Assessing & Selecting AI Pilots for Tactical and Training Skill

2020 NATO
Modelling & Simulation Group (NMSG) Symposium
22-23 October 2020

Jared Freeman, Ph.D., & Eric Watz, Aptima
Winston Bennett, Ph.D., AFRL 711 HPW/RHA

This article 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. This article is cleared for public release as of 28 August 2020, 88ABW-2020-2719.

Overview

- The challenge
  - Tactical expertise is a function of practice, but…
  - Expert human adversaries are scarce
  - Many CGFs are relatively inexpert adversaries
    - Lowest value for the most accomplished pilots
    - Potential for negative transfer for all pilots
- AFRL solution strategy
  - Robust AI pilots
  - Testbed for accelerated AI development
  - Assessment of AI
  - Train with tactical AI

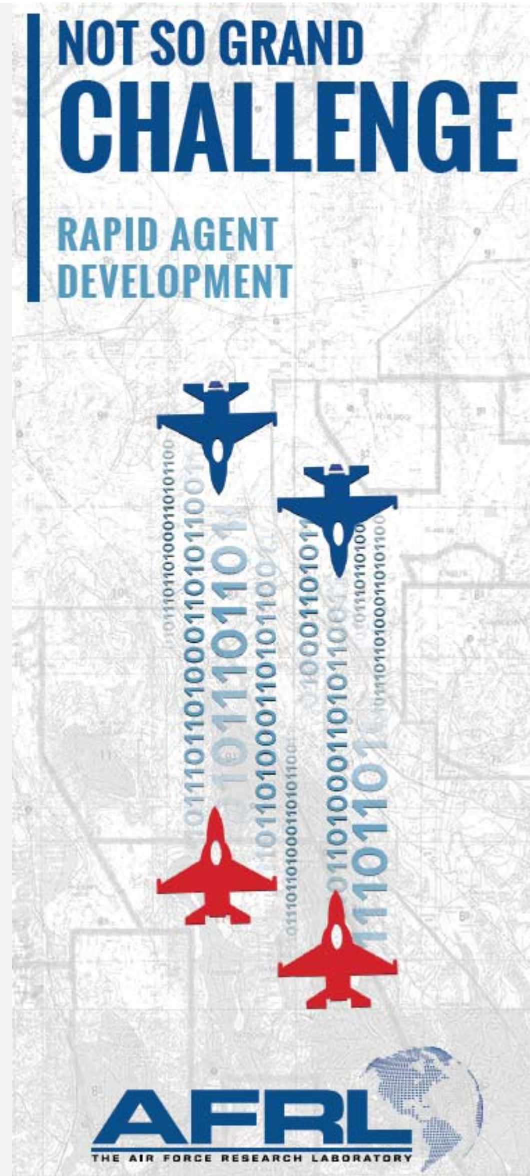

AI Pilots

- Challenge
  - Develop adversary (Red) AI that is tactically proficient & robust to trainee behavior in 1v1 & 2v2 engagements
- Solutions
  - TiER1 Performance Solutions
    - Task network model represents operator goals & functions. Accumulator model controls transitions through the task network.
  - Stottler Henke Associates
    - SimBionic architecture integrates multiple behavior transition networks. Dynamic scripting ML algorithm adapts agents.
  - Aptima
    - Reinforcement learning models tactical state and response. Behavior Definition Language (BDL) represents these plus goals, behavioral constraints, measures.

