university of groningen / faculty of science and engineering / bernoulli institute

Spatiotemporal Distribution Modeling of Cod in Norway Through a Diffusion Model

Bachelor's Project Thesis

Ian Desley Blok, S5227596, ian.blok@gmail.com
Supervisors: J.D. Cardenas Cartagena

Abstract: Cod is a vital species on earth and for Norway. Following the collapse of the Norwegian Spawning Herring, Norway has introduced measures to promote sustainable fishing. Yet these measures do not suffice. Just four out of the fifteen cod populations have a healthy population size. Furthermore, the fishing industry has a significant impact on the climate, having produced 179 million tonnes CO_2 worldwide in 2011. Modeling the distribution of cod could reduce these issues, allowing for more efficient and sustainable fishing. Through a diffusion model, this study tries to forecast catch locations of the next day. The model is compared against a baseline (based on a previous paper) that predicts catch data for a given day by summing the catch data from a 15-day window (7 days before, the current day, and 7 days after) from the previous three years. The study shows that the model outperforms the baseline, generating predictions that are closer to the actual catch locations. The model's predictions outperformed the baseline, with shorter mean distances to true catch locations and lower MSE and MAE scores. The findings reveal that catch and environmental data alone cannot reliably predict cod distribution patterns, especially due to the limited temporal resolution and sparse datasets. However, the study also shows the learning capacity of diffusion models despite high sparsity in data, suggesting that given high quality data, the model is able to make more accurate predictions.

1 Introduction

Cod (Gadus Morhua) is vital for life on earth. 164.6 million tonnes of aquatic animals were consumed in the year 2022, of which cod played an important role. Norway itself produced 3.1 % of the total world catch, being the second largest exporter of marine life next to China (Food and Agriculture Organization of the United Nations, 2024). For Norway itself, the seafood industry is of high importance as well. In 2017, this industry is estimated to have a contribution of 94 billion NOK ($\sim$ 7.96 billion euros) to the gross national product (GNP). Since 1970, the average increase in contribution to GNP per year for the fishing industry is 19%. In comparison, the average increase for Norway is 8%, and 3% for the manufacturing industry in the same period (Johansen et al., 2019).

Following the collapse of the Norwegian Spring Spawning Herring, Norway has introduced measures to promote sustainable fishing (Gullestad et al., 2013). Yet, implemented measures like total allowable catches have not improved fishing populations enough (Anderson, 2012). Cod and other fish species have still not recovered fully and should be fished on less (Kualtska et al., 2024).

Carbon emissions produced by these fishing vessels also have a significant impact. Marine fisheries produced 179 million tonnes CO_2 in 2011, of which 20 million was produced by the EU (Parker et al., 2018). Furthermore, spatial restrictions such as the exclusion of other EU members from the United Kingdom's exclusive economic zone, doubled fuel use intensity and undid $\sim$ 15 years of improved fuel efficiency in Norwegian pelagic fisheries (Scherrer et al., 2024).

This highlights the need for more monitoring and forecasting of the distribution and migration of marine life. Not only would such a technology allow fisheries to fish in a more productive manner, it would also allow for more sustainable fishing measures. One possible measure that could be implemented are marine protected areas. Marine protected areas are a measurement part of ecosystem based management, an approach for treating ocean life responsibly (Halpern et al., 2010). This is what motivated Meeanan et al. (2023), who produced two separate studies on this matter.

Meeanan et al. (2023) showcase the possibility of a better prediction system. Using a combination of environmental, spatial, temporal and fishery predictors, they achieved an accuracy of 86.14% $\pm$ 0.23% in terms of predicted fishing spots for the Andaman Sea near Thailand. Meeanan et al. (2024) extends this results by also including the Gulf of

Thailand, achieving an accuracy of 83.49 ± 2.16%. Previous results show that it is possible to predict the distribution of marine life.

This study is based on the FishAI: Sustainable Commercial Fishing Challenge (Nordmo et al., 2022). It is a challenge published in 2022 by the Nordic Machine Intelligence group. It aims to explore the means of AI for more sustainable fishing. The challenge provides historical catch data and environmental datasets which can be used to predict fishing locations. Solutions to this challenge have attempted to use classical machine learning approaches (Syms, 2022; Dammen et al., 2022; Linkiö et al., 2022), but results show that the available data is insufficient for such an approach. Furthermore, visual approaches to this problem have not led to converging models due to the sparsity present in the data (Róason, 2024; Cronin, 2024). Other models do exist that can predict the distribution of species. But these models either require absence data, high quality presence data or rely on certain assumptions to be made (see Section 2.1). Applying a visual based deep learning approach may be an alternative route to model the distribution of species.

Diffusion models have been used for temporal-geospatial forecasting in previous studies, with promising results (Yang et al., 2024; Liu et al., 2024). Given enough data and model capacity, diffusion models can approximate any continuous data distribution (Song et al., 2021). Diffusion models are therefore a viable option for this distribution task, as they may not be limited by the sparse nature of the datasets unlike more traditional visual deep learning models.

This study explores the usage of diffusion models for temporal-geospatial forecasting for predicting cod in the exclusive economic zone of Norway. It examines the performance of diffusion models to forecast the spatiotemporal distribution of fish with environmental data and catch notes. The study presents a conditioned diffusion model with an extra layer of sparsity bits. This model is compared to a custom baseline based on previous work (Linkiö et al., 2022). Using accuracy, the distance between each prediction and its corresponding true catch location and mean absolute error & mean squared error metrics, this study demonstrates that the model outperforms the baseline. The study shows that the model can learn spatiotemporal patterns in fish migration despite a highly sparse dataset.

2 Literature Review

2.1 AI in the Fishing Industry

Many models are used worldwide within the fishing industry. These types of models typically represent the biomass or abundance of a species and include probability models (Libralato, 2025). One example is the CMSY++ method created by Froese et al., which is a data-limited stock assessment method for fishes and invertebrates. It is a Bayesian implementation of a modified Schaefer model, using an Artificial Neural Network for assistance. The model can predict biomass, maximum sustainable yield and stock status and exploitation rates. The model can predict this based on catch-only data (if this data is accurate enough) or based on absence and presence data of a fish species.

Species Distribution Models (SDM) have existed for a long time, being able to predict all sorts of species (Dhyani et al., 2023). For example, distribution of species can be modelled through a maximum-entropy method. This is a popular machine learning method able to create a distribution map of any species, as long as presence data is given (Dhyani et al., 2023; Phillips et al., 2006). These SDMs have been used in the context of the fishing industry. SDMs have, for instance, been used to predict species found in estuarine fisheries (ting Zhang et al., 2022). Furthermore, by combining multiple SDMs into an ensemble, one can predict distributions of multiple fish species. These models can be used to spatially manage fishing areas which can encourage sustainable fishing (Panzeri et al., 2024). These models make use of assumptions on the environment (Dhyani et al., 2023).

2.2 Diffusion Models for Geospatial and Timeseries Forecasting

Diffusion models (Ho et al., 2020) can not only be applied to generating images, but can also be applied for data purification, molecular graph modeling and anomaly detection (Yang et al., 2023). Another application of diffusion models is forecasting. Research has shown that diffusion models can be applied to this type of problem, with success (Li et al., 2022; Cachay et al., 2023). One example is by Liu et al. (2024), which have used a series of diffusion models to predict the wind power generated for the next three days. All diffusion models are able to "successfully capture the short-term trend of wind power".

Diffusion models generate data by denoising a

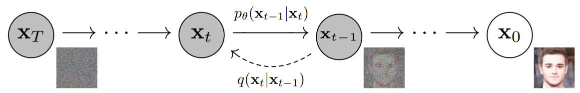

Figure 2.1: Figure showing the processes for a diffusion model. Here $q(x_t|x_{t-1})$ represents the Gaussian noise added and $p_\theta(x_{t-1}|x_t)$ represents the reverse process learned by a diffusion model. The process shows how the model is able to generate an image from a pure noise image X_T. Figure from Ho et al. (2020).

data sample. The data sample is first noisified by iteratively adding Gaussian noise to the sample. This results in a noisified data sample. The diffusion model attempts to denoisify this image by iteratively predicting the noise added, and subtracting that from the data sample. This process is illustrated in figure 2.1 with an image, but a similar process can be applied to for example a time series.

There are many different custom diffusion models (Yang et al., 2024). The diffusion model used in this paper is based on Rasul et al. (2021). In this paper, the authors propose a new model called TimeGrad. This is a conditioned diffusion model able to forecast timeseries. The model is conditioned on the hidden states from either an LSTM or GRU, allowing it to learn the temporal relationships in the data.
Furthermore, by using a U-Net with convolutional layers (Ronneberger et al., 2015), the diffusion model can be trained on images, allowing it to learn the geospatial relationships in the data. This idea has also been used by for example Soni et al. (2025), which used a U-Net to extract road networks from satellite images, a type of geospatial data.

3 Methods

3.1 Datasets

Four different datasets were used for this study. Catch data, salinity and sea surface temperature were suggested by the FishAI: Sustainable Commercial Fishing Challenge Nordmo et al. (2022). Bathymetry data was added as well.
Catch data was sourced from the Norwegian government (Fiskeridirektoratet). This dataset contains about 130 fields of information related to fish caught from different vessels. It includes information about the vessels themselves, personnel involved and information about the fish caught (including where and how much). The data is in tabular form and ranges from 2000 up to and including 2024.
Salinity data consisted of daily salinity values created by applying an 8 day running mean. The data has a granularity of 0.25 degrees longitude and latitude (Remote Sensing Systems (RSS), 2024). Spatiotemporal analysis of cod has shown that salinity is a major influence in the presence of cod (Nian et al., 2021).
Sea surface temperature (SST) data consists of daily average temperature readings from a satellite. The data has a granularity of 0.25 degrees longitude and latitude (Huang et al., 2021). SST is also known to be one of the major influences on cod's behavior Nian et al. (2021).
Bathymetry data was taken from the GEBCO project. More specifically, sub ice bathymetry was taken. The dataset has a granularity of 0.004 degrees longitude and latitude (GEBCO Compilation Group, 2024). Evidence suggests that bathymetry influences the decisions of fishing vessels (Alizadeh Ashrafi et al., 2020), as well as the migration of fish (Hobson et al., 2009; Meeanan et al., 2023).

3.2 Data Processing

Pre-processing. Multiple pre-processing steps were applied before using the data. 1) Incorrect dates in the catch data were removed. 2) Missing specific longitude and latitude coordinates in the catch data were substituted with the given 'Hovedomrade' longitude and latitude coordinates. This is a second coordinate specified with each catch, and was more consistently provided. 3) Any catches with land-coordinates were removed. Natural Earth's 10 meter land polygons were used to do this (Patterson and Kelso, 2024). 4) Remove any coordinates not present in Norway's Exclusive Economic Zone (EEZ). The EEZ dataset from Flanders Marine Institute was used to do this (Flanders Marine Institute, 2024).

Four different maps were created, each map corresponding to a different dataset. All maps had the same longitude and latitude size, derived from the maximum and minimum coordinates in the catch data. An example sample can be seen in Figure 3.1.

Maps representing catch data were created by summing catch weight for a given day. A Gaussian filter is applied to each map (i.e. for each day)($\sigma = 3.5$).
The salinity data consisted of missing values on specific ocean coordinates. These were interpolated by replacing each missing value with the mean of all available observations within 3 degrees longitude or latitude. The same interpolation was applied to the sea surface temperature dataset.
The bathymetry map did not contain any missing coordinate values and was therefore not altered.

