Proceedings of the 2025 Winter Simulation Conference
E. Azar, A. Djanatliev, A. Harper, C. Kogler, V. Ramamohan, A. Anagnostou, and S. J. E. Taylor, eds

HUMAN AI METACOGNITIVE ALIGNMENT IN AGENT-BASED DRONE SIMULATION: A NEW ANALYSIS METHODOLOGY FOR BELIEF UPDATING

Bowen Sun[1], Jiahao Wu[1], and Jing Du, Ph.D.[1]

[1]Dept. of Civil and Coastal Eng., University of Florida, Gainesville, FL, USA

ABSTRACT

Effective decision-making for autonomous drones relies on timely belief updating, yet the mechanisms governing this process in AI under uncertainty are poorly understood. This study presents a framework integrating Bayesian inference, evidence accumulation theory, and Dynamic Time Warping to analyze belief updating in autonomous agents. In simulated urban search and rescue scenarios, we demonstrate that drones consistently prioritize key environmental variables like wind conditions. We also find the duration of belief updates is inversely related to the magnitude of environmental change, mirroring adaptive human cognition. A novel dual-threshold model separates internal belief shifts from observable actions, offering a clearer interpretation of the agent's internal state. Our approach contributes to developing more transparent, adaptive, and cognitively aligned AI for complex, high-stakes environments.

1 INTRODUCTION

Agent based simulations are widely utilized across various domains including disaster response, autonomous navigation, infrastructure management, and industrial systems due to the inherent complexity of real-world scenarios and the necessity of robust decision-making (Chen et al. 2022, Binz and Schulz 2023, Hagendorff et al. 2023). Such simulations uniquely enable the capture of agent behaviors under uncertainty and partial observability, offering insights often missed by traditional methods. However, current analytical approaches largely overlook the internal cognitive processes of simulated agents, particularly their belief updating mechanisms (Gunning and Aha 2019). Belief updating refers to an agent's continuous internal revision of its understanding of the environment based on incoming data, profoundly influencing its decisions. Existing analytical methodologies predominantly emphasize observable agent outcomes or external performance metrics, thus inadequately capturing the complexity of these internal cognitive dynamics (Kawato and Cortese 2021).

With advancements in computational capabilities, breakthroughs in deep reinforcement learning and cognitive modeling, along with progress in high-performance simulation platforms and multi-source data acquisition methods, agent-based simulations have rapidly advanced, achieving substantial improvements in complexity and realism. Nevertheless, a critical gap remains in methodologies explicitly designed to analyze internal agent cognition. Current analytical approaches primarily focus on observable agent outcomes, often neglecting systematic investigation into the timing and triggers of internal belief transitions. This lack of transparency may undermine trustworthiness, reliability, and the effectiveness of human-AI interactions, particularly within safety-critical applications (Rago and Martinez 2024).

This study proposes a simulation based analytical methodology explicitly targeting the hidden belief updating processes within agent based simulations. The proposed method integrates Bayesian inference to quantify uncertainty, Dynamic Time Warping (DTW) to analyze temporal relationships between environmental stimuli and agent responses, and evidence accumulation modeling using Ornstein-Uhlenbeck process (Pisauro et al. 2017). Additionally, a dual threshold Bayesian calibration informed by human subject experimental data effectively differentiates observable behavior changes from subtle internal cognitive recalibrations.

We present a framework that uncovers how AI agents update beliefs, focusing on when internal shifts occur and how they relate to observed actions and environment. To do this, we apply DTW to reveal hidden delays between environmental changes and agent responses, showing when cognitive states start to shift. We also use evidence accumulation modeling with dual Bayesian thresholds to distinguish subtle internal recalibrations from clear behavioral changes. Validated through drone simulations in urban search and rescue scenarios, this approach can improve the clarity of agent based simulations by making formerly inaccessible belief processes more visible. It also increases robustness and reliability in complex environments and supports deeper human AI alignment in high-stakes operational settings. Figure 1 shows the overall framework of our metacognitive calibration research, the paper mainly focuses on belief update stage.

