A testbed for developing & evaluating AI pilots

Jared Freeman[1], Eric Watz[1] and Winston Bennett[2]
[1]Aptima, Inc., Woburn, MA USA; [2]Winston "Wink" Bennett, AFRL / RHA, Dayton, OH USA
Contact: freeman@aptima.com

Abstract — Deficient data encumbers development, execution, evaluation, and revision of CGFs. We describe a testbed designed to provide data with the quality, volume, and variance needed to accelerate development of robust AI pilots that adapt smartly in training and operational applications.

1. Introduction

Computer Generated Forces (CGFs) are used in many training systems to supplement or replace the opposing forces that scarce and costly simulation staff control manually. CGFs thus make complex training more available and economical. However these benefits are often offset by the limited tactical repertoire of CGFs and their inability to respond realistically to trainee innovations and errors.

The root cause of this problem is not, we argue, the sophistication of agent architectures, which are many, reasonably varied, and capable of supporting adaptive behavior. The problem is that agents are built from and for impoverished data environments. Specifically, simulation testbeds typically fail to provide developers with data of sufficient quality, in volume, and variance. Automated performance measurements are also rare.

Data quality -- Simulations deliver raw physics data such as entity position and kinematics. Developers invest much of their time coding software functions that transform these data into tactically meaningful information such as flight formation. Accordingly, developers often have less time than they need to design, program, and test advanced, adaptive agent behaviors.

Data quantity -- Developers typically design agents from tactical documentation and expert advice. They rarely have sufficient flight data with which to learn tactical states and behaviors more rapidly, using statistical and machine learning techniques.

Data variance -- Developers test agents against very few, invariant scenarios. These agents fail unpredictably when trainee pilots stray outside the bounds of those scenarios. We want pilots to test the edge of the tactical envelope, to try and to fail; this is precisely when agents fail.

Measures -- Developers often evaluate agents by observing their performance in simulations; rarely do simulators issue measures of agent behavior that help developers automatically identify and efficiently diagnose agent deficiencies.

These data deficiencies have a profound effect. They slow the development, evaluation, and revision of AI agents for training and operational applications. The resulting agents are often less adaptive and robust than trainees and operational personnel require.

2. Approach

The Air Force Research Laboratory (AFRL) has created a unique testbed for developing and evaluating smart, agile agents. We call this AGENT, the Agent Generation & Evaluation Networked Testbed.

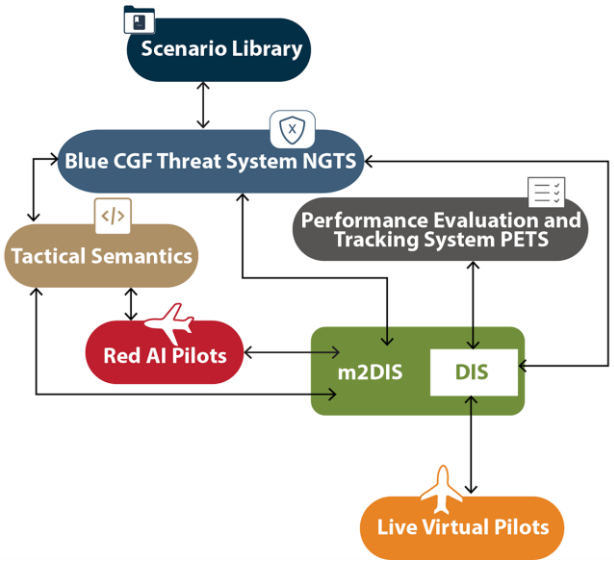

Fig. 1. High level architecture of AGENT.

AGENT provides agent developers with a secure and private environment for exercising their red AI pilots against blue CGF in scenarios defined in a shared library. Blue airframes, sensors, and weapons are modeled in the Next Generation Threat System (developed initially by the Air Force and now by the Naval Air Warfare Center Training Systems Division). AGENT publishes standard entity state and interaction data using the DIS protocol. It computes tactically meaningful information concerning the tactical situation over a custom Model to DIS (m2DIS) interface. Measures of performance and effects are computed automatically by the PETS Performance Evaluation and Tracking System. Users can observe scenario runs on the LNCS LVC Network Control Suite (not shown).

