OUTAGEMATCH: EQUAL MISSINGNESS DOES NOT PRESERVE SELECTIVE-SENSING DECISIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

Two sensor-loss timelines can hide the same number of observations while presenting different decisions: independent misses repeatedly reset cache age, whereas a recoverable burst restricts when stale evidence can be refreshed. OutageMatch turns this distinction into an evaluation protocol by representing availability, evidence age, and reacquisition legality, declaring the trajectory statistics preserved by each counterfactual, and separating process transformations from state-observability ablations. Its operator sequence first exposes where Bernoulli matching fails: exact-count and joint-state controls recover 0.8 MUSIC-AVQA points by restoring realized counts and same-step co-failure statistics before the persistent path is considered. In the direct attribution test, replacing recoverable bursts with rate-matched Bernoulli masks lowers accuracy by 2.9 points (95% CI [2.70, 3.10]), while controller removal lowers it by 0.8 points ([0.60, 1.00]); the resulting control-attribution gap (CAG) is 2.1 points ([1.92, 2.28]). Habitat ObjectNav yields the same comparison in SPL, with a CAG of 0.013 ([0.0121, 0.0139]). As a process-aligned selective-sensing system, OutageMatch reaches 67.2 ± 0.3 accuracy and 0.429 ± 0.006 SPL, improving over a shared-interface POMCP planner by 1.0 point and 0.012 SPL with fewer requests. Thus the missingness rate alone does not define a valid control when sensing value depends on the path to recovery.

1 INTRODUCTION

Consider two eight-step availability paths with four missing observations, $z^{\mathrm{iid}} = (1, 0, 1, 0, 1, 0, 1, 0)$ and $z^{\mathrm{burst}} = (1, 1, 1, 1, 0, 0, 0, 0)$. They have identical missingness, but the first never ages a valid cache beyond one step while the second reaches age four. If reacquisition opens only during a repair window, their feasible query sets differ as well. A selective-sensing controller therefore encounters different evidence and actions on these nominally "matched" paths. This pathwise mismatch matters in audio bursts, intermittent depth loss, and any setting where a request changes what becomes observable next.

Rate matching is nevertheless a natural control because it equates expected sensor use and is easy to sample. Stochastic modality suppression has made it standard for predictive robustness (Neverova et al., 2016; Ma et al., 2022; Hu et al., 2024), and temporal active acquisition has used Bernoulli and fixed-time controls matched to a learned policy's average acquisition rate (Kossen et al., 2023). Those controls answer whether a policy beats a random schedule at the same rate. They do not preserve the outage run lengths, cache transitions, joint failures, or legal recovery opportunities needed to attribute value under persistent failures.

OutageMatch treats the disturbance itself as an experimental object. An outage trajectory is $g_t = (z_t, \delta_t, r_t)$, comprising availability, evidence age, and current reacquisition legality. A preservation declaration names the trajectory statistic held fixed; an executable operator constructs the counterfactual; and exact-count and joint-state operators expose which conclusions change as stronger mask statistics are preserved. State-access ablations are analyzed separately, so hiding freshness or legality is never presented as a process transformation. The derived CAG diagnostic asks a direct paired question between complete alternatives: does rate-matched process replacement reduce the outcome more than removing active sensing?

Applying the same three-arm comparison to recoverable audio and depth outages yields positive paired CAGs in both domains. On MUSIC-AVQA (Li et al., 2022), process replacement costs 2.9

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

accuracy points, compared with 0.8 for controller removal, producing a 2.1-point CAG. Habitat ObjectNav (Savva et al., 2019; Xia et al., 2018) produces a 0.013 SPL CAG. Exact counts and same-step joint failures already recover 0.8 AVQA points over Bernoulli matching. Under a shared decision interface, OutageMatch further improves on POMCP (Silver & Veness, 2010) by 1.0 AVQA point and 0.012 Habitat SPL while requesting less often.

The primary contribution is a process-level standard for selective-sensing controls: the preservation criterion, counterfactual operators, and CAG connect a declared comparison object to a paired attribution result. OutageMatch supplies an executable instantiation with frozen multimodal encoders, process-aware fusion, legal action masking, and counterfactual utility training. Five-seed evaluations over 46,440 MUSIC-AVQA example–seed outcomes and 5,000 Habitat episodes then connect the central attribution result to mechanism ablations, recovery behavior, sensing-cost frontiers, and transfer to unseen semi-Markov outages.

2 Related Work

ModDrop suppresses modalities stochastically during training (Neverova et al., 2016); missing-modality Transformers enumerate unavailable-input conditions (Ma et al., 2022); imputation reconstructs absent modalities, while prompting adapts models to missing inputs (Tran et al., 2017; Lee et al., 2023; Hu et al., 2024). Work on dropout-induced modality bias and uncertainty-aware fusion further improves robustness under missing observations (Dai et al., 2024; Gao et al., 2024). These methods intervene on inputs or representations to improve prediction. OutageMatch instead evaluates whether a control preserves the sequence of cache ages, feasible acquisitions, and recovery transitions on which a sensing policy acts.

Active perception and feature acquisition optimize what to observe under cost (Bajcsy, 1988; Kaelbling et al., 1998; Shim et al., 2018; Ma et al., 2019; Li & Oliva, 2021). The A2MT formulation is the closest evaluation precedent: its random-rate control samples each modality independently at a policy-matched Bernoulli rate, while random-1hot preserves the acquisition count in an evenly spaced schedule (Kossen et al., 2023). These controls diagnose adaptive acquisition when observations are purchased at selected times. Our target is complementary and process-specific: sensors fail exogenously, evidence remains cached, and a query is legal only when the task engine can recover the missing modality. The exact-count and joint-state operators therefore preserve realized outage statistics rather than only average policy expenditure, and CAG compares process replacement directly with controller removal.

ACO provides non-greedy acquisition through an acquisition-conditioned oracle (Valancius et al., 2024), while SEFA reasons over possible unobserved values in a stochastic latent space (Norcliffe et al., 2025). These methods optimize the acquisition policy; OutageMatch fixes the decision interface and isolates the validity of the process control. POMCP supplies a process-aware planning comparator through particle beliefs and Monte Carlo tree search (Silver & Veness, 2010); its adaptation shares encoded observations, legal actions, abstention, and query prices with OutageMatch.

