

Structural Transfer and the Compatibility of Offline Values

Abstract

Source-free graph adaptation preserves structural information, while cross-domain offline learning aligns values collected under different dynamics. Evaluation should identify which relationships are transferred in each method. Graph topology and behavior-policy support require separate diagnostics. Merely combining more target records cannot establish compatible values or supply source information the structural method is not allowed to access. 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.

Graph structure as transferable evidence

Source-free graph adaptation preserves structural information, while cross-domain offline learning aligns values collected under different dynamics. Evaluation should identify which relationships are transferred in each method. Graph topology and behavior-policy support require separate diagnostics. Merely combining more target records cannot establish compatible values or supply source information the structural method is not allowed to access.

A graph encodes relations that a flat case count does not preserve. Transfer can depend on topology, node attributes, neighborhood composition, and the availability of source information. A generic domain label is too coarse to identify which of these mechanisms has changed. The evidence needed for adaptation should therefore be defined in terms of the representation used by the method.

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.

Universal structural-basis distillation studies source-free graph adaptation under changes in graph structure (Y. Wang et al., 2026a). Value alignment and assignment are studied jointly in heterogeneous cross-domain offline reinforcement learning (Z. Qiao et al., 2026a).

Evidence principles behind the proposed comparison

The requirement in “USBD: Universal Structural Basis Distillation for Source-Free Graph Domain Adaptation” is sharpened by this principle: structural transfer requires information about topology that an ordinary tabular benchmark does not provide. 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).

One way to examine graph structure as transferable evidence is to preserve the boundary identified for “Unifying Value Alignment and Assignment in Cross-Domain Offline Reinforcement Learning with Heterogeneous Datasets”: pooling source records can be misleading when their values were learned under

incompatible dynamics or behavior policies. 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).

For “USBD: Universal Structural Basis Distillation for Source-Free Graph Domain Adaptation,” a useful methodological question follows from graph structure as transferable evidence: structural transfer requires information about topology that an ordinary tabular benchmark does not provide. 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).

The requirement in “Unifying Value Alignment and Assignment in Cross-Domain Offline Reinforcement Learning with Heterogeneous Datasets” is sharpened by this principle: pooling source records can be misleading when their values were learned under incompatible dynamics or behavior policies. Gaussian-process prediction uses a covariance function to express relationships among input locations. In a selection application those relationships determine how measurements of one configuration influence the estimated attractiveness of another. A smooth response surface can support useful borrowing, but abrupt feasibility changes or unrepresented regimes can make that borrowing misleading. The model should be assessed at the locations relevant to choice, not only on a random collection of average errors. Posterior uncertainty is conditional on the representation and should not be described as protection against every form of model discrepancy. (Rasmussen and Williams, 2006).

One way to examine graph structure as transferable evidence is to preserve the boundary identified for “USBD: Universal Structural Basis Distillation for Source-Free Graph Domain Adaptation”: structural transfer requires information about topology that an ordinary tabular benchmark does not provide. Low-rank adaptation parameterizes a trainable update more economically than updating every weight of a base model. This creates a resource question about what quality can be preserved or improved under the restricted update. A useful comparison should hold the base model and adaptation evidence fixed while varying the update structure. Trainable parameter count, memory use, optimization work, and inference behavior should remain separate measurements. Favorable task accuracy does not by itself show calibration under shift or the benefit of spending saved resources on additional evaluation. (Hu et al., 2022).

For “Unifying Value Alignment and Assignment in Cross-Domain Offline Reinforcement Learning with Heterogeneous Datasets,” a useful methodological question follows from graph structure as transferable evidence: pooling source records can be misleading when their values were learned under incompatible dynamics or behavior policies. 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 “USBD: Universal Structural Basis Distillation for Source-Free Graph Domain Adaptation” is sharpened by this principle: structural transfer requires information about topology that an ordinary tabular benchmark does not provide. 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 graph structure as transferable evidence is to preserve the boundary identified for “Unifying Value Alignment and Assignment in Cross-Domain Offline Reinforcement Learning with Heterogeneous Datasets”: pooling source records can be misleading when their values were learned under incompatible dynamics or behavior policies. 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).

