HOKKAIDO UNIVERSITY

Title	Quantitative representations of COLREG-compliant behaviour from large-scale AIS trajectory data
Author(s)	Huseby, Bang Arne
Relation	Proceedings of the 10th International Symposium on Reliability Engineering and Risk Management (ISRERM 2026), June 28 - July 1, 2026, Sapporo, Japan
Citation	Proceedings of the 10th International Symposium on Reliability Engineering and Risk Management (ISRERM 2026)
Issue Date	2026-07-01
Doc URL	https://hdl.handle.net/2115/100453
Type	conference paper
File Information	C000119_1.pdf

Hokkaido University Collection of Scholarly and Academic Papers : HUSCAP

ISRERM 2026

10th International Symposium on Reliability Engineering and Risk Management (*ISRERM 2026*),
28 June-01 July 2026, Sapporo, Japan
Kitahara, T., Miyamori, Y. & Beer, M. (editors)

Quantitative representations of COLREG-compliant behaviour from large-scale AIS trajectory data

Arne Bang Huseby[1] *

[1]*Department of Mathematics, University of Oslo, P.O.Box 1053 Blindern, 0316 Oslo, Norway*
E-mail: arne@math.uio.no
**Corresponding Author*

The safe integration of autonomous ships into existing maritime traffic requires reliable models of navigational behaviour that conform to the International Regulations for Preventing Collisions at Sea (COLREGs). In this presentation we give a brief introduction to how machine-learning methods can be used to extract quantitative representations of COLREG-compliant behaviour from large-scale AIS trajectory data. We first indicate how to establish encounter-detection and situation-classification procedures that identifies head-on, crossing, and overtaking scenarios, along with give-way and stand-on roles. Combining trajectory embedding and sequence modelling, this can be used to identify patterns characteristic of human navigators in different COLREG situations. Supervised learning, imitation learning, and inverse reinforcement learning may then be applied to infer decision models that map encounter geometry to heading and speed adjustments. This way one can capture implicit navigational preferences based on real traffic. Finally, we outline how these learned behaviour policies can serve as priors or constraint mechanisms in autonomous-ship decision systems, enabling behaviour that is both predictable to human mariners and consistent with maritime collision-avoidance rules.

Keywords: COLREG; AIS trajectory data; Machine-learning; Autonomous ships.

1 Introduction

The development of autonomous maritime vessels has accelerated in recent years due to advances in sensor technologies, control systems, and artificial intelligence. A central challenge in this development is ensuring that autonomous ships can operate safely within existing maritime traffic environments dominated by human-operated vessels.

The International Regulations for Preventing Collisions at Sea (COLREGs) define the rules governing vessel interactions in potential collision situations. However, practical navigation involves interpretation and contextual judgment, and real-world navigational behaviour often follows statistical patterns rather than strictly deterministic rule implementations.

Automatic Identification System (AIS) data provide an extensive source of information about real-world maritime traffic. AIS messages include time-stamped information on vessel position, speed, and heading, allowing reconstruction of vessel trajectories and interactions.

Recent developments in machine learning provide new opportunities to extract behavioural models from such trajectory data. A prominent direction is the use of deep learning for trajectory prediction. For instance, Liu et al. (2020) propose a sequence-based deep neural network model (LSTM) that captures nonlinear vessel dynamics and significantly improves prediction accuracy compared to traditional kinematic models. While such approaches focus on forecasting vessel motion, they typically do not explicitly incorporate encounter geometry or COLREG rules.

In contrast, maritime navigation is fundamentally governed by relative motion and geometric relationships between vessels. Key quantities include relative bearing, range, and velocity, from which the Closest Point of Approach (CPA) and Time to CPA (TCPA) can be derived. These quantities provide a physically interpretable basis for encounter assessment and are widely used for detecting collision risk. In particular, constant relative bearing combined with decreasing range is a classical indicator of potential collision, directly linking geometric features to navigational decision-making.

