Artificial Intelligence and Social Computing, Vol. 72, 2023, 11–21
https://doi.org/10.54941/ahfe1003269

Automated Decision Support for Collaborative, Interactive Classification

Randolph M. Jones, Robert Bixler, Robert P. Marinier III, and Lilia Moshkina

Soar Technology, 3600 Green Court, Suite 600, Ann Arbor, MI 48105, USA

ABSTRACT

Traditional classification approaches are straightforward: collect data, apply classification algorithms, then generate classification results. However, such approaches depend on data being amply available, which is not always the case. This paper describes an approach to maximize the utility of collected data through intelligent guidance of the data collection process. We present the development and evaluation of a knowledge-based decision-support system: the Logical Reasoner (LR), which guides data collection by unmanned ground and air assets to improve behavior classification. The LR is a component of a Human Directed and Controlled AI system (or "Human-AI" system) aimed at semi-autonomous classification of potential threat and non-threat individuals in a complex urban setting. The setting provides little to no pre-existing data; thus, the system collects, analyzes, and evaluates real-time human behavior data to determine whether the observed behavior is indicative of threat intent. The LR's purpose is to produce contextual knowledge to help make productive decisions about where, when, and how to guide the vehicles in the data collection process. It builds a situational-awareness picture from the observed spatial relationships, activities, and interim classifications, then uses heuristics to generate new information-gathering goals, as well as to recommend which actions the vehicles should take to better achieve these goals. The system uses these recommendations to collaboratively help the operator direct the autonomous assets to individuals or places in the environment to maximize the effectiveness of evidence collection. LR is based on the Soar Cognitive Architecture which excels in supporting Human-AI collaboration. The described DoD-sponsored system has been developed and extensively tested for over three years, in simulation and in the field (with role-players). Results of these experiments have demonstrated that the LR decision support contributes to automated data collection and overall classification accuracy by the Human-AI team. This paper describes the development and evaluation of the LR based on multiple test events.

Keywords: Decision support, Collaborative reasoning, Classification, Unmanned autonomous systems, Mobile sensors, Urban operations, System of intelligent systems

INTRODUCTION

This paper describes the development and evaluation of a knowledge-based system called the Logical Reasoner (LR), which is the primary reasoning component in a larger ***Human Directed and Controlled AI system*** (ISOLATE, Intelligent Sensing to prObe and Localize Adversaries in Threatening Environments) that uses autonomous vehicles and sensors to interact with a population of individuals, gather information, and generate accurate classifications of people exhibiting threat-like behavior.

ISOLATE's mission is uncertainty reduction through threat classification: to help Mission Commanders (MCs) identify individuals in a population who are most or least likely to present imminent threat. The system operates in largely unknown environments on short time frames (1-3 hours), without sufficient time to observe patterns of life (Folsom-Kovarik et al, 2014). Standard approaches to classification acquire data ("evidence") from the population and use a classification algorithm to generate results. But that approach depends on data and evidence being amply available, which is not true for the ISOLATE mission.

As opposed to passive evidence collection, ISOLATE has limited resources (time, assets, sensors, etc.) and must gather evidence actively. While the human MC has the ultimate control over asset management, evidence interpretation, and final threat/non-threat classifications, the Logical Reasoner supports the MC by using knowledge to help make productive decisions about resource allocation.

This paper introduces the system of intelligent, robotic, and sensor systems that comprise the ISOLATE architecture and then discusses the LR's significant role within that architecture. The paper then discusses the LR's challenges, goals, design, and technical details of the eventual solution. There is a particular focus on the encoding and structure of knowledge that the LR uses to compensate for the lack of data accessible to the statistical modules of ISOLATE. The paper concludes with a discussion of the deployment of ISOLATE and the LR in numerous experimental exercises, and lessons learned from the results of those exercises.

