

# Human vs. Autonomous Control of UAV-Based Surveillance: Optimizing allocation of decision-making responsibilities

**Michael Freed**

freed@arc.nasa.gov  
NASA Ames Research Center  
Institute for Human and Machine Cognition

**Michael G. Shafto**

[mshafto@arc.nasa.gov](mailto:mshafto@arc.nasa.gov)  
NASA Ames Research Center

**Robert Harris**

rharris@arc.nasa.gov  
NASA Ames Research Center  
QSS Inc.

Airborne surveillance, widely regarded as a critical application for UAVs, often entails lengthy and largely uneventful periods of observation that strain the vigilance and morale of human remote pilots. Following the rule of “dull, dirty, or dangerous,” surveillance is thus considered an excellent application for autonomous control. Autonomy, however, should not be seen as eliminating human factors concerns but as exchanging one set of concerns for another. For example, an autonomous UAV requires specification of goals, preferences, tactics, and properties of observation targets in much greater detail than would a human pilot, and is much more sensitive to specification errors and omissions. An autonomous UAV would also tend to be less easily interrogated regarding its intentions and more difficult to redirect. And autonomous control does not necessarily imply good performance at the surveillance task. Current algorithms useful for directing surveillance effort – i.e. for deciding which surveillance target to observe next and what measurements to take at the target – vary greatly in the circumstances in which they will tend to perform well. One algorithm may, e.g., do exceptionally well at minimizing the total circuit time (time required to visit all observation targets once), but end up performing poorly in surveillance tasks where differences in the importance of observation targets imply a need to visit some more often than others. Another algorithm may do well optimizing for a small number of targets but scale poorly to tasks with a greater number. These considerations suggest that successful UAV-based surveillance should mix autonomous and human-directed decision-making, eliciting help from human decision-makers where autonomous performance is likely to fall short.

We have developed an approach to allocating responsibility for surveillance decision-making between an autonomous controller and human operator based on an empirical characterization of their relative strengths and weaknesses. The approach has three main elements: (1) a decision-theoretic framework for quantifying human or algorithm surveillance performance in a given operational scenario; (2) a set of 243 test scenarios representing a wide range of possible surveillance missions (5 independent variables, each with 3 values); and (3) software for constructing and managing scenario sets and for acquiring performance data from both algorithms and human subjects. Our initial experiments reveal systematic strengths and weaknesses in human performance relative to that of a reference surveillance algorithm. We will discuss these performance patterns, how they provide a basis for runtime allocation of decision responsibilities and efforts to apply our framework to related problems such as evaluating the effectiveness of decision-aids for human controllers.