GhostField: Simulating an Integrated Information Training Environment at Scale

Jonathan H. Morgan[1], Jake Shaha[2], Rebecca Marigliano[2], Matthew S. Hicks[2], and Kathleen M. Carley[1,2]

[1] Netanomics, Pittsburgh PA 15217, USA,
jonathan.morgan@netanomics.com,
[2] Carnegie Mellon University, Pittsburgh PA 15217, USA,

Abstract. This paper introduces GhostField, a synthetic data generation system for network science training in social media analysis. GhostField approximates cognitive constraints—such as attention, priming, and schema activation—to produce realistic communication and interaction patterns using large language models. It combines narrative design (AESOP), signal detection (NetMapper), and content evaluation (MOMUS) to generate socially structured environments that support reasoning about coordination and influence. Analysts use ORA to examine these networks, applying metrics and visualizations to detect strategic behaviors. We evaluate first- and second-generation network outputs, demonstrating improved modularity, narrative coherence, and structural realism. These features allow analysts to distinguish authentic behavior from artificial manipulation. Future development will emphasize structured defaults and population-level modeling to reduce scenario design overhead, expand narrative fluidity, and support scalable, mission-relevant training.

Keywords: network science, training, social cyber security

1 Introduction

A major strength of network science lies in its ability to reveal the structure of complex systems. By mapping patterns of interaction, networks expose features such as hierarchy, cohesion, and influence, insights that support both interpretation and, in some cases, forecasting. For decades, educators have relied on small-scale case studies to teach these principles. Datasets such as Padgett and Ansell's Florentine Families [1], Sampson's Monastic Factions [2,3], and Zachary's Karate Club [4] illustrate how structure emerges from context, helping students connect abstract metrics to meaningful social processes.

Scenario-based simulations extend this approach by embedding multiple overlapping dynamics in synthetic environments where the ground truth is known. The GhostField system supports this instructional model by allowing learners to detect meaningful patterns, interpret them in light of narrative and organizational context, and reason about how interventions might reshape the network.

However, to support this level of engagement, synthetic environments must go beyond surface-level realism. They must reproduce the structure, dynamics, and narrative coherence of real-world systems in ways that make analytical reasoning possible. By formalizing the cognitive and structural constraints required for meaningful analyst training, GhostField offers a blueprint for designing synthetic information environments that support immersive, interpretable, and mission-relevant learning.

In what follows, we introduce GhostField as a platform for synthetic data generation and network science training. We argue that generating plausible social environments with LLMs - networks that exhibit realistic modularity, hierarchy, and coordination - requires approximating key cognitive constraints such as attention, priming, and schema activation. Although these constraints are not modeled directly, GhostField emulates their behavioral effects through structured prompts and scenario design via AESOP, symbolic pattern detection through NetMapper, and content and structural evaluation with MOMUS. Together, these components ensure that linguistic and symbolic signals (e.g., hashtags, emojis, key phrases) align with group and role identities and remain consistent with the overarching scenario. Without such approximations, large language models produce implausible messaging and interaction patterns that cannot support higher-level analysis. The resulting socially structured environments provide the scaffolding for rational band reasoning, where analysts assess goals, evaluate messaging strategies, and anticipate agent actions. GhostField is therefore more than a content generation tool. It provides a structured framework for the construction of interpretable and behaviorally plausible networks, networks that support learning in the rational band by first emulating the cognitive constraints that lead to coherent group structures in the social band [5].

2 The GhostField Framework

GhostField is a robust framework for simulating an information environment at scale. It combines AI-enabled network science with empirical validation to help planners build complex training scenarios quickly. The AESOP scenario design tool provides comprehensive design control, enabling planners to structure scenarios with the specific features needed to meet well-defined learning objectives. It supports the creation of coherent synthetic environments through consistent agent behavior, synthetic personae, and high-volume media and social media content (e.g., 9,000 agents, 12,000 posts, and 10,000 links). The system also enables automated adjudication, supports human-in-the-loop evaluation, and allows trainees to inject messages and agents into the simulated information environment.