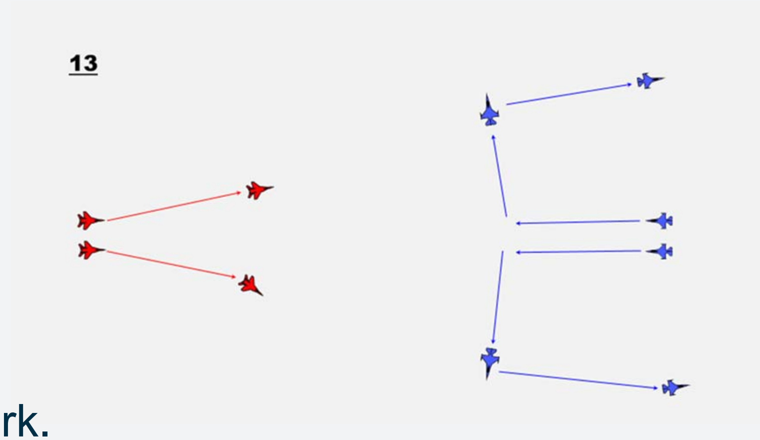

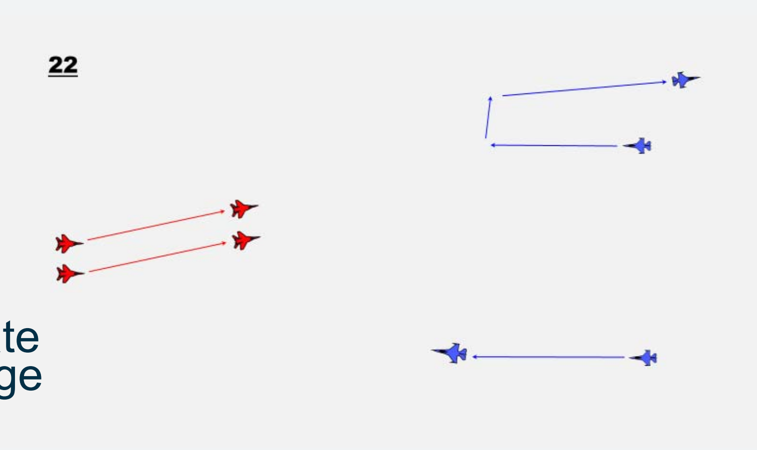

AI Pilots

- Solutions
  - Soar Technology
    - A production system -- the Soar-based cognitive architecture -- dynamically processes environmental states to accomplish deconflicted goals.
  - Eduworks
    - Brahms-Lite models work systems in which humans and technologies interact.
  - Discovery Machine
    - DMInd represents hierarchies of prespecified problem spaces and response strategies, which are retrieved as a function of fit to context.
  - CHI Systems
    - Personality-enabled Architecture for Cognition (PAC) uses narrative threads and personality characteristics (e.g., risk tolerance) to control perception and behavior.
  - Charles River Analytics
    - Hap manages competing goals and generates behaviors for dynamic situations, using rapid loops of information gathering, assessment, and decision-making.

Testbed

- Challenge: Rapid agent development
- Solution:
  - Testbed* enables independent, distributed exercise and assessment of AI pilots on 72 scenarios, with reduced support from SMEs

