Journal of *Marine Science and Engineering*

MDPI

Article

An Adaptive Generalization Method for Complex Shoreline Segments in Nautical Charts Based on Bend Sequence Attention Learning

Haoshi Song [1], Shuaidong Jia [1,*], Xianpeng Liu [1], Zikang Song [2] and Qiqian Sun [1]

[1] PLA Dalian Naval Academy, Dalian 116018, China
[2] Chart Information Center, Tianjin 300450, China
* Correspondence: sky_jsd@163.com

Abstract

To address the challenges posed by significant variations in bend morphology and spatial relationships within complex shoreline segments—where single generalization operators and uniform parameters struggle to balance node compression with morphological preservation—this paper proposes an adaptive shoreline generalization method based on bend sequence attention learning. Treating bend units as local decision objects, the method integrates geometric morphology, proximity distance, and shoreline type features. It employs bidirectional long short-term memory (BiLSTM) and self-attention mechanisms to learn contextual relationships among bends within each shoreline and predict local generalization operators and their parameters. Proximity links between shorelines support spatial consistency checks during constraint-based correction. After operator execution, candidate geometries that violate cartographic constraints undergo parameter rollback or operator reselection; unresolved bends are restored to their original geometry before output. Under the evaluated datasets, scoring procedure, and compression settings, the final generalized shorelines achieved a vertex-reduction ratio of 86.43% and a bend-retention ratio of 94.84%. These results support the feasibility of bend-level adaptive operator selection at the evaluated compression–preservation operating point. Performance across other compression settings and chart series remains to be established.

Keywords: shoreline generalization; nautical cartography; bend unit; self-attention; adaptive operator selection

Check for updates

Academic Editor: Eugen Rusu

Received: 7 September 2026
Revised: 18 October 2026
Accepted: 19 October 2026
Published: 22 October 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

The shoreline is the geometric basis for representing hydrographic and terrestrial elements on nautical charts and for nearshore spatial cognition, and its representation governs the spatial relationships of islands, reefs, harbors, estuaries and navigation channels [1,2]. As chart scale decreases, shoreline generalization must reduce geometric detail while preserving the principal morphology and the shoreline features that support navigational positioning [2]. In nautical cartography, displacement direction is asymmetric. Moving a generalized shoreline landward decreases the charted land area and increases the charted water area, which may overstate the clearance available to a vessel. Moving it seaward increases the charted land area and decreases the charted water area, which is generally the conservative direction. These two directions must therefore be recorded and constrained separately. Additional constraints are still required at harbor entrances,

narrow channels and relations with adjacent elements [3]. It is precisely in estuaries, narrow channels, island-adjacent areas and lagoons, where bends are dense, shorelines lie close together and local enclosures form, that independent processing of single shorelines or the use of uniform parameters can produce morphological distortion and spatial conflicts such as excessive bend deletion, shoreline transgression, island conglomeration and harbor closure [4–6].

Research on linear feature generalization has produced several technical pathways, including point selection, bend processing and spatial relationship constraints [7]. The Douglas–Peucker algorithm [8] controls node retention through a perpendicular-distance criterion, the Visvalingam–Whyatt algorithm [9] deletes nodes progressively according to effective area, and hybrid line simplification [10] and sinuosity-based segmentation [11] provide ways of configuring processing strategies for different local morphologies. In the case of shorelines, studies on the preservation of drowned-valley geomorphological features [12], on the hierarchical expression of estuarine skeletons [5] and on bend skeleton-line indicators [3] have integrated bend structure and geographic characteristics into the simplification process, and Zhang et al. [6] used a hierarchical triangulation network to organize the morphology and proximity relationships of multiple shorelines for collaborative simplification. These studies extend shoreline generalization from the node level to the bend level and to multi-shoreline relationships, and they show that a generalization strategy has to adapt to local morphology and to proximity.

Machine learning and deep learning have opened a further route by learning morphological features and generalization strategies directly from samples [13–15]. Part of the existing work focuses on recognition and representation learning: Gao et al. [16] applied a sequential deep learning model to coastline shape classification, Yu and Chen [17] used a stacked autoencoder to learn polyline bend features and to generate simplified representations, and Chen et al. [18] combined vertex and edge modeling in a graph network for the simplification of roads and shorelines. A parallel line of work predicts generalized geometry directly from imagery or from geometry. Generative adversarial networks have been used to derive generalized mountain road maps [19] and multi-scale electronic map tiles under cartographic requirements [20], whereas graph-based models reformulate simplification as node movement or node selection over a graph [21–23]. Unsupervised autoencoder approaches have also been applied to river network simplification [24]. These methods show that geometric detail can be reproduced from data, but they generate geometry rather than choosing among the generalization operators already implemented in chart production. For tasks that require an existing operator to be invoked, recognizing a morphological category is not enough; the operator and its parameters have to be selected for the local unit. Bends along one shoreline are sequentially correlated, and processing one bend changes the combined morphology of adjacent bends. The question addressed here is therefore narrower: how can the correlation between the bends of a sequence be used to learn local operator and parameter configurations? The aim is not to redraw the line with a network. Recent reviews place this question within the wider agenda of GeoAI for multi-scale cartography and note that integrated, knowledge-aware generalization remains an open problem [25].

To address these gaps, this paper proposes an adaptive generalization method for complex shoreline segments based on bend sequence attention learning. Three points distinguish it from the work reviewed above. First, whereas existing data-driven approaches decide at the level of the whole line, of a fixed tolerance or of generated geometry, the decision here is taken for each individual bend, while the output is kept inside conventional, specification-compliant operators. Second, the sequential dependency between neighboring bends is used for operator and parameter selection: a bidirectional long short-term

memory network and a self-attention layer allow non-adjacent bends to take part in a local decision. Third, proximity between different shorelines is represented explicitly rather than only through a global buffer: links are established between bends whose inter-shoreline separation falls below a scale-dependent threshold, the minimum distance enters the model as a feature, and the links are available as a consistency check during rule-based correction; on the present dataset, however, the measured effect of the link-aware correction was small (Section 3.5), so this element is presented as a design feature rather than as a demonstrated performance gain. On this basis, the method predicts a local generalization operator and its parameter for every bend and then corrects the resulting geometry against cartographic constraints. The practical significance of the formulation is that every operator selection can be traced to a specific bend, so that an unacceptable local result can be rolled back or reselected on its own without reprocessing the whole shoreline; this keeps each decision auditable and locally editable inside the production workflow.

2. Data and Methodology

2.1. Underlying Principles

Shoreline generalization aims to reduce the detail of a source shoreline at a specified target scale while preserving its principal morphology and necessary spatial relationships under cartographic constraints. Bends at estuarine bifurcations, narrow-channel constrictions, island-adjacency zones, and lagoon openings differ in scale, orientation, and structural role. Consequently, simplification with a uniform tolerance cannot readily balance local shape preservation against overall vertex reduction. In the proposed approach, generalization decisions are therefore made at the bend-unit level. Sequence context is used to predict the local generalization method and its parameters, accommodating the requirements of different coastal segments [12,26].

The retention or removal of individual bends changes their combined shape along a shoreline and may also affect the separation between neighboring shorelines. To account for these relationships, bend feature learning is combined with conventional operator execution. A generalization strategy is predicted for each bend, and the resulting geometry is subsequently checked against cartographic constraints. The overall workflow is shown in Figure 1.

The model inputs comprise geometric descriptors, proximity distances, and coastal-segment types for individual bend units. Contextual information along the bend sequence is extracted using bidirectional sequence encoding and self-attention. The operator class and local parameter are predicted by classification and regression branches, respectively. Each selected operator is then applied to the corresponding bend using the predicted parameter, and candidate results that violate cartographic constraints are corrected.

2.2. Shoreline Sample Construction and Bend-Unit Feature Representation

2.2.1. Sample Preprocessing and Bend-Unit Segmentation

The sample data were taken from shoreline features on multi-scale nautical charts covering Chinese waters. Shorelines were identified using the shoreline feature identifier defined in GB 12320-2022 [1], consistent with the IHO chart specifications [2]. Chart sheets were selected so as to cover the five coastal-segment scenarios of interest, and the chart number, scale and edition of each source sheet were recorded; the resulting chart list will be made available through the controlled-access route described in the Data Availability Statement. The segmentation was performed as follows. A seven-point, second-order Savitzky–Golay filter was first applied to smooth the coordinate sequence, and the discrete curvature was then estimated by central differences. A bend unit is the arc between two successive inflection points of the smoothed curve, that is, between two successive

sign changes in the curvature, and the curvature extremum inside the arc is the apex of that bend. When the arc-length separation between two adjacent curvature extrema was smaller than δ, where δ is the scale-dependent displacement tolerance defined in Section 2.2.2, only the extremum of larger absolute curvature was retained as an apex, and the two arcs were treated as a single bend. Units whose bend height was below δ were merged with an adjacent bend of the same orientation or, where the orientations differed, with the neighboring unit of greater chord height. The geometry of a bend unit and the rule that determines its extent are illustrated in Figures 2 and 3; the smoothing step and the screening of adjacent curvature peaks are illustrated in Supplementary Figures S1 and S2.

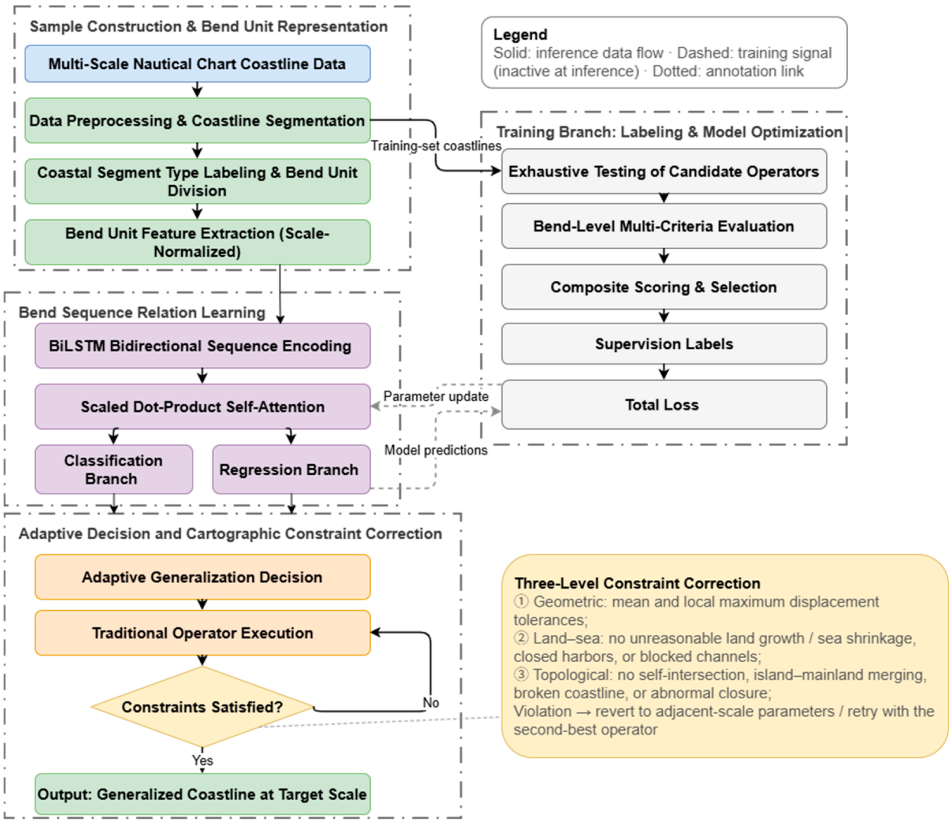


Figure 1. Workflow for adaptive shoreline generalization in complex coastal segments using bend sequence attention learning.

