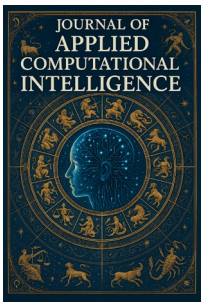
UNIVERSITY OF WEST ATTICA

eISSN: 3136-6937

Journal of Applied Computational Intelligence

Vol. 4, No. 2, April, 2026



GraphSAGE Model with Tether Topology Updates for Underwater Robot Tether Entanglement Prediction

Rawda Al-Rumaithi[1,*], Hana Al-Maqbali[1], Shahla Al-Otaibi[2], Tufaha Al-Sibairi[1,2]

[1] *College of Technological Innovation, Zayed University, Abu Dhabi, PO Box 144534, United Arab Emirates*

[2] *School of Information Technology, Skyline University College, Sharjah, PO Box 63000, United Arab Emirates*

Received 26 January 2026; Revised 20 May 2026; Accepted 27 May 2026; Published 31 May 2026

Abstract

Underwater robot tether geometry changes as material segments move around obstacles, and future geometric wrapping depends on both recent motion and spatial relationships. This study evaluates a GraphSAGE predictor with permanent material-chain edges and hysteretically updated proximity edges around a vertical cylindrical obstacle. Twenty-four ordered tether nodes carry nine features comprising position, recent velocity, radial clearance, material coordinate and discrete bending. Two neighborhood-aggregation layers and graph-level mean and maximum pooling estimate the probability of a geometric risk state two seconds after a six-frame observation history. The generated benchmark contains 1,500 independent trajectories, partitioned into 900 training, 300 validation and 300 test trajectories. A positive target combines sufficient angular sweep, radial proximity and bounded axial extent at the forecast horizon. Across three training seeds, the updated-topology model attains 94.11% accuracy, 72.65% positive-state recall, F1 0.7534 and Brier score 0.0447. Chain-only and frozen-proximity alternatives attain higher F1 values of 0.7945 and 0.8013, respectively, in this trajectory family. Velocity-omitted retraining yields F1 0.5447; zeroing velocity at inference in a fitted complete model yields F1 0.4863. Proximity connectivity changes over the observation interval in 35.47% of trajectories. Independent position perturbations and graph-delay tests expose sensitivity to descriptor construction and temporal alignment. These results identify predictive value in recent motion while limiting claims of benefit from topology refresh under the declared synthetic geometry. Measured tether motion, contact and tension data are needed to link the geometric label to mechanical entanglement.

Keywords: Tether Entanglement Prediction, Topology Updates, GraphSAGE, Underwater Robots, Geometric Risk Assessment

1. Introduction

Graph learning provides a computational representation of interacting spatial elements through node attributes and selected relationships [1], supporting prediction from structured observations. Geometric sequence analysis connects successive object configurations with their recent motion [2], enabling a distinction between persistent proximity and an approaching interaction. Adaptive connectivity allows a predictor to refresh the neighbourhood over which information is aggregated [3], accommodating changing spatial relationships. Reliability assessment links predictive scores with decision thresholds and observable errors [4], establishing how model outputs enter a warning process. These capabilities are relevant to tethered underwater robots operating near submerged infrastructure. A tether can approach an obstacle, wrap around its axis, or bring distant material sections into spatial proximity while the robot

* Corresponding author: rumatithi.r@zu.ac.ae

http://dx.doi.org/10.68406/jaci.2026.vol4iss2no10:137-149

➢ This is an open access article under the CC BY-NC-SA 4.0 license (https://creativecommons.org/licenses/by/4.0/).

© Authors retain all copyrights.

follows its inspection trajectory. Prediction therefore requires both an ordered description of the cable and a representation of relationships that change during movement.

Material orders and Spaces are for the same cable. Consecutive nodes are adjacent along the tether; their neighbours in the material can be far away. Dynamic Graph Representation Research [5] is employed to address the problem of neighbourhood inference. A fixed chain is continuous, but it cannot transmit the information of the wrapped cable section directly. A fully connected graph can freely communicate among the nodes in a cable but has little local interaction. Distance-based updates provide an indirect way to share features of close-up regions in materials directly. Geometric Prediction Methods [6] also carry the motion information of the current shape, as equal configurations can lead to different future states. Therefore, at present, the location of the cable needs to be determined in light of its recent relocation and the surrounding topography.

The value of a revised chart is compared with the models that have been modified. Controlled Model Comparison [7] keeps the source of features, trajectory partition and optimisation budget and prediction horizon fixed in the evaluation of connectivity. A pooled feature network can obtain the global shape statistics without explicit neighbour propagation. A chain graph can also obtain useful information by ordered local aggregation. A particular type of inductive bias for the modified proximity graph is determined by the target and trajectory populations. Probabilistic evaluation [8] also reports the ranking quality and score error of thresholded classification. Positive states are rare, and thus, the overall accuracy of warnings is not reliable. Therefore, the evaluation should show the recall, false alarm rate, probability error and variation of the training seeds, and distinguish between a future risk condition and a current-state observation.

This study evaluates a GraphSAGE predictor that combines permanent material-chain connectivity with spatial proximity edges refreshed over a six-frame observation history. Node descriptors align three-dimensional position, recent velocity, cylinder clearance, material order and local bending in an obstacle-centered frame. The two-second target is a defined conjunction of angular sweep, radial proximity and axial extent in a parameterized trajectory family. Matched updated, frozen, chain and pooled-feature models assess graph representation, while velocity-omitted retraining and fixed-model velocity zeroing assess different forms of motion dependence. Section 2 reviews tether geometry and dynamic graph prediction. Section 3 defines trajectory construction, node attributes, hysteretic topology refresh, neighborhood aggregation and probability and threshold fitting. Section 4 reports trajectory coverage, baseline ranking, node-feature effects, observation perturbations and emerging-state warnings. Section 5 presents three quantitative findings and the scope of this geometric risk model.

2 Related Work

2.1 Tether Geometry and Entanglement Representation

A discrete cable model is used to connect the spatial observations with the learned prediction. Geometric Representation Methods [9] aim to preserve the order of materials and can distinguish between a local bend and the proximity of separate cable sections. A Common Obstacle Frame can be used to uniformly express the clearance and direction of motion for all Positions. For an extended object, observations should also be made in the direction of the obstacle and perpendicular to it along the axis of the object. Different constraints for wrapping are available; a relatively large radial clearance can be used to obtain a large angle of rotation, but a close approach will only cover a small angle of rotation. Structured geometric learning [10] is employed to combine the above complementary descriptors explicitly. The selected obstacle shape, the space-time resolution of the cable, and when to set the future condition, all affect the results of interpretation.

Assessor of Entanglement: Geometric Proximity, Stable Loops and Mechanical Restraint. Task-specific labelling research [11] proposes to define the predicted condition according to the observable features with a particular computational meaning. A geometric wrapping label can be directly derived from the trajectory and is employed as a reproducible target for representation studies. Another is the force of contact or friction that opposes the movement of the robot in the direction of travel. The present purpose is to correct the reported geometric risk at one obstacle. It can still link the forecast back to the original item. Sequential monitoring methods [12] are also used to separate new problems from existing conditions at the forecast origin. The alert that occurs before the first qualified configuration is different from a forecast of the stability of the present situation. The same future-state predictor and a uniform time axis can be employed for observation.

2.2 Graph Learning for Dynamic Interaction Prediction

Neighborhood aggregation maps the local connections of variables to representations with shared node parameters. Inductive graph learning [13] is introduced to use node features and a neighbourhood operator instead of a separate trainable identity for each sampled cable location. A mean aggregator performs regular aggregation of adjacent states, and a separate self-feature path keeps the node's own observations. The two aggregation layers are connected differently in the two figures presented. Dynamic Connectivity Methods [14] change which pairs are involved in propagation

according to the change of geometry. As part of the input construction, the threshold and time of this update should be provided along with the learned network. The two switching states of hysteresis have no change in the presence of a small fluctuation near the switching point. The materials in the update will not change.

