

TECHNICAL UNIVERSITY OF CRETE  
SCHOOL OF ELECTRICAL AND COMPUTER  
ENGINEERING



**Music Generation with Neural Networks**

Diploma Thesis by  
Argiris Papadopoulos

Thesis Committee

Associate Professor Michail G. Lagoudakis (supervisor)  
Associate Professor Georgios Chalkiadakis  
Professor Euripides Petrakis

Chania, February 2022

ΠΟΛΥΤΕΧΝΕΙΟ ΚΡΗΤΗΣ

ΣΧΟΛΗ ΗΛΕΚΤΡΟΛΟΓΩΝ ΜΗΧΑΝΙΚΩΝ & ΜΗΧΑΝΙΚΩΝ

ΥΠΟΛΟΓΙΣΤΩΝ



Δημιουργία Μουσικής με χρήση Νευρωνικών  
Δικτύων

Διπλωματική Εργασία του  
Αργύρη Παπαδόπουλου

Τριμελής Επιτροπή

Αναπληρωτής Καθηγητής Μιχαήλ Γ. Λαγουδάκης (επιβλέπων)

Αναπληρωτής Καθηγητής Γεώργιος Χαλκιαδάκης

Καθηγητής Ευριπίδης Πετράκης

Χανιά, Φεβρουάριος 2022

# Abstract

Machine Learning has been used in many applications in recent years and has produced impressive results. Interestingly, beyond well-known problems of academic, research or commercial interest, machine learning techniques are finding more and more their way into the area of arts, in the sense of generative modeling. In this thesis, we explore how deep neural networks can be used to automatically generate musical sequences. The goal of this work is to construct models that are able to learn the basic patterns of an input dataset (corresponding to a specific music genre) and try to replicate these patterns embedded into new, original samples, under the assumption that modeling and sampling may be more effective in two-dimensional image representations. To this end, we utilize well-known machine learning generative models, namely Variational Autoencoders (VAEs) and Generative Adversarial Networks (GANs), and propose our own deep network architectures maintaining simplicity and low computational power and time requirements. The implemented models are trained using datasets of MIDI files, containing tunes from different music genres, which are first converted to two-dimensional images during preprocessing. After training at image level, the trained models generate new images of a similar kind, which are decoded back to MIDI tunes, following a reverse procedure, and thus to music. MIDI files are ideal for our purposes due to their discrete nature, which facilitates the conversion to images back and forth. In the course of this work, we focused on data engineering issues, namely how to shape and form data in a way that helps our generative models learn easier and faster. We also offer a comparison of the different models and infer results on their effectiveness. The produced music tunes seem to resemble basic features of the original ones, but only in a few cases the outcome was truly interesting in terms of music theory. The proposed approach could potentially help musicians improve and explore original tunes based on preferred genres or types of music. Furthermore, our work can be used as a model for other learning tasks, not necessarily related to music, which may be facilitated, if explored through image representations, as we proposed. While our results do not generate consumer-grade music yet, our work represents a first step in the direction of automated music generation and computational creativity in general.

## Περίληψη

Η Μηχανική Μάθηση έχει χρησιμοποιηθεί σε πολλές εφαρμογές τα τελευταία χρόνια και έχει δώσει εντυπωσιακά αποτελέσματα. Είναι ενδιαφέρον ότι πέρα από τα γνωστά προβλήματα ακαδημαϊκού, ερευνητικού ή εμπορικού ενδιαφέροντος, οι τεχνικές μηχανικής μάθησης βρίσκουν όλο και περισσότερο τον δρόμο τους στον χώρο των τεχνών, με την έννοια της παραγωγικής μοντελοποίησης (generative modeling). Στην παρούσα διπλωματική εργασία, διερευνούμε πώς τα βαθιά νευρωνικά δίκτυα μπορούν να χρησιμοποιηθούν για την αυτόματη δημιουργία μουσικών ακολουθιών. Ο στόχος αυτής της εργασίας είναι η κατασκευή μοντέλων που μπορούν να μάθουν τα βασικά μοτίβα ενός συνόλου δεδομένων εισόδου (που αντιστοιχεί σε κάποιο είδος μουσικής) και να προσπαθήσουν να αναπαραγάγουν αυτά τα μοτίβα ενσωματωμένα σε νέα, πρωτότυπα δείγματα, με την υπόθεση ότι η μοντελοποίηση και η δειγματοληψία μπορεί να είναι πιο αποτελεσματική σε αναπαραστάσεις δισδιάστατων εικόνων. Για το σκοπό αυτό, χρησιμοποιούμε γνωστά μοντέλα μηχανικής μάθησης, συγκεκριμένα Variational Autoencoders (VAE) και Generative Adversarial Networks (GAN), και προτείνουμε τις δικές μας βαθιές αρχιτεκτονικές δικτύων που διατηρούν απλότητα και χαμηλές απαιτήσεις υπολογιστικής ισχύος και χρόνου. Τα μοντέλα που υλοποιήθηκαν εκπαιδεύονται χρησιμοποιώντας σύνολα δεδομένων αρχείων MIDI, που περιέχουν μελωδίες από διαφορετικά είδη μουσικής, τα οποία αρχικά μετατρέπονται σε δισδιάστατες εικόνες κατά την προεπεξεργασία. Μετά την εκπαίδευση σε επίπεδο εικόνας, τα εκπαιδευμένα μοντέλα παράγουν νέες εικόνες παρόμοιου είδους, οι οποίες αποκωδικοποιούνται σε μελωδίες MIDI, ακολουθώντας μια αντίστροφη διαδικασία, και συνεπώς σε μουσική. Τα αρχεία MIDI είναι ιδανικά για τους σκοπούς μας λόγω της διακριτής φύσης τους, η οποία διευκολύνει τη μετατροπή σε εικόνες εμπρός και πίσω. Κατά τη διάρκεια αυτής της εργασίας, εστιάσαμε σε ζητήματα επεξεργασίας δεδομένων, δηλαδή στον τρόπο διαμόρφωσης και στοίχισης δεδομένων, ώστε να βοηθά τα μοντέλα παραγωγής να μαθαίνουν ευκολότερα και γρηγορότερα. Προσφέρουμε επίσης μια σύγκριση των διαφορετικών μοντέλων και εξάγουμε συμπεράσματα σχετικά με την αποτελεσματικότητά τους. Οι μουσικές μελωδίες που παράγονται φαίνεται να προσομοιάζουν σε βασικά χαρακτηριστικά των αρχικών μελωδιών, αλλά μόνο σε λίγες περιπτώσεις το αποτέλεσμα ήταν πραγματικά ενδιαφέρον, όσον αφορά τη θεωρία της μουσικής. Η προτεινόμενη προσέγγιση θα μπορούσε ενδεχομένως να βοηθήσει

τους μουσικούς να βελτιώσουν και να εξερευνήσουν πρωτότυπες μελωδίες με βάση τα προτιμώμενα είδη μουσικής. Επιπλέον, η εργασία μας μπορεί να χρησιμοποιηθεί ως πρότυπο για άλλα προβλήματα μηχανικής μάθησης, που δεν σχετίζονται απαραίτητα με τη μουσική, τα οποία μπορεί να διευκολυνθούν, εάν διερευνηθούν μέσω αναπαραστάσεων εικόνων, όπως προτείναμε. Αν και τα αποτελέσματά μας δεν παράγουν ακόμα μουσική καταναλωτικής ποιότητας, η δουλειά μας αντιπροσωπεύει ένα πρώτο βήμα προς την κατεύθυνση της αυτοματοποιημένης παραγωγής μουσικής και της υπολογιστικής δημιουργικότητας γενικότερα.

## **Acknowledgements**

First and foremost, I would like to give my warmest thanks to my supervisor Dr. Michail G. Lagoudakis, whose guidance and advice gave shape to this work. I would also like to thank the thesis committee members, professors Euripides G. M. Petrakis and Georgios Chalkiadakis for their thoughtful comments and suggestions.

Second, I would like to give special thanks to my partner, for her existence and presence which are the greatest gifts I have ever received, and for standing with me at my best and worst during the most crucial years of my life. My family for their continuous support and understanding while undertaking my research.

Last, but not least, I would like to express my thanks to my friends and classmates whose support was essential for the completion of this work.

# Contents

<b>1</b>	<b>Introduction</b>	<b>12</b>
1.1	Computer-Based Musical Systems . . . . .	13
1.2	History of MIDI . . . . .	14
1.3	Autonomy versus Assistance . . . . .	15
1.4	Deep Learning . . . . .	16
<b>2</b>	<b>Background</b>	<b>17</b>
2.1	Basic Music Theory . . . . .	17
2.1.1	Measure and Bar Lines . . . . .	18
2.1.2	Time Signature . . . . .	19
2.1.3	Tempo . . . . .	21
2.1.4	Note Values . . . . .	22
2.2	Deep Learning . . . . .	23
2.2.1	Deep Learning and the Human Brain . . . . .	23
2.2.2	Artificial Neural Networks . . . . .	24
2.2.3	Shallow Neural Network . . . . .	25
2.2.4	Deep Neural Network . . . . .	27
2.2.5	Back Propagation . . . . .	27
2.2.6	Convolutional Neural Networks . . . . .	28
<b>3</b>	<b>Methods</b>	<b>31</b>
3.1	Preprocessing of training data . . . . .	32
3.1.1	Messages . . . . .	33
3.1.2	Meta Messages . . . . .	36

3.1.3	Music to data . . . . .	36
3.2	Training Models . . . . .	42
3.2.1	Variational Autoencoder (approach and architecture) . . . . .	42
3.2.2	Deep Convolutional GANs (approach and architecture) . . . . .	56
<b>4</b>	<b>Results</b>	<b>69</b>
4.1	Variational Autoencoder training results . . . . .	71
4.1.1	Training the VAE model . . . . .	71
4.1.2	VAE evaluation . . . . .	74
4.2	GAN training results . . . . .	79
4.2.1	Training the GAN model . . . . .	79
4.2.2	GAN evaluation . . . . .	82
<b>5</b>	<b>Discussion</b>	<b>86</b>
5.1	Note probability graphs . . . . .	86
5.2	Metric plots . . . . .	88
5.2.1	VAE metric plots . . . . .	89
5.2.2	GAN metric plots . . . . .	90
5.3	Listening to the results . . . . .	91
5.4	Model comparison . . . . .	93
<b>6</b>	<b>Conclusions</b>	<b>95</b>
<b>7</b>	<b>Future Work</b>	<b>97</b>

# List of Figures

2.1	Measures and pentagram . . . . .	18
2.2	Bar line types . . . . .	18
2.3	Different types of Time Signatures. . . . .	20
2.4	Note values . . . . .	22
2.5	Neuron (Artificial Neural Network) . . . . .	24
2.6	Shallow Neural Network . . . . .	26
2.7	Deep Neural Network . . . . .	27
2.8	Filter application . . . . .	29
3.1	Example of a track with time attributes . . . . .	34
3.2	Example of MIDI to image encoding compared to the same MIDI opened by an external application . . . . .	41
3.3	The Standard Autoencoder . . . . .	43
3.4	MNIST Latent space optimizing only on reconstruction loss . . . . .	44
3.5	MNIST Latent space optimizing only on KL divergence loss . . . . .	47
3.6	MNIST Latent space optimizing on both reconstruction and KL divergence loss . . . . .	48
3.7	Encoder Architecutre . . . . .	49
3.8	ReLU . . . . .	51
3.9	Decoder Architecutre . . . . .	52
3.10	Sigmoid function . . . . .	53
3.11	Adam compared to other optimizers . . . . .	56
3.12	The Generator Model . . . . .	59
3.13	The Discriminator Model . . . . .	60

3.14	The Generative Adversarial Network Architecture . . . . .	61
3.15	The Generator Architecture . . . . .	62
3.16	Rectified Linear Units (Basic and Leaky) . . . . .	64
3.17	The Discriminator Architecture . . . . .	65
4.1	Generated samples for VAE during training . . . . .	72
4.2	VAE Latent space for each dataset . . . . .	73
4.3	VAE Loss history for each dataset . . . . .	73
4.4	Notable samples generated by the VAE . . . . .	75
4.5	Note frequency bar graphs (for the PMD dataset) between real and (1000) generated samples by the VAE . . . . .	76
4.6	Note frequency bar graphs (for the Undertale dataset) between real and (1000) generated samples by the VAE . . . . .	77
4.7	Note frequency bar graphs (for the Classic dataset) between real and (1000) generated samples by the VAE . . . . .	78
4.8	Generated samples for GAN during training. . . . .	80
4.9	GAN loss history for each dataset . . . . .	81
4.10	Discriminator Accuracy history for each dataset . . . . .	81
4.11	Samples generated by the GAN . . . . .	82
4.12	Note frequency bar graphs (for the PMD dataset) between real and (1000) generated samples by the GAN . . . . .	83
4.13	Note frequency bar graphs (for the Undertale dataset) between real and (1000) generated samples by the GAN . . . . .	84
4.14	Note frequency bar graphs (for the Classic dataset) between real and (1000) generated samples by the GAN . . . . .	85
5.1	GAN generating fully black images up to epoch 50 . . . . .	91
5.2	VAE Track 11 (Undertale) sheet . . . . .	93

# List of Tables

3.1	Supported MIDI Messages . . . . .	35
3.2	MIDI Message parameter types . . . . .	35

# Chapter 1

## Introduction

Artificial Intelligence (AI) and specifically machine (or deep) learning has been rising in popularity in the last decades. It is a thriving field whose applications are dedicated to automate laborious routines, conceive the essence of speech or images, identify patterns that can contribute to medical analysis and so on. AI is a buzzword which consists of hundreds subsets and one of them is deep learning. Even deep learning has no consensual definition, it consists of Machine Learning (ML) techniques based on artificial neural networks. Those networks are inspired by the biological human neural system and how they process and weigh information is fairly similar. A set of neurons fire up given that information from our senses as input and the network outputs a feeling, a thought, a memory, or even possible actions or decisions for a matter. Same goes to the process of artificial neural networks, which if trained properly, they can be used for classification, action selection (Deep Q Agent), data compression (Autoencoders), data generation (generative models such as GANs, Variational Autoencoders), and much more (Goodfellow, Bengio, and Courville 2016).

In the case of generative modeling, where this work is based on, the networks are trained upon processed musical files (MIDI) which are translated to images. As the models become trained, they are able to produce images that closely resemble the ones from the input set of data. The goal in such cases is not to produce same results, but to explore and generate original data.

Neural networks are ideal for such tasks as they are able to handle massive amounts of data, while being as efficient as possible with affordable computing power. This case is also presented in this thesis as we offer models that can run even on common CPU environments, once they have been trained upon thousands of samples.

## 1.1 Computer-Based Musical Systems

Researchers have been trying to create a bridge between computers and music as the first computer systems were created. It has been an active topic for many decades with different attempts and approaches (Briot, Hadjeres, and Pachet 2017).

The first music generation program was introduced in the early 1960s by Max Mathews at Bell Labs, a research center run by phone company AT&T. Mathews was mainly researching telephony-related fields, but worked on music software in his spare time. He named this program MUSIC and later on he created different versions indicated by Roman numerals. MUSIC produced its first sound in 1957, playing a 17 seconds long melody named "The Silver Scale". The same year, the first score was composed by a computer named "The Illiac Suite". It was produced by the ILLIAC I computer at the University of Illinois at Urbana-Champaign (UIUC) in the United States by Lejaren A. Hiller and Leonard M. Isaacson, both musicians and scientists. It was the first example of algorithmic composition, by utilizing stochastic models, such as Markov chains, for generation, as well as rules to filter generated material according to desired properties.

In 1983, Yamaha DX7 synthesizer was released by Chowning, which was a model of synthesis based on frequency modulation (FM). The same year the MIDI interface was launched, which created a link between software and instruments (including the DX7 synthesizer). Another case was the development of the Max/MSP environment, by Puckette at IRCAM, which allowed for real-time synthesis and interaction.

In the year of 1962, Iannis Xenakis used computers to compose music with a stochastic approach. He created the ST Program, a stochastic music program,

written in FORTRAN, which used probability distributions and calculated various possibilities designed by the composer in order to generate musical samples. Another approach was that of Ebcio Glu, which followed the direction of "The Illiac Suite", and used certain rules and constraints to specify the style of a given corpus. He created CHORAL, a composition program of a four-part chorale in the style of Johann Sebastian Bach. In the late 1980s, a system called Experiments in Musical Intelligence (EMI), developed by David Cope, extended that approach (Cope 2000). His program learned from a corpus of a composer's scores and was able to create its own constraints and rules.

## 1.2 History of MIDI

The MIDI format played a huge role in the music industry, as it offered a general compatibility between programs and electronic instruments. This work's interaction with music is done solely by the use of MIDI files, as they are excellent in offering discrete information, which is ideal for the deep learning method presented.

MIDI, which is an acronym for Musical Instrument Digital Interface, first saw the light in the early 1980s. Its main purpose was to synchronize electronic musical instruments manufactured by different companies, as each company had their own standards of synchronization, such as CV/gate and Digital Control Bus (DCB). This lack of synchronization led Ikutaro Kakehashi, founder of Roland to bridge this gap, as he felt this limitation would delay the growth of electronic music. Ikutaro went on to propose this idea of creating such a machine of synchronization to Oberheim Electronics found by Tom Oberheim developer of Oberheim Systems, which he thought was difficult to carry and use and wanted a simpler design. Together with Dave Smith from Sequential Circuits, they discussed a simpler and cheaper alternative with American and Japanese companies (Yamaha, Korg, Kawai), and initially, only Sequential Circuits and Japanese companies were interested.

Dave Smith and engineer Chet Wood from Sequential Circuits managed to create a universal interface for communication between equipment from different companies. Together they proposed the standard in a paper, Universal Synthesizer

Interface, presented at the Audio Engineering Society show in October 1981. The standard went under modification by Roland, Yamaha, Korg, Kawai and Sequential Circuits, and Kakehashi showed preference for the name Universal Musical Interface (UMI), pronounced you-me, but Dave Smith felt it was banal. Also Smith liked the word "instrument" instead of "synthesizer" so the proposition Musical Instrument Digital Interface (MIDI) was accepted and announced in the October of 1982 issue of Keyboard.