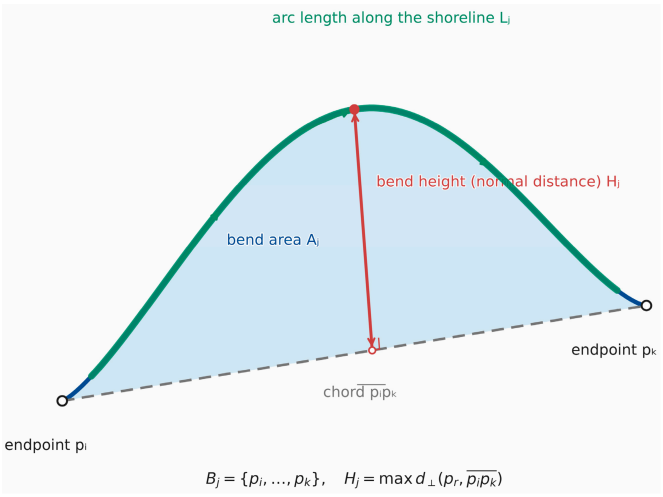


Figure 2. Geometry of a bend unit.

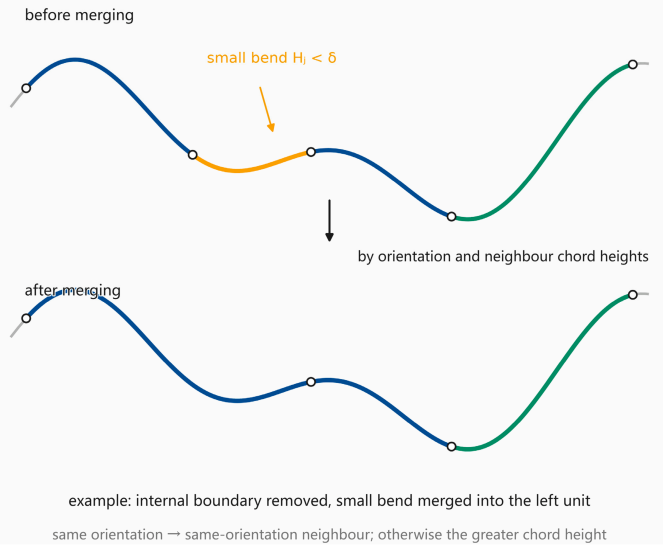


Figure 3. Merging of low chord-height bends.

A bend unit is the arc between two successive inflection points of the smoothed shoreline; p_i and p_k are its endpoints; the dashed line is the chord joining them; H_j is the bend height, measured as the maximum perpendicular distance from the arc to the chord; A_j is the area enclosed by the arc and its chord; and L_j is the arc length accumulated along the shoreline.

Above: before merging, the small bend whose chord height H_j falls below the tolerance δ is identified. Below: after merging, the internal boundary is deleted, and the small bend is combined with a neighboring unit—with the neighbor of the same orientation where one exists, and otherwise with the neighboring unit of the greater chord height. The shoreline coordinates themselves are not altered by the merge.

Samples were organized into five coastal-segment types, referred to below as scenarios, according to their dominant morphological characteristics and spatial relationships: estuarine bifurcation, narrow-channel constriction, island adjacency, semi-enclosed lagoon, and ordinary sinuous shoreline. Coastal-segment types were assigned by nautical-chart compilation experts through visual interpretation of shoreline morphology and geographic context. The resulting labels were encoded as the feature component to incorporate scenario information into the bend representation. Coastal-segment type attributes are also required during inference.

Each source shoreline is represented as a sequence of discrete coordinate points and partitioned into ordered bend units according to local bend morphology. The sampled point sequence is denoted by P_O, as defined in Equation (1), where p_i is the ith sampled point and n is the total number of points. The resulting sequence of bend units is denoted by B, as defined in Equation (2).

$$P_O = \{p_i = (x_i, y_i)\}_{i=1}^{n} \quad (1)$$

$$B = \{b_j\}_{j=1}^{m} \quad (2)$$

In Equation (2), b_j denotes the jth bend unit, which is defined jointly by its endpoints, internal shape-defining points, and adjacent bends; m is the total number of bend units.

2.2.2. Bend-Unit Feature Extraction

Bend structures provide organizational units for analyzing the morphology of continuous curves [3,26]. Numerical features are used to describe the geometry and spatial

proximity of each bend unit. After scale normalization, the feature vector f_j for the jth bend unit is expressed as in Equation (3).

$$\boldsymbol{f}_j = [L_j, H_j, A_j, \kappa_j, \Delta\kappa_j, R_j, D_j, \tau_j]^{\mathrm{T}} \quad (3)$$

Here, L_j is the arc length of bend unit j along the shoreline polyline; H_j is the bend height, defined as the maximum perpendicular distance from the bend curve to the chord connecting its endpoints; and A_j is the bend area enclosed by the curve and its chord. The mean curvature of the discrete points within the bend is denoted by κ_j, whereas $\Delta\kappa_j$ represents curvature dispersion or the magnitude of curvature variation. The concave or convex orientation of the bend is denoted by R_j and can be determined from the sign of the cross product of adjacent line segments. The minimum Euclidean distance from the bend to neighboring shorelines is denoted by D_j, and the coastal-segment type code is denoted by τ_j.

Each shoreline forms an independent bend sequence, within which sequence encoding and self-attention are computed. Spatial relationships between different shorelines are represented through distance features and proximity links. The minimum distance from each bend to neighboring shorelines is included as a model input. A proximity link is established when the distance between shorelines falls below a specified threshold, allowing their spatial relationships to be coordinated during generalization. Thus, attention aggregates bend context within a sequence, whereas distance-based links represent proximity between different shorelines.

Before feature concatenation, the length unit was $L_{unit} = S \times 0.01$ m, the ground length corresponding to 1 cm at target-scale denominator S. Arc length, bend height and proximity distance were divided by L_{unit}; bend area was divided by L_{unit}^2; mean curvature and curvature variation were multiplied by L_{unit}; R_j and τ_j were used directly. The concave/convex orientation R_j enters as a sign (+1 or −1), and the coastal-segment type code τ_j as a scalar index over the five types (0–4). This integer encoding introduces an artificial ordering among the five shoreline types, which is a limitation of the current feature representation. Section 3.5 reports the ablation in which this feature is removed. Operator parameters were normalized separately to [0, 1] using each operator's candidate bounds and inverse-transformed after inference. No other unit system was used.

A proximity link was created when the minimum Euclidean distance between two bends on different shorelines was below 2δ, with $\delta = 0.0003S$ m; sensitivity checks used 1δ, 2δ and 4δ. The reproducible subset produced 5165, 5195 and 5204 linked pairs at 1δ, 2δ and 4δ, respectively, but only 0–1 conflicting pairs, and link-aware correction changed CR/LR by no more than 0.03 percentage points. The links are therefore used as a minimum-distance feature and as an optional consistency check during correction. The ablation results for the proximity links are reported in Section 3.5.

2.3. Bend Sequence Attention Learning Model

As defined in Equations (2) and (3), a shoreline is represented by an ordered sequence of bend features, $\boldsymbol{F} = [\boldsymbol{f}_1, \boldsymbol{f}_2, \cdots, \boldsymbol{f}_m]$. Bend units follow a strict order along the shoreline, and a generalization decision for any unit is typically influenced by neighboring units. To model these sequential dependencies, a network combining bidirectional sequence encoding, self-attention, and multitask prediction is constructed.

The feature sequence $\boldsymbol{F}$ is encoded in both shoreline directions using a bidirectional long short-term memory network (BiLSTM) [27,28]. Context from both directions is incorporated into the hidden-state sequence $\boldsymbol{H}$, as expressed in Equation (4).

$$\boldsymbol{H} = \mathrm{BiLSTM}(\boldsymbol{F}) \quad (4)$$

Here, $H = [h1,h2,\cdots,h_m]$ is the hidden-state sequence, and h_j is the contextual representation of bend unit j. Relationships between bends are captured by propagating information along the sequence through the BiLSTM. To further aggregate information from nonadjacent bends, scaled dot-product self-attention is introduced. The hidden-state sequence H is linearly projected into query, key, and value matrices:

$$Q = HW_Q, K = HW_K, V = HW_V \tag{5}$$

Here, Q, K, and V are the query, key, and value matrices, respectively, and W_Q, W_K, and W_V are the corresponding learnable parameter matrices. Information is aggregated across bends using scaled dot-product attention:

$$\boldsymbol{A} = \mathrm{softmax}\left(\frac{\boldsymbol{Q}\boldsymbol{K}^{\mathrm{T}}}{\sqrt{d}}\right), \boldsymbol{Z} = \boldsymbol{A}\boldsymbol{V} \tag{6}$$

Here, A is the attention weight matrix, and a_{ij} is the relative weight assigned to information from bend j when constructing the representation of bend i. The key-vector dimension is denoted by d, and Z is the representation incorporating context from the current input sequence. Explicit weighted aggregation establishes information links between nonadjacent bends [29].

Following the shared-representation principle of multitask learning [30], the contextual representation z_j is supplied to both the classification and regression branches. The generalization strategy and local parameter are thereby predicted jointly. The classification output is given by

$$\hat{\boldsymbol{y}}_j = \mathrm{softmax}\left(\boldsymbol{W}_c\boldsymbol{z}_j + \boldsymbol{b}_c\right) \tag{7}$$

Here, $\hat{\boldsymbol{y}}_j$ is the algorithm-class probability vector for the jth bend unit, and K_c is the number of candidate algorithm classes. The candidates comprise Douglas–Peucker (DP) [8], Delaunay triangulation (DT) [31,32], double-buffer (DB), and monotone-chain identification (MC) algorithms. The learnable parameters of the classification branch are denoted by $\boldsymbol{W}_c$ and $\boldsymbol{b}_c$.