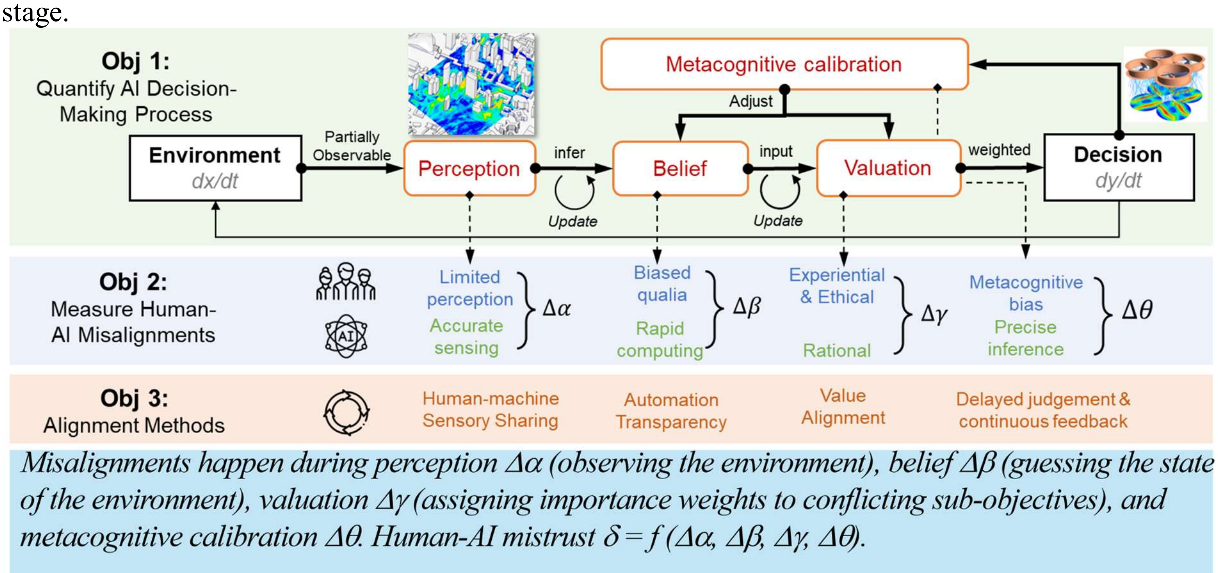

Figure 1: Overview of metacognitive calibration framework.

2 RELATED WORK

Belief updating is a fundamental component of decision making in both human cognition and artificial intelligence (AI), enabling agents to adjust their internal understanding in response to new information (Anderson 1991). Humans accomplish this by intuitively integrating prior knowledge, heuristics, and cognitive biases, which allows effective adaptation under uncertainty (Frömer and Nassar 2023). However, translating such human-like flexibility into AI systems is difficult because AI relies on explicit algorithms, lacks innate intuitive reasoning, and must manually incorporate contextual information (Adadi and Berrada 2018). Bayesian inference offers a rigorous way to model belief updating by combining prior beliefs with incoming evidence (Bissiri et al. 2016, Khalvati et al. 2021), yet it frequently encounters computational challenges in high dimensional environments. Approximate methods such as variational inference or Markov Chain Monte Carlo can mitigate these challenges but may introduce errors (Charniak 1991, Korb and Nicholson 2010), and accurately defining prior distributions remains challenging, especially in scenarios characterized by sparse or noisy data (Murphy 2012).

In dynamic and uncertain environments, belief updating becomes even more demanding. Humans employ heuristics and contextual cues effectively but can experience performance declines under cognitive overload (Byyny 2016). Traditional AI approaches, such as Kalman and particle filters, are useful for state estimation yet often require complete and accurate models of the environment, a condition seldom realized in real world scenarios (Thrun 2002, Simon 2006). These methods are also computationally intensive and tend to be algorithmically rigid, limiting their adaptability compared to human intuition (Kaelbling et al. 1996). Explainable AI (XAI) methods aim to enhance transparency in AI decision making processes, but current techniques like feature attribution, saliency mapping, and rule extraction focus on static or post-hoc explanations (Adadi and Berrada 2018, Rudin 2019). Although recent "chain-of-thought" methodologies illustrate intermediate reasoning steps, they remain primarily descriptive and do not quantitatively capture how beliefs change over time (Kojima et al. 2022).

Conceptualizing AI agents as cognitive entities capable of continuous belief formation, updating, and revision enhances transparency and adaptability, particularly within critical application domains such as USAR mission (Hassabis et al. 2017, Lake et al. 2017). By grounding AI methodologies in cognitive theory, researchers can refine real time belief tracking and enhance the interpretability of AI decision making, ultimately leading to more robust human AI collaboration in complex, rapidly evolving scenarios. In the following section, we build upon these perspectives by introducing a framework that systematically captures belief updates over time, leveraging Bayesian inference, temporal alignment, and evidence accumulation to address the limitations identified in previous research.

3 METHODOLOGY

This paper introduces a methodology aimed at analyzing and interpreting internal cognitive belief updating processes of autonomous agents operating within dynamic, partially observable environments through agent-based drone simulations. The approach integrates several quantitative techniques: Bayesian inference for uncertainty quantification, Dynamic Time Warping (DTW) for temporal alignment analysis, and evidence accumulation modeling via Ornstein-Uhlenbeck processes (Pisauro et al. 2017). Additionally, a dual-threshold Bayesian calibration technique informed by empirical human-subject experimental data is applied to differentiate observable behavioral shifts from latent cognitive recalibrations.

3.1 Environmental Score Formulation

The first step of our methodology involves defining an Environmental Score (ES), a quantitative measure representing the environmental conditions perceived by the agent at any given time. The ES encapsulates multiple environmental features such as wind intensity and obstacle density. It is calculated through a linear combination of normalized environmental factors:

$$ES(t) = \lambda_1 \frac{\Delta\vec{v}_{wind}^{\,t} \cdot \vec{v}_{drone}^{\,t}}{Ref_1} + \lambda_2 \frac{|\Delta\vec{v}_{wind}^{\,t} \times \vec{v}_{drone}^{\,t}|}{Ref_2} - \lambda_3 \frac{|\vec{s}_{bld}^{\,t}|}{Ref_3}. \tag{1}$$

λ_1 measures how significantly crosswind influences drone perception of the environment, λ_2 measures the effect of tailwind, λ_3 measures the drone's perceived importance of distance from buildings. Belief updates are triggered when the environmental score surpasses a predefined threshold (ε_e):

$$B^{t+\Delta t} - B^{t} \neq 0, when \int_{t}^{t+\Delta t} (ES^{t+\Delta t} - ES^{t})\, dt \geq \varepsilon_e. \tag{2}$$

3.2 Bootstrapping and Environment Score Fitting

To obtain the threshold value of epsilon corresponding to the environmental score triggering AI behavior change, we first determine a set of coefficients $\lambda_1, \lambda_2, \lambda_3$. These coefficients correspond to when a change in the drone's behavior is detected. By substituting them into $ES(t)$, the functional trend of $ES(t)$ should align with the drone behavior function. In other words, when a behavior change (BC) occurs, it must correspond to a significant change in the environment.

