drones

MDPI

*Article*

# The Pull–Push Engine: Bidirectional Emotion Regulation for Emotionally Intelligent UAV Traffic Monitoring

**Mohamed Zaidan [1], Nafaâ Jabeur [2,*], Muhammad Aamir Basheer [3] and Ansar-Ul-Haque Yasar [1]**

[1] Transportation Research Institute (IMOB), Hasselt University, 3500 Hasselt, Belgium; mohamed.zidan@uhasselt.be (M.Z.); ansar.yasar@uhasselt.be (A.-U.-H.Y.)
[2] Department of Computer Science, German University of Technology in Oman (GUtech), Muscat 1816, Oman
[3] Department of Architecture and City Design, King Fahd University of Petroleum and Minerals (KFUPM), Dhahran 31261, Saudi Arabia; aamir.basheer@kfupm.edu.sa
* Correspondence: nafaa.jabeur@gutech.edu.om; Tel.: +968-9627-5005

## Highlights

### What are the main findings?

- The Pull–Push Engine (PPE) maintains zero PAD saturation across 36 runs and three urban traffic scenarios—and across all five personality profiles (60 runs)—versus 57–88% saturation when the personality-anchoring term is removed, confirming its central role in stable, generalizable emotion-regulated control.
- Behavioral diversity in Emotionally Intelligent UAV (EI-UAV) control scales with operational complexity through emotional dynamics alone—without any explicit complexity-detection mechanism—with strategy entropy increasing monotonically from 1.885 to 2.033 bits across scenarios.

### What are the implications of the main findings?

- In SUMO-simulated urban scenarios of increasing complexity, the PPE delivers an interpretable, adaptive, and stable regulatory pattern (subject to hardware-in-the-loop and field validation). This pattern is intended to generalize to other domains that require transparent decision-making.
- Every EI-UAV decision is causally traceable from environmental stimulus through emotional processing to action in simulation. This provides the structural transparency typically associated with the interpretability requirements of future safety-certification and human-oversight processes.

## Abstract

Autonomous UAVs for urban traffic monitoring must respond quickly to changing operational conditions while maintaining stable, transparent decision-making. Rule-based controllers respond only at predefined thresholds, while learning-based methods adapt well but lack the certification transparency required for safety-critical deployment. This paper proposes a bio-inspired emotion-regulated decision-control mechanism and introduces the Pull–Push Engine (PPE), a regulatory architecture that balances environmental stimuli against personality-anchored baselines through weighted temporal integration. The PPE is embedded in a three-layer framework combining Big Five personality traits, the Pleasure–Arousal–Dominance (PAD) model, and Ortony–Clore–Collins (OCC) event appraisal. Validation in a SUMO-based simulation across three scenarios of increasing complexity showed that PPE regulation maintained bounded PAD trajectories and zero saturation despite concurrent stressors, whereas removing the pull term caused 57–88% saturation. Behavioral diversity scaled naturally with operational demands: Surprised



Academic Editor: Pablo Rodríguez-Gonzálvez

Received: 26 March 2026
Revised: 6 May 2026
Accepted: 6 May 2026
Published: 17 May 2026

mood dominated across all scenarios (47.8–67.5%), with Anxious and Focused increasing systematically with complexity. Strategy entropy rose monotonically (1.885–2.033 bits). A sensitivity sweep confirmed robust regulation across a stable operating region, with degradation only at the boundary ($p < 0.001$ for all key comparisons). Every simulated decision remains causally traceable from stimulus through emotional processing to action. This ensures interpretability, which is essential for future safety-critical UAV deployment, although hardware implementation and field validation are still required.

**Keywords:** affective computing; UAV autonomy; traffic monitoring; emotional regulation; PAD model; OCC theory; Big Five personality; decision control; Pull–Push Engine; interpretable AI

## 1. Introduction

Autonomous Unmanned Aerial Vehicles (UAVs) are increasingly investigated and deployed for urban traffic monitoring [1–3]. They can dynamically reposition, survey extensive areas, and reach incidents that ground-based infrastructure cannot easily access. In practice, a patrolling UAV must continuously balance competing factors, including which area to check, how long to dwell at each location, and when to prioritize battery margins over coverage objectives. These trade-offs become increasingly challenging during prolonged missions, making effective decision-making approaches essential. This research focuses on missions lasting over one hour of continuous monitoring in congested urban corridors, a duration sufficient for these pressures to accumulate. Most UAV decision-making approaches fall into two broad categories, namely rule-based systems and learning-based systems. Rule-based systems use predefined thresholds, such as switching to intensive monitoring when congestion exceeds a limit or returning to base when battery levels drop [4–6]. Under predictable conditions, these approaches are generally reliable. However, they react only after predefined thresholds are crossed and do not account for recent state history. More recent learning-based approaches, including reinforcement learning and neural networks, are more flexible and can outperform rule-based systems in simulation. However, they typically operate as black boxes [7,8]. This is a known barrier to regulatory acceptance of autonomous systems [9].

Achieving genuinely intelligent autonomy, where decisions are sufficiently transparent to support human oversight and safety auditing, remains a significant challenge [10], and most operational systems continue to rely on rigid rules or opaque learned policies. This work takes a different approach, following the bio-inspired computing paradigm, by transplanting the functional role that emotions play in biological systems, such as directing attention, adjusting risk sensitivity, and integrating recent experiences into ongoing decisions [11–13], and repurposing these mechanisms for operational control rather than social communication. The intent is not to make UAVs "feel" anything, but to test whether emotion-inspired regulation can help balance rapid responsiveness with behavioral stability, a trade-off that current UAV control approaches often struggle to manage effectively.

In this work, an emotional state is defined as a formal engineering construct: a bounded continuous variable in the three-dimensional PAD space [14], with explicit dynamics governing its evolution under environmental and event-driven inputs. The framework does not claim that UAVs experience emotions, nor does it require physiological or perceptual measurement of internal affect. The construct functions as an internal control signal that integrates recent observations into a low-dimensional representation, biasing subsequent decisions in interpretable and bounded ways. The term "emotion" is

retained for consistency with the established affective-computing literature, where it refers to bio-inspired computational architectures rather than subjective experience. Accordingly, the mood labels used throughout this paper (e.g., Surprised, Anxious, Focused) act as human-readable descriptors of regions within this formally defined control-state space.

Three frameworks from the field of affective computing are drawn upon.

The Pleasure–Arousal–Dominance (PAD) model [14] represents emotional state as coordinates in three-dimensional space, well suited for tracking gradual mood shifts. Discrete events (accidents, sensor faults, communication dropouts) require a different treatment; for these, we use the Ortony–Clore–Collins (OCC) model [15] to evaluate events against mission goals. To prevent excessive sensitivity to transient disturbances, the Big Five personality model (OCEAN) [16] supplies trait-based baselines that moderate responsiveness over time. Because environmental inputs can still drive PAD toward extreme values, we introduce the Pull–Push Engine (PPE) as an explicit regulation layer. PAD monitors ongoing emotional state in accordance with PPE regulation, OCC manages event-driven perturbations, and OCEAN stabilizes the system while fostering adaptability.

To the best of our knowledge, no prior affective UAV control framework has formalized a bidirectional regulation mechanism combining personality anchoring with contextual push dynamics. This research presents five key contributions: (1) emotion is reformulated as an autonomous decision-control mechanism from which strategy selection directly arises; (2) the Pull–Push Engine (PPE) balances environmental responsiveness with personality-anchored stability using explicit bounds, damping, and inherent hysteresis; (3) the approach is applied to a safety-critical domain, i.e., autonomous UAV traffic monitoring; (4) empirical validation encompasses three SUMO-based traffic scenarios of increasing complexity; and (5) causal traceability evaluation confirms that each decision is empirically traceable from environmental stimulus through emotional processing to behavioral output, given the system's configured parameters—factor weights (Appendix D), OCC mappings (Appendix C), and mood-strategy table (Appendix B). Empirical validation is structured around five primary components: A per-objective validation summary linking the nine design goals to quantitative evidence (Section 5.6); a parameter sensitivity analysis identifying a stable operating region (Section 5.7); statistical comparisons across scenarios with effect sizes (Section 5.8); cross-personality generalization across five profiles (Section 5.9); and an architectural ablation demonstrating the role of bidirectional regulation (Section 6.2.2). The paper focuses on internal emotional regulation and its effect on decision stability. Perception accuracy and low-level flight control are abstracted in this simulation-based study and assumed to be handled by standard UAV subsystems in any future hardware integration.

## 2. Related Work

### *2.1. UAV Traffic Monitoring Systems*

Rule-based approaches dominate road traffic monitoring owing to their predictability and certification readiness [4–6]. These systems have proved effective for congestion detection and event-triggered surveillance [4,5] but struggle when multiple concurrent factors must be balanced simultaneously. Learning-based approaches [7] can discover non-obvious strategies for coverage and resource use; for example, deep reinforcement learning has been widely applied to autonomous decision problems [7]. However, the black-box nature of neural policies creates obstacles for safety certification [9], since regulators require traceable rationales for autonomous decisions in public airspace. A systematic analysis of over 315 scholarly works confirmed that UAV-based systems enhanced with AI algorithms have advanced traffic flow analysis, safety conflict analysis, and emergency response [17,18]. The same body of work, however, identified two persistent gaps: adaptive decision-making in uncertain conditions and growing pressure for explainability that

supports transparency, trust, and compliance with regulatory standards. Atakishiyev et al. [9] examined explainability approaches for autonomous driving systems, arguing that systems must not only make safe decisions but also supply understandable rationales. This interpretability gap in existing learning-based UAV solutions motivates our interest in emotion-driven control.

### 2.2. Affective Computing and Emotional Models

Mehrabian's PAD model [14] treats emotion as a behavior-regulating mechanism: Pleasure reflects overall positivity or negativity, Arousal captures activation level, and Dominance encodes perceived control over the situation. Affective terminology is clarified first, following conventions from computational affective modeling [11,19]. Emotion refers to short-lived responses elicited by discrete events (handled by OCC appraisal); mood refers to longer-duration states that influence behavioral tendencies (classified regions in PAD space); affect is the umbrella term encompassing both; personality denotes stable trait dispositions (Big Five/OCEAN) that anchor baseline affective profiles. These definitions serve engineering precision rather than strict psychological fidelity. The PAD representation, continuous yet bounded, fits naturally into autonomous control loops and has seen renewed use in autonomous robotics, intelligent vehicles, and human–machine interaction [20,21]. Picard [11] introduced affective computing as a field concerned with representing emotions as functional components in intelligent systems. The OCC model [15] complements PAD by classifying emotion-triggering events via cognitive appraisal, and Broekens et al. [22] formalized this in appraisal models that show how events can be checked against goals to generate affective responses. Pairing the two allows ongoing state tracking alongside disciplined handling of discrete events.

Related studies have examined the role of emotions in shaping agent learning and resource allocation: Marinier et al. [23] reported that emotion-driven reinforcement learning improves goal prioritization in cognitive architectures. Lowe et al. [24] demonstrated that affect-based signals can modulate attentional and behavioral resources in autonomous agents through embodied emotion-appraisal coupling, and Pessoa [25] provided neuroscientific evidence that emotion and cognition are tightly coupled rather than separable, a principle underlying the integrated appraisal approach adopted in this work. Moerland et al. [7] survey emotion in reinforcement learning agents and robots, cataloguing how affective mechanisms have been used to modulate exploration, action selection, and learning rates in varied architectures.

The Big Five model (OCEAN) [16] describes personality along five stable dimensions, namely Openness, Conscientiousness, Extraversion, Agreeableness, and Neuroticism. Mehrabian [26] established linear mappings between these traits and PAD dimensions, later validated by Zhao et al. [27]. Gebhard [19] and Marsella and Gratch [28] combined emotions, mood, and personality in virtual agents, with recent reviews [29,30] and empirical work [31] confirming that socio-emotional attributes improve human–AI collaboration. Yet personality has mostly played a background role, lacking explicit regulatory mechanisms [19,21,28], and most implementations remain in social contexts rather than safety-critical applications, a gap the present work addresses.

On the application side, emotion has served as a scheduling variable [32] and as part of autonomous perception and action selection [20], while PAD-based modulation has demonstrated value in factory settings [21]. Yet principled affective mechanisms for open-ended environments remain scarce [33], and although AI-based affective models have proliferated since 2020 [34], few scale beyond controlled scenarios. Table 1 compares prior approaches relative to this work. Surveys of social emotions in robotics [35] further highlight gaps among computational affect models and real-world operational needs. The

frameworks adopted here—the PAD temperament model [14], the OCC cognitive appraisal theory [15], the Big Five personality model [16], and the neuroscientific evidence coupling emotion and cognition [25]—collectively constitute a bio-inspired affective substrate, grounding the proposed architecture in established biological mechanisms rather than arbitrary engineering choices.

**Table 1.** Comparison of emotion-based control approaches in autonomous systems.

| Approach | Emotion Role | Regulation | Domain | Interpretability |
|---|---|---|---|---|
| ALMA [19] | Mood generation | Sequential pull–push | Virtual agents | Medium |
| EMA [28] | Appraisal-based | Implicit | Social simulation | Medium |
| Gadanho [36] | RL modulation | None explicit | Learning agents | Low |
| Cañamero [20] | Action selection | Implicit | Autonomous robots | Medium |
| Nunnari et al. [21] | PAD behavior | None explicit | Industrial robots | High |
| This Work (PPE) | Operational decision control | Explicit bidirectional (PPE) | UAV traffic monitoring | Full causal trace |

### 2.3. Research Gap and Contribution

Across these studies, a recurring limitation is that emotion is either confined to social interaction or treated as a secondary factor influencing decisions determined by other mechanisms. What is currently lacking is a system that allows emotion to directly control behavior while maintaining operational stability. PAD-based studies often leave unclear how noisy affective signals can drive reliable behavioral control [21], a concern that motivated the Pull–Push Engine (PPE) presented in Section 4. Although Gebhard's ALMA framework [19] employs "pull" and "push" terminology for mood evolution in virtual characters, the underlying purpose and formalization differ substantially. In ALMA, push–pull dynamics shape affective expression and believability. In contrast, the PPE formalizes push–pull as an explicit regulatory mechanism for operational decision control: environmental inputs induce deviation from a personality-anchored baseline, while the baseline contributes a stabilizing feedback term that constrains drift. The responsiveness parameter ($\alpha$) quantitatively governs this balance, enabling controlled trade-offs between adaptability and stability. Unlike past work, where emotion biases decisions made by separate planners, the PPE couples regulation directly to strategy selection. By regulating emotional evolution explicitly, rather than merely biasing downstream decisions, the system accommodates changing conditions while preserving operational stability.

## 3. Proposed Solution

### 3.1. Design Rationale

Effective UAV operations in transportation environments require consistent behavioral tendencies that govern how the system prioritizes safety, responsiveness, and coordination across prolonged missions [4–6]. The term behavioral attributes refers to observable, repeatable decision-making patterns that emerge from the interaction among emotional states, personality baselines, and environmental context. Five attributes are central to the proposed design: (1) responsiveness to emerging conditions without overreaction; (2) behavioral stability under noisy or conflicting inputs; (3) context awareness integrating simultaneous operational factors; (4) persistence balanced with flexibility as circumstances change; and (5) communicability enabling operator trust and oversight.