Several works address decision-making under COLREG constraints using reinforcement learning. In Chen et al. (2022), a safety-aware deep reinforcement learning (DRL) framework is proposed in which COLREG compliance is enforced through reward shaping and constraint mechanisms. However, such approaches typically rely on implicit representations of encounter situations, without explicitly structuring the state space in terms of standardized encounter types such as head-on, crossing, and overtaking.

A complementary line of research focuses on learning navigation behaviour from demonstrations. The work of Brechtel et al. (2014) applies inverse reinforcement learning (IRL) to infer reward functions that explain expert behaviour. This framework is particularly well suited for maritime applications, where navigational decisions are governed by both geometric constraints and regulatory rules.

More recently, IRL has been applied directly to AIS data. In Fang et al. (2022), large-scale trajectory data is used to learn implicit reward structures governing real ship behaviour. However, without an explicit encounter classification framework, the learned policies may conflate fundamentally different navigational situations.

A broad overview of autonomous ship navigation methods is provided by Gunnarsson et al. (2021), highlighting challenges related to perception, decision-making, and COLREG compliance. The survey emphasizes the need for interpretable and rule-consistent models.

In this context, the present paper integrates data-driven learning with a physically grounded encounter representation. Specifically, encounters are first identified and classified using geometric sector rules based on relative bearing and heading differences, distinguishing head-on, crossing, and overtaking situations. This classification is combined with kinematic filters based on CPA and TCPA to ensure that only dynamically relevant interactions are considered. By structuring AIS data according to COLREG-compliant encounter types, the learning problem is decomposed into homogeneous decision contexts. This provides a principled interface between classical maritime navigation theory and modern machine learning methods such as supervised learning, reinforcement learning, and inverse reinforcement learning.

2 Encounter Geometry and COLREG Situation Types

AIS data consist of time-stamped vessel positions, speeds, and headings. From these measurements vessel trajectories can be reconstructed and analysed. A first step is identifying *encounters* between vessels. A central step in modelling navigational behaviour from AIS data is to represent the encounter geometry between two vessels in a mathematically consistent and COLREG-compliant manner. Consider two vessels *i* and *j* with positions:

$$\mathbf{p}_i = (x_i, y_i), \qquad \mathbf{p}_j = (x_j, y_j) \tag{1}$$

and velocity vectors:

$$\mathbf{v}_i = s_i(\cos\psi_i, \sin\psi_i), \qquad \mathbf{v}_j = s_j(\cos\psi_j, \sin\psi_j) \tag{2}$$

where s_i and s_j denotes the speeds, and ψ_i and ψ_j denotes the courses over ground for the two vessels respectively. The relative position and velocity are defined as:

$$\mathbf{r}_{ij} = \mathbf{p}_j - \mathbf{p}_i, \qquad \mathbf{v}_{ij} = \mathbf{v}_j - \mathbf{v}_i \tag{3}$$

The time to closest point of approach (TCPA) is given by:

$$\mathrm{TCPA}_{ij} = -\frac{\mathbf{r}_{ij} \cdot \mathbf{v}_{ij}}{\|\mathbf{v}_{ij}\|^2} \tag{4}$$

and the closest point of approach (CPA) distance is:

$$\mathrm{CPA}_{ij} = \|\mathbf{r}_{ij} + \mathrm{TCPA}_{ij}\mathbf{v}_{ij}\| \tag{5}$$

An encounter is typically considered potentially hazardous when $\mathrm{TCPA_{ij}} > 0$ and $\mathrm{CPA_{ij}}$ is below a safety threshold. Based on the relative geometry, encounters are classified into three canonical COLREG situations: *head-on*, *crossing*, and *overtaking*. A *head-on situation* occurs when two vessels are on reciprocal or nearly reciprocal courses. Both vessels observe each other near the forward direction, implying a symmetric risk. Under COLREG Rule 14, both vessels are required to alter course to starboard. A *crossing situation* occurs when vessel trajectories intersect with collision risk. In this situation the classification depends on whether the target vessel lies on the starboard or port side. COLREG Rule 15 states that if vessel *j* is observed on the starboard side of vessel *i*, then *i* is the give-way vessel. Otherwise, *i* is the stand-on vessel. An overtaking situation occurs when a vessel approaches another from more than $22.5°$ abaft the beam. The overtaking vessel is always the give-way vessel (COLREG Rule 13).

