

An Internal Alert for Errors the Policy Can Exploit

Abstract

Sharpness-based alerts and value-aware learning concern different views of model error. Evaluation should compare alert calibration with errors in state-action regions emphasized by the policy. It should independently measure the consequence of acting on an alert. A correlation with static factuality labels cannot establish that the alert prevents policy exploitation of a learned model. This focused synthesis develops an empirical comparison that preserves the terminal action, observation process, and total verification cost. Its proposed tests are prospective; no new simulation run, model-training result, or intervention effect is reported.

A value model and a behavior distribution

Sharpness-based alerts and value-aware learning concern different views of model error. Evaluation should compare alert calibration with errors in state-action regions emphasized by the policy. It should independently measure the consequence of acting on an alert. A correlation with static factuality labels cannot establish that the alert prevents policy exploitation of a learned model.

In-context activation sharpness is investigated as a signal for detecting and mitigating language-model hallucinations (S. Chen et al., 2024). Robust value-aware model learning uses adaptive weighting to address model exploitation in offline reinforcement learning (Z. Qiao et al., 2026b).

Offline transitions describe actions selected by a behavior process. A learned policy can prefer different actions and visit different states, which makes its value uncertain outside that support. Heterogeneous data add a further problem when values correspond to different dynamics. Good prediction within the stored file does not resolve either issue by itself.

An artificial-intelligence workflow can change parameters, representations, retrieved evidence, actions, or verification effort. Each intervention changes a different part of the information process. An efficient update is not the same as a well-calibrated alert, and a grounded agent is not the same as a fluent answer. A comparison should preserve these distinct objects while evaluating their contribution to an explicitly defined task. Resource savings have practical value when they preserve useful outcomes or permit evidence that was previously too expensive to collect.

From an estimator to a defensible action

For “In-Context Sharpness as Alerts: An Inner Representation Perspective for Hallucination Mitigation,” a useful methodological question follows from a value model and a behavior distribution: an internal signal needs external outcome validation before it can justify a verification or abstention decision. Calibration concerns whether reported confidence corresponds to observed correctness frequencies. It is different from ranking risk and different from achieving a high accuracy score. A confidence adjustment developed on one collection needs evaluation on the population where it will guide actions. In a decision workflow, the calibration unit should match the event used for acceptance or escalation. A score for a segment cannot silently become a probability that an entire response is correct. Confidence-dependent allocation needs both a calibration check and an outcome comparison for the actions it triggers. (Guo et al., 2017).

The requirement in “Model-Based Offline RL via Robust Value-Aware Model Learning with Implicitly Differentiable Adaptive Weighting” is sharpened by this principle: model error should be assessed where it affects policy values rather than only through an average prediction score. Reinforcement learning

defines value through rewards accumulated under actions and transitions. This makes it different from supervised prediction on a static response table. A policy changes the distribution of future observations, so an evaluation needs to consider the states and actions it induces. The horizon and reward definition are substantive parts of the problem. A claimed policy improvement should be supported by an evaluation design appropriate to that interaction, rather than by treating a fitted association on old records as the effect of deploying the learned policy. (Sutton and Barto, 2018).

One way to examine a value model and a behavior distribution is to preserve the boundary identified for “In-Context Sharpness as Alerts: An Inner Representation Perspective for Hallucination Mitigation”: an internal signal needs external outcome validation before it can justify a verification or abstention decision. Offline reinforcement learning uses previously collected interaction data without acquiring new training interactions. This exposes a support problem when the learned policy recommends actions poorly represented in those records. The quality of the dataset therefore cannot be summarized by row count alone. Coverage, behavior policy, and dynamics determine what can be learned and evaluated. When data from different environments are combined, the analyst should identify which relationships are assumed transferable and which are allowed to vary. A static fit statistic cannot fully characterize the value of a new policy outside observed support. (Levine et al., 2020).

For “Model-Based Offline RL via Robust Value-Aware Model Learning with Implicitly Differentiable Adaptive Weighting,” a useful methodological question follows from a value model and a behavior distribution: model error should be assessed where it affects policy values rather than only through an average prediction score. Conservative value learning addresses the danger of optimistic value estimates for actions inadequately supported by offline data. This is a specific response to extrapolation risk, not a general guarantee of safety across every deployment shift. The amount of conservatism can affect which actions remain attractive. Evaluation should compare estimates with realized outcomes in an appropriate independent environment and examine sensitivity to coverage. A procedure can be too optimistic in one region and unnecessarily pessimistic in another; both affect the decision loss and the use of available information. (Kumar et al., 2020).

