

Article

Multi-Output Regression with Generative Adversarial Networks (MOR-GANs)

Toby R. F. Phillips ^{1,*} , Claire E. Heaney ^{1,2} , Ellyess Benmoufok ³, Qingyang Li ³, Lily Hua ⁴,
Alexandra E. Porter ⁵, Kian Fan Chung ⁶ and Christopher C. Pain ^{1,2} 

¹ Applied Modelling and Computation Group, Department of Earth Science and Engineering, Imperial College London, London SW7 2AZ, UK

² Centre for AI-Physics Modelling, Imperial-X, Imperial College London, London W12 7SL, UK

³ Department of Earth Science and Engineering, Imperial College London, London SW7 2AZ, UK

⁴ Department of Chemistry, Imperial College London, London SW7 2AZ, UK

⁵ Department of Materials, Imperial College London, London SW7 2AZ, UK

⁶ Faculty of Medicine, National Heart & Lung Institute, Imperial College London, London SW3 6LY, UK

* Correspondence: t.phillips18@imperial.ac.uk

Abstract: Regression modelling has always been a key process in unlocking the relationships between independent and dependent variables that are held within data. In recent years, machine learning has uncovered new insights in many fields, providing predictions to previously unsolved problems. Generative Adversarial Networks (GANs) have been widely applied to image processing producing good results, however, these methods have not often been applied to non-image data. Seeing the powerful generative capabilities of the GANs, we explore their use, here, as a regression method. In particular, we explore the use of the Wasserstein GAN (WGAN) as a multi-output regression method. The resulting method we call Multi-Output Regression GANs (MOR-GANs) and its performance is compared to a Gaussian Process Regression method (GPR)—a commonly used non-parametric regression method that has been well tested on small datasets with noisy responses. The WGAN regression model performs well for all types of datasets and exhibits substantial improvements over the performance of the GPR for certain types of datasets, demonstrating the flexibility of the GAN as a model for regression.

Keywords: Generative Adversarial Networks; Wasserstein GAN; regression; multi-output regression; multi-modal distributions



Citation: Phillips, T.R.F.; Heaney, C.E.; Benmoufok, E.; Li, Q.; Hua, L.; Porter, A.E.; Chung, K.F.; Pain, C.C. Multi-Output Regression with Generative Adversarial Networks (MOR-GANs). *Appl. Sci.* **2022**, *12*, 9209. <https://doi.org/10.3390/app12189209>

Academic Editor: Grzegorz Dudek

Received: 22 August 2022

Accepted: 8 September 2022

Published: 14 September 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Regression is a statistical technique which aims to find and describe relationships that exist between inputs (the independent variables also known as predictors, covariates, features) and outputs (dependent variables also known as responses, targets, outcomes). An abundance of data has enabled machine learning techniques to be successfully applied to regression modelling. Data from observations or experiments often comes from complex nonlinear systems that are challenging to model, therefore, a regression model that is able to model uni- or multi-modal distributions, single or multi-output regression problems and quantify uncertainty is highly desirable. Borchani et al. [1] highlight two challenges for regression: (1) modelling uncertainty, both handling the uncertainty in the data itself, but also in quantifying the uncertainty in the responses; and (2) identifying co-dependencies between response variables (for multi-output regression problems). Two approaches are commonly used for multi-output regression problems: transforming the problem and applying single-output methods, and developing extensions to single-output regression methods (such as kernel methods, regression trees and support vector regression) so they are capable of analysing multi-output distributions [2]. Although the former is more straightforward, the latter, when possible, gives better results. In this paper, we propose

a generative model for performing regression. This model is flexible as it can be applied (without modification other than hyperparameter tuning) to uni- and multi-modal data; multiple regression problems; single- and multi-output regression tasks (including co-varying responses); and to data with uncertainty or noise. It can also be used to calculate the uncertainty associated with a prediction. We compare this method to Gaussian Process Regression (GPR) which has performed well for regression problems. GPR is a machine learning technique based on Gaussian Processes introduced by Rasmussen and Williams in 1996 [3]. A probability distribution is defined, rather than a single-valued function, which can be applied to data where a range of responses can come from a single point in the regression phase space. Feed-forward neural networks give a single response for a given input, whereas, both GPR and the method proposed here can give multiple responses for a single input enabling the uncertainty in the response to be quantified. This is a highly desirable feature for a regression method.

1.1. Related Work

Generative models were originally developed with the aim of creating a network that could generate realistic examples, that is examples that appear to be drawn from the distribution which was used to train the model. A powerful generative model is the Generative Adversarial Network (GAN) introduced in 2014 by Goodfellow et al. [4]. GANs have quickly become one of the most popular generative models and are widely used in image processing [5] where they are well known for generating images that are capable of tricking the human eye into believing that it is seeing genuine data [6]. Instead of learning a mapping between an input and output determined by training data, these models attempt to learn the distribution underlying the training data (in fact, they learn a mapping from a simple distribution to the more complex distribution which describes the training data). This property is desirable as we would like to avoid extrapolating because it can lead to unreliable results. A GAN consists of two neural networks, a generator and a discriminator that are trained simultaneously according to a min-max game. The generator and discriminator adopt the structure of popular neural networks [7–9]. Although many studies have explored the idea of using GANs when manipulating or identifying images, little research currently exists around implementing GANs to generate non-image data with targeted distributions. One exception is Jolaade et al. [10] who apply GANs to the time series prediction of fluid flow. Furthermore, GANs have shown to be able to perform well even with small samples of data [11], making them a reliable technique and suitable for regression in these circumstances. Since their introduction in 2014, a number of variants have been developed, including the Wasserstein GAN (WGAN) [12,13]. This particular flavour of GAN was introduced to address the problems of mode collapse and vanishing gradients [14] from which the GAN [4] and DCGAN [7] are known to suffer.

GAN methods are not widely used for regression in the literature, with the exception of Aggarwal et al. [15] and McDermott et al. [16]. Aggarwal et al. [15] apply Conditional GAN (CGAN) to a number of datasets, including one which predicted property prices in California and another which predicted the control action on the ailerons given the status of the aeroplane. McDermott et al. [16] apply a semi-supervised Cycle Wasserstein Regression GAN (CWR-GAN) to biomedical applications such as predicting a patient's response to treatment. Both articles showed good results, but both commented on the additional training time and training complexity exhibited by the GAN models in comparison with other methods. The CGAN and the CWR-GAN both have a different structure to the WGAN implemented here. Our WGAN (as with a standard GAN) generates a sample from random values (the input to the WGAN), whereas the CGAN and CWR-GAN have inputs and outputs of the same dimension, although the input can have additional variables corresponding to noise or constraints. Therefore CGANs and the CWR-GAN can be more straightforward to use for regression and time series modelling.

We compare our GAN approach with regression performed by Gaussian Process Regression (GPR). GPR has become an effective, non-parametric Bayesian approach that

can be applied to regression problems and can be utilised in exploration and exploitation scenarios [17]. Instead of inferring the distribution of parameters, non-parametric methods can directly predict the distribution of functions. Gaussian Process Regression starts with a set of prior functions based on a specified kernel. After incorporating some known function values (from the training dataset), a posterior distribution is obtained. The posterior can then be evaluated at points of interest (from the test dataset) [18].

1.2. Contributions and Outline

Due to the structure of GANs, the independent and dependent variables appear in the output of the generator (whereas for feedforward networks, the independent variables would be more likely to appear in the input, and the dependent variables in the output). The input of a GAN is a set of random variables, and it generates a realistic sample from these random variables. For regression problems, although sampling the latent space will give a good idea of the distribution learned by the generator, it can also be desirable to be able to obtain a response at a particular value of the independent variable. In order to do this, we propose a prediction algorithm which involves minimising the difference between the output of the GAN for the independent variable and its desired value. This prediction algorithm has been used previously to enable a GAN to make time series predictions [10,19]. It is somewhat similar to an algorithm presented by Wang et al. [20], which searches the latent space in order to match a given image with an image produced by the generator. The necessity for these algorithms comes about because the output of the GAN contains both the independent and dependent variables. In this paper, we develop a new regression method based on GANs and show how it compares to a state-of-the-art GPR regression method by testing both methods on a range of datasets. We apply the same model (a GAN) to all the datasets in the paper and compare with a standard GPR model. Although specific types of GPR have been developed for particular datasets (for example, Heteroscedastic GPR [21], and GPR for clustered data [22]), here we choose a single type of GPR model as we do not tailor the GAN to the specific datasets (other than optimising the architecture and other hyperparameters as is usual). This enables us to demonstrate the flexibility of the single GAN model.

The contributions of this article are the use of a WGAN to perform regression; the ability to apply this model to multi-modal data and multi-output regression (MOR-GAN) tasks with no modifications required to the GAN; the presentation of a prediction algorithm to be used with the trained GAN in order to predict a response for a given independent variable; the exploitation of the WGAN's critic to provide a confidence level or assessment of reliability for the predictions made by the WGAN's generator.

The remainder of the paper is organised as follows: Section 2 describes the methods used in this paper, Section 3 presents results from the synthetic example problems and Section 4 shows results from an in vitro study. Section 5 gives an overview of the speed of the proposed method. Conclusions are drawn and indications given as to future work in the final section. The notation used in this paper is summarised in Table A1 in the Appendix A.

2. Methods

2.1. Data Generation

We investigate the performance of Gaussian Process Regression (GPR) and Wasserstein GAN (WGAN) models for regression using a number of datasets. Simple functions were used to generate all but one of the datasets, which have different properties, including with and without additive Gaussian noise (which here, represents uncertainty in the data); one- or two-dimensional examples; uni- and multi-modal distributions; single or multi-output regression; and, for the WGAN model, we explore both random inputs and constrained inputs (where input refers to the independent variable or input of the regression problem not the input of the WGAN). The final dataset is taken from an in vitro study and explores the influence of silver nanoparticles on cells taken from the lungs. Following standard practice, preprocessing was applied to all the datasets to ensure that no bias is introduced

due to different variables having different ranges of values. This was done by applying a linear mapping to normalise the values so they were in the range $[-1, 1]$.

2.2. Gaussian Process Regression

GPR is a machine learning technique, based on Bayesian theory and statistical learning which has wide applicability to complex regression problems with multiple dimensions and non-linearity [18]. The basic theory of prediction with Gaussian processes dates back to the 1940s [23,24], and, since then, there have been many developments and insights gained into using Gaussian Processes as a regression technique. For example, Jerome Sacks and Wynn [25] introduced GPR for computer experiments and used parameter optimisation in the covariance function and also applied it to experimental design, i.e., the choice of input that provides the most information. Moreover, Rasmussen and Williams [18] described GPR in a machine learning context, and expressed the optimisation of the GPR parameters in terms of co-variance functions.

A python library GPy was used to perform the GPR [26]. Important to the performance of the GPR is the choice of kernel. Here we use a radial basis function (RBF) kernel which has three hyperparameters; length, kernel variance, and the standard deviation of the Gaussian noise. These hyperparameters are automatically tuned via GPy.

2.3. Generative Adversarial Networks

A Generative Adversarial Network (GAN) consists of two neural networks: a generative model or generator, G , and a discriminative model or discriminator, D . The models are trained simultaneously resulting in a generator that can produce samples which appear to be taken from the same distribution as the training data. During training, the generator tries to fool the discriminator that it is generating real data, see [4]. For each data point the following combined loss function is defined for G and D :

$$L = \min_G \max_D [\log(D(\mathbf{x})) + \log(1 - D(G(\boldsymbol{\alpha})))] \quad (1)$$

where $\mathbf{x} \in \mathbb{P}_r$ is a sample from the real data and $\boldsymbol{\alpha}$ represents the latent variables. The generator and discriminator are essentially playing a two-player min-max game through the corresponding function $V(G, D)$ [4]:

$$\min_G \max_D V(D, G) = \mathbb{E}_{\mathbf{x} \sim \mathbb{P}_r(\mathbf{x})} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_{\boldsymbol{\alpha}}(\boldsymbol{\alpha})} [\log(1 - D(G(\boldsymbol{\alpha})))] \quad (2)$$

GANs are notoriously difficult to train, often reported to suffer from mode collapse and the vanishing gradient problem [14]. Mode collapse occurs when the generator G produces only one solution, or a limited set of solutions, which is/are able to fool the discriminator, and the vanishing gradient problem is described below (Section 2.4).

