Application of Data Linkage Techniques to Pacific Northwest Commercial Fishing Injury and Fatality Data

Jasmine Nahorniak (✉ jasmine.nahorniak@oregonstate.edu)

College of Earth, Ocean, and Atmospheric Sciences, 104 CEOAS Admin Bldg., Oregon State University,
Corvallis OR 97331, USA <https://orcid.org/0000-0003-4241-129X>

Viktor Bovbjerg

Oregon State University College of Health and Human Sciences

Samantha Case

National Institute of Occupational Safety and Health

Laurel Kincl

Oregon State University College of Health and Human Sciences

Original Contribution

Keywords: Fishing, Vessel, Safety, Injury, Fatality, Linkage

Posted Date: January 27th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-151522/v1>

License: This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Version of Record: A version of this preprint was published at Injury Epidemiology on July 5th, 2021. See the published version at <https://doi.org/10.1186/s40621-021-00323-z>.

Abstract

Background

Commercial fishing consistently has among the highest workforce injury and fatality rates in the United States. Data related to commercial fishing incidents are routinely collected by multiple organizations which do not currently coordinate or automatically link data. Each dataset has the potential to generate a more complete picture to inform prevention efforts. Our objective was to examine the utility of using statistical data linkage methods to link these datasets in support of incident surveillance and hazard assessment in the commercial fishing industry.

Methods

In this feasibility study, we identified true matches and discrepancies between de-identified datasets using the Python Record Linkage Toolkit. Four commercial fishing datasets from Oregon and Washington were linked: the Commercial Fishing Incident Database, the Vessel Casualty Database, the Nonfatal Injuries Database, and the Oregon Trauma Registry. The datasets each covered different date ranges within 2000 - 2017, containing 458, 524, 184, and 11 cases respectively. Several data linkage classifiers were evaluated.

Results

The Naïve-Bayes classifier returned the highest number of true matches between these small datasets. A total of 41 true matches and 8 close matches were identified, of which 29 were determined to be duplicates. In addition, linkage highlighted 4 records that were not commercial fishing cases from Oregon and Washington. The optimum match parameters were the date, state, vessel official number, and number of people on board.

Conclusions

Statistical data linkage enables accurate, routine matching for small de-identified injury and fatality datasets such as those in commercial fishing. It provides information needed to improve the accuracy of existing data records. It also enables expanding and sharpening details of individual incidents in support of occupational safety research.

Background

Commercial fishing consistently has one of the highest workforce injury/fatality rates in the United States (Lucas & Case, 2018; Bureau of Labor Statistics, 2018). To gain a clear understanding of the hazards and opportunities for incident and injury prevention in the commercial fishing industry, it is critical to have a more complete picture of injury characteristics and burden in this workforce. Effective safety strategies can be informed by knowledge of the circumstances leading up to the incident, actions taken to prevent or minimize injury, resulting injuries or fatalities, and injury treatment and outcomes.

US data related to commercial fishing incidents are routinely collected by multiple organizations. Linking these datasets is not straightforward. Each dataset has strict restrictions on data use, unique data request procedures and costs, different inclusion criteria, different data collection and maintenance methods, varying update rates and data lags, different storage formats, and differing data elements and data definitions. Safety-related datasets also often come with multiple challenges: small numbers of incidents, inconsistent values across datasets, and missing data. In addition, personally identifiable information, such as the victim's name, is secured by the source and usually not available for matching. Evaluating the level of difficulty of obtaining routine access to and working with these datasets is a critical component of the overall project under which this linking feasibility study was completed.

Statistical data linkage techniques exist that not only help identify true matches in these cases, but also provide statistics describing the linkage confidence. The utility of statistical data linkage for linking de-identified data has been successfully demonstrated in other health-related fields (Conderino et al., 2017). A recent study used a combination of data linkage software and manual verification to link Alaskan commercial fishing data with moderate success (Syron, in press).