The local generalization parameter is predicted by the regression branch as

$$\boldsymbol{\rho}_j = \boldsymbol{W}_r\boldsymbol{z}_j + \boldsymbol{b}_r \tag{8}$$

Here, $\boldsymbol{\rho}_j$ is the predicted parameter value, and $\boldsymbol{W}_r$ and $\boldsymbol{b}_r$ are the learnable parameters of the regression branch. The contextual representation z_j is shared by both branches, so operator selection and parameter prediction are based on the same structural information about the coastal segment.

2.4. Adaptive Generalization Decisions and Correction Under Cartographic Constraints

The local generalization strategy is determined jointly by the predicted class and parameter. For bend unit j, the class with the highest predicted probability is selected:

$$u_j = arg_k max\hat{y}_{j,k} \tag{9}$$

Here, u_j is the generalization operator class for bend unit j. Before operator execution, the predicted parameter is clipped to its lower and upper bounds:

$$\theta_j = \mathrm{clip}(\rho_j,\theta_{\min},\theta_{\max}) \tag{10}$$

Here, θ_j is the local generalization parameter after range restriction; $\theta_{\min}$ and $\theta_{\max}$ are the lower and upper parameter bounds for the selected operator, respectively. Because

parameter meanings and units differ among operators, clipping is performed within the range defined for the corresponding operator.

The operator class u_j and local parameter θ_j jointly define the generalization strategy for bend unit j. During execution, the conventional geometric operator corresponding to the predicted class is applied to the bend using the clipped parameter. Different bends can therefore be processed using different operators, and local geometric changes can be traced to specific operators and parameter values.

Topological consistency in line simplification and the preservation of shoreline buffer relationships have been investigated previously [33–35]. Candidate results generated by the operators are checked against cartographic constraints. These checks cover mean and local maximum displacement relative to the source shoreline, as well as excessive land-area increase (equivalently, water-area decrease), land-area decrease (equivalently, water-area increase), harbor closure, and narrow-channel blockage. Shoreline self-intersections, island–mainland merging, shoreline discontinuities, and abnormal closure relationships are also checked. Violations are corrected by parameter rollback or operator reselection. Bends for which local correction does not resolve a hard-constraint violation are handled by the final fallback described below.

Checks were applied in the fixed order: geometry, area and topology. With $\delta = 0.0003S$ m, mean displacement was limited to δ and local maximum displacement to 2δ. Signed landward shoreline displacement, which decreases charted land area and increases charted water area, was limited to 0.5δ. For each bend, land-area increase (water-area decrease) was limited to $0.50\delta Lc$, whereas land-area decrease (water-area increase) was limited to $0.25\delta Lc$, where Lc is chord length. Self-intersection, island–mainland merging, discontinuity and abnormal closure were hard failures. A protected channel or harbor entrance had to retain at least 80% of its source minimum width and at least 2δ. Local correction stopped when all hard checks passed, the violation count ceased to decrease, or five iterations were reached. In the latter two cases, the violations that remained at that point were recorded as unresolved events, and the affected bends were restored to their original geometry in the final fallback. The event totals reported for the corrected schemes in Sections 3.4 and 3.5 refer to this state, that is, after local correction but before the fallback. After all bends are assembled into the output shoreline, the complete generalized shoreline is re-checked against the four topological conditions. This final validation operates on the assembled geometry rather than on individual bend candidates, and any bend flagged at this stage is likewise restored to its source geometry. Following any restoration at this stage, the shoreline was reassembled once more and re-validated. For the proposed method and the corrected baselines, all geometric metrics and expert-assessment scores were calculated from this final geometry, and intermediate post-operator or pre-fallback geometries were not used; for the uncorrected baselines, the operator output itself is the final geometry, as indicated in the note to the comparison table in Section 3.4.

2.5. Construction of Training Labels and Model Optimization

Supervisory labels are determined by selecting preferred candidate generalization results and conducting manual quality checks. For each bend unit, candidate results are generated using DP, DT, DB, and MC over their respective sets of candidate parameter values. Candidates with topological errors are excluded first. Those whose local maximum displacement exceeds the tolerance are also excluded from scoring. Eligible candidates are evaluated using the vertex-reduction ratio, length retention ratio, and bend-area retention ratio. These metrics are min-max normalized and combined into a weighted composite score. Selection proceeds in two stages. First, the highest-scoring parameter value is selected for each algorithm; the preferred results are then compared across algorithms to

obtain paired algorithm-class and parameter labels. The same normalization reference and weights are used at both stages. Samples with similar scores or ambiguous cartographic acceptability are reviewed manually before the final labels are assigned.

The baseline composite score used equal weights: $w_v = 1/3$, $w_l = 1/3$ and $w_a = 1/3$. In the local geometric reproduction, 4906 bends were evaluated under 231 weight vectors of the score simplex, taken at steps of 0.05. The agreement between the parameter label selected under each of those weight vectors and the label selected under equal weights ranged from 0.8% to 95.6% (mean 43.4%), while the mean selected vertex reduction ranged from 15.03% to 69.11%. In five retraining runs with ± 0.1 perturbations around equal weights, classification accuracy changed by at most 0.62 percentage points, bend retention by 1.15 percentage points and the mean run-level total of topological events recorded after local correction but before final fallback by 1.80.

Parameters were normalized separately by operator as $\theta_{norm} = (\theta - \theta_{min})/(\theta_{max} - \theta_{min})$. A shared scalar regressor was used, followed by class-specific inverse normalization. Parameter loss was computed for every bend with a paired parameter label, not only for correctly classified bends. MSE was first computed within each operator and then pooled across all validation bends, weighting each operator by its validation support; the coefficient of determination was computed on the pooled validation set. Per-operator MSE and RMSE are reported in Section 3.3 and Supplementary Table S1.

On the reproducible DP candidate subset, exhaustive per-bend evaluation of 31,413 bends over the six DP tolerance factors, that is 188,478 candidate results, required 6.27 s, or 0.20 ms per bend, on Windows 11 with a 12-logical-core AMD64 processor, Python 3.12.5 and NumPy 2.5.0. The per-bend value is an amortized cost, obtained by dividing the cost of a segment by its bend count, and is not comparable with the per-chart figure in Appendix A, which additionally includes the segmentation, feature extraction, proximity linking and candidate generation of every bend on the chart.

The classification and regression branches are supervised by the algorithm-class labels and their paired parameter labels, respectively. Stratified sampling and class weighting are used to reduce the effects of sample imbalance. Early stopping is applied on the basis of validation loss.

The overall training loss is defined as

$$\mathcal{L} = \lambda_c \mathcal{L}_{\text{cls}} + \lambda_r \mathcal{L}_{\text{param}} \tag{11}$$

Here, $\mathcal{L}_{\text{cls}}$ is the cross-entropy loss for algorithm classification, quantifying the discrepancy between the predicted probability vector $\hat{\boldsymbol{y}}_j$ and the ground-truth class label y_j^*. The parameter regression loss, $\mathcal{L}_{\text{param}}$, is the mean squared error between the predicted parameter ρ_j and the ground-truth parameter label ρ_j*. The two loss terms are weighted by λ_c and λ_r, respectively. Geometric conditions, land–sea relationships, and topological conditions are excluded from gradient computation. Instead, they are used to check constraints and correct results after the predicted strategies have been executed by the conventional operators.

2.6. Inference Procedure

During inference, the input shoreline is first partitioned into bend units, and their features are extracted. Input sequences and shoreline proximity links are organized as described in Section 2.2. The operator class and local parameter for each bend are obtained by a forward pass through the model. The bend's position on the source shoreline determines the operator's spatial scope, while the target scale determines the applicable cartographic representation requirements. Each selected operator is then applied using its predicted parameter. Cartographic constraints are subsequently checked, and the results

are corrected, with the proximity links used as a consistency check, to produce the output shoreline. Bends with unresolved hard-constraint violations are restored to their original geometry before output. This final fallback is distinct from parameter rollback during local correction.

2.7. Implementation Details

Algorithm 1 summarizes the training-label construction, model training and inference procedure. Appendix A collects the settings used to reproduce the experiments. Table A1 lists the candidate parameters, constraint thresholds and correction rules; Table A2 lists the network configuration, the five fixed seeds and the timing scope; Table A3 lists the frequencies of the correction branches.

Algorithm 1. Training label construction, model training and inference

Input: source shorelines at scale S_s, target scale S_t, operator set O = {DP, DT, DB, MC}, parameter candidate sets $\{P_o\}$.
Step 1. Segment each shoreline into ordered bend units (Section 2.2.1) and extract the feature vector of Equation (3) with the scale normalization of Section 2.2.2.
Step 2. Build proximity links between bend units whose minimum distance falls below the proximity threshold.
Step 3. For each bend and operator o, generate candidates over $\{P_o\}$, reject candidates violating the cartographic constraints, and score the remainder with $w_v = w_l = w_a = 1/3$; select the best parameter within each operator and then the best operator.
Step 4. Train the BiLSTM plus self-attention model of Section 2.3 on the paired class and parameter labels, minimizing the loss of Equation (11) with early stopping on validation loss.
Step 5. At inference, forward-pass each bend sequence, take the arg-max class and the clipped regression parameter, execute the corresponding operator, and apply the constraint checks and correction branches of Section 2.4 with reference to the proximity links. Record the unresolved topological violations after local correction, restore bends with unresolved hard-constraint violations to their original geometry, reassemble the shoreline, and apply final topology validation. The validated final shoreline is the output from which all geometric metrics and expert-assessment scores are calculated.

3. Experiments and Analysis

The method was evaluated in terms of both strategy prediction and final shoreline quality. Classification accuracy and parameter regression metrics were used to assess agreement with the supervisory labels. Geometric metrics and stage-specific topological error counts were used to assess the workflow. For methods with correction, topological errors were counted after local correction, but before final fallback; for methods without correction, they were counted directly after operator execution. DP, DT, DB, and MC were used as baselines. Model ablations and representative coastal-segment cases were also examined to compare the different processing schemes. For the proposed method and the corrected baselines, all geometric and area metrics and expert-assessment scores were calculated from the final shorelines obtained after fallback replacement, reassembly and final topology validation, and intermediate post-operator or pre-fallback geometries were not used for these evaluations; the uncorrected baselines were evaluated at the operator-output stage, as indicated in the note to the comparison table in Section 3.4.

