

Electronic Navigational Chart Change Classification [experiments]

Jacob Arndt 
arndtjw@ornl.gov
Oak Ridge National Laboratory
Oak Ridge, Tennessee, USA

Abhishek Potnis 
potnisav@ornl.gov
Oak Ridge National Laboratory
Oak Ridge, Tennessee, USA

Alexandre Sorokine 
sorokina@ornl.gov
Oak Ridge National Laboratory
Oak Ridge, Tennessee, USA

Abstract

Electronic Navigational Charts (ENCs) are geospatial vector datasets used in maritime navigation systems that represent hydrographic and navigational information such as depths, navigational aids, traffic schemes, and hazards. A major challenge for hydrographic offices is determining whether a given chart change poses a critical or non-critical risk to maritime safety. Existing workflows rely heavily on manual review, which is labor-intensive, scales poorly with the volume of incoming chart updates, and introduces inter-analyst inconsistencies. To address this challenge, we propose a method for automated classification of ENC changes. We establish a baseline encoding scheme to translate complex vector data changes into a structured tabular format for classification models. The two crucial components of the encoding scheme include a spatial context encoder to enrich the change representations with surrounding geographic features, and an ENC attribute encoder to represent nuanced attribute-value descriptions of the modified objects. We evaluate the proposed approach across two distinct operational datasets, comprising 1,308 chart pairs containing over 100,000 individual chart modifications. Tuned gradient-boosted trees leveraging the proposed encoding schemes achieve accuracies of 90% and 94% on the two datasets, yielding a 5-7% improvement over default hyperparameterized models trained on encodings without spatial context and attribute embeddings. These results demonstrate the viability of integrating machine learning into operational geospatial pipelines to improve ENC maintenance and enhance maritime safety. Finally, our experiments demonstrate the effectiveness of simple location and spatial aggregation methods, providing a foundation for evaluating more sophisticated spatial representation learning techniques for this application.

CCS Concepts

• **Information systems** → **Geographic information systems**; • **Computing methodologies** → **Artificial intelligence**.

Keywords

change classification, electronic navigational charts, geospatial vector data, spatial context encoding

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference acronym 'XX, Woodstock, NY

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/2018/06

<https://doi.org/XXXXXXX.XXXXXXX>




Figure 1: Changes in ENCs must be classified in terms of their significance to navigational safety. ENC changes can include additions or deletions of ENC objects, or modifications to an ENC object's attributes, geometry, relationships, or spatial context.

ACM Reference Format:

Jacob Arndt , Abhishek Potnis , and Alexandre Sorokine . 2026. Electronic Navigational Chart Change Classification [experiments]. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 4 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Maintaining Electronic Navigational Charts (ENCs) is a requirement and significant challenge for many hydrographic offices [5]. New incoming data in the form of hydrographic surveys or updated ENCs are received at an increasingly rapid rate and must be assessed for their potential impact to navigational safety at sea. Requirements and guidance for maintenance is described in detail in the International Hydrographic Organization's (IHOs) Standard Reference S-4. A key step in the ENC maintenance workflow is determining whether observed differences between new data and existing data represent critical or non-critical changes (Figure 1). Critical changes, those that are urgent and pose an immediate risk to navigational safety, must be promptly incorporated into the ENC and disseminated using appropriate mechanisms while non-critical changes must be recorded and queued for the next appropriate chart revision. Ensuring that this process is efficient and accurate is absolutely mandatory, as delays or misclassifications can result in serious consequences and risk to safety.

Classifying ENC change criticality is currently performed manually by domain experts with extensive knowledge of maritime data, cartography, and maritime navigation. Upon receiving an ENC change, a typical human review must account for multiple

factors, including: 1) the type of change (e.g., addition, deletion, modification), 2) the affected ENC object class (e.g., light, sounding, depth contour), 3) object attributes (e.g., sector limits, depth value, category of light), 4) spatial relationships (e.g., proximity to other objects, shared boundaries), 5) geometry modifications, 6) collection membership (associations, aggregations, master-slave), and 7) chart metadata (e.g., navigational purpose, compilation scale). To further illustrate the diversity of information and complexity of the decision-making, consider the example in Figure 2.