### 3.2. Stability Challenges in Emotion-Driven Control

Without explicit regulation, emotion-driven control faces four interrelated stability concerns [21]: oscillation caused by rapid mood swings, overreaction to brief events,

threshold chattering near mood boundaries, and context amnesia from ignoring recent state history. These concerns motivated the design of the Pull–Push Engine (Section 4), which addresses each through its recursive formulation, as detailed in Section 4.1.

### *3.3. Operational Constraints*

Decision arbitration follows the priority hierarchy shown in Figure 1: Safety > Operator directives > OCC appraisal > PAD modulation > Personality baseline. Safety-critical constraints always take precedence; for instance, a critical battery or thermal limit triggers immediate return to base regardless of emotional state. Operator directives come next, followed by event-driven affective responses (OCC), continuous affective modulation (PAD), and then personality baselines. This ordering keeps safety first while emotional dynamics continue to shape routine behavior. For traffic monitoring, we selected the Calm personality profile from five predefined profiles (Calm, Confident, Proactive, Investigative, and Cautious) because its low-arousal baseline aligns with routine surveillance demands and sustained monitoring tasks. The remaining profiles are validated and comparatively evaluated in Section 5.9. Trait values and the resulting PAD baseline are reported in Section 5.1.

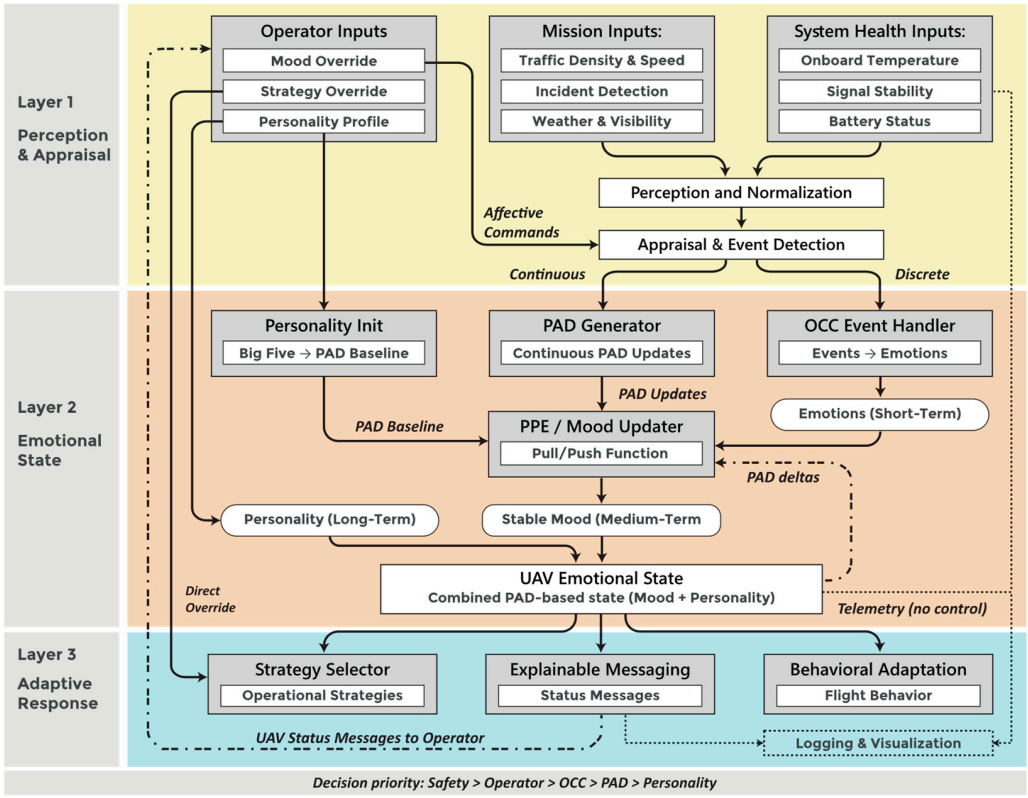


**Figure 1.** Regulated emotional decision-control architecture of the Emotionally Intelligent UAV (EI-UAV). Three layers connect raw inputs to flight behavior through perception, emotional regulation via the PPE, and adaptive strategy selection. Decision priority: Safety > Operator > OCC > PAD > Personality. Solid arrows indicate primary data flow and direct-override commands; dashed arrows indicate feedback paths (PAD deltas to the PPE, UAV status messages returned to the operator); dotted arrows indicate telemetry routed to logging and visualization (no control influence).

### *3.4. System Architecture*

Having established the design requirements, the three-layer architecture that addresses them is presented in Figure 1. Evidence from cognitive architectures [20] and industrial applications [21] indicates that emotional dynamics enhance control performance only when they are systematically regulated, which constitutes the central rationale for the

proposed Pull–Push Engine. The Perception and Appraisal layer performs input processing and contextual evaluation. The Emotional State layer governs affective regulation through the Pull–Push Engine. The Adaptive Response layer is responsible for strategy selection and behavioral execution. The following subsections describe each layer in detail.

#### 3.4.1. Layer 1: Perception and Appraisal

The first layer processes three input streams: operator inputs (mood override, strategy override, and personality profile selection), mission inputs (traffic density and speed, incident detection, and weather and visibility conditions), and system health inputs (onboard temperature, signal stability, and battery status). Raw signals are normalized to consistent scales and subsequently forwarded to the appraisal module, which differentiates continuous contextual factors from discrete events. Continuous signals are transmitted to the PAD Generator in Layer 2, whereas discrete events activate the OCC Event Handler.

#### 3.4.2. Layer 2: Emotional State

Layer 2 represents the UAV's internal emotional state as a point in PAD space. The layer initializes a personality baseline, denoted as $PAD_0$, derived from Big Five trait parameters, and subsequently updates the emotional state through two complementary mechanisms. Gradual contextual influences accumulate over time, while discrete events elicit immediate affective adjustments. Personality initialization is based on the Big Five OCEAN framework, which characterizes personality along five dimensions, i.e., Openness (O), Conscientiousness (C), Extraversion (E), Agreeableness (A), and Neuroticism (N). The Big Five traits have well-established relationships with emotional baselines and behavioral tendencies [16], providing a principled basis for initializing PAD state rather than selecting baseline values arbitrarily. Using the selected personality profile defined in Section 3.1, the Big Five trait values (O, C, E, A, N) are mapped to baseline emotional values in PAD space using the linear relations established by Mehrabian [26] and validated by Zhao et al. [27]. Specific trait values are reported in Section 5.1. The baseline PAD components are computed as

$$\begin{aligned} P_0 &= 0.21E + 0.59A + 0.19N \\ A_0 &= 0.15O + 0.30A - 0.57N \\ D_0 &= 0.25O + 0.17C + 0.60E - 0.32A \end{aligned} \tag{1}$$

Substituting the selected personality values into these equations produces the personality-specific PAD baseline, which anchors emotional regulation throughout the mission. The Pull–Push Engine uses this baseline as its pull reference, against which real-time environmental inputs and event-driven appraisals dynamically modulate emotional state.

**Design note.** The simulation platform provides several predefined personality presets (Calm, Confident, Proactive, Investigative, and Cautious), as well as manual OCEAN configuration. In this study, personality is fixed to the Calm profile across all scenarios—selected for its low-arousal baseline appropriate to routine surveillance and to isolate the effects of emotional regulation without confounding behavioral variations from personality differences.

**Continuous Pathway (Contextual Appraisal).** The continuous pathway models the manner in which traffic conditions and internal UAV system states influence the evolution of the emotional state. Factors such as traffic density, battery status, and communication quality continuously modulate PAD values through small and incremental changes. For example, elevated congestion increases Arousal (encoding urgency), declining battery reduces Dominance (reflecting diminished operational confidence), and sustained suc-

cessful coverage increases Pleasure (representing mission satisfaction). These influences accumulate incrementally through a transparent weighted summation:

$$PAD_t^{cont} = \sum_{i=1}^{n} w_i s_i(t) \tag{2}$$

where $s_i(t)$ denotes the normalized score associated with contextual condition i, and $w_i$ represents its corresponding influence weight. All weights are specified a priori in the system configuration and remain fixed throughout the experiments; the complete factor-weight mapping is provided in Appendix D (Table A4). Discrete events are intentionally excluded from this pathway and are handled separately through OCC-based appraisal.

**Discrete Pathway (Event-Driven OCC Appraisal).** The discrete pathway handles salient events that require immediate emotional response. The Ortony–Clore–Collins (OCC) appraisal framework [15] is applied to evaluate detected events (accidents, communication failures, system faults) in terms of their impact on mission goals. Each event type is mapped to a predefined emotion-specific target within PAD space, with mappings configured a priori according to domain requirements.

The complete OCC event–PAD target mapping is provided in Appendix C (Table A3). The implementation distinguishes two activation modes. Hard OCC events (e.g., accidents, critical failures) override the emotional state directly, setting PAD to the event's target as specified in Algorithm 1 (line 2). Soft OCC events (e.g., moderate congestion changes, non-critical alerts) contribute incremental deltas that are integrated through the PPE's continuous pathway. Equation (3) formalizes the soft event contribution as a weighted delta:

$$\Delta \mathrm{PAD}_{\mathrm{t}}^{\mathrm{occ}} = \sum_{i=1}^{m} \lambda_i \left( \mathrm{PAD}_{\mathrm{i}}^{\mathrm{target}} - \mathrm{PAD}_{\mathrm{t}} \right) \tag{3}$$

Here, PAD target represents the emotion-specific PAD reference associated with the detected event (See Appendix C, Table A3). For high-priority events, the highest-priority event takes precedence and directly overrides the current PAD state (Algorithm 1, lines 1–4). For low-priority events, the weighted deltas defined in Equation (3) accumulate alongside continuous contextual factors in the PPE's REGULATING state (Algorithm 1, line 12). Event priority determines which operational mode is activated when multiple events occur simultaneously. The continuous and discrete pathways jointly constitute the environmental push signal that is subsequently regulated by the Pull–Push Engine.

#### 3.4.3. Layer 3: Adaptive Response

Layer 3 converts the continuous PAD state into discrete operational decisions through a two-stage process comprising mood classification and strategy selection. The current PAD coordinates are mapped to the nearest centroid from a set of ten moods (Anxious, Excited, Satisfied, Distressed, Joyful, Focused, Neutral, Relieved, Surprised, Exhausted), though the set is extensible (six of these ten moods are activated under the Calm profile; see Section 6.4 for a discussion of the four inactive centroids). A safety-level Critical state overrides this classification when battery drops below 20%, forcing immediate return to base. Simple octant membership cannot distinguish moods that share an octant (e.g., Excited vs. Joyful), so centroid-based classification provides the needed granularity. These discrete labels function as an interpretability layer imposed on continuous PAD space and do not affect PPE dynamics. Centroid values are presented in Appendix A (Table A1), and their spatial distribution is illustrated in Figure A1.

**Algorithm 1.** Pull–Push Engine (PPE) Regulation.

Input: $PAD_t$, $PAD_0$, $\alpha$, $\beta_0$, Event, $\Delta$ctx, State
Output: $PAD_{t+1}$, State'

1: if Event $\neq$ NULL then ► Discrete pathway
2: $PAD_{t+1} \leftarrow$ Event.PAD_target
3: State' $\leftarrow$ OVERRIDE
4: return (CLAMP($PAD_{t+1}$), State')
5: if State = OVERRIDE and Event = NULL then ► Event ends
6: State' $\leftarrow$ RECOVERING
7: if State = RECOVERING then ► Recovery pathway
8: $PAD_{t+1} \leftarrow PAD_t + \beta_0 \cdot (PAD_0 - PAD_t)$ ► conceptual form; see Notes
9: if $\|PAD_{t+1} - PAD_0\| < \varepsilon$ then State' $\leftarrow$ REGULATING else State' $\leftarrow$ RECOVERING
10: return (CLAMP($PAD_{t+1}$), State')
11: // REGULATING: Push-Pull balance (Equation (4)) ► Continuous pathway
12: $PAD_{t+1} \leftarrow PAD_t + \alpha \cdot \Delta ctx + (1-\alpha) \cdot \beta_0 \cdot (PAD_0 - PAD_t)$
13: State' $\leftarrow$ REGULATING
14: return (CLAMP($PAD_{t+1}$), State')

Notes:

- $\Delta$ctx denotes the aggregated PAD increment from continuous contextual factors.
- CLAMP(x) = max($-1$, min($+1$, x)).
- $\alpha = 0.57$ (responsiveness), $\beta 0 = 0.30$ (personality coupling constant), $\varepsilon = 0.05$ (recovery threshold). $\varepsilon$ was selected so that convergence corresponds to a PAD deviation below 0.05 from the personality baseline, achievable within a typical recharge cycle (~3000 steps); values between 0.01 and 0.10 produced qualitatively similar recovery behaviour in preliminary testing. Note: the implementation uses a graduated three-stage path (OVERRIDE target $\rightarrow$ Focused $\rightarrow$ Satisfied $\rightarrow$ PAD0) with a smaller effective pull rate (~0.03/step) to produce smooth emotional descent; Algorithm 1 line 8 presents the conceptual update form and $\beta 0 = 0.30$ governs the REGULATING pull strength (Equation (4)), not the RECOVERING pull rate. Event in line 1 refers exclusively to discrete OCC events; continuous contextual factors are aggregated into $\Delta$ctx prior to the update step (line 12). Note that line 7 evaluates the original State value (before the assignment on line 6). The RECOVERING pathway, therefore, activates on the subsequent simulation step, not the current one.

**Strategy selection:** Each mood maps to a primary surveillance strategy from eight defined modes (S0–S7), ranging from emergency return-to-base to congestion-prioritized patrol and operator-directed deployment (Appendix B, Table A2; Table A2 defines the complete mood-to-strategy mapping for all ten moods, including centroids not reached under the Calm profile, which remain defined for completeness and apply under other personality configurations). Five modes (S0–S3 and S6) were consistently activated across all experimental runs and personality profiles. Selection considers mood alongside operational constraints (battery, communication status) and mission priorities (active incidents, operator directives), with safety-critical constraints always overriding emotional preferences.

## 4. The Pull–Push Engine: Affective Regulation Mechanism

### *4.1. The Regulation Problem*

Section 3.2 identified four stability challenges: oscillation from minor traffic variations triggering unnecessary strategy switches; overreaction to brief events such as momentary congestion spikes; threshold chattering near mood-class boundaries; and context amnesia when instantaneous updates discard recent state history. The PPE addresses each through its recursive formulation: $\alpha$-weighted averaging (Equation (4)) attenuates oscillatory behavior; bounded PAD space (Section 4.4) constrains excessive reactions; implicit hysteresis in continuous PAD evolution mitigates threshold chattering; and recursive state updates preserve temporal context. Algorithm 1 details the complete regulation procedure as a three-state machine (REGULATING, OVERRIDE, RECOVERING) implemented via dual-pathway update logic.

### *4.2. The Pull–Push Function*

The Pull–Push Engine (PPE) regulates emotional dynamics through two counterbalancing forces as illustrated in the middle layer of Figure 1. The push originates from environmental conditions. Elevated congestion levels or the occurrence of an accident increase Arousal and reduce Pleasure, thereby encoding operational urgency, whereas sustained mission success gradually elevates Pleasure and Dominance. The pull derives from personality. Baseline PAD values anchor emotional state to the UAV's characteristic operating range and exert a restorative influence when external pressures subside. Excessive push leads to behavioral volatility, while excessive pull results in reduced responsiveness. The PPE maintains this balance within bounded limits, ensuring that emotional responses remain adaptive without becoming unstable. Operating as a three-state machine (Figure 2), the engine's state transitions and update rules are formalized in Algorithm 1.