3.1. Experimental Data and Sample Construction

The experimental data were derived from shorelines delineated on nautical charts of Chinese waters at multiple scales. Shoreline geometries were extracted after the chart projection, reference ellipsoid, central parallel, and other geodetic parameters had been harmonized, and were transformed into a common planar coordinate system. To avoid restricting the dataset to a single chart or shoreline morphology, samples were stratified by coastal-segment type. Types were assigned by nautical-chart compilation experts through visual interpretation of shoreline morphology and geographic context.

The dataset comprised 2571 coastal segments extracted from multiple charts, covering estuarine bifurcation, narrow-channel constriction, island adjacency, semi-enclosed lagoon, and ordinary sinuous shoreline scenarios. These segments were partitioned into 104,619 bend units using local curvature changes, extrema, and scale-dependent thresholds. A feature vector was extracted for each bend unit according to Equation (3). The source shoreline geometry, target scale, candidate generalization results, and algorithm-class and parameter labels obtained through candidate evaluation and manual quality checks were retained.

The annotation protocol used three experienced nautical-chart compilers, who independently assigned the coastal-segment type. Pre-discussion agreement was Fleiss' $\kappa = 0.86$, and 190 samples (7.39%) required joint discussion, after which consensus labels were formed.

The chart-level split was 70%/15%/15%. Training, validation and test sets contained 1798, 388 and 385 segments and 73,377, 15,495 and 15,747 bend units, respectively. The 90 source charts were allocated as 63/14/13 while stratifying by scenario, with no chart crossing sets.

Sample sizes differed among the five coastal-segment types listed in Table 1. Ordinary sinuous shorelines were the most numerous, with 983 segments accounting for 38.23% of the total. Coastal-segment type annotation and its use as an input are described in Section 2.2.1.

Table 1. Composition of the sample dataset.

Coastal-Segment Type	No. of Segments	No. of Source Charts	No. of Bend Units	Mean Bends per Segment	Mean Shoreline Length (m)
Estuarine bifurcation	398	69	18,665	46.90	15,555.8
Narrow-channel constriction	389	81	14,989	38.53	12,118.7
Island adjacency	425	73	21,740	51.15	17,655.5
Semi-enclosed lagoon	376	86	16,489	43.85	13,915.3
Ordinary sinuous shoreline	983	90	32,736	33.30	10,784.6
Total	2571	90	104,619	40.69	13,318.7

The numbers of source charts in the five categories sum to more than the 90 distinct charts of the dataset, because one chart may contain shorelines that belong to several scenarios. The column therefore counts chart-scenario combinations, while the total row reports the number of distinct charts.

3.2. Analysis of Training-Label Construction

The nautical-chart scales included 1:30,000, 1:50,000, 1:250,000, 1:500,000, 1:750,000, and 1:1,000,000. Larger-scale shorelines were used as source data for the candidate operators. Target scales were taken from smaller-scale charts covering the corresponding areas. Smaller-scale shorelines were used for preliminary pairing and reference screening, whereas the supervisory labels were generated using the candidate evaluation procedure in Section 2.5. Generalization quality was assessed by comparing the generalized results with the source shorelines. A displacement tolerance of 0.03 cm on the chart was used,

corresponding to a ground tolerance of the target-scale denominator multiplied by 0.0003 m. For DB, the buffer half-width served as the generalization parameter. Candidate parameter settings were based on a chart half-width of 0.015 cm, corresponding to a reference ground half-width of the target-scale denominator multiplied by 0.00015 m.

When labels were generated according to Section 2.5, each algorithm yielded candidates satisfying the constraints in this experiment. To reflect both simplification intensity and shape preservation, the three indicators were combined with the equal weights $w_v = w_l = w_a = 1/3$ defined in Section 2.5. The distributions of preferred algorithms and manual label correction rates for the different scenarios are reported in Table 2.

Table 2. Distribution of preferred algorithms during label construction by coastal-segment type.

Coastal-Segment Type	DP Preferred (%)	DT Preferred (%)	DB Preferred (%)	MC Preferred (%)	Manual Correction (%)
Estuarine bifurcation	33.47	30.33	6.99	29.22	8.06
Narrow-channel constriction	33.76	2.62	33.37	30.24	5.82
Island adjacency	36.01	24.58	8.71	30.70	8.93
Semi-enclosed lagoon	29.47	2.54	39.39	28.61	7.31
Ordinary sinuous shoreline	54.92	0.36	1.35	43.37	3.47

As shown in Table 2, DP was preferred for 54.92% of ordinary sinuous shoreline samples, whereas DB was preferred for 39.39% of semi-enclosed lagoon samples. These different distributions indicate that operator selection varied with the coastal scenario in the candidate comparison, providing a sample-based rationale for configuring strategies from local features.

The operator-class support counts were obtained directly from the final label file: DP 41,974; DT 11,934; DB 15,137; MC 35,574, giving 104,619 bend units in total. The class counts are 29,439/8370/10,616/24,952 in the training set, 6217/1768/2242/5268 in the validation set and 6318/1796/2279/5354 in the test set. The confusion matrices for the validation and test sets are shown in Figure 4 (rows: true class; columns: predicted class), where diagonal cells are correct predictions and off-diagonal cells are misclassifications.

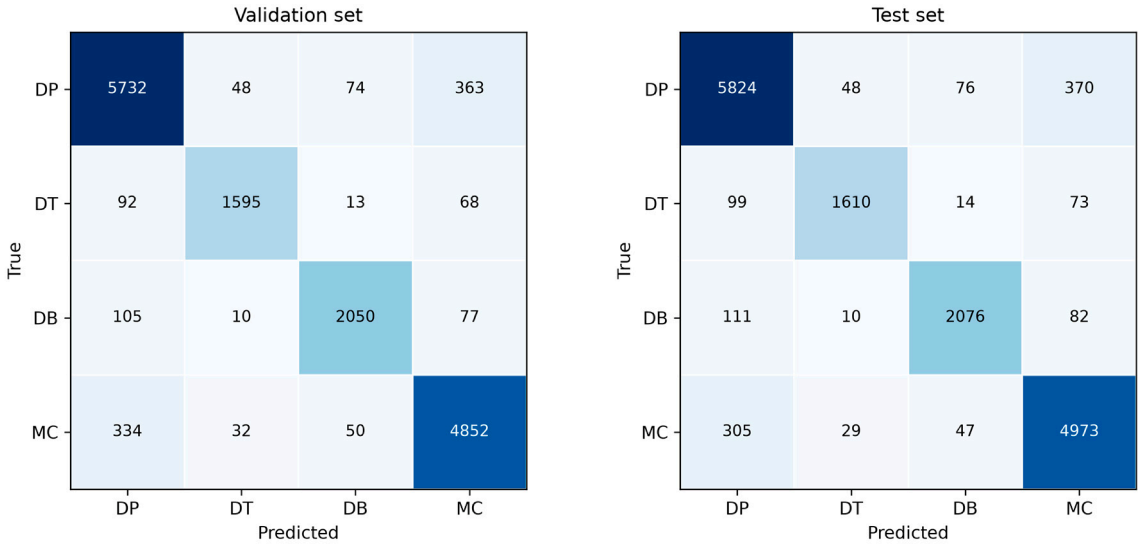


Figure 4. Confusion matrices of operator-class prediction for the validation and test sets.

3.3. Model Training and Validation

The data were stratified by coastal-segment type and split into training, validation, and test sets at the chart level. This reduced the risk of inflated performance estimates

caused by similar samples from the same chart appearing in different sets. The model was optimized according to Equation (11) using Adam with an initial learning rate of 10^{-3}. Sequences within each batch were dynamically padded according to their lengths. Training was stopped when validation loss failed to decrease for 15 consecutive epochs. Validation results are reported in this section; test results are presented in Sections 3.4 and 3.5.

During training, the classification loss and parameter regression loss decreased rapidly over the first six epochs and subsequently approached convergence. Validation classification accuracy was 91.83 ± 0.23%, and the macro-averaged F1 score was 92.03 ± 0.28%. Table 3 reports the per-scenario results on the validation set. For the proposed method and the learned ablations, the ± values denote the sample standard deviation over five independent training runs; for the single-operator baselines, whose output is deterministic for a given parameter, they denote the dispersion of the metric over the evaluation set, obtained by resampling that set.

Table 3. Model performance on the validation set.

Metric	Overall	Estuarine Bifurcation	Narrow-Channel Constriction	Island Adjacency	Semi-Enclosed Lagoon	Ordinary Sinuous Shoreline
Classification accuracy (%)	91.83	90.06	92.84	89.45	91.16	94.23
Macro-averaged F1 (%)	92.03	90.09	92.86	89.41	90.83	94.23
DP recall (%)	92.20	90.08	93.23	89.52	91.22	94.26
DT recall (%)	90.21	90.53	92.92	89.19	91.90	96.92
DB recall (%)	91.44	89.20	92.66	88.37	90.93	95.37
MC recall (%)	92.10	89.76	92.62	89.77	91.32	94.13
Parameter regression MSE (pooled, normalized)	0.0089	0.0059	0.0073	0.0083	0.0098	0.0112
Parameter regression R^2 (pooled)	0.893	0.929	0.910	0.901	0.885	0.864

Note: overall validation classification accuracy is 91.83 ± 0.23% and the macro-averaged F1 score is 92.03 ± 0.28% (mean ± SD over five runs). The overall parameter-regression MSE is the pooled (sample-weighted) value over all 15,495 validation bends; it is not an unweighted mean of the four per-operator MSEs (unweighted mean: 0.0077). The overall R^2 is computed on the pooled validation set. The word "pooled" in the row headings refers to the Overall column only. Scenario-level MSE and R^2 are computed within each scenario and are reported for descriptive purposes; they are not used in the pooled overall calculation. Per-operator MSE and RMSE are reported in Supplementary Table S1.

Validation classification accuracy ranged from 89.45% to 94.23% across scenarios (Table 3). DT recall was 96.92% for ordinary sinuous shorelines, and DB recall was 88.37% for island adjacency. For parameter regression, MSE was computed within each operator class and then pooled across all 15,495 validation bends, weighting each operator by its validation support; the pooled MSE was 0.0089, and the coefficient of determination on the pooled validation set was 0.893. Because the four classes differ in support, this pooled value is not equal to the unweighted mean of the per-operator MSEs (0.0077). Per-operator errors are reported in Supplementary Table S1. The normalized MSE was 0.0139 for DP, 0.0069 for DT, 0.0045 for DB and 0.0055 for MC; the corresponding RMSE values were 0.118, 0.083, 0.067 and 0.074 in normalized units, i.e., 11.78%, 8.30%, 6.70% and 7.42% of the respective parameter spans, or 0.88δ, 0.08δ, 0.03δ and 0.06δ (13.2 m, 1.24 m, 0.50 m and 0.83 m at 1:50,000). Errors were grouped by the ground-truth operator label. Because parameters were normalized to [0, 1] within each operator's candidate range, these percentages are directly comparable across operators, while the δ and meter values use the spans in Table A1.

