

Beyond Binary Opinions: A Deep Reinforcement Learning-Based Approach to Uncertainty-Aware Competitive Influence Maximization

Qi Zhang, Dian Chen, Lance M. Kaplan, *Fellow, IEEE*, Audun Jøsang, Dong Hyun Jeong, *Member, IEEE*, Feng Chen, *Member, IEEE*, Jin-Hee Cho, *Senior Member, IEEE*

Abstract—The *Competitive Influence Maximization (CIM)* problem involves multiple entities competing for influence in online social networks (OSNs). While *Deep Reinforcement Learning (DRL)* has shown promise, existing methods often assume users' opinions are binary and ignore their behavior and prior knowledge. We propose DRIM, a multi-dimensional uncertainty-aware DRL-based CIM framework that leverages *Subjective Logic (SL)* to model uncertainty in user opinions, preferences, and DRL decision-making. DRIM introduces an *Uncertainty-based Opinion Model (UOM)* for a more realistic representation of user uncertainty and optimizes seed selection for propagating true information while countering false information. In addition, it quantifies uncertainty in balancing exploration and exploitation. Results show that UOM significantly enhances true information spread and maintains influence against advanced false information strategies. DRIM-based CIM schemes outperform state-of-the-art methods by up to 57% and 88% in influence while being up to 48% and 77% faster. Sensitivity analysis indicates that higher network observability and greater information propagation boost performance, while high network activity mitigates the effect of users' initial biases.

Index Terms—Competitive influence maximization, deep reinforcement learning, uncertainty, opinion models, influence propagation

1 INTRODUCTION

Online Social Networks (OSNs) have become dominant platforms for information exchange and opinion formation. However, they also facilitate the rapid spread of false information, leading to reputational harm, financial loss, and manipulated public opinion [1]. The *Competitive Influence Maximization (CIM)* problem arises in environments where competing entities, such as political parties or corporations, strategically select seed nodes to maximize their influence in OSNs. This study focuses on two competing parties: The *true party* (TP), disseminating true information, and the *false party* (FP), spreading false information. To mitigate the impact of false information, it is crucial to develop strategies that empower the TP to counteract the FP effectively and enhance the reach of true information in OSNs.

Traditional CIM approaches typically model user opinions as binary (i.e., supporting either the TP or the FP), overlooking the complexity and uncertainty of real-world opinion dynamics. In reality, OSN users often encounter ambiguous, evolving, and multi-dimensional information that shapes their decisions. A binary assumption oversimplifies influence propagation, neglecting belief strength, uncertainty, and adaptability. This limitation highlights the need for an advanced CIM framework that incorporates opinion uncertainty and dynamic belief updates.

Qi Zhang, Dian Chen and Jin-Hee Cho are with Department of Computer Science, Virginia Tech, USA (emails: {qiz21, dianc, jicho}@vt.edu).

Lance M. Kaplan is with U.S. Army DEVCOM Army Research Laboratory, USA (email: lance.m.kaplan.civ@army.mil).

Audun Jøsang is with Department of Informatics, University of Oslo, Norway (email: josang@ifi.uio.no).

Dong Hyun Jeong is with Department of Computer Science, University of the District of Columbia, USA (email: djjeong@udc.edu).

Feng Chen is with Department of Computer Science, University of Texas at Dallas, USA (email: feng.chen@utdallas.edu)

Deep Reinforcement Learning (DRL) [2] has emerged as a powerful tool for real-time decision-making and strategy optimization in CIM tasks. By learning policies from interactions, DRL enables adaptive decision-making that evolves with the changing OSN landscape. A key challenge in DRL-based CIM is balancing exploration and exploitation. The agent must discover new influential nodes (exploration) while leveraging known high-impact nodes (exploitation) to maximize long-term influence. Conventional methods such as Epsilon-greedy [2] and Upper Confidence Bound (UCB) [3] either fail to adapt the exploration rate dynamically or rely solely on reward or state values, ignoring the DRL model's intrinsic uncertainty.

To address these challenges, we propose an uncertainty-aware DRL framework for CIM that integrates *Subjective Logic (SL)* [4], a belief model that explicitly quantifies multiple types of uncertainty, to represent opinion uncertainty. By incorporating uncertainty estimates into the exploration-exploitation trade-off, our approach enables adaptive decision-making that considers belief strength and confidence levels in user opinions. This allows the DRL agent to make more informed decisions under uncertainty, enhancing its ability to counteract false information while maximizing the spread of true information in OSNs.

Our work has made the following key contributions:

- 1) **Beyond Binary Opinion Models:** We advance CIM research by moving beyond traditional binary opinion models, which dominate existing studies [5–8]. Instead, we formulate opinion dynamics using *Subjective Logic (SL)* [4], a belief model that explicitly allows estimating uncertainty and dynamically updating user opinions based on real-world interactions. Unlike static binary models, our approach reflects the evolving nature of user

beliefs in OSNs. Furthermore, we introduce and compare three different opinion models to evaluate how user attitudes toward opinion updates influence the effectiveness of false information mitigation.

- 2) Dual-Agent DRL for CIM: We adopt a dual-agent DRL framework, where both the TP and the FP engage in a CIM process. While DRL has been previously applied to CIM [7–9], prior works have only considered DRL to model a single agent's decision-making. By explicitly modeling the interaction between two competing parties, our framework captures the adversarial nature of influence propagation in OSNs, enabling a more realistic and strategic seed selection process.
- 3) Uncertainty-Aware Decision-Making: We introduce uncertainty-aware decision-making at three critical levels: user opinion modeling, balance between exploration and exploitation in DRL, and network environment simulation. Our framework integrates SL-based opinion modeling with uncertainty estimation to better capture the uncertainties in user beliefs. Moreover, we develop multiple uncertainty-aware exploration-exploitation strategies using Evidential Neural Networks (ENNs) [10, 11], bridging SL with Deep Neural Networks (DNNs) for more informed decision-making. Furthermore, we consider partially observable networks, where network edges are only partially visible, reflecting real-world constraints. To the best of our knowledge, no prior CIM work has incorporated such a multi-faceted uncertainty-aware approach across opinions, decision-making, and environmental modeling.
- 4) Extensive Empirical Validation: We conduct extensive simulation experiments on three real-world social network datasets to rigorously evaluate our proposed uncertainty-aware DRL-based CIM framework. Our experiments assess the impact of different opinion models, uncertainty-aware DRL strategies, and sensitivity analyses across various settings, ensuring the robustness and applicability of our findings.

While preliminary results of this research were presented in [12], this work significantly extends [12] by (1) introducing uncertainty-aware DRL with novel exploration-exploitation strategies, (2) leveraging ENNs to estimate multiple uncertainty types for seed node selection, and (3) validating the proposed framework across three real-world datasets with extensive sensitivity analyses to ensure higher applicability and robustness.

We specifically use the term *false information* instead of *misinformation* or *disinformation* to maintain clarity. Misinformation refers to mistakenly shared information without intent to deceive, while disinformation is deliberately spread to mislead others, characterized by malicious purpose. The nuanced differences between misinformation and disinformation and their separate impacts fall outside the purview of our research focus.

This work focuses on the CIM setting, where two competing parties aim to optimize their seed selection strategies. We defer detailed analysis of how user priors and behavioral parameters affect propagation dynamics to future work, with discussion on the potential impacts of user heterogeneity provided in Appendix E.2. Additionally, since our focus is on information propagation and opinion updating, we

follow the decision-making framework of SL. When users accumulate sufficient evidence, it contributes to belief mass, which becomes the primary driver of their final decision. In such cases, the influence of the base rate becomes negligible.

2 RELATED WORK

This section presents an overview of related work in competitive information maximization, opinion and information propagation models, exploitation-exploration balance in DRL, and evidential deep learning. Additionally, we highlight the distinctions of our proposed approach from these works and summarize them in Table 1 and Appendix C (Tables S-4–S-6) in the supplement for easy reference.

2.1 Competitive Information Maximization

Domingos and Richardson [13] introduced Influence Maximization (IM), focusing on selecting optimal seed nodes to maximize influence in Online Social Networks (OSNs). Kempe et al. [14] formalized IM as a discrete stochastic optimization problem, laying its theoretical foundation.

IM extends to Competitive IM (CIM) by introducing multiple competing entities. Early works applied game theory, such as [15] finding Nash Equilibrium in competitive networks. More recent methods include [16] with Sandwich Approximation, eliminating Monte Carlo simulations [17], and developing Competitive Reverse Influence Estimation (CRIE) [18]. Tsaras et al. [19] proposed Awareness-to-Influence model collective IM, Galante et al. [20] leveraged game theory for influencer competition. Community structure has been explored for CIM [21, 24]. Xie et al. [22], Liu et al. [23] improved influence estimation using community borders and user preferences. These methods assume well-defined community structures, which may not always exist.

Reinforcement Learning (RL) has also been applied. Lin et al. [8] introduced STORM, a RL framework for multi-round CIM. Chung [7] extended this with spectral community detection. Ali et al. [5, 6] refined DRL-based CIM by incorporating network exploration and dynamic graph embeddings, and further integrated Transfer Learning (TL) to reduce training time [34]. He et al. [25] incorporated Nash Equilibrium with Q-learning for strategic robustness. On the other hand, other works address specific constraints in terms of time [26], user behaviors [27, 29], and specific group [28]. Uncertainty [30, 31] and fairness in CIM problem [32, 33] also have been explored.

Despite advancements, CIM research faces key limitations. Many rely on greedy algorithms [18, 21, 22, 28, 30, 32], lacking global optimality. Others assume full knowledge on network and competitor [7, 8, 16, 17], or oversimplify user's and competitor's behavior [5, 6, 19, 20, 25, 27]. Fairness-aware methods use simplistic metrics [32, 33]. Our work improves CIM realism by modeling opinion dynamics and incorporating uncertainty in user beliefs and network structure. Table 1 highlights our contributions by summarizing key differences between our work and existing approaches.

2.2 Opinion Models

The Voter [35], DeGroot [36], and Friedkin-Johnsen (FJ) [37] models focus on opinion dynamics but lack a competitive

TABLE 1
COMPARISON OF RELATED WORKS IN CIM RESEARCH AND OUR PROPOSED APPROACH

Work	Contribution	Difference from Our Work
[13, 14]	Introduced and formalized the IM problem.	Focuses on traditional IM, not competitive settings.
[15]	Applied game theory and Nash Equilibrium to CIM.	Assumes rational competitors and static strategies.
[16–18]	Developed heuristics like SPBA and CRIE for seed selection.	Lacks opinion modeling and uncertainty-awareness.
[19, 20]	Used game theory and probabilistic influence propagation.	Does not incorporate uncertainty-aware opinion evolution.
[21–24]	Proposed community structure-based CIM.	Relies on predefined communities, limiting adaptability.
[5–8]	Applied reinforcement learning (RL) for multi-round CIM.	Models only a single agent, lacks adversarial interaction.
[25]	Integrated budget constraints and multi-stage CIM.	Assumes full network knowledge and static competitors.
[26–29]	Considered user constraints such as time delays and interests.	Does not model uncertainty-aware opinion updates.
[30, 31]	Addressed CIM uncertainty via probabilistic modeling.	Lacks belief-based opinion dynamics and uncertainty modeling.
[32, 33]	Explored fairness-aware seed selection strategies.	Uses simplistic fairness metrics, lacking opinion evolution.
Our Work	Proposes a dual-agent DRL CIM framework with SL for opinion dynamics and DRL-based uncertainty-aware decision-making.	Moves beyond binary opinion models, integrates adversarial RL for CIM, and applies multi-faceted uncertainty modeling.

nature. They adopted a random neighbor's opinion, the DeGroot model averages neighbors' opinions, and FJ extends it by introducing stubborn users. The Deffuant-Weisbuch (DW) [38] and Hegselmann-Krause (HK) [39] models use bounded confidence, emphasizing pairwise interactions and clustering, being less suitable for competitive scenarios.

The Galam model [40] focuses on binary opinions and employs majority rule, oversimplifying CIM. The Sznajd model [41] relies on local interactions, making it unsuitable for global influence spread. Noisy models [42, 43] add randomness to existing models to simulate real-world phenomena, but randomness complicates CIM modeling.

Wildly used CIM propagation models include Independent Cascade (IC) and Linear Threshold (LT) [14, 15, 24, 27, 28], with many variants [5, 6, 16–19, 23, 26, 29–34]. IC assumes independent neighbor influence, while LT activates nodes based on cumulative influence thresholds. Despite their effectiveness, their assumption cannot fully reflect the realistic CIM process. Our work enhances these models by incorporating dynamic opinion dynamics and evaluating their resilience to misinformation.