At the core of GhostField's analytic capability is the Artemis toolkit, a platform-agnostic system that integrates text analytics, machine learning, and network forecasting. Although not exclusive to GhostField, Artemis supports key elements of its pipeline, enabling consistent analysis across synthetic, real, and hybrid information environments. Its modular architecture includes NetMapper [6], which applies computational linguistic techniques to identify emotional cues, frames of reference, syntactic markers, and other textual features. BotHunter [7] and related classifiers further enrich the data by detecting bot-like behavior, hate speech, and geolocation. These annotations are passed to ORA-PRO [8], where the networks are extracted or inferred, and node-level attributes are assigned.

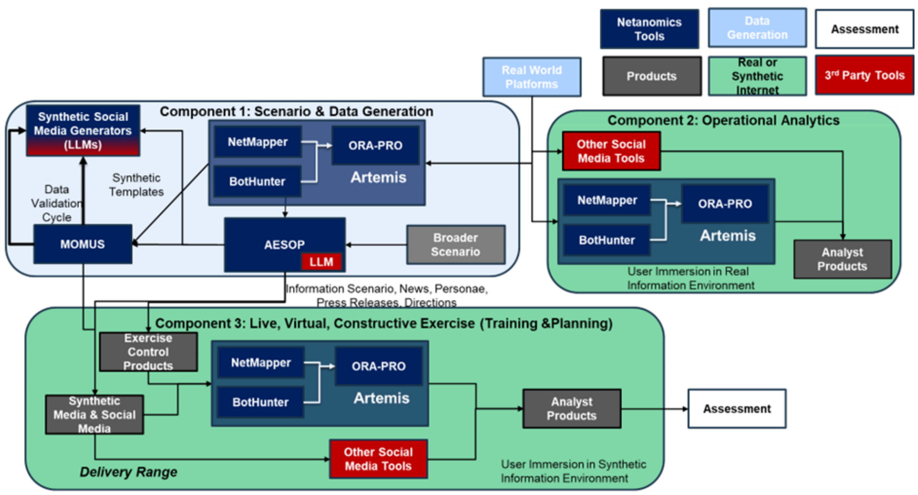

Fig. 1: The GhostField Framework.

Component 1: Scenario and Data Generation forms the foundation of GhostField's training support. Synthetic actors are guided by structured prompts that define attributes such as nationality, organizational affiliation, and position on issues. These prompts shape social media interactions by driving behaviors like hashtag use, mentions, quotes, and replies. SynTelX employs algorithms to construct the network structure and uses the prompts to generate realistic, context-sensitive message content via large language models (LLMs). It then renders this content into platform-specific formats with metadata such as timestamps, user handles, and engagement types (e.g., retweets or quote tweets). To ensure compatibility with external tools, synthetic data are formatted to match the API standards of the modeled platforms, such as X and Telegram, enabling seamless integration into existing analytical pipelines.

NetMapper analyzes LLM-generated messages for social-psychological signals such as emotion, moral content, and institutional references. MOMUS then evaluates the annotated output against more than 200 structural, semantic, and behavioral indicators. Although LLMs can easily generate content that appears human-like, their outputs often diverge from actual human behavior, for instance, humans use emotion-bearing language more frequently than most LLMs do. MOMUS helps verify how closely synthetic outputs align with empirical patterns, highlighting gaps that may not be obvious from surface plausibility alone. Although initial content generation takes only hours, building a fully scalable and instructionally effective scenario –complete with actor prompts, community structures, and media streams –typically requires iterative refinement using MOMUS to ensure messaging fidelity and network plausibility. Still, what once took many months to a year can now be accomplished in weeks.