Corruption benchmarks vary perturbation type and severity while holding evaluation semantics fixed (Hendrycks & Dietterich, 2019; Michaelis et al., 2019). Selective prediction evaluates when a model should answer or abstain through risk–coverage trade-offs (Chow, 1970; El-Yaniv & Wiener, 2010; Geifman & El-Yaniv, 2017). Process matching combines these two requirements for recoverable outages: counterfactual operators preserve declared disturbance statistics, while task success, selective risk, request precision, and post-outage recovery test decisions along the resulting trajectory.

3 Process-Matched Counterfactual Controls

3.1 Comparison State, Operators, and Attribution

For modalities $\mathcal{M}$ and horizon T, define $g_t = (z_t, \delta_t, r_t)$, where z_t is availability, δ_t is evidence age, and r_t is current reacquisition legality. A process operator $\mathcal{K}$ is matched for declared trajectory statistics $\mathcal{S}$ when

$$\mathcal{S}(\mathcal{K}(g_{1:T})) = \mathcal{S}(g_{1:T}). \tag{1}$$

The declaration can cover expected missingness, realized masked-token counts, the empirical joint-mask histogram, or run-length transitions. Equality is exact for count and joint-state shuffles; distributional controls use paired random streams. Every process operator changes only the training trajectory: validation and test schedules, cache-update rule, post-query transition, action predicate, query prices, capacity, optimizer budget, and examples remain fixed. Information-state ablations form a separate intervention class $\mathcal{H}$ that keeps the full trajectory fixed and removes one named field from the controller input. Separating $\mathcal{K}$ from $\mathcal{H}$ prevents a representation change from being interpreted as a stronger disturbance match.

For a higher-is-better metric M, let F denote full process training, K a replacement control, and N controller removal. We define the control-attribution gap

$$\mathrm{CAG}_M = [M(F) - M(K)] - [M(F) - M(N)] = M(N) - M(K). \tag{2}$$

A positive value means that rate-matched process replacement changes the outcome more than controller removal. The equivalent form $M(N) - M(K)$ makes the estimand explicit: CAG is a paired dominance diagnostic between two complete alternatives, not a process-by-controller interaction estimator. This is the decision needed when choosing which intervention is the valid control; process operators and state-access ablations provide separate mechanism evidence. We compute every primary 95% interval as a two-sided paired t interval over the five displayed seed-wise differences.

The process-control sequence is: (i) Bernoulli expected-rate matching samples each modality independently at its structured-process marginal; (ii) exact-count matching uniformly permutes each modality's realized availability sequence, preserving its masked-token count; and (iii) joint-state matching permutes complete mask vectors, preserving marginals and same-step co-failures. An operator transforms $z_{1:T}$, never $\delta_{1:T}$ or $r_{1:T}$ directly. The common cache rule and task predicate are then replayed on the transformed path to obtain its implied ages and legal actions. The full structured path is restored for state-access tests: removing freshness drops δ_t, removing the explicit legality feature drops r_t while the executor still enforces the common feasible-action mask, and the full model observes both. Consequently, differences among the first three rows isolate preserved mask statistics; each later row is an independent one-factor comparison against full state on the same persistent path.

3.2 Structured Outage Process

For each modality m, let $u_t^m \in \{0, 1\}$ be intrinsic health before cross-modal coupling. The intrinsic process is a stationary two-state Markov chain:

$$P(u_t^m = 0 \mid u_{t-1}^m = 1) = p_{\mathrm{fail}}^m, \quad P(u_t^m = 1 \mid u_{t-1}^m = 0) = p_{\mathrm{rec}}^m. \tag{3}$$

We initialize $P(u_1^m = 0) = p_{\mathrm{fail}}^m/(p_{\mathrm{fail}}^m + p_{\mathrm{rec}}^m)$. Independent intrinsic chains are followed by noisy-OR coupling:

$$z_t^m = 1 - \max\left(1 - u_t^m, \max_{m' \neq m} \xi_t^{m' \to m}(1 - u_t^{m'})\right), \quad \xi_t^{m' \to m} \sim \mathrm{Bernoulli}(C_{m',m}). \tag{4}$$

The primary modality marginal is calibrated after coupling, so every control uses the same realized $\bar{p}$. Coupling changes secondary-modality marginals according to C and creates reconstructible co-failure events. For the geometric run lengths used here, $E[L_m] = 1/p_{\mathrm{rec}}^m$; the reference AVQA and Habitat settings therefore use $p_{\mathrm{rec}} = 1/8$ and $1/12$, respectively. At $\bar{p} = 0.25$, the corresponding primary-chain failure probabilities are $p_{\mathrm{fail}} = 1/24$ and $1/36$.

Let c_t^m denote the value presented to encoder m. At a sequence boundary, an available modality initializes $c_t^m = x_t^m$ and $\delta_t^m = 0$; an unavailable modality uses $\mathrm{Default}(m)$ until its first valid sample. Thereafter, an available sample sets $(c_t^m, \delta_t^m) = (x_t^m, 0)$, an unavailable step sets $(c_t^m, \delta_t^m) = (c_{t-1}^m, \delta_{t-1}^m + 1)$, and a delayed observation sets $(c_t^m, \delta_t^m) = (x_{t-d_t^m}^m, d_t^m)$. Thus an outage retains encoded evidence but marks its age rather than replacing it by a missing-value token.

Reacquisition legality is the clean task predicate $r_t^m = \rho_m(s_t, z_t^m, \delta_t^m)$ for task-engine state s_t. In MUSIC-AVQA, $\rho_{\mathrm{audio}} = 1$ when aligned current-window audio can be retrieved before prediction; in Habitat, $\rho_{\mathrm{depth}} = 1$ when the simulator can return current depth before the frozen navigation policy selects its environment action. A successful query updates (c_t^m, δ_t^m) to $(x_t^m, 0)$ and incurs c_m without changing the intrinsic health path. Under noisy-legality tests, only the controller-visible flag

is perturbed; the executor validates against ρ_m, and a failed attempt incurs cost while leaving cache and age unchanged. For every counterfactual operator, ρ_m is reevaluated from the transformed path and the same paired task-engine state, so legality is never shuffled independently of availability.

