ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-TUPDP115

MACHINE LEARNING FOR COMPACT INDUSTRIAL ACCELERATORS

J. P. Edelen, M. J. Henderson, J. Einstein-Curtis and C. C. Hall, RadiaSoft LLC, Boulder CO, USA,
J. A. Diaz Cruz and A. L. Edelen, SLAC, Menlo Park CA, USA

Abstract

The industrial and medical accelerator industry is an ever-growing field with advancements in accelerator technology enabling its adoption for new applications. As the complexity of industrial accelerators grows so does the need for more sophisticated control systems to regulate their operation. Moreover, the environment for industrial and medical accelerators is often harsh and noisy as opposed to the more controlled environment of a laboratory-based machine. This environment makes control more challenging. Additionally, instrumentation for industrial accelerators is limited making it difficult at times to identify and diagnose problems when they occur. RadiaSoft has partnered with SLAC to develop new machine learning methods for control and anomaly detection for industrial accelerators. Our approach is to develop our methods using simulation models followed by testing on experimental systems. Here we present initial results using simulations of a room temperature s-band system.

INTRODUCTION

In recent years machine learning (ML) has been identified as having the potential for significant impact on the modeling, operation, and control of particle accelerators (e.g. see [1, 2]). Specifically, in the diagnostics space, there have been many efforts focused on improving measurement capabilities and detecting faulty instruments. When it comes to diagnostics, developments for beam position monitors have been quite ubiquitous over the years. Relatively recently, ML methods have been utilized to improve optics measurements from beam position monitor data [3]. Additionally, machine learning has been used to identify and remove malfunctioning beam position monitors in the Large Hadron Collider (LHC), prior to application of standard optics correction algorithms [4]. Furthermore, we have developed techniques for automation of noise removal in BPM data using machine learning [5]. On the contrary, there is a real dearth of knowledge when it comes to the application of machine learning for industrial accelerators. Moreover, the developments for using machine learning to improve RF signal processing are considerably further behind than other diagnostics in use at accelerators. The ability to remove noise from RF measurements would greatly improve our ability to extract meaningful information from RF systems especially in an industrial setting.

Machine learning methods such as autoencoders and variational autoencoders (VAEs) are well established for the removal of noise from various signals. For VAEs specifically noise reduction due to the enforcement of a smoothness condition in the latent-space representation. This feature of VAEs has been applied to gravitational wave research [6, 7] and geophysical data [8], for example. Recurrent autoencoders have the added advantage of being well suited to work with data sequences. In this paper we explore the use of Variational Recurrent Autoencoders (VRAEs) to remove different power law spectra (colors) of noise from simulated BPM data in a ring.

Our work utilizes a combination of approaches to understand which is best when considering RF waveform data. Our work has explored the use of model based approaches such as Kalman filters and machine learning approaches such as convolutional neural networks and variational autoencoders. Here we begin with a review of our data generation model followed by an analysis of Kalman filters, convolutional autoencoders, and variational autoencoders for the removal of noise from RF signals.

DATA GENERATION

Our data was generated using a RF simulator that reproduces waveforms as they would be seen in industrial systems. Over the past year, RadiaSoft has been developing a full RF simulation tool that is integrated with EPICS for the development of new control algorithms, developing IOC software, and testing user interfaces. The simulator can be run through various APIs including a command line interface, via a Jupyter notebook, or directly through an EPICS connection. The simulator is based on a linear circuit model that takes into account coupling factor, quality factor, frequency, drive amplitude and phase, pulse duration, detuning, etc. The dynamics of our model are based off of equations derived here [9–11].

The data were generated by varying the RF pulse characteristics and the cavity characteristics. For the pulse the length of the pulse was varied from 3 μs to 7.5 μs which is a reasonable range for industrial accelerator applications operating at S-Band. Additionally we varied the start time of the RF pulse in the data window. While we typically don't expect the RF pulse to vary in position along the DAQ window adding in this flexibility will ensure better generalization when transferring from simulations to measurement.

