

AlphaDogfight Trials: Bringing Autonomy to Air Combat

Christopher R. DeMay, Edward L. White, William D. Dunham, and Johnathan A. Pino

ABSTRACT

The Defense Advanced Research Projects Agency (DARPA) Air Combat Evolution (ACE) program “seeks to increase trust in combat autonomy by using human–machine collaborative dogfighting as its challenge problem. This also serves as an entry point into complex human–machine collaboration” (<https://www.darpa.mil/program/air-combat-evolution>). To set the stage for ACE, the AlphaDogfight Trials program was created to explore whether artificial intelligence (AI) agents could effectively learn basic fighter maneuvers. DARPA contracted the Johns Hopkins University Applied Physics Laboratory (APL) to create an arena to host simulated dogfights—close-range aerial battles between fighter aircraft—where autonomous agents could be trained to defeat adversary aircraft. During the dogfight trials, AI agents competed against each other and the winner competed against a human pilot. By the end of the trials, the program demonstrated that AI agents could surpass the performance of human experts. APL was critical to the success of this program: the Lab created the simulation infrastructure, developed the adversary AI agents, and evaluated the competitors’ AI solutions. This article details APL’s role in advancing combat autonomy through this program.

INTRODUCTION

The vision for the future of air combat is one where unmanned aerial vehicles (UAVs) are operated by artificial intelligence (AI) algorithms in highly complex and dynamic environments. These AI-driven UAVs will need to function seamlessly in teams with manned

platforms, coordinating and collaborating to complete mission objectives. To most effectively team with their autonomous partners, pilots will need to develop trust in their actions and capabilities. The Defense Advanced Research Projects Agency (DARPA) launched its Air

Sponsored by Defense Advanced Research Projects Agency Strategic Technology Office: Air Combat Evolution (ACE) Program; Purchase Request HR0011942236; Program Code MCAC0 under Contract No. HR0011-17-D-0001/Task Order HR001119F0091. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressly or implied, of the Defense Advanced Research Projects Agency or the US government.

Some content in this article is based on C. DeMay, M. Rich, E. White, W. Dunham, J. Pino, W. Li, Z. Akilan, B. Barkley, K. Brady, and C. Cooke, “AlphaDogfight Trials final report,” FPS-R-20-0698 (report to DARPA), Laurel, MD: APL, 2020, Unclassified/For Official Use Only.

Combat Evolution (ACE) program¹ to develop trusted, scalable AI capabilities for air combat. The ACE program seeks to accomplish this goal by building and evaluating AI algorithms first in a modeling and simulation environment and then flying those algorithms on subscale UAVs and finally on full-scale aircraft such as F-16s.

A precursor to ACE, the AlphaDogfight Trials (ADT) program was created as a risk-reduction effort for the larger program. This competition-based activity sought to answer three core questions:

1. Is it possible to teach an AI algorithm to perform within-visual-range air combat?
2. Can we engage a diverse set of companies, including those that are not traditional defense contractors, to capture innovative AI concepts?
3. Can a novel program structure accelerate the development of new AI technologies and applications?

OVERVIEW

Working closely with DARPA, APL played a critical role in shaping and developing the ADT program. In our first responsibility as the trials coordination team (TCT), we assisted DARPA with developing the announcement to solicit proposals and then served as subject-matter experts to advise the DARPA team that evaluated proposals. Competitors were solicited under the Autonomy Research Collaboration Network (ARCNet), a consortium that brings together experts from academia, industry, and the government to develop autonomous technologies for the warfighter. From the proposals, DARPA selected eight technically and organizationally diverse companies, from a university research institute to several small companies to a large defense contractor, to compete in the trials. Throughout ADT, we served as the primary conduit between DARPA and the eight competitor teams, working collaboratively to advance the state of play throughout the program.

For the trials' modeling and simulation environment, we developed a framework called the Coliseum where the AI algorithms were trained and tested in one-versus-one (1-v-1) combat, and where the APL-developed algorithms would face off against the selected competitors. To test the arena and to establish a baseline for the performance of competing autonomous agents, we developed numerous autonomous adversary agents over the course of the

competition. These agents became progressively more complex, with the most highly advanced agents developed using deep reinforcement learning (RL) techniques that trained via self-play without human input. Ahead of the first ADT competition event, these agents were provided to competitors so that they could train and improve their own algorithms. As the agents were continuously updated, the teams could progressively develop and train their AI agents against them.

So that DARPA could identify the most promising AI algorithms, ADT was planned around a series of competition events where the algorithms were put to the test. Between events, APL hosted scrimmages. During the first two competition events, teams faced the APL-developed adversary agents in a series of constructive dogfight engagements. In a third event, added after the program was extended because of the coronavirus pandemic, teams competed against each other's agents. The trials culminated in a 3-day virtual competition, broadcast to the public, where teams competed in simulated combat scenarios against the APL-developed AI agents and each other, and the champion competed against a human pilot. To date, the final event, ADT 3, has garnered over 440,000 views on YouTube, and the game-changing outcome has been recognized by the air and space community² and the Pentagon.³

In this article, we first describe our technical approach for creating a common simulation environment. We then describe the scripted and AI adversary agents we developed for training. And finally, we outline the competition from program inception through the virtual finals.