2.4. Wasserstein Generative Adversarial Networks

The WGAN [12] was developed in order to alleviate the issue of the vanishing gradient problem. To measure the distance between probability distributions, rather than use the Jensen-Shannon (JS) divergence (expressed by Equation (1)) as in the GAN, Arjovsky et al proposed the Earth-Mover (EM) or Wasserstein-1 distance:

$$W(\mathbb{P}_r, \mathbb{P}_g) = \inf_{\gamma \in \Pi(\mathbb{P}_r, \mathbb{P}_g)} \mathbb{E}_{(x,y) \sim \gamma} [\|x - y\|], \quad (3)$$

where $\Pi(\mathbb{P}_r, \mathbb{P}_g)$ denotes the set of all joint distributions $\gamma(x, y)$ whose marginals are respectively \mathbb{P}_r (real data) and \mathbb{P}_g (generated data) [12]. The Wasserstein-1 distance is able to provide a similarity measure between two probability distributions, even when the two probability distributions have no overlap, making it a more sensible cost function. The discriminative model is renamed the **critic** in the WGAN, as it is not explicitly attempting to classify inputs as real or fake, but rather to determine how real an input

is. The WGAN value function is constructed via the Kantorovich-Rubinstein duality as Equation (3) is computationally intractable [12]:

$$\min_G \max_{D \in \mathcal{D}} \mathbb{E}_{x \sim \mathbb{P}_r} [D(x)] - \mathbb{E}_{\tilde{x} \sim \mathbb{P}_g} [D(\tilde{x})] \quad (4)$$

where \mathcal{D} is the set of 1-Lipschitz functions. To enforce the Lipschitz constraint, weight clipping was originally used by Arjovsky et al. [12], who stated that this method of enforcement was terrible, despite it working well for the examples shown in their paper and was, at least, simple. Gulrajani et al. [13] introduced an improvement to weight clipping, by enforcing the Lipschitz constraint with a **Gradient Penalty (GP)** method. By enforcing a soft version of the constraint with a penalty, the new loss function becomes:

$$L = \underbrace{\mathbb{E}_{\tilde{x} \sim \mathbb{P}_g} [D(\tilde{x})] - \mathbb{E}_{x \sim \mathbb{P}_r} [D(x)]}_{\text{loss of the critic}} + \lambda \underbrace{\mathbb{E}_{\hat{x} \sim \mathbb{P}_{\hat{x}}} [(\|\nabla_{\hat{x}} D(\hat{x})\|_2 - 1)^2]}_{\text{gradient penalty}}. \quad (5)$$

Throughout this study we enforce the Lipschitz constraint by using the GP method in our WGAN models.

2.5. Regression with WGAN

GANs are a type of generative network, the aim of which is to generate realistic looking samples that appear to have been drawn from the same distribution as the training data. The input to a WGAN (and GANs in general) is a set of random numbers (not related to the data) and the output contains the generated sample. Consider simple regression, where a relationship is sought between an independent variable x (also known as covariate or feature) and a dependent variable y (also known as a response). The WGAN is trained to produce both the independent and dependent variables of the regression problem (x, y) as its output. In contrast, for feed-forward networks, the input of the network often takes the independent variable and the output, the dependent variable. Therefore, in order to specify a particular value for the independent variable, a prediction algorithm is introduced to the WGAN.

Suppose we have trained the WGAN and wish to use it to predict the output at a given value of the independent variable x_p . First, the latent variables (α) are set to random numbers. The generator is evaluated at these values, $G(\alpha)$, producing a pair of values (x, y) —the input and output or response in our regression problem. The difference between x and x_p is then minimised with respect to the latent variables. Once this is done, we assume that the output of the generator closely approximates (x_p, y_p) . The minimisation can be done efficiently by using the same software libraries that are used for back-propagation during training. This procedure means that we can generate multiple outputs for one input, x_p , by starting from different random states for the latent variables, and we can therefore produce a distribution of values which reflect the uncertainty in the output y_p . This procedure is detailed in Algorithm 1 and introduces a projection operator, Proj, which projects the output of the WGAN onto a space that contains only the variables that are to be constrained. For the example described in this paragraph, the projection operator would be represented by the matrix $[1 \ 0]$ for an output of the generator in the form $(x, y)^T$.

Simple regression (for a single independent variable), multiple regression (for more than one independent variable) and multi-output or multi-variate regression (for more than one dependent variable) can all be performed by the WGAN and demonstrated by the results in this paper. Due to the generative nature of the WGAN, both independent and dependent variables are contained in the output of the generator which means that, when randomly sampling the latent space of the generator to produce an output, we have no control over the particular value of independent variable. In order to specify particular values of independent variable(s), a prediction algorithm is used, described in Algorithm 1. So there are two ways of using the WGAN, either with random values for the independent variable or with constrained values:

- **Random input:** Random variables are assigned to the latent space from which the generator of the WGAN yields a realistic output of a n -tuple of the independent and dependent variables associated with regression problem. By sampling the generator many times, this can be used to assess the probability density function learned by the generator. The value of the independent variable(s) cannot be controlled, however, as they are an output of the generator. Although random inputs allow us to see the distribution learned by the generator, having the facility to constrain the independent variables is an important feature.
- **Constrained input:** An algorithm is used in conjunction with the (trained) WGAN to find predictions for given value(s) of the dependent variable(s). This results in a property similar to a GPR, where, for example, the independent variables are inputs of the GPR (and can be prescribed) and the outputs are dependent variables. An inherent property of a trained WGAN is that both independent and dependent variables are contained within the output the generator. Using the constrained input method described here, a WGAN can therefore make a prediction for any combination of known and unknown variables, with the independent variables being treated in the same way as dependent variables. A GPR, however, can only make predictions for the particular set of dependent variables that it was trained on, given the set of independent variables that it was trained on.

Algorithm 1 Prediction Function. Built to be used in conjunction with the trained WGAN, to constrain the independent variable of the regression problem.

Require: The desired value of the independent variable x_p , initial values of the latent variables $\alpha^{(1)} \sim \mathcal{N}(0, 1)$, trained generator G , number of iterations N .

for $i=1, \dots, N$ **do**

$\tilde{x}^i = G(\alpha^{(i)})$ ▷ Output of GAN from latent space of iteration i

$\varepsilon = \text{Proj}(\tilde{x}^i - x_p)$ ▷ Work out mismatch between GAN output and desired value

$\alpha^{i+1} \leftarrow \text{BackPropagation}(\alpha^i, \varepsilon)$ ▷ Adjust latent space by backpropagating mismatch

end for

2.6. WGAN Architectures

The WGAN models were constructed using Keras [27]. The generator is a four-layered network, as displayed in the top orange box in Figure 1, which takes Gaussian distributed noise from the latent space as an input, and outputs the x and y coordinates in a 1D regression problem. The first dense layer employs batch normalisation and the leaky rectified linear activation functions (LeakyReLU) followed by fully-connected dense layers. The last layer applies the non-linear *tanh* activation function. The structure of the critic is also a four-layered network, with a reduced number of neurons. Its input is data from the training set and data generated from the generator. To reduce the likelihood of overfitting, dropout with a probability of 0.2 is applied to the critic. Layer normalisation is employed for the critic as opposed to batch normalisation, as the latter inhibits the performance of the gradient penalty term in Equation (5).

The MOR-GAN used in Section 3.4 for the co-varying spiral dataset follows the same architecture described in the previous paragraph but with certain fully-connected layers replaced by convolutional layers. In fully-connected or dense layers, every neuron in the input is connected to every neuron in the output. Instead, convolutional layers apply filters to the input where only neurons close to each other are connected to the output.

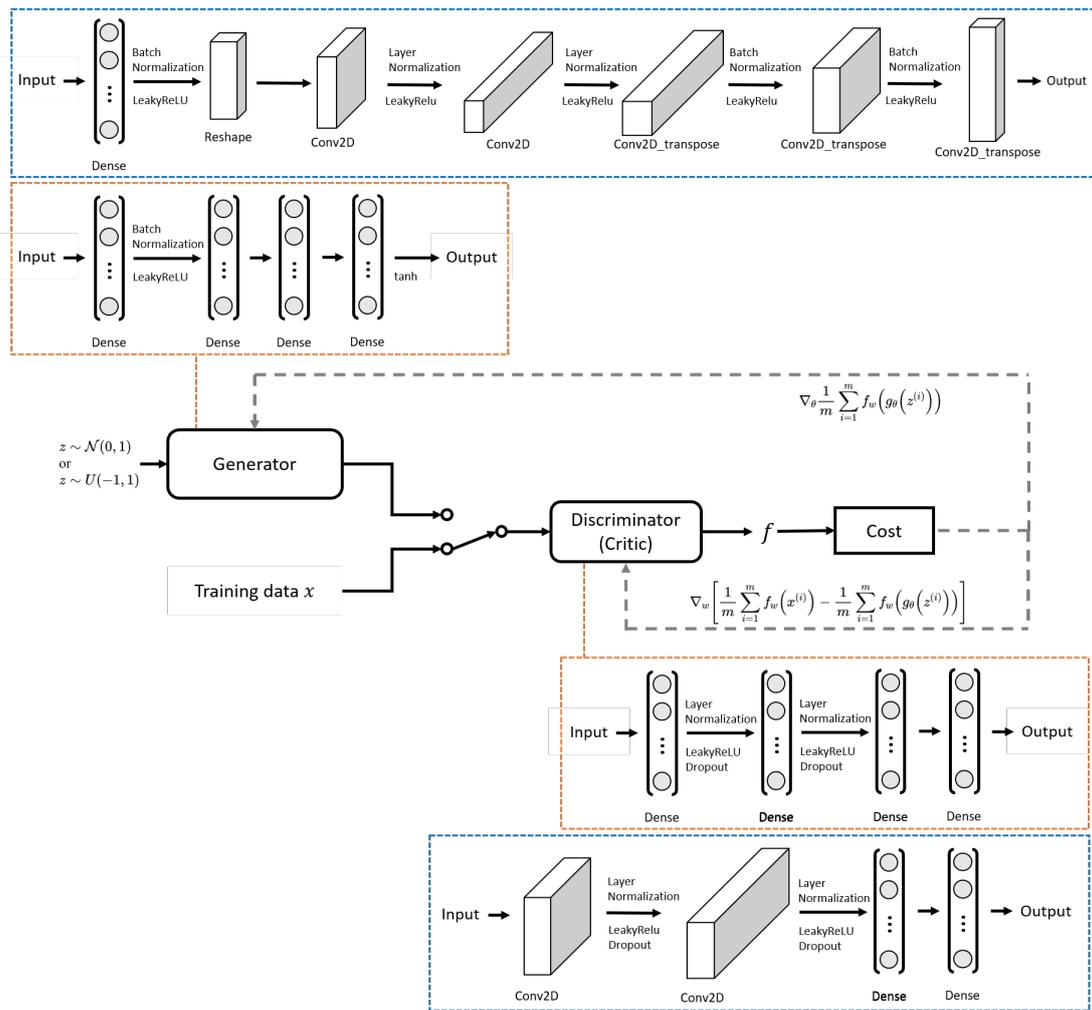


Figure 1. WGAN Structure. The architectures of the Generator and Critic are shown above and below the WGAN structure respectively. The equations displayed are the losses used to update each component. The orange-boxed structure used for the **single-output regression** problems is a multilayer perceptron network. The blue-boxed structure used for **multi-output regression** has convolutional layers.

