Capiti: Improving Biosecurity Screens With Edge AI

Paper

Mackenzie Noon[1], Grace Oualline[1]

[1]Yale School of Medicine, Genetics Department

Video Demonstration

First Place in AIxBio Hackathon
Track 4: Benchtop Synthesizer Security Winner

1. Current DNA Screening Tools struggle to catch AI generated pathogens

Benchtop DNA synthesizers currently produce sequences up to 300bp and are projected to reach 7,000bp within 2–5 years, sufficient to reconstruct the smallest viral genomes (Service, 2023). Unlike centralized synthesis services, benchtop devices bypass existing screening chokepoints and are largely unmonitored (Laird et al., 2025). **Current screening platforms such as SecureDNA and IBBIS commec rely on homology matching, which has three main structural vulnerabilities:**

- AI-designed protein variants preserve biological function while evading homology-based detection. Such variants are accessible via open-source tools like ProteinMPNN (Wittmann et al., 2025)
- Last-resort antimicrobial resistance (AMR) genes such as mcr-1, blaNDM, and blaKPC are absent from major screening frameworks despite representing a credible threat. AMR is projected to cause 39 million deaths by 2050 (WHO, 2024; University of Oxford, 2024)
- Existing tools require internet connectivity, introducing latency and expanding the attack surface

We present Capiti, an edge-AI biosecurity system that intercepts valve control signals on benchtop synthesizers and identifies pathogenic function in real time, fully offline. **We demonstrate that Capiti detects AI-redesigned protein variants that evade IBBIS commec and SeqScreen,** and introduce Capiti-C, extending screening coverage to last-resort AMR genes.

