DRIFT-NCRN: A Benchmark Dataset for Drifter Trajectory Prediction

Khalil Virji
Department of Computer Science
McGill University
khalil.virji@mail.mcgill.ca

Johanna Hansen
Supervisor
Mobile Robotics Lab
McGill University
johanna.hansen@mail.mcgill.ca

Travis Manderson
Supervisor
Mobile Robotics Lab
McGill University
travis@cim.mcgill.ca

David Meger
Supervisor
Mobile Robotics Lab
McGill University
dmeger@cim.mcgill.ca

Gregory Dudek
Supervisor
Mobile Robotics Lab
McGill University
dudek@cim.mcgill.ca

***Abstract*—Influenced by complex interactions at the intersection of air and water, the fate of objects floating in the ocean is difficult to predict even a few days into the future. Despite the complexities, long-term ocean trajectory prediction has many important applications for search and rescue missions, ecological studies, and disaster remediation. Inspired by the DARPA Forecasting Floats in Turbulence challenge, we present an open-source benchmark dataset for measuring progress in ocean trajectory modeling. The dataset is based on a collation of ocean drifter trajectories and archival wind and current data. We hope it will enable further development of models tuned for the complexities of drifter trajectory prediction. In addition to the benchmark dataset, we also provide a baseline solution set built using OpenDrift, an open-source software package for modeling the trajectories of objects in the ocean or atmosphere.**

I. Introduction

Ocean drifters are non-actuated, freely-floating buoys that are deployed in the ocean. They are equipped with instruments to measure physical parameters of the ocean such as temperature, wind speed and direction, currents, and salinity. Since ocean drifters are non-actuated, they are relatively cheap to build and deploy. Consequently, they are widely deployed throughout the world to aid with climate research, oceanographic research, storm and weather forecasting, and oil spill monitoring.

However, since ocean drifters are non-actuated, once deployed, they drift freely with surface wind and ocean currents. If strategic deployment is not employed, a common phenomenon is that drifters end up getting clumped together in the same area, following the same trajectories. This is often an undesirable setting for modelling physical processes because measurements will come from the same regions in the ocean. If somehow we are able to predict where drifters will end up after a given amount of days, we can strategically deploy them to achieve maximum spatial coverage over a given area [1]. Extending further, if we are able to accurately predict drifter trajectories, this would open the door for predicting trajectories of other objects in the ocean. For example, trajectory predictions could be used to aid with search and rescue missions, and for forecasting the effects of disasters in the ocean such as oil spills [2].

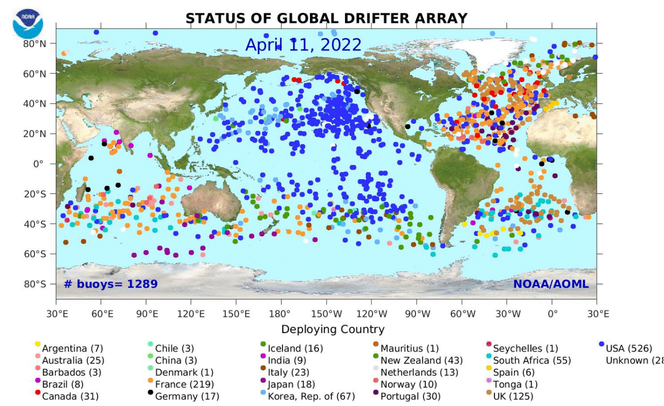

Fig. 1. Global Drifter Array recorded by the Atlantic Oceanographic and Meteorological Laboratory (AOML) Data Assembly Center

To enable further development of models turned for the complexities of drifter trajectory prediction, we built an open-source benchmark dataset, consisting of the trajectories of 90 drifters, and archival wind and current data. We also built a baseline solution set consisting of trajectory predictions for each of the 90 drifters. We hope this baseline solution set will be helpful for evaluating and measuring progress in drifter trajectory prediction.