Before proceeding, both $BC(t)$ and $ES(t)$ are normalized to ensure consistency in scale and remove any bias due to magnitude differences. The normalization is performed by transforming each function $f(t)$ (both $BC(t)$and $ES(t)$) as follows:

$$f_{norm}(t) = \frac{f(t) - f_{min}}{f_{max} - f_{min}}. \tag{3}$$

Next, we detect the significant change in AI decisions. In the context of autonomous drone control for USAR, this corresponds to noticeable changes in the drone's navigation behaviors, specifically the change in heading between consecutive timestamps (denoted $\Delta\theta^t$). A behavior change is recorded if $\Delta\theta^t \geq \varepsilon_b$, where ε_b is the direction angle threshold (derived from the mean value plus or minus twice the standard deviation). Once all significant heading changes are identified, 20% of these points are randomly sampled ($N_s = 0.2\ N$, where N is the total number of significant change points) to ensure a representative subset for coefficient fitting.

As part of the fitting process, we adjust $\lambda_1, \lambda_2, \lambda_3$ so that the normalized $ES(t)$ curve matches the normalized $BC(t)$curve at each sampled point. Formally, this is achieved by minimizing:

$$\min_{\lambda_1\, \lambda_2\, \lambda_3} \left(\sum\nolimits_{t\in s} \left(BC'(t) - ES'(t)\right)^2 + \alpha \cdot (|\lambda_1| + |\lambda_1| + |\lambda_1|) \right). \tag{4}$$

Where S is the set of sampled points, $BC'(t)\ and\ ES'(t)$ are the normalized results of $BC(t)$ and $ES(t)$, α is regularization strength, controlling the impact of L1 regularization, penalize large coefficients to prevent overfitting. After $\lambda_1, \lambda_2, \lambda_3$ for each sampled point, the final values are computed by averaging over all samples:

$$\lambda_1^{final} = \frac{1}{N_s}\sum_{t\in s} \lambda_1(t)\,; \quad \lambda_2^{final} = \frac{1}{N_s}\sum_{t\in s} \lambda_2(t)\,; \quad \lambda_3^{final} = \frac{1}{N_s}\sum_{t\in s} \lambda_3(t)\,. \tag{5}$$

3.3 Evidence Accumulation Estimation

According to evidence accumulation theory, even small, instantaneous environmental changes can be sufficient to update a decision maker's belief (Drugowitsch et al. 2015, Forstmann et al. 2016, Friston et al. 2017). We model this accumulation of evidence (AE) using an Ornstein-Uhlenbeck process (Pisauro et al. 2017) to estimate the time t at which AE leads to a behavior change $BC(t)$. Formally:

$$AE(t) - AE(t-1) = \left[\lambda \cdot AE(t-1) + k \cdot \left(ES(t) - ES(t-1)\right)\right] \cdot dt + N(0,\sigma). \tag{6}$$

Where λ represents the leak strength of the process, k modulates the input from $ES(t)$, and $N(0,\sigma)$ is Gaussian noise. Whenever a behavior change occurs, we reset $AE(t)$ to zero and reapply (6) until $|AE(t) - AE(t-\Delta T)| > \varepsilon_e$. Then the time duration of evidence accumulation is $l^t = \Delta\mathrm{T}$. ε_e is the threshold to determine the amount of AE those triggers BC in our idealized scenario (The action will change once the belief has changed). We use (7) to quantify ε_e. It provides a measure of how different the aligned curves are over time. Where t_0 is the start time point of vertical segment in DTW best path line, the lag is the duration of the vertical segment of the best path in the DTW figure (Figure 2), which means the phase difference of the θ^t curve advance to sc curve. Integer i is the index of each vertical segment in DTW best path.

$$\varepsilon_e = \frac{1}{N}\sum_{i=1}^{N}\left(\int_{t_0^{(i)}}^{t_0^{(i)}+lag_i} ES'(t)\,dt - \left(\min_{t\in\left[t_0^{(i)},t_0^{(i)}+lag_i\right]} ES'(t)\cdot lag_i\right)\right). \tag{7}$$