SIMULATION ENVIRONMENT

For the simulation environment, we developed the Coliseum, an AI arena where AI agents train and

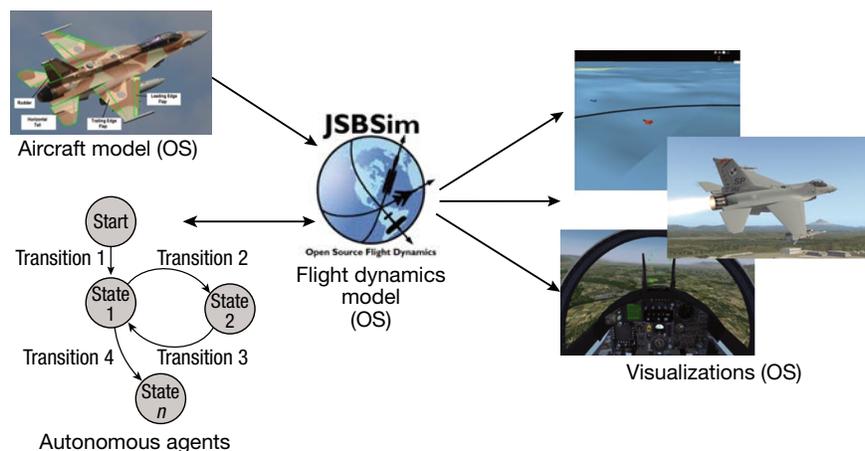


Figure 1. The ADT simulation environment. The framework was developed using a combination of open-source (OS) and APL-developed middleware. ADT competitors were responsible for only their autonomous agents. Leveraging OS software enabled the team to rapidly create a framework to meet program timelines.

compete. It is a constructive, faster-than-real-time framework that combines the open-source flight dynamics software JSBSim⁴ with APL-developed middleware, autonomy algorithms, and visualization software to simulate 1-v-1 dogfights. An overview of the simulation environment is shown in Figure 1. APL-developed and competitor-developed agent algorithms plug in to the ADT simulation environment through the ADT autonomy framework. This framework connects to the ADT Gym-JSBSim environment using OpenAI Gym-inspired standards that allow simple integration of competitors' agent development tools. The Gym-JSBSim environment calculates aircraft states derived specifically for ADT, manages the underlying JSBSim flight dynamics instances, and pushes information to visualization tools.

Visualization

To enable livestreaming of an engagement, the ADT simulation environment contains an ACE viewer application. The ACE viewer uses IEEE 1278 Distributed Interactive Simulation standard protocol data unit packets from the simulation environment to display the dogfights.⁵ The application is built on the open-source CesiumJS world renderer for 3-D geographical display and shows the aircraft in the engagement, trails of the aircraft trajectories, and the weapon engagement zones.⁶ Relevant information, such as engagement time, tactical data about each aircraft, the commanded inputs from the agents and actual flight control system response, and the health of each aircraft, overlays the Cesium display, as shown in Figure 2. The ACE viewer was developed to mimic capabilities used at Nellis Air Force Base and Naval Air Station Fallon to debrief pilots after training missions. For clarity and cohesiveness, a unique color scheme that matches the palette of the overall event scoreboard was used for each competitor.

Virtual Reality

To enable human expert pilots to dogfight the AI

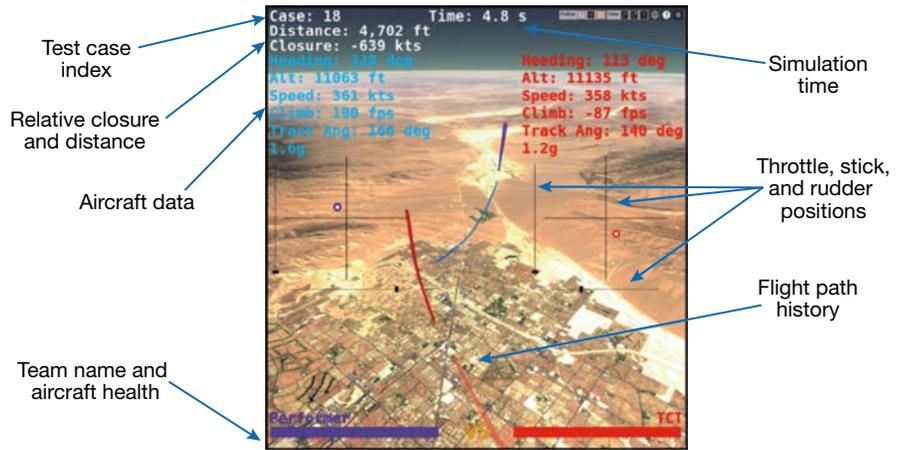


Figure 2. An annotated image of the ACE viewer used in ADT 3. The central image is the ACE viewer, and the information that it displays is labeled.

agents, we created an ADT virtual reality (VR) system for the human pilots. This system provides the pilots with information intended to match the information given to the AI agent. A graphical summary of this engagement information is provided in Figure 3. In addition to providing typical fighter displays, such as a traditional heads-up display, the VR headset also presents ADT simulation-specific components to enhance the pilot's situational awareness of the threat and their relative positioning.

Developing the ADT heads-up display was an iterative process with current and former pilots who helped to optimize not only which information was displayed but also how it was displayed.

