Principled Data-Driven Decision Support for Cyber-Forensic Investigations

Soodeh Atefi, Sakshyam Panda, Manos Panaousis, Aron Laszka

UNIVERSITY OF HOUSTON

UNIVERSITY of GREENWICH

PennState

AAAI 2023

This material is based upon work sponsored by the National Science Foundation under Grant No. CNS-1850510.

Cyber Forensics Investigation

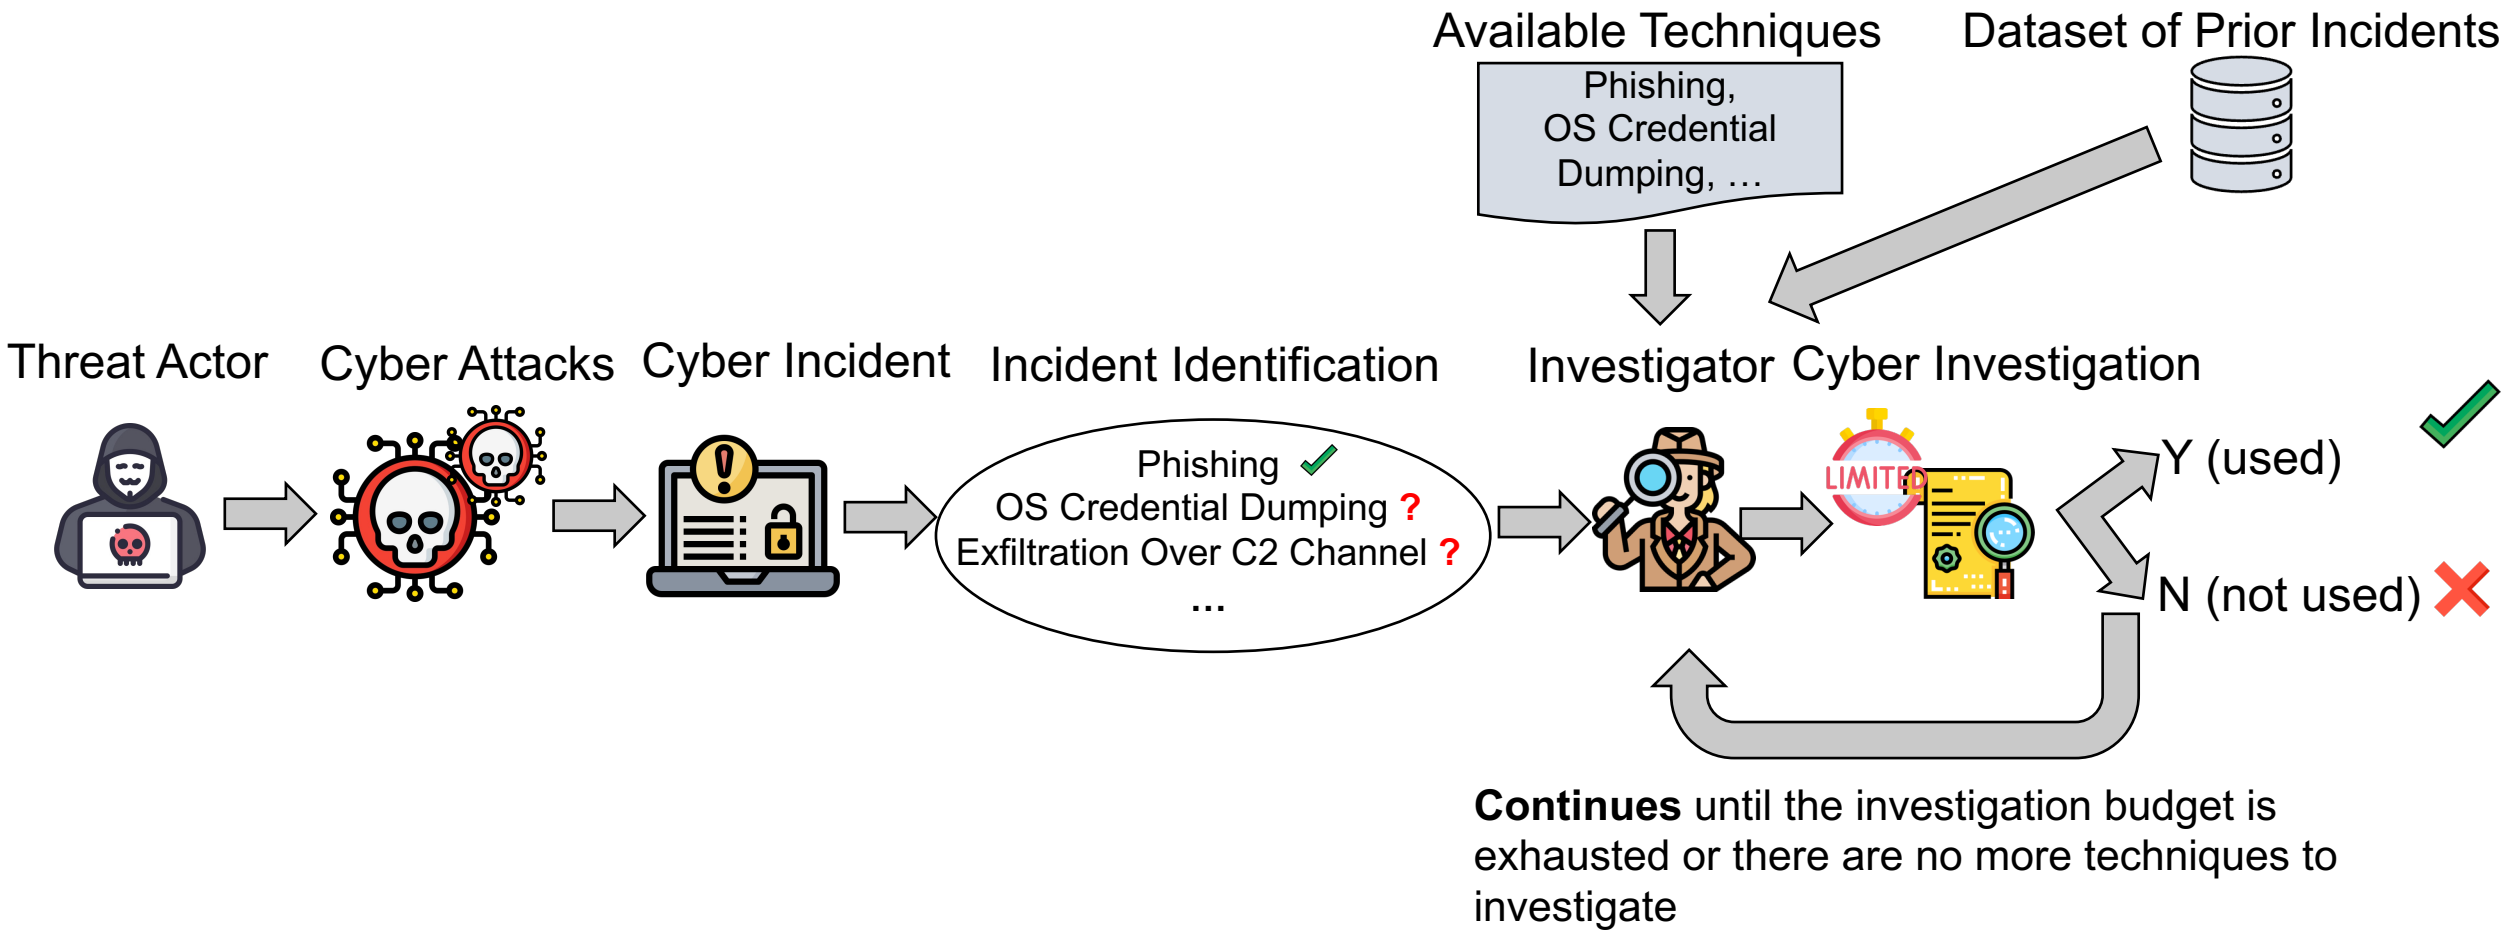

Icons from https://www.flaticon.com/

Goal of Decision Support

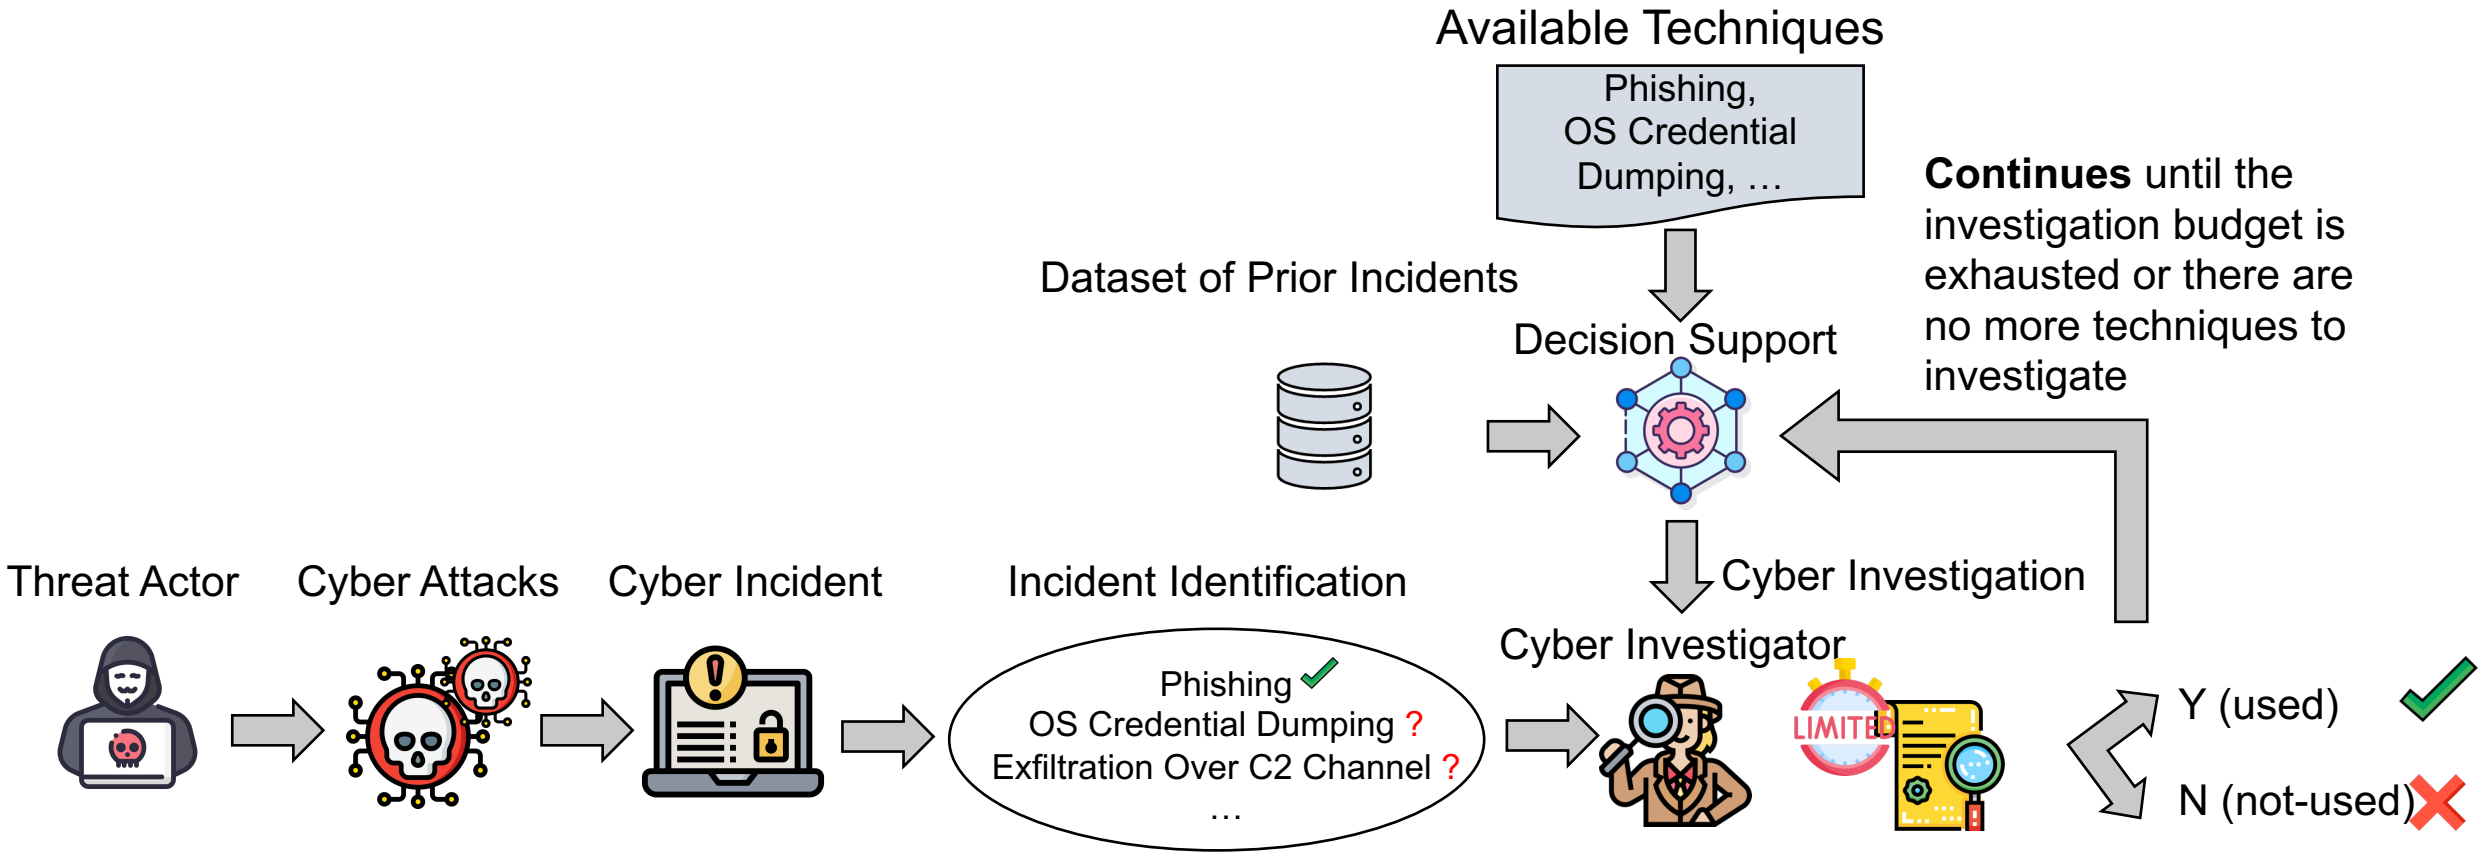

Icons from https://www.flaticon.com/

Limitations of Prior Studies

- Horsman et al. (CBR-FT) provide a probabilistic ranking of file system paths that contain suspicious evidence
  - It is only applicable to file system type of evidence.
  - It only helps the investigator **at the beginning of the investigation** (it is not a **stepwise** type of investigation).
- De Braekt et al. introduce a workflow management framework that guides investigator to increase the efficiency of the investigation (in terms of time and resources)**.**
