IMPERFECT MATES: HUMANS AND AI IN THE COCKPIT

Maj. Richard C. Agbeyibor
Flight Test Engineer, United States Air Force

Dr. Karen M. Feigh
Professor, Georgia Institute of Technology

Abstract

Rapid developments in Artificial Intelligence (AI) are bringing increasingly complex autonomy capabilities to the cockpit. Autonomous electric Vertical Take Off and Landing aircraft, swarms, Collaborative Combat Aircraft, and other new aviation mission constructs are on the horizon. In the last few decades, military and civil aviation have achieved remarkable safety and effectiveness thanks to automation and a deliberate focus on teamwork. As automation gets replaced by autonomy, the challenges of automation could be exacerbated. Effective Human-AI teaming requires both collaborative task work and teamwork which will be critical for continued safety and mission effectiveness. Despite the incredible ability of expert operators to make exceptional judgment calls in highly stressful situations, humans suffer from cognitive biases that may pose a challenge to this teaming. AI brings incredible data processing capabilities to the team but can suffer from a lack of adaptability to its environment and teammates, particularly in collaborative settings. As pilots retrain Crew Resource Management for their new AI mates, AI will also need to learn to adapt to its human mates. System developers can help achieve effective human-AI teaming by providing bidirectional transparency through interface design and system features such as status, feedback, planning mechanisms, and engagement prompts.

Introduction

AI is coming and it is going to be enmeshed in every part of Air Force operations, including inside the cockpit. The field of Artificial Intelligence (AI) has gone through several booms and busts since it first surged in the 1950s and weathered the first AI Winter in the 1970s. Over the decades since, AI has benefited from exponential growth in computational power and storage over the decades, as predicted by Moore's law. Developments in algorithm design, computer engineering, networking and other related fields have led to cloud computing. This and sustained research in machine learning, reinforcement learning, natural language processing, and computer vision have enabled the latest boom that has brought Deep Learning AI to everyday consumers.

The Air Force is investing heavily in the development of AI-enabled aircraft. In his keynote speech at the 2023 Air Force Association (AFA) Symposium, the Secretary of the Air Force unveiled a plan to build and deploy a thousand Collaborative Combat Aircraft (CCA) in the near future [1]. These AI-enabled CCAs will be designed to team with fifth and sixth-generation fighter aircraft and perform a range of different missions – from carrying electronic warfare pods to forward deploying weapons and other sensors. Beyond the immediate vision of CCA, the core enabling autonomy technologies will usher in the realization of additional human-AI missions such as autonomous refuelers, AI copilots, remotely piloted swarms, and other sophisticated forms of airborne robotics.

The field of robotics has learned over the years that AI, computer vision, mechanics, and controls alone are not sufficient. System developers must take into account how people actually do their work for systems to be safe and effective. Just as universities like Georgia Tech teach Human Robot Interaction (HRI) as a fundamental part of robotics, the Air Force should build foundations of human-centric design and human systems integration into its autonomous aircraft development.

This article aims to start conversations about interface design and software practices requisite for safe and effective human-AI teaming in aviation. In the article, the terms AI, autonomy, and robots are used interchangeably to represent agents that can sense, decide, and act independently without human input.

AI in the Cockpit

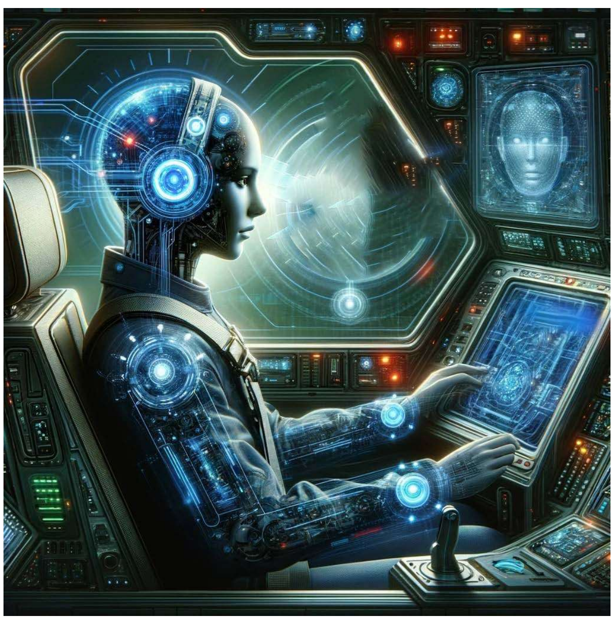

Figure 1. *AI Generated Image: AI Pilot in the Cockpit [A].*

As AI inevitably finds its way into the cockpit, automation that we've come to rely on in the last few decades of aviation for improved safety and efficiency will be replaced by autonomy. Aircraft manufacturers, like car manufacturers and other sectors already have, are keenly looking at ways to take advantage of the latest developments in AI. The requirements of cockpits and avionics systems being more stringent than enterprise systems will certainly impose constraints and require adaptation. In the airline transport sector, manufacturers and airlines are exploring ways to use AI to reduce the crew requirements from two pilots to single pilot operations [2]. In the Remotely Piloted Aircraft (RPA) sector, the Department of Defense (DoD) is exploring ways to fly several aircraft simultaneously with one ground control station crew [3]. In the fighter sector, the Defense Advanced Research Projects Agency (DARPA) and the Air Force are actively experimenting with Uncrewed Combat Aerial Vehicles (UCAV), CCAs, and other concepts [4].

