

Mixture Density Conditional Generative Adversarial Network Models (MD-CGAN)

Jaleh Zand & Stephen Roberts
Machine Learning Research Group & Oxford-Man Institute of Quantitative
Finance, University of Oxford
{jz,sjrob}@robots.ox.ac.uk

Abstract

Generative Adversarial Networks (GANs) have gained significant attention in recent years, with impressive applications highlighted in computer vision in particular. Compared to such examples, however, there have been more limited applications of GANs to time series modelling, including forecasting. In this work, we present the *Mixture Density Conditional Generative Adversarial Model* (MD-CGAN), with a focus on time series forecasting. We show that our model is capable of estimating a probabilistic posterior distribution over forecasts and that, in comparison to a set of benchmark methods, the MD-CGAN model performs well, particularly in situations where noise is a significant component of the observed time series. Further, by using a Gaussian mixture model as the output distribution, MD-CGAN offers posterior predictions that are non-Gaussian.

1. Introduction

Generative Adversarial Networks (GANs) have been one of the many breakthroughs in Deep Learning methods in recent years. Several different variations of the model have been introduced since the method was first introduced in [1]. One of the most popular variations of this work is that of the *Conditional Generative Adversarial Network* (CGAN) [2] in which the *generator* and *discriminator* (we briefly review the GAN process in Section 2) are both conditioned on some observed information. In the application of time series forecasting, future values are conditioned on information observed from the past, either from the time series itself, some set of associated exogenous data, or a combination of the two. This conditioning addition to the GAN formalism makes the CGAN approach particularly useful as a foundational model for time series prediction. Most applications of (C)GANs, however, have been within computer vision and, to a lesser extent, in natural language processing and simulation models [3, 4, 5].

The literature on the application of any form of GAN model to problems associated with time series is, to date, limited. However, some recent literature shows the potential usefulness of the method. For example, the work reported in [6] applies a recurrent GAN to generate realistic, synthetic, medical data

series. In [7] a GAN is used to forecast high-frequency stock data and in [8] GANs are used to generate missing values in incomplete time series. We note that the GAN models used in all these applications make *point estimates* for the forecast. Although a perfectly valid approach, and one that has a long history in time series forecasting, we argue that *probabilistic* forecasts are a prerequisite in many application domains, in which knowledge of the predictive uncertainty is as vital as the prediction itself. In this work, we expand on the CGAN algorithm to allow a full predictive probability distribution, rather than a point value. To obtain richer predictive densities we model the posterior distribution using a finite Gaussian Mixture model (GMM). Although we find only occasional evidence that such non-Gaussian predictions offer significant benefits, we note that producing them is not much more costly than single Gaussian predictive distributions, and so present our approach as a more general multi-component model.

1.1. Related work

We here briefly review recent literature which is close to our approach. We start by noting the work of [9], in which a mixture of GAN models is proposed as a data clustering model. Although clearly related, this is somewhat different to our approach, in which we use a single GAN generator, linked to the hyperparameters of the posterior mixture model, rather than a mixture of GAN models. Furthermore, our goal is forecasting rather than unsupervised data classification. The approaches advocated in [10] and [11] propose a latent space, used for sampling latent vectors in the GAN, formulated via a Gaussian mixture. The latter replaces, in these papers, the single Gaussian distribution used for such generation in standard GAN models. In both these papers, the generator and discriminator have a similar structure to a standard GAN. In [12] a mixture model is used, but for the discriminator alone, with the generator being that of a standard GAN model; we note the difference to our approach, in which the *generator* is a GMM. Finally, [13] compare (standard) GAN models to GMMs for image generation. The authors show that GANs are superior in their ability to generate sharp images, but note that mixture models offer more efficient inference. They propose a combination of the two and introduce a GAN model in which the GAN generator is a mixture model. However, the sample generator still makes point estimates from the multimodal distribution in order to retain a discriminator function the same as that of a standard GAN model. We offer discriminator extensions which allow the GAN process to operate on the full (multi-modal) posterior distribution.

1.2. Paper structure

The rest of this paper is set out as follows: in Section 2 we provide a brief overview of the (C)GAN model, introducing the key concepts. In Section 3 we present the structure of the MD-CGAN model. In Section 4 we test the model on a variety of ‘real-world’ datasets and discuss the results. Finally, in Section 5, we conclude.