  - It does not leverage the assessment of the investigator **at each step** of investigation.

Horsman G, Laing C, Vickers P. A case-based reasoning method for locating evidence during digital forensic device triage. Decision Support Systems. 2014 May 1;61:69-78.
de Braekt RI, Le-Khac NA, Farina J, Scanlon M, Kechadi T. Increasing digital investigator availability through efficient workflow management and automation. In 2016 4th International Symposium on Digital Forensic and Security (ISDFS) 2016 Apr 25 (pp. 68-73). IEEE

Limitations of Prior Studies (DISCLOSE)

- DISCLOSE (Nisioti et al. 2021) is a data-driven decision-support framework
- The goal of an investigator is to **maximize** the **benefit** obtained **during the investigation** without going over a given **budget**.
  - Investigation of each technique has a benefit and a cost (denoted by **B** and **C**).
  - **Budget** is the total cost that the investigator can spend during the investigation.
- Demonstrates that it **outperforms** the prior studies
- **Approach**:
  - Computes conditional probabilistic relations between techniques
  - Computes proximity values (based on attacks' life cycles) between techniques
  - Chooses techniques based on these relations

Nisioti A, Loukas G, Laszka A, Panaousis E. Data-driven decision support for optimizing cyber forensic investigations. IEEE Transactions on Information Forensics and Security. 2021 Jan 28;16:2397-412.

Limitations of Prior Studies (DISCLOSE)

- Computational approach is based on some heuristic likelihood values.
- It only considers the immediate benefit but not subsequent investigated steps (which is a myopic approach).
- It is a heuristic approach that does not approximate optimal decisions.

High-Level Explanation

Modeling cyber-forensic investigation as a Markov decision process (MDP)

Solve MDP by formulating decision support as a **policy** that **maximizes** the expected rewards obtained during the forensic investigation

Solve decision support using **k-NN-based MCTS**

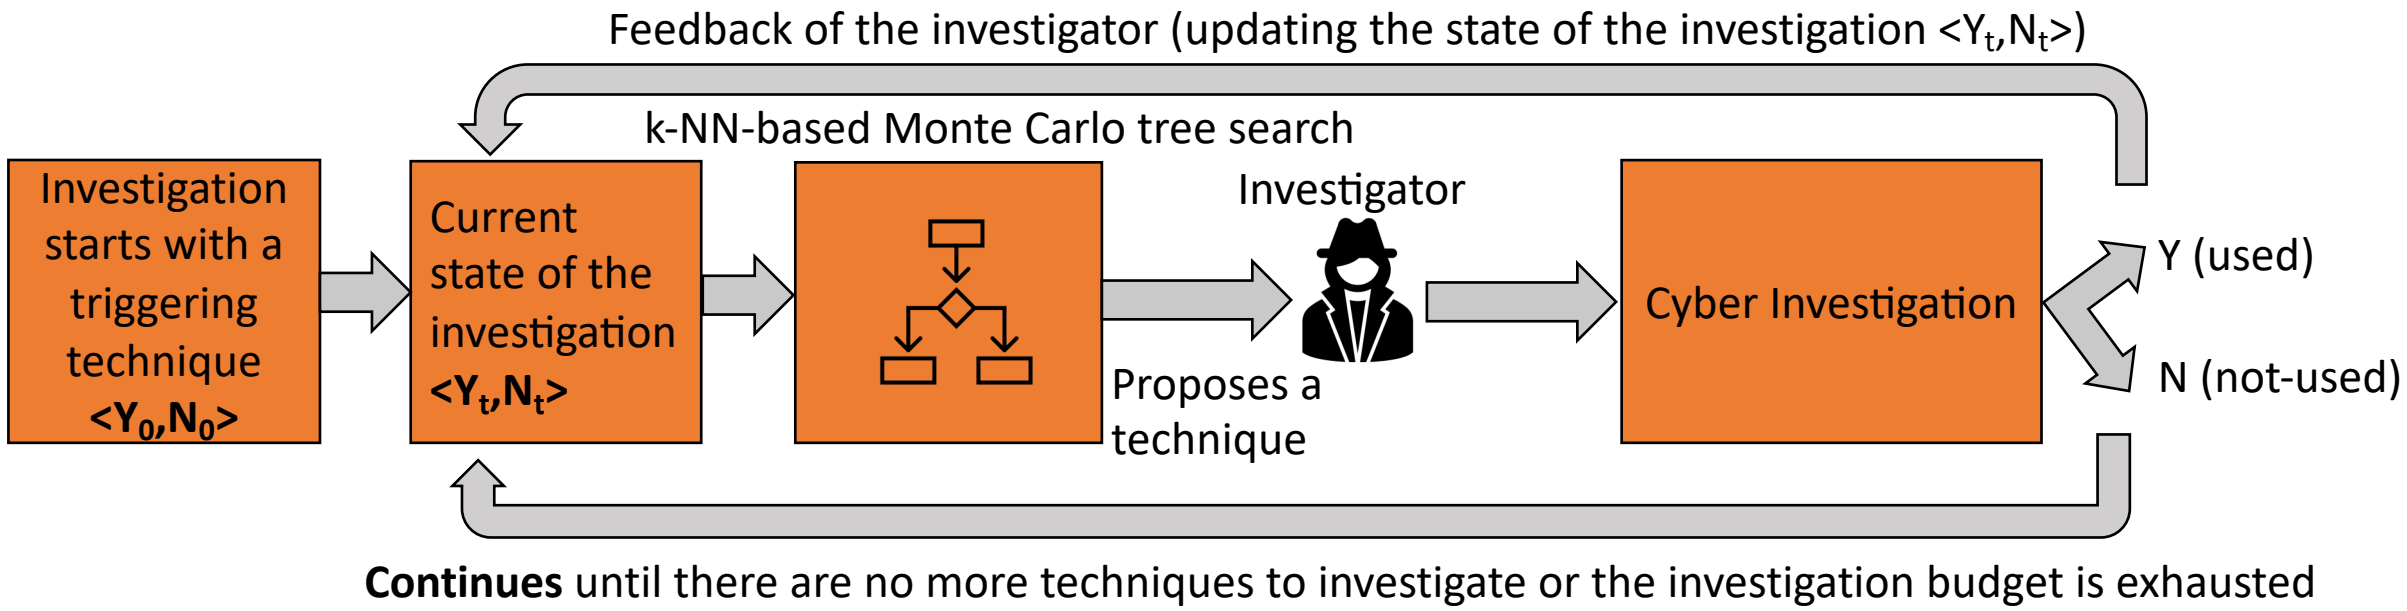

Output =discovered techniques (some are used Y, and some are not used N)

