

Original Article

Toward AI-Augmented Applied Systems: Emerging AI and LLM Directions Across Trajectory Prediction, Cybersecurity, and Engineering Automation

Muhammad Shahbaz¹, Muhammad Emran², Rubiya Hanni Shafiulla Kannavalli³, Farhan Ajmal Khan⁴,
Meer Zafarullah Noohani⁵

¹Department of Computer Engineering, University of Coimbra, Coimbra, Portugal.

²Computer Information Science, Faculty of Computer Information Systems, Higher Colleges of Technology, Al Ain, United Arab Emirates.

³Department of Computer Science, Sapienza University of Rome, Italy.

⁴Department of Artificial Intelligence, Sapienza University of Rome, Rome, Italy.

⁵Department of Electronic Engineering, Sapienza University of Rome, Rome, Italy.

³Corresponding Author : kannavalli.1774436@studenti.uniroma1.it

Received: 03 August 2026

Revised: 05 September 2026

Accepted: 24 September 2026

Published: 08 October 2026

Abstract - This paper reviews fifteen recent publications (2015-2026) spanning trajectory prediction, cybersecurity and trust management, classical machine learning and algorithmic optimisation, and engineering automation, and asks a forward-looking question: where is artificial intelligence — particularly Large Language Models (LLMs), foundation models, and agentic AI systems — positioned to extend each of these research directions next? For each of four thematic clusters we first summarise the current method documented in the reviewed sources, then discuss a concrete, currently-emerging AI direction grounded in recent (2025-2026) literature: LLM- and foundation-model-based trajectory reasoning for Cluster A; agentic, LLM-driven security-operations copilots for Cluster B, including a reported 6.5x improvement in malicious-alert identification over human analysts in one production deployment; learned and graph-neural-network-based approaches to the classical optimisation and graph-isomorphism problems in Cluster C; and natural-language-to-workflow agentic automation for the low-code and embedded-systems work in Cluster D. Throughout, we explicitly separate what the reviewed papers themselves established from this review's own forward-looking analysis, and close with a discussion of the reliability, interpretability, and adversarial-robustness risks that accompany each proposed direction.

Keywords - Large Language Models, Agentic AI, Foundation Models, Trajectory Prediction, Cybersecurity, Trust Management, Graph Neural Networks, Engineering Automation.

1. Introduction

A companion review of these fifteen references organized them retrospectively into a thematic taxonomy, showing that trajectory prediction, cybersecurity, classical ML/optimization, and engineering automation work by an overlapping author group forms coherent applied-systems research programmed spanning 2015-2026. This paper takes a complementary, forward-looking angle: for each of the same four thematic clusters, we ask specifically where large language models (LLMs), foundation models, and agentic AI — the dominant AI development of 2024-2026 — are positioned to extend the methods the reviewed papers document. This is necessarily a different kind of contribution from a backward-looking survey. The reviewed papers report what their authors built and measured; this paper's "emerging

direction" discussion is this review's own analysis of plausible near-term extensions, grounded in currently published AI literature (2025-2026) but not a claim about what the original authors of [1-15] have built or intend to build. We maintain that distinction explicitly throughout, rather than attributing forward-looking claims to the reviewed papers themselves.

The remainder of the paper is organized as follows. Section II describes the methodology. Section III briefly restates the current-state taxonomy as a baseline. Sections IV-VII discuss the emerging AI direction for each cluster in turn. Section VIII presents a consolidated current-to-future map. Section IX synthesizes this into quantitative counts and qualitative comparison tables and figures. Section X discusses risks and open challenges common to AI-augmenting this kind of applied-systems work, and Section XI concludes.



2. Methodology

This review extends the six-stage methodology of the companion taxonomy with a seventh, forward-projection step: for each thematic cluster, recent (2025-2026) AI/LLM literature was searched specifically for work that (a) addresses the same underlying task as the cluster's reviewed papers, and (b) uses LLMs, foundation models, or agentic architectures as its primary method, so that the "emerging direction" discussed for each cluster is grounded in published, citable work rather than speculation. Where such grounding literature was found, it is cited directly (e.g. [16-18]); where a plausible direction is discussed without a directly on-topic AI publication

identified, this is stated explicitly as this review's own extrapolation rather than implied to be an established result.

3. Current State: A Brief Baseline

Table 1 restates, in condensed form, the current method each reviewed paper documents, organized by the same four clusters as the companion taxonomy — trajectory and behavioral modelling (A), cybersecurity and trust (B), classical ML/optimization/algorithms (C), and engineering automation/software systems (D). This serves purely as the baseline against which Sections IV-VII discuss emerging AI directions.

