EMODnet Thematic Lot n°1 – Bathymetry

EASME/EMFF/2019/1.3.1.9/Lot1/SI2.836043

Start date of the project: 20/12/2022 (24 months)

Centralisation Phase

D3.12 - Deep-learning methodology for assessing bathymetry between coastline and foreshore

December 2024

The European Marine Observation and Data Network (EMODnet) is financed by the European Union under Regulation (EU) 2021/1139 of the European Parliament and of the Council of 7 July 2021 establishing the European Maritime, Fisheries and Aquaculture Fund.

Disclaimer

Document info

Title (and reference)	Deep-learning methodology for assessing bathymetry between coastline and foreshore
WP title (and reference number)	WP3
Task (and reference number)	D3.12
Authors [affiliation]	Willem Tromp, Matthijs Gawehn, Martin Verlaan
Dissemination level	Public
Submission date	18/12/2024
Deliverable due date	30/11/2024

Table of Contents

1 Introduction

Despite enormous improvements over the years, there are still numerous coastal inlets and some other smaller regions for which EMODnet-bathymetry has not been able to find surveys and fully achieve the goal of mapping all European seas and coasts. In the EMODnet central portal the bathymetry section also contains a source layer that shows this very well. All coastal areas marked as grey (GEBCO) were not covered with surveys yet. Also, there are inlets where the land-sea mask is not set as sea. This is for example the case where the Loire River runs into the ocean at Saint Nazaire. The current EMODnet grid provides data until the bridge but lacks data inland from the bridge. In summary, there are areas where there is currently a lack of available surveys, that is now filled by interpolation or excluded by the land-sea mask. This is of course a valid representation of the status, since both the source identifiers and the land-sea mask make the lack of data visible to the users. However, for several applications of coastal bathymetry, especially those that make use of numerical models, this may lead to large errors that also deteriorate the results in the surrounding region. For example, a tidal model will experience a reflection of the tidal wave at the Saint Nazaire bridge, where in reality tides will enter into the estuary and to a large extend be dissipated there.
In this report we describe a methodology based on generative AI to generate one (or multiple) plausible realizations of the missing areas. These estimates will necessarily be quite inaccurate but may nonetheless represent several important features realistically. In the future we aim to fill the existing gaps in the gridded bathymetry with such a tool, if only to bridge the gap until surveys become available.
To assess the performance of the inpainting, a case study is presented for the Humber Estuary. For comparison, alternative estimates of the bathymetry are derived based on two methods. The first method uses the propagation of wind waves as visible on Sentinel-2 imagery. The second method uses the Water Index also derived from Sentinel two in combination with data from a tide-gauge.

Figure 1 EMODnet gridded bathymetry (2023 release) for Loire River near saint Nazaire. Note the lack of bathymetry in the estuary.

Figure 2 Looking landward (east) from saint Nazaire bridge, showing the area where we will aim to fill the gap in bathymetry (source Google Streetview)

2 Methods

2.1 Inpainting methods for bathymetry

This task aims to develop a generative machine learning approach, that fills the bathymetry in areas where there is a gap in the data, with the aim of generating a plausible bathymetry. This goal is different from the best guess, which will in general be very smooth since smaller details cannot be predicted in the right locations. The result of best-guess interpolation therefore lacks detail. Here we aim to generate a similar amount of detail. The rationale is that for several applications it is more useful to generate a channel in an estuary in the wrong position, than no channel at all.

In recent years, several generative AI techniques have been developed for images. Since a gridded bathymetry is conceptually not very different from an image, we look towards this pioneering field for possible approaches. In this study, we based our work on the paper by Song etal (Song2020). The key concept has later been popularized as stable diffusion. The method used here is similar to the 'stable diffusion' model, but also different in several ways. The main idea (see also Figure 3) is to gradually add noise to an image, step by step, until there is nothing left but noise. Next, we use these distorted images to train an AI model to step backwards in this sequence towards less noisy images. This process is unstable, but with a generative AI model can be used to create somewhat random samples of images. In its raw form the method will generate random images that have similarities with the images in the training dataset. In the method by Song etal the forward steps of adding noise is described by a stochastic differential equation. The method then uses a formulation for a reversed stochastic differential equation. This equation requires knowledge of the distribution of the noisy images p(t), but instead of trying to model this, the gradient/sensitivities are modelled directly. This term, the 'score function' is modelled with a deep-learning approach using a dataset of noisy images. The generation process starts with a noise only image and move step by step towards an image. The score function tries to identify similarities between the

noise and the learned images and extra noise is added to the process to avoid an early collapse. This implies that if the process is repeated, it will result in a different image.

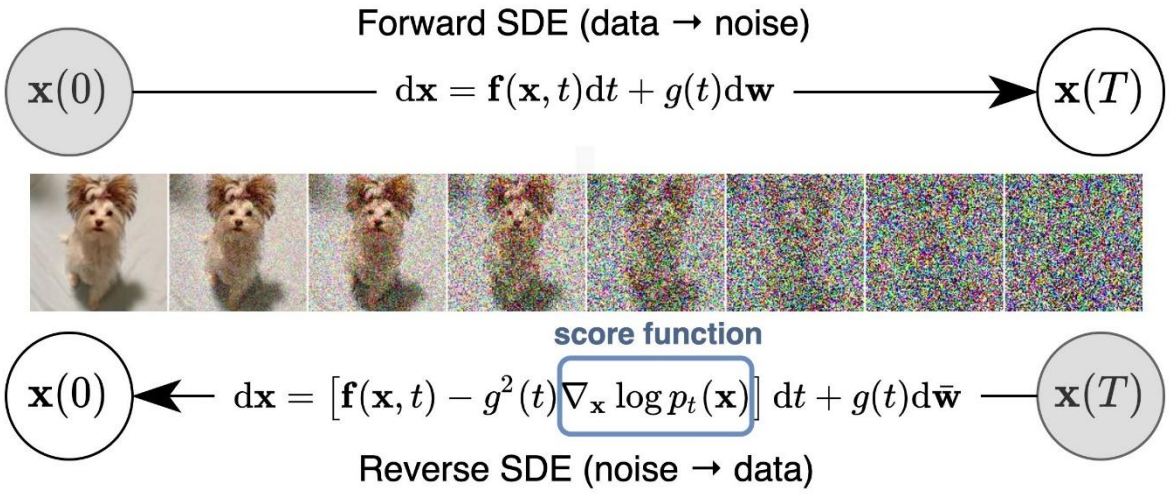

Figure 3 Concept of the diffusion idea and method by Song etal (source Song2020)

In this raw form the method can often not be applied directly, instead we usually want to achieve a secondary goal, for example generate an image of a specific topic like a dog, instead of an arbitrary one. Here, we aim to only generate the missing part of the image but keep the part that is already known. We use a relaxation term towards the known part of the image for this purpose.

