

Genetic Fuzzy based Artificial Intelligence for Unmanned Combat Aerial Vehicle Control in Simulated Air Combat Missions

Nicholas Ernest¹, David Carroll², Corey Schumacher³, Matthew Clark⁴, Kelly Cohen^{5*} and Gene Lee⁶

¹CEO, Psibernetix Inc., USA

²Programming Lead, Psibernetix Inc., USA

³Portfolio Lead, Air Force Research Laboratory, USA

⁴Branch Chief, Air Force Research Laboratory, USA

⁵Aerospace Engineering, University of Cincinnati, USA

⁶(Retired) United States Air Force Colonel, USA

Abstract

Breakthroughs in genetic fuzzy systems, most notably the development of the Genetic Fuzzy Tree methodology, have allowed fuzzy logic based Artificial Intelligences to be developed that can be applied to incredibly complex problems. The ability to have extreme performance and computational efficiency as well as to be robust to uncertainties and randomness, adaptable to changing scenarios, verified and validated to follow safety specifications and operating doctrines via formal methods, and easily designed and implemented are just some of the strengths that this type of control brings. Within this white paper, the authors introduce ALPHA, an Artificial Intelligence that controls flights of Unmanned Combat Aerial Vehicles in aerial combat missions within an extreme-fidelity simulation environment. To this day, this represents the most complex application of a fuzzy-logic based Artificial Intelligence to an Unmanned Combat Aerial Vehicle control problem. While development is on-going, the version of ALPHA presented within was assessed by Colonel (retired) Gene Lee who described ALPHA as “the most aggressive, responsive, dynamic and credible AI (he’s) seen-to-date.” The quality of these preliminary results in a problem that is not only complex and rife with uncertainties but also contains an intelligent and unrestricted hostile force has significant implications for this type of Artificial Intelligence. This work adds immensely to the body of evidence that this methodology is an ideal solution to a very wide array of problems.

Keywords: Genetic Fuzzy Systems; Artificial Intelligence; Unmanned Combat Aerial Vehicles; Verification & Validation; Computational Efficiency

Abbreviations: UCAV: Unmanned Combat Aerial Vehicle; AI: Artificial Intelligence; GFT: Genetic Fuzzy Tree; FIS: Fuzzy Inference System; AFRL: Air Force Research Laboratory; AWACS: Airborne Warning and Control System; SA: Situational Awareness; (L,M,S) RM: (Long, Medium, Short) Range Missile; WEZ: Weapon Engagement Zone; CoA: Course of Action; V&V: Verification and Validation

Introduction

Unmanned Combat Aerial Vehicle (UCAV) research has allowed the state of the art of the remote-operation of these technologies to advance significantly in modern times, though mostly focusing on ground strike scenarios. Within the context of air-to-air combat, millisecond long timeframes for critical decisions inhibit remote-operation of UCAVs. Beyond this, given an average human visual reaction time of 0.15 to 0.30 seconds, and an even longer time to think of optimal plans and coordinate them with friendly forces, there is a huge window of improvement that an Artificial Intelligence (AI) can capitalize upon. While many proponents for an increase in autonomous capabilities herald the ability to design aircraft that can perform extremely high-g maneuvers as well as the benefit of reducing risk to our pilots, this white paper will primarily focus on the increase in capabilities of real-time decision making.

There are a number of obstacles to a system being an effective AI within this context. The main developmental difficulties for this type of problem are the vast number of inputs and outputs to be considered, as well as the uncertainty and randomness inherits in the problem. Additionally, in combat one’s opponent could attempt to not match an AI’s training data; hostiles can and will actively seek to fool and exploit the system. Regarding implementation, the ability to verify and validate

the AI is crucial. Safety specifications and operating doctrines need to be guaranteed to be followed via formal methods. Of course computing systems can crash and sensors can fail, though this is just as true for manned aircraft, and redundancies can be put in place.

Fuzzy control has been proven effective in problems containing almost all of these difficulties. Fuzzy is robust, adaptable, high-performing, computationally efficient, and provides an excellent framework to synthesize formal models for the purposes of verification and validation [1]. However, to date, it suffers heavily from the “curse of dimensionality”; that is, as the scope of the problem increases, the computational cost increases rapidly [2]. A standard fuzzy-based system would be extremely computational intractable for a problem such as air-to-air combat.

This intractability not only comes from the run-time associated with evaluating a fuzzy controller of immense size, but also the creation of the controller. A genetic fuzzy system is a methodology in which a genetic algorithm creates all of the components of the controller [3]. Genetic fuzzy systems have found immense success in developing high-performing controllers in small-scale problems [4].

***Corresponding author:** Kelly Cohen, Professor of Aerospace Engineering, University of Cincinnati, USA, Tel: 15137666586; E-mail: Kelly.Cohen@uc.edu

Received December 18, 2015; **Accepted** March 19, 2016; **Published** March 22, 2016

Citation: Ernest N, Carroll D, Schumacher C, Clark M, Cohen K, et al. (2016) Genetic Fuzzy based Artificial Intelligence for Unmanned Combat Aerial Vehicle Control in Simulated Air Combat Missions. J Def Manag 6: 144. doi:10.4172/2167-0374.1000144

Copyright: © 2016 Ernest N, et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

While this automates the process, the computational cost of the genetic algorithm for searching over a practically infinite solution space causes the traditional genetic fuzzy system to be infeasible in more complex problems. Methodologies have been developed to help mitigate this, but one method in particular is capable of applying fuzzy control to problems of this scope.

