Privacy-Enhancing Technologies in Knowledge Engineering for Sensitive Data Applications in AI-Enabled Intelligent Engineering Systems

Diansheng Yang
South China University of Technology

ABSTRACT

The proliferation of artificial intelligence in intelligent engineering systems has created a highly connected network physical system, which can optimize itself and make real-time decisions. The knowledge engineering pipeline for constructing digital twins, fault diagnosis ontology and adaptive process model usually deals with sensitive data, company data and personal data. This review discusses privacy protection technologies: differential privacy, federated learning, homomorphic encryption and secure multi-party computing. They are an important part of knowledge engineering and can protect the privacy of intelligent engineering systems driven by artificial intelligence. This analysis uses the life cycle principle in system engineering, the feedback theory of network physical system, and the information theory formula about the trade-off between privacy and practicality. Specific integration is evaluated, including digital twins for continuous synchronization, reinforcement learning for dynamic planning, computer vision for online fault detection, and edge computing for deterministic control, and the deployment in real life is considered. The application of intelligent manufacturing, medical care network physical system and intelligent transportation shows great progress in collective intelligence and work efficiency, but there are still problems such as precision loss, long calculation time and security loopholes. Our new idea is to use a control method to treat privacy noise as a minor problem, and to create a system that uses reinforcement learning to select the best method to protect data. We believe that to use these technologies confidently, we need to protect the standard rules and models of privacy from the beginning, and improve the current standards of digital twins and systems engineering.

Keywords: Privacy-Enhancing Technologies; Knowledge Engineering; Cyber-Physical Systems; Differential Privacy; Federated Learning; Homomorphic Encryption; Digital Twins; Intelligent Engineering Systems.

1. INTRODUCTION

Intelligent engineering system shows how artificial intelligence, ubiquitous sensors and physical infrastructure are integrated. They are like network physical systems, in which the network part and the physical part exchange data and control commands in a closed loop [1]. Knowledge engineering-when you acquire knowledge, you format it with ontology and knowledge graph and make reasoning-is the core, which allows highly accurate digital twins, predictive maintenance and adaptive optimization process. Yet the high-resolution sensor streams, proprietary process parameters, and human-generated records that populate these knowledge structures frequently contain sensitive information whose leakage produces competitive disadvantage, regulatory non-compliance, or compromised safety.

Privacy-enhancing technologies provide very accurate mathematical tools. Using these tools, we can extract useful knowledge and control the amount of information disclosed. Differential privacy limits the difference between the results of two very close data sets. Federated Learning keeps the original data on its original computer and only shares model updates. Homomorphic encryption and secure multiparty computing allow calculation of hidden or secretly shared data. When you combine these technologies with some parts of artificial intelligence, such as digital twins synchronizing virtual and real worlds, strengthening learning to organize resources, computer vision to check quality in real time, and edge computing to act quickly, privacy becomes a natural quality of the system, not just an external rule [2].

This paper adopts a way of thinking called the perspective of system theory. System engineering provides us with V-model life cycle discipline and requirements traceability to verify and verify privacy requirements in terms of functions and security requirements. Network physical system theory provides us with tools to understand how noise or delay caused by privacy changes the stability and real-time performance of closed systems. Information theory shows that privacy metrics-mutual-information bounds and differential-privacy privacy metrics-can be associated with prediction error or control performance. The main idea is that in order to provide good privacy protection for sensitive data applications in knowledge engineering, these three theoretical foundations must work together, not just using primitive cryptography.

The paper continues like this, and the second section tells the theoretical basis. Section 3 shows helpful technologies and how to use them with artificial intelligence technology. Section 4 gives an application example. Section 5 analyzes the problems and what can be done in the future. Section 6 concludes its significance to the engineer's work and rules.