Our Approach: Markov Decision Process

- **State Space**: state corresponds to the set of **used** techniques discovered by step t ($\mathbf{Y_t}$) and the set of **not-used** techniques discovered by step t ($\mathbf{N_t}$)
- **Action Space:** the set of actions is the set of techniques $A \setminus (Y_t \cup N_t)$ at step t
  - A is a set of all adversarial techniques (actions in MDP).
- **Transition Probability:**
  - The probability that the investigated technique was actually used by a threat actor
  - Estimated based on k-NN regression (discussed later)
- **Rewards:**
  - $\mathbf{B_a}$ if technique a was **used** (state $\langle Y_t, N_t \rangle$ to state $\langle Y_{t+1}, N_{t+1} \rangle = \langle Y_t \cup \{a\}, N_t \rangle$)
  - **0** if technique a was **not-used** (state $\langle Y_t, N_t \rangle$ to state $\langle Y_{t+1}, N_{t+1} \rangle = \langle Y_t, N_t \cup \{a\} \rangle$)

Cyber Forensic Decision Support Problem

- A policy π, which maps a state $\langle Y_t, N_t \rangle$ to a recommended action $a \in A \setminus (Y_t \cup N_t)$.
- **Objective** is to maximize the expected rewards obtained during the forensic investigation.
- Formulation of the objective:

$$\max_{\pi} \mathbb{E}_{I_Y} \left[\sum_{t=0}^{T_{limit}} 1_{\{a_t \in I_Y\}} \cdot B_{a_t} \;\middle|\; a_t = \pi\left(Y_t, N_t\right) \right]$$

  - where T_{limit} is the last step before the investigation budget G is exhausted

$$T_{limit} = \max_T \sum_{t=0}^{T} C_{a_t} \leq G$$

Computational Approach

- We implement the policy π as an MCTS algorithm, relying on k-NN for estimating transition probabilities.
- **MCTS**
  - In each step of the investigation, we run a Monte Carlo tree search
  - Starting from the current state $<Y_t, N_t>$
  - **Outputs** an action a_t that is **estimated** to result in the **maximum expected discounted sum of rewards**

Updating Objective

- Reformulating the objective as maximizing the expected discounted sum of rewards:

$$\max_{\pi} \mathbb{E}\left[\sum_{t=0}^{|\mathcal{A}|-1} \gamma^t \cdot 1_{\{a_t \in I_Y\}} \cdot B_{a_t}/C_{a_t} \,\middle|\, a_t = \pi\left(Y_t, N_t\right)\right]$$

- Adding a **temporal discount factor**
- Replacing $\mathbf{B_{at}}$ with $\mathbf{B_{at}}$/ $\mathbf{C_{at}}$ (avoid focusing on immediate benefit)

Selection and Expansion (MCTS)

- t is the current step of the investigation
- Exploration rule: UCT with k-NN-based Myopic Pruning formula
- Expected immediate reward: $Pr[a \mid Y, N] \cdot B_a/C_a$

