Research Statement

Abdulrezzak Zekiye

My research lies at the intersection of **distributed systems, blockchain technologies, and machine learning**, with an overarching goal of designing **decentralized, secure, and fairness-aware systems** for large-scale, real-world applications. I am particularly interested in how decentralized architectures and privacy-preserving AI can be used to support critical infrastructures, such as financial systems, energy networks, healthcare platforms, and intelligent transportation, where trust, transparency, and scalability are essential.

Current and Past Research

My current research focuses on **decentralized and AI-enhanced systems for the Internet of Energy (IoE)**. In my doctoral work, I proposed a blockchain-based, dual-layer architecture that combines **off-chain computation** with **on-chain verification** to address the scalability and cost limitations of fully on-chain systems, enabling complex decision-making such as fairness-aware resource allocation and secure coordination in decentralized environments .

A central contribution of this work is the development of **hybrid energy donation and bartering algorithms** that integrate multiple energy sources, including utility grids, prosumers, and peer-to-peer donations, governed via **decentralized autonomous organizations (DAOs)**. To address equity concerns, I proposed a **fairness-aware allocation strategy** that account for both historical access to resources and current needs, validated through extensive simulations and real-world datasets.

In parallel, I have explored **on-chain federated learning** models that enable distributed training of machine learning models without sharing raw data, addressing challenges of **trust, incentive compatibility, and accountability** in decentralized AI systems .

In addition to my core IoE research, I have conducted research in **cryptocurrency analysis and machine learning for financial systems**. Specifically, part of my master's studies focused on data-driven models to analyze blockchain parameters and identify risk patterns in cryptocurrencies, demonstrating the value of machine learning in understanding complex decentralized financial phenomena .

My research also includes **AI and health-care applications**. For example, I contributed to *Decentralized Healthcare Systems with Federated Learning and Blockchain*, which reviews the integration of blockchain and federated learning to enable privacy-preserving collaboration across distributed healthcare data sources. This work highlights the potential to overcome data silos and privacy barriers in medical data analytics by combining decentralized trust with collaborative learning. Moreover, earlier publications such as *Classification of Medical Transcriptions with Explanations* and *Trusting Heart Failure Predictors* reflect my long-standing

interest in **interpretable machine learning for healthcare**, focusing on building models that not only achieve high performance but also provide explanations suitable for clinical contexts .

Future Research Agenda

My future research agenda builds on my work in decentralized learning, privacy-preserving AI, and trustworthy intelligent systems. In the near term, I will focus on **federated and privacy-preserving LLMs for AI safety**, then extend this work to broader decentralized AI systems and high-stakes application domains. Across these directions, my goal is to design AI systems that are secure, privacy-aware, robust, accountable, and deployable in real-world environments where data, authority, and risk are distributed across multiple stakeholders.

1. Federated and Privacy-Preserving LLMs for AI Safety

My near-term research will focus on privacy-preserving and safety-oriented large language model systems for high-stakes domains where sensitive data cannot be centralized, such as healthcare, finance, education, and critical infrastructure. I aim to develop federated and decentralized methods that allow institutions to collaboratively adapt, evaluate, and improve LLMs while preserving data confidentiality, model integrity, and regulatory compliance.

This direction includes federated instruction tuning, parameter-efficient fine-tuning, secure aggregation, differential privacy, and cryptographic techniques for privacy-preserving LLM training, inference, and evaluation. Beyond data protection, I will study the safety risks that arise when LLMs are deployed in decentralized and agentic environments, including memorization, privacy leakage, prompt injection, hallucination, poisoning attacks, backdoor vulnerabilities, gradient inversion, and high-severity misuse.

I am particularly interested in developing empirical safety evaluations, benchmarks, and oversight mechanisms for federated and multi-agent LLM systems. By addressing the trade-offs among privacy, utility, fairness, robustness, and communication efficiency, my goal is to build LLM systems that are trustworthy by design: privacy-aware, robust, accountable, and aligned with human and societal values.

2. Security, Privacy, and Trustworthiness of Decentralized AI Systems

A second direction of my research will investigate the security and privacy foundations of decentralized AI systems. As AI models are increasingly trained, adapted, and deployed across distributed institutions, edge devices, and autonomous agents, they become vulnerable to data leakage, inference attacks, poisoning, adversarial updates, model manipulation, insecure coordination, and failures in decentralized governance.

To address these challenges, I plan to integrate cryptographic techniques, blockchain-based verification, privacy-enhancing technologies, and robust machine learning methods into decentralized AI pipelines. A key goal is to develop mechanisms that make distributed AI systems not only accurate, but also auditable, resilient, and trustworthy under adversarial and

uncertain conditions. This work will also examine how technical safeguards can be combined with governance, evaluation, and accountability mechanisms so that decentralized AI systems can be deployed responsibly in real-world settings.

3. Decentralized AI for Critical Infrastructure

I will also apply these methods to critical infrastructure domains such as healthcare, energy management, and connected vehicle ecosystems. These settings often involve distributed data sources, heterogeneous stakeholders, privacy constraints, safety-critical decisions, and strict regulatory requirements. My work will explore how decentralized learning, distributed intelligence, and secure data-sharing mechanisms can improve resilience, efficiency, fairness, and accountability in automated decision-making.

In healthcare, this may involve privacy-preserving clinical analytics and LLM-assisted decision support across institutional boundaries. In energy systems, decentralized AI can support adaptive resource management, demand forecasting, and resilient grid operation. In connected vehicle ecosystems, distributed intelligence can enhance safety, coordination, and trust among vehicles, infrastructure, and users. Across these applications, my goal is to design AI systems that remain reliable under real-world constraints, including data heterogeneity, limited communication, adversarial behavior, and institutional fragmentation.

Long-Term Vision

My long-term vision is to establish a research program on **decentralized, privacy-aware, and socially responsible AI systems**, with a particular emphasis on federated learning, privacy-preserving LLMs, AI safety, and trustworthy digital infrastructures. I aim to connect technical innovation with societal values such as fairness, transparency, accountability, inclusivity, and human-centered design.

Through interdisciplinary collaboration and engagement with real-world stakeholders, I seek to produce research that is both theoretically rigorous and practically impactful, while mentoring the next generation of researchers and engineers working at the intersection of AI, security, privacy, and decentralized systems.