

Authority Sharing in a Swarm of UAVs: Simulation and Experiments with Operators

François Legras¹, Arnaud Glad², Olivier Simonin², and François Charpillat²

¹ Institut TELECOM / TELECOM Bretagne, Département LUSI, France
LabSTICC, UMR 3192, CNRS, France

² LORIA Laboratoire Lorrain de Recherche en Informatique et ses Applications
UMR 7503, Université Henri Poincaré, CNRS, INRIA, Nancy 2, INPL, France
INRIA Nancy Grand Est, Equipe-projet MAIA, France

Abstract. It is emphasized in numerous prospective studies that the development of swarms of Unmanned Aerial Vehicles (UAV) should be important in the next years. However, the design of these new multi-agent systems involves to take up many challenges. In particular, reducing the number of operators requires to define new interfaces in order to interact with such autonomous multirobot systems. We present an approach that allows one operator to control a swarm of UAVs in the context of simulated patrolling and pursuit tasks. Self-organized control relying on digital pheromones, as well as authority sharing based on several operating modes are defined. Experiments with human operators on the simulated system show that the combination of the two approaches is effective.

1 Introduction

Nowadays several operators (usually one for the platform and another for the payload, not counting support) are required to supervise the mission of a single UAV (Unmanned Air Vehicle). Future systems for surveillance are envisioned using many vehicles that cooperate to perform their mission. In this context the current ratio between the number of operators and the number of vehicles will not be sustainable [9]. The existing approaches for monitoring (close to teleoperation) must be changed to augment autonomy of the system [5] in order to share the authority on various aspects of the mission.

This paper tackles the problem of the operation of a swarm of UAVs. This work was carried out within the SMAART project. This project aims at the surveillance of a military airbase. To perform such a task we consider a system composed of a set of UAVs, *i.e.* of autonomous rotary-wing aircrafts (helicopters). This system is simulated through a continuous environment where all UAVs fly at the same altitude. Each of them can perform some simple actions such as take off, landing or reaching a particular location on the base. They carry a camera that allows them to perceive their vicinity and to detect some possible intruders. We also study the coupling of the swarm of UAVs with a sensor network augmenting the surveillance system.

Providing autonomy to UAVs is a problem that can be addressed following two approaches. The first one, the most classical, consists in planning paths and then setting the UAVs to execute them [14, 11]. On the contrary, the second approach relies on the self-organization of the UAVs in function of their interaction with the environment, leading to the emergence of patrolling paths [17, 3].

The main disadvantage of the first approach is that it builds a fixed solution, that makes difficult to adapt online to changes such as the addition or removal of agents, or to respond to operator commands. Each of these changes implies to re-plan, which is subject to combinatory explosion when the number of agents or the environment size grows. At the opposite, the swarm intelligence approach relies on the autonomy of agents and their indirect cooperation via environment interactions. In this paper we aim at studying how operators can interact with such a swarm of UAVs. On the one hand, this approach allows to envision indirect interactions with the system. On the other hand, such a self-organized system naturally adapts to online changes. We also choose this approach as it is impossible for an external observer to predict the UAV trajectories. We show in this paper how an operator can, for instance, control a sub-part of the swarm without reconsideration of the whole system organization.

The reminder of the paper is organized as follows: section 2 mentions works relative to multiple UAVs control and supervision. Sections 3 and 4 present swarm algorithms to deal with patrolling and pursuit tasks as well as the associated operating modes. In section 5 we present and analyze experiments performed with human operators on a real time simulation of the proposed system. Finally we conclude in Section 6 by discussing results and proposing some perspectives.

2 Approaches for Multi-UAVs Control

On the one hand, there exists a very large amount of literature in the field of multi-agent systems devoted to enable a group of artificial agents to accomplish one or several tasks in cooperation. On the other hand, most of the research on interaction between human and semi-autonomous system focus on “single instance” systems like intelligent cockpits, industrial process control system, etc. But there is few work conducted on the human control of a multi-agent system. We present here some work related to the domain of multi-UAVs control.

Cummings *et al.* work on human supervisory control of multiple unmanned vehicles (tomahawk tactical missiles and UAVs [4]). Their work focuses more on the human factor aspects of this task (workload, number of vehicles, *etc.*) rather than on the control aspects proper. In addition, although several vehicles are involved, they are not interacting, so one cannot speak of a multi-agent system.

