FROM BAYESIAN BELIEF TO GENERATIVE BEHAVIOR: UNCERTAINTY-PRESERVING QUANTUM BAYESIAN FEEDBACK UNDER EVOLVING EVIDENCE

Anonymous authors
Paper under double-blind review

ABSTRACT

Generative systems under incomplete or changing evidence must decide both which conditions to express and how frequently to express them. The challenge is to translate uncertainty over alternative conditions into observable generative behavior. We introduce Quantum Bayesian Feedback (QBF), which maintains a persistent joint belief over task hypotheses and observation noise, recursively updates it with new evidence, and maps the resulting belief to a distribution over controls for a frozen generator. We show that QBF belief updates are equivalent to Bayesian filtering, derive bounds characterizing uncertainty preservation through readout and generation, and quantify information lost under point-valued readouts. Experiments with pretrained diffusion models demonstrate that posterior differences not captured by posterior-mean or maximum a posteriori (MAP) readouts remain reflected in generated regional attributes under posterior sampling. Target-switching experiments further depict how belief revision translates into adaptive changes in generative response across candidate sets. Extending QBF to real-image evidence and native 512×512 semantic generation further reveals trade-offs among accuracy, stability, and recovery under different belief-update dynamics. Measured by the mean energy score, QBF achieves 31.6–34.3% lower regional generation loss than point-valued readouts and semantic scores close to those of adaptive feedback baselines, while preserving posterior distinctions and adapting to target changes. Ideal quantum simulations show that the belief dynamics are induced directly by QBF state evolution and readout. Overall, QBF translates evolving posterior uncertainty into adaptive generative behavior under changing evidence, from regional attribute control to semantic high-resolution generation.

1 INTRODUCTION

Generative systems increasingly rely on evidence that is incomplete and progressively augmented as new observations become available, as in long-context language interaction, interactive image generation, and generative inference from sequential observations (Xiao et al., 2024; Lee et al., 2025; Huang et al., 2026). Retaining the full evidence history can avoid premature information loss, but the amount of information that must be carried forward and processed grows as evidence accumulates (Xiao et al., 2024; Huang et al., 2026). Conversely, compressing the accumulated evidence into a compact summary can reduce this burden, but may also obscure differences in how strongly alternative task conditions are supported, thereby discarding uncertainty that remains relevant to generation (Hoeting et al., 1999). Accordingly, the problem is to summarize evolving evidence without erasing task-relevant uncertainty. The representation must remain compact and updatable while preserving meaningful distinctions among plausible task conditions. Those distinctions should remain available to shape subsequent generation.

A key distinction is between uncertainty over task hypotheses and stochastic variation under a fixed generative condition (Kendall & Gal, 2017; Hoeting et al., 1999). Such randomness can diversify individual samples without preserving how belief is distributed across alternative hypotheses. The resulting question is how an evolving belief can shape generative behavior such that both the supported hypotheses and their changing relative support remain observable in the output distribution.

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.

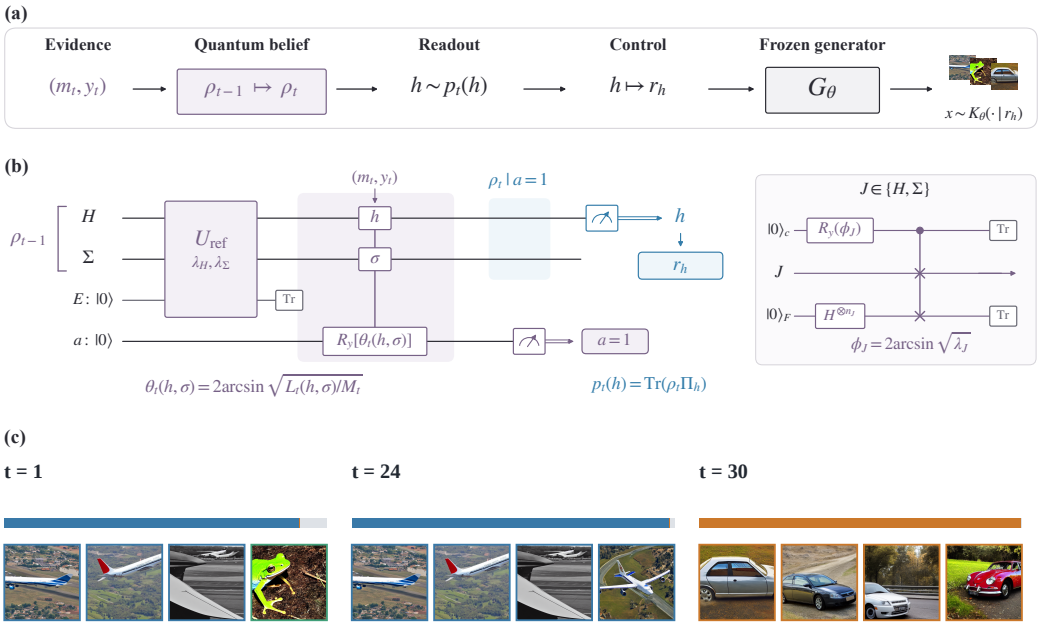

Figure 1: **Quantum Bayesian Feedback.** (a) Evidence updates a joint belief that controls a frozen generator. (b) Ideal quantum update and readout conditioned on $a = 1$; inset: local refresh. Wires denote registers. (c) Generation responds to a target change at $t = 25$ under matched diffusion noise. Bars show posterior mass: airplane (blue), car (orange), others (gray). Generation uses equivalent classical readout probabilities.

