“2, Lead You’re On Fire, and Save the Slow One for Me…”

Communication Lessons from Combat Aviation for Unmanned Wingmen

By Mark Boyer

“2.”, “Lead, you’re on fire!” and “Save the fat one for me!” are professed as the only words that the ultimate wingman would ever say, according to the fighter pilot band Dos Gringos in their ode to the fictional greatest wingman of all time. While the comedy of the song may exaggerate the limited speaking role of a young fighter pilot in the air, there is a kernel of truth regarding how communication happens in an extreme stress situation like combat aviation. One of the key skills that a fighter pilot must learn when going from a ‘regular Joe’ to Robin Olds is how to communicate with an entirely new dictionary of words, grammar of how to string them together, and an awareness of timing, content, and emphasis based upon the situation. As we begin to integrate unmanned aircraft controlled by AI agents into the role of ‘wingman’, we must carefully contemplate these communication standards to align humans and autonomy together. Here I will explore a brief history of fighter pilot communication, frame the challenges AI agents may face, and suggest a methodology for when an AI or autonomous wingman should ‘speak’.

Figure 1: Robin Olds – A legendary fighter pilot revered for his bold leadership, outstanding skill, and tremendous mustache.

A Brief History of the Wingman

In the many ways combat aviation has evolved from the swashbuckling days of the Red Baron to modern stealth fighters, communication has evolved alongside the leaps in aircraft technology. Pilots of WWI relied on visual signals and paint to distinguish friend from foe. WWII brought about a single radio in the cockpit, allowing for formations to coordinate attacks or relay basic information. As aircraft became more sophisticated, they were equipped with multiple radios, allowing pilots to communicate and coordinate at higher levels. This became a blessing and a curse – as complexity of missions grew, so did the amount of information a single pilot was required to digest from potentially 2 or 3 radios going at once. Hence, pilots began to crystallize the importance of "comm brevity" – only saying the minimum necessary. Wingmen had to stay close to their flight lead, and commands were often terse verbal orders like, "2, target West group, 25,000" to which the only response required was "2". The flight lead makes the decision for the entire formation, and the wingman simply executes the order. 'Legacy' wingmen only speak when spoken to and only speak up by exception. The primary job of the legacy wingman was to execute their assignment and watch their flight lead's back.

Around the turn of century, datalinks provided additional channels to receive even more information from a variety of sources, with today's modern fighter like the F-35 or F-22 generally having 3 radios and at least 2 datalinks. Instead of close visual formations, 'modern' wingmen are free (and required) to roam far away from their flight lead, allowing them to make many more decisions than in the past. This freedom meant wingmen must be much better at interpreting the deluge of information and synthesizing it with the complex battlefield. The primary role of the 'modern' wingman today is to digest information and make good decisions.

Wingman Prototypes

So why the brief history of combat aviation? It provides us with two contrasting models for how an autonomous wingman should communicate when placed in an uncertain, highly-time compressed situation, and how human pilots expect to communicate with a wingman. These models can help inform the AI agents we develop and the architecture of our communication standards for the generations of unmanned wingmen to come. For fun, let's call our two communication models **Eagle** and **Panther**.

Eagle is agent based upon the 'legacy' wingman described above. It follows orders perfectly within its capabilities, lets the leader know when it cannot accomplish an action, and executes preplanned actions and maneuvers well. It can minimize communication by providing a limited menu of capabilities to the leader to utilize that are well known in advance and then succinctly notifying when an action has or has not occurred. Communication is generally direct. "Execute this tactic, target this bad guy, then return to a holding pattern." While this may seem rudimentary, it significantly reduces the need for the agent to sense the environment and handle uncertainty by imposing much of the decision-making onto the leader.

There are pros and cons to a wingman like **Eagle**. This type of agent could translate into much cheaper hardware to produce, much simpler code to verify as safe and flightworthy, and much simpler systems to train with -- all good things. While the limited capabilities may hinder any single drone, the ability to deploy many at low cost may offset the relative technical limitations. Flight leads working with agents like this could easily understand, trust, and interpret the actions of Eagle because there is little uncertainty in the decision space. However, this agent cannot function well independently of the flight lead, and adding additional responsibilities to already-busy flight leads may overwhelm even the most competent of pilots. Communication bottlenecks here could focus on limited flexibility of capabilities, degradation of drone capabilities, or limited ability to query the system if anomalies arise. Designers must consider how to balance simplicity of communication with the complexity of the environment and capabilities with which the agent is bestowed.