A widespread approach to multi-robot control is to endow them with elementary behaviors (follow a target, go to a point, patrol, *etc.*) and to allow the operator to assign behaviors to robots on an individual or group basis. The operator supervises the system of vehicles and assign new behaviors according to the context. Though, behavior-based approaches are useful only if the number of

available behaviors remains small and if the operator can anticipate their effects [18].

Control by policy can be used in conjunction with other approaches in order to allow an human operator to restrict the activity of the system through formal constraints [6]. This can be used for example for security reasons (*e.g.* avoid certain zones) or in order to help the artificial reasoning by giving partial solutions.

The *playbook* metaphor refers to a library of plans of action that are available for the operator to instantiate at various levels of detail, hence allowing various levels of autonomy for the agents. This is inspired by football teams' coaches tactics and was studied for the control of tactical ground robots [16] or heterogeneous UAVs [8].

The *Machinetta* framework [15] takes an original stance by assigning a proxy agent to each operator and each vehicle and include these proxy agents in an artificial team. Work is being conducted to apply this approach to the control of large UAV system (dozens of vehicles). But considering a human operator as "just another agent" raises important human factor issues.

3 Multi-agent Patrolling

3.1 Autonomous Patrolling

Patrolling consists in deploying several agents in order to visit at regular time intervals some defined places of an area [11]. In this way, we propose a model relying on digital pheromones. They are the computational model of chemical substances (pheromones) dropped by ants which allows them to interact. Even if ants individual behaviors are very simple, this indirect mean of communication allows them to self-organize in order to accomplish complex tasks (pathfinding, sorting, *etc.* [1, 2, 13]).

Pheromones are bound to two distinct mechanisms. On the one hand, the evaporation process which realizes a progressive fade of the information. On the other hand, the diffusion process which propagates the information across the environment. This process also exhibits the property of building a pheromone gradient usable by the agents.

In order to realize the patrolling task, we adapt to the UAVs specificities the EVAP model [3] which only exploits the evaporation process (*cf.* algorithm 1). This algorithm is initially defined for theoretical discrete environments (grid of cells covering the environment). EVAP relies on the environments marking: an agent drops a fixed pheromone quantity Q_{max} when it visits a cell. The evaporation process makes this value decrease so that the remaining pheromone quantity represents the elapsed time since the last visit of the cell (called *idleness* of the cell). So, on the cells set, local gradients appear giving the direction of the highest idleness cells (*i.e.* with the lowest pheromone quantity). As a consequence, the behavior of the agents is defined as a descent of the pheromone gradient. Therefore, it ensures locally the patrolling of the cells which have not been visited for the longest time (*cf.* Fig. 1).




Fig. 1. 3D illustration of the EVAP algorithm (with only one agent). Pheromone field altitude represents elapsed time.

Algorithm 1 EVAP Algorithm

EVAP AGENT

- A) Find a cell y in $Neighbors(x)$ such that $q(y) = \min_{w \in Neighbors(x)} q(w)$
in case of multiple choices do a random choice
- B) Move to cell y
- C) Set $q(y) \leftarrow Q_{max}$ (drop the Max quantity of pheromone)

EVAP ENVIRONMENT

For every cell x of the environment

If $q(x) \neq 0$ then $q(x) \leftarrow \rho \cdot q(x)$

($\rho \in]0,1[$)

Unlike EVAP agents, SMAART UAVs move in real coordinates according to a given azimuth and a given speed. So, we keep the cells matrix in overlay of the real environment only for the pheromones deposit and perception. Agents can now perceive the environment through a r radius disk. They are thus able to choose their destination among the cells that belong to this disk (which defines $Neighbor(x)$). The pheromone is dropped on all the cells of $Neighbor(x)$.

Moreover, at the environment level, we add to each cell its own evaporation rate ρ_x . It allows us to modify evaporation speed of some cells dynamically, therefore creating priority zones.



Fig. 2. Base patrolled by a set of UAVs dropping a visit pheromone: initial exploration phase (left) and stabilized phase (right). Brighter color means more pheromone.