The requirement in “In-Context Sharpness as Alerts: An Inner Representation Perspective for Hallucination Mitigation” is sharpened by this principle: an internal signal needs external outcome validation before it can justify a verification or abstention decision. Retrieval-augmented generation separates some evidence access from information stored in model parameters. Its reliability consequently involves the retrieval stage as well as the language-generation stage. A relevant retrieved passage can support a correct answer, while an irrelevant or outdated one can introduce a different error path. A workflow comparison should retain the retrieved evidence and count retrieval costs. Source access also does not eliminate the need to check whether the final claim is supported by that evidence. Adaptation and retrieval are distinct interventions whose benefits should be evaluated separately. (Lewis et al., 2020).

One way to examine a value model and a behavior distribution is to preserve the boundary identified for “Model-Based Offline RL via Robust Value-Aware Model Learning with Implicitly Differentiable Adaptive Weighting”: model error should be assessed where it affects policy values rather than only through an average prediction score. Selective classification permits abstention when the system declines to make a sufficiently reliable prediction. This changes both coverage and risk among accepted cases. A low error rate can be obtained by accepting very few cases, so the two quantities must be reported together. In a workflow abstention can trigger review, delay, or another model, each with a cost. An uncertainty score should therefore be evaluated through the decisions it produces, including the rejected cases, rather than reporting only a favorable rate on the subset easiest to accept. (Geifman and El-Yaniv, 2017).

For “In-Context Sharpness as Alerts: An Inner Representation Perspective for Hallucination Mitigation,” a useful methodological question follows from a value model and a behavior distribution: an internal signal needs external outcome validation before it can justify a verification or abstention decision. Predictive uncertainty needs evaluation under changed data distributions because a signal useful in the development population may become misleading elsewhere. Shift can affect inputs, labels, or the relationships connecting them. A single corrupted test file represents only one departure. Applications should state the shifts examined and preserve group-specific outcomes so that broad averages do not conceal local failure. An uncertainty method that reacts appropriately to one stress provides evidence about that stress, while protection against a different mechanism requires its own evaluation. (Ovadia et al., 2019).

The requirement in “Model-Based Offline RL via Robust Value-Aware Model Learning with Implicitly Differentiable Adaptive Weighting” is sharpened by this principle: model error should be assessed where it affects policy values rather than only through an average prediction score. Robust optimization makes the cost of protection visible by comparing decisions under nominal and adverse conditions. The uncertainty description determines which departures are defended against, so conservatism should not be assessed in the abstract. A design can be robust within one description and fragile to a departure excluded from it. Applications can examine how the preferred action changes as protection is increased and identify the constraints responsible for the change. This exposes the tradeoff between nominal performance and the chosen protection without claiming that one uncertainty set is universally correct. (Bertsimas and Sim, 2004).

In-Context Sharpness as Alerts: An Inner Representation Perspective for Hallucination Mitigation

In-context activation sharpness is investigated as a signal for detecting and mitigating language-model hallucinations. Activation sharpness is considered as an internal signal related to hallucination. Its operational value depends on whether it identifies errors that matter in a future response collection. A signal can rank risk without supplying calibrated error probabilities, and mitigation adds a further intervention question. Detection, calibration, and the effect of acting on the alert should therefore be evaluated separately. (S. Chen et al., 2024).

The workflow implication is to identify which object changes: parameters, retrieved evidence, value estimates, visual grounding, or the allocation of review. These interventions can share a budget while requiring different outcome evidence. A static correctness statistic can support one component without validating an interactive policy. Comparison should retain the mechanism and information restriction of each method.

For the present question of a value model and a behavior distribution, this study supplies the starting comparison. Its strongest contribution is the explicit decision object in “In-Context Sharpness as Alerts: An Inner Representation Perspective for Hallucination Mitigation.” The next step would be to identify which observations can change that object at the operating budget. An internal signal needs external outcome validation before it can justify a verification or abstention decision. This is an implication for evaluation design rather than a result obtained by applying the source method to a new system.