2.7. Visualisation

To compare the models, the predictions of a test set are visualised. Assuming we have a regression problem in which x is on the horizontal axis and y on the vertical axis and a point on the graph is represented by (x, y) . The GPR outputs are the randomly sampled x values and the associated y values from the GPR posterior distribution. On the other hand the WGAN outputs a prediction of both the x and y values. After training, the generator will produce an output of (x, y) when given a value(s) of the latent variable(s) α .

2.8. Statistical Analysis

To assess the accuracy of the regression method some statistical analysis is performed on the results. The 1D synthetic datasets have the Kolmogorov-Smirnov (KS) test applied to them [28]. A number of specific coordinates (x_i) are chosen with which to perform the KS test. Within the real and generated data there exists a range where the x -coordinate satisfies the condition $x_i - 0.01 < x < x_i + 0.01$. The corresponding y -coordinates form

a distribution in the real dataset ($\mathbb{P}_{r,i}$) and a distribution in the generated dataset ($\mathbb{P}_{g,i}$). The average p -value is then determined by:

$$\bar{p} = \frac{\sum_{i=1}^I \text{KS}(\mathbb{P}_{r,i}, \mathbb{P}_{g,i})}{I}, \quad (6)$$

where KS is the Kolmogorov-Smirnov test.

The silver data is assessed by using the Mann-Whitney U test [29]. The test is performed on the real data and generated data that corresponds with a given time level (x_i), concentration level (y_j) and surface area (z_k). A number of responses ($r_{i,j,k}$) exists for these three measurements and they form a distribution in the real dataset ($\mathbb{P}_{r,i,j,k}$) and a distribution in the generated dataset ($\mathbb{P}_{g,i,j,k}$). The average p -value is then determined by:

$$\bar{p} = \frac{\sum_{i=1}^I \sum_{j=1}^J \sum_{k=1}^K \text{MW}(\mathbb{P}_{r,i,j,k}, \mathbb{P}_{g,i,j,k})}{IJK}, \quad (7)$$

where MW is the Mann-Whitney U test. Both metrics are implemented using the SciPy package [30].

3. Results from Synthetic Datasets

In this section we present results for the performance of the GPR and WGAN approaches for regression on a number of synthetic datasets. We would like our model to perform well on different types of dataset, so datasets with different properties are used here, including uni- and multi-modal distributions, one- or two-dimensional inputs (or independent variables), single- or multi-output. The WGAN can also be used in two ways: with a random input or a constrained input. These combinations are given in Table 1. For the datasets with noise, 500 samples are taken. All the models here follow the general architecture in the orange box of Figure 1 with slight variations on the number of nodes in each dense layer.

Table 1. Properties of the synthetic datasets.

Dataset	Dimension	Distribution of Dataset		Output	Input Type	Section
		Noise	Type			
sine wave	1D	✓	uni-modal	single output	random	3.2.1
heteroscedastic	1D	✓	uni-modal	single output	random	3.2.1
circle	1D	✓	multi-modal	single output	random	3.2.2
sine wave with lines	1D	✓	multi-modal	single output	random	3.2.2
distance	2D	✗	uni-modal	single output	random	3.2.4
helix	2D	✓	multi-modal	single output	random	3.2.4
sine wave	1D	✓	uni-modal	single output	constrained	3.3
heteroscedastic	1D	✓	uni-modal	single output	constrained	3.3
circle	1D	✓	multi-modal	single output	constrained	3.3
eye	1D	✗	multi-modal	multi-output	constrained	3.4.1
spiral	2D	✗	uni-modal	multi-output	constrained	3.4.2

3.1. Training

The training process of the WGAN is described in Algorithm 2. Training a WGAN can be easier than training a GAN, due to the former's removal of the issues associated with mode collapse and weight clipping. Nonetheless, there are still many factors (neural network architecture and training hyperparameters) that can be optimised during training. See Table 2 for the set of hyperparameters that we use for WGAN training. Some values were found by hyperparameter optimisation; others were informed by the literature. For example, $\lambda = 10$ and $n_{\text{critic}} = 5$ are commonly used settings and have been shown to work well across a range of datasets and architectures [12,31].

Algorithm 2 WGAN with gradient penalty and sample-wise optimisation. All experiments in the paper used the default values $\lambda = 10, n_{\text{critic}} = 5, \alpha^\ell = 0.0001, \beta_1 = 0.5, \beta_2 = 0.9$. This algorithm is a modified version of the one displayed in the paper by Gulrajani et al. [13]

Require: The gradient penalty coefficient λ , the number of critic iterations per generator iteration n_{critic} , the batch size m , Adam hyperparameters $\alpha^\ell, \beta_1, \beta_2$.

Require: initial critic parameters ω_0 , initial generator parameters θ_0 .

```

while  $\theta$  has not converged do
  for  $t = 1, \dots, n_{\text{critic}}$  do
    for  $i = 1, \dots, m$  do
      real data  $x \sim \mathbb{P}_r$ , latent variable  $\alpha \sim p(\alpha)$ , a random number  $\epsilon \sim \mathcal{U}[0, 1]$ .
       $\tilde{x} \leftarrow G_\theta(\alpha)$ 
       $\hat{x} \leftarrow \epsilon x + (1 - \epsilon)\tilde{x}$ 
       $L^{(i)} \leftarrow D_\omega(\tilde{x}) - D_\omega(x) + \lambda(\|\nabla_{\hat{x}} D_\omega(\hat{x})\|_2 - 1)^2$ 
    end for
     $\omega \leftarrow \text{Adam}(\nabla_\omega \frac{1}{m} \sum_{i=1}^m L^{(i)}, \omega, \alpha^\ell, \beta_1, \beta_2)$ 
  end for
  Sample a batch of latent variables  $\{\alpha^{(i)}\}_{i=1}^m \sim p(\alpha)$ .
   $\theta \leftarrow \text{Adam}(\nabla_\theta \frac{1}{m} \sum_{i=1}^m -D_\omega(G_\theta(\alpha)), \theta, \alpha^\ell, \beta_1, \beta_2)$ 
end while

```

Table 2. Hyperparameters used in the construction and training of our WGANs for both the single-output and multi-output distributions.

Hyperparameters	Single-Output	Multi-Output
Learning rate	10^{-3}	10^{-4}
Number of Critic iterations per Generator iterations	5	5
Batch size	100	32
Latent Space Dimension	3	3 (3, 6, 12 used for spiral problem)
Adam optimiser hyperparameters (decay rates of moving averages)	0.5 & 0.9	0.5 & 0.9
Gradient penalty hyperparameter λ	10	10

3.2. Single-Output Regression with Random Input Values

In this section, we sample the posterior of the GPR at random points. For the WGAN, we randomly sample points in the latent space which leads to outputs of n -tuples which are the inputs and responses of the regression problem. We do not control which values are taken by the independent variables(s) or inputs when using regression with randomly generated inputs. The test or sample data is generated by evaluating the functions used to create the training data with randomly generated independent variables. Therefore, the three sets of results have different values of the independent variable(s).

3.2.1. 1D Uni-Modal Examples

To generate a sinusoidal dataset with uncertainty, we use the function

$$y = \sin(x) + \eta\phi \quad \text{where} \quad \phi \sim \mathcal{N}(\mu, \sigma) \tag{8}$$

where \mathcal{N} is a Gaussian distribution with mean $\mu = 0$ and standard deviation $\sigma = 1$. The uncertainty is represented by Gaussian noise through the term ϕ and its magnitude is adjusted by a scalar $\eta \in [0, 1]$. We can see from Figure 2 that the random sampling from WGAN and GPR both match well to the test data. For the sinusoidal dataset, the WGAN structure shown in Figure 3 is used, and for the remaining problems in this section, we increase the number of neurons, see Figure 4.

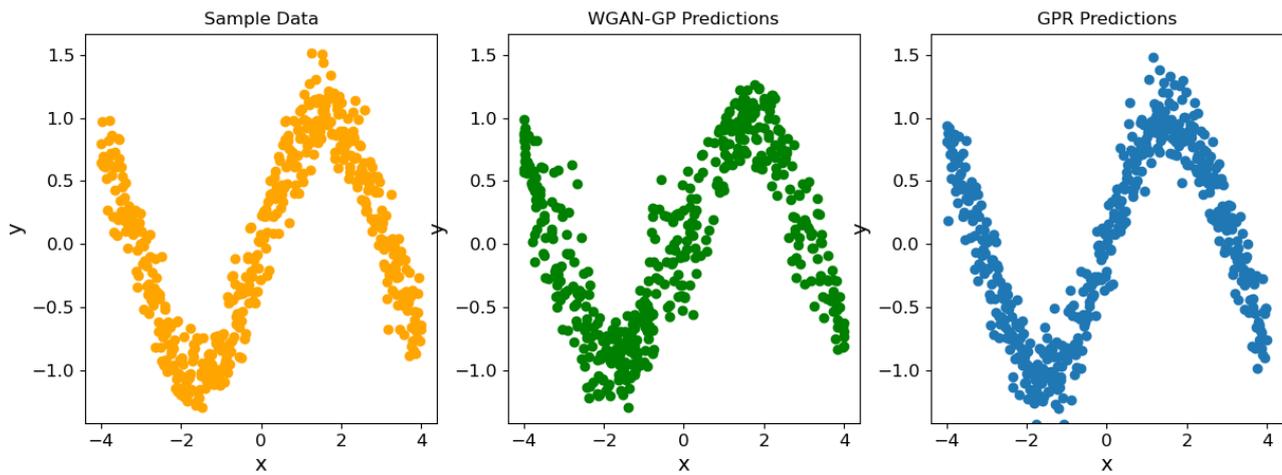


Figure 2. Sinusoidal dataset with added noise ($\eta = 0.2$, see Equation (8)). The test data is shown on the left, sampled points from the WGAN are shown in the middle and sampled points from the GPR posterior are shown on the right.

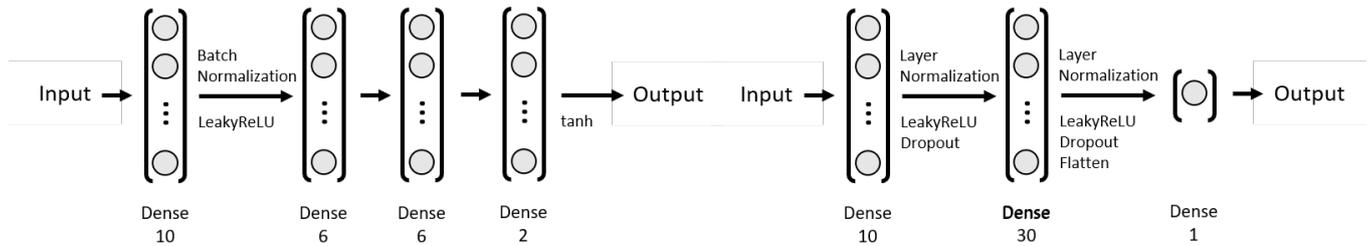


Figure 3. The structure of the generator (left) and critic (right) for the sinusoidal datasets.