USBD: Universal Structural Basis Distillation for Source-Free Graph Domain Adaptation

For the present question of graph structure as transferable evidence, this study supplies the starting comparison. Its strongest contribution is the explicit decision object in “USBD: Universal Structural Basis Distillation for Source-Free Graph Domain Adaptation.” The next step would be to identify which observations can change that object at the operating budget. Structural transfer requires information about topology that an ordinary tabular benchmark does not provide. This is an implication for evaluation design rather than a result obtained by applying the source method to a new system.

Universal structural-basis distillation studies source-free graph adaptation under changes in graph structure. Source-free graph adaptation faces an information restriction as well as a distribution shift: the source data are unavailable during adaptation. Structural-basis distillation is therefore relevant to what can be carried forward from earlier learning. A graph's topology is not interchangeable with a table of task labels. Tests of transfer should preserve the structural information required by the method and specify which aspects of the target graph differ from the source setting. (Y. Wang et al., 2026a).

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.

Unifying Value Alignment and Assignment in Cross-Domain Offline Reinforcement Learning with Heterogeneous Datasets

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.

Value alignment and assignment are studied jointly in heterogeneous cross-domain offline reinforcement learning. Heterogeneous offline datasets can encode values obtained under different dynamics and behavior policies. Joint value alignment and assignment addresses the difficulty of combining them meaningfully. Merely increasing the number of stored transitions can intensify the problem if the values are incompatible. A credible comparison needs to distinguish a benefit from additional coverage from a benefit due to resolving that incompatibility. (Z. Qiao et al., 2026a).

The role of “Unifying Value Alignment and Assignment in Cross-Domain Offline Reinforcement Learning with Heterogeneous Datasets” in this argument is to qualify how the initial comparison should be interpreted. Under graph structure as transferable evidence, its information model deserves separate scrutiny. Pooling source records can be misleading when their values were learned under incompatible dynamics or behavior policies. 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.

An empirical agenda with a declared information boundary

A structural transfer test should vary relation patterns separately from attribute shifts and preserve the information restrictions imposed by source-free adaptation. It should report local performance in changed neighborhoods as well as whole-graph accuracy. If downstream decisions use paths or retrieved relations, those intermediate objects should be checked for relevance. A tabular reweighting exercise can illuminate population composition but cannot by itself validate structural adaptation, because it omits the graph on which the method depends.

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.

From conditional estimates to a selection task

The comparison between “USBD: Universal Structural Basis Distillation for Source-Free Graph Domain Adaptation” and “Unifying Value Alignment and Assignment in Cross-Domain Offline Reinforcement Learning with Heterogeneous Datasets” turns on the following distinction. Structural transfer requires information about topology that an ordinary tabular benchmark does not provide. Pooling source records can be misleading when their values were learned under incompatible dynamics or behavior policies. These qualifications identify different reasons a recommendation can fail. An assessment of graph structure as transferable evidence 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.

Conclusion

Evaluation should identify which relationships are transferred in each method. Graph topology and behavior-policy support require separate diagnostics. Merely combining more target records cannot establish compatible values or supply source information the structural method is not allowed to access. 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. Y. Wang, K. Zhang, M. Wang, S. Gao, and N. Yin, “USBD: Universal Structural Basis Distillation for Source-Free Graph Domain Adaptation,” *Proceedings of the 2026 ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2026.
2. Z. Qiao, J. Lyu, C. Bai, P. Wang, S. Gao, and S. Qiu, “Unifying Value Alignment and Assignment in Cross-Domain Offline Reinforcement Learning with Heterogeneous Datasets,” *Proceedings of the 2026 International Conference on Machine Learning*, 2026.
3. 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>.
4. 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>.
5. 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>.
6. Rasmussen, Carl Edward, and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006. <https://gaussianprocess.org/gpml/>.
7. Hu, Edward J., et al. “LoRA: Low-Rank Adaptation of Large Language Models.” *International Conference on Learning Representations*, 2022. <https://openreview.net/forum?id=nZeVKeeFYf9>.
8. 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>.
9. Sutton, Richard S., and Andrew G. Barto. *Reinforcement Learning: An Introduction*. 2nd ed., MIT Press, 2018. <http://incompleteideas.net/book/the-book-2nd.html>.
10. 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>.