AUTONOMY DEVELOPMENT

For training and evaluation, APL developed numerous adversary autonomous agents. The agents ranged in complexity from rudimentary agents, to scripted opponents using finite-state machines, to highly advanced

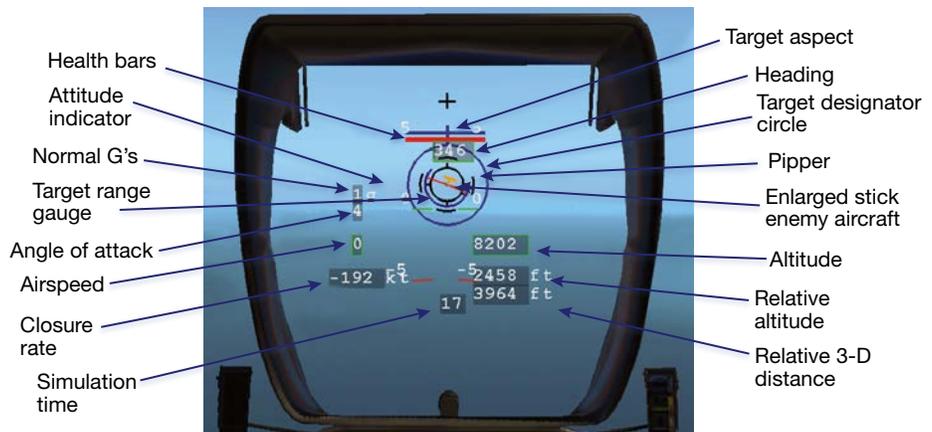


Figure 3. Annotated image of the ADT 3 heads-up display. The information displayed to the pilot is labeled.

| | |
|--|------------------------|
| Rudimentary opponent | Simulates |
| “Zombie”—Straight and level, fixed speed | <i>Cruise missile</i> |
| Basic opponent | Simulates |
| “Rosie”—Executes time-based, minor altitude and speed deviations | <i>One-circle flow</i> |
| Scripted opponent | Simulates |
| “BUD FSM”—Recognizes the state of the engagement and employs a scripted response | <i>Novice pilot</i> |
| RL opponent | Simulates |
| “AlphaMav0”—State-of-the-art agent developed via self-play without human input | <i>Expert pilot</i> |

Figure 4. APL-developed AI agents. APL developed a number of adversary agents throughout ADT using a variety of approaches to simulate tactically relevant opponents.

agents developed using deep RL techniques that trained via self-play without human input. A graphical summary of the APL-developed adversary agents is shown in Figure 4.

Rudimentary Agents

As mentioned, the AI agents span a range of complexity and capability. The least complex (and least capable) are the rudimentary agents. Their purpose is to test and verify that their actions propagate at the correct frequency and interval in the simulation environment. They are not intended to compete against or deceive their opponents and, in fact, do not react to their opponents at all.

Basic Agents

One level above the rudimentary agents are the basic agents. While some of these agents remain nonreactive, others have reactive logic based on rules inspired by simple basic fighter maneuvers (BFM). The basic agents are built on a simple proportional-integral-derivative controller with three loops (inner, middle, and outer). The controller allows the programmer to implement commands in a more discrete fashion (e.g., set a desired altitude) than the “hands on throttle and stick” (HOTAS) inputs that are native to the simulation environment. (Typical fighter aircraft are configured with various buttons and switches on the cockpit throttle lever and flight control stick; pilots can carry out all functions and fly the aircraft by keeping their hands on the throttle and stick and by using these buttons and switches.) The outer loop uses the relative positions, speeds, and headings of the two aircraft to compute a desired speed and change in heading for the agent to intercept its target in a few simple profiles. The computations are derived by establishing a handful of cases of relative distance and

angle and by defining a fixed command that is manually tuned to perform acceptably in each case. For example, at a large distance (10+ miles) and low relative angle ($<60^\circ$), maximum speed with a heading directly at the target is selected. The middle loop exposes set points for absolute heading, altitude, and speed; these commands can come from the profiles in the outer loop or can be overridden in the agent’s logic if desired. The inner loop takes the middle loop commands and maps them to HOTAS commands.

Scripted Agents

Agents in the next level are called scripted agents. These agents contain a small state machine that selects an engagement strategy based on the state of the fight. The agents then use the same set points as the other basic agents to employ or implement that strategy. For example, if an agent detects an opponent closing in from behind, it could select a “flare” strategy and rapidly reduce speed to force its opponent to overshoot.

AI Agents through RL and Adversarial Learning

Our most sophisticated agents were developed using deep RL. RL is a method of generating a desirable sequence of decisions (behaviors) through trial and error. This algorithmic approach has been shown to perform better than human experts in Atari games,⁷ GO,⁸ and StarCraft II.⁹

The general paradigm for RL requires that an agent take one of a number of possible actions in an environment. The environment then provides feedback about the current state of the agent, along with a reward (either positive or negative). The agent can then learn the reward structure of the environment through repeated interactions, and it works to generate an action strategy (policy) that maximizes the cumulative reward at the end of a sequence of interactions with the environment (game). The end result is a policy that maps environment states to agent actions and performs well, as indicated by high rewards.