THE ISOLATE ARCHITECTURE AND MISSION

ISOLATE is a Human-AI system aimed at semi-autonomous classification of potential threat and non-threat individuals in a complex urban setting. The setting provides little to no pre-existing data; thus, the system collects, analyzes, and evaluates real-time human behavior data to determine whether or not the observed behavior is indicative of threat intent. To achieve this mission, ISOLATE comprises a complex, integrated, collaborative system, including numerous interacting components for operator-system interaction, robotic control, sensor deployment, data analysis, sensemaking, planning, execution, and other functions (see Figure 1). The LR (see following sections) adds knowledge to the process and enhances Human-AI collaboration. The overall system operates with limited resources (time, assets, sensors) and actively uses them to gather evidence (data) to support such determination. The Human-AI team must decide:

1. On whom to gather evidence
2. From where to collect the evidence
3. When to gather the evidence
4. What kind of evidence is most beneficial to gather
5. How to combine the collected data into an evidence chain

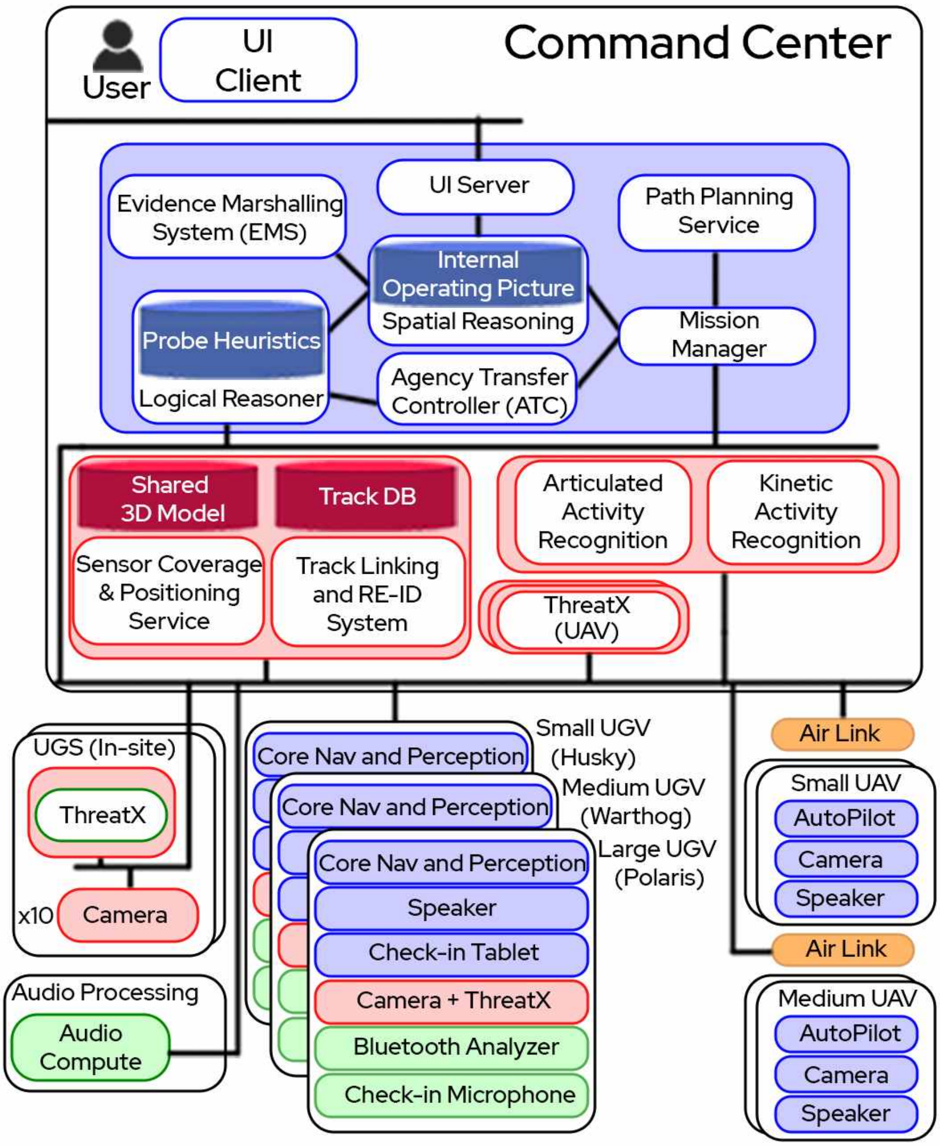

Figure 1: The ISOLATE architecture is an integrated system of software systems, robotic systems, and sensor platforms that are intended to interact with a population, gather data, and ultimately help MCs to classify individuals who may represent threats.