Human-AI Interaction

Human-AI Interaction is the discipline that studies, designs, and evaluates autonomy, robotics, and machine systems for use by or with humans, in various domains. The process of use by or with humans is called interaction, and there are five attributes that affect the interactions between humans and AI [5]:

- Level and behavior of autonomy
- Nature of information exchange
- Structure of the team
- Adaptation, learning, and training
- Nature of the task

Despite astounding advances in computational capabilities and the complexity of tasks AI can handle, it still does not work well with humans. It does very well when pitted against humans in tasks with well-defined constraints and observable environments but not so well when paired with humans in more open environments. Researchers have developed AI agents that have learned to beat human experts in complex strategy games like Starcraft, Quake, Dota, Go, and Chess [6] but they do not do very well when asked to team with humans in simple collaboration games [7]. Carroll et al. and other researchers [7][8] have found that most AI agents naively assume perfect analytic decision making in their teammate and behave as if paired with another AI agent, unless explicitly trained with a model of human behavior.

Both members of the human-AI team are at fault for failures of collaboration. AI can be opaque, inflexible, or brittle, and humans can be too flexible or rely too heavily on heuristics or pattern matching. The human and the AI will need to learn to adapt to each other [9].

The AI Crew Member

Machines are not new to aviation. In the 1950s, a group of researchers led by Paul Fitts investigated ways to use machines for more effective air navigation and traffic control systems [10][11]. Fitts, a former Army Air Forces psychologist, is considered a founder of the Human Factors discipline [12]. In a seminal report on function allocation published in 1951, Fitts et al. "surveyed the kinds of things men can do better than present-day machines, and vice versa" [11]. That list of 11 statements became known as Fitts' list. The 11 skills surveyed in Fitts' list are: judgment, improvisation, simultaneous operations, speed and power, replication, induction, detection, perception, long-term memory, short-term memory, and computation. De Winter and Hancock, in 2015, surveyed 2,941 respondents on each of the statements of Fitts' list. According to their results, present-day humans consider that machines surpass humans in simultaneous operations, speed and power, replication, detection, perception, long-term memory, short-term memory, and computation [12] given the following statements:

- **Simultaneous operations:** "Ability to handle highly complex operations, i.e. to do many different things at once."
- **Speed and power:** "Ability to respond quickly to control signals and to apply great force smoothly and precisely."
- **Replication:** "Ability to perform repetitive, routine tasks."
- **Detection:** "Ability to detect a small amount of visual or acoustic energy."
- **Perception:** "Ability to perceive patterns of light or sound."
- **Long-term memory:** "Ability to store very large amounts of information for long periods and to recall relevant facts at the appropriate time."
- **Short-term memory:** "Ability to store information briefly and then to erase it completely."
- **Computation:** "Ability to reason deductively, including computational ability."

Data Fusion

Machines are incredibly good at processing large amounts of information as programmed. Thanks to advances in computing power and algorithms, data fusion capabilities have exploded. Data fusion is the process of integrating multiple data sources to produce better information than that provided by any individual data source. Data fusion is where AI has truly made advancements in accuracy, insightfulness, and usefulness.

Brittleness

One thing AI suffers from is brittleness – a concept which refers to AI's propensity to break, fail, or produce errors when faced with unexpected inputs or situations. And often, AI fails silently or "hallucinates" and confidently generates incorrect or misleading results [13].

Distribution shifts contribute to AI's brittleness [14]. A distribution shift is what happens when AI faces a real-world situation that is very different from what it was trained for. Types of distribution shifts discussed in AI literature are covariate shifts, concept shifts, and domain shifts.

Covariate shifts occur when the distribution of the input data changes, but the conditional distribution of the output expected remains the same. For example, an AI trained to recognize adversary aircraft on daytime images faces a covariate shift when faced with nighttime images.

Label or concept shifts occur when the distribution of the output labels changes. Say an AI is trained to recognize military installations in an Area of Responsibility (AOR) or theater of operations by a specific set of features. If the adversary changes tactics and starts camouflaging their installations or making them look like civilian structures, the AI faces a concept shift because the features it associates with military installations no longer match the new reality.

A domain shift would be taking an AI agent trained for a semi-arid desert AOR to a dense jungle AOR without retraining on the new set of features and signatures.