Smith demonstrated his design by connecting a Prophet 600 to a Roland JP-6, two different manufactured synthesizers, at the 1983 NAMM Show. Those were the first synthesizers that were released with the MIDI standard (1982), and later in 1983, Roland TR-909, the first MIDI drum machine, and Roland MSQ-700, the first MIDI sequencer, were released. The first computer to support MIDI, the NEC PC-88 and PC-98, was released in 1982.

Later in 1984, at the Summer NAMM Show in Chicago, the MIDI Manufacturers Association (MMA) was formed. The MIDI 1.0 Detailed Specification was then published at the 1985 Summer NAMM show by MMA. The standard advanced in 1991 by adding standardized song files (General MIDI) and adapted to new connection standards, such as USB and FireWire. In 2016, the MIDI Association was formed to continue overseeing the standard. The MIDI 2.0 standard was proposed in January 2019 and was introduced at the 2020 Winter NAMM Show.

### **1.3 Autonomy versus Assistance**

When referring to computer-based music generation, or any generative problem, there are two approaches for tackling such a task.

- Autonomy: developing autonomous music-generating systems aimed at creating musical patterns from scratch.
- Assistance: programs that aim to assist musicians in further exploring musical possibilities, by receiving inputs, rules or restrictions by human hand.

The task of automated music generation is quite interesting, as it can offer explo-

ration about the process of composition. Two recent examples, being the deep-learning based Amper and Jukedeck systems/companies, target the creation of original music for commercials and documentary.

On the other hand, assistance computer-based music programs are developed in order to assist humans in various steps of music creation: composition, arranging, orchestration, production. In practice, composers rarely create their own music from scratch, as they reuse or adapt based on their experience, consciously or unconsciously, while following music principles and rules, such as harmony and scales theory. The assistance program then suggests or complements the various musical steps, which the human composer tackles. Such examples of assistance computer-based music programs are the FlowComposer environment developed at Sony CSL-Paris and the OpenMusic environment developed at IRCAM.

This work is aimed towards automated music generation, more specifically at small two-measure (depending on the time signature) musical parts. Most music generative tasks follow the autonomous approach, although more and more systems begin to address the option of assistance and human interaction.

## 1.4 Deep Learning

The generality that deep learning techniques offer is ideal for such generative tasks. As opposed to handcrafted models, grammar-based, or rule-based music generation systems, a machine learning-based generation system can be agnostic, as it learns a model from an arbitrary corpus of music. Furthermore, a certain model can be used to train upon various and different musical genres and themes. Therefore, as more large scale musical datasets are made available, a machine learning-based generation system can detect the musical style from a corpus and generate new relevant content. Rather than creating a complicated system filled with rules and grammars, there is the approach of a deep learning model, which is able to process raw unstructured data, from which it identifies musical distributions and patterns, as its layers extract higher level representations adapted to the task.

# Chapter 2

## Background

The knowledge required to understand this work is split into two sections, (a) basic music theory and (b) the machine learning technique called neural networks. Basic principles are only required for both of the fields in order to comfortably read this work. Specifically,

- little to no music theory knowledge is required, as with the use of MIDI files the data is discrete, and
- with basic understanding of dense neural networks and training processes, other layers and concepts are easily absorbed.

### 2.1 Basic Music Theory

Music theory examines the fundamentals of music and provides principles and rules, which are a guideline for musicians and allow them to create musical sequences. Music theory has been practiced for thousands of years and it branches to many topics and levels. This section will focus only to the very basics, as it is deemed enough to understand the following sections (Schmidt-Jones 2013).

## 2.1.1 Measure and Bar Lines

Sheet music is mostly read and written in a format of five horizontal parallel lines, also known as the pentagram, which is then sequenced into measures (Figure 2.1). Measures exist to organize long pieces of music into smaller units. Professional musicians read music or perform a piece of music in real-time, as they are able to read each measure at a time.

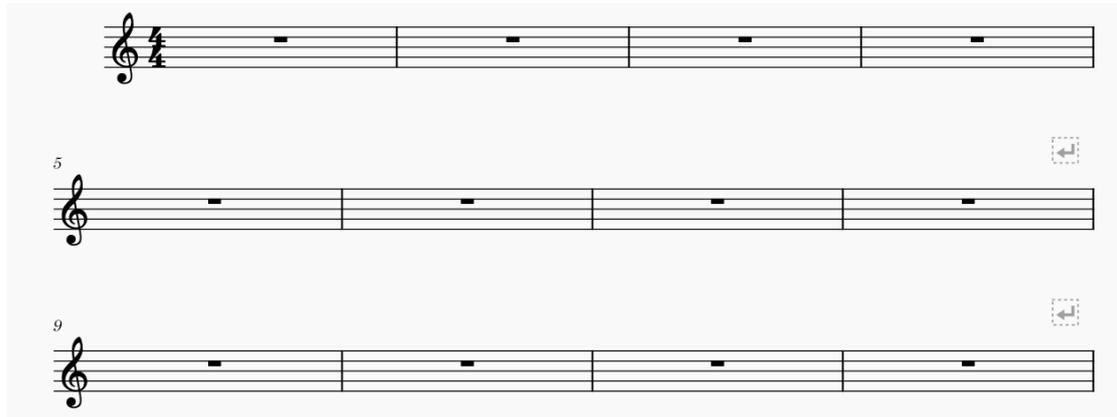


Figure 2.1: Measures and pentagram

Measures are split by the use of vertical bar lines. There are different bar line types that offer musicians easy sight-reading and better overall performance.

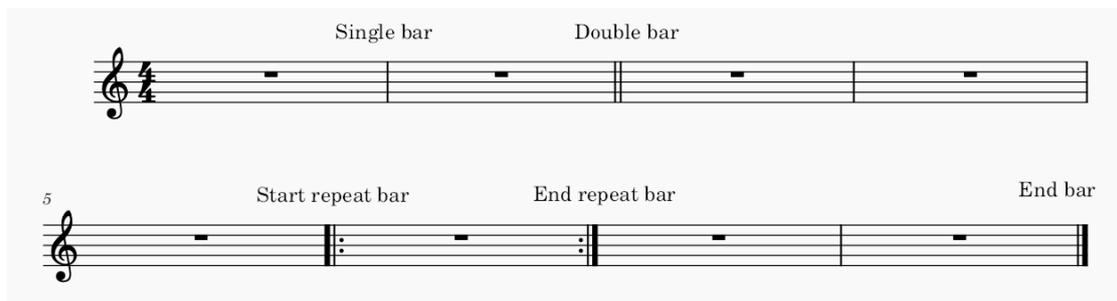


Figure 2.2: Bar line types

Bar line types (Figure 2.2):

- Single bar line, a single vertical line that indicates the end of one measure and the beginning of another.

- Double bar line, two side-by-side vertical lines, indicating the end of one section and the beginning of another.
- End bar lines, two vertical lines, one bolder, which indicate the end of the musical composition.
- Start repeat, two vertical lines, one bolder, with two dots on the right, indicating the following sequence will be repeated.
- End repeat, two vertical lines, one bolder, with two dots on the left, indicating the musician should return to the last start repeat bar lines and replay the sequence once.

For obvious reasons there cannot exist start repeat-type bar lines without end repeat, and the other way around. Furthermore the sequence inside the repeat bar lines is repeated once, unless indicated otherwise, usually with a number next to the end repeat bar lines.

### 2.1.2 Time Signature

A musical time signature is a fraction, which shows the number of beats per measure (the top number in a time signature, numerator) and the duration of each beat (the bottom number in the time signature, denominator), as shown in Figure 2.3. For instance,  $3/4$  time signature indicates there are three beats per measure and each beat has the duration of a quarter note.

The three basic types of time signatures (Figure 2.3):

- Simple: The most common simple time signatures are  $2/4$ ,  $3/4$ ,  $4/4$ , and  $2/2$ . The  $4/4$  time signature indicates that there are four beats per measure (numerator) and each beat counts as a fourth note (denominator). Furthermore for  $2/4$  and  $3/4$ , there are two and three quarter note beats per measure, respectively.
- Compound: Common compound time signatures include  $9/4$ ,  $6/8$ , and  $12/8$ . The beat of a piece of music with a compound time signature is broken into a three-part rhythm. In each of the cases above, quarter or eighth notes are

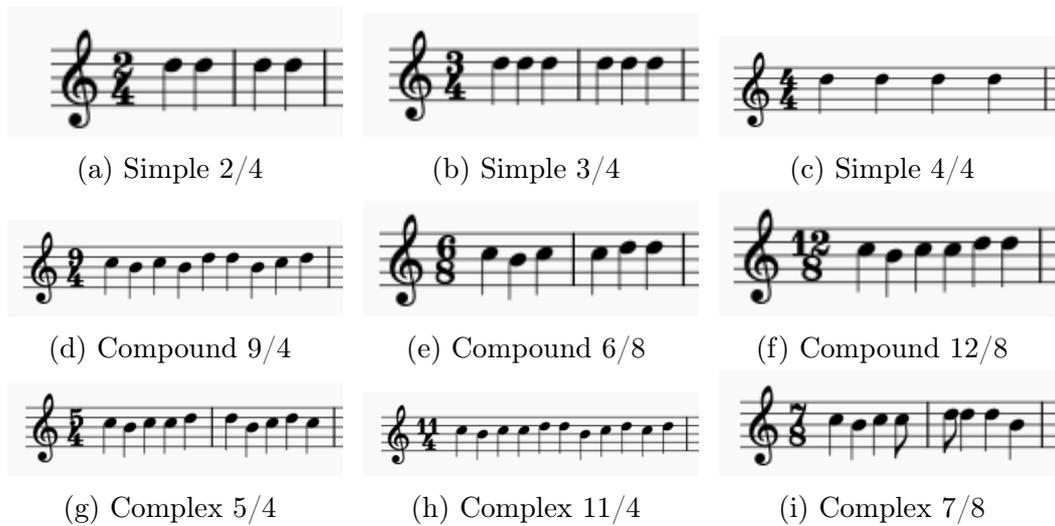


Figure 2.3: Different types of Time Signatures.

combined in multiples of three.

- Complex: Complex time signatures are more common in music written after the nineteenth century. Complex time signatures do not follow typical duple or triple meters. Examples of complex time signatures include: 5/4, 11/4, and 7/8.

While there are many types of time signatures, below are the most commonly used by musicians.

- 2/4, two quarter-note beats per measure.
- 3/4, three quarter-note beats per measure.
- 4/4, four quarter-note beats per measure, also known as common time and also notated as a "C".
- 2/2, two half-note beats per measure, also known as cut time and notated as a "C" with a vertical slash through it.
- 6/8, six eighth-note beats per measure.
- 9/8, nine eighth-note beats per measure
- 12/8, twelve eighth-note beats per measure.

The first note of every bar or measure is called the downbeat. Every measure has strong (loud) and weak (silent) beats. In a time signature like 4/4, the first beat of every measure is the strongest beat, and the third beat is also a strong beat. Beats two and four are weak beats. Regardless, this can vary between different genres and compositions.

### 2.1.3 Tempo

Tempo refers to the speed of a section of music. It can be indicated in metronome markings using beats per minute (BPM) or by using descriptive words (traditionally Italian words describe tempo like, *adagio* or *andante*).

Typically the metric of BPM is self-explanatory. As its name suggests, it indicates the count of beats a minute contains. For example, a musical part with 60 BPM results in one beat per second, and with 120 BPM it is twice as fast resulting in two beats per second.

In terms of musical notation, a beat almost always corresponds with the piece's time signature. This means that the denominator of the time signature counts as one beat. For example in 2/4, 3/4, 4/4 a quarter note counts as a beat, in 6/8, 9/8, 12/8 an eighth note counts as a beat and so on.

BPM is the most precise way of tracking real time of musical sequences, both as a whole duration and as a time offset between notes. It is effective due to its precision, offering parallel play between instruments without temporal confusion in environments like bands or orchestras, as well as time consistency in solo play. If the tempo is not respected, the sound can feel out of place, even in untrained ears. Typically, professional musicians are able to sustain tempo during their play, as well as understand the time of a composition before performing it. Usually the most skilled musicians in tempo are percussion players as their instruments are specialised in rhythm rather than melody.

## 2.1.4 Note Values

Individual notes within a measure last for a specific fraction of the duration of that measure. They indicate both their pitch depending on where they are set in the pentagram, and their duration or time offset from the previous played note.

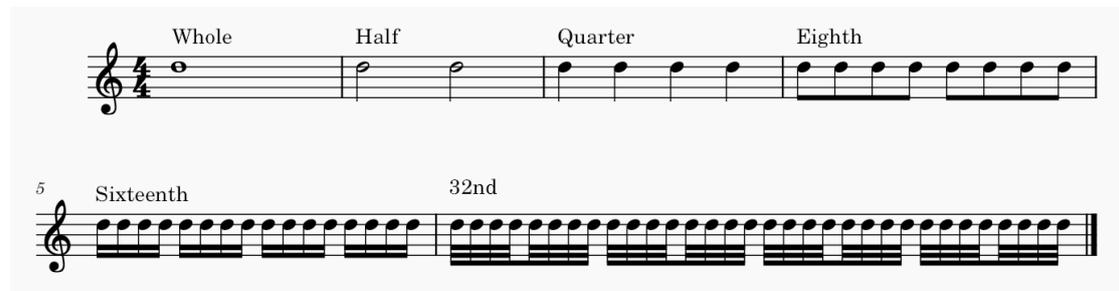


Figure 2.4: Note values

Notes in a part or sheet of music indicate their duration and their offset from the previous note played inside a measure. In a common 4/4 time signatures the types of notes are (Figure 2.4):

- Whole note: a note that covers the whole 4-beat measure, an empty circle
- Half note: covers half of a 4-beat measure, an empty circle with a vertical bar on top right or bottom left
- Quarter note: covers a fourth of a 4-beat measure, filled circle with a vertical bar on top right or bottom left
- Eighth note: covers an eighth of a 4-beat measure, filled circle with a vertical bar on top right or bottom left and a tail
- Sixteenth note: covers a sixteenth of a 4-beat measure, filled circle with a vertical bar on top right or bottom left and a double tail

Some musical parts have smaller subdivisions, such as 32nd, 64th and 128th notes, but musicians typically set tempos that work around the need for such small durations and offsets.

In summary, musical parts have principles and rules that allow musicians to translate and read them. In any case this section is not intended to be a lesson about

music theory. It simply contains some basic principles that are required in order to assist in the engineering of generative modeling.

## 2.2 Deep Learning

Machine Learning has been thriving in the last decades and its techniques and tasks in everyday laborious tasks are showing great results. Their ability to process unstructured data and extract higher level representations adapted to the task suit very well for our music generation purpose. It is a field that branches into many categories, such as classification and regression algorithms (Linear Regression, Logistic Regression, Decision Tree, SVM and so on), which are categorised as supervised or unsupervised learning, agent-based algorithms (Q-learning), neural networks and the list goes on (Goodfellow, Bengio, and Courville 2016). Our main focus is a branch called Deep Learning, which utilizes dense, multi-layered neural networks that are trained upon a certain dataset and are able to identify patterns in order to deliver a decision, usually to solve a classification task.

The aim of this section is to present the very basic concepts of deep neural networks, which will allow a smooth transition to unique and complex techniques used by the generative models (Piyush Madan 2020).

### 2.2.1 Deep Learning and the Human Brain

In order to create systems that learn similarly to how a human is able to learn, the architecture of a deep learning network tends to mimic the human brain and its process in learning. Deep learning essentially was inspired by neurology, as similarly to how neurons form the fundamental building blocks of the brain, its architecture is assembled by layers of neurons, which are computational units allowing modeling of nonlinear functions.

The main point of a deep learning network is the neuron, which receives a list of input signals and transforms them into output signals, inspired by the human neuron, which in a biological manner transmits electrical pulses throughout our nervous system.

The neuron's purpose is to understand data representation by stacking together many layers, each layer responsible for understanding some pattern of the input. A layer is a collection of neurons, with the aim of learning to detect a repeating occurrence of values. Each layer of neurons is responsible for interpreting a specific pattern within the data. The interconnection of neurons in the human brain is replicated by the use of layers forming a network, so this computational architecture is called neural networks (or artificial neural networks).

### 2.2.2 Artificial Neural Networks

An Artificial Neural Network (ANN) processes input data by the use of a large number of neurons. Each layer contains neurons that will in turn feed their outputs to the ones of the next layer, as shown in Figure 2.5.

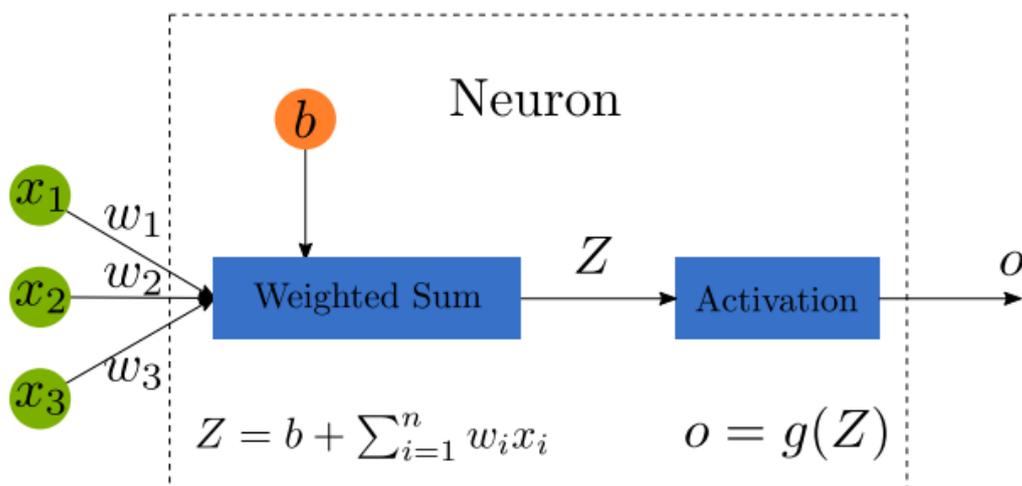


Figure 2.5: Neuron (Artificial Neural Network)

The processing of input data by each neuron is done by two steps of calculation:

- compute a weighted sum of the inputs and an offset, and
- pass the weighted sum through an activation function to deliver an output.