Here, we assess an alternative method to provide routine, accurate, automated data linkage of commercial fishing data. The method is applied to data from Oregon and Washington, two states with active commercial fisheries on the west coast of the United States. While this study focused on commercial fishing incident data from the Pacific Northwest, the statistical matching methods presented here could be easily adapted to other regions, industries and occupations.

Methods

Four datasets were linked in pairs to determine the overlap between the datasets. Data linkage software was used to automatically suggest potential matches by comparing the link probabilities to a threshold, then confirmed manually. The following sections describe the datasets and linkage approach.

Datasets

The following datasets with commercial fishing incident data from Oregon and Washington were studied: the Commercial Fishing Incident Database (CFID), the Vessel Casualty (VC) database, the Nonfatal Injuries (NFI) database, and the Oregon Trauma Registry (OTR).

The Commercial Fishing Incident Database (CFID) contains information regarding commercial fishing vessel disasters and fatalities due to traumatic injuries from the entire United States (CDC/NIOSH, 2019). The CFID definition of a vessel disaster is an event such as sinking that forces the crew to abandon the vessel because it is no longer safe to remain onboard. The types of data collected include the date, time, and location of the incident, vessel details, contributing factors, and personnel injury and fatality information. A single incident may involve injuries to and/or fatalities of multiple personnel. The original sources of data used to populate this database include United States Coast Guard (USCG) reports and news articles. CFID was developed and is actively maintained by the Centers for Disease Control and

Prevention's (CDC) National Institute for Occupational Safety and Health (NIOSH), Western States Division.

The Vessel Casualty database (VC) recorded information about commercial fishing vessel-related incidents in Alaska, Oregon, and Washington that (a) are not classified as vessel disasters and (b) did not involve any fatalities. These incidents tend to be less serious but still present problems with vessel systems that can put crewmembers at risk, such as loss of power, propulsion, or steering (Case & Lucas, in press). This dataset was maintained by NIOSH and merged with CFID in 2020. Cases were originally obtained from USCG reports. This dataset provides information about the incident date, time, location, and circumstances, in addition to the vessel information. Unlike CFID, this dataset does not include any personnel information, focusing instead on vessel damage. For this study, data were requested from Oregon and Washington only.

The Nonfatal Injuries (NFI) database was developed at NIOSH to complement the information recorded in CFID. To date, this dataset has covered Alaska, Washington, Oregon, and California. The NFI records injuries sustained during commercial fishing that are not recorded in CFID, such as those incurred while working on deck. Cases were originally obtained from USCG reports. Variables in NFI are similar to those in CFID; in addition to personnel demographics and injury characteristics, the NFI also contains vessel information. Data were requested from Oregon, Washington, and California only.

The Oregon Trauma Registry (OTR) (Oregon Health Authority Public Health Division, 2019) includes information concerning all Oregon patients who either entered into the trauma system in Oregon or met specific clinical- or admission-based criteria for inclusion in the registry, based either on field entry by EMS responders or by the activation of a trauma team or surgeon at a receiving hospital. Information recorded includes patient demographics (including occupation), date, time and location of incident, emergency service response, injury circumstances and details, medical procedures performed, length of stay, insurance, and costs. Data were requested for patients with work-related injuries and occupations of farming/fishing/forestry. This dataset was further pruned using incident location and narratives to include only fishing-related cases. The OTR is maintained by the Injury and Violence Prevention Program of the Oregon Health Authority Public Health Division.

Each dataset used in this study contains data from different date ranges and regions (Table 1). The date range varied by data source due to lags in data abstraction, coding, and review, or to availability of data elements during specific time periods.

By definition, (a) the CFID dataset should not overlap the VC or NFI datasets, (b) the VC and NFI dataset may overlap, and (c) the OTR dataset could overlap with any of them (Fig. 1).

Data linkage method

Data linkage is a statistical technique used to identify records from two datasets that likely describe the same event. The two datasets must have some parameters in common (i.e., the matching variables) that

can be used to distinguish events and link the records. Every record in one dataset is compared with every record in a second dataset. The likelihood that two records match (their match probability) is determined by comparing the contents of the matching variables for that pair. Match probabilities range from 0 to 1. Any record pair with a match probability above a specified threshold is designated as a link. Those below the threshold are designated as non-links. In our project, this was followed by a manual review process where all links were examined further to identify true matches.