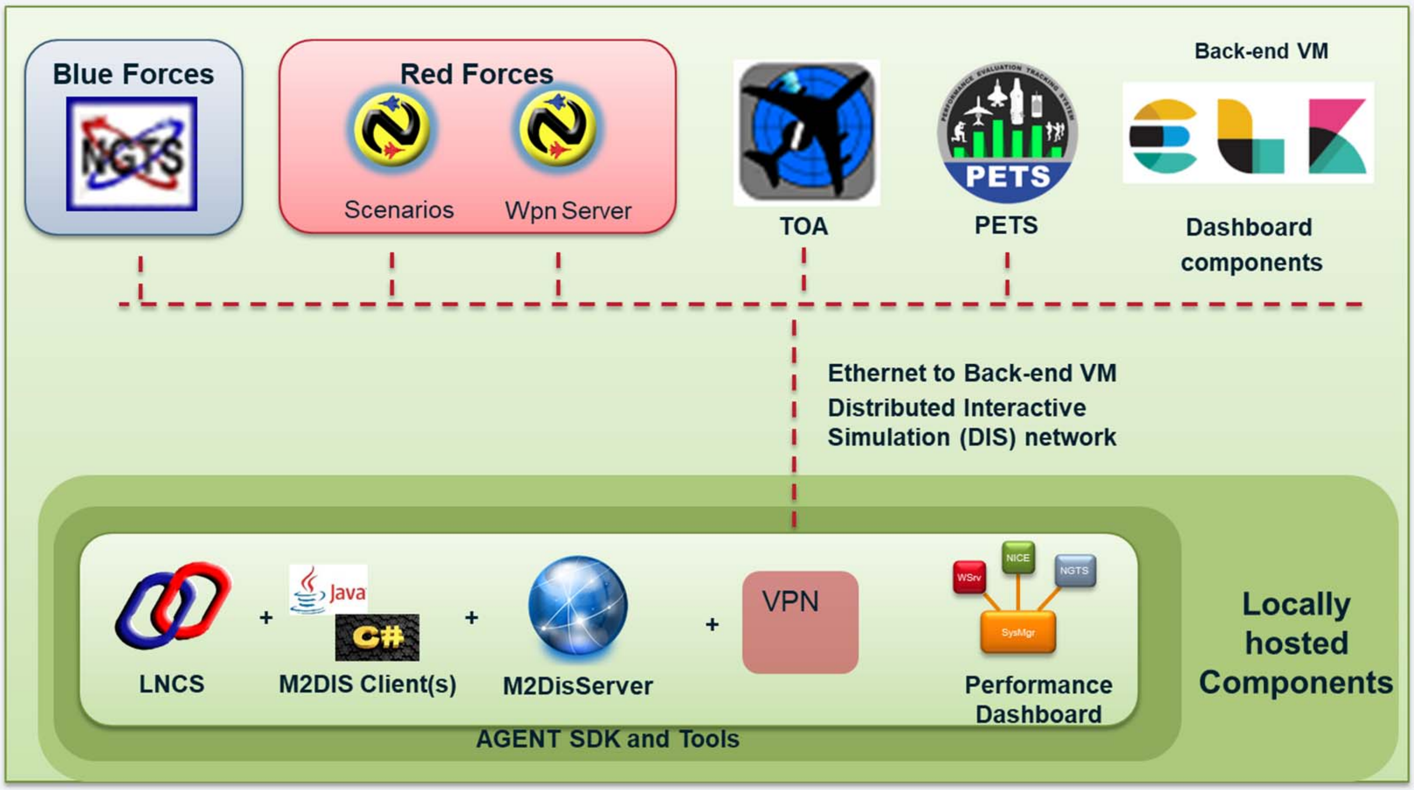

*Key components of the testbed are freely available to partner nations.

Testbed

- Generate constructive players
  - Red -- AFRL's Network Integrated Combat Environment (NICE), driven by AI
  - Blue -- Next Generation Threat System (NGTS)
- Capture data
  - Measure performance -- Performance Evaluation and Tracking System (PETS™)
- Record & play back
  - Simulation protocol -- Distributed Interactive Simulation (DIS) comms AI intent to distributed CGFs
- Tactical data
  - Syntactic -- Wrappers around native M&S data
  - Semantic
    - Tactical Observation Agent (TOA) simulates airborne control operator
    - Fighter Combat-Tactical Awareness Capability (FC-TAC) API relays overall state of and beliefs about the environment
- Data storage -- Data lake architecture

Assessment

- Method: Periodic assessment
  - Automated: Kills & Losses; deconfliction between Red Aircraft; aircraft location relative to adversary weapon lethality range; time in adversary weapon lethality range; Airspeed;
  - Expert judgment: Intercept geometry, adherence to contract, split decisions, spike awareness, post merge maneuvers, fuel & weapons management.
- Findings:
  - Process measures predict outcomes (right)
  - High variance on process and outcome measures within agents between scenarios, & between agents within scenarios

Measure	r(OutcomeScore,*)
OutcomeScore	
ContractAdherence	0.75
Deconfliction	0.21
ElementTargeting	0.86
FuelManagement	0.72
InterceptGeometry	0.71
PostMergeManeuver	0.64
SpikeAwareness	0.88
Split	0.39
TacticalIntelligence	0.90
WeaponsManagement	0.71