2. The GAN and CGAN Model

The goal of the GAN model is to estimate a generative model using an *adversarial* process [1]. This is achieved by simultaneously training two models. Firstly, a generative model G , that in the case of data forecasting learns past patterns in the data and infers the predictive values. Secondly, a discriminative model D , that determines how likely a sample was to originate from the ‘true’ training data, compared to being sampled from the generator. The generator is hence matched against an adversary, the discriminator (whose goal is to detect the difference between a true data sample and one created by the generator). Components of the model are then trained (via an optimization process) to maximize the probability of the discriminator being unable to distinguish true from generated data samples. Typically, including the approach we take here, the generator and the discriminator are both constructed as multilayer perceptrons, with stochastic gradient methods being employed to obtain optimization.

We start by formulating the definitions which we use in common across the GAN models we test. We consider a time series, y_t . Our aim is to estimate the forecast of some $y_{t'}^f$, conditioned on a set of observations which we denote \mathbf{x}_t . The inputs to the generator network are \mathbf{x}_t and \mathbf{z}_g , where \mathbf{z}_g is a collection of samples from a normal distribution, $p(z_n)_g = \mathcal{N}(0, \text{var}_{\text{data}})$. During model training, the output from the generator, $y_{t'}^f$, as well as the true forecast sample $y_{t'}$, are fed to the discriminator, whose role is to discriminate between them i.e. to identify $y_{t'}^f$ as the ‘fake’ sample.

In an unconditioned GAN model, there is no control over the data that is generated. In the CGAN model, in contrast, by conditioning the model on additional information, it is possible to direct the data generation process [2]. Schematics of the GAN and CGAN models are depicted in Figure 1.

3. The MD-CGAN Model Framework

As with the GAN and CGAN methods, we consider a time series, y_t . Our aim now is to infer the *posterior distribution* over some $y_{t'} > t$, conditioned on the set of observations, \mathbf{x}_t . In order to form the posterior distribution we model the *full conditional density* $p(y_{t'}|\mathbf{x}_t)$ as an adversarial network. To achieve this we use a Mixture Density Network (MDN) model [14] for the generator G . The inputs to the generator network are as per the CGAN approach, \mathbf{x}_t and \mathbf{z}_g , where \mathbf{z}_g is, as before, a collection of samples from a normal distribution, $p(z_n)_g = \mathcal{N}(0, \text{var}_{\text{data}})$. The outputs of $G_{t'}(\mathbf{x}_t, \mathbf{z}_g)$ are now, however, the parameters of the Gaussian mixture posterior over the forecast. This mixture has mixing coefficient, standard deviation and mean for the i -th component denoted as α_i , σ_i , and μ_i respectively. As first proposed in [14], we achieve this by using latent variables $\mathbf{s} = \{\mathbf{s}_\alpha, \mathbf{s}_\sigma, \mathbf{s}_\mu\}$, conditioned on the inputs. The mapping from $[\mathbf{x}_t, \mathbf{z}_g] \mapsto \mathbf{s} \mapsto \{\alpha_i, \sigma_i, \mu_i\}$ is modelled via our network. As the mixings must satisfy $\sum_i \alpha_i = 1$, we map \mathbf{s}_α to α via the *softmax* function, where $\alpha_i = \frac{\exp(s_{\alpha,i})}{\sum_{i'} \exp(s_{\alpha,i'})}$. The elements of σ are strictly positive so we adopt, $\sigma_i = \exp(s_{\sigma,i})$. Finally the means can be mapped

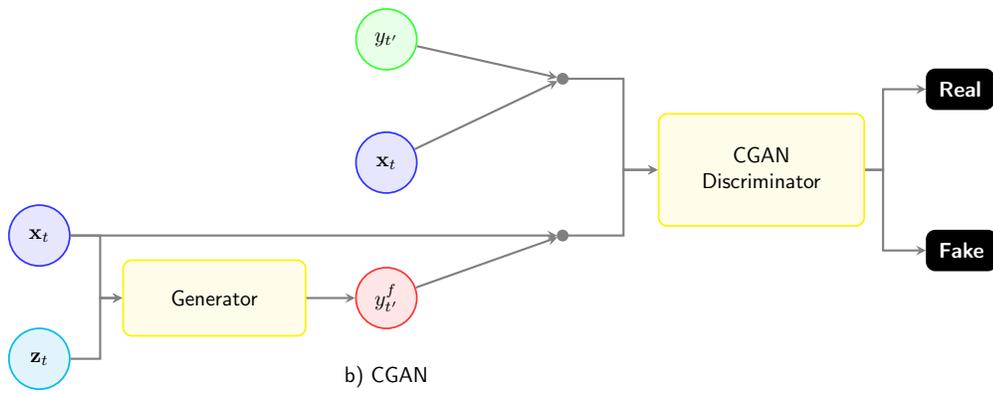
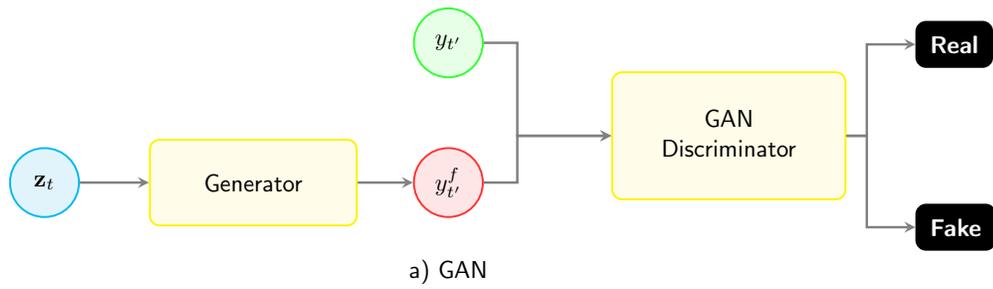