Panther is the agent based upon the 'modern' wingman. It has much greater ability to sense the environment on its own, make decisions about best actions to take, and may suggest actions or plans rather than waiting for instructions to execute best policies. While **Panther** nominally has a set of actions in its 'playbook', it may have the ability to improvise slightly each time to best achieve what it thinks are best rewards within some operating bounds. Communication with **Panther** is much more flexible and based upon established training, norms, and planning, rather than direct communication. Rather than, "Target this guy using this tactic," communication with **Panther** is based upon general responsibilities and meeting the plan's intent such as, "Target anyone on the east side who will cause mission failure." This forces much greater responsibility on the agent to interpret the state of the battlefield, make a decision within (sometimes ill-defined) rules of engagement, and communicate that decision concisely to the rest of the formation.

This type of agent could also have many pros and cons to consider. First, this agent requires much greater understanding of the world, which means much greater sensing of the world. Additional hardware and software to sense the world increases the cost, size, weight, and complexity of a drone that can support internal sensing, even at a level much below the flight lead. Beyond hardware complexity and cost, developing software that can integrate much more information, handle large amounts of uncertainty, and then make safe, reliable decisions will take much more time and effort to develop, test, and certify for flight. With all those considerations in mind, **Panther** provides a large increase in capability to operate with much greater autonomy in the battlespace, a huge advantage for pilots already at their limits of information processing capabilities. Beyond autonomy, integration of high-level intelligence could greatly boost performance of the team for a given number of drones and crewed aircraft. The primary communication challenges here are communicating uncertainty of environment (especially when outside the training data), changing plans based upon new information, relaying degraded sensing capabilities, and communicating risks across the formation. It is expected that these more advanced drones may be equipped with datalink and/or natural language processing, so matching modality, timing, content, length, and urgency in communication will be key challenges for designers.

A Model for 'Modern' Communication

The challenge of developing communication standards for 'legacy' wingmen like **Eagle** are non-trivial and having an arsenal of highly capable but cost-effective drones will be critical for any future conflict. However, the remainder of this paper will focus on developing a general communication framework for **Panther**.

One key advantage that modern fighters like an F-35 have is the ability to rapidly develop shared situation awareness. High speed datalinks can instantly share information between aircraft so that two pilots can view the same targets, threats, and other information simultaneously on their own displays. However, often information can get lost in the clutter, sometimes there are glitches in the system, and sometimes the datalink doesn't work right. In those cases, pilots still must be able to communicate effectively, often verbally and sometimes in plain English. The key questions asked to any 'modern' wingman are: *1) <u>What</u> should you communicate? 2) <u>When</u> should you say it?*

1) *<u>What</u> should you communicate?*

When stepping out the door to a mission, pilots will all have some briefing on expected threats, targets, routes to follow, and more. Essentially, each pilot has an initial *belief* of the situation, an initial *policy* of actions to take in various situations, a *likelihood* of how likely they are to go from certain situations to others, sensors to *observe* the world, and a general *reward* for the mission (ex: BOTOTCHA – Bombs on Target, On Time, Come Home Alive). Hence, it is possible to frame combat aviation as a Partially Observable Markov Decision Process (POMDP). This POMDP could have elements such as:

States – Location of targets, Location of threats, Aircraft status, Aircraft location, Wingman status, etc.

Actions – Move towards location (X,Y,Z), target enemy at (X,Y,Z), Observe, conduct electronic attack, return to base, etc.

Transition – Targets move, ownship location uncertainty, fault likelihood, etc.

Reward – Eliminate threats (good) – get eliminated by threats (bad)

Observations – Sensors in ownship, outside information via datalink/comms

Gamma – How much risk do I want now vs. later? How much do I care about losing my wingman?

So, what should an agent communicate? Looking at our POMDP formulation, there are three good reasons to communicate, which I'll call a cue, a decision, and an outcome:

- Cue -> If a belief in the world has changed more than a certain threshold, that is the cue the agent should communicate. For example, the agent detects a new threat that wasn't part of the intel briefing, now its belief about the world has changed substantially. While it may or may not directly change the plan, this information will help all pilots update their own belief about the world and may impact future planning that could be beyond a fixed time horizon. This could be shown as some change:
$$|b_t - b_{t-1}| > \delta_b$$
- Decision -> Suppose the agent has a certain plan or policy, then it receives an updated observation of a new threat that changes its belief (i.e. a cue). That cue then triggers a change in plan, aka a decision. If this plan deviation is minor or far in the future, perhaps it isn't worth communicating about; if the plan changes substantially, then the agent should communicate this. This could be shown as some change:
$|\tau_t - \tau_{t-1}| > \delta_\tau$ (where τ =trajectory of next k actions/steps)
- Outcome -> Assuming a well-calibrated reward function, reward can change in two different ways. The first would be if belief changes but the trajectory doesn't change, for example now there is likely a hazard along the route making it more costly. The second would be that belief changes *AND* the plan changes, which could result in a new total expected reward or no change in reward if the new plan offsets belief change. In either case, the agent should communicate the associated change in reward or outcome, especially if that will change whether it accomplishes some overarching mission goals. This could be described as:
$$|U_t - U_{t-1}| > \delta_U$$