II. Related Work

A. Forecasting Floats in Turbulence Challenge

The Forecasting Floats in Turbulence (FFT) challenge was a 30 day challenge conducted from November to December 2021 by the Defence Advanced Research Project Agency (DARPA) [3]. In this challenge, competitors were given 20 days of ocean drifter trajectories as a training data, and were tasked with forecasting the exact position of each drifter every

2 days for the remaining 10 days. Despite the straightforward goals of this challenge, the task of predicting ocean drifter trajectories over a 10 day period proved to be remarkably difficult. Fewer than half of the participants (14/31) predicted any of the 90 drifter positions on the last day to within 32km of the true position. The winner of the competition was Second Sign Predictions, a single-person team, Chris Wasson, who is an engineer in Southern California. When interviewed about his thoughts on the challenge, Wasson gave further insights into the difficulties of the challenge:

> *There are still many examples in the challenge float data where the float trajectory disagrees strongly with all available ocean current and wind data at that latitude/longitude. I think this is a strong indication that one of the biggest remaining hurdles in this problem is not in modeling surface effects but in identifying errors in these models, particularly ocean surface current models.*[1]

The DARPA FFT challenge took an important exploratory step into understanding the effects of wind, waves, and currents on objects floating in the ocean. It also sparked further investigation into drifter trajectory prediction and its challenges.

Fig. 2. Drifters used in the DARPA FFT challenge

B. OpenDrift

OpenDrift is an open-source software package for modelling the trajectories of objects floating in the ocean or atmosphere [4]. It is built in Python and consists of four main classes. The first is the `LagrangianArray` class which describes a particle in OpenDrift and its properties. The second is the `Model` class which corresponds to the physics of a trajectory model and takes care of updating particle properties such as position and velocity at each time step. Each model in OpenDrift must include an update function which takes environmental data as input, and will update particle properties accordingly. OpenDrift provides several built-in models such as `OceanDrift`, `OpenOil`, `Leeway`, and `OpenBerg`, however the framework has the flexibility to define custom models as well. The third class is the `Reader` class which is responsible for reading and processing environmental data from a given source, such as a NetCDF file. Readers are also responsible for retrieving environmental variables at a given position and time. Built-in OpenDrift readers use bilinear interpolation to approximate environmental variables between readings, however custom interpolation methods can be also be defined. The fourth and final class is the `Writer` class which takes care of saving the results of a simulation to a file.

Running simulations in OpenDrift can be as simple as initializing a model, defining readers for each environmental data file, seeding particles at some position and time, and then starting the simulation. The results can be visualized as plots, animations, or saved to a file for further processing. OpenDrift proved to be a useful and easy-to-use package for modelling and visualizing trajectories, and we took advantage of its capabilities when building our baseline solution set.

C. Navigating Stratospheric Balloons

Understanding the effects of physical parameters on stratospheric balloons was key to the study of Bellemare et al [5]. In this study, Bellemare et al. used reinforcement learning techniques to keep Loon balloons within 50 kilometers of a ground station for effective communication, a term referred to as *station-keeping*. Loon balloons are similar to ocean drifters in the sense that their trajectories are predominantly dictated by physical parameters, namely wind speed and direction. However, Loon balloons are able to control their vertical motion by pumping air ballast in and out of a fixed-volume envelope, while ocean drifters only have control over initial deployment location. The goal for the study of Bellemare et al. was to build an effective station-keeping flight controller, which would ascend and descend a Loon balloon to different altitudes to follow favourable wind patterns.

A notable challenge throughout this study was working with imperfect, incomplete, and low-resolution environmental data. This was also a challenge in the DARPA FFT competition as well as in our work compiling the dataset. However, despite the challenges, Bellemare et al. were still able to utilize wind forecasts and implement an effective flight controller that outperformed Loon's previous algorithm. We hope our work building this dataset will enable models to achieve similar success in utilizing wind and current forecasts for drifter trajectory prediction.

III. Dataset

The dataset we have built is comprised of ocean drifter trajectory data, observational wind and current data, and forecast wind and current data. The data has been collated from multiple sources such as the Global Forecast System (GFS), the Real-Time Ocean Forecast System (RTOFS), the Global Ocean Forecast System (GOFS), Wave-Watch 3 (WW3), and the Defence Advanced Research Project Agency (DARPA). Detailed descriptions of the dataset are provided below and summarized in Table I.

