Journal of *Marine Science and Engineering*

MDPI

Article

An Accident-Grounded Ontology of COLREGs and Seamanship Practice as a Basis for Explainable Collision Avoidance

Jae-Yong Lee [1] and Ho Namgung [2,*]

[1] Department of Maritime Transportation System, Mokpo National Maritime University, Mokpo 58628, Republic of Korea; naraan42@gmail.com
[2] Department of Maritime Industry Convergence, Mokpo National Maritime University, Mokpo 58628, Republic of Korea
* Correspondence: ngh2009@mmu.ac.kr; Tel.: +82-61-240-7162

Abstract

The International Regulations for Preventing Collisions at Sea (COLREGs) assume knowledge that the written rules do not state. This study builds a collision-avoidance ontology from the COLREGs and their commentary, requirement studies and 4026 Korean and international collision records. In a Web Ontology Language (OWL) 2 knowledge base, 37 inference rules derive the applicable COLREGs rules and case-level obligations from the stated encounter type, visibility, navigational status, and fishing activity, with location and vessel type where reported; each derivation traces to its facts and rule. Twelve accident-derived items form 3 of 56 requirements and refine 7 others. Three accident-derived findings are exploratory: stand-on action under Rule 17(b) (16 rulings), conflicting multi-ship roles (six records of five casualties) and conduct attached to the ordinary practice of seamen (262 rulings). The derived rule sets contain 86.7% of the conduct rules cited in the Korean records (an in-sample check) and 70.8% of those in the international records (64.0% for explicit citations). A lookup from the stated facts gives similar or higher coverage, so these figures measure consistency with the record, not predictive skill. Accident records can add knowledge beyond the rule text, and the ontology may serve as a basis for explainable inference.

Keywords: collision avoidance; COLREGs; ontology; maritime accident investigation; multi-ship encounter; seamanship practice

Check for updates

Academic Editor: Sarat Chandra Mohapatra

Received: 5 October 2026
Revised: 3 November 2026
Accepted: 4 November 2026
Published: 7 November 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 International Maritime Organization (IMO) has reviewed how its instruments, including the collision regulations, apply to maritime autonomous surface ships [1], and collision avoidance by such ships is governed by the International Regulations for Preventing Collisions at Sea (COLREGs) [2]. The rules were written to be applied with human judgment. Several obligations use qualitative terms such as "ample time" and "substantial action" (Rules 8 and 16), whose meaning depends on the circumstances of each encounter [3,4]. An autonomous system therefore needs both the rule text and the knowledge that navigators use when applying the rules.

This raises two problems. First, quantitative models of the rules and of seamanship assess each target ship one at a time against own ship [5], and the rules give no explicit procedure for resolving conflicting obligations toward several ships at once [6]. Second, part of the knowledge that governs conduct is unwritten. Rule 2 refers to the ordinary practice of seamen and to departures required by special circumstances, and tribunal decisions have

been criticized for invoking the ordinary-practice clause to assign fault without naming a specific breach [7]. Both problems call for evidence of how the rules are applied in practice.

Machine-readable collision-avoidance knowledge has developed in three forms. Rule-based systems encode the rules and expert judgment as hand-written, fuzzy, or knowledge-base rules [8–12]. Formal specifications express the COLREGs as behavior protocols, compliance metrics, and temporal-logic formulas [13–17]. Ontologies and knowledge bases represent collision-avoidance knowledge and navigation rules and, in most cases, reason over them with description logics (DL), fuzzy relations, rule languages, or probabilistic models [18–24]. A search in English and Korean combined the terms "collision avoidance", "COLREGs", "ontology", "knowledge graph", "expert system", and "formal specification" and was extended by backward and forward citation searching. It retained 31 studies, published from 1989 to 2026, that represent the conduct rules in a form a machine can reason over. Three test their methods on a few past collisions, and the rest on constructed or illustrative scenarios, consistency checks, sea trials, or automatic identification system (AIS) records of routine traffic. None compares its inferences with the rules cited in a corpus of adjudicated collisions.

A second line of research analyzes maritime accident investigation reports. Studies that coded collision reports with the Human Factors Analysis and Classification System (HFACS) identified recurring human and organizational factors [25], and a word-propensity analysis of Korean adjudications identified latent causes of complex collisions [26]. Later studies build knowledge graphs (KGs) from report corpora with ontology-guided information extraction (IE) [27] and Bidirectional Encoder Representations from Transformers (BERT) models for named-entity recognition (NER) [28]. Large language models (LLMs) now assist report screening and factor identification [29], and a retrieval-augmented generation (RAG) assistant has been demonstrated for question answering over investigation reports [30]. Most of these studies produce cause frequencies, co-occurrence networks, and queryable graphs. The same search retained 34 such studies, published from 2012 to 2026. One of them applies text mining to navigational communication in simulated collision situations, and none links each case in its corpus to the rules that applied.

Few studies combine the two lines. Mantani examined numerical standards for the Japanese collision rules with questionnaires, simulator experiments, and tribunal adjudications without building a machine-readable knowledge base [31]. Shu et al. coupled an ontology with a Bayesian network (BN) to apportion liability [32]. Zhang et al. built and reasoned over a collision KG of autonomous ships with an enhanced BERT model [33]. No reviewed study extracts knowledge from a large set of accident reports, represents it in a COLREGs ontology and then checks the ontology against those accidents. Table 1 summarizes each line of prior work and the limitation that this study addresses.

Table 1. Prior work on collision-avoidance knowledge systems and accident report analysis.

Category	Studies	Method/ Technology	Knowledge Source	Validation Basis	Limitation Addressed Here
Rule-based expert systems	[8–12]	Production rules, Prolog, fuzzy rules, rule-guided search, optimization	Expert knowledge, COLREGs text	Simulation and laboratory scenarios; field tests on two vessels	Hand-coded knowledge; no accident grounding

Table 1. *Cont.*

Category	Studies	Method/ Technology	Knowledge Source	Validation Basis	Limitation Addressed Here
Formal specification of COLREGs	[13–17]	Protocol encoding, LTL/MTL/STL, model checking	COLREGs text	Illustrative scenarios with model checking; simulated encounters; on-water trials; AIS replay of routine traffic	Rules mostly evaluated per target pair; not tested on collision cases
Ontology and knowledge-base systems	[18–24]	OWL, SWRL, RDF, fuzzy relational products, type-2 fuzzy ontology, dynamic BN	COLREGs text, local navigation rules, expert seamanship thresholds, IMO communication phrases	Simulation; illustrative scenarios; consistency checking; one-day AIS replay; replay of one investigated collision	Scenario or routine-traffic evaluation; one accident case replayed; seamanship practice not drawn from adjudications
Accident report analysis and text mining	[25–30]	HFACS coding, word-propensity statistics, ontology-guided IE, BERT NER, LLM extraction, RAG	Investigation reports and adjudications	Descriptive and multivariate statistics; retrieval demonstration; extraction F1; expert ground truth	Descriptive or probabilistic outputs; no rule-level conduct reasoning
Nearest fusion attempts	[31–33]	Manual statistics; ontology + BN; KG link reasoning	Questionnaires and tribunal rulings; investigation reports	Questionnaire statistics and simulator experiment; liability reproduction on sample cases; extraction F1	Non-machine-readable, probabilistic, or associative outputs; no deterministic rule inference
This study	—	OWL 2 DL ontology with SWRL inference layer	Rule text, requirements literature, Korean adjudications, and international reports	4026 adjudicated and investigated collisions; transfer test on the international tier	—

BERT: Bidirectional Encoder Representations from Transformers, BN: Bayesian network, DL: description logic, HFACS: Human Factors Analysis and Classification System, IE: information extraction, IMO: International Maritime Organization, KG: knowledge graph, LLM: large language model, LTL: linear temporal logic, MTL: metric temporal logic, NER: named-entity recognition, OWL: Web Ontology Language, RAG: retrieval-augmented generation, RDF: Resource Description Framework, STL: signal temporal logic, SWRL: Semantic Web Rule Language.

AIS trajectories can flag geometric risk, such as a steady bearing, a closing range, and a small predicted closest point of approach (CPA), but they record only how the ships moved. Investigation reports and adjudications record what led to a collision and why, so they form the corpus of this study.

This study combines three knowledge sources: the collision-avoidance requirements in the literature, the COLREGs with their established interpretation, and 4026 ship-to-ship collision records. The records comprise 709 Korean adjudications published in 2016–2025 and 3317 international investigation reports from 77 authorities and flag administrations, all

processed with the same parsing pipeline. The scope is collision avoidance between vessels under way (Rule 3(i)). Cases involving anchored, moored, or berthing vessels, governed by Rules 30 and 35 and by port practice, are kept in the ontology and the statistics but lie outside this scope. In Section 3.3, accident-derived requirements are drawn from the Korean adjudications and checked against the international reports.

A deterministic inference layer derives the applicable rules and required actions from the stated facts on which the Steering and Sailing Rules depend. Rules 11 and 19 depend on sight and visibility, Rules 13 to 15 depend on encounter type, and Rule 18 depends on navigational status and fishing (Rules 3(c) and 3(d)). Location in a narrow channel or traffic separation scheme (TSS) adds Rules 9 and 10, and vessel type and anchored or drifting status add the lights, signal, and anchored-vessel rules. Every step of the layer can be traced to a class (a named category) or an axiom (a logical statement) of the ontology. The ontology also records knowledge beyond the written rules, such as the conduct for which tribunals invoke the ordinary practice of seamen. The special-circumstances clause of Rule 2(b) is open by design; the ontology marks only where it was invoked (Section 2.3).

This study tests one proposition: an ontology built on the rule text and on adjudicated accidents can be consistent with the rules applied in real collisions and can represent collision-avoidance knowledge that the written rules do not state.

This study makes three contributions. First, a collision-avoidance ontology is built on the rule text, the requirements literature, and a decade of Korean adjudications, with a verified mapping from statutory articles to COLREGs rules. The rules that its deterministic inference layer finds applicable are compared with those cited in the Korean adjudications and, as a transfer test on records that did not inform the layer, in the international investigation reports. Second, conduct knowledge that rule-text ontologies omit is documented as exploratory evidence. The evidence comprises a manually coded taxonomy of the conduct for which tribunals invoke the ordinary practice of seamen (262 rulings), the recorded or adjudged timing of the stand-on vessel's action under Rule 17(b) in 16 rulings, and an initial taxonomy of conflicting roles in six multi-ship records. These findings are turned into accident-derived requirement items. Third, 56 collision-avoidance requirements are consolidated from the literature, the rule text, and the accident record.

2. Materials and Methods

2.1. Data: A Tri-Source Corpus of Collision Knowledge

The three sources were parsed and structured before being loaded into the ontology; Table 2 gives their origin and the units extracted. The first is the rule text: the COLREGs [2] with the indicative values given in the commentary of Cockcroft and Lameijer [3]. The Korean adjudications cite articles of the national statute, which enacts the COLREGs with the same content under its own article numbers; a verified article-to-rule mapping converts each article to the corresponding rule. The mapping also covers the renumbering made when the navigation provisions moved to the Maritime Traffic Safety Act in 2024.

Table 2. Composition of the three knowledge sources.

Knowledge Source	Provenance	Units Extracted
Rule text	COLREGs [2]; A Guide to the Collision Avoidance Rules [3]	20 rule individuals; 5 interpretation-parameter individuals
Requirements literature	Fifteen studies on collision-avoidance requirements and evaluation [15,34–47]	234 requirement items

Table 2. *Cont.*

Knowledge Source	Provenance	Units Extracted
Korean adjudications (knowledge source)	KMST—Busan 249, Mokpo 171, Incheon 124, Central 92, Donghae 73	709 collision adjudications
International reports (knowledge source; inference-layer transfer test)	77 authorities and flag administrations—JTSB 1980; GISIS depository (56 administrations) 538; China MSA 337; DSB 71; NTSB 71; TTSB 53; MAIB 46; CIAIM 31; twenty further authorities 190	3317 collision reports
Accident record, total	—	4026 collision cases

CIAIM: Standing Commission for Maritime Accident and Incident Investigations (Spain), DSB: Dutch Safety Board, GISIS: Global Integrated Shipping Information System, JTSB: Japan Transport Safety Board, KMST: Korea Maritime Safety Tribunal, MAIB: Marine Accident Investigation Branch (UK), MSA: Maritime Safety Administration, NTSB: National Transportation Safety Board (USA), TTSB: Taiwan Transportation Safety Board.

The second source is the literature on collision-avoidance requirements and their evaluation [15,34–47]. It comprises 15 studies, identified by a search separate from the reviews of Section 1. A study was included when it states requirements or evaluation criteria for COLREGs-compliant collision avoidance as its main contribution or as the core of its method, at a level at which each item can be traced to a source sentence. Three of the 31 rule-representation studies of Section 1 meet this criterion and are included [15,35,43]; the others encode the rules without stating requirements, and the 34 studies that mine accident reports identify causes, not requirements. Table A1 in Appendix 5 details each study. Six studies form the core of the set: Ohn and Namgung [34], Woerner et al. [15], Hagen et al. [37,38], Gleeson et al. [35], and Tengesdal and Johansen [36]. Ohn and Namgung [34] is co-authored by one of the present authors; it contributes to 24 of the 56 requirements and is the only literature source of four of them. The others cover collaborative avoidance, reviews, operator tasks and human failure events, formal traffic-rule predicates, avoidance algorithms, and quantitative encounter thresholds [39–47].

