

Referee Report on “Learning to Trace Seiberg Dualities”

arXiv:2607.28628

Recommendation

The manuscript presents an ambitious combination of quiver mutation, graph neural networks, and heuristic graph search. The problem is interesting, the exposition is generally energetic and accessible, and the release of code and data is valuable. However, the central numerical claims currently depend on several definitions and benchmarking choices that are either inconsistent within the paper or insufficiently specified. In particular, the reported target distance is not consistently defined on the same state space as the pathfinder, the anomalous out-of-distribution tests violate the involutivity assumption used by bidirectional search, and the very large efficiency ratios are obtained against an analytic branching-tree estimate rather than a measured search baseline. These are not cosmetic issues: addressing them would require new dataset construction, new training splits, and substantial rerunning of the experiments.

I therefore recommend that the manuscript not be accepted in its present form. A substantially reconstructed version could merit consideration as a new submission. Such a version would need to define a single physical or mathematical state space throughout, repair the anomalous-quiver experiment, and replace the present efficiency claims with controlled, reproducible comparisons.

Synopsis and logical structure

The paper studies the problem of finding mutation sequences between ranked directed quivers. Section 2 encodes a quiver by an antisymmetric adjacency matrix together with a rank vector. Equation (2.1) imposes the anomaly-cancellation condition, equation (2.2) updates the rank at a mutated node, and equations (2.3)–(2.4) implement the usual matrix mutation of the net arrow multiplicities. Repeated mutations generate a graph of theories, called a mutation tree in much of the paper.

Section 3 constructs samples of pairs of quivers and associates a distance to each pair. Equation (3.3) first defines distance as the length of a shortest path in the mutation graph. Equations (3.4)–(3.5), by contrast, obtain a label from the longest common prefix of two paths from a chosen root. Section 4 introduces two learned models. The distance GNN (DGNN) predicts a scalar separation using the node-wise comparison in equation (4.6) and the pooled representation in equations (4.7)–(4.8). The action GNN (AGNN) predicts which node to mutate, using Laplacian positional encodings and the construction in equations (4.11)–(4.17).

Section 5 embeds these models in bidirectional search. The DGNN supplies the heuristic in equation (5.5), the AGNN ranks mutations, and the Hybrid method modifies the accumulated cost according to equation (5.8). The work compares these procedures to a branching-tree estimate of breadth-first search and to a deterministic longest-common-ancestor (LCA) path. Section 6 reports in-distribution and out-of-distribution tests; Section 7 analyzes failures and introduces the aggregate complexity variables in equations (7.1)–(7.3). The principal claim is that learned guidance substantially reduces the number of visited states and continues to work beyond the training regime.

The overall progression from representation, to supervised predictors, to guided search is sensible. The principal definitions and assumptions do not, however, remain aligned across these stages. The following points should be resolved before the reported performance can be interpreted.

Major comments

1. The target distance and the state space are not defined consistently.

Equation (3.3) defines the desired quantity as a shortest-path distance in the mutation graph. The LCA prescription in equations (3.4)–(3.5) instead gives the distance in the particular rooted search tree used to generate the two samples. It is only an upper bound on graph distance whenever two branches merge, a cycle produces a shortcut, or permuted representatives describe the same quiver. Such phenomena are central rather than exceptional here. Indeed, the discussion surrounding footnote 29 acknowledges that a pathfinder can find a shorter path than the nominal “ground truth.” This is incompatible with treating that nominal value as the distance defined in equation (3.3), and it also conflicts with the assertion after equation (5.15) that the corresponding gain is bounded below by one.

There is a second layer to this problem. The pathfinder declares success up to graph isomorphism, whereas the generated states and their stored paths appear to retain node labels. The same pair can therefore have one label in a labeled graph, a shorter route in the quotient by simultaneous relabeling, and yet another LCA length in a selected generation tree. These quantities must not be interchanged in training, termination, and evaluation.