Figure 1: Schematic of (a) GAN and (b) CGAN models, showing Generator and Discriminator components and associated variables.

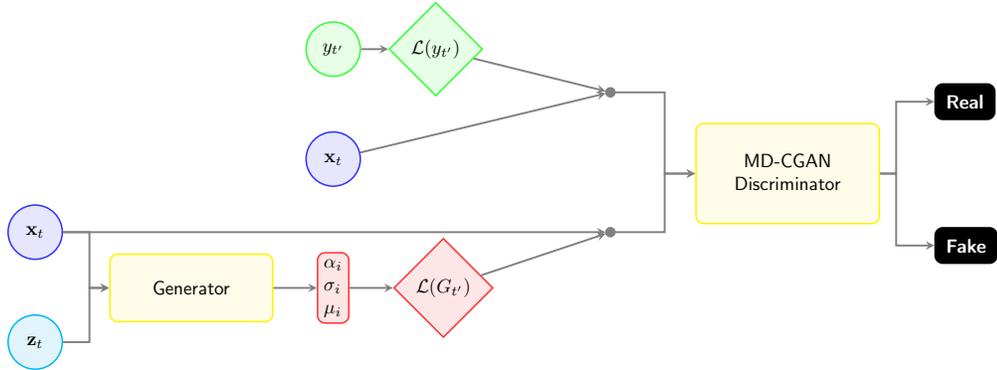


Figure 2: Schematic of the MD-CGAN model. We note that the discriminator in the MD-CGAN model has a different loss function and structure to the GAN and CGAN models.

directly from the latent variables, hence $\mu_i = s_{\mu,i}$. Schematically, the MD-CGAN method is depicted in Figure 2.

The above formalism allows us to directly model the predictive likelihood conditioned on an input, and the likelihood of G , conditioned on the observations \mathbf{x}_t and samples \mathbf{z}_g as:

$$\mathcal{L}(G_{t'}(\mathbf{x}_t, \mathbf{z}_g)) = \sum_{i=1}^m \alpha_i(\mathbf{x}_t, \mathbf{z}_g) \mathcal{N}_i(y_{t'} | \mu_i(\mathbf{x}_t, \mathbf{z}_g), \sigma_i(\mathbf{x}_t, \mathbf{z}_g)) \quad (1)$$

where m is the number of mixture components.

As in the CGAN model, the discriminator, D , is also conditioned on \mathbf{x}_t . The input to the discriminator model is, by design, $\mathbf{x}_t \sqrt{2\pi} \sigma_a \mathcal{L}(y_{t'})$, where σ_a is the standard deviation of the set of observed y_t ¹. For true values of $y_{t'}$, $\sqrt{2\pi} \sigma_a \mathcal{L}(y_{t'})$ is maximized. The generator tries to ‘fool’ the discriminator by generating $G_{t'}$ such that the $\sqrt{2\pi} \sigma_a \mathcal{L}(G_{t'})$ is maximized. The loss function for the generator, L_G is as in Equation 2, The discriminator network, on the other hand, attempts to differentiate between true $y_{t'}$ values and the pseudo-values created by the generator. The loss function for the discriminator, L_D is as in Equation 3. We note that the lowest value of the discriminator loss is achieved when $\sqrt{2\pi} \sigma_a \mathcal{L}(y_{t'})$ is maximal (unity) and $\mathcal{L}(G_{t'}(\mathbf{x}_t, \mathbf{z}_g))$ is minimal (zero).

$$L_G = \mathbb{E}_{z \sim P_z(z)} [-\mathcal{L}(G_{t'}(\mathbf{x}_t, \mathbf{z}_g))] \quad (2)$$

$$L_D = \mathbb{E}_{y \sim P_{\text{data}}(y)} [\|\mathbf{x}_t \sqrt{2\pi} \sigma_a \mathcal{L}(y_{t'}) - \mathbf{x}_t\|^2] + \mathbb{E}_{z \sim P_z(z)} [\|\mathbf{x}_t \sqrt{2\pi} \sigma_a \mathcal{L}(G_{t'}(\mathbf{x}_t, \mathbf{z}_g))\|^2] \quad (3)$$

Our algorithm thus follows the steps in Algorithm 1 below.