For a brief illustration, the standard MNIST dataset of 60000 handwritten images was used. Figure 4 shows some original images on the top left. Part of the image was masked to mimic lacking parts of the bathymetry and on the right two samples are shown. Note that for example the second digit becomes 1 in the first sample, but 2 in the second sample. This is understandable, given that it is also for us humans hard to see the correct digit from the masked image and the original images are not used in the process, but only shown for comparison.

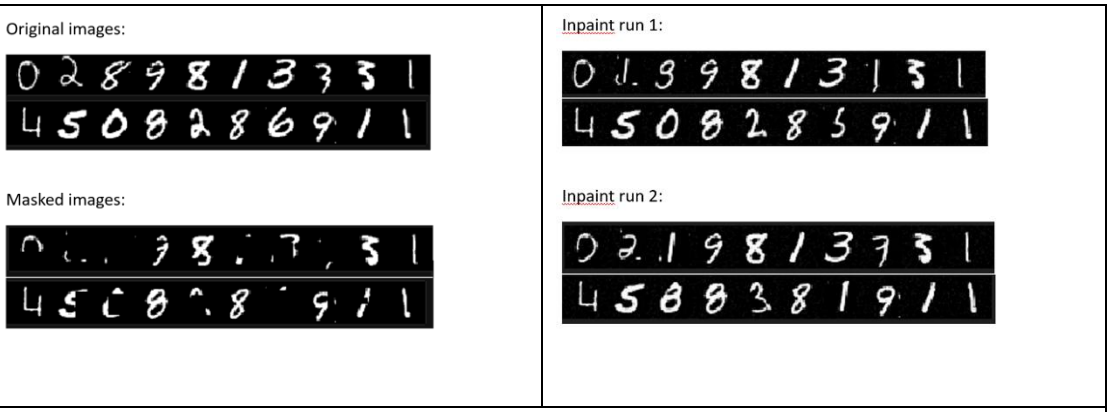

Figure 4 Test of score-based diffusion model on MNIST dataset.

To fill missing parts of the gridded bathymetry, we use small tiles (128x128) for training sampled randomly from the gridded bathymetry. We settled on a normalization method, where each 128x128 image is normalized to the unit range [0,1] individually, instead of applying unit range normalization to the data set as a whole. This way the dynamic range of the images is consistent across the data set. This partly addresses the issue of noisy sampled images. When applying normalization across the full data set, the

parts of the image we are interested in sampling will have a narrow dynamic range comparable to the noise scale added each time step during the diffusion process. Switching to a per-image normalization increases the dynamic range of the area of interest, separating it from the noise scale.

Additionally on the preprocessing side, we added data augmentation to the pipeline. We augment the data by applying 90, 180, 270 deg rotations to the images. This way the training data volume increases by a factor 4. Other augmentation operations can still be added, but the augmentation process needs to be optimized first to avoid memory issues.

Finally, we explored applying a convolutional filter to the landmasks and add it to the original landmasks to have non-zero weight for land pixels close to the coastline. This way we aim to generate more realistic coastlines when we loosen the restrictions during the sampling process. Currently, however, the coastline is treated as a given during the sampling process.

On the sampling side we removed the final traces of noise from the sampled image by adding a final sampling step that, unlike the other sampling steps, does not add noise. With the inclusion of this final sampling step, the sampling process has two-time hyperparameters: the time value of the final noise adding sampling step, and the time value passed to the network for this extra sampling step. We performed a preliminary optimization of these two parameters, balancing sample speed, sample quality, and residual noise level.

Although there are several aspects that we aim to improve further, we will now first perform a few more realistic tests, automate this procedure and evaluate the performance more from a user perspective. The figures below show some test cases, as they are applied routinely during the training cycles. There are still some biases, but it is becoming harder to see which image is generated and which one measured. Note that these examples are small patches (128x128 cells) of the D5 tile of the 2023 release of the EMODnet gridded bathymetry.

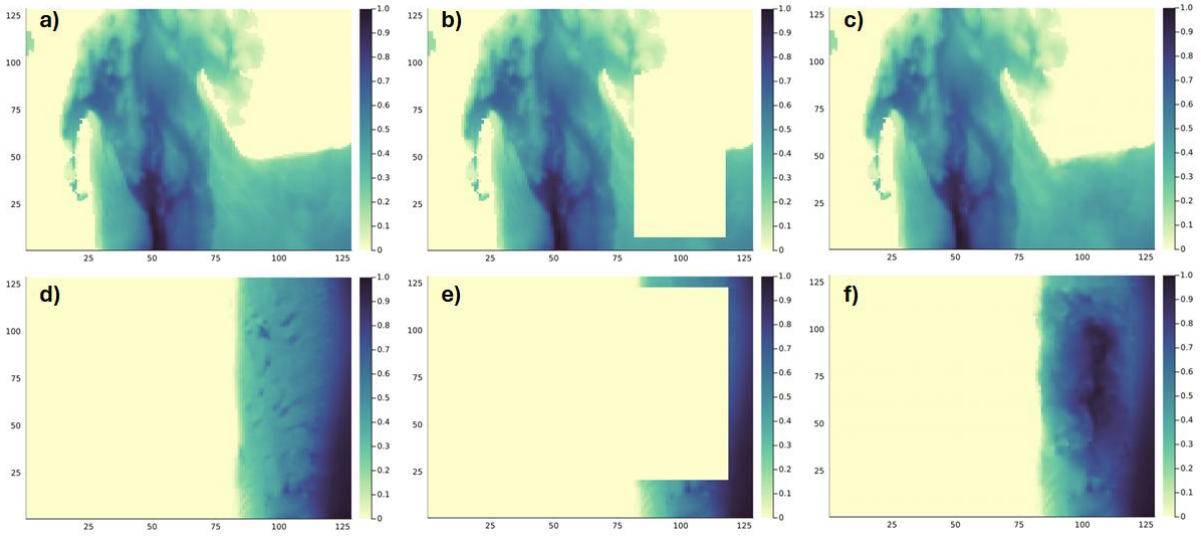

Figure 5 ***Examples of sampled bathymetry images*** *The original images (a,d) the applied masks indicating the areas to sample (b,e) and the final sampled image (c,f). Figures d-f show an example where the sampled bathymetry is further removed from the original. This highlights our intention to primarily generate realistic looking bathymetry rather than reconstructing the unknown bathymetry as accurately as possible.*