Table 1. Condensed current-state baseline by cluster, as documented in the fifteen reviewed references

Cluster	Current method (as documented)	References
A	Transformer/skip-gram and GCN-based sequence models; kinematic template matching	[1, 2, 10]
B	IAM framework design; honeypot deception; WSN trust scoring; federated learning; kernel-level protection	[3, 5, 6, 11, 13]
C	Quadric error metrics; classical supervised classifiers; combinatorial graph algorithms	[4, 7, 15]
D	Rule-based automation frameworks; simulation software; visual low-code tooling; RTOS scheduling	[8, 9, 12, 14]

4. Cluster A: Toward Foundation-Model Trajectory Reasoning

The reviewed trajectory work uses transformer and graph-convolutional architectures trained end-to-end on trajectory-specific datasets: STTFFN [1] on taxi GPS traces, 3D-DCM [2] on pedestrian surveillance data with visual/background context, and kinematic template matching [10] on mouse-cursor motion. A 2025 survey of Large Foundation Models (LFMs) for autonomous-driving trajectory prediction documents a broader shift already underway in this exact task family: integrating large language models and multimodal LLMs to bring linguistic and scene-level semantic reasoning into trajectory prediction, which the survey reports improves interpretability and generalization to long-tail scenarios that conventional deep-learning trajectory models — the category [1, 2] both belong to — are documented to struggle with [16].

Applied to the specific papers reviewed here, this suggests a concrete extension path rather than a generic one: a foundation-model backbone could unify [1]'s vehicle-trajectory forecasting, [2]'s pedestrian trajectory-plus-context modelling, and [10]'s cursor-trajectory bot detection under a single pretrained sequence-reasoning model fine-tuned per agent type, rather than three separately-trained architectures — since all three tasks reduce to the same underlying problem (predicting a continuation of a partially observed 2D or 3D trajectory) that [10] already makes explicit by applying trajectory-style reasoning outside the vehicle/pedestrian

domain. The main open question, consistent with [16]'s own reported limitation, is whether the interpretability gain from language-grounded reasoning offsets the substantially larger pretraining data and compute requirements relative to the task-specific models in [1, 2].

5. Cluster B: Toward Agentic, LLM-Driven Security Operations

The reviewed cybersecurity work is architectural and procedural: IAM extended to bot identities [3], honeypot deception [5], WSN-based vehicular trust scoring [6], federated intrusion detection [11], and kernel-level protection [13].

A concrete, currently-deployed AI direction for this cluster is the emergence of agentic LLM-based Security-Operations-Centre (SOC) copilots, which use an LLM as a planner-executor investigation loop over security telemetry rather than as a classifier: a 2026 study of one such production system reports that its autonomous phishing-triage agent identified 6.5 times more malicious alerts than human analysts working alone, and describes a broader architecture combining a unified activity timeline, schema-validated LLM prompts, and an investigation loop that generates and tests attack-specific hypotheses [17].

Mapped onto the specific reviewed papers, this direction suggests each of Cluster B's systems as a data or action surface

for an LLM agent rather than a replacement for them: an agent could reason over the trust scores [6] produces to decide when a vehicle's traffic should be treated as suspicious, correlate [11]'s federated CAN intrusion alerts across the fleet-wide model that federation otherwise keeps siloed, and orchestrate [5]'s honeypot deployment adaptively in response to observed attacker behavior rather than following a static configuration.

[3]'s IAM-for-bots framework is a precondition for this extension rather than a target of it: an LLM security agent is itself exactly the kind of non-human identity [3] argues conventional IAM was not designed to govern, so a production agentic-SOC deployment of this kind depends on [3]'s governance layer to authenticate and constrain the agent's own actions — including against the kernel-level tampering concern [13] addresses, since a compromised kernel could feed a security agent falsified telemetry.

6. Cluster C: Toward Learned Optimisation and Graph-Neural Methods

The reviewed classical-methods work is deterministic and hand-designed: two-round QEM mesh simplification [4], classical supervised classifiers for face recognition [7], and combinatorial algorithms for graph isomorphism [15]. Each has a documented learning-based counterpart in current AI literature.

For [15]'s graph isomorphism problem specifically, Graph Isomorphism Networks (GIN) — a graph neural network architecture explicitly designed and named around the graph-isomorphism connection, shown to be as powerful as the Weisfeiler-Lehman graph-isomorphism test at distinguishing non-isomorphic graphs — are a directly on-topic, well-established AI method for the exact combinatorial problem [15] addresses algorithmically [18].

