Research Statement

The number and complexity of wirelessly connected devices are growing dramatically, and the current network infrastructure has several significant limitations, including latency and privacy issues, that are not sufficient to fully support applications such as autonomous vehicles, healthcare monitoring, etc. In addition, deep learning and data-analytics-based AI systems are expected to be a major part of Next Generation (NextG) networks [1], and these processes can require tremendous computation and communication resources, potentially causing further significant latency and privacy leakage in training and inference processes. In my research, I focus on addressing these issues as they are expected to arise in the NextG networks. Taking a mathematical approach, I bridge the gap between real-world applications and innovative solutions, generating insights that could fundamentally reshape the future of NextG networked systems. By drawing on techniques from information theory, network science, and machine learning, my work tackles the most pressing challenges in **wireless** and **social networks**, with a strong emphasis on enhancing **data privacy**, **scalability**, and the **efficiency of data transmission**.

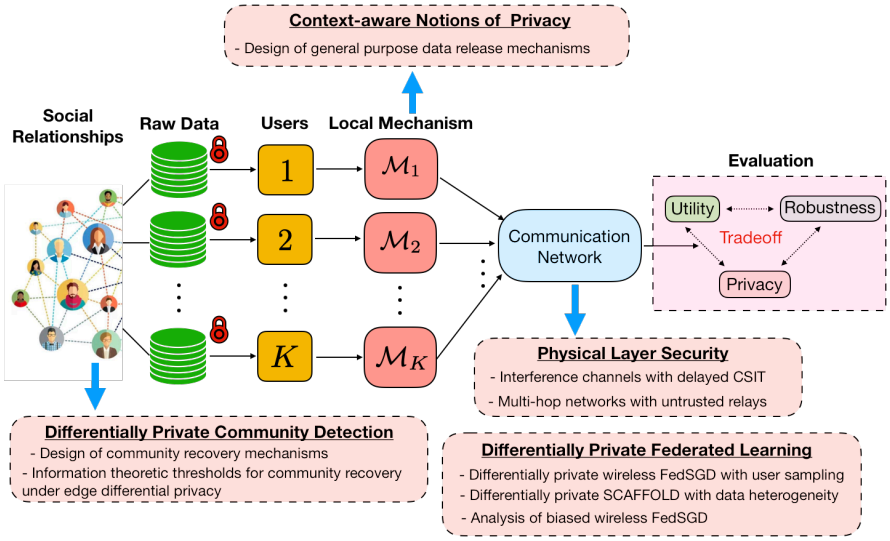

Fig. 1: Designing private, scalable and efficient algorithms over information networks.

In the following, I describe my key contributions and outline a long-term vision for advancing privacy technologies in NextG networks: My research efforts tackle several critical challenges, including enhancing privacy-preserving machine learning at the network edge, developing context-aware privacy mechanisms that adapt to evolving user needs, and advancing privacy-preserving analytics for structured data. Additionally, I explore the potential of information-theoretic approaches to secure wireless communications. These efforts form a foundation for future innovations, which I plan to expand on as my work evolves.

Previous and Current Work

Differential Privacy at the Network Edge. Federated learning (FL) is a collaborative approach that allows multiple users to train a machine learning model collectively, facilitated by a parameter server, without sharing raw data (e.g., [1]). The appeal of FL lies in its ability to parallelize training across increasingly powerful end-user devices while helping preserve data privacy. However, even the exchange of model parameters (i.e., rather than raw data) in FL can unintentionally leak sensitive information and can also lead to high communication costs and latency, highlighting the need for **efficient communications protocols** and means for **safeguarding privacy**. My research interest broadly falls into the design of cross-layer communication platforms for efficient on-device and edge learning in wireless networks. I develop privacy-preserving distributed on-device machine learning solutions for next-generation wireless networks. In my research work, I have examined the deployment of **wireless networks** as platforms for machine learning, specifically focusing on enhancing communication efficiency and strengthening privacy guarantees through advancing federated and decentralized learning algorithms, and analyzing their communication dynamics within the wireless environment.