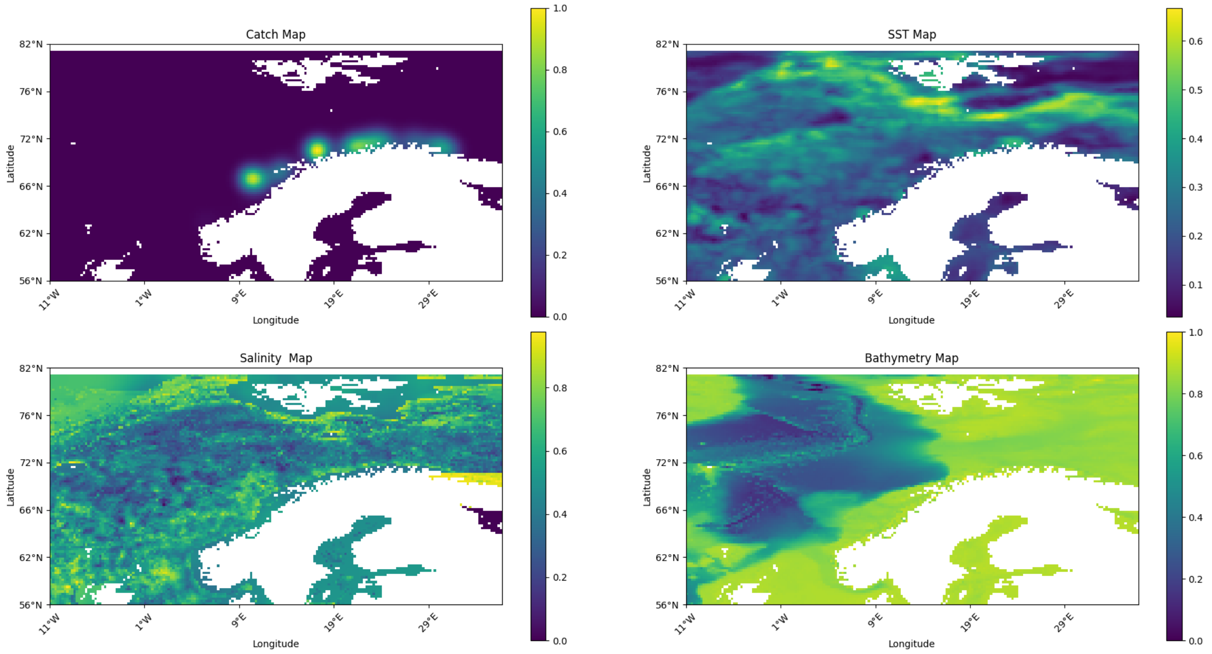

Figure 3.1: An example sample used to train the model on. From left to right: the catch data, sea surface temperature, salinity and bathymetry data. Note the high sparsity present in the catch data.

After these operations, days where no fish was caught (a catch map of only zeros) and salinity maps of only NaNs were removed from the dataset. Any remaining NaN values were then converted to zero. A mask was created and converted any land values to zero (Patterson and Kelso, 2024). Slices of 8 days were created from the catch data: the current day, 6 days behind and the next day. The current day of salinity, SST, and bathymetry were used as environmental data. The next day of catch data is what the model was trained on and what the model predicted. After these slices were created the data was split chronologically into a training, validation and testing dataset. Catch data and bathymetry data were normalized locally while salinity and SST were normalized globally Mazziotta and Pareto (2022).

Post-processing. Predictions of the model and baseline consisted of a map similar to the catch map. In order to derive the exact catch coordinates from the predictions, multiple Gaussian Mixture Models were fitted with a number of means from 1 up to and including 15. The model with the lowest Bayes Information Criterion score was selected out of these 15. This process was applied to the model, baseline and ground truth. The process was also applied to the ground truth as a lot of catch locations did not result in a significant Gaussian blur. Only a few spots had such a high value that you could visually see a Gaussian blur, which is the most important information when it comes to predicted distribution.

3.3 Baseline

A baseline was created to compare the model with. This baseline was created similarly to the winning paper of the FishAI: Sustainable Commercial Fishing Challenge Linkiö et al. (2022). The baseline creates a prediction for a given day by summing the past 7 days, current day and next 7 days of the previous 3 years. A Gaussian filter ($\sigma = 3.5$) is then applied. This is done for all the days also predicted by the model.

3.4 Diffusion Model

Main Architecture. The diffusion model implemented is based on the ADM network (Dhariwal and Nichol, 2021), which is also implemented in (Karras et al., 2024). The architecture is represented in Figure 3.2. The Figure shows the complete architecture used for this study. The model is fed a padded noisy image, which gets passed through multiple encoder and decoder blocks, some of which include an attention mechanism. For each encoder or decoder block, the embedding is injected into the image by multiplying and then summing the embedding element-wise to the image. The resulting values from the model are added to the original noisy image, creating a denoised image. Note that after this process the image may still contain a lot of noise, as the model needs to denoise the same

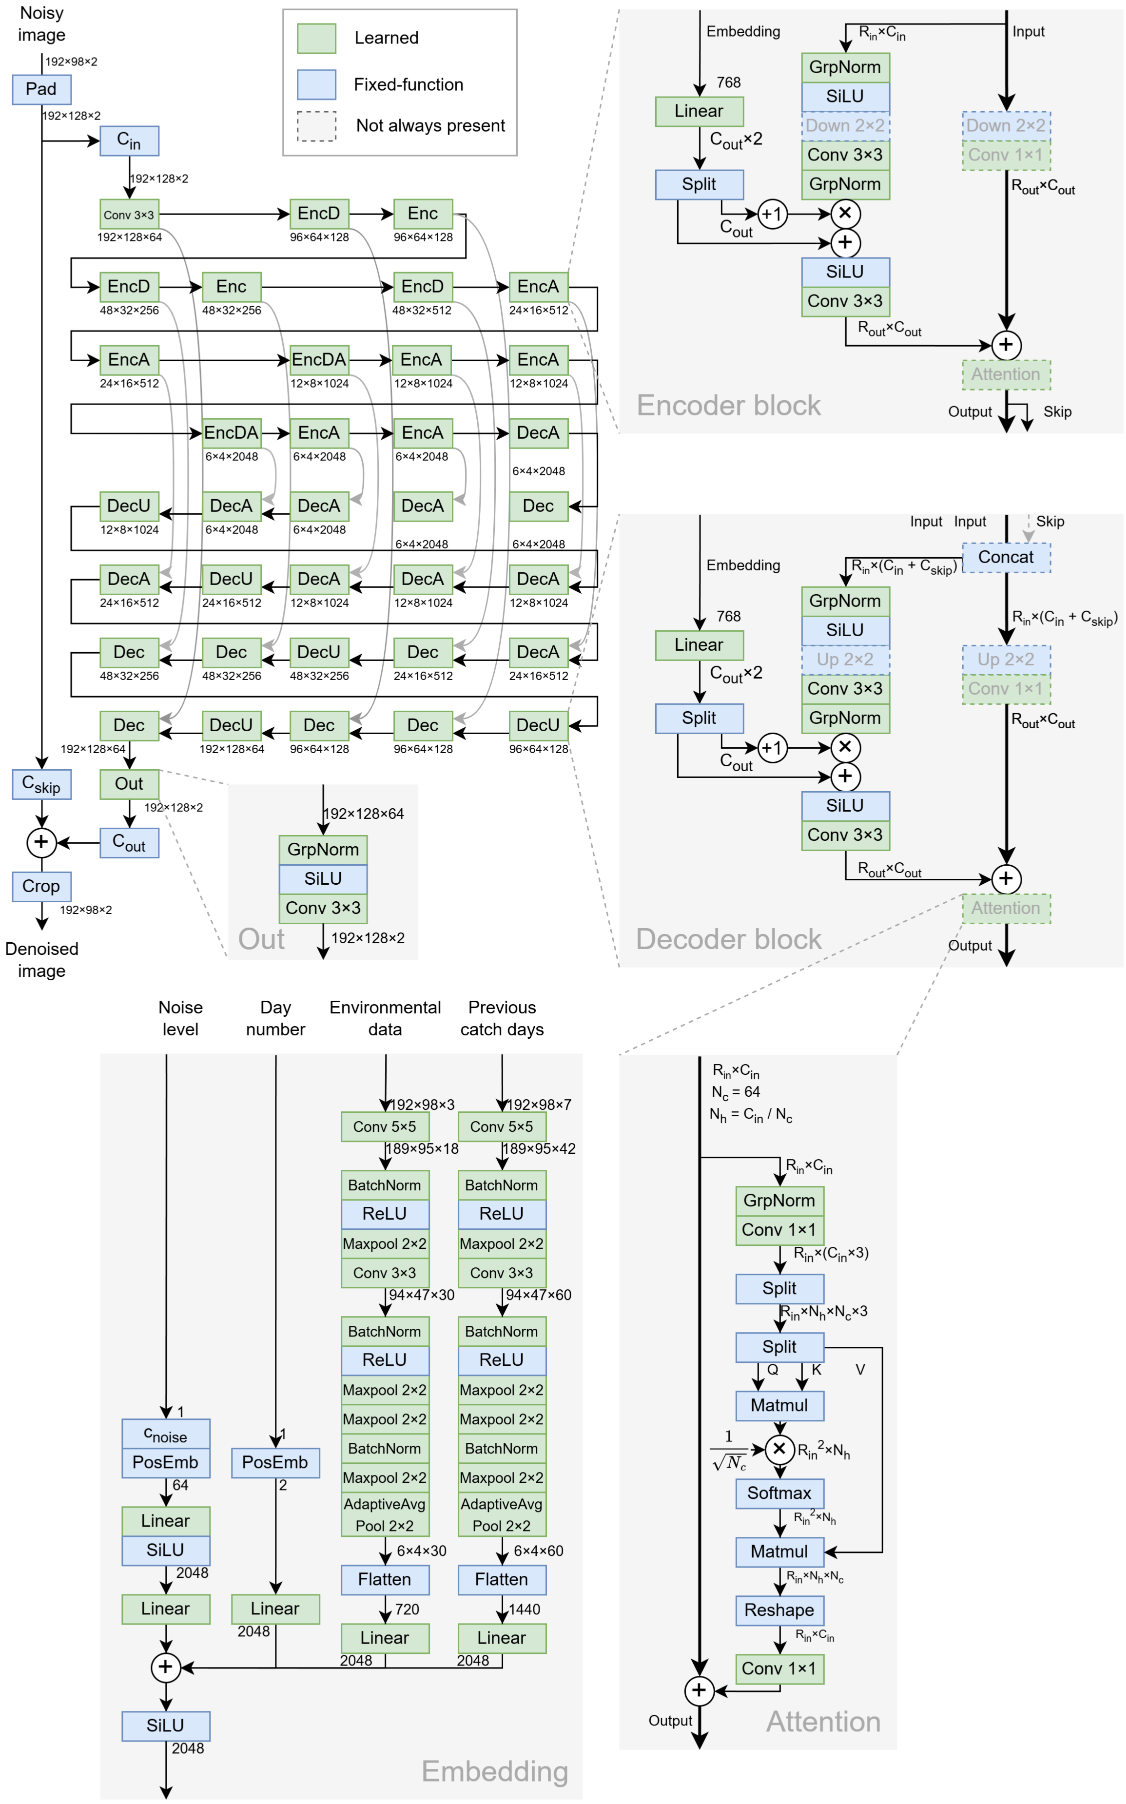

Figure 3.2: The architecture of the model used. It is a conditional diffusion model. A letter 'D' means that a down block was used, 'U' that an up block was used and 'A' that an attention block was used. Image style derived from Karras et al. (2024).

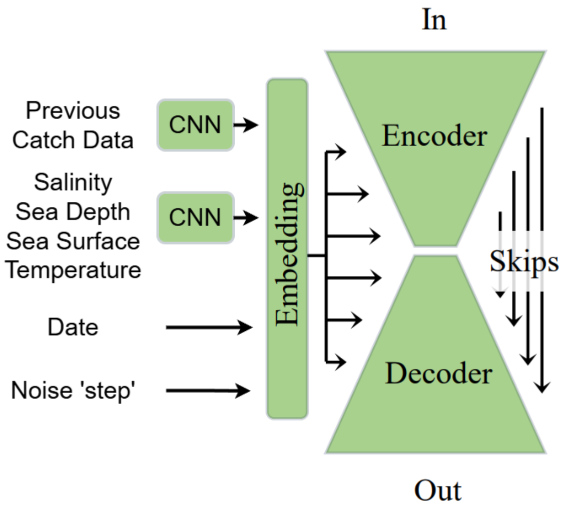

Figure 3.3: A simplified representation of the architecture used in this study. The Figure highlights the encoder/decoder structure as well as the injection of the embedding. Figure style derived from Karras et al. (2024).

image many times for it to become coherent. The architecture used represents an U-Net structure (Ronneberger et al., 2015). Each encoder block passes the encoded image to a corresponding decoder block "which allows the network to propagate context information to higher resolution layers" (Ronneberger et al., 2015). A schematic representation of this can be seen in Figure 3.3.