The Genetic Fuzzy Tree (GFT) has shown an incredible ability to obtain unparalleled levels of performance in very large and complex problems that contain all of the difficulties that alternative intelligent systems have issues coping with. This new subtype of genetic fuzzy system was recently developed during Dr. Ernest's graduate studies, under the guidance of Dr.'s Cohen and Schumacher and supported by the Dayton Area Graduate Studies Institute. The aim of this initial work was to control a flight of ground strikeUCAVs in a low-fidelity simulation environment [5,6]. The success of this study led to Psibernetix Inc. partnering with the US air force research laboratory (AFRL) to apply the GFT methodology to a much more complex problem.

Just as UAVs represented a revolutionary capability for the USAF in the mid-1990s, Manned-Unmanned Autonomous Teaming in an air combat environment will certainly represent a revolutionary leap in capability of airpower in the near future. Air combat, as it is performed by human pilots today is a highly dynamic application of aerospace physics, skill, art, and intuition to maneuver a fighter aircraft and missile against an adversary moving at high speeds in three dimensions. Today's fighters close on each other at speeds in excess of 1,500 MPH while flying at altitudes above 40,000 feet. The selection and application of air-to-air tactics requires assessing a tactical advantage or disadvantage and reacting appropriately in microseconds. The cost of mistakes is high.

The US's "near-peer" adversaries are developing capabilities which could challenge US air superiority. Anti-access environments with modern, overlapping ground and air threats present an existential threat to modern day pilots and fighter aircraft. Future aircraft are likely to employ a high level of coordinated autonomous offensive and defensive capabilities, requiring reaction times which surpass that of a human pilot, in order to survive in such hostile environments. These future fighters may be optionally manned or employ autonomousUCAV wingmen, capable of performing air combat. These fighters would employ with a battle management system on-board which would possess an intelligent agent which would select tactics, manage weapon employment, determine own-ship reactions, and supervise the reactions of its wingmen.

Psibernetix has created ALPHA, a GFT that serves as an AI for flights ofUCAVs in air-to-air combat missions. While currently a simulation tool, ALPHA is aimed towards "increasing autonomous capabilities to allow mixed combat teams of manned and unmanned air fighters to operate in highly contested environments" [7]. This project began in May of 2015, and as of November, already has produced significant results. The GFT methodology has already found great success in various domains, however the level of performance so far obtained here in a realistic simulation environment serves as a tremendous breakthrough in the realm of fuzzy control.

Problem Statement

ALPHA's current primary objective is to serve as an intelligent hostile force for pilots to train against within the AFSIM simulation environment [8]. While the models showcased within this document

for platforms do not mimic any system specifically, AFSIM realistically represents a modern air combat environment with appropriately-behaving models for aircraft, sensors, and weapons. Data is collected from sensors, fused real-time, and sent to ALPHA, complete with noise and potential failures. This data feed is received at and commands sent to unique UDP ports for each platform, causing this problem to contain some of the software complexities that would be present in actual hardware implementation.

For current mission profiles, ALPHA's red forces are handicapped with shorter range missiles and a reduced missile payload than the blue opposing forces. ALPHA also does not have airborne warning and control system (AWACS) support providing 360° long range radar coverage of the area; while blue does have AWACS. The aircraft for both teams are identical in terms of their mechanical performance. While ALPHA has detailed knowledge of its own systems, it is given limited intelligence of the blue force a priori and must rely on its organic sensors for situational awareness (SA) of the blue force; even the number of hostile forces is not given. Both to mirror training exercises and to offset these weaknesses, ALPHA is typically given a numeric advantage over the blue forces. However this is not always the case and ALPHA is capable of controlling any finite number of friendly aircraft. The current problem is focused on purely beyond visual range air-to-air combat missions; no ground targets or friendly platforms requiring escort are considered at this time.

The mission analyzed in this document features two blue fighters vs. four red. The red aircraft begin over a defended coastline and the blues are 54 nautical miles due west. The blues each have 4 long range missiles (LRMs) and 4 acrobatic short range missiles (SRMs), whereas the reds have 4 medium range missiles (MRMs) onboard each platform. The reds' radar is long range with a +/- 70° azimuth angle and a 15 degree elevation angle. The initial state of the mission is displayed below in Figure 1, with the blue AWACS off-screen, due northwest of the blue fighters.

Even in its current developmental stage, ALPHA is taking in an enormous amount of data from the system. Over 150 inputs are considered in this example mission already; some of which ALPHA utilizes both the current value as well as the time history of the input. Examples of inputs include all positional, velocity, and acceleration states, estimated missile range data, visibility of each platform, bogey ratio, and number of shots taken by hostile platforms. ALPHA currently has control of the motion and firing capabilities of each aircraft, with control over more complex sensors planned for future work.

The initial merit for success of ALPHA was to be able to consistently outperform a baseline controller that was previously utilized by AFRL within AFSIM. Within the first months of development this

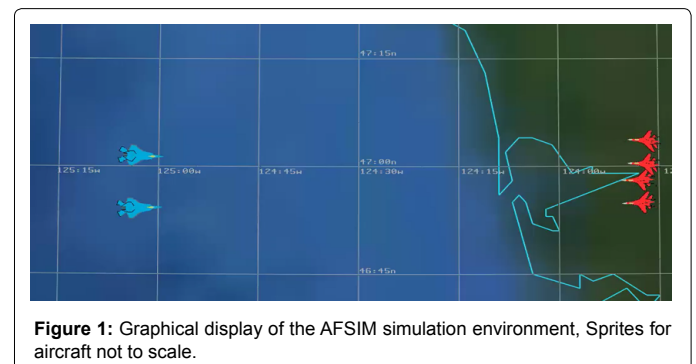


Figure 1: Graphical display of the AFSIM simulation environment, Sprites for aircraft not to scale.

milestone was passed. The current goal of this study is to further push the envelope of how effective the system can be against trained fighter pilots. Preliminary testing against human opponents has been conducted and more formal and thorough ALPHA vs. human testing is planned (Figure 2).