Matching variables must be selected carefully. Ideal matching variables are independent, reliable, and complete. A major aim of this project is to determine the feasibility of using data linkage methods with commercial fishing incident data when personally identifiable information (PII) is not available. Depending on the two datasets involved, the matching variables used for linking in this study were some combination of: Incident Date, Incident State, Vessel Official Number, and Latitude/Longitude. These independent variables were identified during preliminary data linkage analyses to be the strongest indicators of links.

The linking results presented here were derived using components of the Python Record Linkage Toolkit software (De Bruin, 2019). This toolkit includes several data linkage classifiers, which use different methods to separate record pairs into links and non-links. The quality of the performance of each classifier depends on the dataset involved. Each of the classifiers described below were tested to determine the optimum classifier for our datasets.

Classifiers can be divided into two groups: supervised and unsupervised. Supervised classifiers require training using a "golden data set", a subset of the data where the true match status is known. Unsupervised classifiers, on the other hand, do not require training.

For the supervised classifiers, a golden data set was derived for each pair of datasets to be linked. First, a rudimentary approach was used to identify a small list of potential matches. True matches were then verified manually. Next, a set of non-matches was derived by creating fake records from scrambled real records. Finally, a golden data set for the dataset pair was created consisting of a combination of these verified true- and non-matches. For this study, the supervised classifiers used were Naïve-Bayes (NB), Logistic Regression (LR), and Support Vector Machine (SVM). Classifier definitions are provided in the Supplementary Materials.

The unsupervised classifier used in this study was the Expectation/Conditional Maximization Algorithm (ECM), a probabilistic classifier closely related to both the Naïve-Bayes Classifier and the probabilistic Fellegi and Sunter (1969) approach.

Quality metrics

Quality metrics provided by data linkage include TP (the number of True Positives), FP (False Positives), FN (False Negatives), and TN (True Negatives). For supervised classifiers, these describe how well the linkage technique performed on the golden data set. Additional metrics can be derived including

precision, recall and f-score (Eqns. 1-3, respectively). Each of these three metrics (precision, recall, f-score) can range in value from 0 (worst) to 1 (best).

$$precision = \frac{true\ positives}{total\ predicted\ positives} = \frac{TP}{TP + FP} \quad (1)$$

$$recall = \frac{true\ positives}{total\ actual\ positives} = \frac{TP}{TP + FN} \quad (2)$$

$$f-score = 2 * \frac{precision * recall}{precision + recall} \quad (3)$$

The f-score is a summary metric of the classifier's performance that balances precision and recall. Generally, a higher f-score indicates better performance.

Results

Data linkages were assessed for all combinations of the four datasets (a total of six pairs), regardless of whether matches were expected, to determine the efficacy of the method. In addition, the data linkage process was repeated for each of the four classifiers, for a total of 24 runs. The data linkage quality metrics are summarized in Table 2. No matches were found between VC & OTR. After manual review, a total of 41 true matches plus 8 close matches were found (Table 3). Four cases overlapped three datasets (Table 3, Fig. 2).

Three of the classifiers (NB, ECM, LR) provide match probability values for each possible link. The probabilities range from 0 (not a match) to 1 (exact match) and indicate the confidence that the pair is a true match. The probabilities output from the NB classifier for every possible link in the pairs of datasets are shown in Fig. 3.

For the 41 true matches, the values of some common parameters were compared to assess the relative accuracy and completeness of the records (Table 4). Note that not all parameters listed in the table are provided by all datasets.

Discussion

Overlapping information from multiple sources (true matches) can serve multiple purposes. For dataset pairs where overlaps are expected, the matching rows can be compared for accuracy and merged to obtain a more complete story. For dataset pairs where no overlaps are expected yet overlaps are discovered, the overlap highlights possible errors or duplicates. And finally, for two datasets where overlaps between them are expected but not found, novel cases can be identified. Thus, each dataset has

the potential to identify new incidents, assess data accuracy, or provide complementary data to generate a more complete picture.