Addition or modification of a light object's sector attributes. Criticality of the change is determined by the spatial proximity of the sector limit to a danger object. Movement of the sector limit must be plottable by the chart user, which depends on the chart scale and the range of the light. This is unlikely to be less than 1 degree for long-range lights and less than 3 degrees for short-range lights. – S-4 Part B, Section-600 Chart Maintenance

Types of information: Change type, ENC object class, ENC attributes, Spatial relationships, Chart metadata

Figure 2: Summarized excerpt providing guidance on change criticality for a light sector. Color-coding indicates types of information considered in determining change criticality.

One approach to automate ENC change classification is to program expert-defined rules for each possible change scenario. This can be used for simple change scenarios, however, the method does not generalize to the full complexity of ENC data and its maintenance. ENCs contain hundreds of object classes, attributes, and attribute values of diverse types along with multiple geometry types, scale dependencies, defined object-to-object semantic relationships, spatial relationships, and topological rules. Additionally, hydrographic offices often maintain different ENC products that require different criticality rules which further complicates rule development, management, and application. As a result of this complexity, many changes still require human review.

To address these challenges, we introduce a machine learning framework for the automated classification of ENC changes as critical or non-critical. We propose encoding methods that translate complex vector chart modifications into structured tabular representation suitable for machine learning models. Our approach features a spatial context encoder to capture surrounding geographic objects and an attribute encoder to ingest nuanced object properties. We evaluate our encoding methods along with several baseline classification models across two operational datasets and demonstrate the improvement provided by these spatial and attribute representations. To the best of our knowledge, this is the first work to apply machine learning to automated ENC change classification.

The main contributions of this paper are:

- We introduce ENC change classification as a spatial data science and machine learning problem.
- We propose a set of ENC change data encodings that include both spatial context and attribute representations.
- We develop evaluation protocol to benchmark data encodings and classification models for classifying ENC changes as critical or non-critical.

2 Background

The S-57 data standard [3] defines a highly structured theoretical data model and data structure to permit the exchange of digital hydrographic data. Within the S-57 data standard, the ENC product specification ensures consistency in the production of ENC data and its usage in on-vessel navigation systems. ENCs are complex and rich spatial datasets organized into “cells” consisting of: 1) cartographic scale dependency, 2) diverse geometry types (points, lines, areas), 3) hierarchical semantic relationships (e.g. master-slave relationships between objects), 4) spatial topological relationships (e.g. depth contour touches depth area), 5) diverse attribute types (e.g. free-text, number, categorical, list), 6) specific object-attribute relationships (e.g. category of light attribute is applicable to light objects, but not applicable to coastline objects) and 7) consistent updates (ENCs are constantly updated to reflect real-world conditions).

3 Methods

3.1 Preliminaries

Given as input an ENC change instance $I = (C_{old}, C_{new}, \Delta)$, our objective is to predict a *critical* or *non-critical* class label. An ENC change instance consists of an existing ENC cell C_{old} , a new ENC cell C_{new} , and an identified change $\Delta = (\Delta_{old}, \Delta_{new})$. Each ENC cell includes ENC objects $O = \{g_1, g_2, \dots, g_n\}$, where each object g_i is described by an ENC object class g_i^{class} (e.g. sounding, lighthouse, coastline, buoy), a 2D geometry g_i^{geom} such as a point, line, or polygon, a set of key-value(s) attributes $g_i^{attr} = \{(a_1, v_1), \dots, (a_k, v_k)\}$ describing the object, and S-57 relationships ($g_i^{rel} \subseteq \mathcal{R}$) (e.g. aggregations, associations, master-slave) linking the object to other ENC objects in the cell.