Conditional Information. The diffusion model implemented was conditioned on multiple pieces of information. Apart from the standard Gaussian noise schedule, environmental data, previous catch data and date information was added. After embedding all of the information in a 1 dimensional tensor of length 2048, the embedding vector gets directly added in the encoder and decoder blocks. This allows the model to denoise an image based on conditional information (Rasul et al., 2021). The diffusion model implicitly learns the geospatial relationships present in the data, as it needs to denoise images and models the underlying distribution.
Environmental data and previous catch data were encoded through a convolutional neural network (CNN), the structure is loosely based on the LeNet-5 architecture (Lecun et al., 1998).
Time information is encoded through a sin/cos embedding layer after which a linear layers scales it up to the highest channel size of 2048.

Sparsity Bits. The final modification applied is the addition of sparsity bits, a modification proposed by (Ostheimer et al., 2025). Sparsity

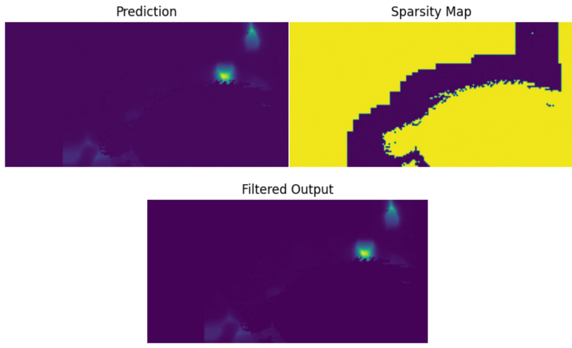

Figure 3.4: Three visualizations. The "Prediction" layer is the standard layer outputted by the model. The "Sparsity Map" is the second channel, indicating which pixels need to be zero and which should not be. "Filtered Output" is the final sample and used for evaluation.

bits are an extra layer representing which pixels in an image are zero and which have a value higher than zero, this is represented through a boolean channel. This additional information allows the model to learn the sparsity patterns in the original image. In practice, it means the model has an input of two channels, one for the map to denoise and one for the sparsity layer to denoise. When sampling, the model outputs two channels, one representing the catch map and one the sparsity channel. Using a threshold function, the sparsity layer is transformed into a boolean layer, which is used to keep or zero the pixels in the catch map. An example of this process is shown in figure 3.4, which shows the prediction and sparsity map (as a boolean layer), followed by the final output.

3.5 Metrics

Multiple metrics were implemented to accurately understand the performance of both the model and baseline. Both the baseline and model were compared with the ground truth datasets from the test set. The metrics not only focus on the performance image-wise, but they also show performance with respect to predictions made on the best fishing locations.

Mean Squared Error. The mean squared error (MSE) is calculated through the following formula:

$$MSE = \frac{1}{N}\sum_{i=1}^{n}(P_{ij} - \hat{P}_{ij})^2 \tag{3.1}$$

where P stands for the ground truth image and P_{ij} represents the pixel at row i, column j. And where $\hat{P}$ stands for the sampled image. The MSE compares the images generated by the model with the ground truth. It shows the overall difference between all pixels, with a stronger focus on larger errors (due to the squaring).

Mean Absolute Error. The mean absolute error (MAE) is calculated through the following formula:

$$MAE = \frac{1}{N} \sum_{i=1}^{n} |P_{ij} - \hat{P}_{ij}| \tag{3.2}$$

The MAE makes the same type of comparison with the MSE, but it shows the average difference. This gives a good overview in general on the performance of the model/baseline.

Haversine Distance. A distance metric was used to compare the fishing spot predicted by the model, with the nearest actual fishing spot. This metric highlights the actual fishing locations and can give a better insight in how the model/baseline predicts. The longitude and latitude coordinate values were used to calculate the distance between prediction and closest actual fishing locations. The haversine distance is used, and is calculated as follows:

$$d = 2R \arcsin \left(\sqrt{\sin^2\left(\frac{\Delta\phi}{2}\right) + \cos\phi_1 \cos\phi_2 \sin^2\left(\frac{\Delta\lambda}{2}\right)} \right) \tag{3.3}$$

Where d is the distance between two points (in kilometers), R is the radius of Earth (6371 km) ϕ_1, ϕ_2 are the latitude of points 1 and 2 (in radians) and λ_1, λ_2 are the longitude of points 1 and 2 (in radians).

Accuracy and Number of Predictions. A prediction is counted as correct if the distance between the prediction and closest catch location is less than or equal to 50 kilometers. This binary classification also allows the computation of accuracy. The amount of predictions and correct and incorrect predictions give more insight into the behavior of the model/baseline. Accuracy allows for an easier comparison with respect to performance.

Inference time. This refers to the time it takes for the model to sample one image given a catch history, environmental data and day of the year value.

Training time. This refers to how long it takes for the model to train for one epoch.

CO_2 emissions and power consumption. How much carbon dioxide was emitted and electricity was consumed for the whole training of the model. And how much environmental impact generating the test samples has. Both of these metrics are generated through CodeCarbon (Courty et al., 2024)

4 Results

Training results

A diffusion model was trained for 500 epochs. Training occurred on a single NVIDIA Tesla v100 GPU along with 1 CPU core. Training took about 20.1 hours in total which includes sampling from the model every 50 epochs (a total of 10 sampled images). Training time per epoch had a mean of 2.88 seconds. Inference time took 152.73 seconds. Training the model cost 7.62 kWh of energy and 2.03 Kg of carbon dioxide. Generating the 490 test samples cost 0.74 kWh and 0.20 Kg of carbon dioxide. This means that to produce 1 test sample, it cost 1.51 Wh and 0.41 grams of carbon dioxide. This means that the model does not consume a lot of electricity and carbon dioxide. Moreover, you only need to produce 1 sample per day, this would result in about 150 grams of CO_2 emitted within a year. This is negligible compared to the millions of tonnes of $CO2$ emitted each year from the fishing industry in the EU Parker et al. (2018).

Samples generated by the diffusion model show that the model has successfully learned the underlying distribution. It can generate many samples without much or any noise. An example image can be seen in Figure 3.4. The quality of the sampled image is of similar quality than the dataset itself. The learned sparsity map also shows the learned spatial layout of the diffusion model. This spatial structure is present in all sampled images.

Metrics can be seen in Table 4.1 and Figures 4.1a and 4.1b. The diffusion model scores equally on accuracy with an accuracy of 0.17 ± 0.25 while the baseline has an accuracy of 0.17 ± 0.19. MSE and MAE scores also indicate that the model produces results that are closer to the ground truth than the model, a lower MSE or MAE score is better. Distance metrics indicate that the model scores significantly better than the baseline. The mean distance is much lower. Minimum and maximum distance metrics were also calculated. These are the minimum distance and maximum distance of all distances for a given day. The baseline has a much higher maximum distance and much lower minimum distance. This behavior is consistent with the amount of predictions. The

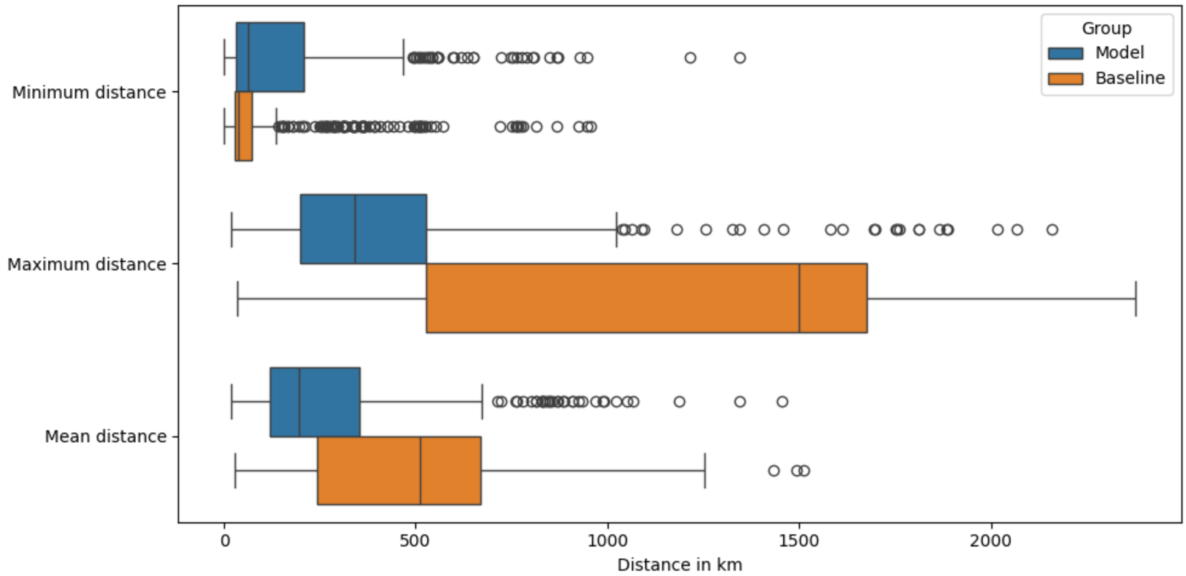

(a) Bar plot showing different distance metrics. Per predicted day, multiple predictions and therefore multiple distances were calculated. The minimum, mean and maximum was taken of these distances for each day and plotted as a bar plot. On average, the model's locations are much closer to the ground truth locations compared to the baseline. The maximum distance between a prediction of the baseline and ground truth is much larger than that of the model as well, indicating higher performance from the model.

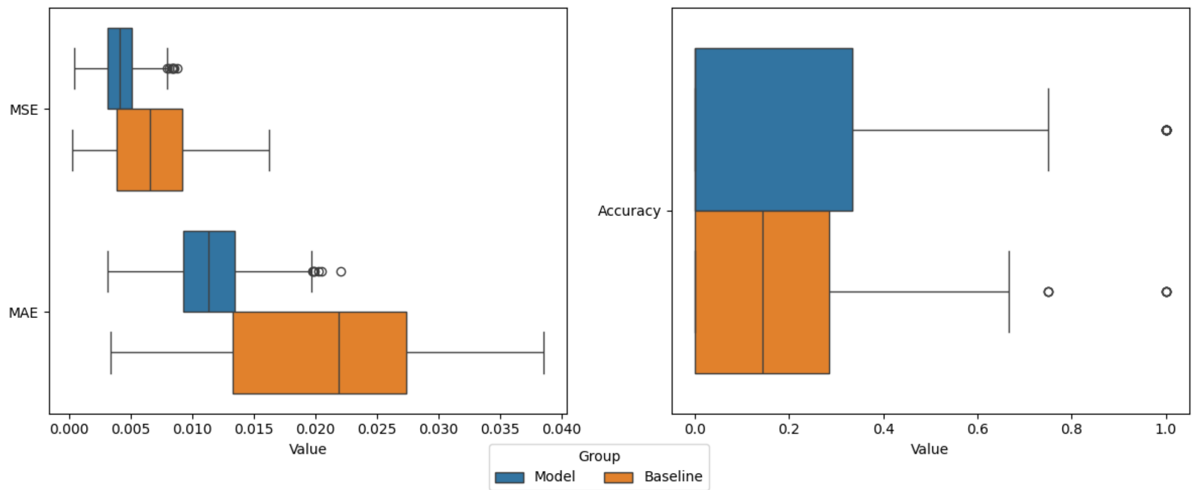

(b) Left: Different metrics showing how the model and baseline performed. The model produces much lower MAE and MSE scores compared to the baseline. Right: A bar plot with a strip plot overlayed showing the accuracy and correct predictions by the model and baseline. The accuracy and amount of correct predictions are very similar.