Hallucinations are when an AI system, usually a Large Language Model (LLM) or other Generative AI, confidently generates outputs that are incorrect or altogether misleading. It is a phenomenon where the AI agent perceives patterns that are nonexistent or imperceptible to human observers. The phenomenon is analogous to human hallucinations where one might sometimes see figures or other patterns in clouds. AI hallucinations occur due to errors in data interpretation, incorrect model assumptions, or over-fitting to the training data [15]. There can be grave consequences in the military context if AI hallucinates military targets where none are present.

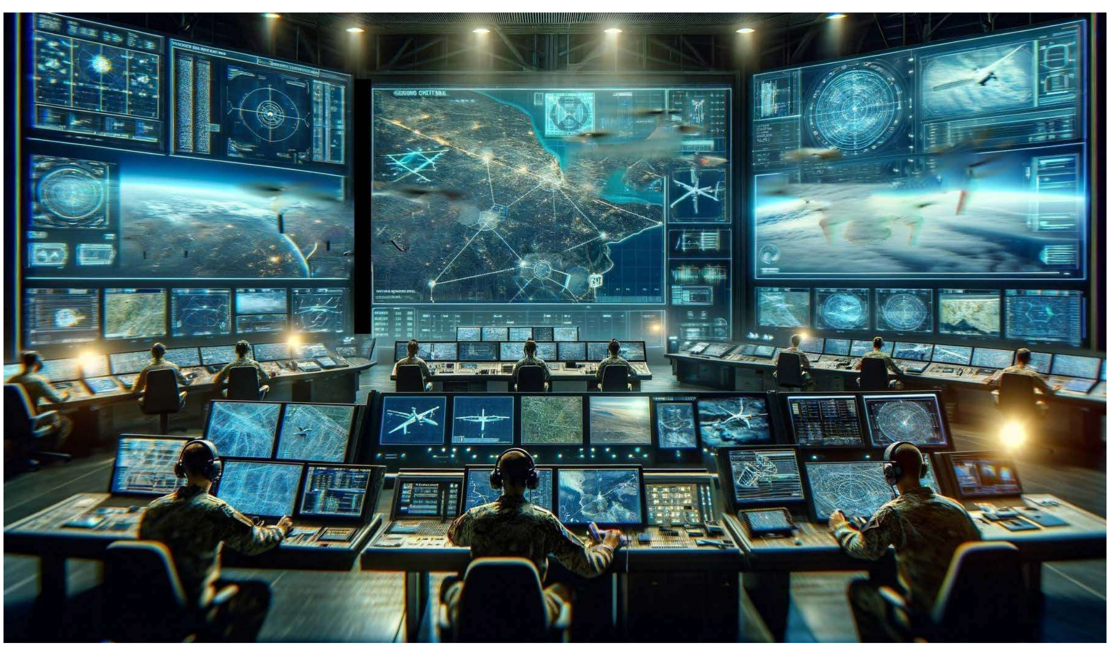

Figure 2. *AI Generated Image: Ground Control Station [B].*

Human Awareness

Although AI has learned to beat expert humans in competitive games, research has found that it generally performs poorly when teamed with humans in collaborative games [7]. There are many reasons for this. Some of the most important are the assumptions AI makes about its teammates and environments when determining its own strategy. Training of AI, particularly reinforcement learning AI, requires thousands of repetitions or examples. Since individual humans are rarely able to provide that number of repetitions, AI is often trained against Oracles (AI or Automated Agents designed to stand in for humans) using Self-play or Population-Based Training [6][8].

When co-trained with other AI, AI learns to expect predictable, analytic, optimizing decisions from its teammate. The behavioral economics, psychology, decision-making fields have shown, however, that humans are not perfectly analytic in their decision-making. Thus, when an optimal AI competes against a sub-optimal human, it can exceed expectations. In collaborative settings, however, when this same AI is teamed with a human, the performance can be drastically worse because it fails to understand and to be understood by the human [7].

The Human Crew Member

According to DeWinter et al., modern day humans consider humans to surpass machines in judgment, improvisation, and induction [12] given the following statements from Fitts' list:

- **Judgment:** "Ability to exercise judgment."
- **Improvisation:** "Ability to improvise and use flexible procedures."
- **Induction:** "Ability to reason inductively."

Aviators are trained to be cognizant of human fallacies and cognitive biases in aeronautical decision making (ADM) [16]. When these cognitive biases are checked, humans make for astonishing aircraft operators who can accomplish extraordinary feats in the most difficult of circumstances – including sparse information environments. Some of the most celebrated examples of this expert airmanship include Captain Sully Sullenberger's landing of U.S. Airways Flight 1549 on the Hudson River.