RL was particularly well suited for generating a highly skilled agent in the ADT environment because the 1-v-1 dogfight simulation was much faster than real time, thus allowing for many interactions in quick succession. In addition, the game has fixed rules that are transferable to a rewards structure well suited for RL, and the software hooks for an open-source learning environment

(AI Gym) were already largely in place. With the right learning architecture, algorithms, and training, an RL agent would have broad knowledge of how to react to other agents given their states in the environment, without actually having trained through trial and error against them.

To create the agent, we successfully used adversarial

RL in the form of self-play in the Coliseum. We sought to develop a balanced learning algorithm that was easy to implement, capable of learning in complex environments, and sample efficient—meaning it required fewer simulation runs to perform well.

The Coliseum: Learning Framework Overview

In 2018, a team in APL's Force Projection Sector began developing a general framework capable of training deep RL algorithms to play both symmetric and asymmetric games. In this independent research and development project, called Athena Inspired, the team developed a set of design principles that combined open-source libraries to create an adversarial learning architecture. The codified implementation of the architecture is called the Coliseum. The Coliseum supported advanced agent development during ADT, including the production of RL agents.

The Coliseum manages agent development using adversarial self-play, as shown in Figure 5. It maintains a repository of historical agents, and the newest agent trains by playing games against these historical agents. The new agent's performance is periodically evaluated via calculation of its win ratio in these games. Once the new agent achieves a desired win ratio, it is copied into the repository as another historical agent, and the agent under training has its version number incremented. The process repeats as the new agents train in self-play against the historical agents. Using this process, the current training agent will always be able to select from a distribution where the best agent in the pool is itself.

The Coliseum selects opponents for the training agent from the repository by using a probability distribution of the opponents based on skill. Opponents that are more skilled are more likely to be selected from the distribution, although the specific logic and distributions are customizable. Likewise, opponents that are not as skilled are typically chosen less frequently. This process ensures that agents are trained against a diverse set of adversaries with similar skill levels; both are necessary to prevent overfitting and to promote continuous skill improvement.

The Coliseum uses the Elo system to estimate the agents' skill levels relative to one another. Elo, a rating

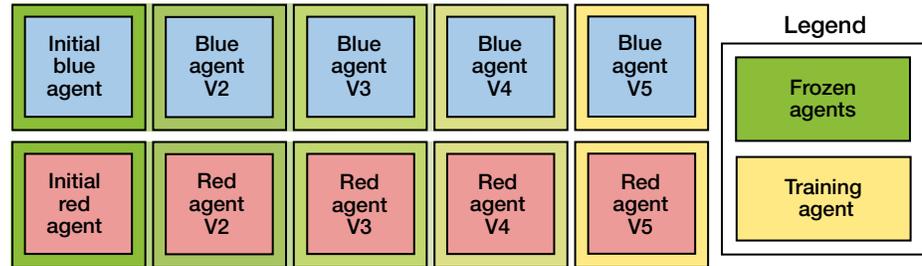


Figure 5. Illustration of adversarial self-play. Self-play is an agent-training pattern in which an agent plays against a mixture of historical copies of itself within the simulation environment.

system invented by Arpad Elo for ranking chess players,¹⁰ has been used for a wide variety of different symmetric zero-sum games (i.e., games where one agent wins and one loses). In the Elo system, a game's result is considered a sample from a Gaussian distribution around the agent's skill. To produce an accurate Elo rating, the agent must sample many games against an opponent to determine relative skill levels. Although there are other ranking systems, Elo remains the most popular for zero-sum adversarial games.

The simulation environment for ADT provides a true score at the end of the game. There are two main techniques for mitigating the effects of sparse reward signals: reward shaping and curriculum learning.

Reward shaping uses the current observation space and expert knowledge to provide a reward signal to the agent more frequently. Although using expert knowledge helps to ensure that a computed reward is frequently provided for “good” behaviors, it is prone to incentivizing behaviors that are often unintended or not in the spirit of the game. Reward shaping also limits the agent's overall ability to find optimal policies, since part of the solution is prespecified by the developer.

In curriculum learning, the agent first learns to solve an easier task to develop a certain behavior or policy that will transfer to a harder task. Curriculum learning requires the game designer to create a set of scenarios that will gradually converge to the real task over multiple iterations. When developing agents, we used a combination of curriculum learning combined with the Coliseum's self-play to achieve strong performance. After each evaluation period, if the desired win ratio is achieved, the Coliseum pushes the curriculum one iteration closer to the final game. Using a well-designed curriculum can prevent the need to use reward shaping.

DEVELOPMENT OF COMPETITION EVENTS

APL developed the ADT competition structure and planned and executed three trials and numerous scrimmages for the eight competitors (see Figure 6). During ADT 1 and ADT 2, 2-day events held at APL in November 2019 and January 2020, respectively, the eight