The authors should choose and state the vertices and edges of the problem precisely. A natural mathematical choice is the quotient graph of valid ranked quivers under an explicit notion of rank-preserving directed-graph isomorphism. Distances should then be computed by exact shortest-path search on that same graph for all samples within the generated radius. If instead labeled quivers are intended, the isomorphism-based stopping condition should be removed. The manuscript should report how many labels change under the corrected definition, retrain both networks, and recompute every table and breaking point. It should also reconcile the text of equations (3.4)–(3.5) with the linked data generator, which appears to calculate all-pairs shortest paths on a generated labeled graph rather than LCA distances.

2. The physical object being mutated is underspecified for the claimed Seiberg-duality interpretation.

A four-dimensional quiver gauge theory is not in general fixed by an adjacency matrix and ranks alone. Seiberg duality also transforms mesons and the superpotential, after which massive pairs must be integrated out. Equations (2.3)–(2.4) operate only on net arrow multiplicities. In particular, netting opposite arrows cannot distinguish a massive two-cycle from an original vector-like pair with different superpotential data. Consequently, a path in the paper’s graph need not by itself establish equivalence of two quantum field theories.

This does not invalidate the machine-learning problem, but it changes its scope. The authors should either (i) restrict the claims to mutation of ranked, 2-acyclic quivers (or cluster seeds) under clearly stated assumptions on a nondegenerate potential, or (ii) include the superpotential and the integration-out step in the state transition and in the equivalence test. The status of all seed families in Appendix A should be stated under the chosen assumptions. Claims about RG trajectories, duality cascades, holographic throats, and computational cost in the physical theory

should then be separated from claims about distance in the combinatorial mutation graph. At present the paper moves too quickly between these notions.

3. **The experimental split and the learned representations do not establish generalization to unlabeled quivers.**

The paper does not give enough information to reconstruct the train, validation, and test sets. The linked generator appears to assign individual pairs randomly to splits. If so, a single quiver, nearby descendants, and many overlapping pair endpoints can occur in both training and test data. This makes the in-distribution task substantially transductive: memorizing state or family information can be useful even when the model has not learned a transferable distance. Hyperparameters are also described in Section 5.2 as selected to maximize efficiency while retaining 100 percent success, without a clearly separate tuning set.

The architecture makes leakage by labeling especially consequential. In equation (4.6) the two graphs are compared node by node, which presupposes an alignment. The ordered concatenation in equation (4.7) does not make the DGNN output manifestly symmetric under interchange of its two arguments. Neither construction demonstrates invariance under independent node permutations, even though the target is stated to be a symmetric distance between quivers modulo permutation. For the AGNN, the Laplacian eigenvectors in equations (4.11)–(4.17) have sign ambiguities and, in degenerate eigenspaces, arbitrary basis rotations. The manuscript does not explain how equivariance is maintained under these ambiguities. Finally, when several first mutations lie on distinct optimal paths, training on one arbitrarily selected action penalizes other correct actions.

At minimum, the authors should construct state-disjoint and preferably mutation-family-disjoint splits before forming pairs; publish exact counts and random seeds; and reserve an independent validation set for every search hyperparameter. They should test the DGNN numerically under independent permutations and under exchange of source and target. The model should be redesigned or symmetrized if those tests fail. For the AGNN, the authors should document an equivariant positional-encoding treatment and use the set of all shortest-path first actions as a multi-label target. Results from the present random-pair split and from the stronger inductive splits should be reported separately.

4. **The pathfinders do not have the correctness properties attributed to them.**

The DGNN prediction in equation (5.5) is neither proved nor empirically shown to be an admissible and consistent heuristic. The statement following equation (5.8) that a positive unit contribution “ensures admissibility” is incorrect: admissibility is the condition that the heuristic never overestimate the remaining optimal cost. Moreover, the Hybrid method adds an action-dependent term to the path cost. Unless that term is used only as a tie breaker, the search is optimizing a different objective from mutation length.

Algorithm 2 returns when the two searches first find isomorphic frontier states. With an arbitrary or inconsistent heuristic, first intersection of bidirectional frontiers does not guarantee a shortest path. A correct termination condition requires comparing the best complete path found with a lower bound from the open queues, together with an appropriate bidirectional heuristic construction. This is consistent with the paper’s own observation in Section 6 that Hybrid produces a path longer than the LCA path in 12.6 percent of cases. The output should therefore not be labeled “Shortest path” in the algorithm.