Methodology and Implementation

The GFT methodology utilizes a collection of Fuzzy Inference Systems (FISs) with varying levels of connectivity. Unlike in a standard genetic fuzzy system, here a genetic algorithm, or other learning system, is utilized to train each system in the Fuzzy Tree simultaneously [5]. Each FIS has membership functions that classify the inputs and outputs into linguistic classifications, such as “far away” and “very threatening”, as well as if-then rules for every combination of inputs, such as “If missile launch computer confidence is moderate and mission kill shot accuracy is very high, fire missile”. By breaking up the problem into many sub-decisions, the solution space is significantly reduced. The cost of this approach is a risk of not accommodating for all sources of coupling. Through optimal design this can be minimized if not entirely mitigated. Unlike in Fuzzy Decision Trees or Fuzzy Networks, the fuzzy nodes of a GFT are not individual components of FISs, but rather are unique FISs themselves [9-11].

A key strength of this methodology is the priority given to flexibility and ease of design. While most of the complex decisions are determined by a collection of FISs, other algorithms can be easily included in the system, as well as common-sense heuristics. For example, past GFTs created by Dr. Ernest have included a Cooperative Task Assignment Algorithm, Fuzzy Clustering Route Solver, and No Communications Fire Control System [12-14]. Where an optimal solution to a sub-problem is known, it is directly utilized. An example of the layout of this type of system is displayed in Figure 3. The ability to input lessons learned from expert knowledge, such as current doctrines and teachings of expert fighter pilots, combined with the ability to fully optimize these concepts via a learning system is a significant factor in the success of the GFT. This architecture, utilization of fuzzy logic, and Psibernetix’s use of the Python programming language allows the development of these systems to be rapid and very cost-effective [15].

ALPHA currently has branches for high-level tactics, firing, evasion, SA, and defensiveness. As an example of this structure, the defensiveness branch of the Fuzzy Tree quantifies each aircraft’s defensiveness relative to each threat it faces. The primary output of this branch is a measure between 0 and 100 percent. This defensiveness

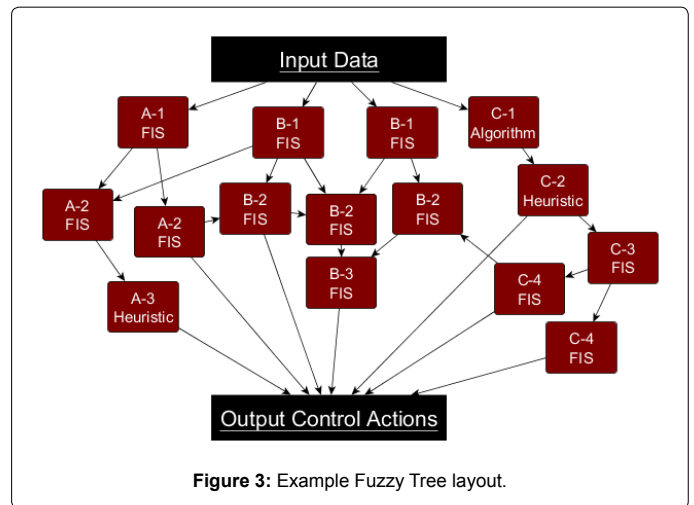


Figure 3: Example Fuzzy Tree layout.

measure is utilized to plan missile evasion maneuvers, alter angles of approach to (or retreat from) hostiles, and determine and adjust higher-level tactics.

This branch is a two-tiered FIS cascade. The first layer measures the effective range of the threat by taking in 2-dimensional distance and altitudes of the hostile and friendly platform. This feeds into the second layer which takes in the effective range and combines it with the threat’s azimuth and the threat’s turning rate in the North-East axis, both with respect to the friendly aircraft. Each of these measurements is broken down into tactically useful membership functions for the FIS. These classifications collectively form a picture upon which ALPHA can make decisions. The defensiveness measure has the sole responsibility for deciding which friendly aircraft should begin evasive maneuvers against hostile missiles whose targets are unknown in the current implementation of ALPHA. It is also solely responsible for determining when defensive missile firings should occur. The SA branch will utilize this measure often, and the effective range FIS itself has applications in many other areas.

ALPHA has inputs to alter behaviors based on mission performance. If the enemy is successfully defeating intended kill shots, ALPHA will adjust the range at which shots are taken in order to not waste the extremely limited ammunition. If the initial estimation of the enemy missiles’ capabilities is inaccurate, defensiveness factors will be adjusted to allow ALPHA to more properly engage the current foe. Additional on-line learning capabilities are planned, such as a continuously updated 3-dimensional modelling of each enemy’s weapon engagement zone (WEZ) for each of their weapon types.

Rather than utilize a standard genetic algorithm as a learning system, Psibernetix’s patent-pending EVE learning system has trained ALPHA [16]. EVE has been shown to have unparalleled learning capabilities in extremely complex problems where obtaining the absolute optimal solution is not necessary or practical [13]. For example, an air-to-air missile does not need to perfectly penetrate the cockpit and physically strike the pilot of an aircraft, and finding a course of action (CoA) that utilizes 0.05 milligrams less fuel to complete an objective is not mission critical. A learning system that could guarantee obtaining these solutions would be computational intractable and would never provide an answer. EVE is a GFT whose objective is the creation and optimization of other GFTs. Through multiple recursive application, EVE has been trained to train other GFTs very effectively.