The RF cavity parameters of interest for this study are Q_0 and β which were varied over a range of 10,000 to 225,000 for the Q_0 and 1 to 3 for β. The detuning was also varied within a range of plus minus one half bandwidth, a fairly typical range seen on industrial systems. In all the parameter range chosen represents a reasonable range of industrial RF systems and will allow us to develop simulation based algorithms that should be readily transferable to measurement when the time arises.

KALMAN FILTERS

First we consider the Kalman filter for noise reduction. Kalman filters, also referred to as linear quadratic estimators,

ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-TUPDP115

use a liner model to predict the dynamics of a system using a state estimator and the provided input signal. The model used id comprised of the 1-D dynamical equations for an RF cavity derived from an equivalent circuit model. Here the dynamical variables for the cavity model are given by Equation 1.

$$\mathbf{x} = \begin{bmatrix} \Re(V_t) \\ \Im(V_t) \end{bmatrix} \quad \mathbf{u} = \begin{bmatrix} \Re(I_{fw}) \\ \Im(I_{fw}) \end{bmatrix} \quad \mathbf{y} = \begin{bmatrix} \Re(V_t) \\ \Im(V_t) \\ \Re(V_r) \\ \Im(V_r) \end{bmatrix} \tag{1}$$

The dynamical equations that describe the system with noise are given by $\dot{\mathbf{x}} = A\mathbf{x} + B\mathbf{u} + \Gamma\tilde{\mathbf{w}}$ and $\mathbf{y} = C\mathbf{x} + D\mathbf{u} + \tilde{\mathbf{v}}$ Here $\tilde{\mathbf{w}}$ and $\tilde{\mathbf{v}}$ are the noise components that show up in the dynamics that we wish to remove. The matrices A, B, C, and D, are defined by the cavity dynamics model as:

$$A = \begin{bmatrix} -\omega_{1/2} & \Delta\omega \\ \Delta\omega & -\omega_{1/2} \end{bmatrix} \tag{2}$$

$$B = \frac{R_L \omega_{1/2}}{m} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \tag{3}$$

$$C = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 1/m & 0 \\ 0 & 1/m \end{bmatrix} \tag{4}$$

$$D = \frac{Z_0}{2} \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \tag{5}$$

.

This continuous representation can be transformed into a discrete representation and then used to define an update formulae that allows us to estimate the system response in the absence of noise. Additionally as part of the Kalman filter algorithm we can estimate the covariance matrix for the system which gives us an uncertainty metric on the prediction in addition to the denoised data. Figure 1 shows an example waveform result of the Kalman filter. Here we predict the both the transmitted and reflected signals in both I and Q domains.

It's clear that the Kalman filter used in this way has some reasonable denoising capabilities. Here the output of the Kalman filter ids shown in blue with the confidence interval in shaded blue. The ground truth is in Black and Red shows the noisy signal. Compared with the baseline noisy signal the Kalman filter does quire good and has no training However because it relies on the input waveform which also has noise there is still substantial noise in the output signal.

CONVOLUTIONAL NEURAL NETWORK

We also developed a 1D convolutional autoencoder for denoising of our waveform data. Convolutional neural networks are adept at feature extraction especially in cases where there is translation invariance. While typical LLRF signals are time synchronized we explored signal translations as described above to improve the generality of our approach. The convolutional network follows a structure very similar to a U-net which is often used for image segmentation and other image to image learning problems. The model architecture consisted of 1-D convolutional layers and max-pooling layers that reduce the feature space down to a latent space of 10. We then used up-sampling and convolutional layers to reconstruct the waveform. The model was trained using noisy waveform data from our simulator in an unsupervised fashion. That is during training the model inputs and outputs both contain noise. The mechanism for noise reduction is due to the fact that there is no information contained in the noise it cannot be modeled by the latent space and the reconstruction will be noiseless. The data were tested on an independent dataset. When training the CNN we treated each waveform as unique to allow the CNN to learn noise rejection regardless of if the data being processed is a measured forward, reflected, or probe signal. This will improve our ability to generalize when considering data collected on different types of machines where probes are not always available or traveling wave structures where the signal envelopes do not follow the normal standing wave profiles. Figure 2 shows the noisy signal in grey, the model prediction in green, and the ground truth signal with no noise in black.

For comparison we see signals from the drive signal, the reflected signal, and the cavity probe. While generally the reconstructed signal is closer to the ground truth than the noisy signal there are cases spurious signals are present in the reconstructed data (top left for example).