A readout function is employed to aggregate the local representations at the graph level for graph-level prediction. Research on feature pooling [15] aims to keep both the generalised evidence and local extreme responses. Mean pooling is the average of all encoded cable states; Max pooling retains the strong channel activations in different regions. Together, they can generate a scalar probability without the need for fixed semantic meanings of the hidden coordinates. The Quality of this probability is an empirical attribute of the trained predictor. Therefore, reliability analysis [16] is employed to examine the discrepancy of binary outcomes and the impact of validation-selected thresholds. Adjustment of the threshold changes recall without altering the ranking quality. Similarly, replacing a graph at inference measures the sensitivity of the fitted mapping, and training a different graph variant evaluates adaptation to a modified representation. The above are the answers to the previous questions.

3 Geometry Conditioned GraphSAGE for Tether State Prediction

3.1 Tether Discretisation and Interaction Encoding

The reference scene contains a vertical cylinder of radius 0.30 m and a tether discretized into 24 ordered material nodes. The obstacle-centered coordinate frame fixes the vertical axis and allows radial clearance to have one geometric meaning for every node. Material coordinate s ranges uniformly from zero to one and identifies the same cable location through time. Six observations at 0.20-s spacing span the one-second interval preceding the forecast origin. The node identities and material coordinates remain associated across those frames. For node i, horizontal position denotes the first two spatial components and R is cylinder radius in the clearance relation below. This obstacle-centered feature is evaluated from current node geometry.

$$c_i = \|p_{i,xy}\| - R \tag{1}$$

Each node has nine observed channels: three position components, three velocity components, cylinder clearance, material coordinate and local bending magnitude. Velocity is the backward difference between the last two observed positions divided by 0.20 s. Bending is the magnitude of the second spatial difference along material order, with zero endpoint values. Clearance measures horizontal distance from the cylinder surface in the obstacle frame. Means and standard deviations are fitted across training nodes and retained across validation and test partitions, with declared lower scale bounds for nearly constant channels. This feature construction distinguishes current shape, recent movement and cable ordering in a common material-node index.

$$x_i = [p_i; v_i; c_i; s_i; b_i] \tag{2}$$

Figure 1 shows the whole sensing-to-prediction process. Ordered tether geometry is in the descriptor branch and the connectivity branch, and they have the same obstacle frame and node IDs. The descriptor branch obtains the final position and recent motion, and the connectivity branch keeps track of cable continuity and updates the spatial proximity of the six observations. The two branches are connected at the graph encoder, and then a pooled probability readout and a validation-selected warning threshold are employed. Thus, the time origin of all the inputs is known. The endpoint at material coordinate one is the robot attachment in the geometric construction, and the other endpoint is the opposite end of the represented cable span.

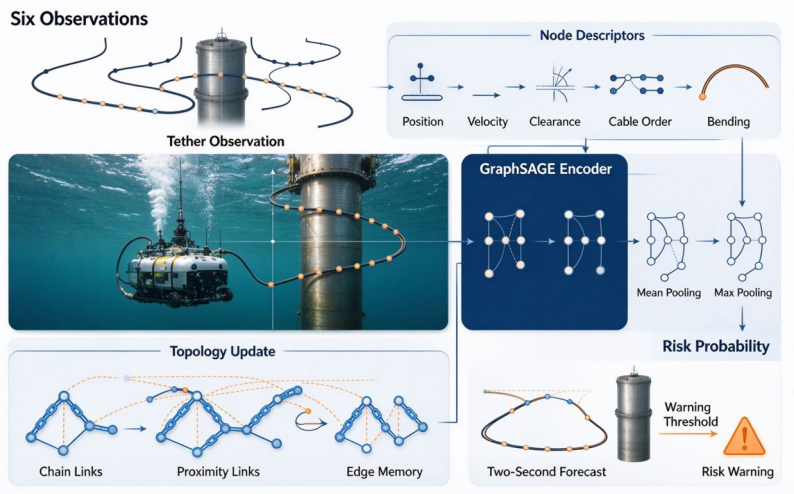


Figure 1. Observation encoding and topology updating framework for tether entanglement risk prediction.

Eight independently sampled construction variables define the kinematic trajectory family. Initial angular sweep is 2.5-6.6 rad, angular velocity is -0.5 to 0.9 rad/s and angular acceleration is -0.08 to 0.08 rad/s squared. Initial central radius is 0.44-1.25 m with radial velocity -0.12 to 0.12 m/s. Axial extent is 0.4-2.2 m, initial phase is -pi to pi and angular ripple amplitude is 0.05-0.25 rad. The forecast state at two seconds is positive when angular sweep reaches 5.4 rad, central radius is at most 0.85 m and axial extent is at most 1.6 m. This conjunction defines geometric wrapping risk for the generated experiment. It is evaluated from future trajectory geometry under the same cylinder frame used by node features.

3.2 Topology Refresh and Neighbourhood Aggregation

Permanent bidirectional chain edges join neighboring material indices. Additional undirected spatial edges are eligible for node pairs separated by at least three indices, capturing proximity between distant cable locations. On the first observed frame, an eligible link activates below 0.75 m Euclidean separation. On each subsequent frame, inactive links use the same entry threshold and active links persist until separation reaches 0.95 m. Current observed positions and prior edge membership determine the updated graph, retaining chain continuity throughout. At each GraphSAGE layer, the message for a node is the arithmetic mean of representations from its current neighbors. This rule combines material order with a temporally updated spatial interaction neighborhood.

$$m_i = mean(h_j : j \in N_i) \tag{3}$$

Each GraphSAGE layer concatenates the node self-representation with the mean of its current neighbors. A shared affine transformation and ReLU produce the updated hidden state. The first layer maps 18 concatenated feature values to 32 channels, and the second maps 64 values to 32 channels. A separate self path preserves the node descriptor as proximity degree changes. Mean aggregation normalizes by the number of neighbors, stabilizing the scale of propagated evidence across sparse and dense graphs. In the expression below, W and b are the layer weights and bias, and the superscript denotes layer output. The same learned parameters process all 24 material nodes.

$$h_i^+ = ReLU(W[h_i; m_i] + b) \tag{4}$$

Figure 2 is the connection diagram of edge refresh, neighbourhood aggregation and graph-level prediction. Permanent chain links are still in the observation window, and eligible spatial links are added or retained according to the hysteresis rule. Connect each modification to the position of the origin in the diagram and draw the final graph via the two aggregation layers. Add an independent self-feature path at each level via neighbor means. The mean and maximum readouts of the second layer are used to obtain an all-around view of all the cables. Therefore, the diagram distinguishes between the geometric update operation and the learned feature transformation, and both are connected to the same forecast origin.