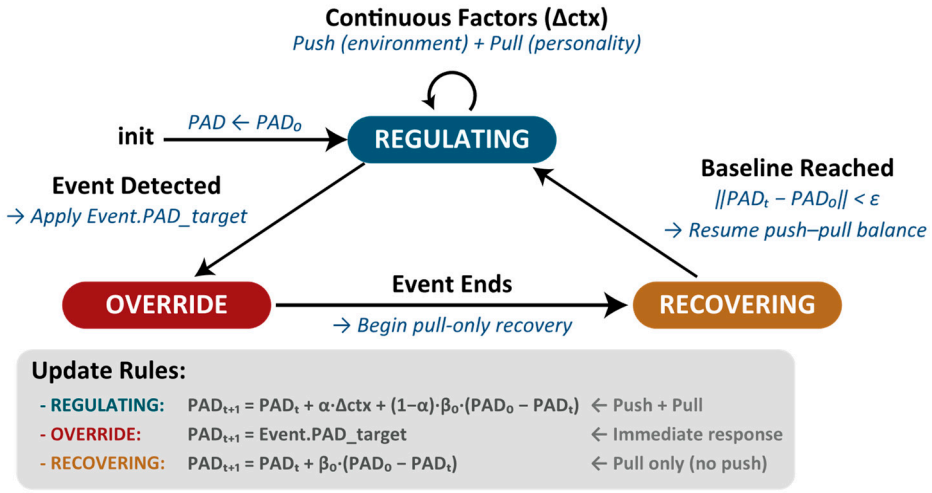


**Figure 2.** Pull–Push Engine (PPE) state machine. Three modes: REGULATING (push–pull balance), OVERRIDE (event-driven), and RECOVERING (return to baseline).

### *4.3. Mathematical Formulation of the PPE*

At each time step, environmental context induces a state change proportional to $\alpha$, while a recovery term draws the emotional state back toward the personality baseline. This preserves responsiveness to situational change while preventing unbounded drift:

$$PAD_{t+1} = PAD_t + \alpha \cdot \Delta_{ctx} + (1 - \alpha) \cdot \beta_0 \cdot (PAD_0 - PAD_t) \quad (4)$$

Here, $PAD_t$ denotes the current regulated emotional state, and $\Delta_{ctx}$ represents the per-step increment from continuous contextual factors as defined in Equation (2). The

parameter $\alpha$ scales the magnitude of the push term, while $\beta_0$ determines the strength of the pull toward the personality baseline. The parameter $\alpha$ takes values in the interval [0, 1] and governs environmental responsiveness, whereas $(1 - \alpha)\cdot\beta_0$ modulates the restorative pull. Higher values of $\alpha$ increase sensitivity to environmental variation, while lower values maintain the emotional state closer to its baseline. Within the proposed framework, personality determines the PAD baseline $PAD_0$ through the Mehrabian mapping (Equation (1)), while $\alpha$ is matched to the behavioral intent of the personality profile—a stability-favoring profile warrants a lower value, a reactivity-favoring profile a higher one. The operating region for the Calm profile is characterized empirically in Section 5.7. In contrast, $\beta_0$ is fixed at 0.30 as a design constant representing the minimum pull strength required to guarantee bounded recovery across all tested scenarios and personality profiles. Values below $\beta_0 = 0.30$ produced either prolonged deviation from $PAD_0$ or slow recovery convergence; $\beta_0 = 0.30$ consistently ensured bounded convergence (effective pull = 0.129/step). Higher values were not pursued to avoid overly aggressive attraction toward the baseline that would suppress legitimate event responsiveness. Figure 3 shows the resulting mean PAD trajectory for Scenario 2 (Sc2), annotated with the three PPE state phases.

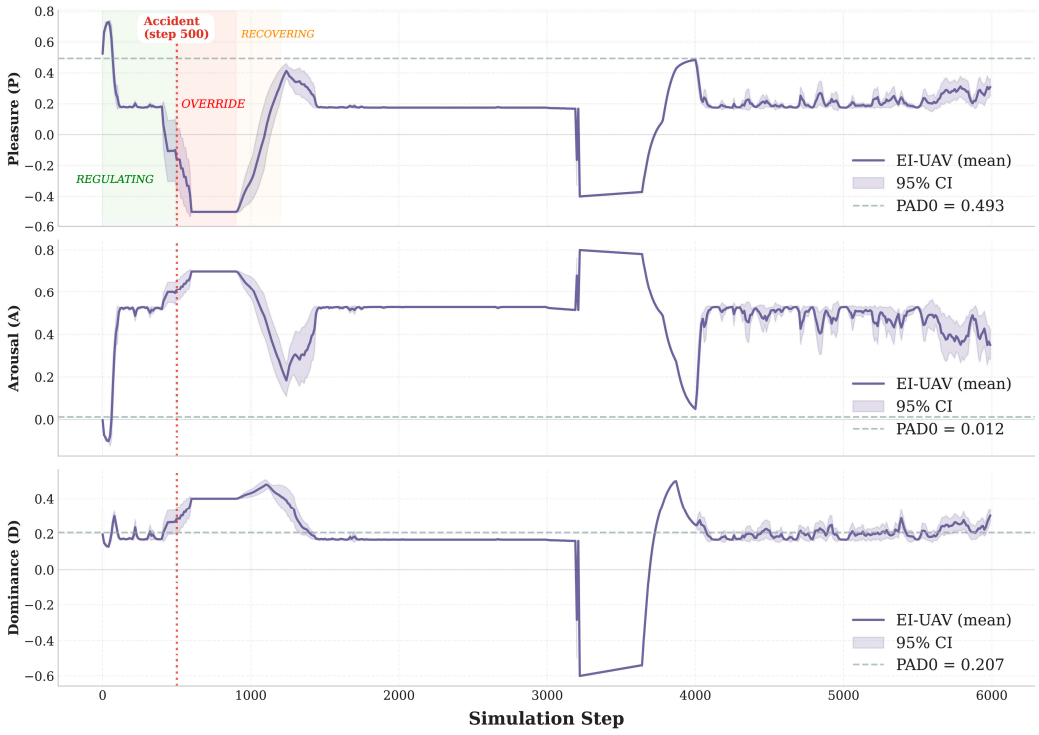


**Figure 3.** PAD trajectory (Sc2) across PPE state phases (EI-UAV, $n$ = 12, 95% CI). Shaded bands: REGULATING (green), OVERRIDE (red, onset step ≈ 490–600 across seeds), RECOVERING (within a ~3000-step window). Dotted lines: Calm personality baselines ($P_0$ = 0.493, $A_0$ = 0.012, $D_0$ = 0.207). Note that battery depletion from 100% to 20% also spans approximately 3000 steps; the two cycles overlap in timing, meaning RECOVERING continues from a shifted PAD starting point after recharge rather than completing uninterrupted. The PPE constrains saturation (0.0%) while preserving full event responsiveness.

### *4.4. PAD Bounds*

To prevent the emotional state from reaching extreme values, the Pull–Push Engine applies limit constraints after each update. Each PAD dimension is clamped to the standard interval [−1, +1] (Equation (5)):

$$PAD_{clamped} = \text{clamp}(PAD_{raw}, -1, +1) \quad (5)$$

where $PAD_{raw}$ denotes the unclamped update. The system cannot leave the defined PAD range regardless of input intensity, keeping emotional state interpretable and operationally valid.

### 4.5. Properties of PPE Regulation

Three properties of the PPE are particularly significant for prolonged missions. First, the recursive update acts as a temporal filter, smoothing short-term fluctuations while permitting sustained trends to shift state over time. Second, PAD bounds prevent extreme values regardless of input intensity. Third, incorporation of recent emotional history within the update mechanism produces implicit hysteresis. As a result, mood transitions occur only under persistent forcing rather than transient perturbations, thereby reducing oscillatory behavior near mood classification boundaries. Lower $\alpha$ values yield more gradual responses (as with the Calm profile), while higher values increase reactivity, yet all dynamics stay bounded. Mathematically, Equation (4) is equivalent to a first-order IIR filter with a forced decay term toward the personality baseline. The architectural contribution lies not in the scalar arithmetic but in the tripartite integration of PAD, OCC, and OCEAN as a regulatory ensemble, and in the three-state behavioral machine that governs which pathway is active at each step.

## 5. Experimental Analysis

### 5.1. Experimental Design

**Simulation Environment.** All experiments used the Simulation of Urban Mobility (SUMO v1.22.0; German Aerospace Center (DLR), Berlin, Germany; https://sumo.dlr.de, accessed on 6 May 2026) platform [37,38] with the Traffic Control Interface (TraCI, bundled with SUMO v1.22.0) [39]. The simulated environment was the Bologna urban network, featuring eight monitoring zones (A–H) with arterial roads, multi-lane segments, and signalized intersections representative of metropolitan traffic. This is depicted in Figure 4, where Zone A is the base station and Zone F is the accident location used in Scenario 2 (Sc2) and Scenario 3 (Sc3). Zone colors indicate instantaneous congestion severity: red (Severe), yellow (High), green (Moderate), cyan (Low). The semi-transparent rectangle marks the EI-UAV's current field of view (FOV), which adapts to its active strategy; the yellow dotted trail shows its flight path. The inset details Zone F under severe congestion.

Experiments were conducted with a single UAV; multi-UAV coordination is reserved for future work. The UAV operated under realistic operational constraints, including strategy-dependent altitude (40–70 m, 120 m maximum safety ceiling), adjustable speed (8–20 km/h), 100% initial battery, and telemetry logging every 10 steps. Each run lasted 6000 steps, representing approximately 100 min of mission time. The UAV is represented using a kinematic point-mass model. In each simulation step, the agent state is described by horizontal position (x, y), altitude (z), heading angle, speed, and yaw rate. Horizontal motion is updated as a first-order kinematic step (position += velocity $\times$ step length), with altitude, speed, and yaw rate bounded by the operational constraints listed above and modulated by the active strategy (Appendix B). Aerodynamic forces, body inertia, and low-level flight-control dynamics are not modeled, as the present study focuses on emotional regulation at the decision-making level rather than on flight-dynamics control. This abstraction is consistent with mission-level UAV simulation studies and is appropriate for the stated scope; physics-fidelity refinements are reserved for hardware-in-the-loop evaluation in future work.

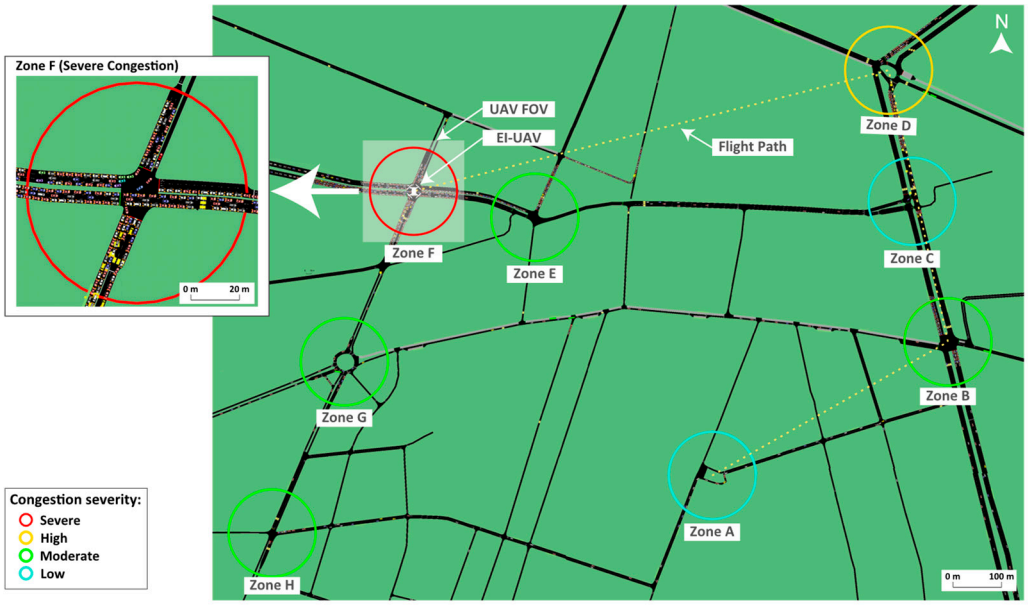


**Figure 4.** SUMO simulation environment based on the Bologna urban network, with the eight monitoring zones (Zone A–Zone H, 80 m radius) used in this study.

**Scenario Design.** Three scenarios test emotional regulation under progressively demanding conditions. Scenario 1, representing routine patrol, reflects standard operations characterized by moderate traffic variability and the absence of major incidents. Congestion fluctuates between low and moderate levels, thereby establishing baseline behavioral patterns. Scenario 2, representing a single accident response, introduces a multi-vehicle accident at step 500. This event generates sustained severe congestion in Zone F for the duration of the accident event (400–600 steps across seeds) and assesses the system's affective response to a discrete emergency together with its capacity for prolonged crisis management. Scenario 3, representing multi-stressor conditions, integrates five concurrent stressors comprising accident response, signal loss at step 1500, thermal warning at step 2500, operator directives at step 2000, and progressive battery depletion, testing whether emotional coherence holds when multiple stressors compete for attention.

**Ablation Design.** To isolate the contribution of PPE regulation, each scenario was supplemented with ablation runs in which $\alpha$ was set to 1.0, eliminating the personality pull term ($(1 - \alpha)\cdot\beta_0 = 0$). Discrete OCC OVERRIDE behavior is identical across conditions; the ablation isolates the continuous personality-anchored recovery term by removing the pull component during REGULATING states, while intentionally preserving discrete event handling. Under ablation, PAD trajectories are driven solely by environmental inputs and OCC events with no recovery toward personality baselines. A fully unregulated variant was also tested without clamping, but strategy switching was so frequent that mission-level evaluation was not feasible, confirming the need for explicit regulation.

**Configuration.** All scenarios used the Calm personality profile—selected for its low-arousal baseline appropriate to routine surveillance conditions, consistent with the rationale in Section 3.1—with Big Five trait values: O = 0.4, C = 0.7, E = 0.3, A = 0.6, N = 0.4. Applying the Mehrabian mapping (Equation (1)) yields the PAD baseline: $P_0 = 0.493$, $A_0 = 0.012$, $D_0 = 0.207$. The PPE responsiveness parameter was set to $\alpha = 0.57$, with personality coupling constant $\beta_0 = 0.30$.