The role of top- k AGNN pruning is also unclear: Section 5.3 discusses filtering the available

mutations, while the later experimental discussion suggests that Hybrid keeps the full frontier. Hard filtering destroys completeness if the required action falls outside the retained set, so it cannot merely cause a graceful degradation to ordinary BFS. The authors should state the exact setting used for each row and provide an ablation. More fundamentally, they should distinguish (a) a complete shortest-path algorithm, (b) a complete but non-optimal path finder, and (c) an incomplete beam or top- k procedure. Every claimed success rate and path length should be reported in the appropriate category, and each returned sequence should be replayed transition by transition to certify its endpoint.

5. The anomalous out-of-distribution experiment is incompatible with the involution assumed by bidirectional search.

For an anomaly-free node, equation (2.1) makes the rank mutation in equation (2.2) involutive after the arrows reverse. This ceases to hold when the authors deliberately relax anomaly cancellation. Let

$$F_{\text{in}} = \sum_j N_j A_{jk}, \quad F_{\text{out}} = \sum_j N_j A_{kj}$$

with positive arrow multiplicities understood. The first mutation gives $N'_k = F_{\text{in}} - N_k$. Applying the same rule after reversing the arrows gives

$$N''_k = F_{\text{out}} - N'_k = N_k + F_{\text{out}} - F_{\text{in}}.$$

Thus $N''_k = N_k$ precisely when the balance condition holds. On the anomalous test set, a step grown backward from the target cannot in general be reversed into a valid forward step. The undirected mutation graph, the factor of $K - 1$ in the BFS estimate, and the bidirectional reconstruction used throughout Section 5 therefore do not apply.

The anomalous results should be removed or regenerated with a directed transition graph and a mathematically correct inverse operation. If the intended motivation is BPS-quiver mutation, the relevant left/right, charge-dependent mutation rules should be defined rather than identified with the anomaly-free Seiberg-duality rule. For every retained anomalous example, the authors should publish a forward replay of the reported sequence and verify every intermediate rank and adjacency matrix. Until this is done, the OOD success claims for anomalous quivers are not evidence for extrapolation.

6. The efficiency ratios do not support the headline comparison with deterministic search.

Equations (5.2)–(5.3) model BFS as expansion of a regular tree with an average branching factor and only immediate backtracking removed. A mutation graph has variable degree, cycles, branch mergers, repeated labeled states, and isomorphic representatives. The formula is therefore not a measured BFS workload on the problem being solved. Dividing by this exponential estimate produces the very large efficiency ratios in Table 2, but it does not show that the learned methods outperform an implemented deterministic algorithm by those factors. The comparison with the actual LCA search in the same table is much more modest: the quoted gains for Hybrid and Hybrid LCA are only about 1.10 and 1.14, respectively.

The revised paper should implement BFS or bidirectional BFS with the same validity checks, isomorphism convention, deduplication, and stopping criterion as the learned methods. A* with standard nonlearned lower bounds and a deterministic action ordering would also be useful. Report expansions, generated states, peak memory, wall-clock time, and model inference time

on identical hardware. The cost of graph-isomorphism calls and neural evaluation cannot be represented by visited-state counts alone. Any analytic branching estimate can remain as intuition, but it should not be the denominator for claims of empirical speedup. Equation (5.16) should also be rederived: at distance one its stated denominator reduces to one even though locating the correct neighbor generally requires candidate generation and evaluation.

7. The statistical support and reproducibility record are insufficient for the extrapolation claims.

The reported samples of 500 pairs do not come with cell counts across K and distance, confidence intervals, variability over data generation and training seeds, or uncertainty on success rates. Dataset sizes, duplicate rates, endpoint overlap across splits, optimizer settings, stopping criteria, model-selection rules, and the exact search budgets are not collected in one reproducible specification. These omissions are important because the DGNN has a mean absolute error of roughly 1.2, its errors are systematically shifted, and nonmonotonic predictions occur for a substantial fraction of paths; the AGNN likewise has appreciable inversion rates.