Component 2: Operational Analysis applies Artemis to both GhostField-generated and externally sourced data. For training applications, this enables the continuity in tools and methods between synthetic immersive training environments and operational settings. Analysts use Artemis to identify key actors, map relationships, and track narrative diffusion. Compatibility with Artemis thus ensures that skills and insights gained in simulation environments translate meaningfully to real-world analytic tasks.

Component 3: Live, Virtual, Constructive (LVC) Execution delivers GhostField scenarios into immersive training environments. In this phase, the synthetic content generated in Component 1 is injected into LVC platforms, supporting analyst workflows, scenario adjudication, and post-exercise evaluation. Artemis processes the injected media in real time, applying the same analytic pipeline used in operational contexts. This ensures that analysts are trained in environments that mirror the data structures, narrative patterns, and coordination challenges found in real-world information operations.

3 Immersive Training Exercises

GhostField has supported a series of training exercises in information operations, public affairs, psychological operations, and PAI analysis [9–12]. These exercises follow a two-phase format: four to five days of structured instruction in analytic workflows using scenario-specific training data, followed by five to seven days of immersive application in a live virtual-constructive (LVC) setting. Participants respond to scenario injects as if supporting a real-world operation.

During the instructional phase, learners are introduced to the foundational network and influence analytics. They learn to identify mission-relevant groups, isolate key actors, assess structural influence, and detect coordinated activity such as bot networks that boost cohesion and amplify strategic messaging. Participants also apply the BEND framework [13] to describe messaging campaigns using sixteen constructive and deconstructive maneuvers, spanning both the narrative and structural domains. Unlike broader models (e.g., Generate, Preserve, Deny, Project), BEND specifies actionable levers of influence. GhostField remains the only system that operationalizes this framework.

These techniques are supported in all three components of GhostField. Component 2 (Operational Analytics) uses Artemis, NetMapper, and ORA-PRO to ingest and analyze data, identifying bridging, echo chambers, and narrative dynamics. Component 1 (Scenario and Data Generation) was introduced eighteen months ago and has since supported three major exercises. It uses AESOP to define scenario parameters, SynTelX to simulate platform activity, and MOMUS to evaluate content realism. Together, they generate high-volume, context-rich data sets that include text, metadata, hashtags, mentions, and links.

Component 3 (LVC Integration), deployed in the two most recent exercises, closes the loop. Participant actions now feed back into the scenario pipeline, prompting new output from AESOP and SynTelX. Injects evolve in simulated 24-hour cycles, transforming static play into dynamic and responsive simulation. This closed-loop design allows participants to test influence maneuvers in a live environment, reinforcing concepts through realistic feedback and iterative practice.

4 Key Findings and Lessons Learned

Achieving data with sufficient realism to focus participants on the scenario - rather than distract them with its artificiality - required a series of innovations. Large Language Models (LLMs), when used with minimal scaffolding, often produce hallucinations, expose prompts, and exhibit a positivity bias due to their optimization for user satisfaction. They also struggle to produce human-like interaction networks because they lack the constraints of embodiment: bounded rationality, limited attention, and resource scarcity in particular. As shown in Figure 2, our first-generation synthetic networks did not reproduce key structural features of the real social media data. In the account-level network (left panels), ties represent retweets, quotes, and replies, overlaid with mentions (pink edges). The message-level network (right panels) connects posts via shared propagation behavior.

The lower panels depict synthetic networks with a hub-and-spoke topology anchored on media accounts and a military periphery. Despite including retweets and replies, the dominant source of ties are co-mentions by three high-degree nodes, yielding an overly dense, undifferentiated structure. In contrast, real-world networks show a community structure: clusters of actors linked by affiliation, language, or role, connected by bridging nodes. These social formations contribute to the reinforcement of norms, group identity, and myside bias. Message-level networks also diverged: synthetic outputs lacked reply chains and shared URLs, eliminating the cross-post linkages necessary for narrative flow. Without these structures, propagation algorithms fail to capture how content spreads, and narratives remain confined within isolated communities.