2.3 Exploitation-Exploration Balance in DRL

A key challenge in reinforcement learning is balancing exploration and exploitation (EE) [2]. One simple method is *Epsilon-greedy* [2], where the agent explores randomly with probability ϵ and otherwise exploits the best-known action. ϵ decays over time to shift towards exploitation.

Reward shaping [44] introduces additional rewards to guide agents to learn, especially in sparse reward environments. Other methods encourage agents to visit new states or choose diverse actions, such as *Entropy Regularization* (ER) [45], which tries to maximize policy entropy. *Random Network Distillation* (RND) [46] uses prediction errors from another network as intrinsic rewards. *Pseudo-counts* [47] and *Never Give Up* (NGU) [48] reward visiting less frequent or novel states. Goal-based methods focus exploration on unknown areas. *Go-Explore* [49] memorizes visited states and selects goals probabilistically for further exploration. *Reverse Curriculum Generation* [50] starts from goal states and works backward to explore the environment.

Uncertainty-based methods balance EE using confidence measures. *Upper Confidence Bound* (UCB) [3] selects actions based on reward estimates plus a confidence term, widely used in multi-armed bandits and RL. *Bootstrapped DQN* [51, 52] uses multiple Q-networks to estimate uncertainty

and guide exploration, effective in continuous control but computationally expensive. *Variational Information Maximizing Exploration* (VIME) [53] uses Bayesian neural networks for information gain by reducing uncertainty. However, existing EE strategies have yet to fully explore the role of different uncertainty types in decision-making.

2.4 Evidential Deep Learning

Evidential Deep Learning (EDL), leveraging Dempster-Shafer Theory (DST) for uncertainty estimation, was refined by [10] to enhance confidence measures via a specialized loss function. This improves performance in noisy environments and out-of-distribution (OOD) detection. Further advancements integrated Wasserstein GANs (WGANS) and Normal Inverse Gamma (NIG) distributions for better OOD detection and stereo matching uncertainty [54, 55].

Deng et al. [56] introduced Fisher Information-based EDL (IEDL), dynamically reweighting loss via the Fisher Information Matrix (FIM) for improved high-uncertainty estimation. Shao et al. [57] proposed Dual-level Deep Evidential Fusion (DDEF), enhancing multimodal learning by linking neural outputs with Dirichlet parameters and fusing evidence across modalities. Ancha et al. [58] developed a Dirichlet-based framework for pixel-wise uncertainty estimation in semantic segmentation, aiding robotic navigation. Duan et al. [59] extended classification uncertainty quantification using total covariance, offering richer insights beyond entropy-based approaches.

EDL has diverse applications. Li et al. [60] adapted it for brain tumor segmentation, generating reliable uncertainty maps for clinical use. Schreck et al. [61] showed EDL quantifies uncertainty effectively in Earth system science, aiding weather and climate modeling. However, careful tuning of its loss function is required to balance accuracy and uncertainty, posing challenges in parameter optimization.

3 BACKGROUND: SUBJECTIVE LOGIC (SL)

We employ *Subjective Logic* (SL) to model user uncertainty in Section 4 and develop an uncertainty-aware DRL-based CIM framework in Section 5. The reason SL was chosen over Bayesian Networks and Fuzzy Logic is because it explicitly models opinion uncertainty using belief, disbelief, and uncertainty mass. Additionally, SL incorporates trust propagation and dynamic opinion updates, making it well-suited for capturing evolving user beliefs in competitive

influence maximization. This section provides the necessary background on SL to facilitate understanding of our proposed opinion model and CIM framework.

3.1 Binomial Opinion in SL

In binary logic, a user's opinion toward something is either 100% believed in or not, which rarely happens. To better capture real-world opinion dynamics, we model user beliefs using Subjective Logic (SL) [4], which quantifies degrees of uncertainty in beliefs. For propositions classified as true or false, we use SL's binomial opinion model to represent user opinions in OSNs, defined as:

$$\omega = (b, d, u, a), \quad (1)$$

where belief b represents agreement with true information, disbelief d denotes agreement with false information (or disagreement with true information), and uncertainty u accounts for insufficient evidence. For simplicity, we omit the user ID (e.g., i) in this opinion representation. These components lie within the range $[0, 1]$ and satisfy the additivity constraint $b + d + u = 1$. These components are:

$$b = \frac{r}{r + s + W}, \quad d = \frac{s}{r + s + W}, \quad u = \frac{W}{r + s + W}, \quad (2)$$

where r and s are the numbers of evidence to support b and d , respectively, and W refers to the number of uncertain evidence that cannot be judged as true or false, supporting neither b nor d . The base rate a represents the prior belief favoring true information, with $1 - a$ corresponding to false information. It is a real number within the range $[0, 1]$. While this paper focuses on two opinions using SL's binomial opinion model, the approach can be readily extended to multiple opinions involving additional parties.

When making opinion-based decisions (e.g., selecting a product), users account for uncertainty by leveraging the *projected probability* in SL, represented by the projected belief ($P(b)$) and projected disbelief ($P(d)$). They are defined as:

$$P(b) = b + a \cdot u, \quad P(d) = d + (1 - a) \cdot u, \quad (3)$$

where $P(b) + P(d) = 1$. Since users choose between b and d , they interpret uncertainty based on their prior belief, $a = \{a, 1 - a\}$. Specifically, user i is considered part of TP if $P(b_i) > 0.5$ and is deemed to believe false information if $P(d_i) > 0.5$. When a user i has $P(b_i) = P(d_i)$ under a uniform prior belief (i.e., base rate a), they are not counted toward TP or FP. However, such cases are extremely rare, so the sum of users in TP and FP is nearly equal to the total number of users in a given network.

3.2 Dissonance in SL

We consider two types of uncertainty in the user's opinion formulated by SL [62]: *vacuity* and *dissonance*. *Vacuity* refers to uncertainty caused by a lack of evidence, while *dissonance* indicates uncertainty due to conflicting evidence.

For a multinomial opinion in SL, the dissonance in the opinion is calculated as [62]:

$$\omega^{\text{diss}} = \sum_{i \in K} \frac{\beta_i \sum_{j \in K/i} \beta_j \text{Bal}(\beta_j, \beta_i)}{\sum_{j \in K/i} \beta_j} \quad (4)$$

where

$$\text{Bal}(\beta_j, \beta_i) = 1 - \frac{|\beta_j - \beta_i|}{\beta_j + \beta_i}.$$

Here, $\text{Bal}(\beta_j, \beta_i)$ measures the degree to which belief masses are evenly distributed, indicating uniformity among them. A higher balance level suggests greater inconclusiveness in the opinion, making decision-making more challenging.

A binomial opinion is a special case of a multinomial opinion, relying on the Beta distribution [4]. For clarity, the dissonance of a binomial opinion is computed as follows:

$$\omega^{\text{diss}} = (b + d) \cdot \text{Bal}(b, d), \quad (5)$$

where $\text{Bal}(b, d)$ represents the relative balance between belief masses b and d and is defined as:

$$\text{Bal}(b, d) = 1 - \frac{|b - d|}{b + d}. \quad (6)$$

We incorporate vacuity (i.e., u in an opinion) and dissonance uncertainty measures in developing an uncertainty-aware opinion model (UOM) in Section 4 and in designing uncertainty-aware exploration-exploitation strategies for DRL, as described in Section 5.

3.3 Uncertainty Maximization

Uncertainty maximization (UM) in SL [4] is computed by adjusting the uncertainty mass (\ddot{u}_i) to its highest possible value while preserving the projected probabilities. To formally put, \ddot{u}_i is estimated by:

$$\ddot{u}_i = \min \left[\frac{P(b_i)}{a_i}, \frac{P(d_i)}{1 - a_i} \right], \quad (7)$$

$$\ddot{b}_i = P(b_i) - a_i \cdot \ddot{u}_i, \quad \ddot{d}_i = P(d_i) - (1 - a_i) \cdot \ddot{u}_i,$$

The UM \ddot{u}_i is determined as the minimum of the scaled projected belief and disbelief, ensuring consistency with the prior base rate (a_i). The belief (\ddot{b}_i) and disbelief (\ddot{d}_i) are then recalibrated by subtracting their respective contributions to uncertainty, maintaining the overall probability distribution. We incorporate this UM in designing uncertainty-aware exploration-exploitation strategies in Section 5.3.

3.4 Discounting Operator in SL

In SL, a *trust filter*, c_i^j , is used to discount j 's opinion because user i accepts user j 's opinion to the extent that user i trusts user j . Each component of $\omega_{i \otimes j} = (b_{i \otimes j}, d_{i \otimes j}, u_{i \otimes j}, a_{i \otimes j})$ is estimated by:

$$b_{i \otimes j} = c_i^j b_j, \quad d_{i \otimes j} = c_i^j d_j, \quad (8)$$

$$u_{i \otimes j} = 1 - c_i^j (1 - u_j), \quad a_{i \otimes j} = a_j.$$

3.5 Consensus Operator in SL

The consensus operator \oplus updates user i 's opinion by combining its belief weighted by j 's uncertainty and vice versa. Disbelief is updated using the same weighting mechanism. User i 's uncertainty decreases as the product of both users' uncertainty levels, and opinion updates continue unless either user's uncertainty reaches zero, indicating complete confidence. User i 's updated opinion, incorporating trust in user j , is obtained by applying the consensus operator \oplus

to its original opinion and the discounted opinion of j . The result is given by $\omega_i \oplus \omega_{i \otimes j}$:

$$\begin{aligned} b_i \oplus b_{i \otimes j} &= \frac{b_i(1 - c_i^j(1 - u_j)) + c_i^j b_j u_i}{\beta}, \\ d_i \oplus d_{i \otimes j} &= \frac{d_i(1 - c_i^j(1 - u_j)) + c_i^j d_j u_i}{\beta}, \\ u_i \oplus u_{i \otimes j} &= \frac{u_i(1 - c_i^j(1 - u_j))}{\beta}, \\ a_i \oplus a_{i \otimes j} &= \frac{(a_i - (a_i + a_j)u_i)(1 - c_i^j(1 - u_j)) + a_j u_i}{\beta - u_i(1 - c_i^j(1 - u_j))}, \\ \beta &= 1 - c_i^j(1 - u_i)(1 - u_j) \neq 0. \end{aligned} \quad (9)$$

Therefore, user i 's updated opinion, ω'_i , is given by:

$$\begin{aligned} \omega'_i &= \omega_i \oplus \omega_{i \otimes j} = (b'_i, d'_i, u'_i, a_i) \\ &= (b_i \oplus b_{i \otimes j}, d_i \oplus d_{i \otimes j}, u_i \oplus u_{i \otimes j}, a_i \oplus a_{i \otimes j}), \end{aligned} \quad (10)$$

where ω_i is user i 's prior opinion, and $\omega_{i \otimes j}$ is j 's discounted opinion based on i 's acceptance level. Using SL's discounting and consensus operators, we propose uncertainty-aware opinion models in Section 4.

4 UNCERTAIN OPINION MODELS

This section outlines the user types and uncertainty-aware opinion models, detailing how uncertain opinions are updated, shared, and interpreted in this work.

4.1 User Types

We categorize users in OSNs into three types:

- *True information propagators* (TIPs) are seed nodes selected by the *true party* (TP) with an initial opinion $\omega = (b, d, u, a) = (b \rightarrow 1, d \rightarrow 0, u \rightarrow 0, a = 1)$. This reflects a strong belief in true information ($b \approx 1$) and no belief in false information ($d \approx 0$).
- *False information propagators* (FIPs) are seed nodes selected by the *false party* (FP), with an initial opinion of $\omega = (b, d, u, a) = (b \rightarrow 0, d \rightarrow 1, u \rightarrow 0, a = 0)$. This indicates FIPs strongly believe in false information ($d \approx 1$) while lacking belief in true information ($b \approx 0$).
- *Legitimate users* in OSNs are regular users with highly uncertain opinions, initialized as $\omega = (b, d, u, a) = (b \rightarrow 0, d \rightarrow 0, u \rightarrow 1, a = 0.5)$.

In this work, we adopt a uniform base rate for all legitimate users as a practical modeling assumption. This choice reflects a neutral starting point, allowing the model to update beliefs based solely on the evidence accumulated through network interactions. While individual users may have diverse initial beliefs, this variability often diminishes over time as repeated interactions and social dynamics drive convergence toward shared perspectives, particularly within topic-focused communities [63]. Thus, a uniform base rate provides a reasonable approximation for capturing the broader trends in social information propagation.

Since TIPs and FIPs hold strong beliefs, they do not change their opinions, while *legitimate users* update their opinions by adjusting (b, d, u, a) based on Eq. (10) and the consensus operation in Eq. (9), describing how opinions shift by supporting evidence during interactions.