3.4. Comparison of Generalization Results

The proposed method was compared on the test set with four schemes that applied DP, DT, DB, or MC as the sole operator. Results for DP with a uniform threshold are also reported in Table 4. Evaluation metrics comprised the vertex-reduction ratio (compression ratio, CR), length retention ratio (LR), bend-retention ratio, mean displacement ($\overline{D}$), mean total unresolved topological events per run, and two directional area components. Mean displacement is the one-way mean distance from the vertices of the source shoreline to their nearest points on the generalized shoreline; it is not a symmetric Hausdorff distance. Land-area increase is equivalent to water-area decrease within the comparison domain, whereas land-area decrease is equivalent to water-area increase. These two signed components are distinct from the unsigned symmetric-difference area, which is their sum. Topological events were counted as the number of unresolved hard-constraint events remaining after local correction and before the final fallback. The reported values are means across independent runs of the total number of unresolved events accumulated over all 385 test shoreline segments in each run; a segment may contain more than one such event.

Table 4. Comparison of generalization results between the proposed method and baseline algorithms.

Method	CR (%)	LR (%)	Bend Retention (%)	$\overline{D}$ (m)	Mean Total Unresolved Topological Events per Run	Land-Area Increase/Water-Area Decrease (m^2)	Land-Area Decrease/Water-Area Increase (m^2)
Global DP	88.45 ± 0.14	94.40 ± 0.11	82.98 ± 0.38	38.96 ± 0.62	45.60 ± 3.05	2629 ± 118	2388 ± 109
Global DT	40.05 ± 0.21	98.18 ± 0.08	93.51 ± 0.31	16.59 ± 0.29	10.80 ± 1.30	1084 ± 72	1011 ± 68
Global DB	32.31 ± 0.24	98.33 ± 0.07	91.32 ± 0.36	22.05 ± 0.34	12.40 ± 1.52	1181 ± 79	1098 ± 75
Global MC	87.00 ± 0.17	95.03 ± 0.10	84.26 ± 0.41	38.32 ± 0.58	45.40 ± 2.88	2543 ± 113	2296 ± 105
Uniform-threshold DP	88.75 ± 0.13	94.32 ± 0.12	82.70 ± 0.37	21.84 ± 0.40	47.80 ± 3.11	1712 ± 96	1551 ± 91
Proposed method	86.43 ± 0.16	97.16 ± 0.09	94.84 ± 0.25	17.12 ± 0.13	6.20 ± 2.28	980 ± 61	870 ± 58
Global DP + correction	87.92 ± 0.18	94.68 ± 0.12	87.61 ± 0.43	34.72 ± 0.55	8.40 ± 1.52	2180 ± 102	1950 ± 94
Global DT + correction	39.72 ± 0.22	98.25 ± 0.08	94.05 ± 0.30	15.98 ± 0.27	5.80 ± 1.30	650 ± 48	610 ± 45
Global DB + correction	31.94 ± 0.25	98.41 ± 0.07	92.18 ± 0.34	20.74 ± 0.31	5.40 ± 1.14	720 ± 52	690 ± 50
Global MC + correction	86.52 ± 0.19	95.36 ± 0.10	88.43 ± 0.39	34.06 ± 0.51	7.60 ± 1.52	2050 ± 97	1870 ± 90
Uniform-threshold DP + correction	88.13 ± 0.15	94.61 ± 0.11	87.24 ± 0.41	20.15 ± 0.36	8.00 ± 1.58	1350 ± 81	1220 ± 76

Note: global DP applies a single tolerance to the whole shoreline, whereas uniform-threshold DP partitions the shoreline into bends and applies the same scale-related threshold to every bend. Rows marked "+ correction" add the constraint-correction stage of Section 2.4 to the corresponding baseline; the proposed-method row includes the correction stage by design; the remaining baseline rows do not include correction. Topological-event values are means ± SD across independent runs for the learned methods; for the deterministic baselines, the ± value denotes the dispersion over the evaluation set obtained by resampling the 385 evaluation segments (Section 3.3), so the ± entries in this column quantify two different sources of variation rather than a single common uncertainty. In each run, events were accumulated over the 385 test shoreline segments. For methods with correction, events were counted after local correction and before final fallback; for methods without correction, they were counted directly after operator execution. These values are diagnostic pre-fallback or post-operator totals and do not represent residual errors in the final shorelines. All geometric and area metrics were calculated from the final shoreline after the applicable correction, fallback replacement, reassembly and final validation. Because $\delta = 0.0003S$ m is defined per segment, the absolute displacement and area values are not referenced to a single target scale.

Bend endpoints were matched within $\delta = 0.0003S$ m, in line order, and the matched bend had to exceed the bend-height threshold δ. The bend-retention ratio is the number of matched source bends divided by the number of source bends. Mean displacement was computed from the vertices of the source shoreline to the generalized shoreline. Each

segment carries its own target scale S, so $\delta = 0.0003S$ m is evaluated per segment. The displacement constraint of Section 2.4 is applied to each bend candidate relative to its source bend, whereas $\overline{D}$ is a shoreline-level one-way mean from the source vertices to the final generalized shoreline; the two are different aggregations, and the absolute values in Table 4 should not be compared with the δ of a single target scale. For the proposed method, this total was 6.20 ± 2.28 unresolved events over the 385 test segments. These events were counted after local correction but before final fallback and therefore serve as a diagnostic measure of the violations that remained at that stage. The corresponding final shorelines were obtained after fallback replacement, shoreline reassembly and final topology validation. The total comprised 2.40 ± 1.14 self-intersections, 0.80 ± 0.45 island–mainland merges, 1.20 ± 0.84 discontinuities and 1.80 ± 0.84 abnormal closures. Candidates that still violated a hard constraint were replaced by their original bend geometry before output. These counts therefore characterize violations unresolved by local correction, rather than errors remaining in the final generalized shoreline. For the proposed method, no residual violations of these four types were detected in the final outputs after fallback. The final outputs were re-checked across all evaluated sets (test set: 385 segments; Bohai Bay hold-out: 214 segments; external comparison: 96 segments). Because each operator affects only its own bend and the source geometry is topologically valid, restoring a bend that fails a hard check removes the local self-topology violation it would otherwise introduce; this argument does not by itself cover inter-shoreline proximity conflicts, which involve two shorelines and are handled through the proximity links of Section 2.2.2. The assembled output is re-validated as described in Section 2.4. The absence of these violations is therefore an empirical outcome of this re-validation rather than a formal guarantee.

Within Table 4, the schemes fall into two groups according to simplification intensity. Global DP, global MC, uniform-threshold DP and the proposed method occupy a narrow compression band: their vertex-reduction ratios span 86.43% to 88.75%, a spread of 2.32 percentage points, and within that band the proposed method retained 94.84% of bends, against 82.70%, 82.98% and 84.26% for the three baselines, an advantage of 10.58 to 12.14 percentage points. With the correction stage applied to every scheme, the same group spans 86.43% to 88.13%, a spread of 1.70 percentage points, and the proposed method retains 94.84% against 87.24%, 87.61% and 88.43%, an advantage of 6.41 to 7.60 percentage points. These four schemes are therefore compared at closely comparable simplification intensity. Global DT and global DB fall outside that band, at 40.05% and 32.31% vertex reduction (39.72% and 31.94% with correction), which is 46.38 and 54.12 percentage points below the proposed method. They are listed in Table 4 for completeness only and are not used here for a direct comparison of shape preservation or displacement. The proposed method had a mean displacement of 17.12 m and a mean run-level total of 6.20 unresolved topological events before the final fallback. Table A3 reports the number of bends assigned to each correction branch over the complete dataset; because a bend may be associated with more than one unresolved event, these bend counts are expressed in a different unit from the run-level event totals of Table 4 and are reported separately. Because the uncorrected rows are counted at the post-operator stage whereas the corrected rows and the proposed method are counted after local correction, the topological-event entries should not be compared across these two groups; the within-method attribution of the correction stage is instead shown by the contrast between 103.70 and 6.20 events in the ablation of Section 3.5. Representative final generalization results for the lagoon, estuary, island-adjacency and narrow-channel examples are compared in Figure 5.

The DP sweep used tolerance factors 0.5, 1, 2, 4, 6 and 8, and shows the expected monotonic compression–retention trade-off. Full parameter-sweep curves for the four operators,

with the proposed method plotted as a single point on the same axes, are presented in Figure 6.

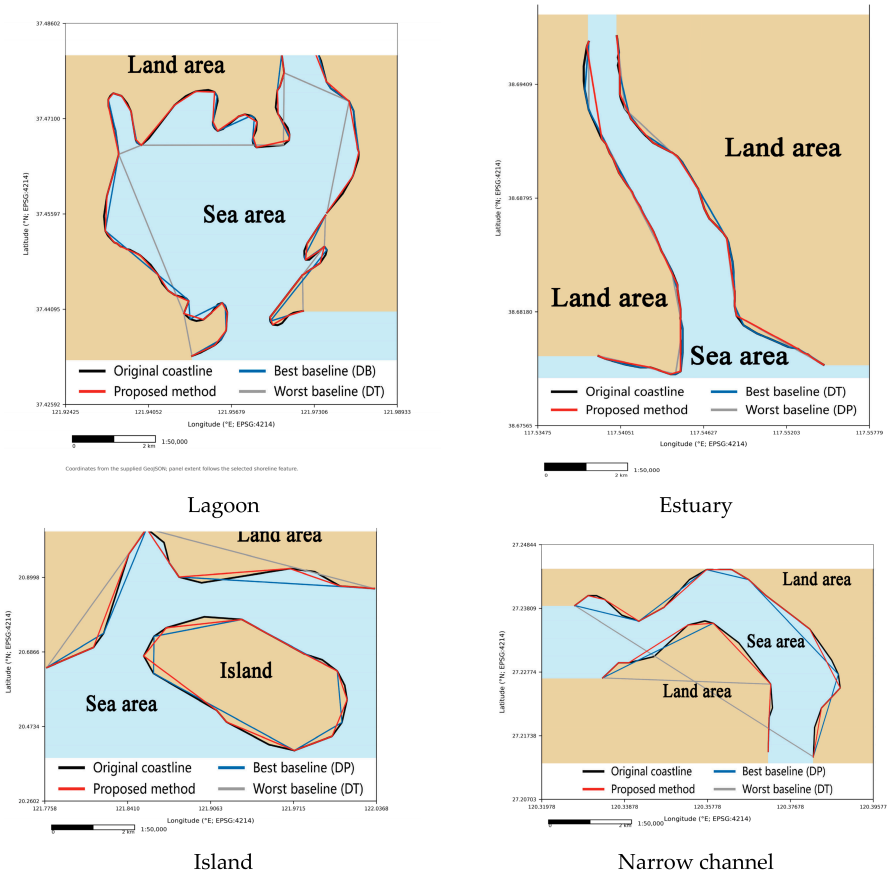


Figure 5. Comparison of final generalization results for representative coastal segments. The comparison baselines are identified in each panel's legend (DP, DT, DB and MC; see Section 2.3 for definitions).