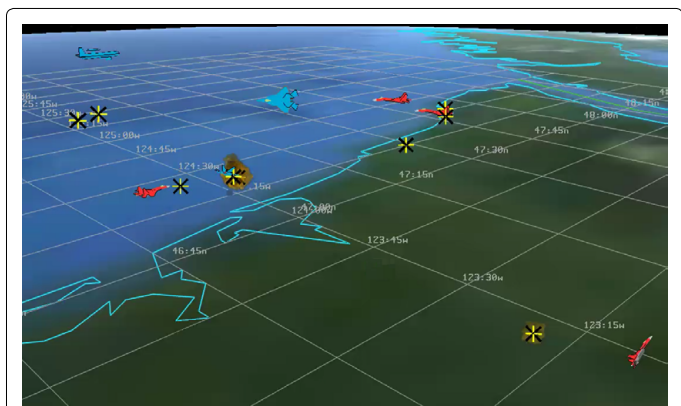


Figure 2: Side-view during active combat. Past and current missile detonation locations marked. Two Blue vs. four Red, all Reds have successfully evaded missiles, one Blue has been destroyed, Blue AWACS in distance.

The computational cost and complexity of training an AI within AFSIM is a limiting factor for many alternative methods. The very lightweight nature of the GFT, in combination with Psibernetix's efficient fuzzy logic module, PsiberLogic, and utilization of the Cython computer language in computationally costly procedures allows Psibernetix to accomplish this task with a budget desktop PC [17-19]. While it is not necessary, EVE can intelligently allocate jobs and utilize a heterogeneous computing network. Genetic algorithms and EVE are perfect examples of parallelizable algorithms; a population of ALPHA controllers are developed, their performance measured entirely independently of each other, a new population of ALPHAs are created, and the process repeats for some duration.

EVE training for ALPHA initially took place with it fighting against the baseline controller utilized previously by AFRL, but now occurs against different static versions of it. EVE here is optimizing FISs as well as other parameters such as thresholds and formations. Each ALPHA in the EVE generation calls a separate instance of AFSIM, with unique port numbers assigned to each platform. Stops are put in place to prevent extremely long missions from hanging the system. Each ALPHA and AFSIM combination takes up only one core, easing the difficulty in implementing this process. AFSIM can operate in an event-stepped fashion, running through the mission as fast as the processor allows. Figure 4 depicts the training process.

The fitness function, or equation which measures the performance of each instance of ALPHA, has been designed to completely capture the meaning of good behavior. Merely rewarding friendly kill shots and punishing friendly misses and deaths could produce a competent controller. This can be improved upon though, as it is both possible to complete a mission with the enemy just barely being unable to get kill shots off and alternatively finish missions in which no red was ever close to a blue WEZ. The two versions of ALPHA that produce these performances should be valued significantly differently. To that effect, the defensiveness of each red aircraft is constantly being taken away from the controller's performance measure. This makes EVE seek ALPHAs that complete missions in both a time and safety optimal fashion, with optimal friendly missile accuracy and minimized enemy missiles fired. Future work could contain different sensors and missions that would allow EVE to train ALPHA to defeat the entire blue force without ever being detected.

This training has led to tactical and well-performing decision making by ALPHA. More training missions to form a training portfolio are planned, along with additional static blue variants of ALPHA such as some focusing on aggression with others emphasizing defensive tactics. The fact that the intended opponent is a team of humans is an interesting complexity that is difficult to account for in the typical machine learning setting. The pilots could perform extremely strange, errant, and apparently suboptimal behaviors just to confuse or exploit ALPHA. Very lengthy, inefficient, and even suicidal maneuvers can differ significantly than the opponents ALPHA faces through training with EVE. To account for this and enable post-training testing and adjustments, a simple User Interface was created to allow humans to fly against ALPHA in-house. The human operators have been able to achieve a modest, but reasonable, amount of control over two blue aircraft to compete against ALPHA.

Results and Conclusions

As a result of EVE training, implementing lessons learned from expert fighter pilots, and adjustments from preliminary human testing, the current version of ALPHA is already a deadly opponent to face.

Formal results will be presented in future publications; this section will focus on displaying preliminary capabilities and findings of ALPHA. Again, ALPHA is currently trained to utilize a force of superior numbers, but weaker capabilities. In order to complete missions safely, the blue forces must be manipulated into poor positions, as their kill range is farther than the reds'.

As one example, ALPHA can perform lethal cooperative tactics if the opposing force allows ALPHA to pincer it. The flow of this tactic is shown in Figure 5, in 3 distinct phases. The northernmost ALPHA is designated WOLF-1 and WOLF-4 the southernmost. In the first phase, ALPHA seeks to obtain the flank by having WOLF-1 and WOLF-4 climb altitude and approach the blues at opposite and wide angles. WOLF-2 and WOLF-3 reduce velocity and climb altitude, to maintain range from the incoming blues.

The second phase begins as WOLF-1 fires a MRM to evoke a defensiveness response by blue, having no intent of actually killing its target. This missile is shot at a range in which the blue aircraft will need to evade away from WOLF-1 or be hit, but will be able to do so successfully. If fired pre-emptively, the blues can take an alternative evasion route, and if WOLF-1 delays this shot, it will be past the abort range in the incoming blues' superior WEZ. If done correctly, this forces a situation that WOLF-4 can capitalize. Shortly into the second phase, WOLF-4's launch computer reports that the enemy could easily evade a missile, but this does not take into consideration the fact that the optimal evasion route has been cut off. Two kill shots are fired from WOLF-4, and then the final phase of the tactic begins. WOLF-1 and

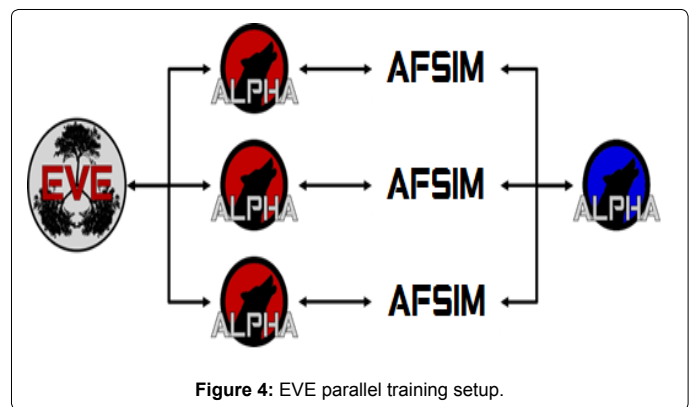


Figure 4: EVE parallel training setup.