2.2 Satellite derived bathymetry using wave kinematics

In addition to the inpainting method, also a physics-based approach was applied, which estimates local water depths from wave kinematics, using satellite imagery. For this approach, the inter-band time shift between Sentinel2 color bands is used to estimate wave frequencies, ω (Bergsma et al., 2019). In combination with the spatial properties of the waves and specifically their wave number, k, a depth, d, can be estimated via the linear dispersion relationship:

$$\omega^2 = gk \tanh kd \,,$$

where g is the gravitational acceleration. For a group of frequency-wavenumber pairs, this equation can be solved for *d* using non-linear regression (Young et. al, 1985). Here, a standard Levenberg Marquardt regression is used. The pixel resolution of the RGB bands is Δx_{pixel} = 10 m, which means that the Nyquist limit for wave lengths is 20 m. In standard signal processing practice typically at least 2 wavelengths are desired for Fourier analysis, which translates to waves with lengths larger than 40m. For the specific case of analyzing Sentinel2 images Bergsma et al., 2019 suggest wave lengths of at least 70-80 m. Here, we set the limit to 40 m.

In this study the blue and the red band are used to derive wave frequencies. The time shift between these two bands is Δt = 1.005 s. The local phase-shift ϕ' between the blue and the red band, divided by the time shift, provides the local wave frequency, ω = ϕ'/Δt. Local phase-shifts ϕ' are computed through Gabor analysis. This involves the construction of a computational grid, of which each cell is analyzed using local two-dimensional Fourier transforms. These local transforms yield a wavenumber spectrum for the blue band and a wavenumber spectrum for the red band, and the difference between these two spectra gives the phase-shift ϕ' for each wavenumber k.

In this study a spacing of Δx_{grid} = 200 m is used between grid points, which means that the resolution of the computed bathymetry map is 200 m. Each grid point is analyzed using a grid cell of size 480x480 m^2, which represents a balance between staying local on the depth estimate while allowing sufficiently long waves to be captured in a grid cell.

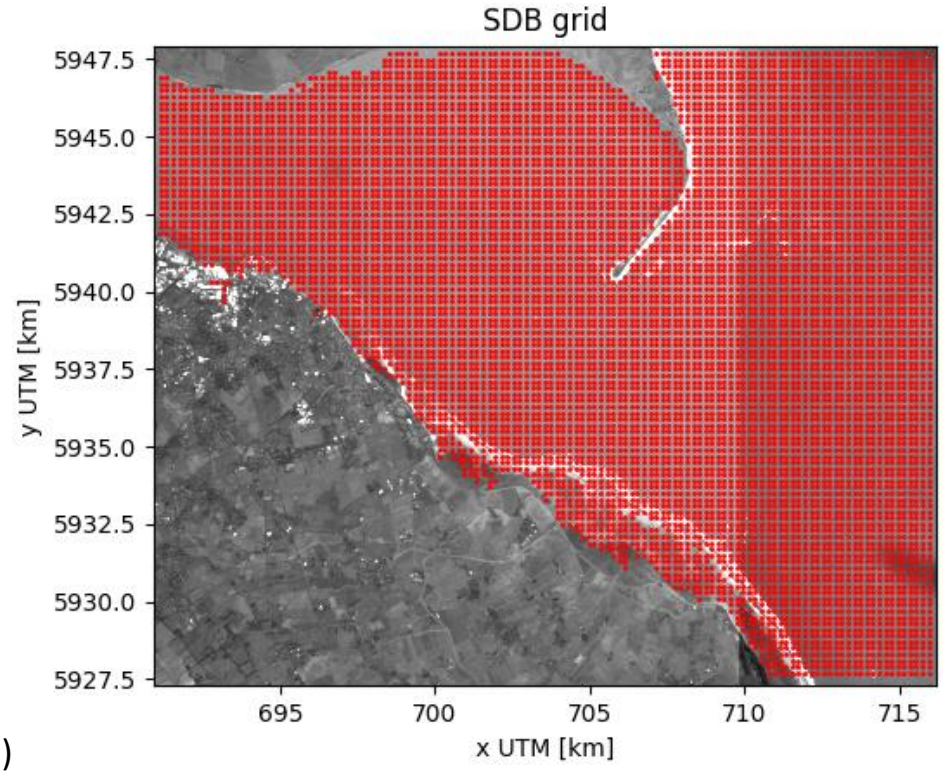

a)

b)

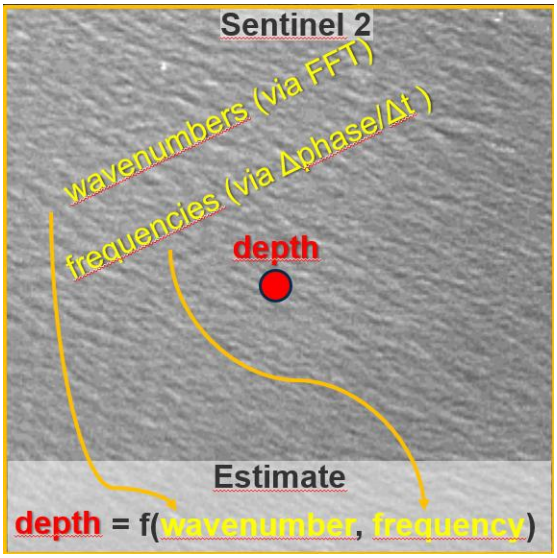

Figure 6 Gabor analysis for SDB from wave-kinematics. *a) A computational grid of points is used to b) locally determine wavenumbers and wave frequencies from the blue and red bands of a Senintel2 image. Using the linear dispersion relationship, these wavenumbers and wave frequencies yield a representative local depth estimate at each grid point.*

3 Use case: Humber

3.1 Status and goals

As a demonstration area of the deep learning methods, we choose the Humber estuary in the United Kingdom. This choice is motivated by the fact that large sections of the bathymetry data in that area are sourced from GEBCO (grey areas in Figure 8). The lower resolution of the source data makes these sections good candidates for reassessing the bathymetry. The areas are easily identifiable through the source index layer of the bathymetry data product, as they are already included in the release of the data product. This provides us with a mask indicating on a grid cell level which areas are to be reassessed. The goal of this use case is to demonstrate the capabilities of the deep learning methodologies and to provide a guide for further improvements.

