Hybrid Approaches for Spatial Data Interpolation

Marcos Domingues Barata

Abstract

Spatial data interpolation is a problem that crosses a panoply of natural and social sciences. The goal is to estimate a surface based on a sample of locations. Traditional spatial interpolation techniques are based on mathematical approaches and, in some cases, combined with geographic concepts. The growth of computational power consequently resulted in the evolution of machine learning approaches, especially neural networks. One of the most relevant steps in terms of these techniques was the development of neural networks capable of processing images and capturing features that characterize them. When applied to the spatial interpolation problem, neural networks, like the multi-layered perceptron, can equalize the performance of more traditional techniques, but the neural networks based on image processing and relying on convolutional layers can overcome the traditional techniques. Traditional techniques can be applied in every context, while neural networks face the problem of generalization where a network trained within a particular context can underperform when put in a different context. This paper presents a hybrid approach. Based on a generative adversarial network, where we combine a neural network with a traditional spatial interpolation approach to overcome the traditional approaches and to achieve a higher generalization degree. This study shows that our approach is a viable option to solve the spatial interpolation problem since the results of interpolating the digital elevation model of mainland Portugal and the islands of the Portuguese autonomous regions evidence a solid performance and generalization ability.

1 Introduction

Spatial Data Interpolation is a problem that covers many subjects and fields going from the natural to the social phenomena. This problem consists in the estimation of new spatial information from the already known information, and can be formally defined by the following d-variate function that satisfies the condition for each N known points:

$$F(r_j) = z_j \; , \; j = 1, \ldots, N \tag{1}$$

where z_j, is the observed value at the r_j discrete point.

As can be seen, an infinite number of functions can fulfil this constraint, so additional conditions have to be added, differentiating the interpolation techniques. Traditional techniques restrain the solution based on a local neighbourhood approach (e.g., Inverse Distance Weighting - IDW), geostatistical approach (e.g., Simple Kriging) or a varionational approach (e.g., Thin Plate Spline). Although these methods are easy to understand/implement and very adaptable to all contexts, they are limited when it comes to capture complex geospatial features and contexts, affecting negatively their accuracy. Oppose to traditional techniques, modern approaches based on machine learning condition the solution by going through a set of input-output pairs in order to capture the geospatial contexts of the available data. This solutions started by having the potential to match the traditional approaches performance, but nowadays have proved to overcome them by relying in advanced image processing techniques like convolutional neural networks. Although neural networks can overcome traditional techniques in terms of performance, they face the disadvantage of generalization since a network trained within a specific context can underperform when put in a different context. This limitation contrasts with the traditional ones since these techniques do not rely on a learning phase over input data to capture the spatial contexts. Summarizing, we have a modern technique that can overcome the traditional techniques, but since it is data dependent it faces a high probability of failure when applied to different contexts than the ones it was trained in.

In this paper, we report on experiments with an hybrid spatial interpolation technique that combines the ideas of a traditional spatial interpolation technique and a modern approach based on image processing neural networks applied to the digital elevation models of mainland Portugal and the islands that compose the Portuguese autonomous regions. Our objective was to create a hybrid approach that takes advantage of the neural networks capability to capture deep spatial features and to achieve a higher generalization degree when combined with traditional techniques since they can be applied in every context. The traditional technique we selected is the IDW method because it is very easy to comprehend/implement and is still one of the most utilized techniques to solve the spatial interpolation problem. The modern approach we followed was proposed by Zhu et al. (2019) and consists in a generative adversarial network.

The rest of this paper is organized as follows: Section 2 presents fundamental concepts and important related work. Section 3 describes the considered hybrid spatial interpolation approach. Section 4 details the results and major observations. Finally, Section 5 presents the main conclusions, and highlights possible directions for future work.

2 Fundamental Concepts and Related Work

In this section we approach the spatial interpolation solving approaches that constitute the basis of our study. In addition to exploiting the approaches we also focus on work related to those approaches and their main conclusions.

2.1 Inverse Distance Weighting

One of the simplest methods to interpolate data is IDW, which is the mathematical representation of Tobler's first law of Geography, "Everything is related to everything else, but near things are more related than distant things." Tobler (1970), given by:

$$V_j = \frac{\sum_{i=1}^{n} v_i \frac{1}{(d_i)^p}}{\sum_{i=1}^{n} \frac{1}{(d_i)^p}} \tag{2}$$

Where the estimation of the value V_j in the position j, takes into account all the n known values (v_i) and the distance (d_i) between the known point i and unknown point j, giving more weight to closest values, as we can see in the following equation. The value p represents the power parameter and it's responsible to control the impact of known points on the interpolated values based on their distance. This value is a positive real number with the default of 2.