Figure 2 shows how SMAART UAVs patrol over the environment (the simulator was developed in java using Madkit/Turtlekit framework [12]). Agents first start an exploration phase of the environment. When all the cells has been visited once, the system tends to stabilize towards an average performance where agents are homogenously distributed inside the environment (see EVAP algorithm performances study in [3]). Moreover it was shown in [7] that the system self-organizes in (sub)optimal cycles of same length. UAVs individual behavior although remains unpredictable, which is desirable in the frame of the surveillance of a military area.

3.2 Operating Modes for Patrolling

The main task of the operator in his/her day-to-day activity is to supervise the surveillance of the airbase by the patrol of the UAVs. The objective is to make sure that every point of the airbase is regularly scanned by an UAV. This can be done homogeneously or with some emphasis on certain zones of the airbase according to the operational context. In this activity, the operator has to evaluate the state and evolution of the coverage of the airbase (potentially neglected zones, *etc.*) and act accordingly on the UAV system. He/she has several operating modes for this observation task:

- observation of the position of the UAVs across the area;
- reading on the HCI (Human–Computer Interface) of the computed average and maximum idleness values for particular zones (landing strip, hangars, tower, *etc.*);
- analyzing a color-gradient representation of the idleness grid (it can be configured so that only idleness values above a certain threshold are represented, *e.g.* show points with an idleness above 3 minutes).

According to the operational context, the operator can assign different priorities to certain zones of the airbase (landing strip, hangars, tower, *etc.*). If the operator sets a higher (lower) priority for a zone, the evaporation value is locally raised (lowered) which leads to a quicker (slower) disappearance of the pheromone and incites UAVs to visit the zone more (less) often.

The trajectories of the UAVs for patrolling can be determined according to two operating modes:

- the UAVs can follow the modified EVAP algorithm and decide their direction according to the local patrol pheromone level (this is their default behavior);
- the operator can assign a subset of the UAVs to a set of positions specified by the operator on the airbase. In this case, the UAVs are dealt over the positions and adopt their default behavior once they reach their respective position.

4 Pursuit

4.1 Autonomous Pursuit

Patrolling the environment is not sufficient to ensure the interception of intruders. This task aims at carrying on the search over a limited area in the case that an intruder has been spotted and contact has been lost. When an UAV perceives an intruder, it drops a second type of pheromone (alarm pheromone) which diffuses locally. This diffusion represents a disk of probability of the intruder presence which all the UAVs may use in order to find it again (see Fig. 3). So, as soon as a UAV perceives some alarm pheromone, it climbs its gradient and consumes it. As a consequence, UAVs move towards the signal origin first to consume the isolines of the alarm pheromone field.



Fig. 3. Illustration of the use of the alarm pheromone by an UAV.

The propagation of the information through the environment allows to attract other nearby UAVs towards the search zone. It aims at improving the interception probability of the intruders before they reach their objective. It is necessary to tune correctly the diffusion coefficient in order to avoid attracting too many UAVs and letting parts of environment without surveillance (the diffusion tuning is conducted empirically for each environment).

The evaporation process determines the duration of the signal. If the signal vanishes before the UAVs have consumed all the pheromone, they revert to the patrolling task without having taken advantage of the pheromone trace. On the contrary, it is useless (and even penalizing) to keep obsolete information that may mobilize some agents to find an already gone intruder. We were able to establish analytically the Q_{max} value and the evaporation rate in order to size up the alarm pheromone propagation.

We define here a second surveillance mode, joint to the UAVs, and based on the use of a sensor network (as proposed in [10]). In fact, the number of required UAVs to ensure a good patrolling rate may be important and therefore too expensive. So, sensors may be placed randomly on the environment or along the border. Each sensor is able to trigger an alert by dropping a given quantity of the same alarm pheromone. UAVs can either use their alarms or the sensor ones for the tracking task.

We do not present here a testbed for the fully autonomous system but rather focus our attention in the next section on the human–UAVs swarm interaction

4.2 Operating Modes for Pursuit

When an alarm is raised in the perimeter of the airbase, the operator switches his/her activity: surveillance/patrolling become secondary, while it becomes crucial to intercept the potential intruders that have triggered the alarm(s). An intruder is considered intercepted if the UAVs manage to detect him/her several times in a row (otherwise it gives just another contact and an indication on the position of the intruder). The role of the operator is twofold:

- he/she has to analyze, identify, interpret and classify the contacts (alarms). A contact can be a false alarm (animal, malfunction of a sensor) but also the sign of a well-prepared multiple coordinated intrusion. In such a case, it is vital to correctly interpret the pattern of contact;
- he/she has to supervise the deployment of the UAVs that have to search and intercept intruders.