¹We note that the GAN approach is not sensitive, within reason, to this value (it is, in effect, a constant in the update equations) and we discuss its setting later in the paper.

Algorithm 1 MD-CGAN Algorithm

- 1: **for** number of training iterations **do**
- 2: **for** j steps **do** **do**
- 3: Sample N noise samples, $\{\mathbf{z}^1, \dots, \mathbf{z}^N\}$ from $p_g(\mathbf{z})$
- 4: Sample N data points, $\{\mathbf{x}^1, \dots, \mathbf{x}^N\}$ from $p_{\text{data}}(\mathbf{x})$
- 5: Update the discriminator by descending its stochastic gradient:

$$\nabla_{\theta_d} \sum_{n=1}^N [\|\mathbf{x}^{(n)} \sqrt{2\pi}\sigma_a \mathcal{L}(y^{(n)}) - \mathbf{x}^{(n)}\|^2 + \|\mathbf{x}^{(n)} \sqrt{2\pi}\sigma_a \mathcal{L}(G(\mathbf{z}^{(n)}, \mathbf{x}^{(n)}))\|^2]$$

- 6: **end for**
- 7: Sample N noise samples, $\{\mathbf{z}^1, \dots, \mathbf{z}^N\}$ from $p_g(\mathbf{z})$
- 8: Update the generator by descending its stochastic gradient:

$$\nabla_{\theta_g} \sum_{n=1}^N -\mathcal{L}(G(\mathbf{z}^{(n)}, \mathbf{x}^{(n)}))$$

- 9: **end for**
-

4. Experiments

4.1. Comparison with other Learning Models

To provide a range of comparisons to methods related to this work, we compare our MD-CGAN model to the following baseline methods: the Mixture Density Network model (MDN), chosen to baseline mixture density outputs against [14]; the CGAN model, chosen as a well-known GAN approach [2]; and a “standard” Multi-Layer Perceptron (MLP) neural network (SNN) as a simple, yet effective, baseline. To promote as fair comparison as possible, we use a *common* neural network architecture across core components in the models, choosing the neural network structure commonly used for GAN models in recent literature [2]. We do not alter this structure throughout our experiments and we keep to fixed hyperparameter settings (guided by those used in [2]). Figure 3 provides a schematic of the structure used. We note that, whilst the lengths of the input and output vectors are dependent on the model, the structure of all networks remains constant.

4.2. Details of implementation

All models were coded in the Python language and we use the Keras library [15] to build the neural networks. The neural networks in Figure 3 follow the structure of the CGAN model detailed in [2]. The hyperparameters in the models were set as follows: for the discriminator, both in CGAN and MD-CGAN, the dropout parameter is set to 0.4, and alpha for the leaky ReLU is set to 0.2. For the generator the dropout rate is set to 0.5 and alpha for the leaky ReLU is, again, 0.2. The parameter governing the number of neurons, n , in the dense layers of the neural network modules is set to 20. The variance parameter, σ_a ,