Table 1. *Modern Day Attribution of Fitts' List [12].*

Characteristic	Fitts' List Statement [11]	Modern Day Attribution [12]
Simultaneous operations	"Ability to handle highly complex operations, i.e. to do many different things at once."	Machine
Speed and power	"Ability to respond quickly to control signals and to apply great force smoothly and precisely."	Machine
Replication	"Ability to perform repetitive, routine tasks."	Machine
Detection	"Ability to detect a small amount of visual or acoustic energy."	Machine

Perception	“Ability to perceive patterns of light or sound.”	Machine
Long-term memory	“Ability to store very large amounts of information for long periods and to recall relevant facts at the appropriate time.”	Machine
Short-term memory	“Ability to store information briefly and then to erase it completely.”	Machine
Computation	“Ability to reason deductively, including computational ability.”	Machine
Judgment	“Ability to exercise judgment.”	Human
Improvisation	“Ability to improvise and use flexible procedures.”	Human
Induction	“Ability to reason inductively.”	Human

Table 1. *Modern Day Attribution of Fitts’ List [12].*

Expert Decision-Making

Intuitive decision-making by experts has been studied by multiple academic disciplines since the 1940s [17][18]. There exist many schools of thought on how it works, the pros and cons of so-called ”professional intuition.”

The field of Naturalistic Decision Making (NDM) conducts field studies on subject-matter experts who make decisions under complex conditions. They have found that some experts are able to “successfully attain vaguely defined goals in the face of uncertainty, time pressure, high stakes, team and organizational constraints, shifting conditions, and action feedback loops that enable people to manage disturbances while trying to diagnose them” [19]. This ability is often required of aircrews and has been colloquially linked to intuition and judgement of human beings.

In sharp contrast to NDM, researchers in the field of Heuristics and Biases (HB) favor a skeptical attitude toward expertise and expert judgment. In laboratory experiments, they have found that intuitive judgments are less likely to be accurate and are prone to systematic biases [19]. It is not that intuitive judgments are always incorrect, but that the noisiness, inconsistency, and unpredictability of human judgement could lead to fatal errors in a military mission.

Cognitive Biases & Heuristics

In their aeronautical decision-making training, pilots are taught to recognize and mitigate five hazardous attitudes to aviation safety: antiauthority, impulsivity, invulnerability, macho, and resignation [20]. These are but a subset of cognitive biases that manifest from the utilization of heuristics that can affect safety and mission effectiveness.

Cognitive biases are predictable but flawed patterns in people’s responses to various situations. Not all biases and heuristics are bad. Some cognitive biases and heuristics are adaptive and may lead to more effective actions in a given context by enabling fast decision-making which can be desirable when timeliness is more valuable than accuracy. On the other hand, cognitive biases may lead to perceptual distortion, inaccurate judgment, illogical interpretation, or broad irrationality [21].

Other cognitive biases that can manifest in aviation are expectation bias, confirmation bias, plan continuation bias, automation bias, and automaticity [21]. These cognitive biases, unchecked, can lead to hazardous incidents and accidents.

Bias	Definition
Expectation Bias	When we have a strong belief or mindset towards something we expect to see or hear, and act according to those beliefs
Confirmation Bias	When we only look for, listen to, or acknowledge information that confirms our own preconceptions
Plan Continuation Bias	The unconscious cognitive bias to continue with the original plan in spite of changing conditions
Automation Bias	when we over-rely on automated aids and decision support systems, or become complacent in assuming the technology is always correct
Automaticity	when routine tasks lead to an automatic response without any real consideration to what is being said or done.

Table 2. *Cognitive Biases that can manifest in Aviation [21].*

Ironies of Automation

One of the cognitive biases that can ironically lead to hazards in aviation is automation bias. Automation bias is when we over-rely on automated aids and decision support systems or become complacent in assuming the technology is always correct [21]. When automation is working correctly, people tend to become easily bored or occupied with other tasks and fail to attend well to automation performance. This is one of the ironies of automation from the operator's view of the system.

From an automation system designer's view, they may think that the human is unreliable and inefficient so should be eliminated from the system [22]. There are two ironies of this attitude. One is that the designer's own errors can become a major source of operating problems. The other is that the designer who tries to eliminate the operator still leaves him/her/them to do the task which cannot be easily automated, often without adequate support [22].

Crew Resource Management

To operate well with AI, human crew members will need to focus on these challenges in their Aeronautical Decision Making (ADM) and crew resource management (CRM) training [23].

CRM is a set of training procedures recommended by the National Transportation Safety Board (NTSB) for improving aviation safety and focuses on situation awareness, communication, leadership, and decision making in aircraft cockpits. CRM training will need to evolve to prepare human crew members for integration of AI into the crew.