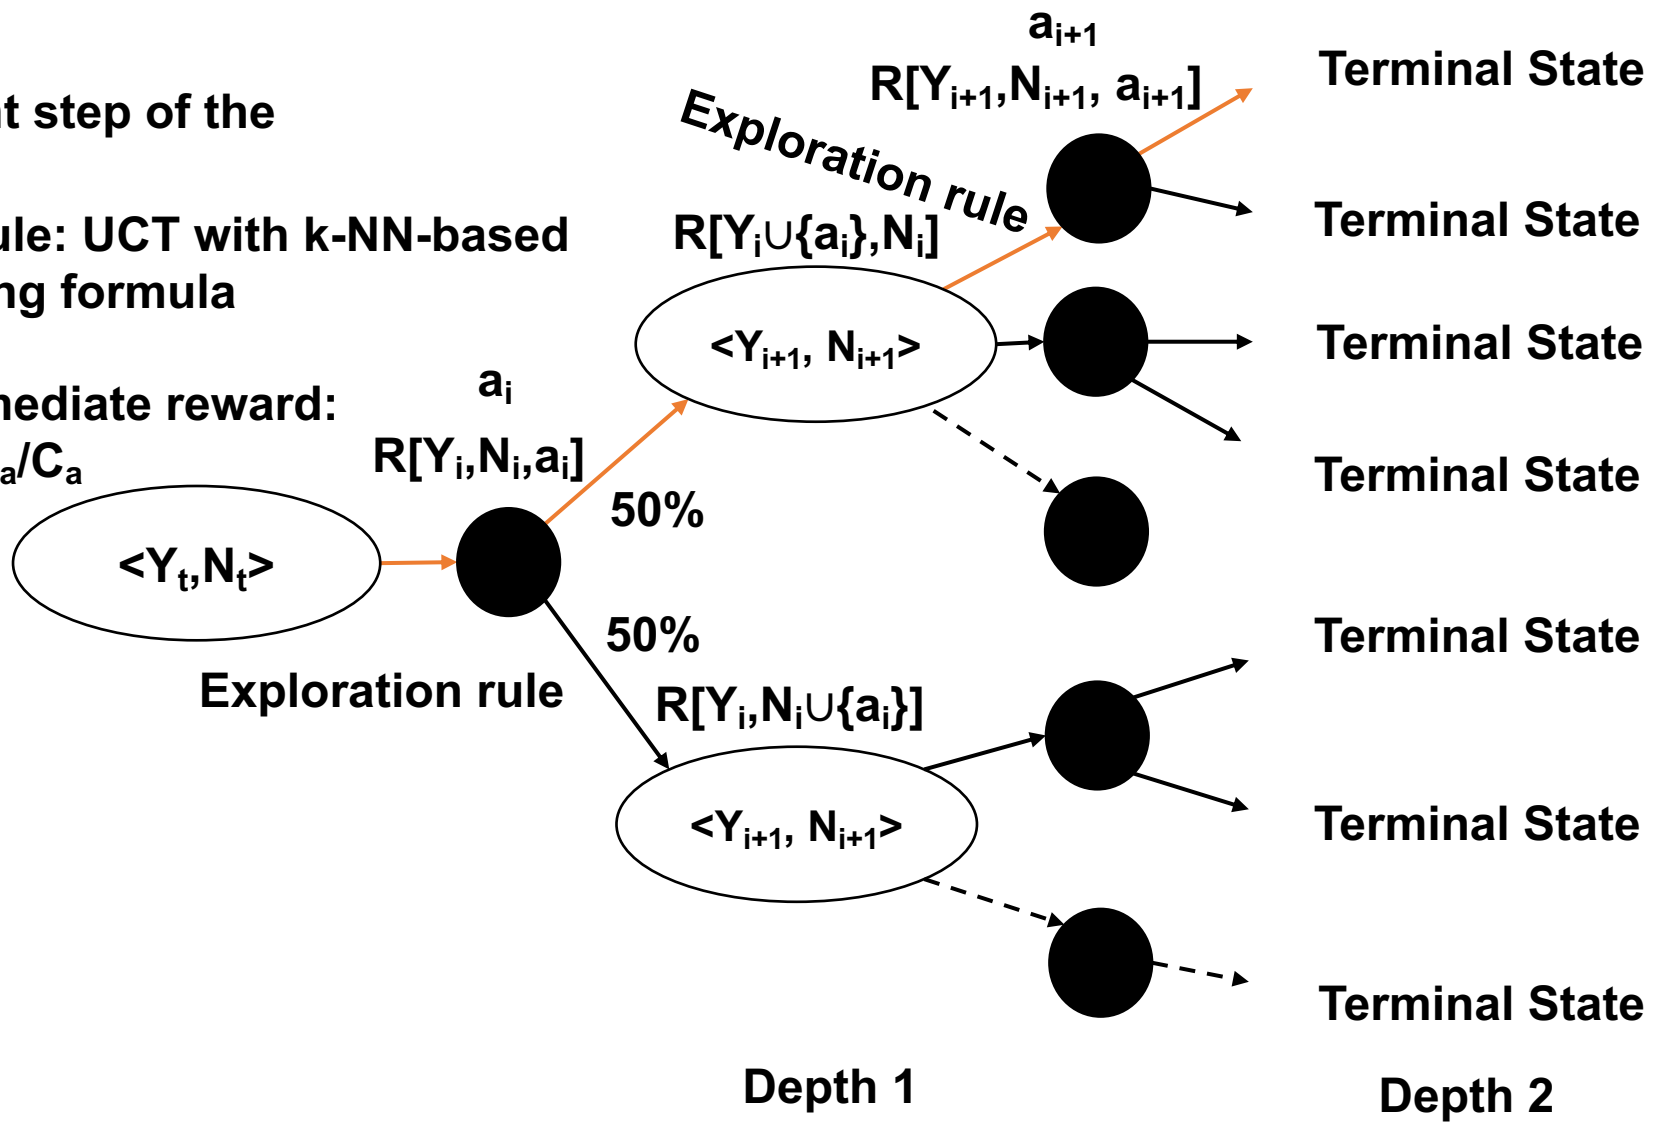

Backpropagation (MCTS)