**Data Collection.** Telemetry, covering the PAD state, mood classification, active strategy, and operational variables, was logged every 10 simulation steps. Recharge pauses reduce usable observations to approximately 556–563 per run depending on scenario (Table 2). Each scenario was repeated 12 times with different random seeds ($n$ = 36 EI-UAV

runs). Traffic micro-dynamics vary across seeds; PPE regulation compresses output variability, which itself evidences the mechanism's stabilizing effect. These were supplemented by 12 ablation runs per scenario ($n$ = 36, $\alpha$ = 1.0), matching the EI-UAV sample size to support rigorous statistical comparison. Setting $\alpha$ = 1.0 yields a maximally environment-responsive agent, not a pathological configuration, providing the strongest possible contrast with the personality-regulated condition. Scenario-level effects within the EI-UAV condition were tested using one-way ANOVA with $\eta^2$ as effect size (Section 5.8); comparisons between PPE-regulated and ablation runs used Mann–Whitney U tests because near-zero variance in ablation saturation distributions violates the normality assumption required for parametric tests. Because the Mann–Whitney U test is based on rank ordering rather than distributional assumptions, it remains valid even under low-variance distributions. Effect sizes are reported alongside $p$-values throughout to support robust statistical interpretation. Saturation is defined as any step where $|\text{PAD\_d}| \geq 0.999$ for any dimension d $\in$ {P, A, D}, reflecting practical loss of affective headroom rather than exact clipping at $\pm$1.0. The prominence of S6 in Scenario 1 (Sc1) routine patrol follows directly from the Calm profile's near-zero Arousal baseline ($A_0$ = 0.012): routine traffic variability persistently elevates Arousal into the high-Arousal region occupied by the Surprised centroid, which maps to S6 as its primary strategy (Appendix B, Table A2).

**Table 2.** Simulation Run Statistics.

| Scenario | Zones Visited | Strategy Changes | Events Detected | Recharge Cycles | Data Points |
|---|---|---|---|---|---|
| Sc1 (routine patrol) | 66.8 $\pm$ 3.0 | 17.8 $\pm$ 1.9 | 0 | 1 | 6682 |
| Sc2 (Accident) | 58.2 $\pm$ 3.9 | 17.6 $\pm$ 1.5 | 1 | 1 | 6675 |
| Sc3 (Multi-Stressor) | 48.5 $\pm$ 2.9 | 29.3 $\pm$ 2.7 | 5 | 1 | 6753 |

**Reproducibility.** The experimental setup is designed to ensure reproducibility. It uses twelve fixed random seeds (1–12) to control traffic generation, event timing, and stochastic perturbations across three SUMO-based scenarios of increasing complexity. The resulting dataset comprises 36 EI-UAV runs and 36 ablation runs ($\alpha$ = 1.0), corresponding to 12 runs per scenario. All simulations were executed using identical configuration settings across runs, ensuring consistency in comparative evaluation. Controller parameters ($\alpha$, $\beta_0$, $\varepsilon$, PAD bounds, decay rates, OCC mappings, and the mood-strategy table) are fully specified in Appendix B, Appendix C, and Appendix D. The simulation framework is the open-source SUMO microsimulator [37,38] coupled to the controller through TraCI [39], using the publicly documented Bologna urban network. Complete simulation logs and per-run summary statistics are available from the corresponding author on reasonable request.

Values are mean $\pm$ SD per run across 12 seeds (Zone-A entries exceeding 50 dwell steps excluded as charging stops). Observations per run vary slightly by scenario (557, 556, and 563 for Sc1, Sc2, and Sc3) due to recharge timing. Events detected match scenario design (0, 1, 5). Strategy changes remain comparable between Sc1 and Sc2 (17.8 and 17.6 per run), then rise sharply in Sc3 (29.3 per run), reflecting the additional behavioral demands introduced by concurrent stressors.

### *5.2. PAD Trajectory Analysis*

All three PAD dimensions remained within operational bounds (Table 3) with 0.0% saturation across all 36 regulated runs despite the same $\pm$1.0 clamp applied in both conditions. Ablation runs ($\alpha$ = 1.0, removing personality pull) exhibited 57–88% saturation, with values driven to the clamp boundary and held there. The contrast is explicit: PPE regulation actively prevents saturation through continuous, personality-anchored pull, whereas clamping alone merely truncates already extreme values. Each scenario leaves

a distinct PAD signature consistent with its operational load, with bounded oscillations near baseline in Sc1, pronounced shifts during emergency monitoring in Sc2, and varied dynamics under competing influences in Sc3. Between events, PAD values evolve smoothly as Equation (4) filters out short-term noise. At event onset, abrupt transitions occur consistent with the OVERRIDE state, reflecting immediate affective adjustment to salient stimuli. The system responds to genuine events within 10 timesteps (e.g., accident detection modulating arousal) while filtering transient disturbances, and the emotional state gradually returns toward personality-anchored baselines once pressure subsides. Figure 5 shows the mean PAD trajectories aggregated across all runs ($n$ = 12 per scenario, with 95% confidence intervals).

**Table 3.** Mood Distribution Across Scenarios (EI-UAV with PPE).

| Mood State | Sc1 (Routine Patrol) | Sc2 (Accident) | Sc3 (Multi-Stressor) |
|---|---|---|---|
| Anxious | 1.4% ± 0.0% | 11.7% ± 0.1% | 15.7% ± 0.4% |
| Excited | Not observed | Not observed | Not observed |
| Satisfied | 15.1% ± 2.8% | 7.8% ± 0.9% | 12.3% ± 1.3% |
| Joyful | Not observed | Not observed | Not observed |
| Distressed | 0.9% ± 0.0% | 0.9% ± 0.1% | 4.6% ± 1.0% |
| Focused | 8.0% ± 1.3% | 9.4% ± 1.5% | 15.5% ± 1.1% |
| Relieved | 8.5% ± 2.4% | 2.7% ± 0.3% | 4.1% ± 0.2% |
| Neutral | Not observed | Not observed | Not observed |
| Surprised | 66.0% ± 1.5% | 67.5% ± 1.6% | 47.8% ± 1.9% |
| Exhausted | Not observed | Not observed | Not observed |

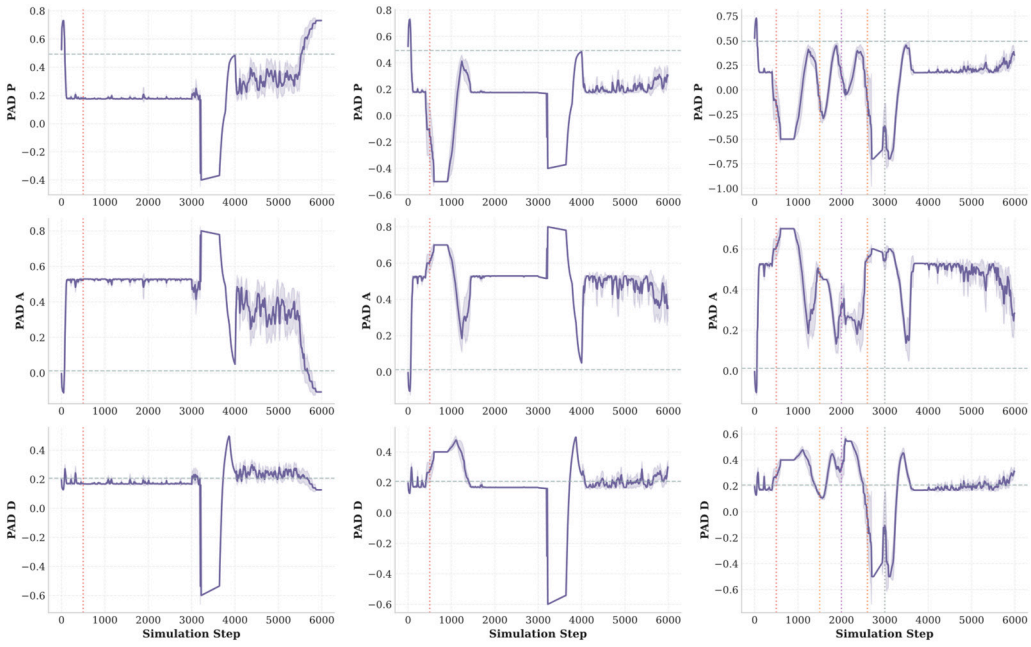


**Figure 5.** Mean PAD trajectories across three scenarios (EI-UAV, $n$ = 12 per scenario; shaded bands: 95% confidence interval). Dotted red lines: recharge cycles; dotted orange lines: event onsets. Dotted grey lines: personality baselines. Sc1 shows stable near-baseline oscillations; Sc2 shows pronounced event response in Arousal and Pleasure; Sc3 shows repeated multi-stressor dynamics.

### *5.3. Dual-Pathway Emotional Response*

The accident illustrates discrete event processing. The PPE-regulated response set PAD directly to the Anxious OCC target at onset (A = 0.70), then gradually returned toward the personality baseline ($A_0$ = 0.012) through the RECOVERING phase, which follows a graduated path (Focused → Satisfied → $PAD_0$) within the ~3000-step recharge window. In Sc1, where no follow-on events occur, recovery episodes generally complete within the

available window; in Sc2 and Sc3, subsequent events or recharge cycles frequently interrupt RECOVERING before full convergence, meaning the pull toward the personality baseline is sustained and directional rather than a guaranteed episode completion. The continuous pathway shows a complementary pattern. Battery depletion proceeds over the same ~3000-step window—declining from 100% to ~20% before recharge—meaning the two cycles frequently overlap in timing. The Battery–Dominance relationship varies by scenario (Sc1: r = +0.16, Sc2: r = +0.26, Sc3: r = −0.17): in lower-complexity scenarios, higher battery coincides with elevated Dominance during active patrol, while multi-stressor conditions progressively compress Dominance independent of battery level, reflecting the accumulated weight of concurrent stressors rather than depletion alone. This asymmetry—immediate PAD jumps for discrete events versus incremental drift for gradual conditions—is precisely the regulatory behavior the dual-pathway design produces, as illustrated in Figure 6.

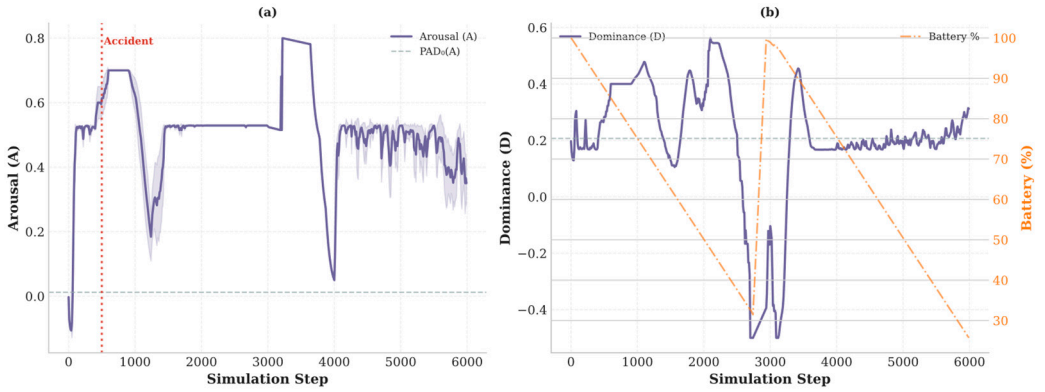


**Figure 6.** Dual-pathway emotional response. (**a**) Discrete pathway: mean Arousal trajectory in Sc2 ($n$ = 12; shaded band: 95% confidence interval) showing the OVERRIDE phase (onset A = 0.70) and RECOVERING phase (within a ~3000-step window). (**b**) Continuous pathway: Battery–Dominance relationship per scenario. Grey dashed lines: personality baselines ($PAD_0$(A) = 0.012 in panel (**a**), $PAD_0$(D) = 0.207 in panel (**b**)); orange dash-dotted line: battery percentage trajectory (right axis). Correlation varies by scenario (Sc1: r = +0.16, Sc2: r = +0.26, Sc3: r = −0.17), reflecting scenario-dependent rather than uniform depletion dynamics.

### *5.4. Mood Distribution Analysis*

Table 3 shows how mood states distribute across scenarios, and Figure 7 presents these patterns as a heatmap (mean across 12 runs per scenario). Six moods were active; Excited, Joyful, Neutral, and Exhausted were not observed under PPE regulation with the Calm profile and are marked 'Not observed' in Table 3. Surprised dominated across all three scenarios (66.0%, 67.5%, 47.8%), contracting sharply in Sc3 as stress-related moods expanded. Anxious rose from 1.4% in Sc1 to 11.7% in Sc2 and 15.7% in Sc3; Focused increased from 8.0% to 9.4% to 15.5%. Together, Anxious and Focused accounted for 31.2% of Sc3 steps, reflecting concentrated crisis-oriented attention. Distressed increased more than fivefold to 4.6% in Sc3 but remained well below pathological levels. The non-monotonic pattern in Satisfied (15.1% → 7.8% → 12.3%) warrants brief comment: the partial recovery in Sc3 reflects structured goal-achievement episodes created by operator focus directives at step 2000, which elevate Pleasure between acute stressor events and temporarily reactivate the Satisfied region. These shifts were consistent across all 12 runs.

Under routine patrol (Sc1), Surprised accounted for 66.0% of mission time, consistent with a Calm profile monitoring steady traffic. Anxious and Distressed together remained at 2.3%, confirming the absence of runaway negative affect. When an accident occurred (Sc2), Surprised remained dominant at 67.5% while Anxious rose sharply to 11.7% and Focused emerged at 9.4% as attention concentrated on the incident zone. Under maximum complexity (Sc3), Surprised contracted to 47.8% as stress-related moods expanded: Anxious

reached 15.7%, Focused 15.5%, and Distressed rose more than fivefold to 4.6%. These shifts were consistent across all 12 runs.

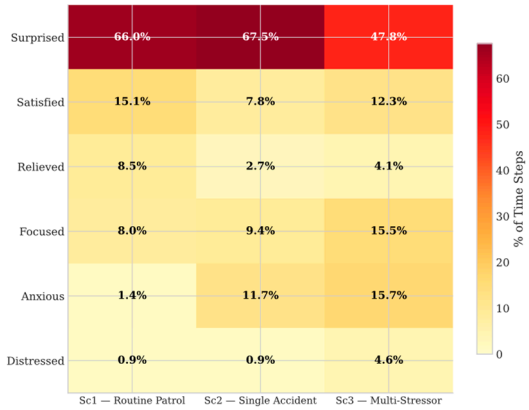


**Figure 7.** Mood distribution heatmap (EI-UAV, % of simulation steps, $n$ = 12 per scenario). Surprised dominates all scenarios (66.0%, 67.5%, 47.8%). Anxious and Focused expand progressively with operational complexity. Distressed increases more than fivefold from Sc1 to Sc3. Four mood states (Excited, Joyful, Neutral, Exhausted) were not observed under the Calm profile and are omitted from the figure.

The dominance of Surprised across all scenarios reflects the dynamic behavior of the PAD space under PPE-damped patrol conditions. During routine monitoring, traffic variability continuously elevates Arousal above its near-zero baseline ($A_0$ = 0.012), shifting the operating point into the high-Arousal region of PAD space where the Surprised centroid is located. The PPE's damping effect prevents sustained excursions into other mood regions under moderate stimulation, keeping the system within the Surprised classification boundary during routine patrol, and its contraction to 47.8% in Sc3 confirms this is not a degenerate absorbing state.

### *5.5. Behavioral Diversity and Strategy Adaptation*

Table 4 presents strategy diversity measures, and Figure 8 visualizes the trends (mean ± SD across repeated runs). Shannon entropy H is computed as $H = -\Sigma\, p_i \log_2 p_i$, where $p_i$ is the proportion of steps using strategy i. Panel (a) shows entropy increasing from 1.885 to 2.033 bits with effective strategy count rising from 3.69 to 4.09; panel (b) shows unique strategies constant at 5 across all scenarios, while dominant strategy concentration decreases from 39.0% to 36.4%.