Recent studies on IDW (Greenberg et al. (2011); Stachelek and Madden (2015)) proved that we can add additional information to the IDW model, by exchanging the euclidean distance for a non-euclidean distance obtained by traversing a search tree for source i to sink j. This informed search allows us to add useful knowledge to the model by establishing barriers and highlighting known relations between points. The use of a non-euclidean distance has proved better than the classic euclidean distance, but only when the distance is defined properly and in problems where the presence of barriers is necessary.

2.2 Generative Adversarial Networks

Generative Adversarial Neural Networks (GAN) were introduced by Goodfellow et al. (2014) revolutionizing the neural networks training paradigm with a zero-sum game between its two modules. These networks are mainly used in image generation and manipulation but its applications can go from simulations to classification. These networks are composed by two sub models: a Generator, G, and a Discriminator, D. The Generator goal is to capture the data distribution and generate outputs as close as possible to the real data distribution. The Discriminator function is to estimate the probability of its input being generated data or real data. To learn the Generator distribution p_g over data x, a prior on input noise variables $p_z(z)$ is defined and then a mapping to data space is represented as $G(z;\theta_g)$, where G is a differentiable function represented by a multilayer perceptron with parameters θ_g. The second sub model is again defined by a multilayer perceptron $D(x;\theta_d)$ that outputs the probability that x comes from the data rather than p_g. The discriminator is trained to maximize the probability of assigning the correct label (from the data or from the Generator), while the generator is trained to minimize the difference between real data distribution and the generator data distribution $(log(1 - D(G(z))))$. The minmax adversarial game that these models play is represented by the following equation defined by Goodfellow et al. (2014):

$$\min_{\theta_g} \max_{\theta_d} V(G, D) = \mathbf{E}_{x\ p_{data}(x)}[log(D(x))] + \mathbf{E}_{z\ p_z(z)}[1 - log(D(G(z)))] \tag{3}$$

The training of the network can be achieved by simultaneously updating the weights (θ_d,θ_g). The update of the Discriminator weights can be conducted by ascending its stochastic gradient of the loss function, i.e.

$$\nabla_{\theta d} \frac{1}{n} \sum_{i=1}^{n} [logD(x^{(i)}) + log(1 - D(G(z^{(i)})))] \tag{4}$$

and the update of the Generator weights can be conducted by descending its stochastic gradient of the loss function, i.e.

$$\nabla_{\theta g} \frac{1}{n} \sum_{i=1}^{n} log(1 - D(G(z^{(i)}))) \tag{5}$$

where n is the number of samples in each data batch during training.

Since the spatial interpolation problem relies on sampled data in order to interpolate the solution, the noise vector z defined above needs to be replaced by the sampled data, which the generator takes as input to interpolate. Having that, we need to extend the GAN model to a conditional version named Conditional GANs (CGANs), since the generator and discriminator are both conditioned by the same auxiliary information, which can restrict both generation and discrimination processes.

For spatial interpolation problems, the traditional adversarial strategy needs to be modified to ensure the stability of conditional generations. The random noise vector $\mathbf{z}$ should be removed such that the conditional generation could be considered to be determined by the sampled data as the only constraint. When training an adversarial spatial interpolation net, the generator requires the sampled image $f(x)$ as input in order to output a generated fake image $G(f(x))$ as similar as possible to the real image x. Adding to that, the discriminator needs to be trained to distinguish a fake image $G(f(x))$ from real image x based on the sampled image f(x). The min max game is then altered to:

$$\min_{\theta_g} \max_{\theta_d} V(G, D) = \mathbf{E}_{x\ p_{data}(x|f(x))}[log(D(x, f(x)))] + \mathbf{E}_{z\ p_{data}(x|f(x))}[1 - log(D(G(f(x)), f(x)))] \tag{6}$$

Where G is a differentiable equation representing the Generator's structure and it's parameters θ_g and D is a differentiable equation representing the Discriminator's structure and it's parameters θ_d. G attempts to approximate $p_g(G(f(x))|f(x))$ to $p_{data}(x|f(x))$ in the real dataset, minimizing the second term of the equation above. At the same time, D evaluates if the spatial image come from $p_g(G(f(x))|f(x))$ or $p_{data}(x|f(x))$, maximizing both terms of the equation. These changes make the adversarial spatial interpolation learning to approximate the conditional generative probability distribution given a set of spatial sampled images $p_g(x|f(x))$ rather than the probability distribution of real existing data(p_{data}).