This question involves two distinct mappings: from belief to control, and from control to observable outputs.

Existing approaches provide important components of this process. Conditional and guided generative methods can shape the sampling process, while sequential inference methods can update the belief state as new observations arrive (Wang et al., 2023; Chung et al., 2023; Wei et al., 2026; Huang et al., 2026). Integrating these capabilities requires the updated belief to guide subsequent generation. As new evidence changes the relative support among task hypotheses, generation should adapt accordingly while retaining the uncertainty that remains unresolved. Existing methods do not jointly characterize how such belief revisions should be reflected in the resulting output distribution.

We introduce Quantum Bayesian Feedback (QBF), a framework in which an evolving belief state governs subsequent generation. As shown in Figure 1, QBF maintains a joint belief over task hypotheses and observation noise. The joint belief is recursively updated as new evidence arrives, and its task marginal then determines the probability assigned to each control condition of a frozen generator. Here, feedback refers to the recurrent influence of new evidence on generative control through the evolving belief state, rather than feedback from generated outputs themselves. In QBF, the joint belief is encoded in a quantum state, with prediction and likelihood conditioning realizing the belief update. The full joint belief is retained across successive updates, allowing uncertainty about observation noise to continue influencing subsequent belief revision rather than being discarded after each observation.

Our results show that QBF can carry task-relevant posterior uncertainty into generative control. As evidence changes, the generated distribution responds to the corresponding revision of belief. Theoretical analysis identifies when this uncertainty is retained through readout and generation. Under the mean energy score R, QBF achieves 0.019697 across 96 regional-control histories, compared with 0.028790 for MAP readout and 0.029967 for mean readout; its semantic score of 0.073242 is close to those of window, discount, and learned-hazard controllers. Ideal quantum simulations further show that belief updating and control readout can be realized through quantum-state evolu-

tion and measurement. These results illustrate how quantum representations can serve as structured interfaces between evolving uncertainty and classical generative models.

2 Related Work

Conditional sampling and distributional steering. DDNM enforces linear measurement constraints (Wang et al., 2023), DPS approximates posterior sampling for noisy inverse problems (Chung et al., 2023), and TFG unifies training-free guidance objectives (Ye et al., 2024). Feynman–Kac steering combines intermediate potentials with particle resampling (Singhal et al., 2025); latent diffusion enables high-resolution text conditioning (Rombach et al., 2022). QBF maintains an evolving distribution over such conditions and studies its transmission to output attributes.

Sequential inference and belief-based control. Belief-space planning (Kaelbling et al., 1998) and control as probabilistic inference (Levine, 2018) connect uncertainty to sequential decisions. SURGE combines diffusion surrogates with particle filtering (Wei et al., 2026), and sequential flow matching learns belief transport under incoming observations (Huang et al., 2026). Bayesian Flow Networks integrate Bayesian updates into generation (Graves et al., 2023), while PriorGuide adapts amortized diffusion inference to changing test-time priors (Yang et al., 2026). QBF selects a distribution over generative conditions and characterizes how the resulting outputs preserve the evolving belief. Our comparisons adapt constant-unknown hazard learning to this finite-state setting (Wilson et al., 2010).

Model averaging and interpretable control. Posterior-sampled conditioning implements Bayesian model averaging (Hoeting et al., 1999). Concept bottlenecks expose variables for interpretation and intervention (Koh et al., 2020), including in generative models (Kulkarni et al., 2025). QBF connects belief probabilities to measurable output frequencies and characterizes when mean or MAP controls lose distinctions that remain observable through posterior sampling.

Quantum inference and generation. Our construction builds on quantum Bayesian updating and rejection sampling (Soklakov & Schack, 2006; Low et al., 2014). Amplitude amplification uses coherent preparation and inverse access (Brassard et al., 2002); general black-box updating faces computational constraints (Wiebe & Granade, 2017). Whereas quantum generative adversarial models place quantum components in generation or discrimination (Lloyd & Weedbrook, 2018), QBF uses a quantum belief representation to control a classical generator.

3 Quantum Bayesian Feedback

3.1 Sequential distributional control

We formulate generative control under evolving evidence as the selection of a distribution over conditions. Let $\mathcal{D}_t = \{(m_\tau, y_\tau)\}_{\tau=1}^{t}$ denote measurement indices and observations. The latent state $S_t = (H_t, \Sigma_t)$ comprises a task hypothesis and an observation-noise variable. A transition T and likelihood $L_t(s) = p(y_t \mid m_t, s)$ define the recursive belief

$$\begin{aligned} \bar{b}_t(s) &= \sum_{s'} T(s \mid s') b_{t-1}(s'), \\ b_t(s) &= \frac{L_t(s)\bar{b}_t(s)}{\sum_u L_t(u)\bar{b}_t(u)}. \end{aligned} \tag{1}$$

The filter uses observations without access to latent target labels or change times. Its task marginal is $p_t(h) = \sum_\sigma b_t(h, \sigma)$.

