2026 IMPACT CIRCLE

Project Title: OASIS (Operational AI Safety and Inference System): Engineering Safe AI for Biological Simulation

Investigator(s) and collaborations:
James Yurkovich, PhD, Emma Glass, Research Scientist

Unmet Need/Primary Question:
Through our existing DARPA award, we are constructing an AI-driven simulator of E. coli capable of designing genetic variants—Project SIMBA. AI technologies could design new organisms that produce biomolecules of interest or help combat emerging antimicrobial resistance threats. However, we are faced with an uncomfortable truth: these AI tools could accidentally design the very threats we seek to combat. AI technologies are developing so rapidly that we do not yet have safeguards. Here, we propose to design safeguards for emerging generative AI technologies for biological design. Just as a modern car automatically brakes if it detects a crash is imminent, our safety system will monitor computer simulations of biological cells. If the software detects that a user is designing a biological agent that is becoming dangerous or unstable—whether accidentally or on purpose—our system will automatically detect the threat and stop the simulation before the design can be finalized or sent to a lab to be physically built.

Novel Hypothesis:
We hypothesize that AI safety can be integrated into AI tools by replacing rigid historical databases and lookup tables (the current state of the art) with safety requirements that follow strict engineering process flows. We will apply Functional Safety—the same rigorous systems engineering logic used to keep nuclear power plants and airplanes safe—directly to biological code. Unlike standard antivirus software that scans for known viruses, we will design a knowledge curation platform that can be used to police the behavior of a biological simulation and is fully integrated with AI technologies.

Project Proposal:
Today, biological safety relies almost entirely on looking backward. We verify new genetic designs by checking them against static databases of known, historical diseases. However, as generative AI becomes powerful enough to design entirely new, complex biology, these old checklists are no longer sufficient. Because AI can generate biological structures we've never seen before, we are left completely vulnerable to “Unknown Unknowns”—novel, emergent biological failures that aren’t on any warning list because they haven’t existed until now. To

safely harness the immense power of AI for biology and medicine, we urgently need a fundamental shift in how we approach biosecurity. Early results from SIMBA, our DARPA-funded AI-driven platform for simulating complex microbial behaviors, show this powerful simulation architecture is extensible to human cellular models, an advancement that OASIS for AI Safety — Yurkovich Lab will open doors for aging research and precision medicine. However, there lies a major barrier to the successful translation and commercialization of this technology: the lack of an automated safety system. Designing a safety system is outside the scope of our DARPA contract, meaning that the future of this research direction is limited.

Here, we propose to design OASIS, a proactive "digital safety brake" for AI biological design. We will adapt the same rigorous systems engineering principles used to prevent catastrophic failures in nuclear power plants and commercial aerospace, translating ratified industrial safety standards into biological code in the form of a knowledge graph. We will not be starting from scratch; we will build upon a foundation of well-defined, ratified standards: IEC 61508 for functional safety, MIL-STD-882E for system safety, and the emerging ISO/IEC TS 22440 for AI safety. While the FDA provides oversight for health applications, their current framework is not well-suited for simulation-based verification. With this funding, we will build the foundation for this safety system.

The basis for many generative AI simulators for biological design leverage Knowledge Graphs—a mathematical representation of a network. Traditionally, encoding biological knowledge into a graph is challenging because of our incomplete knowledge of biology. We propose to take a fundamentally different approach: instead of limiting ourselves to only encoding specific observations (e.g., "Mutation X on Gene A results in hazardous behavior Z"), we describe more general mechanisms (e.g., "Any secreted protein that self-assembles into a ring structure on a lipid bilayer is a hazard"). This framework allows our safety system to handle *Unknown Unknowns*. Specifically, we will:

- Map existing standards to our SIMBA simulation framework. This process involves applying classical systems engineering principles for Functional Safety Systems to a biological process. A fairly straightforward process, this is a high-impact / low-risk proposition.
- Identify important mechanisms of the bacterial design space. We will curate publicly available databases in the literature to collect annotated mechanisms for pathogenic bacteria, as well as known interventions (i.e., antimicrobial compounds like penicillin).
- Encode these behaviors into a Knowledge Graph. Once we have our list of mechanisms, we will work to transform these "biological facts" into the graph. We do this using natural language to allow for large language models like ChatGPT to search this space and develop new insights.

At the end of OASIS, we will have a prototype demonstrating a framework for AI safety in biological simulation that will drive future funding opportunities. Further, we will have a manually curated graph database of biological mechanisms that will be used to begin implementation of the full system—a requirement for commercial translation of these ideas.

Description of Potential Impact:

By funding OASIS, you are not just supporting aging research—you are stepping forward as pioneers in the critical, emerging field of safe AI design. This investment builds the essential digital guardrails required to safely transfer our highly successful single-cell simulations into predictive models for precision medicine. Expanding this technology will revolutionize personalized medicine, accelerating the discovery of tailored therapeutics that could profoundly impact your own health, wellness, and lifespan. Furthermore, successfully demonstrating this first-of-its-kind safety framework will position the Buck as the undisputed leader in safe and ethical biological AI, making us highly competitive for massive, dedicated federal funding to scale this technology and transform healthcare.