In real AIS data, however, encounter classification cannot rely on idealized geometry alone. Measurements are noisy, vessel motion is non-uniform, and interactions often evolve over time. Robust classification therefore combines kinematic filtering, geometric rules, and temporal consistency constraints.

2.1 *Kinematic filtering*

The first step is to identify candidate vessel pairs that may constitute an encounter. This is done using distance- and motion-based filters derived from relative kinematics. Thus, we let:

$$d_{ij} = \|\mathbf{r}_{ij}\| \tag{6}$$

denote the inter-vessel distance for vessels *i* and *j*. Candidate encounters are typically selected for vessels where:

$$d_{ij} \leq d_{\max} \tag{7}$$

where $d_{\max}$ is a predefined detection radius (e.g., 3–6 nautical miles depending on traffic density). Collision risk is further assessed using TCPA and CPA. More specifically, the vessel pair (i, j) is retained if:

$$\mathrm{TCPA}_{ij} > 0, \qquad \mathrm{CPA}_{ij} \leq d_{\mathrm{CPA}} \tag{8}$$

That is, the vessel pair (i, j) is retained if the time to closest point of approach has not yet passed, and the closest point of approach is below a given safety threshold, denoted d_{CPA} (e.g., 0.5–1.0 nautical miles).

Additional filters may include:

- Minimum speed threshold (to exclude drifting or anchored vessels)
- Maximum relative speed (to remove implausible AIS artefacts)
- Course stability (to avoid erratic trajectories)

2.2 *Geometric sector rules*

Once candidate encounters are identified, the encounter type is determined using relative geometry. The relative bearing of vessel *j* seen from vessel *i* is defined as:

$$\beta_{ij} = \mathrm{atan2}(y_j - y_i, x_j - x_i) - \psi_i \tag{9}$$

Thus, the relative bearing of vessel *j* seen from vessel *i* is the angle between the velocity vector $\mathbf{v}_i$ and the relative position vector $\mathbf{r}_{ij}$. This is illustrated in Figure 1. Note that in this figure the relative bearing of vessel *k* seen from vessel *i* is negative since vessel *k* is on starboard side, while the relative bearing of vessel *j* seen from vessel *i* is positive since *j* is on port side.

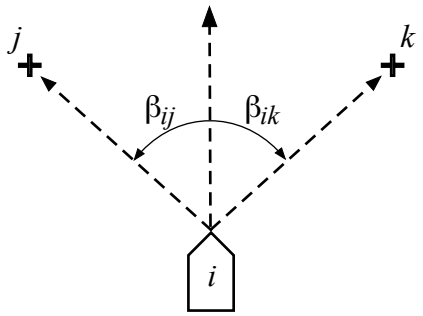


Figure 1. Relative bearing of vessels *j* and *k* seen from vessel *i*.

Classification is then based on angular sectors:

Head-on:

$$|\psi_i - \psi_j| > \pi - \epsilon_\psi, \quad |\beta_{ij}| < \epsilon_\beta \tag{10}$$

Crossing:

$$\epsilon_\beta < |\beta_{ij}| < \pi - \epsilon_{\mathrm{aft}} \tag{11}$$

The give-way/stand-on roles determined by the sign of the relative bearing.