Figure 4.1: Metrics showing the performance of the model and baseline.

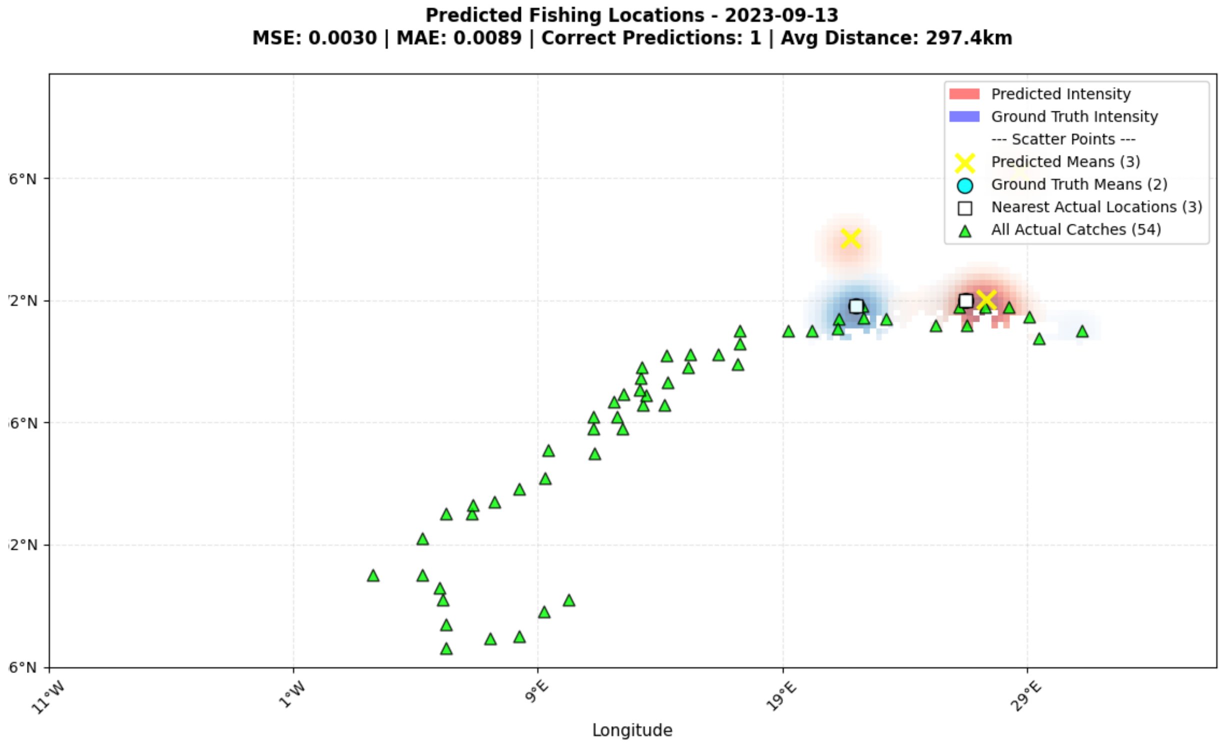

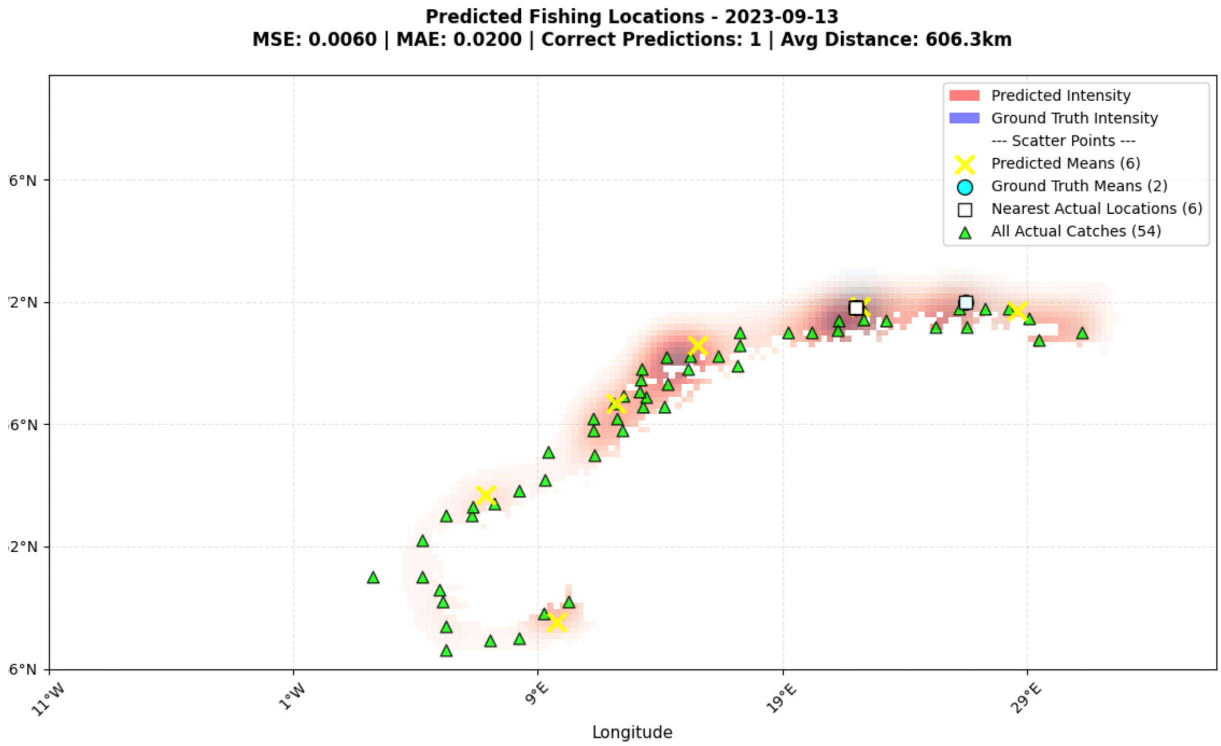

Figure 4.2: A sample from both the model (top) and baseline (bottom). The model's prediction are much closer to the actual fishing locations while those of the baseline are much more distributed. Note that while there are a lot of catch locations (54), most of these catch locations have an insignificant catch amount, meaning they are not recognized by the post-processing described in Section 3.2

Metric	Model (mean ± SD)	Baseline (mean ± SD)
MSE ↓	0.00408 ± 0.00147	0.00680 ± 0.00364
MAE ↓	0.01144 ± 0.00315	0.02066 ± 0.00833
Minimum distance ↓	151.73 ± 192.21	105.77 ± 165.96
Maximum distance ↓	418.25 ± 348.06	1177.03 ± 649.08
Mean distance ↓	270.90 ± 224.80	485.39 ± 289.60
Accuracy ↑	0.16985 ± 0.24979	0.17049 ± 0.18964
Total predictions	2.90 ± 1.32	5.98 ± 2.25
Correct predictions	0.50 ± 0.70	0.88 ± 0.87
Incorrect predictions	2.40 ± 1.28	5.10 ± 2.34

Table 4.1: Comparison of model versus baseline performance different metrics. The model achieves significantly better spatial accuracy, with predictions averaging 271 km from actual fishing locations compared to 485 km for the baseline. The model also shows superior performance in predicting the distribution, with an MAE of 0.011 versus 0.021 for the baseline, indicating that the distribution predicted by the model is superior to that of the baseline.

baseline creates more predictions than the model, most of which are incorrect.

Figure 4.3 plots the distance along the accuracy of the model and baseline. As the distance increases, more predictions are counted as correct predictions, increasing the accuracy. The Figure shows that as you increase the accepted distance, the accuracy of the model increases more than that of the baseline. The model makes predictions which are closer to the ground truth catch locations compared to the baseline.

5 Discussion

Key Findings. A diffusion model was created with the intent of predicting fish catch locations. A baseline was created to compare the model with. The baseline itself did not produce very meaningful results. It produced a lot of predictions of which a lot were incorrect with respect to the ground truth data. The model itself produced better results. This indicates that the model does have the capability to predict future fishing locations and generate a distribution map on cod.

The diffusion model used was able to produce meaningful maps despite the sparsity present in the data. This can occur because of three reasons. First of all, in theory, a diffusion model is able to learn any probability distribution, if the model capacity is large enough, there is enough data and the number of timesteps used is infinity (Song et al., 2021). The second reason for this convergence is that the model does not try to predict zero-valued numbers, it tries to predict the added noise, which is almost never zero. This means the loss function can accurately evaluate the models performance. Thirdly, the sparsity bits added allow the model to learn spatial information better. Figure 3.4 shows that the model used the extra layer to learn spatial information, as this layer contains the outline of Norway.

Whether the predictions produced by the model could actually be used is difficult to tell. The accuracy metric used in the study does not perfectly show how well the model performs, as the post-processing used has some error. Some predicted spots will not have been marked as a prediction, and vise versa. This means that the actual accuracy of the model may be higher or lower. Using the current metrics, you would need to cover a circular area of 200 km^2 to get an accuracy of 50%. Whether this is feasible is dependent on many factors. It should be noted that the coastline of Norway is about an order of magnitude larger than this Central Intelligence Agency (2024). This means that traveling to certain locations may only be profitable if it is near 100% correct. This would also be dependent on the selling price of cod, which is not consistent through the year Alizadeh Ashrafi et al. (2020). It would also make the implementation of a marine protected area more difficult, as you would need to cover a large area. To draw more decisive conclusions, future research should compare the accuracy of fisheries with that of the model.

The results generated by the model are in line with (Meeanan et al., 2023). This study also researched the influence of different features, and found that salinity, sea surface temperature, bathymetry, latitude, longitude, date and total catch was accountable for 17.4% of the predictive power. It should be noted that the study did not use visual based deep learning techniques but instead used classical machine learning techniques such as a Random Forest or Decision Tree model. Nevertheless, it explains the relatively low accuracy in our study.

This shows one of the major limitations of this project: the quality of the data. First of all, as

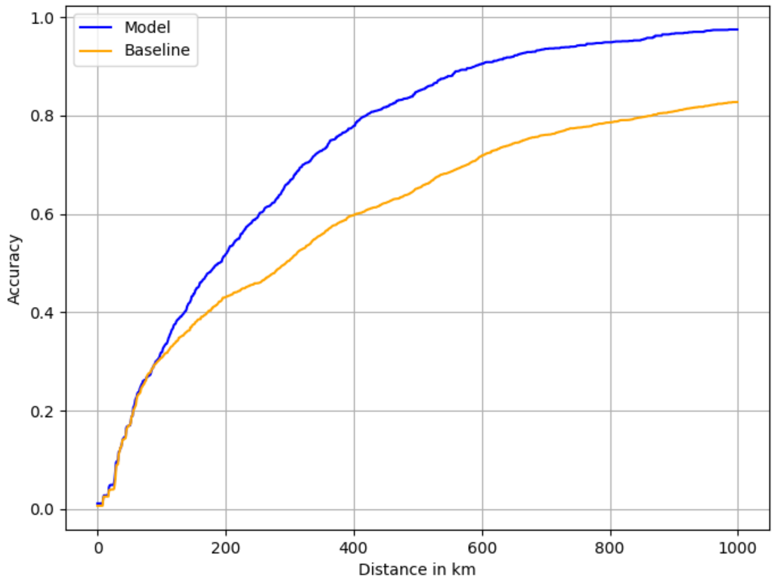

Figure 4.3: A plot plotting how much distance there can be between a prediction and catch location before it is counted as an incorrect prediction. The model and baseline perform similarly, with the baseline being slightly more accurate overall.

noted by other studies, there is a high bias in data. The catch locations are heavily determined by the decisions made by the vessels. These decisions are not only made based on the highest fish caught, but also for example on the weather, fuel consumption and price of fish. Furthermore the temporal granularity was very low, also limiting predictive power. Meeanan et al. (2023) found that information about the vessel and its behavior (such as vessel length, average speed, etc.) along with a day and night variable, accounted for about 69.9% of predictive power. The main contributions of these data would be that it would reduce sparsity, which is one of the major difficulties of this problem (Róason, 2024; Cronin, 2024). It would also allow for absence data to be incorporated, which is crucial information in most standard models and reduces their performance drastically if this is not available (Dhyani et al., 2023; Froese et al., 2023).