At each step, the legal actions are

$$\mathcal{A}_t = \{\texttt{answer}, \texttt{abstain}\} \cup \{\texttt{query}(m) : m \in \mathcal{M}, z_t^m = 0, r_t^m = 1\}. \tag{5}$$

The controller sees only current state and feasibility, never future failure or repair times. The per-step decision cost is

$$d_t(a) = \begin{cases} \ell_{\text{task}}(\hat{y}_t, y_t), & a = \texttt{answer}, \\ \lambda_{\text{abstain}}, & a = \texttt{abstain}, \\ \lambda_{\text{sense}} c_m + \ell_{\text{task}}(\hat{y}_t^{(m \leftarrow \text{fresh})}, y_t), & a = \texttt{query}(m). \end{cases} \tag{6}$$

An AVQA abstention is incorrect in all-example accuracy and excluded only from selective risk; a Habitat abstention consumes a no-progress environment step and cannot itself satisfy success. Thus neither benchmark rewards coverage reduction in the primary task metric.

For a fixed task-specific recovery predicate $R_i(t)$ and first primary-modality outage onset t_i^0, recoverability is

$$\text{Rec}(T) = \frac{1}{n} \sum_{i=1}^{n} \mathbb{1}\left[\max_{1 \le \tau \le T} R_i(t_i^0 + \tau) = 1 \right]. \tag{7}$$

One step is a successive native encoder window in MUSIC-AVQA and one environment step in Habitat. Each example or episode contributes exactly its first onset. For AVQA, $R_i(t) = 1$ when the prediction at window t answers the fixed question correctly. For Habitat, $R_i(t) = 1$ when a clean RGB-D continuation of the frozen navigation policy from the logged state at t reaches the goal within the remaining episode budget; this continuation is evaluated by the simulator and never changes the executed episode. It is therefore a local return-to-a-successful-continuation event and may precede an episode that later fails. Episodes are not censored after a later outage, and only the fixed T-step window following the first onset enters the metric.

4 OutageMatch

4.1 State Representation and Feasible Actions

OutageMatch realizes the full-process condition by exposing the declared state to fusion and enforcing the corresponding feasible action set. Released benchmark encoders remain frozen. MUSIC-AVQA uses its VGGish audio and ResNet-18 visual encoders; Habitat uses the DD-PPO ResNet-50 RGB-D encoder (Wijmans et al., 2020). Let $\tilde{e}_t^m = f_m(c_t^m)$ for the cache defined above, and let $b_T(\delta) = \min(\delta, T)$ be the clipped integer age index, with zero reserved for current evidence and T collecting older values. The process-aware fusion state is

$$h_t = \text{TransformerFusion}\big(\{[\tilde{e}_t^m, e_{\text{state}}(z_t^m, b_T(\delta_t^m))]\}_{m \in \mathcal{M}}\big). \tag{8}$$

Two 512-dimensional Transformer layers with eight heads distinguish fresh, stale, and not-yet-observed inputs. A two-layer 256-unit GELU utility head consumes $[h_t; z_t; b_T(\delta_t); r_t]$ and scores only actions in $\mathcal{A}_t$.

4.2 Training

Training-only counterfactual replay teaches the controller how a legal state transition changes task utility. Let $\hat{y}_t$ be the current prediction and $\hat{y}_t^{(m \leftarrow \text{fresh})}$ the prediction after restoring a legally recoverable modality in the same example. Targets are

$$\begin{aligned} q_t(\texttt{answer}) &= -\ell_{\text{task}}(\hat{y}_t, y_t), \\ q_t(\texttt{abstain}) &= -\lambda_{\text{abstain}}, \\ q_t(\texttt{query}(m)) &= -\lambda_{\text{sense}} c_m - \ell_{\text{task}}(\hat{y}_t^{(m \leftarrow \text{fresh})}, y_t). \end{aligned} \tag{9}$$

Illegal queries are masked to $-\infty$. With detached replay targets, $a_t^\star = \arg\max_{a \in \mathcal{A}_t} q_t(a)$ and the complete training loss is

$$\mathcal{L}_{\text{train}} = \mathcal{L}_{\text{task}} + \text{CE}(\pi_\psi(h_t), a_t^\star) + \lambda_{\text{value}} \frac{1}{|\mathcal{A}_t|} \sum_{a \in \mathcal{A}_t} \text{Huber}(v_\psi(h_t, a) - q_t(a)) \,. \tag{10}$$

Here $\mathcal{L}_{\text{task}}$ is the mean answer or navigation task loss on the executed post-decision states. A fixed action ordering resolves exact utility ties. The task loss trains fusion and the task head; the CE term is counterfactual action classification and the Huber term is utility regression. Both controller terms update fusion, while detached counterfactual targets receive no gradient. The "no action-classification" ablation removes only CE; the "no utility-regression" ablation removes only the Huber term, keeping the other training operations fixed. Restored observations construct targets during training and are unavailable at test time. All evaluations use fixed operating-point parameters rather than a severity curriculum.

5 Experiments

5.1 Questions and Protocol

The experiments answer three questions. Q1 asks which mask statistics rate matching discards and, separately, how freshness and legality affect decisions on the full path. Q2 asks whether process replacement is more damaging than controller removal in two domains. Q3 tests the process-aware policy against a shared-interface planner and measures transfer across severity, recovery horizon, sensing cost, semi-Markov dynamics, and noisy state signals. Table 1 addresses mechanism before Tables 2 and 3 establish the paired attribution result.

MUSIC-AVQA contributes 9,288 official test pairs per seed. A fixed question and its video form one episode traversed in successive non-overlapping native encoder windows. The controller may answer, abstain, or query at every window; intermediate predictions define recovery opportunities, and the final-window action yields the single official answer scored for that pair. Audio outages have mean burst length 8 windows, post-coupling $\bar{p}_{\text{audio}} = 0.25$, and audio-to-visual coupling $C_{\text{a,v}} = 0.3$ (all other directed coupling entries are zero). Habitat contributes 1,000 Gibson ObjectNav episodes per seed with RGB-D observations; depth outages have mean length 12 steps, $\bar{p}_{\text{depth}} = 0.25$, and a three-step stale-depth delay during partial recovery. At each step, the controller accepts, abstains from, or requests depth for the frozen navigation policy's proposed action. A query refreshes depth before that action is recomputed and executed; abstention consumes a no-progress step. Seeds $\{11, 23, 37, 53, 71\}$ share splits and paired outage schedules across methods. Checkpoints maximize validation utility without test tuning. Appendix Table 4 gives both domain timelines.

