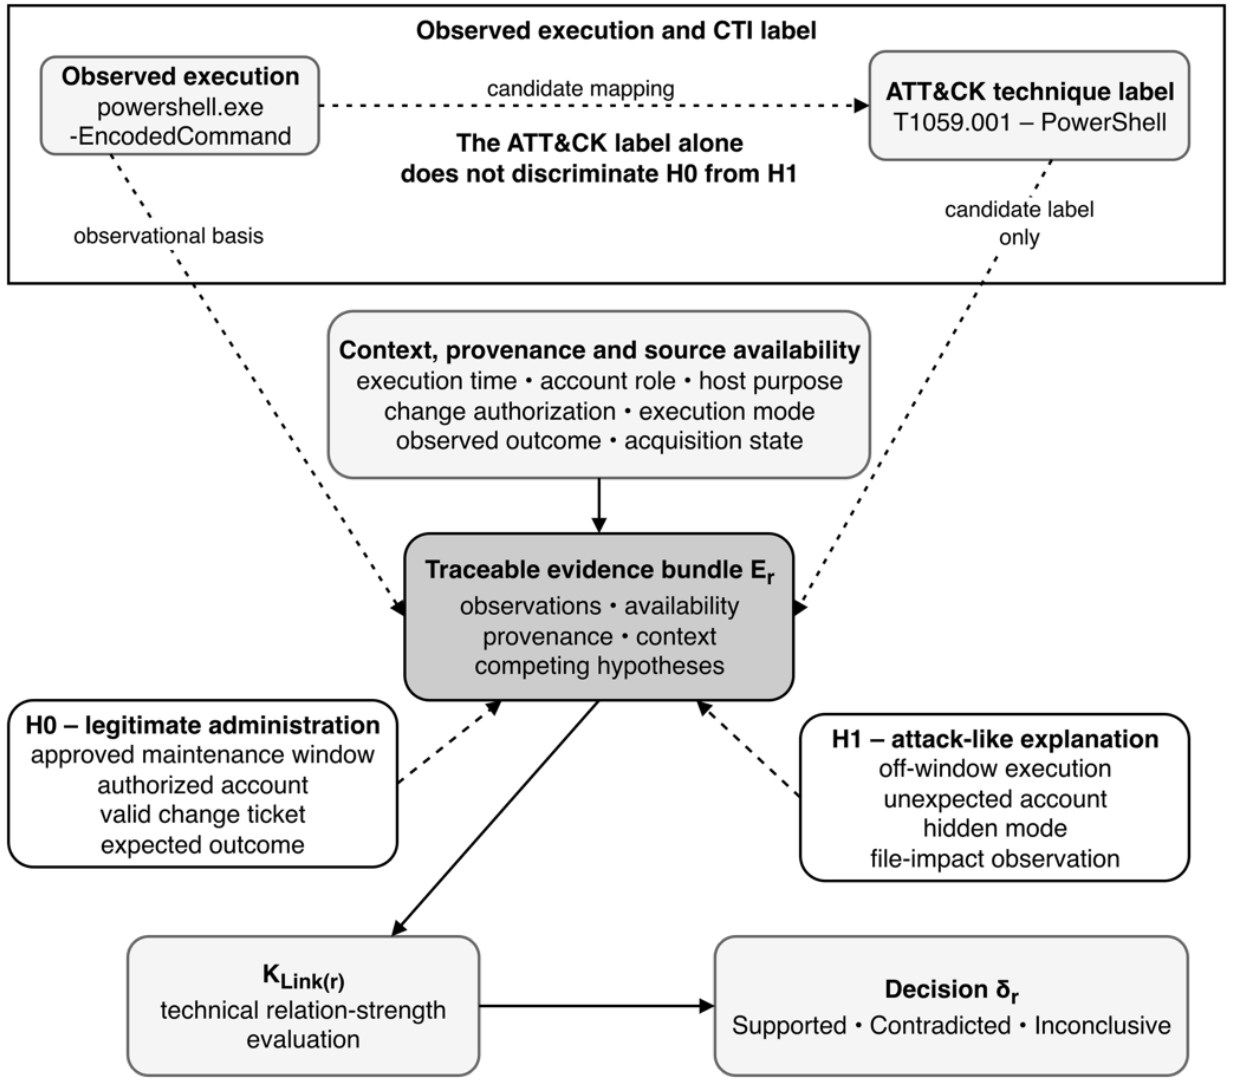


Observed entities:

Session
Account
Device
Process

Human identity remains outside scope

Dual use and context



Execution context:
Timing
Account role
Authorization

Threat scenarios with encoding signals malicious intent:

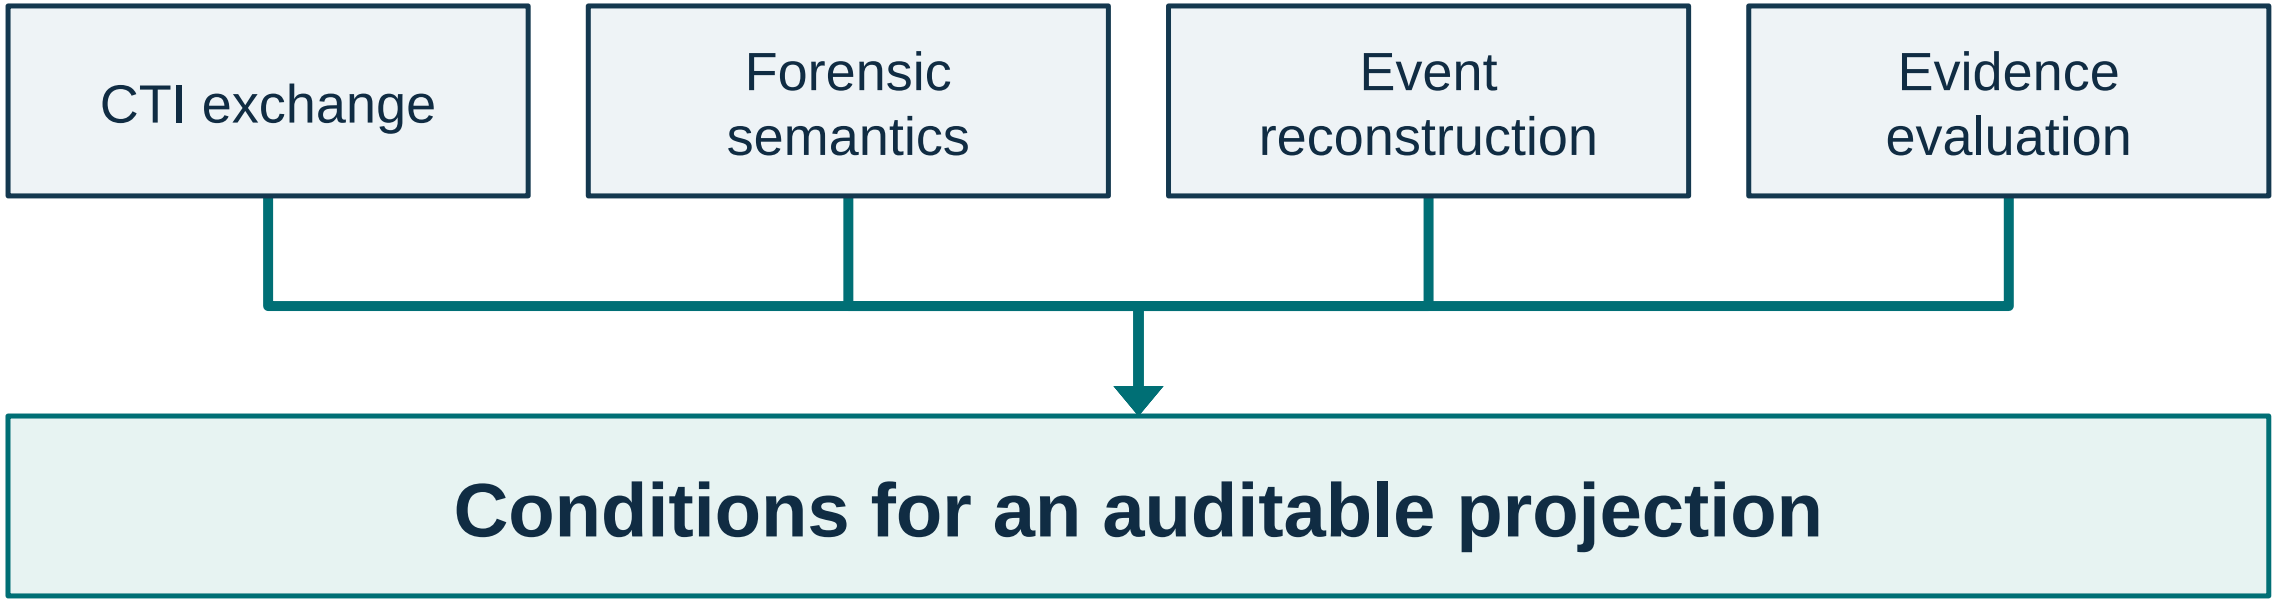
- Office macros dropping or executing inline encoded stagers
- Phishing attachments spawning powershell.exe as a child of winword.exe
- Metasploit/Cobalt Strike post-explotation modules
- Ransomware reconnaissance scripts and pre-deployment staging
- Encoded persistence mechanisms from supply chain compromises

Legitimate scenarios that produce the sane alert:

- Intune script execution via Management Service
- Vendor remote support tools (Kaseya, ConnectWise)
- CI/CD pipeline runners (Azure DevOps, GitHub Actions)
- EDR/AV update scripts and IT automation frameworks (Ansible WinRM)
- Scheduled tasks created by legitimate software installers.