2. Graphical Abstract

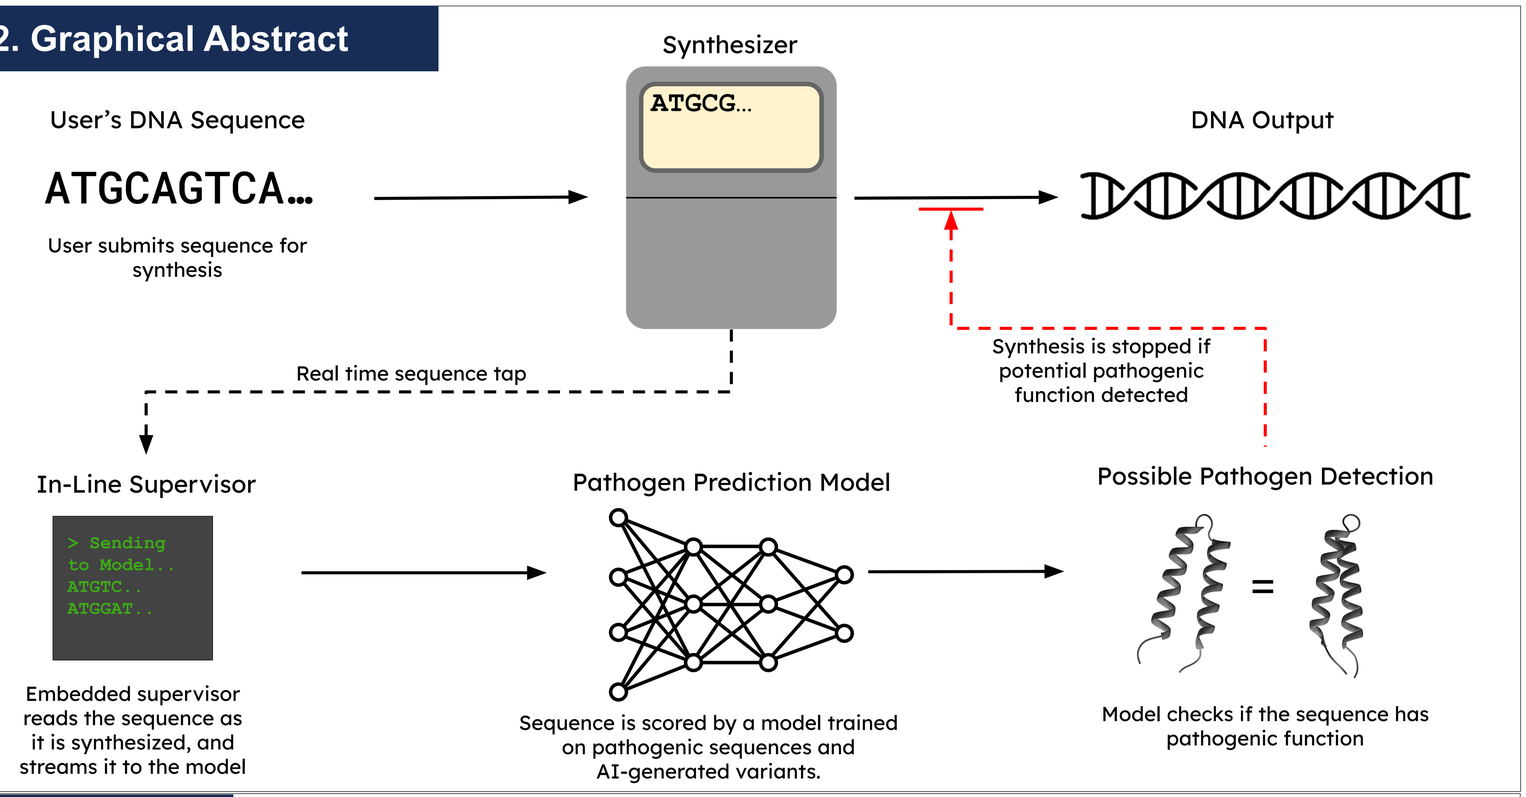

3. Capiti learns pathogenic function and runs on realistic hardware

1. Sequence variant generation

We assembled a dataset of proteins with known pathogenic function. Key functional residues were fixed and **variants generated using ProteinMPNN** (Dauparas et al., 2022), forming the positive set. The negative set comprised six categories:

- Active site knockouts via alanine scanning and combined knockout
- Sequence scrambling and 30% perturbation
- Randomly generated decoy sequences
- Nonpathogenic sequences from the same protein family as positive examples

Protein truncation was applied during test set generation, training the model to recognize pathogenic function from incomplete fragments and reflecting realistic synthesis interception conditions.

OXA-48 (confers resistance to carbapenems): **Wild type (yellow)**, & 15x in-silico **generated variants (red)**. Variant structures predicted with ESMfold, visualized with ChimeraX.

2. Hardware

The test rig consists of an Arduino UNO emulating a DNA synthesizer (Emusynth) and a Raspberry Pi 4 running Capiti, connected via a five-wire nucleotide bus and two control lines:

- Four wires carry nucleotide valve signals; one carries the Tetrazole activator signal
- An end-of-synthesis line pulses at run completion
- **An interrupt line pulses when Capiti makes a high-confidence pathogenic sequence call, aborting synthesis**

Emusynth emulates Applied Biosystems 3400 DNA Synthesizer and generalizes to most other models (Applied Biosystems, 2003).

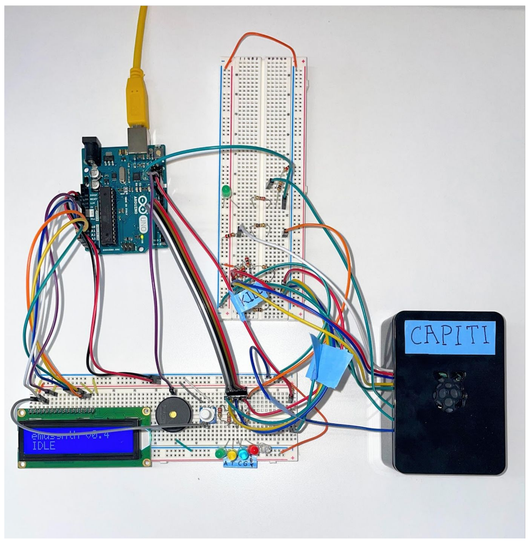

3. Model training

Capiti's architecture is designed to learn pathogenic function while remaining light enough for embedded deployment. The model is a simple CNN with two key features:

- Dilated convolutions capture long-range sequence features without expensive attention heads
- Residual connections enable block specialization and mitigate vanishing gradients (Bai et al., 2018)

Separate models were trained for pathogenic function detection (Capiti-E) and last-resort AMR genes (Capiti-C) on an NVIDIA RTX 5000 Ada Generation GPU.