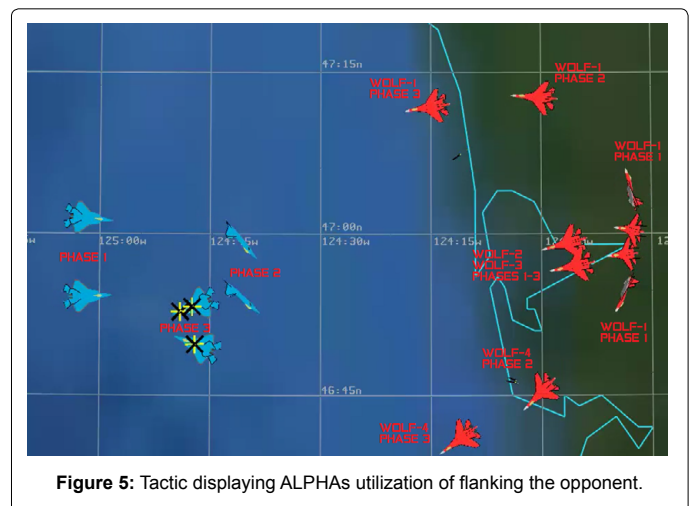


Figure 5: Tactic displaying ALPHAs utilization of flanking the opponent.

WOLF-4 maintain the flank, but keep wide approach angles in order to be able to evade any potential blue shots. If they are not fired upon, but WOLF-4's missiles miss, they will be in position to fire weapons again. By now, WOLF-2 and WOLF-3 have climbed to high altitude and increased in speed. If the flankers fail their tactic or are fired upon, the middle group can then advance and engage the blues.

In a different example, the blues shoot first, preventing the reds from getting an early positional advantage. This mission is lengthier and will be broken out into 6 less distinct phases. Phases 1 and 2 are shown in Figure 6. Due to blue's longer range, ALPHA must evade these missiles, and let blue advance. In phase 1, the southern blue, BLUE-2, fires a LRM at WOLF-4. In phase 2, the northern blue, BLUE-1, fires a LRM at WOLF-2 and WOLF-1, and BLUE-2 fires again at WOLF-4.

These initial missiles are of enough danger to ALPHA that after determining who is the target of each hostile missile, the reds must perform optimal evasive maneuvers to survive this wave. Afterwards, the blues split, with BLUE-1 heading to the north. As BLUE-2 is aggressively heading towards it, WOLF-3 maintains distance from the blues even though it was not fired at, which will prove to be a vital component to the success of the mission. This aggressive range presentation in the first phases by blue does not allow ALPHA to end the mission quickly utilizing the pincer tactic from Figure 5.

ALPHA has the ability to continue to seek positional dominance even during evasive maneuvers. For example, if evading missiles from approximately the same direction, two reds will elect to perform maneuvers in differing directions if possible. Additionally, ALPHA does not stop tracking defensiveness to all other hostiles if evading a missile nor does performing an evasive maneuver stop ALPHA from considering taking shots at all threats. ALPHA also does its best to prevent the opposing force from acquiring positional dominance during these times; it will take defensive shots to keep hostile aircraft from obtaining such an advantage.

All the initial blue missiles were defeated in before phases 3 and 4, displayed in Figure 7 below. ALPHA then takes advantage of an improper move by the blues. In phase 3, BLUE-2 has managed to get a close shot at WOLF-3, but had to over-extend to do so. This allows WOLF-4 to turn and fire a MRM. BLUE-2 begins to evade, and BLUE-1 comes to assist his wingman in phase 3, getting a shot off on WOLF-1

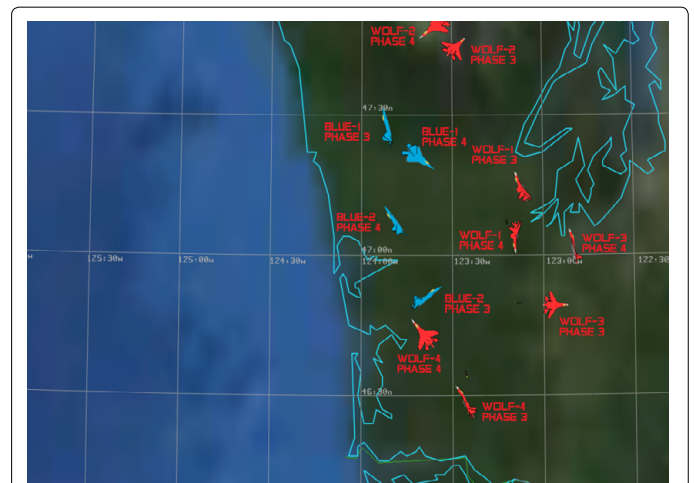


Figure 7: Phases 3 and 4 of example mission.

