

Simulation of an Elevator Group Control Using Generative Adversarial Networks and Related AI Tools

Tom Peetz Sebastian Vogt Martin Zaefferer Thomas Bartz-Beielstein
Institute for Data Science, Engineering, and Analytics
Technische Hochschule Köln
Steinmüllerallee 1, 51643 Gummersbach

September 4, 2020

Abstract

Testing new, innovative technologies is a crucial task for safety and acceptance. But how can new systems be tested if no historical real-world data exist? Simulation provides an answer to this important question. Classical simulation tools such as event-based simulation are well accepted. But most of these established simulation models require the specification of many parameters. Furthermore, simulation runs, e.g., CFD simulations, are very time consuming. Generative Adversarial Networks (GANs) are powerful tools for generating new data for a variety of tasks. Currently, their most frequent application domain is image generation. This article investigates the applicability of GANs for imitating simulations. We are comparing the simulation output of a technical system with the output of a GAN. To exemplify this approach, a well-known multi-car elevator system simulator was chosen. Our study demonstrates the feasibility of this approach. It also discusses pitfalls and technical problems that occurred during the implementation. Although we were able to show that in principle, GANs can be used as substitutes for expensive simulation runs, we also show that they cannot be used “out of the box”. Fine tuning is needed. We present a proof-of-concept, which can serve as a starting point for further research.

1 Introduction

Modern high-rise buildings require a lot of space for elevators. Classic elevator systems generally use central machine rooms and steel ropes for moving the elevator cars. Therefore, a shaft can only hold one car or a fixed stack of cars. Steel cables also limit the amount of reachable levels for the elevator, because the cables are too heavy [Miyamoto and Yamaguchi, 2008]. Recent approaches suggest elevator cars with independent linear motors, because they do not require any cables. Multiple cars can be operated independently in every shaft. These cars need to be controlled more carefully due to possible collisions, which cannot happen in single car shafts. They have to be controlled efficiently to maximize utility while reducing energy consumption, maintenance costs, and waiting times for passengers. Modeling such complex systems accurately is a demanding task. Additional elevator behavior data could improve the required algorithms for optimization and control. Simulation software is available, but is limited in features and general simulation accuracy. Generally, existing simulators focus on specific building and elevator types. We propose to generate additional data for elevator systems with

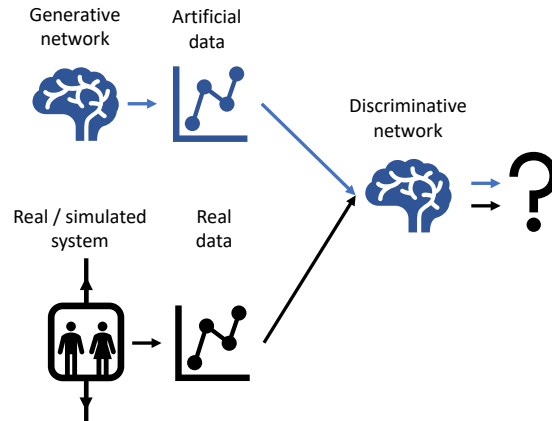


Figure 1: Architecture of a GAN. The discriminative network attempts to determine whether data is real or fake, whereas the generative network attempts to generate data that is not determined to be fake by the discriminative network.

the application of Generative Adversarial Networks (GANs). They are promising tools mainly used for generating images. Less frequent applications include text and language processing.

The goal of this paper is to explore the possible applications of GANs to create machine readable logs for the described elevator systems based on log data from an existing simulator. As previous work has shown, generative networks such as standalone Recurrent Neural Networks (RNNs) suffer from issues like exposure bias resulting in the generation of data which are never observed in training data [Bengio et al., 2015]. GANs offer a solution to this by adding a neural network, the discriminator, which tries to classify the generated data as “real” or “fake”.

The remainder of this report is structured as follows: section 2 presents related approaches. Section 3 describes materials and methods. Results from our first experiments are shown in Section 4 and discussed in Section 5. Further steps are presented in Section 6. A conclusion is given in Section 7.