True matches and discrepancies between four commercial fishing datasets were successfully identified using the Python Record Linkage toolkit with different classifiers. The match results highlight the relative accuracy of parameters common to these datasets.

The small sizes of the commercial fishing datasets, ranging from 11 to 1316 cases, provides both advantages and challenges. The small number of incidents, spread over many years, means that Incident Date is a strong matching variable (the likelihood of two incidents occurring on the same day is small). On the flip side, the small datasets make it difficult to create a representative golden data set for the supervised classifiers.

Due to the small number of cases, it is important to capture every match. Hence, the optimal data linkage approach errs on returning false positives rather than false negatives. Ideally (a) all of the matches previously identified by the golden data set are discovered by the linkage tool, which implies also that (b) no false negatives are returned. In addition, it is helpful if the absolute number of false positives returned is low to limit the amount of manual review required to identify true matches.

Four classifiers were evaluated (ECM, SVM, NB, and LR) for each combination of datasets. The classifier that was the least successful at identifying matches for these small datasets was SVM. The binary nature of this classifier meant that the threshold cannot be changed from 0.5, so improvement of these results was not possible. The ECM classifier, on the other hand, performed very well; only one match was missed using a threshold of 0.5. The strongest performers were NB and LR, both of which successfully found all TPs but only after lowering the thresholds to 0.01 or 0.005. The optimum classifier was NB which returned an additional true match between NFI and OTR that was not found by any of the other classifiers. This particular match was difficult to find as the incident state information was missing.

Typically, a threshold of 0.5 is used to separate possible matches from non-matches. However, it can be seen from Fig. 3 that a much lower threshold (0.005) was necessary to capture all of the true matches in these particular datasets. This speaks to the inherent data variability both within and between the datasets. For example, while date was found to be a relatively reliable field, the date of the incident and the date of entry into a trauma registry may not always coincide due to the remoteness of the work.

Of the four datasets, the OTR was the most challenging to link. While OTR identifies the industry and occupation of injury cases, the latitude and longitude fields are often blank and vessel information is not recorded. This limited matching variables for OTR to Incident Date and Incident State.

The low number of records in OTR (eleven) could indicate that (a) the number of commercial fishing injuries that were classified as trauma was extremely low, (b) not all commercial fishing related traumas were captured by the OTR variables "work-related" and "farming / fishing / logging", and/or (c) some traumatic injuries did not make it into the trauma registry.

These OTR results and conclusions may not be representative of other state trauma registries. Trauma registry datasets differ from state to state in their format, variables captured, access policies, and data request procedures. This is also true of many other sources of occupational health and safety data. Accommodating and accounting for this variety is a significant challenge to studies that span multiple states.

Another challenge occurs when multiple personnel are involved in a single incident. In this situation, a single record in one dataset may link correctly with multiple records in another dataset if the matching variables do not allow individual personnel to be distinguished. This results in false positives. For example, if three personnel are injured in a single incident and are all recorded in both CFID and NFI, it may not be possible to identify true patient-level matches with certainty unless the records contain enough information to distinguish them. This study identified eight injury cases that could not be matched with certainty down to the person level; they were designated as close matches (Table 3).

Data linkage is also useful for identifying extraneous records. For example, if overlaps are found between datasets that should be distinct, such as CFID and VC, the results can be used to purge the dataset(s) of the errant records. CFID stores vessel disasters and fatalities while VC stores nonfatal vessel casualties. By definition, there should be no overlap between VC and CFID, however nine true matches were found. These particular cases were unintentionally recorded twice; once in CFID classified as vessel disasters and again in VC classified as vessel casualties. This highlighted a potential issue in the vessel disaster classification process which was then corrected.

Similarly, there should be no overlap between CFID and NFI, however 12 true matches and an additional 8 close matches were found. The close matches involved incidents where multiple personnel were involved and could not be distinguished. All 20 incidents involved vessel disasters and/or fatal incidents and hence they should have only been recorded in CFID and not NFI. In addition to the previous example, this illustrates the utility of data linkage in the data cleaning process.