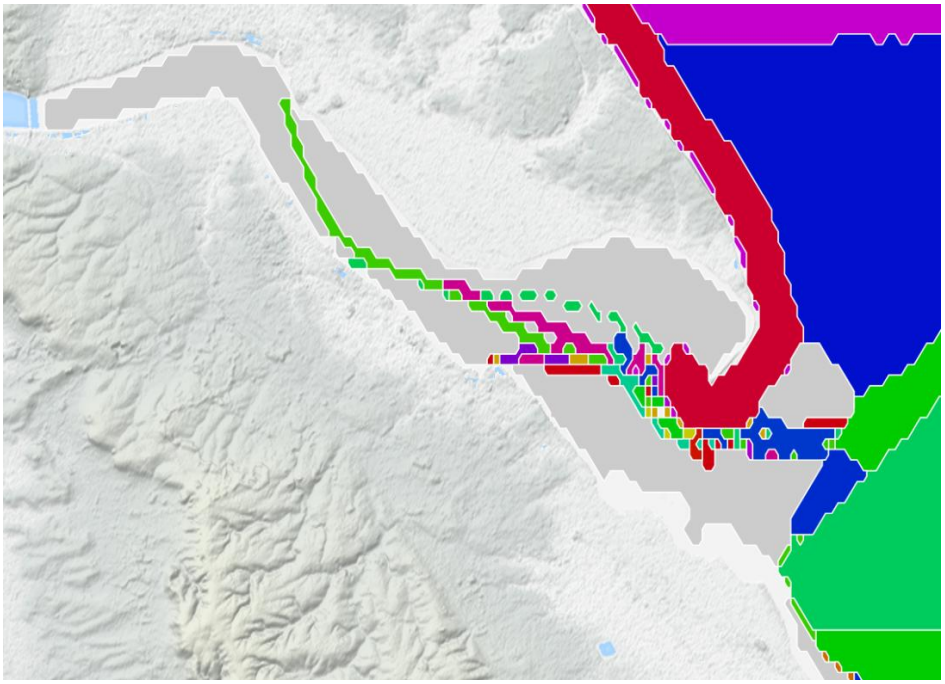

Figure 7: *Map of the bathymetry source index in the Humber estuary. The grey polygons indicate areas sourced from GEBCO. Screenshot taken from the online EMODnet viewer.*

3.2 Inpainting results

As an input, the inpainting network requires a sampling mask indicating which pixels are to be sampled and which pixels act as context for the sampling process. The sampling mask is a combination of a land-sea mask and a mask indicating GEBCO pixels, both of which are extracted from the source index layer. The mask is created by setting all pixels sourced from GEBCO to 1 (indicating it should be sampled), while all others are set to 0 (indicating it should act as context). The mask used is shown in Figure 8.

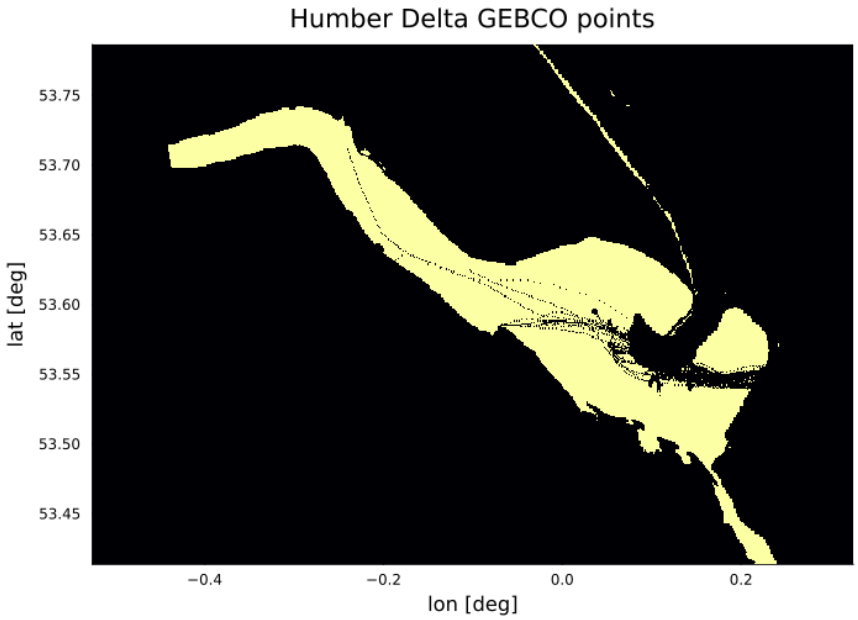

Figure 8 Sampling mask: *The mask indicating which pixels need to be sampled, derived from the bathymetry source layer.*

The network used for inpainting is trained on 128x128 pixel patches of bathymetry data. This sets the size of both the network input and output during sampling. The target area of the use case is

larger than this predetermined size, so it needs to be broken down into smaller 128x128 patches. We create these patches with a certain amount of overlap. The overlap aims to create a sense of continuity over the whole area as pixels sampled as part of one patch are then used as context information for the neighbouring patches (if the pixel lies within the overlap). How much neighbouring patches overlap is set beforehand and acts as another hyperparameter to the sampling process. Each patch is normalized individually to [0,1], after which the pixels indicated by the sampling mask are sampled. The patch is then rescaled to its original dynamic range and is inserted back into the overall bathymetry data. The sampling mask is updated to 0 everywhere each time a patch is sampled. This makes sure that each pixel is sampled only once (the first time that pixel is encountered when running over the patches) and acts as context for subsequent patches.

Next the order in which the patches are sampled needs to be set. We do this by first finding the patch with the least amount of GEBCO pixels (any patch with 0 GEBCO pixels is removed from the list). All pixels in this patch are set to 0 in the sampling mask and the number of GEBCO pixels for all other patches is recounted. We then move to the patch with the next lowest amount and skip any patch we already sampled when recounting the number of GEBCO pixels. This constant recounting ensures that at any moment during sampling we are always looking at the patch with the least number of unknown pixels, or in other words the patch with the most available context information. This way we aim to have maximum continuity or cohesion across the overall sampling area. A typical example of the sampling order is shown in Figure 9.