2.3 Neural Networks and the Spatial Interpolation Problem

Neural networks evolved from the simplest unit, the perceptron, to much more complex models like the GANs. In 2001, Rigol et al. (2001), demonstrated the potential of neural networks by applying a multilayered perceptron to interpolate air temperatures in the United Kingdom. This study proved that neural networks could match and even surpass the performance of more traditional techniques like kriging and IDW. More recently Zhu et al. (2019) used GANs to overcome the traditional approaches fulfilling the prediction Rigol et al. (2001) made 18 years before that neural networks would surpass the traditional techniques. This achievement was possible because of the evolution in neural network models that are now able to process images and capture their features by relying on convolutional layers. When transposed to the spatial interpolation problem this ability proves very useful, since deep spatial features and patterns can be captured, counteracting one of the main disadvantages of traditional techniques. Despite being able to overcome the traditional techniques the work done by Zhu et al. (2019) also focused on the generalization problem where a network trained in a specific context underperformed when put into a different spatial context.

Marcos et al. (2018) and Deng et al. (2019) presented two works focusing on the generalization problem. Marcos et al. (2018) proposed Rotation Equivariant Vector Field Networks (RotEqNet) Modules into the CNN architecture. RotEqNet involves rotating CNN filters and pooling across orientation space, retrieving maximal activations per filter, hence train the network to capture the same features but in different positions, which allows the network to be more robust to different contexts. Deng et al. (2019) came up with Getis-Ord $G{*}_i$ pooling, which is a pooling method based on spatial Getis-Ord $G{*}_i$ analysis of CNN feature maps. Getis Ord $G{*}_i$ analysis is a technique for geo-spatial clustering that is used to encapsulate the fundamental rules of geography, especially the first law of geography. While incorporating the IDW principles into the pooling phase of a convolutional layer, Deng et al. (2019) were able to generalize the capturing of spatial features providing a higher generalization degree for the network.

3 A Hybrid Approach for Spatial Data Interpolation

In the previous Related Work section we exposed the IDW technique and how variations in the calculation of the distance parameter in equation 2 can be used in certain cases to overcome or attenuate some of the IDW disadvantages when using euclidean distance. In a more modern approach, we detailed the potential neural networks could have in solving the spatial interpolation problem. The potential provided by the multi-layered perceptron in solving the spatial interpolation problem was confirmed when new neural networks were able to capture complex features and patterns in images, allowing the models to apprehend and generate/interpolate new images conditioned by sampled information. Machine Learning went from being as good as traditional techniques, to overcome them. The last subsection of the Related Work focuses on the challenge of generalization of neural network used to interpolate spatial data and to maintain its performance in contexts different from the ones the network was trained in.

Our approach is based in the neural network proposed by Zhu et al. (2019). Like Zhu et al. (2019), we propose a conditional encoder decoder GAN combined with the IDW technique. This hybrid approach has the goal of overcoming the IDW technique and achieve a good generalization degree. The network of our approach is divided into two sub-models, the Generator and the discriminator, which are detailed in the subsections below.

3.1 The Proposed Generator

The generator can be divided into an encoder and a decoder. The encoder takes the sampled patch as input and forwards it to three two-dimensional convolutional layers (conv 1, 2 and 3). The decoder part has three two-dimensional transposed convolutional layers (deconv 1, 2 and 3), that upsample the previously encoded feature maps. After each of the convolutional or transposed convolutional layers a ReLU activation function is applied, as we can observe in Figure 1.

The settings used for each encoder layer are the same ones used by each decoder layer, which means they share the same convolving 5x5 kernel and stride, valued as 1, sizes and perform a zero-padding with the given kernel and stride.