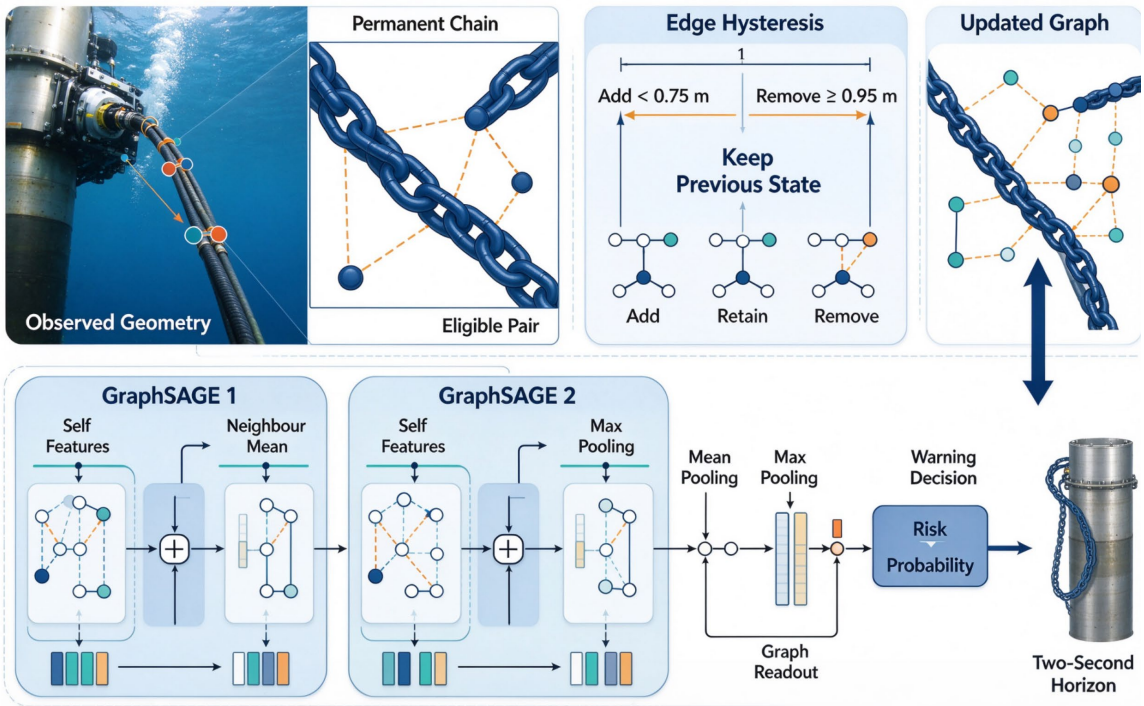


Figure 2. Distance hysteresis and GraphSAGE neighbourhood propagation for tether risk estimation.

The final node embeddings are pooled channel-wise by both mean and maximum. Mean pooling summarizes distributed cable geometry, whereas maximum pooling retains strong local responses. Each operation returns 32 values, giving a 64-value graph representation g independent of the edge count. All neighbors contribute to the exact mean in this 24-node implementation. Two aggregation layers define the node-level propagation range before the pooled graph readout. The combination of material and proximity edges can change which local states are exchanged, while pooling provides a common graph-level input across updated, frozen and chain variants.

$$g = [mean(h_i); max(h_i)] \tag{5}$$

3.3 Prediction Readout and Learning Objective

A 64-to-24 affine-ReLU layer and final scalar affine layer transform the pooled graph representation into a logit. Sigmoid activation yields probability p for the geometric risk conjunction at the two-second horizon. The readout function f includes both affine maps and their biases. Binary cross-entropy fits probability predictions on independent training trajectories. A separate threshold selected on validation data converts the continuous score into a warning. Thus F1, recall and false-alarm rate describe thresholded decisions, while Brier score evaluates probability error on the common held-out trajectories. This separation is important for comparing representations under a 13% positive prevalence.

$$p = \sigma(f(g)) \tag{6}$$

Seed 2820 generates 1,500 independent parameterized trajectories, and a seed-2821 permutation assigns 900 to training, 300 to validation and 300 to testing. All six observed frames and the two-second future target of one trajectory remain in its assigned partition. Feature means and scales are fitted only to training nodes. Lower scale bounds are 0.01 for position, clearance, material coordinate and bending channels and 0.10 for velocity channels. These fitted statistics apply unchanged to validation, test and matched graph baselines. Complete-trajectory separation and common normalization ensure that model comparisons assess the same geometry and label construction.

Three independent fitting runs use seeds 282, 283 and 284. AdamW has learning rate 0.002 and weight decay 0.0001. Each run evaluates up to 70 epochs with 15 shuffled minibatches of 60 trajectories per epoch and gradient-norm clipping at one. The checkpoint with the smallest validation binary cross-entropy is retained. A decision threshold is then chosen on validation predictions to maximize positive-class F1 over 0.10-0.90 in increments of 0.025; exact ties retain the lower threshold. Each selected checkpoint and threshold remain fixed for its test and inference-perturbation assessments. The updated GraphSAGE contains 4,273 learned parameters.

At each time, a smooth radial, angular and axial profile over material coordinate generates the 24 cable positions. Angular sweep evolves from initial sweep through velocity times time and half acceleration times squared time. Central radius follows initial radius and radial velocity, with a 0.34-m lower bound and a material-dependent radial term. Initial phase, current sweep and ripple define node angle; axial extent and a small sinusoid define node height. The same trajectory law supplies observed frames, future positions, radial clearance and the two-second risk criterion. Its explicit kinematic construction supports controlled variation of shape and motion while keeping the obstacle frame and cable ordering fixed.

4 Geometric Sensitivity and Entanglement Prediction Performance

4.1 Trajectory Coverage and Baseline Comparison

Figure 3 is the input geometry and the corresponding risk population. The three are angular progression, radial approach and vertical extension, and they are related to the present and the future differently. Controlled Distribution Design [17] is used to prevent the retention of independent variables in the evaluation of a shape-and-motion representation. The training, validation and test sets have 137, 36 and 39 positive future states respectively. Therefore, the proportion of this feature is 13.00%, and a general error has been found. The current-state positives in these partitions are 104, 30 and 29. Their differences from the future show both new developments and changes in the required movement over time.

Figure 3(a) is the distribution of angular sweep, central radius and axial extent, as well as some cable arrangements. The three amounts are the necessary components of the desired state. It may be outside the range of angles and far from the obstacle, or it may be close to the obstacle and high up. Coverage analysis [18] is used to find the joint distribution of classification performance. Independent Constructions can have different Angular and Radial Velocities for the same current shape. Thus, the predictor has a reason to employ motion evidence. As shown in the configuration, the obstacle axis and material order have been kept; therefore, a compact visual projection of the three-dimensional shape can be obtained.

Figure 3(b) is the future-state count and proximity-edge distribution of the partitions. The new graph has an average of 44.50 undirected proximity edges and 23 permanent chain edges, and the range of the proximity count is from 0 to 164. Changes in connectivity at the beginning and end of the observation interval for 35.47% of trajectories. Partition-aware evaluation [19] is used to save the entire history of a trajectory because neighbouring frames have similar geometry and parameters. The construction is based on independent Trajectories for division. Due to the change at the edge, graph renewal may affect propagation; however, whether the prediction is altered can be verified by the matched baseline experiments mentioned above.

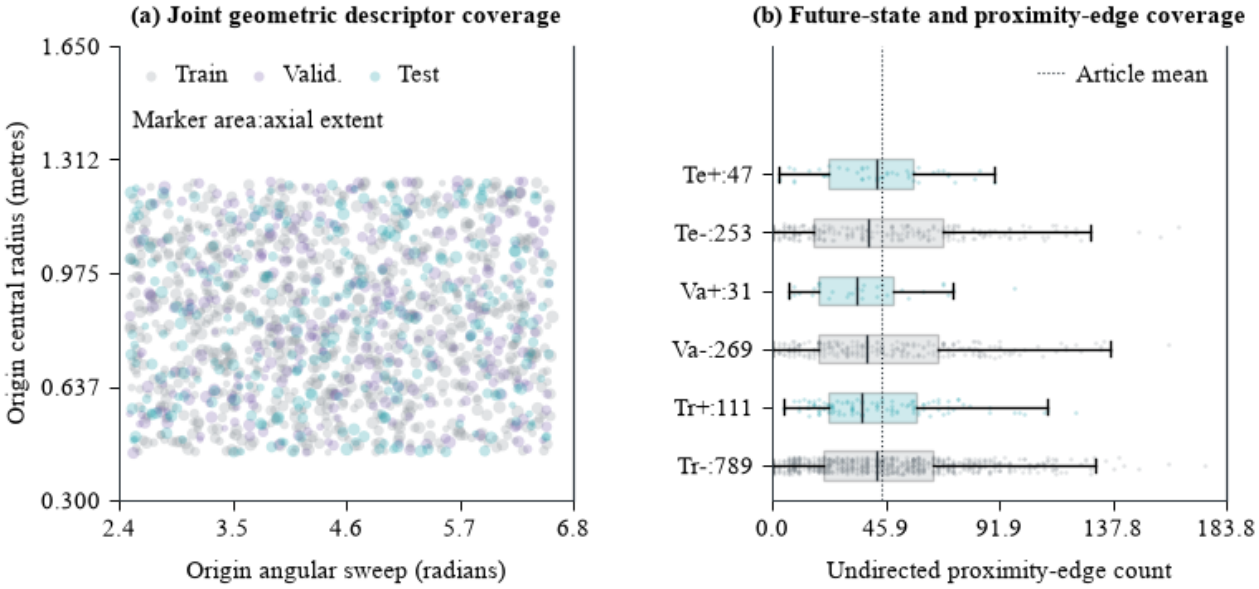


Figure 3. Trajectory geometry and evaluation coverage. (a) Geometric descriptors and representative configurations. (b) State distribution and proximity connectivity.