Items were extracted one study at a time by an LLM (Section 2.2) with a fixed prompt that returns each stated requirement with its decision stage, mapped rule, any quantitative threshold, and the verbatim source sentence. Candidates not supported by that sentence were discarded, leaving 323 items (Table A1); normalizing the terminology and merging duplicates gave 234 distinct items. Each item is stored in the ontology as an instance, termed a requirement individual, with its source and the extracted attributes.

The items were then consolidated into collision-avoidance requirements: statements of what a collision-avoidance system, or the navigator it supports, must perceive, decide or do in an encounter, as given by the literature, the rule text and the adjudications. Seventy-one items do not concern how vessels perceive or act in ship-to-ship encounters; most deal with system architecture and communications (18), the mechanics of evaluation programs (13), and the reward design of learning methods (9). The other 163 items were merged when they shared the same COLREGs provision, obligation, and ontology class, or described a single determination or monitoring task. The admitted accident-derived items of Section 3.3 were merged in the same way. The rule text added the restricted-visibility regime (Rule 19(a)) and the duty not to impede (Rule 8(f)). The multi-ship rulings of Section 3.3 added early action by the vessel better able to maneuver, a practice of seamanship that is not a COLREGs priority rule.

The 163 items form 50 requirements. Three accident-derived items form new requirements (CA20, CA26, and CA27), and nine join seven existing ones; the rule text adds two more and the multi-ship rulings add one, giving 56 requirements (Appendix 5, Table A2).

The requirements are grouped into four decision stages: situation awareness (SA), risk assessment (RA), encounter determination (ED), and collision avoidance (CA), with 8, 9, 7, and 32 requirements, respectively.

Each requirement is divided into evaluation items, the parts of its duty that a test of an autonomous system would score separately. The 158 items include six for CPA-based risk (RA3).

The third source is the accident record, which has two tiers. The Korean tier contains every ship-to-ship collision adjudication published by the Korea Maritime Safety Tribunal (KMST) [48] in 2016–2025 and retrieved from its public website: 709 adjudications (also called rulings below) from the central tribunal and four district tribunals. When a casualty was heard twice, both adjudications are kept because the appeal may add findings; 25 casualties appear at both stages, and Section 3.2 reports the effect of counting each once. Cases were selected by publication date because an investigation may take several years.

The international tier comprises 3317 published investigation reports of ship-to-ship collisions issued over the same period by 77 foreign authorities and flag administrations. They were retrieved from the report registers of the authorities and from the Global Integrated Shipping Information System (GISIS) of the IMO [49]. The Japan Transport Safety Board (JTSB) supplies 1980 because it investigates and publishes every collision, including those involving small craft. The rest are 538 reports deposited in GISIS by 56 flag and coastal administrations, 337 reports of the Chinese maritime safety administrations and 462 reports of other authorities (Table 2). When several authorities investigated the same casualty, all reports are kept and linked, since investigators may attribute different rules to the same facts. Table A3 gives the vessel types involved. A fishing vessel is a party to 72.2% of the Korean and 42.9% of the international collisions, and 36.4% and 13.5%, respectively, are collisions between two fishing vessels.

Both tiers are stored in the ontology as case individuals and support the accident-derived requirements of Section 3.3. The inference axioms of Section 2.4 were written from the rule text and the Korean tier only, so in Section 3.2, the international tier also tests whether the axioms remain valid for reports not used to write them. The tiers are kept apart because the documents differ in purpose. Korean adjudications are liability rulings that must assign stand-on and give-way roles and cite the governing rule, whereas international reports are no-blame safety investigations that usually avoid such attribution. Pooling them would dilute the explicit role and rule assignments of the rulings and remove the transfer test.

2.2. Keyword Parsing and Structured Extraction

The records are written mainly in Korean, Japanese, English, and Chinese, and the remainder is written in 11 other languages. A lexicon of 20 concept categories was defined as regular-expression patterns in 15 languages, each based on the official terminology of the national COLREGs enactment in that language. The categories span the SA, RA, and CA stages, from look-out and risk of collision to the Rule 2 concepts of seamanship practice and special circumstances. A second pattern set captures rule citations in the COLREGs, the Korean statutes, the Chinese collision regulations, and the Japanese enactment. Each pattern starts from the name of the instrument and collects every article number in the same sentence, so that lists of provisions are captured in full. The JTSB seldom cites article numbers, and when it does, the citation appears in a background or recommendation passage, not in the causal finding. For JTSB reports, the rules are therefore inferred from the wording of the causal sections (clause-heading attribution), using the criteria in Appendix 5, Table A4.

Article IX of the Convention makes the English and French texts equally authentic [2]. The rule text and the interpretation parameters of the ontology follow the English text

and one English-language commentary [3], whereas detection uses the terminology of each national enactment. Explicit citations do not depend on wording because each is resolved to a COLREGs rule number through the article-to-rule correspondence of the enactment, whereas the JTSB attributions rely on the Japanese wording of the causal findings (Appendix 5, Table A4). Concept detection and the derived requirements can be affected where a national term departs from the English text, as with the Korean term for the Rule 17(b) action (Section 3.3.2). The national enactments and the authentic French text were not compared systematically.

The concept and citation counts come from regular-expression matching and the mapping table; the other counts come from the consolidation of requirements and from the ontology. Multi-ship cases, accident-derived items, the taxonomies of Section 3.3, and the reading check of the keyword screen were reviewed and coded by one author. Claude Opus 5.0 and Claude Fable 5.0 (Anthropic, 2026) assisted in retrieving the international reports, parsing the accident-report corpus and extracting the literature requirement items. Model outputs were checked and corrected before use. For a concept category c with pattern set P_c and a document d with extracted text T_d, Equation (1) gives the match count as

$$m_{c,d} = \sum_{p \in P_c} count(p, T_d) \quad (1)$$

Equation (2) then gives the document frequency within each corpus as

$$DF_s(c) = \frac{|\{ d \in D_s : m_{c,d} > 0 \}|}{|D_s|} \quad (2)$$

where D_s is the document set of corpus $s \in \{KR, INTL\}$, which denote the Korean and international tiers, respectively.

Figure 1 shows the pipeline. Both tiers fill the assertion box (ABox), which holds the individuals of the ontology. The parsed output of each case is stored in a standard record of 34 fields.

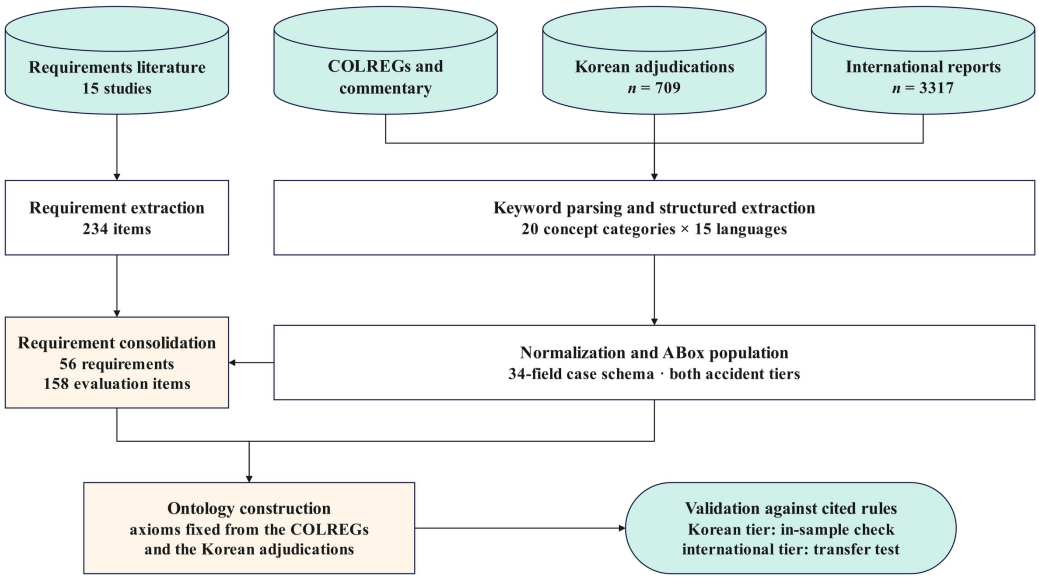


Figure 1. Pipeline from knowledge acquisition to ontology construction and validation.

Three observations shaped the design. Table A5 in Appendix 5 gives the number of cases that mention each category in the two tiers. First, SA and RA concepts dominate both tiers: look-out is mentioned in 89% of the Korean adjudications and 68% of the international reports, and Rule 5 leads the citations at 44% and 55%. Second, the tiers differ most in stand-on conduct (69.5% versus 9.1%), seamanship practice (38% versus 5%), and Rule 2 citations

(11% versus 2%), and much of the knowledge that Rule 2 leaves to seamanship is therefore found in the Korean adjudications. Third, narrow-channel, TSS, and restricted-visibility provisions are rarely mentioned in either tier, so Sections 2.3 and 2.4 encode them explicitly as axioms.

Two procedures verified the parsing. In a stratified sample of 50 documents, each match was re-extracted with its surrounding text, and 98.3% of the matches were correct (instance-level precision). Twenty fully annotated documents, 10 from each tier, gave a document-level precision of 0.97 and a recall of 0.81 with a 95% confidence interval (CI) of 0.60–0.93; recall was 0.80 for the Korean and 0.82 for the international documents, so missed mentions may lower the concept frequencies. Three of the case-level fields that drive the inference (encounter type, visibility, and fishing state) agreed with the source text in 47, 47, and 46 of the 50 documents. A separate screen searches each Korean ruling for 14 fault keywords, each mapped to one of the four decision stages of Section 2.3 (Appendix 5, Table A6). To check the screen, the causal findings of all 709 rulings were also coded into the 14 fault categories by reading, each counted once per ruling. Of the 2438 findings in the four stages, 39.4% are SA, 14.3% are RA, 1% are ED, and 45.3% are CA, within 4 percentage points of the screen shares (Section 3.1).

2.3. Ontology Design

The ontology was built in the Web Ontology Language (OWL) 2 with the Protégé 5.6.9 editor. It extends a COLREGs class model with three layers derived from the parsed evidence (Table 3). Classes, properties, and axioms form the terminological box, which defines the vocabulary of the ontology. Figure 2 shows the core classes and properties of the four layers and their links to the case individuals.

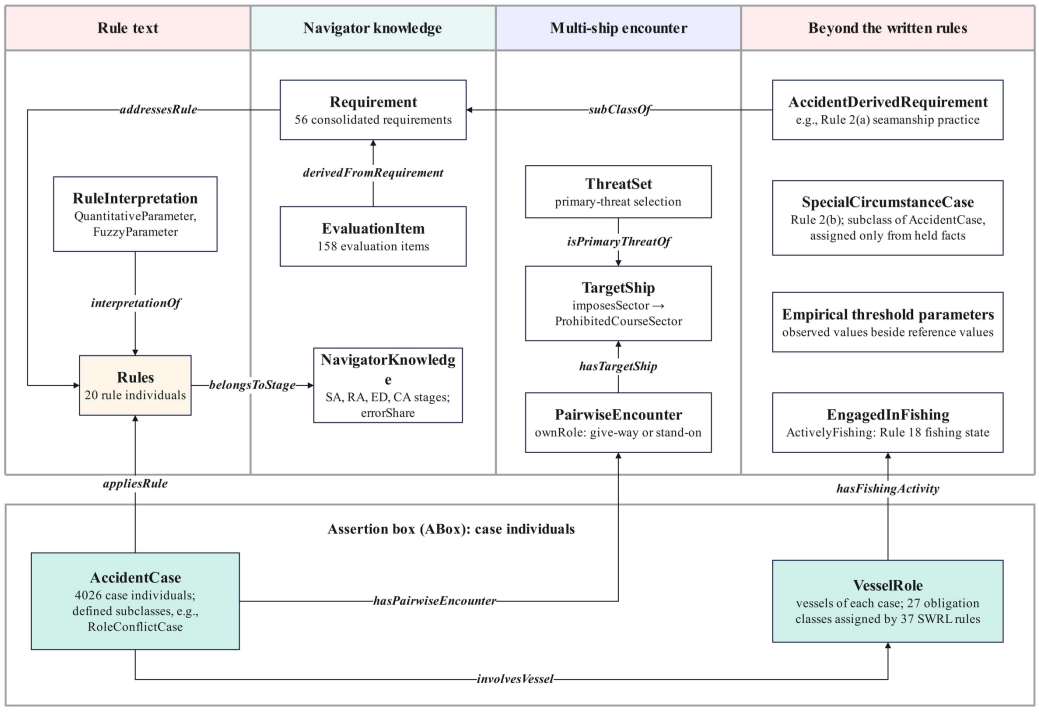


Figure 2. Core classes and properties of the four ontology layers. Italic labels are object properties; teal boxes hold the case and vessel individuals of the ABox.

Table 3. Layers of the collision-avoidance ontology.