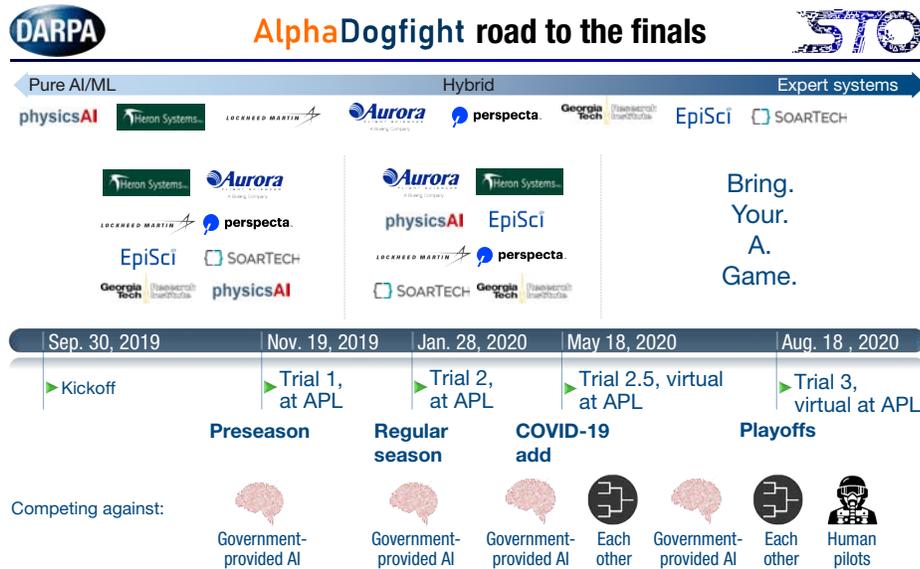


Figure 6. ADT competition overview. The competitors selected by DARPA were Aurora Flight Systems, EpiSci, Georgia Tech Research Institute, Heron Systems, Lockheed Martin, Perspecta Labs, PhysicsAI, and SoarTech. (Source: DARPA ACE program.)

teams—Aurora Flight Systems, EpiSci, Georgia Tech Research Institute, Heron Systems, Lockheed Martin, Perspecta Labs, PhysicsAI, and SoarTech—competed against APL-developed AI agents in a series of constructive dogfight engagements. During ADT 2.5, a virtual scrimmage added in May 2020 after a program extension resulting from the coronavirus pandemic, competitors submitted their own agents to APL and competed against one another in a round-robin competition. The finale, ADT 3, was a 3-day virtual competition broadcast live from APL via ZoomGov and YouTube in August 2020. Events included a flight brief, a series of simultaneous batch runs, live performance tracking via a dashboard, and playback/debrief of individual runs.

ADT 3—Virtual Finals

The finals were designed to provide critical insight into the future of AI's role in air combat. The event was originally planned to take place at the US Air Force's AFWERX innovation hub and then at Nellis Air Force Base (both in Nevada) but had to be postponed because of the COVID-19 pandemic. Rather than continue to postpone the event, DARPA decided to host it virtually at APL. The APL and DARPA teams collaborated to ensure that the 3-day public event was seamless and engaging for competitors and spectators alike.

Held on August 18–20, 2020, ADT 3 was broadcast live from APL's Intelligent Systems Center via ZoomGov webinar and YouTube. On day 1, the eight teams competed against the APL-developed AI agents. Day 2 featured a competitor-versus-competitor round-robin face-off. Day 3 began with the top four teams competing

in a single-elimination bracket competition, and the champion then competed against an F-16 pilot in an AI-versus-human matchup.

Production Elements

As the TCT, the APL team executed a wide range of activities to develop, advertise, implement, and broadcast ADT 3. The team, which included staff members from various APL sectors and departments with an array of skills and expertise, created a comprehensive branded ADT 3 registration website that included a description of the event, information on the competitors, media resources (pictures, videos, and press releases), and the ADT visualization display. Leading up to the event, APL and DARPA coordinated to maximize publicity, with both APL and DARPA issuing public media releases before and during the event. APL communications professionals also created visual materials, including a promotional video, a competition logo, and competition photos from the production at APL.¹¹ Figure 7 shows some of these elements. In addition, the APL team provided technical support to ADT 3 registrants before and during the competition.

The layout of the space for the event, in APL's Intelligent Systems Center Gym, was completely reworked to accommodate the shift from an in-person event to a virtual live broadcast event. Central to this improved layout was the ADT Studio, which housed the "Control Zone" stage and displays (discussed in more detail below



Figure 7. Some of the branded elements for ADT. The APL team created various visual materials, including a branded registration website, a promotional video, a competition logo, and competition photos of the event at APL.



Figure 8. ADT 3 studio where the virtual live event was broadcast. In the background are commentators in the Control Zone, where experts in air combat and autonomy provided analysis and commentary.

and shown in Figure 8) as well as the VR competition space (Figure 9). Signage, lighting, and other professional production elements were included in the layout to create an immersive competition environment for a broadcast using multiple cameras. The dogfights were run and displayed in a section the APL production team called the AI Arena (Figure 10). Although ADT's original goals did not include demonstrating how to host a polished virtual event, the team certainly rose to the unforeseen challenge.

Control Zone

As noted, one of ADT's primary goals was to engage and excite partners beyond traditional Department of Defense contractors. To help realize that goal, ADT was structured as an e-sports competition (usually a competitive video gaming event that includes spectators and commentators, similar to traditional competitive sporting events). At ADT 2, we introduced the Control



Figure 9. ADT 3 VR competition space where human expert pilots would dogfight the best AI agent. Shown is a pilot wearing the VR headset in combat with an AI agent. The image shown on the display is a pilot's cockpit view of the dogfight.

Zone, modeled on ESPN's *SportsCenter*, where experts in air combat and autonomy provided commentary. The Control Zone was named after the area in space where pilots can counter any defensive maneuver and remain in their controlling position. In ADT's Control Zone, experts discussed the basics of AI and dogfighting, and how AI and human pilots train. Commentators in the Control Zone were both educational and entertaining as they provided analysis and commentary on the dogfighting engagements.