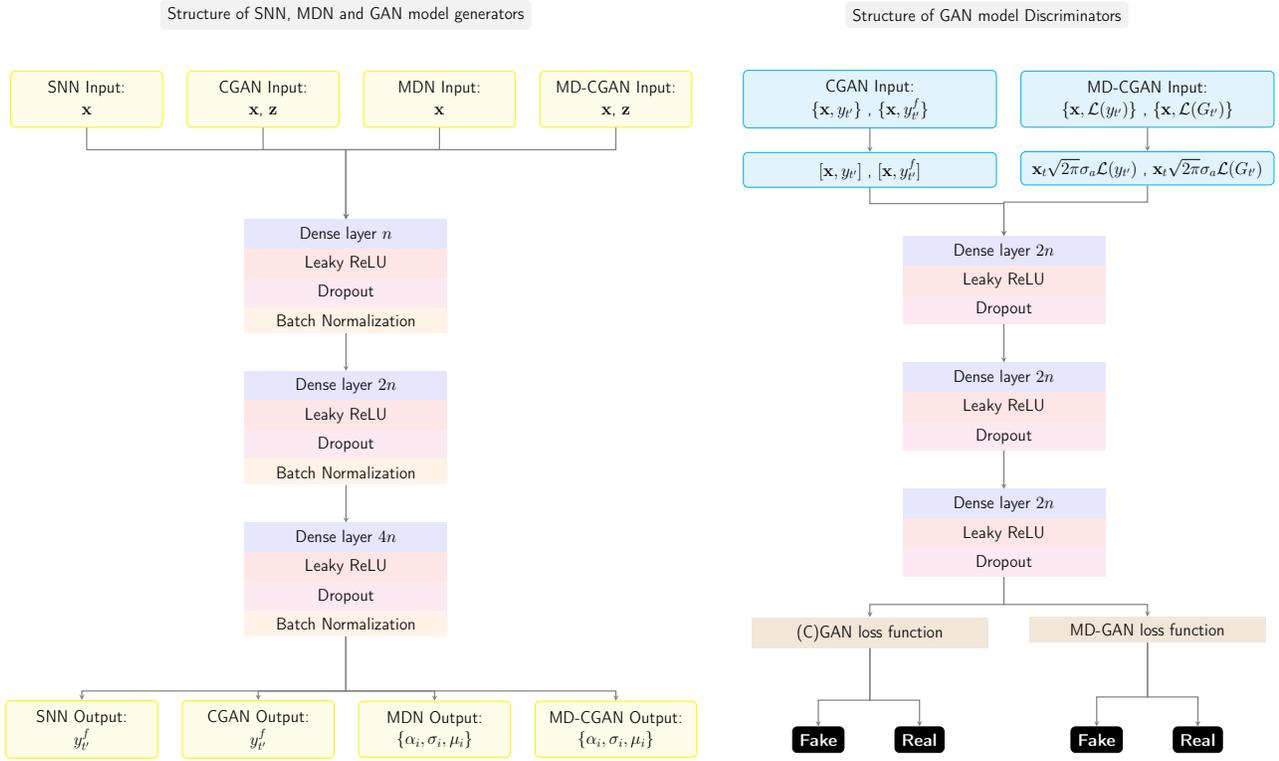


Figure 3: *Common Neural Network structure used across all models.*

in the MD-CCGAN models is set to 0.2. During training, we follow the steps of algorithm 1, and set j to 1 for all datasets. The number of training iterations is, however, specific to model and dataset. With the exception of the GAN models we monitor the errors, until they reach saturation during training. For the GAN models, we have more than one error and they are not good indicators to use ². For these two models we monitor the average error from each model’s estimation of the last 400 points of the training data and continue iteration until the errors reach saturation. In all models we optimize parameters using the Adam optimizer [16], with learning rate set to 0.001, exponential decay for the first momentum set to 0.9, exponential decay for the second momentum set to 0.999, and epsilon set to 10^{-7} .

4.3. Data

We perform initial experiments on four datasets with differing provenance; the Mackey-Glass chaotic dataset [17]; the sunspot dataset [18]; US initial jobless

²The generator and the discriminator in the GAN models, each have a loss value per iteration.

claims (USIJC, weekly intervals, [19]); and EURUSD foreign exchange daily rates (EURUSD FX rate, [20]). For each dataset, the time series is split into training and out-of-sample test sets. The training data sets in all our experiments comprise 2000 samples and all test sets consist of 400 sequential data points post the training set. All algorithms are provided as input the last k data points, set fixed to $k = 5$ for the purpose of all our experiments. We note that this value is not optimized, however, and is chosen to allow simple comparisons across methods. All data sets are pre-normalized to the $[0,1]$ interval, again to allow for simpler comparison across data and methods. We further note that both the CGAN and SNN models make *point estimate* predictions, whilst MD-CGAN and MDN estimate posterior *distributions*. To enable a simple comparison, across models, we therefore report the mean-square error (MSE) for all methods. The number of mixture components, m , in both MD-CGAN and MDN is set to unity (we vary this in Section 4.6), to further ease comparison. The most-likely value (which for $m = 1$ is merely the posterior mean) of the predictive distribution is taken as the forecast value for both the MDN and MD-CGAN models for the purposes of point-prediction error reporting.

4.4. One-step forecasting

Mackey-Glass and Sunspot time series: We start our experiments looking at one-step ahead forecasts on two well known data sets, the Mackey-Glass chaotic time series [17] and the Sunspot data set [18]. One-step forecast error comparisons are indicated in Table 1. We note that both GAN models (CGAN and MD-CGAN) are outperformed, in terms of MSE, for these datasets. We attribute this to the high signal to noise associated with these data sets and thus the low levels of observed noise in the test data. In these circumstances, the ‘adversarial advantage’ [2] of GAN approaches, which offers robustness to input perturbations (in this case noise), is less important.