Broader Impacts. In this study, cod has been treated as a single stock due to the data provided. This resulted in a model treating cod as a single species. However, there are some considerations that need to be taken when treating cod in this way. One important observation is that cod consists of multiple different subspecies. This is also modeled by fisheries and institutions themselves Kualtska et al. (2024). However, treating cod as a single species can have detrimental effects. According to Hutchinson (2008), the treatment of the Atlantic Cod as a single stock can lead to loss on genetic diversity, failure of conservation measures, and threats to the ecosystem. By treating cod as a single stock, sustainable fishing may not be guaranteed despite a fully working model.

6 Conclusion

Being able to predict the presence and absence of fish is crucial not only for the efficient travel of fishing vessels, but is also required for certain measures taken against the overfishing of fish. In this study, a diffusion model was used to forecast the catch locations of cod one day ahead. By conditioning the model with the last 7 days of catch data, last day of salinity, sea surface temperature and bathymetry data along with the date, a properly trained model was created. Furthermore, a baseline was created originating from previous studies. This study shows that the model outperforms this baseline. The model performs lower in terms of mean squared error and mean absolute error with respect to the ground truth when compared to the model, which means it more closely resembles the ground truth than the baseline. This result is also highlighted by the minimum, maximum and mean distance of both models. While the minimum distance of all predictions made by the model is larger, the mean and maximum show that the model outperforms the baseline. As you increase the accepted distance for a correct prediction, the model's accuracy increases more than that of the baseline.. This study has shown that despite a large presence of sparsity in the data, the diffusion model can predict the presence of Cod. Although the model's performance may or may not be useful, this study suggest that by adding other features, the quality of data would be increased, allowing the model to make more accurate predictions. It is to be hoped that a model like this can be used to allow for more sustainable fishing, decreasing carbon dioxide emissions and allowing for healthy cod and fish stock sizes.

References

Tannaz Alizadeh Ashrafi, Arne Eide, and Øystein Hermansen. Spatial and temporal distributions in the norwegian cod fishery. *Natural Resource Modeling*, 33(4):e12276, 2020. doi:https://doi.org/10.1111/nrm.12276. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/nrm.12276.

Lee G. Anderson. Handbook of marine fisheries conservation and management. *Marine Resource Economics*, 27(4):360–366, 2012. ISSN 07381360, 23345985. URL http://www.jstor.org/stable/10.5950/0738-1360-27.4.397.

Salva Rühling Cachay, Bo Zhao, Hailey Joren, and Rose Yu. Dyffusion: A dynamics-informed diffusion model for spatiotemporal forecasting, 2023. URL https://arxiv.org/abs/2306.01984.

Central Intelligence Agency. Coastline – the world factbook. https://www.cia.gov/the-world-factbook/field/coastline/, 2024. Accessed: 2025-07-11.

Benoit Courty, Victor Schmidt, Sasha Luccioni, Goyal-Kamal, Marion Coutarel, Boris Feld, Jérémy Lecourt, Liam Connell, Amine Saboni, Inimaz, supatomic, Mathilde Léval, Luis Blanche, Alexis Cruveiller, ouminasara, Franklin Zhao, Aditya Joshi, Alexis Bogroff, Hugues de Lavoreille, Niko Laskaris, Edoardo Abati, Douglas Blank, Ziyao Wang, Armin Catovic, Marc Alencon, Michał Stęchły, Christian Bauer, Lucas Otávio N. de Araújo, JPW, and MinervaBooks. mlco2/codecarbon: v2.4.1, May 2024. URL https://doi.org/10.5281/zenodo.11171501.

Luc Cronin. Fish catch optimisation using variational autoencoder, july 2024. URL https://fse.studenttheses.ub.rug.nl/33749/. Supervisor: J. D. Cardenas Cartagena.

Jonas Dammen, Åsmund Brekke, Kristian Andersen Hole, Ludvig Løddesøl, Tomas Roaldsnes, and Julia Ortheden. Fishai: The lodestar fishing platform. *Nordic Machine Intelligence*, 2(2): 10–12, November 2022. doi:10.5617/nmi.9931.

Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis, 2021. URL https://arxiv.org/abs/2105.05233.

Shalini Dhyani, Dibyendu Adhikari, Rajarshi Dasgupta, and Rakesh Kadaverugu, editors. *Ecosystem and Species Habitat Modeling for Conservation and Restoration*. Springer Nature Singapore, Singapore, 1st edition, 2023. ISBN 978-981-99-0131-9. doi:10.1007/978-981-99-0131-9.

Fiskeridirektoratet. Åpne data: Fangstdata (seddel) koblet med fartøydata. URL https://www.fiskeridir.no/statistikk-tall-og-analyse/data-og-statistikk-om-yrkesfiske. Accessed: 2025-02-25.

Flanders Marine Institute. Union of the esri country shapefile and the exclusive economic zones (version 4). Available online at MarineRegions.org, 2024. URL https://www.marineregions.org/. Consulted on 2025-03-06.

Food and Agriculture Organization of the United Nations. *The State of World Fisheries and Aquaculture 2024: Blue Transformation in Action*. State of World Fisheries and Aquaculture. Food and Agriculture Organization of the United Nations, Rome, 2024. ISBN 978-92-5-138763-4. doi:10.4060/cd0683en. URL https://openknowledge.fao.org/items/06690fd0âĂŠd133âĂŠ424câĂŠ9673âĂŠ1849e414543d.

Rainer Froese, Henning Winker, Gianpaolo Coro, Maria-Lourdes "Deng" Palomares, Athanassios C. Tsikliras, Donna Dimarchopoulou, Konstantinos Touloumis, Nazli Demirel, Gabriel M. S. Vianna, Giuseppe Scarcella, Rebecca Schijns, Cui Liang, and Daniel Pauly. new developments in the analysis of catch time series as the basis for fish stock assessments: The cmsy++ method. *Acta Ichthyologica et Piscatoria*, 53:173–189, 2023. ISSN 0137-1592. doi:10.3897/aiep.53.105910. URL https://doi.org/10.3897/aiep.53.105910.

GEBCO Compilation Group. Gebco 2024 grid, 2024. URL https://doi.org/10.5285/1c44ce99-0a0d-5f4f-e063-7086abc0ea0f.

Peter Gullestad, Asgeir Aglen, Åsmund Bjordal, Geir Blom, Sverre Johansen, Jørn Krog, Ole Arve Misund, and Ingolf Røttingen. Changing attitudes 1970–2012: evolution of the norwegian management framework to prevent overfishing and to secure long-term sustainability. *ICES Journal of Marine Science*, 71(2):173–182, 07 2013. ISSN 1054-3139. doi:10.1093/icesjms/fst094. URL https://doi.org/10.1093/icesjms/fst094.

Benjamin S. Halpern, Sarah E. Lester, and Karen L. McLeod. Placing marine protected areas onto the ecosystem-based management seascape. *Proceedings of the National Academy of Sciences*, 107(43):18312–18317, 2010. doi:10.1073/pnas.0908503107. URL https://www.pnas.org/doi/abs/10.1073/pnas.0908503107.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models, 2020. URL https://arxiv.org/abs/2006.11239.

Victoria J. Hobson, David Righton, Julian D. Metcalfe, and Graeme C. Hays. Link between vertical and horizontal movement patterns of cod in the north sea. *Aquatic Biology*, 5:133–142, 2009. doi:10.3354/ab00144. URL https://www.int-res.com/abstracts/ab/v5/ab00144.

Boyin Huang, Chengfei Liu, Viva Banzon, Eric Freeman, Gary Graham, Benjamin Hankins, Thomas Smith, and Huai-Min Zhang. Improvements of the daily optimum interpolation sea surface temperature (doisst) version 2.1. *Journal of Climate*, 34:2923–2939, 2021. doi:10.1175/JCLI-D-20-0166.1.

William F. Hutchinson. The dangers of ignoring stock complexity in fishery management: the case of the north sea cod. *Biology Letters*, 4(6):693–695, December 2008. doi:10.1098/rsbl.2008.0443.

Ulf Johansen, Heidi Bull-Berg, Lars H. Vik, Arne M. Stokka, Roger Richardsen, and Ulf Winther. The norwegian seafood industry – importance for the national economy. *Marine Policy*, 110:103561, 2019. ISSN 0308-597X. doi:https://doi.org/10.1016/j.marpol.2019.103561. URL https://www.sciencedirect.com/science/article/pii/S0308597X1830914X.

Tero Karras, Miika Aittala, Jaakko Lehtinen, Janne Hellsten, Timo Aila, and Samuli Laine. Analyzing and improving the training dynamics of diffusion models, 2024. URL https://arxiv.org/abs/2312.02696.

Nataliia Kualtska, Daniel Howell, Peter J. Wright, and Ingibjörg G. Jónsdóttir. *Biology and Ecology of Atlantic Cod*, pages 271 – 292. Taylor & Francis Group, 2024. doi:https://doi.org/10.1201/9781003120872.

Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11): 2278–2324, 1998. doi:10.1109/5.726791.

Hui Li, Zhouyang Ren, Miao Fan, Wenyuan Li, Yan Xu, Yunpeng Jiang, and Weiyi Xia. A review of scenario analysis methods in planning and operation of modern power systems: Methodologies, applications, and challenges. *Electric Power Systems Research*, 205:107722, 2022. ISSN 0378-7796. doi:https://doi.org/10.1016/j.epsr.2021.107722. URL https://www.sciencedirect.com/science/article/pii/S0378779621007033.

S. Libralato. Numerical models for monitoring and forecasting ocean ecosystems: a short description of the present status. *State of the Planet*, 5-opsr:13, 2025. doi:10.5194/sp-5-opsr-13-2025. URL https://sp.copernicus.org/articles/5-opsr/13/2025/.

Visa Linkiö, Kristian Lahtinen, and Juho Kolmonen. Clusters and traveling fisherman. *Nordic Machine Intelligence*, 2(2):13–15, 2022. doi:10.5617/nmi.9930. URL https://journals.uio.no/NMI/article/view/9930. FishAI Competition.

Junqi Liu, Hao Liu, Tianyu Hu, and Huimin Ma. Wind speed forecasting based on denoising diffusion model and multivariate-aware spatiotemporal transformer. In *2024 IEEE 8th Conference on Energy Internet and Energy System Integration (EI2)*, pages 145–150, 2024. doi:10.1109/EI264398.2024.10991433.

Matteo Mazziotta and Adriano Pareto. Normalization methods for spatio-temporal analysis of environmental performance: Revisiting the min–max method. *Environmetrics*, 33(5):e2730, 2022. doi:https://doi.org/10.1002/env.2730. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/env.2730.

Chonlada Meeanan, Pavarot Noranarttragoon, Piyachoke Sinanun, Yuki Takahashi, Methee Kaewnern, and Takashi Fritz Matsuishi. Estimation of the spatiotemporal distribution of fish and fishing grounds from surveillance information using machine learning: The case of short mackerel in the andaman sea, thailand. *REGIONAL STUDIES IN MARINE SCIENCE*, 62, SEP 2023. doi:10.1016/j.rsma.2023.102914.