VARIATIONAL RECURRENT AUTOENCODERS

Next we considered variational recurrent autoencoders, VRAEs. VRAEs are an excellent tool for data reduction and noise elimination due to the fact that they enforce a smoothness criterion in the latent space. This combined with the fact that the noise cannot propagate through the latent space as discussed previously makes them a prime candidate for removing noise from RF signals. Moreover, by utilizing recurrent layers we can effectively translate time dynamics to the principle components of the simulation which will be represented in the latent space.

Our implementation of the VRAE architecture is based on [12] and uses Long Short-Term Memory (LSTM) units for both the encoder and decoder. The loss function is composed of two terms: the Kullback–Leibler divergence [13] — which acts as a regularization term — and the reconstruction loss. For the reconstruction loss, mean squared error between the encoder input and decoder output is used. The VRAE is trained and tested with each of the waveforms (forward, reflected, and probe) treated as features for the dataset. The goal here is that the VRAE will be able to learn a latent space representation of the waveform data and by extension the cavity model parameters. This is in contrast to the CNN

ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-TUPDP115

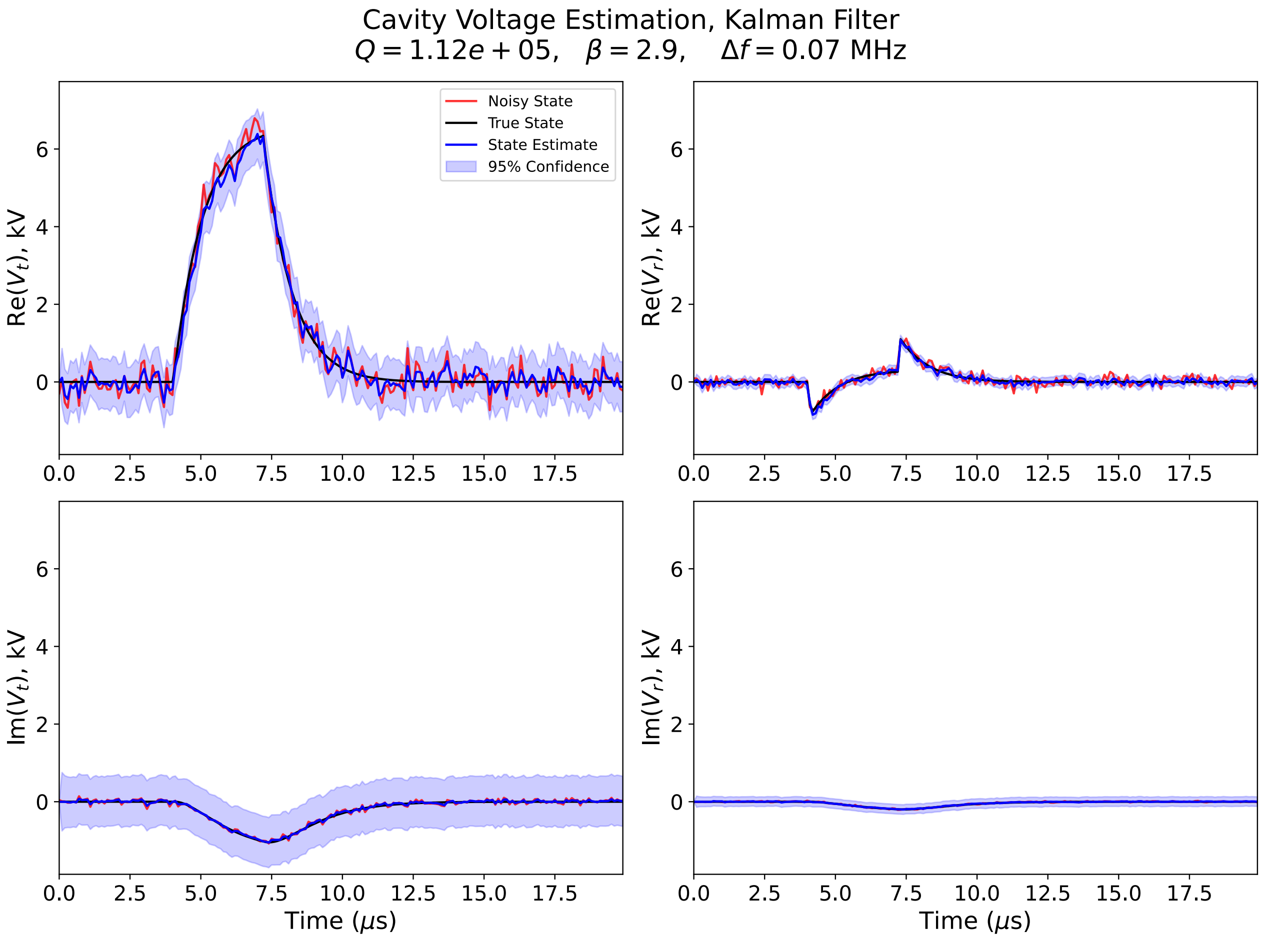

Figure 1: Comparison of the reconstructed waveform from the Kalman filter (red).

which was trained to process one waveform at a time and is largely learning to remove noise as opposed to a more generalized representation of the system.

The prediction for the VRE is generally quite a bit better in terms of the ability to remove noise but the profile of the reconstructed waveform is not always correct. This is likely due to the fact that there is some degeneracy in the dataset due to the relationship between Q and β.

COMPARISON