Assign each hypothesis attributes $r_h \in \mathcal{A}$, measurable in generated outputs by $f : \mathcal{X} \to \mathcal{A}$. The belief induces the target attribute law

$$\nu_t = \sum_h p_t(h)\delta_{r_h}. \tag{2}$$

A control law maps the current belief to condition probabilities $\tilde{p}_t(h)$. With a frozen generator whose conditional output distribution is $K_\theta(\cdot \mid r)$, these probabilities induce

$$Q_t = \sum_h \tilde{p}_t(h) K_\theta(\cdot \mid r_h). \tag{3}$$

The control objective is to match the output-attribute law $f_\# Q_t$ to ν_t. For a probability metric D, we call the interface δ_t-uncertainty-preserving if

$$\mathsf{D}(f_\# Q_t, \nu_t) \le \delta_t. \tag{4}$$

QBF uses the task posterior as its ideal control law, $\tilde{p}_t = p_t$; Sections 3.4–3.5 characterize its fidelity under readout and generation errors. The criterion concerns declared output attributes. With θ fixed, feedback follows $\mathcal{D}_t \to b_t \to p_t \to r_h \to f(X)$: external evidence revises the belief and thereby the distribution of subsequent controls.

3.2 Quantum belief dynamics

Encode the joint belief in the computational-basis diagonal of a density operator ρ_t. Registers are ordered as (Σ, H), with k noise levels and n hypotheses. For refresh rates $\lambda_H, \lambda_\Sigma \in [0, 1]$, define

$$\begin{aligned} \mathcal{E}_H(\rho) &= (1 - \lambda_H)\rho + \lambda_H \ \mathrm{Tr}_H(\rho) \otimes |u_H\rangle\langle u_H|, \\ \mathcal{E}_\Sigma(\rho) &= (1 - \lambda_\Sigma)\rho + \lambda_\Sigma \, |u_\Sigma\rangle\langle u_\Sigma| \otimes \mathrm{Tr}_\Sigma(\rho), \end{aligned} \tag{5}$$

where $|u_H\rangle$ and $|u_\Sigma\rangle$ are uniform superpositions. The prediction is $\bar{\rho}_t = \mathcal{E}_\Sigma(\mathcal{E}_H(\rho_{t-1}))$. Each completely positive, trace-preserving channel refreshes one variable while retaining the other's marginal. The resulting state is generally mixed. Pure uniform replacements, used here for coherent channel preparation, have the same basis probabilities as maximally mixed replacements.

For $M_t > 0$ satisfying $M_t \ge \max_s L_t(s)$, likelihood conditioning uses

$$\begin{aligned} B_t &= \sum_s \sqrt{L_t(s)/M_t} \, |s\rangle\langle s|, \\ \rho_t &= \frac{B_t \bar{\rho}_t B_t^\dagger}{\mathrm{Tr}(B_t \bar{\rho}_t B_t^\dagger)}. \end{aligned} \tag{6}$$

The denominator is the success probability, assumed positive.

Proposition 1 (Bayesian-filter equivalence) *If $\mathrm{diag}(\rho_{t-1}) = b_{t-1}$, the successful update has diagonal b_t for the factored transition $T(s \mid s') = T_H(h \mid h')T_\Sigma(\sigma \mid \sigma')$, where $T_H(h \mid h') = (1 - \lambda_H)\mathbf{1}[h = h'] + \lambda_H/n$ and T_Σ is analogous. The likelihood scale cancels from the posterior.*

Channel diagonals give the classical prediction; conditioning multiplies them by the likelihood (Appendix A.1). Retaining the joint belief preserves hypothesis–noise dependence that can affect later updates even when current task marginals agree (Appendix A.4).

3.3 Measurement and the generative response

Measuring $\Pi_h = I_\Sigma \otimes |h\rangle\langle h|$ selects h with probability $\mathrm{Tr}(\rho_t \Pi_h)$. Mapping this outcome to r_h defines

$$\mathcal{G}_\theta[\rho] = \sum_h \mathrm{Tr}(\rho \Pi_h) K_\theta(\cdot \mid r_h). \tag{7}$$

Evidence changes the frequencies of conditional responses; interventions on the attribute map change the conditions those responses express.

Corollary 1 (diagonal sufficiency). For computational-basis dephasing Δ, $\mathcal{G}_\theta[\rho] = \mathcal{G}_\theta[\Delta(\rho)]$. Under the specified channels and diagonal instruments, dephasing the initial state also preserves every subsequent control law.

Thus measurement probabilities determine the output law. Because projective readout generally disturbs the state, repeated outputs and subsequent filtering require equivalent preparations or a persistent-readout protocol. Our simulations assume independently available successful-state snapshots; Appendix C details the corresponding quantum constructions.

3.4 FROM STATE READOUT TO OUTPUT ATTRIBUTES

The next bound relates uncertainty preservation to the accuracy of readout and conditional generation.