My contributions to this domain center around pioneering the use of the **superposition property of wireless channels** for bandwidth-efficient aggregation of users' models or gradients, with strong local differential privacy (LDP) guarantees. I was the first to show that privacy leakage per user **decreases** as the number of users increases, a key insight that surpasses traditional orthogonal transmission schemes where leakage remains constant. This significant

result demonstrates how leveraging the superposition effect during transmission allows gradients to be aggregated naturally, without additional processing, while safeguarding privacy. Through my work, I also introduced the novel concept of using optimized power control factors in conjunction with channel noise as a perturbation mechanism. This approach not only enhances differential privacy guarantees but also ensures that individual user contributions are effectively obscured within the aggregated transmission, making it exceptionally difficult for adversaries to isolate specific signals. This work is the first to show that this strategy outperforms existing privacy-preserving methods, offering a more scalable and efficient solution for wireless aggregation with robust privacy protection.

Context-aware Privacy. Contextual information is pivotal in various applications, from location-based services to mobile health data. This information can include data distribution, correlations within the data, and user privacy expectations, among other factors. For example, in location-based services, people are more likely to be near famous landmarks than lesser-known locations. Similarly, mobile health data often contains background knowledge, like the prevalence of certain diseases based on prior studies. Finally, **prompt history** for **text-to-image generation** applications. Moreover, privacy-sensitive attributes, such as a user's web browsing history, may correlate with raw data. While the privacy research community has generally avoided modeling background knowledge explicitly, there is a recent trend to incorporate such contextual or partial knowledge into privacy notions. My research endeavors to construct a robust **theoretical foundations** for context-aware information privacy, coupled with the design of **optimal privacy mechanisms** tailored for **localized** settings. This work focuses on unraveling and mitigating the vulnerabilities of differential privacy to Bayesian inference attacks, which remain insufficiently explored within the privacy research community.

As I showed in my research work, context-aware privacy notions offer improved utility-privacy trade-offs compared to context-free approaches like local differential privacy (LDP), in which LDP guarantees the worst case indistinguishability for any two different inputs, which significantly deteriorate the utility. I have studied and proposed a new privacy notion called local information privacy (LIP) that guarantees the ratio of the posterior and the prior data are bounded. I have also explored the relationships of LIP with different context-free (e.g., LDP) and other context-aware privacy notions (e.g., mutual information privacy). The proposed notion is a natural bridge between **average** and **worst-case** privacy notions. The main advantages of LIP are that (a) it can be seamlessly adapted to account for a variety of scenarios with different contextual/prior knowledge and (b) it leads to simple and modular mechanism designs with low complexity (needs only $O(n^2)$ constraints for all input-output combinations in contrast to $O(n^3)$ for LDP-like notions). I have also designed optimal general-purpose context-aware privacy mechanisms, including Laplacian, Gaussian, and sampling-based mechanisms, and explored the associated trade-offs between utility and privacy. Notably, the primary advantage lies in the adaptability of noise or perturbation levels, which can be adjusted based on differing input instances or prior data distributions.

Private Data Analytics on Graphs. Graph datasets are ubiquitous, representing nodes as individuals and edges as relationships, such as social or romantic connections. The sensitive nature of these datasets enlightens a fundamental question, *"What information can be learned on graphs without violating the privacy of individual relationships and interactions?"* To address this, two common variants of differential privacy have emerged for graph datasets: edge differential privacy, which aims to protect relationships, and node differential privacy, which focuses on minimizing the impact of user participation on the learning algorithm. In this research area, I am interested in developing private algorithms for several fundamental problems in graph mining and network science.