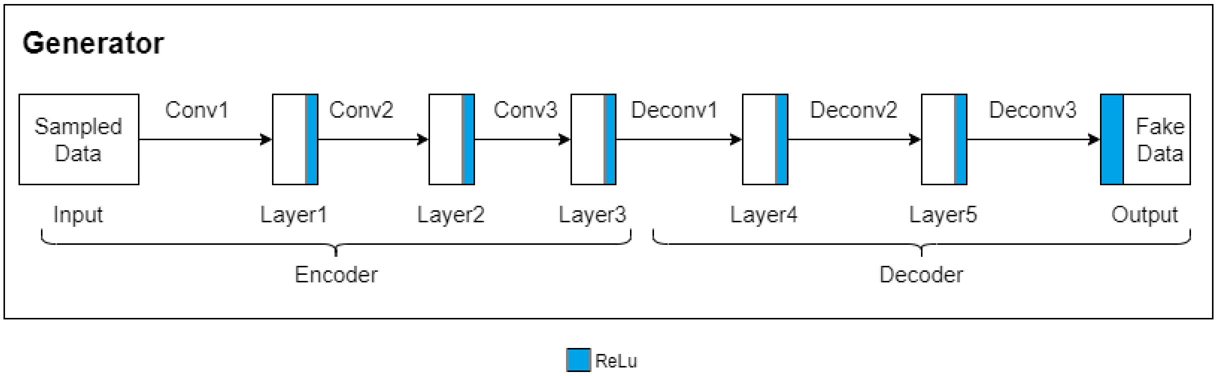

Figure 1: Architecture of the generator sub-model.

3.2 The Proposed Discriminator

The discriminator takes two inputs, the sampled patch and the corresponding real patch or the patch generated from the corresponding sampled patch. Each input is then forwarded into an independent two-dimensional convolutional layer (conv 1.1 and 1.2). The feature maps that outcomes the convolutional layers (conv 1.1 and 1.2) are merged using a concat operation. The outcome of the concat operation is then passed through three two-dimensional convolutional layers (conv 2, 3 and 4). After each one of the convolutional layers the ReLU activation function is applied, except in the last convolutional layer, where a Sigmoid activation function is applied, in order to obtain a scalar from 0 to 1 indicating that values closer to 1 represent a strong belief on the discriminator that the patch is real and values closer to 0 express the discriminator belief that the patch is fake/interpolated. The discriminator sub-model, can be represented as Figure 2 demonstrates.

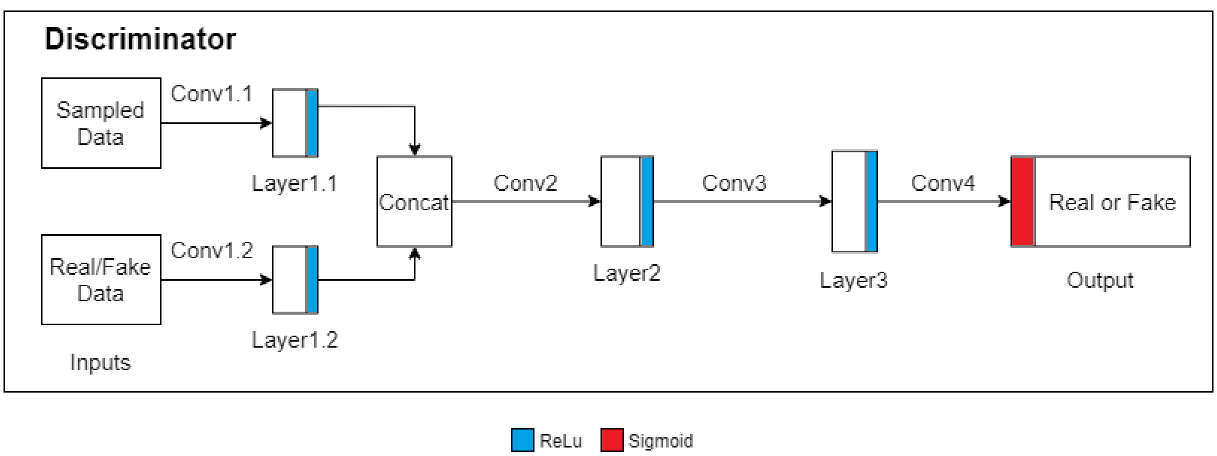

Figure 2: Architecture of the discriminator sub-model.

The layers 1.1, 1.2, 2 and 3 share the same kernel and stride sizes and perform a zero-padding convolution with the convolving 4x4 kernels and a stride of 2. The last layer, 4, does not pad the feature maps derived from the layer 3 and has a stride, different from the previous layers, valued as 1.