The UAV exhibited five strategies across all three scenarios. In Sc1, S2 (Weighted Priority) and S6 (Operator Directed/Emergency) are nearly equally distributed at 39.0% and 38.5%, respectively, reflecting a mixed patrol regime in the absence of discrete events. Under accident conditions (Sc2), S6 becomes the dominant strategy at 38.9% as emergency dynamics shift the mood distribution toward Surprised and Focused states. In Sc3, S6 remains dominant at 36.4% with S2 at 31.2%, and S0 (Safe Mode) rises to 11.6% as concurrent stressors activate Distressed and Anxious mood states that route to conservative operational postures. Strategy entropy increases monotonically with complexity: 1.885 bits (Sc1), 1.937 bits (Sc2), and 2.033 bits (Sc3), corresponding to effective strategy counts of 3.69, 3.83, and 4.09. Strategy changes per run remain comparable between Sc1 and Sc2 (17.8 and 17.6), then rise sharply in Sc3 (29.3), consistent with the increased frequency of competing stressor events requiring behavioral reallocation. These patterns were consistent across all 12 runs (Section 5.8).

**Table 4.** Behavioral Diversity Measures.

| Metric | Sc1 (Routine Patrol) | Sc2 (Accident) | Sc3 (Multi-Stressor) |
|---|---|---|---|
| Unique Strategies Observed | 5 | 5 | 5 |
| Strategy Entropy (bits) | 1.885 ± 0.055 | 1.937 ± 0.057 | 2.033 ± 0.062 |
| Effective Strategy Count (2^H) | 3.69 | 3.83 | 4.09 |
| Most Used Strategy | S2 | S6 | S6 |
| Most Used Percentage | 39.0% | 38.9% | 36.4% |
| Strategy Changes/Run | 17.8 ± 1.9 | 17.6 ± 1.5 | 29.3 ± 2.7 |

*Note:* In Sc1 (routine patrol), S6 (Emergency Response) accounts for 38.5% of steps, co-dominant with the leading strategy S2 (39.0%).

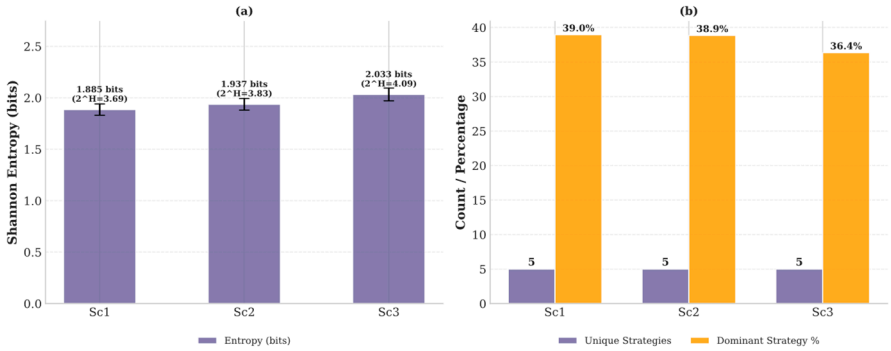


**Figure 8.** Behavioral diversity measures (EI-UAV, mean ± SD, $n$ = 12 per scenario). (**a**) Shannon entropy increases monotonically: 1.885 → 1.937 → 2.033 bits; effective strategy count $2^H$ rises from 3.69 to 4.09. (**b**) Dominant strategy concentration decreases modestly (39.0% → 38.9% → 36.4%) while unique strategy count remains constant at 5 per run, indicating diversity driven by frequency redistribution rather than new strategy activation.

### *5.6. Decision Interpretability*

Every behavioral decision can be traced from environmental stimulus through emotional processing to strategy execution. Unlike black-box learning approaches, the full causal chain is open to inspection. The following box illustrates this process using actual data from Scenario 2.

As Box 1 illustrates, operators can review decision rationale, engineers can adjust settings without redesigning the system, and regulators can inspect causal chains, something black-box approaches cannot provide. The trace shown in Box 1 is qualitatively consistent across all 12 seeds for Scenario 2; seed-specific variation appears only in accident onset timing (400–600 steps), not in the decision logic itself. Table 5 summarizes how the experimental results map to the system's design objectives. All nine objectives were met across the three test scenarios, with scenario-level trends confirmed by one-way ANOVA, as summarized in Table 6.

**Table 5.** Validation Summary.

| Design Objective | Validation Evidence |
|---|---|
| Bounded emotional evolution | PAD saturation 0.0% (36 runs) vs. 57–88% ablation; PAD_P max = +0.731 |
| Smooth state transitions | Continuous trajectories; recursive update filters short-term fluctuations |
| Event responsiveness | Arousal adjusted to A = 0.70 at accident onset; discrete OCC response confirmed |
| Gradual condition adaptation | Battery–Dominance r = +0.16/+0.26/−0.17 across scenarios; incremental response verified |
| Context-appropriate moods | Surprised dominant (66.0%/67.5%/47.8%); Anxious and Focused rise with complexity |
| Absence of pathology | Exhausted not observed; Distressed peaks at 4.6% in Sc3 |
| Behavioral diversity scaling | Entropy 1.885→1.937→2.033 bits; strategy changes 17.8/17.6→29.3/run |
| Adaptive strategy selection | S2/S6 co-dominant Sc1; S6 leads Sc2–Sc3; S0 rises to 11.6% in Sc3 |
| Complete interpretability | Complete causal traceability confirmed across all decision steps |

**Table 6.** Scenario-Level Confirmation of Key Trends.

| Metric | Directional Trend | $\eta^2$ |
|---|---|---|
| Strategy Entropy | Increases monotonically with scenario complexity | 0.55 |
| PAD Dominance | Increases monotonically with scenario complexity | 0.92 |
| Surprised Mood | Contracts under multi-stressor conditions | 0.97 |
| Strategy Changes | Stable Sc1→Sc2, sharp increase in Sc3 | 0.88 |

Note: $\eta^2$ = proportion of variance explained (ANOVA; see Section 5.8 for full results and Figure 9).

**Box 1.** Decision Trace: Accident Response in Scenario 2 (Steps 490–500+).

Step 490: [PRE-EVENT STATE]
├── PAD = (+0.18, +0.53, +0.17)
├── Mood: Surprised
├── Strategy: S6 (Operator Directed/Emergency)
└── Congestion level: Moderate

Step 500: [ENVIRONMENTAL STIMULUS]
├── Accident detected in Zone F
├── Traffic density: sharp increase
├──Average speed: sharp decrease
└── Congestion level: Moderate → Severe

Step 500: [OCC EVENT APPRAISAL]
├── Event Type: Accident (High-Severity Emergency)
├── Priority: 1 (highest)
├── Intensity: λ = 0.8 (high-priority event)
└── PAD_target: (P = −0.50, A = +0.70, D = +0.40)

Step 500: [PPE OVERRIDE—Algorithm 1, lines 1–4]
├── Discrete OCC event detected
├── PAD directly set to event target
├── PAD = (−0.50, +0.70, +0.40)
└── PPE State → OVERRIDE

Step 500: [MOOD CLASSIFICATION]
├── PAD = (−0.50, +0.70, +0.40)
├── Nearest Centroid: Anxious (P = −0.15, A = +0.55, D = +0.25)
└── Classification: "Anxious"

Step 500: [STRATEGY SELECTION]
├── Mood: Anxious → Primary: S2, Secondary: S1
├── Priority Override: Discrete OCC Priority-1 events override
├── mood→strategy mapping and trigger S6 for direct
├── emergency deployment (Operational Constraints, Section 3, priority hierarchy)
├── Context: Incident Active (Zone F); Battery adequate; Comms normal
└── Selected: S6 (Operator Directed/Emergency)

Step 500+: [BEHAVIORAL EXECUTION]
├── Target Zone: F (accident location)
├── Strategy: S6—emergency deployment
├── OVERRIDE persists while accident remains active (onset 400–600 steps across seeds)
└── PPE enters RECOVERING once accident clears.

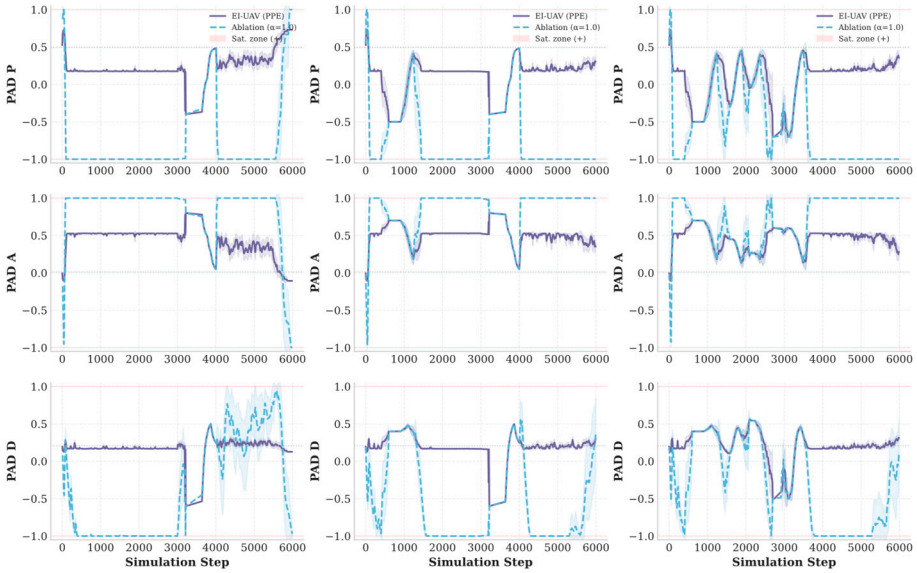


**Figure 9.** PAD trajectories: EI-UAV (PPE-regulated) versus ablation ($\alpha = 1.0$) across three scenarios ($n = 12$ EI, $n = 12$ ablation; shaded bands: 95% confidence interval, purple for EI-UAV and blue for ablation). Ablation runs exhibit 57–88% PAD saturation across scenarios; PPE-regulated runs maintain 0.0% saturation. The contrast confirms that homeostatic pull, not clamping, prevents saturation.

### *5.7. Parameter Sensitivity Analysis*

To characterize the operating region of $\alpha$ for the Calm personality profile, a sensitivity analysis was conducted on Scenario 3 (multi-stressor), the most demanding condition. A stable operating region identified under maximum stressor load is guaranteed to remain stable under lighter conditions; Sc3 is therefore the conservative and sufficient choice for characterizing the operating boundary. Five values were tested, $\alpha \in \{0.40, 0.57, 0.63, 0.70, 0.75\}$, spanning the range from pull-dominant to push-dominant regulation. Each was run 12 times with $\beta_0 = 0.30$ fixed. Table 7 summarizes the results.

**Table 7.** $\alpha$ Sensitivity Analysis Results—Scenario 3 (Multi-Stressor).

| Metric | $\alpha = 0.40$ | $\alpha = 0.57$ | $\alpha = 0.63$ | $\alpha = 0.70$ | $\alpha = 0.75$ |
|---|---|---|---|---|---|
| PAD Saturation (%) | $0.0 \pm 0.0$ | $0.0 \pm 0.0$ | $0.0 \pm 0.0$ | $0.0 \pm 0.0$ | $46.4 \pm 2.4$ |
| Strategy Entropy (bits) | $1.953 \pm 0.085$ | $2.033 \pm 0.062$ | $2.054 \pm 0.068$ | $2.044 \pm 0.060$ | $1.953 \pm 0.079$ |
| Effective Strategies ($2^H$) | 3.87 | 4.09 | 4.15 | 4.12 | 3.87 |
| Strategy Transitions/Run | $37.2 \pm 1.9$ | $29.3 \pm 2.7$ | $28.5 \pm 2.5$ | $29.2 \pm 2.1$ | $35.0 \pm 2.0$ |
| Unique Moods Observed | 5 | 6 | 6 | 6 | 6 |

The sweep identifies a stable operating region across $\alpha \in \{0.57, 0.63, 0.70\}$, with zero PAD saturation and consistent entropy (2.03–2.05 bits) throughout this range. Performance degrades clearly at $\alpha = 0.75$ (46.4% saturation), marking the boundary of the stable region. Within the stable region, lower $\alpha$ values favor stability (reduced environmental reactivity) while higher values favor responsiveness. For the Calm profile—designed for sustained surveillance with a high-Conscientiousness, low-arousal baseline—$\alpha = 0.57$ represents the stability-consistent operating point at the lower bound of the stable region, appropriate to its behavioral intent. $\alpha = 0.40$ also achieved zero saturation but exhibited lower entropy (1.953 bits) and higher transition rates (37.2 transitions/run), indicating excessive pull-dominance that suppresses environmental responsiveness—the opposite failure mode to $\alpha = 0.75$. When the pull is too strong, PAD returns to the personality baseline so rapidly that any minor perturbation is sufficient to cross a mood boundary, increasing strategy switching.

### *5.8. Statistical Validation*

This subsection reports statistical comparisons across the three scenarios. Table 8 summarizes one-way ANOVA outcomes with $\eta^2$ effect sizes for each PAD dimension, and Figure 9 shows the corresponding trajectory comparisons against the ablation baseline.

**Table 8.** Statistical Comparisons Across Scenarios.

| Metric | Sc1 vs. Sc2 | Sc2 vs. Sc3 | Sc1 vs. Sc3 | ANOVA |
|---|---|---|---|---|
| Strategy Changes | d = +0.10 ns | d = −5.36 *** | d = −4.94 *** | F(2,33) = 123.22, $\eta^2$ = 0.88 *** |
| Strategy Entropy | d = −0.93 ns | d = −1.61 ** | d = −2.51 *** | F(2,33) = 19.93, $\eta^2$ = 0.55 *** |
| Mood: Surprised (%) | d = −0.95 ns | d = +11.18 *** | d = +10.39 *** | F(2,33) = 503.85, $\eta^2$ = 0.97 *** |
| Mood: Focused (%) | d = −1.00 ns | d = −4.71 *** | d = −6.28 *** | F(2,33) = 115.87, $\eta^2$ = 0.88 *** |
| PAD P (mean) | d = +18.43 *** | d = +5.57 *** | d = +16.47 *** | F(2,33) = 1270.72, $\eta^2$ = 0.99 *** |
| PAD A (mean) | d = −10.62 *** | d = +8.77 *** | d = −5.02 *** | F(2,33) = 405.26, $\eta^2$ = 0.96 *** |
| PAD D (mean) | d = −7.16 *** | d = −2.20 *** | d = −6.50 *** | F(2,33) = 184.56, $\eta^2$ = 0.92 *** |

*Note 1:* ** $p < 0.01$, *** $p < 0.001$ (Bonferroni-corrected). Cohen's d (pairwise), $\eta^2$ (ANOVA). $n = 12$ per scenario. *Note 2:* PAD P $\eta^2 = 0.99$ reflects near-zero within-group variance under PPE regulation, a direct consequence of personality anchoring.