My current research centers on **community detection** over graphs, where the goal is to find hidden network patterns. In particular, I study and design new algorithms for privately classifying users, such as by their political affiliations. In my recent work [5], I examined the fundamental limits of community detection in graphs generated from a stochastic block model (SBM). This research work involved deriving information theoretic sufficient conditions on the separation between intra-community and inter-community connection probabilities, differential privacy budget, and computational complexity for the exact recovery of community labels. Further, I have proposed the first differentially private community recovery algorithms and studied the fundamental tradeoffs between **privacy requirements** and **computational complexity**. I have comprehensively analyzed three distinct private community detection methods: stability-based, sampling-based, and graph-perturbation. I have shown that both the sampling and stability-derived approaches increased computational complexity and longer computational time but a good tradeoff between privacy and data recovery.

Statistical Learning for LLMs. Statistical learning for large language models (LLMs) is emerging as a powerful framework for building reliable and efficient AI systems. Recent work has explored principled approaches such as conformal inference, hypothesis testing, and sequential decision-making to quantify uncertainty, improve safety, and

guide interactions with LLMs. These developments highlight the importance of treating LLMs as stochastic systems that can be analyzed and controlled using statistical tools.

My work contributes to this direction through **trustworthy and efficient inference for LLMs**. In particular, I have studied the problem of detecting and localizing edits in watermarked LLM outputs. While watermarking enables identification of AI-generated text, it becomes less reliable under post-generation edits. To address this, I developed a combinatorial, pattern-based watermarking framework that supports both global verification and fine-grained localization of edits. This work led to a pending patent and represents a step toward robust post-hoc inference for generated content.

More broadly, I am interested in **statistical learning for LLM-based data reduction**, especially document summarization. From a statistical perspective, summarization can be viewed as a form of structured dimensionality reduction, where the goal is to preserve key semantic information while reducing length and computational cost. This viewpoint opens new opportunities to design **efficient and communication-aware summarization methods** with principled performance guarantees. Looking ahead, my goal is to develop unified frameworks that combine **statistical inference, efficiency, and trustworthiness** in LLM systems, bridging ideas from information theory and modern machine learning.

Other Work

Towards Understanding Privacy Attacks in Split Learning. I am currently deepening my expertise in the innovative domain of Split Learning and its potential applications, especially in wireless communication technologies. In contrast to federated learning, split learning emerges as a more versatile framework, presenting itself as particularly advantageous when data is needed to remain anchored to the client's device or when the transmission of large model parameters becomes a challenge. Establishing theoretical convergence guarantees within the differential privacy landscape is a cornerstone of my exploration. This initiative is driven by a vision to shed light on the seamless integration of robust privacy mechanisms within the split learning framework. One particular research aim I am interested in is understanding the trade-off between privacy leakage and local model complexity that will help the research community determine the right choice of privacy parameters for potential privacy-preserving release mechanisms. Shallow models, while being computationally efficient, tend to reveal more about the input data. Conversely, more complex models offer better privacy protection at the expense of computational demands. Due to the limitations in existing theoretical metrics, most works in the privacy field rely on empirical measures to demonstrate the privacy properties against a limited set of attacks without solid theoretical arguments, specially for the data reconstruction at the adversay's side. I want to point out that this problem is of vital interest, especially for complex DNN models, e.g., generative models.

AI/ML based for Wireless Technology Recognition. I explored the problem of automatic wireless technology recognition by proposing a distributed machine-learning approach that offered faster and more effective results. This work was part of my collaboration with **imec** in Europe, where we tackled key challenges in wireless communications. One major issue was the interference and performance degradation caused by 5G New Radio Unlicensed (NR-U) sharing its frequency band with WiFi. Common solutions, such as GPU-based AI/ML methods, often led to slow processing and increased latency. To address these limitations, we focused on a distributed machine-learning strategy to improve technology coexistence in shared spectrum environments. As a result, our proposed transmission scheme significantly enhanced communication efficiency, reducing the data transfer between edge devices and the server by 96% compared to the previous centralized based approaches, with only a slight degradation in accuracy performance of around 2%.