Layer	Representative Classes and Properties	Knowledge Encoded	Provenance
Rule text	Rules (20 individuals), RuleInterpretation, QuantitativeParameter and FuzzyParameter (5 parameter individuals)	The COLREGs, together with the reference thresholds of A Guide to the Collision Avoidance Rules	COLREGs [2]; A Guide to the Collision Avoidance Rules [3]
Navigator knowledge	NavigatorKnowledge with SituationAwarenessKnowledge, RiskAssessmentKnowledge, EncounterDetermination-Knowledge and CollisionAvoidanceKnowledge; errorShare	The four decision stages of collision avoidance, each weighted by the error share measured in the causal findings	Causal findings of the Korean adjudication corpus (2078 findings)
Multi-ship encounter	PairwiseEncounter, ThreatSet, isPrimaryThreatOf, ProhibitedCourseSector; defined classes MultiShipEncounterCase and RoleConflictCase	Simultaneous pairwise encounters, primary-threat selection, and the coexistence of conflicting roles	Requirements literature [15,34–47]; multi-ship cases
Beyond the written rules	Rule2b SpecialCircumstanceCase, seamanship practice class, activity-based fishing state, empirical threshold parameters	The ordinary practice of seamen, special circumstances (assigned from held facts only), and the decision thresholds observed in adjudged encounters	Korean and international accident record; parsed thresholds

Twenty rule individuals represent the COLREGs. Five interpretation-parameter individuals store the indicative values of Cockcroft and Lameijer [3] with the values observed in the records. These values are stored for comparison and evaluation and do not classify encounters because the deterministic layer takes the encounter type as stated in the record (Section 2.4). The head-on criteria of Rule 14 remain under discussion [50]. A kinematic classifier must combine the course difference with the relative bearing (requirement ED1), and classifiers for autonomous ships use head-on bands from ±5° to ±13° [15,37,43,51]. The doubt clause of Rule 14(c) is held at the requirement level (ED5) and not in the deterministic layer because a record does not state a navigator's doubt. The 135° stern sector that defines overtaking, more than 22.5° abaft the beam, comes from the rule text.

The navigator knowledge layer arranges the rules into the four stages of Figure 3: SA covers Rule 5, RA Rules 6 and 7, ED Rules 13 to 15 and 18, and CA Rules 8, 16, and 17. The staging follows the rule sequence of Namgung [51] and the four steps of Ohn and Namgung [52]: detection, RA, ED, and evasive action. Comparable schemes use perception, situation recognition, decision, and control [53]; information, diagnosis and decision, and action [41]; or motion prediction, conflict detection, and conflict resolution [6]. RA and ED are kept separate because the COLREGs assign them different rules. CA is assigned the requirements governing conduct under the situation rules (Rules 9, 10, and 19), Rule 2 and the signal rules. A descriptive data property of each stage records its share of causal findings in the keyword screen (Section 3.1).

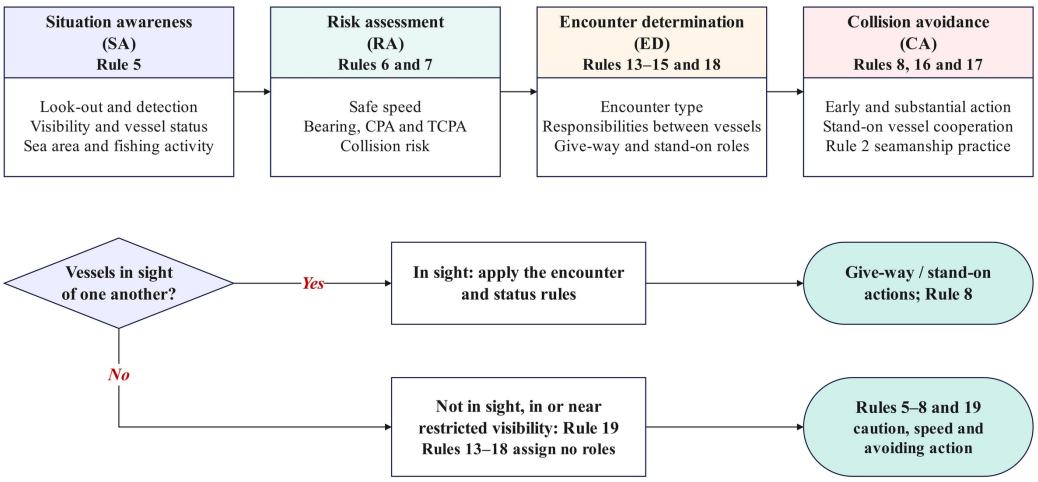


Figure 3. Decision stages of collision avoidance and the rules assigned to them. CPA: closest point of approach, TCPA: time to CPA.

Each stage lists the rules it applies and the case facts it uses. SA establishes whether the vessels are in sight of one another (requirement SA8), which decides between the in-sight branch and the Rule 19 branch of Section 2.4.

The multi-ship layer represents an encounter as a set of pairwise encounters between own ship and each target, stored as individuals. ThreatSet and isPrimaryThreatOf implement the multi-target requirements (ED7, RA8, CA28, and CA29; Table A2), and ProhibitedCourseSector holds the prohibited sectors of the ship-domain requirement (RA4). The reasoner (HermiT; Section 2.4) assigns a case to the defined class MultiShipEncounterCase when the case contains two or more pairwise encounters. It assigns a case to the defined class RoleConflictCase when own ship holds the give-way role in one pairwise encounter and the stand-on role in another, as recorded by the ownRole property.

The layer beyond the written rules holds knowledge that the rules leave to seamanship: a Rule 2(a) seamanship-practice class, supported by the citation rates of Section 2.2, and a Rule 2(b) special-circumstances class. The Rule 2(b) class is assigned only from facts the ontology already holds, so it marks where the clause was invoked and is not counted as formalized knowledge. The Rule 18 fishing state is taken from the record where it states that gear was being worked and otherwise from the reported speed (Table A4 in Appendix 5). This surrogate may include a waiting fishing vessel, miss trawlers towing at 3–5 kn and pot vessels setting gear at 5–8 kn (the working speeds stated in the KMST rulings), and cannot separate working, waiting, and transiting vessels within a fishing fleet. In the 50-document sample, it disagreed with the text in four cases. The observed encounter geometry is held in separate individuals marked as descriptive.

The class model was written for this study. Published navigation-rule ontologies model rule-based ship behavior [21], flag rule-violating behavior [23], or classify encounters and assign give-way and stand-on roles [24]; none of them holds accident-derived evidence.

Four questions drove the design: for a given case, which rules apply, which vessel must act, what action is required, and what evidence supports the requirement. Each layer of Table 3 answers one of them.

Accident-derived knowledge is added to the ontology by a fixed procedure (Figure 4). The evidence is a passage, found by pattern screening and reading (Sections 3.3.1–3.3.4), in which a ruling states what a vessel should have done or avoided. Categories of conduct are induced from these passages, and each becomes a requirement individual linked to its rule. A category is admitted only if its condition can be expressed with facts recorded in the ontology and if at least 10 cases support it. As a further screen, the lower bound of the 95%

Wilson interval for its prevalence must exceed 1%. The prevalence is computed over the whole record or, for conduct limited to some situations, over the relevant subset, such as active-fishing cases or confirmed multi-ship encounters. For conduct established only from the narrative, the base is the 1125 cases read in full or, for conduct counted in the Korean rulings only, the 709 Korean adjudications; one item (CA27) keeps the whole record as the more conservative base. Section 3.3.5 lists the base of each item.

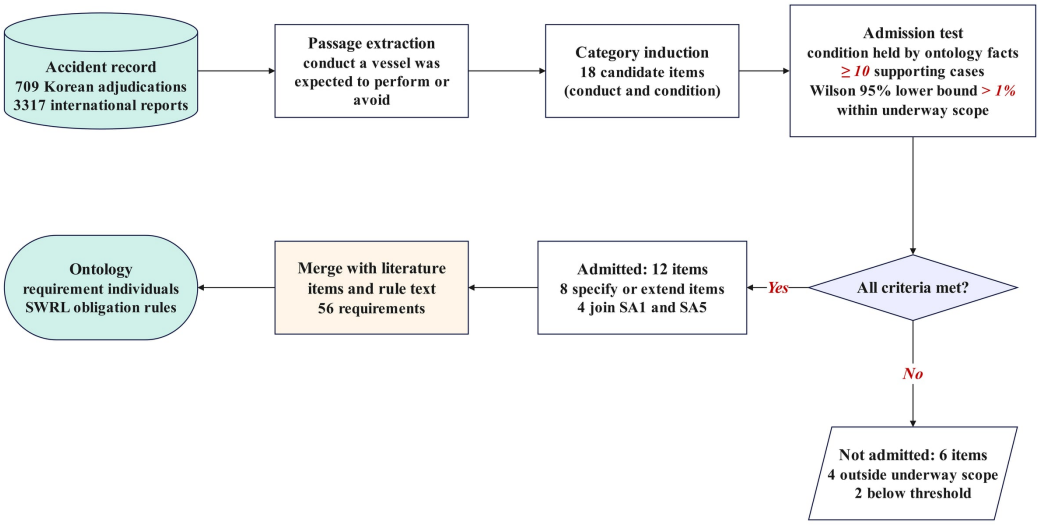


Figure 4. Derivation of requirements from the adjudications and investigation reports.

The count threshold excludes items supported by fewer than 10 cases, for which the Wilson interval is uninformative, and the 1% bound filters infrequent conduct. The bound does not estimate casualty-level prevalence because separate reports can describe the same collision. Section 3.3.5 reports how the results change with both thresholds. An LLM-proposed label was accepted only when the passage supported it.

For an admitted item, a Semantic Web Rule Language (SWRL) rule of the form used in Section 2.4 links the condition to an obligation class. Items outside the underway scope or failing the test are kept in the ontology as non-admitted candidates.

2.4. Deterministic Inference Layer

The inference layer works in three steps: it finds the applicable rules, derives the obligations, and classifies the case. Given the encounter type, visibility state, navigational status, and fishing activity, with location and vessel type where reported, the applicable rule set is computed deterministically. Rules 5 to 8 apply in every case. The in-sight rules (11 to 18) and Rule 19 form separate branches. Rule 19 applies when vessels are not in sight of one another in or near restricted visibility, and a disjointness axiom ensures that no case belongs to both InSightCase and RestrictedVisibilityCase. Three further disjoint classes divide the cases by the navigational status of the vessels at impact: both under way, under way against an anchored or moored vessel, and under way against a drifting vessel. The applicable branch of the rules differs among the three.

Rules 9 and 10 are added in either branch when their location facts are stated. Rule 18 is added only in the in-sight branch, and the lights and signal rules (20 to 27, 30, 34, and 35) are added from their own facts. When one vessel is engaged in fishing and the other is a power-driven vessel, Rule 18 assigns the give-way role in place of Rules 14 and 15, except where Rules 9, 10, and 13 otherwise require. The layer adds Rule 18 beside the encounter rules without removing them because the fishing state is a surrogate (Section 2.3); the worked case in Section 3.4 shows the consequence. Rules 9 and 10 bring their own

obligations, but the layer does not let them modify the obligations of the encounter. This is a limit of the implementation because Rules 9(d) and 10(c) restrict the crossing of a channel or lane, and Rules 9(b), 9(c), 10(i), and 10(j) impose duties not to impede. Rule 6 is treated as applicable in every case, as its text requires, without assessing whether the speed was safe, although safe speed is mentioned in 40% of the Korean adjudications (284 cases, Table A5). That assessment needs speed, traffic, and visibility facts that the records do not state consistently (requirement RA9).

The 37 SWRL rules link COLREGs rules to the case classes and assign the required actions, represented by 27 obligation classes such as MustAlterCourseToStarboard and MustKeepCourseAndSpeed. Under the open-world semantics of OWL, an unstated fact is unknown, not false. A definition that counts encounters (a cardinality restriction) therefore classifies a case only when its pairwise encounters are stated to be mutually distinct, so each case is given an explicit AllDifferent axiom. The obligation classes are subclasses of VesselRole. A SWRL rule assigns an obligation to the vessel that holds the role, and the inference script records it on the case through the property requiresObligation, so a case carries the obligations of both vessels. Give-way and stand-on obligations coexist without contradiction because no disjointness is asserted among the obligation classes.

Three of the 37 rules illustrate the form; ?c and ?v are SWRL variables for a case and one of its vessels, ∧ joins the conditions of the antecedent, and → leads to the consequent:

(i) RestrictedVisibilityCase(?c) → appliesRule(?c, R19)

(ii) CrossingInSightCase(?c) ∧ hasGiveWayVessel(?c, ?v) → MustGiveWayToStarboardTarget(?v)

(iii) AccidentCase(?c) ∧ hasObstacle(?c, FishingGear) ∧ involvesVessel(?c, ?v) → MustKeepClearOfFishingGear(?v)

Rule (i) attaches Rule 19 to a case in restricted visibility, rule (ii) assigns the crossing give-way obligation to the give-way vessel, and rule (iii) implements the accident-derived requirement CA20, derived from the fishing-gear conduct of Section 3.3.1. The consequent of rule (i) is a property assertion on the case, whereas those of rules (ii) and (iii) are class assertions on the vessel, which assign the obligation classes.

The knowledge base lies within OWL 2 DL, whose underlying DL is SROIQ(D), and is extended by DL-safe SWRL rules. It comprises 162 classes, 33 object properties, 49 data properties, and 4938 individuals (144,516 triples in total). The inference script implements Equation (3) and computes the applicable-rule sets. By default, Steps 3 and 4 of Algorithm A1 read a stated encounter type as an in-sight fact when the report gives no visibility because Rules 13 to 15 presuppose vessels in sight. This applies to 385 cases, or 9.6% of the record. Two of these cases, both international, cite Rule 19, so the default rarely contradicts the record. An encounter observed by radar alone in fog would nevertheless receive the in-sight rules in error. For use with sensor data, the default must be replaced by an explicit visibility state, obtained by fusing a visibility measurement with the way each target is detected, by sight or by radar alone. A target detected by radar alone in restricted visibility then falls in the Rule 19 branch, and a case without a visibility state must remain undecided. A case in which the vessels came into sight shortly before the collision is placed in one branch only; the transition from Rule 19 to the in-sight rules is not modeled.