### *5.9. Multi-Personality Validation*

All five personality profiles maintained zero PAD saturation and bounded trajectories across all 60 runs (Table 9), confirming that PPE regulation generalizes beyond the Calm baseline used in the primary experiments. Oscillation rates remained low and consistent across profiles (4.6–5.3%; Figure 10b), indicating that personality-induced differences in PAD baselines do not destabilize the regulation mechanism. With $\alpha$ held constant at 0.57 across all profiles, two distinct behavioral clusters emerged (Figure 10a,c), organized by the Dominance dimension of the personality-derived PAD baseline ($D_0$). Low-$D_0$ profiles—Calm ($D_0$ = 0.207) and Cautious ($D_0$ = 0.224)—produced Surprised-dominant mood distributions (46–48% of steps) and higher strategy entropy (2.02–2.03 bits), consistent with continuous PAD movement in the low-Dominance region under dynamic patrol conditions. High-$D_0$ profiles—Investigative ($D_0$ = 0.382), Confident ($D_0$ = 0.512), and Proactive ($D_0$ = 0.614)—produced Focused-dominant distributions (~66% of steps) and consistent lower entropy (1.75–1.76 bits), as their elevated Dominance baseline concentrates PAD trajectories in assertive mood regions under identical stimulation. This clustering confirms that with $\alpha$ controlled, the Dominance component of $PAD_0$ determines which emotional region the PPE trajectory occupies, which in turn governs strategy diversity. All five profiles activated the same five strategies (S0–S3 and S6), confirming a stable reachable strategy set under the present PPE configuration. These results demonstrate that the PPE does not impose uniform behavior: personality modulates behavioral character through the PAD baseline without compromising regulation stability. These results are based on Sc3 alone; full cross-scenario statistical evaluation across personality profiles remains for future work, consistent with the scope acknowledged in the Limitations.

**Table 9.** PPE Performance Across Personality Profiles.

| Personality | PAD Baseline (P, A, D) | PAD Bounded | Oscillation | Strategies | Entropy (bits) |
|---|---|---|---|---|---|
| Calm | (0.493, 0.012, 0.207) | Yes | 5.2% | 5 | 2.03 ± 0.06 |
| Confident | (0.499, 0.069, 0.512) | Yes | 4.6% | 5 | 1.76 ± 0.18 |
| Proactive | (0.520, 0.084, 0.614) | Yes | 4.6% | 5 | 1.76 ± 0.19 |

**Table 9.** *Cont.*

| Personality | PAD Baseline (P, A, D) | PAD Bounded | Oscillation | Strategies | Entropy (bits) |
|---|---|---|---|---|---|
| Investigative | (0.455, 0.042, 0.382) | Yes | 4.6% | 5 | $1.75 \pm 0.18$ |
| Cautious | (0.512, −0.045, 0.224) | Yes | 5.3% | 5 | $2.02 \pm 0.05$ |

Notes: Strategies = total unique strategies observed in Sc3 runs for each profile; Table 4 reports per-scenario counts for the Calm profile across all three scenarios, which may be lower when not all strategies are triggered in a given scenario. Oscillation = strategy change rate per 100 simulation steps. All profiles verified against $n = 12$ runs (Sc3, $\alpha = 0.57$; 60 runs total across five profiles).

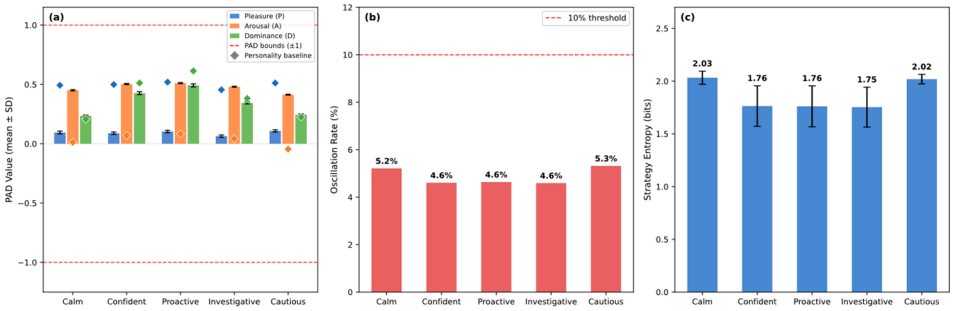


**Figure 10.** Multi-personality PPE validation across five profiles. (**a**) Mean PAD values per dimension (grouped bars, mean ± SD; diamonds = personality baseline). (**b**) Strategy oscillation rate per profile; dashed line = 10% stability threshold. (**c**) Mean strategy entropy with standard deviation ($n = 12$ runs per profile, Sc3, $\alpha = 0.57$).

## 6. Discussion

### *6.1. Emotion as Functional Decision Control*

The results support the claim: emotion can serve as a functional decision-control mechanism when appropriately regulated. This extends Cañamero's [20] argument that affect enables flexible adjustment to surroundings and addresses the gap identified by Nunnari et al. [21], who demonstrated PAD-based behavioral modulation without the explicit regulation needed for sustained autonomous control. What distinguishes the present work is that strategy selection follows directly from the regulated PAD state rather than from a separate planning layer. Traffic congestion, battery depletion, incident severity, and communication status all feed into a single affective representation whose temporal evolution indicates how these factors interact, persist, and resolve, something difficult to replicate with threshold logic, where each factor is handled independently. The mood distributions across scenarios illustrate this pattern. During routine patrol, Surprised accounts for 66.0% of mission time, appropriate for a Calm profile monitoring uneventful traffic, while Distressed and Anxious together remain at 2.3%. As conditions deteriorate, the distribution shifts to match operational demands, despite the absence of explicit rules linking the number of stressors to a predefined stress level. These patterns emerge from the weighted integration performed by the PPE and arise as natural consequences of its underlying regulatory dynamics. The dominance of Surprised warrants direct comment. Its centroid (P = +0.15, A = +0.70, D = +0.10) maps primarily to S6 (Operator Directed/Emergency), which may appear to imply a perpetual quasi-alert posture. This is a calibration consequence of the Calm profile's near-zero Arousal baseline ($A_0 = 0.012$): routine traffic variability continuously elevates Arousal into the high-Arousal region where Surprised resides, even without any discrete event. Critically, Surprised contracts to 47.8% in Sc3, confirming it is not a degenerate absorbing state but rather the natural operating point of a low-baseline profile under moderately dynamic conditions. This does not indicate state collapse, as evidenced by increasing entropy and the activation of multiple strategies across scenarios. Profiles with a higher Dominance baseline (e.g., Proactive, $D_0 = 0.614$) exhibit Focused-dominant mood

distributions rather than Surprised-dominant ones, as shown in Section 5.9, confirming that Surprised dominance is profile-specific to low-Dominance-baseline profiles and not a structural deficiency of the PPE. Recent work by Hernández-Marcos and Ros [40] further validates this direction, presenting a self-learning emotional framework where temporal patterns in reward signals serve as emotion-like regulatory signals, conceptually paralleling the PPE's use of PAD dynamics to modulate decision-making.

### *6.2. The Key Role of PPE Regulation*

The PPE architecture combines two design choices: bidirectional regulation (in which a personality-anchored pull term continuously balances environmental push) and stateful temporal evolution (in which each PAD update incorporates the previous state recursively). Each choice is examined here through a separate ablation, isolating its individual contribution. Section 6.2.1 removes the bidirectional regulation by setting $\alpha = 1.0$, eliminating the pull term while preserving the recursive update structure. Section 6.2.2 removes the stateful temporal evolution by replacing the recursive update with a stateless graded controller of equivalent interpretability (Mamdani fuzzy inference). Section 6.2.3 synthesizes the two findings.

#### 6.2.1. Bidirectional Regulation Ablation ($\alpha = 1.0$)

Ablation results provide clear empirical evidence addressing a key design question: whether clamping alone suffices to keep emotion-regulated control well-behaved. It does not. With the personality pull term disabled, identical clamping produced 57–88% saturation; once PAD values reach the clamp boundary and remain there, the emotional state carries no useful information, and downstream strategy selection degenerates. The decrease in ablation saturation from Sc1 (88.3%) to Sc3 (56.6%) is counterintuitive but mechanistically explained: Sc3 contains more discrete OCC OVERRIDE events (five concurrent stressors versus one in Sc2 and none in Sc1), each of which sets PAD directly to a predefined event-specific target rather than accumulating at the clamp boundary. More override events therefore produce more time spent at interior targets, reducing measured boundary saturation despite higher overall stress. The PPE counteracts this effect not by imposing harder limits, but by continuously drawing the state back toward personality-anchored baselines. The sensitivity analysis (Section 5.7, Table 7) confirms that both suppression and saturation are simultaneously avoided within the validated stable region. The implicit hysteresis embedded in Equation 4 warrants separate consideration. Because each update incorporates the current state, transitions between mood regions occur only under sustained forcing rather than transient perturbations. This property efficiently mitigates the chattering problem identified in Section 4.1. During the seed-dependent emergency monitoring period in Scenario 2 (onset 400–600 steps across seeds), the system maintained stable, crisis-appropriate behavior because transient PAD variations could not overcome the hysteresis inherent in the recursive update.

Behavioral diversification emerges directly from emotional dynamics rather than explicit complexity detection. The system does not measure how many stressors are active. Instead, more demanding conditions generate broader PAD excursions that activate a broader range of moods and associated strategies. This distinction is critical for generalization. A threshold-based complexity detector would be required to anticipate and encode every possible combination of stressors. In contrast, the emotion-driven approach produces appropriate shifts, such as the move from S2 (Weighted Priority) to S6 (Emergency Response) under multi-stressor conditions, as direct consequences arising from accumulated PAD dynamics.

6.2.2. Stateless Graded-Control Baseline and Ablation (FLC-UAV)

The $\alpha = 1.0$ ablation isolates the bidirectional regulation component while keeping the recursive PAD update intact. A complementary question concerns the stateful evolution itself: Does graded interpretability suffice on its own, or is the stateful PAD evolution that links successive decisions a necessary element of the architecture? Mamdani fuzzy inference is a natural choice for testing this question because it provides graded interpretability comparable to the PPE's continuous mood representation, but achieves this through a stateless input–output mapping rather than through recursive temporal accumulation. Comparing the two architectures on identical operational scenarios isolates the contribution of temporal evolution while controlling for graded representation.

The FLC-UAV variant—our concrete realization of the Mamdani fuzzy inference architecture—replaces the PPE update with a six-rule fuzzy inference engine operating on two scalar inputs (Equations (6) and (7)) and producing a strategy index through centroid defuzzification (Equation (8)):

$$TS = \frac{\max(s_{zone}, s_{event})}{4} \tag{6}$$

where $s_{zone}$ is the highest severity classification (1 = Low, 2 = Moderate, 3 = High, 4 = Severe) observed across monitored zones within the freshness window, and $s_{event}$ is the highest severity of any active OCC event; the divisor 4 normalizes $TS$ to the unit interval.

$$WL = \frac{N_{active}}{5} \tag{7}$$

where $N_{active}$ is the count of currently active emergency events; the divisor 5 normalizes $WL$ to the unit interval under the maximum five-concurrent-event design budget.

$$Strategy = round(centroid(min - \max(\mu_{TS}, \mu_{WL}, R))) \tag{8}$$

where $\mu_{TS}$ and $\mu_{WL}$ are trapezoidal membership functions over the antecedent variables, R is the six-rule base over $(TS, WL) \in$ {Low, Moderate, High} $\times$ {Low, High}, min-max denotes Mamdani inference (the min operator for rule activation, the max operator for rule aggregation), centroid$(\cdot)$ is the centroid-of-area defuzzification operator, and round$(\cdot)$ snaps to the nearest integer in $\{0, 1, \ldots, 7\}$ corresponding to strategies S0 through S7.

The variant is otherwise architecturally identical to the regulated baseline: same SUMO scenarios, same twelve seeds, same operational safety overrides for signal loss, severe weather, overheating, and battery depletion (the S0 wrapper is shared across all controllers in this study and is not part of the FLC inference path). The only architectural change is the strategy selection mechanism. PAD coordinates, OCC event appraisal, personality coupling, and the recursive update of Equation (4) are all absent by construction, replaced by a memoryless input–output mapping.

The behavioral consequences manifest within the first patrol cycle and persist across all scenarios. Across 36 ablation runs, FLC-UAV strategy entropy averaged $0.46 \pm 0.14$ bits in Sc1, $0.19 \pm 0.02$ bits in Sc2, and $0.41 \pm 0.05$ bits in Sc3, compared with $1.60 \pm 0.03$ bits, $1.43 \pm 0.04$ bits, and $1.71 \pm 0.08$ bits for the PPE-regulated baseline in Sc1, Sc2, and Sc3, respectively, indicating a collapse to low-diversity strategy selection in the stateless controller, consistent with the response-time degradation reported below. These entropies are computed over per-step strategy distributions, consistent with the steady-state collapse claim of this section; the per-transition entropies reported in Table 4 (1.885–2.033 bits) characterize switching diversity rather than steady-state occupancy, and the two metrics are complementary. The fuzzy inference engine itself is mathematically responsive: per-

step traces show Traffic Stress evolving correctly through 0.0 → 0.75 → 1.0 as the UAV samples zones of increasing congestion, and the rule-firing patterns shift as expected. However, once any zone is sampled at Severe classification, the maximum operator in the TS computation saturates the input space at 1.0 within the first patrol cycle. The graded output then collapses to a near-binary regime: 92–99% of all simulation steps produce S6 (Emergency Response), with the remaining 1–8% distributed across S0 (safety wrapper) and S1 (initial pre-saturation steps). The controller produces a graded output, but operates within an input regime where stateless inference cannot recover the temporal context that the recursive PAD update encodes.

The dual-pathway response characterized in Section 5.3 is reflected operationally in event-response timing. Removing the stateful PPE update produced substantial response-time degradation on state-aware events. Accident response increased from 44.6 ± 20.5 s to 205.8 ± 72.7 s (4.6× slower; $n$ = 24 events per condition; Mann–Whitney U = 576, $p < 0.001$, r = 1.00—complete rank separation, matching the effect-size convention used in Section 6.2.1's Table 10), and operator-focus events from 50.8 ± 31.3 s to 201.8 ± 80.3 s (4.0× slower; $n$ = 12), exhibiting the same complete rank separation. In contrast, safety-wrapper events (signal loss, bad weather, overheating) showed only minor differences across conditions, since the S0 override path is shared by all controllers in this study and does not depend on the PPE update mechanism. The asymmetry is informative: events that require contextual reprioritization through accumulated PAD dynamics—accidents and operator focus—degrade severely without temporal evolution, whereas events that bypass the strategy selection layer entirely show no comparable effect.

**Table 10.** PAD saturation rates: EI-UAV (PPE regulated) versus ablation (α = 1.0), identical clamping.

| Scenario | EI-UAV (%) | Ablation (%) | U | $p$ | r |
|---|---|---|---|---|---|
| Sc1 (Routine) | 0.0 ± 0.0 | 88.3 ± 0.2 | 0.0 | <0.001 | 1.00 |
| Sc2 (Accident) | 0.0 ± 0.0 | 73.7 ± 0.2 | 0.0 | <0.001 | 1.00 |
| Sc3 (Multi-Stressor) | 0.0 ± 0.0 | 56.6 ± 1.9 | 0.0 | <0.001 | 1.00 |

*Note:* Saturation defined as $|PAD\_d| \geq 0.999$ for any $d \in \{P, A, D\}$. $n$ = 12 EI-UAV runs, $n$ = 12 ablation runs per scenario. U = Mann–Whitney U; r = rank-biserial effect size. PAD_P peaked at +0.731 across all 36 EI-UAV runs. r = 1.00 reflects complete rank separation between conditions (zero overlap in saturation distributions).