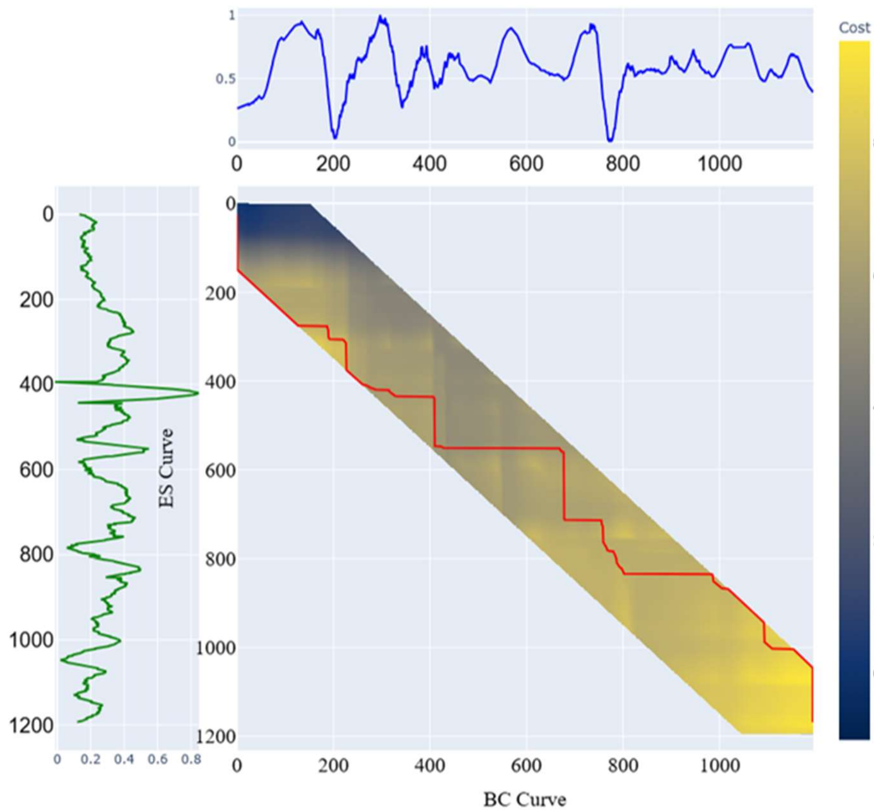

Figure 2: Schematic representation of lag periods in dynamic time warping analysis.

3.4 Bayesian Threshold Adjustment

The threshold ε_e , which determines when the AI updates its belief and makes a decision, was estimated using a Bayesian model. Given that not all belief updates lead to observable changes in the drone's actions, the Bayesian approach assumes that the conditional probability of a belief update leading to a change in action $A(t)$, denoted by $A(t)|B(t)$, is analogous to how humans update beliefs. Through human subject experiments, participants were asked to verbalize their belief updates during decision making scenarios. From these observations, we estimated the AI's probability of belief update $P(B^{(t)})$as follows:

$$P\left(B^{(t)}\right) = \frac{P\left(A^{(t)}B^{(t)}\right)}{P\left(A^{(t)}\middle|B^{(t)}\right)}. \tag{8}$$

Finally, the true value of ε_e was adjusted based on the calculated probability of belief update:

$$\varepsilon_e' = \varepsilon_e P\left(B^{(t)}\right). \tag{9}$$

4 EXPERIMENT

4.1 Simulation Environment

We developed a sophisticated Unity based simulation environment, featuring a detailed DJI Mavic 2 Pro drone navigating a realistically modeled urban landscape of Manhattan based on our previous studies (Sun et al. 2025, Wu et al. 2025). The simulation accurately represents drone physics, including mass, velocity, and aerodynamic interactions, and incorporates dynamic environmental conditions such as variable wind fields and complex urban obstacles. Specifically, wind fields were simulated using a simplified aerodynamic representation method derived from computational fluid dynamics (CFD) models. This approach quantifies wind interactions around urban structures by creating distinct wind zones with varying intensities based on proximity to buildings, enabling efficient simulation of realistic wind scenarios. Reinforcement learning methods, specifically utilizing Physics Informed Neural Networks (PINN), were integrated into the simulation to enable autonomous drone navigation and generate robust and realistic interaction data for analysis.

4.2 Autonomous Drone Control and Perception Modules

The drone's autonomous navigation capabilities were developed using a Multi-Objective Reinforcement Learning approach based on the Proximal Policy Optimization (PPO) algorithm. This approach optimizes the drone's decision making by integrating real time environmental feedback to balance multiple conflicting objectives such as path efficiency, obstacle avoidance, and wind resistance. To enhance human drone interaction, we incorporated a Perception Sharing (PS) module using virtual reality (VR) and haptic feedback systems. This allowed human operators to intuitively perceive the drone's environmental conditions, such as wind strength and direction, facilitating real time informed decision making and interactive control adjustments during drone operations.

4.3 Human Experiment

A diverse sample of 30 participants interacted with the simulation, marking environmental changes, and influencing drone strategies, thus contributing to threshold estimation ($\varepsilon_e{}'$) based on human cognitive patterns.

Figure 3: Experimental setup and experimental equipment.

The demographic information presented in Table 1 illustrates the diversity and composition of our participant sample, ensuring transparency about the subject population involved. Although demographics were not explicitly analyzed as influencing factors in belief updating, clearly reporting these details demonstrates the representativeness and robustness of our experimental procedure.

Table 1: Demographic information about the subjects.

		Number	Percentage
Gender	**Male**	**22**	**73.33%**
	Female	**8**	**26.67%**
Age Group	**18-24**	**8**	**26.67%**
	25-30	**21**	**70.00%**
	31 and older	**1**	**3.33%**
Major	**Civil, Transportation, Construction**	**18**	**60.00%**
	Computer Science, Computer Engineering	**7**	**23.33%**
	Psychology	**5**	**16.67%**

5 RESULTS

This section investigates the AI's internal belief update mechanism by examining how Environmental Scores (ES) relate to observable Drone Behavior Changes (BC). Experimental data were collected from a drone agent pre-trained via the Multi-Objective Reinforcement Learning (MORL) framework within a Unity based simulation environment. The drone was tasked with navigating to a target location in a highly detailed, 1:1 scale simulated urban environment of Manhattan, chosen specifically for its complexity and realistic representation, which requires effective obstacle avoidance and adaptability to severe conditions, especially strong wind. Each drone trajectory was repeated 30 times across four distinct testing scenarios to ensure statistical stability and consistency in AI behavioral responses. The subsequent analysis aims to elucidate the temporal dynamics of the AI's belief updating, clarify the relationship between environmental perception (ES) and action decisions (BC), and provide detailed statistical insights into these cognitive processes.