The identified change $\Delta = (\Delta_{old}, \Delta_{new})$ contains information about changed ENC objects in the old and new ENC cells. $\Delta_{old} = \{o_1^{old}, o_2^{old}, \dots\} \subseteq C_{old}$ contains references to ENC objects that have been either deleted or modified in the old ENC cell, and $\Delta_{new} = \{o_1^{new}, o_2^{new}, \dots\} \subseteq C_{new}$ contains references to ENC objects that have been either added or modified in the new ENC cell. All ENC objects referenced in an ENC change Δ share the same ENC object class. Five change types are possible depending on the number of objects in the identified change. These change types include *addition*, *deletion*, *modification*, *multi-feature modification*, and *duplicate feature modification*. For example, a *deletion* change type will consist of one object in Δ_{old} and zero objects in Δ_{new} .

3.2 ENC Change Encoding

We encode ENC changes into numeric feature vectors h_{change} designed to be representative of the identified change. Our encoding options include ENC object class, the type of change, ENC relationships for the change objects, the ENC attributes of the change objects, and the spatial context of the change object in the old and new ENC cells.

Object class. We use a simple one-hot encoding to represent the ENC object class for the change. With 172 possible ENC object categories, the ENC object class one-hot encoder E_{class} produces a binary vector h_{class} with dimension 172.

Change type. We use a one-hot encoding to represent ENC change type. The ENC change type one-hot encoder E_{type} produces a binary vector h_{type} with dimension 5.

Relationships. To obtain a representation of S-57 ENC standard relationships for an ENC change, we develop a simple aggregation function. We consider five relationship types including *aggregation peers*, *association peers*, *slave peers*, *slave objects*, and *master object*. We define $\mathcal{R}_r(x)$ as the set of ENC objects related to the change object x via relationship type r . For each relationship type r , we one-hot encode the object class of all objects in $\mathcal{R}_r(x)$ and aggregate the resulting vectors using an element-wise sum. To form the final ENC relationships vector h_{rel} for the ENC change, we concatenate the vectors obtained from each relationship type r and for the old and new contexts resulting in h_{rel} of dimension 1720.

$$h_r(x) = \sum_{x_i \in \mathcal{R}_r(x)} E_{\text{class}}(x_i), \quad (1)$$

Spatial context. We develop a simple spatial context encoder that encodes and aggregates ENC objects that neighbor a given ENC change instance. Specifically, we treat the identified change object as a target node and its neighboring ENC objects as source nodes in a spatial neighborhood. We define $\mathcal{N}(\Delta_c)$ as the set of ENC objects within a spatial neighborhood of the change for context $c \in \{\text{old}, \text{new}\}$. For an ENC change instance Δ , the old neighborhood $\mathcal{N}(\Delta_{\text{old}})$ and new neighborhood $\mathcal{N}(\Delta_{\text{new}})$ are defined as all ENC objects intersecting a 500-meter radius buffer around the centroid of the change instance geometry. Objects in $\mathcal{N}(\Delta_{\text{old}})$ are queried from the old ENC cell, while objects in $\mathcal{N}(\Delta_{\text{new}})$ are queried from the new ENC cell.

For each neighbor $x_i \in \mathcal{N}(\Delta_c)$, the message function is defined as a static one-hot encoding of the ENC object class, $E_{\text{class}}(x_i)$. The target node then computes its context representation h_c by applying an element-wise sum as the aggregation function, accumulating the incoming messages:

$$h_c = \sum_{x_i \in \mathcal{N}(\Delta_c)} E_{\text{class}}(x_i) \quad (2)$$

This aggregation yields a localized bag-of-words (or histogram) representation of ENC object classes surrounding the change, a simple case of an aggregation location encoder [4] inspired by the graph neural network framework. Finally, the spatial representations from both old and new contexts are concatenated to form the complete spatial context encoding h_{context} of dimension 344.

Attributes. For attribute encoding we convert all key-value(s) attributes for an ENC object into strings and concatenate them. All numeric codes and values are translated into their natural language description found in Appendix A of the S-57 data standard. We separate each key and value with a colon and space and separate each key-value pair with a comma and space. For example, an attribute set {BOYSHP : 4, COLOUR : [3], CATLAM : 2, OBJNAM : “Little peconic bay lighted buoy 18”} is translated into: “Buoy shape: Pillar, Colour: Red, Category of lateral mark: Starboard-hand lateral mark, Object name: Little peconic bay lighted buoy 18”. After these strings are constructed, we tokenize and embed them using the DistilBERT transformer architecture pretrained

on Wikipedia and the Toronto Book Corpus. We embed the attributes in this way for the the old change object and new change object and concatenate the resulting vectors together to form a single attribute vector h_{attr} with dimension 1536.