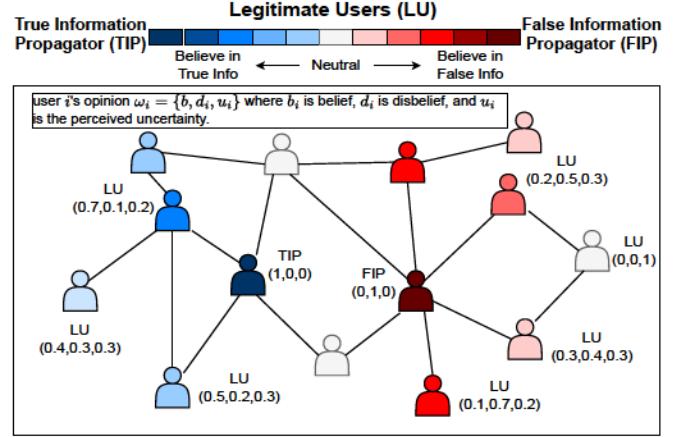


Fig. 1. Overview of the CIM problem with Subjective Logic. An *SL-based competitive network* visualizes belief strength via a color gradient where lighter shades indicate higher uncertainty, with white for neutrals. Blue represents true information supporters (darkest: TIP), and red denotes false information supporters (darkest: FIP). Each user i 's opinion is represented by $\omega_i = \{b_i, d_i, u_i\}$, we ignore prior belief a_i because it is predefined as a constant for all users.

To illustrate SL-based opinion propagation in a competitive OSN with FP and TP, Fig. 1 presents an overview of the CIM problem with Subjective Logic. The color gradient represents the level of belief, where lighter shades indicate lower certainty, and white denotes neutral users who are totally uncertain. Blue signifies belief in true information, with the darkest blue representing the TIP chosen by the True party (TP). Red represents belief in false information, with the darkest red indicating the FIP selected by the False party (FP).

4.2 Updating, Sharing, and Reading Opinions

User i 's opinion is represented by $\omega_i = (b_i, d_i, u_i, a_i)$. The user's behavior, uc_i , comprises opinion updating, sharing, and reading, expressed as $uc_i = \{updating_i, reading_i, sharing_i\}$, as detailed below.

4.2.1 Opinion Updating

Recall that SL is chosen for its capability to formulate subjective opinions while handling multiple types of uncertainty causes, namely *vacuity* and *dissonance*. When user i encounters user j and reads their information (i.e., ω_j), user i updates their opinion accordingly. The *consensus operator* in SL [4] is utilized to compute user i 's updated opinion, as detailed in Section 3.5.

This study presents three opinion models (OMs): uncertainty-based trust (UOM), homophily-based trust (HOM), and no-trust (NOM), each with a distinct *trust filter* definition (see Section 3.4):

$$c_i^j = \begin{cases} uc_i^j, & \text{if UOM} \\ hc_i^j, & \text{if HOM} \\ nc_i^j, & \text{if NOM} \end{cases} \quad (11)$$

where c_i^j refers to the trust filter used in Eq. (8).

- *Uncertainty-based OM (UOM)*: This model represents users who seek new information when lacking sufficient

evidence to decide [64–66]. UOM applies uncertainty-based trust, where the trust filter uc_i^j is defined as:

$$uc_i^j = (1 - u_i)(1 - u_j), \quad (12)$$

where $u_j \neq 1$ and $u_i \neq 1$, as the initial uncertainty of a legitimate user is close to 1 but not exactly 1 (see Sections 4.1 and 6.2). To prevent zero β in Eq. (9), we apply uncertainty maximization (UM)(Section 3.3) when a user's uncertainty is low, but conflicting evidence prevents a decision. Given thresholds TUM_v and TUM_d , if $u_i < TUM_v$ (low vacuity) and $\omega_i^{\text{diss}} > TUM_d$ (high dissonance), user i 's opinion is updated as $\tilde{\omega}_i = (\tilde{b}_i, \tilde{d}_i, \tilde{u}_i, a_i)$, where \tilde{b}_i , \tilde{d}_i , and \tilde{u}_i follow Eq. (7). The updated uncertainty \tilde{u}_i enables user i to accept new information and refine its opinion.

- **Homophily-based OM (HOM):** Homophily, the tendency of like-minded individuals to associate, influences opinion updates [67]. The homophily trust filter hc_i^j is computed using *cosine similarity*, measuring alignment between users' beliefs and disbeliefs:

$$hc_i^j = \frac{b_i b_j + d_i d_j}{\sqrt{b_i^2 + d_i^2} \sqrt{b_j^2 + d_j^2}}. \quad (13)$$

This captures opinion similarity through belief (b_i, b_j) and disbelief (d_i, d_j) , indirectly incorporating uncertainty since $b + d + u = 1$.

- **No-Trust-based OM (NOM):** Without a trust filter, the No-Trust-based OM sets $nc_i^j = 1$, allowing user i to fully accept user j 's information without any discount in opinion updates.

4.2.2 Reading & Sharing Opinions

We model a user's reading and sharing behavior following [66], which derived these behaviors from a survey on social media usage patterns. A user's reading probability, P_r , indicates how frequently they read information from neighbors, with values randomly assigned as 1 (multiple times per day), 0.5 (daily), 0.25 (weekly/monthly), or 0.1 (rarely) to reflect diverse reading habits. Similarly, a user's sharing probability, P_s , determines how often they share opinions ω , with P_s randomly assigned as 1 (always/mostly), 0.5 (half the time), 0.25 (sometimes), or 0.1 (rarely).

Since reading encourages opinion sharing [68], users share only after reading. When user i shares, it transmits its updated opinion ω_i , not the original 'True' or 'False' information. Only seed nodes, *True* and *False* information propagators (TIPs and FIPs), share the original information.

5 UNCERTAINTY-AWARE DRL-BASED CIM

5.1 Seed Set Selection Process

We use DRL to optimize strategy selection for maximizing influence spread in a multi-round process. In each round, the FP selects a seed node (turns it into an FIP), and then the FIP starts sharing its opinion with its neighbors and propagates further until all nodes are reached. The TP next selects a TIP to propagate true information. Information spreads through the network via *Breadth-First Search*, with nodes deciding to update and propagate based on their reading and sharing frequency. An episode consists of T rounds, matching the number of seed nodes.

5.2 Formulation of DRL-based Best Seed Selection

In this section, we describe the formulation of a DRL agent's states, actions, and rewards in this work.

States. At each round t , the state s_t is defined as:

$$s_t = \left\{ \sum_{i,j \in \mathcal{U}} e_{i,j}, \max_{i \in \mathcal{U}} \deg_i \right\}. \quad (14)$$

Here, *free nodes*, defined as $\{j \in \mathcal{U} | u_j \geq 0.5\}$, represent users with high uncertainty who have not aligned with either party, where u_j denotes j 's uncertainty (i.e., vacuity). The term $\sum_{i,j \in \mathcal{U}} e_{i,j}$ counts edges among free nodes, while $\max_{i \in \mathcal{U}} \deg_i$ gives the highest degree among them.

Actions. The action space at round t is $a_t = \{a_t^{AF}, a_t^{BF}, a_t^{SGF}, a_t^{CF}\}$, focusing on user behavior and centrality in the OSN, and each action means as follows:

- **Active First (AF, a_t^{AF})** prioritizes the most active user, determined by the highest $P_r \times P_s$, where P_r and P_s represent the user's reading and sharing probabilities.
- **Blocking First (BF, a_t^{BF})** targets neighbors of the opponent's party with the highest *free degree* (connected to *free nodes*). A user belongs to the TP if $P(b_i) > 0.5$, or the FP if $P(d_i) > 0.5$ where $P(b)$ refers to an expected probability for a given belief b , as discussed in Eq. (3).
- **SubGreedy First (SGF, a_t^{SGF})** selects the node with the most neighbors within k -hops [7, 8] (e.g., $k = 2$ to balance efficiency and effectiveness).
- **Centrality First (CF, a_t^{CF})** selects the user with the highest degree centrality. We use degree as a centrality metric, considering its efficiency. Investigating the performance of using a different centrality metric is beyond the scope of our work.

Rewards. We use immediate rewards to enhance the learning process. At each round t , the rewards for each party, TP or FP, are the sum of users' beliefs aligned with each party by:

$$R_t^{TP} = \sum_{i \in TP} b_t^i, \quad R_t^{FP} = \sum_{i \in FP} d_t^i, \quad (15)$$

where b_t^i and d_t^i refer to user i 's belief mass at round t .

The respective accumulated rewards over an episode from $t = 0$ to $t = T$ are:

$$R_T^{TP} = \sum_{t=1}^T \gamma^{T-(t-1)} R_t^{TP}, \quad R_T^{FP} = \sum_{t=1}^T \gamma^{T-(t-1)} R_t^{FP}, \quad (16)$$

where $\gamma \in [0, 1]$ is the discount factor. We integrate model prediction uncertainty to enhance decision-making in the EE trade-off. After obtaining the action distribution from the policy network's final layer, we interpret it as belief mass β_i for each action, ensuring $\sum_{i \in K} \beta_i = 1$ with initial uncertainty mass $\nu = 0$. To introduce positive uncertainty mass $\tilde{\nu}$, we apply the uncertainty maximization (UM) technique [4], detailed in Section 3.3. We assume no prior knowledge for belief propositions, assigning an equal base rate probability.

In addition to vacuity (uncertainty mass $\tilde{\nu}$) for assessing evidence sufficiency, we compute dissonance to detect conflicting decisions, where a flatter action distribution indicates lower confidence. For example, with belief masses 0.24, 0.25, 0.24, 0.27, the system's highest probability (0.27) does not imply confidence in selecting the optimal action

due to the distribution's uniformity. Dissonance (see Eq. (4)) quantifies this distinction, implying that the flatter the distribution, the lower the dissonance. These uncertainty estimates guide EE decisions, as discussed below. Fig. 2 illustrates the proposed Uncertainty-Aware DRL-based node selection process, where each party's DRL agent processes OSN states through a policy network to determine action distributions.

5.3 Exploration-Exploitation (EE) Strategies

A DRL agent looks for the best action by randomly choosing an action called *exploration* and relying on their experiences, called *exploitation*. In this work, we evaluate the following EE strategies: the first three are our proposed uncertainty-aware approaches, and the rest are state-of-the-art methods. Their performance will be compared to highlight the strengths of each strategy. The considered EE strategies are:

- EE Using Vacuity (VAC-EE): This strategy balances EE based on vacuity (\bar{v}), which quantifies uncertainty from insufficient evidence. If vacuity exceeds a threshold (T_v), the agent explores due to limited knowledge. Otherwise, it exploits by selecting the best-known action. This approach explores under high uncertainty from a lack of information and exploits otherwise.
- EE Using Dissonance (DIS-EE): This strategy balances EE based on dissonance (ω^{diss}), which quantifies conflicting evidence in decision-making. If dissonance exceeds a threshold (T_d), the agent explores to resolve inconsistencies. Otherwise, it exploits the best-known action. This approach adapts decisions by exploring under conflict and exploiting otherwise.
- EE Using Vacuity and Dissonance (VD-EE): This strategy dynamically balances exploration and exploitation using vacuity (\bar{v}) and dissonance (ω^{diss}). If vacuity exceeds a threshold (T_v), the agent explores due to insufficient knowledge. Otherwise, it evaluates dissonance; if above (T_d), exploration occurs to address conflicting information. When both measures are low, the agent exploits the best-known action. This adaptive approach optimizes decision-making by exploring under high uncertainty and exploiting when confidence is sufficient.
- EE Using Entropy (ENT-EE): This strategy measures the randomness in the agent's action selection using the entropy of the action probability distribution [69]. High entropy indicates greater exploration, while low entropy reflects a more deterministic policy, favoring exploitation.
- EE Using Epsilon Greedy (EPS-EE) [2]: A widely used exploration strategy in reinforcement learning, epsilon-greedy balances exploration and exploitation through random action selection. With probability ϵ , the agent explores by selecting a random action, while with probability $1 - \epsilon$, it exploits by choosing the action it currently estimates to be the best based on gathered information.
- EE Using Entropy Regulation (ER-EE) [45]: This strategy incorporates entropy into the reward function by adding an entropy term scaled by a coefficient (which determines the importance of entropy or exploration), the modified objective encourages the agent not only to maximize the expected reward but also to maintain a level of randomness in its action selection.

• Upper Confidence Bound (UCB-EE) [3]: UCB balances exploration and exploitation by adjusting exploration based on action uncertainty. The agent estimates expected rewards from past observations and computes an uncertainty measure, which decreases as more data is gathered. Less-explored actions are prioritized for exploration, increasing the likelihood of discovering optimal choices.

Algorithm 1 details the first three proposed uncertainty-aware EE strategies.