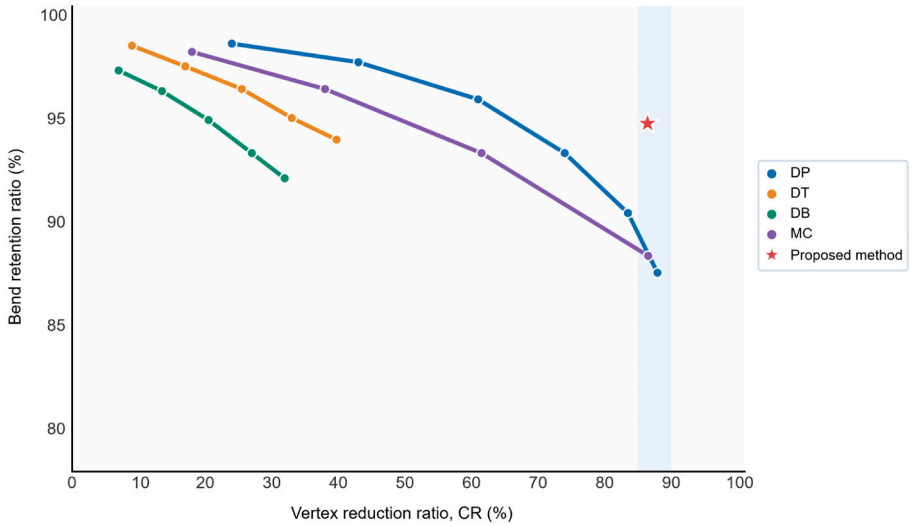


Figure 6. Compression–retention trade-off at a target scale of 1:50,000. All operator-sweep curves were generated through the applicable constraint-correction, fallback, shoreline-reassembly and final-topology-validation steps and use the constraint thresholds and stopping rule of Table A1; they therefore represent final shoreline outputs. The proposed method is shown as a single operating point on the same axes rather than as a curve.

Local morphological differences among four coastal scenarios are illustrated in Figure 6. In the lagoon example, several inward bends were replaced by longer straight segments in the DT result, whereas more interior shoreline detail was retained by the proposed method. In the island-adjacency example, the closed island outline was retained, and a water gap remained discernible near the mainland shoreline. In the narrow-channel example, the major directional changes in the source shoreline were followed by the proposed method, whereas long straight segments crossed several bends in the DT result. In the estuary example, the overall alignment was similar across schemes, with differences concentrated at local turns. These visible differences complement the aggregate metrics in Table 4.

3.5. Ablation Experiments

The complete model and its variants were compared on the same test set (Table 5). The complete model achieved a classification accuracy of 91.97% and a parameter regression R^2 of 0.895. The variants altered the encoding direction, attention mechanism, parameter regression, constraint correction, or coastal-segment type input. In the variant without multitask learning, the parameter regression branch was removed and fixed parameters were used during execution. Its results therefore reflect the combined effects of these two changes.

Table 5. Results of the ablation experiments.

Model Variant	Classification Accuracy (%)	Macro-Averaged F1 (%)	Bend Retention (%)	Mean Total Unresolved Topological Events per Run	$\bar{D}$ (m)
Complete model	91.97 ± 0.24	92.07 ± 0.29	94.84 ± 0.25	6.20 ± 2.28	17.12 ± 0.13
Unidirectional LSTM instead of BiLSTM	89.11 ± 0.31	88.31 ± 0.36	92.43 ± 0.39	15.20 ± 2.07	25.79 ± 0.54
Without attention	87.20 ± 0.36	86.29 ± 0.41	90.53 ± 0.44	18.20 ± 2.39	27.34 ± 0.61
Without multitask learning (fixed parameters)	90.74 ± 0.28	90.04 ± 0.33	91.89 ± 0.42	17.80 ± 2.17	28.37 ± 0.66
Without constraint correction	91.97 ± 0.24	92.07 ± 0.29	90.25 ± 0.47	103.70 ± 5.12	28.63 ± 0.73
Without coastal-segment type code τ_j	88.57 ± 0.33	87.80 ± 0.38	91.50 ± 0.41	17.60 ± 2.07	27.21 ± 0.59
Without proximity links	91.42 ± 0.25	90.91 ± 0.32	94.21 ± 0.31	7.00 ± 1.58	17.35 ± 0.22

Note: For variants with constraint correction, topological events were counted after local correction but before final fallback. For the variant without constraint correction, events were counted directly after operator execution. The reported values are run-level totals accumulated over the 385 test segments. The ± values denote the sample standard deviation over five independent training runs.

When the BiLSTM was replaced by a unidirectional LSTM, classification accuracy decreased from 91.97% to 89.11%. Removing attention reduced classification accuracy by 4.77 percentage points and the bend-retention ratio from 94.84% to 90.53%. Removing the coastal-segment type code τ_j reduced classification accuracy by 3.40 percentage points. These results characterize the predictive performance associated with changes in encoding, attention, and type inputs. When constraint correction was omitted, classification accuracy remained unchanged, but the mean run-level total of unresolved topological events increased from 6.20 after local correction to 103.70. This indicates that constraint correction primarily affects the geometry produced after operator execution.

Table 5 reports means ± SD over five runs. Without proximity links, accuracy was 91.42 ± 0.25%, macro-F1 90.91 ± 0.32%, bend retention 94.21 ± 0.31%, the mean run-level total of unresolved topological events 7.00 ± 1.58 and mean displacement 17.35 ± 0.22 m. The changes are small across all five metrics, and their direction is consistent with the local geometric reproduction, in which the link-aware branch altered CR/LR by no more

than 0.03 percentage points; that reproduction reports CR and LR only and is therefore not directly comparable with the retrained variant in Table 5.

In a perturbation experiment on 385 test segments, the eight-dimensional feature vector at the highest mean-attention position was replaced by its within-scenario median and compared with perturbation of a random position. The predicted operator changed for 31.4% versus 12.6% of segments, and the median normalized parameter change was 0.083 versus 0.029; paired-permutation tests (10,000 random sign flips of the 385 segment-level differences; random seed 20260906) gave $p < 0.001$ for both contrasts. The attention weight matrices of four representative segments are provided in Supplementary Figure S3.

3.6. Summary of Experimental Results

The distribution of preferred algorithms differed among coastal scenarios, providing a sample-based rationale for configuring operators and parameters at the bend level. Under the experimental settings used here, the proposed method reported a vertex-reduction ratio close to that of global DP while retaining more bend detail. In the ablation experiments, changes to encoding, attention, and coastal-segment type inputs were accompanied by differences in predictive performance. Omitting constraint correction left classification accuracy unchanged but produced more topological violations at the operator-output stage than remained after local correction in the complete model. Strategy prediction and result correction therefore affect different stages of the generalization workflow. Both predictive performance and final shoreline quality should consequently be assessed when evaluating the method. For global DP, global MC, uniform-threshold DP and the proposed method, the comparison is made at closely comparable simplification intensity, whereas global DT and global DB operate far below it and are reported for reference only. The results therefore characterize the proposed method at the compression level of Table 4 rather than across the full compression range.

3.7. Independent and Expert Validation

The evaluations above measure agreement with the supervisory labels, which were generated by the scoring procedure described in Section 2.5. Such agreement indicates how well the model reproduces the internal candidate-selection procedure, but it does not by itself establish that the resulting shoreline is superior from an independent cartographic perspective. We therefore conducted three complementary checks, including a geographic hold-out evaluation, an external comparison with official smaller-scale shorelines, and a blinded expert assessment. All geometric metrics and expert-assessment scores reported in this section were calculated from the final shoreline outputs obtained after the applicable correction, fallback replacement, reassembly and final topology validation; intermediate post-operator and pre-fallback geometries were not used for these evaluations.

For the geographic hold-out evaluation, the Bohai Bay–Liaodong Bay region was excluded from model development, leaving 214 shoreline segments and 8742 bends for evaluation. Across five random-seed runs, the classification accuracy was 89.76 $\pm$ 0.41%, the bend-retention rate was 92.88 $\pm$ 0.52%, and the mean run-level total of unresolved topological events after local correction but before final fallback was 7.40 $\pm$ 1.52.

The external comparison used 96 shoreline segments for which official smaller-scale shorelines from the same general area were available. Compared with the official products, the per-segment differences (mean $\pm$ SD across the 96 segments) were 1.8 $\pm$ 3.6 percentage points in vertex-reduction rate and 2.3 $\pm$ 4.1 percentage points in bend-retention rate, while the root-mean-square error of the per-segment mean-displacement differences was 4.6 m. The survey dates differed by approximately 2–11 years, and the source accuracy and compilation objectives were not identical. Therefore, the official shorelines were treated

as contextual external references rather than point-wise ground truth. With 96 external segments and 60 segments rated by five experts, these independent evaluations rest on relatively small samples; the corresponding estimates should therefore be read as indicative rather than definitive, and the confidence intervals reported with them should be interpreted with this limitation in mind.

Finally, five experienced nautical-chart compilers independently evaluated 60 representative shoreline segments drawn from the test set using a five-point scale. The three evaluation dimensions were morphological preservation, navigational legibility, and target-scale acceptability. The five methods—the proposed method and the four global baselines (global DP, DT, DB and MC), each processed with the constraint-correction stage of Section 2.4—were presented using anonymous method codes and in randomized order. After the ratings were completed, the method identities were unblinded for analysis. The mean scores and standard deviations are summarized in Table 6. Inter-rater reliability was assessed using a two-way random-effects, absolute-agreement, average-measures intraclass correlation coefficient, ICC(2,k), in Table 7. Because global DT and global DB operate at a substantially lower compression level than the proposed method and the other two baselines (Section 3.4), the corresponding comparisons in Table 8 are reported for completeness and should not be read as like-for-like quality comparisons. Pairwise differences between the proposed method and each baseline were evaluated using paired-permutation tests. Four paired-permutation comparisons of the overall expert score (the proposed method versus each of the four baselines; 60 paired segments) were made, and Holm–Bonferroni correction was applied across these four comparisons; the mean differences, 95% confidence intervals, effect sizes (Cliff's δ) and the raw and Holm-adjusted *p*-values are reported in Table 8.