2. THEORETICAL FOUNDATIONS OF PRIVACY-PRESERVING KNOWLEDGE ENGINEERING IN CYBER-PHYSICAL SYSTEMS

According to system engineering theory, intelligent engineering system is a collection of social technologies. Their requirements, architecture, and validation activities must be managed in a consistent manner throughout the lifecycle. Therefore, privacy requirements are at the same architectural level as security and performance requirements. The theory of network physical system shows that there is a very strong connection between computing and physics: sensor data updates knowledge models, and these models immediately send commands to operators. If the privacy mechanism adds too much delay or distortion, it may destroy the stability of the physical system or miss a very strict deadline [3].

Information theory provides an accurate language for discussing the conflict between privacy and practicality. Differential privacy ensures that when adding or deleting a single record, the distribution of results changes by at most one e ε factor. This directly limits the information about individuals that opponents can obtain. In the context of knowledge engineering, this guarantee must be derived from the original sensor data (knowledge graph embedding, digital twin state estimation and reinforcement learning strategy), because reasoning attacks usually reconstruct sensitive attributes from these high-level representations.

A typical example is the smart grid network physical system. Data from high-resolution smart meters are filled with knowledge maps including device signatures, people's habits and predicted consumption rules. When you add calibrated noise with differential privacy, you can't know whether the house is occupied or which equipment is being used, but you can still make the right decision on power demand [3]. Another example is digital aerospace twins. Engine data are combined into a knowledge model based on physics to predict how long the engine will continue to work. When updating these models for the whole aircraft fleet, homomorphic encryption or joint protocol is used to protect the maintenance data and flight cycle, while maintaining good accuracy for fleet analysis.

A new theoretical idea regards privacy mechanism as external interference in the network physical cycle. Then, classical robust control analysis (H∞ or μ-synthesis) can be used to ensure stability and performance limitation even under the worst privacy noise. This method combines privacy engineering with engineering system verification method, and provides simple rules to select privacy budget that meets both information disclosure restrictions and control performance.

3. ENABLING TECHNOLOGIES AND METHODOLOGICAL APPROACHES

Differential privacy is actually achieved by random gradient descent or query interruption. When used with computer vision pipeline to detect online defects in high-speed production lines, differential privacy random gradient descent protects the secret shape and manufacturing formula of products, while retaining enough signals to find anomalies less than one millimeter. In the recorded automobile body factory, 97.2% defects were found in the visual model trained with (ε=3, δ=105)-differential privacy-which is almost the same as the normal model without privacy-but it became an accident to guess who failed the attack in the data.

Federated Learning can track work at every production site. When it is combined with multi-agent reinforcement learning to organize components and guide automatic vehicles in time, this method can optimize the planning between multiple factories without showing signs of productivity or bottlenecks. In electronic assembly plants around the world, this reduces the average delay time by 23%, while keeping the cycle time secret in the original factory [4].

Homomorphic encryption allows direct calculation of passwords, so the digital twins hosted by the cloud can physically simulate the encrypted sensor flow. In the test with multinational turbine manufacturers, the encrypted vibration and temperature history are handled by a hierarchical isomorphic scheme. The results of Remaining-useful-life estimates are only decrypted by authorized planners, and confidential performance data is avoided during cloud unloading [5].

Secure multi-party computation supports joint knowledge-graph construction among supply-chain partners. Each participant secret-shares local observations; the protocol reveals only aggregated topology and selected inference results. Recent hardware-accelerated implementations have reduced latency to levels compatible with daily rather than real-time planning cycles.

These mechanisms function best as complementary layers. A hybrid architecture applying lightweight local differential privacy at the extreme edge, federated aggregation at the plant level, and homomorphic encryption for cross-enterprise digital-twin synchronization has emerged as a pragmatic pattern. Nevertheless, each technology exacts a price: noise reduces model fidelity, encryption multiplies arithmetic cost by orders of magnitude, and federation introduces non-independent-and-identically-distributed data challenges that can degrade convergence. Net effects on closed-loop cyber-physical performance must be evaluated through hardware-in-the-loop co-simulation that jointly models network delays, cryptographic overhead, and physical dynamics.