A case is assigned to the class RestrictedVisibilityCase when its report states restricted visibility or an encounter in restricted visibility. The class approximates the Rule 19(a) condition, and none of its 442 cases receives Rule 18. A case stating neither fact receives only Rules 5 to 8 and its situation and status rules. The resulting sets and the obligations stated by the 37 SWRL rules are written into the ABox. The HermiT reasoner, run through the owlready2 0.51 Python 3.13 (Python Software Foundation, Wilmington, DE, USA) library, checks consistency and assigns cases to the defined classes of Section 2.3.

Equation (3) summarizes the applicable-rule set $A(x)$ for case x in the two stated-visibility branches. $R_{sit}(x)$ contains the rules added by situation and status facts: Rules 9 and 10, Rule 18 in the in-sight branch, and the lights and signal rules. $R_{enc}(e_x)$ contains the conduct rules of the encounter type e_x, and v_x is the stated visibility. Algorithm A1 in Appendix 5 gives the inference procedure in pseudocode. The applicable-rule sets are compared quantitatively against the record in Section 3.2. The multi-ship classifications are given in Section 3.1, and the derived obligations are shown in Section 3.4, because the record does not systematically state the required actions.

$$A(x) = \{R5,\ R6,\ R7,\ R8\} \cup R_{sit}(x) \cup \begin{cases} R_{enc}(e_x), & v_x = \textit{in sight} \\ \{R19\}, & v_x = \textit{not in sight},\ \textit{restricted visibility} \end{cases} \tag{3}$$

3. Results

3.1. Ontology Statistics and Multi-Ship Identification

The 4938 individuals comprise 4026 case records, 528 multi-ship structure individuals, 234 literature items, 18 accident-derived items, 56 consolidated requirements, 24 individuals for the tribunal and the investigating authorities, 20 rules, 5 interpretation parameters, 4 stages, and 23 situation and condition values used by the class restrictions and SWRL rules. The 528 structures include 205 pairwise encounters and their own-ship, target, and threat-set individuals. HermiT found no inconsistency and no unsatisfiable class (a class that cannot have any member), so the knowledge base is logically consistent. Three true/false data properties (reqInScopeUnderway, reqInScope, and reqAdmitted) mark knowledge outside the underway or requirement scope and items below the admission threshold.

The reasoning cost was measured with HermiT on a computer with two 2.1 GHz Intel Xeon cores and 8 GB of memory. For a single encounter, the reasoner derives from the stated facts the classes of the case, the rules that the SWRL rules attach to it, and the obligations of every vessel. With the reasoner running, this took at most 0.10 s per case, both for 20 randomly selected pairwise cases and for the 59 multi-ship cases, which have two to 29 targets, and 0.72 s for a synthetic encounter with 200 targets. Starting the reasoner and loading the ontology, which are needed only once, took about 3 s.

Each decision stage is assigned its share of the causal findings in the 709 Korean adjudications. A finding is one keyword match from the 14-keyword screen of Section 2.2 in the judgment and cause sections of the 706 rulings matched to the parsed text. Of the rulings, 636 contain at least one match and many contain several, giving 2078 findings. The shares are SA 38.2%, RA 18.2%, ED 0.1%, and CA 43.5% (Table A6). The screen has one ED keyword (mistaken encounter status or role), which matched three rulings, and reading finds such a misjudgment in 25 rulings (1.0% of findings; Section 2.2). ED is the smallest stage in both counts; failures to recognize a crossing situation may instead be recorded as a missed bearing in SA or an omitted action in CA. Within the multi-ship subset, RA rises to 27.3%. Figure 5 compares the stages with their treatment in the requirements literature.

Multi-ship encounters were identified within the 1125 cases read in full: all Korean cases and 416 international reports examined before the international tier was extended. A case was confirmed when, before impact, one vessel was simultaneously in encounters with at least two other independently navigating units. Of 163 screened candidates, 59 were confirmed (31 Korean and 28 international records) and 104 rejected: 61 counted a tug and tow twice, 30 involved a vessel that played no part in the encounter, and 13 were parsing artifacts or unreadable. Seven of the 30 involved fishing vessels or small boats with no encounter relationship to the vessel whose encounters were counted (two were its own fishing fleet): they distracted the look-out, spoke by radio, or lay far away. Counting parsed

vessel names would have overstated prevalence by a factor of 2.8. The 416-report subset was taken in reading order and over-represents Chinese and United States authorities, so its prevalence figures and most common patterns apply to this subset only. A fishing vessel is a party in 22 of the 31 Korean multi-ship cases (71%), close to its share of all Korean collisions (72%). The reasoner assigns the same 59 cases to MultiShipEncounterCase; because it applies the same criterion, this is not an independent test.

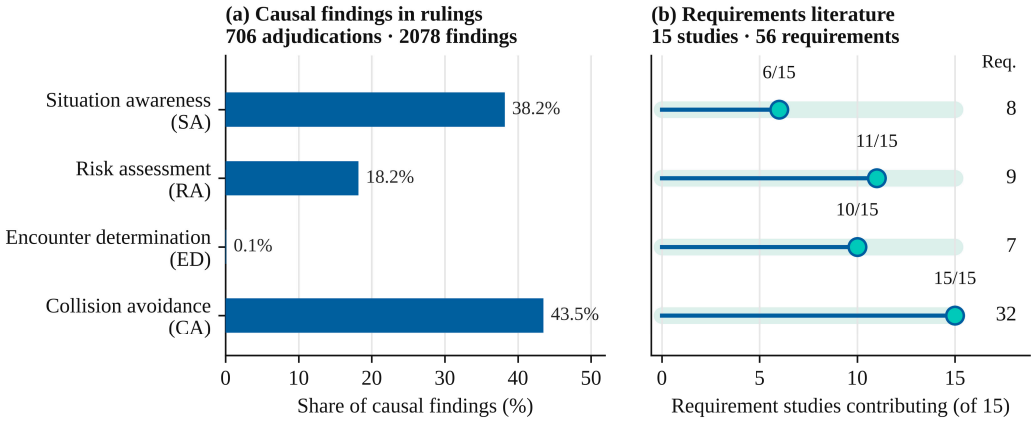


Figure 5. Decision stages in the accident record and the requirements literature: (**a**) share of causal findings; (**b**) studies contributing to the requirements of each stage.

Figure 5a counts how often rulings fault each stage. Figure 5b counts the studies contributing to at least one consolidated requirement in each stage, with the number of requirements at the right; each requirement retains the sources of all items it consolidates. Studies are counted instead of items, so that a study that lists its requirements in finer detail does not weigh more (4–40 items per study; Table A1). SA accounts for 38.2% of findings but is addressed by only 6 studies, whereas CA, with the largest share (43.5%), is addressed by all 15.

3.2. Validation Against Investigated Collisions

Because the axioms were written with all Korean rulings in view, the Korean comparison is an in-sample check; a temporal split is reported later in this section. The international comparison is the transfer test. The validation distinguishes what a report states from what the inference adds. The first factor is decidability: whether the stated facts are enough to decide which conduct rules apply. A record is decidable when it states an encounter type together with visibility or explicitly establishes restricted visibility. Of the 4026 case records, 48.7% are decidable. A further 9.6% state an encounter type without visibility, 26.7% state only situation facts such as anchoring or fishing, and 15.1% invoke only Rules 5 to 8. Decidability measures the completeness of the report and of its extraction.

The second factor is agreement with the rules cited by investigators, taken as the reference. The scored conduct rules are the nine rules of Table 4 (Rules 9, 10, and 13 to 19); Rules 5 to 8 are excluded as always applicable. The Korean tier holds 544 cases that cite at least one rule, 391 of which cite a conduct rule. Together, they contain 594 instances; an instance is a citation of one of the nine conduct rules in one case. The international tier holds 2273 such cases with 1304 instances. Of these cases, 416 cite articles explicitly, and 1857 are Japanese reports whose instances come from the clause-heading patterns of Section 2.2.

Table 4. Agreement of the inferred rule sets with the cited rules, with two lookup baselines.

Rule	Internal *n*	P	R	F1	External *n*	P	R	F1
R9	44	0.53	0.93	0.68	72	0.26	0.65	0.37
R10	15	0.37	0.87	0.52	36	0.27	0.97	0.42
R13	36	0.51	0.67	0.58	175	0.25	0.76	0.37
R14	28	0.36	0.96	0.52	142	0.63	0.51	0.57
R15	130	0.69	0.96	0.80	286	0.26	0.79	0.39
R16	63	0.22	0.71	0.34	179	0.09	0.69	0.17
R17	87	0.32	0.75	0.45	275	0.14	0.66	0.23
R18	129	0.46	0.88	0.61	39	0.04	0.41	0.08
R19	62	0.85	1.00	0.92	100	0.34	0.87	0.48
All, layer	594	0.45	0.87	0.60	1304	0.18	0.71	0.29
All, lookup A	594	0.42	0.60	0.49	1304	0.20	0.76	0.31
All, lookup B	594	0.43	0.89	0.58	1304	0.19	0.84	0.31

Internal: Korean tier, in-sample; External: international tier, transfer test; *n*: number of cited conduct-rule instances; P: precision; R: recall; R9–R19: Rules 9–19; All: micro-averaged over the nine rules. P, R, and F1 are computed per rule over the cases citing at least one rule (Section 3.2), and F1 from unrounded P and R.

Equation (4) defines coverage as the share of cited instances contained in the inferred sets, and Equation (5) gives the standard precision (P), recall (R), and F1 for each rule r over the cases of one tier:

$$Coverage = \frac{|C \cap A|}{|C|} \tag{4}$$

$$P_r = \frac{TP_r}{TP_r + FP_r},\ R_r = \frac{TP_r}{TP_r + FN_r},\ F1_r = \frac{2\,P_r R_r}{P_r + R_r} \tag{5}$$

where C and A are the sets of cited and inferred rule instances, and TP_r, FP_r, and FN_r are the true positives, false positives, and false negatives for rule r. On the Korean tier, the inferred sets contain 86.7% of the cited instances (515 of 594; bootstrap 95% CI 83.4–90.0). In the transfer test, they contain 70.8% of international instances (923 of 1304), and 75.8% of the 1898 instances over both tiers. On the international tier, agreement is 76.8% for heading-attributed instances (530 of 690) and 64% for explicit citations (393 of 614).

Where a report states the necessary encounter facts, agreement reaches 93.5% on the Korean tier (348 of 372 instances over 227 cases) and 72.6% on the international tier (690 of 951 over 582 cases). Without Rules 13 to 15, which follow directly from the stated encounter type, agreement is 84.8% and 70.0%. Korean agreement is 90.1%, 86.9%, and 82.2% for 2016–2018, 2019–2021, and 2022–2025. The decline partly reflects the case mix: Rule 19, derived for every case stating restricted visibility, falls from 34 cited instances from 2016–2018 to 10 from 2022–2025, whereas Rule 13, the rule with the lowest recall on the Korean tier, rises from 12 to 15. The share of decidable cases also falls from 45.5% to 38.6%. Table 4 gives the per-rule scores.

Most of the remaining disagreement lowers precision: the layer derives every rule that applies, whereas investigators cite only the breached rules. P is therefore a lower bound on agreement. Macro-averaged F1, the unweighted mean over the nine rules, is 0.60 for Korea and 0.34 internationally. The lowest recalls are 0.67 for Rule 13 on the Korean tier and 0.41 for Rule 18 on the international tier ($n = 39$). Precision for Rule 16 is 0.09 internationally because the layer derives the rule for every give-way vessel, whereas the reports cite it in only 179 instances (86 explicit citations, 19 attributed from a term of art and 74 from the wording of a causal finding).

Rule 18 recall depends on the speed surrogate in Section 2.3. If the speed test is applied only to cases whose parties include a fishing vessel, as coded in Table A3, Rule 18 precision on the Korean tier rises from 0.46 to 0.49, and recall falls from 0.88 to 0.84. Internationally,

only 39 of the 1422 cases involving a fishing vessel cite Rule 18, so precision stays near 0.04 with or without this restriction. Of these 1422 cases, 447 are collisions between two fishing vessels, to which Rule 18 applies only when one is engaged in fishing and the other is not. The international reports, as no-blame investigations, also seldom name the rule that a vessel breached (Section 2.1). Counting each of the 25 casualties heard on appeal only once leaves Korean agreement at 86.7%.

Two lookup baselines show how much of this agreement the stated facts alone explain (Table 4). A direct lookup from the stated encounter type (lookup A) maps crossing to Rules 15 to 17, head-on to Rule 14, overtaking to Rules 13, 16, and 17, and restricted visibility to Rule 19; it contains 59.8% of the Korean and 76.5% of the international instances. Adding Rules 9, 10, and 18 from stated location and fishing (lookup B) raises its coverage to 89.1% and 84.4%, above that of the inference layer; the micro-averaged precision of lookup B is 42.9% and 19.0%, against 45.4% and 18.1% for the layer. The lookup gains coverage in restricted-visibility cases whose reports also cite in-sight rules (16 Korean and 116 international cases) because the layer does not apply the in-sight rules when Rule 19 applies. On the international tier, even lookup A scores slightly above the layer (F1 0.31 against 0.29). Coverage therefore measures consistency between the derived sets and the record, not predictive skill; what the layer adds is the branch structure required by the COLREGs, the obligations, and the trace.