2 Related Work

GANs have been first proposed by Goodfellow et al. [2014]. They describe a framework for generative modeling which consists of two separate networks, a generative network, G , and a discriminative network, D . Both networks are trained in an adversarial process. That means, the network G tries to counterfeit data that deceives D and D is trained to discriminate between real training data and the ‘fake’ data generated by G . Figure 1 illustrates this general idea.

This process results in both G and D improving their ability to fulfill their task, which in theory should create a generative network which creates data samples that are very hard to distinguish from real samples.

GANs have mostly been used to train generators for image data sets [Goodfellow et al., 2014, Yu et al., 2017]. Images are a well-suited application for GANs, because the generative network can make continuous adjustments to the generated data [Berges et al., 2017]. This can easily be achieved for images, because a numerical color value might be changed by a fraction of 1%. This is not possible if discrete data such as text have to be processed. In this case, small changes may generate text sequences that have an entirely different meaning [Yu et al.,

2017].

Yu et al. [2017] propose a possible solution called Sequential GAN (seqGAN) for this issue by adding feedback for unfinished sequences to the generator based on possible outcomes for a finished sequence. As the discriminator is only able to calculate rewards for finished sequences, an additional step is taken by spanning a search tree for n possible sequences based on the current time step and getting feedback for these from the discriminator. The next step is then chosen from the best candidate. This process is repeated until a whole sequence has been generated. Thereby, long-term rewards are taken into account for each time step instead of just focusing on short term rewards, resulting in a better quality of the generated sequences as a whole. seqGAN showed promising results when tested on the generation of Chinese poems and political speeches, outperforming other approaches for sequence generation [Yu et al., 2017]. A more detailed description of seqGANs can be found in Section 3.2.5.

3 Methods and Tools

3.1 MceSim

The training data for this project was generated using the Multi-Car Elevator Simulator (MceSim). MceSim is an open-source application for simulating multi-car elevator systems developed by Miyamoto and Yamaguchi [2008]. The simulation is designed for testing control patterns of elevator systems. It allows for different building configurations containing the amount of shafts in the building, cars per shaft, and the different types of floors in a building. Internally, the simulation software processes the calls and prevents overlapping (colliding) cars. Elevator calls can be created randomly by the simulation (as used in our experiments) or by providing a file that contains a list of the calls.

The code of the simulator was modified in order to generate the output of the internal states when processing calls into a text file. An example is shown in listing 1. Each line consists of a timestamp (the internal time of the simulation) and an operation which is executed on a call.

```
1 114 - New call: call_1 from 4 to 1 guests 4
2 114 - Assign call: call_1 on car_01_01
3 120 - Load call: call_1 on car_01_01
4 131 - Unload call: call_1 on car_01_01 overtravel 0
```

Listing 1: MceSim Custom Log File

During the simulation, many of these calls are handled in parallel. The life cycle for each call is composed of the following operations, which always occur in this order:

New call Passengers call an elevator inputting target location and the amount of passengers.

Assign call The call is assigned to a car.

Load call The car arrives at the start level and loads the passengers.

Unload call The car arrives at the destination level and unloads the passengers.

To create a basic log file with entries that can be used as training data for the GAN, a single log was generated with the following parameters:

```
1 java -jar MceSim2018.jar -r -t 1000000
```

Listing 2: basic log command

The command specifies to use random calls and to run the simulation for $1e6$ time units with the default configuration of five shafts, three cars each, and thirty floors. This generates a log file with approximately 21,000 lines, each containing an operation on a call. For the purposes of testing the approach, this single log was chosen as a simple dataset to test whether the application of GANs is feasible in the first place. In addition to the simple dataset, a more complex dataset can be created by iterating through all provided configurations and running the simulation for $1e4$ time units.

3.2 GAN implementation

3.2.1 Tooling

The programs used in our experiments were written in the Python programming language. The Keras framework as provided by Tensorflow was used to implement the neural networks. The following software versions were used:

- Python 3.8
- Tensorflow 2.2.0
- Numpy 1.19

Tensorflow and numpy can be installed using `pip` or by creating a `virtualenv`¹. It is recommended to utilize GPU acceleration for training and inference, as especially the generative network is very resource and time consuming. To enable GPU acceleration, `tensorflow-gpu` and and required drivers and libraries like `Cuda` and `cudnn` need to be installed².