3.3 The proposed Hybrid Algorithm

After defining the network structure we present, in this subsection, the proposed hybrid algorithm. As mentioned before this technique combines the network we just described with IDW. This combination is achieved by replacing the one-channel input of the generator and the corresponding input of the discriminator by a two-channel input composed by the sampled patch and the corresponding patch interpolated using IDW. Despite the interpolated patch on the second channel can be labeled as inaccurate information we believe it will improve the network ability to capture the spatial contexts, since some of the spatial features can be found in the interpolation by IDW, and the network can then learn new ones and improve the capture of the spatial features exposed in the patch interpolated by IDW. With this combination we provide even more information, hoping to overcome the IDW technique and the network that requires only a one-channel input. Besides overcoming other techniques we also focus the training behavior of the network and its generalization ability.

3.4 Sources of Ancillary Data

From all the innumerable variables that can be represented in a spatial context, we selected the elevation of the terrain as a variable to be displayed in space. This choice was due to the amount of available data, since the land in all of the globe is available with a high resolution (25 meters)[1] and due to its appearance in the literature related with spatial data interpolation.

Since we live in Portugal, we decided to focus on the digital elevation model (DEM) of mainland Portugal and its autonomous regions, Azores and Madeira. Being mainland Portugal located in Europe and its autonomous regions placed in the Atlantic North we relied on a DEM from the European Union's Earth observation programme[2].

After obtaining the rasters corresponding to the Portugal mainland and islands, we merged them using gdal-merge[3]. To split the merged raster into the mainland Portugal and each individual island that composes the archipelagos of Azores and Madeira, without considering the islands of Desertas , Selvagens and small islanders, we used gdalwarp[4] with the shapefiles (i.e., a popular format for geographic information system information) of their boundries provided by Direção-Geral do Território[5].

Having extracted and processed the data, we were left with a dataset of twelve digital elevation models corresponding to each island of the Portuguese autonomous regions (nine Azores islands and two Madeira Islands) and to mainland Portugal.

4 Experimental evaluation

In this section, we approach several aspects of our experimental evaluation. Firstly, we explain the evaluation metrics used to measure and compare the performance of our hybrid approach with the other approaches. Secondly, we describe the sampling process and IDW parameters we selected for this study. Lastly we detail the training procedure of the networks and the overall results, summing up the main observations and conclusions.

4.1 Methodology and Evaluation Metrics

Spatial Interpolation is an estimation of the unknown values based on a sample of known values. Like in every estimation the goal is to get as close as possible to the real values. To evaluate how close the interpolated values are to the real ones, we selected two metrics that are widely used in the literature regarding this kind of problems. The root mean squared error (RMSE) and the mean absolute error (MAE) with their mathematical representation as follows:

$$RMSE = \sqrt{\frac{\sum_{i=1}^{n}(\hat{y}_i - y_i)^2}{n}} \tag{7}$$

[1] `https://earthexplorer.usgs.gov/`
[2] `https://land.copernicus.eu/imagery-in-situ/eu-dem/eu-dem-v1.1`
[3] `https://gdal.org/programs/gdal_merge.html`
[4] `https://gdal.org/programs/gdalwarp.html`
[5] `https://www.dgterritorio.gov.pt/dados-abertos`

$$MAE = \frac{\sum_{i=1}^{n} |\hat{y}_i - y_i|}{n} \tag{8}$$

The equations described above take the n number of predictions and calculate the average of the absolute difference or squared average of the squared difference between the $\hat{y}_i$ representing a predicted value and y_i corresponding to a true value. As we can observe, the equations 7 and 8 emphasize different aspects of the model performance. While the MAE assigns the same weight to all errors, the RMSE penalizes variance by giving errors with bigger absolute values more weight when compared to smaller absolute values. In order to obtain a better analysis, we, like in various studies regarding this problem, selected this two metrics since they can provide a varied picture of the error distribution.