ISOLATE collaborates with a human Mission Commander (MC), whose goal is to observe the situation as it unfolds, direct data collection (in conjunction with the AI teammate) and produce a final classification of potential threat/non-threat intent for each observed person. The system collects data through a combination of stationary and mobile cameras on autonomous air and ground platforms. These robots not only serve as sensor platforms, but also interact with members of the population (e.g., providing news updates, handing out water, asking for ID, etc.) to better understand their behaviors and actions. The MC determines which vehicle actions are most appropriate for reducing uncertainty within a given context. Data collected by the mobile and stationary platforms and sensors is processed to determine spatial relationships, individual activities, and ultimately allows the MC to make a threat/non-threat classification of each observed individual.

ISOLATE's AI core is a collection of mobile and stationary sensors, gathering data from observations of a populated environment, and a set of modules to process this data and make sense of it. Activity recognition module parses observed human movements into kinetic and articulated "activities" such as "move away", "move toward", "run", etc. which are in turn aggregated into more complex spatio-temporal and contextual pattern activities. In addition to carrying sensors, the mobile platforms have multiple ways in which they can interact with the population. These interactions are called "probes". Examples are to broadcast a news report, play music, say hello, ask to see an ID, give orders to clear the area, etc. The mobile platforms are autonomous ground and rotary-wing air vehicles, and their probing capabilities can involve movement, audio, taking pictures, dispensing items, etc.

With these probing capabilities, classification of data migrates from being a passive to an active process. Instead of just observing the behaviors and activities of population members, ISOLATE creates population interactions and collects observations of how individuals respond to those interactions. This interactive data collection process feeds into an "Evidence Marshalling System", which uses decision-theoretic Bayesian statistics to estimate the evidence and uncertainty that each individual exhibits behavior indicative of a potential threat (Shafer, 1976; van Oijen & Brewer, 2022).

When a decision has been made to perform a particular probe in a particular location, there are additional modules for path planning, scheduling, and monitoring to move the mobile platform to the desired location and execute the probe actions. The combined set of maneuver and probe actions, that a mobile platform needs to perform, is called a "probe plan".

The job of selecting which probe plans to execute is performed collaboratively by the LR and the MC using a Graphical User Interface (GUI). The goal of the collaborative team is (in the mission time allotted) to perform probes that maximize the quantity and quality of data collected to correctly categorize as many individuals as possible (with a bias against incorrectly classifying potential non-threats as threats). In the current version of ISOLATE, MCs can use a high-level delegation and monitoring tool (Agency Transfer Controller, ATC) to delegate one or more vehicles to be under system control. The ATC supports the DoD *governability* principle of Ethical AI, allowing the MC to maintain meaningful control and situation awareness of the system while delegating certain functions to it. In this manner, assets can receive probing commands both from the MC and from the system. The LR provides probe plan recommendations to the MC and the ATC. The ATC then applies delegation and other ethical constraints to its probe plan selection using LR recommendations.

It is these interactive and dynamic features that distinguish the ISOLATE mission from a "standard" empirical classification task. The collection of statistical evidence is still important to classification, so there are components of ISOLATE that focus on that task. But the sensors cannot provide full instantaneous coverage; they can only collect observations in their local areas, so they must be allocated with care to increase the quality and quantity of collected data. Because data collection depends in part on the nature of each type of probe, the types of interactions each probe can provoke, and expectations about how individuals with and without threat intent might react to different probes, this allocation problem depends on knowledge of the probes themselves and how populations interact with them. It is the job of the LR to encode this knowledge and employ it collaboratively to assist the MC in selecting probe plans and interpreting the evidence they collect.