Scenario 13d

Average of Score Column

Row Labels	ContractAdherence	DecisionExplanation	Deconfliction	ElementTargeting	FuelManagement	InterceptGeometry	OutcomeScore	PostMergeManeuver	SpikeAwareness	Split	TacticalIntelligence	WeaponsManagement	Grand Total	S.D.
Agent1	1.0	1.0	2.5	1.0	1.5	1.0	1.0	1.0	1.0		1.0	1.0	1.2	0.44
Agent2	3.5	3.0	4.0	4.0	3.5	4.0	3.0	4.0			2.5	2.0	3.4	0.67
Agent3	2.0	3.0	2.5	3.0	3.5	2.5	4.0	3.0	4.0		2.5	3.0	3.0	0.60
Agent4	3.0	3.0	3.0	3.5	3.5	3.0	4.0	2.0	3.0	1.0	3.0	2.5	3.1	0.74
Agent5	1.0	1.0	3.5	1.0	2.0	1.5	1.0	1.0	1.0	1.0	1.0	1.5	1.5	0.71
Agent6	1.0	2.0	2.5	1.0	1.5	1.5	1.0	1.0	1.0	1.0	1.0	1.5	1.3	0.47
Agent7	4.0	2.0	4.0	4.0	3.0	4.0	4.0	3.0			3.0	3.5	3.6	0.65
S.D.	1.19	0.83	0.64	1.34	0.87	1.13	1.40	1.12	1.26	0.00	0.89	0.83		

Applying tactical AI to training

- Challenge
  - How can we best train human pilots with AI that are tactically proficient, not instructionally expert?
- Solution Strategy
  - Create an automated librarian that selects the best AI for the training scenario and trainee, based on a characterization of the AI

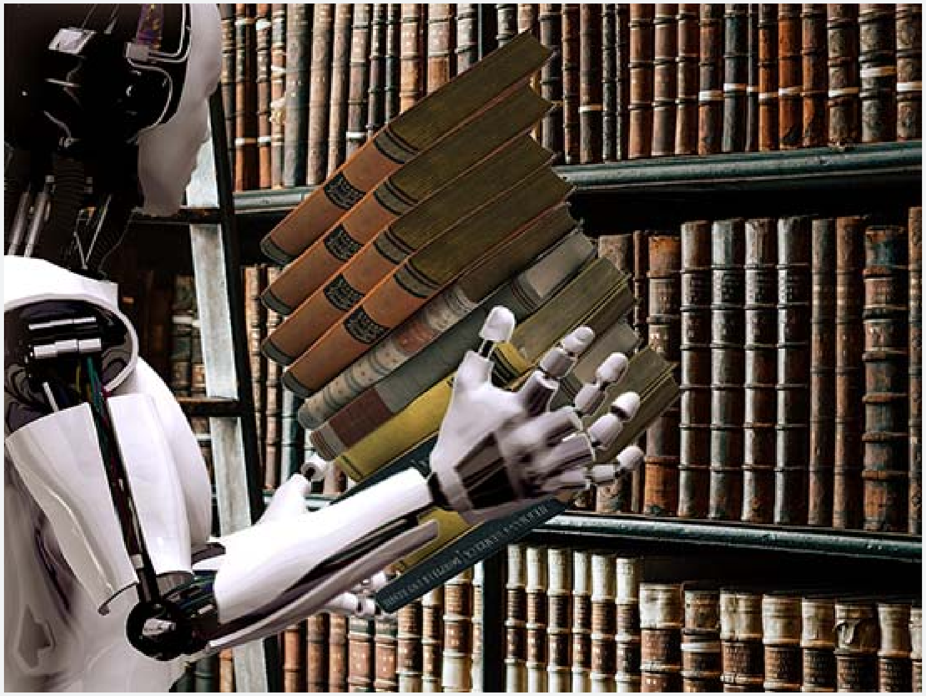

APTIMA®
Human-Centered Engineering®

Applying AI to Tactical Training

- Select by rule
  - Analyst computes the performance of all AI in all 72 scenarios
  - Trainer selects the next scenario in the curriculum
  - Librarian selects the AI that is (almost) the most tactically proficient
- Select by expert judgment
  - SME estimates the training effectiveness of each AI per scenario or vignette (within scenarios) for trainees at n levels of pilot expertise
  - Trainer estimates the expertise of the pilot
  - Librarian selects the most effective AI (with some systematic experiments to assess alternative AI)