Proposition 2 (Feedback fidelity) *Equip $\mathcal{A} \subseteq [0,1]^d$ with $d_{\mathcal{A}}(u,v) = \|u-v\|_2/\sqrt{d}$ and diameter $D \leq 1$. For $X \sim K_\theta(\cdot \mid r_h)$, let $\epsilon_h = \mathbb{E}[d_{\mathcal{A}}(f(X), r_h)]$. Then*

$$W_1(f_{\#}Q, \nu_p) \leq \sum_h \tilde{p}(h)\epsilon_h + D \ \mathrm{TV}(\tilde{p}, p). \tag{8}$$

The first term measures conditional-generation error; the second measures distortion in readout probabilities. Coupling each output to its selected condition and transporting unmatched probability mass proves the bound (Appendix A.2). It gives an explicit tolerance for $\mathsf{D} = W_1$, with $f_{\#}Q = \nu_p$ under exact readout and realization.

For categorical controls $r_h = h$, the analogous bound is $\mathrm{TV}(f_{\#}Q, p) \leq \sum_h \tilde{p}(h)e_h + \mathrm{TV}(\tilde{p}, p)$, where $e_h = \Pr[f(X) \neq h \mid h]$. Safety flags are included as an additional outcome with zero target mass. For readout from states $\rho, \tilde{\rho}$, the readout term in Proposition 2 is at most $D\|\tilde{\rho} - \rho\|_1/2$. Finite output sets also incur empirical sampling error (Appendix A.2). These bounds characterize transmission relative to the specified posterior; inference accuracy additionally depends on the observation model and hypothesis bank.

Our regional diagnostic uses empirical energy discrepancy. For $N \geq 2$ paired conditions and generated attributes with mean condition error ϵ, generation changes the unbiased discrepancy to a common reference by at most 4ϵ. It changes the gap between a point method m and posterior sampling by at most $4(\epsilon_m + \epsilon_{\mathrm{sample}})$ (Appendix D.3), connecting conditional fidelity directly to this finite-sample diagnostic.

3.5 INFORMATION LOST BY A POINT SUMMARY

A point estimate can be accurate while discarding distinctions relevant to distribution matching.

Proposition 3 (Indistinguishable summaries) *Suppose a generator's only belief-dependent input is a deterministic summary $F(b)$. If admissible beliefs b, b' satisfy $F(b) = F(b')$, its common output-attribute law μ satisfies*

$$\max\{W_1(\mu, \nu_b), W_1(\mu, \nu_{b'})\} \geq \tfrac{1}{2} W_1(\nu_b, \nu_{b'}). \tag{9}$$

The triangle inequality proves the claim. Admissible beliefs are allowed interface inputs for a control intervention, or reachable posteriors for a trajectory claim. Mean and MAP readouts can map distinct target laws to the same summary.

Combining the propositions gives a sufficient condition for full-belief control to improve worst-case fidelity: for a same-summary pair separated by Δ, if each full-belief output error is at most $e < \Delta/2$, then its worst-case error is smaller than that of every summary-only interface on the pair (Appendix A.3). This conclusion concerns posterior fidelity; true-task accuracy also depends on inference.

Finite-set coverage is a related objective with different optimal allocations (Appendix D.2). We therefore compare independent posterior sampling with posterior-proportional quotas in the semantic study.

4 EXPERIMENTS

We test uncertainty transmission, adaptation to revised evidence, and transfer across conditional generators. Generation uses equivalent classical readout probabilities; quantum simulations test the state dynamics and control readout.

4.1 EXPERIMENTAL DESIGN

Tasks and generators. Regional control infers 16 image hypotheses and eight noise levels from noisy RGB measurements. A frozen CIFAR-10 DDPM (Ho et al., 2020; Krizhevsky, 2009) imposes

12 regional RGB means through 200-step DDNM–DDIM sampling (Wang et al., 2023; Song et al., 2021). Semantic control infers classes from corrupted real images using frozen ResNet-20 likelihoods (He et al., 2016); category prompts drive frozen Stable Diffusion 1.5 (Rombach et al., 2022), with native 512×512 outputs assessed by CLIP (Radford et al., 2021). These tasks instantiate continuous and categorical attribute spaces (Appendices B and D.5).

Comparisons. Readout comparisons share the posterior and vary MAP, mean, sampling, or quotas. Update comparisons share the generator and sampled readout: frozen belief stops updating after observation 24; target-static inference uses $\lambda_H = 0$ with unchanged noise transitions. Semantic baselines add task-defined window and discount controls, plus learned hazard (Wilson et al., 2010) (Appendix D.5.1); a pilot varies DDNM, DPS, and TFG (Section 4.5). Comparisons match evidence, output counts, evaluation times, and generation randomness where applicable. Controllers observe neither latent targets nor change times.