By authority, international agreement is 76.8% for the JTSB, 63% for the GISIS depository, 45.2% for the UK Marine Accident Investigation Branch (MAIB), 14.3% for the Chinese maritime safety administrations (seven instances), and 75.4% for the remaining authorities (Appendix 5, Table A7). The international figure is thus weighted toward Japanese reports, which supply 1857 of the 2273 citing cases and 690 of the 1304 instances. The two low values reflect extraction, not inference. Rule 14 is cited in 17 MAIB reports but derived in three because the parsed facts state a head-on encounter in only four of them. In the Chinese reports, most citations map to Rule 6, which is not scored; only seven map to the nine conduct rules, six of them to Rule 9 in reports whose parsed facts do not state a narrow channel. The 14.3% therefore rests on too few instances to interpret.

To test whether agreement depends on period-specific citation patterns, the 391 Korean cases citing at least one conduct rule were split by publication year into a fitting period (2016–2020, 235 cases) and a hold-out period (2021–2025, 156 cases and 246 instances). A lookup fitted to 2016–2020 alone serves as one comparator. It predicts a rule for a given combination of facts when at least a threshold share of the 2016–2020 cases with that combination cite the rule; the threshold was tuned for F1 on the same period. The fixed lookup B of Table 4 is a second comparator. Layer coverage falls from 88.5% to 84.1% (difference 4.4 percentage points, case-level bootstrap 95% CI −2.5 to 11.1 points); with the 2016–2020 mix of cited rules, the hold-out coverage would be 86.9% (Mantel–Haenszel odds ratio of containment between the periods 1.15, 95% CI 0.61–2.16). Layer F1 is 0.60 and 0.59 in the two periods, whereas the fitted lookup falls from 0.66 to 0.48, so that in the hold-out, the layer exceeds it by 0.11 (95% CI 0.04–0.17; Figure 6). The layer therefore does not depend on period-specific citation frequencies, and lookup B, which is not fitted to citation frequencies either, is as stable as the layer (0.58 in both periods). On the international tier, the layer scores above the fitted lookup (F1 0.29 against 0.23; 95% CI of the difference 0.03–0.08) and below the fixed lookup (0.31).

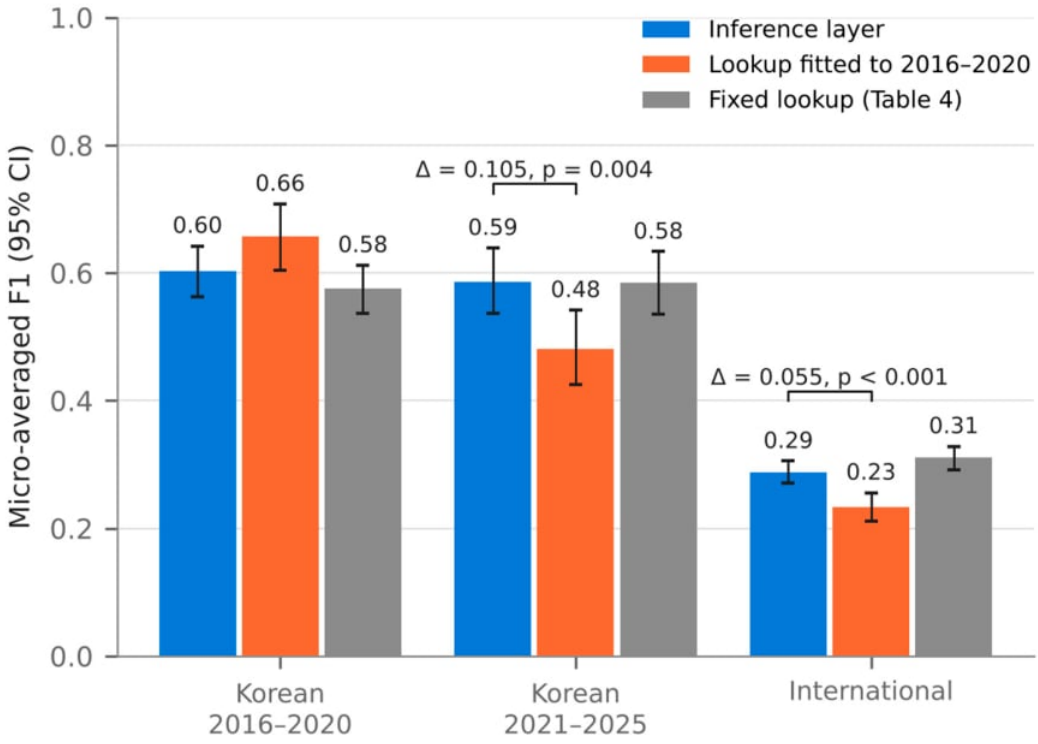


Figure 6. Temporal split on the Korean tier and comparison on the international tier: micro-averaged F1 of the inference layer and two lookups, with 95% bootstrap CIs; Δ is the paired difference between the layer and the lookup fitted to 2016–2020, and p is its two-sided case-level paired bootstrap p-value. I: confidence interval.

3.3. Knowledge Beyond the Written Rules: Accident-Derived Requirements

Four findings concern knowledge that the written rules state only in general terms or not at all. The first three produce accident-derived conduct requirements, and the fourth breaks down the SA failures in the causal findings. Sections 3.3.1–3.3.4 present the findings, and Section 3.3.5 presents the admission of the items they produce.

3.3.1. The Ordinary Practice of Seamen

The first finding concerns the ordinary practice of seamen, which the requirements literature states only in general terms. The pattern detects the phrase in 272 Korean adjudications (Section 2.2); a usable passage was extracted from 266, and 262 link the phrase to specific conduct. Table 5 presents the taxonomy applied to every case in either tier that invokes the phrase. Cases may fall into several categories, so the total exceeds the case count. Table 5 counts only conduct linked directly to the phrase. Fishing vessels are parties to most Korean collisions (Table A3). The statutory equivalent of Rule 2 is nevertheless cited in 13.6% of the 258 collisions between two fishing vessels and in 9.2% of the 434 other collisions of known vessel type (Fisher's exact test, $p = 0.08$). It is cited in 3 of the 31 multi-ship cases and in 72 of the other 678 ($p = 1.0$). Neither difference is significant at the 5% level, so the invocation of Rule 2 is not confined to fishing-vessel or multi-ship encounters.

Table 5. Conduct for which the ordinary practice of seamen is invoked.

Conduct Category	Cases
Conduct toward anchored, moored, or drifting vessel	187
Stationary vessel's watch, lights, and cooperation	92
Warning and attention signals; VHF communication	61

Table 5. *Cont.*

Conduct Category	Cases
Berthing, unberthing, and shifting maneuvers	42
Residual mutual care where steering rules assign no roles	34
Early, positive, and substantial avoiding action	27
Conduct around fishing vessels and gear	24
Keeping a safe passing distance	20
Preventive speed reduction	15
Other (port-to-port custom, no sudden alteration into another's path, engine readiness, joining traffic, look-out skill)	71

Between vessels under way, tribunals apply the doctrine to five kinds of conduct. There are three concern actions expressed toward another vessel: keeping clear of vessels engaged in fishing and their gear (24 cases), using attention-attracting signals and very high-frequency (VHF) radio while risk persists (61), and reducing speed preventively when not bound to give way (15). The signal-and-VHF conduct reflects Korean tribunal practice. Requirement CA22 admits a VHF call only as a supplement to the early action that Rule 8 requires, not as a condition for it, because negotiation by radio can delay that action [54]. There are two restraint concerns: residual mutual care where the steering rules assign no roles (34 cases), and avoiding sudden alterations into the path of a following or passing vessel. International reports mainly use good seamanship to judge the quality of avoiding action, which should be early, positive, and substantial (27 cases). The adjudications thus state what conduct was expected, and the investigation reports describe how well it was carried out.

The largest category, conduct toward anchored, moored, or drifting vessels (187 cases), lies mostly outside the underway scope. Conduct toward drifting vessels lies within the scope because a drifting vessel remains under way in law. Sixty-eight Korean rulings apply the ordinary practice of seamen to a vessel stopped and making no way, not at anchor, without a Rule 18 priority; most of these vessels were angling or fishing boats. By parsed navigational status, drifting vessels appear in 22.7% of international and 3.8% of Korean collisions. The parser does not separate the Korean terms for lying at anchor and lying stopped, so it places 24 of the 68 rulings in the anchored class, 11 in the drifting class, and 33 in neither. This conduct forms requirement CA27.

3.3.2. Timing of the Stand-On Vessel's Action Under Rule 17(b)

The second finding concerns the timing of the action that Rule 17(b) requires of a stand-on vessel when the action of the give-way vessel alone can no longer avoid collision. The Korean enactment names this duty cooperative action and separates it from the discretionary action of Rule 17(a)(ii). The analyzed passages come from the 406 rulings, counted by reading, in which Korean tribunals identify failure to take that action as a contributory cause; the keyword count of Table A6 (446) is a screen and is not used here. Each ruling was read to determine whether the vessel owing the duty had the other vessel in sight. In 92 rulings, the vessel had the other in sight but did not act. Another 162 rulings record that the other vessel was detected late or not at all; these are look-out failures because Rule 17 presupposes that the other vessel is in sight. The remaining 152 lack sufficient facts or concern a vessel not under way.

Some rulings state when action was taken or when the tribunal held that it was due. Of 42 quantities confirmed against their source sentences, 35 were retained after screening for scope. Each was then checked against the action of the stand-on vessel itself, which left 16 rulings with an action time, counted once per ruling. In 12 of them, the time is

that of recorded action, with a median of 2.5 min before collision (range 1–15 min). In the other four, it is the time at which the tribunal held that action was due, with a median of 3.5 min (range 2–6 min). The pooled median is 3 min: seven values are at most 2 min, five are 3–5 min, and four exceed 5 min. Two of the 16 rulings concern a vessel at anchor; the tribunal assessed its signals and engine use as its cooperative action under the ordinary practice of seamen. The longest time, 15 min, is one of these two, and without them, the pooled median is 2.5 min (14 rulings). Only rulings that state a time enter the sample, and every case ended in collision. The statistic therefore describes when Rule 17(b) action was recorded or adjudged due in these cases; it is not a safe-action threshold because the sample cannot show whether action at this time would have sufficed. The finding is exploratory.

The rulings also state how ample time is judged. The maneuver reserve ratio (MRR) is the inter-vessel range divided by own ship's tactical diameter. Two rulings accept that time was ample at an MRR of at least 2, and another 203 invoke ample time without a numeric test. MRR ≥ 2 is therefore retained as a candidate evaluation measure, not as a general obligation under Rule 8 or 16. Five rulings hold that the steering rules apply from about 15 min before collision, which leaves time for a second attempt if the first avoiding action fails.

3.3.3. Conflicting Roles in Multi-Ship Encounters

The third finding concerns conflicting roles in multi-ship encounters. Five Korean rulings and one report of the US National Transportation Safety Board (NTSB), covering five casualties, show one vessel holding opposite obligations toward two others simultaneously. Two of the rulings concern the same casualty at first instance and on appeal, and two records contain no explicit resolution. None of the records ranks one pairwise obligation above the other. Table 6 lists these records, from which four initial patterns emerge: (i) a maneuver for one encounter must not aggravate another; (ii) if no steering rule assigns roles, Rule 2 governs the conduct, through the ordinary practice of seamen or, where the limitations of the vessels so require, the special-circumstances clause; (iii) in port, a vessel traffic service (VTS) may sequence agreed passings; and (iv) the vessel better able to maneuver should act early at ample distance, as requirement CA29 states. Each pattern rests on one or two of these six records and is retained as an exploratory taxonomy. One requirement recurs in 11 cases and is admitted as item 5 of Table 7 (requirement CA28): a single maneuver must resolve all encounters at once.

Table 6. Records in which a vessel held conflicting roles, and their resolution.

Case	Conflict	How It Was Resolved	Principle
KMST Central 2018-006	Stand-on vessel altering to port for a fishing fleet while holding for a crossing vessel	The fleet-avoidance track was faulted because it increased the risk to the vessel held for	Maneuver must not aggravate the other encounter
KMST Incheon 2016-029	Tug-and-tow meeting an anchored vessel with its port turn blocked by overtaking boats	No steering rule assigns roles; Rule 2 applied: early action at ample distance and an extra look-out on the blocked side	Special circumstances, Rule 2(b)
KMST Central 2017-005	Same casualty on appeal	The blocked-turn defense was dismissed without role ranking; fault placed on look-out and unreduced speed	No explicit resolution

Table 6. *Cont.*

Case	Conflict	How It Was Resolved	Principle
KMST Busan 2019-059	Vessel unable to turn to starboard because another ship was overtaking on that quarter	Faulted for entering the channel under schedule pressure and for not seeking the overtaking ship's consent by radio	Sea room secured in advance; radio coordination
NTSB MAB-21/16	Vessel overtaken while meeting another tow in a bend	No rule cited and no priority determined; cause placed on radio communication and inaccurate tow dimensions	No explicit resolution
KMST Incheon 2023-030	Starboard-to-starboard passing agreed with one vessel while another required port-to-port	The traffic service sequenced the two passings in time, and the agreed passing displaced the statutory keep-right duty	VTS-sequenced agreed passing under port-traffic authority; no general license for agreements displacing the rules

Table 7. Accident-derived requirement items admitted within the underway scope.