Chonlada Meeanan, Pavarot Noranarttragoon, Piyachoke Sinanun, Wirat Sanitmajjaro, Yuki Takahashi, Methee Kaewnern, and Takashi Fritz Matsuishi. Using vessel surveillance data to estimate spatiotemporal patterns in the short mackerel purse-seine fishery: Implications for time–area closure management in thai waters. *Marine*

Policy, 170:106382, 2024. ISSN 0308-597X. doi:https://doi.org/10.1016/j.marpol.2024.106382. URL https://www.sciencedirect.com/science/article/pii/S0308597X24003804.

Rui Nian, Qiang Yuan, Hui He, Xue Geng, Chi-Wei Su, Bo He, and Amaury Lendasse. The identification and prediction in abundance variation of atlantic cod via long short-term memory with periodicity, time-frequency co-movement, and lead-lag effect across sea surface temperature, sea surface salinity, catches, and prey biomass from 1919 to 2016. *FRONTIERS IN MARINE SCIENCE*, 8, JUN 7 2021. doi:10.3389/fmars.2021.665716.

TorArne Schmidt Nordmo, Ove Kvalsvik, Svein Ove Kvalsund, Birte Hansen, Pål Halvorsen, Steven Hicks, Dag Johansen, Håvard D. Johansen, and Michael Alexander Riegler. Fish AI: Sustainable Commercial Fishing Challenge. *Nordic Machine Intelligence*, 2(2):1–3, June 2022. ISSN 2703-9196. doi:10.5617/nmi.9657. URL https://journals.uio.no/NMI/article/view/9657.

Phil Ostheimer, Mayank Nagda, Jean Radig, Carl Herrmann, Stephan Mandt, Marius Kloft, and Sophie Fellenz. Sparse data generation using diffusion models, 2025. URL https://arxiv.org/abs/2502.02448.

Diego Panzeri, Tommaso Russo, Enrico Arneri, Roberto Carlucci, Gianpiero Cossarini, Ivana Isajlović, Sanja Krstulović Šifner, Claudia Manfredi, Francesco Masnadi, Marco Reale, Giulia Scarcella, Cosimo Solidoro, Maria Teresa Spedicato, Nikola Vrgoč, Walter Zupa, and Stefano Libralato. Identifying priority areas for spatial management of mixed fisheries using ensemble of multi-species distribution models. *Fish and Fisheries*, 25:187–204, 2024. doi:10.1111/faf.12802.

Robert W. R. Parker, Julia L. Blanchard, Caleb Gardner, Bridget S. Green, Klaas Hartmann, Peter H. Tyedmers, and Reg A. Watson. Fuel use and greenhouse gas emissions of world fisheries. *Nature Climate Change*, 8(4):333–337, Apr 2018. ISSN 1758-6798. doi:10.1038/s41558-018-0117-x. URL https://doi.org/10.1038/s41558-018-0117-x.

Tom Patterson and Nathaniel Vaughn Kelso. Natural earth 10m land (physical vector). Natural Earth, public domain, 2024. URL https://www.naturalearthdata.com/downloads/10m-physical-vectors/10m-land/.

Steven J. Phillips, Robert P. Anderson, and Robert E. Schapire. Maximum entropy modeling of species geographic distributions. *Ecological Modelling*, 190(3):231–259, 2006. ISSN 0304-3800. doi:https://doi.org/10.1016/j.ecolmodel.2005.03.026. URL https://www.sciencedirect.com/science/article/pii/S030438000500267X.

Kashif Rasul, Calvin Seward, Ingmar Schuster, and Roland Vollgraf. Autoregressive denoising diffusion models for multivariate probabilistic time series forecasting, 2021. URL https://arxiv.org/abs/2101.12072.

Remote Sensing Systems (RSS). Smap sea surface salinity products. ver. 6.0. PO.DAAC, CA, USA, 2024. URL https://doi.org/10.5067/SMP60-3SPCS. Dataset accessed 2025-03-03.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation, 2015. URL https://arxiv.org/abs/1505.04597.

Jón Róason. Fish location forecasting in the norwegian ocean using deep learning techniques, August 2024. URL https://fse.studenttheses.ub.rug.nl/34256/. Supervisor: Juan Diego Cárdenas-Cartagena.

Kim J.N. Scherrer, Tom J. Langbehn, Gabriella Ljungström, Katja Enberg, Sara Hornborg, Gjert Dingsør, and Christian Jørgensen. Spatial restrictions inadvertently doubled the carbon footprint of norway's mackerel fishing fleet. *Marine Policy*, 161:106014, 2024. ISSN 0308-597X. doi:https://doi.org/10.1016/j.marpol.2024.106014. URL https://www.sciencedirect.com/science/article/pii/S0308597X24000125.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations, 2021. URL https://arxiv.org/abs/2011.13456.

Yashwant Soni, Uma Meena, Vikash Kumar Mishra, and Pramod Kumar Soni. AM-UNet: Road Network Extraction from high-resolution Aerial Imagery Using Attention-Based Convolutional Neural Network. *Journal of the Indian Society of Remote Sensing*, 53(1):135–147, 2025. ISSN 0974-3006. doi:10.1007/s12524-024-01974-3. URL https://doi.org/10.1007/s12524-024-01974-3.

Craig Syms. Satellite ocean data can inform precision fishing. *Nordic Machine Intelligence*, 2(2): 4–6, November 2022. doi:10.5617/nmi.9945.

Ting ting Zhang, Zhi Geng, Xiao rong Huang, Yu Gao, Si kai Wang, Tao Zhang, Gang Yang, Feng Zhao, and Ping Zhuang. Mapping essential habitat of estuarine fishery species with a mechanistic sdm and landsat data. *Ecological Indicators*, 142:109196, 2022. ISSN 1470-160X. doi:https://doi.org/10.1016/j.ecolind.2022.109196.

URL https://www.sciencedirect.com/science/article/pii/S1470160X22006689.

Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM Comput. Surv.*, 56(4), 2023. ISSN 0360-0300. doi:10.1145/3626235. URL https://doi.org/10.1145/3626235.

Yiyuan Yang, Ming Jin, Haomin Wen, Chaoli Zhang, Yuxuan Liang, Lintao Ma, Yi Wang, Chenghao Liu, Bin Yang, Zenglin Xu, Jiang Bian, Shirui Pan, and Qingsong Wen. A survey on diffusion models for time series and spatio-temporal data, 2024. URL https://arxiv.org/abs/2404.18886.

A Appendix

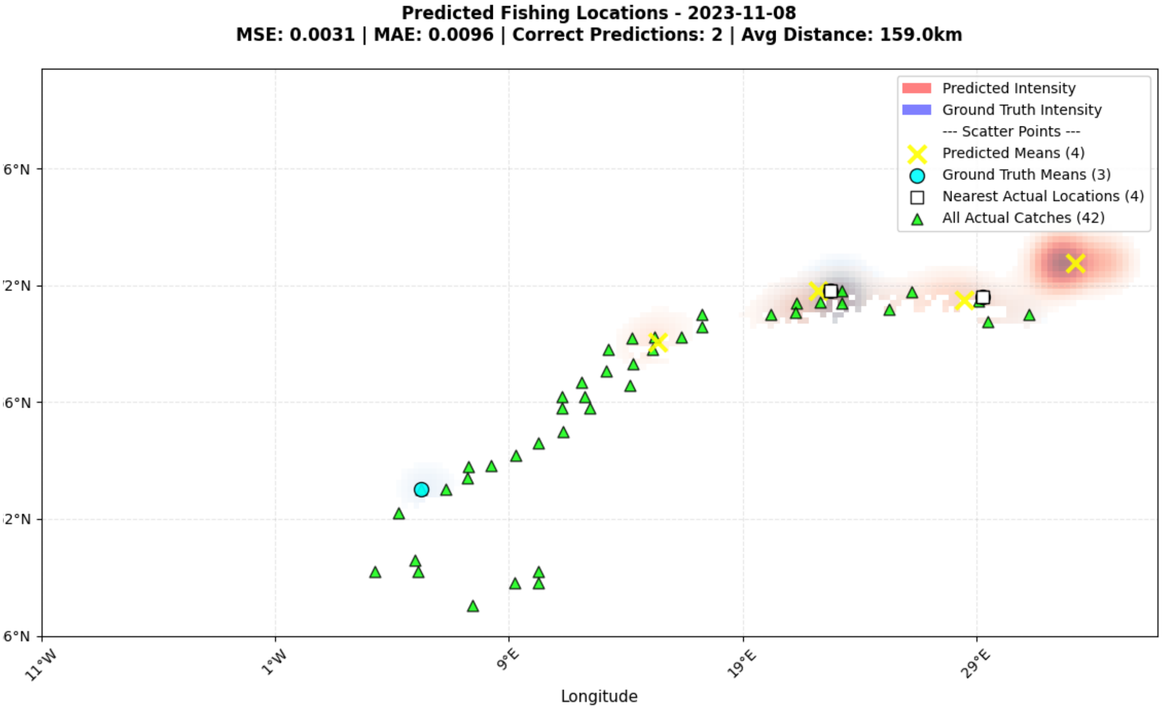

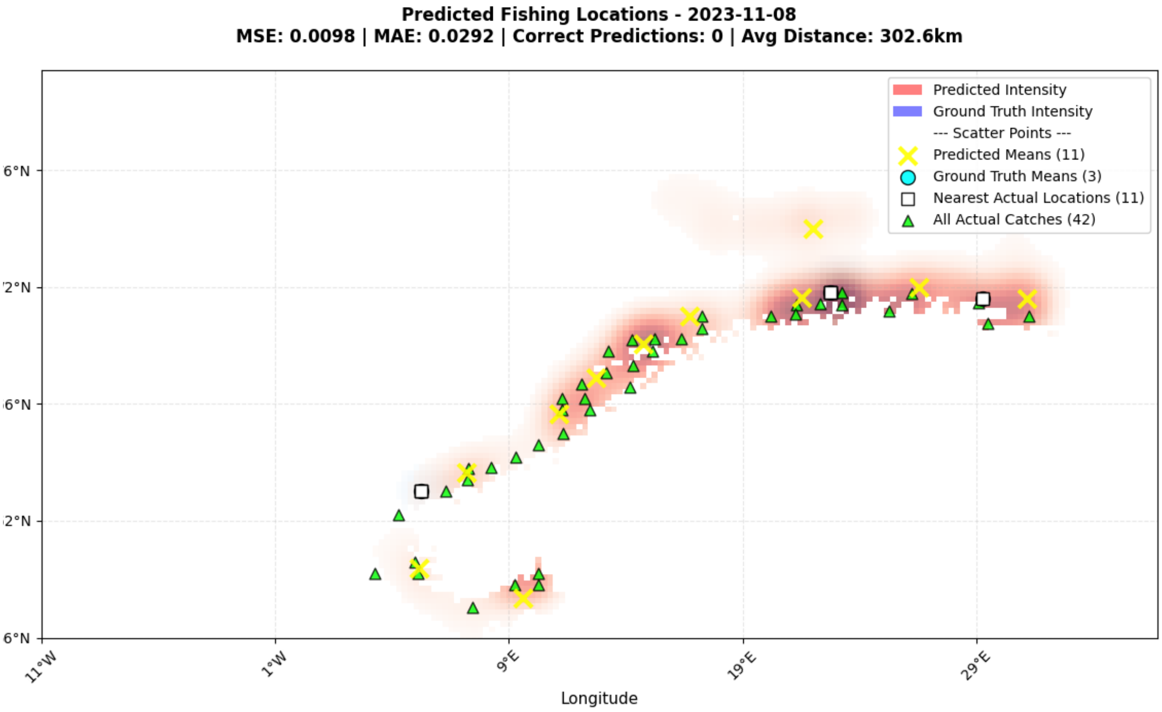

Figure A.1: Samples generated by the model (top) and baseline (bottom).