3.2.2 Data Preprocessing

Because neural networks internally use numerical values such as integers or floats, data usually must be converted before they can be passed to the network. For images, this step is fairly simple because a 2D image can be represented as a 3D tensor of which the first two dimensions describe the position of a pixel while the third dimension describes its color. For textual input, a dictionary needs to be defined. It is used to translate human readable text into its integer representation. This was achieved by first creating the vocabulary `vocab` of the dataset containing each individual character. After that, each vocabulary entry is mapped to an integer. The resulting dictionary `char2idx` is used to translate text to its integer representation, which can be processed by the network.

A character-based vocabulary was chosen because of the structure of the data. This might appear counter-intuitively at the first sight, because only a few keywords are used in the log files. However, as can be seen in listing 1, ascending IDs and timestamps result in increasingly large vocabularies, almost proportional to the length of the logfile. A character-based logfile results in only 36 vocabulary entries or even only 32 when converting all letters to lower case. Furthermore, the character based vocabulary is needed for the generator to be able to create new timestamps and IDs, which are not present in the training data and would as a consequence not appear in any vocabulary.

The dataset uses the format `[BATCH_SIZE] [SEQUENCE_LENGTH] [1]`. Because of the variable length of the input text and the required uniform data shape, some remaining characters had to be cut off at the end of longer texts. For data loading, shaping, and preprocessing, the resources from Chollet [2020], Brownlee [2020], Liang [2018], and Keras [2020] were used.

¹see <https://virtualenv.pypa.io/en/latest/>

²see for details: <https://www.tensorflow.org/install/gpu>

Table 1: Exemplary output Tensor of the generator

a	b	c	d	e	f	g	i
0	0	0.05	0	0	0.05	0	0.9

3.2.3 Recurrent Neural Networks

The first part of the GAN is the generative network G . Based on their performance for text generation and the promising results of the seqGAN approach [Yu et al., 2017], RNNs were chosen for the generator. By utilizing Long Short Term Memory (LSTM) layers, issues like the “vanishing gradient problem” or the “exploding gradient problem” are mitigated [Yu et al., 2017]. Vanishing gradients describe situations in which the model is not actually learning anything because the gradients are too small to have an effect. The exploding gradient problem describes the opposite issue where gradients are too large and result in very unstable network behaviors [Pykes, 2020]. LSTMs are commonly used to prevent this and are capable of handling large input sequences of varied length [Yu et al., 2017], which is required for our experiments. The RNN is created as a sequential model in Keras, containing the following layers:

- Embedding layer, to break down high dimensional data into low dimensional data, to make it more understandable for the model.
- LSTM layer as the recurrent layer with added long term memory, to enhance the network’s capabilities for handling long sequences.
- Dense layer, which reduces the dimensions so that the output of the layer represents the probability of each item in the vocabulary for the next word.

The model was kept simple in order to keep the focus on the development of the total GAN architecture. A more sophisticated architecture might be useful, when more statistics about general GAN performance are available and parameters of the individual networks need to be tuned.

For inference, a sequence can be fed into the network, which results in the network predicting the next character for the sequence in form of a tensor containing the probability for each individual entry of the vocabulary. For example, if the network is trained to predict the next character in the alphabet and the sequence *abcdefgh* is given as input, the network might produce an output tensor as shown in Table 1. In practice, it is recommended to include some randomness in the selection process for the next word, because otherwise results might include loops of the most common vocabulary entries [TensorFlow Core Tutorials, 2020]. For this model, samples are drawn randomly from the weighted output distribution.

3.2.4 Convolutional Neural Networks

The next important part of the GAN is the discriminator, for which a Convolutional Neural Network (CNN) was chosen. Part of the selection process involved an analysis of different network types for text classification. Maheshwari [2018] evaluated three different types of networks for text classification. The results show a slight advantage of the chosen CNN compared to a RNN or a Hierarchical Attention Network (HAN) when looking at the training and validation accuracy, with the HAN coming close to or surpassing the CNN in some instances. The main benefit, why this network was chosen for the GAN, was the time factor. The CNN needs only a fraction of the time per epoch when compared to the RNN and HAN. While these benefits might be subject to how each network was built, the CNN was chosen with the option to replace it later if needed.