APTIMA®
Human-Centered Engineering®

Applying AI to Tactical Training

- Select by probabilistic model
  - SME estimates the training effectiveness of each AI per scenario or vignette (within scenarios) for trainees with varied measurements of scenario performance
  - Analyst defines POMDP model of AI effects given scenario & measured trainee performance
  - Analyst computes a training policy
  - Librarian applies the policy to select the AI and scenario that most reliably advance the trainee furthest to expertise
- Select by empirical effects
  - Researcher conducts training experiments that cross scenario x AI x pilot expertise
    - 5 scenarios x 5 AI x 3 expertise x 8 subjects per level = 600 trials
  - Analyst identifies empirically best AI given scenario & expertise
  - Librarian selects the best AI for the training task

APTIMA®
Human-Centered Engineering®

Summary

- Accomplishments
  - Robust AI pilots
  - Testbed for accelerated AI development
  - Assessment of AI
- Future research
  - Parallelize the testbed for efficient development, big data volume
  - Develop data and librarian to train with tactical AI

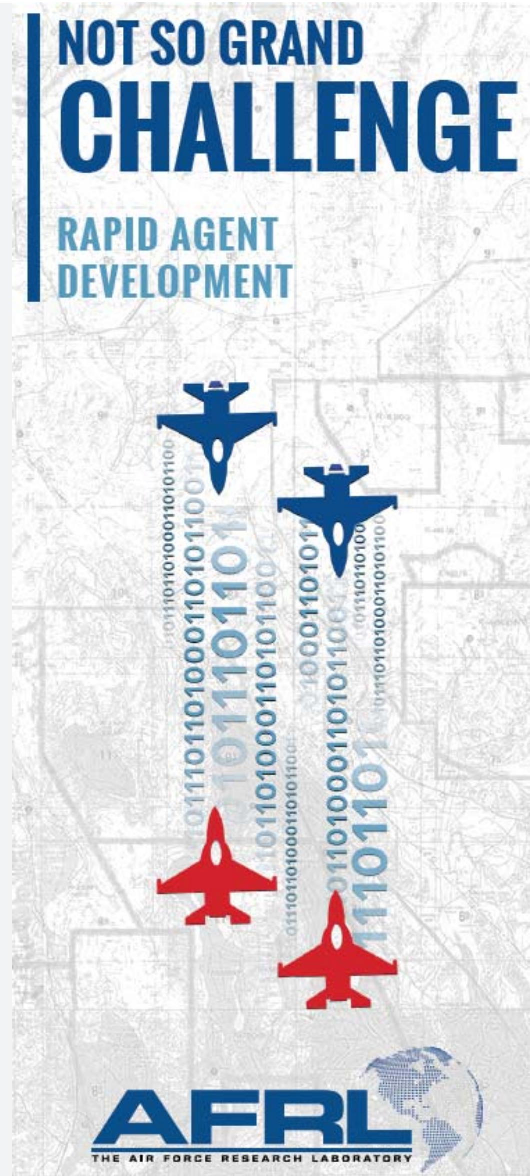

Jared Freeman, Eric Watz
Aptima, Inc., 1010 North Glebe Road, Arlington, VA USA
freeman@aptima.com
www.aptima.com

Winston Bennett
AFRL, Airman Systems Directorate
Warfighter Readiness Research Division, 711 HPW/RHA
2620 Q Street, Wright Patterson AFB, OH USA