In all three cases the approaches are capable of reducing noise in the waveform data. Figures 4 and 5 show a direct comparison between the three approaches and the ground truth in addition to the original noisy signal. Here we can see that the CNN consistently does well compared to the Kalman Filter. The VRAE does quite well on the second example but does not reconstruct the signal properly for the first example. The advantage to the Kalman filter is that it does perform some noise reduction and because it relies on the physics of the system as opposed to being trained on the dataset it is less likely to produce spurious signals which can be seen in the second example using the CNN right before the cavity turns off.

To evaluate the performance of each technique across the whole dataset we computed the sum squared error between the reconstructed signal and the ground trough for each example waveform. We then computed a histogram of these results and compared it to histogram of the sum-squared-error between noisy signal and the noiseless signal. Figure 6 shows the result of this comparison.

The comparison across the whole test dataset shows that each method is capable of reducing the noise levels of the signals. The VRAE does the best job as can be seen by the prominent spike in the histogram near zero and a significant reduction in the noise contribution between 0.1 and 0.2. The CNN and the Kalman filter perform similarly overall with slightly better noise reduction coming from the CNN.

CONCLUSIONS

We have explored three possible approaches for noise reduction on industrial RF signals using machine learning. Specifically we have utilized Kalman filters, convolutional autoencoders, and variational autoencoders. Our test dataset was generated using a cavity circuit model that has been well benchmarked against measured data. While each method

ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-TUPDP115

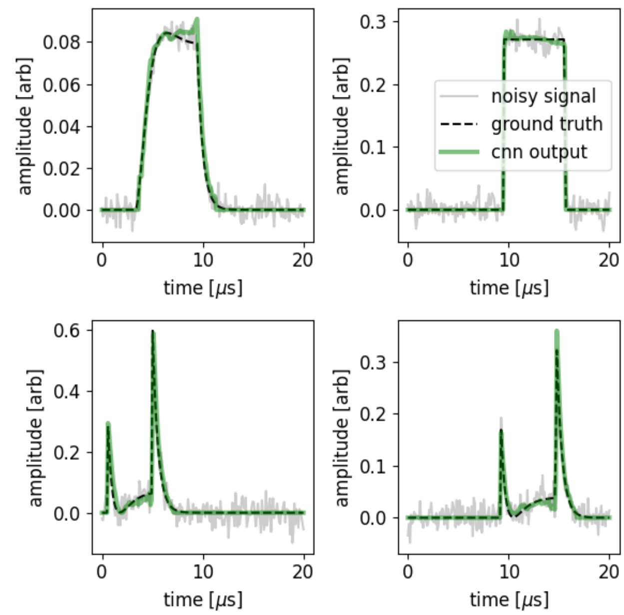

Figure 2: Comparison between the 1D-CNN reconstruction (green) of four randomly chosen waveforms and the ground truth (black) and the noisy signal (gray).

