

# GAN-Based Prediction of Time Series

Sven FESTAG<sup>a,1</sup> and Cord SPRECKEISEN<sup>a</sup>

<sup>a</sup>*Institute of Medical Statistics, Computer and Data Sciences,  
Jena University Hospital, Germany*

**Abstract.** The study aims at generating initial and directional insights in the applicability of conditional generative adversarial nets for the imputation and forecasting of medical time series data. Our experiment with blood pressure series showed that a generative recurrent autoencoder exhibits significant individual learning progress but needs further tuning to benefit from joint training.

**Keywords.** GAN, time series, patient care, imputation, forecasting

## 1. Introduction

Time series data accrue in many medical contexts. First of all, bedside monitors gather patient-related time series data like vital signs, EEG waves and alike. Moreover, hospital logistics and stock management depend on temporal inventory data as well. Thus, accurate predictions of missing and future values can help to improve patient care. Several generative adversarial nets (GAN) have been introduced for the generation of synthetic time series data [1] or for forecasting one step [2]. However, to our knowledge, there are no published results about a system based on recurrent conditional GANs for the multi-step imputation and prediction task in the biomedical domain. We want to investigate the suitability of such a system and have developed an initial approach to do so.

## 2. Methods

In the present work we combine two recurrent sequential autoencoders (SAE) into a conditional GAN trained to enable the generator to predict time series data. A similar setup had proven to be well suited for text infilling [3]. For the supervised training procedure there exist two versions of every time series, the original one and a masked one. For the latter, the masked steps are replaced by a special character. The generator and discriminator both consist of two long short-term memory (LSTM) networks, one for encoding the other one for decoding. The bidirectional encoder LSTMs work in the same way in both AEs. They process the time series stepwise and output one so-called context vector each, which summarises all unmasked time series steps. In the generator, the context is attached to all values of the masked series before they are handed over to the decoder producing estimates at the missing positions. The discriminator's context is attached to each of these outputs of the generator. Afterwards, its decoder decides for every input position whether it is a fake (produced by generator) or an original time series value.

---

<sup>1</sup> Corresponding Author, Sven Festag, Germany; E-mail: sven.festag@med.uni-jena.de.

### 3. Results

For the experiment, a data set comprising time series of ICU patients' arterial blood pressure over intervals of 10 s was used. All data points were extracted from MIMIC-III Waveform [4]. During the masking process, a mask of 20 consecutive samples was used that was applied to every time series at a random position. The generator was pretrained on the full training set for 50 epochs to minimise the mean squared error. Afterwards, the GAN training proposed in [5] was started. The generator's optimizer kept the same learning rate as during pretraining. The rate for the discriminator was set to half of this. During the combined training, discriminator and generator were trained in alternation and both got half of the batch as training samples in every step. In all training phases, teacher forcing was applied, i.e. previous target values were passed to the decoder instead of predicted ones, as was the case during inference. During pretraining, the generator's mean squared error dropped from 2609 to 489. In the first epoch of the adversarial training, the generator's error increased slightly and decreased again very slowly over the following 25 epochs. However, the generator never performed better than directly after pretraining. The discriminator showed the opposite behaviour. Its accuracy increased rapidly before it levelled out to 0.5. The latter phase of the combined training led to oscillations in the performance measures of both parts.

### 4. Discussion

The results of the pretraining exhibit that the SAE topology used for the generator and the discriminator is expedient. However, only the pretraining showed the expected behaviour. In the presented experiments, the adversarial training did not lead to further improvement of the generator, since it destabilised the model after around 30 epochs.

### 5. Conclusion

We developed a GAN-based imputation and prediction approach for time series data. While the generator performed well on its own in our experiment, the current version does not achieve performance increase during the adversarial training. In the next stages, we would like to add approaches that counteract the destabilisation by the discriminator (Wasserstein loss) and some that might improve the generator (attention mechanism).

### References

- [1] Esteban C, Hyland SL, Rätsch G. Real-Valued (Medical) Time Series Generation with Recurrent Conditional GANs. arXiv:170602633 [cs, stat]. 2017 Dec.
- [2] Koochali A, Dengel A, Ahmed S. If You Like It, GAN It. Probabilistic Multivariate Times Series Forecast With GAN. arXiv:200501181 [cs, eess]. 2020 May.
- [3] Fedus W, Goodfellow I, Dai AM. MaskGAN: Better Text Generation via Filling in the \_\_\_\_\_. arXiv:180107736. 2018 Mar.
- [4] Moody B, et al. MIMIC-III Waveform Database Matched Subset (version 1.0). PhysioNet; 2020.
- [5] Goodfellow I, et al. Generative Adversarial Nets. In: Ghahramani Z, et al, editors. Advances in Neural Information Processing Systems 27. Curran Associates, Inc.; 2014. p. 2672-2680.