Overtaking:

$$|\beta_{ij}| > \pi - \epsilon_{\text{aft}}, \quad |\psi_i - \psi_j| < \epsilon_\psi \tag{12}$$

Typical threshold values are

$$\epsilon_\psi \approx 15^\circ, \quad \epsilon_\beta \approx 10^\circ, \quad \epsilon_{\text{aft}} \approx 22.5^\circ \tag{13}$$

These sector-based rules approximate the COLREG definitions while remaining computationally efficient for large AIS datasets.

2.3 *Temporal consistency constraints*

AIS-based classifications based on instantaneous geometry are highly sensitive to noise and small fluctuations in course or position. Therefore, temporal filtering is applied to ensure consistency over time. An encounter classification is assumed valid if the following conditions are met:

- The classification remains unchanged over a minimum duration (e.g., 2–5 minutes)
- TCPA decreases over an initial phase (indicating that the vessels are getting closer)
- CPA remains below threshold throughout the encounter window

It is also common to apply various smoothing techniques on the AIS data, such as moving averages of heading and speed, majority voting over sliding windows, and hidden Markov models for state stabilization.

Encounter start times are defined by the point of time when two vessels enter within the detection radius, $d_{\max}$, while end times are defined either by the point of time when the distance between the vessels become larger than $d_{\max}$, or when TCPA flips negative indicating that the vessels are moving away from each other.

3 Inverse Reinforcement Learning for AIS-Based Behaviour Modelling.

A central objective in modelling maritime navigation behaviour is to infer the latent decision-making principles that govern human-operated vessels. While conventional reinforcement learning (RL) assumes a known reward function and seeks an optimal policy, *inverse reinforcement learning* (IRL) addresses the inverse problem: given observed behaviour, infer the underlying reward function that renders such behaviour approximately optimal. See Ng and Russel (2000). For a recent survey on IRL-related reserach, see Arora and Doshi (2021).

In order to formalize this we use familiar concepts from reinforcement learning, such as the state space, denoted S, and the action space denoted $\mathcal{A}$. We then introduce the notion of a (discretized) vessel *trajectory* as a sequence:

$$\tau = (s_0, a_0, s_1, a_1, \ldots, s_T) \tag{14}$$

where $s_t \in S$ denotes the system state at time t, and $a_t \in \mathcal{A}$ the corresponding control action. AIS data provides observations of vessel motion from which such trajectories can be reconstructed

We assume that human navigators follow a policy π_{human} that is (approximately) optimal with respect to an unknown reward function $R(s, a)$. The IRL problem is to estimate the reward function R such that the induced optimal policy π^*_R satisfies:

$$\pi_{\text{human}} \approx \pi^*_R \tag{15}$$

Obviously, the space of possible reward functions is very large. Thus, in the general case it may be virtually impossible to obtain a stable estimate. In order to simplify the problem, a common technique within IRL is to assume that the reward function admits a linear parametrisation:

$$R(s, a) = \theta^\top \phi(s, a) \tag{16}$$

where $\phi(s, a)$ is a vector of predefined features, and θ is an unknown parameter vector to be estimated based on the available AIS data.

3.1 *MDP Representation from AIS Data*

To apply IRL, AIS trajectories must ideally be embedded into a Markov decision process (MDP). The state representation typically includes:

- **Own-ship kinematics:** Position, speed, and heading
- **Relative motion features:** Relative position and velocity to nearby vessels
- **Encounter indicators:** CPA and TCPA
- **Situational context:** Encounter class (head-on, crossing, overtaking) and COLREG role (give-way or stand-on)

Actions are inferred from successive AIS observations and typically represented as changes in heading and speed, i.e.,

$$a_t = (\Delta\psi_t, \Delta v_t) \tag{17}$$

It is usually necessary to keep the number of possible actions as low as possible. Thus, the action space $\mathcal{A}$ is typically strongly discretized. However, this simplification often works well, especially if this is balanced with short timesteps.