Table 6. Results of the blinded expert assessment.

Method	Morphological Preservation	Navigational Legibility	Target-Scale Acceptability	Overall Score
Proposed method	4.42 ± 0.36	4.35 ± 0.40	4.38 ± 0.37	4.38 ± 0.34
DP	3.96 ± 0.48	3.82 ± 0.52	3.91 ± 0.46	3.90 ± 0.43
DB	3.61 ± 0.55	3.52 ± 0.58	3.57 ± 0.53	3.57 ± 0.51
DT	3.48 ± 0.60	3.44 ± 0.57	3.51 ± 0.56	3.48 ± 0.54
MC	3.39 ± 0.62	3.35 ± 0.61	3.42 ± 0.59	3.39 ± 0.58

Table 7. Inter-rater reliability of the blinded expert assessment.

Evaluation Dimension	ICC(2,k)	95% Confidence Interval	Interpretation
Morphological preservation	0.84	0.79–0.88	Good agreement
Navigational legibility	0.81	0.75–0.86	Good agreement
Target-scale acceptability	0.79	0.72–0.85	Good agreement
Overall score	0.83	0.78–0.87	Good agreement

Table 8. Pairwise paired-permutation comparisons between the proposed method and each baseline for the overall expert score (n = 60 paired segments). Mean differences are given on the five-point scale; Holm–Bonferroni correction was applied across the four comparisons. Effect sizes are reported as Cliff's δ.

Comparison	Mean Difference	95% CI	*p* (Raw)	*p* (Holm)	Effect Size (Cliff's δ)
Proposed vs. DP	0.48	[0.38, 0.58]	<0.001	<0.001	0.62
Proposed vs. DB	0.81	[0.69, 0.93]	<0.001	<0.001	0.81

Table 8. *Cont.*

Comparison	Mean Difference	95% CI	p (Raw)	p (Holm)	Effect Size (Cliff's δ)
Proposed vs. DT	0.90	[0.78, 1.02]	<0.001	<0.001	0.84
Proposed vs. MC	0.99	[0.86, 1.12]	<0.001	<0.001	0.86

Note: Each expert rated every segment independently on a five-point scale. For each segment and method, the five expert ratings were averaged within each dimension, and the overall score was the unweighted mean of the three dimension scores (equal weights). All means, standard deviations and mean differences reported here are computed over the 60 segment-level values, with the segment rather than the individual rating as the unit of analysis. Paired differences were tested with a two-sided paired-permutation test using 10,000 random sign flips of the 60 within-segment differences; $p = (1 + \#\{|T^*| \geq |Tobs|\})/(B + 1)$, with B = 10,000 and random seed 20260906. The 95% confidence intervals were obtained by a paired percentile bootstrap with 10,000 resamples of the 60 segments. Effect sizes are two-sample Cliff's δ computed between the two sets of 60 segment-level scores (all 3600 cross-pairs) as $\delta = P(X > Y) - P(X < Y)$, where X and Y are randomly selected segment-level scores under the proposed method and the baseline, respectively, and ties contribute zero to both probabilities. This marginal effect size is reported alongside the paired mean difference, paired confidence interval and paired-permutation test and is not a within-segment contrast. Holm–Bonferroni adjustment was applied across the four comparisons.

4. Discussion

The experiments address the problem stated in Section 1, namely that a single operator with a uniform parameter may be inadequate for a segment in which bend morphology and inter-shoreline proximity vary along the line. The label distribution in Table 2 shows that no single operator dominates: DP is preferred for a majority of ordinary sinuous bends (54.92%), DB for the largest share of lagoon bends (39.39%), and DT and MC take substantial shares in the estuarine and island-adjacency scenarios. The learned strategy reproduces this differentiation with a validation accuracy of 91.83%. Because the model uses bend features, sequence context and the supplied coastal scenario, the result is consistent with an association between these inputs and the operator labels generated by the scoring procedure; it does not demonstrate that they causally determine an independently optimal cartographic operator. Table 4 reports a vertex-reduction ratio close to that of global DP, 86.43% against 88.45%, with a bend-retention ratio of 94.84%, against 82.98% for global DP. Its mean displacement of 17.12 m is well below the 38.96 m and 38.32 m of global DP and global MC at comparable compression; global DT reports 16.59 m, but at 40.05% vertex reduction it is not a like-for-like comparison. Within the matched compression band, the data identify a compression–preservation operating point that none of the three comparable baselines reaches, but they do not establish behavior outside that band or against global DT and global DB.

Among the corrected baselines, corrected global MC retained 88.43% of bends, 6.41 percentage points below the proposed method. Corrected global DT retained 94.05%, only 0.79 points below the proposed method, but its vertex reduction was 39.72%. The proposed method's mean run-level total of 6.20 unresolved topological events after local correction but before final fallback was not the lowest (corrected DB: 5.40). These values characterize violations unresolved by local correction rather than errors retained after final fallback. The defensible interpretation is therefore a different compression–retention operating point, not across-the-board superiority over every corrected baseline.

Transferability is constrained by national charting practice. The admissible displacement, the minimum size of a surviving bend and the treatment of channels and harbor entrances differ between chart series, so a model trained on one series encodes the label-scoring and constraint conventions of that series. The architecture itself is not specific to nautical charts, and the same bend-level formulation could be applied to other linear features once the operator set, the parameter ranges, and the constraint thresholds have

been redefined for the target specification, but applying the method to another series would require re-deriving the labels and re-validating the thresholds, and should be expected to require adaptation rather than to transfer without change.

Future work will extend the method in four directions. The candidate operator set and the parameter ranges will be widened, and the label-scoring sensitivity analysis of Section 2.5 will be extended from the local geometric reproduction to full retraining across the complete score simplex, in order to test how much of the learned strategy is an artifact of the current scoring rules. Geographic and cross-series validation will be extended to a larger set of regions and to at least one non-Chinese chart series. The coastal-segment type will be replaced or supplemented by an automatic pre-classification of the same attributes, using the sequential classification approaches already available for coastline shape [16], so that the method no longer depends on expert input. Finally, the correction stage will be made more explicit by learning which constraint is binding, so that the interaction between the learned and the rule-based components becomes measurable rather than reconstructed after the fact.

5. Conclusions

An adaptive generalization method for complex shoreline segments based on bend sequence attention learning was proposed. Bend units were used as the basis for local decisions, and paired operator-class and parameter labels were constructed by screening candidate results against cartographic constraints and applying a two-stage selection procedure. Geometric descriptors, proximity distances and coastal-segment type features were combined with bidirectional sequence encoding and self-attention to predict a local generalization strategy for every bend, which conventional operators then executed and which the correction stage subsequently adjusted. Under the evaluated settings, the workflow reported a vertex-reduction ratio of 86.43%, a bend-retention ratio of 94.84% and a one-way mean displacement of 17.12 m; the validation classification accuracy was 91.83%, and the parameter regression R^2 was 0.893. Without constraint correction, the mean run-level total of unresolved topological events directly after operator execution was 103.70 over the 385 test segments. After local correction but before final fallback, the corresponding total was 6.20. These are diagnostic quantities at different processing stages. Following fallback replacement, shoreline reassembly and final topology validation, no residual violations of the four checked topological types were detected in the evaluated final shorelines. The reported geometric metrics and expert-assessment scores were calculated from these final shorelines. The reported comparisons define performance within the present scoring procedure, datasets and simplification settings; they do not establish independent cartographic quality or universal superiority. The method is further limited because the evaluation covers only Chinese nautical-chart series and the five complex-segment morphologies examined here, and because the coastal-segment type currently has to be supplied at inference time. Extending the method to other chart series, replacing the expert type attribute with an automatic classifier, and widening the operator and parameter space are the next steps.

Supplementary Materials: The following supporting information can be downloaded at: https://www.mdpi.com/article/10.3390/jmse14211985/s1, Figure S1: Smoothing of the digitized shoreline and identification of curvature zeros and candidate peaks; Figure S2: Screening of adjacent curvature peaks; Figure S3: Attention weight matrices of four representative shoreline segments; Table S1: Per-operator MSE and RMSE of parameter regression; Table S2: Mean scores from the completed blinded expert assessment; Table S3: Inter-rater reliability of the completed blinded expert assessment.

Author Contributions: Conceptualization, H.S. and S.J.; methodology, H.S. and S.J.; software, H.S.; validation, H.S., S.J., X.L. and Z.S.; formal analysis, Q.S.; investigation, Q.S.; resources, Z.S.;

data curation, X.L.; writing—original draft preparation, H.S.; writing—review and editing, S.J.; visualization, H.S.; supervision, S.J. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The source shorelines are derived from licensed nautical charts, and the original geometry cannot be released publicly. The code, the parameter configuration, the split manifest, the source-chart list and an anonymized example will be provided through a controlled-access procedure once the license review has been completed.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A

Tables A1–A3 collect the candidate parameters, constraint thresholds, network configuration, repeated runs and timing scope used in this study.

Table A1. Candidate parameter ranges, constraint thresholds and correction settings of the four operators.

Item	DP	DT	DB	MC
Candidate parameter values	0.50δ, 1.00δ, 2.00δ, 4.00δ, 6.00δ, 8.00δ	0.50δ, 0.75δ, 1.00δ, 1.25δ, 1.50δ	0.50b, 0.75b, 1.00b, 1.25b, 1.50b (b = 0.00015S m)	0.25δ, 0.50δ, 0.75δ, 1.00δ
Clipping bounds (θ_{min}, θ_{max})	0.50δ–8.00δ	0.50δ–1.50δ	0.50b–1.50b	0.25δ–1.00δ
Displacement tolerance (mean/local max)	δ/2δ	δ/2δ	δ/2δ	δ/2δ
Landward shoreline-displacement limit (land-area decrease/water-area increase)	0.5δ	0.5δ	0.5δ	0.5δ
Land-area-increase/water-area-decrease limit (m^2)	0.50δL*c*	0.50δL*c*	0.50δL*c*	0.50δL*c*
Land-area-decrease/water-area-increase limit (m^2)	0.25δL*c*	0.25δL*c*	0.25δL*c*	0.25δL*c*
Channel/harbor entrance protection rule	min width ≥ max ($0.8W_0$, 2δ)	min width ≥ max ($0.8W_0$, 2δ)	min width ≥ max ($0.8W_0$, 2δ)	min width ≥ max ($0.8W_0$, 2δ)
Correction stopping condition	all hard checks pass, no improvement, or 5 iterations	all hard checks pass, no improvement, or 5 iterations	all hard checks pass, no improvement, or 5 iterations	all hard checks pass, no improvement, or 5 iterations

Table A2. Network architecture, training settings, repeated runs, runtime and quality gap.