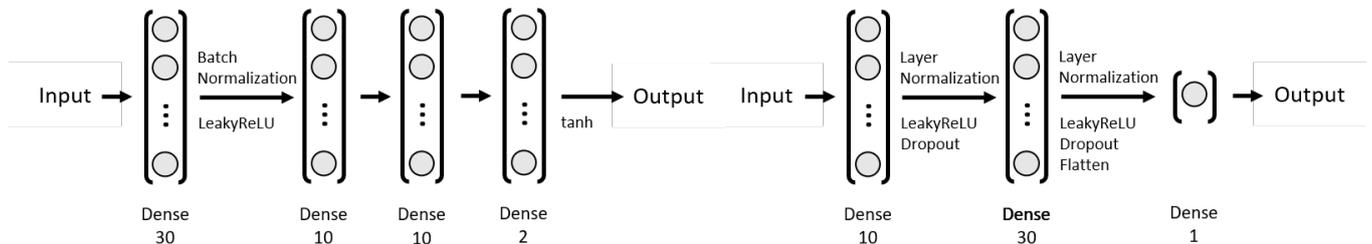


Figure 4. The structure of the generator (left) and critic (right) for the majority of the problems in this section.

The previous example modelled uncertainty by using noise that was independent of x . To test the WGAN model more thoroughly, a heteroscedastic dataset is used where the noise increases with increasing x . Figure 5 shows that the WGAN model is capable of modelling the variation in noise accurately, whereas the GPR, with a single kernel size representing the probability density function, is unable to do so. We note that there is a variant of GPR called Heteroscedastic GPR [21], which has been designed to handle intricate changes in noise. Implementing this method would result in a better performance of the GPR. However, here we aim to avoid tailoring methods to different datasets, so that we can demonstrate the flexibility of the single WGAN model.

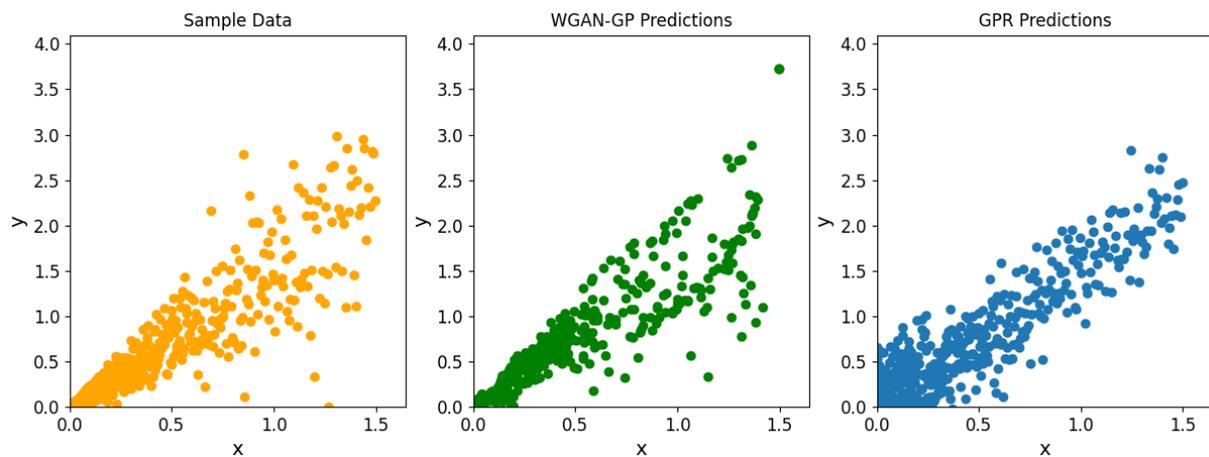


Figure 5. Heteroscedastic dataset. The test data is shown on the left, sampled points from the WGAN are shown in the middle and sampled points from the GPR posterior are shown on the right.

3.2.2. 1D Multi-Modal Examples

Here we explore the use of WGAN and GPR to perform regression of multi-modal distributions. The WGAN models in this section use the architecture displayed in Figure 4. For the first multi-modal distribution, a uniform distribution of data points is generated within an annulus (i.e. between two concentric circles) as shown in Figure 6 (left). There is a significant difference in the performance of the GPR and WGAN. Whilst the WGAN captures the distribution very well (see Figure 6 (middle)), the GPR is unable to represent it (see Figure 6 (right)), predicting an almost uniform distribution of points.

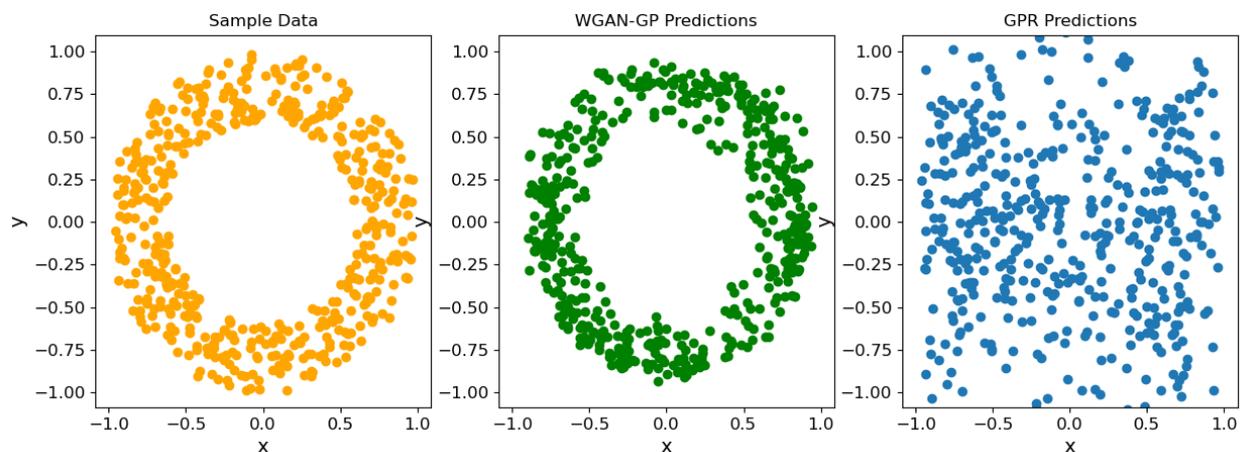


Figure 6. Annulus dataset. The test data is shown on the left, sampled points from the WGAN are shown in the middle and sampled points from the GPR posterior are shown on the right.

The second multi-modal distribution is a sinusoidal wave with several intersecting lines. The same trends appear as seen when using the annulus dataset: the WGAN outperforms the GPR, which is unable to detect the gaps that exist in the dataset, see Figure 7. The overall profile of the data is visible, but within the bounds of the minimum and maximum y values there is no gap. Although GPR struggles with these complex functions, it has been used and built upon to work on clustering complex functions [22], so there is the capability of modelling these types of complex functions. However, we wish to compare the WGAN against one model, without tailoring it for different types of data.

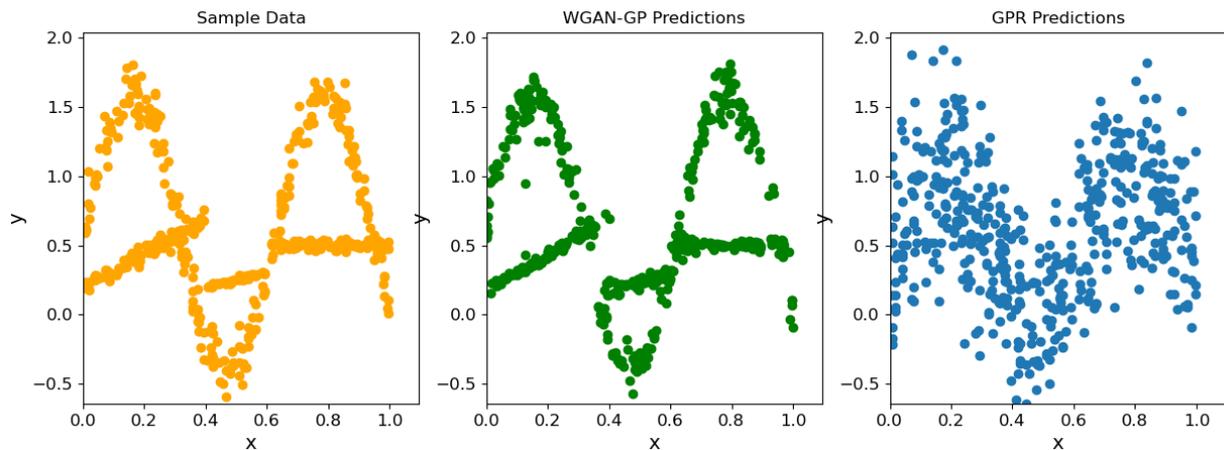


Figure 7. A sine wave intersected by several lines. The test data is shown on the left, sampled points from the WGAN are shown in the middle and sampled points from the GPR posterior are shown on the right.

3.2.3. Confidence of Solutions from the Critic

Sections 3.2.1 and 3.2.2 show how sampled points produced by the generator of the WGAN match the distribution seen in the test data (or sample data). During the training of the WGAN, the critic learns to determine how real an sample is. This section demonstrates how the critic can be used to determine the confidence in a sample produced by the generator, which is an indication of how reliable the method’s predictions are.

Figure 8 shows the value taken by the critic for predictions or responses made throughout domain for both the sinusoidal and annulus datasets. These are produced by finding the value of the critic for each point on a 100×100 grid that covers the same domain as the original data. As previously stated, the critic of a WGAN does not explicitly determine whether a sample is real or fake, but instead, how real a sample is. Therefore the larger the value produced by the critic, the more confidence the model has in the prediction. The critic values shown here are normalised to be between 0 and 1.

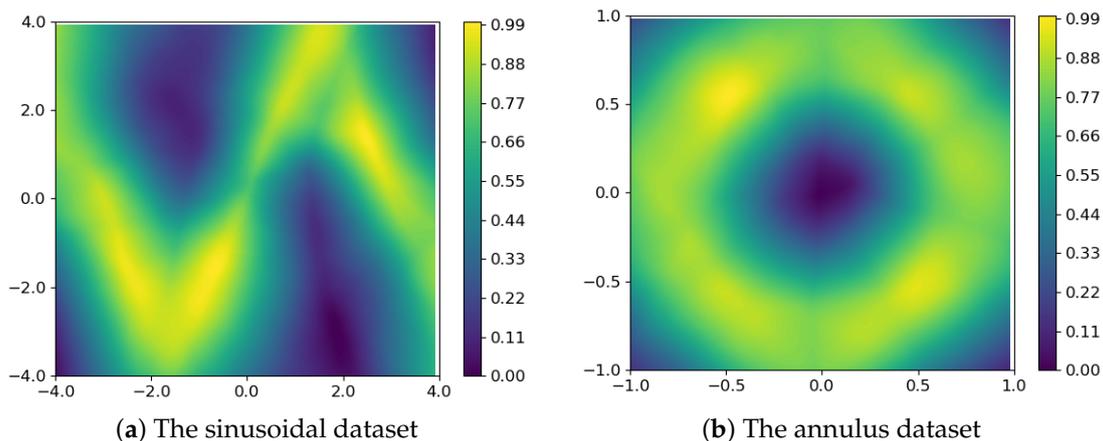


Figure 8. Contour plots showing the values of the two critics for the sinusoidal and annulus datasets. These indicate the confidence in or reliability of the predictions and also indicate where extra training data may be required.

Figure 8a shows the values of the critic produced for the sinusoidal dataset. It can be observed that the values of the critic are higher where the data of the noisy sine curve occurs (see Figure 2), which corresponds to the region mostly occupied by the training data. These values are higher where there is a higher concentration of data points, particularly

around $x = -2.0$ and $x = 2.0$. Outside of where the sine wave is located the critic value sharply decreases, therefore confidence in any prediction made here is low.

Figure 8b shows the values of the critic produced by the annulus dataset. It can be observed that the values are higher within the annulus, which corresponds to the region occupied by the training data (see Figure 6). We can see that the critic produces lower values for coordinates predicted outside of the annulus, meaning that the confidence in predictions or responses that occur here is low.