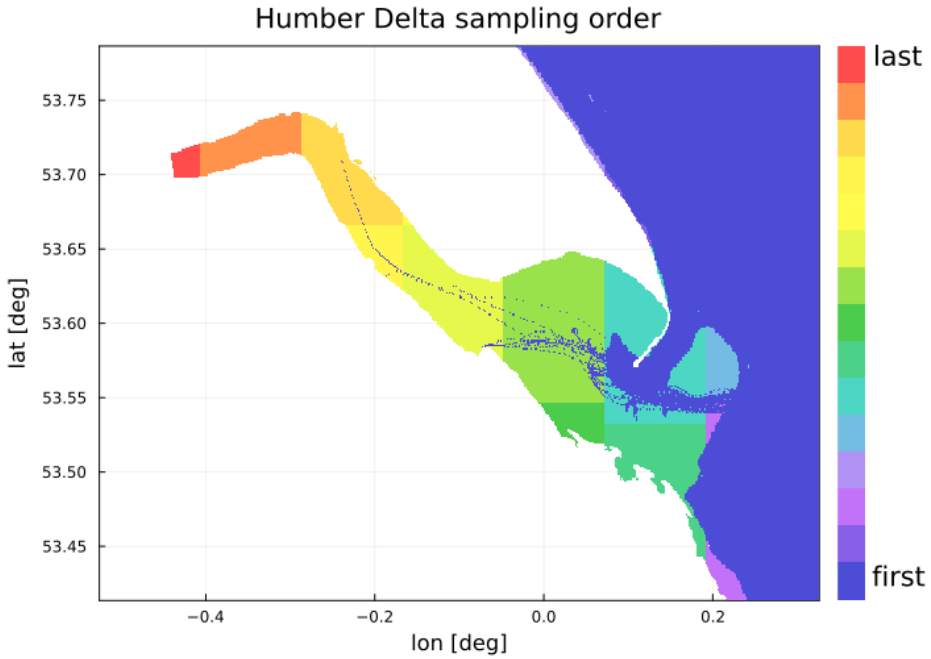

Figure 9 Patch sampling order: *The overall area is subdivided into smaller overlapping patches, after which the sampling order this determined through the number of GEBCO pixels per patch. An overlap of 14 pixels on either side was used for this order. The dark blue areas are not sampled at all since these have non-GEBCO bathymetry data.*

Data generation through diffusion allows us to easily generate multiple realizations by drawing new starting points of the sampling process which by design follow a Gaussian distribution. In Figure 10 we show 5 examples of sampled bathymetry for the Humber estuary. For these examples, the network was trained on 4000 images from the D5 tile of the bathymetry data, meaning that the Humber estuary is an out-of-sample area as it is part of the D4 tile. The 128x128 sampling

patches were created with an overlap of 14 pixels on either side, with the sampling order given in Figure 9.

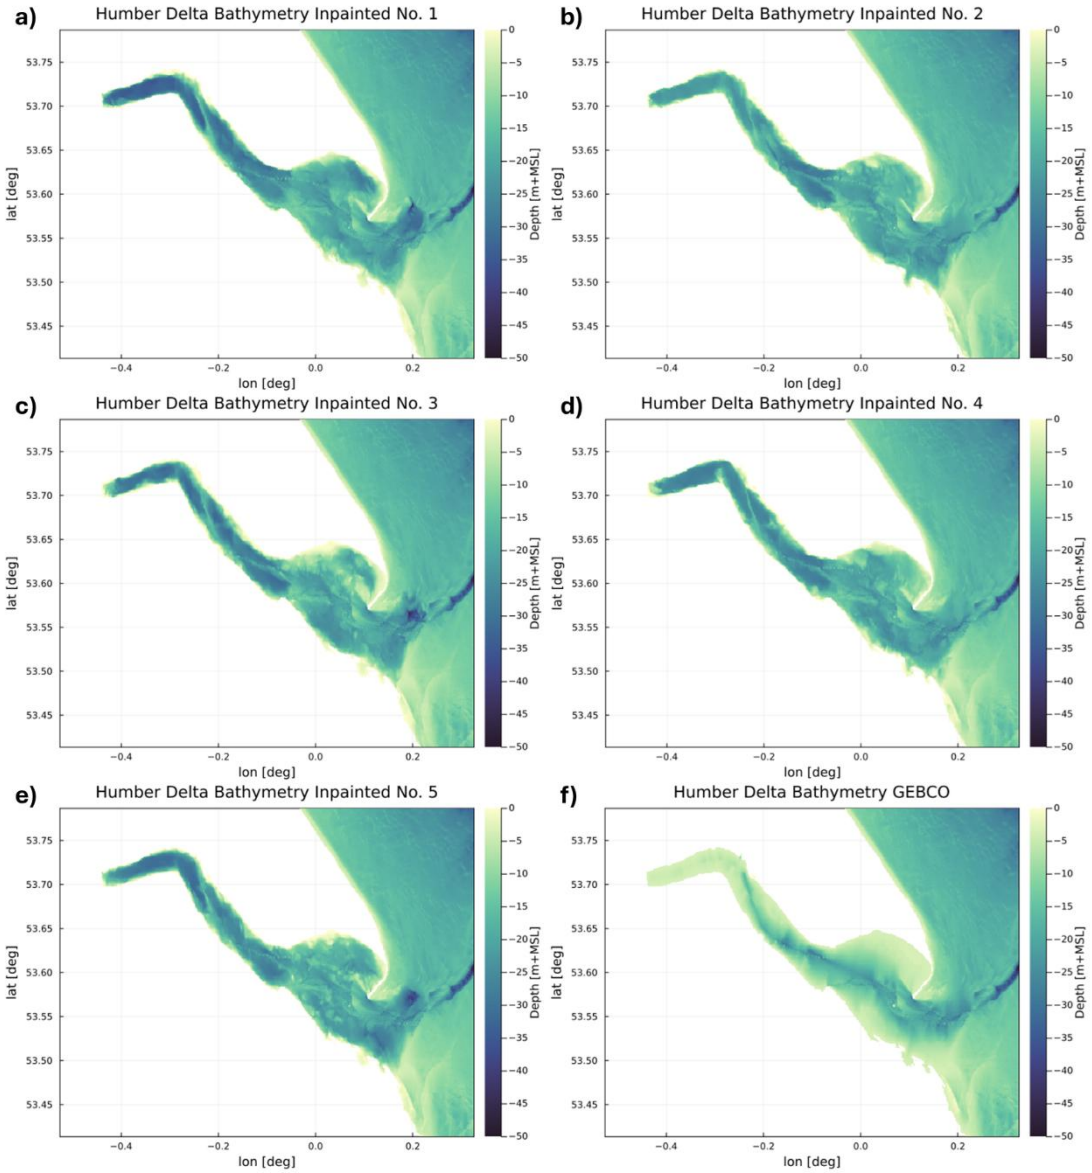

Figure 10 Sampled bathymetry: *Examples of sampled bathymetry data for the Humber estuary (a-e) with the bathymetry data including GEBCO in f).*

From the examples shown in Figure 10a-e, it is clear that the intent of generating bathymetry with a higher level of detail than the low resolution GEBCO data is achieved. The intent of generating more realistic looking bathymetry is not fully reached as certain features we expect to see in an estuary like the Humber are absent. For example, we expect a deeper channel across the entire length of the estuary, with potentially some branches and large tidal flats that have an increasing elevation towards the coast. This result is not unexpected as no additional information on estuary characteristics was provided during training. Furthermore, the training set was relatively small (4000 128x128 images from tile D5) meaning that the number of estuaries like the Humber in the training set is very low, if any at all. This points towards a need to train on a larger and more diverse dataset.