Similar to the RNN, a sequential Keras Model was used for the CNN. This model implemented the following layers:

- Embedding layer, to break down high dimensional data into low dimensional data, to make it more understandable for the model.
- 1D Convolutional layer, which uses a window of a certain size and calculates descriptive numbers for each possible window position in the dataset. This layer provides the name for the network.
- GlobalMaxPooling layer, which acts similar to the Convolutional layer, but takes maximum value from each window position. Here, ‘global’ indicates that the maximum of the complete sample is used.
- Dense layer with ReLu (Rectified Linear Unit) [Google Developers, 2020] activation, which filters out negative values and sets them to zero. Dense layers, which are also known as fully connected layers, are “hidden layer[s] in which each node is connected to every node in the subsequent hidden layer.” [Google Developers, 2020]
- Dropout layer to prevent overfitting
- Dense layer on sigmoid activation, to translate calculated predictions into probability numbers between 0 and 1

These layers were chosen because of their prominence in text classification in a CNN. Further evaluation of an improved network setup is needed. A CNN works as a classification network. This means, that the network is being fed data with corresponding labels, in the case of this GAN “real” for the data from the simulator and “fake” for the data from the discriminator. The network then learns to classify each set of input data into one of these classes and tries to maximize accuracy and minimize loss. A standalone CNN needs multiple datasets for training, validation, and testing, with corresponding label datasets which must be in such a shape, that the first dimension is equal to the corresponding text dataset. It is recommended to select the second dimension to be either [1] or [amount of label classes] when using one hot encoding [Sharma, 2020].

3.2.5 Sequential GANs

GANs work efficiently on continuous data, but not on discrete data. Especially long sequences of discrete data, such as text, are hard to process for GANs. The main benefit of using GANs on continuous data can be found when analyzing how GANs operate. GANs process training data for the discriminator and use approaching methods to minimize loss, with approaching steps sometimes being as small as one percent when looking at one single value. Small changes like this can be made on continuous data like images without changing the depicted object or even without humanly visible changes at all. Figure 2 illustrates this effect. When incrementing the hex value of the displayed pixel by just one, there is no visible change at all to the color. However, when the same strategy is applied to discrete data like text, differences can get much more pronounced. Figure 3 illustrates how incrementing the ASCII value of the letter “A” by one results in “B” and therefore in a changed meaning. Changes like this may have different consequences like spelling mistakes, creating completely unreadable sentences or even texts with completely different meanings. For word based vocabularies, changes might even be more drastic, because slight changes might either change nothing at all (so the word remains the same) or a completely unrelated word might get chosen.

To tackle this issue, further research provided possible solutions. The most promising solution found is the idea of using seqGANs as proposed by Yu et al. [2017]. They propose “modeling the data generator as a stochastic policy in reinforcement learning” [Yu et al., 2017],



Figure 2: Small changes in a pixel result in not human perceptible change

Letter "A"	Letter "B"
<u>Hex: 41</u>	<u>Hex: 42</u>
<u>Dec: 65</u>	<u>Dec: 66</u>

Figure 3: Small changes in ASCII values resulting in a changed meaning

meaning that the generator creates partial sequences, which then receive an expected result value, forcing the generator to learn and improve while building a dataset, instead of only optimizing the generator after finishing one dataset or epoch.

Due to the discriminator only being able to evaluate complete sequences, a strategy for evaluating incomplete sequences needed to be developed. The seqGAN approach achieves this through building a Monte Carlo Search Tree describing n possible complete sequences based on the current incomplete sequence. Each possible next set of partial sequences is then discriminated by the discriminator, which allows for ranking the possible next parts of the sequences by their long term value. To achieve a relatively accurate discrimination, the discriminator needs to be pre-trained with fake data from an untrained generator and real data.

The seqGAN algorithm can be described as follows:

1. Import and shape data
2. Build the network models (as described in Section 3.2.3 and 3.2.4)
3. Pre-train the discriminator with training data
4. Loop until the generator can reliably fool the discriminator

Generator inner loop

- (a) Generate a partial sequence
- (b) Calculate the expected results for possible following sequences
- (c) Optimize the generator with gradients drawn from the expected results

Discriminator inner loop

- (a) Discriminate real data and the fake data from the first inner loop

- (b) Calculate losses
- (c) Calculate gradients
- (d) Train either discriminator or both networks

A more detailed description can be found in Yu et al. [2017].