### Step 1: Compute weighted sum

$$Z = b + \sum_{i=1}^n w_i \cdot x_i \quad (2.1)$$

- Inputs  $x_1$  through  $x_n$ , are denoted by a vector  $X$  of size  $n$ .  $X_i$  represents the  $k$ th entry in the dataset. Each entry in the dataset contains  $n$  dependent variables.
- Weights  $w_1$  through  $w_n$ , are denoted as a matrix  $W$ . They represent how strong a connection is between a certain input and the neuron, indicating this part of the input affects little or much the outcome of the neuron.
- A bias term  $b$ , which is a constant.

### Step 2: Activation Function

$$o = g(Z) \quad (2.2)$$

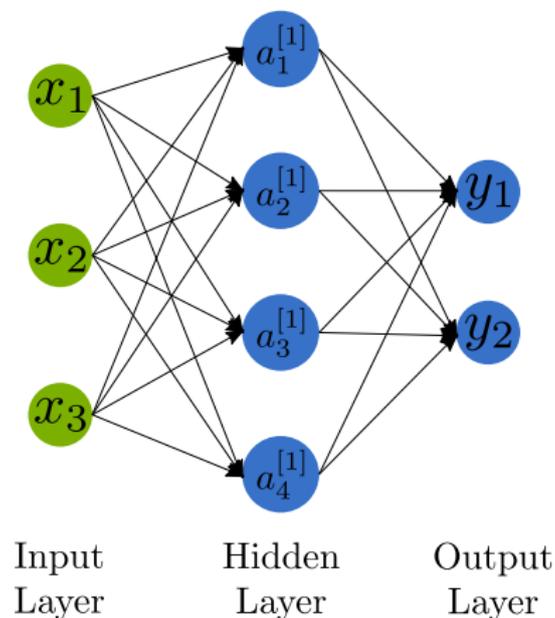
The output of step 1,  $Z$ , is passed through an activation function resulting in the output of neuron  $o$ . The function  $g$  is a mathematical function that allows transformation of the outputs to a desired non-linear format, before it is passed to the next layer. It essentially normalizes the summation result to a desired range, which helps identifying whether the neuron "fired", meaning whether its output will affect the next layer.

### 2.2.3 Shallow Neural Network

The most basic form of a neural network is the shallow neural network, which essentially consists of three layers: input layer, hidden layer and output layer. The input layer refers to the data sent by the input set. Hidden is the layer that consists of neurons each being connected to all of the inputs, dedicated in computations presented in the previous section. The last layer, the output layer, consists of a certain number of neurons (depending on the task), that determine the output of the network, in other words the prediction. A neural network with just one hidden

layer is called shallow, as shown in Figure 2.6.

Figure 2.6: Shallow Neural Network



The direction in which the data is passed once through the network (usually presented from left to right) is called a forward propagation. After one forward pass is completed, the output layer compares its results to the actual ground truth labels and adjust the weights in a reverse way based on the differences between the ground truth and the predicted values. The process of a backward pass through the neural network is also known as back propagation.

Each layer consists of  $n_i$  neurons,  $i$  indicating the label of each layer, total  $l$  layers. Between two layers there are connections based on the number of neurons each layer consists of. For example, between layers  $l_1$  and  $l_2$  there are  $n_1 \cdot n_2$  connections as each neuron of  $l_2$  connects with every neuron of  $l_1$ . Usually this helps determine the load of the network, as it indicates the number of trainable parameters there are in the network, by considering all the connections. The weights of each connection are called trainable parameters, because their value will be changed during back propagation based on the error of prediction.

## 2.2.4 Deep Neural Network

Deep Neural Networks resemble the architecture of a shallow neural network, with the difference being that they consist of more than one hidden layer (Figure 2.7). As its name suggests, they have depth based on the number of extra hidden layers. Each neuron in a hidden layer is connected to all the others of the next layer. Selecting the number of hidden layers is determined by the nature of the task and/or the size of dataset.

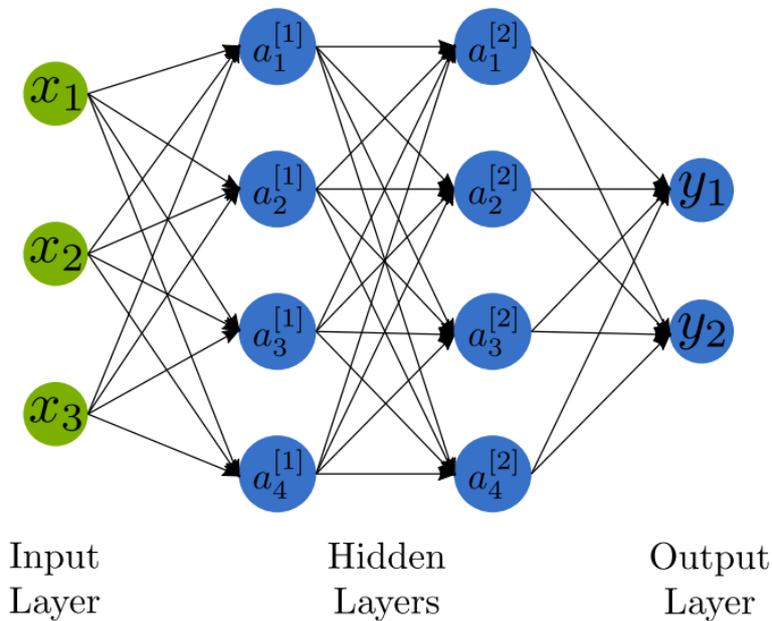


Figure 2.7: Deep Neural Network

The calculations and processes of forward pass or backward pass for each layer and neuron remain the same as suggested in the previous sections.

## 2.2.5 Back Propagation

The Back Propagation algorithm (LeCun et al. 1988) is used to effectively train a neural network through a method called chain rule. After each forward pass, the algorithm performs a backward pass, while adjusting the trainable parameters

(weights and biases). The algorithm is used to effectively train a neural network through a method called chain rule. In simple terms, after each forward pass through a network, backpropagation performs a backward pass while adjusting the model's parameters (weights and biases).

The basics of the process of back propagation follow these steps:

- The network works to minimize an objective function, for example, the error incurred across all points in a dataset.
- At the output layer, the network must calculate the total error (difference between actual and predicted values) for all data points and take its derivative with respect to the weights at that layer. The derivative of error function with respect to the weights is called the gradient of that layer.
- The weights for that layer are then updated based on the gradient. This update can be the gradient itself or a fraction of it. This fraction is determined by a factor known as the learning rate, and it controls how large the steps that are taken to change the weights are.

This process is then repeated for the previous layer and continues until the first layer is reached. During this process, values of gradients from previous layers can be reused, making the gradient computation efficient.

## 2.2.6 Convolutional Neural Networks

Convolutional layers (Brownlee 2019b) are the basis of creating convolutional neural networks. A convolution is an application of filters to an input data, resulting in a feature map, indicating the location of certain features in an input, such as an image.

Convolutional neural networks learn automatically a large number of filters specified to identify patterns and features among the training dataset, usually predicting such features in image data. This feature detection is usually applied in classification tasks. CNNs are specialized to work with two-dimensional image data most of the times, but the input can be one-dimensional and three-dimensional data.

The convolutional layer performs an operation called a "convolution", that is a linear operation that involves a multiplication of a set of weights with the input, like a typical neural layer. The main distinction is that the set of weights, also called filter or kernel, and the input data are both two-dimensional.

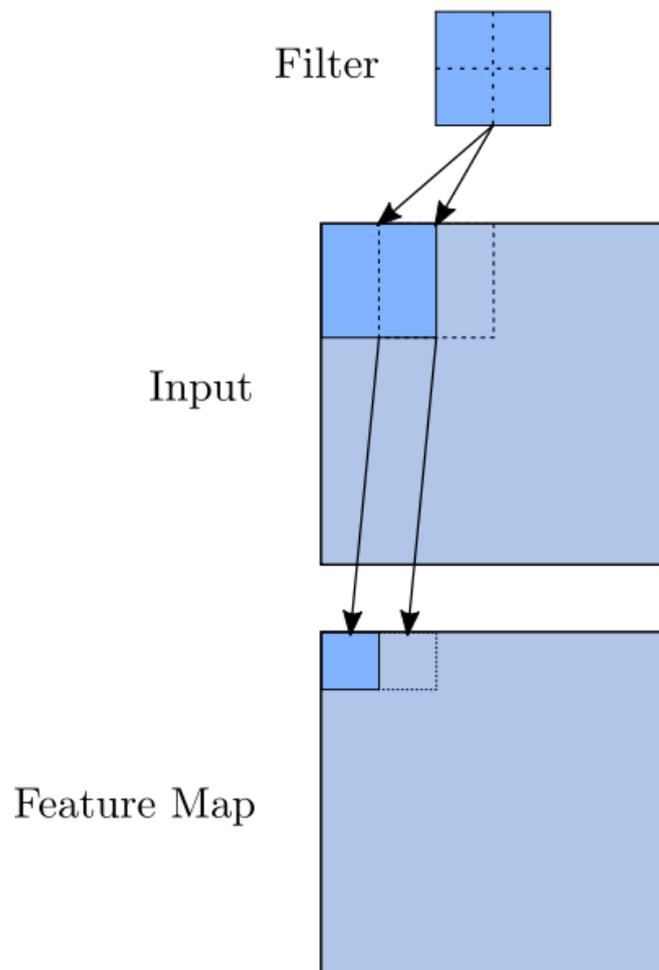


Figure 2.8: Filter application

The operation between the data and the filter is a dot product. Each element of the filter is multiplied with each element of a filter-size patch of the input image, then those values are added resulting into a single value (this operation is also referred to as a "scalar product"). This process repeats for each filter-size patch

of the input image, resulting in a set of dot products which is the feature map (Figure 2.8). This means that the filter size must be smaller than the input image, to allow detection in multiple regions in the image. The filter is applied using this operation from left to right and top to bottom of the image.

Convolutions have been a common technique used in computer vision. Filters were designed to be applied to images for feature map extraction. For example, there are filters that are able to identify vertical or horizontal lines. Those filters are dragged through the whole input resulting in a feature map indicating the presence of vertical or horizontal lines. The application of both of those filters result in a feature map highlighting the edges of the input image.

Utilizing multiple filters offers a greater feature detection in image data. The use of multiple weighted filters that are able to be altered for desired feature extraction is what makes CNNs successful. The network will naturally learn what types of features to detect from the input dataset. Specifically, training under stochastic gradient descent, the network learns to extract features from the image that minimize the loss for the task the network is designed to solve, e.g. extract features that are most useful for distinction between images of cats and dogs.

Convolutional layers are often stacked allowing for a hierarchical decomposition of the input. The filters that operate directly on the raw pixel values will learn to extract low-level features, such as lines. The next set of layers are able to extract features that are combinations of lower-level features, such as ones that compromise multiple lines to express shapes. This process continues until the deepest layers are able to identify complex shapes, such as faces, animals etc. This is the reason that CNNs are desired in solving tasks that often require complex feature detection.

# Chapter 3

## Methods

In this section the steps taken to achieve the goal of this work will be presented. The whole process is split in two sub-sections, first the preparation and preprocessing of the training data (MIDI files) is presented, and then the architecture of the training models (neural networks).

The most essential part for this work regarding the preprocessing of musical data, is tempo, time signatures and note duration. In practice, it will be apparent that MIDI files contain information about note values and duration in a discrete way. Regardless, most time attributes are set into high time values, or ticks, and we aim to reduce that value for the image translation, e.g. a note being played with a duration of 256 ticks, needs to be reduced in order to have smaller images (less pixels) with the same information. Furthermore the images that will be created need to have enough pixels for at least two measures, while at the same time being able to contain different time signature types, as well as normalizing the time durations and offsets in a way that information is not lost. It is also apparent that other concepts of music theory, such as keys, scales and so on are not present due to the fact that they are not needed. This work is based on training upon musical parts that already follow music principles and rules. Regardless understanding more musical aspects could help evaluate the trained generative models, as in a more complex and better performing environment such concepts could be helpful.

The training part by itself contains two solutions, using a Variational Autoencoder (VAE) and a Generative Adversarial Network (GAN), for the purpose of research and comparison. Both of these models are trained with the same type of training data.

### 3.1 Preprocessing of training data

MIDI files are used broadly today in the music industry. If opened with certain applications, there can be seen sequences of notes in a tabular form. The current work utilizes a python library, called Mido (Bjørndalen 2013), which offers easy access to such files and their information, as well as the creation of MIDI files, which is handy for generation purposes. Aside from Mido, music21 (Cuthbert 2006) is also a great alternative to extract information from MIDI files.

With Mido, by opening a .mid (.midi) file from any directory and assigning it to a Mido object, this object contains a `tracks` attribute, that is a list of tracks. Each track on its own is a list of messages and meta messages. Each message has a `time` attribute. Time is set to delta time (in ticks), which will be explained later in detail.

The following sections determine the type of a MIDI file:

- type 0 (single track): all messages are saved in one track
- type 1 (synchronous): all tracks start at the same time
- type 2 (asynchronous): each track is independent of the others

When creating a new file, the type must be selected by the creator, but can be changed anytime.

Also, each Mido object contains a playback `length` item, the total musical sequence time in ticks. It should be noted that it's not possible to access the `length` property of a type 2 file, as a `ValueError` will be raised, because the playback time of an asynchronous file cannot be computed.

### 3.1.1 Messages

As mentioned previously, each track is a list of messages. Those messages hold the information of the musical sequences. Most common messages are the ones that press or release notes:

```
>>> mido.Message('note_on')
<message note_on channel=0 note=0 velocity=64 time=0>
```

Arguments can be passed like this:

```
>>> mido.Message('note_on', note=100, velocity=3, time=6.2)
Message('note_on', note=100, velocity=3, time=6.2)
```

where the message type can be:

- 'note\_on': a note is pressed
- 'note\_off': a note is released

Note refers to all the possible notes as integers in the range of 0 to 127. This data value  $d$  corresponds to a frequency  $f$  in Hz by the MIDI Tuning Standard (MTS), which is a specification of precise musical pitch agreed to by the MIDI Manufacturers Association in the MIDI protocol. To get the frequency  $f$  of any data value  $d$  the following formula is used:

$$f = 2^{(d-69)/12} \cdot 440 \text{ Hz} \quad (3.1)$$

Inversely, the conversion from frequency  $f$  to a midi note number  $d$  can be done with this formula:

$$d = 69 + 12 \log_2 \left( \frac{f}{440 \text{ Hz}} \right) \quad (3.2)$$

Velocity is the volume at which the note will be played (high value of velocity is loud, low value is silent).

As far as the time attribute is concerned, it is specified as delta time in ticks. This essentially means that each message is accompanied by that attribute and

determines how much time has passed since the previous message. In theory music there is also the beat, which is essentially a quarter note. In a signature of a 4/4 song, each measure has 4 beats. Each MIDI file has a value stored, called `ticks_per_beat`, which declares how many ticks a beat counts. The last time-related attribute that a track requires is the BPM (beats per minute) of a song (usually known in music as tempo, although in MIDI tempo refers to something else), and it counts how many beats a minute contains. For example a track that has 4 beats per minute and 3 ticks per beat is presented in Figure 3.1:

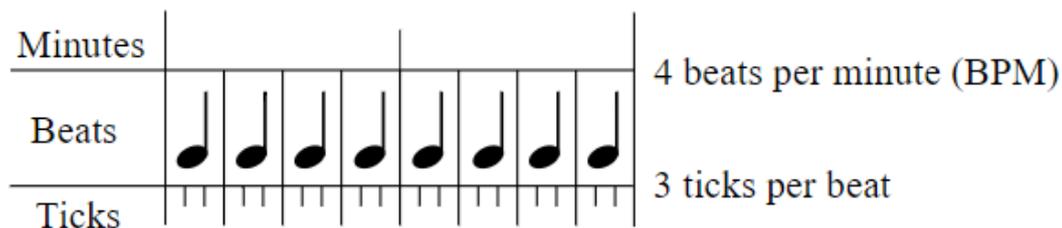


Figure 3.1: Example of a track with time attributes

Tempo in MIDI is not given as beats per minute, but rather in microseconds per beat. The default tempo is 500000 microseconds per beat, which is 120 beats per minute. There can be used some methods, like `bpm2tempo()` and `tempo2bpm()`, for conversions to and from beats per minute.

There are many other types of messages that refer to program, control, channel (used for parallel instruments and note effects), but for simplicity and due to the fact that the model is trained only on one instrument, the piano, the MIDI files are filtered and simplified to contain only 'note\_on' and 'note\_off' messages. The full list of messages is shown in Table 3.1, whereas Table 3.2 provides information about the various parameters.

Table 3.1: Supported MIDI Messages

<b>Name</b>	<b>Parameters / Arguments / Attributes</b>
note_off	channel note velocity
note_on	channel note velocity
polytouch	channel note value
control_change	channel control value
program_change	channel program
aftertouch	channel value
pitchwheel	channel pitch
sysex	data
quarter_frame	frame_type frame_value
songpos	pos
song_select	song
tune_request	
clock	
start	
continue	
stop	
active_sensing	
reset	

Table 3.2: MIDI Message parameter types

<b>Name</b>	<b>Valid Range</b>	<b>Default Value</b>
channel	0..15	0
frame_type	0..7	0
frame_value	0..15	0
control	0..127	0
note	0..127	0
program	0..127	0
song	0..127	0
value	0..127	0
velocity	0..127	64
data	(0..127, 0..127, ...)	() (empty tuple)
pitch	-8192..8191	0
pos	0..16383	0
time	any integer or float	0

### 3.1.2 Meta Messages

Meta Messages in MIDI behave like normal messages and can be created in the usual way. They store information either related to the song or musician in general, like copyright information, text, track name, lyrics, or attributes that affect the track directly, like tempo, time signature.

For the pre-processing of the training data, meta messages were not utilized in any way. Meta messages were implemented and inserted at the generation of the musical sequences, purely for presentation reasons.

### 3.1.3 Music to data

This section presents how the MIDI files were processed into proper training data for the models. The generative models chosen for this task were VAE (Variational Autoencoder) and GAN (Generative Adversarial Network). Due to their nature of utilizing convolutional layers, translating MIDI files to images is a rational approach. The goal of the pre-processing procedure was to convert MIDI files, or message sequences, to images.

#### 3.1.3.1 Message filtering

Firstly, as stated before, the messages need to be filtered in order to contain only message types of note pressing or releasing. This means that each message that is not type of "note\_on" or "note\_off" is discarded. Regardless, irrelevant typed messages are still processed and their time attribute is taken into account, otherwise the normalized track would not be the same.