In addition, we did not filter the data based on its quality, but there are several indicators available for the EMODnet gridded bathymetry that potentially allow to give a larger weight to reliable data. Figure 11 shows an example with both the source index and the quality index, which in this example align with the artefacts observed in the bathymetry.

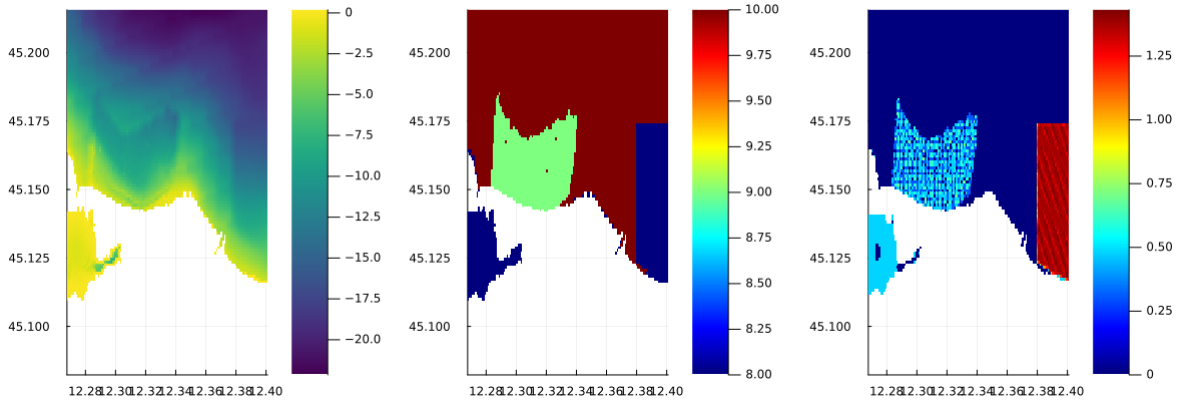

Figure 11 Sample bathymetry tile (left), with source (middle) and quality index (right). The transitions between the different sources are clearly visible. Note, that we searched for a clear example.

We also note that the tendency to sample the full dynamic range we reported on earlier is clearly visible in the eastern most parts of the estuary. These are areas where a large majority of the bathymetry data is sourced from GEBCO. These areas then are sampled with little context, the only context coming from the overlap with the previous patch, itself being largely resampled GEBCO pixels.

This lack of context is also linked to a related bias, the tendency for outliers or exaggerated depths. This tendency is most pronounced in patches with little context to reign in outliers, and in patches at the end of the sampling order. The reason for the latter is that outliers can propagate through the overall sampling area. Note that the value of a sampled pixel can be outside the strict [0,1] interval the patch is normalized to, implying that the dynamic range of the sampled patch is larger than that of the original patch. If such an outlier is generated in the overlap between two patches, the dynamic range of the other patch will as a result also be increased. This leads to a tendency for patches later in the sample order to have exaggerated dynamic range.

Finally, we note that in the eastern and middle parts of the estuary there are lines of bathymetry data not coming from GEBCO, instead coming from research vessel surveys. These single pixel lines act as context when sampling the surrounding bathymetry. In most generated samples these lines are clearly visible, albeit slightly broadened. This indicates that there was not enough context provided to properly incorporate these bathymetry data into the newly sampled data. As such there is a minimal context data volume needed to steer the model during sampling towards realistic bathymetry. Data below this threshold seem to be largely ignored, only broadened to avoid sharp discontinuities with the newly sampled data.

Next, we'll compare the AI generated bathymetry to some satellite derived products to gain more insight into the true nature of the bathymetry in the Humber Estuary.

3.3 Satellite Derived Bathymetry using wave kinematics

Initial tests using SDB from wave kinematics on the Humber case showed large differences in the estimated bathymetries between different years and time periods. The major causes for this were identified to be:

- A limited amount of usable Sentinel 2 images due to cloudiness

- Limited penetration of significant wave signals through the Humber entrance
- The large tidal range

Since wave penetration into the estuary is limited, an area of ~20 km^2 around the lower estuary was chosen for analysis, where some wave footprint was deemed to be present.

A pragmatic approach to alleviate these issues simultaneously was to generate an average bathymetry for the complete period 2017-2024 for which Sentinel 2 images were available. It meant that all usable data in terms of cloudiness and significant wave footprint could incorporated, and it aimed to average out tidal fluctuations on the depth results.