5.1 Data Preprocessing and Preliminary Results

Environmental and behavioral data were recorded at a frequency of 50 Hz in Unity and underwent preprocessing to enhance analytical clarity. Initially, drone behavioral change (BC) and environmental interaction variables in (1) were smoothed using the Savitzky-Golay filter to mitigate high frequency noise while preserving essential data trends. Subsequently, the data were Min-Max normalized within a [0,1] range to standardize magnitude variations.

We optimized a linear combination model of environmental scores (ES) using L1 regularization, determining optimal weight coefficients $(\lambda_1, \lambda_2, \lambda_3)$ to best fit ES to BC. Figure 4 presents comparison between environment score and drone behavior Change. L1 regularization prevented over fitting, ensuring model stability across diverse scenarios. Figure 5 presents the optimized coefficients for ES across all trajectories, indicating stable coefficients within a certain range (mean values: λ_1=0.51, λ_2=0.33, λ_3=0.23). λ_1 exhibited the greatest influence (47.9%), followed by λ_2 (30.6%), and λ_3 (21.5%). The dominance of λ_1 suggests that sidewind speed significantly influences the ES, overshadowing obstacle complexity and other environmental factors.

The temporal relationship between ES and BC (after optimization and normalization) revealed that peaks in ES consistently precede significant behavioral changes. This pattern underscores the AI's capability to recognize environmental shifts before making behavioral adjustments, highlighting a reaction lag between environmental perception and behavior change. To clarify the dynamics of this lag, DTW analysis was conducted. DTW adeptly handles temporal misalignment analysis, identifying the optimal alignment path between ES and BC. The method also emphasizes “lag periods”, time intervals in which the AI accumulates evidence but does not immediately change its behavior, suggesting that there is a phase of evidence accumulation before action is taken.

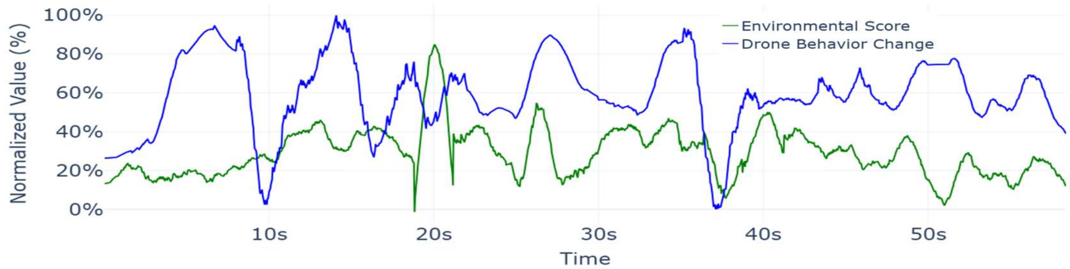

Figure 4: Example of comparison between environment score and drone behavior change.

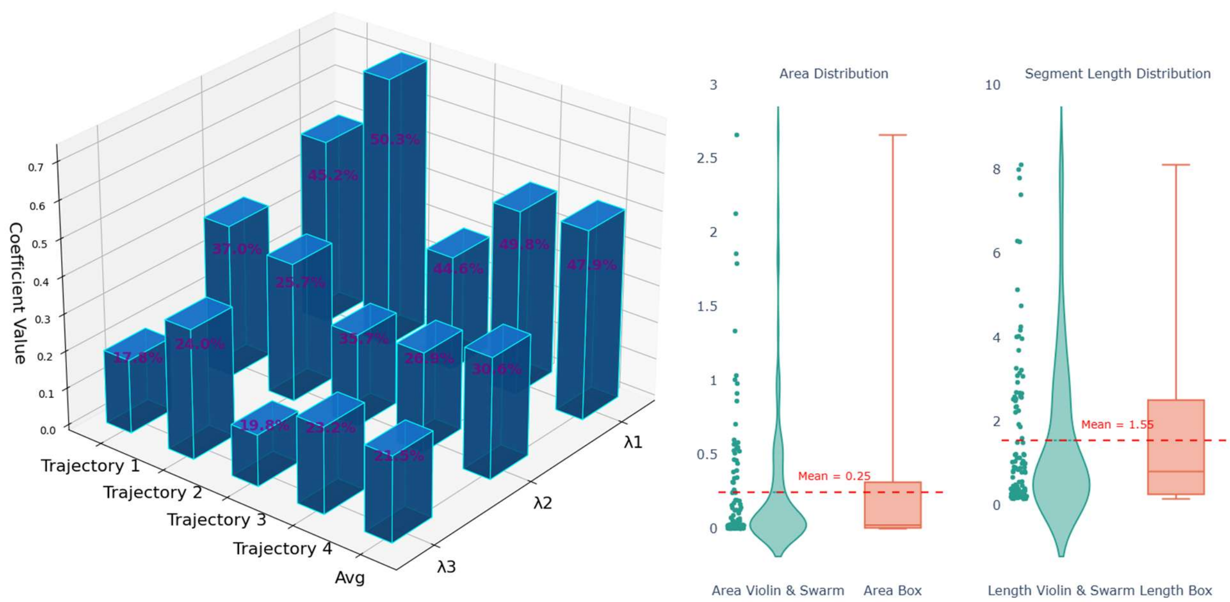

Figure 5: Distribution of optimal weight coefficients (λ1, λ2, λ3) and lagged periods indicated by DTW analysis on given weight coefficients.

5.2 Statistical Analysis and Threshold Determination

Statistical properties of vertical segments (as lagged periods indicated by DTW analysis) are presented in Figure 5. Length of segments had mean values of 1.55s, mean value of net area of the corresponding segments length in ES is 0.2672, respectively. Most segments clustered near these means, suggesting typical lag durations around 1.55 seconds. However, some segments extended beyond 7 seconds, reflecting variability in the AI’s responsiveness to diverse environmental challenges.