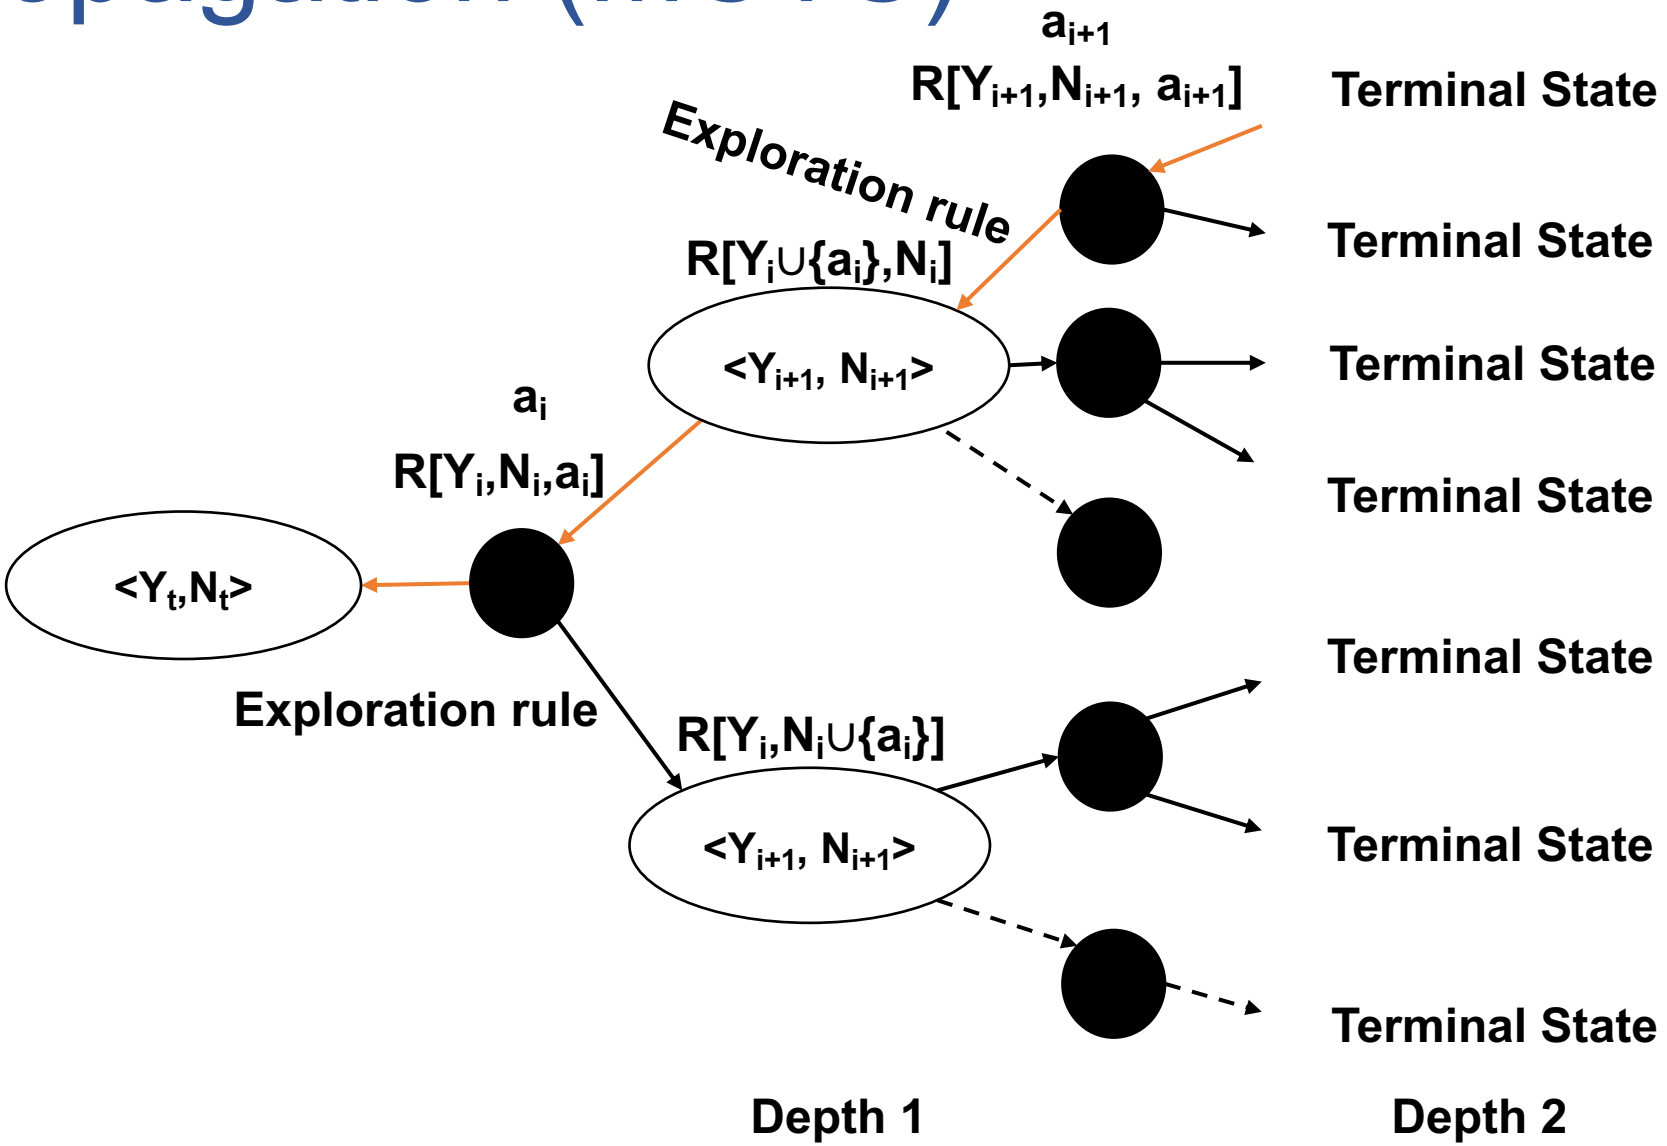

Backpropagation (MCTS)

$$R[Y_i, N_i, a_i] \leftarrow \Pr[a_i|Y_i, N_i] \cdot (B_{a_i}/C_{a_i} + \gamma \cdot R[Y_i \cup \{a_i\}, N_i]) + (1 - \Pr[a_i|Y_i, N_i]) \cdot (\gamma \cdot R[Y_i, N_i \cup \{a_i\}]).$$