RESEARCH, DEVELOPMENT, AND ENGINEERING GOALS

As ISOLATE is a research system, the design and implementation of the LR was aimed to achieve research goals, as well as mission goals. As described above, at the mission level the most important goal is to reduce uncertainty by improving MC threat/non-threat classifications. The LR's role in achieving this goal is to make effective recommendations about how assets should be allocated and which probe plans should be selected as the mission progresses. An additional goal is for the LR to support the Human-AI team collaboration.

However, because we did not know at the beginning of the project which knowledge-based methods would most effectively improve outcomes, the design was also constrained and informed by the need to make it as flexible and extensible as practical to explore a variety of modes of collaboration, interaction, and value-added decision making. The following subsections present some of the engineering goals, together with a description of our solution approach during LR research and development.

Supporting Reusability and Extensibility

From its beginnings, the LR was not just a research system, but a component in a much larger and complex, engineered system of systems. This implied that we needed to attend to software engineering goals, to enable research and exploration without needing to redo the engineering work when exploring different approaches. Thus, the design of the LR is reusable and extensible. We implemented the LR in the Soar Cognitive Architecture (Laird, 2012; Laird, 2022) to exploit decades of prior research in engineering knowledge-based systems. We designed the knowledge and decision-making representations so we could easily incorporate new heuristic knowledge or change existing knowledge. We also engineered the LR's role within ISOLATE so that it serves as a "value-added" module, rather than a critical bottleneck. That is, the LR can be entirely disabled, and the rest of the ISOLATE system continues to function; it is just a "smarter" system when the LR is enabled.

Supporting Alternative Modes of Collaboration

One of the key research goals of the LR was to explore modes of Human-AI collaboration. Recall that the job of the LR is to use knowledge to make useful recommendations for resources allocation in an ever-changing situation.

There were three primary variations we explored for LR collaboration with the MC:

- Fully manual: The LR would provide recommendations to the MC, but all final decisions would be made by the MC.
- Hybrid: The MC would be the primary decision maker, but the LR would be allowed to initiate probe plans autonomously in contexts for which the MC gave permission using the ATC module.
- Fully autonomous (proof-of-concept only): All decisions to launch new probes would be made by the LR, without involving the MC.

We approached these goals by allowing the LR to *contribute* to decision making, without necessarily being *critical* to decision making. The LR continuously computes its estimates of the most effective probe plans to execute. When given permission through ATC, it launches the best rated probe plans with ATC's high-level constraint monitoring. Otherwise, it provides recommendations to the MC. Whether the choices are made by the LR or the MC, the LR monitors situation and mission status to update its recommendations as the mission progresses.

LR DESIGN AND IMPLEMENTATION

To meet the combined mission and engineering goals, we developed a modular design with reusable *task patterns* that can be employed across the types of decisions the LR makes. One key aspect of this design was to identify which of the LR's decisions would always be made autonomously and which could involve Human-AI collaboration. We also kept in mind future proofing, in case requirements ever were to change. The resulting design consists of a simultaneous collection of decision-making tasks that run continuously and cascade their decision-making results to produce recommendations. The set of implemented LR tasks is depicted in **Figure 2.** The tasks highlighted