An implementation of this approach by Yu et. al. can be found on GitHub³ [Yu, 2018] and was tested during this project, but proved to be outdated, using Python 2 and old Tensorflow libraries. Additionally, no simple way to feed custom training data into the network was found without significantly modifying the implementation.

Considering the general focus of this project, it was decided not to use that implementation and instead focus on integrating the seqGAN approach into the GAN after finishing and evaluating it as a non-sequential GAN. This decision was also motivated by the fact that the seqGAN implementation uses Tensorflow directly without the Keras framework.

4 Results

The following subsections describe the results of standalone generator G and discriminator D , and preliminary results of the GAN.

4.1 Standalone generator

In order to gather information about the performance of the generator G without any feedback from the discriminator, the model was trained for ten epochs using the previously generated log files containing data of 1e6 simulation time units. After that, the input “1 - New Call:” was passed to the model, as each log starts with this sequence. The model was used to generate 1e6 characters. Generating the whole sequence took about 2.5 hours on a Nvidia RTX 2080Ti GPU, but the run time should be taken with some amount of uncertainty since the model is not optimized for performance and utilization of the GPU. As a comparison, a similar amount of lines generated using MceSim takes about ten minutes while only running on a single core of a CPU. The first part of the resulting logfile is shown in Listing 3.

```

1 1 - New call: c198196 from 19 to 27 guests 6
2 997953 - Assign call: c199257 on car_03-03
3 968675 - Load call: c192360 on car_02-01
4 979819 - Load call: c198112 on car_03-02
5 999762 - Unload call: c194291 on car_02-01 overtravel 0
6 993170 - New call: c198744 from 22 to 16 guests 8
7 979225 - Assign call: c197549 on car_04-02 overtravel 0

```

Listing 3: Log generated by G

Some observations can be made in this output which can be generalized to the whole log:

- The model was able to learn the general structure of the logfile. This includes features like lines beginning with a timestamp followed by the type of operation on the call. Furthermore, the structure for each type of operation seems to be learned correctly.
- Timestamps are not ascending and most are in the high end of the range given by the input log file (near 1e6). Otherwise they seem to be random.
- IDs for calls and cars seem to be completely random.

³See <https://github.com/LantaoYu/SeqGAN/>

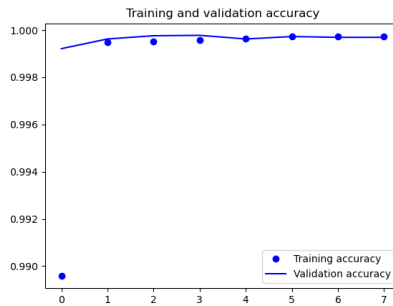


Figure 4: Standalone discriminator. Accuracy plotted against process steps. Larger values are better.

- The “life cycle” for calls is not taken into account. Otherwise, all lines in the presented outtake of the log should start with *new call* because IDs do not repeat.

Additionally, some incomplete lines like listing 4 (missing the additional information which car is assigned) or lines with random content like listing 5 can be found in the log.

```
994352 - Assign call: c1879902
```

Listing 4: Incomplete line

```
94_03
```

Listing 5: Line with random content

4.2 Standalone Discriminator

The discriminator as a standalone network was given “real” data from the simulator and “fake” data from the standalone generator as described in Chapter 4.1. The discriminator was quickly able to differentiate between real and fake data, which an accuracy of almost 99% and a loss below 0.005 after the first epoch as seen in Figures 4 and 5 with slight improvements in the following epochs.

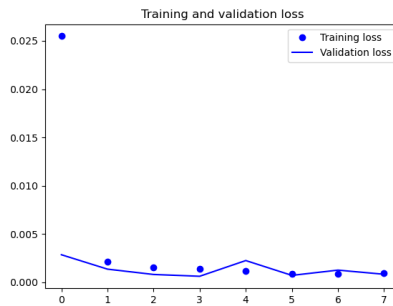


Figure 5: Standalone discriminator. Loss plotted against process steps. Smaller values are better.