Healthcare and Computational Epidemiology. My previous research concentrated on the complexities of processing count queries on genomic data within biomedical databases while upholding **perfect** privacy constraints, a critical consideration in genome-wide association studies and similar explorations. In our landmark paper [8], my colleague and I pioneered optimal privacy mechanisms while maintaining perfect information-theoretic privacy for **sensitive genotypes**, particularly those indicating rare diseases or specific health traits. Furthermore, our work established information-theoretic lower bounds for error probabilities in privacy-preserving query responses, demonstrating that our approaches approximate or even match these limits under certain conditions. In expanding my research from genome privacy to healthcare analytics, I am motivated to contribute to epidemiological research with a novel approach. I intend to leverage computational epidemiology and data analytics to craft sophisticated, real-time surveillance models that utilize anonymized data, enabling the swift identification and mitigation of public health risks.

AI-Enabled Closed-Loop Control for 5G/6G Networks. In collaboration with Ghent University–imec, we developed and demonstrated an AI-native closed-loop control system at the EuCNC & 6G Summit 2025. The system ingests real-time telemetry from a live 5G O-RAN deployment and leverages containerized microservices to support scalable inference pipelines. At its core, a large language model (LLM) interprets multi-source network data to make dynamic slice-level resource allocation decisions. The framework integrates with a digital twin environment and enables end-to-end control by coupling observability with LLM-driven actuation. While the demo targeted 5G/6G network slicing, the architecture generalizes to other cyber-physical systems, including smart grids and industrial IoT.

Future Directions

Along with continuing my research in the previously mentioned areas, I am also focused on diversifying my research interests and leadership roles. In the coming four to five years, I plan to spearhead and oversee new research endeavors that resonate with my areas of expertise, specifically in NextG communication systems and social data analytics. These projects will aim to make meaningful contributions to these sectors. I aim to introduce innovative approaches and solutions to the complex challenges inherent in these essential domains, drawing upon my knowledge and experience.

Thrust 1: Semantic Communications. In the digital communications arena, the need for robust, secure, and efficient information transmission is paramount. I am focused on exploiting the potential of language and masked image models, underexplored areas in semantic communications [13]. Despite progress, research gaps in channel coding, privacy, and security persist, particularly against adversarial attacks. Current studies often treat deep neural networks as fuzzy '**black boxes**,' lacking performance and privacy guarantees — crucial for communications systems. These networks are also commonly trained under conditions that differ from their actual use, making consistent performance guarantees elusive. My research aims to bridge these gaps by forging theoretical frameworks that provide reliable benchmarks for neural network models throughout training and deployment in semantic communications. My research interest extends to enhancing security and privacy in integrated sensing and communication **(ISAC)** amidst the growing use of sensing technologies and their associated vulnerabilities. I aim to address the risk of unauthorized data interception through resilient system design. Drawing on my expertise in differential privacy, I intend to contribute to advancing secure Wi-Fi network standards, focusing on the emerging **IEEE 802.11bf** landscape. Another goal is to enable **multi-user generative AI-based** transmission schemes over wireless networks. A significant challenge in this field is that while recent studies have demonstrated the effectiveness of these approaches compared to traditional transmission schemes, even when incorporating error correction codes like LDPC, the results are often limited to simple scenarios, such as point-to-point communications or environments with linear interference. These conditions do not accurately reflect the complexities of practical wireless communications, where non-linear interference and multi-user interactions are prevalent. My research will focus on overcoming these limitations by developing more sophisticated models and techniques that can be applied in real-world, multi-user wireless environments.

Thrust 2: Privacy-preserving Inference over Network Edge. Collaborative inference in NextG networks can significantly enhance AI applications, such as autonomous driving, personal identification, and activity classification. This approach involves three key stages: a) data acquisition through sensing, b) feature extraction, and c) feature encoding for transmission. However, the transmission of extracted features introduces the risk of exposing sensitive personal data. To mitigate this issue, I have developed a novel privacy-preserving collaborative inference mechanism. Under this mechanism, each edge device in the network protects the privacy of extracted features before transmitting them to a central server for inference. This method is designed to achieve two main objectives while ensuring effective inference performance: 1) reducing communication overhead, and 2) maintaining strict privacy guarantees during feature transmission.