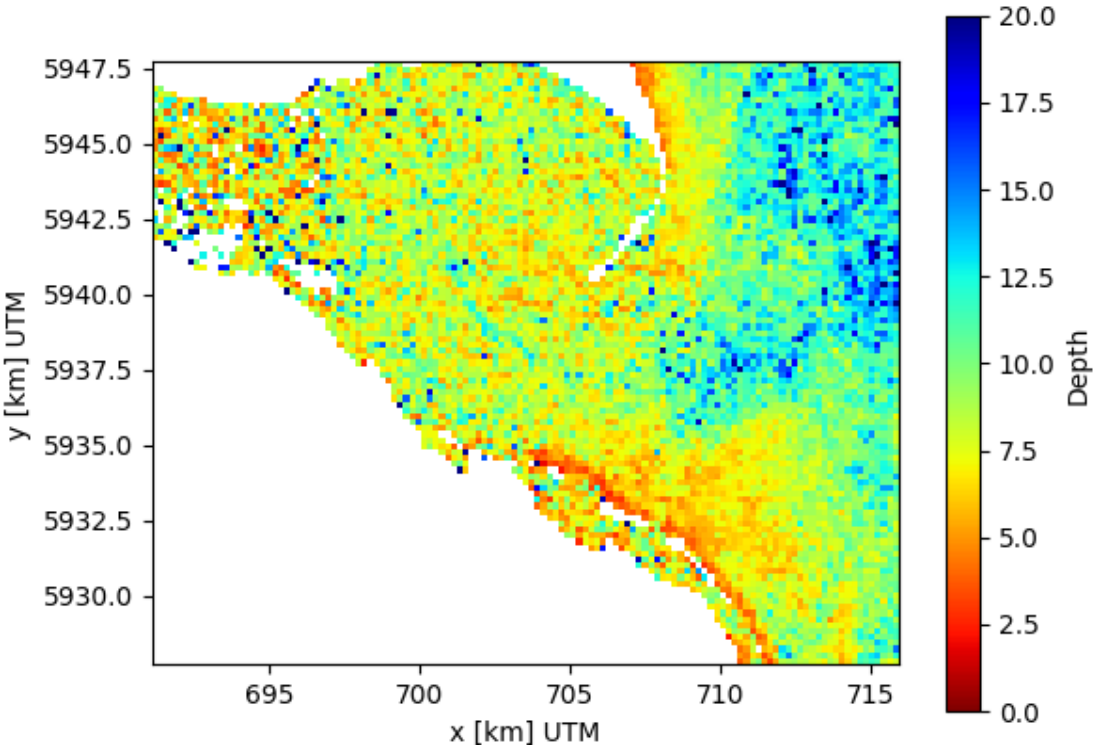

Figure 12 Humber average SDB 2017-2024

Visual comparison with a Humber nautical chart (ABP Humber Estuary Services, 2024), suggests reasonable agreement in the area offshore of the harbour dam “Spurn Head”. The Characteristic shoals offshore of Spurn Head, but also the large shoal in the south is captured. The official elevations of these shoals approximately range between -0 to -6m. The SDB results appear to overestimate these depths to some degree as they range more between 4-7 m, while the overall feature shapes are captured. The channel elevations typically range from -10 to -20 m depth. This is generally captured by the SDB, albeit a direct comparison was not feasible in this initial study.

Bathymetry estimates on the inside of the entrance at Spurn Head are noisy and suggest a shallow estuary. This is partly correct, yet the deep channel is not captured, probably due to insufficient wave signal and the intertidal flats are overestimated in depth. Hence, while the inside regions of Humber are hard to estimate from wave properties, the SDB from wave kinematics shows potential

for estimating bathymetry in data scarce regions on the ocean side of Spurn Head.

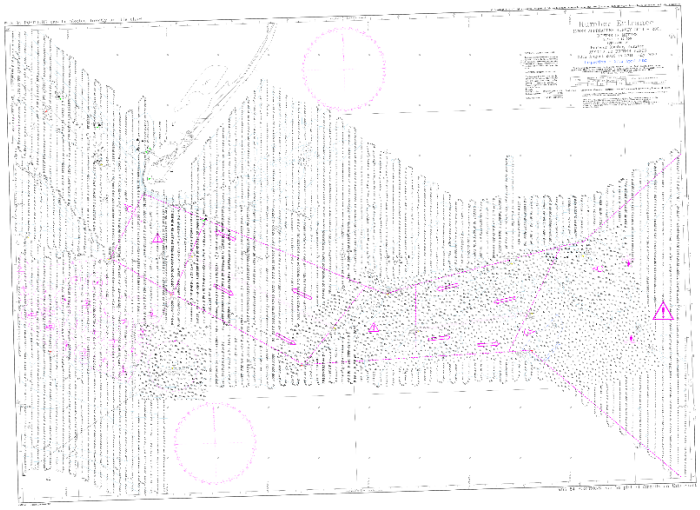

Figure 13 Nautical chart of Humber estuary inlet (for detail see: ABP Humber Estuary Services, 2024)

3.4 Satellite Derived Bathymetry from land-water detection

In the shallow inter-tidal parts of the estuary, one can also use the Water index (NDWI) derived from a large stack of Sentinel-2 images to assess the bathymetry. Figure 14 clearly shows the small channels and increasing slope towards the coast. The highest parts of the area (in dark blue) only become wet a small fraction of the time. A comparison with the data from a nearby tide-gauge at Immingham, shows that the upper slope is likely to be around 3 meters above mean sea level. The upper slope in the EMODnet 2023 release was only just above mean sea level. The Highest Astronomical Tide (HAT) values for this region are around 3.5 meter above mean sea level and could here provide a proxy for the heights near the coastline.

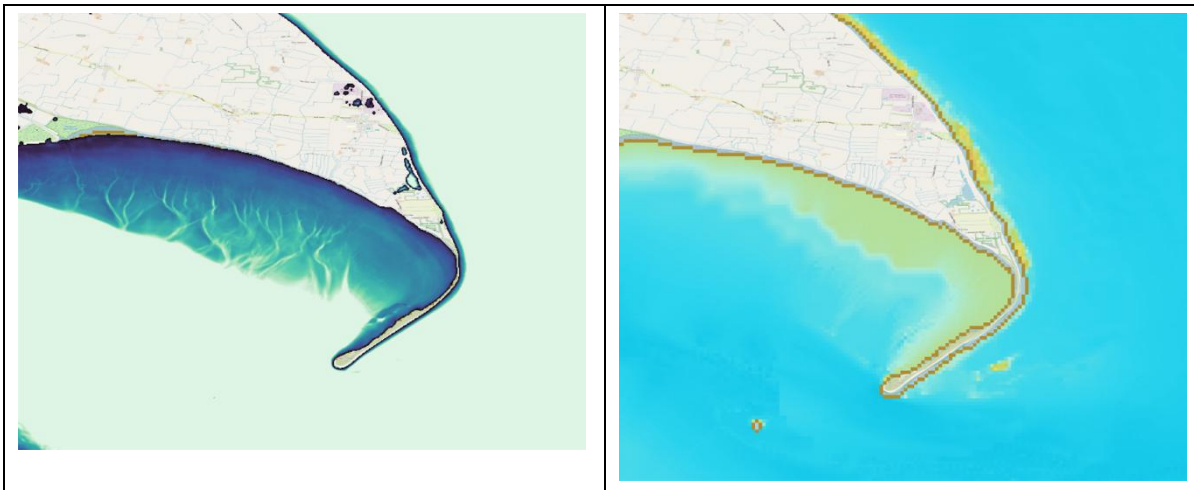

Figure 14 Left panel shows Water index and right panel the EMODnet gridded bathymetry for the northeastern part of the Humber Estuary.