The upper panels show real data with the expected structural complexity of operational environments. The account-level network contains modular communities, cross-group links, and interactions aligned with organizational proximity. The message-level network shows narrative clusters linked by frames and cues. Although hashtags are omitted, modularity is evident in propagation patterns alone. These features reflect underlying cognitive constraints. Their absence in earlier synthetic output led to implausible groupings and insufficient bridging. Training tools must emulate these constraints across Newell's cognitive bands [5] to generate realistic social environments that support critical thinking.

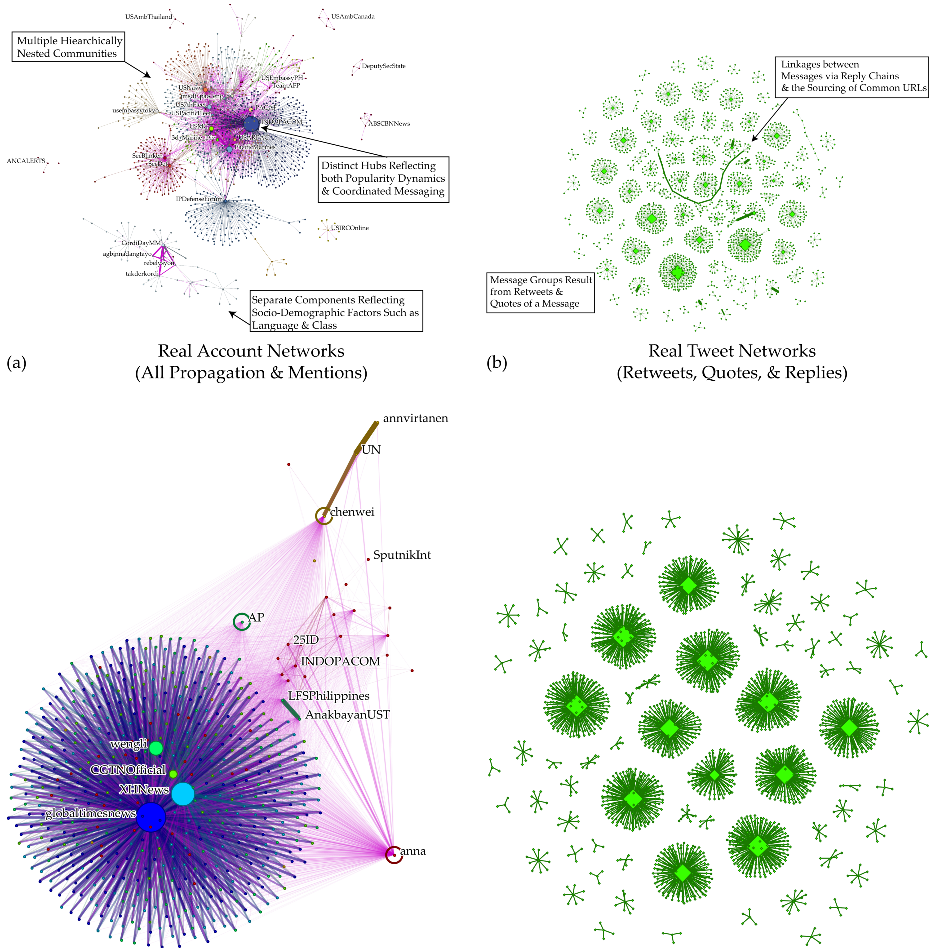

(a) Real Account Networks (All Propagation & Mentions)

(b) Real Tweet Networks (Retweets, Quotes, & Replies)

(c) First Generation Synthetic Account Networks

(d) First Generation Synthetic Tweet Networks

Fig. 2: Real versus Synthetic Networks, Scenario-Based Interactions.

GhostField approximates these constraints through behavioral emulation. The cognitive band - which includes attention, memory, perception, and language - shapes how agents process and act on information. GhostField models these processes not directly, but via design mechanisms that produce realistic interaction structures. Three approximations are especially important: **attention and selective processing**, **priming and cue activation**, and **chunking and schema activation**. Together, these approximations support the semantic coherence and group structure necessary for meaningful network analysis in the rational band.