No.	Accident-Derived Requirement	Req.	Rule	Cases (Base)	Quantitative Value Where Measured
1	Attention-attracting signal and VHF communication to confirm intentions (VHF as a supplement to early action)	CA22	R34(d), R36	64 (4026)	—
2	Residual mutual care where the steering rules assign no roles	CA26	R2(a)	34 (1872)	—
3	Keep clear of vessels engaged in fishing and their gear	CA20	R2(a)	24 (859)	—
4	Timing of the stand-on vessel's action under Rule 17(b)	CA14	R17(b)	164 (1125)	Median 3 min before collision (recorded 2.5 min, adjudged 3.5 min; 16 rulings; exploratory)
5	Multi-ship: maneuver serving all encounters simultaneously	CA28	R2, R8(c)	11 (59)	Early speed reduction or stop preferred
6	Proper lights and shapes; correct identification	SA3	R5, R18, R20	53 (1125)	—
7	Ample-time action; maneuver reserve as a candidate measure	CA1	R8(a), R16	205 (709)	MRR = range ÷ tactical diameter ($\leq$ five ship lengths, IMO MSC.137(76) [55]); MRR $\geq$ 2 stated in two rulings; candidate evaluation measure
8	Conduct toward a drifting vessel (stopped and making no way, not at anchor)	CA27	R2(a)	68 (4026)	Drifting vessel in 22.7% of international and 3.8% of Korean collisions (parsed status)

Cases are supporting records across both tiers; item 4 comprises 92 Korean and 72 international records, and item 1 includes three rulings that state the conduct without invoking the ordinary practice of seamen. The base in parentheses is the denominator of the admission test (Sections 2.3 and 3.3.5). Req.: requirement code in Table A2. Item 8 counts the 68 rulings of Section 3.3.1 against the whole record, the more conservative base; over the 1125 cases read in full, its lower bound would be 4.8%. MRR: maneuver reserve ratio, MSC: Maritime Safety Committee, R: Rule, VHF: very high frequency.

3.3.4. Situation-Awareness Failures

The fourth finding breaks down the SA failures of Section 3.1. International counts come from the 416 reports read in full. The shares are counted by reading and are not the keyword counts of Table A6. Sleep or fatigue appears in 10.4% of Korean adjudications and 4.8% of the international subset. Unlit or misidentified lights and shapes occur in 6.8% and 1.2%, respectively, and often involve small fishing vessels. Alcohol appears in 3.5% and 0.2%, respectively. Telephone and secondary-task distractions are relatively more common internationally. These failures bear on the look-out requirement (SA1; Section 3.3.5).

3.3.5. Admission of the Accident-Derived Items

Table 7 lists the eight admitted items and their case bases: the whole record (4026), the cases read in full (1125), all Korean adjudications (709), the cases without assigned steering roles (1872), the active-fishing cases (859), and the confirmed multi-ship encounters (59).

The four findings produced 18 candidate individuals in the AccidentDerivedRequirement class, each assessed by the procedure in Section 2.3. The fishing-gear item shows how a passage becomes an axiom. Passages that fault a vessel for passing close to working gear become a requirement individual linked to Rule 2(a); Rule 18 imposes a duty toward the fishing vessel, not toward its gear. The matching SWRL rule assigns MustKeepClearOfFishingGear to every vessel of a case in which fishing gear is an obstacle, including the fishing vessel itself, because the parsed facts do not identify which vessel approached the gear. Twelve candidates passed the admission procedure within the underway scope. Eight form three new requirements or extend existing literature-derived items (Table 7). Four join existing SA requirements: bridge manning, fatigue prevention and distraction control join the look-out requirement (SA1), and systematic radar observation joins SA5. The other six candidates were not admitted: four concern stationary vessels or berthing maneuvers, and two fall below the evidence threshold.

The Wilson lower bounds of the admitted items range from 1.2% to 25.7%, with per-tier values retained in the ontology. Lowering the count threshold to five cases does not change the admitted set. Raising it to 20 removes the multi-ship item supported by 11 cases. Lowering the prevalence bound to 0.5% would also admit the item on sudden alteration into another vessel's path, whose lower bound is 0.7%. Raising it to 2% would remove the signal-and-VHF item (1.2%), residual mutual care (1.3%), the drifting-vessel item (1.3%), and the fishing-gear item (1.9%). The eight items are therefore stable for prevalence bounds between 0.7% and 1.2%. Six items specify conduct that the rules or the literature state only in general terms, one extends pairwise treatment to simultaneous encounters, and one, conducting toward a drifting vessel, has no counterpart in the steering rules. The admitted set thus depends on the thresholds and on the bases chosen for each conduct condition. The four items closest to the 1% bound (the signal-and-VHF item, residual mutual care, the drifting-vessel item and the fishing-gear item) are the most sensitive to them.

3.4. Explainability Demonstration

Three pairwise cases illustrate the inference chain and its limits. In KMST Busan 2019-064, a pusher-barge unit and a fishing vessel met in a crossing situation while the latter was recorded, from her reported speed, as engaged in fishing. The asserted facts are encounter type, good visibility, and this active fishing state. The case is classified as CrossingInSightCase, and the SWRL rules assign the crossing obligations to the give-way and stand-on vessels. The fishing branch also requires the power-driven vessel to keep out of the way of the vessel engaged in fishing (Rule 18(a)(iii)) and, under Rule 2(a), every vessel to keep clear of her gear (CA20). The derived rule set is {5, 6, 7, 8, 15, 16, 17, 18, 26, 34}. The tribunal cited Rules 5, 7, 8, 15, 16, and 17. Rules 18, 26, and 34 follow from the parsed fishing state and vessel type; the ruling states these facts but does not cite the rules. The tribunal decided the case on crossing roles. The parsed fishing state may therefore overstate Rule 18, which is consistent with the speed-surrogate error described in Section 2.3. The layer keeps both the crossing and the fishing branches unranked, although Rule 18 would displace the crossing roles if the fishing state were confirmed.

KMST Donghae 2020-013 illustrates the restricted-visibility branch. The vessels approached in fog, and no in-sight fact is stated. The layer derives {5, 6, 7, 8, 19, 26, 35}. MustNavigateWithExtremeCaution and MustProceedAtSafeSpeed replace the role assignments of the in-sight branch. MustSoundFogSignal follows from Rule 35. The tribunal cited Rules 5, 6, 7, 8, 19, and 35.

KMST Incheon 2023-030 illustrates a remaining limit. The vessels agreed to pass starboard to starboard despite the statutory keep-right duty. The parser could not assign the pairwise roles, so the case was not classified as RoleConflictCase. The conflict reported in Section 3.3.3 and Table 6 was found only by reading the ruling.

Table 8 extends the demonstration to multi-ship encounters and traces each case from asserted facts to rule classifications and obligations. In KMST Donghae 2016-030, avoiding one vessel created a new crossing with a following vessel. Its two pairwise encounters make it a MultiShipEncounterCase, and the tribunal decided the new situation under the seamanship clause, not the crossing rule. KMST Busan 2019-043 combines reciprocal courses, recorded as head-on, with restricted visibility and an anchored vessel, so the layer derives Rules 19, 30, and 35 and withholds Rule 14. The fishing vessels that blocked the agreed alteration appear only in the narrative. The JTSB and NTSB reports on the Fitzgerald–ACX Crystal collision describe a look-out divided among three targets. The layer derives the crossing and stand-on obligations, and the linked reports supply the Rule 17(b) action and the bridge-fatigue evidence of Section 3.3.

Table 8. Multi-ship encounters traced through the ontology.

Case and Pattern	Asserted Facts → Classification	Derived Rules and Key Obligations	Requirements Engaged (Table 7 Items; Table A2 Codes)
KMST Donghae 2016-030. Avoiding one vessel created a new crossing with a following vessel	Crossing · in sight · both under way; two pairwise encounters (ThreatSet) → MultiShipEncounterCase	{5–8, 15–17, 34}; MustGiveWayToStarboardTarget, MustKeepCourseAndSpeed	Item 2, CA26 (mutual avoidance): the tribunal held the crossing rule inapplicable to the emergent situation and applied the seamanship clause; motivates item 5 (CA28)

Table 8. *Cont.*

Case and Pattern	Asserted Facts → Classification	Derived Rules and Key Obligations	Requirements Engaged (Table 7 Items; Table A2 Codes)
KMST Busan 2019-043. Fishing vessels to starboard blocked the agreed alteration	Head-on · restricted visibility · anchored vessel; four pairwise encounters → MultiShipEncounterCase	{5–8, 19, 30, 35}; MustNavigateWithExtremeCaution, MustKeepClearOfAnchoredVessel, MustSoundFogSignal	Target situation of the ship-domain requirement with prohibited sectors (RA4); the blocked sector itself is stated only in prose
JTSB MA2019-8-1 with linked NTSB MAR-20-02 (Fitzgerald–ACX Crystal). Look-out divided among three targets; the masked one struck	Crossing geometries with three targets; three pairwise encounters → MultiShipEncounterCase	{5–8, 15–17, 34}; MustGiveWayToStarboardTarget, MustKeepCourseAndSpeed	Item 4, CA14 (stand-on action under Rule 17(b), three stand-on vessels) and bridge fatigue (SA1, Section 3.3.4), evidenced in the linked NTSB report

4. Discussion

Because the ontology is grounded in 4026 investigated collisions, it can also represent some conduct that reports describe but the written rules leave implicit. The two tiers respond differently to complete inputs. When the encounter facts are stated, Korean agreement rises by about seven percentage points, but international agreement rises by less than two. The lower international agreement therefore reflects differences in legal purpose, reporting practice, extraction error, and case composition more than missing facts.

The method keeps knowledge acquisition apart from inference. Requirements from the literature, the rule text, and the accident record are merged by a single criterion, so each stays tied to a provision and an obligation class. The deterministic layer returns the same rule set for the same facts. Evaluation frameworks such as those of [15,35,37] score rule compliance from an observed trajectory, in real time, or after the event. The ontology states which rules and obligations govern a case before any maneuver. Disagreement is concentrated where an input is a surrogate, as with the speed-based fishing state behind Rule 18, or where reports cite only the rule that was breached, as with Rule 16. The accident-derived knowledge lies mostly in situations that scenario-based evaluations seldom include: the timing of stand-on action, conflicting roles among several vessels and encounters with drifting vessels.

For autonomous navigation, the ontology is intended to sit between perception and planning as a normative layer that turns the facts of an encounter into applicable rules and traceable obligations. These outputs may support supervision by an operator and assessment by an authority. The reasoning time for pairwise and multi-ship encounters alike (Section 3.1) appears to be adequate for on-board use. The 56 requirements, their evaluation items, and the per-rule scores may also inform test scenarios and pass criteria for collision-avoidance systems.

These results have limits. Agreement was measured against the rules cited in the same documents from which the inputs were parsed, and the Korean figures are in-sample (Section 3.2). The derived obligations were illustrated with worked cases but not scored against an independently annotated reference, and the passages were coded by one reader. The accident-derived knowledge comes mainly from Korean liability rulings, which judge conduct with hindsight, and every record ends in a collision, so the empirical values describe encounters in which avoidance failed.

5. Conclusions

The ontology developed in this study joins the rule text, the requirements literature, and the accident record in one OWL 2 knowledge base. Its deterministic layer turns the stated facts on encounter type, visibility, navigational status, fishing, location, and vessel type into the applicable rules and case-level obligations, and each derivation follows named classes and axioms.

Against the rules cited by investigators, the layer covered 86.7% of the Korean instances (in-sample) and 70.8% of the international instances (64.0% for explicit citations), and its F1 was stable between 2016–2020 and 2021–2025 (0.60 and 0.59). Because a lookup from the stated encounter type, location and fishing facts reached similar or higher coverage and was as stable across the two periods, the value of the layer lies in its rule branches, obligations, and traces, which a lookup does not provide.

The accident records also yielded conduct knowledge beyond the rule text. In 16 rulings, Rule 17(b) action was recorded or adjudged to a median of 3 min before collision; six records of five casualties showed conflicting duties without any ranking; and 262 rulings invoked the ordinary practice of seamen for specific conduct, including encounters with drifting vessels. These findings are exploratory. The literature, the rule text, and the accident evidence were consolidated into 56 requirements with 158 evaluation items; eight accident-derived items form three of these requirements or extend five others, and four more strengthen the SA requirements, the stage the literature addresses least.

Connected to sensor-derived facts, ontology may serve as a basis for explainable collision avoidance in which each avoiding maneuver of an autonomous ship is tied to a named rule and obligation. Future work will address the following:

- Sensor-based derivation of the encounter type from course difference and relative bearing, the risk of collision from range and CPA, the visibility and the state of wind, sea, and current, and the fishing state from lights, shapes, or track instead of speed;
- A model of the transition from Rule 19 to the in-sight rules, connection of the layer to a planner, and tests in simulation and at sea;
- Representation of the obligations of Rules 9 and 10 that modify those of the encounter, such as the restrictions on crossing a channel or lane and the duties not to impede;
- Independent annotation of the obligations and taxonomies by a second reader, and review of the requirement set by practicing navigators;
- Tests of the knowledge beyond the written rules on adjudications from other jurisdictions and on a prospective hold-out, with the axioms frozen before later rulings are read, and a comparison of the national enactments with the authentic English and French texts;
- Collision avoidance in ice-covered waters, including the ice conditions that restrict the available maneuvers and decision support for avoiding action [56].

Author Contributions: Conceptualization, J.-Y.L. and H.N.; methodology, H.N.; software, J.-Y.L. and H.N.; validation, J.-Y.L. and H.N.; formal analysis, J.-Y.L. and H.N.; investigation, J.-Y.L. and H.N.; resources, H.N.; data curation, J.-Y.L. and H.N.; writing—original draft preparation, J.-Y.L.; writing—review and editing, H.N.; visualization, J.-Y.L.; supervision, H.N.; project administration, H.N. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data that support the findings of this study are available from the corresponding author upon reasonable request.