This is a case where the emerging AI direction is not merely analogous but formally connected to the reviewed paper's problem statement: [15]'s target-cell-selection algorithm and GIN-style GNNs are two different approaches (deterministic vs. learned) to structurally compare graphs.

For [4]'s mesh simplification and [7]'s face recognition, the corresponding directions are less formally tight but well documented in the broader literature this review's companion piece already surveys: learned, data-driven mesh-simplification networks that predict which edges to collapse rather than following a fixed error metric, and deep convolutional face-recognition models that have, in the broader literature, substantially displaced the classical classifiers [7] compares. In both cases, the reviewed paper

represents the classical baseline against which a future learned method's gain would need to be measured, rather than a method already superseded within the scope of what [4] and [7] themselves report.

7. Cluster D: Toward Agentic, Natural-Language-Driven Automation

The reviewed engineering-automation work spans a toll-collection framework [8], P2P energy-trading simulation software [9], low-code/no-code AI/ML tooling [12], and real-time OS concepts [14].

Of these four, [12] is already positioned closest to the emerging AI direction discussed here: low-code/no-code platforms are, by [12]'s own framing, about broadening access to intelligent automation beyond traditional software engineers, and the natural extension of that goal — replacing visual drag-and-drop configuration with natural-language-to-workflow generation via an LLM agent — is a direct continuation of [12]'s stated motivation rather than a departure from it.

For [8, 9], an agentic extension is more speculative but structurally plausible: an LLM-based negotiation agent could plausibly automate bilateral price discovery in the P2P energy market [9] simulates, and an agentic layer over [8]'s toll framework could handle exception cases (disputed charges, malformed vehicle identification) that a rule-based system routes to manual review today.

For [14]'s real-time OS concepts, the connection is the weakest of the four: RTOS scheduling has a hard correctness requirement (bounded worst-case latency) that current LLM-based approaches do not provide guarantees for, so we note this as an open research question — whether a learned scheduler could meet real-time guarantees at all — rather than a direction with existing grounding literature, in contrast to the concrete citations available for Clusters A-C.

8. Consolidated Current-To-Future Map

Figure 1 and Table II consolidate Sections IV-VII into a single view. The rightmost column's citation status is reported explicitly: Clusters A and B have a directly on-topic, currently-published AI-literature citation grounding the proposed direction ([16, 17]); Cluster C has one formally-connected citation for the graph-isomorphism sub-case ([18]) and two directionally-consistent but less formally tight extensions; Cluster D has one direction that is a direct continuation of a reviewed paper's own stated goal ([12]) and two to three that are this review's own speculative extrapolation, disclosed as such.




Fig. 1 Current method (as documented in the reviewed references) mapped to an emerging AI direction (this review's discussion) for each of the four clusters

Table 2. Consolidated emerging-AI-direction map with grounding status disclosed per row

Cluster	Emerging AI direction	Grounding
A	LLM/foundation-model trajectory reasoning unifying [1, 2, 10]	Directly cited survey [16]
B	Agentic LLM SOC copilots orchestrating [3, 5, 6, 11]	Directly cited production system [17]
C	GNN-based graph isomorphism testing (GIN) extending [15]	Directly cited, formally connected [18]
C	Learned mesh simplification / deep face recognition extending [4, 7]	Directionally consistent, general literature
D	NL-to-workflow agentic automation extending [12]	Direct continuation of [12]'s stated goal
D	Agentic negotiation/exception-handling for [8, 9]; learned RTOS scheduling for [14]	This review's own extrapolation; RTOS direction flagged as an open question

9. Results: Quantitative and Qualitative Synthesis

This section consolidates the analysis in Sections III-VIII into quantitative counts and qualitative comparisons. Because this paper is a literature analysis rather than an original

experimental study, the figures below are of two kinds only: real counts derived directly from the fifteen reviewed references (Figure 2), and this review's own explicit qualitative classifications (Figures 3-4, Table III), each labelled as such rather than presented as measured experimental results.




Fig. 2 Distribution of the fifteen reviewed references across the four thematic clusters (real count)

Cluster B (cybersecurity and trust) is the largest single cluster at five references, consistent with Section V's observation that this cluster's papers ([3, 5, 6, 11, 13]) address

different layers of a single, multi-layer trust problem; Clusters A and C are smallest at three references each.