A. Ocean Drifter Trajectory Data

Ocean drifter trajectory data is taken from the DARPA FFT challenge archives. It contains the locations of 90 drifters, updated every hour, from November 2^{nd} to December 2^{nd} 2021. All the drifter trajectories over the 30 days are enclosed in the region [103-175]° E × [22.5-50]° N which we ensure is covered by all our environmental data. We provide the drifter trajectory data as a comma-separated values (.csv) file.

[1] https://www.darpa.mil/news-events/2021-12-13

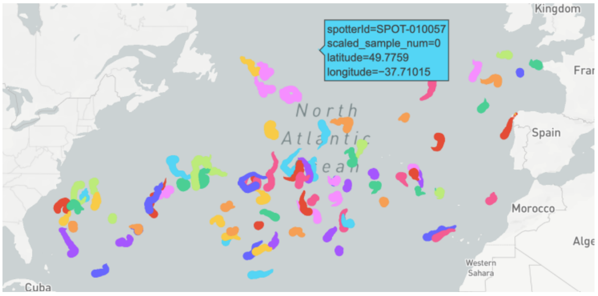

Fig. 3. DARPA drifter trajectories included in our dataset

B. Observational Data

Observational wind data is taken from the GFS Analysis dataset. It has a spatial resolution of 0.5° E × 0.5° N and a temporal resolution of 6 hours. Observational current data is taken from the GOFS 3.1 dataset. It has a spatial resolution of 0.08° E × 0.04° N and a temporal resolution of 3 hours. Both the wind and current data cover time periods from November 2nd to December 2nd 2021. We provide this observational data as NetCDF (.nc) files with global spatial coverage. Due to the large nature of these files, we also provide a processed version as NumPy data (.npy) files with spatial coverage [93-185]° E × [12.5-60]° N. We chose these bounds to cover the region containing all drifter trajectories, plus 10° of padding.

C. Forecast Data

Forecast current data is taken from the global RTOFS dataset. RTOFS is an 8-day forecast system with a spatial resolution of 0.08° E × 0.08° N and a temporal resolution of 1 hour for days 1 to 3, and 3 hours for days 3 to 8. The RTOFS data has global spatial coverage and covers the time period from November 21st to November 30th 2021. For forecast wind data, we selected data from two different sources. The first is from GFS, a 16-day forecast system with a spatial resolution of 0.5° E × 0.5° N and a temporal resolution of 3 hours. The GFS data spatially covers [0-360]° E × [5-90]° N from October 29th to December 6th 2021. The second is from WW3, a 7-day forecast system with a spatial resolution of 0.5° E × 0.5° N and a temporal resolution of 3 hours for days 1 to 3 and 6 hours for days 3 to 7. The WW3 data covers [0-360]° E × [5-60]° N from Nov 1st to November 28th 2021. In addition to forecast wind data, the WW3 data also includes wave forecasts with the same spatial and temporal resolutions. We provide all forecast data as NetCDF (.nc) files.

IV. Baseline Solutions

Our baseline solution set contains trajectory predictions for all 90 drifters, from Nov. 2nd to Dec. 2nd 2021. We generated this solution set using OpenDrift and observational wind and current data. For the trajectory model, we used the built-in model `OceanDrift`, which seemed appropriate for our application because it is a buoyant particle trajectory model. To generate the results, we added readers for each of our observational data files, seeded 100 particles in a small radius around each drifters initial location, and let `OceanDrift`'s update function take care of updating the particles positions at every time step.

We provide our baseline solution set as an array of NumPy data files, one for every drifter. Each file contains the predicted locations of a drifter for every hour between Nov. 2nd and Dec. 2nd 2021. Each file also contains the interpolated wind and current parameters at the predicted location, which will be used by the `OceanDrift` model to generate the next predicted location. In addition to the data files, we provide plots and animations to visualize the results. These plots contain the trajectories of the 100 simulated particles, and contrast them with the true trajectories of the drifters.