Acknowledgments: During the preparation of this study, the authors used Claude Opus 5.0 and Claude Fable 5.0 (Anthropic, 2026) to assist in retrieving the international reports, parsing the accident-

report corpus and extracting the literature requirement items, as described in Section 2.2. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

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

Appendix A

This appendix contains the tables and pseudocode supporting Sections 2 and 3. Table A1 lists the 15 requirement studies, Table A2 lists the 56 consolidated requirements with their COLREGs provisions, Table A3 lists the vessel types involved in the collisions, Table A4 lists the definitions of case-level facts and the rule-attribution criteria, Table A5 lists the 20 concept categories of the lexicon with their case counts, and Table A6 lists the fault keywords behind the stage shares. Algorithm A1 gives the inference procedure in pseudocode, and Table A7 reports international agreement by authority.

Table A1. Requirement studies of the second knowledge source.

Study	Focus	Stages	Extracted	Items	Req.	Distinctive Requirements (Examples)
Woerner et al. [15]	Evaluation metrics for autonomous collision avoidance	SA, RA, ED, CA	43	40	17	CPA safety and pose scores (RA3, CA4, CA5); encounter change handling (ED3); restricted-visibility conduct (CA16)
Ohn and Namgung [34]	Requirements for local route planning	SA, RA, ED, CA	40	37	24	Radar observation (SA5); bearing change check (RA2); course alteration in sea room (CA3); status-based give-way (CA19)
Gleeson et al. [35]	Scenario classification and compliance scoring with multi-vessel awareness	CA	16	14	8	Delayed and non-apparent maneuver penalties (CA1, CA2); consistent multi-ship maneuver (CA28)
Tengesdal and Johansen [36]	Simulation framework for evaluating avoidance algorithms	CA	6	6	3	Apparent maneuver (CA2); CPA safety score (CA4)
Hagen et al. [37]	Safety and COLREGs evaluation	RA, ED, CA	32	30	15	Encounter stage (ED2); encounter per target (ED7); safe passing angle (CA5)
Hagen et al. [38]	Grounding hazard in evaluation	CA	4	4	2	Compensation of maneuver penalties near grounding hazards (CA2, CA12)
Akdağ et al. [39]	Review of collaborative avoidance	SA, RA, ED, CA	22	21	10	Lights and shapes identification (SA3); AIS as an aid (SA4); reactive fallback (CA31)

Table A1. *Cont.*

Study	Focus	Stages	Extracted	Items	Req.	Distinctive Requirements (Examples)
Campbell et al. [40]	Review of unmanned surface vehicle autonomy	SA, RA, ED, CA	32	32	12	Target classification (SA2); doubt resolved conservatively (ED5); reactive fallback (CA31)
Ramos et al. [41]	Operator tasks and human failure events	SA, CA	21	21	5	Own-ship and target data (SA6); intention confirmation (CA22)
Öztürk et al. [42]	Review of path-planning algorithms	SA, RA, ED, CA	22	21	16	Sound and light signals (CA21); restricted-visibility conduct (CA16); narrow-channel and traffic-separation conduct (CA17, CA18)
Krasowski and Althoff [43]	Formal traffic-rule predicates	RA, ED, CA	18	17	11	Emergency detection (RA7); encounter priority (ED4); stand-on required action (CA14)
Ahmed et al. [44]	Fuzzy avoidance for multiple ships	RA, ED, CA	16	16	9	Non-dimensional CPA distance and time (RA3); per-target collision risk in multi-ship encounters (RA8)
Zaccone [45]	Dynamic-programming avoidance	RA, ED, CA	13	12	9	Minimum and maximum course alteration (CA2); crossing-order constraint (CA10); maneuver within limits (CA30)
Lan et al. [46]	Multi-stage fuzzy avoidance model	RA, ED, CA	22	22	8	Four-ring distance model (RA4); encounter zones and stages (ED2)
Li et al. [47]	Critical ship arena from accident reports	RA, CA	16	15	4	Critical ship arena radii (RA4); close-quarters criterion (RA5); early-warning bands (RA6)

Extracted: items verified for the study before duplicates were merged (323 in total). Items: distinct literature items (of 234) attributed to the study after merging; an item stated by several studies is counted for each. Req.: requirements of Table A2 to which the study contributes. AIS: automatic identification system.

Table A2. Consolidated collision-avoidance requirements.

Code	Requirement	COLREGs
SA1	Proper look-out	R5
SA2	Target classification	R5
SA3	Lights and shapes identification	R5; R18; R20
SA4	AIS as an aid	R5; R7(c)
SA5	Radar observation	R7(b)

Table A2. *Cont.*

Code	Requirement	COLREGs
SA6	Own-ship and target data	R7
SA7	Target track monitoring	R7; R8(d)
SA8	Restricted visibility	R3(k); R3(l); R19(a)
RA1	Collision risk determination	R7(a); R7(c)
RA2	Bearing change check	R7(d)
RA3	CPA-based risk	R7
RA4	Ship domain	R8(c)
RA5	Close-quarters distance	R8(c); R19(d)
RA6	Avoidance start distance	R8(a); R16
RA7	Emergency detection	R17(b)
RA8	Multi-ship risk ranking	R7
RA9	Safe speed	R6
ED1	Encounter classification from relative bearing and contact angle	R13; R14; R15; R16; R17; [3]
ED2	Encounter stage	R17; [3]
ED3	Encounter change handling	R13(d)
ED4	Encounter priority	R13(a)
ED5	Doubt resolved conservatively	R7(a); R13(c); R14(c)
ED6	Responsibilities between vessels	R18
ED7	Encounter per target	R7; R8
CA1	Early action	R8(a); R16
CA2	Large alteration	R8(b); [3]
CA3	Course alteration in sea room	R8(c)
CA4	Safe passing distance	R8(c); R8(d)
CA5	Safe passing angle	R8(c)
CA6	Speed reduction	R8(e)
CA7	Not impeding	R8(f); R9(b–d); R10(i–j)
CA8	Overtaking give-way	R13(a); R13(d)
CA9	Head-on starboard turn	R14(a)
CA10	Crossing give-way	R15; R16
CA11	Give-way action	R16
CA12	Stand-on course keeping	R17(a)(i)
CA13	Stand-on permitted action	R17(a)(ii)
CA14	Stand-on required action	R17(b)
CA15	No turn to port	R14(a); R17(c); R19(d)(i)
CA16	Restricted-visibility conduct	R19(b); R19(d); R19(e)
CA17	Narrow-channel conduct	R9(a); R9(e); R9(f); R34(c); R34(e)
CA18	TSS conduct	R10
CA19	Status-based give-way	R18
CA20	Fishing gear avoidance	R2(a)
CA21	Sound and light signals	R34; R35
CA22	Intention confirmation supplementing early action	R34(d); R36
CA23	Good seamanship	R2(a)
CA24	Rule deviation to avoid danger	R2(b)
CA25	Non-cooperative target	R2(a); R17(a)(ii)
CA26	Mutual avoidance	R2(a)
CA27	Drifting-vessel avoidance	R2(a)
CA28	Consistent multi-ship maneuver	R2; R8(c)
CA29	Maneuverability-based give-way	R2(a); R8(a)
CA30	Maneuver within limits	R8(b)
CA31	Reactive fallback	R8
CA32	Integrated avoidance decision	Parts A–C

The code letters give the decision stage. The COLREGs column gives the provisions, with [3] where the commentary supplies the interpretation. TSS: traffic separation scheme.

Table A3. Vessel types in the collisions of the two tiers.

Vessels Involved	Korean, *n*	Korean, %	International, *n*	International, %
Fishing vessel and fishing vessel	258	36.4	447	13.5
Fishing vessel and merchant vessel	115	16.2	525	15.8
Fishing vessel and other vessel	139	19.6	450	13.6
No fishing vessel	180	25.4	1711	51.6
Unknown	17	2.4	184	5.5
Total	709	100.0	3317	100.0

Merchant: cargo, tanker, container, bulk, and other trading vessels. Other: passenger vessels, tugs and barges, pleasure craft, angling boats, and government vessels. Korean types are taken from the vessel particulars of each ruling. International types are taken from the JTSB report index (1955 reports), the report headers (379) and a keyword screen of the report text (971), so the international shares are approximate.

Table A4. Definitions of case-level facts and clause-heading attribution.

Item	Definition
Fishing engagement	Recorded as engaged in fishing when the record states that gear was being worked, or when the vessel is a fishing vessel whose reported speed over the ground is below 3 kn on the fishing ground (a surrogate; Section 2.3).
JTSB clause heading—admitted terms	Two classes of evidence, both taken from the causal sections only. (i) A term of art that names a rule: the give-way vessel, the stand-on vessel, the crossing rule, whistle and maneuvering signals. (ii) The wording of a causal finding that attributes an avoiding or holding action to an identified vessel. The two classes are counted separately, and the split is reported in Section 3.2.
JTSB clause heading—excluded wording	Generic operational wording not tied to an identified vessel (a speed reduction, a measure to raise attention or the mere mention of radar), and any occurrence outside the causal sections, in particular the statute appendix and the standing safety-recommendation passages, which carry the same vocabulary verbatim.
JTSB clause heading—basis	Across the 437 Japanese collision reports parsed in full when the attribution criteria were fixed (of 1980 in the final tier), the domestic enactment of the COLREGs is named in 7.3%, and an article number appears in 6.9%, falling to 1.8% once statute appendices are excluded. These citations sit in background or recommendation passages, not in the causal findings.

Table A5. Concept categories of the lexicon and the number of cases in which each is mentioned.

Concept Category	Rule	Korean Cases (n = 709)	International Cases (n = 3317)
Look-out	R5	634	2244
Risk of collision	R7	638	1062
Give-way conduct	R16	481	704
Stand-on conduct	R17	493	302

Table A5. *Cont.*

Concept Category	Rule	Korean Cases ($n = 709$)	International Cases ($n = 3317$)
Crossing	R15	186	448
Head-on	R14	85	140
Overtaking	R13	65	358
Restricted visibility	R19	73	190
Narrow channel	R9	85	132
Traffic separation scheme	R10	39	168
Safe speed	R6	284	358
Course alteration	R8	422	713
Starboard action	R8	255	1248
Speed reduction	R8	266	746
Sound signals	R34	463	920
Radar use	R7	604	1651
VHF communication	—	376	663
Seamanship practice	R2(a)	272	155
Special circumstances	R2(b)	82	117
Fishing operation	R18	313	491

A case counts once per category when the category is mentioned at least once (Section 2.2). The rule column gives the provision to which the category is mapped.

Table A6. Fault keywords of the stage-share screen.

Fault Keyword	Stage	Rulings Matched
Neglect of look-out	SA	559
Lights and shapes	SA	130
Asleep on watch	SA	67
Alcohol	SA	22
Neglect of watchkeeping	SA	12
Radar not used	SA	3
Unsafe speed	RA	225
VHF-related fault	RA	85
Failure to judge risk of collision	RA	68
Mistaken encounter status or role	ED	3
Failure of cooperative action under Rule 17(b)	CA	446
Failure to take avoiding action	CA	330
Sound-signal fault	CA	99
Large-alteration fault	CA	29
Total		2078

Keywords are Korean fault phrases matched in the judgment and cause sections (first 4000 characters) of 706 rulings; the English terms translate them. A ruling counts once per keyword. Stage totals are SA 793, RA 378, ED 3, and CA 904 of 2078 findings.

Table A7. Agreement with the cited conduct rules by international investigating authority.