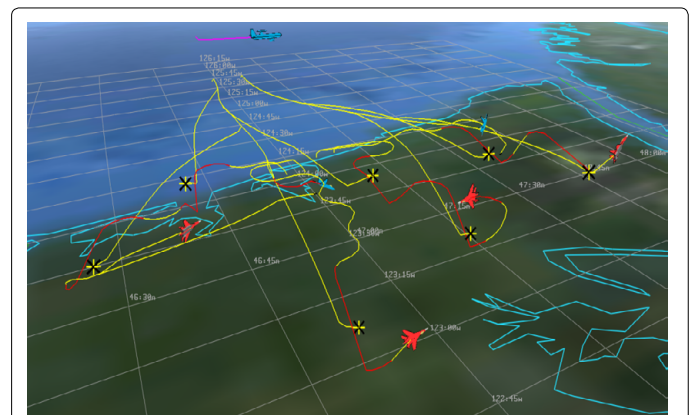


Figure 8: Side-view of example mission showing traces to phase 4.

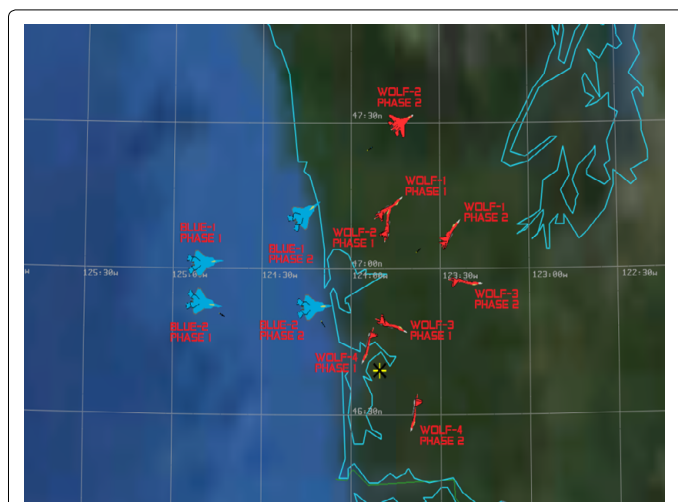


Figure 6: Phases 1 and 2 of example mission.

that would otherwise have secured a deadly pincer against BLUE-2. As BLUE-1 has had to turn back to the south, WOLF-2 does an about-face and secures the escape route of the blues. WOLF-3 successfully defeats BLUE-2's missile, and turns northward, pursuing the blues while WOLF-1 is evading BLUE-1's missile. WOLF-4 can advance and pursue the blues uncontested, and moving into phase 5, fires a MRM at BLUE-1, forcing it to turn again and evade to the north.

Note that ALPHA ends phase 4 being in a very favorable position, despite blue's early and extreme aggression. In order for the blue force to have accomplished the amount of disruption as it has thus far into the mission, a significant amount of ammunition has been utilized. Even if ALPHA were to be in a less favorable position come phase 5, the blue forces could not maintain the current pace of aggression before running out of LRMs and being at a severe disadvantage. When any of the red aircraft have a level of defensiveness above a threshold, they attempt to utilize their MRMs in a defensive manner. While ALPHA has been shot at multiple times, as can be seen clearly in Figure 8, it has taken no defensive shots. Effectively, ALPHA has not felt threatened thusly throughout this mission. While the first example mission in Figure 5 shows ALPHA taking an early strike and ending a mission quickly, this mission thus far highlights ALPHA's ability to face a highly aggressive enemy. The key to red's success going into phase 5 has been consistently forcing the blues to choose between two negative

outcomes. Had BLUE-2 not pushed aggressively towards WOLF-3, it would have had to expend even more resources to keep both WOLF-3 and WOLF-4 away; a pace it could not keep for much longer anyway. BLUE-1 could have abandoned his teammate, allowing BLUE-2 to be killed by the end of phase 4, but had it done so, red would have BLUE-1 surrounded come phase 5 regardless.

Figure 9 below displays the final phases of this mission. By phase 5, BLUE-2 has been destroyed by a second shot from WOLF-4, and WOLF-2 is to the north, keeping a safe distance but continuing to block an easy route for BLUE-1 to escape. WOLF-4's missile aimed at BLUE-1 was able to be defeated. WOLF-3 is the closest come phase 6, and once within lethal range fires the shot that kills the remaining blue fighter, ending the mission with a red victory.

ALPHA's ability to defeat AI-flown enemies is only one measure of success; it must also be able to defeat highly trained and experienced fighter pilots. ALPHA was assessed by Colonel (retired) Gene "Geno" Lee. As a former United States Air Force Air Battle Manager, Geno is a United States Air Force Fighter Weapon School graduate and Adversary Tactics (Aggressor) Instructor, and has controlled or flown in thousands of air-to-air intercepts as a Ground Control Intercept officer, as a Mission Commander on AWACS, and in the cockpit of multiple fighter aircraft.

Geno noted how the first generation of red ALPHA held its own against the blue variant of ALPHA, but the resulting engagements often ended with heavy losses for both sides. Psibernetix and Geno worked together to develop tactics, techniques, and procedures to overcome red ALPHA's payload and no-AWACS disadvantage, capitalize on blue's mistakes, and take advantage of numeric platform superiority (when the situation presented itself). The current revised red ALPHA model presented the blue adversary with credible offensive tactics and timely defensive reactions that challenged blue's radar sort logic, compressed their engagement timeline, and rapidly put blue into a defensive position from which they could not escape. The net result after a prolonged engagement was blue's total defeat with no or minimal losses by red ALPHA.

When Geno took manual control of the blue aircraft against the reds controlled by the baseline controller AFRL had previously been utilizing, he could easily defeat it. However, even after repeated attempts against the more mature version of ALPHA, not only could he not score a kill against it, he was shot out of the air by the reds every time after protracted engagements. He described ALPHA as "the most aggressive, responsive, dynamic and credible AI (he's) seen-to-date."