Each fact retains its source and procedure

The verification interface



Origin or Provenance

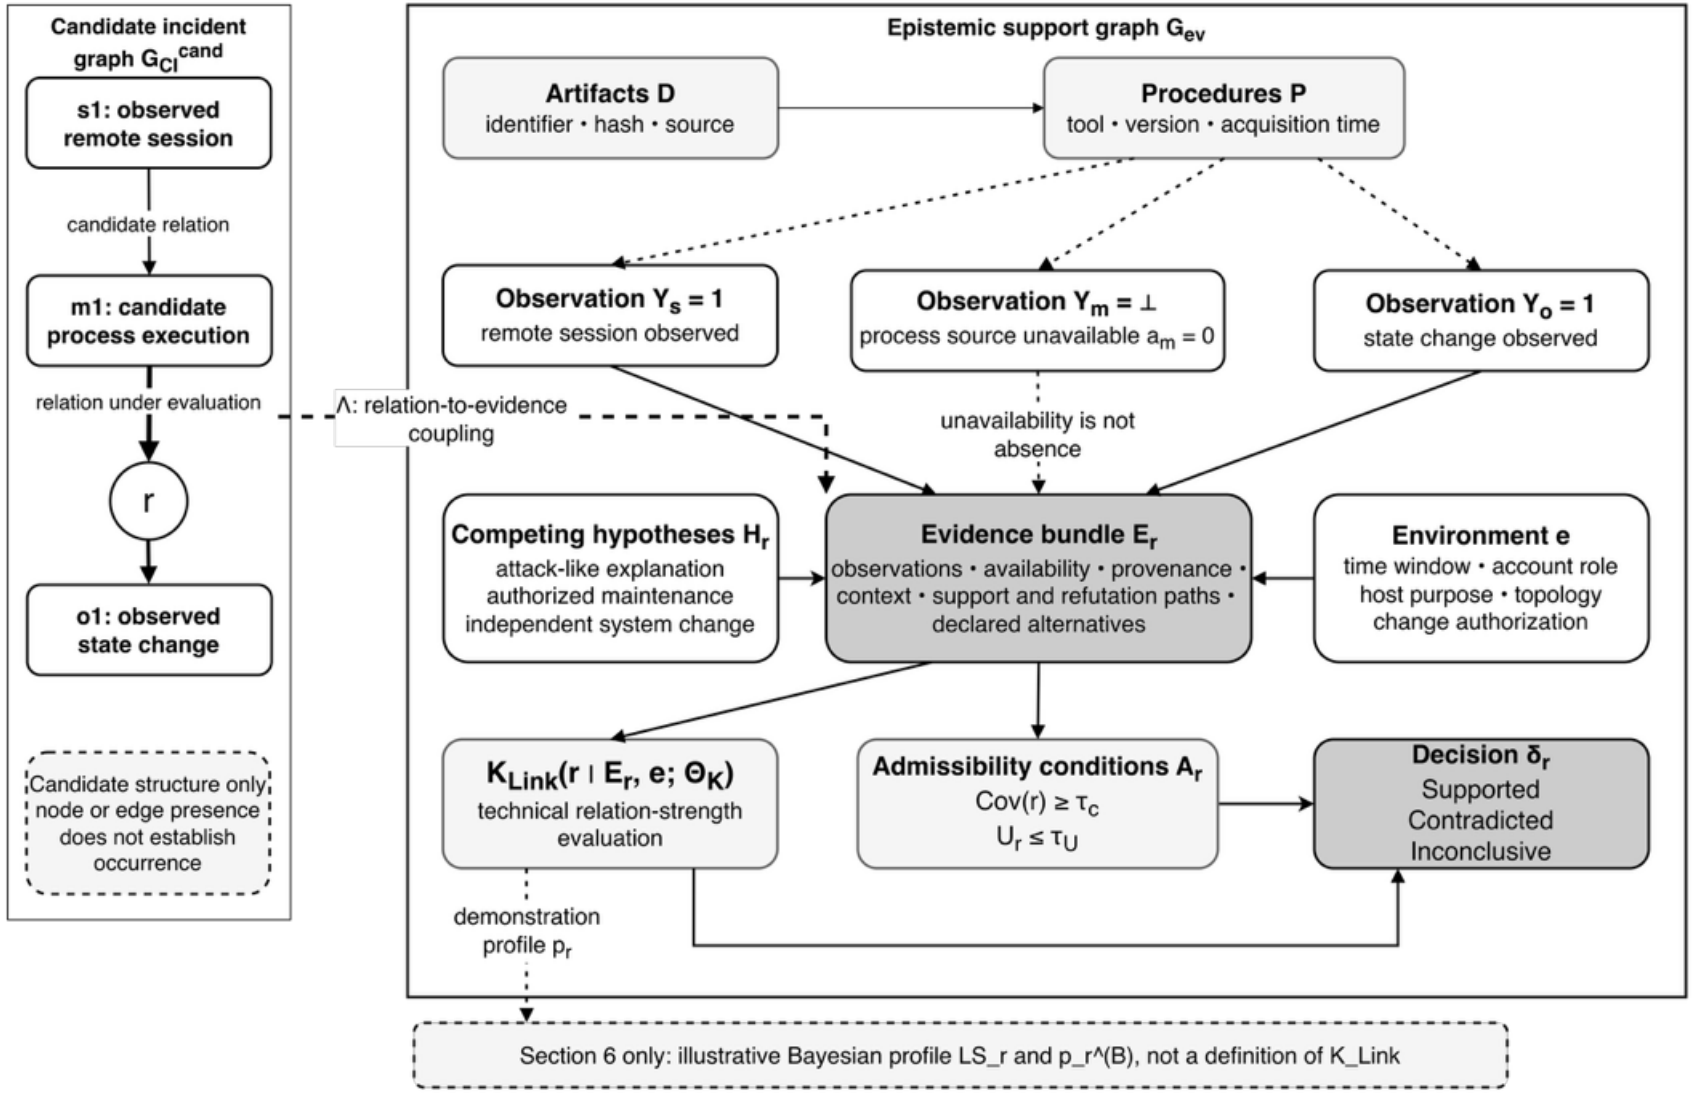
Source availability

Alternatives

Abstention

The incident tuple is a previously developed input schema

Two graphs and relation scoring



Graphs

Traceable evidence bundle E_r

K_Link

Separate relation strength evaluation

$$\kappa_r = K_{Link}(r \mid E_r, e; \Theta_K)$$