Selective attention shapes user engagement under constraint. Retweet counts, trending topics, and follower counts act as heuristics for prioritization, producing hub-and-spoke structures. GhostField models these dynamics through preferential attachment, time-weighted recency, and group-linked information exposure, scaffolding the emergence of topic-oriented groups without scripting every connection. Attention constraints also influence tie formation. For example, reposts are the most common tie on X, as they allow users to endorse content or signal affiliation with minimal effort. Original posts remain frequent because they allow users to frame issues in their own terms and initiate discourse, offering the rare possibility of achieving narrative dominance. Replies and quotes are less common, require more cognitive effort than reposting, and offer less chances of going viral than original posts. When used, they typically reflect stronger user investment, whether to elaborate, affirm, or contest a position. These asymmetries arise not only from platform design but also from broader patterns of bounded cognition and anticipated utility. GhostField emulates these regularities through group-linked decision logic and exposure filtering, producing realistic distributions of interaction types.

Priming and cue activation use compact signals—hashtags, slogans, and emojis—to activate emotional and semantic frames. These cues foster convergence and shared narrative structure, which in turn create analyzable co-occurrence patterns. In GhostField, hashtags and URLs are linked to roles and communities via AESOP, allowing coherent factional behavior and some contestation, though adversarial cue use remains limited.

Schemas, preexisting frames that guide interpretation, are modeled through account archetypes and prompts that encode identity, role, and communicative style. Narrative consistency is enforced through scenario-linked constraints: A post about a drone strike must match the time, location, and group involved. High-profile agents function as narrative anchors, while peripheral accounts amplify their messages. Together, NetMapper and MOMUS are used to assess cue coherence, emotional tone, and symbolic alignment against empirical baselines. This scaffolding helps learners infer the intent of the actor and the narrative strategy. However, group-based coherence can also limit variation and suppress cross-group dynamics. Future iterations will use population modeling and more adaptive rules to simulate cross-cutting affiliations and persuasive reach.

GhostField emulates the behavioral conditions necessary for reasoning in the rational band. Strategic behaviors, such as bot deployment, message targeting, and narrative maneuvering, are approximated through templates, scenario logic, and structured prompting rather than emergent autonomy. The goal is not to simulate cognition, but to construct environments where strategic coordination is both plausible and analytically tractable. By aligning behavior with known cognitive constraints, attention, priming, and schema coherence, the system enables learners to recognize coordination patterns, infer intent, and evaluate influence maneuvers. The BEND framework guides this reasoning by linking observable signals to strategic maneuver categories, allowing analysts to both interpret coordination and assess the effectiveness of messaging strategies.