The general idea is, each message gets parsed, if its type is neither "note\_on" nor "note\_off" its time attribute is saved to a variable and the message is deleted from the sequence. That variable is utilized and its value is added to the last message that was saved. This basically means that each message of type "note\_on" or "note\_off" is kept as is, and the other types are deleted, but their time values are added to their previous messages, so the sequences remain the same in terms of total duration. This filtering allows for an easier mapping of the musical se-

quence to an image. Also messages that are "note\_on" with velocity value 0, or are "note\_off", but there was not a previous "note\_on" with the same "note" value, are also filtered out. Some examples are presented below which indicate the application of such filtering.

```
"Filtering of a note_off with no prior note_on"
```

```
"Initial sequence:"
```

```
note_on      channel=0  note=74    velocity=81  time=0
note_off     channel=0  note=70    velocity=30  time=512
note_off     channel=0  note=74    velocity=0   time=256
...
```

```
"Filtered sequence:"
```

```
note_on      channel=0  note=74    velocity=81  time=0
note_off     channel=0  note=74    velocity=0   time=768
...
```

```
"Filtering of a note_on with velocity=0"
```

```
"Initial sequence:"
```

```
note_on      channel=0  note=74    velocity=81  time=0
note_on      channel=0  note=80    velocity=0   time=512
note_off     channel=0  note=74    velocity=0   time=256
...
```

```
"Filtered sequence:"
```

```
note_on      channel=0  note=74    velocity=81  time=0
note_off     channel=0  note=74    velocity=0   time=768
...
```

```
"Filtering anything besides note_on, note_off message types"
```

```
"Initial sequence:"
```

```
note_on      channel=0  note=74    velocity=81  time=0
control_change channel=0  control=6  value=0     time=128
control_change channel=0  control=101 value=0     time=128
set_tempo    tempo=800000  time=128
set_tempo    tempo=850000  time=128
```

```

note_off      channel=0  note=74      velocity=0   time=256
...

"Filtered sequence:"
note_on       channel=0  note=74      velocity=81   time=0
note_off      channel=0  note=74      velocity=0   time=768
...

```

### 3.1.3.2 Time normalization

Besides message filtering, while parsing through messages, their time attribute is reduced. This is essential due to the fact that each time tick will be treated as a pixel when converted to an image. Most MIDI files contain different ticks per beat values (which is how many ticks a beat counts, Figure 3.1), and those values are high (for example: ticks\_per\_beat=1024) which would produce high scale resolution images. It is impossible to find perfect MIDI files for this task, with small value of ticks per beat, due to the multiple range of applications or programs dedicated to creating digital music, plus the various musicians and creators who use them. Each new time attribute is calculated given the file's ticks per beat and its value is set to 12 (the reason is given in the next section), in order to attain uniformity among the tracks.

$$new\_time = round\left(\frac{current\_time * 12}{ticks\_per\_beat}\right) \quad (3.3)$$

### 3.1.3.3 Message sequences to Images

Images prove essential for generational purposes using VAEs and GANs. The converted images have width and height set to 96 by 96 respectively. Each image's  $x$  axis refers to time and  $y$  axis refers to every possible piano pitch or note.

The ticks per beat is set to 12 for each track, so  $\frac{96}{12} = 8$  quarter notes are contained in each image. In a 4/4 time signature (which is the most common case) this counts as two measures. The ticks per beat attribute is set to 12 to cover for most time signatures and speed of note presses. If the ticks per beat is set to 1, and

a track contains eighth or sixteenth notes, then information is lost and the track gets corrupted. Speed exceeding sixteenth notes is rather rare so a ticks per beat set to 4 is acceptable, but some songs contain triplet notes (3 notes counted in a quarter) so  $4 \cdot 3 = 12$  is the safest choice. This value essentially covers cases where information is not lost, if the track contains sixteenth triples. Regardless, files can still be found with unexpectedly low message time values (grace notes) and by rounding, information can also be lost.

The conversion process is done by creating an image of a whole song at first. An empty list, or image, is initialized with dimensions 96 height and `song_duration` width with values set to 0 (black pixels). Also a dictionary is used to hold the active notes and each note is deleted from this dictionary when a `note_off` message with the same pitch value is presented. Each message is parsed and if its type is `note_on`, it is added to the dictionary, plus its time value is saved to a `message_time` variable. For each item in the dictionary (active notes), pixels are set to 255 (white) depending on the `message_time` variable. After each note from the dictionary is assigned to the list (white pixel), if the message was type `note_off`, the note gets deleted from the dictionary, meaning it will not change the respective pixel to white in the next iteration, but it will remain black.

For example, say that the note dictionary contains the active notes: 45, 50, 55. Current message type is `note_on`, `note=70` and `time=12` and current time value is 120 (which is the pointer that points to the current  $x$  value of the image). The image is filled like this:

- `image[96-45][120 to 120+12]` from 0.0 to 255.0
- `image[96-50][120 to 120+12]` from 0.0 to 255.0
- `image[96-55][120 to 120+12]` from 0.0 to 255.0
- `current_time=120+12`
- `note_dict=45, 50, 55, 70`

The values of note 70 are not yet set to white, but in the next iteration, the pixels will be swapped, as it is present in the note dictionary.

A second example presents the process with a type `note_off` 50 message (current `time=120`, `time=12` of message and same note dictionary):

- `image[96-45][120 to 120+12]` from 0.0 to 255.0
- `image[96-50][120 to 120+12]` from 0.0 to 255.0
- `image[96-55][120 to 120+12]` from 0.0 to 255.0
- `current_time=120+12`
- `note_dict=45, 55`

Pixels for note 50 are changed, but the same note is discarded from the dictionary, and in the next iterations its pixels will remain black.

In the resulting image after conversion, as shown in Figure 3.2, the rows determine the note value ( $y$  axis) and columns refer to a specific time tick ( $x$  axis). It is noted that the image's row value (pitch/note) is  $96 - pitch$  purely for presentation purposes, otherwise it would be inverted, but the results would be the same.

It is also noted that some files contain more than one tracks, which partition the whole song (low and high notes or left and right hand), meaning the file type is synchronous (datasets do not contain asynchronous files), so the images created for each track are then concatenated into one. After a whole song is converted to an image, it is then segmented into a series of smaller images with width value equal to 96, so each file produces  $track\_duration/96$  images.

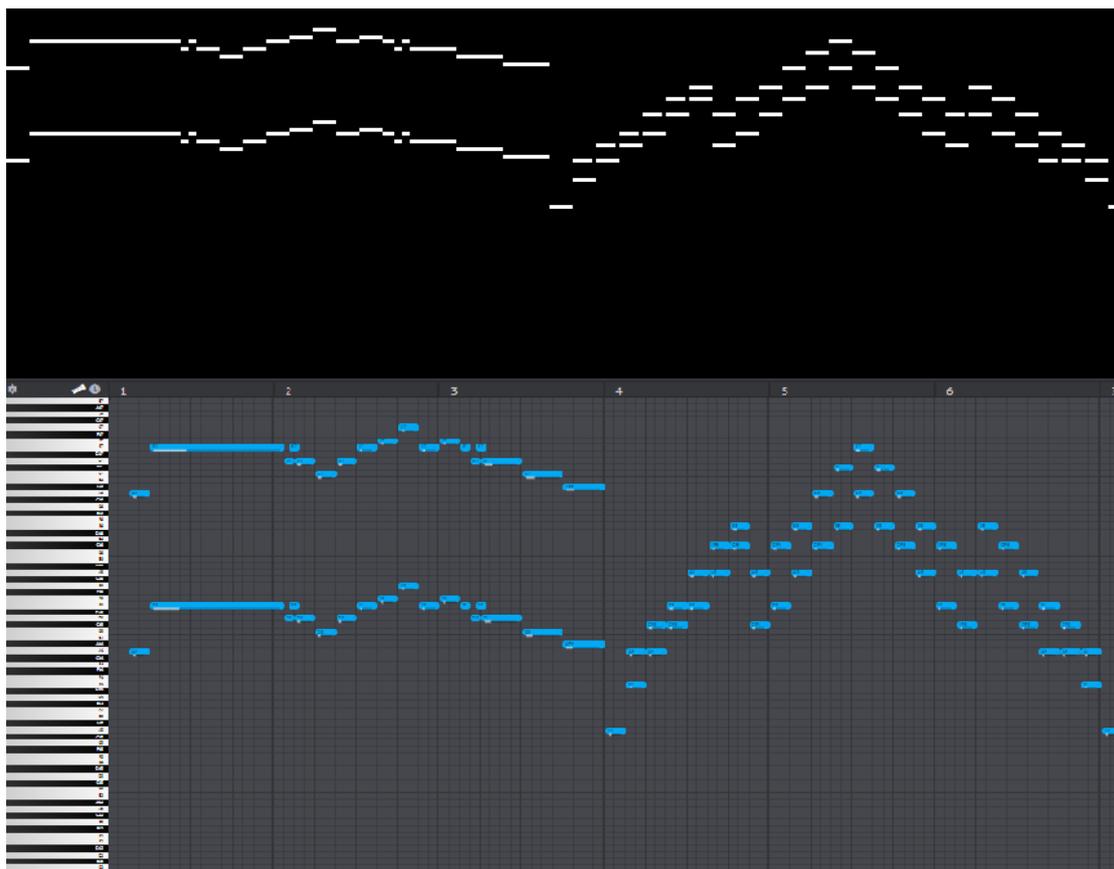


Figure 3.2: Example of MIDI to image encoding compared to the same MIDI opened by an external application

## 3.2 Training Models

In this section the two approaches for solving the music generation task are presented. Two models were used, as stated previously, a Variational Autoencoder (VAE) and a Generative Adversarial Network (GAN). Both models have powerful generative applications with different pros and cons. They are similar in their structure, as each one consists of two sub-networks, one dedicated to generating images and another assisting the generator network. Also they take advantage of convolutional layers which convert images to a more compact dense representation. After the training process is completed, the sub-network that translates data to images is used by passing random or specific input and it outputs images that are then converted to MIDI files. In the next sub-sections each approach, as well as the architectures used with their respective neural layers, are explained.

### 3.2.1 Variational Autoencoder (approach and architecture)

Variational Autoencoders are very well suited to generate random data that kind of resemble the initial dataset, but they work better at altering and exploring data at a specific non-random direction (Shafkat 2018).

#### 3.2.1.1 Standard Autoencoder

An Autoencoder consists of two connected networks, the encoder and the decoder. An encoder receives an input which is then translated to a small dense representation, which in turn is used by the decoder to convert it back to its original input (Figure 3.3).

The encoder is similar to Convolutional Neural Networks (CNNs). CNNs reduce a tensor of height-width-channel to a dense representation, for example a rank 1 tensor of size 1000. Then this representation is ready to get passed to the classifier. The only difference of an encoder is that it produces a representation of a much smaller size, the encoding, enough to get passed to the rest of the network and attain a desired output. Usually, the encoder is trained with the rest of the network's parts, optimized by back propagation, to attain encodings desired

for a certain task. In summary, the general idea is making the encoder generate encodings used for reconstructing the initial input.

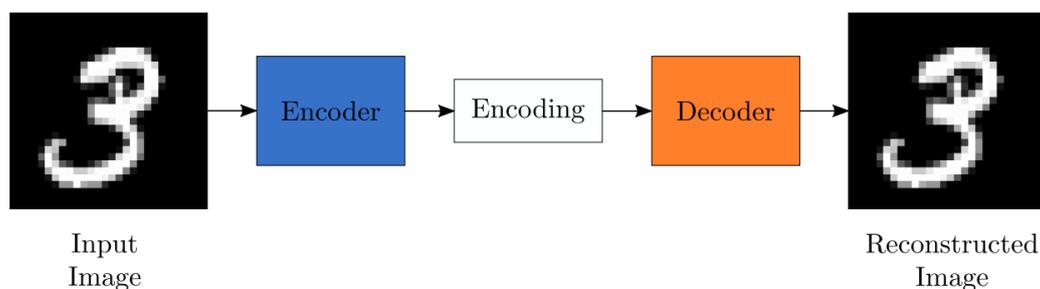


Figure 3.3: The Standard Autoencoder

The entire autoencoder network is trained as a whole and the loss function is either mean-squared error or cross-entropy between the output and the input, known as the reconstruction loss, which penalizes the network for creating outputs different from the input.

As the encoding (which is the output of the hidden layer of the encoder) is much smaller than the input, the encoder inevitably will discard information. The encoder learns to choose or preserve as much of the relevant information as possible in the limited encoding, and intelligently discards irrelevant parts. The decoder receives the encoding and is trained to properly reconstruct it into a full image. Together, the encoder and the decoder form an autoencoder.

### 3.2.1.2 Weakness of standard Autoencoders

Standard autoencoders learn to generate dense representations and then reconstruct their inputs fairly well. Nevertheless, besides from a few applications like denoising autoencoders, their usage is fairly limited.

Their main issue, for generation, is the latent space that they convert their inputs to and where the encoded vectors lie, may not be continuous and can contain gaps.

For example a visualized 2D latent space (Figure 3.4) of a trained network to the MNIST dataset reveals certain distinct clusters of each class (Shafkat 2018). This

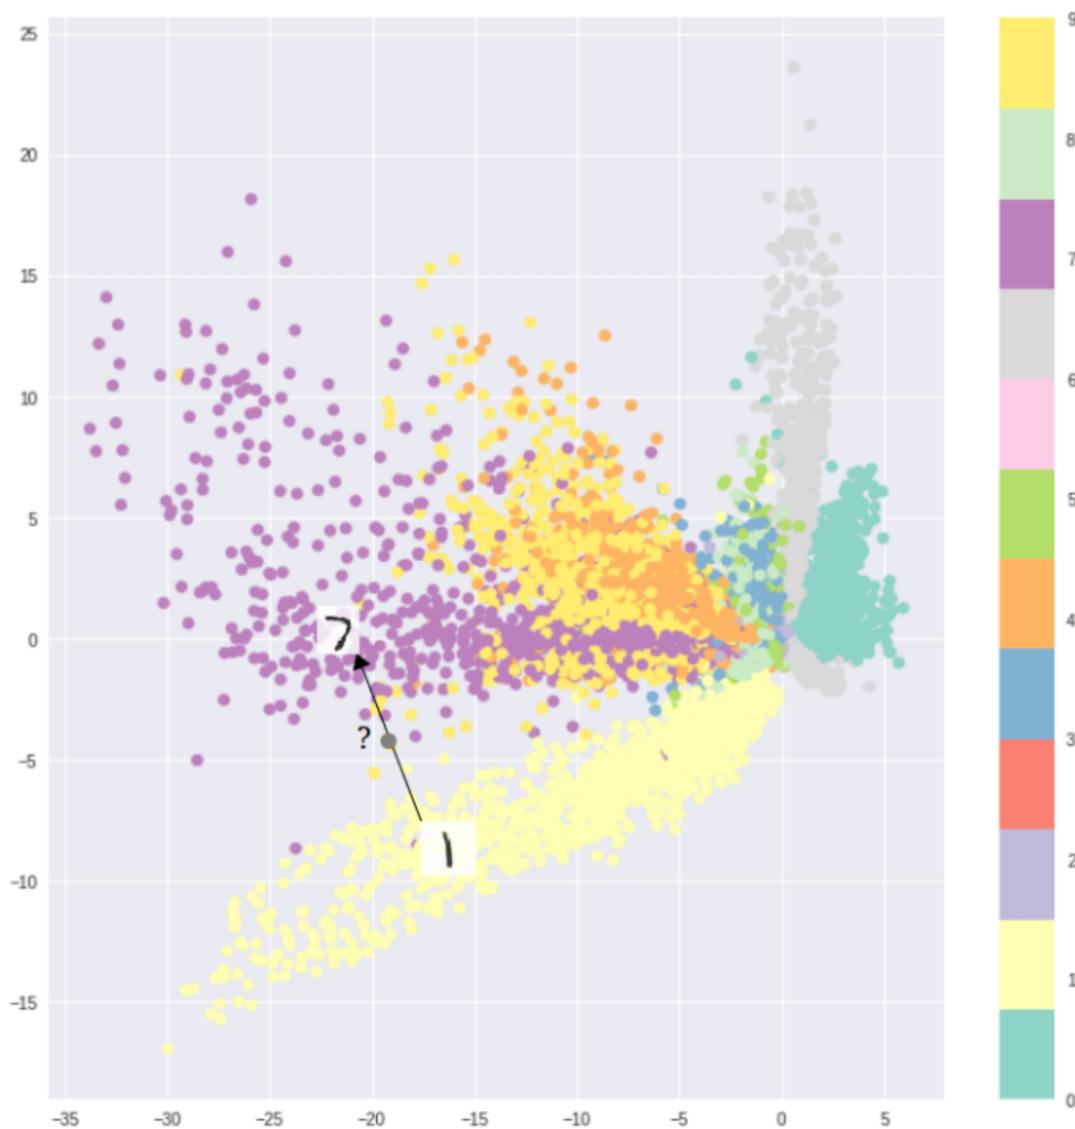


Figure 3.4: MNIST Latent space optimizing only on reconstruction loss

is fair, as distinct encodings for each image type or class makes it far easier for the decoder to decode them. This is fine for replication of the same images, yet it is not desired for generation purposes. The goal is to randomly sample from the latent space, or generate variations on an input image, from a continuous latent space.

If the space has discontinuities (eg. gaps between clusters), and the decoder sam-

ples encodings from those gaps, it will simply generate an unrealistic output, because it has no idea how to deal with that region of the latent space. This happens because during training, it never saw encoded vectors coming from that region of the latent space.

### 3.2.1.3 Variational Autoencoders

Variational Autoencoders (VAEs) have one fundamentally unique property that distinguishes them from standard autoencoders, and it is this property that makes them so useful for generative modeling: they are designed to produce latent spaces that are continuous, which allow for easy random sampling and interpolation.

This is achieved by outputting two vectors of size  $n$  instead of one: a vector of means,  $\mu$ , and another vector of standard deviations,  $\sigma$ . They represent the parameters of a vector of random variables of length  $n$ , with the  $i$ -th element of  $\mu$  and  $\sigma$  being the mean and standard deviation of the  $i$ -th random variable,  $X_i$ , from which the sampled encoding is obtained, which in turn is passed onward to the decoder.