Algorithm 1 Uncertainty-Aware Exploration-Exploitation Strategies (VAC-EE, DIS-EE, and VD-EE)

```

1: Given an opinion formulated based on the actions taken by a DRL
   agent for seed node selection,
2:  $\bar{v} \leftarrow$  vacuity,  $\omega^{\text{diss}} \leftarrow$  dissonance
3:  $T_v \leftarrow$  vacuity threshold,  $T_d \leftarrow$  dissonance threshold
4: procedure VACUITY-BASED EE (VAC-EE)
5:   if  $\bar{v} > T_v$  then
6:     explore
7:   else
8:     exploit
9:   end if
10: end procedure
11: procedure DISSONANCE-BASED EE (DIS-EE)
12:   if  $\omega^{\text{diss}} > T_d$  then
13:     explore
14:   else
15:     exploit
16:   end if
17: end procedure
18: procedure VACUITY AND DISSONANCE-BASED EE (VD-EE)
19:   if  $\bar{v} > T_v$  then
20:     explore
21:   else
22:     if  $\omega^{\text{diss}} > T_d$  then
23:       explore
24:     else
25:       exploit
26:     end if
27:   end if
28: end procedure

```

5.4 Partially Observable Network

Following [70], a partially observable network is defined as an undirected graph $G = (V, E)$ with an observable subset $G' = (V, E')$, where $E' \subset E$. Only a subset of edges is visible to the DRL agent when observing the state in OSNs.

In an OSN, two parties compete to maximize user belief in their opinion, starting with all users neutral. In each round, FP selects a seed node based on the policy network's action distribution to maximize influence. The chosen node (FIP) propagates opinions as per Section 4.2. TP then selects a seed node to counter FP's influence. This process repeats until the predefined seed limit is reached or no further selections are possible. TP's influence is evaluated by the number of nodes adopting its opinion.

6 EXPERIMENT SETUP

6.1 Algorithm Settings

We employ Proximal Policy Optimization (PPO) [71] in our DRL framework for optimal seed node selection, as PPO enables multiple updates with the same data batch, reducing computational complexity given CIM's long data collection time. TP utilizes a trained DRL model for node selection,

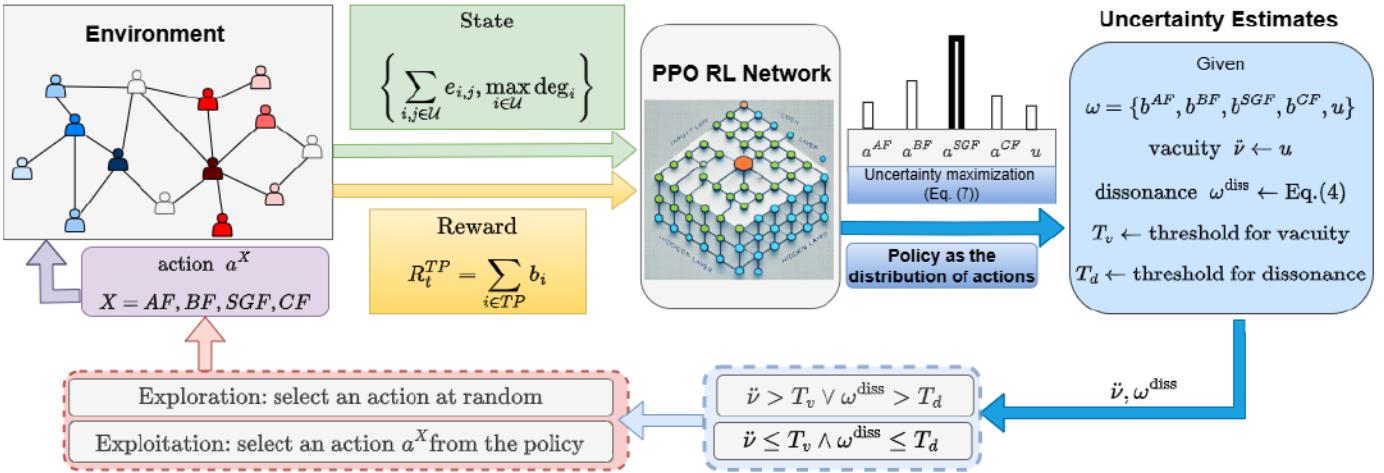


Fig. 2. Overview of the Proposed Uncertainty-Aware DRL-based Node Selection: Each agent (i.e., DRL agents for TP and FP, respectively) processes OSN states through a policy network to determine the action distribution. Vacuity and/or dissonance guide the exploration-exploitation decision-making: *low uncertainty* favors exploitation (selecting the highest-probability action for maximum influence), while *high uncertainty* promotes exploration (random action selection). The policy network is iteratively updated based on rewards.

while FP can choose from six strategies, such as DRL, Active First (AF), Blocking First (BF), SubGreedy First (SGF), and Centrality First (CF), as described in Section 5.2. We train TP's DRL agent against each FP strategy and evaluate their performance in Section 7.

The opinion propagation and update processes in Section 4 were implemented using custom Python scripts developed for this study. The code, written in Python 3.9, leverages standard scientific libraries including NumPy, SciPy, and NetworkX for efficient graph handling and matrix operations. The source code is available at our GitHub repository here.

6.2 Parameterization

We initialize SL-based opinions for legitimate users using Eq. (2) with $(r, s, W) = (1, 1, 101)$, indicating high uncertainty. TIPs start with $(100, 1, 2)$, reflecting strong confidence in true information, while FIPs begin with $(1, 100, 2)$, showing confidence in false information. In UOM, when conducting Uncertainty Maximization toward individual user's opinion, we set vacuity and dissonance thresholds at $TUM_v = 0.01$ and $TUM_d = 0.6$ to optimize influence and maintain effectiveness. The effect of varying these thresholds is discussed in Appendix A.1. In SubGreedy First, $k = 2$ balances efficiency and effectiveness, while DRL employs a discount factor $\gamma = 0.99$ for long-term influence.

For uncertainty-aware EE strategies, we set vacuity and dissonance thresholds via sensitivity analysis to optimize influence. Table 2 summarizes the selected thresholds. Vacuity (T_v) captures overall uncertainty in the action distribution, while dissonance (T_d) reflects conflicting preferences among top actions. We present an example of how to select thresholds as well as sensitivity analysis on these thresholds in the supplement document at Appendix A.2.

For FP strategies like AF, BF, CF, and DRL, the TP's DRL agent performs best with a low vacuity threshold ($T_v = 0.3$), enabling confident, decisive actions. Further lowering T_v is counterproductive, as it risks missing high-reward nodes. For SGF, a higher T_v is preferred to respond early to FP's rapid influence over dense local clusters.

TABLE 2
THRESHOLDS SUMMARIZATION

FP strategy	T_v	T_d
AF	0.3	0.9
BF	0.3	0.5
SGF	0.9	0.7
CF	0.3	0.7
DRL	0.3	0.9

Dissonance thresholds are typically moderate to high. High T_d values suit fast-changing scenarios (e.g., DRL, AF), tolerating internal conflict to enable timely actions. For BF, CF, and SGF, moderate T_d balances decisiveness and caution. Very low T_d makes the agent overly conservative, risking failure to capture critical or contested nodes.

FP and TP alternately select 50 seeds for propagation, with FP initiating the process, followed by TP. This models real-world false information mitigation, where true information counters ongoing false information. TP propagates twice after each seed selection, while FP propagates once. Given the critical role of human judgment and public awareness, true information is assumed to propagate more effectively in OSNs. All results are averaged over 50 runs.

Experiments ran on an HPE Apollo 6500 with AMD EPYC 7742 processors (2.25 GHz base, 3.4 GHz boost). Key design parameters, definitions, and defaults are in Table S-1, Appendix A of the supplementary document. Unless stated otherwise, performance evaluations assume perfect network observability, except when analyzing its variations.

6.3 Network Datasets

We validate our framework using three datasets: the URV Email Network [72], an undirected graph of email communications at *Universitat Rovira i Virgili*, Spain, a Facebook social circles dataset [73] (FBN, Facebook network); and a Facebook page network (FBPN) [74]. These networks effectively model CIM, capturing real user interactions. The URV Email and Facebook datasets reflect how individuals share ideas, opinions, and information, while the Facebook page-page network represents content spread by entities like brands or political campaigns via influential hubs. The

TABLE 3
CHARACTERISTICS OF NETWORK DATASETS

Dataset	Nodes	Edges	Degree	Betweenness
URV Email [72]	1133	5452	7	703.94
Facebook [73]	4039	88234	25	47.57
FB Page [74]	22470	171002	7	5496.91

datasets' diverse structures and sizes enable a comprehensive evaluation of our framework's generalizability on complex OSNs. Table 3 provides detailed statistics. For centrality measures, we report unnormalized medians instead of means, as social networks are scale-free, where most nodes have low degrees and a few have very high degrees, making the median more representative.

Due to the lack of datasets that jointly provide network structure, user communication behaviors, and message-level content, we use real network topologies and synthetically parameterize key behavioral variables to assess their impact on information propagation, as in Section 4.2.

6.4 Metrics

We evaluate the proposed approaches using the following metrics: (1) Percentage (%) of nodes in TP, measured by the number of users aligned with TP (Eq. (3)). Users with $P_i(b_i) > 0.5$ are classified under TP, indicating its influence, denoted as n^{TP} ; (2) Algorithmic efficiency, assessed by the simulation's runtime per round (Section 7.5, Table 5); (3) Uncertainty estimates, evaluated by the measures of vacuity, dissonance, and entropy.

6.5 Comparing Schemes

To evaluate our framework, we compare its performance and runtime against existing DRL-based CIM frameworks. Unlike [5–8], our approach employs an SL-based dynamic opinion model, enabling opinion evolution through interactions. For a fair comparison, all models operate under identical conditions, including the same opinion models (HOM, UOM, or NOM), seed selection settings (FP moves first, equal propagation times, $T = 50$ seed nodes), and opinion update mechanisms (Section 5). We use ER-EE as the general EE trade-off strategy across all schemes. While adapting these frameworks may affect their original performance, this standardization ensures a realistic evaluation.

We evaluate the following CIM algorithms: (1) DRIM-A: Our proposed framework DRIM, where 'A' denotes the AF strategy (Section 5.1), distinguishing it from DRIM-NA. (2) DRIM-NA: A variant excluding the AF strategy (No-AF) from the action space to assess AF's impact in DRIM. (3) STORM [8]: Originally designed with binary opinions for node occupation, we adapt it by defining *free nodes* ($\{j|u_j \geq 0.5\}$) as unoccupied and merging max-weight and max-degree actions, as our datasets are unweighted. (4) C-STORM [7]: An extension of STORM incorporating a preliminary community detection step for seed selection. We adapt it using our opinion model and the same definition of a free node as used in STORM.

7 EXPERIMENTAL RESULTS & ANALYSES

7.1 Performance Analyses of Opinion Models

We first analyze the effect of opinion models by comparing performance across three models using four entropy

TABLE 4
TRUE PARTY (TP)'S INFLUENCE IN n^{TP} UNDER VARIOUS CIM ALGORITHMS WITH DIFFERENT OMS

Scheme / OM	AF	BF	SGF	CF	DRL
DRIM-A/UOM	1077.74	463.79	827.05	374.1	691.37
DRIM-A/HOM	15.7	263.38	4.18	4.39	15.74
DRIM-A/NOM	27.59	354.64	5.85	4.86	17.56
DRIM-NA/UOM	1078.26	529.52	959.63	525.7	576.47
DRIM-NA/HOM	38.24	274.32	5.18	4.58	4.15
DRIM-NA/NOM	60.64	331.34	6.26	4.78	6.51
C-STORM/UOM	1070.75	1082.53	524.39	229.15	551.44
C-STORM/HOM	26.88	475.27	4.57	4.55	533.17
C-STORM/NOM	15.47	509.03	6.38	4.82	442.93
STORM/UOM	1072.45	1058.80	461.19	63.61	424.09
STORM/HOM	60.61	477.76	4.96	4.92	375.17
STORM/NOM	60.58	487.43	6.42	4.69	377.3
DRIM-A-VDEE/UOM	1077.64	646.94	960.66	581.9	794.33
DRIM-A-VDEE/HOM	38.23	195.8	5.05	4.47	4.68
DRIM-A-VDEE/NOM	38.06	241.91	6.23	4.82	6.91

regulation schemes and our uncertainty-aware approaches, with TP facing various FP strategies. Table 4 presents results on the URV Email Network, where TP's influence (n^{TP}) is measured across five FP node selection strategies.