Evaluation. We average the standard energy score (Gneiting & Raftery, 2007) over M histories and evaluation times $\mathcal{T}$. This mean R evaluates saved output sets against common hidden targets:

$$R = \frac{1}{M|\mathcal{T}|} \sum_{r=1}^{M} \sum_{t \in \mathcal{T}} \left[\frac{1}{N} \sum_{i} d(z_{rti}, z^{\star}_{rt}) - \frac{1}{2N^2} \sum_{i,j} d(z_{rti}, z_{rtj}) \right]. \tag{10}$$

Here z is a generated attribute, $z^{\star}$ the hidden target, and N the output count. Regional distance is Euclidean distance divided by $\sqrt{12}$; categorical mismatch yields half the Brier loss of output frequencies. Lower R is better. Appendix G gives the retrospective protocol and exploratory paired intervals. Computation, posterior fidelity, and image quality are assessed separately.

4.2 Transmission of posterior uncertainty

The regional study has 32 histories per scenario: stable, noise shift, and target shift. Changes occur at observation 25; methods produce 16 outputs at observations 1, 24, 30, and 48. Common posteriors isolate the readout effect.

Table 1: **Regional generation under the common score R.** All 96 histories and four snapshots, with 16 outputs per cell. Overall weights the three scenarios equally. Lower is better; paired score intervals appear in Appendix G. This retrospective score evaluates the delivered output sets against truth.

Method	Overall $R \downarrow$	Stable	Noise shift	Target shift
QBF sampling	0.019697	0.018252	0.021671	0.019169
MAP readout	0.028790	0.021363	0.031074	0.033932
Mean readout	0.029967	0.026083	0.034095	0.029723
Frozen belief	0.053630	0.018246	0.018393	0.124251

Across 96 histories, posterior sampling achieves $R = 0.019697$, compared with 0.028790 for MAP readout, 0.029967 for mean readout, and 0.053630 for frozen belief (Table 1). These correspond to 31.6% and 34.3% lower loss than the two point readouts. Their paired differences are 0.009092 (exploratory 95% interval $[0.005168, 0.013052]$) and 0.010270 ($[0.009020, 0.011563]$), respectively; 62 of 96 history averages favor sampling over MAP.

Common-posterior diagnostics isolate the readout effect: sampling improves posterior fidelity over both point readouts in all 96 history averages (Table 21). The MAP discrepancy contrast is 0.030042 at the requested conditions and 0.030025 at the outputs; the corresponding mean contrasts are 0.024652 and 0.024627. Their close agreement shows that conditional generation preserves the readout difference.

Figure 2 illustrates concentration under point controls and variation across hypotheses under sampling. An intervention with six belief pairs sharing both mean and MAP yields distinguishable generated regional attributes in three pairs (intervals excluding zero). This supports information transmission through the declared attribute channel (Appendix D.4.2).

The effect tracks remaining uncertainty: in stable histories, the MAP-minus-sampling discrepancy decreases from 0.106646 at observation 1 to 0.000015 at observation 48 as beliefs concentrate. Its

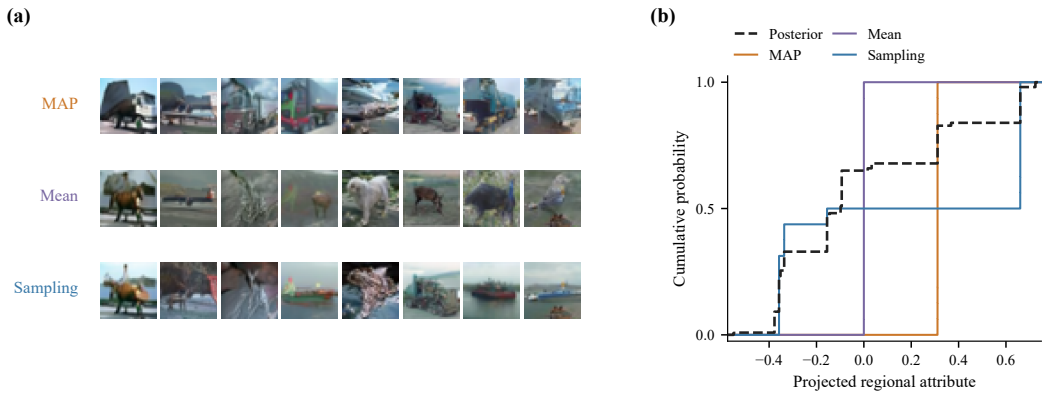

Figure 2: **Point readouts versus distributional feedback.** (a) First eight paired outputs from the first stored history at $t = 1$. (b) CDFs of all 16 outputs, projected onto the leading posterior attribute direction (64.4% of reference variance), selected from the reference alone. Table 1 evaluates all 12 attributes and 96 histories.

association with attribute variance persists after excluding observation 1 (Appendix D.1), linking the readout benefit to uncertainty over alternative conditions.

4.3 From belief revision to changing outputs

At the common observations $\mathcal{T} = \{30, 48\}$, target-shift loss is 0.005382 for dynamic sampling, 0.215546 for frozen belief, and 0.147259 for cumulative target-static sampling (Table 19). All 32 switching histories favor dynamics over target-static control. Because both controllers assimilate every observation and share noise transitions and readout, this comparison isolates target-transition modeling. The complementary strength of target-static control is its lower average loss under stable targets and noise-only changes.