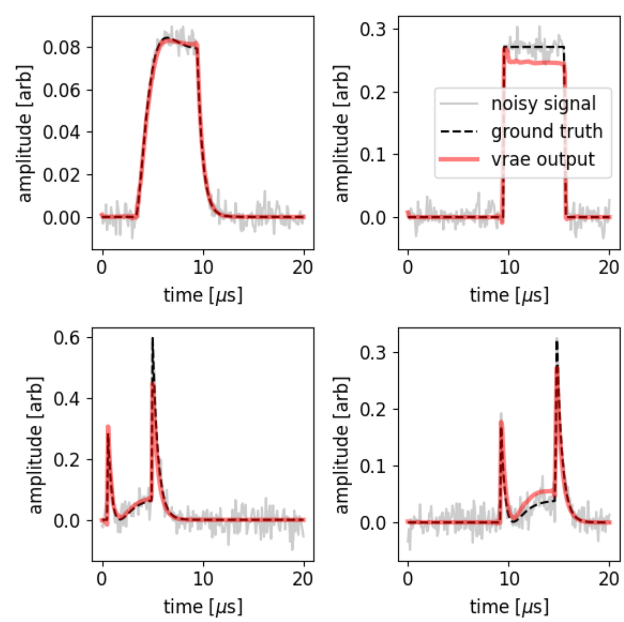

Figure 3: Comparison between the VRAE reconstruction (red) of four randomly chosen waveforms and the ground truth (black) and the noisy signal (gray).

has significant noise reduction capabilities, the VRAE is generally the best tool for this task.

ACKNOWELDGEMENTS

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Accelerator R&D and Production Award Number DE-SC0023641.

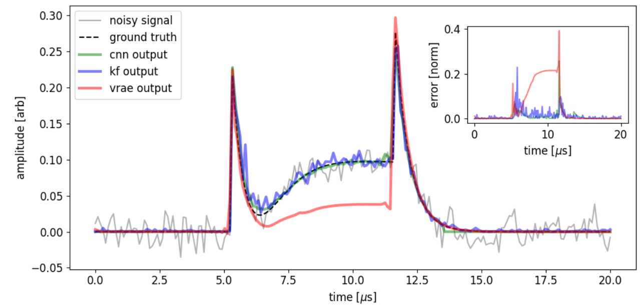

Figure 4: Reflected power waveform reconstruction using the Kalman Filter (blue), the 1D-CNN (green) and the VRAE (red) with the respective errors shown in the inset plot.

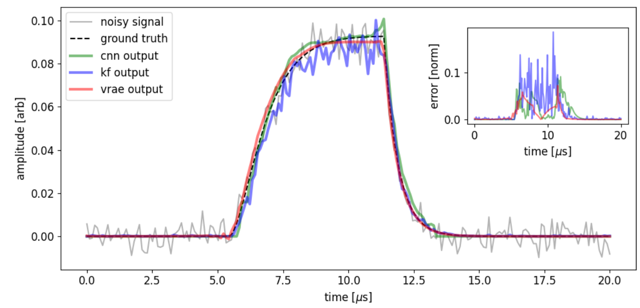

Figure 5: Cavity power waveform reconstruction using the Kalman Filter (blue), the 1D-CNN (green) and the VRAE (red) with the respective errors shown in the inset plot.

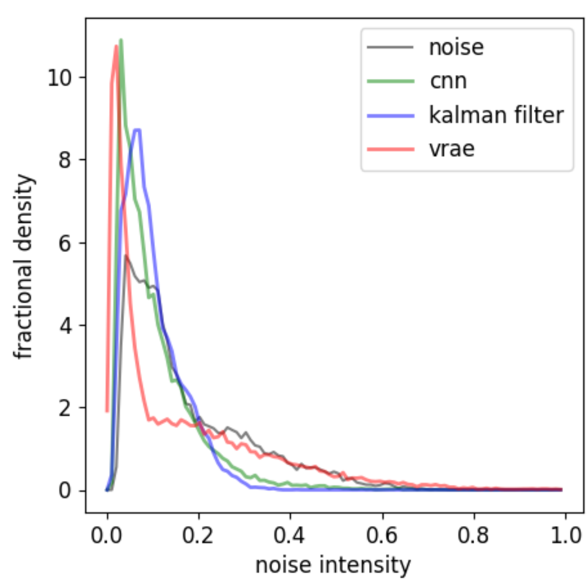

Figure 6: Histogram of the sum-squared-error between the reconstructed signal and the noiseless signal for the VRAE (red), 1D-CNN (green), and Kalman filter (blue).