The mechanistic explanation is direct. State-aware response in the PPE-regulated controller arises from the recursive PAD update accumulating contextual evidence step by step: a sequence of moderate stressors gradually shifts the mood region toward S2 or S6, and an explicit OCC override sets PAD directly to a target consistent with the event's emotional significance. The stateless variant has no such accumulator. Each step's strategy is determined solely by the instantaneous (TS, WL) reading, with no memory of prior conditions and no mechanism for graduated escalation. When the input space is saturated by a continuous severe traffic load, the controller cannot distinguish a routine high-stress moment from an injected emergency overlaid on the same baseline; both produce identical TS, identical rule firings, identical strategy. The graded inference machinery is intact; what is missing is the temporal context that makes graded reasoning behaviorally meaningful. Table 11 summarizes the principal EI-UAV versus FLC-UAV contrasts under identical experimental conditions.

**Table 11.** EI-UAV (PPE) versus FLC-UAV (Stateless Graded Baseline) Under Identical Conditions.

| Metric | EI-UAV (PPE) | FLC-UAV | Statistical Comparison |
|---|---|---|---|
| Strategy entropy, Sc1 (bits) | 1.60 ± 0.03 | 0.46 ± 0.14 | - |
| Strategy entropy, Sc2 (bits) | 1.43 ± 0.04 | 0.19 ± 0.02 | - |
| Strategy entropy, Sc3 (bits) | 1.71 ± 0.08 | 0.41 ± 0.05 | - |
| S6 (Emergency Response) occupancy, Sc3 (% of steps) | 56.5 ± 3.2 | 93.4 ± 0.9 | - |
| Accident response time (s) | 44.6 ± 20.5 | 205.8 ± 72.7 | U = 576, $p < 0.001$, r = 1.00 |
| Operator-focus response time (s) | 50.8 ± 31.3 | 201.8 ± 80.3 | $p < 0.001$, r = 1.00 |
| PAD saturation across Sc1–3 (% of steps) | 0.0 | not applicable | - |

*Note:* Both controllers were run under identical SUMO scenarios with the same twelve fixed seeds and event injection schedule per condition. Strategy entropies are step-occupancy values (per-step strategy distributions, consistent with the methodology used in Section 6.2.2), reported as mean ± SD across $n = 12$ runs per scenario. Accident comparison: $n = 24$ events per condition; operator-focus comparison: $n = 12$ events per condition. r denotes rank-biserial effect size; r = 1.00 indicates complete rank separation between distributions. PAD saturation is not applicable to the FLC-UAV, which has no internal PAD state by construction.

#### 6.2.3. Synthesis

The two ablations identify failure modes that are architecturally distinct rather than overlapping. Removing bidirectional regulation in Section 6.2.1 leaves the recursive update structure intact but eliminates the personality-anchored pull, producing unbounded internal state dynamics: PAD accumulates against the clamp boundary across 57–88% of simulation steps, and mood semantics degenerate because saturated coordinates carry no informational content. Removing stateful temporal evolution in Section 6.2.2 leaves graded representation intact but eliminates the recursive accumulation of context, producing collapsed behavioral diversity: strategy entropy drops to 0.19–0.46 bits and event-response time degrades 4–5× on state-aware operational events, while internal state remains nominally bounded only because there is no internal state to saturate. The two failure modes are non-substitutable. Bidirectional regulation prevents internal-state pathology; stateful temporal evolution produces behavioral diversity from that bounded internal state. The PPE architecture combines both because each component addresses a failure mode the other cannot.

### *6.3. Comparison with Alternative Approaches*

The validated properties permit a qualitative contrast with alternative approaches (Table 12). The rule-based and learning-based columns reflect general tendencies in the literature [4–7,9] rather than specific implementations; individual systems may deviate from these tendencies. The central trade-off in autonomous control (adaptability versus interpretability) assumes a distinct form in the proposed framework compared with the available alternatives. Rule-based systems are transparent but rigid: they respond to what their designers anticipated, in exactly the way prescribed. Learning-based systems adapt to what their training data contains, through mechanisms that resist human inspection. The PPE-regulated approach provides an intermediate solution. Behavior adapts continuously to environmental conditions through a mechanism whose operation is fully traceable at every step.

A reasonable question is whether the Pull–Push Engine could be implemented using fuzzy sets [41], since both approaches map continuous inputs to graded outputs. The empirical evidence from Section 6.2.2 addresses this question directly: the FLC-UAV variant—a Mamdani fuzzy controller of comparable interpretability—run under identical SUMO scenarios with the same twelve seeds, produced strategy entropy of 0.19–0.46 bits versus 1.43–1.71 bits for the PPE-regulated baseline, and event-response time degraded 4–5× on state-aware operational events with complete rank separation. The fuzzy inference engine

itself was mathematically responsive—membership functions activated correctly and rule firings shifted as input values changed—but the absence of stateful temporal evolution prevented the controller from accumulating the contextual history that PPE-regulated decisions depend on. This indicates that the two approaches are not interchangeable for the specific architectural objective of this work: graded representation alone is not sufficient to reproduce PPE behavior, because the recursive update of Equation 4 carries information across decision steps that a memoryless mapping cannot encode. The PPE's structural features—bidirectional regulation, bounded continuous evolution, implicit hysteresis described in Section 4.5, and direct linkage to established affective models (PAD [14], OCC [15], and Big Five [16])—collectively address a temporal regulation problem rather than a static mapping problem.

**Table 12.** Qualitative Comparison of Control Approaches.

| Property | PPE-Regulated | Stateless Graded Variant | Rule-Based | Learning-Based | Evidence |
|---|---|---|---|---|---|
| Regulation | Temporal filtering | Static graded mapping | Threshold-driven | Reward-optimized | - |
| State transitions | Smooth (hysteresis) | Smooth (graded) | Abrupt (binary) | Variable | Figures 4 and 5 |
| Diversity | 1.885–2.033 bits | 0.19–0.46 bits | Fixed rule set | Unconstrained | Table 4 |
| Interpretability | Full causal trace | Rule + MF inspectable | Transparent | Low (black-box) | Section 5.6 |
| Bounded dynamics | Guaranteed ($\pm 1$) | Output range bounded | Implicit | Not explicitly bounded | Section 4.4 |
| Context sensitivity | Personality-anchored | Static rule base | Predefined | Learned | Section 5.9 |

Beyond the architectural ablations reported in Section 6.2 (the bidirectional-pull ablation in Section 6.2.1 and the stateless graded ablation in Section 6.2.2), comprehensive multi-controller benchmarking against learning-based and additional rule-based controller families falls outside the scope of the present paper, which is intentionally limited to establishing and validating the PPE regulation mechanism itself. Such broader comparisons represent a complementary direction for future work.

### *6.4. Limitations*

#### 6.4.1. Scope of Validation

All experimental evidence in this paper was obtained in the SUMO microsimulation environment. The present study reports no real UAV flight tests, no hardware-in-the-loop validation, and no field-collected data. Physical flight dynamics, aerodynamic effects, sensor noise, GPS error, communication latency and packet loss, weather, and battery thermal/discharge nonlinearities are not modeled; zone-level congestion observations are treated as exact rather than subject to onboard sensor noise, partial occlusion, or detection uncertainty; and inter-UAV or UAV-to-ground communication is assumed reliable and synchronous. Any claims regarding practical deployability, certification readiness, or operational effectiveness in real complex urban environments are therefore conditional on the full validation pathway: the staged plan progresses from the present microsimulation, through the embedded feasibility testing already initiated (preliminary ATmega328P timing reported below), to hardware-in-the-loop evaluation, and ultimately to controlled field deployment, which collectively constitute necessary directions for future research.

### 6.4.2. Parameter and Architectural Scope

The responsiveness parameter $\alpha$ was held constant at 0.57 across all profiles to isolate the effect of personality-derived PAD baselines; profile-specific tuning of $\alpha$ is reserved for future investigation. $\beta_0$ was held fixed at 0.30 throughout (the rationale, including convergence behavior for values below 0.30, is detailed in Section 4.3); sensitivity analysis of $\beta_0$ is reserved for future work. The two architectural ablations reported in Section 6.2 ($\alpha$ = 1.0 in Section 6.2.1, Table 10; stateless graded control in Section 6.2.2, Figure 11) collectively serve as the internal baselines aligned with the present manuscript's scope of validating the bidirectional regulation mechanism and the role of stateful temporal evolution; broader cross-controller benchmarking against learning-based approaches and additional rule-based variants represents a complementary direction for future work. While multi-profile validation used $n$ = 12 runs per profile on Sc3 only, full cross-scenario statistical evaluation across personality profiles remains future work. Each scenario was executed for approximately 100 min of simulated time; extended mission durations may expose gradual drift or long-term stability effects not observable within the present evaluation window. Four mood centroids (Excited, Joyful, Neutral, Exhausted) were not reached across the 36 runs of the primary experiment, indicating the ten-mood scheme is over-specified for the operational PAD range associated with the Calm profile; this affects only the granularity of interpretability labels, not control correctness, and Section 5.9 shows higher-baseline profiles expand the reachable mood region.

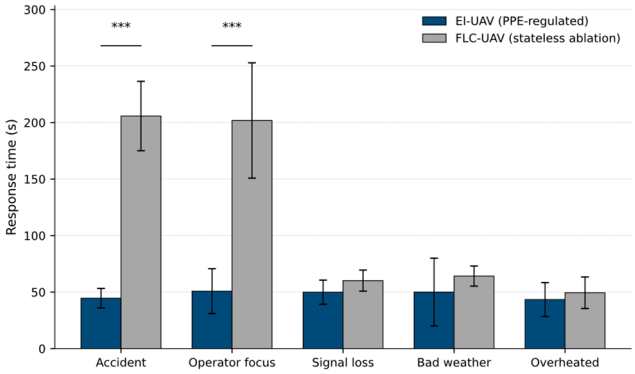


**Figure 11.** Response-time degradation across event types. Mean response time (seconds) for EI-UAV (PPE-regulated; blue) versus FLC-UAV (stateless ablation; grey) across five injected event categories. Error bars: 95% confidence interval. Significance markers: *** $p < 0.001$ (one-sided Mann–Whitney U). Sample sizes per condition: $n$ = 24 (Accident); $n$ = 12 (OperatorFocus, SignalLoss, Overheated); $n$ = 4 (EI-UAV)/$n$ = 11 (FLC-UAV) for BadWeather (asymmetric, scenario-dependent injection windows). Table 11 consolidates the principal EI-versus-FLC contrasts under identical experimental conditions, summarizing the entropy, occupancy, and response-time evidence reported above.

### 6.4.3. Embedded-Platform and Mapping Notes

The PPE update involves only scalar arithmetic per step with low memory requirements; preliminary testing on an ATmega328P microcontroller (16 MHz) indicated that the regulation loop executes in under 1 ms with a memory footprint below 600 bytes, suggesting suitability for resource-constrained embedded platforms, although full hardware-in-the-loop validation remains future work. The OCC PAD-target mappings (Appendix C) were defined through domain knowledge and expert judgment rather than learned from or validated against empirical human affective data; coordinate values are broadly consistent with established PAD region descriptions [14,27], but formal empirical calibration from annotated affective datasets remains a valuable direction for future work.

## 7. Conclusions

This work investigated emotion as a decision-control mechanism for autonomous UAV traffic monitoring, with the Pull–Push Engine (PPE) as the main contribution. The PPE addresses a practical challenge, namely the coupling of environmental stimuli to behavioral adaptation through an intermediate emotional state while avoiding the instability that such coupling typically entails. By weighing environmental push against personality-anchored pull through weighted temporal integration (Equation (4)), the engine filters transient fluctuations while preserving responsiveness to genuine operational changes. Across three scenarios of increasing complexity, PPE-regulated dynamics kept PAD trajectories bounded, produced mood distributions that tracked operational demands, and maintained zero saturation despite concurrent stressors, whereas removing the pull term resulted in 57–88% saturation. The dual-pathway architecture handled discrete events and gradual conditions through appropriately different temporal responses. Behavioral diversity scaled with operational demands without any mechanism explicitly detecting complexity; strategy selection emerged from emotional dynamics alone, with S6 becoming the leading strategy from Sc2 onward and entropy increasing monotonically as stressor load grew. Every decision remains causally traceable from environmental stimulus through each processing stage to strategy execution, given the factor weights documented in Appendix D, the OCC mappings in Appendix C, and the mood-strategy table in Appendix B, providing the structural interpretability that future safety-critical deployment will need to demonstrate, contingent on the hardware-in-the-loop and field validation outlined in Section 6.4.1. More broadly, the PPE offers a reusable regulatory pattern for any domain where adaptive behavior must coexist with stability and transparency. Future work will address physical UAV deployment under real-world sensor noise, multi-UAV coordination with emotional contagion mechanisms, and hybrid architectures combining emotion-driven regulation with learning-based optimization.

**Author Contributions:** Conceptualization, M.Z. and N.J.; Methodology, M.Z.; Software, M.Z.; Validation, M.Z., N.J. and A.-U.-H.Y.; Formal analysis, M.Z.; Investigation, M.Z.; Data curation, M.Z.; Writing—original draft preparation, M.Z.; Writing—review and editing, M.Z., N.J., M.A.B. and A.-U.-H.Y.; Visualization, M.Z.; Supervision, N.J. and A.-U.-H.Y.; Project administration, N.J. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research received no external funding.

**Data Availability Statement:** Simulation logs, configuration files, and analysis scripts are available from the corresponding author upon request. The SUMO simulation platform (v1.22.0) used in this study is open-source and freely available (https://eclipse.dev/sumo/, accessed on 6 May 2026). Personality profiles, OCC mappings, and PPE configuration parameters are fully specified in the Appendices A–D.

**Acknowledgments:** During the preparation of this manuscript, the authors used Claude Sonnet 4.5 (Anthropic, San Francisco, CA, USA) for language editing and text refinement. The authors have reviewed and edited all output and take full responsibility for the content of this publication.

**Conflicts of Interest:** The authors declare no conflicts of interest.

## Appendix A. Mood Classification Centroids

**Table A1.** PAD Centroids for Mood Classification.

| Mood State | Pleasure (P) | Arousal (A) | Dominance (D) |
|---|---|---|---|
| Anxious | −0.15 | +0.55 | +0.25 |
| Excited | +0.50 | +0.65 | +0.55 |

**Table A1.** *Cont.*

| Mood State | Pleasure (P) | Arousal (A) | Dominance (D) |
|---|---|---|---|
| Satisfied | +0.55 | +0.20 | +0.55 |
| Distressed | −0.45 | +0.55 | −0.15 |
| Joyful | +0.65 | +0.55 | +0.60 |
| Focused | +0.15 | +0.45 | +0.50 |
| Neutral | +0.00 | +0.00 | +0.00 |
| Relieved | +0.40 | −0.20 | +0.40 |
| Surprised | +0.15 | +0.70 | +0.10 |
| Exhausted | −0.30 | −0.45 | −0.35 |

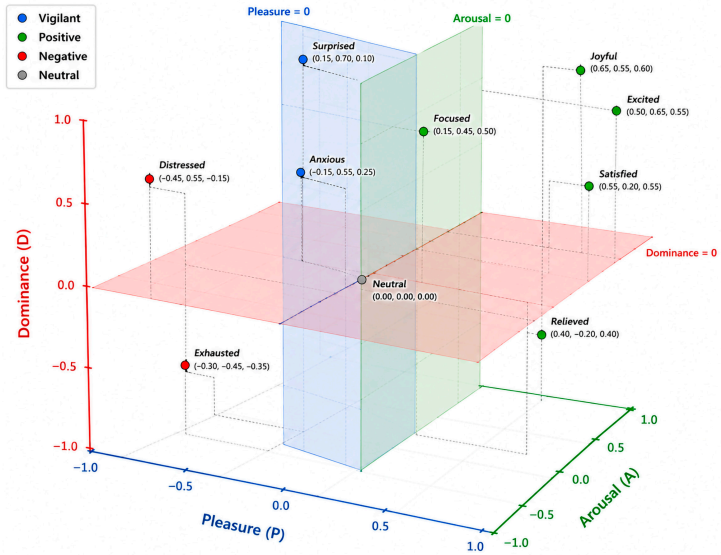


**Figure A1.** Three-dimensional visualization of mood classification centroids in PAD space. Centroid coordinates are in Table A1. Marker colors group the centroids by affective category (see in-figure legend): positive, vigilant, negative, and neutral.