A significant gap in the current literature on distributed learning and inference lies in its limited consideration of the interplay between **wireless data acquisition**, latency, and privacy. As data dimensionality grows, wireless acquisition becomes a critical bottleneck, particularly for **fast** edge learning, where high mobility and device unreliability exacerbate the challenges. These issues highlight the pressing need for advanced radio resource management techniques specifically tailored for edge learning. To address these challenges, I aim to leverage a novel metric, Age-of-Information **(AoI)**, to design adaptive, **real-time** data acquisition schemes that not only optimize performance but also enhance privacy preservation.

Thrust 3: Privacy-preserving and Computationally Efficient Graph Mining Algorithms. Adopting privacy-preserving solutions in such applications is often hindered by the lack of scalability to large-scale problems with billions of nodes and edges. Building on my previous pioneering work, I proposed integrating graph sketching to further improve the computational complexity of semidefinite programming in community detection. Specifically, in the context of binary symmetric SBMs, I have examined the impact of graph sketching on separation conditions and proposed efficient algorithms that achieve the derived information-theoretic necessary conditions for exact community recovery. In my current research, I have focused on the challenges of preserving privacy in graph datasets, primarily through community detection methods. Utilizing differential privacy measures at both edge and node levels, I have provided actionable insights applicable to various domains, such as social networks, healthcare, government administrative data, and mobility studies. Recognizing the broader relevance of these privacy-preserving techniques, I see the need for solutions that extend beyond community detection to address other critical challenges in graph analytics, such as link prediction, node classification, and influence maximization, where safeguarding sensitive information is equally vital.

Building on this foundation, my next goal is to expand my research by incorporating differential privacy principles into various aspects of graph analytics, particularly in the context of online data analysis. The online setting offers valuable opportunities for applications such as health monitoring and epidemiological policy-making. Yet, it also presents two critical challenges: the demand for computationally efficient algorithms and the capacity to manage complex, uncertain, and potentially corrupted data distributions.

Thrust 4: Hardware-based Privacy Mechanisms for Ultra-Low Power IoT Systems. The growing complexity and widespread adoption of Internet of Things (IoT) devices have introduced significant challenges to network infrastructures, particularly in terms of energy consumption, latency, and privacy risks. These issues are especially pronounced in IoT applications like smart home automation, industrial IoT, and healthcare monitoring, where the integration of deep learning and AI technologies demands extensive computational resources, leading to increased latency and privacy vulnerabilities.

One critical area of concern is **privacy threats** within wireless IoT systems. The distributed nature of these devices and the vast amount of data they handle raises the potential for unauthorized access and misuse of personal information. My research focuses on developing robust privacy-preserving techniques for **ultra-low power IoT devices**, leveraging DP mechanisms to address these concerns. Existing **software-based** DP mechanisms, such as Laplacian and Gaussian noise generation, are **computationally intensive** and unsuitable for ultra-low power IoT devices due to the continuous sampling procedures involved. For example, the noise generation process can introduce significant delays and increased energy consumption. Additionally, these mechanisms can be vulnerable to side-channel attacks, where attackers might infer details about the privacy mechanisms based on varying power consumption patterns.

I aim to design computationally efficient, **lightweight** noise generation techniques tailored to the constrained processing power and energy limitations of IoT devices. Traditional privacy mechanisms, such as Laplacian and Gaussian distributions, are not optimal for these devices due to their computational demands and energy consumption. Instead, I explore alternatives that are better suited to the specific constraints of IoT environments. Additionally, I propose leveraging the **inherent** noise generated by IoT devices as a privacy-preserving mechanism, allowing for more efficient privacy protection without extensive resource requirements. My work aims to improve privacy in IoT systems while minimizing computational overhead, energy consumption, and latency, ensuring secure data processing in increasingly distributed and resource-constrained environments.