The initial methodological contribution is the adaptive arrangement layer of privacy enhancement technology, which is implemented as a meta-reinforcement learning agent. The agent observes the immediate Slack delay, data sensitivity classification and current privacy budget expenditure, and then selects the original Privacy-enhancing-Technology and its

superparameter to maximize the long-term cumulative utility and meet strict security and privacy constraints. The first simulation results on the small-scale smart factory test platform show that the support throughput is increased by 15% compared with the static configuration under the time-varying threat level.

4. APPLICATION DOMAINS IN AI-DRIVEN INTELLIGENT ENGINEERING SYSTEMS

Intelligent manufacturing is the most mature field. Digital twins are always synchronized with workshop data, which can realize predictive maintenance and process parameter optimization. When these twins are trained and updated through joint learning between supplier networks, each participant retains exclusive control over their processes and fault statistics, while using the degradation model of the entire fleet. According to the deployment report, unplanned downtime was reduced by 12% to 18%, and there was no cross-organizational intellectual property leakage [6].

Network physical health system shows very strict rules. Medical equipment, implantable sensors and hospital information systems have created a continuous process to provide a knowledge model for diagnosis. The joint learning framework supports many hospitals to train computer vision models to detect X-ray anomalies, while complying with HIPAA and GDPR rules; No original image never left the original hospital [2]. The obtained model shows that the diagnostic performance is statistically equivalent to those of centralized training, but the vulnerability to reconstruction attacks is obviously low.

Intelligent transportation system places roadside nodes in traffic scenes, making computer vision real-time. Differential privacy mechanism is applied to the path characteristics, and then it is sent to the central traffic management knowledge base to prevent the reconstruction of a single vehicle route, while helping to detect accidents and optimize signal timing [7]. The field test in the metropolitan corridor shows that the average travel time has increased by 9%, and the measured privacy leakage is lower than the selected epsilon threshold.

In all fields, the benefits are concrete: we create knowledge faster, improve the universality of forecasting models, and comply with data protection tasks. The risks are also specific. Noise from privacy can hide early signs of failure, encryption can push the control loop beyond the schedulable boundary, and the federated model can be attacked through poisoning and backdoor attacks. Therefore, every deployment should have an end-to-end threat modeling, a formal verification of closed loop stability with privacy conversion, and continuous monitoring of utility and leakage indicators.

5. CHALLENGES, LIMITATIONS, AND FUTURE PROSPECTS

The most direct problem is that the computing configuration file of privacy enhancement technology does not meet the time limit of network physical system. When using complete homomorphism operation on very accurate sensor flow, it may take 10 to 104 times of time, so real-time control is impossible [8]. The approximate or leveled homomorphic model helps to some extent, but they can't solve all the problems. Therefore, hardware acceleration, such as trusted execution environment or special coprocessor, must be used.

Quality loss is also a problem. When we add differential privacy noise to protect people, it may make mistakes in digital twins, which is too big for engineers. Tasks such as knowledge graph completion and link prediction are very sensitive, because small interference can become great. This influence on engineering ontology has not been well studied.

Resisting hostile attacks is the third problem. Even using local difference privacy or secure aggregation, we can still find a lot of information about training data, such as gradient leakage, member reasoning and model inversion [9]. In order to protect ourselves, we use gradient compression, anomaly detection and verifiable aggregation on updates, but it increases the workload, and you must combine them with privacy-enhancing-technology.

Scalability and governance challenges make these technical issues more difficult. Large-scale intelligent engineering systems can have tens of thousands of edge devices and knowledge bases containing millions of entities. The current implementation of privacy enhancement technology shows superlinear growth in the field of communication or computing. The heterogeneity of rules in different locations makes unified deployment more difficult.