## Appendix B. Mood-Strategy Mapping

**Table A2.** Strategy Selection by Mood State.

| Mood State | Primary Strategy | Secondary Strategy | Behavioral Emphasis |
|---|---|---|---|
| Anxious | S2 (Weighted Priority) | S1 (Sequential Patrol) | Heightened vigilance |
| Excited | S7 (Prioritize Unvisited) | S2 (Weighted Priority) | Active engagement |
| Satisfied | S2 (Weighted Priority) | S1 (Sequential Patrol) | Routine coverage |
| Distressed | S3 (Prioritize Highest) | S2 (Weighted Priority) | Crisis response |
| Joyful | S3 (Prioritize Highest) | S2 (Weighted Priority) | Exploration enabled |
| Focused | S6 (Operator Directed) | S2 (Weighted Priority) | Target concentration |
| Neutral | S1 (Sequential Patrol) | S2 (Weighted Priority) | Standard operations |
| Relieved | S2 (Weighted Priority) | S1 (Sequential Patrol) | Recovery mode |
| Surprised | S6 (Operator Directed) | S2 (Weighted Priority) | Situation assessment |
| Exhausted | S4 (High/Severe Only) | S2 (Weighted Priority) | Resource preservation |

## Appendix C. OCC Event-PAD Target Mapping

**Table A3.** Discrete OCC Events with PAD Targets (Discrete Pathway).

| Event Type | PAD Target (P, A, D) | Priority | Intended Affect |
|---|---|---|---|
| Accident | (−0.50, +0.70, +0.40) | 1 | Anxious |
| Signal Loss | (−0.30, +0.45, +0.10) | 1 | Distressed |
| Critical Battery | (−0.40, +0.80, −0.60) | 1 | Distressed |
| Overheated | (−0.70, +0.60, −0.50) | 1 | Distressed |
| Low Battery | (−0.60, −0.40, −0.60) | 2 | Exhausted |
| Operator Focus | (−0.05, +0.25, +0.55) | 2 | Focused |
| Bad Weather | (−0.70, +0.60, −0.50) | 2 | Distressed |
| Recharge Complete | (+0.40, −0.50, +0.30) | 3 | Relieved |
| Accident Cleared | (+0.60, −0.10, +0.70) | 8 | Satisfied |

**Note:** Overheated and Bad Weather share the same PAD target, reflecting that both represent external system-level threats that trigger an identical urgency response regardless of cause; the operational difference (thermal failure vs. environmental hazard) is expressed through their distinct priority levels (1 vs. 2) and subsequent strategy routing, not through different emotional states. Priority 8 is assigned to Accident Cleared to ensure this housekeeping event is always pre-empted by any active operational event; values 4–7 are reserved for future event types.

## Appendix D. Continuous Factor Weights

**Table A4.** Per-Step PAD Deltas for Continuous Factors (Continuous Pathway). Arrows indicate the direction of change in the named PAD dimension: ↑ = increase, ↓ = decrease.

| Condition | ΔP | ΔA | ΔD | Primary Effect |
|---|---|---|---|---|
| Congestion Severe | −0.054 | +0.081 | −0.027 | ↓P, ↑A (stress) |
| Congestion High | −0.036 | +0.054 | +0.018 | ↑A, ↑D (alert) |
| Congestion Moderate | +0.018 | −0.018 | +0.018 | Balanced |
| Congestion Low | +0.018 | −0.009 | 0.000 | ↑P (positive) |
| Battery Drain | −0.0000225 | −0.0000225 | −0.0000225 | Gradual fatigue |
| Optimal Traffic | +0.036 | −0.018 | −0.018 | ↑P, ↓A (calm) |

Values in the table represent per-step base deltas prior to physics scaling; congestion deltas are further multiplied by the vehicle count proportion (capped at vehicle count/50), and battery drain is multiplied by the physics drain rate (0.025). Notes: Base PAD deltas are applied per simulation step and scaled by the physics drain rate (0.025). Congestion deltas are further scaled by vehicle count proportion. Battery depletion applies a continuous nudge toward the Exhausted centroid (P = −0.30, A = −0.45, D = −0.35), reflecting gradual energy fatigue rather than acute stress; the acute arousal response to critically low battery is handled separately via the discrete OCC pathway (CriticalBattery event, OVERRIDE).

## References

1. Coifman, B.; McCord, M.; Mishalani, R.G.; Iswalt, M.; Ji, Y. Roadway traffic monitoring from an unmanned aerial vehicle. *IEE Proc. Intell. Transp. Syst.* **2006**, *153*, 11–20. [CrossRef]
2. Salvo, G.; Caruso, L.; Scordo, A. Urban traffic analysis through UAV. *Procedia Soc. Behav. Sci.* **2014**, *111*, 1083–1091. [CrossRef]
3. Butila, E.V.; Boboc, R.G. Urban traffic monitoring and analysis using unmanned aerial vehicles (UAVs): A systematic literature review. *Remote Sens.* **2022**, *14*, 620. [CrossRef]
4. Bisio, I.; Garibotto, C.; Haleem, H.; Lavagetto, F.; Sciarrone, A. A systematic review of drone based road traffic monitoring system. *IEEE Access* **2022**, *10*, 101537–101555. [CrossRef]
5. Menouar, H.; Guvenc, I.; Akkaya, K.; Uluagac, A.S.; Kadri, A.; Tuncer, A. UAV-enabled intelligent transportation systems for the smart city: Applications and challenges. *IEEE Commun. Mag.* **2017**, *55*, 22–28. [CrossRef]
6. Shakhatreh, H.; Sawalmeh, A.H.; Al-Fuqaha, A.; Dou, Z.; Almaita, E.; Khalil, I.; Othman, N.S.; Khreishah, A.; Guizani, M. Unmanned aerial vehicles (UAVs): A survey on civil applications and key research challenges. *IEEE Access* **2019**, *7*, 48572–48634. [CrossRef]
7. Moerland, T.M.; Broekens, J.; Jonker, C.M. Emotion in reinforcement learning agents and robots: A survey. *Mach. Learn.* **2018**, *107*, 443–480. [CrossRef]
8. Santoni de Sio, F.; van den Hoven, J. Meaningful human control over autonomous systems: A philosophical account. *Front. Robot. AI* **2018**, *5*, 15. [CrossRef]

9. Atakishiyev, S.; Salameh, M.; Yao, H.; Goebel, R. Explainable artificial intelligence for autonomous driving: A comprehensive overview and field guide for future research directions. *IEEE Access* **2024**, *12*, 101603–101625. [CrossRef]
10. Cao, H.; Wu, Z.; Yu, W. Drone for intelligent traffic monitoring: Current status and future trends. In *Computational and Experimental Simulations in Engineering, Proceedings of ICCES 2023*; Li, S., Ed.; Mechanisms and Machine Science; Springer: Cham, Switzerland, 2024; Volume 146, pp. 1133–1150. [CrossRef]
11. Picard, R.W. *Affective Computing*; MIT Press: Cambridge, MA, USA, 1997.
12. Breazeal, C. Emotion and sociable humanoid robots. *Int. J. Hum. Comput. Stud.* **2003**, *59*, 119–155. [CrossRef]
13. Arkin, R.C. *Behavior-Based Robotics*; MIT Press: Cambridge, MA, USA, 1998.
14. Mehrabian, A. Pleasure-arousal-dominance: A general framework for describing and measuring individual differences in temperament. *Curr. Psychol.* **1996**, *14*, 261–292. [CrossRef]
15. Ortony, A.; Clore, G.L.; Collins, A. *The Cognitive Structure of Emotions*; Cambridge University Press: Cambridge, UK, 1988.
16. John, O.P.; Srivastava, S. The Big Five trait taxonomy: History, measurement, and theoretical perspectives. In *Handbook of Personality: Theory and Research*, 2nd ed.; Pervin, L.A., John, O.P., Eds.; Guilford Press: New York, NY, USA, 1999; pp. 102–138.
17. Obaid, L.; Hamad, K.; Al-Ruzouq, R.; Abu Dabous, S.; Ismail, K.; Alotaibi, E. State-of-the-art review of unmanned aerial vehicles (UAVs) and artificial intelligence (AI) for traffic and safety analyses: Recent progress, applications, challenges, and opportunities. *Transp. Res. Interdiscip. Perspect.* **2025**, *33*, 101591. [CrossRef]
18. Ambeth Kumar, V.D.; Ramachandran, V.; Rashid, M.; Javed, A.R.; Islam, S.; Al Hejaili, A. An intelligent traffic monitoring system in congested regions with prioritization for emergency vehicle using UAV networks. *Tsinghua Sci. Technol.* **2025**, *30*, 1387–1400. [CrossRef]
19. Gebhard, P. ALMA: A layered model of affect. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), Utrecht, The Netherlands, 25–29 July 2005*; ACM: New York, NY, USA, 2005; pp. 29–36. [CrossRef]
20. Cañamero, L. When emotional machines are intelligent machines: Exploring the tangled knot of affective cognition with robots. In *Emotional Machines: Perspectives from Affective Computing and Emotional Human-Machine Interaction*; Misselhorn, C., Poljanšek, T., Störzinger, T., Klein, M., Eds.; Springer: Wiesbaden, Germany, 2023; pp. 135–158. [CrossRef]
21. Nunnari, F.; Tsovaltzi, D.; Lavit Nicora, M.; Beyrodt, S.; Prajod, P.; Chehayeb, L.; Brdar, I.; Delle Fave, A.; Negri, L.; André, E.; et al. Socially interactive industrial robots: A PAD model of flow for emotional co-regulation. *Front. Robot. AI* **2024**, *11*, 1418677. [CrossRef]
22. Broekens, J.; DeGroot, D.; Kosters, W. Formal models of appraisal: Theory, specification, and computational model. *Cogn. Syst. Res.* **2008**, *9*, 173–197. [CrossRef]
23. Marinier, R.P.; Laird, J.E.; Lewis, R.L. A computational unification of cognitive behavior and emotion. *Cogn. Syst. Res.* **2009**, *10*, 48–69. [CrossRef]
24. Lowe, R.; Herrera, C.; Morse, A.; Ziemke, T. The embodied dynamics of emotion, appraisal and attention. In *Attention in Cognitive Systems. Theories and Systems from an Interdisciplinary Viewpoint*; Paletta, L., Rome, E., Eds.; Lecture Notes in Computer Science; Springer: Berlin/Heidelberg, Germany, 2007; Volume 4840, pp. 1–20. [CrossRef]
25. Pessoa, L. On the relationship between emotion and cognition. *Nat. Rev. Neurosci.* **2008**, *9*, 148–158. [CrossRef]
26. Mehrabian, A. Analysis of the Big Five personality factors in terms of the PAD temperament model. *Aust. J. Psychol.* **1996**, *48*, 86–92. [CrossRef]
27. Zhao, Y.; Xie, D.; Zhou, R.; Wang, N.; Yang, B. Evaluating users' emotional experience in mobile libraries: An emotional model based on the Pleasure-Arousal-Dominance emotion model and the Five Factor Model. *Front. Psychol.* **2022**, *13*, 942198. [CrossRef] [PubMed]
28. Marsella, S.; Gratch, J. EMA: A process model of appraisal dynamics. *Cogn. Syst. Res.* **2009**, *10*, 70–90. [CrossRef]
29. Pei, G.; Li, H.; Lu, Y.; Wang, Y.; Hua, S.; Li, T. Affective computing: Recent advances, challenges, and future trends. *Intell. Comput.* **2024**, *3*, 0076. [CrossRef]
30. Seyitoğlu, F.; Ivanov, S. Robots and emotional intelligence: A thematic analysis. *Technol. Soc.* **2024**, *77*, 102512. [CrossRef]
31. Kolomaznik, M.; Petrik, V.; Slama, M.; Jurik, V. The role of socio-emotional attributes in enhancing human-AI collaboration. *Front. Psychol.* **2024**, *15*, 1369957. [CrossRef] [PubMed]
32. Kowalczuk, Z.; Czubenko, M.; Merta, T. Interpretation and modeling of emotions in the management of autonomous robots using a control paradigm based on a scheduling variable. *Eng. Appl. Artif. Intell.* **2020**, *91*, 103562. [CrossRef]
33. Belkaid, M.; Pessoa, L. Modeling emotion to enable intelligent behavior in robots. *Intellectica* **2023**, *79*, 109–128.
34. Spitale, M.; Axelsson, M.; Jeong, S.; Tuttosı, P.; Stamatis, C.A.; Laban, G.; Lim, A.; Gunes, H. Past, present, and future: A survey of the evolution of affective robotics for well-being. *arXiv* **2024**, arXiv:2407.02957. [CrossRef]
35. Hieida, C.; Nagai, T. Survey and perspective on social emotions in robotics. *Adv. Robot.* **2022**, *36*, 17–32. [CrossRef]
36. Gadanho, S. Reinforcement Learning in Autonomous Robots: An Empirical Investigation of the Role of Emotions. Ph.D. Thesis, University of Edinburgh, Edinburgh, UK, 1999.

37. Lopez, P.A.; Behrisch, M.; Bieker-Walz, L.; Erdmann, J.; Flötteröd, Y.P.; Hilbrich, R.; Lücken, L.; Rummel, J.; Wagner, P.; Wießner, E. Microscopic traffic simulation using SUMO. In *Proceedings of the IEEE Intelligent Transportation Systems Conference (ITSC), Maui, HI, USA, 4–7 November 2018*; IEEE: Piscataway, NJ, USA, 2018; pp. 2575–2582. [CrossRef]
38. Behrisch, M.; Bieker, L.; Erdmann, J.; Krajzewicz, D. SUMO-Simulation of Urban MObility: An overview. In *Proceedings of SIMUL 2011: The Third International Conference on Advances in System Simulation, Barcelona, Spain, 23–29 October 2011*; IARIA: Wilmington, DE, USA, 2011; pp. 63–68.
39. Wegener, A.; Piórkowski, M.; Raya, M.; Hellbrück, H.; Fischer, S.; Hubaux, J.-P. TraCI: An interface for coupling road traffic and network simulators. In *Proceedings of the 11th Communications and Networking Simulation Symposium, Ottawa, ON, Canada, 14–17 April 2008*; SCS: San Diego, CA, USA, 2008; pp. 155–163. [CrossRef]
40. Hernández-Marcos, A.; Ros, E. A generic self-learning emotional framework for machines. *Sci. Rep.* **2024**, *14*, 25858. [CrossRef]
41. Zadeh, L.A. Fuzzy sets. *Inf. Control* **1965**, *8*, 338–353. [CrossRef]