Human-AI Teaming

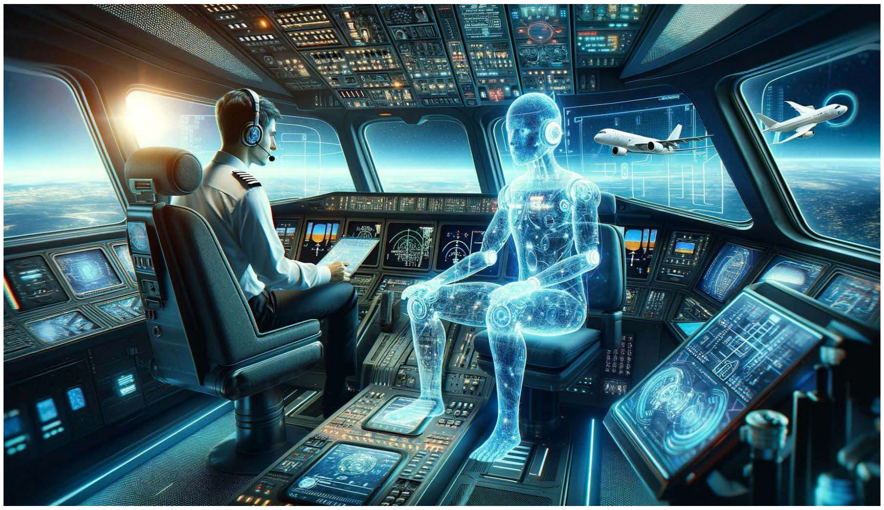

Figure 3. *AI Generated Image: Human AI CRM [C].*

According to the National Academy of Sciences (NAS), Human-AI teams can be more effective than either humans or AI systems operating alone [24]. Human-AI Teaming is a necessary construct for the cockpit environment as crew cohesion has been a key element of performance and safety in both civil and military aviation. Recognizing the critical role of crew cohesion, military protocols even relax rank-based customs and courtesies to foster seamless teamwork among crew members.

AI agents have the capacity to offer a much richer interaction mechanism than automation. With the sophistication of the information exchange and learning attributes of AI, the interaction paradigm should be changed to Human-AI Teaming [24]. As the sophistication of AI increases, so does the criticality of the functions it performs. With increased criticality of the function, consequences of errors can become catastrophic particularly since AI sometimes fails silently [13]. To help mitigate the consequences of failure, the AI's teammates must be familiar with its nominal and off-nominal behaviors.

The challenges humans have attending to automation are also applicable to AI systems, and AI systems must provide humans a mechanism for [24]:

- understanding and predicting the behaviors of the AI system
- developing appropriate trust relationships with the AI system
- making accurate decisions based on input from the AI system
- exerting control over the AI system in a timely and appropriate manner.

Transparency

The human-AI requirements enumerated by the NAS study on human-AI teaming point to a requirement for transparency [24]. Transparency represents the means of providing insightful information from the machine to the human operator and vice versa [25]. Achieving transparency can be a challenging endeavor, particularly as the complexity of the system increases. Transparency involves a bidirectional process between human and AI for mutual understandability. Joseph Lyons proposes a two element model of transparency for human AI teams: robot-to-human and robot-of-human transparency [25]. Robot-to-human transparency is information that the system needs to present to users. Robot-of-human transparency is information on the humans that the robot needs awareness of.

System designers can optimize for transparency by providing system transparency at the design phase or training the team to operate efficiently and effectively. Four system features can provide increased transparency: status, feedback, planning mechanisms, and engagement prompts [26].

Status

Status incorporates the ‘what’ of transparency, by providing the state of the human operator or system at a particular point in time. Various types of information can be provided in a status to help determine whether strategy changes by the human operator or AI agent must be initiated to accomplish a task [26].

Feedback

Feedback incorporates the ‘why’ of transparency by providing explanation, insights into actions, potential uncertainties, reliability of recommendations, and supplementary information from the human operator or machine system. Various modalities (visual, auditory, tactile) can be used to provide feedback and optimize communication within the team [26].

Planning Mechanisms

Planning mechanisms incorporate the ‘how’ of transparency and encompass the allocation of resources and task assignments among an organization’s members. Planning occurs at all mission stages and is necessary for the human-AI team to maximize its desired outcomes [26].

Engagement Prompts

Engagement prompts are cues, alerts, or warnings that encourage the human operator’s involvement. They encompass all three aspects of transparency (what, why, and how) by indicating to the human operator what must be done to resume the task, why they become disengaged, and how to identify different strategies that can be implemented to fulfill a task [26].

AI System Design

Earlier, the article discussed how incorrect assumptions by the AI about the human can lead to drastically poor team performance. To different degrees, these assumptions can be replaced by real-world information about the human teammates.

Training Paradigms

AI training can be accomplished in various ways. Some paradigms are more human-centric or human aware. Human-aware training embeds a model of a human inside the training environment. Within this paradigm, the AI is trained with awareness of the decision-making strategy of the human.

The challenge with AI training is that it requires thousands of examples spanning the full range of situations the agent may encounter. Since individual humans often are not able to provide the requisite number of decision-making examples for input to AI training, AI must often be trained against Oracles – Agents designed to stand in for humans. Oracles can either be designed to mimic humans in specific, predictable ways (like providing only correct answers 80% of the time) or can be trained using techniques like Learning from Demonstration (LfD) or Imitation Learning.

Imitation Learning

Imitation learning is a paradigm in which AI acquire new skills by learning from human demonstration [27]. Behavioral cloning, one of the simplest approaches to Imitation Learning, learns a singular deterministic policy from several expert demonstrations by directly learning a mapping from observations to actions with standard supervised learning methods. In this way, the AI can be trained repeatedly with an Oracle or model representation of human decision-making strategy.