Figure 4 compares updated, frozen, chain and pooled-feature predictors under the same nine node channels and two-second target. Matched-input comparison [20] supports interpreting differences as representation outcomes under the declared training procedure. Figure 4(a) reports updated-graph accuracy 94.11%, positive recall 72.65%, false-alarm rate 2.68% and precision 0.8021 on the shared test population. Figure 4(b) adds F1 and ranking comparisons: mean F1 is 0.7534 for updated connectivity, 0.7945 for chain, 0.8013 for frozen proximity and 0.7659 for pooled features. The topology-updating model therefore yields a compliant risk predictor, while simpler tested structures attain higher F1 in this synthetic geometry. Both class-specific detection and ranking measures remain relevant to this comparison.

Figure 4(c) compares probability scores: mean Brier error is 0.0447 for updated connectivity, 0.0401 for chain and 0.0416 for frozen proximity. Lower values for the two simpler structures accompany their higher F1 rankings. Figure 4(d) shows run-specific thresholds and outcomes. Repeated-training assessment [21] motivates reporting thresholds 0.725, 0.525 and 0.350 alongside corresponding F1 values 0.6250, 0.8205 and 0.8148. With only 36 positive validation trajectories, decision selection varies among fits. The first run has a higher threshold and lower positive detection on the common test group. Continuous probability quality and thresholded warning outcomes therefore provide complementary views of the reported variation.

4.2 Contributions of Topology and Node Features

Table 1 shows the numerical comparison of the representation variant and inference intervention. The first five rows are the independently trained architectures or feature configurations, and the last three use the trained updated model and change its test input. Method of ablation [22] needs the above two comparisons. Retraining measures the achievable mapping in the alternative representation; inference intervention measures the dependence of a fitted predictor. Accuracy, positive-state recall, F1, false-alarm rate, and Brier score show different effects of the modifications. Average the three runs for all entries in the same 300 test trajectories. A positive prevalence and a fixed validation threshold for the interpretation of each row are available, as well as the conditions under which high accuracy corresponds to a reduction in the detection of future risk.

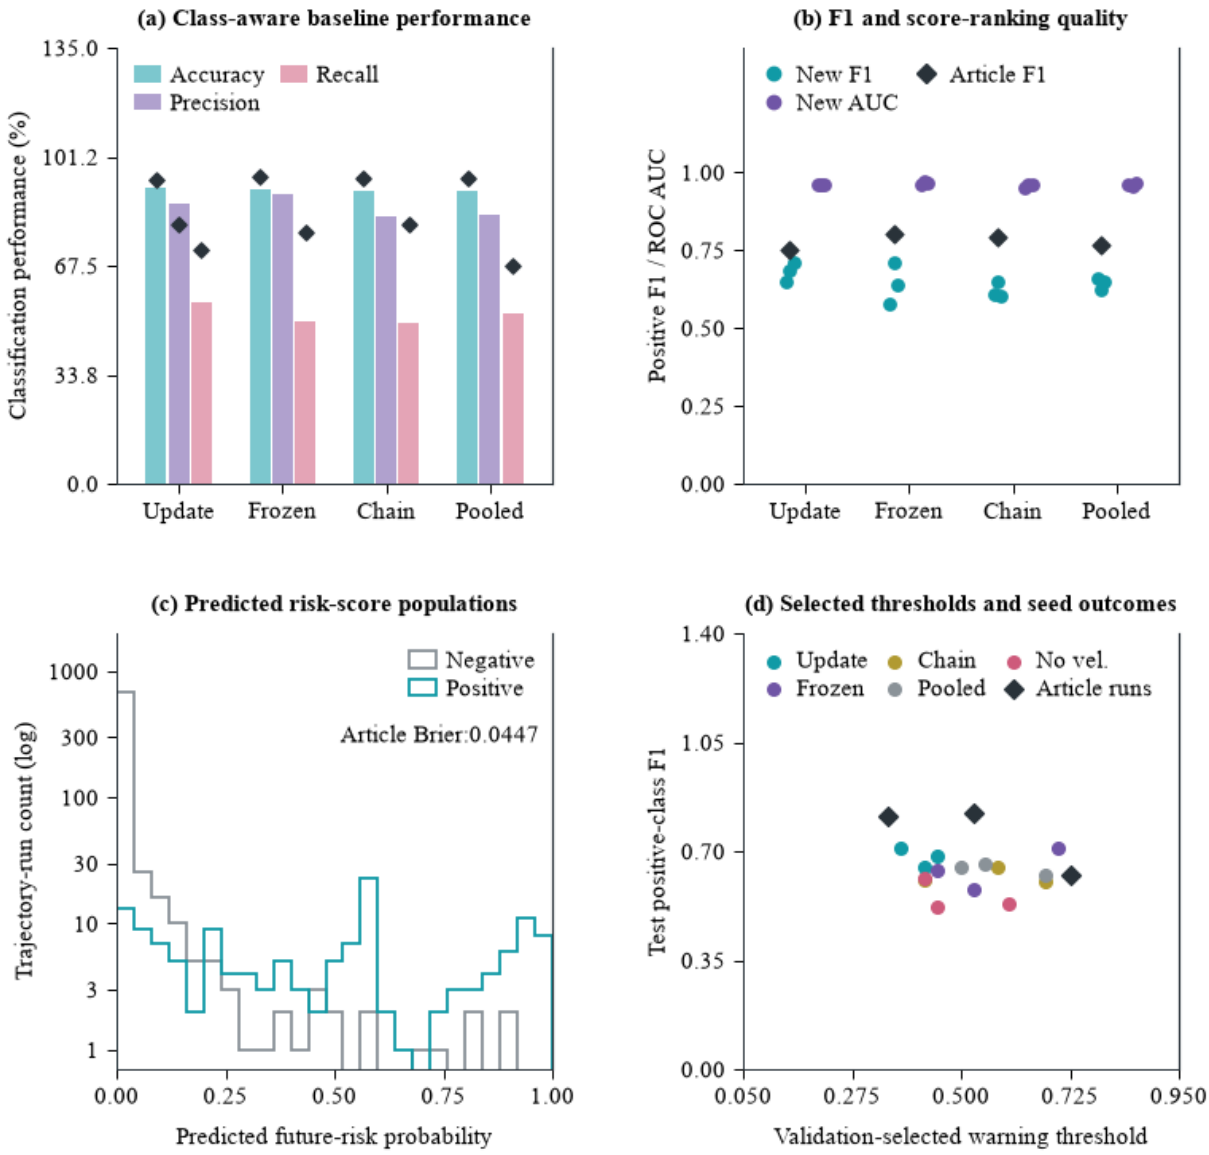


Figure 4. Baseline prediction and probability assessment. (a) Accuracy and class-specific performance. (b) F1 and ranking quality. (c) Probability error. (d) Training-seed and threshold variation.