5.3 Evidence Accumulation

To investigate how environmental change affects belief update process, we analyzed segments in which drones begin to accumulate evidence until a behavioral change occurs. Figure 6 illustrates this process, labeling the different starting points of the evidence accumulation period (T0), the belief update point (T') and the end point (T), with the dashed shaded area quantifying the total amount of accumulated evidence. The blue dashed shaded area represents the total amount of evidence required to update the beliefs, and the green dashed shaded area indicates that the beliefs are updated and continue to accumulate evidence until the agent makes an action, which is directly represented as a dramatic change in the drone behavior curve BC. The analysis shows that changes in accumulation duration correlate with the intensity of environmental change. Shorter durations were associated with more significant environmental change, indicating faster evidence accumulation and behavioral adjustment.

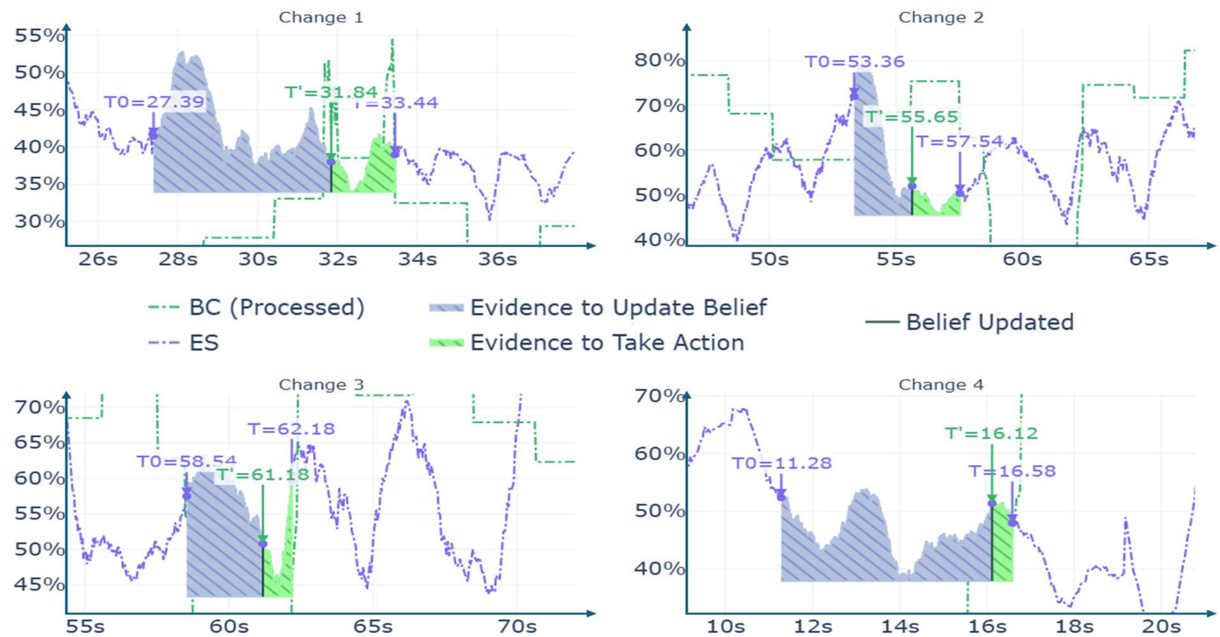

Figure 6: Evidence accumulation and belief update.

5.4 Regression Analysis and Corrections of Belief Updates

By categorizing evidence accumulation fragments and comparing their duration (ΔT) to the strength of environmental changes (ES Diff), an inverse relationship was identified (Figure 7). Specifically, stronger environmental disturbances correspond to faster evidence accumulation, resulting in more rapid belief updating. Initially, a linear regression was applied to the original dataset (red line). Subsequently, outliers were identified and excluded through threshold correction using the ε_e parameter, calibrated from human experimental data. The corrected dataset (blue dots) provided an enhanced linear regression fit (blue line), demonstrating improved accuracy and robustness. This refined model highlights the UAV's capability to rapidly update its beliefs when faced with significant environmental changes, thereby reflecting effective adaptability.

The Bayesian approach (9) also underscores that not all belief updates immediately lead to observable behavioral changes, mirroring human cognitive processes. The calibrated ε_e' threshold enhances the AI agent's ability to discriminate effectively between significant and insignificant environmental changes, thereby balancing sensitivity with operational stability.

Moreover, as illustrated in Figure 6, evidence accumulation continues even after the belief has been updated (green shaded area), preparing the drone for subsequent behavioral adjustments. This two stage accumulation process, including intervals both before and after the belief update, emphasizes the necessity for dynamically adjusting the $\varepsilon_e{}'$ threshold. Such dynamic calibration ensures timely and stable drone responses in rapidly changing scenarios. For instance, as demonstrated in Figure 7, when an environmental change of approximately 30% is detected within a given timeframe, the refined regression model provides guidance for accurately timing belief updates and associated behavioral responses, which is an essential aspect of successful mission accomplishment in complex operational environments, such as search and rescue scenarios.