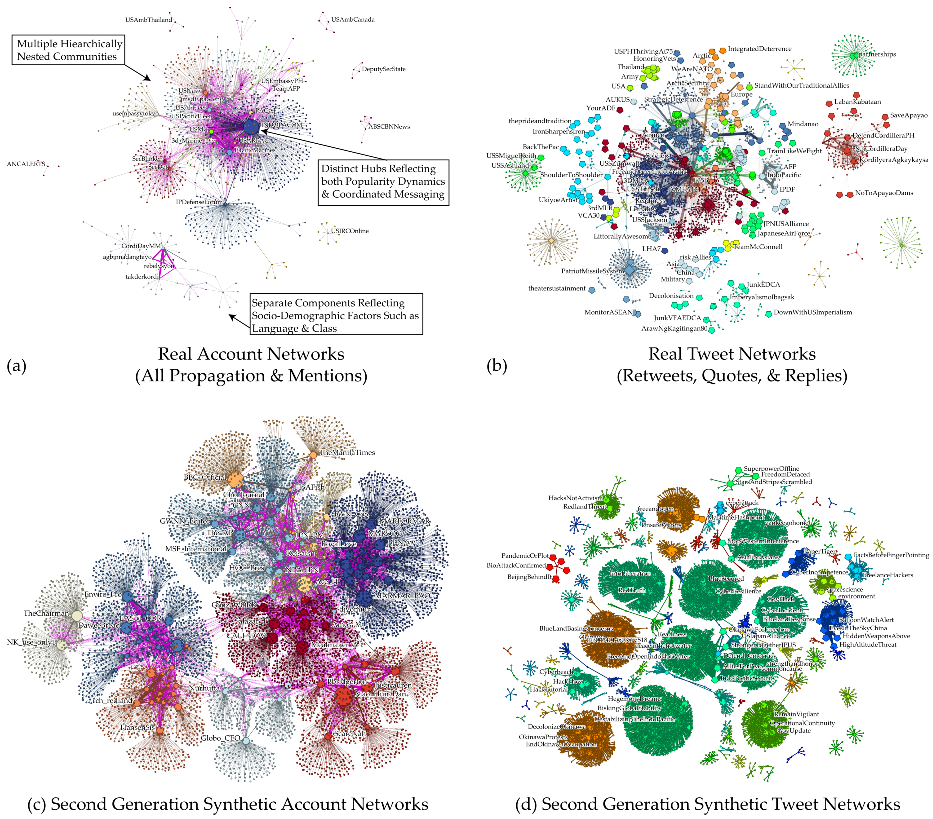

Fig. 3: Second Generation Synthetic Networks, Scenario-Based Interactions.

Second-generation networks illustrate these advances. Figure 3, panel c, shows account-level networks with multiple differentiated communities and plausible bridging, replacing the earlier hub-and-spoke artifact. Edges represent propagation (retweets, quotes, replies), with pink mentions overlaid. Node size reflects total degree; node color indicates Leiden community membership. These improvements stem from refinements in approximating the effects of selective attention. Agents now prioritize information streams based on event salience and group membership, and modulate their responses according to archetypal roles. High-profile accounts still attract engagement, but peripheral agents behave more selectively, reinforcing group structure while supporting inter-group diffusion.

Panel d shows the corresponding gains in the tweet network. In the earlier version, clusters of posts were internally coherent but lacked meaningful inter-group connections either through replies or shared sourcing, but largely due to the overuse of generic cues, which produced spurious links and obscured narrative structure. In contrast, the revised network displays distinct narrative clusters, where posts are connected not only by retweets, replies, and quotes, but also by shared hashtags, slogans, and URLs that meaningfully signal narrative alignment. The message nodes are sized by their degree in the retweet network and colored by the Leiden group. These structures emerge from approximated priming and schema logic: agents use shared signals and group-based frames to generate coherent, group-aligned discourse. To support pedagogical clarity, the number of active topics is intentionally constrained and generic hashtags (e.g. #Peace, #PacificOcean) are limited to prevent noise. Although this simplification improves interpretability, it also limits opportunities for bridging across unaligned publics. Enabling more fluid narrative movement without sacrificing coherence is a key objective of GhostField's next iteration.

5 Conclusion and Future Directions

GhostField provides a structured platform for generating synthetic information environments that support immersive network-based training. By approximating cognitive constraints such as selective attention, priming, and schema activation, large language model outputs can exhibit behaviorally plausible patterns of communication. These patterns - reflected in modular community structures, realistic narrative diffusion, and coherent group/role alignment - scaffold higher-level reasoning without requiring emergent autonomy. Through a layered architecture that integrates narrative design (AESOP), signal detection (NetMapper), and content evaluation (MOMUS), GhostField creates environments where learners can detect coordination, infer intent, and evaluate messaging strategies using the BEND framework. The system, thus, bridges the gap between content generation and analytic instruction, offering a scalable foundation for reasoning in complex socio-technical domains.