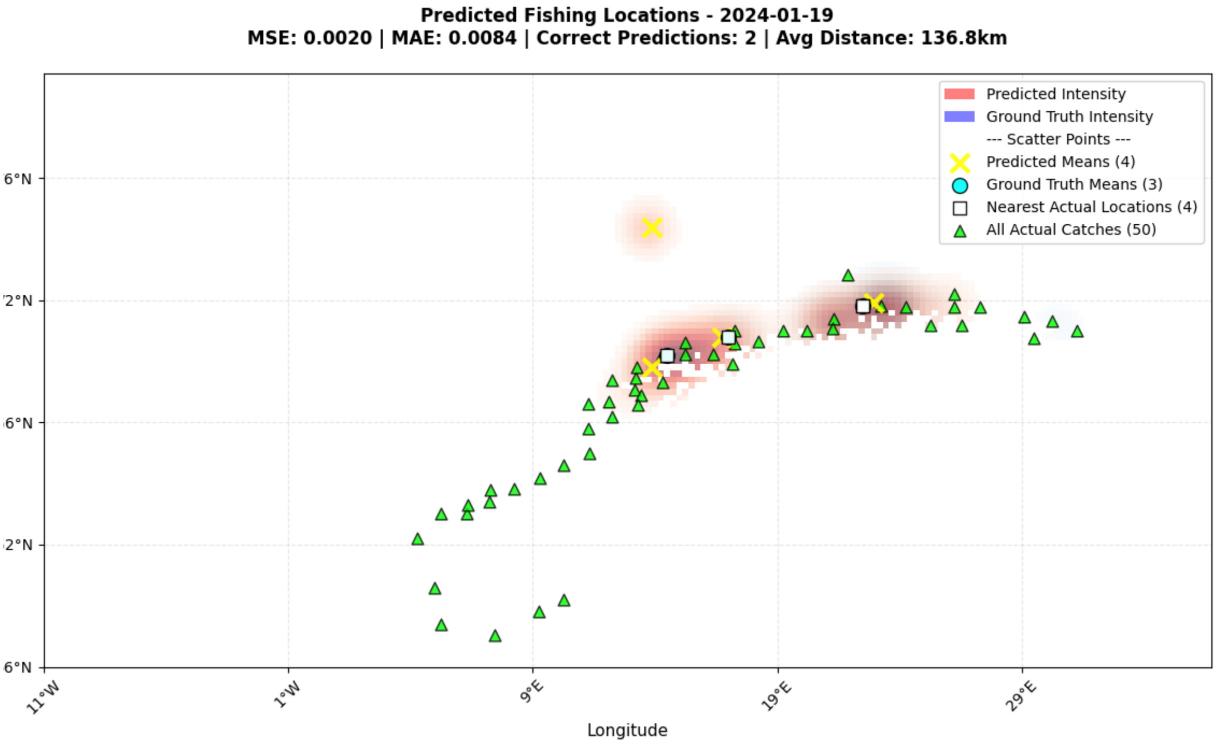

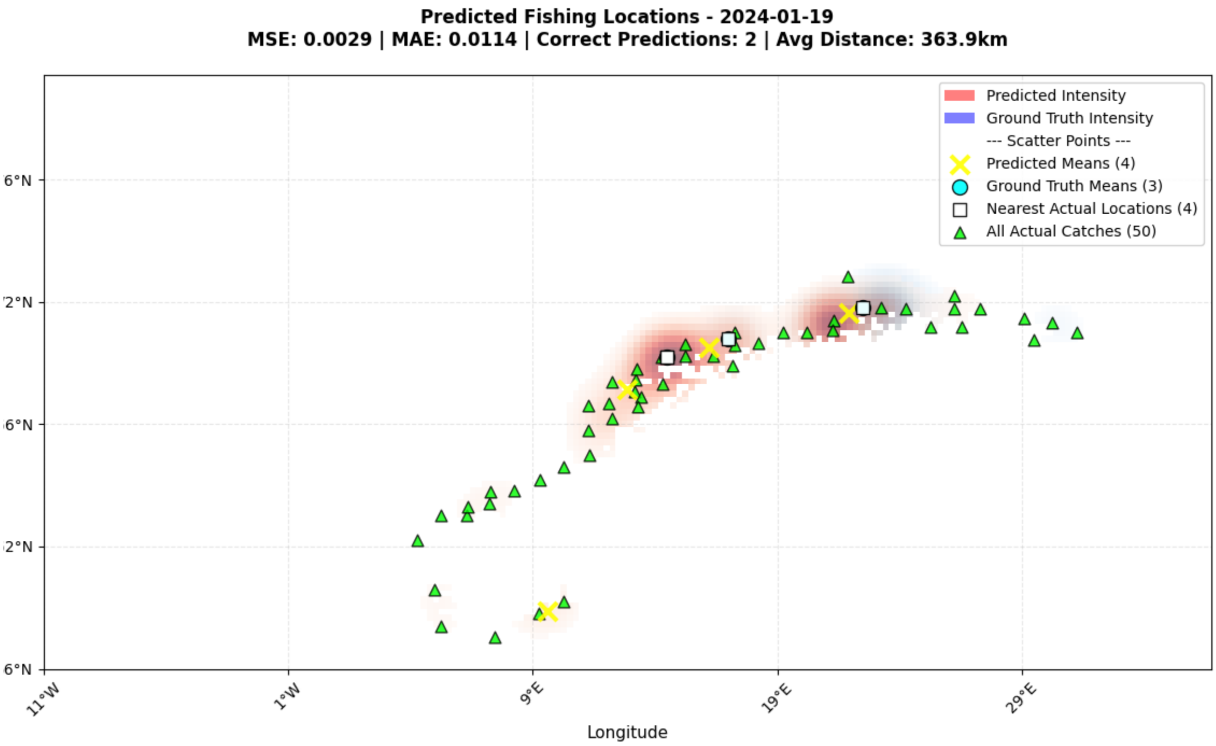

Figure A.2: Samples generated by the model (top) and baseline (bottom).

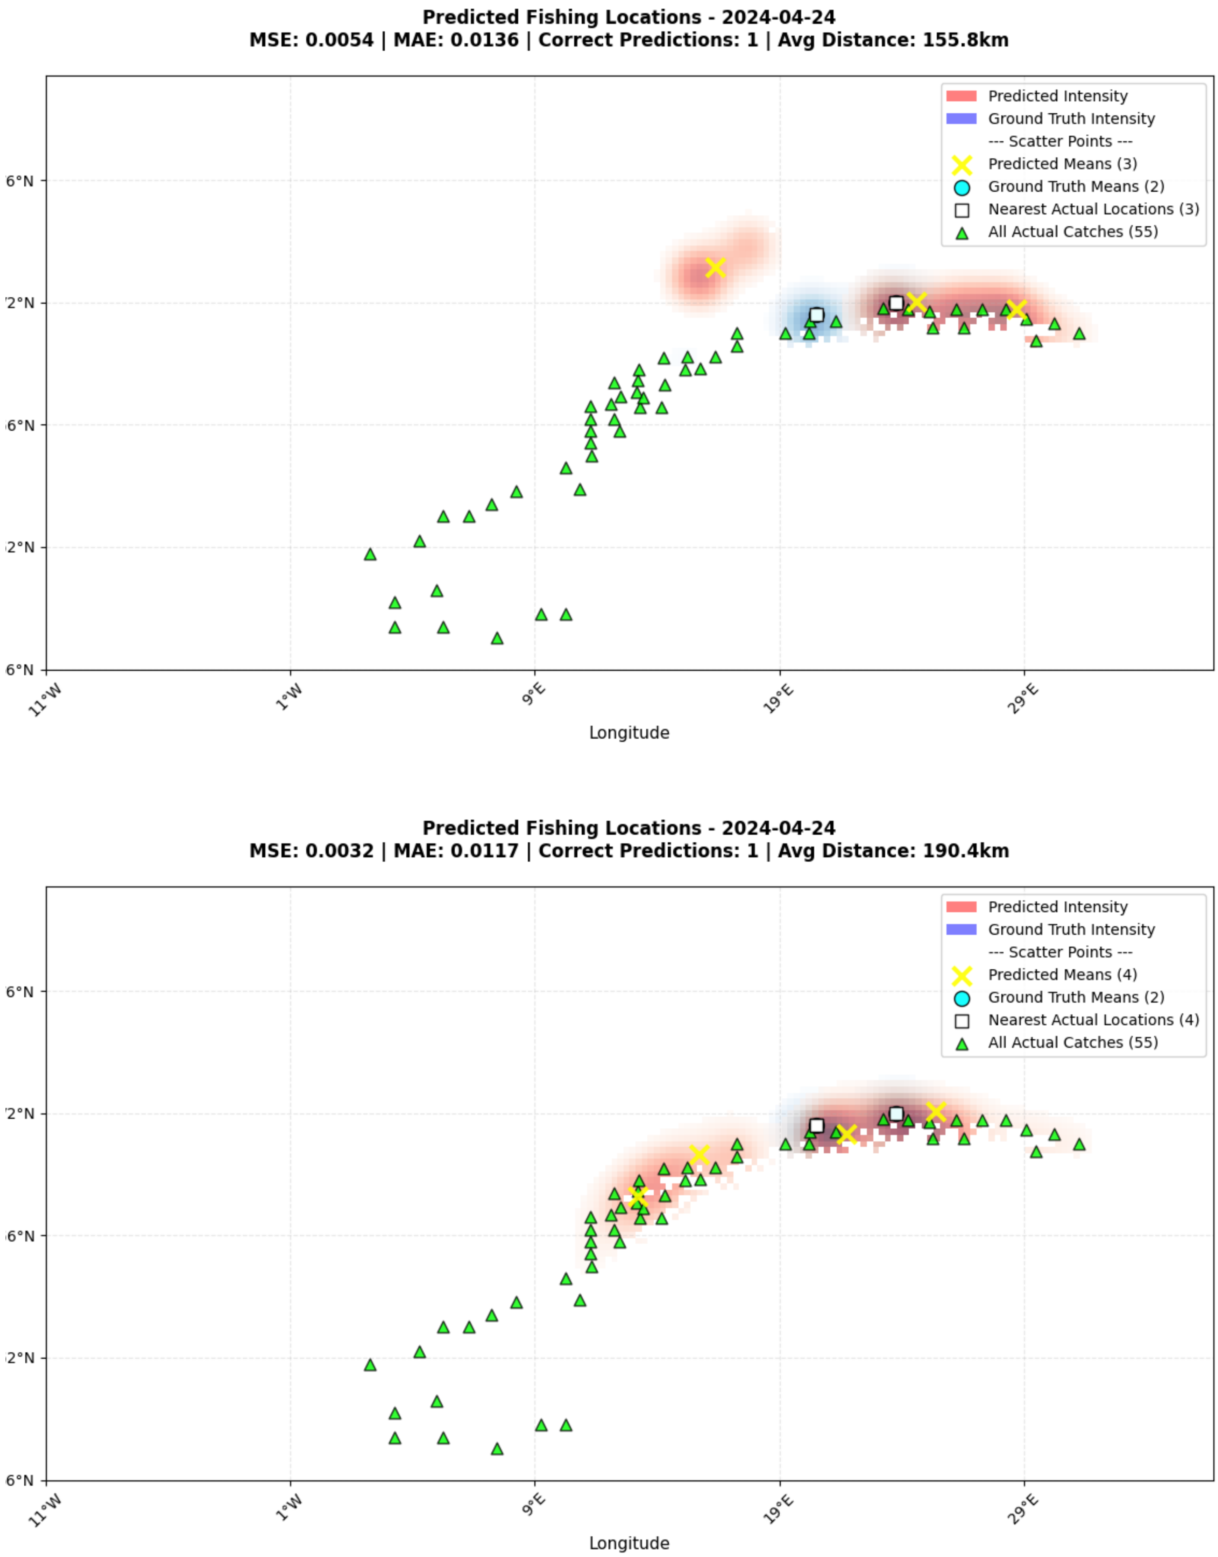

Figure A.3: Samples generated by the model (top) and baseline (bottom).