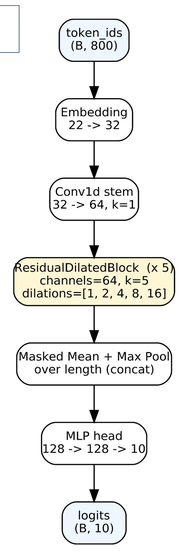

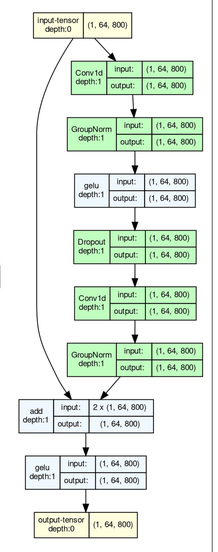

4. Capiti Detects AI redesigned protein variants that evade current screening platforms

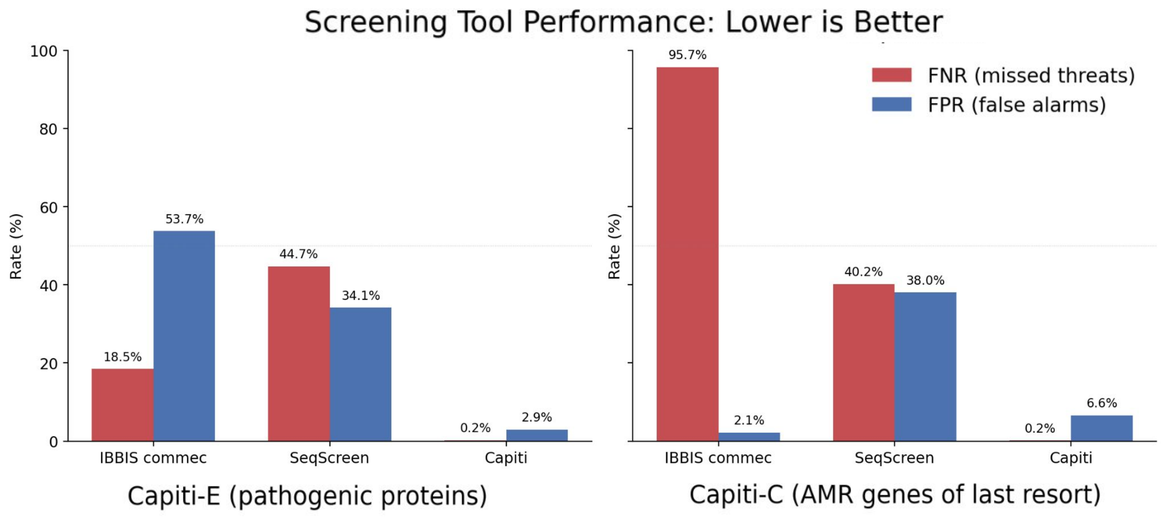

To evaluate existing tools, test set proteins were reverse-translated to DNA using S. aureus codon frequencies (Kazusa, 2025), then run through SeqScreen (fast mode) and IBBIS commec (light biorisk). False negative rates (**FNR, functional variants passed undetected)** and false positive rates (**FPR, negative controls incorrectly flagged**) are measured.

- **IBBIS commec missed 18% of pathogenic variants and 96% of AMR variants,** consistent with AMR genes being absent from its training set
- **SeqScreen missed 44.69% of pathogen variants and 40.17% of AMR variants** despite incorporating functional annotations and FunSoC-based pathogenicity predictions
- **Capiti missed only 0.24% of pathogenic variants and 0.19% of AMR variants,** with FPRs of 2.93% and 6.60% respectively

Both tools showed moderately high FPRs, predominantly on active site knockouts and perturbed variants, suggesting reliance on homology rather than function. Capiti detects variants missed by existing tools, and achieved substantially lower error rates than both existing tools across both threat categories.

5. Conclusions

- **Homology-based screening has a structural vulnerability.** AI-redesigned sequences preserve function while evading sequence-level detection
- **IBBIS Commec and SeqScreen** improve on best-hit approaches but **remain bounded by fixed training distributions** and category sets (Arango-Argoty et al., 2018; Balaji et al., 2022)
- **Capiti's function-based approach is robust to obfuscation** because biological function is constrained by biochemistry regardless of sequence identity

6. Future Directions

Capiti's current architecture addresses remote and network-based threat models. Several extensions would substantially broaden its coverage:

- **Tamper-proofing measures** and hardware attestation to address physical access threat models
- **Cryptographic logging** for post-hoc forensic attribution of synthesis events
- Validation on a live synthesizer beyond the Emusynth emulation environment
- **Minimal-surface heartbeat webhook architecture** for network reporting without full phone-home dependency