Day 3—ADT Champion and Human-versus-AI Matchup

Day 3 of the trials was expected to have the largest viewing audience and was therefore the most highly produced.¹² Pregame activities provided a look ahead at the final day of competition, a recap of days 1 and 2, team videos for four competitor teams that had already been eliminated (EpiSci, Georgia Tech Research Institute, Perspecta, and SoarTech), and live interviews with all eight teams. Unlike on days 1 and 2, all matches were shown sequentially with live commentary from the Control Zone throughout the competition. In the semifinal 1 matchup, PhysicsAI, the number-three team, competed against Lockheed Martin, the number-two team. Lockheed showed a commanding performance against



Figure 10. ADT 3 AI Arena and production spaces. Top, The TCT executed the competition from the AI Arena. Bottom, The APL production team during production and broadcast of the show.

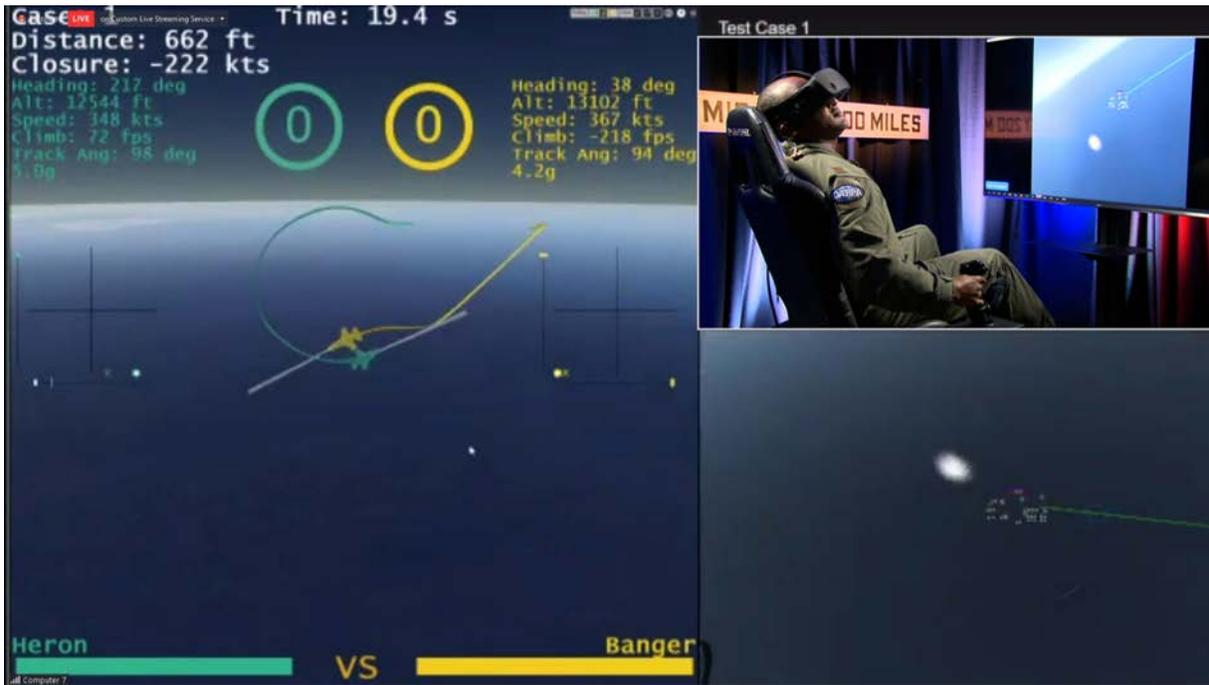


Figure 11. ADT 3 human-vs.-AI screen capture from YouTube. This snapshot from the first engagement shows the ACE viewer, which provided an overview of the fight while also showing the human pilot, Banger, in the competition space and the view from the F-16 pilot's perspective. In this final matchup, the AI agent's highly accurate weapons employment proved too difficult for the human pilot, and Heron Systems won 5 to 0.

PhysicsAI with a 13-to-2 kill ratio over 20 engagements. In semifinal 2, the number-one team, Heron Systems, faced the number-four team, Aurora, with Heron soundly defeating Aurora with a 17-to-2 kill ratio. The finals came down to the number-one and number-two teams, Heron and Lockheed Martin, for the right to be crowned ADT champion. Heron proved victorious, achieving a 16-to-4 kill ratio. Key to Heron's success was its ability to execute aggressive and highly accurate forward gun attacks against its opponent.

The final portion of day 3 was the main event—the showcase matchup between the ADT champion Heron and an F-16 Weapons Instructor Course graduate, call sign Banger. Before the matchup, Dr. Tim Grayson, director of DARPA's Strategic Technology Office (STO), and Dr. Peter Highnam, who was then the acting DARPA director, provided live commentary on the impact of ADT. Analysts in the Control Zone recapped the finals, gave a mission brief of the event, and interviewed Banger from the VR competition space.

The human-versus-AI matchup was a set of five engagements initiated from neutral starting conditions at varying altitudes. Figure 11 is a screen capture from YouTube during the first engagement. The ACE viewer provided an overview of the fight while also showing Banger in the competition space and the view from the F-16 pilot's perspective.