In the next experiments we thus add (30% by amplitude) normally distributed noise to the test data (from a GAN perspective, these input perturbations are, in effect, treated as adversarial attacks). We note that no noise is added to the training dataset. MSE errors are presented in Table 1, and referenced as "Mackey-Glass with Noise" and "Sunspot with Noise" respectively. We note that, under this noisy-observation paradigm, GAN models perform best, confirming the adversarial robustness of the GAN models.

US initial jobless claims (USIJC) and Euro-Dollar foreign exchange (EURUSD FX) daily rate: Financial times series are highly stochastic, and we expect GAN approaches to be well-suited to forecasting in these circumstances. We consider two financial time series, namely the USIJC and EURUSD FX datasets. Test set forecasts are shown in Figure 4 and the associated MSEs in Table 1. We note that, for these data sets, the GAN approaches perform well and our method, MD-CGAN has lowest errors.

4.5. Financial forecasts over longer-horizons

One-step forecasts were presented in Subsection 4.1. Here we extend analysis of the financial data over longer horizons. All models were used to make estimates

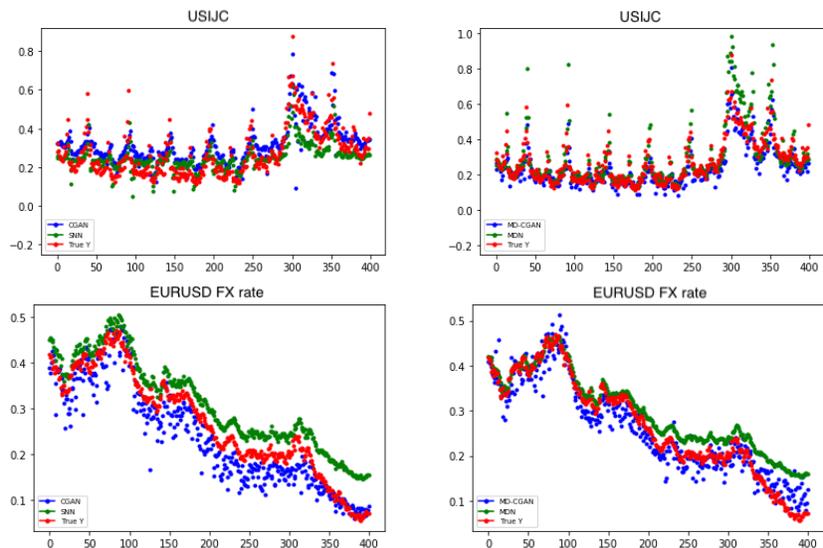


Figure 4: Comparative forecasts over USIJC and EURUSD FX time series. Left plots: Comparison of CGAN (blue) & SNN (green) predictive means and true data (red). Right plots: Comparison of MD-CGAN (blue) & MDN (green) predictive means and true data (red).

over a horizon of ten weeks for an extended set of financial time series; the USIJC, EURUSD FX rate (as previous), WTI crude oil spot prices (WTI, [21]), Henry Hub Natural Gas spot prices (Nat Gas, [21]), CBOE Volatility Index (VIX [22]), New York Harbor No. 2 Heating Oil Spot Price (Heating Oil [21]), Invesco DB US Dollar Index Bullish Fund (USD Index [23]), and iShares MSCI Brazil Small-Cap ETF (EM ETF [24]). A ten-week horizon represents a 50 step forecast for the daily datasets (FX, WTI, Nat Gas, VIX, Heating Oil, USD Index & EM ETF) and 10 steps for the weekly USIJC dataset. Further we perform comparisons against standard econometric linear models, namely a 5-th order autoregressive, AR(5), model and the martingale, or AR(0) model, in which the forecast is the last observed datum. Taking the martingale model as a baseline, we present in Table 2 the mean-square errors as the ratio to the martingale model error. We note that the MD-CGAN approach delivers ratios below unity and provides the lowest error of almost all models in this scenario.

4.6. Multi-modal posterior predictions

Finally, we compare the performance of MD-CGAN over varying numbers of mixture components. In all the previous experiments we set $m = 1$ (hence the model produced a single predictive Gaussian posterior). Here we briefly present the results for the five finance datasets with $m \in \{1, 2, 3\}$. We report (negative) log-likelihood measures (as we do not compare against point-value models in this section) and consider one-step forecasts on all the data sets. Table 3 presents the performance across data sets for varying numbers of mixture components in

Table 1: MSEs for all experiments, with standard deviations in brackets.