Table 1. Representation comparisons and input sensitivity for future tether risk prediction.

Model / input setting	Accuracy (%)	Recall (%)	F1	False alarms (%)	Brier score
Pooled MLP	94.67	67.52	0.7659	1.28	0.0406
Chain GraphSAGE	94.56	80.34	0.7945	3.32	0.0401
Frozen GraphSAGE	95.00	77.78	0.8013	2.43	0.0416
Updated GraphSAGE	94.11	72.65	0.7534	2.68	0.0447
No velocity	89.89	47.01	0.5447	3.70	0.0729
Position noise	66.33	76.92	0.3723	35.25	0.2933
Delayed topology	94.11	72.65	0.7534	2.68	0.0439
Velocity dropout	89.67	38.46	0.4863	2.68	0.0876

Figure 5 is the connection and movement descriptors of the predictor. The new graph has changed the space of neighbourhoods; the old graph still has the original proximity connections, and the chain graph is based only on material continuity. Graph Representation Assessment [23] is motivated by the above structure comparisons. A pooled network is added to introduce an additional reference through shared-node transformations, and then the same mean-pooling and max-pooling schemes are employed. The reduced number of parameters is due to the omission of neighbour

concatenation. All the versions of the graph have 4,273 parameters and can compare updated and frozen connectivity at the same learned capacity. Thus, in combination with the topology-change rate and the geometric structure of the target, the performance pattern can be observed.

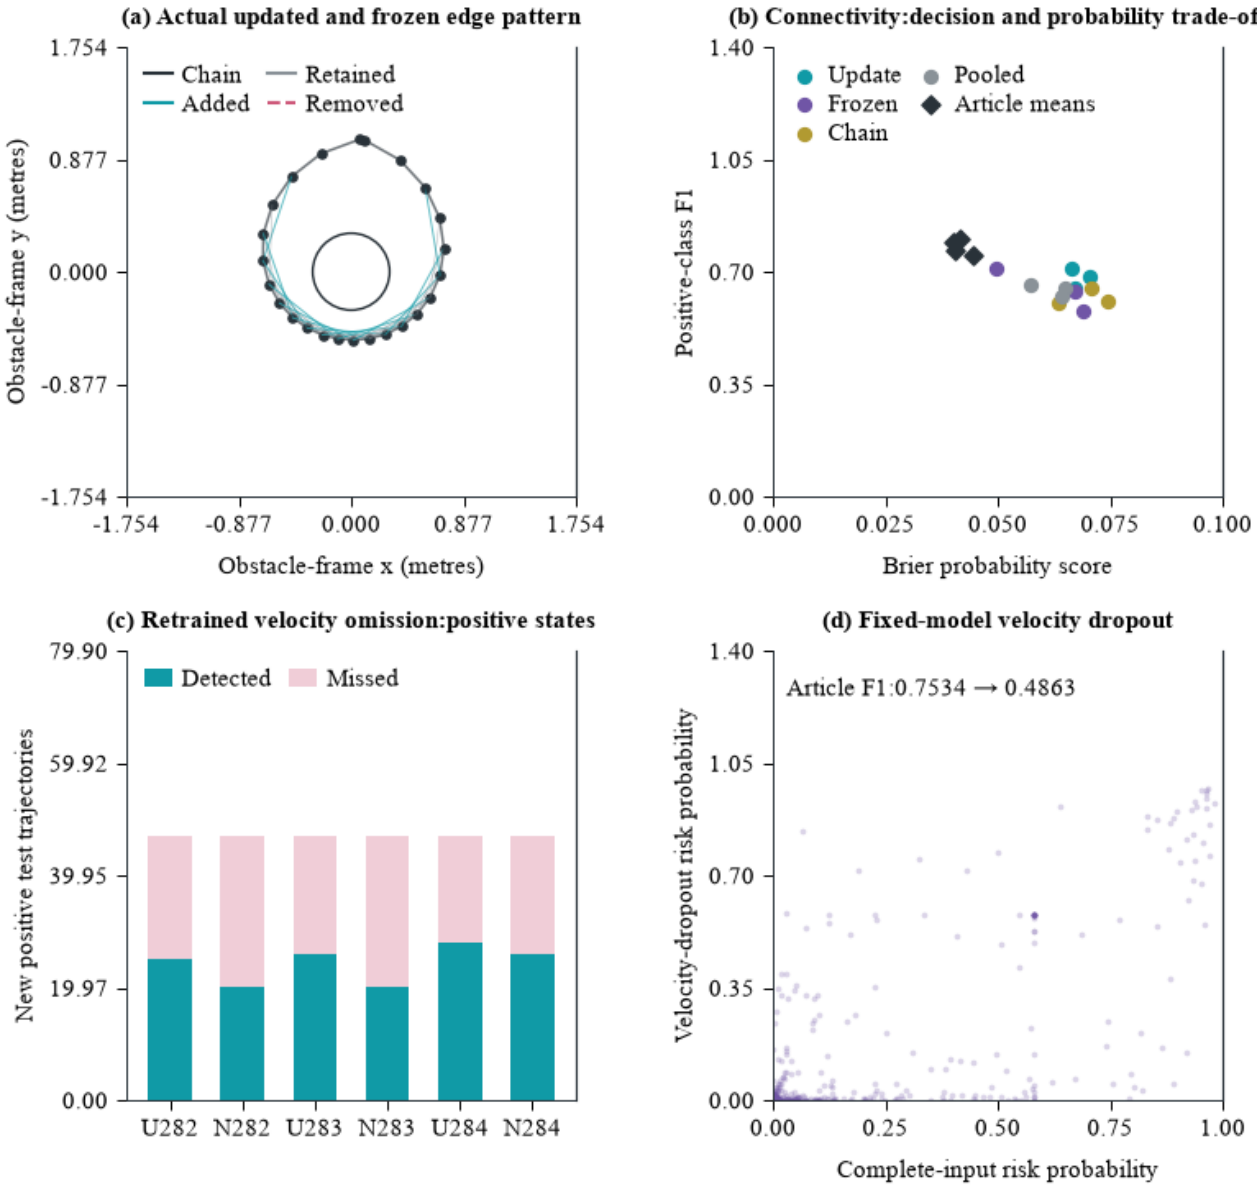


Figure 5. Contributions of connectivity and motion evidence. (a) Updated and frozen neighbourhoods. (b) Connectivity comparison. (c) Retraining without velocity. (d) Inference velocity removal.

Figure 5(a) is an updated and frozen edge pattern at a certain time of observation, along with modifications to the spatial connection and breaks in the chain. Compare the identical material nodes in the paired views to differentiate edge renewal from the actual movement of these nodes. Figure 5(b) is a graph of the variants and their F1 and Brier scores. Connectivity analysis [24] separates the change in the visible graph from the improvement of prediction. Updated connectivity in this benchmark is 0.7534 F1, and the frozen and chain alternatives are 0.8013 and 0.7945, respectively. The objectives of the world-sweep, radius and axial range are also available in node features and pooling. The above feature access provides a reason for the good performance of simpler graph models. The comparison shows that proximity renewal in the task-oriented role of the tested model is not always an improvement achieved by adding new changing connections.

Figure 5(c) evaluates a model retrained with the three velocity channels omitted while retaining position, clearance, material coordinate, bending and the updated graph. This model has 47.01% positive recall and F1 0.5447, compared with 72.65% and 0.7534 for the complete updated model. Feature sufficiency research [25] motivates examining recent motion when the target concerns a future geometric conjunction. Current shape describes present configuration, while the final-frame backward difference informs the direction and speed of subsequent change. The paired architecture and trajectory partition make this retraining result an estimate of performance from the remaining feature set under the same learning budget. The separate fixed-model intervention provides a complementary assessment.