This stochastic generation means, that even for the same input, while the mean and standard deviations remain the same, the actual encoding will somewhat vary on every single pass simply due to sampling.

Intuitively, the mean dictates the center of the position of an encoding, while the standard deviation controls the area of how much a sample can deviate from the mean or center. As encodings are generated randomly by the distribution, given the two vectors, the decoder learns that all nearby points close to the center refer to a certain class, and not just a single one in the latent space. This very property offers the desired exploration aspect, as samples will vary, while still being part of the distribution, and in turn fill the discontinuous gaps. The latent space at which the decoder is trained upon is denser with no discontinuities or gaps.

As long as the network offers encodings whose latent space is denser, it ideally provides the possibility of interpolation between classes, which the standard autoencoders are usually unable to do. However, since there are no limits on what values vectors  $\mu$  and  $\sigma$  can take on, the encoder can learn to generate very differ-

ent  $\mu$  for different classes, clustering them apart, and minimize  $\sigma$ , making sure the encodings themselves do not vary much for the same sample (less uncertainty for the decoder). This allows the decoder to efficiently reconstruct the training data. In this thesis however, interpolation between classes is irrelevant, as our approach utilizes a single class.

What is ideally desired are encodings, that are as close as possible to each other while maintaining distinction, allowing smooth interpolation, and enabling the construction of new samples.

This can be enforced by utilizing the Kullback–Leibler divergence in the loss function. The KL divergence between two probability distributions simply measures how much they diverge from each other, therefore by minimizing the KL divergence, the probability distribution parameters ( $\mu$  and  $\sigma$ ) are optimized to closely resemble that of the given distribution.

$$\sum_{i=1}^n \sigma_i^2 + \mu_i^2 - \log(\sigma_i^2) - 1 \quad (3.4)$$

The KL loss (for VAEs) is computed by the sum of all the KL divergences between  $X_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$  in  $X$  and the standard normal distribution. This is minimized, when  $\mu_i = 0$ ,  $\sigma_i = 1$ .

Intuitively, this forces the encoder to distribute all encodings from different classes to the center of the latent space. The encoder is also penalized, if it tries to cluster apart the encodings.

Using purely KL loss, a latent space is attained (Figure 3.5) whose encodings appear densely randomly near the center regardless of similarities among nearby encodings (Shafkat 2018). Then, the decoder is unable to distinguish anything meaningful.

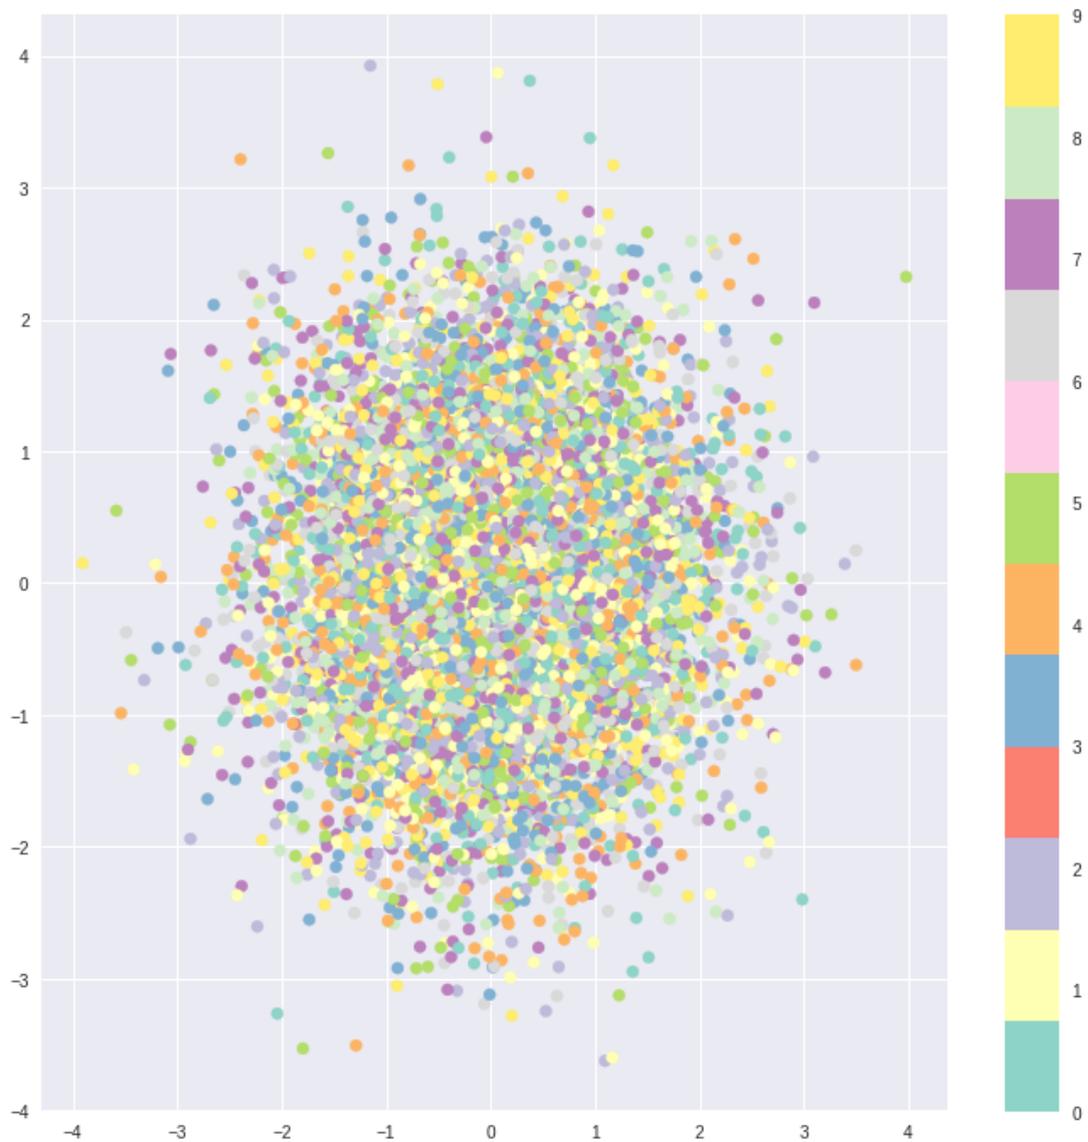


Figure 3.5: MNIST Latent space optimizing only on KL divergence loss

Optimizing the decoder as well (reconstruction loss along with KL divergence), results in the generation of a latent space (Figure 3.6), which maintains the similarity of nearby encodings on the local scale via clustering, and it is fairly dense packed near the latent space origin (Shafkat 2018).

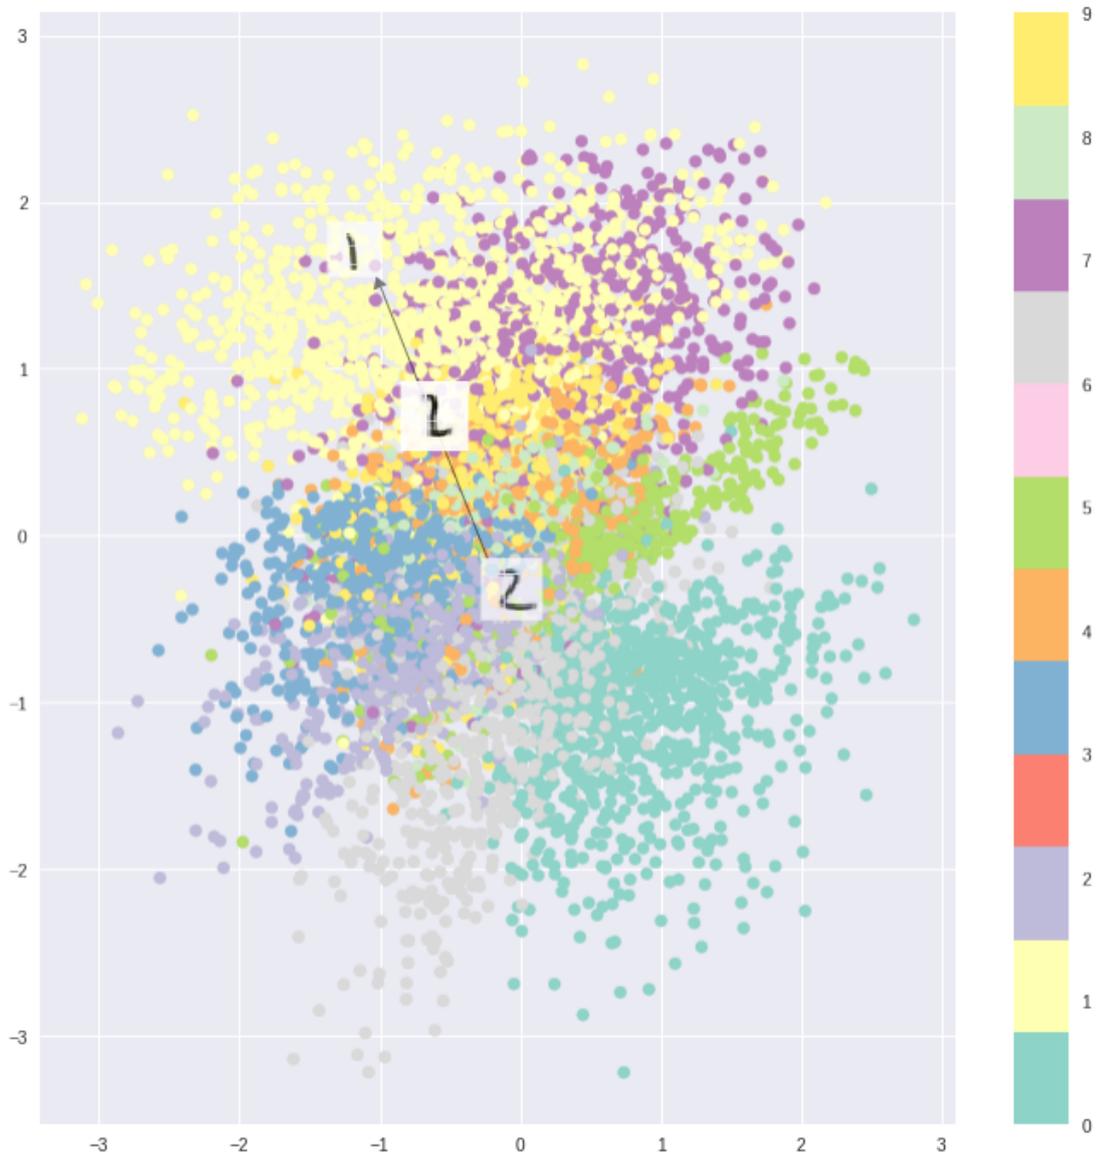


Figure 3.6: MNIST Latent space optimizing on both reconstruction and KL divergence loss

Intuitively, the result of optimizing both reconstruction and KL loss, is that the

decoder can successfully decode encoded vectors from a certain distribution, as well as interpolate and mix features due to the lack of gaps in the latent space.

### 3.2.1.4 VAE architecture

In this section, the VAE architecture used for generation of images that contain musical sequenses is presented and explained. The layers of both networks (encoder and decoder) are also given in detail. This machine learning task is written in Python, which provides exceptional libraries and wrappers, like Tensorflow and Keras.

#### Encoder Model Architecture

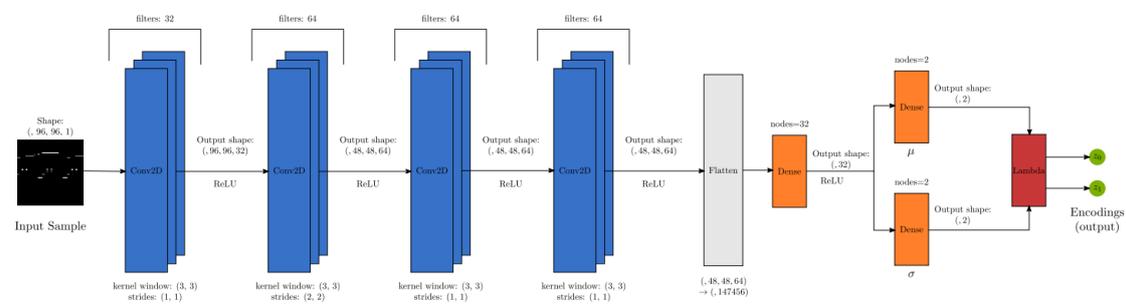


Figure 3.7: Encoder Architecutre

The encoder architecture is presented in Figure 3.7. This particular model receives an input image and translates it to a dense representation. The image data it receives is 96 by 96 by 1 pixel dimensions (height-width-channels), and the latent dimensions are 2, that represent the mean and the standard deviation of the latent distribution.

All the neural layers of the encoder are presented below:

- Conv2D: a regular convolution layer that executes a convolution through every pixel of the image, the number of filters is equal to the number of output channels (input has 32, rest of the layers have 64 filters), kernel size is  $3 \times 3$  which is the window of the convolution, and the strides refer to the reduction of the output image (e.g. default of strides (1, 1) output image of

same size, stride tuple set to (2, 2) outputs image of half height and width). Output of the convolutions are  $96 \times 96 \times 1 \rightarrow 96 \times 96 \times 32 \rightarrow 48 \times 48 \times 64$  followed by two more convolutions of the same size ( $48 \times 48 \times 64$ )

- relu: refers to the activation function of the layers, allows positives to pass through and filters out negatives to 0.
- Flatten: matrix conversion to a vector, in this case flattening the image of shape  $48 \times 48 \times 64$  to a vector of  $48 \cdot 48 \cdot 64 = 147456$  items or network nodes. This is usually done as a preparation for the dense layer.
- Dense: a fully connected layer that has input times output connections, each connection has a weight and bias, that are multiplied and added to the input respectfully.
- Lambda: this layer exists so that arbitrary expressions can be used as a network layer, in this case the computed expression contains  $\mu$  and  $\sigma$  (mean and sigma) from the previous layers.

Specifically this last sub-network executes convolutions, before feeding data through dense layers. It provides two outputs, mean  $\mu$  and standard deviation  $\sigma$  of encoded inputs. We use these to sample random variables in the latent space to which inputs are mapped. Then we define a sampling function to sample from this distribution. The sample is "reparameterized" based on the process defined by Gunderson and Huang into the shape of  $\mu + \sigma \cdot \epsilon$  where  $\epsilon$  is the sample from the distribution. This is done to allow for accurate gradient descent estimation. The Lambda layer contains this exact expression added for gradient descent calculations using  $\mu$  and  $\sigma$ .

The rectified linear unit (ReLU) (3.5) (Figure 3.8) is used as the main activation function. Activation functions, such as ReLU, are used to address the vanishing gradient problem in deep convolutional neural networks and promote sparse activations (e.g. lots of zero values). It is responsible for transforming the summed weighted input from the node into the activation of the node or output for that input. It essentially serves as filtering. If the input is positive, it is passed as is, and if not, the output is 0.

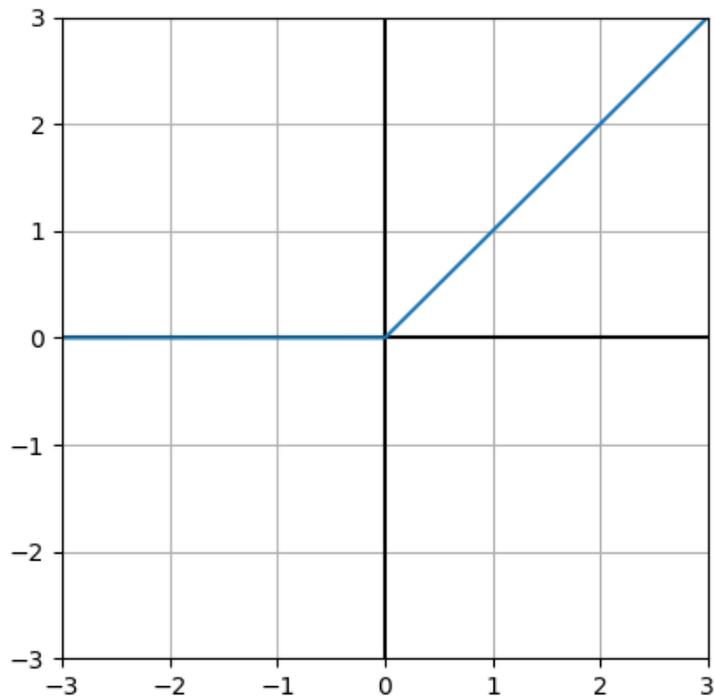


Figure 3.8: ReLU

ReLU function:

$$f(y) = \begin{cases} 0 & \text{if } y < 0 \\ y & \text{if } y \geq 0 \end{cases} \quad (3.5)$$

### Decoder Model Architecture

The decoder architecture is presented in Figure 3.9. This model translates the output of the encoder, which is a sampled latent vector, to an image of shape  $96 \times 96 \times 1$ .

- Dense: a fully connected layer, contains input times output connections, each connection has a weight and bias which are multiplied and added to the input respectfully. This is the input layer of the decoder network.
- Reshape: reshapes the input to an output with desired dimensions (error

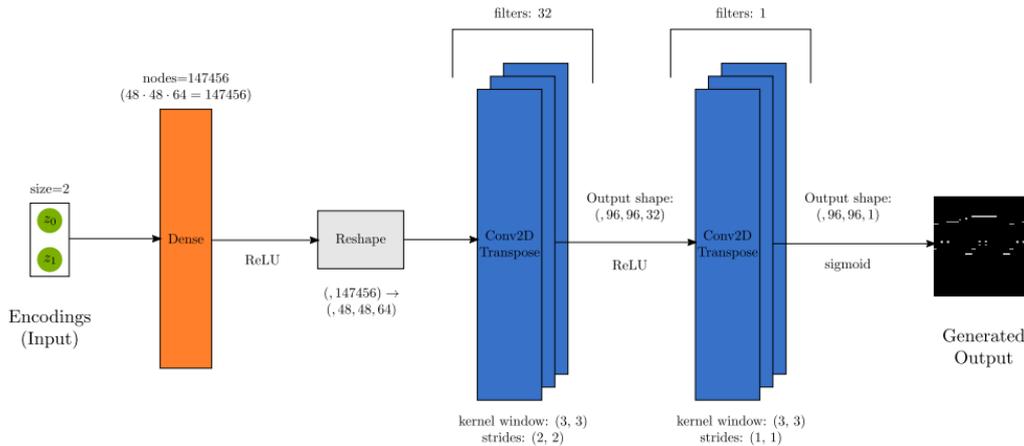


Figure 3.9: Decoder Architecture

will be raised, if the total number of items between input and output are not equal). In this case the input from the dense layer is 147456 nodes (or items) and the reshaped output is a 48 by 48 by 64 matrix, which  $48 \cdot 48 \cdot 64 = 147456$ .