Figure 8a,b demonstrate how the critic can be used in conjunction with the generator to produce a confidence level in the predictions made by the generator. A lower critic value, and therefore a lower confidence, in a prediction made may indicate that extra training data is required there. A possible location requiring extra data for the sinusoidal dataset is $x = 0.5$ and for the annulus dataset is $x = 0.0$. The critic value can be used to remove solutions generated that are not realistic, thereby improving the results. The solutions shown in Figures 2 and 6 had their average p -value improved from <0.2 to <0.1 by removing the 10% of solutions generated that had the lowest value after being passed through the critic.

Thus, the confidence level might help us to determine where to collect more experimental data or where to observe the system. It also suggests where the neural network is not predicting well, which might not be because of lack of data. Ultimately, this confidence level should be combined with the importance of the region where the confidence is being determined. This importance could be set according to how much or little influence this region may have on the final results. If applying the GAN approach to regression to optimisation, importance could, for example, be determined from sensitivities (or adjoints) of what is important with respect to the independent variables.

3.2.4. 2D Uni- and Multi-Modal Examples

Increasing the dimensions in the inputs of the regression problem means the need for a larger neural network, thus the following problems use the structure displayed in Figure 9.

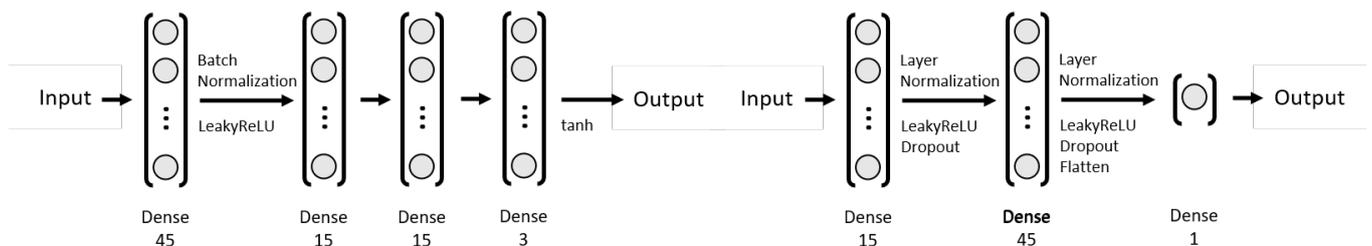


Figure 9. The structure of the generator (left) and critic (right) for the two-dimensional problems.

The performance of the WGAN regression method for data with a single input has been shown to be very reliable. We now test the GPR and WGAN methods on two-dimensional data with a distance function $h = \sqrt{x^2 + y^2}$. The GPR performs exceptionally well, outputting predictions very close to the true model, see Figure 10. The WGAN also performs well, although some deviation from the distance function can be seen.

Having demonstrated that both models are capable of performing regression on datasets with multiple inputs, a more complicated problem is defined as a 2D multi-modal function in the form of a helix with additive Gaussian noise. Figure 11 shows that WGAN is capable of generating data similar to the true model, whereas the GPR struggles to recognise the variation in h (on the z axis) and fills the hole in the circle, looking at the xy plane.

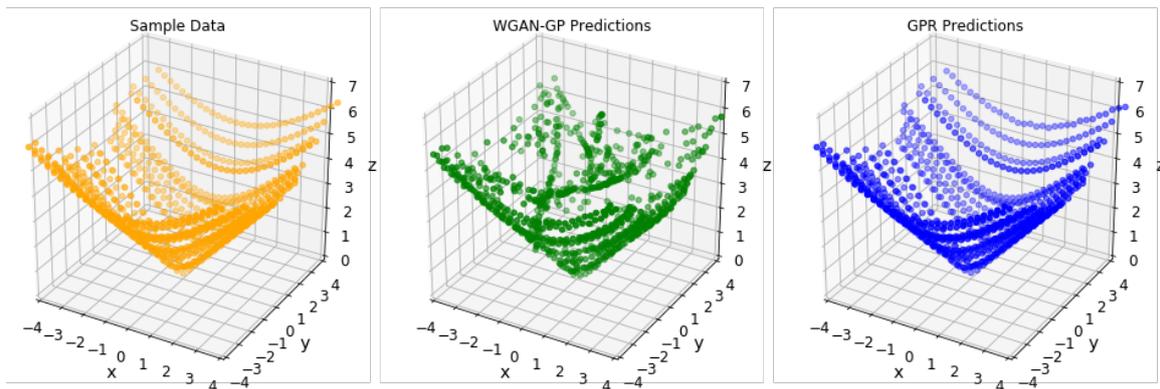


Figure 10. 2D distance function. The test data is shown on the left, sampled points from the WGAN are shown in the middle and sampled points from the GPR posterior are shown on the right.

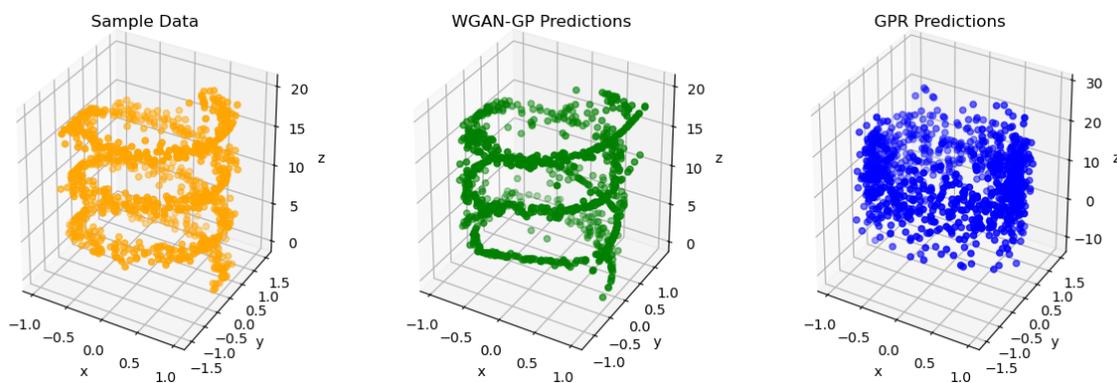


Figure 11. The helix dataset. The test data is shown on the left, sampled points from the WGAN are shown in the middle and sampled points from the GPR posterior are shown on the right.

3.3. Single-Output Regression with Constrained Input Values

A key benefit of using the WGAN for regression is its capability of producing a latent space that with a constrained input, can be optimised to produce multi-modal responses. In Figure 12 we can see the function displaying a few of the potential responses y , at differing fixed x . The WGANs used for the constrained input regression are the same ones used in Sections 3.2.1 and 3.2.2 for their respective datasets.

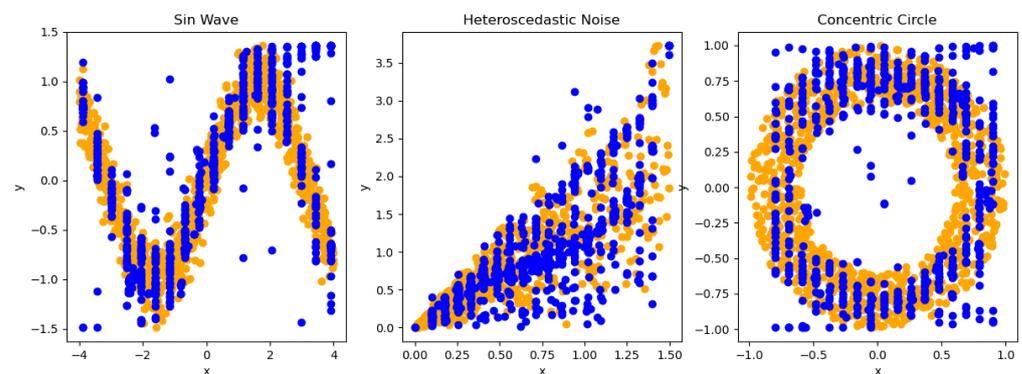


Figure 12. The sinusoidal wave dataset, heteroscedastic noise dataset and annulus dataset predicted at a given values of the x coordinate using the WGAN prediction method. Prediction displays potential responses at the given x coordinates.

The way this optimisation is performed is to first randomly generate a latent input vector of the generator. Then from this initial condition point in latent space we apply our

optimiser to minimise the least squares functional, see Algorithm 1, which aims to match the latent space with the specified x coordinate. We repeat this multiple times in order to obtain a probability density function for this fixed x coordinate but with differing initial latent space inputs. The average p -value for all three solutions generated this way is <0.05 .

3.4. Multi-Output Regression with MOR-GAN

3.4.1. 1D Eye Dataset with Covariance

By taking a digitised, hand-drawn eye and adding a second eye which is obtained by a rotation of 90° and a reflection of the first eye, we produce a distribution which is multi-modal and multi-output or multivariate, see Figure 13. This forms the dataset for the first multi-output regression test. The WGAN is trained to produce two pairs of coordinates, (x_1, y_1) and (x_2, y_2) .

Figures 14 and 15 contains the structure of the generator and discriminator respectively used for the WGAN model in this section.

To provide a challenge for the algorithm which enables the WGAN to make predictions a particular values for the independent variable (Algorithm 1), we constrain the value of x_1 (for the non-rotated eye) and predict the corresponding values for y_1 (non-rotated eye), x_2 and y_2 (rotated eye). We repeat this process for every point in the eye dataset to form the image shown in Figure 16. Similarly, we constrain the value of x_2 (for the rotated eye) and predict the corresponding values for y_2 (rotated eye), x_1 and y_1 (non-rotated eye). This is done for every point in the dataset and the result can be seen in Figure 17. The predictions using the MOR-GAN method take into account the known or learned covariance information between the images, enabling the model to determine the second image from all the points in the first image and vice versa. The agreement between the real data and the predicted data using the constrained input is excellent.

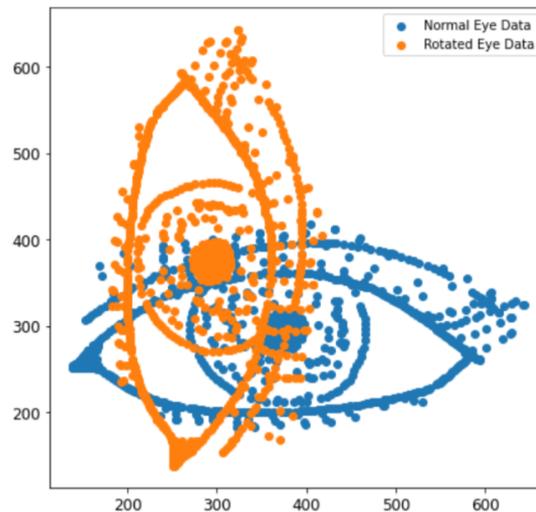


Figure 13. The eye dataset which contains two eyes. One eye is rotated and reflected to produce a second eye.

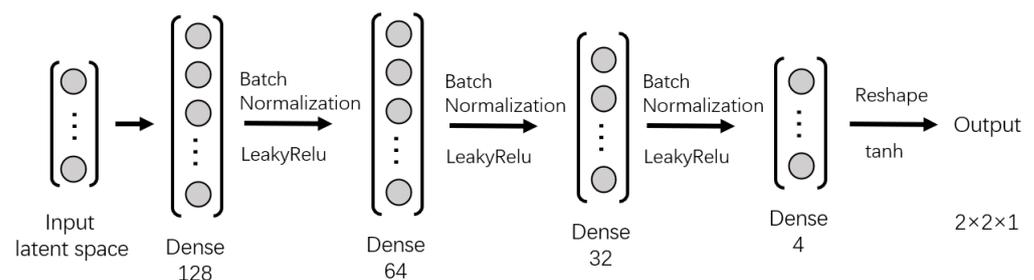


Figure 14. The structure of the generator for the eye dataset.