3.2 *Reward Feature Design*

The feature vector $\phi(s, a)$ encodes relevant navigational objectives and constraints. Typical feature classes include:

- **Safety:** Penalties for small CPA or high collision risk
- **COLREG compliance:** Adherence to prescribed manoeuvres (e.g., starboard alteration in head-on situations)
- **Efficiency:** Deviation from intended route or speed profile
- **Smoothness:** Penalties for abrupt or erratic manoeuvres

This structured feature design enables interpretable reward learning consistent with maritime regulations.

3.3 *Maximum Entropy IRL*

To account for variability and suboptimality in human behaviour, we recommend adopting the maximum entropy IRL framework. See Ziebart et al. (2008). Under this model, the probability measure of a trajectory is given by:

$$P(\tau \mid \theta) \propto \exp\left(\sum_{t=0}^{T} R(s_t, a_t)\right) = \exp\left(\sum_{t=0}^{T} \theta^\top \phi(s_t, a_t)\right) \tag{18}$$

The parameter vector θ can then be estimated by maximizing the likelihood of observed trajectories. This formulation naturally captures stochasticity in navigation behaviour and avoids assuming strict optimality.

3.4 *Multi-Agent Considerations*

Maritime navigation inherently involves multiple interacting agents. In practice, several approximations are employed. The most obvious solution to this poroblem is simply to apply a single-agent approach where other vessels are treated as part of the environment. Another option is conditional modelling where policies are learned separately for different encounter classes and roles. The most advanced approach is a full multi-agent IRL formulations based on equilibrium concepts and game-theory. Among these, conditional modelling provides a robust and scalable approach for large AIS datasets.

3.5 *Learning pipeline*

The proposed methodology for applying IRL to AIS data can now be summarized as follows:

1. **Data preprocessing:** Cleaning, interpolation, and temporal alignment of AIS trajectories
2. **Encounter detection:** Identification and classification of head-on, crossing, and overtaking scenarios
3. **Trajectory segmentation:** Extraction of decision-relevant segments (e.g., within a TCPA threshold)
4. **State–action construction:** Computation of state vectors and inferred control actions
5. **Feature computation:** Evaluation of the vector $\phi(s, a)$ for all samples
6. **IRL optimisation:** Estimating the parameter vector θ using the maximum entropy method
7. **Policy reconstruction:** Solving the forward RL problem under the learned reward

IRL provides a principled framework for extracting interpretable decision models from AIS data. By recovering the implicit reward structure underlying human navigation, it enables the development of autonomous systems that exhibit realistic, safe, and regulation-compliant behaviour. Moreover, the approach facilitates quantitative analysis of trade-offs between safety, efficiency, and rule compliance in real-world maritime traffic.

4 Integration into Autonomous Ship Decision Systems

The behavioural models derived from AIS data can be incorporated into autonomous ship navigation systems in several ways. Learned behavioural models can serve as behavioural priors guiding the decision-making process of autonomous navigation systems. This ensures that autonomous vessels behave in ways that are predictable to human mariners. Moreover, learned policies can be incorporated as constraints in trajectory planning algorithms, ensuring that generated manoeuvres remain consistent with typical COLREG-compliant behaviour. Finally, behavioural models can be used to generate realistic traffic simulations for testing and validating autonomous navigation algorithms. Combining rule-based reasoning with data-driven behavioural models provides a promising approach for safe operation in mixed human autonomous maritime traffic environments.

5 Conclusions

While AIS-based machine learning offers powerful tools for modelling maritime behaviour, several challenges remain. AIS data may contain noise, missing messages, or incomplete trajectory segments. Furthermore, not all observed manoeuvres strictly comply with COLREG rules, which may introduce ambiguity into the learning process. Another important issue concerns interpretability. For safety-critical systems such as maritime navigation, decision models must be transparent and verifiable. Future research will therefore need to combine data-driven methods with formal rule-based reasoning and safety verification techniques.