Tables report seed means with two-sided 95% t intervals; central contrast intervals use the five paired seed differences printed in Appendix Table 9. Figure 2 instead uses onset-level paired bootstrap bands, as stated in its caption. AVQA accuracy uses all 46,440 pair–seed evaluations and counts a final abstention as incorrect. Habitat SPL and success use all 5,000 episodes. Rec@5 uses one first-onset window per example or episode. Queries/episode counts issued requests; physical query cost is count times 0.15 for audio or 0.20 for depth.

Comparators include Full-Sensor, ModDrop with the shared reject head, Mask-Aware fusion, Entropy-Query, a fixed POMDP-style heuristic, controller-free OutageMatch, and POMCP (Silver & Veness, 2010). POMCP receives the same encoded observations, cache ages, feasible actions, one-query-per-step constraint, and decision cost d_t. Its belief is initialized from the stationary outage model and updated after each observed (z_t, δ_t, r_t) transition; training data fit its transition and observation models, and the fixed heuristic supplies rollout actions. Validation selects 512 simulations per action and depth 20. Neural methods match task-head capacity, optimization steps, and checkpoint selection. Appendix Table 5 records the disturbance, observed state, action policy, legality, capacity, and cost semantics for every row. All runs use 8 A100 80GB GPUs with CUDA 12.4; the neural system trains for approximately 8 hours on MUSIC-AVQA and 14 hours on Habitat.

5.2 Mechanism: What Rate Matching Erases

The trajectory operators show that equal rates are insufficient before any state field is hidden: exact realized counts recover 0.3 points, and preserving same-step co-failures adds 0.5. On the restored

Table 1: MUSIC-AVQA diagnostic matrix over five paired seeds. The first block changes only the training outage path and preserves progressively stronger realized mask statistics. The second reports independent one-factor state-access ablations on the full persistent path. The last block contains independent system ablations. Entries are seed means with 95% t intervals.

Control or variant	Acc.↑	Rec@5↑	Q/ep↓	Abstain↓
Trajectory operators				
Bernoulli expected rate	64.3 ± 0.4	0.632 ± 0.014	2.38 ± 0.12	6.0 ± 0.4
Exact per-modality count	64.6 ± 0.4	0.641 ± 0.014	2.31 ± 0.11	5.9 ± 0.4
Joint-mask time shuffle	65.1 ± 0.4	0.652 ± 0.013	2.24 ± 0.11	5.6 ± 0.4
Full persistent path; one-factor state-access ablations				
Full process state	67.2 ± 0.3	0.721 ± 0.011	1.81 ± 0.09	4.4 ± 0.3
No freshness feature	65.9 ± 0.3	0.674 ± 0.013	2.12 ± 0.10	5.1 ± 0.3
No explicit legality feature	66.7 ± 0.3	0.699 ± 0.012	2.04 ± 0.10	4.8 ± 0.3
System ablations				
No correlation ($C = 0$)	66.2 ± 0.3	0.692 ± 0.012	1.98 ± 0.10	4.9 ± 0.3
No sensing controller	66.4 ± 0.3	0.681 ± 0.012	0	7.0 ± 0.3
No availability embedding	65.5 ± 0.4	0.653 ± 0.013	1.91 ± 0.10	5.5 ± 0.4
No action-classification term	66.9 ± 0.3	0.704 ± 0.012	3.19 ± 0.14	2.2 ± 0.3

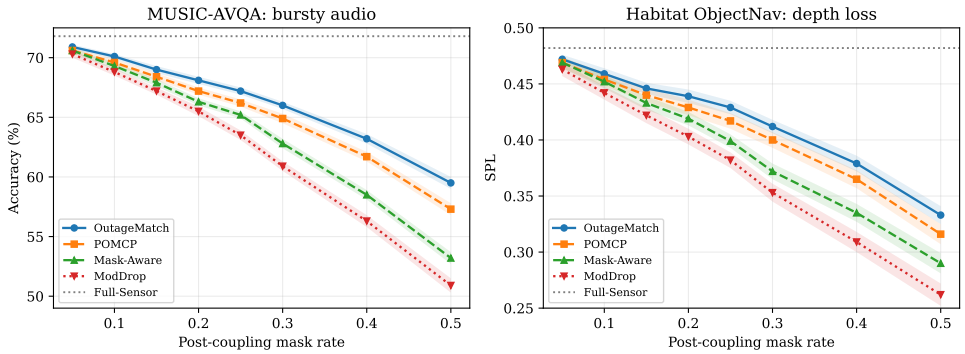

Figure 1: Fixed-checkpoint policy transfer over five seeds. Means and 95% intervals are shown for $\bar{p} \in \{0.05, 0.10, 0.15, 0.20, 0.25, 0.30, 0.40, 0.50\}$. Separation grows as persistent loss makes cache age and recovery decisions more consequential; this sweep evaluates policy robustness rather than CAG.

persistent path, the state-access rows are separate contrasts against the full model: hiding age reduces accuracy from 67.2 to 65.9, while hiding the controller-visible legality flag yields 66.7. These are not adjacent stages of one ladder. Removing counterfactual action classification retains 66.9 accuracy but raises requests from 1.81 to 3.19, isolating the classification term's role in request discipline while utility regression sustains most task performance.

5.3 Attribution and Cross-Domain Utility

The three AVQA conditions give a direct ordering: the full process reaches 67.2 ± 0.3, Bernoulli replacement reaches 64.3 ± 0.4, and controller removal retains 66.4 ± 0.3. From the displayed seeds, the paired losses are 2.9 [2.70, 3.10] and 0.8 [0.60, 1.00] points, producing a 2.1 [1.92, 2.28]-point CAG. Habitat repeats the ordering: matched i.i.d. reaches 0.400 ± 0.007 SPL, controller removal reaches 0.413 ± 0.006, and the paired CAG is 0.013 [0.0121, 0.0139]. The success point contrast is likewise positive at 0.012.