Several attributes of this testbed directly address the data deficiencies we enumerated above.

Data quality -- The testbed delivers data that span the range from raw data concerning entity state and interactions to semantically rich summaries of the tactical state. The TOA describes the adversary formation and location, much as an AWACS operator would do for pilots in flight. The FC-TAC [1] provides other data fundamental

to situation awareness, by responding to requests such as "Am I in the adversary's weapons engagement zone? Where is my wingman in relation to me?" This information stream potentially spares agent developers coding situational awareness functions, so that they can invest more time making agents adaptive and robust.

Data quantity -- The testbed enlarges the volume of data available to developers, because it runs scenarios in batches and at high speed. It stores the data from all runs by all agents by all developers in a common store available to every agent developer. This provides sufficient data for reliable statistical analysis, and for some data-hungry machine learning techniques.

Data variance -- Scenario runs in this testbed vary in ways that sample the large tactical space that trainees may eventually explore. Several features account for this. First, developers can parameterize batch runs so that each run employs a different weapon load, fuel load, and starting position. Second, developers' agents fight unusually intelligent, responsive CGFs; these are programmed as complex behavioral transition networks within the Next Generation Threat System. Third, the agents developed in this program are themselves highly adaptive. The net effect is that any two runs of a scenario can produce usefully varied tactical situations that evoke varied tactical behaviors. The data captured from many runs samples the large scenario space well. Thus, developers can identify regions of that space in which agents fail, and modify agents to be more robust and adaptive.

Measures -- Finally, measures of agent performance and effects are baked into the testbed, in the form of the PETS measurement system [2]. Agent developers can use these measures to rapidly identify the tactical settings in which agents succeed and fail. This, in turn, accelerates revision of agents.

Eight American agent development firms are developing agents using this testbed now, and experienced tactical aviators will assess the performance of those agents, using the testbed, in 2019 [3].

3. Future Directions

Design and development of this testbed has produced a number of recommendations concerning the data output requirements of future Air Force simulators, and of the black boxes on operational aircraft. The reader can infer some of those requirements -- for performance measures, for tactical semantics -- from the presentation above. We are consolidating those recommendations now into a data requirements document. Simulators and operational systems that address these data requirements will support rapid development of adaptive, robust agents for training and decision support.

Some of the most significant features of this testbed have only recently been completed. These include parameterized batch control of scenarios, automated performance measurement, responsive CGFs, and a shared data store. Only if agent developers use these features can we test whether they deliver the benefits we describe, above.

Acknowledgments

This presentation is based upon work supported by the United States Air Force Research Laboratory, Warfighter Readiness Research Division 711 HPW/RHA, under Contract FA8650-16-C-6698.

References

[1] E. Watz, E., M. Doyle. Fall Simulation Interoperability Workshop (SIW) (2014).
[2] E. Watz, L. Keck, L., B. Schreiber. Proceedings of the 2004 Spring Simulation Interoperability Workshop (2004).
[3] J. Freeman, E. Watz. Proc. Adaptive Instructional Systems (in press).

Author/Speaker Biographies

Jared Freeman, Ph.D., is Chief Scientist of Aptima, Inc. His research explores learning and problem solving by individuals and teams in real-world settings.

Eric Watz is Chief Engineer at Aptima, Inc. He leads technical research and development for a number of Air Force Research Laboratory initiatives that develop innovative capabilities to train operational personnel in domains that include tactical fast jet and remotely piloted aircraft.

Winston Bennett, Ph.D., is Readiness Product Line Lead, Airman Systems Directorate, Warfighter Readiness Research Division, 711 HPW/RHA, USAF. Dr. Bennett is an AFRL Fellow whose programs have created and deployed new methods for defining training requirements, and new technologies for training and measuring performance.