Three further banks reproduce the overall ordering: with 96 histories, eight outputs, and three snapshots, R is 0.029551 for sampling, 0.040413 for MAP, 0.044173 for mean, and 0.080087 for target-static control (Table 20). ImageNet-pretrained ResNet-18 features (He et al., 2016) corroborate the switching response (Table 9). This replication is evaluated separately because its output budget and observation times differ.

When the original target returns, both controllers recover in all 32 histories, with mean completion-step delays of 7.78 observations for dynamic and 4.72 for target-static inference (Appendix D). Together, these results connect belief revision to output adaptation and show how transition assumptions favor different temporal settings.

4.4 Semantic response and finite-output readout

The semantic study uses corrupted CIFAR-10 streams under stable and switching targets, with changes at observation 25. Calibrated likelihoods update the class–noise belief. Configurations selected on 40 development histories are fixed for 60 test histories; learned hazard adapts Wilson et al. (2010) (Appendix D.5.2). Generation uses the original 20-history subset, with 16 outputs at observations 1, 24, 30, and 48. Five update controllers plus QBF MAP and quota readouts yield seven policies (Appendix D.5.1).

On the common four-snapshot grid, QBF sampling achieves $R = 0.073242$, equal to quota readout and close to learned hazard (0.073291), discounting (0.073682), and windowing (0.076855; Table 2). The learned-hazard-minus-QBF difference is 0.000049, with exploratory paired 95% interval $[-0.003564, 0.002930]$; intervals against quota, discount, window, and MAP also include zero. These observed scores place QBF alongside the adaptive alternatives, while the paired intervals leave their ordering unresolved.

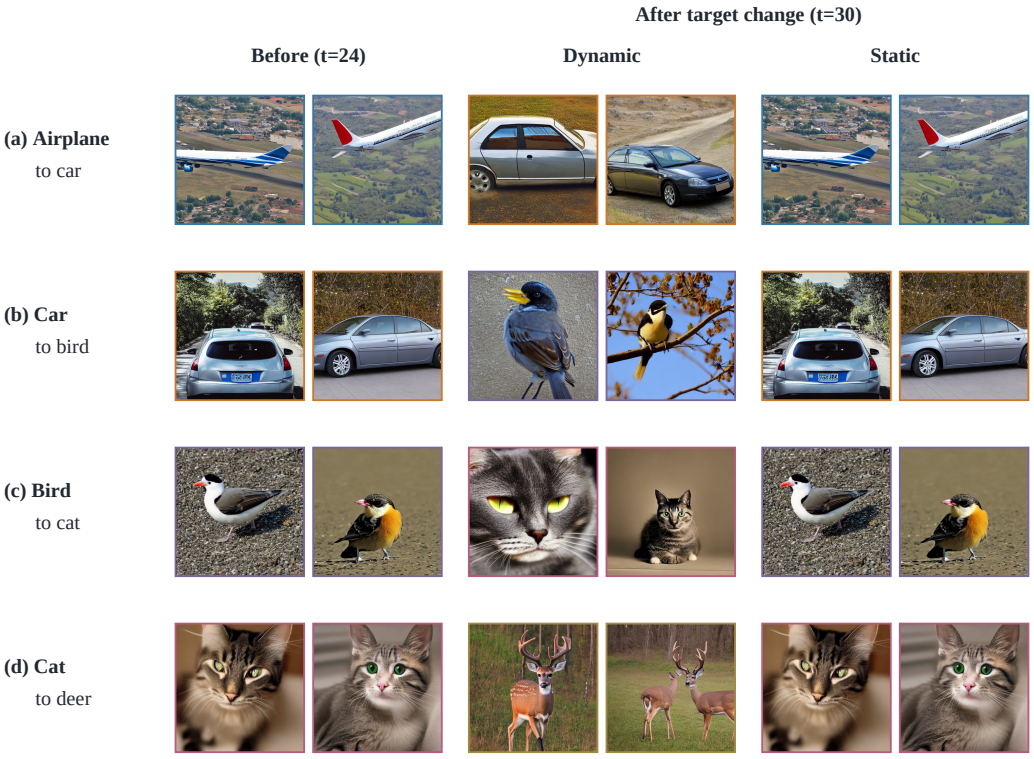

Figure 3: **Evidence redirects generation.** Each row contrasts a shared pre-change output with dynamic and target-static responses after the change, using the same generator and two noise slots. Rows (a–d) show the first four class-ordered histories (24 display positions), without image-based selection; Appendix D.5 gives the full comparison.

Table 2: **Semantic generation under the same score R.** Twenty histories, four snapshots, 16 outputs per cell; failures remain an eleventh category. Overall weights stable and switching scenarios equally. The last four controllers use sampled readout; MAP and quota use the QBF belief. Lower is better. Paired intervals are in Appendix G. Learned hazard adapts Wilson et al. (2010); the other controls are defined in Appendix D.5.1.