The SMAART HCI assists the operator in these roles. First, in order to avoid alarm proliferation on the HCI, alarms are aggregated into *contacts* on a temporal and spatial basis (time and distance thresholds) *i.e.* alarms that are raised almost simultaneously at the same location are considered as one contact generated by the same intruder or group of intruder. That way, if a group of three intruders are detected by an UAV and a ground sensor at the same time, only one contact is generated instead of six distinct alarms.

Second, a module in the HCI allows to organize contacts temporally to represent intrusions *i.e.* successive contacts can be linked to represent the hypothesis that they have been generated by the same intruder or group of intruder. Figure 4.a illustrates how this is represented on-screen. Several operating modes are available to the operator for this task:

- the system can classify a new contact automatically by ranking the different hypotheses (create a new intrusion for this contact or affect it to an existing one) and choosing the highest ranked. The ranking is computed by linking the likelihood of an hypothesis to the speed needed by the hypothetical intruder;
- the system gives a time delay to the operator to perform the classification manually before applying the highest ranked hypothesis;
- the operator can perform the classification manually but the system assists him/her by presenting the hypotheses ranked by likelihood.

Third, the trajectories of the UAVs for patrolling can be determined according to three operating modes:

- the UAVs can follow the modified EVAP algorithm and decide their direction according to the local patrol or alarm pheromone level (this is their default behavior);
- the operator can assign a subset of UAVs to a set of positions specified by the operator on the airbase. In this case, the UAVs are dealt over the positions and adopt their default behavior once they reach their respective position.




Fig. 4. a. An intrusion illustration: the larger the pie part, the more recent the contact. b. Intrusion scenario.

- the operator can also use contacts or intrusions as intermediary objects to dispatch UAVs rather than to specify positions on the map. UAVs are dealt over selected contacts or along a selected intrusion (privileging more recent contacts).

5 Experiments with Operator

5.1 Protocol

In order to evaluate the interaction between human operator and UAV swarm, that is allowed by the different operating modes of the SMAART HCI, a series of experiments with human operator were conducted on a simulator.

Subjects The subjects are eight cadets of the French Naval Academy (École Navale de Lanvéoc-Poulmic) aged 20-23. They are anonymized and referred to as X1a, X1b, X2a, X2b, X3a, X3b, X4a, X4b according to the number of their session (1-4) and their console (a or b).

Scenario The objective of the experiments is to evaluate the quality of the SMAART-operator interaction and the usage of the different operating modes in the context of the two main activities: surveillance and pursuit. The subjects are confronted with a three-parts fifteen minutes scenario: (1) 20' of pure surveillance without intrusion; (2) 5' consisting of a false alarm that allows operators to become familiar with pursuit operations; and finally (3) 25' consisting in a coordinated intrusion by three two-men commandos toward the aircraft hangars, as shown in fig 4.b. The airbase is equipped with a linear sensor network along the border and 10 UAVs with one hour autonomy.

Qualitative Data After each fifty minutes experiment, the subject is interviewed by the experimenter along a pre-defined questionnaire. This is the occasion to collect information about the subjects (knowledge of different domains

like UAV systems or real-time strategy games) and their evaluation of different characteristics of the system. Answers to specific questions are coded as a number between -1 and $+1$, but all free commentaries and remarks are recorded and transcribed afterwards. These are precious in human-centered design loop, but we will not present them in this paper but rather hint at some of them in section 6.

Quantitative Data During the operator’s activity within the scenario, several types of data are logged by the system:

- IMI (Instantaneous Matrix Idleness) at each step in order to evaluate the quality of the surveillance;
- actions of the operator on the HCI;
- intruders progress: reaching waypoints and objectives, being intercepted.

In order to evaluate the performance of the operator on IMI and intruders progress, we averaged the results from 12 runs of the system without operator with the same scenario.

5.2 Surveillance Performance

Figure 5 shows the relative performance (idleness) of the system in autonomous mode (thick line) and of the operators (continuous and dotted thin lines).




Fig. 5. Surveillance performance with and without operator: thick line was obtained without human operator (lower idleness values are better).