4.2 Sampling Process and Configurations

In order to study the behaviour of the IDW technique, the GAN and our proposed hybrid approach, we sampled the DEM referring to mainland Portugal and the Portuguese autonomous regions using three different sampling configurations. These configurations consider only the points with values above the water level (i.e., height ≥ 0 meters) and distribute the known points uniformly, where the percentage of known points relative to the total of points in the first configuration is 5%, in the second configuration is 10% and 15% in the third configuration. Like in the work of Zhu et al. (2019), the network take as input patches with a size of pixels 32x32. For practical reasons we created one mask per configuration with the same size, where the known points are valued as 1 and the unknown points are valued as 0. With this set we have a total of 1024 (32*32) points per mask and the amount of known points for each configuration is:

- 49 known points, where $49/1024 = 0.0478515625 \approx 5\%$
- 100 known points, where $100/1024 = 0.09765625 \approx 10\%$
- 144 known points, where $144/1024 = 0.140625 \approx 15\%$

To uniformly distribute the known points m in the mask of size W x H, we value each coordinate (c_i, r_j) to 1, where:

$$\begin{aligned} c_i &= c_1 + (i-1)\delta_w, \\ r_j &= r_1 + (j-1)\delta_h \\ \forall i, j &= 1, ..., \sqrt{m} \end{aligned} \tag{9}$$

where the initial observed point (c_i, r_j) is (0,0), $\delta_w = (W-1)(\sqrt{m}-1)$ and $\delta_h = (H-1)(\sqrt{m}-1)$.

Having formally defined the masks corresponding to each sampling configuration we illustrate the known points in white and the unobserved points in blue in Figure 3,

The sampling process is only complete when the masks are combined with the rasters that constitute the dataset. This combination is achieved by sliding the 32x32 masks through the rasters, keeping only the real observed values and valuing the unobserved as 0. This operation can be described as a convolution where the 32x32 kernel/mask is sliding through the raster using a (32,32) stride. After performing this operation for each raster and using all the three masks, the sampling process is complete and per each real raster we have three sampled rasters, one for each sampling configuration.

4.3 Interpolation using Inverse Distance Weighting

In this study we implemented the IDW approach as defined in equation 2, where we used the euclidean distance for the parameter d and the power parameter p with the value of 2, because it is the most common value used in the literature and to simplify the calculations of the euclidean distance.

In addition to the parameters described above, we set the number of known points used to interpolate an unknown point to the 100 closest points to the unknown point we want to interpolate. This measure

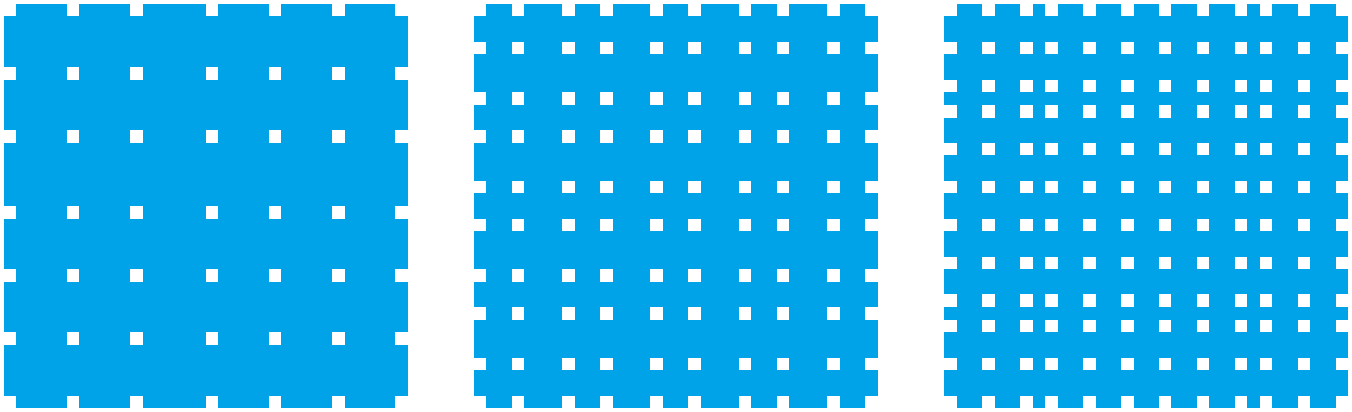

Figure 3: Illustration of the three sampling masks described above, where, from left to right, we have the 49 known points configuration, the 100 known points configuration and the 144 known points configuration.

was taken since in the bigger rasters, especially the one regarding mainland Portugal the amount of known points is in the scale of millions and we could not interpolate all the hundreds of millions of unobserved points using the millions of observed points in an available time.

After interpolating the rasters for each sampling configuration, the new interpolated rasters were stored so they could be used in the two channel input and enable the study of our hybrid approach.