As shown in Table 4, UOM consistently achieves the highest impact for true information, as users with high uncertainty (u) are more receptive to new evidence, enabling continuous updates and broader true information spread. This aligns with human intuition – when uncertain, people seek more information before forming conclusions. Thus, even with FP's first-mover advantage, true information can still propagate effectively in the OSN. In contrast, HOM and NOM yield poor TP performance, as users exposed to false information first tend to form rigid beliefs that true information struggles to correct. This underscores the importance of early intervention and swift responses to counteract misinformation. Under HOM and NOM, STORM and C-STORM outperform DRIM-A-based schemes against BF and DRL strategies. BF is not a primary concern in false information mitigation, as false parties rarely use it due to their first-mover advantage. DRL dominance occurs when both parties employ identical DRL designs, favoring FP as the initial actor. However, against other strategies, DRIM-A performs comparably to STORM and C-STORM under HOM and NOM, demonstrating competitive efficacy.

For all subsequent experiments, we adopt UOM as the opinion model, as it facilitates true information spread despite FP's first-mover advantage. In contrast, HOM and NOM inherently block true information, leaving minimal room for improvement.

7.2 Performance Analyses of CIM Schemes

Comparative Analysis under Three Datasets: This section evaluates the performance of existing CIM algorithms and uncertainty-aware exploration-exploitation schemes. All non-uncertainty-aware schemes use entropy regulation to balance exploration and exploitation. Fig. 3 compares the performance of four ER-DRL-based TP agents under five FP seed selection strategies. The *x*-axis denotes FP's strategy, while the *y*-axis shows the percentage of nodes aligned with TP, normalized across varying network sizes.

Fig. 3 shows that across all three datasets, our DRIM-based schemes outperform under SGF, CF, and DRL strategies, remain competitive against AF, and underperform against BF. When FP uses SGF or CF, DRIM-based TPs

achieve greater influence than STORM and C-STORM, indicating that DRIM-based approaches can surpass state-of-the-art algorithms when FP selects nodes via traditional centrality measures maximizing local reach.

In the dual DRL agent case, the four DRIM-based schemes (i.e., DRIM-A-ER, DRIM-A-VDEE, DRIM-NA-ER, and DRIM-NA-VDEE) consistently outperform STORM and C-STORM. Among them, DRIM-A-VDEE is the most effective, achieving a peak influence of 70.11% in the Email dataset, while DRIM-NA-ER, though the least effective, still reaches 50.88%. In comparison, C-STORM and STORM achieve 48.67% and 37.43%, respectively. DRIM-A-VDEE surpasses them by 44.05% and 87.31%. Similar trends hold for the Facebook and Facebook Page networks, where DRIM-A-VDEE achieves 70.31% and 74.75% influence, compared to 44.82% and 28.23% for C-STORM and STORM on Facebook, and 46.26% and 40.12% on Facebook Page. DRIM-A-VDEE thus outperforms these schemes by at least 56.87% and 61.59% across datasets.

When FP employs AF, all schemes achieve nearly 100% influence. Since AF does not prioritize nodes effectively, it behaves similarly to random selection, allowing all approaches to gain high influence. When FP uses BF, DRIM-based schemes perform worse due to BF's defensive nature, which prioritizes blocking over proactive influence. This suggests that DRIM-based approaches are better suited against aggressive strategies rather than defensive ones.

Across datasets, performance variations reveal differences in adaptability. The gap between STORM and C-STORM is notably larger in the Facebook dataset than in the Email and Facebook Page datasets. When FP employs BF or DRL, C-STORM significantly outperforms STORM on Facebook but only slightly on the other two datasets. This is due to C-STORM's community-based approach, which aligns well with Facebook's social structure, while the other two datasets exhibit more uniform connectivity. In contrast, DRIM-based schemes maintain stable performance across all datasets, demonstrating their robustness and adaptability to diverse network structures.

Sensitivity Analyses Under CIM Schemes: This section examines how the number of true information propagations, network observability, and users' prior beliefs impact CIM algorithm performance. Both TP and FP use identical DRL-based agents for seed selection, enabling an assessment of TP's effectiveness against "smart" opponents. Due to space constraints, Fig. 4 presents results from the URV Email dataset, while similar findings for the other two datasets are provided in Appendix B.2 of the supplement.

Fig. 4(a) examines the impact of increasing TP's information propagations (IPs) from 1 to 3 while FP remains at a single propagation per seed. As expected, all schemes improve with more IPs, demonstrating that additional propagation enhances influence spread. Notably, the four DRIM-based schemes perform best with a single propagation per round, aligning with Fig. 3(a), effectively countering false information with minimal resources. However, beyond two IPs, influence gains plateau as TP nears its maximum reach.

In Fig. 4(b), a clear trend shows improved performance across all schemes as network observability increases from 0.7 to 1. This indicates that greater visibility of the network structure significantly enhances performance. Under

limited observability, all schemes exhibit reduced influence, highlighting the DRL agent's dependence on network state knowledge. The sharp performance increase between 0.9 and 1 suggests a high sensitivity to network observability.

Fig. 4(c) examines the impact of varying prior beliefs across different schemes. Ideally, as prior belief in TP rises (from 0.3 to 0.8), influence maximization performance should improve, as users are more inclined to accept information from TP. However, this trend is not always evident. According to Eq. (3), prior belief primarily influences outcomes when users have nonzero vacuity u , allowing initial biases to shape their acceptance of information. In high-activity settings, frequent information exchange leads to greater certainty over time, reducing u and diminishing the effect of initial beliefs. Thus, despite varying prior beliefs, differences in influence maximization remain subtle. Nevertheless, DRIM-based schemes, particularly DRIM-A-VDEE, consistently achieve the highest influence, even when initial user beliefs do not favor true information.

These results suggest that DRIM-based schemes exhibit robust and stable performance across varying network conditions, showing particular effectiveness in scenarios with limited information propagation and diverse prior beliefs.

7.3 Performance Analyses of EE Strategies

Comparative Analysis of EE Strategies Across Three Datasets: Fig. 5 compares the performance of various EE strategies within the DRIM-A framework across three datasets. In Fig. 5, we observe that TP dominates when FP employs AF, SGF, or DRL strategies while maintaining strong performance against BF and CF. This underscores the impact of FP's strategic choices on DRIM-A's influence. VD-EE consistently performs well across datasets and FP strategies. Unlike ER-EE, it benefits from incorporating uncertainty by considering both vacuity and dissonance. By accounting for missing (vacuity) and conflicting (dissonance) information, VD-EE optimizes exploitation timing. In contrast, traditional uncertainty-based strategies such as EPS-EE, UCB-EE, and ENT-EE show inconsistent results due to limited uncertainty modeling.

VAC-EE, relying solely on vacuity, excels under DRL-based node selection, even surpassing VD-EE despite lacking dissonance modeling. Since vacuity directly signals information gaps, it enables more targeted exploration in networks with uneven information distribution. Conversely, VD-EE's dual consideration of vacuity and dissonance offers deeper insights but increases decision complexity. In rapidly changing environments requiring quick adaptation, this complexity may reduce its effectiveness in balancing exploration and exploitation.

Sensitivity Analyses Under Various EE Strategies: Fig. 6 examines the effects of true information propagation, network observability, and users' prior beliefs on EE strategies. Both TP and FP employ DRL-based agents for seed selection, allowing evaluation against "smart" opponents. Due to space constraints, results for the URV Email dataset are shown, while analyses for FB networks appear in Appendix B.3 of the supplement.

Fig. 6(a) shows that all EE strategies gain influence as information propagations (IP) increase from 1 to 3, confirming that additional propagation enhances effectiveness.

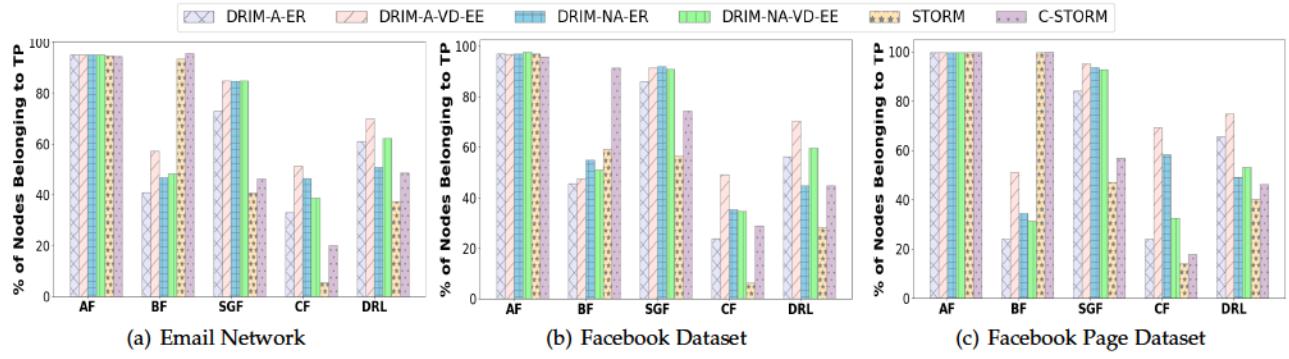


Fig. 3. Performance comparison of the considered CIM algorithms with respect to % of nodes in the true party (TP), representing TP's influence.

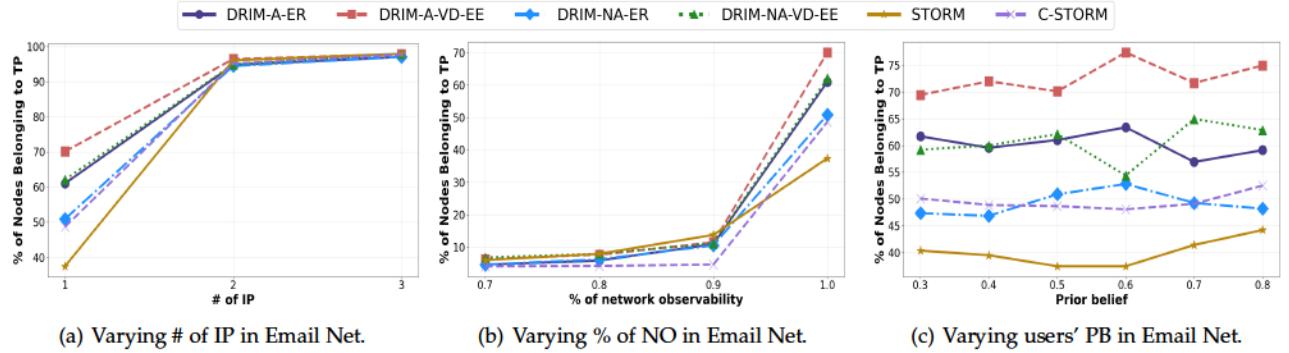


Fig. 4. TP's influence across CIM algorithms when TP and FP use DRL for seed selection in the URV Email network. #IP denotes the number of information propagation by TP, 'NO' represents network observability %, and PB indicates users' prior belief in true information (i.e., a). In all cases except (c), we set $a = 0.5$.

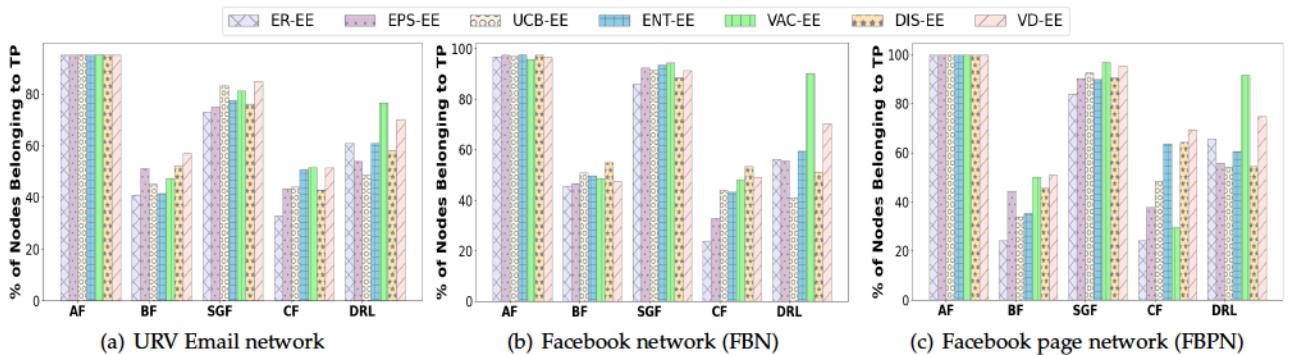


Fig. 5. TP's influence under various EE strategies under the three networks.

Fig. 6(b) reveals a significant boost for all strategies as network observability rises from 0.7 to 1, underscoring their reliance on visibility. VAC-EE and VD-EE improve steeply, excelling in well-observed networks. Fig. 6(c) explores users' prior beliefs. As noted in Section 7.2, high user activity mitigates belief influence. Still, VAC-EE and VD-EE outperform others, benefiting from true information acceptance and low uncertainty. DIS-EE and ENT-EE remain stable, suggesting greater dependence on network structure than user biases.