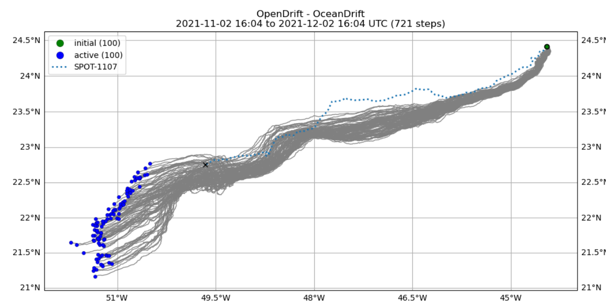

Fig. 4. Trajectory prediction for drifter 1107 included in our baseline solution set. Initial locations are shown in red, predicted trajectories are shown in grey, and predicted final locations are shown in blue. The true trajectory is shown as a dotted blue line, and the true final location is shown as a black *x*.

Although our solution set is relatively naive, in the sense that we use a preexisting model and simply pass in data to generate predictions, we hope it will provide a useful baseline for measuring progress in drifter trajectory prediction.

V. Challenges

During the process of compiling the dataset and generating baseline solutions, we encountered several challenges. The main challenge was building the dataset from data spread across multiple sources, each with different resolutions, projections, time ranges, and spatial coverage. The second major challenge was verifying that OpenDrift was interpreting our data files correctly and generating the correct trajectories.

A. Building the dataset

Once we had compiled drifter trajectory data from the DARPA FFT challenge archives, obtaining high-resolution environmental data to match the time range and region of the challenge proved to be a difficult task. Not only was the data spread across multiple sources, but each source had different spatial and temporal resolutions, different projections, and different data formats. For example, we found high resolution observational current data from GOFS, with a spatial resolution of 0.08° E × 0.04° N and a temporal resolution of 3 hours. However, for observational wind data, the highest resolution data we could find was through GFS, with a spatial resolution of 0.5° E × 0.5° N and a temporal resolution of 6 hours.

TABLE I
DATASET CONTENTS

Source	Type	Parameters	Spatial Resolution	Temporal Resolution	Spatial Coverage	Temporal Coverage
GFS Analysis	observational	wind	0.5° E × 0.5° N	6 hours	global	Nov. 2^{nd} - Dec 2^{nd} 2021
GOFS 3.1	observational	current	0.08° E × 0.04° N	3 hours	global	Nov. 2^{nd} - Dec 2^{nd} 2021
RTOFS	8-day forecast	current	0.08° E × 0.08° N	1 hour (days 1-3), 3 hours (days 3-8)	global	Nov. 21^{st} - Nov 30^{th} 2021
GFS forecast	16-day forecast	wind	0.5° E × 0.5° N	3 hours	(0-360)° E × (5-90)° N	Oct. 29^{th} - Dec 6^{th} 2021
WW3	7-day forecast	wind, waves	0.5° E × 0.5° N	3 hours (days 1-3), 6 hours (days 3-7)	(0-360)° E × (5-60)° N	Nov. 1^{st} - Nov 28^{th} 2021
DARPA	observational	drifter trajectories	N/A	1 hour	(103-175)° E × (22.5-50)° N	Nov. 2^{nd} - Dec 2^{nd} 2021

For forecast data, in addition to the data sources having different resolutions, we found that forecast data is not retained for long periods after its publishing date. This made obtaining archival forecasts a difficult, and in some cases, impossible task. Fortunately, at the time of the DARPA FFT competition, we had downloaded and saved some forecasts, and have included those in the dataset. However for days we didn't record, we could not find any source that had retained current and wind forecasts, which is why there are some gaps in the dataset. Specifically, we are missing current forecasts from Nov. 2^{nd} to Nov 20^{th} and from Nov. 31^{st} to Dec. 2^{nd}. We are also missing WW3 forecasts from Nov 29^{th} to Dec. 2^{nd}.

Another challenge we encountered while building the dataset was dealing with different formats of data across different sources. For example, GFS stores observational wind data as single time files, which means there is one output file for every 6 hours, resulting in 120 files to cover the dates of the competition. This is in contrast to GOFS, which stores observational current data as multi-time files, and just one file is needed to cover the dates of the competition. For forecast data, RTOFS provides current forecasts as unprojected data, meaning that the data has not been projected down to a 2D grid. This is in contrast to all other sources which provide projected data. To deal with these formatting differences and to make the dataset as consistent as possible, we employed tools such as the Climate Data Operators (CDO) [6] which are a collection of operators used to manipulate and analyze environmental data.