Finally, our complete ENC change encoding is a concatenation of h_{class} , h_{type} , h_{rel} , h_{context} , and h_{attr} resulting in the final change vector h_{change} of dimension 3777.

3.3 Classification Models

We use XGBoost (XGB) as our primary baseline model to classify each identified change as critical or non-critical. XGBoost, and more broadly, ensemble decision tree based models achieve strong performance in tabular machine learning on numerous datasets [2] and are often chosen as a competitive baseline for benchmarking [1]. In addition to XGBoost, we train and evaluate several other deep learning and non-deep learning models including logistic regression (LR), random forest (RF), a multilayer perceptron (MLP), and a ResNet.

3.4 Evaluation

Datasets. We consider two datasets in our evaluation. The first dataset, hereafter referred to as the *Critical/Non-Critical* dataset, consists of ENC changes that have been automatically categorized as critical or non-critical by a programmed rule-based system. The second dataset, hereafter referred to as the *Eyes-On* dataset, consists of ENC changes that do not have a programmed rule. Each change in this dataset was reviewed by human analysts and categorized as *critical* or *non-critical* depending on the change’s significance to navigation. These ENC changes are more complex than those found in the *Critical/Non-Critical* dataset, requiring multi-step decisions, and including scenarios where written rules do not exist and are not easy to formulate. There are approximately 140,000 change instances in the *Critical/Non-Critical* dataset and approximately 9,000 change instances in the *Eyes-On* dataset. The ENC change instances in both datasets are derived from 356 ENC cells and a total of 1308 old-to-new chart pairs covering waters in the United States, Canada, Australia, New Zealand, Norway, Netherlands, and Germany.

Metrics and Protocol. All models are evaluated using standard metrics for binary classification including overall accuracy and macro F1-score. We use five-fold cross-validation, partitioning each dataset into five distinct train-validation splits and reporting the average performance across all folds. We train, evaluate, and report results using both the *Critical/Non-Critical* and *Eyes-On* datasets.

4 Experiments and Results

Spatial Context and Attribute Encoding. Here, we compare results obtained using a baseline data encoding consisting of only ENC object class, ENC relationships, and ENC change type to results obtained with the improved ENC change encoding that includes attributes and spatial context (Table 1).

Hyperparameter Tuning. In Table 2, we present results comparing XGBoost trained with default hyperparameters to XGBoost after hyperparameter tuning. For tuning, we use a tree-structured parzen estimator and median pruning, implemented using the Optuna library, and use a similar search space as [1, 2] for the models

Table 1: XGBoost trained on baseline ENC change encoding and performance impact with spatial context and attribute encoding.

Encoding	Eyes-On		Crit./Non-Crit.	
	Accuracy	F1-score	Accuracy	F1-score
Baseline	.828 ± .004	.790 ± .005	.891 ± .002	.891 ± .002
+ Attributes	.857 ± .007	.837 ± .008	.922 ± .001	.922 ± .002
+ Spatial Context	.901 ± .004	.889 ± .004	.900 ± .002	.900 ± .002
+ Spatial Context and Attributes	.892 ± .005	.879 ± .005	.935 ± .001	.935 ± .001

and training. For both default and tuning, we use the complete ENC change encoding that includes spatial context and attribute encoding. As shown in Table 2, careful hyperparameter tuning of the model improves classification accuracy and F1-score on both datasets. The effect of hyperparameter tuning is consistent with results in [2] where they advocate for carefully tuned baselines.

Table 2: Hyperparameter tuning results with XGBoost. Results are reported as Overall Accuracy/F1-Score.