Item	Value
BiLSTM hidden dimension	128
Number of BiLSTM layers	1
Attention heads/key dimension d	4 heads/d = 32
Classifier and regressor dimensions	shared 256 → 128; classifier 128 → 64 → 4; regressor 128 → 64 → 1
Loss weights (λc, λ_r)	λc = 1.0, λ_r = 0.5
Batch size and sequence padding rule	16-segment sequences; dynamic padding with an attention mask
Optimizer and learning rate schedule	Adam, lr = 1×10^{-3}; ReduceLROnPlateau factor 0.5, patience 5, min lr = 1×10^{-5}
Early-stopping patience	15 epochs
Number of independent runs and seeds	5 runs; 20260906–20260910
Hardware (CPU/GPU, memory)	Measured geometry: Windows 11, 12-core-logical AMD64, 32 GB RAM; network: RTX 3060 12 GB
Runtime of exhaustive candidate evaluation per representative chart	0.77 s per representative chart (including segmentation, feature extraction, proximity linking and candidate generation of every bend on the chart)
Runtime of one forward pass for the same chart	0.12 s per representative chart after warm-up (measured as the mean of repeated runs after warm-up on the hardware specified in this table; the value excludes segmentation, feature extraction and candidate generation, which are included in the exhaustive-evaluation figure above, and it also excludes operator execution and the constraint-correction stage. The two runtime rows are therefore time-different stages of the workflow and should not be compared directly with each other).
Quality gap between network output and exhaustive optimum (vertex reduction, bend retention, mean displacement)	CR −0.31 pp; bend retention −0.42 pp; mean displacement +0.28 m

Table A3. Frequency of the constraint-correction outcomes.

Correction Branch	Number of Bends	Share of Corrected Bends (%)	Share of All Bends (%)
Parameter rollback	13,533	77.80	12.94
Operator reselection	2698	15.51	2.58
Original bend retained (final fallback after local correction failed)	1164	6.69	1.11
Total corrected	17,395	100.00	16.63

Note: Table A3 reports bend-level correction outcomes over the complete dataset of 2571 shoreline segments and 104,619 bend units. These counts are not topological-event totals and are not directly comparable with the test-set run-level event totals in Tables 4 and 5.

References

1. *GB 12320-2022;* Compilation Specifications for Chinese Nautical Charts. State Administration for Market Regulation. Standardization Administration of China: Beijing, China, 2022. Available online: https://openstd.samr.gov.cn/bzgk/gb/newGbInfo?hcno=608288A05714A50E3015708F9BEC9A53 (accessed on 7 September 2026).
2. International Hydrographic Organization. *Regulations of the IHO for International (INT) Charts and Chart Specifications of the IHO*, 4.10.0 ed.; S-4; International Hydrographic Organization: Antoine, Monaco, 2026. Available online: https://iho.int/uploads/user/pubs/standards/s-4/S-4%20Ed%204.10.0_FINAL.pdf (accessed on 7 September 2026).
3. Chen, H.; Peng, R.; Zheng, Y.; Li, S. Coastline generalization based on skeleton line of curve bends. *Geomat. Inf. Sci. Wuhan Univ.* **2011**, *36*, 1418–1422.
4. Huang, Y.; Ai, T.; Liu, Y.; Zhang, H. Geographic-feature oriented ria coastline simplification. *Acta Geod. Cartogr. Sin.* **2013**, *42*, 595–601.
5. Du, J.; Wu, F.; Li, J.; Xing, R.; Gong, X. A progressive simplification method for the estuary coastline. *Acta Geod. Cartogr. Sin.* **2018**, *47*, 547–556. [CrossRef]
6. Zhang, L.; Tang, L.; Jia, S.; Dai, Z. A collaborative simplification method for multiple coastlines based on the hierarchical triangulation network partition. *Acta Geod. Cartogr. Sin.* **2019**, *48*, 520–531. [CrossRef]
7. Wu, F.; Gong, X.; Du, J. Overview of the research progress in automated map generalization. *Acta Geod. Cartogr. Sin.* **2017**, *46*, 1645–1664. [CrossRef]
8. Douglas, D.H.; Peucker, T.K. Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. *Cartographica* **1973**, *10*, 112–122. [CrossRef]
9. Visvalingam, M.; Whyatt, J.D. Line generalisation by repeated elimination of points. *Cartogr. J.* **1993**, *30*, 46–51. [CrossRef]
10. Park, W.; Yu, K. Hybrid line simplification for cartographic generalization. *Pattern Recognit. Lett.* **2011**, *32*, 1267–1273. [CrossRef]
11. García Balboa, J.L.; Ariza López, F.J. Sinuosity pattern recognition of road features for segmentation purposes in cartographic generalization. *Pattern Recognit.* **2009**, *42*, 2150–2159. [CrossRef]
12. Ai, T.; Zhou, Q.; Zhang, X.; Huang, Y.; Zhou, M. A simplification of ria coastline with geomorphologic characteristics preserved. *Mar. Geod.* **2014**, *37*, 167–186. [CrossRef]
13. Wu, F.; Du, J.; Qian, H.; Zhai, R. Overview of research progress and reflections in intelligent map generalization. *Geomat. Inf. Sci. Wuhan Univ.* **2022**, *47*, 1675–1687. [CrossRef]
14. Yan, X.; Yang, M.; Ai, T. Deep learning in automatic map generalization: Achievements and challenges. *Geo-Spat. Inf. Sci.* **2025**, *28*, 2905–2926. [CrossRef]
15. Harrie, L.; Touya, G.; Oucheikh, R.; Ai, T.; Courtial, A.; Richter, K.F. Machine learning in cartography. *Cartogr. Geogr. Inf. Sci.* **2024**, *51*, 1–19. [CrossRef]
16. Gao, A.; Ai, T.; Yu, H.; Xiao, T.; Chen, Y.; Li, J.; Huang, H. A vector-based coastline shape classification approach using sequential deep learning model. *Int. J. Appl. Earth Obs. Geoinf.* **2024**, *129*, 103810. [CrossRef]
17. Yu, W.; Chen, Y. Data-driven polyline simplification using a stacked autoencoder-based deep neural network. *Trans. GIS* **2022**, *26*, 2302–2325. [CrossRef]
18. Chen, S.; Hu, A.; Xu, Y.; Wang, H.; Xie, Z. VE-GCN: A geography-aware approach for polyline simplification in cartographic generalization. *ISPRS Int. J. Geo-Inf.* **2025**, *14*, 64. [CrossRef]
19. Courtial, A.; Touya, G.; Zhang, X. Deriving map images of generalised mountain roads with generative adversarial networks. *Int. J. Geogr. Inf. Sci.* **2023**, *37*, 499–528. [CrossRef]
20. Zhu, W.; Guo, Q.; Yang, N.; Tong, Y.; Zheng, C. An improved generative adversarial network for generating multi-scale electronic map tiles considering cartographic requirements. *ISPRS Int. J. Geo-Inf.* **2024**, *13*, 398. [CrossRef]
21. Yan, X.; Yang, M. A deep learning approach for polyline and building simplification based on graph autoencoder with flexible constraints. *Cartogr. Geogr. Inf. Sci.* **2024**, *51*, 79–96. [CrossRef]
22. Xiao, T.; Ai, T.; Yu, H.; Yang, M.; Liu, P. A point selection method in map generalization using graph convolutional network model. *Cartogr. Geogr. Inf. Sci.* **2024**, *51*, 20–40. [CrossRef]
23. Xiao, T.; Ai, T.; Burghardt, D.; Liu, P.; Yang, M.; Gao, A.; Kong, B.; Yu, H. A road generalization method using graph convolutional network based on mesh-line structure unit. *Geocarto Int.* **2024**, *39*, 2413549. [CrossRef]
24. Wang, D.; Qian, H.; Wang, X.; Xie, L.; Chen, Y.; Kong, L. An intelligent simplification method for river networks with an unsupervised variational autoencoder. *Int. J. Digit. Earth* **2025**, *18*, 2495736. [CrossRef]
25. Zhou, Z.; Fu, C.; Feng, Y.; Touya, G.; Sester, M.; Weibel, R. GeoAI for map generalization in multi-scale cartography: Foundations, a research agenda, and interdisciplinary perspectives. *Int. J. Geogr. Inf. Sci.* **2026**, *40*, 3089–3117. [CrossRef]
26. Wang, Z.; Müller, J.C. Line generalization based on analysis of shape characteristics. *Cartogr. Geogr. Inf. Syst.* **1998**, *25*, 3–15. [CrossRef]
27. Schuster, M.; Paliwal, K.K. Bidirectional recurrent neural networks. *IEEE Trans. Signal Process.* **1997**, *45*, 2673–2681. [CrossRef] [PubMed]

28. Hochreiter, S.; Schmidhuber, J. Long short-term memory. *Neural Comput.* **1997**, *9*, 1735–1780. [CrossRef] [PubMed]
29. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, Ł.; Polosukhin, I. Attention is all you need. In *Advances in Neural Information Processing Systems*; Guyon, I., von Luxburg, U., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2017; Volume 30. Available online: https://papers.nips.cc/paper/7181-attention-is-all-you-need (accessed on 7 September 2026).
30. Caruana, R. Multitask learning. *Mach. Learn.* **1997**, *28*, 41–75. [CrossRef]
31. Van der Poorten, P.M.; Jones, C.B. Characterisation and generalisation of cartographic lines using Delaunay triangulation. *Int. J. Geogr. Inf. Sci.* **2002**, *16*, 773–794. [CrossRef]
32. Ai, T.; Ke, S.; Yang, M.; Li, J. Envelope generation and simplification of polylines using Delaunay triangulation. *Int. J. Geogr. Inf. Sci.* **2017**, *31*, 297–319. [CrossRef]
33. Yang, H.; Li, L.; Hu, H.; Wu, Y.; Xia, H.; Liu, Y.; Tan, S. A coastline generalization method that considers buffer consistency. *PLoS ONE* **2018**, *13*, e0206565. [CrossRef] [PubMed]
34. Saalfeld, A. Topologically consistent line simplification with the Douglas–Peucker algorithm. *Cartogr. Geogr. Inf. Sci.* **1999**, *26*, 7–18. [CrossRef]
35. Liu, B.; Liu, X.; Li, D.; Shi, Y.; Fernandez, G.; Wang, Y. A vector line simplification algorithm based on the Douglas–Peucker algorithm, monotonic chains and dichotomy. *ISPRS Int. J. Geo-Inf.* **2020**, *9*, 251. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.