Three observation states

$$Y_q = \begin{cases} 1, & \text{if } q \text{ is present;} \\ 0, & \text{if } a_q = 1 \wedge q \text{ is absent;} \\ \perp, & \text{if } a_q = 0. \end{cases}$$

1 Present

0 Examined and absent

⊥ Unavailable or unexamined

$a_q = 1$ A capable source is available and examined

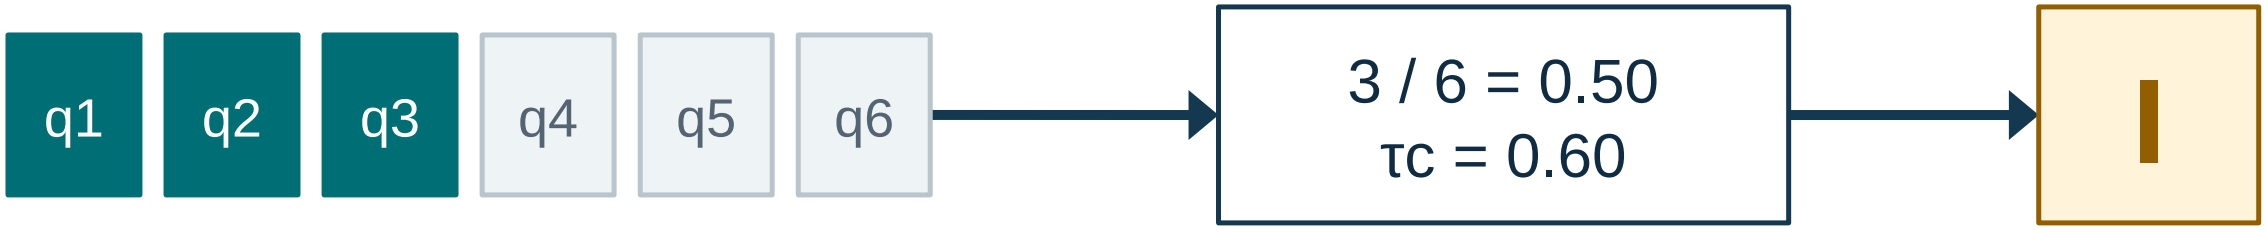
Unknown supplies neither positive nor negative evidence