One challenge to using techniques like Behavioral Cloning is that humans exhibit significant individual differences, i.e. individuals don't suffer from the same cognitive biases, and don't exhibit the same preferences. Another is that, often, humans teach differently based on the kind of feedback the AI is capable of taking in [28]. A significant challenge for future AI is that it will need to "get to know" its teammates and adapt to their preferences and ways of accomplishing the mission.

These are outstanding challenges for system designers as they develop AI agents to accomplish the various future mission constructs that bring AI into the cockpit.

Human-AI Mission Constructs

Mission requirements will determine how AI is implemented in the cockpit. Current crew composition and structure will form a basis for this evolution from automation to autonomy. The nature of the task, the type of information exchange required, and the availability of suitable autonomy will determine how drastic the interaction paradigm will change. This will in turn affect the way human-AI crews are structured and composed.

AI'ing Betty

Today's autopilots, voice alerting systems like 'Betty', and pilot assistance systems like the Automatic Ground Collision and Avoidance System (Auto GCAS) may merge and gain additional AI-enabled autonomy capabilities. In the near future, we may have an AI pilot assistant that collaboratively shares control of the aircraft with the human pilot during high or low workload parts of the mission.

This paradigm involves collaborative control by one human and one AI agent. This dyad relationship is the least complex and most studied human-AI team structure, however the task of controlling the same aircraft will require careful design of the interaction mechanisms.

Lessons learned from aviation incidents caused by mode error in automation should inform design of transparency features such that the pilot has suitable understanding and situation awareness when the AI agent is controlling the aircraft. The phase of flight and phase of mission will also be an important criteria in choosing robot-to-human transparency requirements.

Complementarily, the AI agent should also be aware of its own state, the environment, the phase of flight, phase of the mission, and the state of its human crewmate [9] in order to be effective. With such information, particularly information on the beliefs, desires, and intents of the human, the AI can best adapt its task work and teamwork.

Collaborative Combat Aircraft

The ongoing research efforts into today's Off-Board Sensor Station (OBSS), Off-Board Weapon Station (OBWS), and other programs are working towards the realization of collaborative combat aircraft [29]. Fighter missions may soon be accomplished with heterogeneous teams of human piloted aircraft and AI piloted wingmen.

As the number of wingmen increases, so does the cognitive load of managing the formation. Lessons learned in fighter pilot instruction, and the designated skills required of 2-ship and 4-ship Flight Leads can serve as blueprints for designing the human-AI interaction in this mission construct. The human will be unable to control each aircraft at a low level, so they will learn to command high level behaviors and learn to calibrate their expectation of the AI piloted wingmen in each mission scenario.

The AI agent piloting each wingman is going to need to simultaneously work with other AI and a human(s) in a heterogeneous, multi-agent construct. It is possible that the human crewed aircraft will be piloted by a human assisted by an AI Betty or two humans executing the role of pilot and Combat System Officer (CSO). The AI agent will need to learn to interpret commands contextually as the potential mission scenarios may be too large to explicitly enumerate and may potentially need to deconflict instructions or actions from the pilot and the CSO.

Remotely Piloted Swarms

The evolution of the Remotely Piloted Aircraft (RPA) mission into a Remotely Piloted Swarm mission is not too far away. The U.S. Army is already test launching smaller drones from bigger drones. The Air Launched Effects [30], as they're called, are controlled by the larger drone and can be numerous. Current efforts are testing singular digit numbers, but it is envisioned that one day it will increase to swarm numbers (>50). Swarm control can be a complex task depending on the type, number, and mission of the swarm [26].

Swarms are composed of large numbers of robots or drones that cooperate to achieve a goal. Swarm control can be challenging because of human capability limitations, emergent behaviors as the entities interact with each other and the environment, and constraints on communication abilities. Trade-offs exist between the number of individual swarm entities a human operator can manage and the duration of time the human operator can influence the entities [26].

Figure 4. *AI Generated Image: F-35 with CCA [D].*

The majority of human-swarm interaction literature has focused on robot-to-human transparency through visualization types and human operator influence over a swarm [26]. Research on transparency through visualization types has investigated the effect of different displays, latencies, geometries, and abstraction levels on the human operator's ability to perceive, understand, and predict swarm motion.

Robot-of-human transparency involves both control interaction and bi-directional communication. Control can be achieved through various methods of conveying operator intent, such as the use of forms of leader, predator, and mediator influence mechanisms [31]. As the human monitors status information from the swarm, individual swarm members will need status information on the human controlled mothership to

Figure 5. *AI Generated Image: Drone Swarm [E].*

execute lost link and other communication dependent behaviors in a potentially contested information spectrum.

Conclusion

We stand on the brink of a new era in aviation. The seamless integration of AI into cockpit operations and Air Force missions represents not just an advancement in technology, but a fundamental shift in the paradigm of flight operations. The potential for AI to enhance safety, efficiency, and mission effectiveness is immense, but realizing this potential requires a nuanced understanding of the delicate balance between human judgment and machine intelligence.