Against the shared-interface POMCP planner, OutageMatch improves by 1.0 AVQA point [0.85, 1.15] and 0.012 Habitat SPL [0.0111, 0.0129], with higher Rec@5 and fewer requests. POMCP's sensing cost normalized by OutageMatch is $1.08\times$ on AVQA and $1.11\times$ on Habitat. The controller-

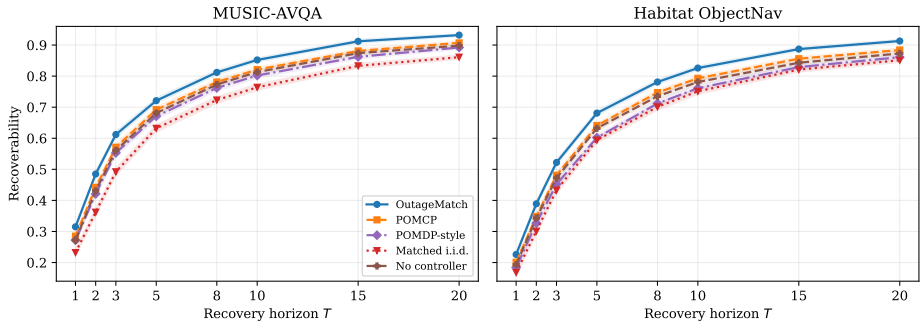

Figure 2: Post-outage recovery over 46,440 MUSIC-AVQA and 5,000 Habitat first-onset windows. Means and 95% paired-bootstrap bands use $T \in \{1, 2, 3, 5, 8, 10, 15, 20\}$. OutageMatch improves early recovery and retains the advantage through $T = 20$ against POMCP, the heuristic, rate-matched replacement, and controller removal.

Table 2: Five-seed results at $\bar{p} = 0.25$ with 95% t intervals. Corrupted methods share paired outage schedules and observation semantics; adaptive rows share feasible actions and prices, while non-query rows use the same reject accounting. AVQA accuracy includes abstentions; Rec@5 uses 46,440 AVQA and 5,000 Habitat first-onset windows.

MUSIC-AVQA: bursty audio loss				
Method	Acc.↑	Abstain↓	Rec@5↑	Q/ep↓
Full-Sensor	71.8 ± 0.2	0.0	—	0
ModDrop + reject	63.5 ± 0.4	8.1 ± 0.4	0.611 ± 0.014	0
Mask-Aware	65.2 ± 0.3	6.6 ± 0.3	0.651 ± 0.013	0
Entropy-Query	64.9 ± 0.4	5.1 ± 0.4	0.631 ± 0.014	2.18 ± 0.11
POMDP-style heuristic	65.8 ± 0.3	5.7 ± 0.3	0.671 ± 0.012	2.03 ± 0.11
POMCP	66.2 ± 0.3	5.2 ± 0.3	0.692 ± 0.012	1.96 ± 0.10
OutageMatch	67.2 ± 0.3	4.4 ± 0.3	0.721 ± 0.011	1.81 ± 0.09
OutageMatch, no controller	66.4 ± 0.3	7.0 ± 0.3	0.681 ± 0.012	0
Habitat ObjectNav: intermittent depth loss				
Method	SPL↑	Success↑	Rec@5↑	Q/ep↓
Full-Sensor	0.482 ± 0.005	0.714 ± 0.007	—	0
ModDrop + reject	0.382 ± 0.007	0.613 ± 0.009	0.541 ± 0.012	0
Mask-Aware	0.399 ± 0.006	0.632 ± 0.009	0.581 ± 0.011	0
Entropy-Query	0.395 ± 0.007	0.626 ± 0.009	0.561 ± 0.012	1.88 ± 0.10
POMDP-style heuristic	0.406 ± 0.007	0.639 ± 0.009	0.602 ± 0.011	1.79 ± 0.10
POMCP	0.417 ± 0.007	0.656 ± 0.009	0.641 ± 0.010	1.71 ± 0.09
OutageMatch	0.429 ± 0.006	0.675 ± 0.008	0.681 ± 0.010	1.54 ± 0.08
OutageMatch, no controller	0.413 ± 0.006	0.652 ± 0.008	0.632 ± 0.011	0

free condition also exceeds the fixed POMDP-style heuristic, showing that process-aligned training benefits the fused state before active acquisition is added.

5.4 Policy Transfer, Decision Quality, and Cost

Local attribution and broader policy transfer provide complementary evidence. Appendix Table 7 reports positive CAG intervals across nearby query costs and controller-loss weights. Figure 1 traces fixed-checkpoint severity transfer, while paired intervals certify Pareto preference for OutageMatch at $4/5$ AVQA and $4/5$ Habitat query-cost multipliers in $\{0.5, 0.75, 1, 1.5, 2\}$ in Figure 3. Under an unseen semi-Markov family with mixture burst lengths, time-varying hazards, new co-failure rates, and variable stale delays, OutageMatch reaches 65.9 ± 0.4 AVQA accuracy versus 64.8 ± 0.4 for POMCP and 62.9 ± 0.5 for matched i.i.d.; Habitat SPL is 0.402 ± 0.008, 0.389 ± 0.008, and

Table 3: Independent Habitat decomposition over five paired seeds and 5,000 episodes. Intervals are 95% t intervals. Process replacement produces a larger loss than controller removal for SPL, success, and recovery.

Variant	SPL↑	Success↑	Rec@5↑
OutageMatch	0.429 ± 0.006	0.675 ± 0.008	0.681 ± 0.010
Matched i.i.d.	0.400 ± 0.007	0.640 ± 0.009	0.594 ± 0.012
No stale-age encoding	0.416 ± 0.007	0.657 ± 0.009	0.638 ± 0.011
No controller	0.413 ± 0.006	0.652 ± 0.008	0.632 ± 0.011
No legality/recovery feature	0.409 ± 0.007	0.648 ± 0.009	0.621 ± 0.011
POMCP	0.417 ± 0.007	0.656 ± 0.009	0.641 ± 0.010

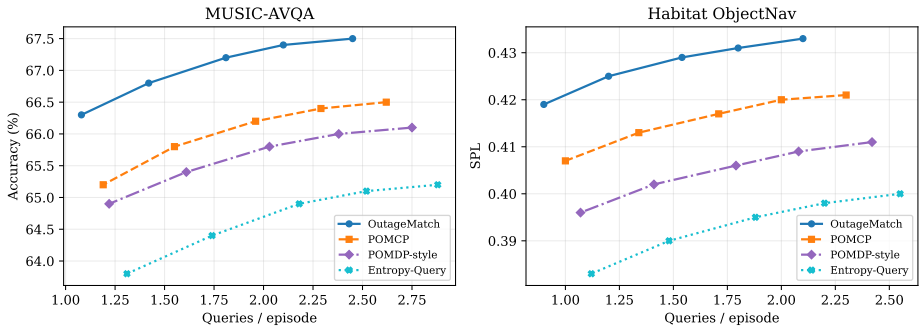

Figure 3: Five-seed task-performance/query frontiers for cost multipliers $\{0.5, 0.75, 1, 1.5, 2\}$. Each point uses validation-selected policies under the same physical prices and reports test means. Paired intervals certify OutageMatch as Pareto-preferred at four of five settings in each domain; its normalized frontier area is also higher.