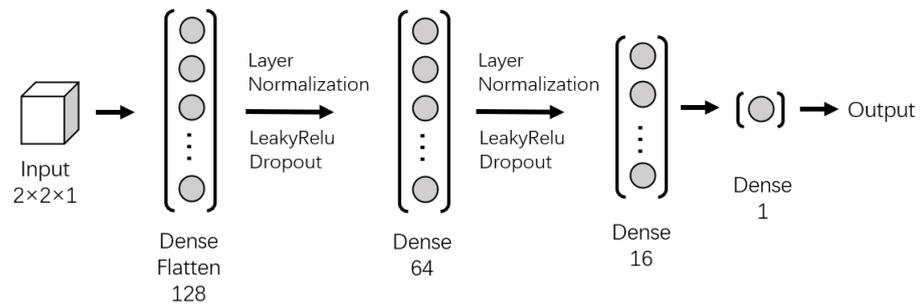


Figure 15. The structure of the discriminator for the eye dataset.

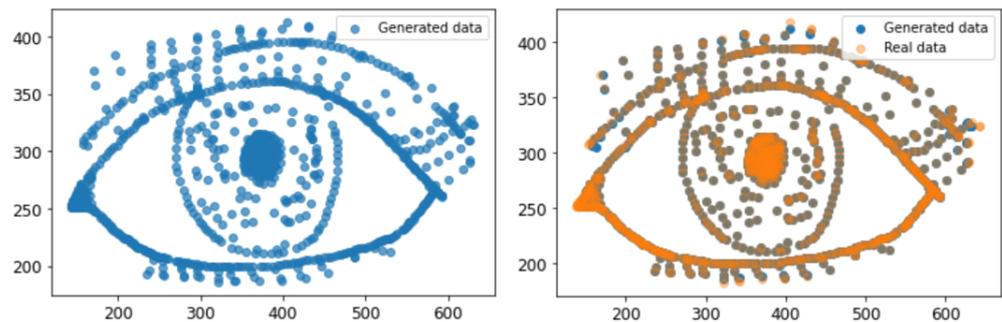


Figure 16. The eye generated by the WGAN (left) and the comparison between the real data and the generated data (right) using the constrained input method described in Algorithm 1.

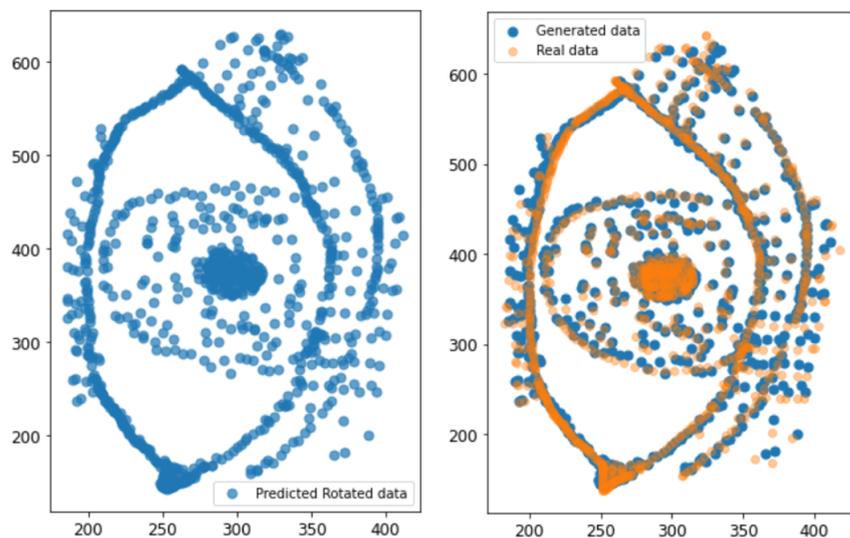


Figure 17. The rotated eye predicted using Algorithm 1 (left) and the comparison between the real rotated data and the predicted data (right).

3.4.2. Co-Varying Spiral Dataset

In many applications, variables often co-vary, in other words, a change in one variable is typically reflected by a change in another variable. In this work, we use a two dimensional spirals dataset as a benchmark to compare the capability of both GAN and WGAN. x and y are the variables that define the spiral at 20 different z levels which are equally spaced with $z \in [0, 4]$. Thus there are 20 pairs of x, y coordinates as the output of the MOR-GAN.

The structure of the model and the hyperparameters of each layer used in this section are displayed in Figures 18 and 19, and Table 3.

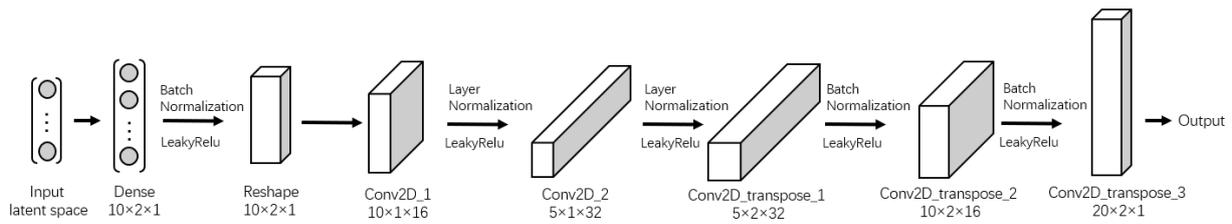


Figure 18. The structure of the generator for the Co-Varying Spiral problem.

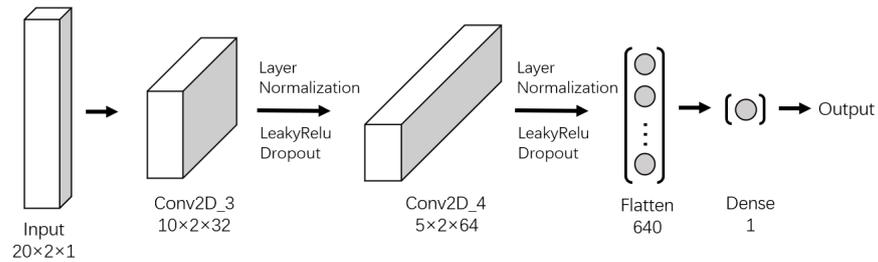


Figure 19. The structure of the discriminator for the Co-Varying Spiral problem.

Table 3. Hyperparameters used in the construction of the convolutional neural network.

Layer	Kernel Size	Strides	Padding	Use Bias
Conv2D_1	(8,2)	(1,2)	same	True
Conv2D_2	(8,2)	(2,1)	same	True
Conv2D_transpose_1	(8,2)	(1,2)	same	False
Conv2D_transpose_2	(8,2)	(2,1)	same	False
Conv2D_transpose_3	(8,2)	(2,1)	same	False
Conv2D_3	(8,2)	(2,1)	same	True
Conv2D_4	(8,2)	(2,1)	same	True

The three-dimensional spiral curves dataset is generated based on the equations below:

$$x = r \sin \theta, \tag{9}$$

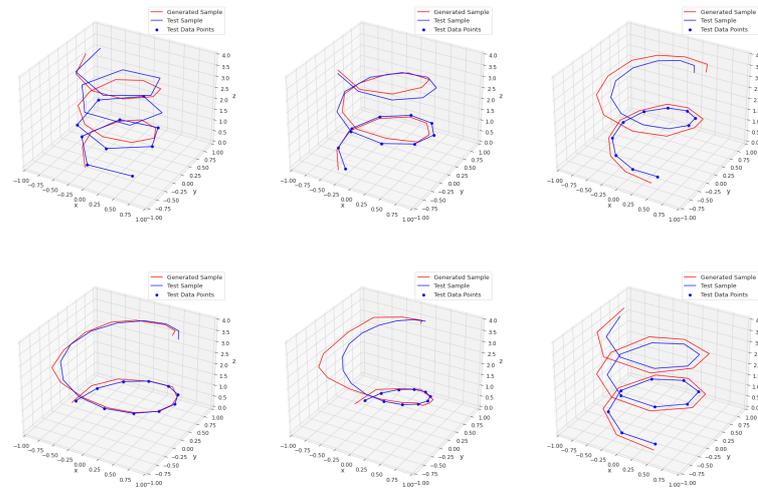
$$y = r \cos \theta, \tag{10}$$

$$z = 4 \left(\frac{\theta - a}{b - a} \right), \tag{11}$$

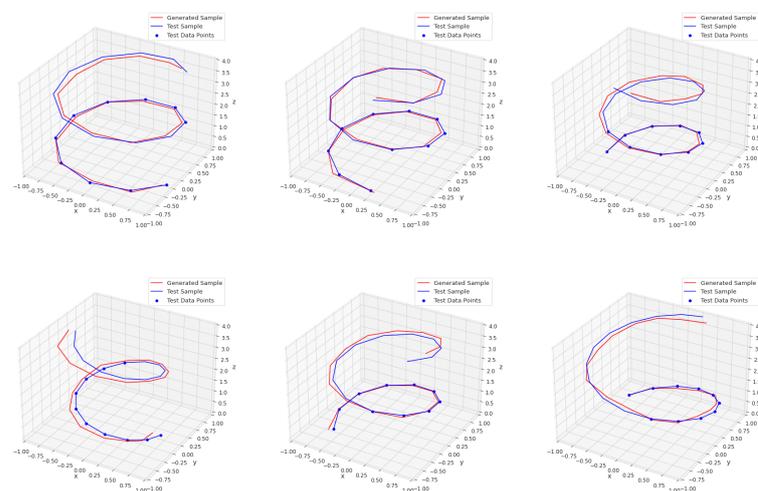
where $\theta \in [a, b]$, $a = 4\pi x_1 - 2\pi$ and $b = 4\pi x_2 + 2\pi$ for x_1, x_2 chosen randomly from the unit interval, and the radius r is chosen randomly from the interval $[0.6, 1]$. For each spiral, r, x_1 and x_2 are chosen at random, and 20 equally-spaced values for θ are chosen from the interval $[a, b]$ to generate the curves shown in Figure 20.

Figure 20 shows the predictions made by the MOR-GAN. The first 10 data points in the spiral, shown as solid blue dots, are used to predict the next 10 data points in the spiral, produced by constraining the output using Algorithm 1. The real spiral is given by the blue line and the spiral generated by the algorithm constraining the first 10 samples is given by the red line. Three different sizes of latent spaces are used and it can be observed that all latent spaces give reasonable reconstructions of the real spirals, therefore demonstrating that the reconstruction reliability of the shape of the curve does not vary much with the increasing dimension of latent space. Figure 20 also shows that the MOR-GAN can learn the structure of the input data and can recreate the shapes (which are spirals in this case) with approximate distributions (which are annular distributions representing the start of the spirals in this case).

Latent space dimension = 3



Latent space dimension = 6



Latent space dimension = 12

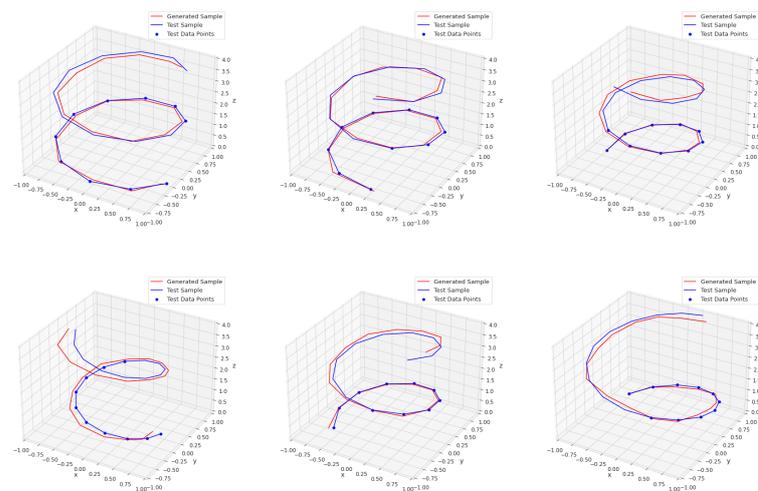


Figure 20. The figures above show the generated data using the prediction function on real samples (rows 1, 3, 5) and test samples (rows 2, 4, 6) when the size of the latent space is 3, 6 and 12 respectively. The blue lines indicate the real spiral, the solid blue dots show the 10 data points that are constrained using Algorithm 1 and the red lines show the spiral produced by the generator for these 10 constrained points.