Model-Based Offline RL via Robust Value-Aware Model Learning with Implicitly Differentiable Adaptive Weighting

The role of “Model-Based Offline RL via Robust Value-Aware Model Learning with Implicitly Differentiable Adaptive Weighting” in this argument is to qualify how the initial comparison should be interpreted. Under a value model and a behavior distribution, its information model deserves separate scrutiny. Model error should be assessed where it affects policy values rather than only through an average prediction score. A later comparison should retain that distinction and assess the resulting action on evidence that did not determine the preferred estimate. Otherwise methodological sophistication can hide an unresolved application assumption.

Robust value-aware model learning uses adaptive weighting to address model exploitation in offline reinforcement learning. Value-aware model learning focuses on model errors that matter to policy evaluation, while adaptive weighting changes the emphasis placed on training observations. This is a different objective from minimizing an unweighted prediction loss. The risk is that a policy exploits inaccuracies in the learned model. Assessing the approach therefore requires evidence about the states and actions emphasized by the resulting policy, not only a global fit statistic. (Z. Qiao et al., 2026b).

The second issue is transfer under shift. Task composition, graph topology, transition dynamics, and internal representations describe different departures from the development setting. Protection against one does not establish protection against the others. Domain-specific and mechanism-specific checks should accompany aggregate quality so that an apparently stable score does not conceal a failure in the cases that determine the action.

Testing the remaining decision uncertainty

A test should declare the behavior data available to learning and preserve a separate policy-evaluation environment. State-action coverage diagnostics can show where the candidate policy relies on extrapolation. Dynamics changes should be varied independently of sample-size changes. Conservative estimates can be compared with realized return in the evaluation environment without assuming that one penalty scale works everywhere. If no environment or valid evaluation design exists, the result should remain a modeling proposal rather than a measured policy benefit.

A prospective assessment would fix the base system and information available to each method, then vary the intervention under a common task and resource definition. Domain-specific outcomes would accompany aggregate scores. Training cost, inference work, retrieval, search, and verification would be recorded separately before an end-to-end comparison was made. Where a method changes actions in an environment, independent interactive outcomes would be required. A static annotation file could support outcome definitions and a diagnostic analysis without demonstrating the benefit of an unrun intervention. Grounding, support coverage, calibration, and shift sensitivity would be tested at the level relevant to the method. The final interpretation would identify which mechanism has been examined and would preserve uncertainty about mechanisms absent from the evaluation.

The information boundary of the comparison

The comparison between “In-Context Sharpness as Alerts: An Inner Representation Perspective for Hallucination Mitigation” and “Model-Based Offline RL via Robust Value-Aware Model Learning with Implicitly Differentiable Adaptive Weighting” turns on the following distinction. An internal signal needs external outcome validation before it can justify a verification or abstention decision. Model error should be assessed where it affects policy values rather than only through an average prediction score. These qualifications identify different reasons a recommendation can fail. An assessment of a value model and a behavior distribution should preserve each reason rather than combine them through a generic accuracy score. The useful common outcome is a declared action evaluated on independent information at a counted resource limit. The common outcome may require new evidence not observed in one of the original settings. That gap is a prospective requirement rather than information supplied by the literature comparison itself.

A formal way to retain this distinction is to write the expected outcome of alternative a in context x as $\mu(a, x)$, and to define the admissible alternatives by the requirements in that context. A prediction model estimates a relationship; a selection procedure uses estimated relationships to recommend an action. The selected object can instead be a retained group or an acceptable set. Its loss depends on the consequence of an exclusion, acceptance, or wrong commitment. These objects should be stated before a budget is optimized. An accurate coefficient or a precise posterior does not define the loss, the requirements, or the operating population on its own.

An observational file may contain outcomes under only the actions historically taken. Its conditional association therefore cannot generally evaluate every alternative a . A released simulation table can cover designed alternatives while still omit the intended field environment. Independent replications from a validated simulator offer a different measurement opportunity from resampling either file. A future comparison needs to state which opportunity it uses, the variables available before action, and the information withheld for terminal checking. This permits estimation evidence and selection evidence to be connected without treating a secondary projection as a policy experiment.

