College of Engineering and Computer Science
Office of the Dean
777 Glades Road, EE96, Room 308
Boca Raton, FL 33431
561.297.3400

Announces the Ph.D. Dissertation Defense of

Tiffany A. Ceasor

for the degree of Doctor of Philosophy (Ph.D.)

Neurosymbolic Behavioral Contract Verification: Formal Decision Traces and Constraint Patterns for Mission-Critical AI Assurance

November 13, 12pm-1:30pm
EE 96, Room 405
777 Glades Road
Boca Raton, FL

DEPARTMENT: Department of Electrical Engineering & Computer Science

ADVISOR: Dr. Borko Furht

PH.D. SUPERVISORY COMMITTEE:
Dr. Borko Furht, Chair
Dr. Fernando Koch
Dr. Xingquan Zhu
Dr. Nataliia Neshenko
Dr. Hari Kalva

LINK TO DISSERTATION DRAFT:
https://docs.google.com/document/d/1ERh_zYj8iuykOh1jW85QJjyMVm0Qke2u0IytxwBW9Ho/edit?usp=sharing

ABSTRACT OF DISSERTATION

Mission-critical artificial intelligence systems present a verification challenge that neither statistical testing nor purely formal methods can address alone. Statistical test and evaluation cannot prove that undesirable behaviors never occur; formal verification of arbitrary properties of large neural networks remains intractable; and probabilistic guardrails built on the same learned substrate as the systems they protect offer no deterministic guarantee. This dissertation develops a neurosymbolic verification architecture that places Satisfiability Modulo Theories (SMT) reasoning in a pre-execution gating position above probabilistic components, producing machine-checkable evidence that each released output satisfies a declared policy.

Three contributions are made. First, Formal Decision Traces (FDTs) are introduced as structured verification artifacts that bind constraint specifications, output encodings, proof certificates, and cryptographic attestations into a single machine-checkable record, validated across three safety-critical domains at millisecond-scale latency. Second, a family of behavioral contracts for autonomous agents is specified and empirically validated across 165 trials against a production large language model, revealing a structural asymmetry between attacks that neural safety training covers and attacks that only symbolic verification can cover. Third, the IRON COMPASS methodology

College of Engineering and Computer Science
Office of the Dean
777 Glades Road, EE96, Room 308
Boca Raton, FL 33431
561.297.3400

partitions mission-critical AI assurance into four tiers — formally verified, statistically bounded, empirically tested, and documented assumption — each bound to the strongest feasible verification method with an articulated independence argument.

An auditability layer for AI-enhanced positioning, navigation, and timing (PNT) demonstrates the approach in a representative cyber-physical setting, with measured SMT decision times of 1.2 to 5.6 ms across policy sizes up to 12 sources and 16 exclusion zones. A regulatory cross-walk positions the framework as accountability infrastructure for e-governance under the EU AI Act, the NIST AI Risk Management Framework, OMB M-25-21, and the Colorado AI Act. The compositional structure — rather than any single verification layer — carries the assurance argument, and provides a path to production deployment of AI components whose correctness must be demonstrable and replayable over decade-scale liability horizons.

BIOGRAPHICAL SKETCH

Tiffany A. Ceasor entered university at thirteen, earned her B.S. at seventeen and her M.S. at nineteen, and was recruited to Microsoft at nineteen, where she was selected for the Future of Microsoft intern cohort chosen by Bill Gates (two of three thousand). She spent five years at Microsoft as a Research Data Scientist, Machine Learning Engineer, and Full-Stack Engineer; her research there served as founding research for the Microsoft Africa Research Institute (MARI). At twenty-five she architected, won, and led a $200M+ homeland-security AI mission program with a ninety-person organization, an automated mission system leading to safer skies for all travelers; by twenty-seven she was a defense technology vice president, and she serves as Vice President of Growth and Technical Mission Strategy (Cyber & AI) at a national security, defense, and intelligence contractor. She is Founder and Principal Researcher of the Technological Society of Applied Research (TSAR), a 501(c)(3) based in Seattle focused on neurosymbolic AI, formal verification, AI auditing, and pre-execution behavioral assurance for mission-critical autonomous systems. Her research develops SMT-based behavioral contract verification, Formal Decision Traces, and constraint pattern frameworks for AI systems operating in defense, intelligence, and regulated civilian environments; the algorithms and frameworks she builds for deterministic, trustworthy AI are implemented at the highest levels of government, and she is a trusted AI advisor at senior levels of the U.S. national security and intelligence community. She currently serves as Program Head and Professor of Practice for AI (Intelligence and National Security) at The Academy, as a board advisor to the AI Task Force at Children's National Hospital in Washington, D.C., and on the board of the Seattle YMCA youth outdoor leadership program. An international keynote speaker (link to one of her talks youtu.be/h5qnDZoNh7A), she has led and executed numerous global service and disaster-relief programs since age thirteen across Africa, the Middle East, the Caribbean and Latin America, and Asia.

CONCERNING PERIOD OF PREPARATION & QUALIFYING EXAMINATION

Time in Preparation: Fall 2021 – Fall 2026

Qualifying Examination Passed: Spring 2022

Published/Accepted Papers:

Ceasor, T. A., & Furht, B. (2026). Compositional assurance for mission-critical AI: A four-tier methodology for IRON COMPASS test and evaluation. The ITEA Journal of Test and Evaluation (accepted). Companion presentation at DATAWorks 2026, Alexandria, VA.

Ceasor, T. A., & McNeil, S. (2026). Designing the Sentinel Policy Assistant: Building ethical, AI-enabled tools for national security education. (Published; available at https://www.academia.edu/168598094/)

Zia, L., Ceasor, T. A., & McNeil, S. (2026). Cybersecurity Data Privacy in the Age of Digital Communication (in press).