Thrust 5: Towards Quantum-Enabled Wireless Intelligence. Looking ahead, I am interested in exploring the intersection of wireless communication and quantum information systems, particularly in hybrid settings where quantum information must be transmitted, processed, or inferred over classical networks. Recent work has highlighted the fundamental challenges of communicating quantum states due to their exponential complexity, and proposed efficient alternatives based on tools such as shadow tomography and unequal error protection [14]. This perspective resonates strongly with my background in information theory and wireless systems, and opens up new directions for designing communication-efficient and robust protocols for quantum-enhanced networks, including applications in distributed sensing, inference, and next-generation (6G) systems. I am particularly interested in developing statistically grounded and resource-efficient methods that bridge classical and quantum communication, enabling scalable and reliable learning and inference in hybrid quantum-classical environments.

References (a complete list of my publications is available on my Google Scholar profile):

[1] Y. Eldar, A. Goldsmith, D. Gündüz and H. V. Poor, *Machine Learning and Wireless Communications*. Cambridge University Press, 2022.

[2] M. Seif, R. Tandon, and M. Li, "Wireless federated learning with local differential privacy," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT),* 2020, pp. 2604–2609.

[3] M. Seif, W.-T. Chang, and R. Tandon, "Privacy amplification for federated learning via user sampling and wireless aggregation," *IEEE Journal of Selected Areas in Communications,* 2021.

[4] M. Bertran, N. Martinez, A. Papadaki, Q. Qiu, M. Rodrigues, G. Reeves, and G. Sapiro, "Adversarially learned representations for information obfuscation and inference," in *Proceedings of the International Conference on Machine Learning (ICML*), 2019, pp. 614–623.

[5] M. Seif, D. Nguyen, A. Vullikanti, and R. Tandon, "Differentially private community detection for stochastic block models," *Proceedings of the International Conference on Machine Learning (ICML),* 2022.

[6] M. Seif, R. Tandon, and M. Li, "On the secure degrees of freedom of 2 x 2 x 2 multi-hop network with untrusted relays," in *Proceedings of the IEEE International Conference on Communications (ICC),* 2018, pp. 1–7.

[7] M. Seif, R. Tandon, and M. Li, "Secure retrospective interference alignment," *Journal of Entropy*, vol. 21, no. 11, p. 1092, 2019.

[8] B. Jiang, M. Seif, R. Tandon, and M. Li, "Answering count queries for genomic data with perfect privacy," *IEEE Transactions on Information Forensics and Security*, 2023.

[9] M. Seif, A. J. Goldsmith, and H. V. Poor, "Differentially private community detection over stochastic block models with graph sketching," in *Proceedings of the 57th Annual Conference on Information Sciences and Systems (CISS),* 2023, pp. 1–6.

[10] R. Saha, M. Seif, M. Yemini, A. J. Goldsmith, and H. Vincent Poor, "Collaborative mean estimation over intermittently connected networks with peer-to-peer privacy," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT),* 2023, pp. 174–179.

[11] B. Jiang*, M. Seif *, R. Tandon, and M. Li, "Context-Aware Local Information Privacy", IEEE Transactions on Information Forensics & Security, vol. 16: 3694-3708, June 2021 (*:co-first authors, alphabetical order)

[12] M. Seif, L. Xie, A. J. Goldsmith, and H. V. Poor, "Differentially Private Online Community Detection for Censored Block Models: Algorithms and Fundamental Limits", IEEE Transactions on Information Forensics & Security, July 2025

[13] M. Naseri, P. Ashtari, M. Seif, E. De Poorter, H. V. Poor, and A. Shahid, "Deep learning-based image compression for wireless communications: Impacts on reliability, throughput, and latency," *npj Wireless Technology*, 2025.

[14] L. Hanzo, Z. Babar, Z. Cai, D. Chandra, I. B. Djordjevic, B. Koczor, S. X. Ng, M. Razavi, and O. Simeone, "Quantum information processing, sensing, and communications: Their myths, realities, and futures," Proceedings of the IEEE, 2025.