The shift from automation to autonomy requires not just a revolution in task capabilities but also in interaction and teamwork capabilities. Achieving effective human-AI teaming will require a collaborative effort among engineers, system designers, pilots, and AI developers to ensure that AI systems are not only capable but also compatible with human operators. This partnership must prioritize mutual understanding, adaptability, and above all, safety. As AI becomes a more integral part of the aviation ecosystem, continuous CRM training and adaptation by human operators will be essential. In turn, AI must undergo continuous learning and adjustment to effectively address issues like brittleness and hallucination, ensuring its suitability for military missions.

As the landscape of aviation evolves, so too must our strategies, tools, and mindsets. The collective goal must be to maintain the highest standards of safety and effectiveness, preserving our proud legacy of aviation while embracing the possibilities of the future. With careful planning, rigorous testing, and thoughtful integration, the synergy between humans and AI has the potential to usher in a new era of aviation.

Acknowledgements

The views expressed in this article are those of the authors and do not reflect the official policy or position of the Department of the Air Force, Department of Defense, or the U.S. Government.

Picture References

[A] OpenAI. “AI Pilot in the Cockpit.” DALL-E, openai.com. Accessed 13 March 2024.

[B] OpenAI. “RPA Ground Control Station.” DALL-E, openai.com. Accessed 13 March 2024.

[C] OpenAI. “Human AI Crew Resource Management.” DALL-E, openai.com. Accessed 13 March 2024.

[D] OpenAI. “F-35 with Collaborative Combat Aircraft.” DALL-E, openai.com. Accessed 13 March 2024.

[E] OpenAI. “Drone Swarm over City.” DALL-E, openai.com. Accessed 13 March 2024.

References

[1] Reardon, Patrick. “Watch, Read: Secretary Kendall on ‘One Team, One Fight.’” Air & Space Forces Magazine, 14 Mar. 2023, https://www.airandspaceforces.com/watch-readsecretary-kendall-on-

one-team-one-fight/.

[2] Harris, D. “Single-Pilot Airline Operations: Designing the Aircraft May Be the Easy Part.” The Aeronautical Journal, vol. 127, no. 1313, July 2023, pp. 1171–91. Cambridge University Press, https://doi.org/10.1017/aer.2022.110.

[3] Roza, David. “AFSOC Flies 3 Reapers With One Crew In First-of-Its-Kind Exercise.” Air & Space Forces Magazine, 26 Jan. 2024, https://www.airandspaceforces.com/afsoc-multiple-mq9-reapers/.

[4] Air Combat Evolution. https://www.darpa.mil/program/air-combat-evolution.

[5] Goodrich, Michael A., and Alan C. Schultz. “Human–Robot Interaction: A Survey.” Foundations and Trends® in Human–Computer Interaction, vol. 1, no. 3, Jan. 2008, pp. 203–75. www.nowpublishers.com, https://doi.org/10.1561/1100000005.

[6] Silver, David, et al. “Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm.” arXiv:1712.01815, arXiv, 5 Dec. 2017. arXiv.org, http://arxiv.org/abs/1712.01815.

[7] Carroll, Micah, et al. “On the Utility of Learning about Humans for Human-AI Coordination.” Advances in Neural Information Processing Systems, vol. 32, Curran Associates, Inc., 2019. Neural Information Processing Systems, https://proceedings.neurips.cc/paper/2019/hash/f5b1b89d98b7286673128a5fb112cb9a-Abstract.html.

[8] Choudhury, Rohan, et al. “On the Utility of Model Learning in HRI.” 2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI), IEEE, 2019, pp. 317–25. DOI.org (Crossref), https://doi.org/10.1109/HRI.2019.8673256.

[9] Feigh, Karen M., et al. “Toward a Characterization of Adaptive Systems: A Framework for Researchers and System Designers.” Human Factors, vol. 54, no. 6, Dec. 2012, pp. 1008–24. SAGE Journals, https://doi.org/10.1177/0018720812443983.

[10] Fitts, Paul M. PSYCHOLOGICAL RESEARCH ON EQUIPMENT DESIGN.

[11] Fitts, Paul M. HUMAN ENGINEERING FOR AN EFFECTIVE AIR-NAVIGATION AND TRAFFIC-CONTROL SYSTEM, AND APPENDIXES 1 THRU 3. https://apps.dtic.mil/sti/citations/tr/ADB815893.

[12] de Winter, J. C. F., and P. A. Hancock. “Reflections on the 1951 Fitts List: Do Humans Believe Now That Machines Surpass Them?” Procedia Manufacturing, vol. 3, Jan. 2015, pp. 5334–41. ScienceDirect, https://doi.org/10.1016/j.promfg.2015.07.641.

[13] “Spinning Up as a Deep RL Researcher.” Spinning Up Documentation. https://spinningup.openai.com/en/latest/spinningup /spinningup.html.

[14] Russell, Stuart Jonathan, et al. “Artificial Intelligence: A Modern Approach.” Prentice Hall, 2010.