	Mackey-Glass	Sunspot	Mackey-Glass with Noise	Sunspot with Noise	USIJC	EURUSD FX rate
SNN	0.0014 (0.0020)	0.0154 (0.0260)	0.1640 (0.1939)	0.0536 (0.0777)	0.0070 (0.0154)	0.0022 (0.0017)
CGAN	0.0036 (0.0048)	0.0196 (0.0358)	0.0360 (0.0525)	0.0266 (0.0380)	0.0074 (0.0157)	0.0018 (0.0026)
MDN	0.0002 (0.0003)	0.0105 (0.0239)	0.1402 (0.1597)	0.0758 (0.1144)	0.0080 (0.0238)	0.0016 (0.0020)
MD-CGAN	0.0026 (0.0030)	0.0172 (0.0293)	0.0264 (0.0392)	0.0203 (0.0359)	0.0041 (0.0089)	0.0008 (0.0011)

Table 2: Ratio of model MSE to martingale baseline model over 10-week forecast horizon.

	USIJC	EURUSD FX rate	WTI	Nat Gas	VIX index	Heating Oil	USD Index	EM ETF
AR(5)	0.78	1.91	0.85	1.01	0.71	0.82	1.24	0.89
SNN	0.79	1.25	0.89	0.94	0.71	0.93	1.34	0.82
CGAN	0.77	0.85	1.53	1.07	0.91	0.54	1.37	0.69
MDN	0.84	3.48	1.48	1.13	0.77	0.89	0.68	0.81
MD-CGAN	0.73	0.76	0.80	0.82	0.66	0.59	0.54	0.65

the posterior prediction. Figure 5 shows the predicted distribution for the test datasets for Nat Gas and Vix Index.

We note performance improvement for some datasets for $m > 1$. We do not attempt to infer m , though choosing its value based on performance on a set of cross-validation data would be an option, as would enforcing regularization over the mixture model posterior, through more extensive use of Bayesian inference. We leave these extensions for future research.

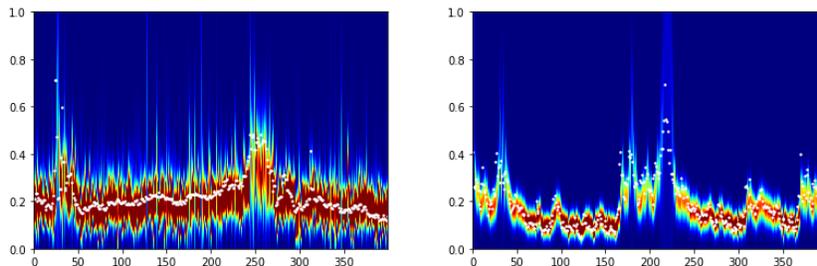


Figure 5: Estimated distributions for (left) Nat Gas with $m = 2$ and (right) VIX index with $m = 1$ over out-of-sample test datasets. True samples are shown as white dots. Red indicates high data likelihood and blue data low likelihood under the MD-CGAN model.

Table 3: Negative log-likelihood variation with m , with standard deviations in brackets.

	USIJC	EURUSD FX rate	WTI	Nat Gas	VIX index	Heating Oil	USD index	EM ETF
$m = 1$	-1.01 (0.65)	-1.79 (0.42)	-1.75 (0.27)	-1.28 (0.72)	-1.50 (0.94)	-1.63 (0.23)	-0.65 (0.38)	-1.26 (0.41)
$m = 2$	-1.05 (0.39)	-0.98 (0.33)	-1.24 (0.32)	-1.33 (0.66)	-1.39 (0.78)	-1.37 (0.24)	-0.83 (0.32)	-1.10 (0.32)
$m = 3$	-1.09 (0.34)	-0.67 (0.18)	-1.33 (0.37)	-1.25 (0.63)	-1.48 (0.92)	-1.35 (0.26)	-0.86 (0.12)	-1.09 (0.27)

5. Conclusion

In this paper we present the MD-CGAN model, which offers extensions to the CGAN [2] methodology, particularly to allow for GAN inference of a (multi-modal) posterior distribution over forecast values. In the experiments considered, we find the MD-CGAN approach outperforms other methods on all datasets in which noise is prevalent, including all the financial time series investigated, over long term forecast horizons. As a GAN model, our approach retains adversarial robustness in forecasting, which we find is most notable when noise is extensively present in data. Our method is thus particularly well suited to dealing with financial data. Furthermore, MD-CGAN can effectively estimate a flexible posterior distribution, in contrast to standard GAN models which (almost without exception) produce point value outputs. Exploiting this rich, multi-model, posterior distribution is not reported in detail here but will feature in follow-up work. In summary, the MD-CGAN model combines the advantageous features of both flexible probabilistic forecasting and GAN methods. We see this as a particularly useful approach for dealing with time series in which noise is significant and for providing robust, long-term forecasts beyond simple point estimates.