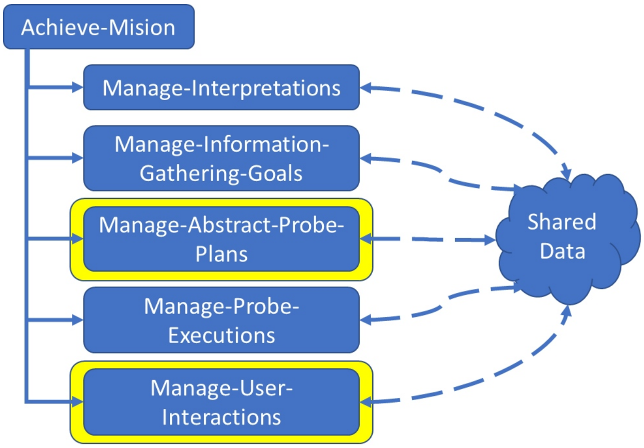

Figure 2: The LR continuously recomputes information in the service of multiple tasks that inform each other. Tasks highlighted in yellow currently allow input/preferences from the human MC. The others are currently autonomously updated but could involve human input in the future.

in yellow are those that are currently achieved in collaboration with the MC, the others being autonomous and updated continuously as the mission progresses.

To summarize, each time the LR receives a new "snapshot" containing information about the current state of the mission, it revises a set of knowledge structures in response to the most recent changes:

- Based on the history of activities and collected evidence, the LR generates new interpretations about the behaviors of members of the population (e.g., "This person tends to avoid interacting with our probes.")
- Based on current interpretations and evidence, as well as other mission factors, the LR identifies the best information-gathering goals to pursue (e.g., "Try to collect more evidence against individual I23 being a potential threat", "Try to reduce the uncertainty associated with individual I942", "Gather more information from the marketplace")
- Combinations of information-gathering goals can most effectively be pursued with different types of probe plans, so the LR selects and ranks which probe plans are predicted to be most effective. An "abstract" probe plan is a plan that has not yet had its details filled in (such as which mobile platform should execute the plan, what the planned route is, etc.).
- After a probe plan has been assigned and initiated, the LR monitors its execution and factors its status into future rankings (e.g., a probe plan may increase in desirability if it is close to a current probe).
- At all times, the LR may receive new information from the MC and may have new information to pass to the MC (through the ISOLATE GUI). There is a task dedicated to managing these interactions. (e.g., "The MC gives permission to initiate probes involved the UGV42", "The MC wants particular attention paid to individual I547").

At the end of each snapshot cycle, the LR updates the current set of suggested probe plans, together with numerical rankings of their desirability. Each recommendation is associated with evidence, interpretations, goals, and user interactions, so when the system is in semi-autonomous mode, the MC can view the explanation of each autonomous probe conducted via ATC.

The knowledge employed by the LR consists of numerous heuristics for each type of decision, implemented as a set of relational patterns. The patterns are efficiently recognized and triggered by a truth-maintenance system (de Kleer, 1986) that is part of the core of the Soar Cognitive Architecture.

LR Knowledge Content

As we have discussed, a typical machine-learning or decision-theoretic approach would compute accurate classifications from large amounts of data. A key feature of the ISOLATE mission is that data is too sparse to generate purely statistical classifications with low uncertainty. Thus, it is necessary to use an interactive approach for data collection, which suggests the need to introduce knowledge into the process. This does not mean we ignore statistics completely, because statistical results from how individuals react with the different types of probes are a useful source of knowledge. But these statistics are insufficient, so we supplemented the knowledge in LR with expertise acquired from the subject matter experts (SMEs) and from heuristic assumptions concerning the hypothesized reactions of individuals presenting potential threat vs not-threat to the different types of probes.

We acquired and engineered expert knowledge into the LR using a set of knowledge-acquisition and knowledge-engineering methods developed for past projects using large knowledge-based systems (e.g., Jones et al., 1999; Moshkina et al., 2019). These methods tease out the expertise of the SMEs who have performed similar intelligence missions in operational environments. We were able to identify several important factors, such as what types of goals the MCs want to achieve at different phases of the mission, what types of factors they use to decide which assets and probe plans to select, what their favored modes of interaction are with the user interface and automated systems, how to trade off data collection with the potential for escalating threats, etc. By intensively interviewing several SMEs, we were able to "round out" the heuristic knowledge base with goals, decisions, and recommendations that are informed by the MCs' collective expertise.