After accounting for the duplicates, 16 true matches remain out of the original 49. Five are matches between OTR and CFID, two are matches between OTR and NFI, and nine are matches between NFI and VC. Since each data source provides information about different aspects of an incident, these matches provide a wealth of additional information about each case that would otherwise be unavailable.

Valuable information can also be gleaned from the absence of expected true matches. For example, four of the eleven OTR cases were not found in any of the other commercial fishing incident datasets, prompting closer examination. Two of the four OTR cases had an incident location of Crescent City, California, which is outside of the area covered by the other datasets. Based on their narratives, the remaining two cases may have involved recreational fishing rather than commercial fishing; CFID, VC, and NFI do not record recreational fishing incidents. Hence, non-matches with these four OTR cases are justified.

Arguably the greatest challenge in the analysis of commercial fishing injury and fatality data is unrecorded incidents. While vessel disasters and fatalities are well-captured via mandated USCG reporting requirements, nonfatal injuries are likely to be under-reported. This may result from concerns such as medical cost, future employment, fear of reprimand, insurance impact, and liability (Pransky et al., 1999). Vessel owners/operators may also be unaware of the mandatory casualty reporting requirements; a Maine study found that more than 40% of commercial fishing vessels were out of compliance with regulatory safety requirements (Davis, 2011). Trauma registries are also incomplete sources of injury information; only those cases that meet specific clinical- or admission-based criteria, based either on field entry by EMS responders or by the activation of a trauma team or surgeon at a receiving hospital, are included.

Conclusions

Effective safety measures depend on accurate and complete information about potential hazards. Data linkage is a valuable tool that enables information from various sources to be merged, potentially yielding a more detailed picture of incidents from inception to outcome. In this study, four de-identified commercial fishing datasets were successfully linked using the Python Record Linkage Toolkit. Various classifiers were tested; the optimum classifier for this particular study was found to be the Naïve-Bayes Classifier (NB). A total of 41 true matches and 8 close matches were identified.

Data linkage also provides a means to assess the relative accuracy of common parameters. Of the parameters examined, the most reliable across the commercial fishing datasets were Incident Date, Incident State, Vessel Official Number, and the Number of People on Board. Knowledge of parameter reliability is essential for guiding appropriate matching variable choices for future data linkage analyses.

This effort is currently being expanded to include other geographic areas along the West Coast and additional data sources. This approach could further be tailored to a national level for commercial fishing, and/or to other occupational injury settings.

The outcomes of this study, the true matches, are also being assessed to better understand the injury causes, contributors, and outcomes to help inform prevention efforts.

Abbreviations

CDC: Centers for Disease Control and Prevention

CFID: Commercial Fishing Incident Database

ECM: Expectation/Conditional Maximization Algorithm

EMS: Emergency Medical Services

FN: False negatives

FP: False positives

IRB: Institutional Review Board

LR: Logistic Regression

NB: Naïve-Bayes

NFI: Nonfatal Injuries Database

NIOSH: National Institute for Occupational Safety and Health

OTR: Oregon Trauma Registry

SVM: Support Vector Machine

TN: True negatives

TP: True positives

USCG: United States Coast Guard

VC: Vessel Casualty Database

Declarations

Ethics approval and consent to participate

This project did not involve collecting primary data nor contacting participants. All of the data for this project came from existing data sources and did not include individually identifiable data. Oregon State University Institutional Review Board (IRB) approval was obtained prior to obtaining and working with these data (Study ID 7633). Secure data practices were followed throughout this research.

Consent for publication

Not applicable

Availability of data and materials

This study brought together existing data obtained upon request from a number of different sources. These datasets contain sensitive information and are not publicly available. These data may be requested from the original data providers (CDC/NIOSH and the Oregon Trauma Registry).

Competing interests

The authors declare that they have no competing interests.