References

- [1] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial nets, in: Advances in neural information processing systems, 2014, pp. 2672–2680.
- [2] M. Mirza, S. Osindero, Conditional generative adversarial nets, arXiv preprint arXiv:1411.1784.
- [3] H. Wu, B. Gu, X. Wang, V. Pickert, B. Ji, Design and control of a bidirectional wireless charging system using GAN devices, in: 2019 IEEE Applied Power Electronics Conference and Exposition (APEC), IEEE, 2019, pp. 864–869.
- [4] J. A. Hodge, K. V. Mishra, A. I. Zaghloul, Joint multi-layer GAN-based design of tensorial RF metasurfaces, in: 2019 IEEE 29th International

- Workshop on Machine Learning for Signal Processing (MLSP), IEEE, 2019, pp. 1–6.
- [5] C. B. Barth, P. Assem, T. Foulkes, W. H. Chung, T. Modeer, Y. Lei, R. C. Pilawa-Podgurski, Design and control of a GAN-based, 13-level, flying capacitor multilevel inverter, *IEEE Journal of Emerging and Selected Topics in Power Electronics*.
 - [6] C. Esteban, S. L. Hyland, G. Rätsch, Real-valued (Medical) Time Series Generation with Recurrent Conditional GANs, arXiv preprint arXiv:1706.02633.
 - [7] X. Zhou, Z. Pan, G. Hu, S. Tang, C. Zhao, Stock market prediction on high-frequency data using generative adversarial nets, *Mathematical Problems in Engineering* 2018.
 - [8] Y. Luo, X. Cai, Y. Zhang, J. Xu, et al., Multivariate time series imputation with generative adversarial networks, in: *Advances in Neural Information Processing Systems*, 2018, pp. 1596–1607.
 - [9] Y. Yu, W.-J. Zhou, Mixture of GANs for Clustering., in: *IJCAI*, 2018, pp. 3047–3053.
 - [10] S. Gurumurthy, R. Kiran Sarvadevabhatla, R. Venkatesh Babu, Deligan: Generative adversarial networks for diverse and limited data, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 166–174.
 - [11] M. Ben-Yosef, D. Weinshall, Gaussian mixture generative adversarial networks for diverse datasets, and the unsupervised clustering of images, arXiv preprint arXiv:1808.10356.
 - [12] H. Eghbal-zadeh, W. Zellinger, G. Widmer, Mixture density generative adversarial networks, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 5820–5829.
 - [13] E. Richardson, Y. Weiss, On GANs and GMMs, in: *Advances in Neural Information Processing Systems*, 2018, pp. 5847–5858.
 - [14] C. M. Bishop, *Pattern recognition and machine learning*, Springer, 2006.
 - [15] Chollet, François and others, Keras, <https://keras.io> (2015).
 - [16] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, arXiv preprint arXiv:1412.6980.
 - [17] M. C. Mackey, L. Glass, Oscillation and chaos in physiological control systems, *Science* 197 (4300) (1977) 287–289.
 - [18] F. Clette, et al., WDC-SILSO, <http://www.sidc.be/silso/> (2015).

- [19] Bureau of Labor Statistics, United States Department of Labor, <https://www.bls.gov/> (2018).
- [20] Investing.com, Euro US Dollar Daily Price, <https://www.investing.com/currencies/eur-usd-historical-data> (2018).
- [21] U.S. Energy Information Administration, <https://www.eia.gov/> (2020).
- [22] Microsoft Corp. (MSFT), Yahoo!Finance, CBOE Volatility Index Historical Data, <https://finance.yahoo.com/quote/%5EVIX/history?p=%5EVIX> (2020).
- [23] Microsoft Corp. (MSFT), Yahoo!Finance, Invesco DB US Dollar Index Bullish Fund Historical Data, <https://finance.yahoo.com/quote/UUP/history?p=UUP> (2020).
- [24] Microsoft Corp. (MSFT), Yahoo!Finance, iShares MSCI Brazil Small-Cap ETF Historical Data, <https://finance.yahoo.com/quote/EWZS/history?p=EWZS> (2020).