Fig. 3 Grounding strength of each proposed emerging-AI direction, classified on a 0-3 scale by this review (0 = speculative/open question, 1 = directionally consistent with general literature, 2 = direct continuation of a reviewed paper's own stated goal, 3 = directly supported by a cited, on-topic AI publication)

Three of the six proposed directions reach the highest grounding level (score 3): foundation-model trajectory reasoning [16], agentic SOC copilots [17], and GNN-based graph isomorphism testing [18]. One direction (learned mesh simplification / deep face recognition) scores lowest among the non-zero entries (1), reflecting that it is consistent with

well-established broader-literature trends without a single directly matching citation identified in this review's search. The agentic negotiation/RTOS-scheduling direction for Cluster D scores 0, which this review treats as an explicit disclosure that no supporting AI-literature citation was identified, not as a claim that the direction is infeasible.




Fig. 4 The one directly reported quantitative result among the grounding sources: production agentic SOC copilot vs. human-analyst baseline, malicious-alert identification rate, as reported in [17]

Figure 4 reproduces the single quantitative benchmark figure available across the three directly cited grounding sources ([16-18]); [16, 18] are surveyed/theoretical contributions rather than sources of a single headline metric, which is why only [17]'s reported 6.5x figure is charted here rather than a multi-source comparison.

Table 3 extends Table I with a qualitative maturity comparison across four dimensions — autonomy level, interpretability, data dependency, and deployment maturity — rated Low/Medium/High by this review for both the current method and the proposed emerging direction in each cluster.

Table 3. Qualitative maturity comparison, current method vs. emerging AI direction, by cluster. Ratings (Low/Medium/High) are this review's own classification, not measured data.

Cluster	Dimension	Current (as documented)	Emerging AI direction
A	Autonomy / Interpretability / Data dependency / Deployment maturity	Low / Med / Med / High	High / High / High / Low
B	Autonomy / Interpretability / Data dependency / Deployment maturity	Low / High / Low / High	High / Med / Med / Med
C	Autonomy / Interpretability / Data dependency / Deployment maturity	Low / High / Low / High	Med / Med / High / Low
D	Autonomy / Interpretability / Data dependency / Deployment maturity	Low / High / Low / High	Med / Med / Med / Low

The pattern in Table 3 is consistent across all four clusters: the current, reviewed methods score High on deployment maturity and (except Cluster A) High on interpretability, while every proposed emerging direction scores Low on deployment maturity — none of the six directions in Figure. 3 describes a system already deployed at the scale of the reviewed papers' own methods. This is the same trade-off Section X discusses in prose, shown here as a structured comparison: the emerging directions trade deployment maturity and (for B-D) some interpretability for higher autonomy and, for A and C, higher data-driven adaptability.

10. Risks and Open Challenges

Three risks recur across the AI-augmentation directions discussed above, and are more acute here than in a purely retrospective review because each concerns a system that does not yet exist. First, reliability and hallucination: an LLM-based trajectory-reasoning agent (Section IV) or SOC copilot (Section V) that generates plausible but incorrect output is a materially different failure mode from a deterministic algorithm like [1]'s STTFN or [15]'s graph-isomorphism algorithm producing a wrong-but-reproducible answer, and the production SOC system reporting a 6.5x detection improvement [17] does not, on its own, establish the system's false-positive rate or behavior under adversarial input crafted specifically to mislead the agent's reasoning loop.

Second, the same non-human-identity governance problem [3] that raises for RPA/AI bots applies with more force to an autonomous security agent with investigation and response privileges (Section V): an agentic SOC copilot is a higher-value target for compromise than the RPA bots [3]

originally scoped, since a compromised security agent could actively suppress or misdirect its own alerts.

Third, the data and compute requirements of foundation-model approaches (Section IV, VI) are substantially higher than the task-specific models in [1, 2, 4, 7], which may be a poor trade-off for organizations operating at the scale the original papers target, particularly where — as [16, 19, 20] itself reports for trajectory prediction — the primary benefit is interpretability rather than raw accuracy.

11. Conclusion

This paper took a forward-looking angle on the same fifteen references reviewed in a companion taxonomy, asking specifically where LLMs, foundation models, and agentic AI are positioned to extend each cluster's current methods. Two directions have direct, currently-published grounding — foundation-model trajectory reasoning [16] for Cluster A and agentic SOC copilots [17] for Cluster B — and one has a formal, named connection already established in the AI literature: Graph Isomorphism Networks [18] for the exact combinatorial problem [15] addresses algorithmically. Other directions, particularly for Cluster D, are this review's own extrapolation and are disclosed as such rather than attributed to the reviewed papers.

Across all four clusters, reliability, non-human-identity governance, and the data/compute cost of foundation-model approaches are the open challenges that would need to be resolved before any of the discussed directions could responsibly extend the deployed, production-grade systems the reviewed papers describe.