4. Silver Nanoparticle Data

We now explore the application of WGAN to real-world data regression. Reference [32] explores the effects on cells from the lungs of four types of silver nanoparticles (AgNPs): silver nanospheres (AgNS) of diameter 20 nm and 50 nm; short silver nanowires (s-AgNWs) of length 1.5 μm and diameter 72 nm; and long silver nanowires (l-AgNWs) of length 10 μm and diameter 72 nm. Silver nanoparticles are increasingly used in consumer products and reports state that up to 14% of products containing AgNPs will release these nanoparticles into ambient air [33,34] where they can be inhaled into the lungs of workers and consumers. The work in [32] explores the influence of the nanoparticles on airway smooth muscle (ASM) cells, which are an important component of the airways in the lungs, being responsible for narrowing the airways in conditions such as asthma. Bronchi and tracheas from transplant donor lungs were dissected to obtain the cells. These cells were serum-starved overnight and then incubated with 20 nm or 50 nm AgNSs, or s-AgNWs (5 μg mL⁻¹ or 25 μg mL⁻¹) or Ag⁺ ions (0.25 μg mL⁻¹ or 25 μg mL⁻¹) for 24 or 72 h. Change in cell viability assessed by a reduction assay and change in cell proliferation assessed by the rate of DNA synthesis were both measured, and the results are reproduced in Figure 21. Cell viability is defined as the number of live, healthy cells in a sample.

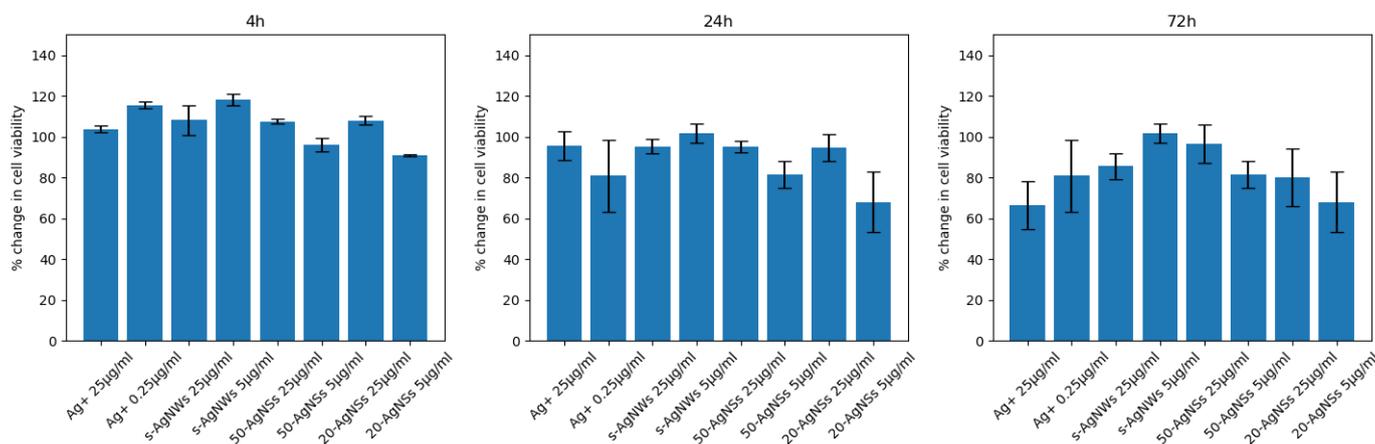


Figure 21. Concentration and time-dependent effect of AgNSs and AgNWs, and Ag⁺ ions on ASM cell viability after 4 h, 24 h and 72 h. The bars represent mean values of 3 ASM cell donors and the whiskers indicate standard error of the mean (SEM). The data is expressed as percentage change with respect to the untreated control. This plot was formed from the dataset also reported in [32].

The data from [32] contains four different molecules analysed at two concentrations at three different times. The molecules were given numerical values based on their specific surface area, defined as the total surface area of a material per unit of mass. This can be seen in Table 4:

Table 4. Specific surface area of the particles formed from different molecules which form independent variables for the WGAN regression.

Molecule	Specific Surface Area m ² g ⁻¹
Ag ⁺	4.4
s-AgNWs	4.6
50 nm AgNSs	6
20 nm AgNSs	40.4

The generator part of the WGAN was trained to produce four outputs: the specific surface area of the particles containing a specific molecule, the concentration level, the time level (we sample the response at 3 time levels: 4 h, 24 h and 72 h) and the response

(change in cell viability). All four outputs were scaled to be between 0 and 1. The WGAN architecture can be seen in Figure 22.

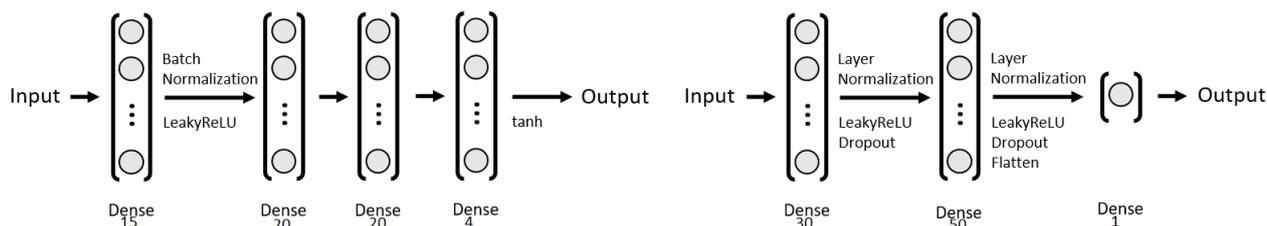


Figure 22. The structure of the generator (left) and critic (right) for modelling the silver data in this section.

Figure 23 contains the predictions made by the WGAN for cell viability, given time level, concentration level and surface area taken from the original study. For each combination of parameters (time level, concentration level and surface area), 10 predictions are made using the prediction Algorithm 1, minimising the error in the numerical value associated with a molecule, the concentration level and the time of interest. It can be observed that the mean of the predictions is close to the mean of the assessment. The average p -value for these predictions made is <0.2 .

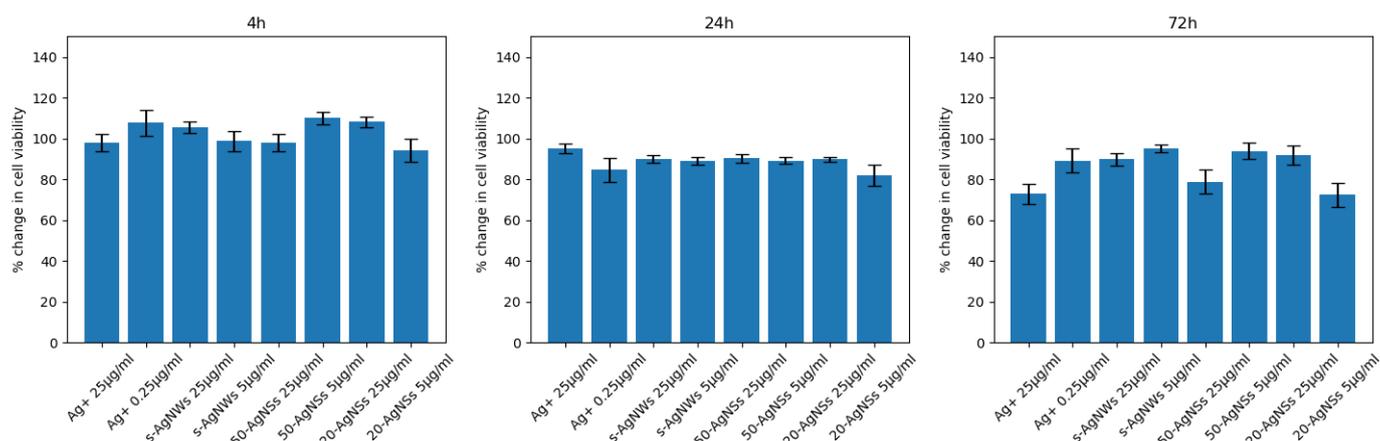


Figure 23. Concentration and time-dependent predictions of AgNSs and AgNWSs, and Ag⁺ ions on ASM cell viability after 4 h, 24 h and 72 h. The bars represent mean values of 10 predictions made using a WGAN and the whiskers indicate standard error of the mean (SEM).

5. Execution Time of Method

Presented in Table 5 is an overview of the execution time of the method.

Table 5. First column contains the type of dataset, second column contains how long randomly sampling 4000 of the posterior of the GPR took in seconds, third column contains how long it took to randomly sample 4000 points of the latent space of the WGAN took in seconds and fourth column is how long the WGAN took to run the prediction Algorithm 1 for 1000 iterations.

Dataset	GPR	WGAN—Random	WGAN—Constrained
sine wave	0.0325 s	0.328 s	2.412 s
heteroscedastic	0.0642 s	0.144 s	2.737 s
circle	0.0444 s	0.198 s	2.543 s
helix	0.0774 s	0.231 s	3.528 s
silver nanoparticle	0.0623 s	0.261 s	4.601 s

Table 5 contains the time taken, in seconds, for randomly sampling 4000 points of the GPR posterior and the WGAN latent space for different datasets. It can be observed that

the time taken to randomly sample does not increase significantly as the number of input parameters increases but sampling the GPR posterior is an order of magnitude faster than sampling the WGAN latent space.

The third column of Table 5 contains the time taken to run the prediction algorithm 1 for 1000 iterations. The values here are meant as a form of comparison, Algorithm 1 incorporating a convergence criteria would reduce the amount of iterations but would make comparison less clear. There is a notable increase in time taken for the algorithm to be applied to the datasets with a larger number of independent variables.

6. Conclusions and Future Work

In this paper, we demonstrate that Generative Adversarial Networks (GANs) can perform well for a number of regression tasks, sometimes outperforming a model based on state-of-the-art Gaussian Process Regression (GPR). The particular model used is a Wasserstein GAN (WGAN), which can be easier to train than a standard GAN. For simple regression and multiple regression tasks, both GAN and GPR perform well, although for the dataset which has variable uncertainty (modelled as heteroscedastic noise), the GPR fails to learn any variation in uncertainty, whereas the GAN captures this variation well. Also, for the more challenging problem of multi-modal distributions, the GPR struggles to learn the distribution whereas the GAN is able to reproduce the distribution very well. Furthermore, for multi-output regression, the WGAN also demonstrated good performance, showing that the GAN is able to capture the covariance information between all the output variables (which includes the independent and dependent variables of the regression problem).

Although the GPR can be modified for improved performance on specific types of data (such as heteroscedastic noise and multi-output regression), we wanted to highlight, here, that the WGAN needs no modification for these problems: one single WGAN model can perform well for all the datasets with which we tested the models.

Novelties of the work include using a GAN for regression; being able to apply this model to multi-modal data and multi-output regression (MOR-GAN) tasks with no fundamental modifications; the presentation of a prediction algorithm to be used with the trained GAN in order to predict a response for a given independent variable; using the critic to provide a confidence level of the predictions made by the generator, which could ultimately be used to help determine where more data is needed.