Funding

This work was supported by the Centers for Disease Control and Prevention CDC/NIOSH Cooperative Agreement Number U54OH007544-16.

Authors' contributions

JN developed the methodology, performed the data linkage analysis, and wrote the original draft. VB procured and extracted the OTR data. SC provided assistance with the CDC/NIOSH data. LK conceptualized and managed the project. All authors contributed to writing the manuscript.

Acknowledgements

We would like to thank the diligent efforts by all personnel involved in recording and providing access to the commercial fishing incident data used in this study. This includes the United States Coast Guard (USCG), the National Institute for Occupational Health (NIOSH), and the Oregon Health Authority. We would also like to thank Alice Shumate (CDC/NIOSH) for valuable comments on the manuscript.

References

1. Bureau of Labor Statistics. 2018 Census of Fatal Occupational Injuries. 2018.
2. Case SL, Lucas DL. Predicting commercial fishing vessel disasters through a novel application of the theory of man-made disasters. J Safety Res. 2020; in press.
3. CDC/NIOSH. Commercial Fishing Incident Database. 2019.
4. Conderino S, Fung L, Sedlar S, Norton JM. Linkage of traffic crash and hospitalization records with limited identifiers for enhanced public health surveillance. Accid Anal Prev. 2017;101:117–23. <http://dx.doi.org/10.1016/j.aap.2017.02.011>.
5. Davis ME. Occupational safety and regulatory compliance in US commercial fishing. Arch Environ Occup Heal. 2011. <https://doi.org/10.1080/19338244.2011.564237>.
6. De Bruin J. Python Record Linkage Toolkit: A toolkit for record linkage and duplicate detection in Python. Zenodo; 2019. <https://doi.org/10.5281/zenodo.3559043>.
7. Dempster AP, Laird NM, Rubin DB. Maximum Likelihood from Incomplete Data Via the EM Algorithm. J R Stat Soc Ser B. 1977. <https://doi.org/10.1111/j.2517-6161.1977.tb01600.x>.
8. Fellegi IP, Sunter AB. A Theory for Record Linkage. J Am Stat Assoc. 1969;64(328):1183–210.
9. Lucas DL, Case SL. Work-related mortality in the US fishing industry during 2000-2014: New findings based on improved workforce exposure estimates. Am J Ind Med. 2018. <https://doi.org/10.1002/ajim.22761>.
10. Meng X-L, Rubin DB. Maximum Likelihood Estimation via the ECM Algorithm: A General Framework. Biometrika. 1993. <https://doi.org/10.2307/2337198>.
11. Oregon Health Authority Public Health Division. Oregon Trauma Registry. 2019.

12. Pransky G, Snyder T, Dembe A, Himmelstein J. Under-reporting of work-related disorders in the workplace: A case study and review of the literature. *Ergonomics*. 1999.
<https://doi.org/10.1080/001401399185874>.
13. Syron L. Linking datasets to characterize injury and illness in Alaska’s fishing industry. *J Agromedicine*. 2020; in press.
14. United States Coast Guard, 46 CFR 4.05-1

Tables

Table 1 Number of commercial fishing incidents recorded in each of the datasets used in this study.

Dataset	Date Range (YYYY-MM-DD)	Region	Total ^a	OR/WA ^a
CFID	2000-01-04 — 2017-12-04	All USA	1315 (2966)	194 (458)
VC	2010-08-15 — 2014-12-31	OR/WA	524 (0)	524 (0)
NFI	2002-01-12 — 2016-10-19	OR/WA/CA	232 (232)	184 (184)
OTR	2009-05-08 — 2016-07-06	OR	11 (11)	11 (11)

^a Each incident may involve multiple personnel. The number between parentheses is the total number of personnel cases.

Table 2 Linkage metrics TP, FP, FN, TN, and f-score for each classifier per pair of datasets.