- Conv2DTranspose: a transpose convolution or deconvolution layer, used for upsampling the input determined by the strides of the convolution, which in this case strides are set to a (2, 2) tuple which essentially doubles the  $x$  and  $y$  axes of the input (width and height).
- relu: an activation function, which simply allows positives to pass through and filters out negatives to 0.
- sigmoid: an activation function, also known as logistic function, which is essentially a smooth step function around 0.5.

The decoder is the component that will be used after training for generating images. The main goal is to essentially train the combined model (encoder plus decoder) in order to achieve a decoder that can sample data from the target distribution, so the image generation is as close to the desired results as possible.

The sigmoid function (3.6) (Figure 3.10) (Han and Moraga 1995), or logistic function, is applied to the output of the decoder network. This function adds non-linearity in machine learning modeling, by mapping real values between 0 and 1.

It is usually added to networks that perform binary classification tasks, and in this case each pixel is mapped between 0 and 1, which determines whether each pixel (or note) will be "played" (in musical terms) or not.

$$f(y) = \frac{1}{1 + e^{-y}} \quad (3.6)$$

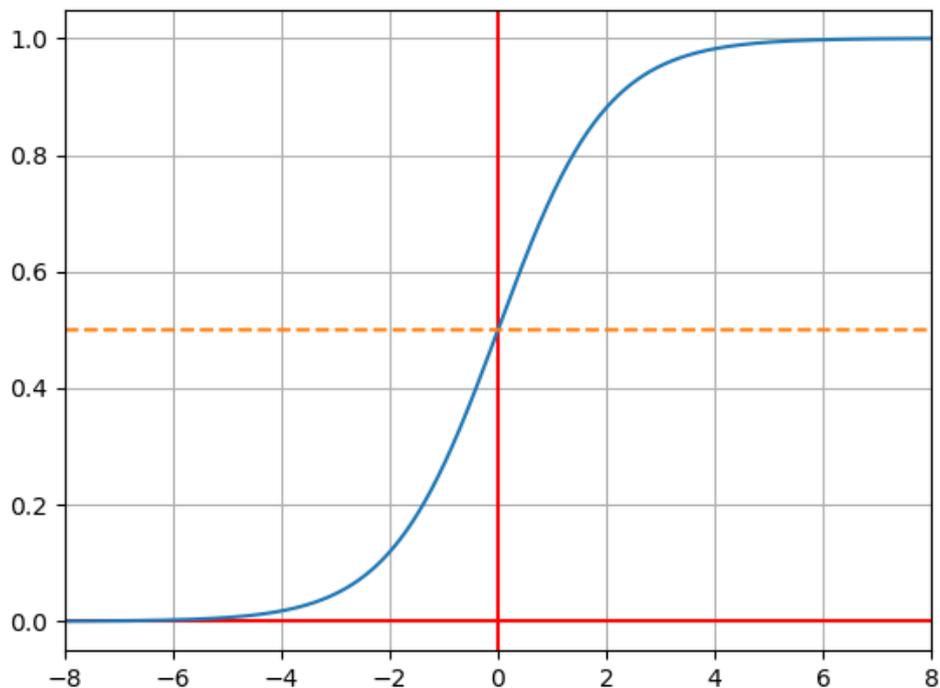


Figure 3.10: Sigmoid function

### The combined model

The two models, the encoder and decoder, are then combined with the help of a custom layer class, which is essential, if we want to define custom loss. The complete VAE is trained using two loss functions, reconstruction loss and KL divergence.

Reconstruction loss refers to the binary cross-entropy loss between the input image

(from the initial dataset) and the decoded image (provided by the decoder). Binary cross-entropy compares each of the predicted probabilities, in this case for each pixel, to the actual output, which can be either 0 or 1 (black or white pixel, note off or note on). It then calculates a score that penalizes the probabilities, based on the distance from the expected value (how close or far).

The loss-function of binary cross-entropy:

$$H_p(q) = -\frac{1}{N} \sum_{i=1}^N y_i \cdot \log(p(y_i)) + (1 - y_i) \cdot \log(1 - p(y_i)) \quad (3.7)$$

Kullback-Leibler (KL) divergence is used to attain sampled encodings that are as close as possible to each other, while maintaining distinction, allowing smooth interpolation, and enabling the construction of new samples. The KL loss is computed as shown in (3.4) and then is added to the reconstruction loss.

## Compilations and Optimizers

Compilations and optimizers of neural network models often refer to the adjustment of trainable parameters, the weights. They are essential processes that often determine the stability of the network and offer a smooth learning curve.

The combined model, encoder along with decoder, are compiled with the custom loss, reconstruction and KL loss. The optimizer of choice is Adam.

Adam is the optimization algorithm used instead of a classic stochastic gradient descent for weight updating. Its name is derived from adaptive moment estimation. Classic stochastic gradient descent has a single term named alpha (learning rate) for updating all the weights and remains unchanged during the training process. Adam utilizes a learning rate as well, which is separately adapted for each network weight as learning unfolds. Adam is produced by combining the advantages of two other extensions of stochastic gradient descent, specifically:

- **Adaptive Gradient Algorithm** (AdaGrad) that maintains a per-parameter learning rate that best performs on problems with sparse gradients (e.g. NLP and computer vision).

- **Root Mean Square Propagation** (RMSProp) that also maintains per-parameter learning rates that are adapted based on the average of recent magnitudes of the gradients for the weight (e.g. how quickly it is changing), meaning the algorithm does well on online and non-stationary problems (e.g. noisy).

Adam utilizes the advantages of both AdaGrad and RMSProp. Instead of adapting the parameter learning rates based on the average first moment (the mean) as in RMSProp, Adam also makes use of the average of the second moments of the gradients (the uncentered variance). Specifically, it calculates an exponential moving average of the gradient and the squared gradient, and parameters  $\beta_1$  and  $\beta_2$  control the decay rates of these moving averages. Adam is preferred among other optimizers in similar tasks, because it is effective. Empirical results demonstrate the superiority of Adam in terms of training cost among other optimizers, as shown in Figure 3.11 (Kingma and Ba 2015).

Adam configuration parameters with their respective empirical good default examples:

- **alpha**. This is known as learning rate or step size. The proportion that weights are updated (e.g. empirical good default value is 0.001). Larger values result in faster initial learning before the rate is updated and smaller values in slow learning.
- **beta1**. The exponential decay rate for the first-moment estimates (e.g. 0.9).
- **beta2**. The exponential decay rate for the second-moment estimates (e.g. 0.999). This value should be set close to 1.0 on problems with sparse gradient (e.g. NLP and computer vision).
- **epsilon**. A very small number to prevent any division by zero in the implementation (e.g. 10E-8).

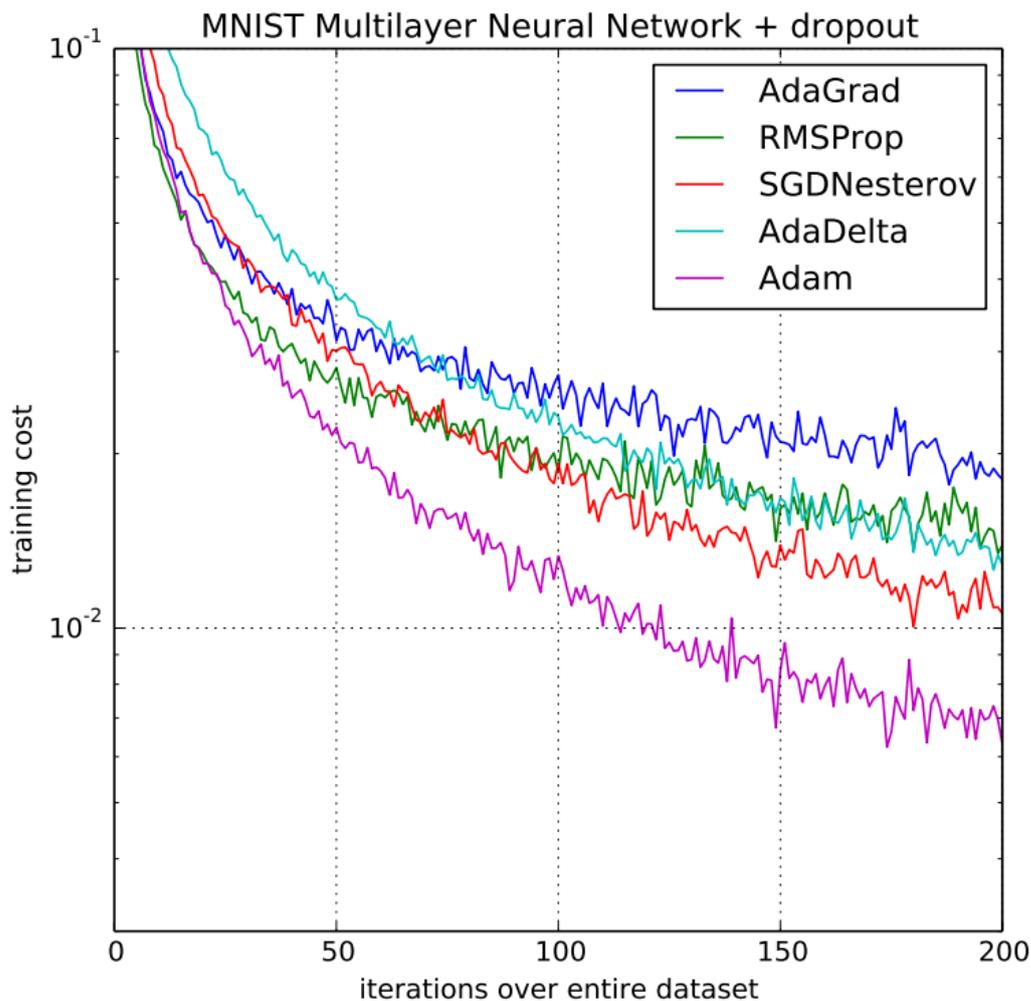


Figure 3.11: Adam compared to other optimizers

### 3.2.2 Deep Convolutional GANs (approach and architecture)

Generative Adversarial Networks (GANs) are broadly used for generative purposes. Generative modeling is an unsupervised learning task in machine learning that involves automatically identifying and learning the similarities or patterns in input data, in such a way that the model can be used to generate or output new examples that plausibly could have been drawn from the original dataset (Brownlee 2019a). They are desired and sought after due to their interesting design of one of their

sub-networks, the generator.

The generator is very interesting due to its nature that translates random noise into desired data. This is accomplished via the second sub-network, the discriminator, which receives the generator's output and decides whether this data is true or false. Due to parallel training, those two networks compete, in order to improve themselves. The generator's purpose is to fool the discriminator and the latter is trained to distinguish real from fake images, that originate from the dataset or the generator. This results in an adversarial zero-sum game, and the ideal situation in theory would be the discriminator to be fooled about half the time, meaning that the generator attains plausible examples. In practice, the expected accuracy of a discriminator from a well-trained GAN hovers around 80%, and this is explained in detail in the Results and Discussion sections.

### 3.2.2.1 Discriminative vs. Generative Modeling

Usually, in supervised learning, models are constructed for predictive purposes, to classify an example of input variables to a specific class. This predictive modeling task is called classification. It is also known as discriminative modeling. From the training data, the goal is to attain a discriminant function  $f(x)$  that maps each  $x$  onto a class label, therefore combining the inference and decision stages into a single learning problem. In summary such a model's goal is to discriminate or classify sets of input variables across classes. The factor of choice or decision, as to what class a given example belongs to, is presented.

On the other hand, unsupervised models that identify the patterns of input variables are able to be used for generative purposes based on the input distribution. Those models are referred to as generative models.

For example, a single input variable may have a known data distribution, such as a Gaussian distribution, or bell shape. A generative model may be able to identify patterns and summarize this distribution, as well as generate new data points that could have been part of the initial dataset. The ideal scenario would be a generative model that produces data that is indistinguishable from real samples.

### 3.2.2.2 Standard GANs

The GAN model architecture consists of two sub-models, a generator model capable of producing new examples and a discriminator model for identifying whether input data originates from the domain (is real), or is fabricated by the generator model (is fake). GANs are based on a game theoretic scenario in which the generator competes against an adversary.

### 3.2.2.3 Generator Model

The generator model receives as input a fixed-length random vector and generates a sample in the domain (Figure 3.12).

The vector is produced randomly by a Gaussian distribution and is used to seed the generative process. After the training process, points in this multidimensional vector space correlate with points in the problem domain, forming a dense representation of the data distribution.

This vector space is also known as latent space, because it contains latent variables. Those latent or hidden variables are important in neural network modeling, they are understood by the specific model, but are not directly observable by humans. Furthermore, latent variables are often referred to as a projection or compression of a data distribution. In this specific case of GANs, the generator model learns to apply meaning to points in a chosen latent space, such as when new points are drawn from it, to be provided to the generator model as input and to be used to generate new and different output examples. When the process of training is deemed complete, the generator is saved and used to generate new samples, hence the point of this work.

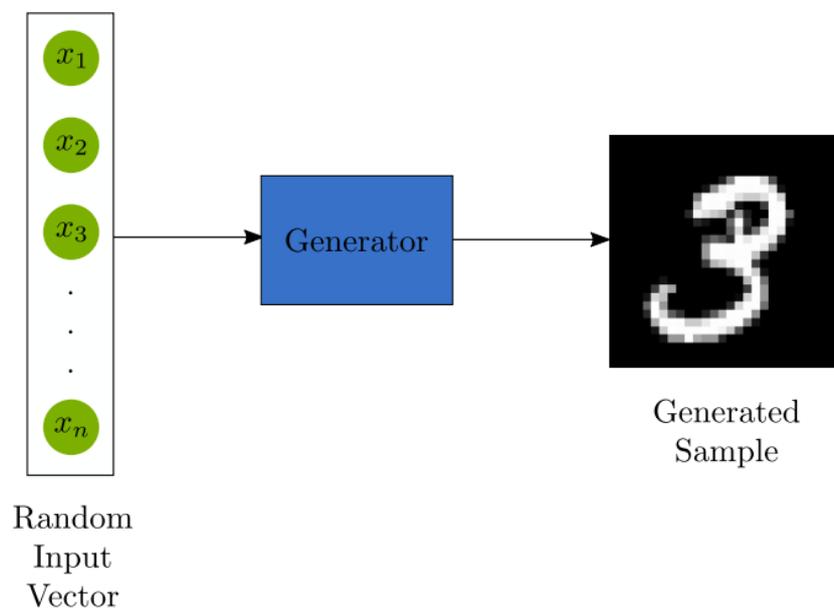


Figure 3.12: The Generator Model

### 3.2.2.4 Discriminator Model

The discriminator model receives as input an example from the dataset (real) or a generated sample (fake) and predicts a binary class label of real or fake (Figure 3.13). The real examples are drawn from the original dataset and the fake ones are fabricated by the generator model. The discriminator is a typical classification model, which in this case classifies images to be true or fake, meaning they originate from the initial dataset or they are produced by the generator. After the training of the GAN is completed, the discriminator is discarded, as the model of interest is the generator purely for generative purposes. Regardless, it could still be used for other tasks as some of the feature extraction layers can be used in transfer learning applications.

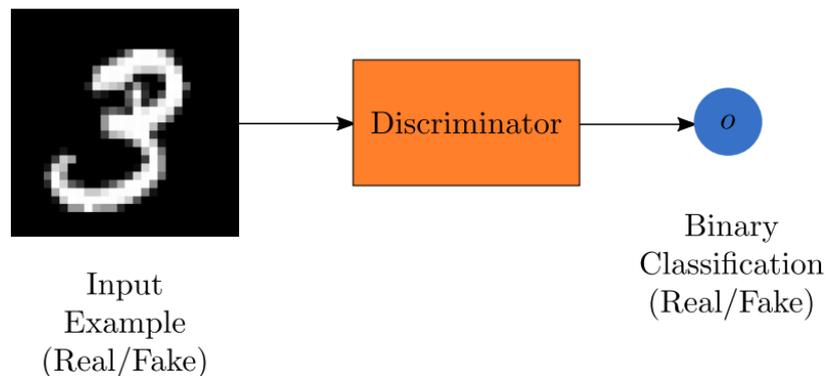


Figure 3.13: The Discriminator Model

### 3.2.2.5 GAN - A two player competition

As mentioned in a previous section, generative modeling is an unsupervised learning problem, but the training of the generative model in GANs is framed as a supervised learning problem.

The two models, generator and discriminator, are trained together (Figure 3.14). The generator generates samples trying to imitate the dataset, and these, along with the real examples, are received by the discriminator and classified as real or fake.

Then the discriminator is updated to become better at discriminating real and fake samples in the next round, and the generator is updated based on how well the generated samples imitated the ones from the dataset and fooled the discriminator.