7.4 Uncertainty Analyses of DRL-based Seed Node Selection Decision-Making

This section evaluates agent performance by analyzing vacuity, dissonance, and entropy estimates at each decision point. For brevity, we present average values across all steps using the URV Email dataset, with analyses of the Facebook

and Facebook Page Network datasets in Appendices B.4 and B.5 of the supplement.

Uncertainty Analyses of DRL-based Seed Node Selection Decision-Making Under Various CIM Schemes: First, we analyze vacuity, dissonance, and entropy across CIM schemes, with results in Fig. 7. When FP employs AF or BE, no significant correlation emerges between CIM scheme performance and uncertainty measures. Under AF, all schemes reach nearly 100% influence regardless of uncertainty. With BF, despite higher vacuity, dissonance, and entropy, STORM and C-STORM outperform others, indicating that confidence in decisions does not guarantee better performance. Effectiveness depends on the opponent's strategy. AF is ineffective in highly active OSNs, while DRIM-based schemes struggle against BF, unlike STORM and C-STORM. However, under SGF, CF, or DRL, a clear relationship appears:

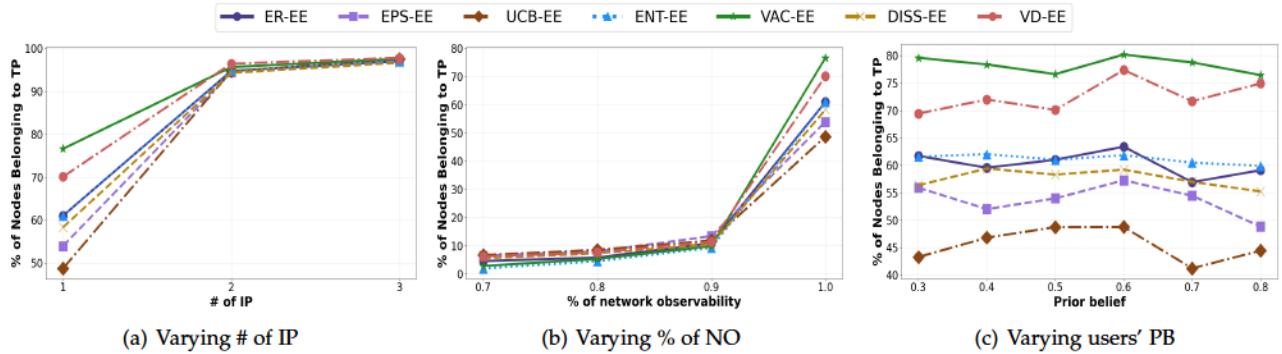


Fig. 6. TP's influence across EE strategies when both TP and FP use DRL for seed selection in the URV Email Network. #IP is TP's information propagation count, 'NO' is network observability (%), and PB is users' prior belief in true information (a). Except (c), we set $\alpha = 0.5$.

lower performance corresponds to higher vacuity. For SGF and CF, STORM and C-STORM exhibit higher vacuity and lower performance, indicating uncertainty in node selection, which hinders effectiveness.

Comparing DRIM-A-ER and DRIM-A-VDEE under DRL, the higher-performing scheme consistently shows lower vacuity and dissonance. A similar trend holds between DRIM-NA-ER and DRIM-NA-VDEE, as well as STORM and C-STORM. Within the same framework, performance improves with greater decision confidence, reinforcing the impact of uncertainty-aware EE algorithms. DRIM-A-VDEE maintains the lowest vacuity, dissonance, and entropy across all FP strategies while achieving superior performance, indicating an optimized balance between exploration and exploitation.

Uncertainty Analyses of DRL-Based Seed Node Selection Under Various EE Strategies: Fig. 8 evaluates uncertainty estimates in DRL-based seed selection across EE strategies. Strategies with lower vacuity generally perform better, as more information enables wiser node selection. For instance, ER-EE exhibits high vacuity and lower influence when FP employs SGF or CF. However, since all EE strategies already maintain low vacuity, a moderate level is not necessarily detrimental. VAC-EE, despite having the highest vacuity, achieves the best performance when FP uses DRL, suggesting that beyond a certain confidence threshold, reducing uncertainty alone does not guarantee better performance. Higher dissonance and entropy often correlate with varied performance. Strategies with high dissonance, such as ER-EE and DIS-EE, face conflicting information, complicating decision-making and influence spread, leading to fluctuating performance.

7.5 Running Time Analyses of CIM Algorithms

Table 5 demonstrates that DRIM-based schemes outperform C-STORM and STORM in terms of running time. We measured the time required to select a node and complete one round of information propagation, averaging these times over 5,000 cycles of node selection and propagation. While larger datasets naturally take longer, the trend within the same dataset remains clear. Specifically, in the URV Email network, DRIM-based strategy, especially DRIM-A-VDEE operates 6.8% faster than C-STORM and 44.5% faster than STORM. This gap widens to 47.8% and 77.1% respectively in the denser Facebook network (FBN), and 15.1% and 61.6% in the larger Facebook Page network

TABLE 5
SIMULATION RUNNING TIME (IN SEC.) OF THE CIM ALGORITHMS

Dataset	DRIM-A-VDEE	DRIM-A-ER	DRIM-NA-ER	C-STORM	STORM
URV [72]	0.352	0.354	0.308	0.380	0.638
FBN [73]	11.248	10.912	10.19	20.890	47.624
FBPN [74]	28.400	24.612	19.546	29.006	64.092

(FBPN). Consequently, our DRIM-based approach proves to be significantly quicker, particularly in denser networks. Moreover, when comparing DRIM-A-VDEE with DRIM-A-ER, the incorporation of an uncertainty-based exploration-exploitation trade-off slightly increases running times, indicating that the efficiency impact of incorporating uncertainty measurement is acceptable.

7.6 Node Selection Analyses of DRL Agent

The DRL agent selects seed nodes based on a policy distribution that evolves during an episode as propagation progresses. We present the Email Network results in the main paper; results for the other two networks are in the supplement document.

Fig. 9 shows the Email Network action distribution. In Fig. 9(a), when FP selects highly active users (AF), TP initially prefers CF to gain broad structural control. Around step 30, TP shifts to AF, likely because high-degree nodes have been selected, prompting TP to disrupt FP's active users. In Fig. 9(b), where FP blocks TP's influence (BF), TP initially mirrors BF to cut off potential FP pathways. As FP's blocking intensifies, TP transitions to CF to regain reach via remaining high-degree nodes, then back to BF to isolate FP clusters. AF and SGF remain rarely selected, as BF better disrupts FP's influence in these scenarios.

In Fig. 9(c), facing FP's SGF strategy, TP starts with CF to secure central nodes, then shifts to BF around step 20 to block clustered influence. The agent alternates between CF and BF, balancing structural control with local blocking. A brief AF spike around step 30 suggests targeted disruption of active nodes. In Fig. 9(d), when FP uses CF, TP quickly adopts BF to block central nodes. A minor AF spike between steps 30–40 indicates brief exploration of active users. CF and SGF are largely absent, reflecting the FP's strong central advantage. Figs. 9(e) and 9(f) show both TP and FP using DRL. TP favors CF throughout, focusing on long-term structural control, with AF gaining prominence after step 15 for strategic disruption or reinforcement. SGF and BF are rarely used, suggesting a proactive strategy.

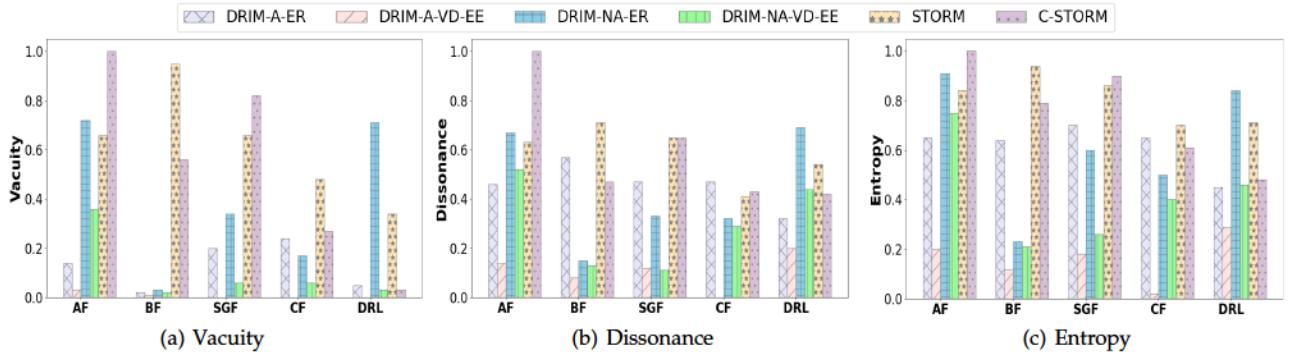


Fig. 7. TP agent's uncertainty measurement under various CIM algorithms in the URV email network.

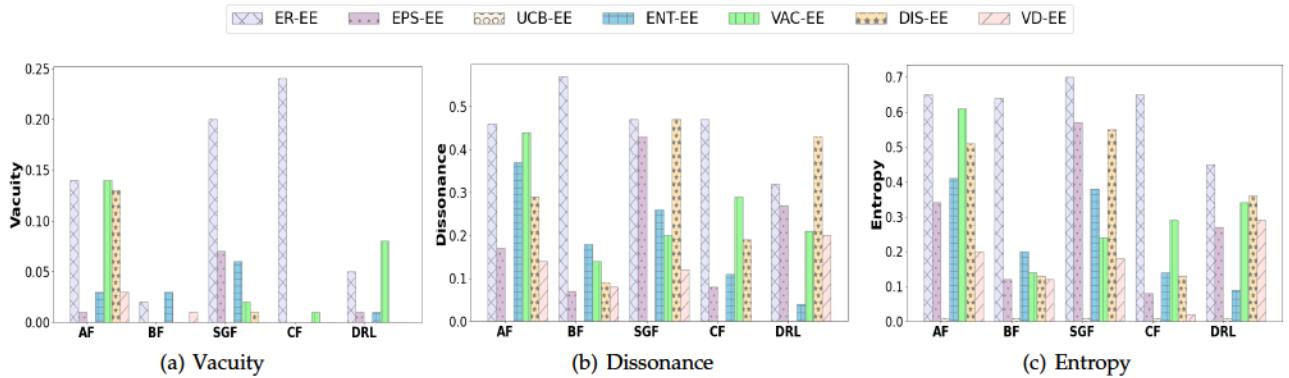


Fig. 8. TP agent's uncertainty measurement under various EE strategies in URV email dataset.

The FP agent initially prefers SGF to exploit dense clusters, then oscillates between SGF and BF after step 25, likely reacting to TP's structural dominance by isolating central regions. AF and CF are rarely used, indicating a shift toward localized tactics.

8 CONCLUSIONS & FUTURE WORK

We highlight the key contributions of our approach as follows. This paper comprehensively analyzes opinion models, CIM schemes, and uncertainty-aware EE strategies for combating misinformation. We introduce DRIM-based CIM schemes that integrate uncertainty-aware decision-making to enhance influence maximization. A comparative evaluation against state-of-the-art methods demonstrates the superiority of DRIM-based schemes in adaptability, efficiency, and robustness across various network conditions.

Our study reveals the following key findings: First, our study highlights the effectiveness of opinion models, with UOM achieving the highest influence, emphasizing the importance of timely and accurate information. High vacuity (u values) increased receptiveness, while HOM and NOM underscored the challenge of countering early false information exposure. Second, DRIM-based schemes demonstrated superiority, consistently outperforming STORM and C-STORM, particularly against SGF, CF, and DRL strategies, showcasing their adaptability and efficiency. Third, we observed the impact of key parameters, where increased propagations improved influence spread, though DRIM-based schemes remained effective even with fewer propagations. Higher network observability significantly boosted performance, while prior beliefs had minimal impact due to fre-

quent information exchanges. Fourth, uncertainty measures correlated with performance, as lower uncertainty consistently led to better outcomes. DRIM-A-VDEE exhibited the lowest uncertainty and highest influence, proving its robustness. Finally, regarding computational efficiency, DRIM-based schemes ran significantly faster than STORM and C-STORM, particularly in dense networks. The exploration-exploitation trade-off in DRIM-A-VDEE had a negligible impact on runtime, ensuring efficiency.

Our current implementation presents a binomial opinion model; this binary assumption provides a practical and computationally feasible approach for our initial model. However, the underlying framework is designed to be extensible. It can be expanded to support multinomial opinions, allowing for a more granular information classification and adapting to more complex classification schemes [75, 76] as needed. We propose the possible extension design in Appendix E.1 for researchers interested in modeling multi-party CIM using SL.