Model	Eyes-On		Crit./Non-Crit.	
	Default	Tuned	Default	Tuned
XGB	.892/.879	.903/.891	.935/.935	.942/.942

Comparing Baseline Models. Here, we compare XGBoost with the other baseline models for ENC change classification. We include two simple rule-based majority vote classifiers as naive baselines, establishing a lower bound for expected performance. These include *Critical Always*, which predicts the critical class for every change instance, and *Majority by ENC Class*, which predicts the critical class if that ENC object class has more critical examples than non-critical in the training dataset, and the non-critical class otherwise. For the *Critical/Non-Critical* dataset, we also include the performance of the *Rules* system, the deterministic procedure used to auto-classify changes. These rules generated the ground-truth labels for the *Critical/Non-Critical* dataset and therefore achieves 100% accuracy and serves only as a point of reference rather than a meaningful predictive model.

Results are reported for default and tuned baseline models on both datasets in Table 3. All baseline models when trained on the complete ENC data encoding can effectively distinguish between critical and non-critical changes with accuracy greater than 85%. Furthermore, these models achieve much better performance than the two naive rule-based classifiers.

5 Conclusion

Maintaining ENC is a critical and resource-intensive task to accurately reflect navigation and safety-critical information to mariners. In this work, we introduced ENC change classification as a novel Geospatial AI task and introduced an encoding method that successfully translates complex bi-temporal ENC vector data into a structured format for machine learning. By integrating spatial context and attribute encodings into the ENC change representation and hyperparameter-tuning an XGBoost classifier, we achieved 90%

Table 3: Evaluation of baseline models. Results are reported as Overall Accuracy/F1-Score.

Model	Eyes-On		Crit./Non-Crit.	
	Default	Tuned	Default	Tuned
LR	.857/.838	.860/.841	.888/.888	.889/.889
RF	.885/.871	.894/.881	.933/.932	.939/.938
XGB	.892/.879	.903/.891	.935/.935	.942/.942
MLP	.890/.876	.898/.886	.918/.918	.920/.920
ResNet	.892/.877	.904/.893	.928/.928	.929/.929
Critical Always ^a	.654/.395	–	.463/.317	–
Majority by ENC ^a	.829/.794	–	.801/.798	–
Rules ^b	–	–	1.000/1.000	–

^a Naive majority vote classifiers. ^b Rule-based system.

and 94% overall accuracy in classifying ENC changes as critical or non-critical across two operational datasets. These baselines provide a strong foundation for evaluating more sophisticated spatial representation learning techniques. Future work will explore location encodings for heterogeneous geometry types, complex spatial aggregation methods, and spatial graph neural networks to further enhance maritime safety pipelines.

Acknowledgments

We acknowledge that this manuscript has been authored by UT-Battelle, LLC under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>). Research sponsored by the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory, managed by UT-Battelle, LLC, for the U. S. Department of Energy.

References

- [1] Yury Gorishniy, Ivan Rubachev, Valentin Khurlov, and Artem Babenko. 2021. Revisiting deep learning models for tabular data. *Advances in neural information processing systems* 34 (2021), 18932–18943.
- [2] Léo Grinsztajn, Edouard Oyallon, and Gaël Varoquaux. 2022. Why do tree-based models still outperform deep learning on typical tabular data? *Advances in neural information processing systems* 35 (2022), 507–520.
- [3] International Hydrographic Organization. 2000. *Transfer Standard for Digital Hydrographic Data (S-57), Edition 3.1*. International Hydrographic Organization, Monaco. <https://iho.int/uploads/user/pubs/standards/s-57/31Main.pdf> Accessed June 30, 2025.
- [4] Gengchen Mai, Krzysztof Janowicz, Yingjie Hu, Song Gao, Bo Yan, Rui Zhu, Ling Cai, and Ni Lao. 2022. A review of location encoding for GeoAI: methods and applications. *International Journal of Geographical Information Science* 36, 4 (2022), 639–673.
- [5] Giuseppe Masetti, Tyanne Faulkes, and Christos Kastrisios. 2018. Automated Identification of Discrepancies between Nautical Charts and Survey Soundings. *ISPRS International Journal of Geo-Information* 7, 10 (2018), 392. doi:10.3390/ijgi7100392

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009