Despite the challenges, we were still able to compile the data into a relatively consistent and usable format. We used the dataset for building baseline solutions, and believe it can be used as is to help model drifter trajectories. We are also actively looking for ways to further process and improve the dataset, and we discuss some of those ideas below.

B. Verifying Baseline Solutions

After generating our baseline solution set using OpenDrift, the results were certainly not ideal. The predicted trajectories were significantly different than the true trajectories, and often in the opposite direction. At first, we thought that there was a problem with how OpenDrift was interpreting our data. To test this we ran verification steps where we got wind and current values at a given position and time from the dataset, and verified that they were indeed the values used during the OpenDrift simulation. To our surprise they were infact the same for all of our test points. Furthermore, we tried running shorter, 10-day predictions using OpenDrift, and the results were much closer to the true trajectories. We came to the conclusion that the quality of our solution set is not due to issues within OpenDrift, but rather due to the difficulties of long-term drifter trajectory prediction.

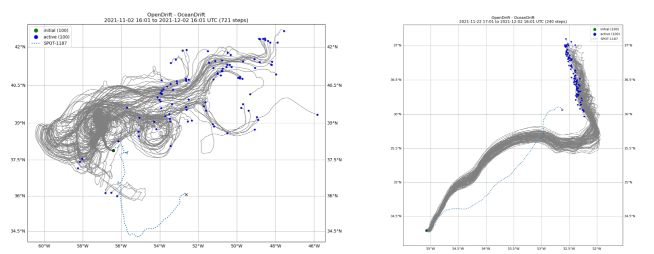

Fig. 5. 30 day prediction (left) versus 10 day prediction (right) for drifter 1187

VI. FUTURE WORK

A clear next step after building this dataset is to implement machine learning techniques to train a model for drifter trajectory prediction, and to improve upon our baseline solution set. In addition to this, we would also like to categorize drifter trajectories in our dataset, build a framework to quantify and evaluate results, and build a secondary dataset for Argo drifters.

A. Categorizing Drifter Trajectories

A common practice when compiling a dataset is categorizing the data [7]. Categorizing data not only makes the dataset more readable, but can also help with developing and evaluating models. By categorizing drifter trajectories in our dataset, we would have a concrete measure to evaluate the performance of a trajectory model, and would know which drifter categories it models well, and which categories it needs improvement.

When building our baseline solution set, we noticed that the `OceanDrift` model is able to model straight drifter trajectories more accurately than drifter trajectories with large curves or loops. Thus we would like to categorize and rank drifter trajectories based on the *straightness* of the trajectory. We could employ techniques such as identifying loops in discrete drifter trajectories [8], and measuring deviance from a straight trajectory to accomplish this. In addition, we could categorize drifter trajectories by distance travelled, distance to land, and by region.

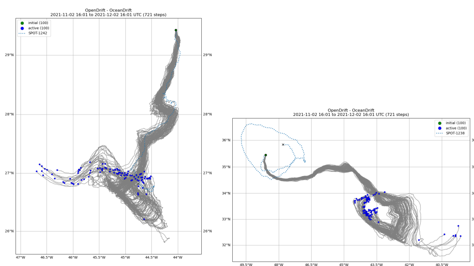

Fig. 6. A relatively straight trajectory prediction (left) versus a curved and looped trajectory prediction (right)

B. A Framework to Quantify and Evaluate Predictions

When building our baseline solution set, we did not have a framework in place to quantitatively evaluate and score a trajectory prediction. We had evaluated predictions by manually analyzing plots and animations, which going forward, is not a reliable evaluation scheme. To fix this, we would like to build a framework which would use quantitative measures such as position error to score trajectory predictions. This would allow us to evaluate and compare trajectory models in a less subjective way, and would allow us to get quantitative results for our baseline solution set.