This capability has both pedagogical and strategic value. In operational contexts, adversaries exploit the same heuristics, frames, and diffusion pathways that GhostField emulates. Training analysts to recognize these features requires environments where normal behavior provides a stable contrast against coordinated manipulation. GhostField makes that contrast visible, allowing learners to distinguish between plausible human behavior and the artificial signatures of BOT networks. As training exercises scale in complexity and tempo, the need for such synthetic scaffolding will only grow. Systems like GhostField enable learning that is not only immersive but interpretable, aligning instructional goals with the realities of modern information conflict.

Future development will focus on reducing reliance on manual scenario construction by introducing structured defaults informed by empirical data and theory. Rather than developing fully autonomous agents, which risks introducing unwanted stochasticity through localized decision-making, GhostField will extend its algorithmic logic to emulate plausible cross-group behavior through population-level modeling. Scenario planners will gain greater support through system libraries and templates for agents, topics, and events, streamlining routine setup while preserving the option for fine-grained control when needed.

To address current limitations in narrative diffusion and inter-group variability, forthcoming enhancements will introduce agents with overlapping affiliations and interests, as well as parameters governing an agent's susceptibility to influence. These improvements aim to preserve scenario coherence while increasing behavioral realism - ensuring that GhostField can meet the growing demands of instructional scalability, operational tempo, and strategic relevance.

Acknowledgments

This work was supported by the Office of Naval Research under Grant Numbers N00014-21-1-2765 and N00014-21-1-2229. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Office of Naval Research or the U.S. Government.

References

1. Padgett, J.F., Ansell, C.K.: Robust action and the rise of the Medici, 1400–1434. Am. J. Sociol. **98**(6), 1259–1319 (1993). https://doi.org/10.1086/230190
2. Batagelj, V., Mrvar, A.: Pajek Datasets. (2006). Available at: http://vlado.fmf.uni-lj.si/pub/networks/data/
3. Sampson, S.F.: Crisis in a Cloister. Unpublished doctoral dissertation, Cornell University (1969), pp. 165–178.
4. Zachary, W.W.: An information flow model for conflict and fission in small groups. J. Anthropol. Res. **33**(4), 452–473 (1977).
5. Newell, A.: *Unified Theories of Cognition*. Harvard University Press, Cambridge, MA (1990). ISBN: 978-0674920996.
6. Carley, L.R., Reminga, J., Carley, K.M.: *ORA & NetMapper*. In: International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation, vol. 1, no. 3.3, pp. 2–1. Springer (2018, July).
7. Ng, L.H.X., Carley, K.M.: *BotBuster: A Mixture of Experts Approach to Social Bot Detection*. In: Proceedings of the 2023 International AAAI Conference on Web and Social Media (2023). https://doi.org/10.1609/icwsm.v17i1.22179
8. Carley, K.M.: *ORA: A Toolkit for Dynamic Network Analysis and Visualization*. In: Alhajj, R., Rokne, J. (eds.) *Encyclopedia of Social Network Analysis and Mining*. Springer, New York (2017). https://doi.org/10.1007/978-1-4614-7163-9_309-1
9. Morgan, J.H., Carley, K.M.: The One True OMEN Limited Objective Experiment (TOTO LOE) – "TOTO 24". Netanomics Technical Report #20240513-001 (2024)
10. Carley, K.M., Morgan, J.H.: Service Level Training Exercise (SLTE) 2-24: Social Cybersecurity Component. Netanomics Technical Report #20240819-002 (2024)
11. Carley, K.M., Morgan, J.H.: The One True OMEN After Action Review – "TOTO 25". Netanomics Technical Report #20250522-001 (2025)
12. Carley, K.M., Morgan, J.H.: PS25 After Action Review. Netanomics Technical Report #20250617-002 (2025)
13. Carley, K.M.: Social cybersecurity: an emerging science. Comput. Math. Organ. Theory **26**(4), 365–381 (2020). https://doi.org/10.1007/s10588-020-09322-9