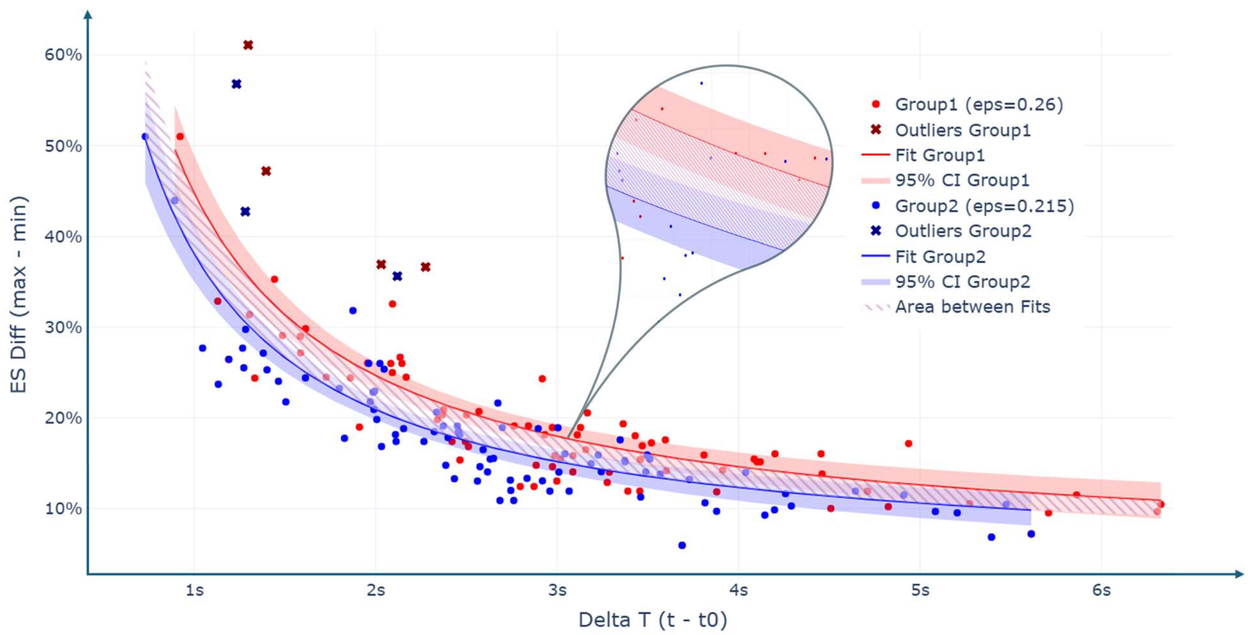

Figure 6: Inverse relationship between evidence accumulation duration and environmental change intensity with ε_e' threshold correction.

6 CONCLUSION

This study investigates the belief update processes of autonomous AI systems, particularly drones navigating dynamic environments, by integrating Bayesian inference, evidence accumulation models, and Dynamic Time Warping (DTW) analysis. A critical finding is the consistent significance of the environmental score (ES) coefficient λ_1, linked to wind conditions, underscoring the model's robustness in prioritizing vital environmental factors. Such stability is indispensable for high risk scenarios like search and rescue, where accurate and timely decisions significantly enhance operational effectiveness. Furthermore, analysis of the evidence accumulation process revealed an inverse relationship between the intensity of environmental disturbances and the duration of belief updates. Specifically, substantial environmental changes triggered rapid belief updates, while minor disturbances resulted in slower, more measured evaluations, effectively balancing responsiveness and stability. Despite these strengths, outliers in the data exposed the AI's limitations in managing abrupt or unfamiliar conditions, suggesting the need for improved generalization capabilities.

Addressing limitations inherent in single threshold models, this study introduced a dual threshold framework by incorporating a secondary threshold (ε_e') calibrated with human experimental data. This innovative approach differentiates internal belief updates from externally observable behaviors, thus aligning AI cognitive patterns more closely with human reasoning processes. By clearly distinguishing significant environmental changes, the calibrated dual threshold framework enhances the AI system's ability to respond efficiently without unnecessary or premature actions, which is crucial for maintaining operational stability and effectiveness.

Despite its promising outcomes, the study's reliance on simulation environments and human derived calibration data introduces limitations regarding real world generalizability. Therefore, future research should prioritize empirical validation in diverse, real world settings, further refine adaptive threshold mechanisms, and expand the framework to multi-agent systems for improved scalability and reliability. Overall, this research significantly advances our understanding of AI belief update mechanisms by explicitly capturing the distinction between internal cognitive processes and observable actions. The proposed framework establishes a strong foundation for developing intelligent, transparent, and trustworthy AI systems, enhancing their capability for meaningful and effective collaboration in complex, dynamic, and uncertain environments.

ACKNOWLEDGEMENTS

This material is supported by the Air Force Office of Scientific Research (AFOSR) under grant FA9550-22-1-0492. Any opinions, findings, conclusions, or recommendations expressed in this article are those of the authors and do not reflect the views of the AFOSR.

REFERENCES

Anderson, J.R. 1991. "Is Human Cognition Adaptive?" *Behavioral and Brain Sciences* 14(3):471–485.

Charniak, E. 1991. "Bayesian Networks without Tears." *AI Magazine* 12(4):50–50.

Kaelbling, L.P., M.L. Littman and A.W. Moore. 1996. "Reinforcement Learning: A Survey." *Journal of artificial intelligence research* 4:237–285.

Thrun, S. 2002. "Probabilistic Robotics." *Communications of the ACM* 45(3):52–57.

Simon, D. 2006. *Optimal State Estimation: Kalman, H Infinity, and Nonlinear Approaches.* Hoboken, New Jersey: John Wiley & Sons.

Korb, K.B. and A.E. Nicholson. 2010. *Bayesian Artificial Intelligence.* Boca Raton, Florida: CRC press.