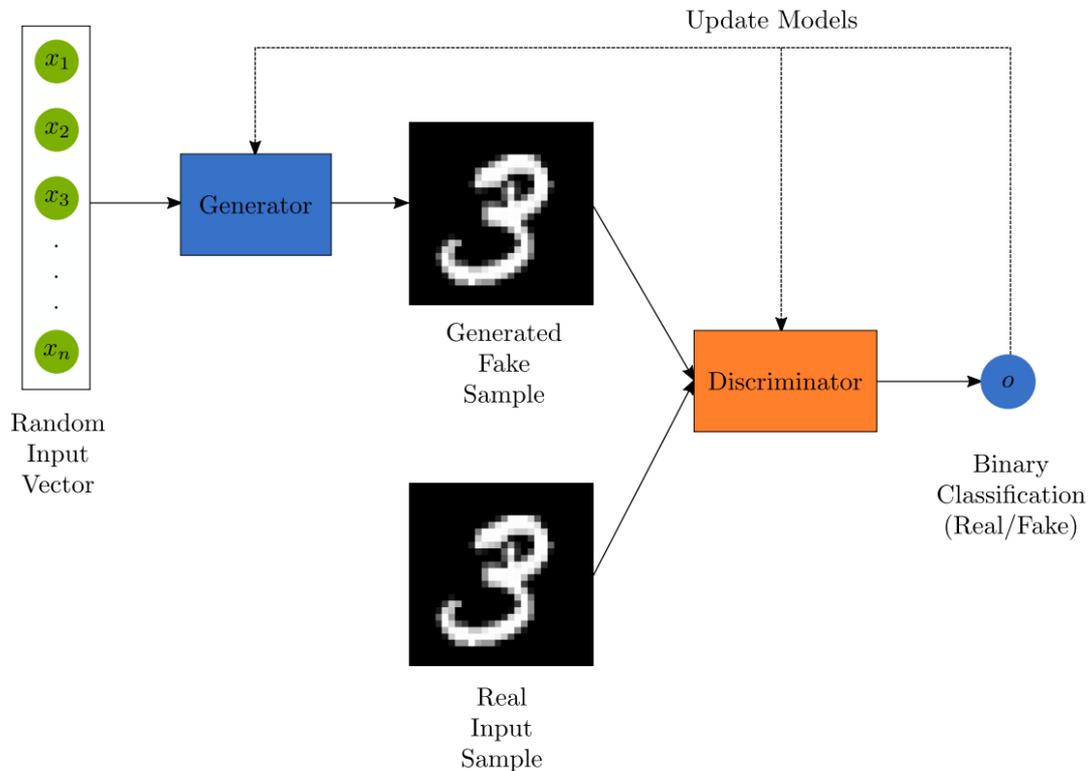


Figure 3.14: The Generative Adversarial Network Architecture

### GANs and Convolutional Neural Networks

In most applications GANs utilize the use of convolutional neural networks or CNNs. Both the discriminator and generator consist of such layers. CNNs are very useful in computer vision tasks, as they are able to compress an image of 3 dimensions (width, height and channels, eg.  $64 \times 64 \times 3$ ) to a vector. This vector is a dense representation of the input that extracts its main features. They are used by the generator to translate the latent space to an image sample, and for the discriminator to identify the validity of the image, whether it is real (from the dataset) or fake (produced by the generator).

Modeling image data refers to that exact process of creating a latent space, the

input of the generator, that provides a dense or compressed representation of the set of images from the dataset used to train the model. The generator learns from training to receive random latent points from a distribution and generate new images.

CNNs are also applied to this learning task, as convolutions and transpose convolutions are essential. The discriminator model utilizes convolutions to break down an image input to a compressed representation and classify whether it's real or fake, and the generator model uses transpose convolutions or deconvolutions which is the opposite direction of a convolution, transformation of a supposed compressed representation (random points from a distribution) to an image, that still maintains connectivity and patterns as a regular convolution.

### 3.2.2.6 GAN Architecture for Music Generation

In this section, the architecture used for generation of images that contain musical sequences is presented and explained. The layers of both networks (generator and discriminator) are also given in detail. This machine learning task is written in Python, which provides exceptional libraries and wrappers, like Tensorflow and Keras.

#### Generator Model Architecture

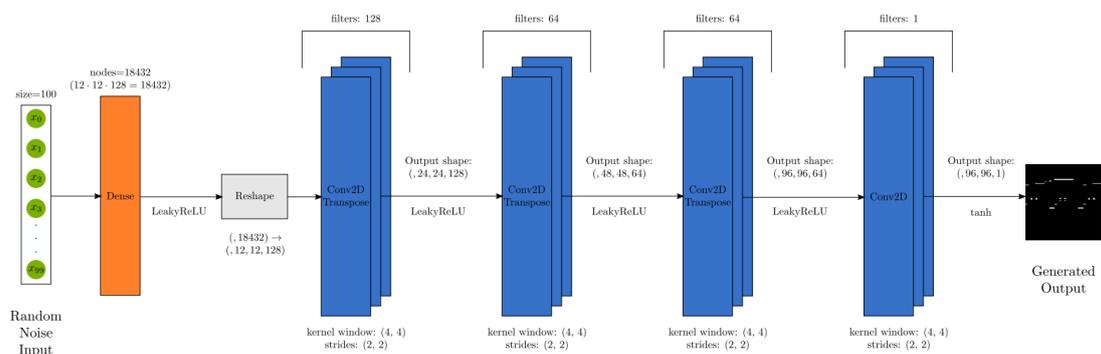


Figure 3.15: The Generator Architecture

The generator architecture is presented in Figure 3.15. The model receives an input of latent points with dimension 1 by 100. This noisy information is then

passed through a fully connected or dense layer, that has  $24 \cdot 24 \cdot 32 = 18432$  nodes. This exact number of nodes is essential, as the desired output of the model will be  $96 \times 96 \times 1$  dimensions. Then, a number of upsamples take place, as the use of deconvolutional layers take place for the transformation of the output of the dense layer to an image.

All the neural layers of the Generator are presented below:

- Dense: a fully connected layer, which contains input times output connections, and each connection has a weight and bias, that is multiplied and added to the input respectfully. This is used to convert the latent vector (noisy input) to an item that is then ready to be upsampled to an image with desired dimensions.
- Reshape: simply reshapes the input to an output with desired dimensions (error will be raised, if the total number of items between input and output are not equal). For example, in this case the output of the dense layer is 18432 nodes, which is then reshaped to an output of  $24 \times 24 \times 32$  (width-height-channels), that is  $24 \cdot 24 \cdot 32 = 18432$ .
- Conv2DTranspose: a transpose convolution or deconvolution layer, used for upsampling the input determined by the strides of the convolution, which in this case strides are set to a (2, 2) tuple, which essentially doubles the  $x$  and  $y$  axis of the input. Also the kernel window size is  $4 \times 4$  at which the convolution is calculated. Filters sizes vary.
- LeakyReLU: serves as the main activation function, a *leaky* rectified linear unit
- Conv2D: a regular convolution layer that shifts through every pixel and executes a convolution, height and width are unchanged as strides default to (1, 1), and serves as the output of the generator model
- tanh: activation function of the final Conv2D layer

The output of the reshaped dense layer is passed through two upsampling layers (Conv2DTranspose). While training, generator models with smaller number

of dense nodes and more upsampling layers produced images with all the pixels valued 0 (black images), which makes sense, since the more dense nodes, the more information the model is able to save and imitate the dataset. Nevertheless the training process becomes slower, as the total trainable parameters increase, as well as the load is heavier for the machine used for training.

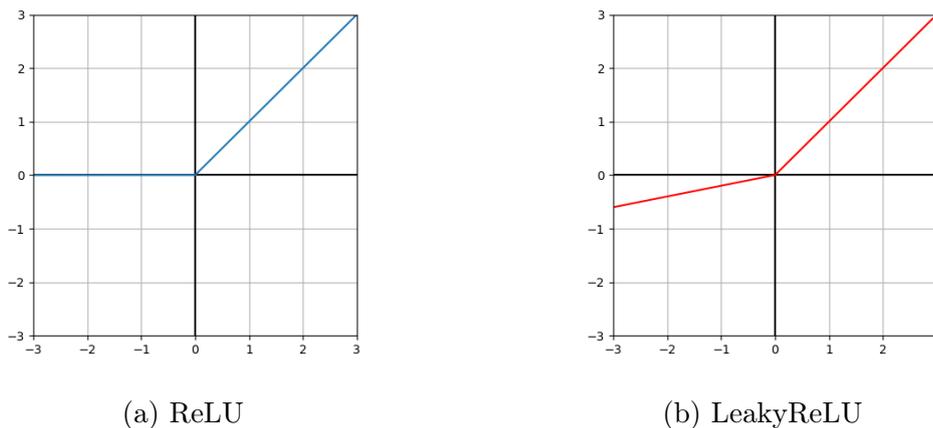


Figure 3.16: Rectified Linear Units (Basic and Leaky)

The rectified linear unit (ReLU) (3.5) (Figure 3.8) is used as the main activation function. Basic ReLU is used in VAE iteration referenced in the previous section, whilst the GAN utilizes LeakyReLU. LeakyReLU (Figure 3.8) (Xu et al. 2015) serves the same purpose as basic ReLU, but instead of converting every negative value to 0, it allows some negative values based on the slope variable that is set by the user, hence the name "leaky". While in the generator, regular ReLU performed just as well as LeakyReLU, in the discriminator the deviation of the results was greater, so the leaky rectified linear unit was kept as default for both of the models with a slope value of  $a = 0.2$  (example shown in Figure 3.16).

LeakyReLU function:

$$f(y) = \begin{cases} a * y & \text{if } y < 0 \\ y & \text{if } y \geq 0 \end{cases} \quad (3.8)$$

The deconvolutions in the generator model were preferred compared to other up-sampling layers (Upsampling2D). Strided convolutions allow the generator to learn

its own spatial upsampling. In this case, the generator upsamples the reshaped dense representation twice, resulting in the desired dimensions of the images, paired with LeakyReLU. Also the filters were set to 64, which were enough to sustain features, and the kernel size to (4, 4) (which is the convolution window). Furthermore, the input upsamples from  $24 \times 24 \times 32$  to  $48 \times 48 \times 64$  and then to  $96 \times 96 \times 64$  to attain the desired height and width. The last layer, or the output layer of the model, is a regular convolution that converts the filters or channels to 1 (since the model is trained upon and generates images in grayscale) with a training function of tanh (output is normalized to a range of  $[-1, 1]$ ), a kernel window (7, 7) and with no strides, since the dimensions of the image match the desired result. Most GANs utilize a regular convolution layer, instead of a fully connected dense layer, and by practice this output layer massively improved the results. Other parameters that refer to bias and kernel constraints or initial values were left to default.

It is noted that the generator model is compiled with an optimizer, when combined with the discriminator model.

### Discriminator Model Architecture

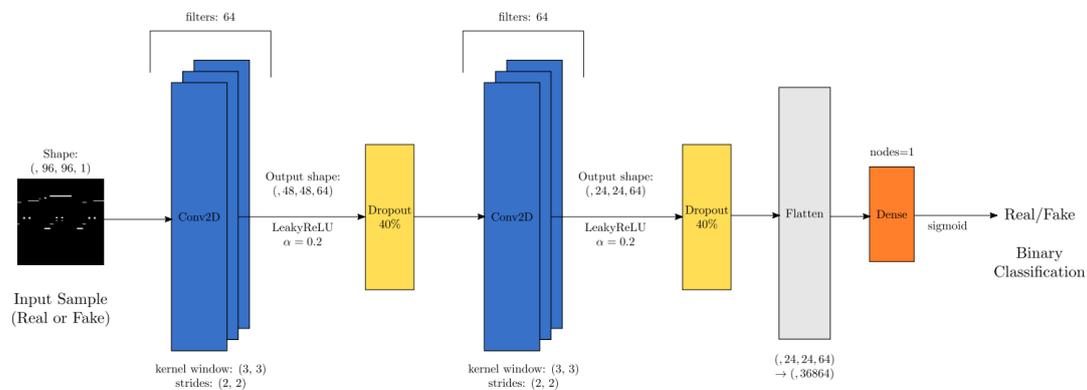


Figure 3.17: The Discriminator Architecture

The architecture of the discriminator, as shown in Figure 3.17, resembles a typical binary classification model using CNNs. As stated before, the process is to feed the discriminator with images of shape  $96 \times 96 \times 1$ , either sampled from the dataset

or generated by the generator model, and determine whether they are real or fake.

The neural layers used by the discriminator model are presented below:

- Conv2D: regular convolution layer, with 64 filters, kernel window size is  $3 \times 3$  at which the convolution is calculated, and the strides are set to  $(2, 2)$  as they determine the size of the output image (strides of size 2 reduce the image in half) eg. output of the convolutions are  $96 \times 96 \times 1 \rightarrow 48 \times 48 \times 64$  and then  $48 \times 48 \times 64 \rightarrow 24 \times 24 \times 64$ .
- LeakyReLU: activation function that lets all positive values to pass through and allows some negatives (regular ReLU filters all negative values to 0) based on the value of the slope set to  $\alpha=0.2$ .
- Dropout: randomly sets input values to 0 based on the frequency or rate set by the user, in this case, rate is 0.4 which means 4 out of 10 neurons will be terminated randomly. This helps with the overfitting problem of neural network classification and it is noted that dropout only takes place during training. Also inputs not set to 0 are scaled up by  $1/(1 - \text{rate})$ , so that the sum over all inputs is unchanged.
- Flatten: converts a matrix to a vector, in this case flattening the image of shape  $24 \times 24 \times 64$  to a vector of  $24 \cdot 24 \cdot 64 = 36864$  items. This is usually done for preparation to the dense output layer.
- Dense: a fully connected layer that has input times output connections, in this case  $36864 \cdot 1 = 36864$  connections, each one weighted and biased calculating the outcome, which is a binary classification. The activation function used is sigmoid, which is the most common and best performing choice for binary classification tasks.

The general idea of the discriminator is to catch the generator and identify fake images against real ones from the dataset. This model uses two downsampling convolutional layers, each with the same number of filters, strides and kernel windows, to break down its input, and a dense layer as the output for classification.

The activation function used, sigmoid, also called the logistic function, for binary

classification problems is like a step function around 0.5, but smoother, allowing values between 0 and 1 to pass through. Sigmoid is also presented in the VAE iteration (3.6) (Figure 3.10).

## Compilation and Optimizers

Compilations and optimizers of neural network models often refer to the adjustment of trainable parameters, the weights. They are essential processes that often determine the stability of the network and offer a smooth learning curve.

While the compilation of the generator takes place after the two models are combined, the discriminator is compiled separately. More specifically the discriminator is compiled with binary cross-entropy as its loss, its performance is observed by the accuracy metric and utilizes an Adam optimizer.

Binary cross-entropy compares each of the predicted probabilities to actual class output, which can be either 0 or 1. It then calculates the score that penalizes the probabilities based on the distance from the expected value. That means how close or far are from the actual value.

The loss-function of binary cross-entropy:

$$H_p(q) = -\frac{1}{N} \sum_{i=1}^N y_i \cdot \log(p(y_i)) + (1 - y_i) \cdot \log(1 - p(y_i)) \quad (3.9)$$

Where  $y$  is the label (1 for real images and 0 for fake images) and  $p(y)$  the predicted probability of the input being real for all  $N$  samples. This means that for each real image originating from the dataset ( $y = 1$ ) it adds  $\log(p(y))$  to the loss, which is the log probability of it being real, and conversely it adds  $\log(1 - p(y))$  for each fake image generated by the generator ( $y = 0$ ), which is the log probability of it being fake.

The metric of accuracy is fairly straightforward, as it provides a percentage of how many times the discriminator was correct at its prediction. In GANs, usually this metric fluctuates in the beginning and ideally converges around 80%, at which point the discriminator is fooled by the generator, indicating a desirable generator

model. More details for accuracy are presented in the Results and Discussion sections.

Adam is the optimization algorithm used instead of a classic stochastic gradient descent for weight updating. Adam is explained in detail in the VAE iteration referenced in the previous section.

# Chapter 4

## Results

In this section, we present the outcome of the generative modeling we propose, presented in the Methods section. Both the VAE and GAN were trained on three different datasets.

- Classical MIDI Collection: 295 tracks converted to 16400 images, all of the classical genre, with great variance on composition (based on different composers from different eras and cultures).
- Undertale video game soundtrack: 110 tracks converted to 3207 images, video game music, fair variance between tracks.
- Pokemon Mystery Dungeon (PMD) Blue/Red Rescue Team video game soundtrack: 11 tracks converted to 298 images, ambient music with little variance between tracks.

By the term variance between tracks we refer to the difference between melodies and sounds. For example in the Pokemon soundtrack, the tracks have all fairly similar ambient rhythm, in the Undertale dataset some tracks are similar to the ambient type, but there is also presence of faster rock type tracks, and in the Classical dataset the difference is greater, due to different styles of composers included. Furthermore, it should be interesting to determine whether that factor affects network performance.

Training two models on three different datasets (six different iterations) is believed to be enough in order to infer accurate results, plus we get the option of comparing the models with themselves.

In an ideal situation, we would want each model to generate images (that are converted back to MIDI tracks), which closely resemble the ones from the datasets with the addition of exploration and originality. In practice, this is not the case as the results are not melodic enough or consumer grade. This work focuses on extracting information as efficiently as possible (due to the limited training time and resources) and obtaining a model that generates images as close to the desired result as possible. This distinction is important though, as good quality images, with little noise, does not necessarily mean that their conversion to music is going to be melodic.

Furthermore, this section is split between two parts, each for every model presented in the Methods section. Both models, VAE and GAN, will be analyzed based on metrics, losses, and on some statistical analysis that can help explain the nature of the models or pose questions that will be addressed in the Discussion section.

In general, metrics, such as loss, can sometimes be misleading. In a generative task the loss between epochs is expected to behave in certain ways, nevertheless such a task is subjective and the results cannot be based solely on that factor.

The statistical analysis presented simply measures the probability of each note in each initial dataset. Then, the same analysis is applied to a set of generated images from the models and are then compared. The idea is if the behavior matches, the respective model will have a tendency to generate results close to the desired ones.

In conclusion, good results in such a task are subjective, and good sound or melodies may vary between even humans. In practice, the outcome of this exploratory work does not generate consumer-quality music, although there are some notable results.

## 4.1 Variational Autoencoder training results

First, the training process of the VAE will be presented, shown by examples of some epochs along with their respective latent space representation (which is crucial for generating results close to the dataset), followed by a loss plot. At last, we present some generated examples, along with a statistical analysis of note frequency between the images of the dataset and the generated ones. This whole process is replicated for all three datasets.

It can be expected that the Undertale dataset will perform better, as it is rather simplistic musically, compared to the Classical dataset, and it contains more samples compared to the PMD dataset.

### 4.1.1 Training the VAE model

The VAE model proved to be heavier computation-wise, due to the high number of trainable parameters (around 5 million), compared to the GAN (around 2 million). Regardless, the training time is fairly low, as it reaches convergence after few epochs, based on the dataset provided.

For each epoch the model trains upon batches, meaning that it executes back propagation, considering a number of inputs, a batch, rather than updating the variables after each input is fed through. This number is set to 32, as in typical training situations it is set to 32, 64, or 128. In addition, the model reserves a portion of the dataset for validation at the end of each training epoch and the validation loss is saved for plotting. For practical reasons, the dedicated script also saves the complete model, as well as the decoder model, offering the ability to retrain or generate samples.