Method	Overall $R \downarrow$	Stable	Target shift
QBF sampling	0.073242	0.083594	0.062891
MAP readout	0.087695	0.100293	0.075098
Quota readout	0.073242	0.082812	0.063672
Target-static	0.248242	0.082520	0.413965
Window	0.076855	0.082520	0.071191
Discount	0.073682	0.083301	0.064062
Learned hazard	0.073291	0.085156	0.061426

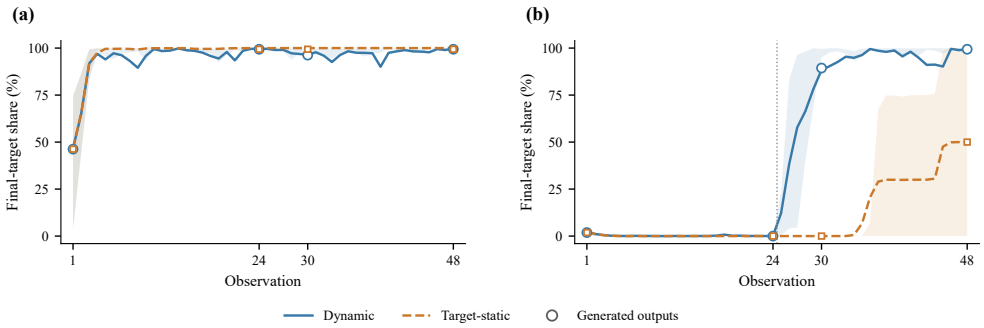

Figure 4: **Belief updates and generated outputs.** (a) Stable target; (b) target change at $t = 25$. Lines show final-target posterior mass over ten histories; bands are interquartile ranges. Circles and squares show CLIP-labelled output frequencies for dynamic and target-static control at four snapshots (160 output positions each). In (b), the tracked target appears at $t = 25$.

Scenario scores reveal complementary behavior: QBF has lower loss under stability (0.083594 versus 0.085156 for learned hazard), whereas learned hazard has lower switching loss (0.061426 versus 0.062891). QBF also improves on target-static control overall (0.248242). Figures 3–4 connect revised beliefs to changing outputs. Additional diagnostics show faster inference recovery for learned hazard and improved early posterior-frequency fidelity for quotas (Table 22).

4.5 Alternative conditional generators

A development pilot combines MAP, mean, and sampled readouts with DDNM (Wang et al., 2023), DPS (Chung et al., 2023), and TFG (Ye et al., 2024). Within each readout, generators share the posterior, conditions, initial noise, and frozen backbone. The untuned pilot comprises three histories (one per scenario), four snapshots, and two images per configuration, totaling 216 images.

Sampled DDNM, DPS, and TFG achieve $R = 0.017279$, 0.018070, and 0.100899, respectively. Sampling gives the lowest aggregate score within each generator, while MAP is favored in the stable scenario (Table 15). These descriptive results extend the interface to multiple conditioning methods at a two-output budget; Appendix F reports the configurations and computational costs.

4.6 Quantum state dynamics and control readout

Ideal density channels, gate-derived instruments, and an 18-qubit two-observation history reproduce QBF's posterior conditioned on success. One amplification iteration raises joint success from 0.05534 to 0.42729 while preserving the conditional state. Matched classical and quantum readouts produce identical pixels in an eight-image control, establishing consistency between the simulated dynamics and the generative control law. Appendix C details preparation, postselection, and repeated readout; Appendix E examines attribute-map interventions and factored belief transitions.

5 Discussion and Conclusion

QBF formulates generative control as matching output-attribute distributions to evolving beliefs, with quantum state dynamics and measurement determining the control law. Our analysis separates readout and generation errors and characterizes distinctions lost by point summaries. Regional experiments show lower mean energy scores and adaptation after target changes; semantic generation extends the interface to high-resolution outputs, with scores close to adaptive feedback baselines.

Distributional readout retains supported alternatives, while transition modeling balances persistence and change. Richer attributes and broader image-quality assessments can extend this connection between uncertainty and generative behavior. The quantum construction yields a classically reproducible control law; computational advantage and persistent operation remain questions for state preparation and repeated-readout protocols.

REFERENCES

Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp. Quantum Amplitude Amplification and Estimation. In *Quantum Computation and Quantum Information*, volume 305 of *Contemporary Mathematics*, pp. 53–74. American Mathematical Society, 2002. URL https://arxiv.org/abs/quant-ph/0005055.

Hyungjin Chung, Jeongsol Kim, Michael T. McCann, Marc L. Klasky, and Jong Chul Ye. Diffusion Posterior Sampling for General Noisy Inverse Problems. In *International Conference on Learning Representations*, 2023. URL https://arxiv.org/abs/2209.14687.

Tilmann Gneiting and Adrian E. Raftery. Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007. doi: 10.1198/016214506000001437.

Alex Graves, Rupesh Kumar Srivastava, Timothy Atkinson, and Faustino Gomez. Bayesian Flow Networks. arXiv:2308.07037, 2023. URL https://arxiv.org/abs/2308.07037.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016. URL https://openaccess.thecvf.com/content_cvpr_2016/html/He_Deep_Residual_Learning_CVPR_2016_paper.html.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. In *Advances in Neural Information Processing Systems*, 2020. URL https://arxiv.org/abs/2006.11239.