4.4 Adversarial training procedure

The network was trained using mini-batch stochastic gradient descent (SGD) with a batch size of 64. The selected training dataset was restricted to the 32x32 patches composed uniquely by values above the sea level (i.e., height ≥ 0) of the mainland Portugal raster and we saved the other regions including the mainland Portugal raster to validate the trained generator. Having that, we split the training dataset of 135577 patches into a training set of 115240 patches (85 % of the dataset) and into a test set of 20337 (15 % of the dataset). We used the Adam optimizer, where β_1=0.9, β_2=0.999, and the learning rate, α, for backpropagation was set to 0.00001. The update of the Discriminator and Generator weights was achieved using the binary cross entropy loss function following equations 4 and 5. We trained the network for 200 epochs and after each epoch we tested the generator using the test set. To measure the performance of the generator after each input we calculated the average RMSE of the test set patches and if the performance of the generator did not improve for 10 consecutive epochs we halted the training process and consider the generator trained. For each sampling configuration we trained two networks the network where the input is composed by a sampled patch and the network where the input is composed by the sampled patch in one channel and the correspondent patch interpolated using the IDW technique in the other channel. Since we had to train 6 networks, two networks for each sampling configuration, and using a dataset with 135577 patches we halted the process when the generator did not improve for 10 consecutive epochs and only used the RMSE, instead of using the MAE and the RMSE, in order to reduce the time necessary to train the network.

4.5 Validation of the trained generator

To validate the generator ability to interpolate images and generalize, we applied the trained generator to the rasters of each Portuguese island and to the raster of mainland Portugal. In the following Table 1 we can see the average of the errors, MAE and RMSE, associated to each 32x32 patch composed uniquely by values with height above the sea level of the different regions we studied. These results were produced by the three techniques studied: IDW, 1 channel GAN and our 2 channel GAN hybrid approach.

After analysing the obtained results, the main observations are:

Regions	Sampling Configuration	IDW		1 channel input GAN		2 channel input GAN	
		avg(MAE)	avg(RMSE)	avg(MAE)	avg(RMSE)	avg(MAE)	avg(RMSE)
	49	15.111	18.308	**5.591**	9.060	5.770	**7.790**
Ao̧res - Corvo	100	10.020	12.313	**5.646**	**7.479**	5.862	7.863
(15 patches)	144	7.760	9.598	4.849	7.716	**4.203**	**6.074**
	49	5.594	7.022	3.316	4.844	**3.222**	**4.193**
Ao̧res - Faial	100	3.837	4.901	4.357	5.389	**3.696**	**4.798**
(226 patches)	144	3.077	3.953	2.972	4.367	**2.397**	**3.289**
	49	11.176	13.886	5.241	7.710	**4.775**	**6.340**
Ao̧res - Flores	100	7.765	9.742	5.701	7.299	**5.240**	**6.903**
(183 patches)	144	6.275	7.879	4.334	6.491	**3.617**	**5.049**
	49	6.150	7.981	2.964	4.374	**2.946**	**3.972**
Ao̧res - Graciosa	100	4.253	5.597	**2.699**	3.661	2.765	**3.647**
(72 patches)	144	3.426	4.520	2.385	3.625	**2.032**	**2.821**
	49	4.399	5.624	3.958	5.681	**3.787**	**4.832**
Ao̧res - Pico	100	3.059	3.970	5.718	6.910	**4.674**	**6.055**
(607 patches)	144	2.443	3.186	3.572	5.037	**2.839**	**3.850**
	49	8.946	11.418	**4.815**	7.223	4.910	**6.467**
Ao̧res - São Jorge	100	6.122	7.910	5.880	7.419	**5.260**	**6.943**
(285 patches)	144	4.869	6.292	4.315	6.684	**3.518**	**4.940**
	49	6.241	7.857	3.150	4.700	**2.747**	**3.675**
Ao̧res - Santa Maria	100	4.593	5.814	3.152	4.122	**2.869**	**3.814**
(114 patches)	144	3.832	4.832	2.564	3.836	**2.143**	**3.002**
	49	7.412	9.393	3.954	5.760	**3.532**	**4.649**
Ao̧res - São Miguel	100	5.339	6.804	4.510	5.712	**3.929**	**5.138**
(1037 patches)	144	4.418	5.620	3.295	4.815	**2.704**	**3.727**
	49	4.503	5.764	2.973	4.236	**2.628**	**3.413**
Ao̧res - Terceira	100	3.199	4.140	3.911	4.762	**3.102**	**4.024**
(556 patches)	144	2.631	3.416	2.490	3.522	**2.042**	**2.773**
	49	18.724	22.931	8.208	12.023	**7.468**	**9.875**
Madeira - Madeira	100	12.776	15.745	9.163	11.813	**8.352**	**10.936**
(1048 patches)	144	10.228	12.625	7.138	10.488	**5.715**	**7.952**
	49	7.745	9.860	**2.906**	**4.181**	3.528	4.681
Madeira - Porto Santo	100	5.181	6.606	**2.715**	**3.719**	2.885	3.739
(39 patches)	144	4.129	5.286	2.557	3.873	**2.160**	**2.921**
	49	4.940	6.092	2.981	4.206	**2.428**	**3.122**
Mainland Portugal	100	3.551	4.388	4.012	4.897	**3.009**	**3.918**
(135577 patches)	144	2.944	3.630	2.438	3.353	**2.002**	**2.723**