In the future, the methods developed here could be applied to imaging, for example, where, when there is missing data from an image or video, we could attempt to re-construct the missing parts. Being able to reconstruct this image with specified uncertainties would be useful. In modelling, the approach could be applied in high-dimensional space (with applications across computational physics e.g. Computational Fluid Dynamics) to perform data assimilation and analyse remaining uncertainties in the modelling, see Silva et al. [35]. Using the confidence level provided by the discriminator in such applications could determine where better models are needed or where coarser models (that are faster) can be used. Performing a sensitivity analysis of the discriminator could also indicate where the model is most error prone and thus where it needs to be improved.

Author Contributions: Conceptualisation, C.C.P., A.E.P. and K.F.C.; methodology, C.C.P., T.R.F.P. and C.E.H.; software, T.R.F.P., E.B., Q.L. and L.H.; data curation, A.E.P.; writing—original draft preparation, T.R.F.P., E.B. and C.E.H.; writing—review and editing, C.C.P., C.E.H. and K.F.C.; funding acquisition, C.C.P. and K.F.C. All authors read and agreed to the published version of the manuscript.

Funding: The authors would like to acknowledge the following EPSRC grants: INHALE, Health assessment across biological length scales (EP/T003189/1); RELIANT, Risk Evaluation fAst iN-telligent Tool for COVID19 (EP/V036777/1); MUFFINS, MUltiphase Flow-induced Fluid-flexible structure InteractioN in Subsea applications (EP/P033180/1); the PREMIERE programme grant (EP/T000414/1) and MAGIC, Managing Air for Green Inner Cities (EP/N010221/1).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Acknowledgments: We would like to thank the reviewers for their comments which have improved the article.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Nomenclature

See Table A1 for a description of the nomenclature used in this article.

Table A1. Nomenclature used in the paper.

Section 2 and Algorithm 2 from Section 3	
G, D	generator and discriminator (or critic) networks (for GANs, D is referred to as the discriminator, for WGANs it is referred to as the critic)
α	latent variables
L, V	the loss function and function describing the two player min-max game
x, \tilde{x}	samples from real and generated data
$\mathbb{P}_r, \mathbb{P}_g$	distributions for the real data and the generated data
p_α	distribution of the latent variables
W	Wasserstein distance between distributions
$\gamma(x, y)$	a joint distribution
\mathcal{D}	set of 1-Lipschitz functions
\hat{x}	a linear combination of a real sample and a generated sample (at which the gradient penalty will be imposed)
λ	gradient penalty
ε	mismatch between desired (partial) output of GAN and actual (partial) output of GAN
x, y	independent and dependent variables
x_p, y_p	particular values of the independent and dependent variables
ε	random number
\mathcal{U}	Uniform probability distribution
α^ℓ	learning rate
β_1, β_2	optimiser hyperparameters
n_{critic}	number of iterations of the critic
m	batch size
N	number of iterations
Section 3	
x, y, z	independent and dependent variables
x_1, x_2, y_1, y_2	independent and dependent variables
θ	angle
η	a scalar controlling the amount of noise
$\phi \sim \mathcal{N}(\mu, \sigma)$	random variable (noise) sampled from a Gaussian distribution \mathcal{N} with mean μ and standard deviation σ
h	distance function

References

1. Borchani, H.; Varando, G.; Bielza, C.; Larrañaga, P. A survey on multi-output regression. *WIREs Data Min. Knowl. Discov.* **2015**, *5*, 216–233. [CrossRef]
2. Xu, D.; Shi, Y.; Tsang, I.W.; Ong, Y.S.; Gong, C.; Shen, X. Survey on Multi-Output Learning. *IEEE Trans. Neural Netw. Learn. Syst.* **2020**, *31*, 2409–2429. [CrossRef] [PubMed]
3. Rasmussen, C.E. Gaussian Processes in machine learning. In *Advanced Lectures on Machine Learning*; Springer: Berlin/Heidelberg, Germany, 2004; Volume 3176, pp. 63–71. [CrossRef]
4. Goodfellow, I.J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; Bengio, Y. Generative Adversarial Nets. Technical report. *arXiv* **2014**, arXiv:1406.2661v1.

5. Kazemina, S.; Baur, C.; Kuijper, A.; van Ginneken, B.; Navab, N.; Albarqouni, S.; Mukhopadhyay, A. GANs for Medical Image Analysis. *Artif. Intell. Med.* **2020**, *109*, 101938. [[CrossRef](#)] [[PubMed](#)]
6. Wang, K.; Gou, C.; Duan, Y.; Lin, Y.; Zheng, X.; Wang, F.Y. Generative adversarial networks: Introduction and outlook. *IEEE/CAA J. Autom. Sin.* **2017**, *4*, 588–598. [[CrossRef](#)]
7. Radford, A.; Metz, L.; Chintala, S. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *arXiv* **2015**, arXiv:1511.06434. [[CrossRef](#)]
8. Kunfeng, W.; Yue, L.; Yutong, W.; Fei-Yue, W. Parallel imaging: A unified theoretical framework for image generation. In Proceedings of the 2017 Chinese Automation Congress, CAC 2017, Jinan, China, 20–22 October 2017; pp. 7687–7692. [[CrossRef](#)]
9. Zhang, K.; Kang, Q.; Wang, X.; Zhou, M.; Li, S. A visual domain adaptation method based on enhanced subspace distribution matching. In Proceedings of the ICNSC 2018—15th IEEE International Conference on Networking, Sensing and Control, Zhuhai, China, 27–29 March 2018; pp. 1–6. [[CrossRef](#)]
10. Jolaade, M.; Silva, V.L.; Heaney, C.E.; Pain, C.C. Generative Networks Applied to Model Fluid Flows. In Proceedings of the International Conference on Computational Science, London, UK, 21–23 June 2022; Springer: Berlin/Heidelberg, Germany, 2022; pp. 742–755. [[CrossRef](#)]
11. Salimans, T.; Goodfellow, I.; Zaremba, W.; Cheung, V.; Radford, A.; Chen, X. Improved Techniques for Training GANs. Technical report. *arXiv* **2017**, arXiv:1606.03498.
12. Arjovsky, M.; Chintala, S.; Bottou, L. Wasserstein GAN. Technical report. *arXiv* **2017**, arXiv:1701.07875
13. Gulrajani, I.; Ahmed, F.; Arjovsky, M.; Dumoulin, V.; Courville, A. Improved Training of Wasserstein GANs Montreal Institute for Learning Algorithms. Technical report. *arXiv* **2017**, arXiv:1704.00028
14. Barnett, S.A. Convergence Problems with Generative Adversarial Networks (GANs) A dissertation presented for CCD Dissertations on a Mathematical Topic. Technical report. *arXiv* **2018**, arXiv:1806.11382
15. Aggarwal, K.; Kirchmeyer, M.; Yadav, P.; Keerthi, S.S.; Gallinari, P. Regression with Conditional GAN. Technical report. *arXiv* **2019**, arXiv:1905.12868. [[CrossRef](#)]
16. McDermott, M.B.A.; Yan, T.; Naumann, T.; Hunt, N.; Suresh, H.; Szolovits, P.; Ghassemi, M. Semi-Supervised Biomedical Translation with Cycle Wasserstein Regression GANs. In Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence, New Orleans, LA, USA, 2–7 February 2018.
17. Schulz, E.; Speekenbrink, M.; Krause, A. A tutorial on Gaussian process regression: Modelling, exploring, and exploiting functions. *J. Math. Psychol.* **2018**, *85*, 1–16. [[CrossRef](#)]
18. Rasmussen, C.; Williams, C. *Gaussian Process for Machine Learning*; MIT Press: Cambridge, MA, USA, 2006.
19. Silva, V.L.; Heaney, C.E.; Li, Y.; Pain, C.C. Data Assimilation Predictive GAN (DA-PredGAN): Applied to determine the spread of COVID-19. *arXiv* **2021**, arXiv:2105.07729.
20. Wang, S.; Tarroni, G.; Qin, C.; Mo, Y.; Dai, C.; Chen, C.; Glocker, B.; Guo, Y.; Rueckert, D.; Bai, W. Deep generative model-based quality control for cardiac MRI segmentation. In Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention, Lima, Peru, 4–8 October 2020; Springer: Berlin/Heidelberg, Germany, 2020; pp. 88–97.
21. Le, Q.V.; Smola, A.J.; Canu, S. Heteroscedastic Gaussian process regression. In Proceedings of the ICML 2005—The 22nd International Conference on Machine Learning, Bonn, Germany, 7–11 August 2005; ACM Press: New York, NY, USA, 2005; pp. 489–496. [[CrossRef](#)]
22. Kim, H.C.; Lee, J. Clustering based on Gaussian processes. *Neural Comput.* **2007**, *19*, 3088–3107. [[CrossRef](#)]
23. Kolmogorov, A.N. Interpolation and extrapolation of stationary random sequences. In *Selected Works of A. N. Kolmogorov*; Springer: Dordrecht, The Netherlands, 1992.
24. Wiener, N. *Extrapolation, Interpolation and Smoothing of Stationary Time Series*; MIT Press: Cambridge, MA, USA, 1949.
25. Sacks, J.; William, J.; Welch, T.J.M.; Wynn, H.P. *Design and Analysis of Computer Experiments*; Institute of Mathematical Statistics: Hayward, CA, USA, 1989. [[CrossRef](#)]
26. GPY. GPY: A Gaussian Process Framework in Python. 2012. Available online: <http://github.com/SheffieldML/GPY> (accessed on 20 December 2020).
27. Chollet, F. Keras. 2015. Available online: <https://github.com/fchollet/keras> (accessed on 20 December 2020).
28. Smirnov, N.V. On the estimation of the discrepancy between empirical curves of distribution for two independent samples. *Bull. Math. Univ. Moscou* **1939**, *2*, 3–14.
29. Mann, H.B.; Whitney, D.R. On a test of whether one of two random variables is stochastically larger than the other. *Ann. Math. Stat.* **1947**, *18*, 50–60. [[CrossRef](#)]
30. Virtanen, P.; Gommers, R.; Oliphant, T.E.; Haberland, M.; Reddy, T.; Cournapeau, D.; Burovski, E.; Peterson, P.; Weckesser, W.; Bright, J.; et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat. Methods* **2020**, *17*, 261–272. [[CrossRef](#)]
31. Gulrajani, I.; Ahmed, F.; Arjovsky, M.; Dumoulin, V.; Courville, A.C. Improved Training of Wasserstein GANs. In *Advances in Neural Information Processing Systems 30*; Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R., Eds.; Curran Associates, Inc.: Nice, France, 2017; pp. 5767–5777.
32. Michaeloudes, C.; Seiffert, J.; Chen, S.; Ruenraroengsak, P.; Bey, L.; Theodorou, I.G.; Ryan, M.; Cui, X.; Zhang, J.; Shaffer, M.; et al. Effect of silver nanospheres and nanowires on human airway smooth muscle cells: Role of sulfidatio. *Nanoscale Adv.* **2020**, *2*, 5635–5647. [[CrossRef](#)]

33. Quadros, M.E.; Marr, L.C. Silver nanoparticles and total aerosols emitted by nanotechnology-related consumer spray products. *Environ. Sci. Technol.* **2011**, *45*, 10713–10719. [[CrossRef](#)]
34. Benn, T.; Cavanagh, B.; Hristovski, K.; Posner, J.D.; Westerhoff, P. The Release of Nanosilver from Consumer Products Used in the Home. *J. Environ. Qual.* **2010**, *39*, 1875–1882. [[CrossRef](#)]
35. Silva, V.L.S.; Heaney, C.E.; Pain, C.C. GAN for time series prediction, data assimilation and uncertainty quantification. *arXiv* **2021**, arXiv:2105.13859. [[CrossRef](#)]