Two phases appear: (1) during the first 20’ there is no intrusion on the airbase and the objective is only to optimize the surveillance (minimize idleness); (2) during the second phase intruders must be intercepted, surveillance becomes secondary.

During the first phase, we observe that the operators’ intervention does not improve the surveillance: on the contrary on average operators induce a -3.2%

decrease in performance (individuals vary between +6.5% and -13%). One can make a striking relation with the following facts from the interviews: (1) the subjects judge the autonomous system’s performance on surveillance only as “somewhat good” (+0.5 in a $[-1, +1]$ range); and (2) they are unanimous to affirm that human intervention is critical for the surveillance, even though on average it was measured that it had an adverse effect!

In the second phase of the scenario, the intrusions and subsequent alarms disturb the surveillance, which worsens the idleness. This phenomenon is visible with or without operator (see Fig. 5). One can note that the deterioration is much worse for three of the operators (X3a, X4a, X4b, with dotted thin lines on Fig. 5). These operators have massively dispatched the UAVs for the pursuit. Even though, the log files show us that while X3a is the best interceptor (he caught all 6 intruders), X4a and X4b are among the worst despite their intense activity (number of actions on the HCI). It seems that massive intervention by the operator is not a positive factor, neither for the surveillance nor for the pursuit of intruders.

5.3 Pursuit Performance

During the sessions with human operators, 2.1 intruders out of 6 (on average, with a mean deviation of 0.7) reached their objectives. During the 12 runs without operator (autonomous mode) 3.4 intruders out of 6 (on average, with a mean deviation of 1.1) did the same. This corresponds to a 20% increase in pursuit performance.

We observe here a very positive effect of the actions of the operator on the system. The global analysis capabilities of the human make a good combination with the local processing of the UAVs. The operator is able to act on a strategic level (affecting UAVs between the intrusions and the patrol task, anticipating the intruder’s next move, *etc.*) while the UAVs perform the actual interception once they are on-site. As a side-note, only two out of the eight operators used contacts or intrusion objects to dispatch UAVs (see Sec. 4.2), the others specified positions manually (using the mouse on the map).

6 Analysis and Perspectives

This paper presents only a part of the experimental results obtained in the SMAART project. Most notably, it lacks a thorough analysis of the individual interviews. Those gave insight in the human–system interaction and spurred recommendations and ideas for future systems. Hereafter, we address two broad topics that arose during these experiments.

6.1 Interpretative Complexity for the Operator

Human operators experience difficulties to judge the performance of the pheromone-based patrol algorithm. This is clearly shown by the discrepancy

between the performance measured and evaluated by the subjects themselves. However human operators are needed as they have a very positive impact on pursuit, and the algorithm can face some problems *e.g.* temporary pheromone islets can appear that could be ignored by the UAVs if the system is not supervised.

These difficulties can likely be lessened by instruction and training of the operators, but the representation gap between a human operator and a pheromone-based swarm seems to call for new interaction and representation tools for the HCI, beyond the thresholded view implemented in SMAART.

6.2 Towards an Extension of Control Modes

The human subjects were able to significantly increase interception rates by positioning UAVs across the airbase, notably on the predicted path of intruders. But operators experienced frustration due to the default behavior of the UAVs *i.e.* upon reaching their designated position, they revert to patrol in the absence of alarm pheromone. The operators had to resort to repeating the same orders to keep UAVs on position.

It would be interesting to add the possibility of confining subsets of the UAVs to specific zones. This would allow a kind of control by policy (see Sec. 2) combined with the current modes. These zones could be manually defined by the operator, computed based on a contact and on an estimation of an intruder's speed, or by a combination of these two modes.

7 Conclusion

This paper dealt with authority sharing between human operator and an UAVs swarm for patrol and pursuit tasks. The swarm algorithm was based on the environment marking by digital pheromones (*i.e.* the EVAP algorithm [3]) and perceptions were augmented with a sensor network. The EVAP model was extended (i) to allow an operator to influence the swarm behavior and (ii) to simulate UAVs in continuous space. Then a number of different operating modes at various levels of autonomy were implemented as an approach to authority sharing *via* adjustable autonomy. Experiments with human operators have shown that although the human has a positive role to play in the control and supervision of the automation, the representation gap between human and swarm intelligence calls for more advanced HCI tools.