References

- [1] Zain Ul Abideen et al., "Improved Deep Learning-based Approach for Spatial-temporal Trajectory Planning Via Predictive Modeling of Future Location," *KSII Transactions on Internet and Information Systems*, vol. 18, no. 7, pp. 1726-1748, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [2] Zain Ul Abideen et al., "Advancing Pedestrian Trajectory Prediction with Interaction-Aware 3D-Dual Contextualized Modeling," *Egyptian Informatics Journal*, vol. 31, pp. 1-17, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [3] Syed Talha Hussain et al., "Securing the Bot Eco System: IAM Framework for RPA and AI BOTS," *2026 IEEE 5th International Conference on Computing and Machine Intelligence (ICMI)*, pp. 1-9, 2026. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [4] Han Chang et al., "Two-Round Optimization Algorithm based on Quadric Error Metrics," *IEEE Access*, vol. 13, pp. 30021-30035, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [5] Muhammad Shahbaz et al., "Deployment of Honeypot System for Threat Detection and Intelligence in Modern Cybersecurity Architectures," *2026 1st International Conference on Innovations in Information and Communication Technology (IICT)*, Jamshoro, Pakistan, pp. 1-5, 2026. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [6] Muhammad Haleem Junejo, Muhammad Shahbaz, and Muhammad Emran, "Trust-based Management for Internet of Vehicles Via Wireless Sensor networks: Safeguarding Trust During Journeys," *2026 IEEE 5th International Conference on Computing and Machine Intelligence (ICMI)*, Al-Ahsa, Saudi Arabia, 2026. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [7] Nisar Ahmed et al., "Face Recognition Comparative Analysis using Different Machine Learning Approaches," *Advances in Science and Technology Research Journal*, vol. 15, no. 1, pp. 265-272, 2021. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [8] Abdul Mujeeb et al., "Electronic Toll Collector Framework," *Advances in Science and Technology Research Journal*, vol. 16, no. 1, pp. 294-302, 2022. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [9] Nisar Ahmed et al., "A Software Capable of simulating Peer-to-Peer Energy Trading," *International Journal of Applied Physical Sciences*, vol. 7, pp. 18-26, 2021. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [10] Hafiz Shafiq Ur Rehman Khalil, Nisar Ahmed, and Yijun Yang, "Mouse Pointing Endpoint Prediction to Distinguish between Human and Bot using Kinematic Template Matching," *International Conference on Human-Computer Interaction*, Montreal, QC, Canada, pp. 343-364, 2026. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [11] AHCENE Bounceur et al., "Privacy-Preserving Multiclass Intrusion Detection in In-vehicle CAN Networks using Federated Learning," *Scientific Reports*, 1-53, 2026. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [12] Nisar Ahmed et al., "The Synergy of Low-code/No-code and AI/ML: Enhancing Intelligent Automation," *World Journal of Engineering and Technology*, vol. 13, no. 4, pp. 754-763, 2025. [[Google Scholar](#)]
- [13] Nisar Ahmed et al., "How to Protecting Kernal Code and Data," *International Journal of Computer and Organization Trends (IJCOT)*, vol. 5, no. 4, pp. 20-23, 2015. [[CrossRef](#)] [[Publisher Link](#)]
- [14] Nisar Ahmed, "Understanding Real Time OS Concepts," *International Journal of Scientific and Engineering Journal*, vol. 6, no. 9, pp. 1-6, 2015. [[Google Scholar](#)]
- [15] Utku Cem Sankal et al., "Target Cell Selection Algorithms in Graph Isomorphism Problem," *International Journal of Scientific and Technical Research in Engineering (IJSTRE)*, vol. 6, no. 4, pp. 35-47, 2021. [[Google Scholar](#)] [[Publisher Link](#)]
- [16] Wei Dai et al., "Large Foundation Models for Trajectory Prediction in Autonomous Driving: A Comprehensive Survey," *arXiv Preprint*, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [17] Scott Freitas, and Amir Gharib, "GenAI-Driven Threat Detection with Microsoft Security Copilot," *arXiv Preprint*, pp. 1-8, 2026. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [18] Keyulu Xu et al., "How Powerful are Graph Neural Networks?," *arXiv preprint*, pp. 1-17, 2019. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [19] Falak Ali, "Real-Time Decision Processing in Artificial Intelligence Systems: A Survey of Scheduling, Inference, and Security Foundations," *International Journal of Emerging Trends in Engineering and Development (IJETED)*, vol. 16, no. 5, pp. 1-4, 2026. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [20] Falak Ali, "Divide-and-Conquer Decomposition Strategies for Real-Time Artificial Intelligence Decision Systems: A Comparative Survey," *International Journal of Emerging Trends in Engineering and Development (IJETED)*, vol. 16, no. 5, pp. 1-6, 2026. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]