As stated prior, ALPHA's high performance is not its only strength;

the underlying methodologies lend themselves very well to verification and validation. Verification and validation can best be described as the process by which you can argue a particular system both satisfies its design requirements (verification) as well as accomplishes the goals of the overall system it is designed for (validation). More commonly, verification asks the question, "did I build it right?" and validation asks the question, "did I build the right thing?" As higher levels of autonomy become a reality, tests and evaluations for the purposes of verification and validation (V&V) become a primary barrier to the fielding of more advanced algorithms. The DoD Autonomy TEVV Strategy states "the most difficult and challenging component of [Autonomous Systems] is the intelligent, learning, and adaptive software embedded within them." and then states "Autonomous systems are characteristically adaptive, intelligent, and/or may incorporate learning. For this reason, the algorithmic decision space is either non-deterministic, i.e. the output cannot be predicted due to multiple possible outcomes for each input, or is intractably complex. Because of its size, this space cannot be exhaustively searched, examined, or tested; it grows exponentially as all known conditions, factors, and interactions expand. Therefore there are currently no established metrics to determine various aspects of success or comparison to a baseline state enumerated" [20]. Therefore, a key challenge to the V&V of intelligent software is to design an adaptive learning algorithm that can simultaneously provide transparency into its decisions, behaviors, and factorial affects in real time.

Fuzzy logic, by design, provides a mechanism to algorithmically reason about the world while still maintaining a clear and understandable linkage between the human designer and the system under design by leveraging the capability to synthesize linguistic, natural language requirements into real world models. Another particularly special advantage to fuzzy logic is the ability to provide real time transparency into the internal layers during and after training. Unlike other learning methods, there is essentially no "hidden layer" in specific instantiations of fuzzy logic. At any time, the GFT can be transformed into a first order logic model, translating what is traditionally referred to as "degree of truth" into a bounded range of "facts, objects, and relations." Specifically, we have developed the ability to run the FIS structures within the ALPHA GFT with the exact same algorithm that provides the transformation into a first order logic model [21]. Using a Satisfiable Modulo Theory based model checker, these first order logic models of the FIS rule structure can be verified against a set of defined properties for all possible initial conditions at once. Additionally, the FIS model can also be transformed into a hybrid system, representing both the discrete behavior as well the continuous representation of the fuzzification process within each set of active rules. Analysis to date on hybrid systems has been effective in demonstrating continuous system properties like stability [22]. However, within the ALPHA GFT, properties like stability are not as interesting as say, sensitivity or robustness. The fuzzy logic GFT to hybrid conversion can provide insight into the degree to which any given input, or n-factorial inputs, affect the output both during training and execution.

These preliminary capabilities showcase the most complex implementation of a fuzzy-logic based AI in the realm of UCAV control. Utilizing the GFT methodology, ease of verification and validation and extremely high mission performance do not come partnered with a high computational cost. Again, ALPHA can currently operate alongside AFSIM on a single 3.2 GHz core of a CPU. In this fashion, ALPHA can operate at an average frequency of 154 Hz; every 6.5 milliseconds ALPHA can take in the entirety of sensor data, organize the data and create a complete mapping of the scenario, analyze its current CoAs and make changes, or create an entirely new

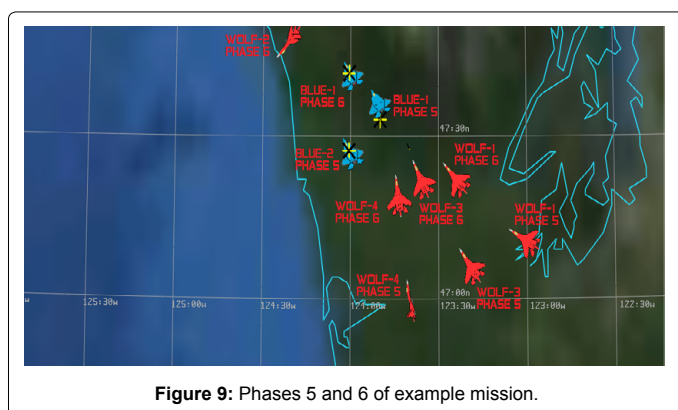


Figure 9: Phases 5 and 6 of example mission.

CoAs for a flight of four aircraft. Many opportunities still exist within ALPHA's code for increased speed optimization. A significant portion of this time is spent running UDP servers and interpreting data. If a separate process completed this task for ALPHA, such as a direct feed from sensors with their own computational capabilities, ALPHA could function even faster. This can be emulated by timing only the post data manipulation computations, which results in ALPHA averaging over 1,100 Hz for the same flight of four fighters.

Thus, ALPHA can currently operate in the domain of microseconds; one of the side-benefits of fuzzy control. This is utilizing low-budget, consumer grade products. The human mind is an extremely powerful machine that will likely always have unbeatable performance in certain areas. However, the speeds at which ALPHA can intelligently operate serve as a distinct advantage within the context of air-to-air combat.

Combining these strengths in a mixed manned and unmanned fighter squadron could prove to be an extremely effective fighting force. ALPHA controlled aircraft would happily volunteer to take risky tactics and have the manned craft perform safer support roles. Due to the transparency and linguistic nature of fuzzy logic, ALPHA would be an extremely easy AI to cooperate with and have as a teammate. ALPHA can continuously determine the optimal ways to perform tasks commanded by its manned wingmen, as well as provide tactical and situational advice to the rest of its flight.