Source coverage

$$Cov(r) = \frac{\sum_{q \in Q_r} w_{rq} a_q}{\sum_{q \in Q_r} w_{rq}}, \quad 0 \le Cov(r) \le 1$$

Expected observations and weights are predeclared

Equal-weight example

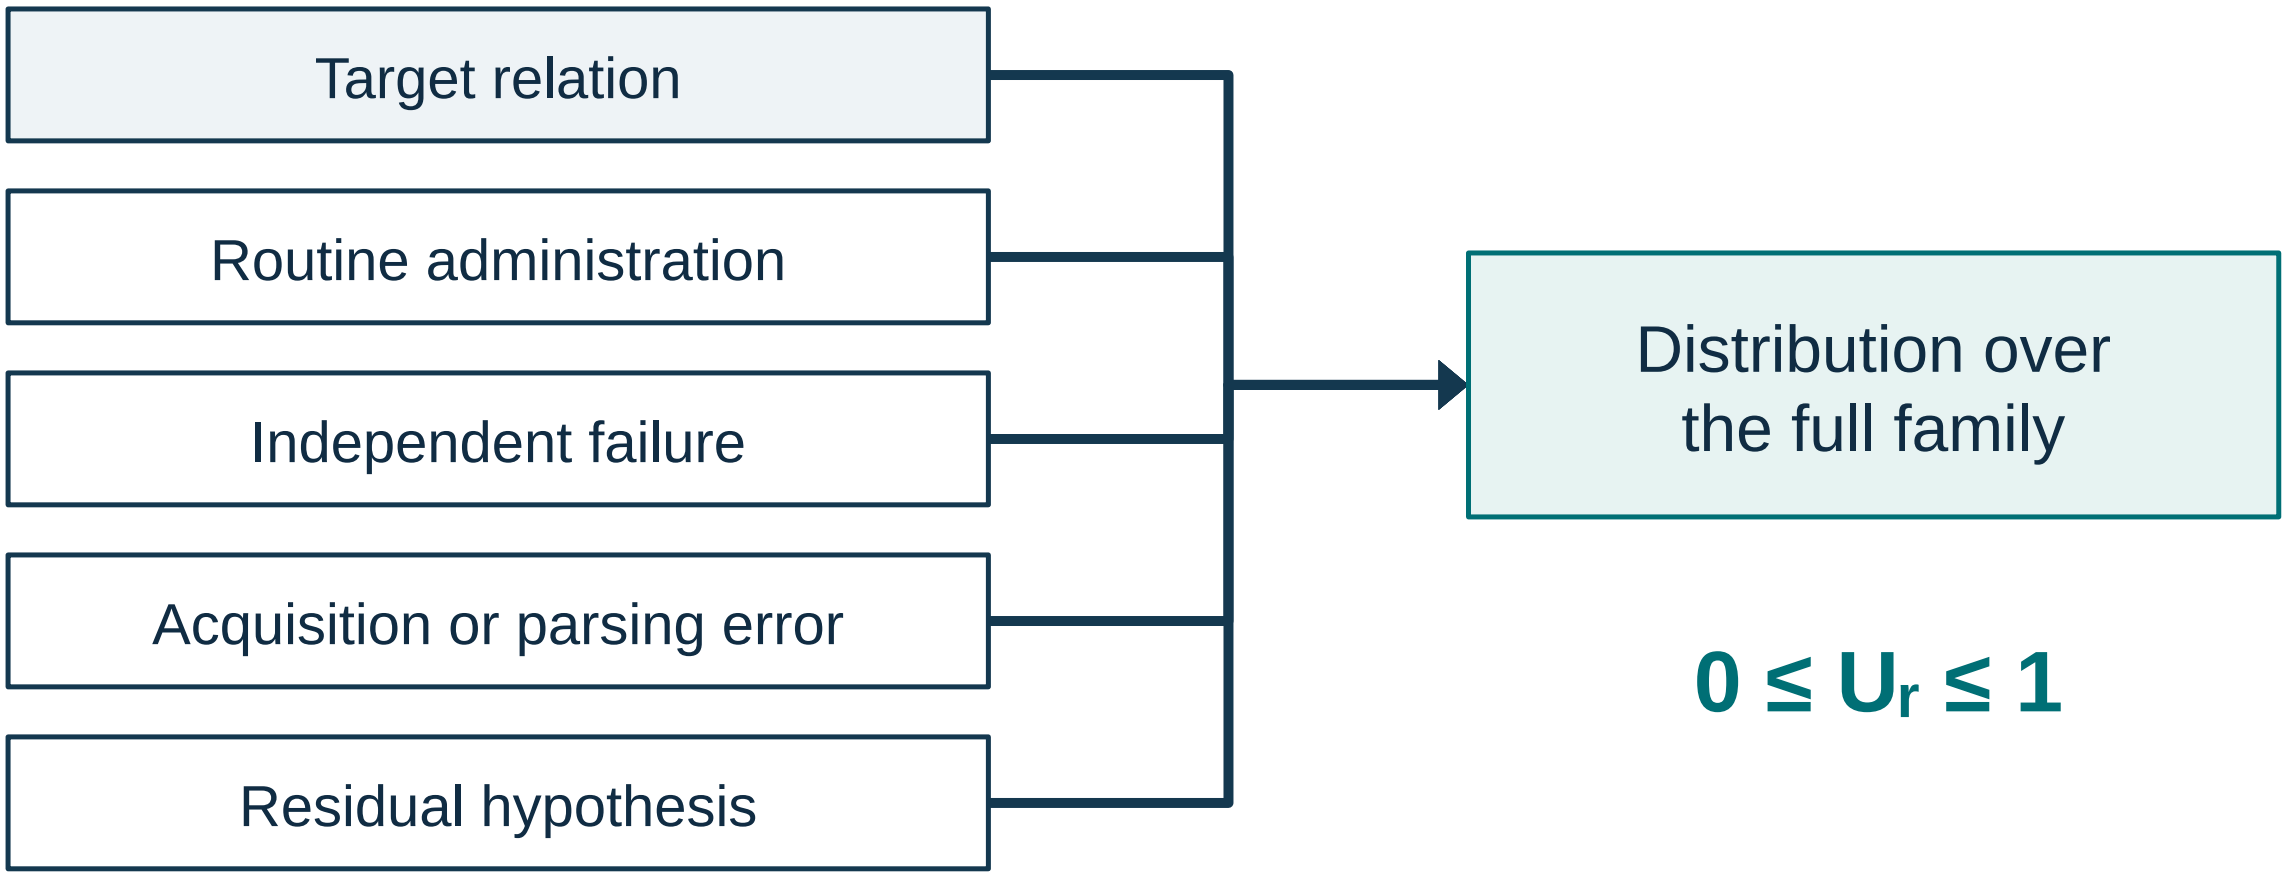


Available and examined

Unavailable

A negative observation still counts toward coverage

Competing hypotheses



Mutually exclusive hypotheses are declared before deciding

Conditions for a categorical decision

$$A_r = [Cov(r) \geq \tau_c] \wedge [U_r \leq \tau_U]$$

$$\delta_r = \begin{cases} S, & if \quad p_r \geq \tau_+ \wedge A_r; \\ C, & if \quad p_r \leq \tau_- \wedge A_r; \\ I, & otherwise. \end{cases}$$