CFID & OTR Match Parameters: Incident Date, Incident State Combinations (2966 * 11): 32,626 Golden Matches: 5						
Classifier	Threshold	TP	FP	FN	TN	f-score
ECM	0.5	4	6	1	32,615	0.53
SVM	0.5	0	0	5	32,621	0
NB	0.005	5	29	0	32,592	0.26
LR	0.005	5	29	0	32,592	0.26
CFID & VC Match Parameters: Incident Date, Vessel Official Number, Latitude/Longitude Combinations (1315 * 524): 689,060 Golden Matches: 9						
Classifier	Threshold	TP	FP	FN	TN	f-score
ECM	0.5	9	3	0	689,048	0.86
SVM	0.5	8	0	1	689,051	0.94
NB	0.005	9	3	0	689,048	0.86
LR	0.005	9	7	0	689,044	0.72
CFID & NFI Match Parameters: Incident Date, Vessel Official Number, Latitude/Longitude Combinations (2966 * 232): 688,112 Golden Matches: 12						
Classifier	Threshold	TP	FP	FN	TN	f-score
ECM	0.5	12	52	0	688,048	0.32
SVM	0.5	0	0	12	688,100	0
NB	0.005	12	52	0	688,048	0.32
LR	0.005	12	52	0	688,048	0.32
NFI & VC Match Parameters: Incident Date, Vessel Official Number, Latitude/Longitude Combinations (232 * 524): 121,568 Golden Matches: 10						
Classifier	Threshold	TP	FP	FN	TN	f-score

ECM	0.5	10	13	0	121,545	0.61
SVM	0.5	9	1	1	121,557	0.90
NB	0.01	10	2	0	121,556	0.91
LR	0.01	10	13	0	121,545	0.61
NFI & OTR Match Parameters: Incident Date, Incident State Combinations (232 * 11): 2,552 Golden Matches: 4						
Classifier	Threshold	TP	FP	FN	TN	f-score
ECM	0.2	4	3	0	2,545	0.73
SVM	0.5	0	0	4	2,548	0
NB	0.005	4	3	0	2,545	0.73
LR	0.005	4	7	0	2,541	0.53
VC & OTR Match Parameters: Incident Date, Incident State Combinations (524 * 11): 5,764 Golden Matches: 0						

Table 3 Number of true matches found for each dataset combination.

Combination				Matches ^a
Two-Dataset				
CFID	OTR			5
CFID		VC		9
CFID			NFI	12 (20)
		VC	NFI	10
	OTR		NFI	5
	OTR	VC		0
Total				41 (49)
Multi-Dataset				
CFID	OTR		NFI	3
CFID		VC	NFI	1
CFID	OTR	VC		0
	OTR	VC	NFI	0
CFID	OTR	VC	NFI	0
Total				4

^a The numbers in parentheses include close matches.

Table 4 Relative accuracy and completeness of data within the 41 true match pairs.

Parameter	Relative Accuracy	Completeness	CFID ^a	VC	NFI	OTR
Incident Date	+/- 1 day	complete	x	x	x	x
Incident Time	varied from complete agreement to no agreement	AM/PM designation and time zone often missing	x	x	x	x
Incident State	always agreed	complete	x	x	x	x
Latitude/Longitude	+/- 0.5 degrees (50 km)	often missing from OTR	x	x	x	x
Miles from Shore	never agreed	complete	x	x	x	
Vessel Official Number	always agreed	sometimes unavailable; state number used instead	x	x	x	
# People on Board	always agreed	only occasionally missing	x	x	x	
Narrative	consistent stories; matching provides additional details	rarely missing	x	x	x	x

^a The rightmost columns indicate the datasets that provide the listed parameters.

Figures

Figure 1

Schematic of the expected overlap between the various datasets (CFID, OTR, NFI, and VC) for commercial fishing incidents (not to scale).

Figure 2

Schematic illustrating the number of true and close matches found between the four datasets (not to scale).

Figure 3

Match probabilities from the NB classifier for all possible links found across different dataset pairs. The threshold of 0.005 is shown as a solid line. A probability of 0.5 is indicated by a dashed line. True matches are highlighted with red circles. The same information is plotted on a log scale (upper panel) and linear scale (lower panel).

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [ClassifierDefinitions.docx](#)