Alerts should be checked against content-stable policy outcomes

Zhou and colleagues keep original replies unchanged while altering their surrounding presentation and checking safety-judge verdicts against repeat disagreement. The resulting sensitivity is judge-specific. Their limited wrapper set and human-validation scope constrain generalization; the study concerns safety judgments rather than factuality labels (Y. Zhou et al., 2026c).

The internal-alert argument concerns errors a policy can exploit. An evaluator that rewards the appearance of caution can obscure those errors even when the response still contains the same consequential content. Holding the reply fixed isolates the evaluator's contribution to that discrepancy. This suggests a staged assessment: establish the error and its action consequence independently, then ask whether the alert anticipates it and whether the judge recognizes it. The source provides a safety-evaluation warning, not a demonstration that the policy or alert mechanism discussed here has been tested with wrappers.

Retain independent action or content adjudication for disputed cases. Compare false reassurance and unnecessary intervention separately under paired presentations. Judge version and grading rubric belong in the validation record because they can alter the apparent alert benefit. No new wrapper judgments were collected for this article. The proposed audit concerns a future evaluation instrument, and does not alter the completed results or establish that the same sensitivity occurs in this task.

Conclusion

Evaluation should compare alert calibration with errors in state-action regions emphasized by the policy. It should independently measure the consequence of acting on an alert. A correlation with static factuality labels cannot establish that the alert prevents policy exploitation of a learned model. The decisive next evidence is an independently checked action under the declared information restriction. Its resource accounting should include the work required to verify the recommendation and the assumptions responsible for that recommendation. This defines a concrete research priority while keeping the application-specific value open to measurement.

References

1. Z. Qiao, J. Lyu, B. Lyu, Y. Shu, S. Gao, and S. Qiu, “Model-Based Offline RL via Robust Value-Aware Model Learning with Implicitly Differentiable Adaptive Weighting,” *Proceedings of the 2026 International Conference on Learning Representations*, 2026.
2. S. Chen, M. Xiong, J. Liu, Z. Wu, T. Xiao, S. Gao, and J. He, “In-Context Sharpness as Alerts: An Inner Representation Perspective for Hallucination Mitigation,” *Proceedings of the 41st International Conference on Machine Learning*, 2024.
3. Zhou, Yongxi, Wenbo Ye, Yuanzhe Liu, Zihan Dong, and Junwei Yao. “Style Over Substance: Content-Invariant Wrappers Flip LLM Safety-Judge Verdicts.” *arXiv*, 2026, arXiv:2609.08236. <https://arxiv.org/abs/2609.08236>.
4. Guo, Chuan, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. “On Calibration of Modern Neural Networks.” *Proceedings of the 34th International Conference on Machine Learning*, PMLR, vol. 70, 2017, pp. 1321-1330. <https://proceedings.mlr.press/v70/guo17a.html>.
5. Sutton, Richard S., and Andrew G. Barto. *Reinforcement Learning: An Introduction*. 2nd ed., MIT Press, 2018. <http://incompleteideas.net/book/the-book-2nd.html>.
6. Levine, Sergey, Aviral Kumar, George Tucker, and Justin Fu. “Offline Reinforcement Learning: Tutorial, Review, and Perspectives on Open Problems.” *arXiv*, 2020, arXiv:2005.01643. <https://arxiv.org/abs/2005.01643>.
7. Kumar, Aviral, Aurick Zhou, George Tucker, and Sergey Levine. “Conservative Q-Learning for Offline Reinforcement Learning.” *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 1179-1191. <https://arxiv.org/abs/2006.04779>.
8. Lewis, Patrick, et al. “Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks.” *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 9459-9474. <https://arxiv.org/abs/2005.11401>.
9. Geifman, Yonatan, and Ran El-Yaniv. “Selective Classification for Deep Neural Networks.” *Advances in Neural Information Processing Systems*, vol. 30, 2017. <https://arxiv.org/abs/1705.08500>.
10. Ovadia, Yaniv, et al. “Can You Trust Your Model’s Uncertainty? Evaluating Predictive Uncertainty Under Dataset Shift.” *Advances in Neural Information Processing Systems*, vol. 32, 2019. <https://arxiv.org/abs/1906.02530>.
11. Bertsimas, Dimitris, and Melvyn Sim. “The Price of Robustness.” *Operations Research*, vol. 52, no. 1, 2004, pp. 35-53. <https://doi.org/10.1287/opre.1030.0065>.