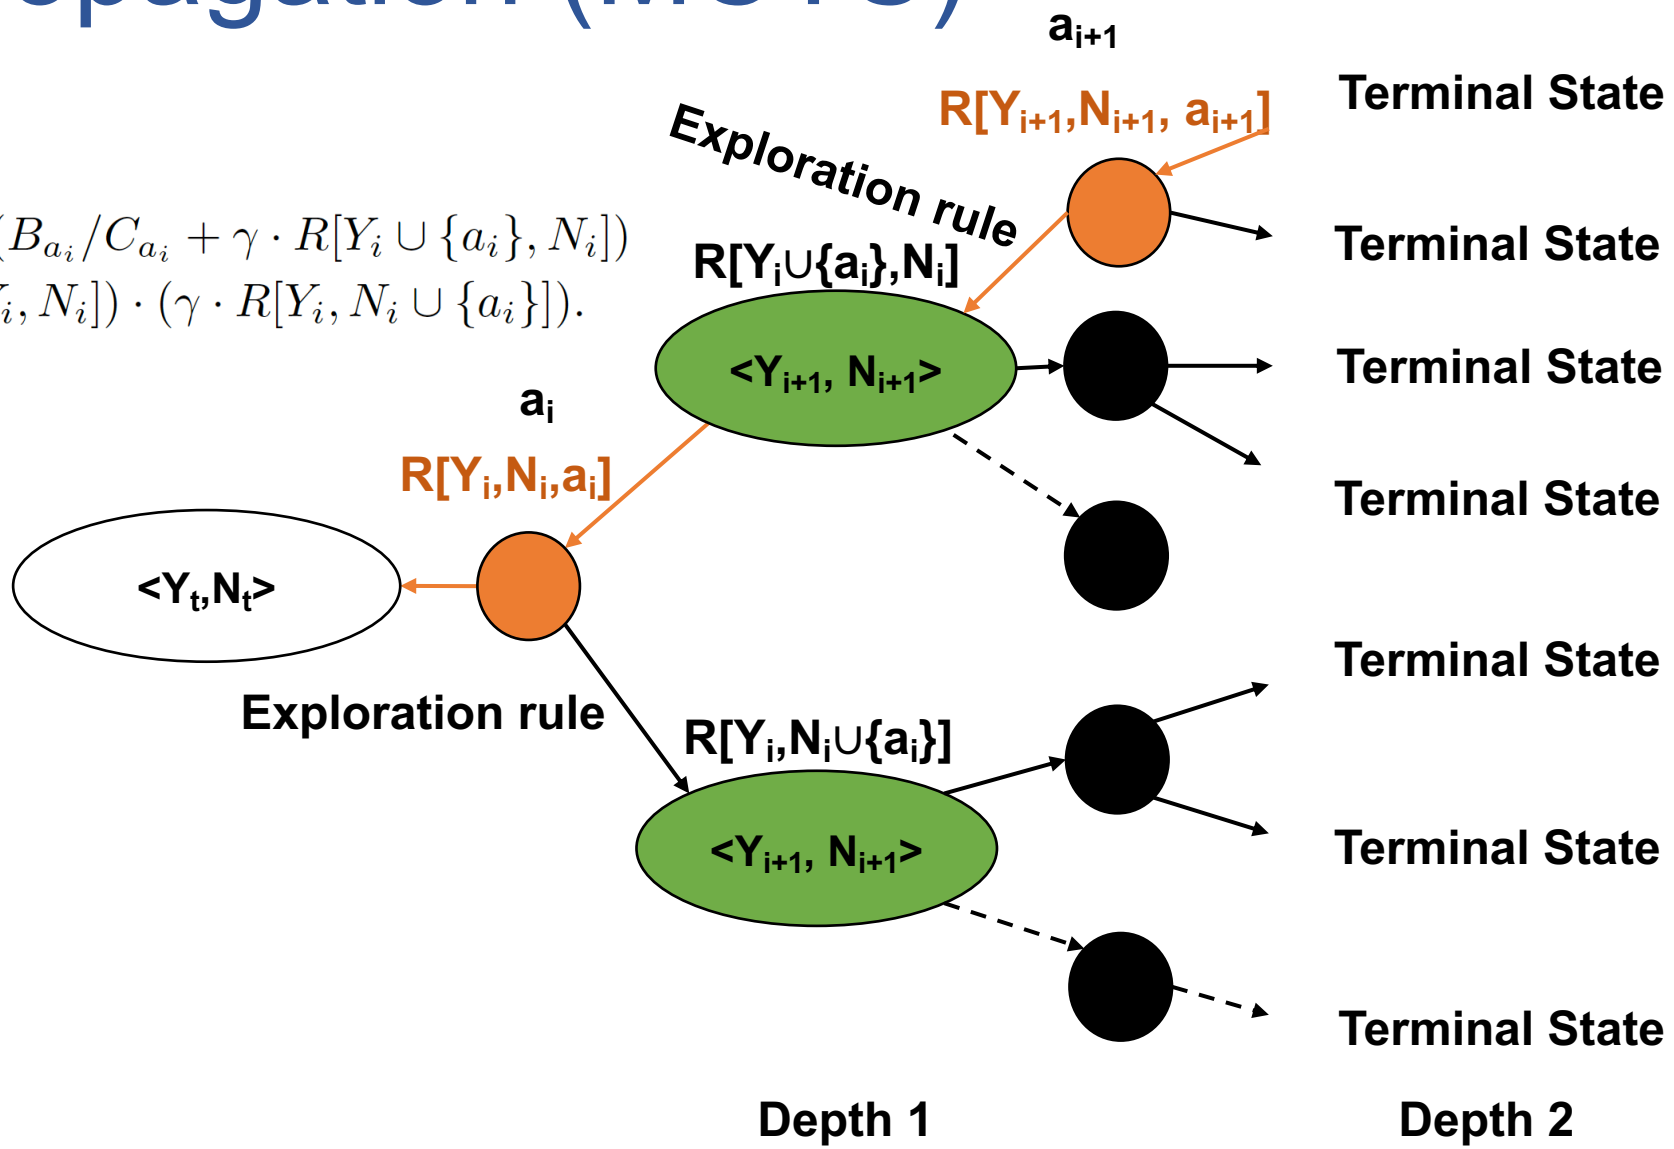

Backpropagation (MCTS)

Quick estimate when the depth limit is reached:

$$R[Y,N] \leftarrow \sum_{a \in \mathcal{A} \setminus (Y \cup N)} \sum_{j=0}^{|\mathcal{A} \setminus (Y \cup N)|-1} \gamma^j \frac{\Pr[a|Y_t, N_t] \cdot B_a / C_a}{|\mathcal{A} \setminus (Y \cup N)|}$$