C. Building an Argo Drifter Dataset

Argo is an international program, similar to DARPA, that deploys ocean drifters to measure physical parameters of the ocean. However, Argo floats only report their location once every 10 days, in contrast to the drifters in the DARPA FFT challenge which reported their location every hour. Taking in account what we learned while building this dataset, we would like to build a secondary dataset for Argo drifters. The dataset would include the trajectories of Argo drifters over a longer period of time, accompanied by relevant wind and current data. Since Argo drifters report their location much less frequently, this dataset could pose new challenges for drifter trajectory prediction and could help tune models to deal with sparse data.

VII. Conclusion

In this paper, we presented a benchmark dataset consisting of the trajectories of 90 drifters, accompanied with relevant

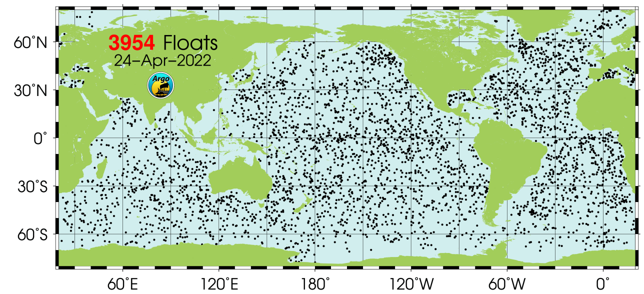

Fig. 7. Global Argo drifter array on April 24th 2022

wind and current data. We also presented a baseline solution set consisting of trajectory predictions for each of the 90 drifters, built using OpenDrift [4]. We hope that this dataset will enable further development of models tuned for the complexities of drifter trajectory prediction, and we hope the baseline solution set will provide a concrete measure to evaluate the progress and performance of a model.

VIII. Acknowledgements

A special thank you to Johanna Hansen, Travis Manderson, David Meger and Gregory Dudek for the opportunity to conduct this research and for their help and supervision throughout.

References

[1] J. Hansen and G. Dudek, "Coverage optimization with non-actuated, floating mobile sensors using iterative trajectory planning in marine flow fields," 2018. [Online]. Available: https://www.cim.mcgill.ca/∼mrl/pubs/jhansen/IROS2018.pdf

[2] D. Simecek-Beatty, "Chapter 11 - oil spill trajectory forecasting uncertainty and emergency response," in *Oil Spill Science and Technology*, M. Fingas, Ed. Boston: Gulf Professional Publishing, 2011, pp. 275–299. [Online]. Available: https://www.sciencedirect.com/science/article/pii/B9781856179430100115

[3] DARPA. (2021) Darpa forecasting floats in turbulence challenge. [Online]. Available: https://oceanofthings.darpa.mil

[4] K.-F. Dagestad, J. Röhrs, Ø. Breivik, and B. Ådlandsvik, "Opendrift v1.0: a generic framework for trajectory modelling," *Geoscientific Model Development*, vol. 11, no. 4, pp. 1405–1420, 2018. [Online]. Available: https://gmd.copernicus.org/articles/11/1405/2018/

[5] M. G. Bellemare, S. Candido, P. S. Castro, J. Gong, M. C. Machado, S. Moitra, S. S. Ponda, and Z. Wang, "Autonomous navigation of stratospheric balloons using reinforcement learning," *Nature*, vol. 588, p. 77–82, 2020. [Online]. Available: https://doi.org/10.1038/s41586-020-2939-8

[6] *Climate Data Operators*. Max Planck Institute for Meteorology, 2022. [Online]. Available: https://code.mpimet.mpg.de/projects/cdo

[7] T.-Y. Lin, M. Maire, S. Belongie, L. Bourdev, R. Girshick, J. Hays, P. Perona, D. Ramanan, C. L. Zitnick, and P. Dollár, "Microsoft coco: Common objects in context," 2014. [Online]. Available: https://arxiv.org/abs/1405.0312

[8] C. Dong, Y. Liu, R. Lumpkin, M. Lankhorst, D. Chen, J. McWilliams, and Y. Guan, "A scheme to identify loops from trajectories of oceanic surface drifters: An application in the kuroshio extension region," *Journal of Atmospheric and Oceanic Technology*, vol. 28, 09 2011.