If we have a perfect reward function, theoretically all you would need to know is reward, but that is nearly impossible for real-world scenarios. Giving information of state belief, plan, and reward will help to alert the leader when it might be necessary to intervene ("No, I want you to do B not A"), help to build shared situation awareness, and foster trust among human and agent. This can all be described in Figure 2:

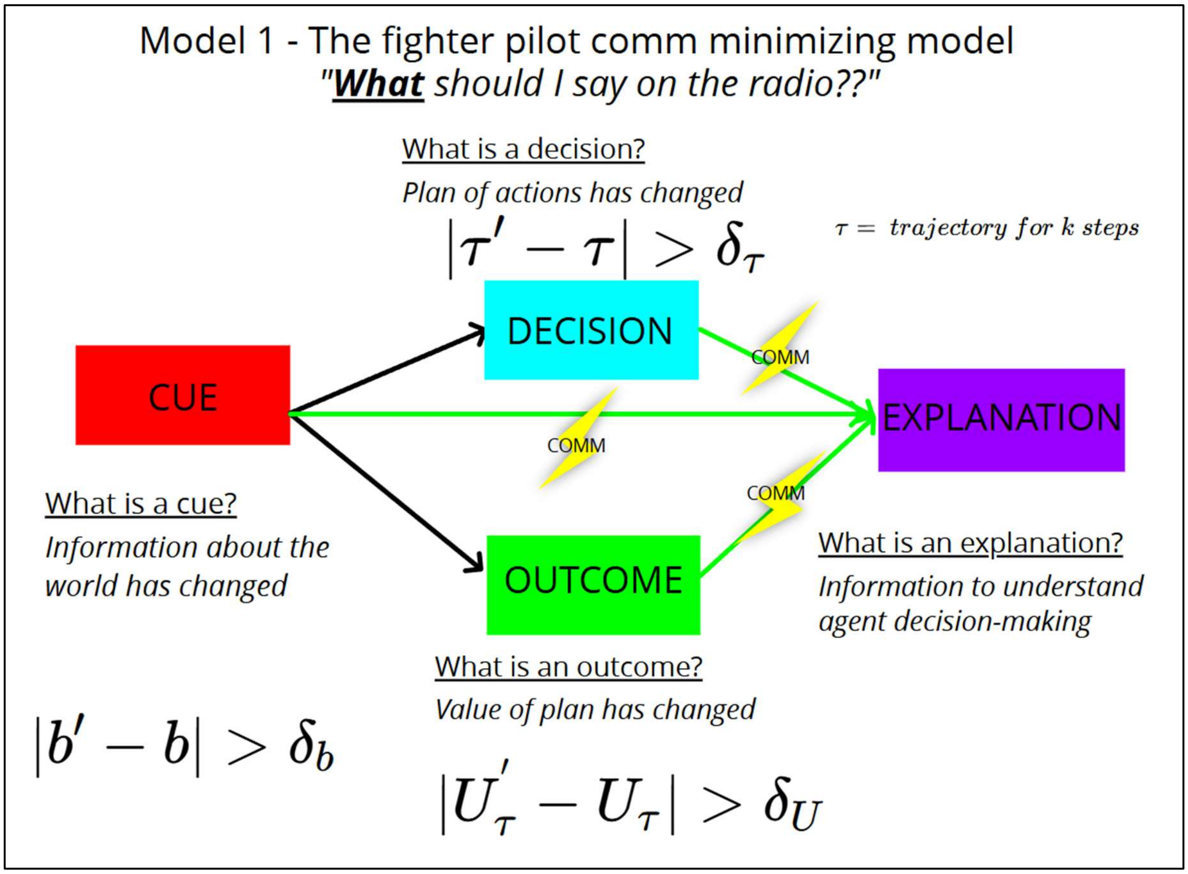

Figure 1: Model 1 of *what* can generate an explanation by agent

2) *When should you say it?*

Now that we have established what to say, the key to being a *good* wingman is knowing when to say it. Agents will be constantly updating their beliefs, trajectories, and utilities, but communicating that all the time risks overwhelming the human operator. Since we now have some notional δ for state, action, and reward at each time step, we could set some thresholds for how much change is required before speaking up. This could be a fixed value based upon heuristics, trial and error, or user preferences. The best wingmen can intuit how busy their flight lead is, how important their information is, and how much detail to use. An adaptive agent, therefore, would be able to adjust the delta thresholds based on context like communication frequency, estimated operator cognitive state, or some other metric. This feedback loop can be seen in Figure 3:

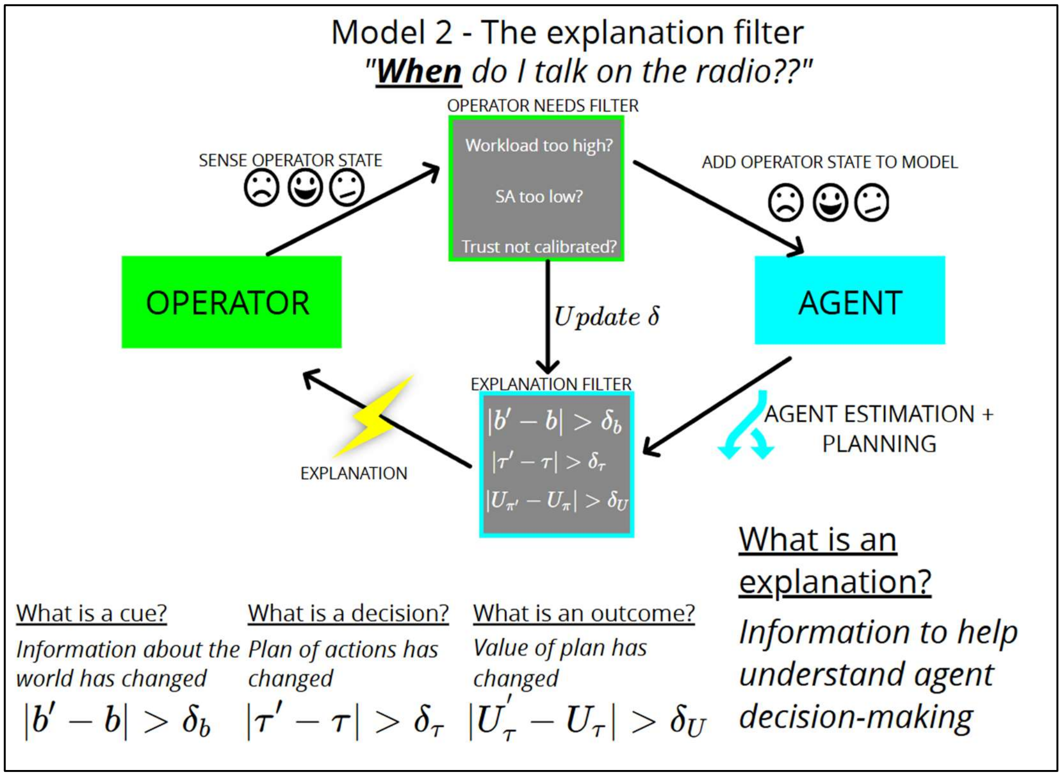

Figure 2: Model 2 of when an explanation should be delivered

An algorithm for explanation

Putting these pieces together, a general algorithm for what and when to explain can be seen below. This assumes an online planner that creates a trajectory τ with *k* steps at some discreet time step t→t+1. The agent would then generate some explanation of action using the *Explain!* function, which generates communication to the human operator.

$$
\begin{aligned}
&Initialize\ s_0,\ b(s)\forall s \in\ S,\ U\\
&s_t \leftarrow s_0\\
&while\ t < T_{end}\\
&\quad \pi \leftarrow Plan(b(s_t))\\
&\quad \tau_t \leftarrow argmax(\pi(b_t)\ for\ k\ next\ steps\\
&\quad r_t,\ s_{t+1},\ o_t \leftarrow Act(s_t, \tau_1)\\
&\quad U_{t+1} = Rollout(\tau_t)\\
&\quad b_{t+1} = UpdateBelief(b_t, o_t, \tau_1, s_{t+1})\\
&\quad \Delta_b = \parallel b_t - b_{t+1} \parallel\\
&\quad \Delta_\tau = \parallel \tau_{t_{1:k-1}} - \tau_{t-1_{2:k}} \parallel\\
&\quad \Delta_U = \parallel U_t - U_{t+1} \parallel\\
&\quad if(\Delta_b > \delta_b\ or\ \Delta_\tau > \delta_\tau\ or\ \Delta_U > \delta_U)\\
&\qquad Explain!\\
&\quad b_t = b_{t+1},\ s_t = s_{t+1},\ U_t = U_{t+1}\\
&\quad t \leftarrow t+1\\
&end
\end{aligned}
$$

Algorithm 1: Proposed explanation generation and timing algorithm

"You can be my wingman anytime"

Combat wingmen have evolved from a second set of eyes to a loyal helper to a detached teammate, and the next evolution will include unmanned wingmen operating in chaotic, information-dense environments. Tradeoffs between simplicity and capability should be at the forefront of designers' minds when considering how we develop AI agents for aerial combat. Presented here was a basic framework and algorithm for communication that takes the framing of a POMDP to answer the 'what' and 'when' of communication. Hopefully, we'll have agents that can say more than "2" "Lead you're on fire" and "Save the slow one for me!"