$$\max_{a \in \mathcal{A} \setminus (Y_i \cup N_i)} R[Y_i, N_i, a]$$

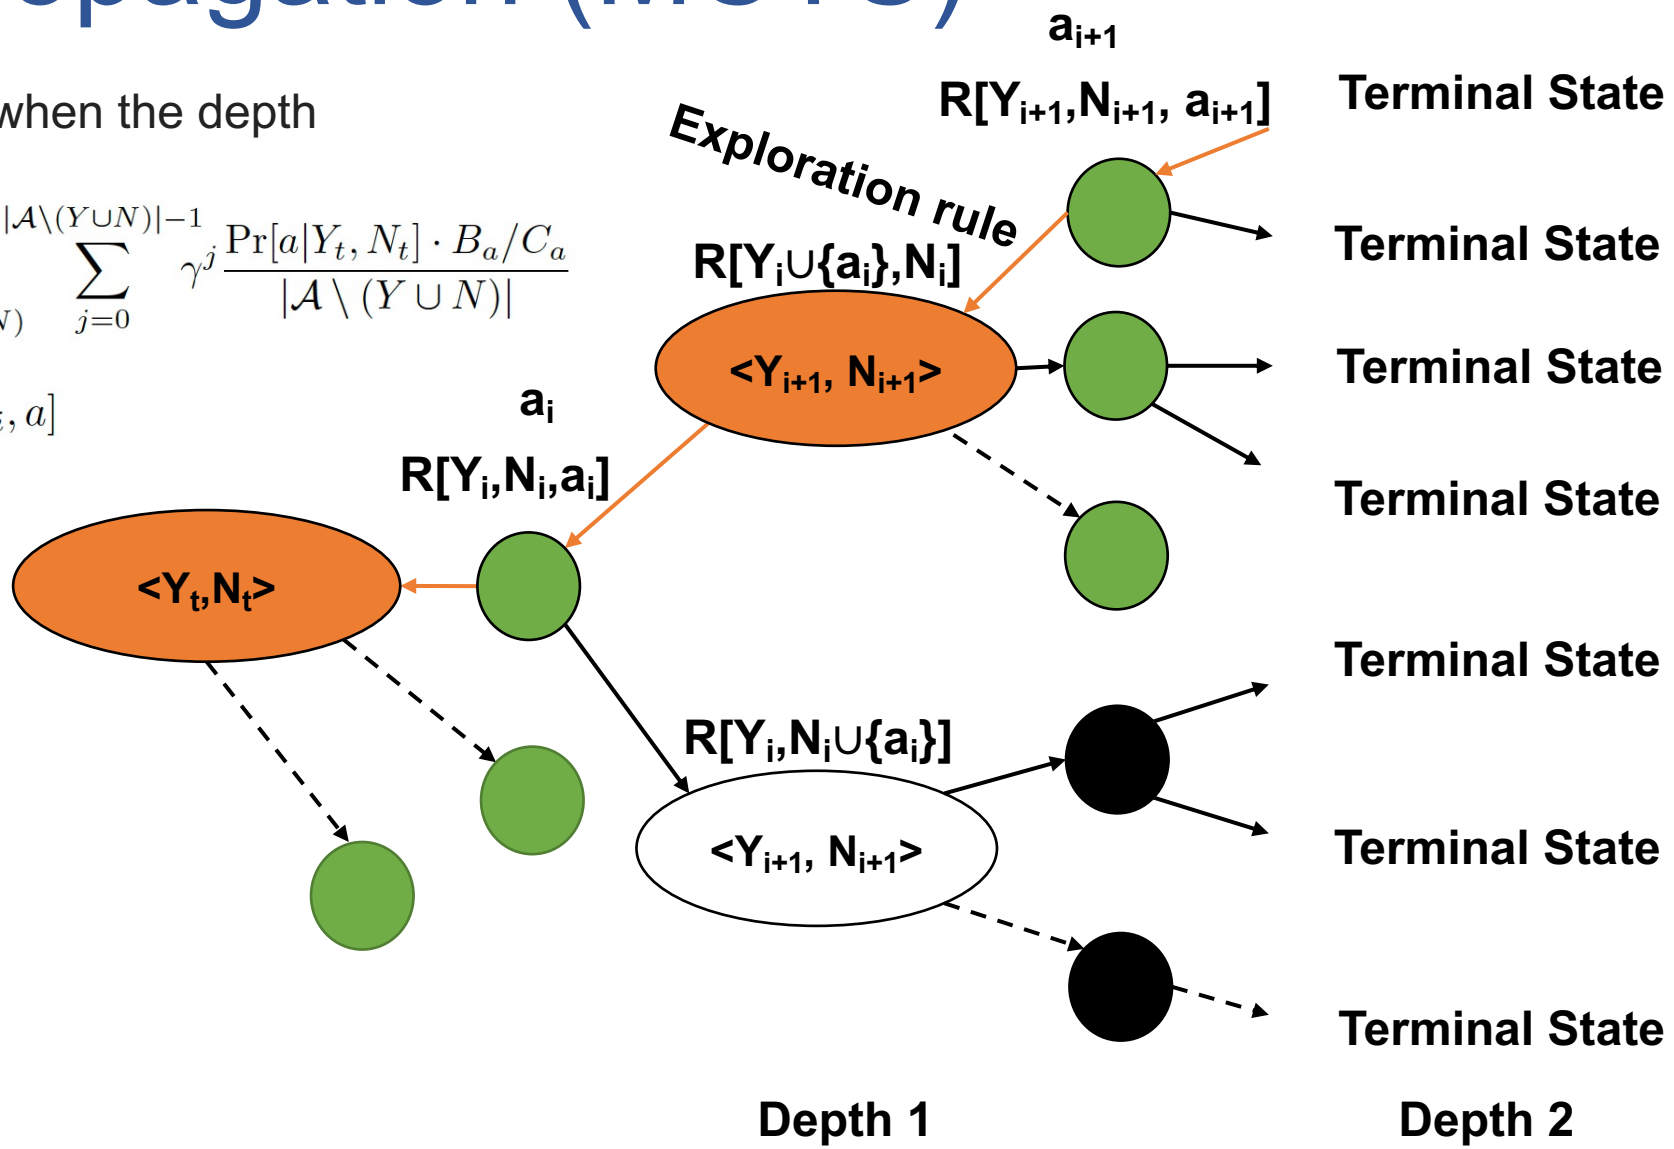

Probability Estimation

- Estimating the state-transition probability Pr[a | Y_t,N_t] as the **exact conditional empirical probability**

$$\Pr[a|Y_t, N_t] \equiv \Pr\left[a \in I_Y | Y_t \subseteq I_Y \wedge N_t \cap I_Y = \emptyset\right] \approx \frac{\left|\left\{\hat{I} \in \mathcal{I}_{\langle Y_t, N_t \rangle} \,\middle|\, a \in \hat{I}_Y\right\}\right|}{\left|\mathcal{I}_{\langle Y_t, N_t \rangle}\right|}$$

  - where $\mathcal{I}_{\langle Y_t, N_t \rangle} = \left\{\hat{I} \in \mathcal{I} \,\middle|\, Y_t \subseteq \hat{I}_Y \wedge N_t \cap \hat{I}_Y = \emptyset\right\}$
  - $\hat{I}$ denotes a prior incident ($\hat{I} \in \mathcal{I}$)
- **Weakness**:
  - The number of prior incidents that "match" the current state of the investigation may reach zero.
- **Solution**:
  - Considering the set of k prior incidents that are closest with respect to metric d

k-Nearest Neighbors (k-NN)

- **Distance Calculation:**
  - Breaking ties arbitrarily
  - Measuring distance by counting techniques that differ

$$d(\langle Y_t, N_t \rangle, \hat{I}) = \left| Y_t \cap \hat{I}_N \right| + \left| N_t \cap \hat{I}_Y \right|$$

- Selecting k incidents with lowest distance (i.e., k-NN)
- $k = \beta_1 + \beta_2 \cdot t$
  - β_1 and β_2 are hyper-parameters
  - t is step of the investigation

Dataset

- Three versions of MITRE ATT&CK Enterprise dataset (v6.3, v10.1, and v11.3 latest).
- Our approach can be applied to newer versions without any changes.
- For fair comparison, we use 31 techniques the same as DISCLOSE.
- Mapping some techniques to equivalent ones for later versions
- Benefit and Cost of each technique (same as DISCLOSE):
  - **Benefit**: Common Vulnerability Scoring System (CVSS)
  - **Cost:** interviews with cyber forensic experts

MITRE Corporation. 2022. MITRE Cyber Threat Intelligence Repository. https://github.com/mitre/cti, accessed on August 25th, 2022.

Experimental Setup

- **Baselines:**
  - DISCLOSE
  - Static policy (most frequent techniques across all prior incidents)
- **Simulation Setup and Metrics**
  - Leave-one-out cross-validation
  - We treat all other incidents in our dataset as prior incidents.
  - We terminate an investigation when the cumulative effort cost reaches **45, and 65.**
  - We measure the area under the benefit-effort curve (AUCBE).

Numerical Results

- Our approach outperforms the baselines in both scenarios (45, and 65) for three versions of the dataset.
- We optimized the hyper-parameters for datasets and different investigation budgets separately.

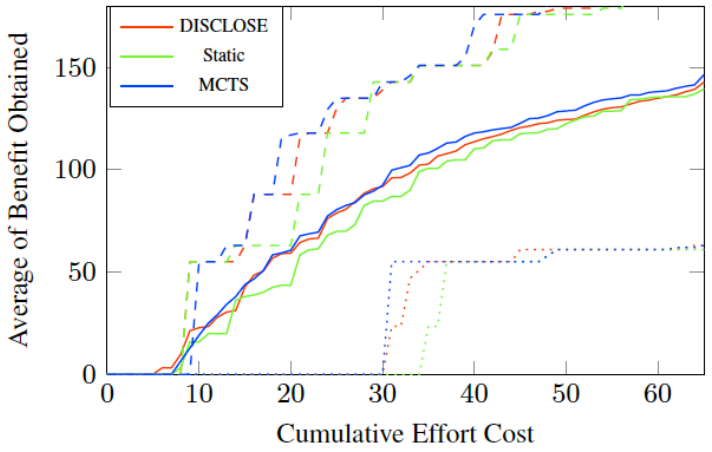

Figure 2: Average benefit obtained as a function of cumulative effort cost (**up to budget 65**) on **v6.3**.

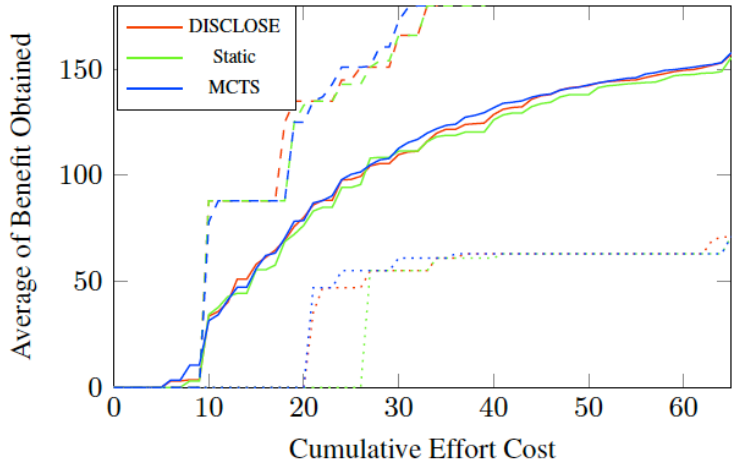

Figure 4: Average benefit obtained as a function of cumulative effort cost (**up to budget 65**) on **v11.3**.

Conclusion

- To address the limitations of DISCLOSE, we propose an MCTS and k-NN-based computational approach.
- Our approach outperforms baselines.
- **Advantages**:
  1. It works directly with the data.
  2. It approximates optimal decisions based on the dataset.

Thank you for your attention

Other References

- Puterman ML. Markov decision processes: discrete stochastic dynamic programming. John Wiley & Sons; 2014 Aug 28.
- Kocsis L, Szepesvári C. Bandit based monte-carlo planning. In European conference on machine learning 2006 Sep 18 (pp. 282-293). Springer, Berlin, Heidelberg.
- Cover T, Hart P. Nearest neighbor pattern classification. IEEE transactions on information theory. 1967 Jan;13(1):21-7.