ALPHA requires further work before its initial embodiment is considered complete. Significant strengths and capabilities still remain to be built, and further training needs to be completed. However, these early results showcase its current strength and the ability to play defensively against a stronger, but outnumbered opponent and wait for key moments to strike and become aggressive. This behavior is mostly tied to the utilization of our defensiveness measure. A different offensiveness classifier is planned to further refine when ALPHA should perform certain tactics. Developments further into the future can include capabilities for different aircraft, weapons, sensors, and mission objectives.

EVE and Psibernetix have created GFTs for problems containing extremely large solution spaces, but primarily in the bioinformatics domain of drug effectiveness prediction. The ability to bring all the strengths of the sub-methods to problems of enormous dimensionality has enabled these successes. The largest solution space EVE has trained a GFT over thus far has been 2.97×10^{961} . For this problem, a standard genetic fuzzy system would have a solution space of $10^{(3.464 \times 10^{106})}$, or many times larger than many googolplexes. The computational ceiling of the GFT method has not been discovered yet on even a single computer; as the problems ALPHA solves become more complex, the GFT will continue to scale well with the problem. The computational efficiency of both the methodologies utilized and technologies employed plays a significant part in this capability. As a less technical but still noteworthy result, ALPHA's current state as of the writing of this document was created in just months.

There is no quality of the GFT methodology that specializes it for aerial combat. Numerous potential defense applications exist for this type of AI, such as surface and undersea naval operations. While many potential application areas lack a hostile intelligent force, there are numerous applications that contain the same volume of complexities and uncertainties. The GFT is especially desirable when the problem has any of the following qualities; a need to be verified and validated for either safety or performance assurance, computational efficiency as

to not require an entire computing cluster, an inability to pre-script control due to uncertainties in the environment, or a requirement to be able to transparently monitor the system and its intent at all times. Examples include robotic surgery, design automation, and cyber security, among many others.

References

1. Clark M, Rattan K (2015) "Hybrid representation of rule-based systems". In Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control (HSCC '15). ACM, New York, NY, USA, 287-288.
2. Bellman R (1961) Adaptive Control processes: a guided tour. Princeton University Press.
3. Cordon O, Herrera F, Hoffman F, Magdalena L (2002) Genetic Fuzzy Systems: Evolutionary Tuning and Learning of Fuzzy Knowledge Bases. World Scientific Publishing Company, Singapore.
4. Cordon O (2011) A historical review of evolutionary learning methods for Mamdani-type fuzzy rule-based systems: Designing interpretable genetic fuzzy systems. *International Journal of Approximate Reasoning* 52: 894-913.
5. Ernest N, Cohen K, Kivelevitch E, Schumacher C, Casbeer D (2015) "Genetic Fuzzy Trees and their Application Towards Autonomous Training and Control of a Squadron of Unmanned Combat Aerial Vehicles". *Unmanned Systems*.
6. The Dayton Area Graduate Studies Institute, <http://www.dagsi.org/>
7. Distribution A: Approved for public release; distribution unlimited. 88ABW Cleared 08/17/2015; 88ABW-2015-4034.
8. Zeh J, Birkmire B (2014) Advanced Framework for Simulation, Integration, and Modeling (AFSIM) Version 1.8 Overview. http://www.ndia.org/Divisions/Divisions/SystemsEngineering/Documents/NDIA-SE-MS_2014-08-19_Zeh-Birkmire.pdf
9. Hanebeck U, Schmidt G (1996) "Genetic Optimization of Fuzzy Networks". *Fuzzy Sets and Systems* 79: 59-68.
10. Gegov A (2010) "Fuzzy networks for complex systems: a modular rule base approach". Springer, Berlin.
11. Liu X, Feng X, Pedrycz W (2013) "Extraction of fuzzy rules from fuzzy decision trees: An axiomatic fuzzy sets (AFS) approach". *Data and Knowledge Engineering*, 84:1-25.
12. Ernest N, Cohen K, Garcia E, Schumacher C, Casbeer D (2015) "Multi-agent Cooperative Decision Making using Genetic Cascading Fuzzy Systems". AIAA SciTech Conference, Kissimmee, FL.
13. Ernest N, Cohen K, Schumacher C (2013) "UAV Swarm Routing Through Genetic Fuzzy Learning Methods". AIAA Infotech@Aerospace Conference, Boston, MA.
14. Ernest N, Cohen K, Schumacher C, Casbeer D (2014) "Learning of intelligent controllers for autonomous unmanned combat aerial vehicles by genetic cascading fuzzy methods". SAE Aerospace Systems Technology Conference, Cincinnati, OH.
15. Van Rossum G The Python Programming Language, <https://www.python.org/>
16. Psibernetix Inc. EVE System, <http://www.psibernetix.com/about/>
17. Ernest N (2015) "Genetic Fuzzy Trees for Intelligent Control of Unmanned Combat Aerial Vehicles". Dissertation, University of Cincinnati, OH.
18. Behnel S, Bradshaw R, Citro C, Dalcin L, Seljebotn DS, et al. (2011) "Cython: The Best of Both Worlds". *Computing in Science Engineering* 13: 31-39.
19. PsiberLogic Python Package Index Page, <https://pypi.python.org/pypi/PsiberLogic/2.1.0>
20. Assistant Secretary of Defense / Research and Engineering (ASD/R&E), Autonomy Community of Interest (COI), "DoD Autonomy Test and Evaluation," Verification and Validation (TEVV) Investment Strategy, Alexandria, VA, 2015
21. Barrett C, Sebastiani R, Seshia SA, Tinelli C (2008) "Satisfiability Modulo Theories." *Handbook of satisfiability* 185: 825-885.
22. Rattan K, Clark M, Hoffman J (2015) "Design and analysis of a multistage fuzzy PID controller." *American Control Conference (ACC) 2015*. IEEE.