4.3 GAN

In the context of this project, a fully working implementation of the GAN for text generation was not completed. As described before, stand alone networks for the RNN generator and the CNN discriminator have been developed and tested. When combining both, an issue for passing the data between generator and discriminator was discovered: in the training loop, the generator generates a batch of predictions in the form of probability distributions as described in Section 3.2.3. However, these cannot be used by the discriminator, because the discriminator expects an integer sequence as input. As before, the solution is to sample from these batches to get actual text predictions, which in turn are now compatible to the discriminator.

Listing 6 shows an excerpt from the training loop for the network. The generator makes predictions for a complete batch, resulting in a tensor of the following shape: [BATCH_SIZE] [SEQ_LENGTH] [VOCAB_SIZE], representing the output distribution. Next, a list for the predictions is created and a loop samples for each batch item, so actual sequences are obtained. The resulting list of sequences is then transformed into an array of the shape [BATCH_SIZE], [SEQ_LENGTH] which can be passed to the discriminator.

```
1 batch_predictions = generator(input, training=False)
2
3 fake_list = []
4 for batch_prediction in batch_predictions:
5     sampled_indices = tf.random.categorical(batch_prediction, num_samples=1)
6     sampled_indices = tf.squeeze(sampled_indices, axis=-1).numpy()
7     fake_list.append(sampled_indices)
8
9 fake_dataset = np.asarray(fake_list)
10 decision = discriminator(fake_dataset)
```

Listing 6: Training loop

The process works fine for passing data between the two networks, but when the gradients of both networks are calculated, the gradients for the generator seem to be lost due to the conversion from the output tensor to an array. Descriptions for this issue can be found in the project issues for Tensorflow⁴. In order to complete the implementation, a method for passing the data without losing gradients is needed. Tests without the generator optimizer seem to indicate, that the network is otherwise functional.

Even though the complete GAN could not be implemented as planned, some recommendations from the literature will be presented: most authors note that GANs suffer from the issues presented in section 3.2.5. For example, Berges et al. [2017] explored different models in a GAN setup and concluded that they were not able to create a generator which was able to deceive the discriminator. Issues seem to be caused by the problem of sampling the probability distributions before passing sequences to the discriminator.

To alleviate these issues, the seqGAN approach proposed by Yu et al. [2017] seems to be a promising solution. They present a way to give feedback to the generator while prioritizing the long term outcome for the generated sequence. Their results show that the approach can be successfully used for natural language (political speeches and Chinese poems).

5 Discussion

The results of a standalone RNN generator look promising. It quickly learns the basic structure of the log file, the possible line types and correct spelling, even though it is trained on character

⁴See <https://github.com/tensorflow/tensorflow/issues/36596>

basis. However, other issues, like incorrect timestamps and random IDs, indicate that the model is not able to learn the broader contexts in the log file.

Moreover, the time, energy, and computation resources required to generate log files similar to those of the MceSim, are much higher. The example sequence of 1e6 characters requires on a GPU accelerated network about 15 times longer than the simulation utilizing only a single CPU core, making it much more expensive without any additional benefit in this concrete use case.

An evaluation of the discriminator is difficult at this point, without evaluating the generator at the same time. It is possible that we observed rather good accuracy and loss values, because the discriminator learns fast and efficient, or else, because the generator creates text which is easy to detect as fake. The discriminator would also have to be trained with more and a bigger variety of datasets to produce a more telling output. As the difference between real and fake logs can be very easily determined by a human, it is assumed that the discriminator learns to check for similar features. Unordered timestamps, no recurring IDs and incomplete lines might be artifacts which are picked up by the discriminator quickly.

Completing and testing the GAN with a combination of both networks should give better insight into how the networks performed. As described before, the implementation was not completed due to vanishing gradients when sampling from the output distribution of the generator. Testing the completed implementation can possibly give more insight about the usefulness of GANs for generating machine readable data like MceSim log files.

Some observations made on the single tested networks are also true for GAN in general. GANs consume much more resources and power than the simulation software when only used to create more of the same data like in this project. Additionally, the literature suggests in part, that the quality of the generated data might never be good enough to deceive the discriminator[Berges et al., 2017], which needs to be tested with the completed implementation. Because of the simple structure of the training data compared to natural language, results for this application might be more favorable. Additions like the seqGAN approach [Yu et al., 2017] might improve on this and could be adapted to the simpler structure of machine readable data.