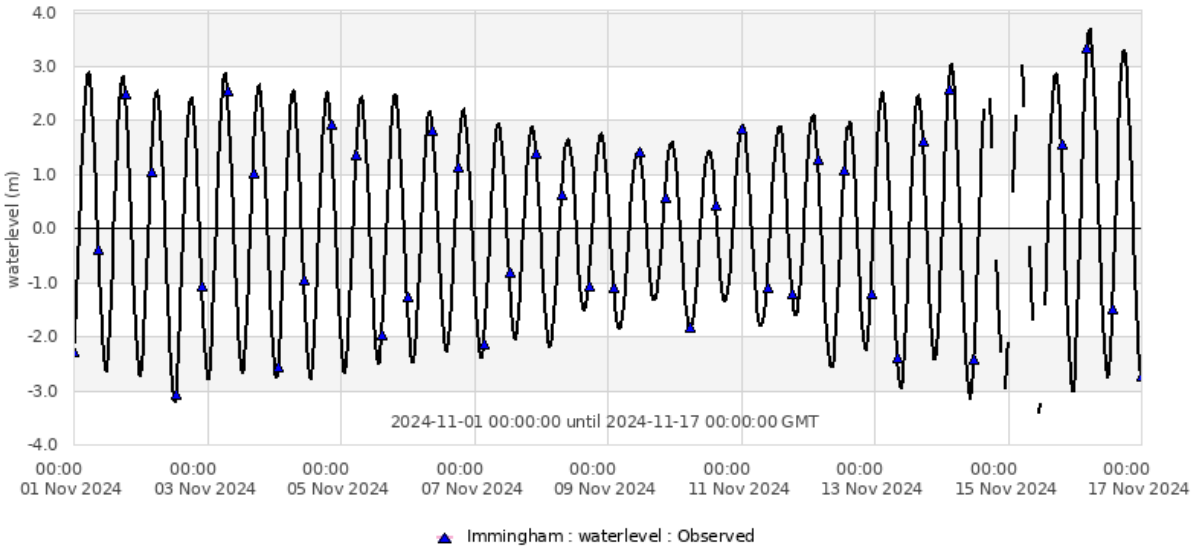

Figure 15 measured water-level at Immingham

4 Conclusions and recommendations

4.1 Bathymetry inpainting

In this and previous reports we outlined a methodology for generating bathymetry data using score-matching diffusion machine learning models. We have outlined a strategy for training these models on the EMODnet bathymetry data, how to sample new data from a trained model, and how to apply this to larger areas where reassessment of the bathymetry data is desirable. Following this outline we have implemented a pilot, trained on a small subset of the EMODnet bathymetry data and applied it to the use case of the Humber estuary. The results from this use case provide an initial indication of performance, pitfalls, and directions for improvement. Overall, the methodology is promising, creating samples at a higher level of detail than a best guess approach. However, the level of realism currently has room for improvement, and the generated bathymetry shows some clear biases. Based on the result illustrated in this report we provide a series of recommendations for improvement, roughly split into two categories: data requirements and technical/architectural requirements.

Data requirements:

- Train the model on the full EMODnet bathymetry dataset. So far, the model has only been trained on a subset of 4000 128x128 images from the D5 tile of the bathymetry data. This relatively low data volume yields a quicker iteration speed when developing the model, but also raises issues with representation of the use case area in the training data.
- Data augmentation techniques. The pilot model was trained using data augmentation of rotations over multiples of 90 degrees of the images. However, since the data is provided in a geographical coordinate system, this method of data augmentation is not entirely valid. Alternate methods of data augmentation that respect the geographical coordinate system are a better fit for this goal.
- Additional data sources. We have seen that there appears to be a minimal amount of context data needed to properly guide the model during sampling. In areas where there is not enough non-GEBCO bathymetry data present in the EMODnet dataset, additional sources of context are needed. These could be bathymetry estimates from other

techniques, or for example data on estuary characteristics. The latter class, i.e. context data that is not bathymetry, will have to be provided during training as data labels.

Technical requirements:

- Network architecture. There is considerable freedom when designing a network architecture for score-based diffusion models, one that has so far largely been unutilized is this pilot. Especially when the training data increases, either in volume by using the entirety of EMODnet bathymetry, or in data type by providing data labels, the network architecture needs to change accordingly to properly utilize the new data. Additionally, it is worth exploring architectures that are more flexible regarding input size as an alternative to the patch-based approach for sampling larger areas used here.
- Sampling strategies. Currently the sampling process for an individual patch uses relatively simple approach as far as sampling diffusion models goes. There has been a lot of research into both increasing the sampling speed and the final sample quality for diffusion models. It is worthwhile investigating the benefits of more state-of-the-art sampling approaches. Additionally, a for the case of the network architecture, if the dataset is expanded to include data labels, the sampling process will need to adapt accordingly.

4.2 Satellite wave kinematics

In this report a methodology was described to estimate bathymetry from wave kinematics observed in Sentinel2 satellite imagery. The method involves Gabor analysis of the blue and red bands in the Sentinel2 imagery. An inherent time shift of ~1s between these bands allows estimation of wave frequencies, which together with estimated wavenumbers yield an estimate of local depths, as per the linear dispersion relationship of freely propagating ocean waves. On the ocean side of Humber estuary, the results correctly suggest the presence of bathymetric features like shallow shoals and the deep channel.

In this initial analysis, depth errors have not been directly quantified, yet comparison with existing navigational charts suggests depth estimates of shoals to be slightly overestimated by ~2-3 m, while estimated channel depths cover the correct depth range of approximately 10-20m. On the inside of the estuary, depth estimates are noisier and flatter, losing the ability to distinguish intertidal flats from the deep channel. Major reasons for this are limited wave penetration into the Humber estuary and the large tidal range, whose effects have not been accounted for on case-by-case basis.

Recommendations for future improvement are:

- Incorporation of tidal effects
- Higher time resolution of SDB bathymetry to estimate also morphodynamics
- Improvement of image preselection based on wave footprint

4.3 Satellite derived intertidal bathymetry and coastline

The large tidal range in the Humber Estuary results in tidal flats and coastal slopes that could be estimated using satellite derived intertidal bathymetry. This would provide additional input for the AI model to generate a more plausible solution also for the deeper parts.
At a larger scale, one can use the high-water level at the coast from a tide model as an indicator of the elevation at the coast to guide normal interpolation. Since this procedure is probably simple to implement, it can serve as a temporary solution, until better ones are ready for production.

A. References

- **ABP Humber Estuary Services, 2024**. https://www.humber.com/Estuary_Information/Marine_Information/Chart_Catalogue/Current_Humber_Charts/, file: Humber Inner Approaches Annual 2016-2021 with Haile Sand inspection - Surveyed 27th April 2023, accessed on 16.12.2024.
- **Bergsma2019** Bergsma, E. W., Almar, R., & Maisongrande, P. (2019). Radon-augmented sentinel-2 satellite imagery to derive wave-patterns and regional bathymetry. Remote Sensing, 11(16), 1918.
- **Song2020** Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., & Poole, B. (2020). Score-based generative modeling through stochastic differential equations. arXiv preprint arXiv:2011.13456.
- **Young1985** Young, I. R., Rosenthal, W., & Ziemer, F. (1985). A three-dimensional analysis of marine radar images for the determination of ocean wave directionality and surface currents. Journal of Geophysical Research: Oceans, 90(C1), 1049-1059.