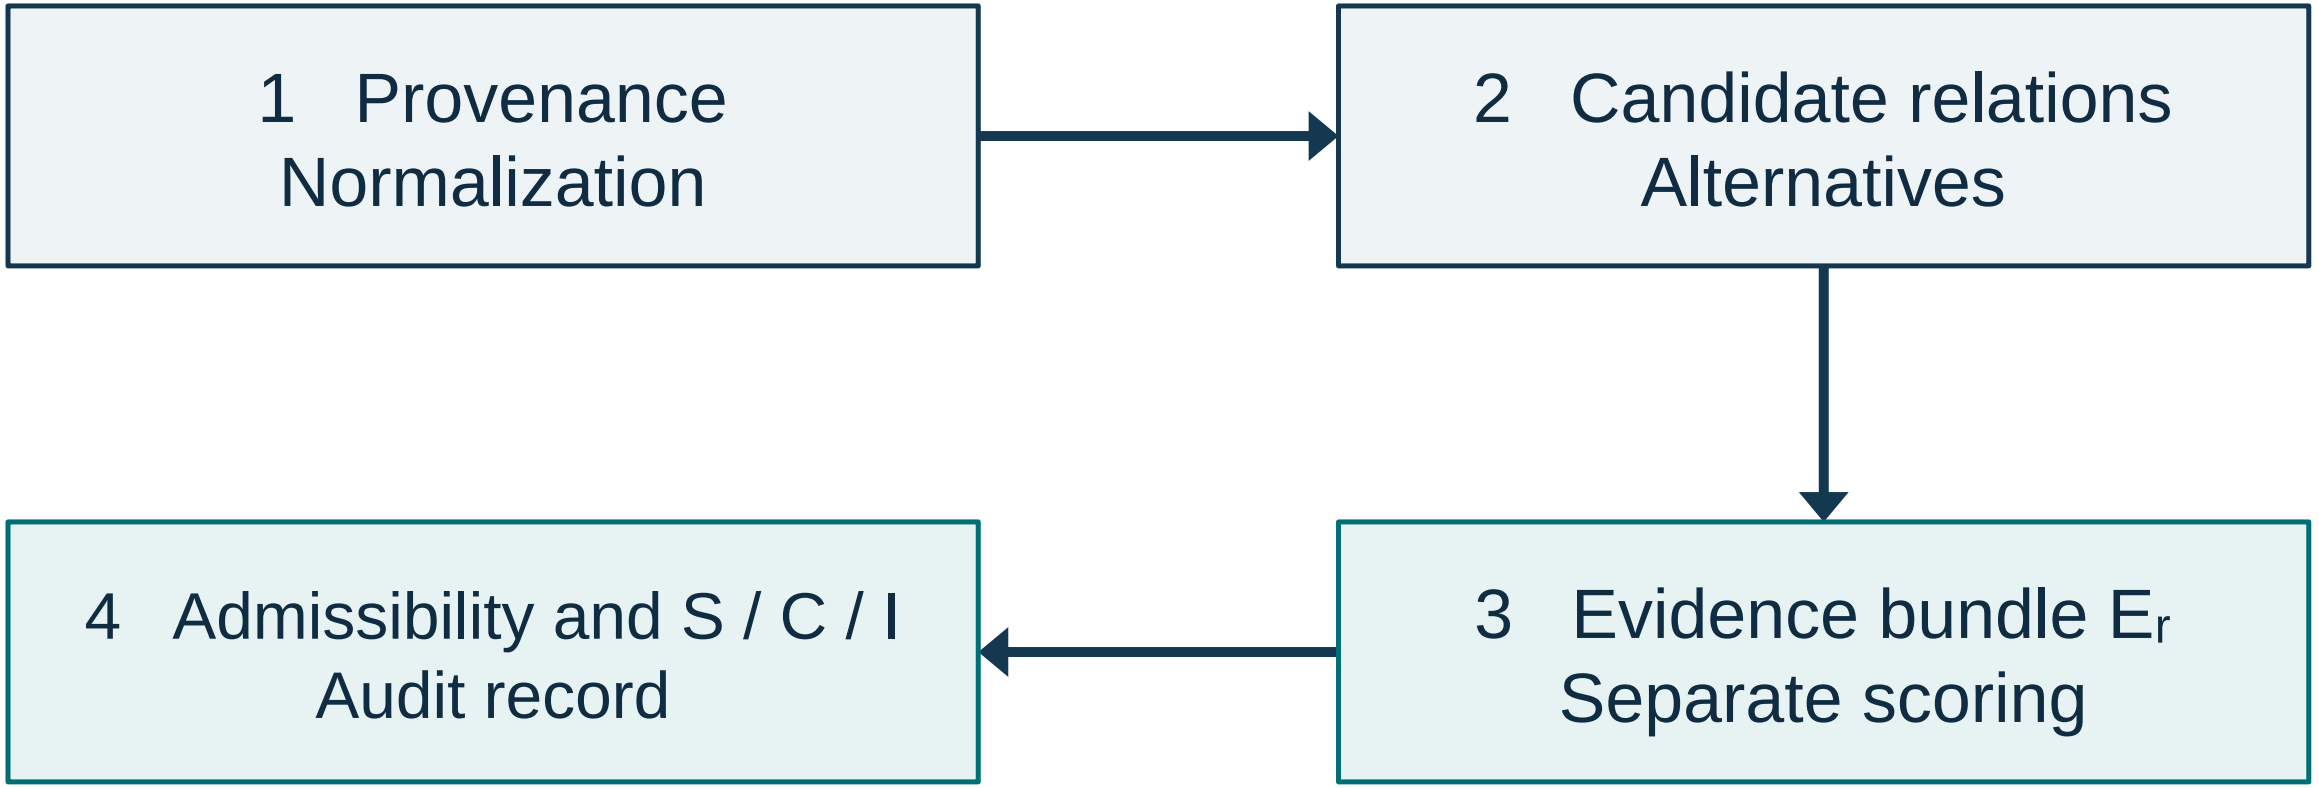
S Relation supported

C Composite alternative supported

I Insufficient grounds

The demonstration uses Bayesian p. A probability interpretation of K_Link requires independent calibration.