We engineered the resulting knowledge into the LR design to create pattern matchers that infer interpretations, information-gathering goals, and suitable probe plans to achieve those goals, as well as to rank and prioritize probe plan recommendations. We also engineered this heuristic knowledge to be as modular and as data driven as practical, which paid dividends in allowing us to revise and update the knowledge base in response to lessons learned from experiments. There is not space in this paper to describe the final knowledge base in detail, but Table 1 gives a representative sample of the types of information-gathering goals the LR considers and the combined evidence and interpretations that trigger their pursuit.

LR SYSTEM EVALUATION

We evaluated the LR across numerous evaluation events over the three-year period of the project. The evaluation events were of two types: simulation testing in Unreal simulation environment and live test events with human role players in urban environments. Evaluations consisted of post-hoc analyses of data across completed missions per event. As the LR evolved and mixed autonomy levels varied, we used different metrics to evaluate performance.

We initially evaluated the LR as a recommender system, measuring mean probe rank percentile computed for each probe selected by the MC as the percentage of probe recommendations ranked lower than the selected probe. We also used a subjective MC questionnaire after each mission to rank usefulness and frustration of the delegation using a Likert scale from 1-7, with 7 being very useful/frustrating.

As we developed the system further and introduced alternative mixes of collaboration, we introduced new metrics. We observed that better coverage of an area could lead to more information gathered and more accurate classifications, so we also computed the percentage of area covered by probes

Table 1. Sample of LR information-gathering goals together with a description of the situations in which the LR considers pursuing the goals.

Information-Gathering Goal	Triggering interpretations and conditions
Allow-low-escalation Allow-medium-escalation Allow-high-escalation	Computed from current population harm, mission time remaining, and how close entities are to potential threat identification.
Establish-baseline	Computed from population information and time since mission start.
Provide interaction opportunity	Triggered when entity has high uncertainty.
Provide compliance opportunity	Triggered when entity has threat evidence above a medium threshold or a record of avoiding interaction.
Flush target from crowd	Triggered when an entity close to potential threat identification is in a crowd.
Flush target from building	Triggered when an entity close to potential threat identification is not visible and seen going into a building.
Resolve entity uncertainty	Triggered when entity has low uncertainty but is not classified.
Probe multiple targets in area	Triggered when multiple targets in one place have the same candidate probe plan.

Table 2. LR system evaluation metrics and values.

Measure	Probe Rank Percentile	Mean Usefulness / Frustration (MC rating)	Mean Time Processing	Area Coverage
Result	65	3.66 / 2.54	< 6s	65%

sent autonomously. Results for these metrics are included in Table 2 (based on the last event).

The probe rank percentile was not as high as desired, with similarity across different probe types being a potential cause. The probe rank percentile tended to be higher toward the beginning of the mission and fall with mission duration. This suggests there is less variability in decision making at mission start (which the LR has captured), but variability increases as novel situations unfold. The mean usefulness is slightly above average, suggesting autonomous probe launching decisions are useful without causing undue frustration, though with room for improvement. We also measured processing speed: the amount of time it took to process a snapshot and produce the ranked list of probes; based on the results, the LR was close to the desired mean processing time of 5 seconds or less. This matches the generation time of the information snapshots. If the LR takes longer, probe recommendations may be slightly "stale" – up to 5–6 seconds is acceptable. As for coverage, it is difficult to determine empirically the ideal coverage of a mission area, which depends on its size, number of robotic systems tasked to that area, and entity density. However, a coverage of 65% seems close to the ideal, which we estimate to be in the 75-85% range, as greater resource allocation likely results in diminishing returns and increased population frustration.

Delegation of assets to the system through ATC, with LR providing probe recommendations, also resulted in more evidence collected, which helped improve overall uncertainty reduction, producing more accurate MC classifications in the final field evaluation event. Figure 3 presents this information for a subset of mission runs where assets were delegated to the system for some portion of the time: the greater ATC usage, the more accurate the picture produced by the MCs.