This work was conducted in the context of a Exploratory Research and Innovation contract (*Recherche Exploratoire et Innovation: REI*) of the French Defense Procurement Agency (*Délégation Générale pour l'Armement: DGA, Mission pour la Recherche et l'Innovation Scientifique: MRIS*).

References

1. R. Beckers, O.E. Holland, and J.-L. Deneubourg. From local actions to global tasks: stigmergy and collective robotics. In *Artificial Life IV: Proc. of the 4th Int.*

Workshop on the synthesis and the simulation of living systems, third edition, MIT Press, 1994.

2. E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. New York, Oxford University Press, 1999.
3. H. Chu, A. Glad, O. Simonin, F. Sempe, A. Drogoul, and F. Charpillet. Swarm approaches for the patrolling problem, information propagation vs. pheromone evaporation. In *ICTAI'07 IEEE International Conference on Tools with Artificial Intelligence*, pages 442–449, 2007.
4. M. L. Cummings and P. J. Mitchell. Operator scheduling strategies in supervisory control of multiple UAVs. *Aerospace Science and Technology*, 2006.
5. S. Dixon and C. Wickens. Control of multiple-UAVs: A workload analysis. In *Proceedings of the 12th International Symposium on Aviation Psychology*, 2003.
6. M.C. Dorneich, S. D. Whitlow, C. A. Miller, and J. A. Allen. A superior tool for airline operations. *Ergonomics in Design*, 12(2):18–23, 2004.
7. A. Glad, O. Simonin, O. Buffet, and F. Charpillet. Theoretical study of ant-based algorithms for multi-agent patrolling. In *Proceedings of the 18th European Conference on Artificial Intelligence ECAI*, pages 626–630, 2008.
8. R. Goldman, C. Miller, P. Wu, H. Funk, and J. Meisner. Optimizing to satisfy: Using optimization to guide users. In *Proceedings of the American Helicopter Society's International Specialists Meeting on Unmanned Aerial Vehicles*, Chandler, AZ, January 18-20 2005.
9. C. Johnson. Inverting the control ratio: Human control of large, autonomous teams. In *Proceedings of AAMAS'03 Workshop on Humans and Multi-Agent Systems*, 2003.
10. J. M. Kahn, R. H. Katz, and K. S. J. Pister. Next century challenges: Mobile networking for "smart dust". In *International Conference on Mobile Computing and Networking (MOBICOM)*, pages 271–278, 1999.
11. F. Lauri and F. Charpillet. Ant colony optimization applied to the multi-agent patrolling problem. In *IEEE Swarm Intelligence Symposium*, 2006.
12. F. Michel, G. Beurier, A. Gouaich, and J. Ferber. The turtlekit platform : application to multi-level emergence. In *ABS 4 Agent-Based Simulation 4*, 2003.
13. L. Panait and S. Luke. A pheromone-based utility model for collaborative foraging. In *Proc AAMAS'04*, pages 36–43. ACM, 2004.
14. H. Santana, G. Ramalho, and R. Bohndana V. Corruble. Multi-agent patrolling with reinforcement learning. In *the 3th international Joint Conference on autonomous Agents and Multi-Agent Systems*, pages 1122–1129, 2004.
15. Paul Scerri, Katia Sycara, and M. Tambe. Adjustable autonomy in the context of coordination. In *AIAA 3rd "Unmanned Unlimited" Technical Conference, Workshop and Exhibit*, 2004.
16. R. Simmons, D. Apfelbaum, D. Fox, R. Goldman, K. Haigh, D. Musliner, M. Pelican, and S. Thrun. Coordinated deployment of multiple heterogeneous robots. In *Conference on Intelligent Robotics and Systems*, Takamatsu Japan, 2000.
17. I. Wagner, M. Lindenbaum, and A. Bruckstein. Cooperative covering by ant-robots using evaporating traces. In *Technical report CIS-9610, Center for Intelligent Systems, Technion, Haifa*, 1996.
18. M.S. Wilson and M.J. Neal. Diminishing returns of engineering effort in telerobotic systems. *IEEE Transactions on Systems, Man and Cybernetics - Part A:Systems and Humans*, 31(5):459–465, September 2001. Special Issue on Socially Intelligent Agents - The Human in the Loop.