Figure 5(d) separates velocity-omitted retraining from zeroing velocity at test time in the already fitted complete model. In the latter condition, weights and validation-selected warning thresholds remain fixed, and all three standardized velocity channels are set to zero. Intervention analysis [26] identifies this as dependence of the trained mapping on its expected motion input, including a changed feature combination. The fixed-model condition gives F1 0.4863, recall 38.46% and Brier score 0.0876. The retrained model can adapt its weights to remaining geometry channels and reaches F1 0.5447. Comparing both outcomes distinguishes available information after adaptation from sensitivity of a deployed predictor to altered input.

4.3 Perturbation Response and Warning Behaviour

Figure 6 assesses observation and graph-time perturbations using the retained updated-model checkpoints and validation thresholds. Observation robustness research [27] motivates applying position error before descriptor and graph construction. In Figure 6(a), independent zero-mean Gaussian noise with 0.03-m standard deviation is added to coordinates in all six observed frames. Position, backward-difference velocity, clearance, bending and hysteretic proximity edges are recomputed from those perturbed records. Mean F1 becomes 0.3723, recall 76.92%, Brier score 0.2933 and false-alarm rate 35.25%. Differencing independently perturbed positions over 0.20 s gives approximately 0.212 m/s coordinate-wise velocity-error standard deviation. This joint feature-and-connectivity disturbance explains why recall alone gives an incomplete account of the degraded predictions.

Figure 6(b) substitutes the proximity graph from the beginning of the one-second observation interval for the final updated graph, while current node features and fitted model weights remain fixed. The mean F1 remains 0.7534 and Brier score becomes 0.0439. Time-alignment assessment [28] distinguishes this test-time graph delay from the independently trained frozen-graph baseline, whose parameters adapt to its connectivity convention. Equal aggregate F1 can accompany a change in continuous probability error. The paired result therefore measures sensitivity of the retained predictor to graph staleness on the common trajectories, and the Brier score captures a difference beyond thresholded warning counts.

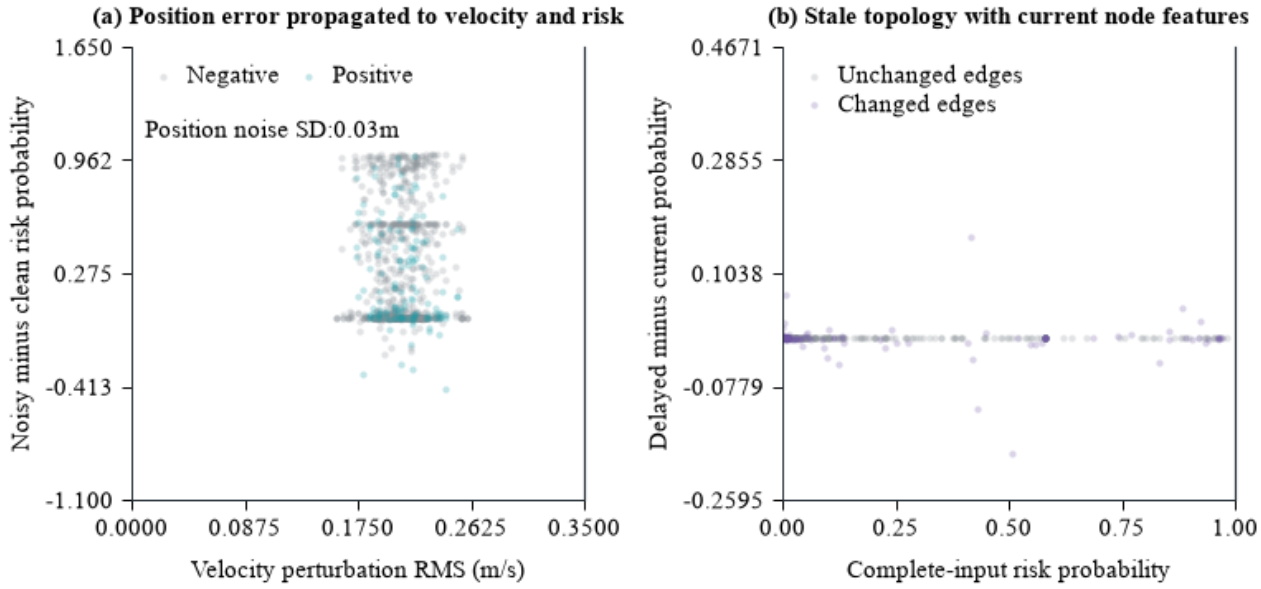


Figure 6. Response to observation and connectivity perturbations. (a) Position uncertainty propagated through the input pipeline. (b) Delayed proximity connectivity.

Figure 7 examines emerging geometric states and model workload. At the test origin, 20 trajectories are currently negative and become positive within the subsequent two-second interval. Event-oriented evaluation [29] treats this emerging subset separately from the full future-positive class. Figure 7(a) locates the first time the geometric conjunction is satisfied using 0.05-s spacing. Mean onset across these 20 trajectories is 1.225 s after the forecast origin. This quantity is an available geometric interval for advance assessment; achieved warning performance depends on the fitted probability and selected threshold. Reporting onset and detections separately preserves the distinction between potential lead and issued warnings.

Figure 7(b) reports emerging-state recall of 30.0%, 75.0% and 80.0% for the three training seeds, corresponding to six, fifteen and sixteen detections among the same 20 trajectories. These differences show substantial fitted-model and threshold variation for emerging warnings. Figure 7(c) places threshold choice in the complete test population, which has 39 future positives and 261 negatives. Warning evaluation methods [30] support reporting this denominator alongside recall and false alarms. Validation F1 selection balances these two decision outcomes under the original

threshold grid. Emerging-state recall thus describes one targeted subset, while overall F1 and probability error characterize the complete held-out population.

Figure 7(d) summarizes graph-processing and model dimensions. Two mean-aggregation layers each output 32 channels, and the graph-level classifier maps pooled 64 channels to 24 and then one output, for 4,273 learned parameters. Computational analysis [31] motivates reporting pair-count workload separately from parameter count. With 24 nodes, each frame has 276 unique pairs and six observations yield 1,656 possible pair distances before material-index exclusion. The reported mean graph has 23 permanent chain edges and 44.50 undirected proximity edges. Deployment-oriented evaluation [32] can retain node count, graph density, update interval and threshold identity with each prediction record. Actual processor latency remains an implementation-specific measurement.