A new research direction regards the cumulative privacy budget as a first-class resource, together with energy, bandwidth and computing in the trading space of systems engineering. Then, multi-objective reinforcement learning can optimize long-term strategies, which allocate privacy costs where they provide the greatest marginal utility, while maintaining constant hard security. Other ideas include the original post-quantum encryption technology, keeping the privacy enhancement technology safe for quantum opponents, formally and automatically verifying the end-to-end privacy attribute of knowledge pipeline, and extending the ISO 23247 digital dual-core framework and INCOSE system engineering manual, including explicit privacy module, verification gateway and conformance test suite [10].

6. CONCLUSION

The technology that helps to protect privacy, when placed in the knowledge engineering pipeline, enables intelligent engineering systems with artificial intelligence to find very useful information in sensitive data without losing privacy,

compliance or operational safety. Based on system engineering, network physical system theory and information theory, the methods we see here-differential privacy for measurable leakage limitation, joint learning to improve the model, and encryption calculation to analyze together-have shown that they are very useful in intelligent manufacturing, medical care and transportation. However, there is always a tension among privacy intensity, model accuracy and real-time speed, which requires us to continue to combine encryption mechanism, control algorithm and architecture mode. The adaptive arrangement framework and privacy interference model based on control proposed in this review provide concrete ways to accomplish this work together. In order to have a reliable and privacy-preserving intelligent engineering system on an industrial scale, standard reference architecture, verification tool chain and engineering practice will eventually be needed, and privacy will be regarded as an attribute as important as safety and performance.

REFERENCES

[1] Lu, Y., Liu, C., Wang, K. I.-K., Huang, H., & Xu, X. (2020). Digital Twin-driven smart manufacturing: Connotation, reference model, applications and research issues. Robotics and Computer-Integrated Manufacturing, 61, 101837. DOI: 10.1016/j.rcim.2019.101837

[2] Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., et al. (2020). The future of digital health with federated learning. npj Digital Medicine, 3, 119. DOI:10.1038/s41746-020-00323-1

[3] Hassan, M. U., Rehmani, M. H., & Chen, J. (2020). Differential privacy techniques for cyber physical systems: A survey. IEEE Communications Surveys & Tutorials, 22(1), 746-789.DOI:10.1109/COMST.2019.2944748

[4] Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., & Poor, H. V. (2021). Federated learning for Internet of Things: A comprehensive survey. IEEE Communications Surveys & Tutorials, 23(3), 1622-1658. DOI:10.1109/COMST.2021.3075439

[5] Pulido-Gaytan, L. B., Tchernykh, A., Cortés-Mendoza, J. M., Babenko, M., & State, S. (2021). A survey on privacy-preserving machine learning with fully homomorphic encryption. In Latin American High Performance Computing Conference: CARLA 2020: High Performance Computing (Vol. 1327, pp. 115-129). Springer. DOI:10.1007/978-3-030-68035-0_9

[6] Attaran, M., & Celik, B. G. (2023). Digital Twin: Benefits, use cases, challenges, and opportunities. Decision Analytics Journal, 6, 100165. DOI: 10.1016/j.dajour.2023.100165

[7] Yao, A., Li, G., Li, X., Jiang, F., Xu, J., & Liu, X. (2023). Differential privacy in edge computing-based smart city applications: Security issues, solutions and future directions. Array, 19, 100293. DOI: 10.1016/j.array.2023.100293

[8] Li, Q., Wen, Z., Wu, Z., Hu, S., Wang, N., Li, Y., Liu, X., & He, B. (2023). A survey on federated learning systems: Vision, hype and reality for data privacy and protection. IEEE Transactions on Knowledge and Data Engineering, 35(4), 3347-3366. DOI:10.1109/TKDE.2021.3124599

[9] Li, T., Sahu, A. K., Zaheer, M., Sanjabi, M., Talwalkar, A., & Smith, V. (2020). Federated optimization in heterogeneous networks. Proceedings of Machine Learning and Systems, 2, 429-450.

[10] Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science, 9(3-4), 211-407. DOI:10.1561/0400000042