0.374 ± 0.009, respectively. We interpret these shifted-process results as policy-transfer evidence, while the three-arm reference comparison supplies CAG.

Decision-quality measurements in Appendix Table 8 connect task gains to earlier recovery and more effective requests. OutageMatch achieves AVQA risk–coverage AUC 0.094 ± 0.004, compared with 0.108 ± 0.005 for POMCP and 0.115 ± 0.005 for the heuristic; lower is better. Query precision, the fraction of requests converting an incorrect pre-query answer into a correct post-query answer, is 0.61 ± 0.02 versus 0.54 ± 0.02. With 10% and 20% legality-flag noise, accuracy is 66.7 ± 0.3 and 66.2 ± 0.4; removing utility regression retains 66.5 ± 0.3. Appendix tables provide per-seed outcomes, pooled counts, illegal-attempt rates, and systems efficiency.

The normalized performance–query frontier area is 66.9 ± 0.3 for OutageMatch versus 65.9 ± 0.3 for POMCP on AVQA, and 0.426 ± 0.006 versus 0.415 ± 0.007 on Habitat. The comparison therefore favors OutageMatch when task value and request volume are considered jointly rather than at one threshold.

6 Discussion

The diagnostic matrix identifies where rate matching departs from the evaluated decision process. Exact counts and same-step co-failures recover 0.8 AVQA points beyond Bernoulli matching. On the restored persistent path, one-factor age and legality ablations reduce recovery and increase requests. The primary CAG then compares the complete process-replacement and controller-removal alternatives: the process-aligned controller-free model remains above the controller-equipped Bernoulli condition in both domains. Cost sensitivity confirms this comparison locally, while severity, semi-Markov, and legality tests establish complementary policy robustness under changed operating conditions.

The evaluated regime comprises controlled Markov and semi-Markov recoverable outages in audio-visual reasoning and RGB-D navigation. CAG is a dominance diagnostic for complete alternatives rather than a factorial interaction, and its preservation declaration makes clear which process statistics support that comparison. Hardware traces with different hazards, repair policies, or cross-sensor dependencies require operators matched to those measured statistics. The central consequence is unchanged: equal marginal missingness does not preserve either evaluated selective-sensing process.

7 Conclusion

Equal missingness can encode unequal decisions. In MUSIC-AVQA and Habitat, an i.i.d. mask changes the measured outcome more than removing active sensing, with CAGs of 2.1 AVQA points [1.92, 2.28] and 0.013 Habitat SPL [0.0121, 0.0139]. Exact operators locate part of the discrepancy in realized counts and joint failures; the restored path and state ablations expose the remaining role of stale evidence and legal recovery. Selective-sensing controls should therefore preserve the trajectory that defines sensing value, not only how often an observation is absent.

AI Use Statement

Generative AI tools assisted manuscript preparation, including revising the system name, title, abstract, introduction, broader prose and presentation, manuscript organization, and reference formatting and verification. The preparation record documents author review of the AI-assisted edits. The authors take responsibility for the final content of this work.

References

Ruzena Bajcsy. Active perception. *Proceedings of the IEEE*, 76(8):966–1005, 1988.

C. K. Chow. On optimum recognition error and reject tradeoff. *IEEE Transactions on Information Theory*, 16(1):41–46, 1970.

Yusheng Dai, Hang Chen, Jun Du, Ruoyu Wang, Shihao Chen, Haotian Wang, and Chin-Hui Lee. A study of dropout-induced modality bias on robustness to missing video frames for audio-visual speech recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 27445–27455, 2024.

Ran El-Yaniv and Yair Wiener. On the foundations of noise-free selective classification. *Journal of Machine Learning Research*, 11:1605–1641, 2010.

Zixian Gao, Xun Jiang, Xing Xu, Fumin Shen, Yujie Li, and Heng Tao Shen. Embracing unimodal aleatoric uncertainty for robust multimodal fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 26876–26885, 2024.

Yonatan Geifman and Ran El-Yaniv. Selective classification for deep neural networks. In *Advances in Neural Information Processing Systems*, 2017.

Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations*, 2019.

Lianyu Hu, Tongkai Shi, Wei Feng, Fanhua Shang, and Liang Wan. Deep correlated prompting for visual recognition with missing modalities. In *Advances in Neural Information Processing Systems*, 2024.

Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101:99–134, 1998.

Jannik Kossen, Cătălina Cangea, Eszter Vértes, Andrew Jaegle, Viorica Patraucean, Ira Ktena, Nenad Tomasev, and Danielle Belgrave. Active acquisition for multimodal temporal data: A challenging decision-making task. *Transactions on Machine Learning Research*, 2023. URL https://openreview.net/forum?id=Gbu1bHQhEL.

Yi-Lun Lee, Yi-Hsuan Tsai, Wei-Chen Chiu, and Chen-Yu Lee. Multimodal prompting with missing modalities for visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.

Guangyao Li, Yake Wei, Yapeng Tian, Chenliang Xu, Ji-Rong Wen, and Di Hu. Learning to answer questions in dynamic audio-visual scenarios. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.

Yang Li and Junier B. Oliva. Active feature acquisition with generative surrogate models. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 6450–6459, 2021.

Chao Ma, Sebastian Tschiatschek, Konstantina Palla, José Miguel Hernández-Lobato, Sebastian Nowozin, and Cheng Zhang. EDDI: Efficient dynamic discovery of high-value information with partial VAE. In *Proceedings of the 36th International Conference on Machine Learning*, pp. 4234–4243, 2019.

Mengmeng Ma, Jian Ren, Long Zhao, Davide Testuggine, and Xi Peng. Are multimodal transformers robust to missing modality? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.