This paper has presented an overview of machine-learning approaches for extracting COLREG-compliant navigational behaviour from AIS trajectory data. The proposed methodological framework includes encounter detection, trajectory representation learning, behavioural pattern discovery, and policy inference. By leveraging large-scale maritime traffic data, these methods enable the development of behavioural models that reflect real-world navigational practice. Such models can support the development of autonomous ship navigation systems capable of operating safely and predictably in mixed maritime traffic environments.

Acknowledgments

This research is carried out as part of the project Norwegian Maritime AI Centre (MAI) funded by the Research Council of Norway.

References

Bakdi, A., I. K. Glad, E. Vanem, Testbed Scenario Design Exploiting Traffic Big Data for Autonomous Ship Trials Under Multiple Conflicts With Collision/Grounding Risks and Spatio-Temporal Dependencies, *IEEE Trans. on Intelligent Transportation Systems*, **22**:12, 7914–7930, 2021.

Brechtel, S., T. Gindele, and R. Dillmann, Learning navigation behaviors from expert demonstrations, In *IEEE International Conference on Robotics and Automation*, 409–415, 2014.

Chen, Y., H. Wang, and J. Li, Safety-aware deep reinforement learning for colreg-compliant ship navigation. *Ocean Engineering* **259**, 111848, 2022.

Fang, R., J. Sun, and Y. Dong, Inverse reinforcement learning for ship navigation using ais data*, Expert Systems with Applications* **195**, 116569, 2022.

Gunnarsson, V., T. Porathe, and M. Baldauf, Autonomous ship navigation: A review of methods, *IEEE Access,* **9**, 140000–140025, 2021.

Huang, Y., P.H.A.J.M. van Gelder, Time-Varying Risk Measurement for Ship Collision Prevention, *Risk Analysis*, **40**:1, 24–42, 2020.

Huang, Y., P.H.A.J.M. van Gelder, M. B. Mendel, Imminent ships collision risk assessment based on velocity obstacle, In *Risk, Reliability and Safety: Innovating Theory and Practice – Walls, Revie & Bedford (Eds)*, 693–700, 2017.

Huang, Y., P.H.A.J.M. van Gelder, Y. Wen, Velocity obstacle algorithms for collision prevention at sea, *Ocean Engineering*, **151**, 308–321, 2018.

Huang, Y., L. Chen, R. R. Negenborn, P.H.A.J.M. van Gelder, A ship collision avoidance system for human-machine cooperation during collision avoidance, *Ocean Engineering*, **217**, 308–321, 107913, 2020.

Huang, Y., P.H.A.J.M. van Gelder, Collision risk measure for triggering evasive actions of maritime autonomous surface ships, *Safety Science*, **127**, 104708, 2020.

Liu, Z., W. Zhang, A deep learning-based method for ship behavior prediction using ais data, *Ocean Engineering,* **209**, 107478, 2020.

Ng, A. Y., S. Russel, Algorithms for Inverse Reinforcement Learning, In *Proceedings of the Seventeenth International Conference on Machine Learning (ICML),* 663–670, 2000.

Arora, S., P. Doshi, A Survey of Inverse Reinforcement Learning: Challenges, Methods and Progress, *Artificial Intelligence*, **297**, 103500, 2021.

Vestre, A., A. Bakdi, E. Vanem, Ø. Engelhardsen, AIS-base near-collision database generation and analysis of real collosion avoidance manoeuvres, *J. of Navigation*, **74**:5, 985–1008, 2021.

Wróbel, K., M. Gil, Y. Huang, R. Wawruch, The Vagueness of COLREG versus Collision Avoidance Techniques—A Discussion on the Current State and Future Challenges Concerning the Operation of Autonomous Ships, *Sustainability*, **14**, 16516, 2022.

Ziebart, B. D., A. Maas, J. A. Bagnell, A. K. Day, Maximum Entropy Inverse Reinforcement Learning, In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI)*, 1433–1438, 2008.