Jennifer A. Hoeting, David Madigan, Adrian E. Raftery, and Chris T. Volinsky. Bayesian Model Averaging: A Tutorial. *Statistical Science*, 14(4):382–417, 1999. doi: 10.1214/ss/1009212519.

Yinan Huang, Hans Hao-Hsun Hsu, Junran Wang, Bo Dai, and Pan Li. Accelerated Sequential Flow Matching: A Bayesian Filtering Perspective. arXiv:2602.05319, 2026. URL https://arxiv.org/abs/2602.05319.

Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and Acting in Partially Observable Stochastic Domains. *Artificial Intelligence*, 101(1–2):99–134, 1998. doi: 10.1016/S0004-3702(98)00023-X.

Alex Kendall and Yarin Gal. What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? In *Advances in Neural Information Processing Systems*, volume 30, 2017.

Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and Percy Liang. Concept Bottleneck Models. In *Proceedings of the 37th International Conference on Machine Learning*, 2020. URL https://proceedings.mlr.press/v119/koh20a.html.

Alex Krizhevsky. Learning Multiple Layers of Features from Tiny Images. Technical report, University of Toronto, 2009. URL https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf.

Akshay Kulkarni, Ge Yan, Chung-En Sun, Tuomas Oikarinen, and Tsui-Wei Weng. Interpretable Generative Models through Post-hoc Concept Bottlenecks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025. URL https://arxiv.org/abs/2503.19377.

Jaerin Lee, Daniel Sungho Jung, Kanggeon Lee, and Kyoung Mu Lee. SemanticDraw: Towards Real-Time Interactive Content Creation from Image Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13021–13030, 2025.

Sergey Levine. Reinforcement Learning and Control as Probabilistic Inference: Tutorial and Review. arXiv:1805.00909, 2018. URL https://arxiv.org/abs/1805.00909.

Seth Lloyd and Christian Weedbrook. Quantum Generative Adversarial Learning. *Physical Review Letters*, 121:040502, 2018. doi: 10.1103/PhysRevLett.121.040502.

Guang Hao Low, Theodore J. Yoder, and Isaac L. Chuang. Quantum Inference on Bayesian Networks. *Physical Review A*, 89:062315, 2014. doi: 10.1103/PhysRevA.89.062315.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, pp. 8748–8763, 2021. URL https://proceedings.mlr.press/v139/radford21a.html.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-Resolution Image Synthesis with Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10684–10695, 2022. URL https://arxiv.org/abs/2112.10752.

Raghav Singhal, Zachary Horvitz, Ryan Teehan, Mengye Ren, Zhou Yu, Kathleen McKeown, and Rajesh Ranganath. A General Framework for Inference-time Scaling and Steering of Diffusion Models. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 55810–55827, 2025. URL https://proceedings.mlr.press/v267/singhal25b.html.

Andrei N. Soklakov and Rüdiger Schack. Bayesian Updating of a Probability Distribution Encoded on a Quantum Register. *Physical Review A*, 74:032315, 2006. doi: 10.1103/PhysRevA.74.032315.

Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising Diffusion Implicit Models. In *International Conference on Learning Representations*, 2021. URL https://arxiv.org/abs/2010.02502.

Gábor J. Székely and Maria L. Rizzo. Energy statistics: A class of statistics based on distances. *Journal of Statistical Planning and Inference*, 143:1249–1272, 2013. doi: 10.1016/j.jspi.2013.03.018.

Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-Shot Image Restoration Using Denoising Diffusion Null-Space Model. In *International Conference on Learning Representations*, 2023. URL https://arxiv.org/abs/2212.00490.

Lifu Wei, Yinuo Ren, Naichen Shi, and Yiping Lu. SURGE: Approximation and Training Free Particle Filter for Diffusion Surrogate. In *International Conference on Machine Learning*, 2026. URL https://arxiv.org/abs/2605.18745.

Nathan Wiebe and Christopher Granade. Can small quantum systems learn. *Quantum Information and Computation*, 17(7&8):568–594, 2017. doi: 10.26421/QIC17.7-8-2.

Robert C. Wilson, Matthew R. Nassar, and Joshua I. Gold. Bayesian Online Learning of the Hazard Rate in Change-Point Problems. *Neural Computation*, 22(9):2452–2476, 2010. doi: 10.1162/NECO_a_00007.

Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. Efficient Streaming Language Models with Attention Sinks. In *International Conference on Learning Representations*, 2024.

Yang Yang, Severi Rissanen, Paul E. Chang, Nasrulloh Loka, Daolang Huang, Arno Solin, Markus Heinonen, and Luigi Acerbi. PriorGuide: Test-Time Prior Adaptation for Simulation-Based Inference. In *International Conference on Learning Representations*, 2026. URL https://arxiv.org/abs/2510.13763.

Haotian Ye, Haowei Lin, Jiaqi Han, Minkai Xu, Sheng Liu, Yitao Liang, Jianzhu Ma, James Zou, and Stefano Ermon. TFG: Unified Training-Free Guidance for Diffusion Models. In *Advances in Neural Information Processing Systems*, volume 37, 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/2818054fc6de6dacdda0f142a3475933-Abstract-Conference.html.