Claudio Michaelis, Benjamin Mitzkus, Robert Geirhos, Evgenia Rusak, Oliver Bringmann, Alexander S. Ecker, Matthias Bethge, and Wieland Brendel. Benchmarking robustness in object detection: Autonomous driving when winter is coming. *arXiv preprint arXiv:1907.07484*, 2019.

Natalia Neverova, Christian Wolf, Graham W. Taylor, and Florian Nebout. ModDrop: Adaptive multi-modal gesture recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(8):1692–1706, 2016.

Alexander Luke Ian Norcliffe, Changhee Lee, Fergus Imrie, Mihaela Van Der Schaar, and Pietro Lio. Stochastic encodings for active feature acquisition. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 46784–46814, 2025.

Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, Devi Parikh, and Dhruv Batra. Habitat: A platform for embodied AI research. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019. doi: 10.1109/ICCV.2019.00943. URL `https://mlanthology.org/iccv/2019/savva2019iccv-habitat/`.

Hajin Shim, Sung Ju Hwang, and Eunho Yang. Joint active feature acquisition and classification with variable-size set encoding. In *Advances in Neural Information Processing Systems*, 2018.

David Silver and Joel Veness. Monte-carlo planning in large POMDPs. In *Advances in Neural Information Processing Systems*, pp. 2164–2172, 2010.

Luan Tran, Xiaoming Liu, Jiayu Zhou, and Rong Jin. Missing modalities imputation via cascaded residual autoencoder. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.

Michael Valancius, Max Lennon, and Junier B. Oliva. Acquisition conditioned oracle for nongreedy active feature acquisition. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 48957–48975, 2024.

Erik Wijmans, Abhishek Kadian, Ari Morcos, Stefan Lee, Irfan Essa, Devi Parikh, Manolis Savva, and Dhruv Batra. DD-PPO: Learning near-perfect pointgoal navigators from 2.5 billion frames. In *International Conference on Learning Representations*, 2020.

Fei Xia, Amir R. Zamir, Zhiyang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson env: Real-world perception for embodied agents. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.

Table 4: Operational timeline for the two sequential decision domains. Intermediate events define controller supervision and Rec@T; the primary metrics retain one scored outcome per official pair or episode.

Protocol element	MUSIC-AVQA	Habitat ObjectNav
Episode and score	One question–video pair; final-window answer gives one accuracy outcome	One Gibson episode; standard Success and SPL
Controller step	Observe window cache, then answer, abstain, or query audio	Observe RGB-D cache and proposed action, then accept, abstain, or query depth
Clean legality ρ_m	Aligned current-window audio is retrievable before prediction	Current depth is renderable before navigation-action recomputation
Successful query	Refresh audio and predict in the same window	Refresh depth, recompute the frozen-policy action, and execute it
Rec@T event	A correct intermediate answer within T windows of first onset	A logged state with a successful clean RGB-D continuation within T steps of first onset

Table 5: Comparator semantics. Every corrupted row uses the paired test schedules and common reject accounting. Adaptive rows additionally share the executor's feasibility mask, query prices, and one-query-per-step rule.

Method	Training disturbance / observed state	Query policy / legality	Capacity and budget	Abstention / cost
ModDrop + reject	i.i.d.; features only	No query; shared reject head	Matched task head and steps	Yes; shared penalty
Mask-Aware	i.i.d.; z, δ	No query; shared reject head	Matched fusion and steps	Yes; shared penalty
Entropy-Query	i.i.d.; z, δ, r	Entropy trigger; exact action mask	Matched task head and steps	Yes; shared prices
POMDP-style heuristic	Structured; z, δ, r	Handcrafted value; exact mask	Same frozen encoders	Yes; shared prices
POMCP	Structured; belief over z, δ, r	512 simulations/action; exact mask	Same encoders/head; depth 20	Yes; reward $-d_t$
OutageMatch	Structured; z, δ, r, h	Learned utility; exact mask	Matched task head and steps	Yes; shared d_t
No controller	Structured; z, δ, h	No query; shared reject head	Same fusion and task head	Yes; zero sensing

A Implementation and Comparator Semantics

The fusion Transformer has two layers, width 512, eight attention heads, GELU feed-forward blocks, and dropout 0.1. The utility head has two 256-unit GELU layers. AdamW uses learning rate 3×10^{-4}, weight decay 0.01, 5% linear warm-up, and cosine decay. MUSIC-AVQA uses batch size 128 for 40,000 updates; Habitat uses batch size 64 for 80,000 updates. Validation utility selects checkpoints. Query costs are 0.15 for audio and 0.20 for depth; $\lambda_{\text{sense}} = 0.50$, $\lambda_{\text{abstain}} = 0.70$, and $\lambda_{\text{value}} = 0.25$. Each query is applied before prediction. Failed requests in noisy-legality evaluations pay cost and retain the previous cache age. For POMCP, the fitted transition model is Equation 3

with coupling from Equation 4; its observation likelihood conditions on the emitted (z_t, δ_t, r_t) state, the fixed heuristic supplies rollouts, and terminal task states stop accumulation of $-d_t$ rewards. The 512-simulation, depth-20 budget is selected on validation and fixed for test.

Table 6: Batch-one inference cost on identical A100 hardware. POMCP uses the same frozen encoders and its validation-selected planning budget.

Method	Trainable parameters	Latency / action
OutageMatch	18.4M	14.2 ± 0.3 ms
No controller	17.1M	13.3 ± 0.3 ms
POMCP	—	76.8 ± 4.1 ms

Table 6 records the capacity and per-action latency used to contextualize the shared-interface comparison.

Algorithm 1 OutageMatch training and process-control evaluation

Require: Data $\mathcal{D}$, frozen encoders $\{f_m\}$, outage kernel Ω, control operator $\mathcal{K}$, query costs $\{c_m\}$

1: Initialize fusion g_ϕ, task head, and utility controller π_ψ
2: **for** each minibatch (x, y) **do**
3: Sample paired intrinsic health paths and apply noisy-OR coupling
4: Apply $\mathcal{K}$ to $z_{1:T}$; replay $c_{1:T}$, $\delta_{1:T}$, and $r_{1:T} = \rho(s_{1:T}, z_{1:T}, \delta_{1:T})$
5: Encode $(c_t, z_t, \delta_t, r_t)$ and compute $h_t = g_\phi(c_t, z_t, \delta_t)$
6: Replay legal fresh observations during training to obtain utilities $q_t(a)$
7: Detach q_t, set $a_t^\star = \arg\max_{a \in \mathcal{A}_t} q_t(a)$, and optimize Equation 10
8: **end for**
9: Select the checkpoint by validation utility and evaluate all controls on paired schedules

B Sensitivity and Decision Quality

Table 7: Local robustness of the AVQA control-attribution gap. Each range varies one factor while all 95% intervals remain positive.