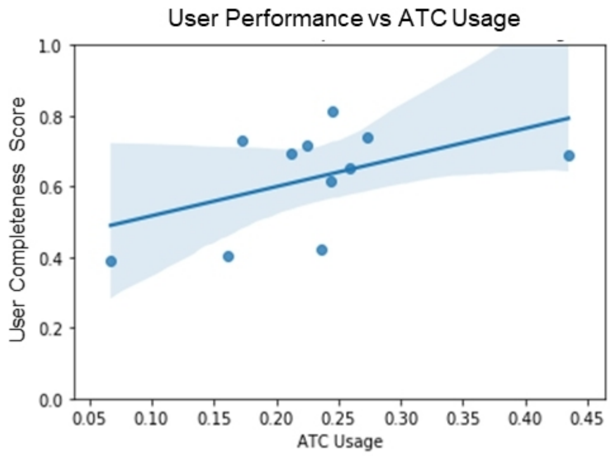

Figure 3: ATC usage (delegation of assets to system) plotted against overall uncertainty reduction based on MC-finalized classifications. Completeness score (Y axis) reflects the percent of correctly classified entities (threat or non-threat) minus those incorrectly identified entities.

CONCLUSIONS AND FUTURE WORK

Over the course of the entire ISOLATE project, we learned several valuable lessons from our efforts in prototyping, redesigning, and reconfiguring the LR:

- For research, development, and deployment, it is essential to have a reconfigurable reasoning framework that enables easy additions and/or changes to the knowledge base.
- The scarcity of available data inherent in this classification mission requires the addition of heuristic knowledge.
- A collaborative decision support system such as the LR must be smoothly and easily adaptable to varying levels of mixed autonomy.

Next steps for the LR will focus primarily on improving heuristic knowledge, as well as extending the use cases for collaborative mission execution. To improve the heuristic knowledge, we can broaden the knowledge relevant to how individuals interact with different probes and use that knowledge to explore new types of probes that may gather more or better information. With each experiment we run with ISOLATE, we are also collecting new statistics reflecting how people interact with the probes "in the wild". Further analyses of this data will help us fine-tune the heuristics to accommodate not just the theory behind individual interactions, but the practical empirical realities that the experiments reveal.

ACKNOWLEDGMENT

The research reported in this document was performed under Defense Advanced Research Projects Agency (DARPA) contract #HR001120C0180, Urban Reconnaissance through Supervised Autonomy (URSA). The views, opinions, and/or findings expressed are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon. Many thanks to Kris Kearns and the entire ISOLATE R&D team.

REFERENCES

de Kleer, Johan (1986). An assumption-based TMS. *Artificial Intelligence*, vol. 28, no. 2, pp. 127–162.

Folsom-Kovarik, J. T., Schatz, Sae, Jones, Randolph M., Bartlett, Kathleen, & Wray, Robert E. (2014). AI challenge problem: Scalable models for patterns of life. *AI Magazine,* vol. 35, no. 1, pp. 10–14.

Jones, Randolph M., Laird, John E., Nielsen, Paul E., Coulter, Karen J., Kenny, Patrick, and Koss, Frank V. (1999). Automated intelligent pilots for combat flight simulation. *AI Magazine*, vol. 20, no. 1, pp. 27–41.

Laird, John E. (2012). *The Soar cognitive architecture*. MIT Press. Cambridge, MA.

Laird, John. E. (2022) Introduction to Soar.https://arxiv.org/abs/2205.03854

Moshkina, Lilia, Saucer, Timothy, Spinola, Matt, and Crossman, Jacob. (2019). *Variable fidelity simulation and replay for unmanned autonomous ground vehicles*. SAE Technical Paper 2019-01-1074.

Shafer, Glenn (1976). *A mathematical theory of evidence*. Princeton University Press,

van Oijen, Marcel, and Brewer, Mark (2022). Probabilistic Risk Analysis and Bayesian Decision Theory. Springer.