The projection workflow



The audit record retains rule versions, parameters and support paths

Numerical demonstration

H_0: routine administration H_1: attack-like scenario

6 conditionally independent observations, 64 vectors, $P(H_1) = 0.10$

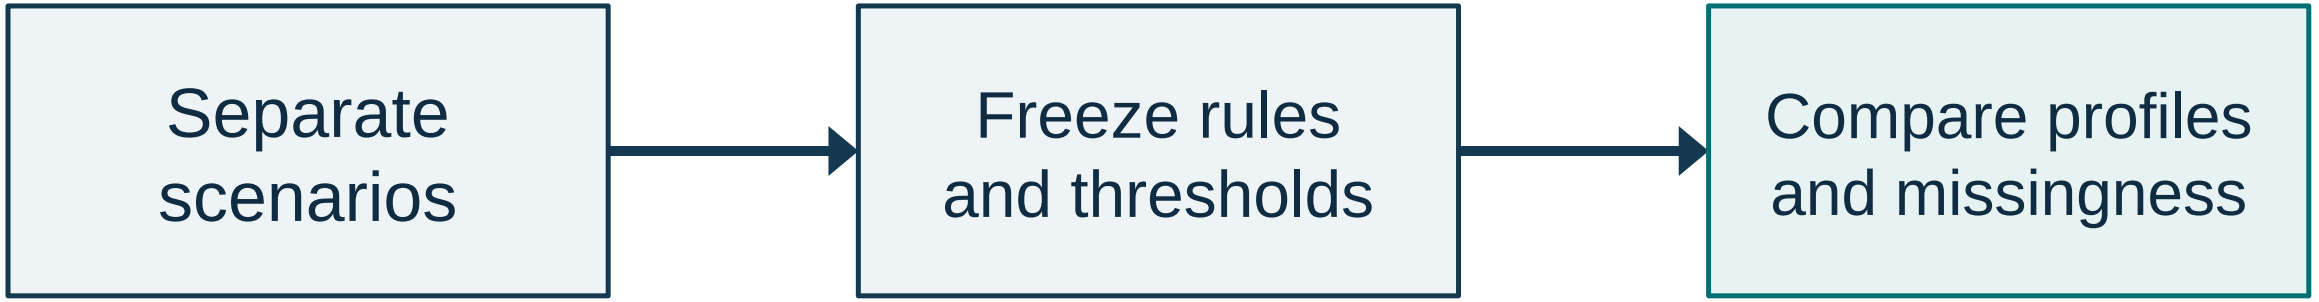
Profile	S under H_0	S under H_1	I under H_1
PowerShell rule	**25%**	**70%**	**No abstention**
Complete profile	**0.284%**	**76.716%**	**15.657%**
3 sources unavailable	**0%**	**0%**	**100%**

Thresholds $\tau_+ = 0.80$ $\tau_- = 0.20$ $\tau_c = 0.60$ $\tau_U = 0.75$

All values are model-implied. Zero FPR with 100% abstention does not mean better accuracy.

External validation protocol

Independent scenario ground truth

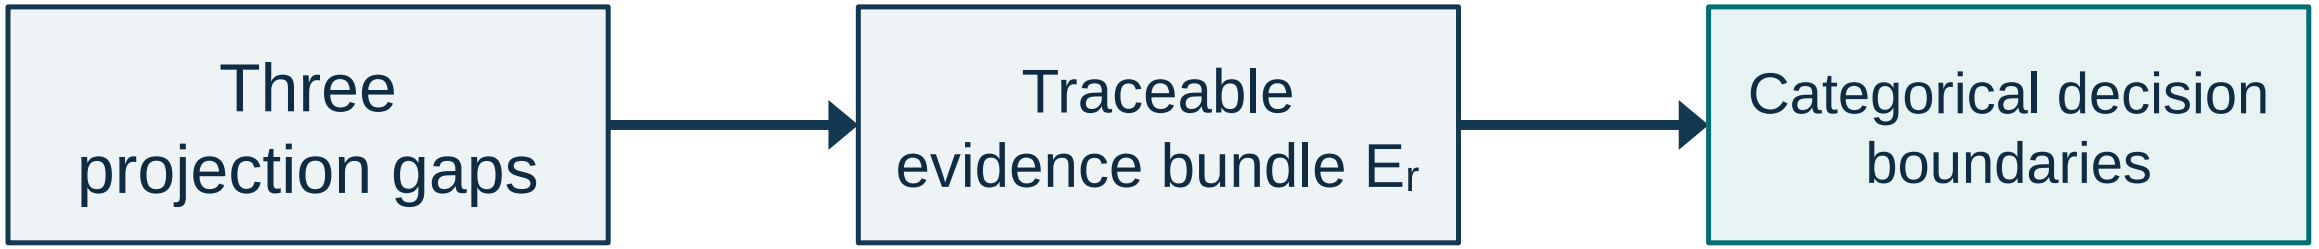


Errors, support, abstention, runtime and memory

The paper reports no real-data benchmark or production prototype

Main result

Auditable conditions for a technical claim



Next stage: independent external validation