Our future research will address current limitations and extend the DRIM framework in several directions: (1) Develop counter-blocking or defensive strategies to improve performance when FP blocks TP's influence. (2) Incorporate techniques to infer or approximate user behavior and improve resilience under limited network observability (Fig. 6(b)). (3) Design lightweight uncertainty-aware strategies to reduce computational overhead in dynamic environments, addressing VD-EE latency (Fig. 6(c)). (4) Integrate graded truth scales, credibility scores, and other nuanced opinion models to capture real-world belief complexity, and analyze the sensitivity of influence outcomes to users'

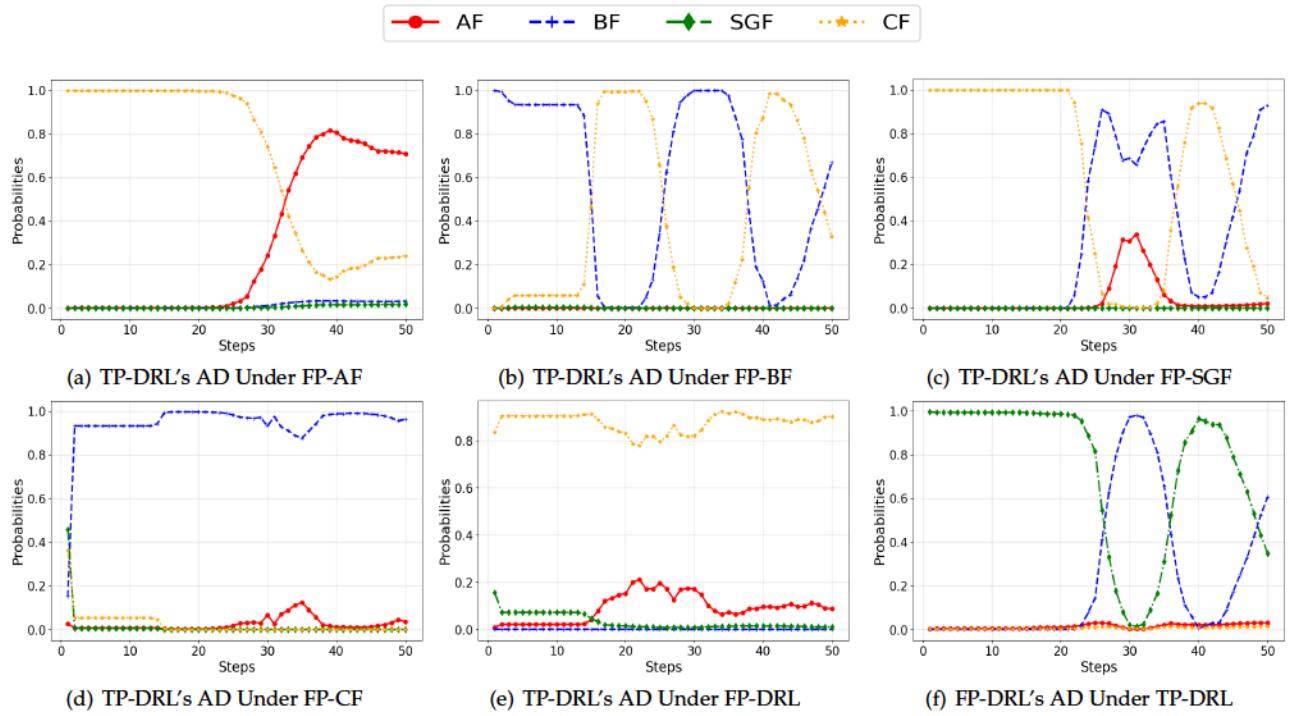


Fig. 9. DRL agent's actions distribution (AD) during a representative episode in the Email Network[72]: TP-DRL denotes the True Party (TP) using deep reinforcement learning (DRL) for seed node selection, while FP-X denotes the False Party (FP) employing strategy X for its seed node selection.

prior beliefs and sharing/reading behaviors, building on insights from prior SL-based models [65, 66]. (5) Scale DRIM to larger, real-world networks to validate effectiveness. (6) Model heterogeneous factors such as dynamic trust, temporal behaviors, and evolving user interactions. (7) Explore multi-agent reinforcement learning to support strategic adaptability in multi-party influence settings. (8) Extend the framework to support real-time deployment by adding a module that periodically updates the network topology to reflect user churn (e.g., account creation or deletion) in dynamic OSNs. (9) Scale the framework to larger and more complex environments by incorporating advanced graph techniques such as graph sampling, hierarchical modeling, and distributed DRL training for efficient influence computation in high-volume networks. (10) Integrate anomaly detection to pre-filter adversarial users and ensure DRL agents operate on trustworthy behavior signals. (11) Extend the framework to support coordinated or multi-agent adversarial behavior by simulating multiple FPs acting in concert. Although our current experiments involve a single FP modeled as a DRL agent, the architecture supports expansion to model coordinated attacks and adaptive adversaries.

ACKNOWLEDGMENT

This work is partly supported by NSF under grants 2107449, 2107450, and 2107451.

REFERENCES

[1] Z. Guo, J. Cho, R. Chen, S. Sengupta, M. Hong, and T. Mitra, "Online social deception and its countermea-

sures: A survey," *IEEE Access*, vol. 9, pp. 1770–1806, 2020.

[2] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.

[3] T. L. Lai and H. Robbins, "Asymptotically efficient adaptive allocation rules," *Advances in Applied Mathematics*, vol. 6, no. 1, pp. 4–22, 1985.

[4] A. Jøsang, *Subjective Logic: A Formalism for Reasoning Under Uncertainty*. Springer, 2016.

[5] K. Ali, C. Wang, M. Yeh, and Y. Chen, "Addressing competitive influence maximization on unknown social network with deep reinforcement learning," in *2020 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 2020, pp. 196–203.

[6] K. Ali, C. Wang, M. Yeh, C. Li, and Y. Chen, "Nedrl-cim: Network embedding meets deep reinforcement learning to tackle competitive influence maximization on evolving social networks," in *2021 IEEE 8th International Conference on Data Science and Advanced Analytics (DSAA)*, 2021, pp. 1–9.

[7] T. Chung, "Deep reinforcement learning-based approach to tackle competitive influence maximization," in *Proc. MLG Workshop*, 2019.

[8] S. Lin, S. Lin, and M. Chen, "A learning-based framework to handle multi-round multi-party influence maximization on social networks," in *Proc. the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, New York, NY, USA, 2015, p. 695–704.

[9] J. Li, T. Cai, K. Deng, X. Wang, T. Sellis, and F. Xia, "Community-diversified influence maximization in social networks," *Information Systems*, vol. 92, p. 101522, 2020.

[10] M. Sensoy, L. Kaplan, and M. Kandemir, "Evidential deep learning to quantify classification uncertainty," *Advances in Neural Information Processing Systems*, vol. 31, 2018.

[11] X. Zhao, F. Chen, S. Hu, and J.-H. Cho, "Uncertainty aware semi-supervised learning on graph data," *Advances in Neural Information Processing Systems*, vol. 33, pp. 12 827–12 836, 2020.

[12] Q. Zhang, L. M. Kaplan, A. Jøsang, D. H. Jeong, F. Chen, and J.-H. Cho, "Uncertainty-aware influence maximization: Enhancing propagation in competitive social networks with subjective logic," in *2024 IEEE International Conference on Big Data (BigData)*, 2024, pp. 2065–2071.

[13] P. Domingos and M. Richardson, "Mining the network value of customers," in *Proc. the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2001, pp. 57–66.

[14] D. Kempe, J. Kleinberg, and É. Tardos, "Maximizing the spread of influence through a social network," in *Proc. 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2003, pp. 137–146.

[15] H. Li, S. S. Bhowmick, J. Cui, Y. Gao, and J. Ma, "Ge-treal: Towards realistic selection of influence maximization strategies in competitive networks," in *Proceedings of the 2015 ACM SIGMOD international conference on management of data*, 2015, pp. 1525–1537.

[16] C. Pham et al., "Competitive influence maximization within time and budget constraints in online social networks: An algorithmic approach," *Applied Sciences*, vol. 9, no. 11, 2019.

[17] W. Liu, L. Chen, X. Chen, and B. Chen, "An algorithm for influence maximization in competitive social networks with unwanted users," *Applied Intelligence*, vol. 50, no. 2, pp. 417–437, 2020.

[18] W. Hong, C. Qian, and K. Tang, "Efficient minimum cost seed selection with theoretical guarantees for competitive influence maximization," *IEEE transactions on cybernetics*, vol. 51, no. 12, pp. 6091–6104, 2020.

[19] D. Tsaras, G. Trimponias, L. Ntaflos, and D. Papadias, "Collective influence maximization for multiple competing products with an awareness-to-influence model," *Proceedings of the VLDB Endowment*, vol. 14, no. 7, pp. 1124–1136, 2021.

[20] F. Galante, M. Garetto, and E. Leonardi, "Competition of influencers: A model for maximizing online social impact," in *Proceedings of the 16th ACM Web Science Conference*, 2024, pp. 343–353.

[21] A. Bozorgi, S. Samet, J. Kwisthout, and T. Wareham, "Community-based influence maximization in social networks under a competitive linear threshold model," *Knowledge-Based Systems*, vol. 134, pp. 149–158, 2017.

[22] X. Xie, J. Li, Y. Sheng, W. Wang, and W. Yang, "Competitive influence maximization considering inactive nodes and community homophily," *Knowledge-Based Systems*, vol. 233, p. 107497, 2021.

[23] J. Liu, Z. You, Z. Liang, and H. Du, "User-driven competitive influence maximization in social networks," *Theoretical Computer Science*, vol. 1018, p. 114813, 2024.

[24] E. Bagheri and R. S. Mirtalaei, "Community-based influence maximization in social networks under a competitive linear threshold model considering positive and negative user views," *International Journal of Modern Physics C*, vol. 35, no. 01, p. 2450003, 2024.

[25] Q. He, L. Zhang, H. Fang, X. Wang, L. Ma, K. Yu, and J. Zhang, "Multistage competitive opinion maximization with Q-learning-based method in social networks," *IEEE Transactions on Neural Networks and Learning Systems*, 2024.

[26] H. Li, L. Pan, and P. Wu, "Dominated competitive influence maximization with time-critical and time-delayed diffusion in social networks," *Journal of Computational Science*, vol. 28, pp. 318–327, 2018.

[27] J. Tong, L. Shi, L. Liu, J. Panneerselvam, and Z. Han, "A novel influence maximization algorithm for a competitive environment based on social media data analytics," *Big Data Mining and Analytics*, vol. 5, no. 2, pp. 130–139, 2022.

[28] Z. Liang, Q. He, H. Du, and W. Xu, "Targeted influence maximization in competitive social networks," *Information Sciences*, vol. 619, pp. 390–405, 2023.

[29] J. Dai, J. Zhu, and G. Wang, "Competitive net influence maximization on intergroup debate effect," *Information Sciences*, vol. 680, p. 121139, 2024.

[30] Y. Lin and J. Lui, "Analyzing competitive influence maximization problems with partial information: An approximation algorithmic framework," *Performance Evaluation*, vol. 91, pp. 187–204, 2015.

[31] Z. Liang, J. Liu, H. Du, and C. Zhang, "Mecim: Multi-entity evolutionary competitive influence maximization in social networks," *Expert Systems with Applications*, vol. 270, p. 126436, 2025.

[32] S. Gao, Z. Zhang, S. Su, J. Wen, and L. Sun, "Fair-aware competitive event influence maximization in social networks," *IEEE Transactions on Network Science and Engineering*, vol. 7, no. 4, pp. 2528–2540, 2020.

[33] C. Zhang, J. Zhou, J. Wang, J. Fan, and Y. Shi, "Fairness-aware competitive bidding influence maximization in social networks," *IEEE Transactions on Computational Social Systems*, vol. 11, no. 2, pp. 2147–2159, 2023.

[34] K. Ali, C.-Y. Wang, and Y.-S. Chen, "Leveraging transfer learning in reinforcement learning to tackle competitive influence maximization," *Knowledge and Information Systems*, vol. 64, no. 8, pp. 2059–2090, 2022.

[35] T. M. Liggett, *Stochastic interacting systems: contact, voter and exclusion processes*. Springer Science & Business Media, 2013, vol. 324.

[36] M. H. DeGroot, "Reaching a consensus," *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974.

[37] N. E. Friedkin and E. C. Johnsen, "Social influence networks and opinion change," *Advances in group processes*, vol. 16, no. 1, pp. 1–29, 1999.

[38] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, "Mixing beliefs among interacting agents," *Advances in Complex Systems*, vol. 3, no. 01n04, pp. 87–98, 2000.