Murphy, K.P. 2012. *Machine Learning: A Probabilistic Perspective.* Cambridge, Massachusetts: MIT press.

Drugowitsch, J., G.C. DeAngelis, D.E. Angelaki and A. Pouget. 2015. "Tuning the Speed-Accuracy Trade-Off to Maximize Reward Rate in Multisensory Decision-Making." *elife* 4:e06678.

Bissiri, P.G., C.C. Holmes and S.G. Walker. 2016. "A General Framework for Updating Belief Distributions." *Journal of the Royal Statistical Society Series B: Statistical Methodology* 78(5):1103–1130.

Byyny, R.L. 2016. "Information and Cognitive Overload." *Pharos* 79(4):2–13.

Forstmann, B.U., R. Ratcliff and E.-J. Wagenmakers. 2016. "Sequential Sampling Models in Cognitive Neuroscience: Advantages, Applications, and Extensions." *Annual review of psychology* 67(1):641–666.

Friston, K.J., T. Parr and B. de Vries. 2017. "The Graphical Brain: Belief Propagation and Active Inference." *Network neuroscience* 1(4):381–414.

Hassabis, D., D. Kumaran, C. Summerfield and M. Botvinick. 2017. "Neuroscience-Inspired Artificial Intelligence." *Neuron* 95(2):245–258.

Lake, B.M., T.D. Ullman, J.B. Tenenbaum and S.J. Gershman. 2017. "Building Machines That Learn and Think Like People." *Behavioral and Brain Sciences* 40:e253.

Pisauro, M.A., E. Fouragnan, C. Retzler and M.G. Philiastides. 2017. "Neural Correlates of Evidence Accumulation during Value-Based Decisions Revealed Via Simultaneous Eeg-Fmri." *Nature Communications* 8(1):15808.

Adadi, A. and M. Berrada. 2018. "Peeking inside the Black-Box: A Survey on Explainable Artificial Intelligence (Xai)." *IEEE Access* 6:52138–52160.

Gunning, D. and D. Aha. 2019. "Darpa's Explainable Artificial Intelligence (Xai) Program." *AI Magazine* 40(2):44–58.

Rudin, C. 2019. "Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead." *Nature Machine Intelligence* 1(5):206–215.

Kawato, M. and A. Cortese. 2021. "From Internal Models toward Metacognitive Ai." *Biological Cybernetics* 115(5):415–430.

Khalvati, K., R. Kiani and R.P. Rao. 2021. "Bayesian Inference with Incomplete Knowledge Explains Perceptual Confidence and Its Deviations from Accuracy." *Nature Communications* 12(1):5704.

Chen, W., X. Meng, J. Liu, H. Guo and B. Mao. 2022. "Countering Large-Scale Drone Swarm Attack by Efficient Splitting." *IEEE Transactions on Vehicular Technology* 71(9):9967–9979.

Kojima, T., S.S. Gu, M. Reid, Y. Matsuo and Y. Iwasawa. 2022. "Large Language Models Are Zero-Shot Reasoners." *Advances in Neural Information Processing Systems* 35:22199–22213.

Binz, M. and E. Schulz. 2023. "Using Cognitive Psychology to Understand Gpt-3." *Proceedings of the National Academy of Sciences* 120(6):e2218523120.

Frömer, R. and M.R. Nassar. 2023. "Belief Updates, Learning and Adaptive Decision Making." *Retrieved from psyarxiv. com/qndba.*

Hagendorff, T., I. Dasgupta, M. Binz, S.C. Chan, A. Lampinen, J.X. Wang et al. 2023. "Machine Psychology." *arXiv preprint arXiv:2303.13988*

Rago, A. and M.V. Martinez. 2024. "Advancing Interactive Explainable Ai Via Belief Change Theory." *arXiv preprint arXiv:2408.06875*

Sun, B., J. Wu, H. You and J. Du. 2025. "Estimating Belief Updates in Ai-Driven Drone Controls for Urban Search and Rescue Operations." *Journal of Construction Engineering and Management* 151(9):04025128.

Wu, J., B. Sun, H. You and J. Du. 2025. "Enhancing Human-Ai Perceptual Alignment through Visual-Haptic Feedback System for Autonomous Drones." *International Journal of Industrial Ergonomics* 109:103780.

AUTHOR BIOGRAPHIES

BOWEN SUN is a Ph.D. student at the Informatics, Cobots, and Intelligent Construction Lab in the Department of Civil and Coastal Engineering at the University of Florida in Gainesville, Florida, USA. His research interests include drone simulation, metacognitive calibration of autonomous robots, human-AI alignment, and reinforcement learning for autonomous robots. His email address is bowensun@ufl.edu.

JIAHAO WU is a Ph.D. student at the Informatics, Cobots, and Intelligent Construction Lab in the Department of Civil and Coastal Engineering at the University of Florida in Gainesville, Florida, USA. His research interests include drone simulation, human-robot collaboration, and reinforcement learning for autonomous robots. His email address is jiahaowu@ufl.edu.

JING DU is a Professor in the Department of Civil Engineering, the Department of Mechanical and Aerospace Engineering, and the Department of Industrial and System Engineering at the Herbert Wertheim College of Engineering, University of Florida. His primary area of research is human-robot collaboration for complex industrial operations. Dr. Du is the elected Secretary of the American Society of Civil Engineers (ASCE) Visualization, Information Modeling and Simulation (VIMS) committee and serves on the editorial board of three journals. His email address is eric.du@essie.ufl.edu, and his website is https://faculty.eng.ufl.edu/ericdu/.