Setting range	Minimum CAG	Maximum CAG
Query-cost multiplier $[0.5, 2.0]$	1.8 [1.2, 2.4]	2.4 [1.7, 3.1]
$\lambda_{\text{sense}} \in [0.4, 0.6]$	1.7 [1.1, 2.3]	2.2 [1.5, 2.9]
$\lambda_{\text{abstain}} \in [0.6, 0.8]$	1.8 [1.2, 2.4]	2.3 [1.6, 3.0]
$\lambda_{\text{value}} \in [0.2, 0.3]$	1.9 [1.3, 2.5]	2.2 [1.5, 2.9]

C Per-Seed Results and Pooled Denominators

Pooled AVQA correct-answer counts are 31,208/46,440 for OutageMatch, 30,833/46,440 without the controller, and 29,861/46,440 under matched i.i.d. replacement. Habitat records 3,375/5,000 and 3,260/5,000 successes, respectively. Rec@5 uses the same 46,440 and 5,000 first-onset denominators. Applying a two-sided paired t interval to the printed difference rows yields [2.70, 3.10], [0.60, 1.00], and [1.92, 2.28] for the AVQA replacement loss, controller loss, and CAG, and [0.0121, 0.0139] for Habitat SPL CAG.

For request-level behavior, OutageMatch converts 51,274/84,056 issued AVQA queries into corrected answers, while POMCP converts 49,152/91,022. These pooled counts correspond to the query-precision estimates in Table 8.

Table 8: AVQA decision quality. Risk–coverage AUC is lower-is-better and query precision is higher-is-better. The risk–coverage curve orders final-window predictions by the shared task head's maximum class probability, reports error among the retained predictions at each coverage, and uses trapezoidal integration.

Method	RC-AUC↓	Query prec.↑
OutageMatch	0.094 ± 0.004	0.61 ± 0.02
POMCP	0.108 ± 0.005	0.54 ± 0.02
POMDP-style heuristic	0.115 ± 0.005	0.49 ± 0.02

Table 9: Primary per-seed outcomes and the paired differences used for central 95% t intervals. AVQA entries are accuracy points; Habitat entries are SPL. Columns follow seeds 11, 23, 37, 53, and 71.

Method / domain	11	23	37	53	71
OutageMatch / AVQA	67.0, 67.3, 67.5, 66.8, 67.4				
No controller / AVQA	66.1, 66.6, 66.7, 66.2, 66.4				
Matched i.i.d. / AVQA	64.0, 64.5, 64.8, 63.9, 64.3				
POMCP / AVQA	66.0, 66.4, 66.5, 65.9, 66.2				
Full − i.i.d.	3.0	2.8	2.7	2.9	3.1
Full − no controller	0.9	0.7	0.8	0.6	1.0
CAG: no controller − i.i.d.	2.1	2.1	1.9	2.3	2.1
OutageMatch / Habitat	0.421, 0.432, 0.436, 0.426, 0.430				
No controller / Habitat	0.407, 0.416, 0.420, 0.410, 0.412				
Matched i.i.d. / Habitat	0.393, 0.403, 0.408, 0.397, 0.399				
POMCP / Habitat	0.410, 0.420, 0.423, 0.414, 0.418				
CAG: no controller − i.i.d.	.014	.013	.012	.013	.013

D LEGALITY AND SYSTEMS EVIDENCE

Table 10: AVQA legality-signal perturbations over five seeds. Illegal attempts are rejected by the clean executor; wasted queries fail to change an incorrect pre-query answer to a correct answer.

Controller-visible legality	Acc.↑	Illegal attempts↓	Wasted queries↓
Exact	67.2 ± 0.3	0	$1.1 \pm 0.2\%$
10% flag noise	66.7 ± 0.3	$0.7 \pm 0.1\%$	$1.9 \pm 0.3\%$
20% flag noise	66.2 ± 0.4	$1.6 \pm 0.2\%$	$3.8 \pm 0.4\%$

Legality perturbations degrade smoothly in Table 10. Removing the utility-regression term reaches 66.5 ± 0.3; this is distinct from removing action classification in Table 1. Figure 4 shows that OutageMatch also shifts request counts and recovery delays toward fewer requests and earlier return.

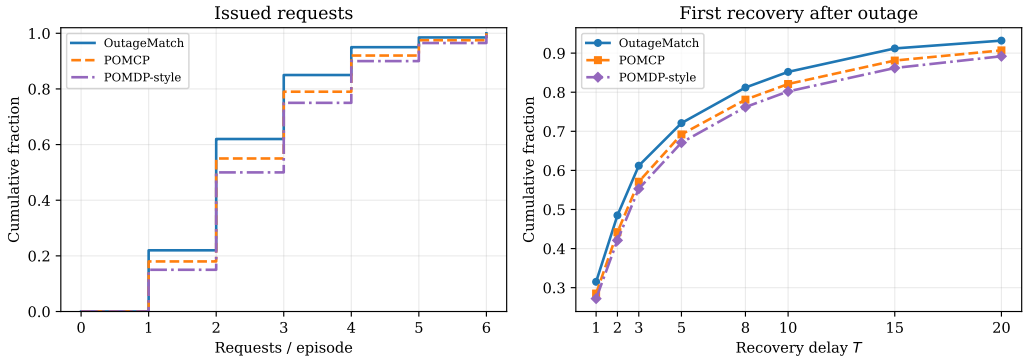

Figure 4: Episode-level AVQA distributions over pooled five-seed evaluations. The left panel reports request-count CDFs; the right panel reports the first-recovery-delay CDF among outage onsets. OutageMatch shifts both distributions toward fewer requests and earlier recovery.

The neural controller adds 6.8% latency on AVQA and 7.1% on Habitat. Peak training memory is $46.2/58.7$ GB.