[39] R. Hegselmann and U. Krause, "Opinion dynamics and bounded confidence models, analysis, and simulation," *Journal of Artificial Societies and Social Simulation (JASSS)* vol. vol. 5, no. 3, 2002.

[40] S. Galam, "Sociophysics: A review of galam models," *International Journal of Modern Physics C*, vol. 19, no. 03,

pp. 409–440, 2008.

[41] K. Sznajd-Weron and J. Sznajd, "Opinion evolution in closed community," *International Journal of Modern Physics C*, vol. 11, no. 06, pp. 1157–1165, 2000.

[42] S. Grauwin and P. Jensen, "Opinion group formation and dynamics: Structures that last from nonlasting entities," *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, vol. 85, no. 6, p. 066113, 2012.

[43] W. Su, G. Chen, and Y. Hong, "Noise leads to quasi-consensus of hegselmann-krause opinion dynamics," *Automatica*, vol. 85, pp. 448–454, 2017.

[44] R. Devidze, P. Kamalaruban, and A. Singla, "Exploration-guided reward shaping for reinforcement learning under sparse rewards," *Advances in Neural Information Processing Systems*, vol. 35, pp. 5829–5842, 2022.

[45] B. D. Ziebart, *Modeling purposeful adaptive behavior with the principle of maximum causal entropy*. Carnegie Mellon University, 2010.

[46] Y. Burda, H. Edwards, A. Storkey, and O. Klimov, "Exploration by random network distillation," *arXiv preprint arXiv:1810.12894*, 2018.

[47] M. Bellemare, S. Srinivasan, G. Ostrovski, T. Schaul, D. Saxton, and R. Munos, "Unifying count-based exploration and intrinsic motivation," *Advances in neural information processing systems*, vol. 29, 2016.

[48] A. P. Badia, P. Sprechmann, A. Vitvitskyi, D. Guo, B. Piot, S. Kapturowski, O. Tieleman *et al.*, "Never give up: Learning directed exploration strategies," *arXiv preprint arXiv:2002.06038*, 2020.

[49] A. Ecoffet, J. Huizinga, J. Lehman, K. O. Stanley, and J. Clune, "Go-explore: a new approach for hard-exploration problems," *arXiv preprint arXiv:1901.10995*, 2019.

[50] C. Florensa, D. Held, M. Wulfmeier, M. Zhang, and P. Abbeel, "Reverse curriculum generation for reinforcement learning," in *Conference on Robot Learning*. PMLR, 2017, pp. 482–495.

[51] I. Osband, C. Blundell, A. Pritzel, and V. B. Roy, "Deep exploration via bootstrapped DQN," *Advances in Neural Information Processing Systems*, vol. 29, 2016.

[52] G. Kalweit and J. Boedecker, "Uncertainty-driven imagination for continuous deep reinforcement learning," in *Conference on Robot Learning*. PMLR, 2017, pp. 195–206.

[53] R. Houthooft, X. Chen, Y. Duan, J. Schulman, F. De Turck, and P. Abbeel, "VIME: Variational information maximizing exploration," in *Advances in Neural Information Processing Systems*, vol. 29, 2016.

[54] Y. Hu, Y. Ou, X. Zhao, J.-H. Cho, and F. Chen, "Multidimensional uncertainty-aware evidential neural networks," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 9, 2021, pp. 7815–7822.

[55] C. Wang, X. Wang, J. Zhang, L. Zhang, X. Bai, X. Ning, J. Zhou, and E. Hancock, "Uncertainty estimation for stereo matching based on evidential deep learning," *Pattern Recognition*, vol. 124, p. 108498, 2022.

[56] D. Deng, G. Chen, Y. Yu, F. Liu, and P.-A. Heng, "Uncertainty estimation by fisher information-based evidential deep learning," in *International Conference on Machine Learning*. PMLR, 2023, pp. 7596–7616.

[57] Z. Shao, W. Dou, and Y. Pan, "Dual-level deep evidential fusion: Integrating multimodal information for enhanced reliable decision-making in deep learning," *Information Fusion*, vol. 103, p. 102113, 2024.

[58] S. Ancha, P. R. Osteen, and N. Roy, "Deep evidential uncertainty estimation for semantic segmentation under out-of-distribution obstacles," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 6943–6951.

[59] R. Duan, B. Caffo, H. X. Bai, H. I. Sair, and C. Jones, "Evidential uncertainty quantification: A variance-based perspective," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2024, pp. 2132–2141.

[60] H. Li, Y. Nan, J. Del Ser, and G. Yang, "Region-based evidential deep learning to quantify uncertainty and improve robustness of brain tumor segmentation," *Neural Computing and Applications*, vol. 35, no. 30, pp. 22 071–22 085, 2023.

[61] J. S. Schreck, D. J. Gagne, C. Becker, W. E. Chapman, K. Elmore, D. Fan, G. Gantos, E. Kim, D. Kimpara, T. Martin *et al.*, "Evidential deep learning: Enhancing predictive uncertainty estimation for earth system science applications," *Artificial Intelligence for the Earth Systems*, vol. 3, no. 4, p. 230093, 2024.

[62] A. Jøsang, J. Cho, and F. Chen, "Uncertainty characteristics of subjective opinions," in *21st FUSION*, 2018, pp. 1998–2005.

[63] X. Li and L. Wu, "Herding and social media word-of-mouth," *Mis Quarterly*, vol. 42, no. 4, pp. 1331–1352, 2018.

[64] J. Cho, "Dynamics of uncertain and conflicting opinions in social networks," *IEEE transactions on computational social systems*, vol. 5, no. 2, pp. 518–531, 2018.

[65] J. Cho and S. Adali, "Is uncertainty always bad?: Effect of topic competence on uncertain opinions," in *2018 IEEE International Conference on Communications (ICC)*. IEEE, 2018, pp. 1–7.

[66] J. Cho, S. Rager, J. O'Donovan, S. Adali, and B. Horne, "Uncertainty-based false information propagation in social networks," *ACM Transactions on Social Computing*, vol. 2, no. 2, pp. 1–34, 2019.

[67] L. Li, A. Scaglione, A. Swami, and Q. Zhao, "Consensus, polarization and clustering of opinions in social networks," *IEEE Journal on Selected Areas in Communications*, vol. 31, no. 6, pp. 1072–1083, 2013.

[68] V. Karnowski, L. Leonhard, and A. Kümpel, "Why users share the news: A theory of reasoned action-based study on the antecedents of news-sharing behavior," *Communication Research Reports*, vol. 35, no. 2, pp. 91–100, 2018.

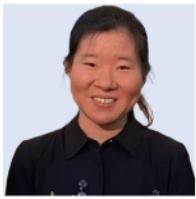
[69] A. Namdari and Z. Li, "A review of entropy measures for uncertainty quantification of stochastic processes," *Advances in Mechanical Engineering*, vol. 11, no. 6, p. 1687814019857350, 2019.

[70] M. Nasim, R. Charbey, C. Prieur, and U. Brandes, "Investigating link inference in partially observable networks: Friendship ties and interaction," *IEEE Transactions on Computational Social Systems*, vol. 3, no. 3, pp. 113–119, 2016.

[71] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms,"

arXiv preprint arXiv:1707.06347, 2017.

- [72] R. Rossi and N. Ahmed, "The network data repository with interactive graph analytics and visualization," in *Proc. the AAAI Conference on Artificial Intelligence*, vol. 29, no. 1, 2015.
- [73] J. Leskovec and J. Mcauley, "Learning to discover social circles in ego networks," *Advances in Neural Information Processing Systems*, vol. 25, 2012.
- [74] B. Rozemberczki, C. Allen, and R. Sarkar, "Multi-scale attributed node embedding," *Journal of Complex Networks*, vol. 9, no. 2, p. cnab014, 2021.
- [75] E. Shushkevich, J. Cardiff, and A. Boldyreva, "Detection of truthful, semi-truthful, false and other news with arbitrary topics using bert-based models," in *2023 33rd Conference of Open Innovations Association (FRUCT)*. IEEE, 2023, pp. 250–256.
- [76] C. Li, S. A. Chun, and J. Geller, "Enhanced multi-class detection of fake news," in *The International FLAIRS Conference Proceedings*, vol. 37, 2024.



Qi Zhang (S'24) received the B.S. degree in Electronic Science and Technology from Beijing Institute of Technology, Beijing, China, in 2019, and the M.S. degree in Computer Science from New York University, Brooklyn, NY, in 2021. She is currently pursuing a Ph.D. in Computer Science at Virginia Tech, Arlington, VA, USA. Her research interests include network science, social cybersecurity, and decision-making under uncertainty.



Dian Chen is currently a Ph.D. student in the Department of Computer Science at Virginia Tech since 2022. She received a B.S. degree in Computer Science from Worcester Polytechnic Institute in 2020 and an M.S. degree in Computer Science from Northwestern University in 2022. Her research interests include network security in cyber-physical systems.



Lance M. Kaplan received the B.S. degree with distinction from Duke University in 1989 and the M.S. and Ph.D. degrees from the University of Southern California in 1991 and 1994, respectively, all in Electrical Engineering. From 1987–1990, he was a Technical Assistant at the Georgia Tech Research Institute. He held NSF and USC Dean's Merit Fellowships from 1990–1993 and was a Research Assistant at USC's Signal and Image Processing Institute from 1993–1994. He then worked in the Reconnaissance Systems Department at Hughes Aircraft Company (1994–1996). From 1996–2004, he was a faculty member in Engineering and Senior Investigator at the Center of Theoretical Studies of Physical Systems, Clark Atlanta University (CAU), where he received the Electrical Engineering Instructional Excellence Award (1999–2001). He is currently a team leader in the Context Aware Processing branch at the US Army DEVCOM Army Research Laboratory (ARL). Dr. Kaplan serves as VP Conferences for the International Society of Information Fusion (ISIF, 2014–Present). He previously served as VP Publications for the IEEE AES Society (2021–2023), Editor-in-Chief of IEEE Trans. AES (2012–2017), and on the Board of Governors for IEEE AES (2008–2013, 2018–2020) and ISIF (2012–2014). He is a Fellow of IEEE, ARL, and the Military Sensing Symposia. His research interests include information fusion, uncertainty reasoning, network science, resource management, and signal/image processing.



Audun Jøsang is professor of cybersecurity at the University of Oslo, as well as adjunct professor at QUT in Australia. Prof. Jøsang is widely known for his research on trust and reputation systems. One of Jøsang's key contributions is the development of Subjective Logic for reasoning under uncertainty, which has applications in areas like risk assessment, decision-making, artificial intelligence and the analysis of trust in social networks.



Dong Hyun Jeong is a Full Professor and Chair of Computer Science and Information Technology at the University of the District of Columbia. He received his Ph.D. in Computer Science from the University of North Carolina at Charlotte in 2010, specializing in information visualization and visual analytics. His research focuses on interactive visual representations and analytical techniques to support decision-making across diverse data types, emphasizing the integration of machine and deep learning. His work has appeared in leading journals and conferences, and he actively serves on editorial boards and committees in the global research community. Dr. Jeong is a Senior Member of IEEE, a member of ACM, a lifetime member of Phi Kappa Phi, and a full member of Sigma Xi.



Feng Chen is an Associate Professor of Computer Science at the University of Texas at Dallas. He received his Ph.D. in Computer Science from Virginia Tech in 2012. His research covers large-scale data and network mining, machine learning, and applications in anomaly detection, uncertainty quantification, and fairness-aware learning. His recent work includes hidden activity signal analysis, graph convolutional networks for anomaly detection, fairness-aware stream learning, uncertainty-aware deep learning, and large-scale software vulnerability analysis. Dr. Chen has published over 150 peer-reviewed papers and received funding from NSF, NIH, ARO, IARPA, and the U.S. Department of Transportation. He is a recipient of the NSF CAREER Award (2018).



Jin-Hee Cho (M'09; SM'14) is an Associate Professor of Computer Science at Virginia Tech and director of the Trustworthy Cyberspace Lab (tclab). She was a computer scientist at the U.S. Army Research Laboratory (USARL) from 2009 to 2018. Her research spans cybersecurity, decision-making under uncertainty, and network science, with publications in top venues and multiple best paper awards (IEEE TrustCom'2009, BRIMS'2013, GLOBECOM'2017, CogSima 2018). She received the IEEE ComSoc William R. Bennett Prize (2015), the CNOM Test of Time Award (2023), and the PECASE Award (2013). At Virginia Tech, she received the 2022 Faculty Fellow Award and the Tyndall Moore Junior Faculty Fellowship. Dr. Cho earned her Ph.D. in Computer Science from Virginia Tech in 2008 and serves as Associate Editor for IEEE TNSM and IEEE TSC. She is a Senior Member of IEEE and a member of ACM.