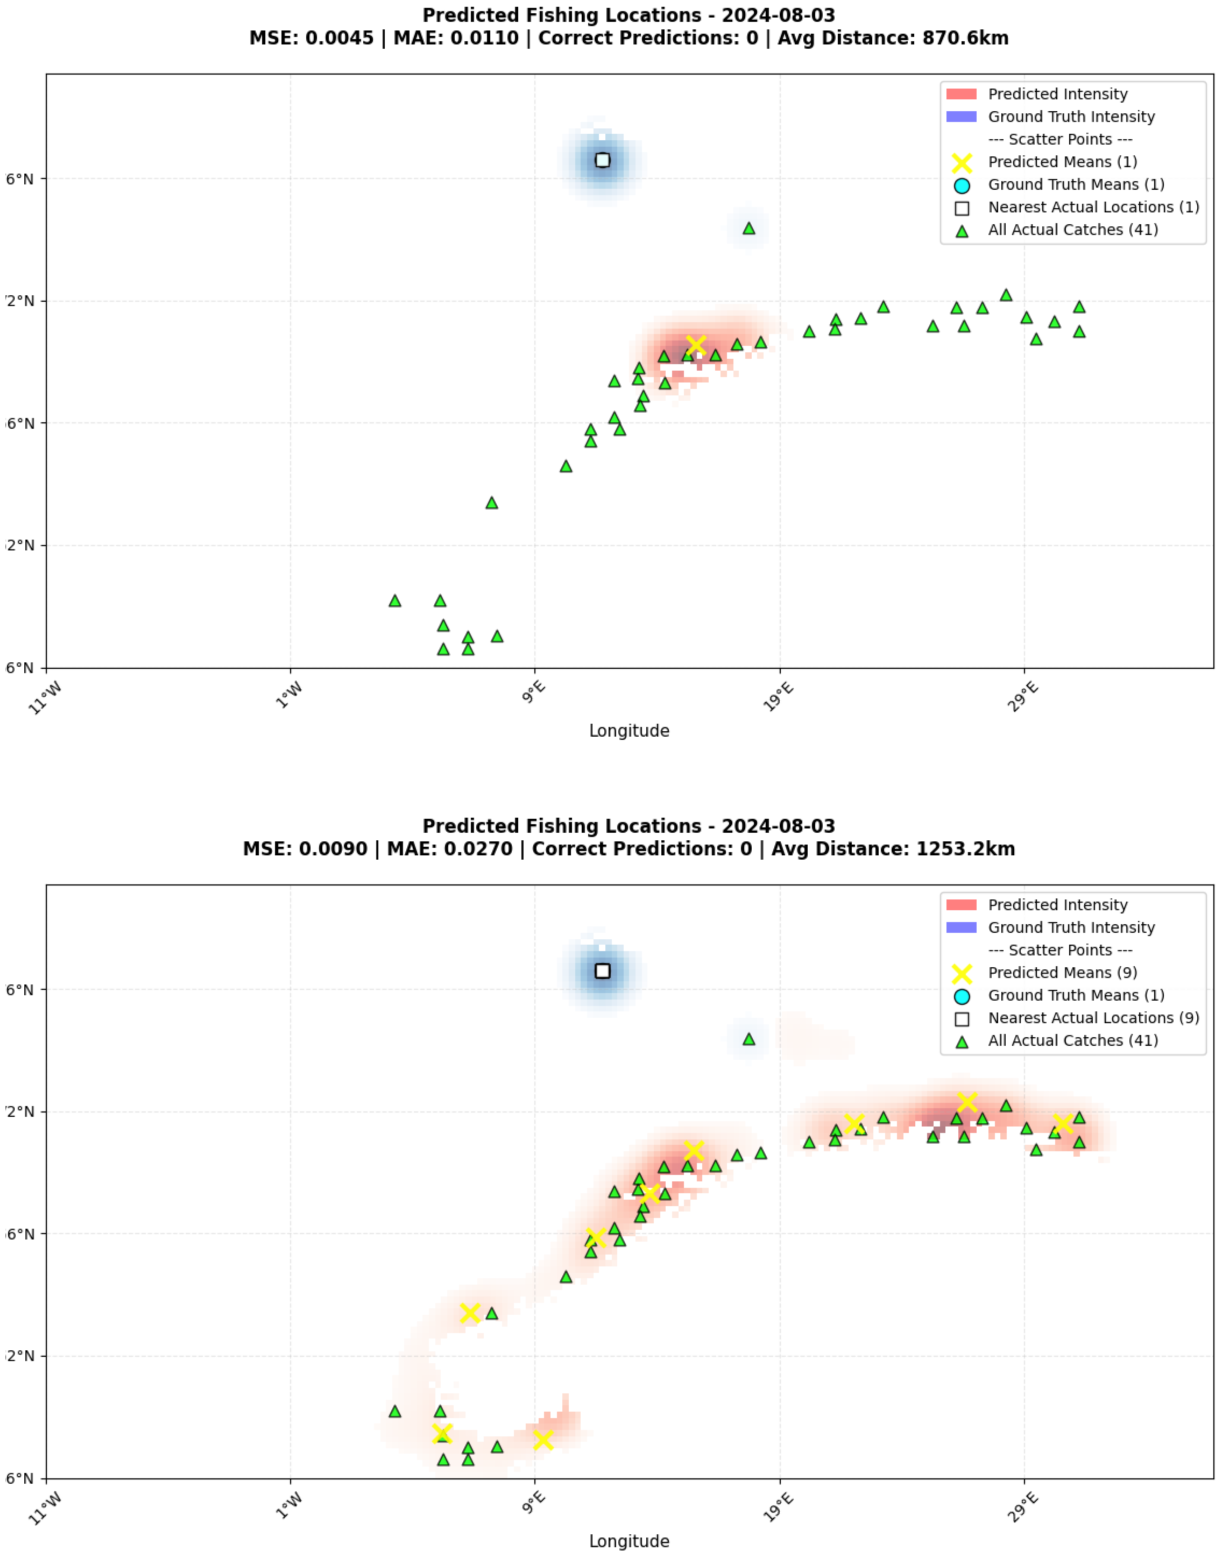

Figure A.4: Samples generated by the model (top) and baseline (bottom).

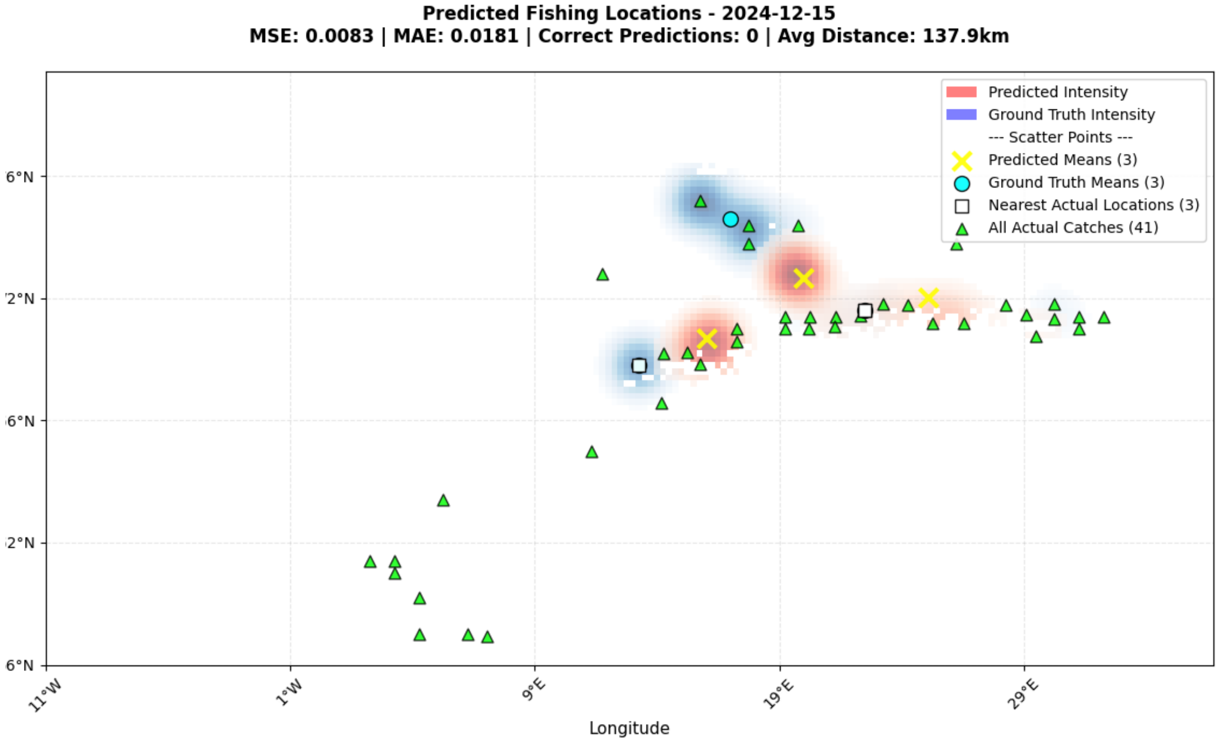

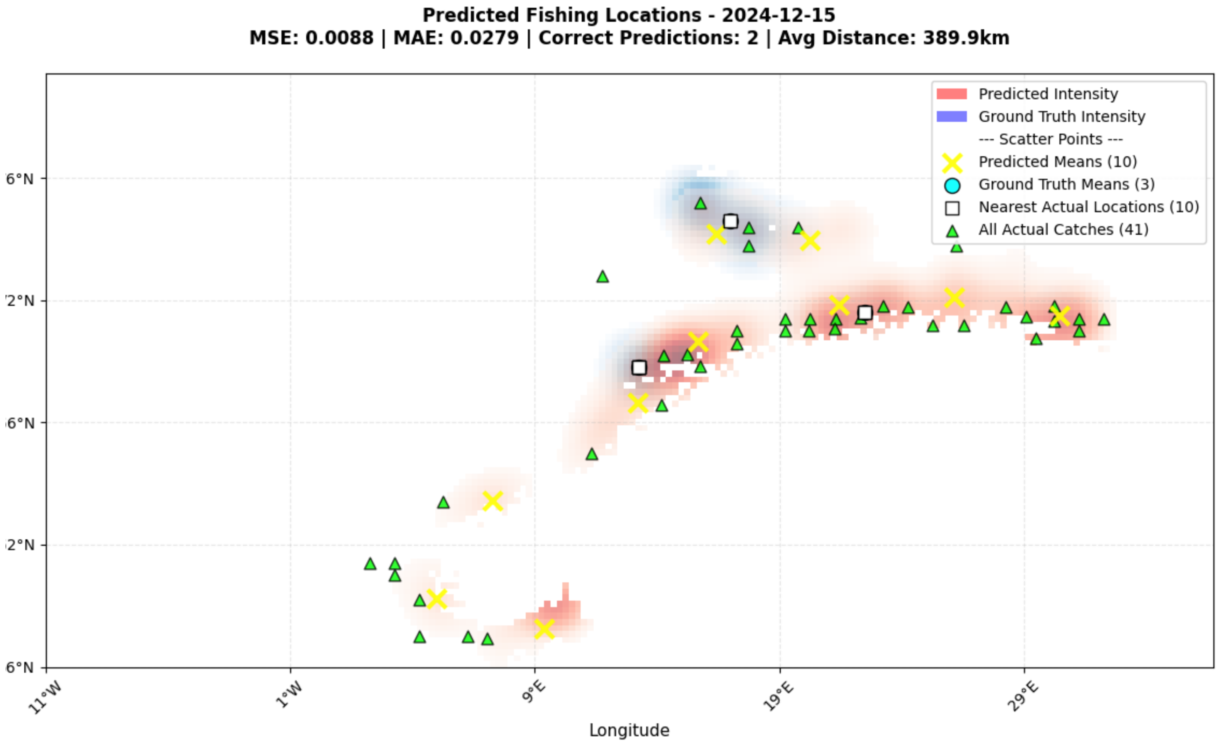

Figure A.5: Samples generated by the model (top) and baseline (bottom).