6 Outlook

This section presents suggestions for future work.

Finishing the GAN implementation: Most importantly, the implementation of the GAN prototype should be finished to enable an extensive performance evaluation. While the literature suggests that GANs are not necessarily a good approach for generating text data, there are approaches which seem to be promising. Additionally, the data handled in this project is of much simpler structure than natural language which is usually considered when talking about the difficulties of text processing.

Improving the networks: After finishing and evaluating the current GAN configuration, improving the networks might be the next step. Lucena [2020] suggested using Structured Categorical Decision Treess (SCDTs) instead of numerical or one hot encoding for improving Gradient calculation. This might accelerate learning for the networks. An evaluation of SCDTs for GANs might provide some benefits regarding learning speed.

As a majority of the available time was consumed understanding and building the networks and combining them into the GAN, no further evaluation of specific network layers or combination of layers was made. With a more thorough understanding of neural networks and their corresponding layers, their parameters and combination and order of layers, the generator and discriminator could be improved. Further testing and evalu-

ation might even suggest using different networks for the generator and discriminator instead of the RNN and CNN that were used in our experiments.

Evaluating GANs and output data: Another issue that should be investigated in subsequent studies is the evaluation of the GAN, the finished output networks, and the generated data. While this might give suggestions on how to improve the networks, it is also important to evaluate before using the finished generator or generated data in critical fields such as elevator planning, vehicle programming or medical analysis. Using faulty or inaccurate data might lead to failures that could affect human safety, health or lives. While no solution to evaluate data for safe usage in critical fields was found during this project, the research of Esteban et al. [2017] suggests two ways of evaluating and optimizing GANs. The first approach is to check whether the GAN “implicitly learns the distribution of the true data” [Esteban et al., 2017] via a medium mean discrepancy, while the second approach suggests training a second GAN with the generated data of the first GAN and testing it with real data. If the second GAN is trained with generated data from the first GAN, it can successfully detect real data as real, then we can consider the first GAN a success.

Implementing a sequential GAN: Another possible addition after finishing the GAN implementation may be the seqGAN approach as described in section 3.2.5, which is expected to work better on discrete data such as text compared to a more standard GAN design. Since the architecture of the seqGAN foresees an additional search algorithm, as well as having the discriminator calculate losses for multiple partial sequences per sequence generated, the resource cost concerning CPU, GPU and time usage is expected to increase. It is to be evaluated if the benefits outweigh the additional cost or if similar results can be achieved by using the non-sequential GAN. An additional idea might be to evaluate the used search algorithm. Yu et. al. use a Monte Carlo Search Tree, but there might be more cost efficient alternatives for this scenario.

Adding building configuration input into the GAN: To enhance the usefulness of a complete application around the GAN, it would be helpful to be able to pass a building configuration describing the amount of floors, shafts, and cars per shaft. The GAN needs to be able to process this input and generate logs which apply to exactly this building configuration. The feature is required to create data which can be used for optimization of elevator systems in specific scenarios. This might be achievable by generating datasets containing only logs for the appropriate building and training the network with this specific dataset. To reduce training times, it might be appropriate to pre-train a general model. This can be adapted using problem specific data, if available.

7 Conclusion

This article investigates the applicability of GANs for imitating complex and expensive real-world simulations. We compared the simulation output of a technical system with the output of a GAN. To exemplify this approach, a well-known multi-car elevator system simulator was chosen. Our study demonstrates the feasibility of this approach. It also discusses pitfalls and technical problems that occurred during the implementation. Although we were able to show that in principle, GANs can be used as substitutes for expensive simulation runs, we also show that they cannot be used “out of the box”. Fine tuning is needed. We presented a proof-of-concept, which can serve as a starting point for further research.

Overall, our research provided much insight in the application of GANs for the generation of machine-readable data, even though a working implementation could not be finished. GANs are powerful tools for creating data, but also come with many difficulties when working on

discrete data. Some of these issues are due to the conceptual design of models that were created with image manipulation in mind, while others are more technical like the gradient issue for the implementation of this project.