REFERENCES

[1] A. L. Edelen, S. G. Biedron, B. E. Chase, D. Edstrom, S. V. Milton, and P. Stabile, “Neural networks for model-

ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-TUPDP115

ing and control of particle accelerators", *IEEE Trans. Nucl. Sci.*, vol. 63, no. 2, pp. 878–897, 2016. doi:10.1109/TNS.2016.2543203

[2] A. Edelen *et al.*, "Opportunities in machine learning for particle accelerators", 2018. doi:10.48550/ARXIV.1811.03172

[3] E. Fol, F. Carlier, J. C. de Portugal, A. Garcia-Tabares, and R. Tomás, "Machine Learning Methods for Optics Measurements and Corrections at LHC", in *Proc. IPAC'18*, Vancouver, BC, Canada, 2018, pp. 1967–1970. doi:doi:10.18429/JACoW-IPAC2018-WEPAF062

[4] E. Fol, J.C. de Portugal, and R. Tomás, "Unsupervised Machine Learning for Detection of Faulty Beam Position Monitors", in *Proc. IPAC'19*, Melbourne, Australia, 2019, pp. 2668–2671. doi:doi:10.18429/JACoW-IPAC2019-WEPGW081

[5] J. Edelen, J. Einstein-Curtis, C. Hall, M. Henderson, and A. Romanov, "Removing Noise in BPM Measurements with Variational Autoencoders", in *Proc. IBIC'22*, Kraków, Poland, 2022, paper MOP10, pp. 43–46. doi:10.18429/JACoW-IBIC2022-MOP10

[6] R. Ormiston, T. Nguyen, M. Coughlin, R. X. Adhikari, and E. Katsavounidis, "Noise reduction in gravitational-wave data via deep learning", *Phys. Rev. Research*, vol. 2, no. 3, p. 033 066, 2020. doi:10.1103/PhysRevResearch.2.033066

[7] G. Vajente *et al.*, "Machine-learning nonstationary noise out of gravitational-wave detectors", *Phys. Rev. D*, vol. 101, no. 4, p. 042 003, 2020. doi:10.1103/PhysRevD.101.042003

[8] D. Bhowick, D. K. Gupta, S. Maiti, and U. Shankar, "Stacked autoencoders based machine learning for noise reduction and signal reconstruction in geophysical data", *arXiv*, 2019. doi:10.48550/arXiv.1907.03278

[9] T. Czarski, K. T. Pozniak, R. S. Romaniuk, and S. Simrock, "Tesla cavity modeling and digital implementation in fpga technology for control system development", *Nucl. Instrum. Methods Phys. Res. A*, vol. 556, no. 2, pp. 565–576, 2006. doi:10.1016/j.nima.2005.10.122

[10] P. Echevarria, E. Aldekoa, J. Jugo, A. Neumann, A. Ushakov, and J. Knobloch, "Superconducting radio-frequency virtual cavity for control algorithms debugging", *Rev. Sci. Instrum.*, vol. 89, no. 8, p. 084 706, 2018. doi:10.1063/1.5041079

[11] P. Echevarria, E. Aldekoa, J. Jugo, J. Knobloch, A. Neumann, and A. Ushakov, "SRF Cavity Simulator for LLRF Algorithms Debugging", in *Proc. IPAC'18*, Vancouver, BC, Canada, 2018, pp. 2118–2120. doi:10.18429/JACoW-IPAC2018-WEPAK013

[12] O. Fabius and J. R. van Amersfoort, "Variational recurrent auto-encoders", *arXiv*, 2015. doi:10.48550/arXiv.1412.6581

[13] D. P. Kingma and M. Welling, "Auto-encoding variational bayes", *arXiv*, 2014. doi:10.48550/arXiv.1312.6114