Research Statement and Five-Year Research Plan

Jeffrey Alan Young

Research Vision

My research focuses on the design, analysis, and deployment of trustworthy artificial intelligence systems, with emphasis on machine learning security, voice assistant and Internet-of-Things (IoT) ecosystems, and predictive analytics for complex cyber-physical environments. The overarching goal of my work is to develop scalable, data-driven methodologies that enable intelligent systems to operate safely, transparently, and robustly in real-world settings.

Modern AI-driven platforms increasingly mediate critical aspects of daily life, yet they remain vulnerable to misuse, policy violations, and emergent security risks. My research addresses these challenges by combining machine learning, data mining, and systems security to identify hidden behaviors, predict failures, and enforce policy compliance at scale. This work is inherently interdisciplinary and draws from computer security, artificial intelligence, data science, and applied systems engineering.

Current and Prior Research

My prior research has established a strong foundation in AI security and large-scale data analysis. I am a co-author of peer-reviewed publications in top-tier venues such as USENIX Security and ACM CCS, where my work introduced automated methods for detecting policy violations in voice assistant platforms and empirically evaluating trust mechanisms in AI ecosystems. These studies leveraged natural language processing, deep learning, and large-scale empirical analysis to expose systemic weaknesses in widely deployed platforms.

In parallel, I have conducted extensive research in predictive analytics and time-series modeling, including recurrent neural network frameworks for financial and industrial systems. This work has demonstrated how data-driven models can capture complex temporal dependencies and provide actionable insights for risk prediction and decision support. Collectively, my research demonstrates a sustained trajectory of scholarly productivity, methodological rigor, and real-world impact.

Five-Year Research Agenda

My five-year research plan is structured around three interrelated thrusts that build on my prior work while enabling new directions suitable for collaborative, externally funded research.

Thrust I: Trustworthy AI and Machine Learning Security

The first thrust focuses on developing principled methods for securing AI-driven systems against misuse, adversarial behavior, and policy violations. Planned research includes automated compliance auditing for AI agents, robust learning methods for untrusted data sources, and explainable security mechanisms for machine learning models. This work will extend my prior contributions

on voice assistant security to broader classes of AI-enabled applications, including conversational agents and autonomous services.

Expected outcomes: peer-reviewed publications in security and AI venues, open-source tools for compliance analysis, and NSF Secure and Trustworthy Cyberspace (SaTC) proposals.

Thrust II: Intelligent IoT and Cyber-Physical System Analytics

The second thrust investigates scalable analytics for IoT and cyber-physical systems, emphasizing anomaly detection, device profiling, and predictive maintenance. Building on my previous work in IoT security and industry experience deploying production ML systems, this research will explore hybrid models combining statistical learning, deep neural networks, and graph-based representations to capture system behavior across heterogeneous devices.

Expected outcomes: collaborative grants with industry partners, interdisciplinary publications, and applied research projects suitable for undergraduate and graduate student involvement.

Thrust III: Data-Driven Risk Modeling and Decision Support

The third thrust focuses on predictive modeling and decision support in complex environments, including financial systems, infrastructure monitoring, and large-scale networks. Research will emphasize interpretable machine learning, temporal modeling, and uncertainty quantification to support responsible decision-making. This thrust provides a natural bridge between theoretical advances and applied deployments in both academic and industry contexts.

Expected outcomes: publications in data science and applied AI venues, externally funded projects, and integration into advanced data science curricula.

Timeline and Milestones

Years 1–2: Establish research group, mentor undergraduate researchers, publish extensions of prior work, submit initial NSF and industry grant proposals.

Years 3–4: Expand interdisciplinary collaborations, supervise graduate theses, secure external funding, and produce sustained publication output.

Year 5: Consolidate research programs into mature, externally funded initiatives; graduate student completions; and leadership in collaborative research projects.

Student Involvement and Mentoring

Student engagement is integral to my research program. Undergraduate and graduate students will participate in all phases of research, including data collection, model development, experimentation, and dissemination. My mentoring philosophy emphasizes early research exposure, ethical responsibility, and professional development, preparing students for both graduate study and industry careers.

Broader Impact and Institutional Fit

My research agenda aligns well with Georgia Southern University's emphasis on applied scholarship, student success, and community impact. By integrating research with teaching and service, my work will contribute to a vibrant research culture while providing meaningful educational experiences for students. I am committed to building a sustainable, inclusive, and externally visible research program that enhances the Department of Computer Science and the University as a whole.