The breaking-point analysis in Section 7 further compresses the problem into $C = d_{\text{true}} \log_{10} K$ in equation (7.1). This is the logarithm of a crude K^d tree count, not a measured computational complexity, and it suppresses topology, rank and arrow magnitudes, quotient identifications, and branching variation. Equation (7.3) then extrapolates an AGNN breaking point near $C = 48$, far beyond the displayed test regime, where success has already become negligible. The claim that the combined methods operate 1.5–2 times beyond training should be replaced by estimates with uncertainty on genuinely held-out families. Multiple independent training runs, bootstrap intervals for search metrics, and explicit failure accounting are needed. Predictions far outside the measured range should be labeled as speculative rather than numerical results.

8. The novelty and significance should be stated more precisely.

The combination of a learned distance surrogate, a learned local action policy, and quiver-specific pathfinding is potentially useful. This appears to be the most interesting contribution. At present, however, broad statements about a first use of machine-learning pathfinders in theoretical physics obscure the relation to learned heuristic search, neural algorithmic reasoning, and the earlier quiver-learning work cited as reference [13]. The authors should specify what is new at the level of the state representation, loss, search algorithm, and physics application, and compare to a standard learned-policy-plus-search baseline. The physical significance will ultimately depend on resolving the state-space issue above: efficiently navigating a precisely defined quotient of ranked quivers would be a useful computational result, whereas a short path in an underspecified proxy graph does not yet imply a new statement about Seiberg-dual quantum field theories.

Minor comments

1. State the admissible domain of every mutation explicitly. Are zero or negative ranks discarded? Are disconnected quivers, self-arrows, and two-cycles allowed? The paper alternates between non-negative and positive ranks, which is physically material.
2. Define the average branching factor used in equations (5.2)–(5.3) operationally and give its measured distribution for every benchmark family. Removing only the inverse of the previous mutation does not account for other duplicate states.

3. Reconcile equation (5.15), which presents the gain as at least one, with footnote 29 and the reported values below one. If the reference length is only an upper bound, the symbol d_{true} should not be used for it.
4. Clarify that the Weisfeiler–Lehman hash used by the implementation is only a prefilter and that endpoint equality is confirmed by an exact, rank-aware directed-graph isomorphism test. A hash by itself is not a canonical representative.
5. Give a complete hyperparameter table for both networks and all pathfinders, including hidden dimensions, depth, positional-encoding dimension, optimizer, learning rate schedule, batch size, number of epochs, early stopping, loss weights, heuristic weights, beam width, top- k values, and search limits.
6. Report hardware and software versions, training time, inference time per state, and total wall time for every search comparison. Archive the exact code and data version corresponding to the paper. The repository README should also be checked: its displayed arXiv identifier does not appear to match this manuscript.
7. Several claims would be clearer if “mutation sequence” replaced “duality sequence” until physical equivalence has been established. Similarly, “mutation graph” should replace “mutation tree” whenever cross edges and repeated states are retained.
8. Please correct typographical errors including “generic cas,” “Hybric,” and “inerpreted,” and revise the malformed sentence in the discussion near page 46. A careful language pass is needed throughout the results section.
9. Some plots contain dense legends and small annotations. Increase label and legend sizes and ensure that line styles remain distinguishable in grayscale. Tables should state whether lower or higher values are favorable and define every abbreviation in the caption.
10. The conclusion should distinguish demonstrated results from proposed applications. The remarks about holographic throat length and field-theoretic computational complexity are interesting motivations, but no couplings, RG scale ordering, or physical operation cost is included in the present state representation.

Final assessment

The manuscript addresses a worthwhile computational question and contains ingredients that could form the basis of a useful contribution. Its present conclusions, however, are stronger than the definitions and experiments support. A publishable revision would need a unified quotient state space, exact labels on that state space, leakage-resistant splits, permutation- and exchange-consistent models, correct treatment of directed anomalous mutations, and empirical comparisons against implemented deterministic search. Because these changes affect essentially all central tables and figures, I regard them as requiring a new submission rather than a limited revision. My recommendation is therefore rejection in the present form, with encouragement to resubmit after the core analysis has been rebuilt and independently validated.