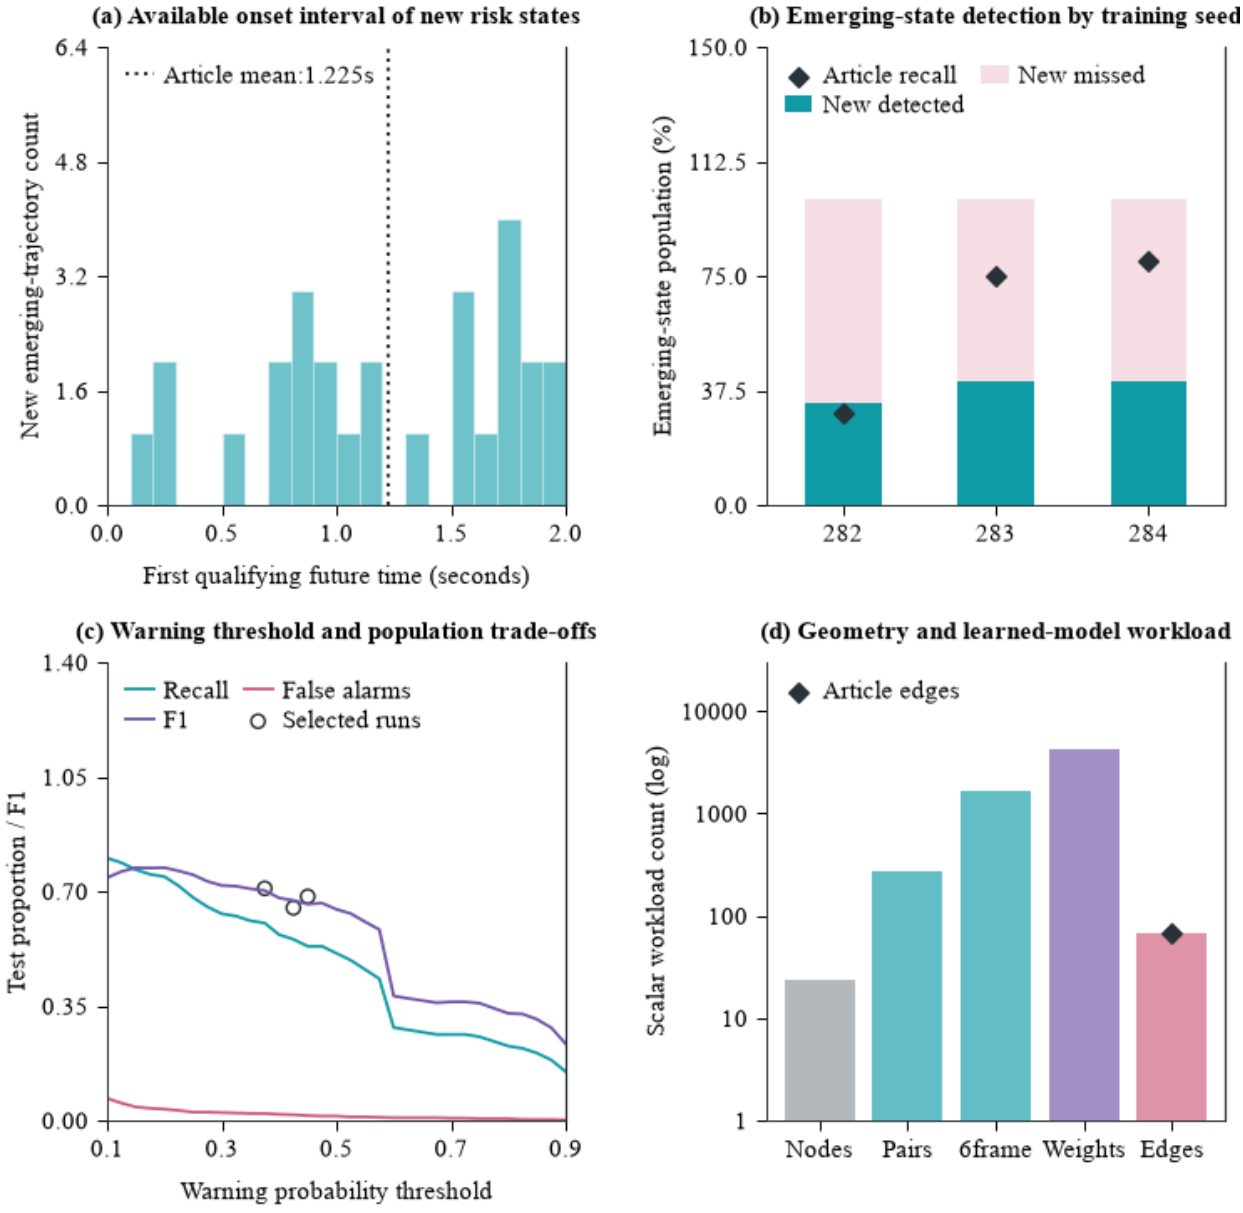


Figure 7. Warning behaviour and computational scale. (a) Onset time of emerging states. (b) Emerging-state recall. (c) Thresholded warnings and population context. (d) Model and graph workload.

5 Conclusion

This study evaluated topology-updating GraphSAGE for predicting a two-second geometric tether-wrapping state around a vertical cylinder. Twenty-four ordered material nodes carry position, recent velocity, clearance, order and bending features. Permanent chain edges preserve tether continuity, hysteretic proximity edges update spatial interaction and two neighborhood layers feed a pooled probability readout. Validation-selected thresholds convert continuous geometric-risk scores into warnings.

First, across 300 test trajectories the updated model attains accuracy 94.11%, positive recall 72.65%, F1 0.7534 and Brier score 0.0447; chain and frozen alternatives attain higher F1 values 0.7945 and 0.8013. Second, velocity-omitted

retraining lowers F1 to 0.5447, while fixed-model velocity zeroing gives 0.4863, distinguishing motion information from an input intervention. Third, 0.03-m position perturbation lowers F1 to 0.3723 with a 35.25% false-alarm rate; the 20 emerging-state trajectories have mean geometric onset 1.225 s and seed-specific recall 30.0%, 75.0% and 80.0%. These findings separate baseline ranking, feature dependence and warning reliability within the generated trajectory family.

The results apply to 1,500 generated trajectories around a fixed-radius cylinder under a prescribed angular, radial and axial wrapping criterion. The 35.47% of histories with changing proximity connectivity establish that graph refresh occurs in the population, while matched baseline scores define its observed predictive value. Position uncertainty propagates through velocity, clearance, bending and edges, and emerging-state warnings vary across training seeds. A physical extension should acquire synchronized cable geometry, robot maneuver history, obstacle contact, cable tension and friction measurements. Such data can test whether this geometric warning target predicts mechanical entanglement and supports calibrated operational decisions.

6. Declarations

6.1 Author Contributions

Rawda Al-Rumaithi contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. Hana Al-Maqbali contributes to validation, analysis, investigation, data collection, draft preparation. Shahla Al-Otaibi contributes to conceptualization, validation, analysis. Tufaha Al-Sibairi contributes to validation, analysis, investigation. All authors have read and agreed with the manuscript before its submission and publication.

6.2 Funding and Acknowledgments

This research received no specific financial support from any funding agency.

6.3 Institutional Review Board Statement

Not applicable.

6.4 Informed Consent Statement

Not applicable.

6.5 Date Availability

Not applicable.

6.6 Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

6.7 Declaration of generative AI and AI-assisted technologies

During the preparation of this manuscript, the authors utilized Grammarly for linguistic polishing, logical organization and draft revision. The generative AI tool was not involved in research design, data collection, data analysis, or the equation of core conclusions. All outputs generated by AI were reviewed, revised and verified manually by the authors. The authors take full responsibility for all content of this manuscript.

7. References

[1] Khoshraftar, S., & An, A. (2024). A survey on graph representation learning methods. ACM Transactions on Intelligent Systems and Technology, 15(1), Article 19, 1–55. https://doi.org/10.1145/3633518

[2] Han, J., Cen, J., Wu, L., Li, Z., Kong, X., Jiao, R., Yu, Z., Xu, T., Wu, F., Wang, Z., Xu, H., Wei, Z., Zhao, D., Liu, Y., Rong, Y., & Huang, W. (2025). A survey of geometric graph neural networks: Data structures, models and applications. Frontiers of Computer Science, 19(11), 1911375. https://doi.org/10.1007/s11704-025-41426-w

[3] Zheng, Y., Yi, L., & Wei, Z. (2025). A survey of dynamic graph neural networks. Frontiers of Computer Science, 19(6), 196323. https://doi.org/10.1007/s11704-024-3853-2

[4] Dong, J., Jiang, Z., Pan, D., Chen, Z., Guan, Q., Zhang, H., Gui, G., & Gui, W. (2025). A survey on confidence calibration of deep learning-based classification models under class imbalance data. IEEE Transactions on Neural Networks and Learning

Systems, 36(9), 15664–15684. https://doi.org/10.1109/TNNLS.2025.3565159

[5] Yang, L., Chatelain, C., & Adam, S. (2024). Dynamic graph representation learning with neural networks: A survey. IEEE Access, 12, 43460–43484. https://doi.org/10.1109/ACCESS.2024.3378111