Authority	Reports	Citing Reports	Cited Instances	Contained	Coverage
JTSB (clause-heading attribution)	1980	1857	690	530	76.8%
GISIS depository	538	242	427	269	63.0%
MAIB	46	33	42	19	45.2%
China MSA	337	45	7	1	14.3%
Other authorities	416	96	138	104	75.4%
International tier	3317	2273	1304	923	70.8%

Algorithm A1. Deterministic inference from stated facts to rule obligations

Input: case x with facts: encounter type e_x, visibility v_x, vessel navigational status, fishing state, location and vessel type where reported
Output: applicable rules $A(x)$; obligations $O(x)$; case classifications; inference trace
1: assert individuals: vessels, pairwise encounters (AllDifferent), situation facts
2: $A(x) \leftarrow$ {R5, R6, R7, R8} ▷ Rules 5–8 always apply
3: if v_x = in sight, or e_x is stated and v_x is not (record input only), then $A(x) \leftarrow A(x) \cup R_{enc}(e_x)$ ▷ R13–R17, Equation (3)
4: if not in sight and in or near restricted visibility then $A(x) \leftarrow A(x) \cup$ {R19}
▷ excludes R13–R18; neither stated: Rules 5–8 and situation rules only
5: if narrow channel/TSS then $A(x) \leftarrow A(x) \cup$ {R9/R10}; if in sight and fishing or status conditions apply then $A(x) \leftarrow A(x) \cup$ {R18} ▷ lights and signal rules (R20–R27, R30, R34, R35) added from their own facts
6: $O(x) \leftarrow$ required actions of $A(x)$ ▷ 27 obligation classes via requiresObligation
7: run HermiT: classify MultiShipEncounterCase, RoleConflictCase; check consistency
8: emit trace: the named class or axiom used at every step

References

1. IMO. *Outcome of the Regulatory Scoping Exercise for the Use of Maritime Autonomous Surface Ships (MASS)*; MSC.1/Circ.1638; IMO: London, UK, 2021.
2. IMO. *Convention on the International Regulations for Preventing Collisions at Sea (COLREGs)*; IMO: London, UK, 1972.
3. Cockcroft, A.N.; Lameijer, J.N.F. *A Guide to the Collision Avoidance Rules*, 7th ed.; Butterworth-Heinemann: Oxford, UK, 2012.
4. Belcher, P. A sociological interpretation of the COLREGS. *J. Navig.* **2002**, *55*, 213–224. [CrossRef]
5. He, Y.; Jin, Y.; Huang, L.; Xiong, Y.; Chen, P.; Mou, J. Quantitative analysis of COLREG rules and seamanship for autonomous collision avoidance at open sea. *Ocean Eng.* **2017**, *140*, 281–291. [CrossRef]
6. Huang, Y.; Chen, L.; Chen, P.; Negenborn, R.R.; van Gelder, P.H.A.J.M. Ship collision avoidance methods: State-of-the-art. *Saf. Sci.* **2020**, *121*, 451–473. [CrossRef]
7. Kim, I. Analysis on the responsibility and exemption clause of COLREG Rule 2. *J. Korean Soc. Mar. Environ. Saf.* **2022**, *28*, 54–63. (In Korean) [CrossRef]
8. Coenen, F.P.; Smeaton, G.P.; Bole, A.G. Knowledge-based collision avoidance. *J. Navig.* **1989**, *42*, 107–116. [CrossRef]
9. Lee, H.-J.; Rhee, K.P. Development of collision avoidance system by using expert system and search algorithm. *Int. Shipbuild. Prog.* **2001**, *48*, 197–212.
10. Hasegawa, K.; Kouzuki, A. Automatic collision avoidance system for ships using fuzzy control. *J. Kansai Soc. Nav. Archit. Jpn.* **1987**, *205*, 1–10. (In Japanese)
11. Pietrzykowski, Z.; Wołejsza, P.; Borkowski, P. Decision support in collision situations at sea. *J. Navig.* **2017**, *70*, 447–464. [CrossRef]
12. Constapel, M.; Koch, P.; Burmeister, H.-C. On the implementation of a rule-based system to perform assessment of COLREGs onboard maritime autonomous surface ships. *J. Phys. Conf. Ser.* **2022**, *2311*, 012033. [CrossRef]
13. Benjamin, M.R.; Leonard, J.J.; Curcio, J.A.; Newman, P.M. A method for protocol-based collision avoidance between autonomous marine surface craft. *J. Field Robot.* **2006**, *23*, 333–346. [CrossRef]
14. Kreutzmann, A.; Wolter, D.; Dylla, F.; Lee, J.H. Towards safe navigation by formalizing navigation rules. *TransNav Int. J. Mar. Navig. Saf. Sea Transp.* **2013**, *7*, 161–168. [CrossRef]
15. Woerner, K.; Benjamin, M.R.; Novitzky, M.; Leonard, J.J. Quantifying protocol evaluation for autonomous collision avoidance: Toward establishing COLREGS compliance metrics. *Auton. Robots* **2019**, *43*, 967–991. [CrossRef]
16. Krasowski, H.; Althoff, M. Temporal logic formalization of marine traffic rules. In Proceedings of the 2021 IEEE Intelligent Vehicles Symposium (IV), Nagoya, Japan, 11–17 July 2021; pp. 186–192. [CrossRef]
17. Torben, T.R.; Glomsrud, J.A.; Pedersen, T.A.; Utne, I.B.; Sørensen, A.J. Automatic simulation-based testing of autonomous ships using Gaussian processes and temporal logic. *Proc. Inst. Mech. Eng. Part O J. Risk Reliab.* **2023**, *237*, 293–313. [CrossRef]
18. Kang, S.-S.; Lee, Y.-I.; Jung, H.; Kim, Y.-G. Collision avoidance system for intelligent ship based on BK-products and COLREGs. *J. Korea Inst. Marit. Inf. Commun. Sci.* **2007**, *11*, 181–190. (In Korean)
19. Ali, F.; Kim, E.K.; Kim, Y.-G. Type-2 fuzzy ontology-based semantic knowledge for collision avoidance of autonomous underwater vehicles. *Inf. Sci.* **2015**, *295*, 441–464. [CrossRef]
20. Wen, Y.; Zhang, Y.; Huang, L.; Zhou, C.; Xiao, C.; Zhang, F.; Peng, X.; Zhan, W.; Sui, Z. Semantic modelling of ship behavior in harbor based on ontology and dynamic Bayesian network. *ISPRS Int. J. Geo-Inf.* **2019**, *8*, 107. [CrossRef]
21. Zhong, S.; Wen, Y.; Huang, Y.; Cheng, X.; Huang, L. Ontological ship behavior modeling based on COLREGs for knowledge reasoning. *J. Mar. Sci. Eng.* **2022**, *10*, 203. [CrossRef]

22. Hatłas-Sowińska, P.; Wielgosz, M. Ontology based approach in solving collision situations at sea. *Ocean Eng.* **2022**, *260*, 111941. [CrossRef]
23. Zhou, C.; Wen, K.; Zhao, J.; Bian, Z.; Lu, T.; Ko Ko Latt, M.; Wang, C. Ontology-based method for identifying abnormal ship behavior: A navigation rule perspective. *J. Mar. Sci. Eng.* **2024**, *12*, 881. [CrossRef]
24. Sabatino, N.; Porello, D.; Zaccone, R. An ontology to classify marine encounters according to COLREG. In *Proceedings of the FOIS 2025 Satellite Events, Catania, Italy, 10–12 September 2025*; CEUR-WS.org: Aachen, Germany, 2025; Volume 4111. Available online: https://ceur-ws.org/Vol-4111/SC_Paper_3.pdf (accessed on 18 October 2026).
25. Chauvin, C.; Lardjane, S.; Morel, G.; Clostermann, J.-P.; Langard, B. Human and organisational factors in maritime accidents: Analysis of collisions at sea using the HFACS. *Accid. Anal. Prev.* **2013**, *59*, 26–37. [CrossRef]
26. Hwang, T.; Youn, I.-H. Latent-cause extraction model in maritime collision accidents using text analytics on Korean maritime accident verdicts. *Appl. Sci.* **2022**, *12*, 914. [CrossRef]
27. Gan, L.; Ye, B.; Huang, Z.; Xu, Y.; Chen, Q.; Shu, Y. Knowledge graph construction based on ship collision accident reports to improve maritime traffic safety. *Ocean Coast. Manag.* **2023**, *240*, 106660. [CrossRef]
28. Liu, D.; Cheng, L. MAKG: A maritime accident knowledge graph for intelligent accident analysis and management. *Ocean Eng.* **2024**, *312*, 119280. [CrossRef]
29. Hao, Y.; Li, H.; Fu, S.; Gu, S.; Mao, W. Enhanced human factor and causation analysis in maritime accidents using large language models. *Ocean Eng.* **2026**, *356*, 125125. [CrossRef]
30. Escribano Arias, D.; Gomez-Lendinez, D.; Navas de Maya, B.; Velasco-Gallego, C. Retrieval-augmented generation for maritime accident report analysis: Evaluating large language models on performance and cybersecurity. *J. Mar. Sci. Eng.* **2026**, *14*, 983. [CrossRef]
31. Mantani, S. Study on the Differences between the Operators of Fishing Vessels and General Power-Driven Vessels in Their Interpreting the Applicable Provisions of the Act on Preventing Collisions at Sea: Verification of Standard Values for the Current Law, Looking Ahead to Autonomously Operated Vessels. Ph.D. Thesis, Kobe University, Kobe, Japan, 2021. (In Japanese)
32. Shu, Y.; Dong, A.; Ye, B.; Liu, C.; Gan, L.; Song, L. Liability division for ship collision accidents based on ontology model and Bayesian network. *Ocean Coast. Manag.* **2025**, *269*, 107824. [CrossRef]
33. Zhang, Z.; Yang, X.; Sun, L.; Sun, Y.; Kang, J. Research on constructing and reasoning the collision knowledge graph of autonomous navigation ship based on enhanced BERT model. *Expert Syst. Appl.* **2025**, *278*, 127429. [CrossRef]
34. Ohn, S.-W.; Namgung, H. Requirements for optimal local route planning of autonomous ships. *J. Mar. Sci. Eng.* **2023**, *11*, 17. [CrossRef]
35. Gleeson, J.; Dunbabin, M.; Ford, J.J. COLREG scenario classification and compliance evaluation with temporal and multi-vessel awareness for collision avoidance systems. *Ocean Eng.* **2024**, *313*, 119552. [CrossRef]
36. Tengesdal, T.; Johansen, T.A. Simulation framework and software environment for evaluating automatic ship collision avoidance algorithms. In Proceedings of the 2023 IEEE Conference on Control Technology and Applications (CCTA), Bridgetown, Barbados, 16–18 August 2023; pp. 186–193. [CrossRef]
37. Hagen, I.B.; Vassbotn, O.; Skogvold, M.; Johansen, T.A.; Brekke, E.F. Safety and COLREG evaluation for marine collision avoidance algorithms. *Ocean Eng.* **2023**, *288*, 115991. [CrossRef]
38. Hagen, I.B.; Murvold, M.N.; Johansen, T.A.; Brekke, E.F. Grounding hazard considerations in evaluation of COLREGS collision avoidance algorithms. *Ocean Eng.* **2024**, *308*, 118204. [CrossRef]
39. Akdağ, M.; Solnør, P.; Johansen, T.A. Collaborative collision avoidance for Maritime Autonomous Surface Ships: A review. *Ocean Eng.* **2022**, *250*, 110920. [CrossRef]
40. Campbell, S.; Naeem, W.; Irwin, G.W. A review on improving the autonomy of unmanned surface vehicles through intelligent collision avoidance manoeuvres. *Annu. Rev. Control* **2012**, *36*, 267–283. [CrossRef]
41. Ramos, M.A.; Utne, I.B.; Mosleh, A. Collision avoidance on maritime autonomous surface ships: Operators' tasks and human failure events. *Saf. Sci.* **2019**, *116*, 33–44. [CrossRef]
42. Öztürk, Ü.; Akdağ, M.; Ayabakan, T. A review of path planning algorithms in maritime autonomous surface ships: Navigation safety perspective. *Ocean Eng.* **2022**, *251*, 111010. [CrossRef]
43. Krasowski, H.; Althoff, M. Provable traffic rule compliance in safe reinforcement learning on the open sea. *IEEE Trans. Intell. Veh.* **2024**, *9*, 7617–7634. [CrossRef]
44. Ahmed, Y.A.; Hannan, M.A.; Oraby, M.Y.; Maimun, A. COLREGs compliant fuzzy-based collision avoidance system for multiple ship encounters. *J. Mar. Sci. Eng.* **2021**, *9*, 790. [CrossRef]
45. Zaccone, R. A dynamic programming approach to the collision avoidance of autonomous ships. *Mathematics* **2024**, *12*, 1546. [CrossRef]
46. Lan, Z.; Gang, L.; Zhang, M.; Xie, W.; Wang, S. A multi-stage collision avoidance model for autonomous ship based on fuzzy set theory with TL-DDQN algorithm. *Ocean Eng.* **2024**, *311*, 118912. [CrossRef]

47. Li, Y.; Cheng, Z.; Zhang, Y.; Wu, B.; Guedes Soares, C. Quantitative critical ship arena based on collision accident reports to pre-warn close-quarters situations. *Ocean Eng.* **2025**, *322*, 120470. [CrossRef]
48. Korea Maritime Safety Tribunal (KMST). Adjudications of Marine Accidents (In Korean). Available online: https://www.kmst.go.kr (accessed on 18 October 2026).
49. IMO. GISIS: Marine Casualties and Incidents Module. Available online: https://gisis.imo.org/Public/ (accessed on 18 October 2026).
50. García Maza, J.A.; Poo Argüelles, M.R. COLREGs and their application in collision avoidance algorithms: A critical analysis. *Ocean Eng.* **2022**, *261*, 112029. [CrossRef]
51. Namgung, H. Local route planning for collision avoidance of maritime autonomous surface ships in compliance with COLREGs rules. *Sustainability* **2022**, *14*, 198. [CrossRef]
52. Ohn, S.-W.; Namgung, H. Interval type-2 fuzzy inference system based on closest point of approach for collision avoidance between ships. *Appl. Sci.* **2020**, *10*, 3919. [CrossRef]
53. Xie, W.; Gang, L.; Zhang, M.; Liu, T.; Lan, Z. Optimizing multi-vessel collision avoidance decision making for autonomous surface vessels: A COLREGs-compliant deep reinforcement learning approach. *J. Mar. Sci. Eng.* **2024**, *12*, 372. [CrossRef]
54. Maritime and Coastguard Agency. *Navigation: Watchkeeping Safety—Use of Very High Frequency (VHF) Radio and Automatic Identification System (AIS)*; MGN 324 (M+F) Amendment 2; MCA: Southampton, UK, 2022. Available online: https://www.gov.uk/government/publications/mgn-324-mf-amendment-2-navigation-watchkeeping-safety-use-of-very-high-frequency-vhf-radio-and-automatic-identification-system-ais (accessed on 18 October 2026).
55. IMO. *Standards for Ship Manoeuvrability*; Resolution MSC.137(76); IMO: London, UK, 2002.
56. Mohapatra, S.C.; Amouzadrad, P.; Guedes Soares, C. Recent developments in the nonlinear hydroelastic modeling of sea ice interaction with marine structures. *J. Mar. Sci. Eng.* **2025**, *13*, 1410. [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.