In the human-versus-AI matchup, Heron's AI agent dominated, using quick kills at the merge, and won

5 to 0. The Heron AI agent's highly accurate weapons employment proved too difficult for Banger, who was employing F-16 BFM-like tactics. Only on the final engagement did Banger manage to survive those early merges by using aggressive out-of-plane maneuvering to draw out the fight, but the F-16 Weapons Instructor Course graduate ultimately lost.

Although the Heron Systems AI agent swept Banger, we should note that the goal is not ultimately to replace human pilots with automated ones. Rather, it is to incorporate AI alongside humans to build a more effective fighting force—and the results of ADT represent one way to build humans' confidence in their future automated counterparts.

CONCLUSION

ADT aimed to develop intelligent autonomous agents capable of defeating an adversary aircraft in a simulated dogfight and to demonstrate their capabilities in simulated air combat. In doing so, ADT engaged both established partners and new ones, increasing technical and organizational diversity. We at APL, as the TCT, were responsible for evaluating the program; developing the simulation infrastructure and adversary autonomy; and developing, planning, and hosting three ADT competition events.

This competition-based risk-reduction activity sought to answer three core questions:

1. Is it possible to teach an AI algorithm to perform within-visual-range air combat?
2. Can we engage a diverse group of partners to capture innovative AI concepts?
3. Can a novel program structure accelerate the development of new AI technologies and applications?

At the conclusion of the trials, we can confidently answer yes to all three of those questions:

1. By ADT 2, the leading teams had produced algorithms that resembled BFM after just 4 months. By ADT 3, most of the teams had produced agents capable of executing a diverse set of offensive, neutral, and defensive BFM. In addition, the top teams demonstrated the ability to go head-to-head with expert human pilots and win, albeit with some unrealistic characteristics, such as perfect state information.
2. The winning team, Heron Systems, was not a traditional Department of Defense partner, nor was PhysicsAI, which ranked in the top four. Both teams developed novel approaches to train and develop their algorithms, allowing them to compete and excel against their peers. Lockheed Martin, a large, established defense partner, showed that the big companies could also assemble teams able to rapidly develop new technology.
3. A competition structure with a common simulation environment allowed the government to assess performance throughout the program. Despite the lack of prize money, the competition structure itself was a significant motivator to drive team performance. A more traditional approach with a final delivery at the program's conclusion probably would have limited discovery and performance.

In short, ADT exceeded its stated goals, and ADT 3 reached an audience vastly wider than imagined at the program's inception. The keys to success in executing the trials were the drive for constant improvement via iteration and the collaborative spirit among APL, DARPA, and the competitors. After the event, Dr. Tim Grayson noted that the "outcome shows great promise for future airborne combat systems and concepts involving human-machine symbiosis."¹³ But the work of bringing AI to air combat has only begun. As impressive as ADT was, there are still technical and cultural challenges to overcome before this technology is ready.

In the follow-on DARPA ACE program, APL will serve as the Experimentation Integration Team. We will leverage ADT development and analysis to reduce technical risk by evaluating autonomous control algorithms across multiple program phases, beginning with modeling and simulation and progressing to subscale live UAVs and then ultimately to full-scale live aircraft, with the goal to develop trusted, scalable AI capabilities for air combat.

The future of warfare will be defined by advanced AI and autonomous systems that outthink, outmaneuver, and outperform conventional manned forces. These systems will not replace humans but will work in tandem with them. Future battles will be fought at machine speeds with compressed timelines and more maneuverable and agile systems. With sophisticated behaviors and decision-making, these systems will have near-optimal decision-making and complex collaborative dynamic behaviors. ADT provided a first look at what is possible at the intersection of AI and air combat. As these capabilities evolve, lessons learned from ADT will serve as the foundation for future developments, ensuring that AI air combat solutions are robust, secure, effective, and trusted.

ACKNOWLEDGMENTS: The authors acknowledge the AlphaDogfight Trials technical and communications team members who were instrumental in making the effort a success—Matthew Rich, Zackariah Akilan, Brett Barkley, Kelly Brady, Kyle Casterline, Christian Cooke, Jenny Gebhardt, David Handelman, Emma Holmes, William Li, Galen Mullins, Khang Ngo, John O'Brien, Corban Rivera, Robert Shearer, Thomas Urban, and Lee Varanyak.

Distribution Statement A: Approved for public release; distribution is unlimited.