[15] “What Are AI Hallucinations?” IBM. https://www.ibm.com/topics/ai-hallucinations.

[16] FAA Pamphlet P-8470-69: “Aeronautical Decision Making.” https://www.faasafety.gov/files/gslac/library/documents/ 2011/Aug/56413/FAA%20P-8740-69 %20Aeronautical%20Decision%20 Making %20%5Bhi-res%5D%20branded.pdf

[17] Kahneman, Daniel, and Amos Tversky. “Subjective Probability: A Judgment of Representativ eness.” Cognitive Psychology, vol. 3, no. 3, July 1972, pp. 430–54. ScienceDirect, https://doi.org/10.1016/0010-0285(72)90016-3.

[18] Klein, Gary. “Naturalistic decision making: Implications for design.” Wright Patterson Air Force Base, OH: CSERIAC, 1993.

[19] Kahneman, Daniel, and Gary Klein. “Conditions for Intuitive Expertise: A Failure to Disagree.” American Psychologist, vol. 64, no. 6, 2009, pp. 515–26. APA PsycNet, https://doi.org/10.1037/a0016755.

[20] FAA Advisory Circular 60-22: “Aeronautical Decision Making.” https://www.faa.gov/sites/faa.gov/files/2022-11/ AC60-22 Chap%201-3.pdf

[21] Chao, Elaine L., et al. “FAA Safety Briefing - July August 2020.” 2020.

[22] Bainbridge, Lisanne. ”Ironies of automation.” Analysis, design and evaluation of man–machine systems. Pergamon, 1983. 129-135.

[23] Helmreich, Robert L., et al. “The Evolution of Crew Resource Management Training in Commercial Aviation.” Human Error in Aviation, Routledge, 2009.

[24] National Academies of Sciences, Engineering. Human-AI Teaming: State-of-the-Art and Research Needs. 2021. nap.nationalacademies.org, https://doi.org/10.17226/26355.

[25] Lyons, Joseph B. “Being Transparent about Transparency: A Model for Human-Robot Interaction.”

[26] Roundtree, Karina, et al. “Transparency: Transitioning From Human–Machine Systems to Human-Swarm Systems.” https://doi.org/10.1177/1555343419842776.

[27] Ravichandar, Harish, et al. “Recent Advances in Robot Learning from Demonstration.” Annual Review of Control, Robotics, and Autonomous Systems, vol. 3, no. 1, 2020, pp. 297–330. Annual Reviews, https://doi.org/10.1146/annurev-control-100819063206.

[28] Krening, Samantha, and Karen M. Feigh. “Interaction Algorithm Effect on Human Experience with Reinforcement Learning.” ACM Transactions on Human-Robot Interaction, vol. 7, no. 2, Oct. 2018, p. 16:1-16:22. ACM Digital Library, https://doi.org/10.1145/3277904.

[29] “AFRL’s XQ-67A Makes 1st Successful Flight.” Air Force Materiel Command, 29 Feb. 2024, https://www.afmc.af.mil/News/ArticleDisplay/Article/3691289/afrls-xq-67a-makes-1stsuccessful-flight

[30] “Army Successfully Demonstrates Launched Effects System.”, https://www.army.mil/article/272675/army successfully demonstrates launched effects system.

[31] Brown, Daniel S., et al. “Human-Swarm Interactions Based on Managing Attractors.” Proceedings of the 2014 ACM/IEEE International Conference on Human-Robot Interaction, Association for Computing Machinery, 2014, pp. 90–97. ACM Digital Library, https://doi.org/10.1145/2559636.2559661.

About the Authors

Richard C. Agbeyibor is a Major and a Flight Test Engineer in the United States Air Force. He is currently pursuing a Ph.D. in Robotics at Georgia Tech. His follow-on assignment will be as an Instructor FTE at the USAF Test Pilot School. He holds a B.S. in Electrical Engineering and Computer Science from MIT, an M.S. in Electrical Engineering from AFIT, and a Master's degree in Experimental Flight Test Engineering from ISAE-SUPAERO. He is a graduate of EPNER, the French Test Pilot School. His research focuses on Human AI Interaction in Autonomous Aerial Vehicles. Contact him at richard.agbeyibor.1@us.af.mil.

Richard C. Agbeyibor

Flight Test Engineer

United States Air Force

richard.agbeyibor.1@us.af.mil.

Dr. Karen M. Feigh is a professor at Georgia Tech's Daniel Guggenheim School of Aerospace Engineering. She holds a B.S. in aerospace engineering from Georgia Tech, an MPhil in aeronautics from Cranfield University, UK, and a Ph.D. in industrial and systems engineering from Georgia Tech. Her lab, the Cognitive Engineering Center, examines human-machine interaction using tools from aerospace engineering, computer science, robotics, industrial engineering, education, and the cognitive sciences.

Karen M. Feigh

Professor

Guggenheim School of Aerospace Engineering

karen.feigh@gatech.edu.