Table 1: Obtained results for the IDW, 1 channel GAN and our hybrid approach considering only above the sea level 32x32 patches for each region.

- both networks overcame the traditional IDW technique, especially in the islands where the networks largely surpassed the IDW i.e. Madeira-Madeira island.
- when comparing the two GANs, the 1 channel GAN only fully overcame (in terms of MAE and RMSE) our hybrid apprach in the Porto Santo island for the 49 and 100 observed points sampling configurations. Having this, we can claim that our hybrid approach surpassed the more simple GAN showing a higher ability to adapt to different contexts.
- focusing on the results of our hybrid approach we highlight the performance of the network for the 149 sampling configuration where it was never surpassed by the simpler GAN. Looking at the 49 and 100 sampling configurations we can observe that the 49 configuration overcame the 100 configuration for every region and for the MAE and RMSE. Despite being an odd result, the ability of the network produce more with less is very good and can save a lot of resources when transposed to the real world.

These observations portray a positive result since our approach surpassed the IDW technique and the 1 channel GAN approach. Besides showing the capability to overcome traditional and modern techniques our approach also demonstrated a higher degree of generalization when compared to the simpler GAN, as we can observe in the performance of our model on the different contexts from the one it was trained, the Portuguese islands.

5 Conclusions and Future Work

Overall we can assert that this study and our proposed approach was able to surpass the traditional and modern techniques in terms of performance. In terms of generalization we cannot compare with the traditional technique, but comparing with the simpler GAN we can also state that our hybrid approach achieved a higher generalization degree, thus fulfilling everything we hoped to achieve. Although our study focuses on a spatial interpolation problem that is not relevant since the elevation land covered area is known with a good precision, we believe that our study using DEMs can be transposed to the ocean surface where the available information is inferior with less precision and, thus contributing to many relevant sciences to our world.

Looking at the future we propose to apply this model to other spatial interpolation problems. In addition to study our approach in other contexts we intend to study how the network will behave when the network structure is altered, for example, when batch normalization layers are added to the sub-models and if that regularization will increase performance and generalization ability. Finally, we propose to study how the various spatial interpolation techniques affect the performance of our approach when used to interpolate the patch that is combined with the sampled patch in our 2 channel input.

References

Deng, X., Zhu, Y., Tian, Y., and Newsam, S. (2019). Generalizing deep models for overhead image segmentation through getis-ord gi* pooling.

Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial networks.

Greenberg, J., Rueda, C., Hestir, E., Santos, M., and Ustin, S. (2011). Least cost distance analysis for spatial interpolation. *Computers Geosciences*, 37:272–276.

Marcos, D., Volpi, M., Kellenberger, B., and Tuia, D. (2018). Land cover mapping at very high resolution with rotation equivariant cnns: Towards small yet accurate models. *ISPRS Journal of Photogrammetry and Remote Sensing*.

Rigol, J., Jarvis, C., and Stuart, N. (2001). Artificial neural networks as a spatial interpolation. *International Journal of Geographical Information Science*, 15:323–343.

Stachelek, J. and Madden, C. (2015). Application of inverse path distance weighting for high-density spatial mapping of coastal water quality patterns. *International Journal of Geographical Information Science.*

Tobler, W. (1970). A computer movie simulating urban growth in the detroit region. *Economic Geography*, pages 234–240.

Zhu, D., Cheng, X., Zhang, F., Yao, X., Gao, Y., and Liu, Y. (2019). Spatial interpolation using conditional generative adversarial neural networks. *International Journal of Geographical Information Science*, pages 1–24.