[6] Mamedov, S., Geist, A. R., Swevers, J., & Trimpe, S. (2024). Pseudo-rigid body networks: Learning interpretable deformable object dynamics from partial observations. In 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 9542–9548). IEEE. https://doi.org/10.1109/IROS58592.2024.10802305

[7] Reuel, A., Hardy, A., Smith, C., Lamparth, M., Hardy, M., & Kochenderfer, M. J. (2024). BetterBench: Assessing AI benchmarks, uncovering issues, and establishing best practices. Advances in Neural Information Processing Systems, 37, 21763–21813. https://doi.org/10.52202/079017-0685

[8] Terven, J., Cordova-Esparza, D.-M., Romero-González, J.-A., Ramírez-Pedraza, A., & Chávez-Urbiola, E. A. (2025). A comprehensive survey of loss functions and metrics in deep learning. Artificial Intelligence Review, 58, 195. https://doi.org/10.1007/s10462-025-11198-7

[9] Mamedov, S., Geist, A. R., Viljoen, R., Trimpe, S., & Swevers, J. (2025). Learning deformable linear object dynamics from a single trajectory. IEEE Robotics and Automation Letters, 10(7), 7635–7642. https://doi.org/10.1109/LRA.2025.3577421

[10] Saranti, A., Pfeifer, B., Gollob, C., Stampfer, K., & Holzinger, A. (2024). From 3D point-cloud data to explainable geometric deep learning: State-of-the-art and future challenges. WIREs Data Mining and Knowledge Discovery, 14(6), e1554. https://doi.org/10.1002/widm.1554

[11] Hiniduma, K., Byna, S., & Bez, J. L. (2025). Data readiness for AI: A 360-degree survey. ACM Computing Surveys, 57(9), Article 219, 1–39. https://doi.org/10.1145/3722214

[12] Akasiadis, C., Kladis, E., Kamberi, P.-F., Michelioudakis, E., Alevizos, E., & Artikis, A. (2024). A framework to evaluate early time-series classification algorithms. In Proceedings of the 27th International Conference on Extending Database Technology (pp. 623–635). OpenProceedings.org. https://doi.org/10.48786/edbt.2024.54

[13] Ju, W., Fang, Z., Gu, Y., Liu, Z., Long, Q., Qiao, Z., Qin, Y., Shen, J., Sun, F., Xiao, Z., Yang, J., Yuan, J., Zhao, Y., Wang, Y., Luo, X., & Zhang, M. (2024). A comprehensive survey on deep graph representation learning. Neural Networks, 173, 106207. https://doi.org/10.1016/j.neunet.2024.106207

[14] Feng, Z., Wang, R., Wang, T., Song, M., Wu, S., & He, S. (2026). A comprehensive survey of dynamic graph neural networks: Models, frameworks, benchmarks, experiments and challenges. IEEE Transactions on Knowledge and Data Engineering, 38(1), 26–46. https://doi.org/10.1109/TKDE.2025.3621291

[15] Wang, Y., Hou, W., Sheng, N., Zhao, Z., Liu, J., Huang, L., & Wang, J. (2024). Graph pooling in graph neural networks: Methods and their applications in omics studies. Artificial Intelligence Review, 57, 294. https://doi.org/10.1007/s10462-024-10918-9

[16] Lane, R. O. (2026). A comprehensive review of classifier probability calibration metrics. Data Science, 9(1), Article 24518492261436231. https://doi.org/10.1177/24518492261436231

[17] Wang, A., Tong, T. P., Mizera, A., & Pang, J. (2024). Benchmarking structural inference methods for interacting dynamical systems with synthetic data. Advances in Neural Information Processing Systems, 37, 135129–135185. https://doi.org/10.52202/079017-4295

[18] Zhang, K., Liu, S., Wang, S., Shi, W., Chen, C., Li, P., Li, S., Li, J., & Ding, K. (2026). A survey of deep graph learning under distribution shifts: From graph out-of-distribution generalization to adaptation. ACM Transactions on Knowledge Discovery from Data, 20(2), Article 31, 1–38. https://doi.org/10.1145/3785475

[19] Semmelrock, H., Ross-Hellauer, T., Kopeinik, S., Theiler, D., Haberl, A., Thalmann, S., & Kowald, D. (2025). Reproducibility in machine-learning-based research: Overview, barriers, and drivers. AI Magazine, 46(2), e70002. https://doi.org/10.1002/aaai.70002

[20] Mathieu, T., & Preux, P. (2024). Statistical comparison in empirical computer science with minimal computation usage. In Proceedings of the 2nd ACM Conference on Reproducibility and Replicability (pp. 20–24). ACM. https://doi.org/10.1145/3641525.3663618

[21] Pecher, B., Srba, I., & Bielikova, M. (2024). A survey on stability of learning with limited labelled data and its sensitivity to the effects of randomness. ACM Computing Surveys, 57(1), Article 19, 1–40. https://doi.org/10.1145/3691339

[22] Chan, H. T. J., Šimić, I., & Veas, E. (2025). Explainability and importance estimate of time series classifier via embedded neural network. Scientific Reports, 15, 34515. https://doi.org/10.1038/s41598-025-17703-w

[23] Gravina, A., & Bacciu, D. (2024). Deep learning for dynamic graphs: Models and benchmarks. IEEE Transactions on Neural Networks and Learning Systems, 35(9), 11788–11801. https://doi.org/10.1109/TNNLS.2024.3379735

[24] Zhou, P., Yin, K., Zhu, H., Liang, Y., Li, L., & Chen, Y. (2026). A survey on graph structure learning. Neurocomputing, 696, 134114. https://doi.org/10.1016/j.neucom.2026.134114

[25] Klankers, M. K., & Steil, J. J. (2025). Modeling of deformable linear objects under incomplete state information. In 2025 IEEE International Conference on Robotics and Automation (ICRA) (pp. 10965–10972). IEEE. https://doi.org/10.1109/ICRA55743.2025.11127734

[26] Li, X., Wang, J., & Yan, Z. (2025). Can graph neural networks be adequately explained? A survey. ACM Computing Surveys, 57(5), Article 131, 1–36. https://doi.org/10.1145/3711122

[27] Dai, E., Zhao, T., Zhu, H., Xu, J., Guo, Z., Liu, H., Tang, J., & Wang, S. (2024). A comprehensive survey on trustworthy graph neural networks: Privacy, robustness, fairness, and explainability. Machine Intelligence Research, 21(6), 1011–1061. https://doi.org/10.1007/s11633-024-1510-8

[28] Zhang, W., Zhang, L., Han, J., Liu, H., Fu, Y., Zhou, J., Mei, Y., & Xiong, H. (2024). Irregular traffic time series forecasting based on asynchronous spatio-temporal graph convolutional networks. In Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (pp. 4302–4313). ACM. https://doi.org/10.1145/3637528.3671665

[29] Salles, R., Lima, J., Reis, M., Coutinho, R., Pacitti, E., Masseglia, F., Akbarinia, R., Chen, C., Garibaldi, J., Porto, F., & Ogasawara, E. (2024). SoftED: Metrics for soft evaluation of time series event detection. Computers & Industrial Engineering, 198, 110728. https://doi.org/10.1016/j.cie.2024.110728

[30] Ghorbani, R., Reinders, M. J. T., & Tax, D. M. J. (2024). PATE: Proximity-aware time series anomaly evaluation. In Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (pp. 872–883). ACM. https://doi.org/10.1145/3637528.3671971

[31] Kose, H. T., Nunez-Yanez, J., Piechocki, R., & Pope, J. (2024). A survey of computationally efficient graph neural networks for reconfigurable systems. Information, 15(7), 377. https://doi.org/10.3390/info15070377

[32] Lu, H., Wang, L., Ma, X., Cheng, J., & Zhou, M. (2025). A survey of graph neural networks and their industrial applications. Neurocomputing, 614, 128761. https://doi.org/10.1016/j.neucom.2024.128761