REFERENCES

- ¹K. Plaks. "Air Combat Evolution (ACE)." DARPA. <https://www.darpa.mil/program/air-combat-evolution> (accessed May 26, 2021).
- ²G. Warwick, "Aerospace technology highlights of 2020: Artificial intelligence," *Aviation Week*, <https://aviationweek.com/aerospace-defense-2021/aerospace/aerospace-technology-highlights-2020> (accessed May 26, 2021).
- ³S. J. Freedberg Jr., "AI to fly in dogfight tests by 2024: SecDef," *Breaking Defense*, Sep. 9, 2020, <https://breakingdefense.com/2020/09/ai-will-dogfight-human-pilots-in-tests-by-2024-secdef/>.
- ⁴"jsbsim." GitHub. <https://github.com/JSBSim-Team/jsbsim> (accessed May 26, 2021).
- ⁵*IEEE Standard for Distributed Interactive Simulation—Application Protocols*, IEEE Standard 1278.1-2012 (revision of IEEE Standard 1278.1-1995), Dec. 19, 2012, <https://doi.org/10.1109/IEEESTD.2012.6387564>.
- ⁶"CesiumJS." Cesium. <https://cesium.com/platform/cesiumjs/> (accessed May 26, 2021).
- ⁷A. P. Badiar, B. Piot, S. Kapturowski, P. Sprechmann, A. Vitvitskiy, D. Guo, and C. Blundell, "Agent57: Outperforming the Atari human benchmark," arXiv, Mar. 30, 2020, <https://arxiv.org/abs/2003.13350>.
- ⁸D. Silver, A. Huang, C. Maddison, A. Guez, L. Sifre, et al., "Mastering the game of Go with deep neural networks and tree search," *Nature*, vol. 529, pp. 484–489, 2016, <https://doi.org/10.1038/nature16961>.
- ⁹O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, et al., "Grandmaster level in StarCraft II using multi-agent reinforcement learning," *Nature*, vol. 575, pp. 350–354, 2019, <https://doi.org/10.1038/s41586-019-1724-z>.
- ¹⁰A. Elo, *The Rating of Chessplayers, Past and Present*, 2nd ed. London: Arco, 1986.
- ¹¹"JHU Applied Physics Laboratory." YouTube. <https://www.youtube.com/user/jhuapl>.
- ¹²Day 3 of ADT 3 can be viewed on YouTube at <https://www.youtube.com/watch?v=NzdhIA2S35w>. Note: The video starts just prior to semifinals 1.
- ¹³DARPA, "AlphaDogfight trials foreshadow future of human-machine symbiosis," Aug. 26, 2020, <https://www.darpa.mil/news-events/2020-08-26>.



Christopher R. DeMay, Force Projection Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Christopher R. “Disco” DeMay is the intelligent combat platforms program manager in APL’s Precision Strike Mission Area. Chris is a graduate of the US Naval Test Pilot School and Electronic Attack Weapons School and holds a BS in aerospace engineering from the US Naval Academy and an MS in aeronautical engineering from the Naval Postgraduate School. He leads applied research and advanced technology development for the Office of the Secretary of Defense, US Air Force, and Defense Advanced Research Projects Agency (DARPA) programs applying autonomy and manned–unmanned teaming to air combat. Before assuming his current role, he was a senior defense analyst in APL’s National Security Analysis Department. Prior to joining APL, he served for over 21 years in the US Navy as a naval aviator, completing multiple operational tours and serving in roles including command, Joint and Naval staff, and executive assistant. His email address is christopher.demay@jhuapl.edu.



Edward L. White, Force Projection Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Edward L. White is a section supervisor in the Intelligent Combat Platforms Group in APL’s Force Projection Sector. He earned BS degrees in aerospace engineering and mechanical engineering from the University of Arizona, an MS in mechanical engineering from the University of Arizona, an MBA from the Eller College of Management at the University of Arizona, and a PhD in mechanical engineering from Purdue University. Edward leads a research portfolio focused on building the infrastructure to transition artificial intelligence concepts from the laboratory to the field. He is the technical lead for the Mission Termination System (MTS) and Skyborg programs, principal investigator or co-principal investigator on several independent research and development projects, acting principal investigator for the Defense Advanced Research Projects Agency (DARPA) Air Combat Evolution (ACE) program, and co-lead for the Superhuman Decision-Making and Autonomous Action technical vector of the AI Futures team, and he supports multiple other programs. Edward is a member of the Applied Biomedical Engineering program faculty at the Johns Hopkins Whiting School of Engineering, where he co-developed and co-taught a class

on bio-inspired robotics for biomedical applications. His doctoral research focused on soft robotic systems and generated more than 20 peer-reviewed journal articles and conference papers on a range of topics from fundamental materials science to applied control for nonlinear robotic systems. His email address is edward.white@jhuapl.edu.



William D. Dunham, Force Projection Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

William D. Dunham is supervisor of the Modeling, Simulation, and Analysis Section in APL’s Force Projection Sector. He has a BS in mechanical engineering from the University of Maryland, an MS in aerospace engineering from the University of Michigan, and a PhD in flight dynamics and controls, also from the University of Michigan. Will is the lead developer for the within-visual-range engagement simulation environment and dogfight autonomy competitor coordination for the Defense Advanced Research Projects Agency (DARPA) Air Combat Evolution (ACE) program. He joined the Lab in the summer of 2019. His email address is will.dunham@jhuapl.edu.



Johnathan A. Pino, National Security Analysis Department, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Johnathan A. Pino is a project manager and assistant section supervisor of the Algorithmic Warfare Analysis Section in APL’s National Security Analysis Department. He has a BS in aerospace engineering from the University of Maryland and an MS in applied and computational mathematics from Johns Hopkins University. He is the adversary autonomy development lead for the Defense Advanced Research Projects Agency (DARPA) Air Combat Evolution (ACE) program and held the same role during AlphaDogfight Trials. In addition to autonomous systems, his focus areas include live, virtual, and constructive mission- and engagement-level modeling, simulation, and analysis (MS&A) in the surface, air, and space domains and model development and integration. In addition to ACE, he supports software design and development, human subjects research, and MS&A efforts across the Laboratory. His email address is johnathan.pino@jhuapl.edu.