Applying the GAN framework to text files, which can also be generated using the actual simulation itself, requires more resources and time. Especially, the generative network takes large amounts of training data and time to be able to generate data which is still easily distinguishable from real training data, by human eye as well as the discriminative network. However, many improvements of the individual networks and the overall GAN can be made to optimize the quality of the outcome. Building more sophisticated network models and adding concepts like the sequential approach might improve the output of the generative network sufficiently. Then, logs could mimic those of the simulation or even formats like those produced by real world elevator systems.

Even though this approach may seem cost- and performance ineffective at the first glimpse, there might be some additional benefits when looking beyond the scope of this project, like the creation of real world like data which must not include real details, e.g., for privacy protection in medicine or enhancing real-world datasets generated by actual elevator systems.

References

- Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. Scheduled sampling for sequence prediction with recurrent neural networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'15, page 11711179, Cambridge, MA, USA, 2015. MIT Press.
- Vincent-Pierre Berges, Isabel Bush, and Pol Rosello. Text generation using generative adversarial networks. Technical report, Stanford University, 2017. URL <http://prosello.com/papers/text-gans-w17.pdf>. Accessed: 2020-09-02.
- Jason Brownlee. Your First Deep Learning Project in Python with Keras Step-By-Step, 2020. URL <https://machinelearningmastery.com/tutorial-first-neural-network-python-keras/>. Accessed: 2020-08-01.
- Francois Chollet. Introduction to Keras for Engineers, 2020. URL https://keras.io/getting_started/intro_to_keras_for_engineers/#data-loading-amp-preprocessing. Accessed: 2020-08-01.
- Cristbal Esteban, Stephanie L. Hyland, and Gunnar Rtsch. Real-valued (Medical) Time Series Generation with Recurrent Conditional GANs. *arXiv*, 2017. URL <http://arxiv.org/abs/1706.02633>. arXiv: 1706.02633.
- Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Networks. *arXiv*, 2014. URL <http://arxiv.org/abs/1406.2661>. arXiv: 1406.2661.
- Google Developers. Machine Learning Glossary, 2020. URL <https://developers.google.com/machine-learning/glossary/>. Accessed: 2020-08-02.
- Keras. Keras API reference: Text data preprocessing, 2020. URL <https://keras.io/api/preprocessing/text/>. Accessed: 2020-08-01.
- Xu Liang. How to Preprocess Character Level Text with Keras, 2018. URL <https://towardsdatascience.com/how-to-preprocess-character-level-text-with-keras-349065121089>. Accessed: 2020-08-01.

- Brian Lucena. StructureBoost: Efficient Gradient Boosting for Structured Categorical Variables. *arXiv*, 2020. URL <http://arxiv.org/abs/2007.04446>. arXiv: 2007.04446.
- Akshat Maheshwari. Report on Text Classification using CNN, RNN & HAN, 2018. URL <https://medium.com/jatana/report-on-text-classification-using-cnn-rnn-han-f0e887214d5f>. Accessed: 2020-08-01.
- Toshiyuki Miyamoto and Shingo Yamaguchi. MceSim: A Multi-Car Elevator simulator. *IEICE Transactions on Fundamentals of Electronics Communications and Computer Sciences*, E91-A(11):3207–3214, 2008. doi: 10.1093/ietfec/e91-a.11.3207.
- Kurtis Pykes. The Vanishing/Exploding Gradient Problem in Deep Neural Networks, May 2020. URL <https://towardsdatascience.com/the-vanishing-exploding-gradient-problem-in-deep-neural-networks-191358470c11>. Accessed: 2020-08-03.
- Aditya Sharma. (Tutorial) Convolutional Neural Networks with TensorFlow, 2020. URL <https://www.datacamp.com/community/tutorials/cnn-tensorflow-python>. Accessed: 2020-08-01.
- TensorFlow Core Tutorials. Text generation with an RNN, 2020. URL <https://www.tensorflow.org/tutorials/text/text-generation>. Accessed: 2020-07-25.
- Lantao Yu. SeqGAN. GitHub, 2018. URL <https://github.com/LantaoYu/SeqGAN>. Accessed: 2020-08-04.
- Lantao Yu, Weinan Zhang, Jun Wang, and Yong Yu. SeqGAN: Sequence generative adversarial nets with policy gradient. *arXiv*, 2017. URL <http://arxiv.org/abs/1609.05473>. arXiv: 1609.05473.