The training history of VAE for each dataset is presented in Figure 4.1 with some generated examples. Those samples are the result of utilizing the decoder model, feeding it with random values in range of  $-4$  to  $4$ . As it was stated in Methods, the VAE modeling builds a latent space during training, which indicates the position of the clusters of each class in space. In this case, the class is one, so we observe a single cluster, as well as the random variables are two, so the space will be a

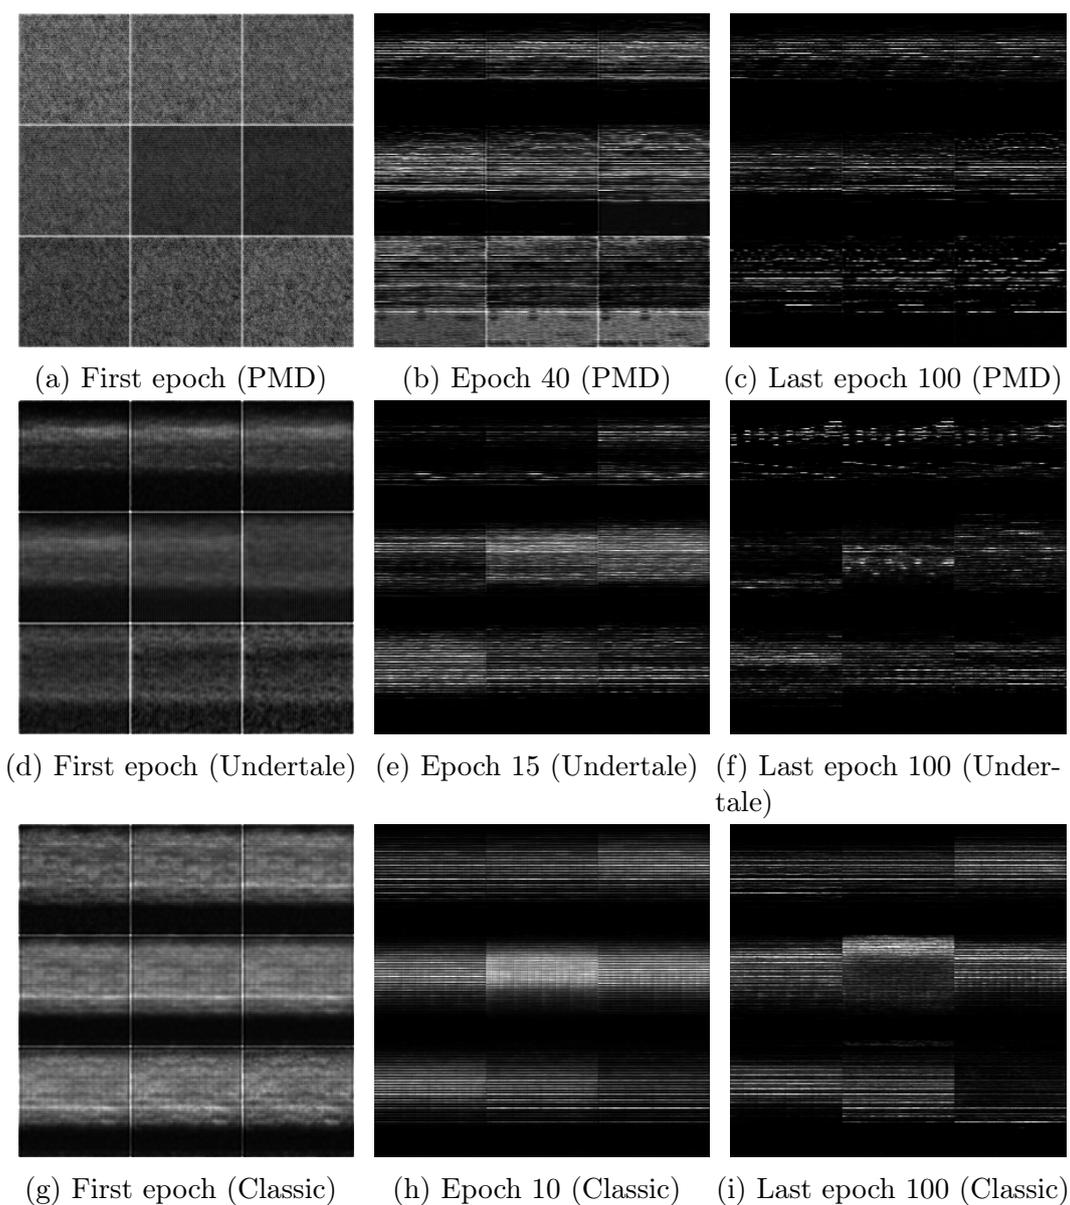


Figure 4.1: Generated samples for VAE during training

two-dimensional plot. The range of the random variables is set in range of  $-4$  to  $4$ , as we expect the cluster to be a circle positioned around  $(0, 0)$  with a radius of  $4$ .

It is observed that initially, epoch 0, the model generates images with random noise for each dataset. Afterwards, it is able to identify some patterns of the

initial datasets. For PMD, the model at epoch 40, is observed to have a tendency to generate musical samples, although noise is still present. Same principle can be seen in the Undertale dataset at epoch 15, as well as the Classical dataset at epoch 10. At last the training is stopped at epoch 100 for all three datasets for comparison.

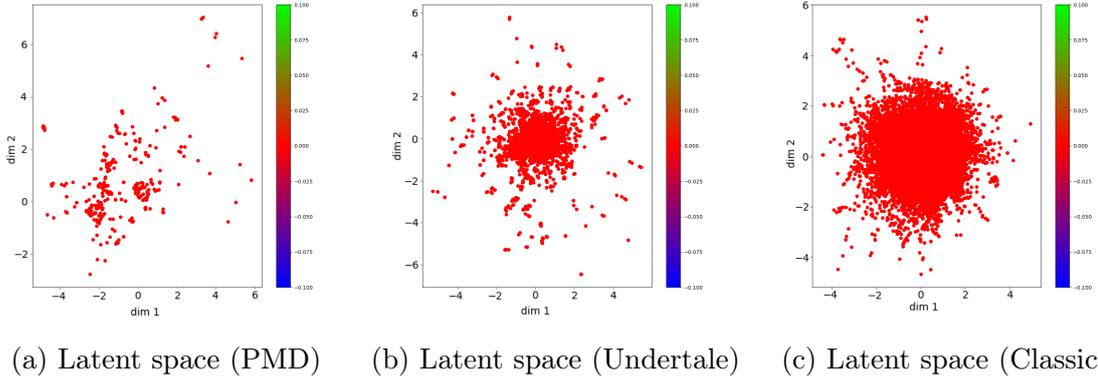


Figure 4.2: VAE Latent space for each dataset

Figure 4.2 presents the said latent spaces for each respective epoch of the generated results.

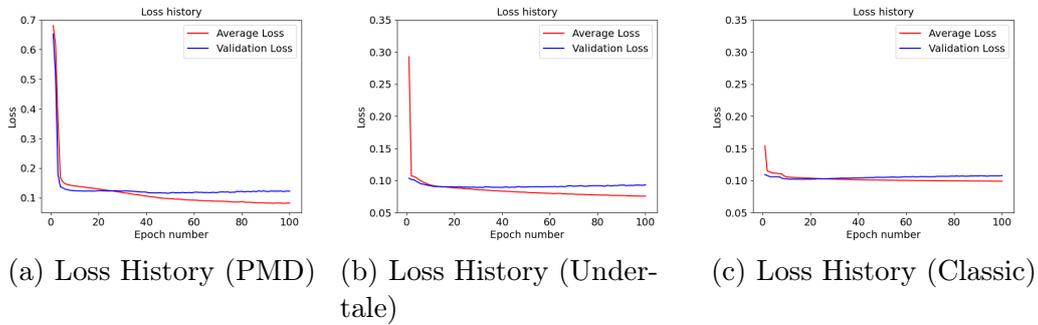


Figure 4.3: VAE Loss history for each dataset

Figure 4.3 presents the losses in a graph based on each epoch. The loss with color red corresponds to the average loss of the samples the model is trained upon, and the loss with color blue, to the validation loss, which is calculated based on a small portion of the dataset (20%), which is calculated at the end of each epoch.

Finally, it is obvious that the number of epochs required to generate images that resemble samples from the initial pool vary for each dataset, due to the fact that each dataset contains a different number of input samples. It is expected that less samples will execute less back propagation calculations for each epoch, resulting in requiring a greater number of epochs. In this case, all three iterations are trained for 100 epochs for comparison reasons, presented in the Discussion section.

### 4.1.2 VAE evaluation

Evaluating a generative task is not easy, as the results are highly subjective. There is the choice of evaluation based on musical theory, but the approach of this work is to be as simplistic as possible.

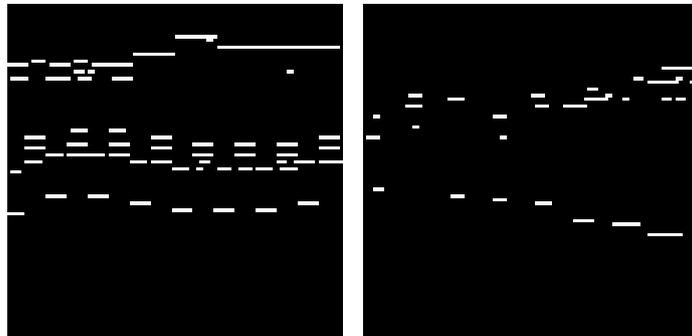
Due to the nature of music, it is not possible to describe in words a musical part, or whether it sounds melodic, although some samples are presented in Figure 4.4. These images have been processed to only contain black or white pixels, in order to convert them to MIDI, meaning the actual result is noisier.

In Figures 4.5, 4.6, 4.7, we present a statistical analysis on note frequency between tracks from the dataset and generated tracks. Each iteration generates 1000 samples and calculates a frequency for each note presented. Then we present a second figure referencing the difference between original and generated frequencies for each note for presentation purposes. This analysis aims to indicate the behavior of each iteration, whether the model is able to generate samples that replicate each original dataset. This feature will be further analyzed in the Discussion section.



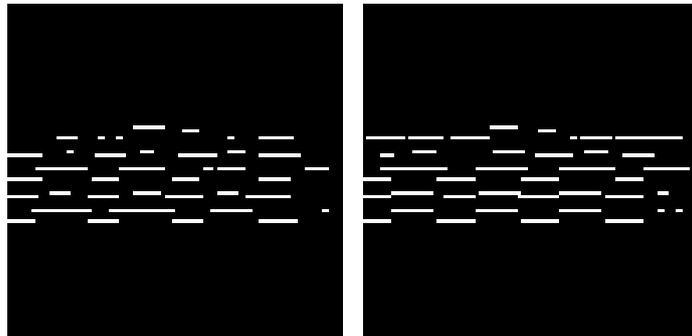
(a) VAE PMD track 11  
(encodings: 2.2, 2.0)

(b) VAE PMD track 28  
(encodings: 2.6, 2.9)



(c) VAE Undertale track 11  
(encodings: 3.3, -3.6)

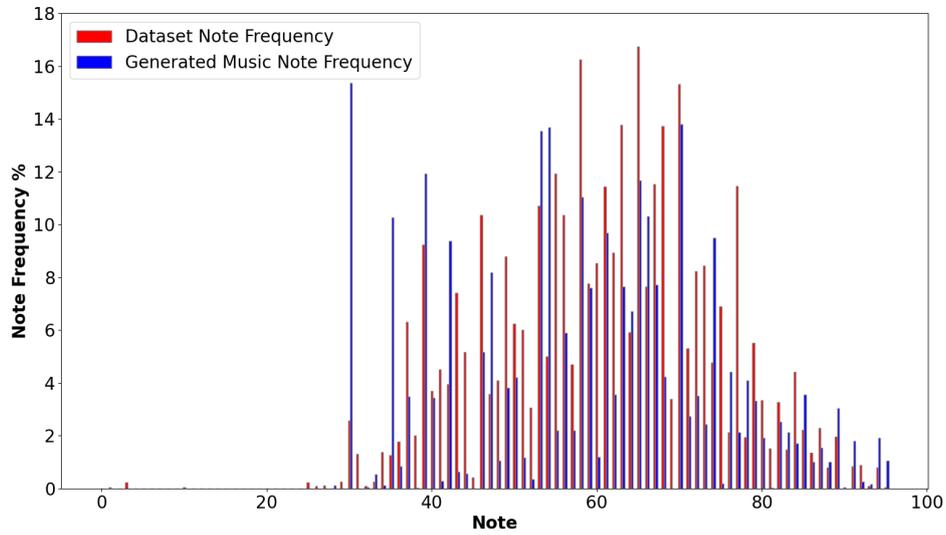
(d) VAE Undertale track 41  
(encodings: -2.0, 1.7)



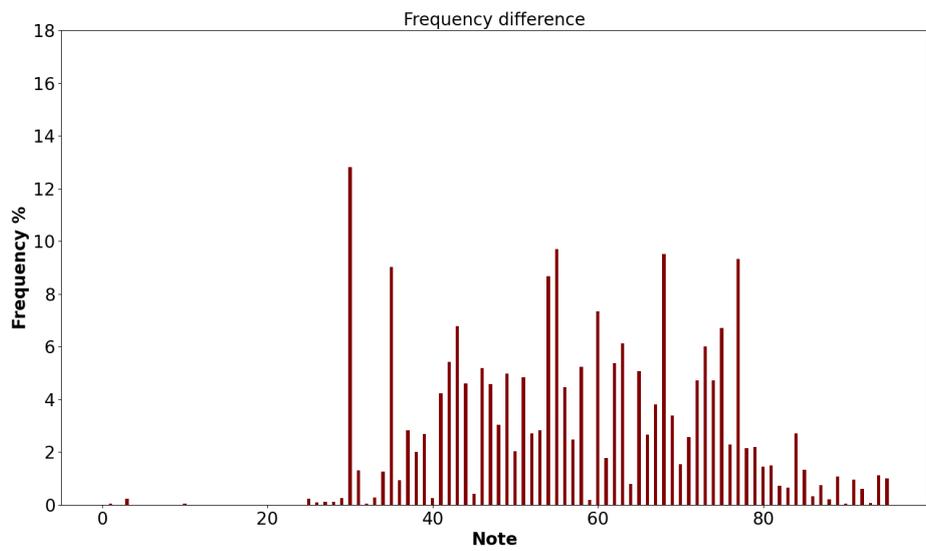
(e) VAE Classic track 209  
(encodings: -3.3, 3.3)

(f) VAE Classic track 212  
(encodings: -2.8, 3.5)

Figure 4.4: Notable samples generated by the VAE

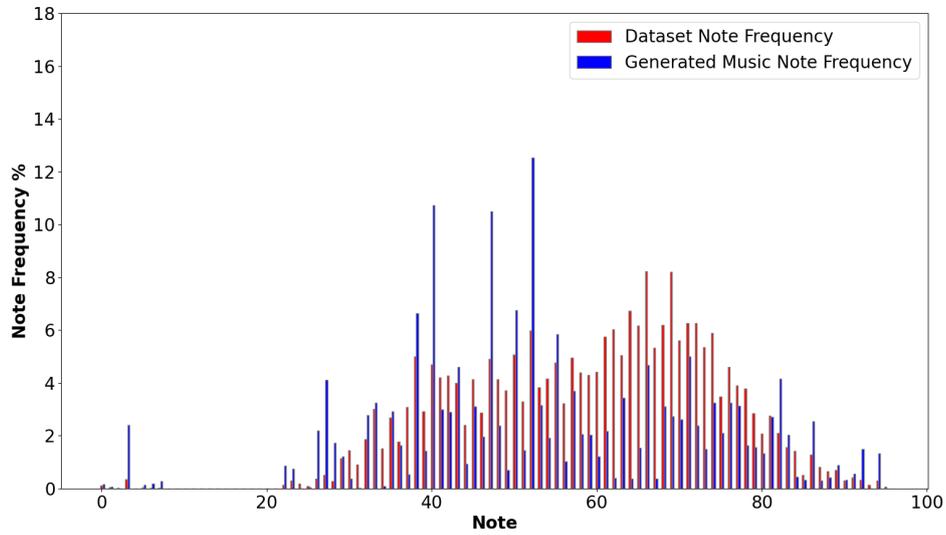


(a) Note Frequency (PMD)

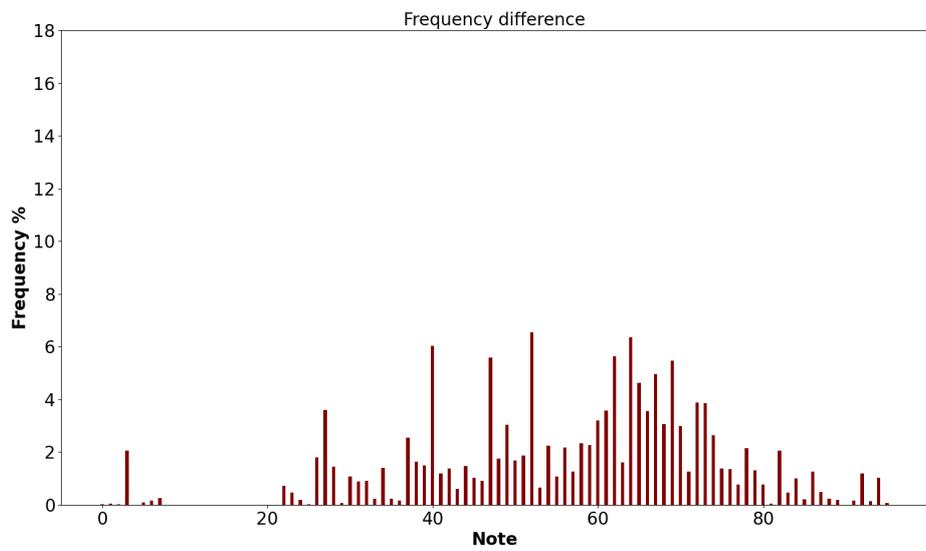


(b) Frequency Difference (PMD)

Figure 4.5: Note frequency bar graphs (for the PMD dataset) between real and (1000) generated samples by the VAE

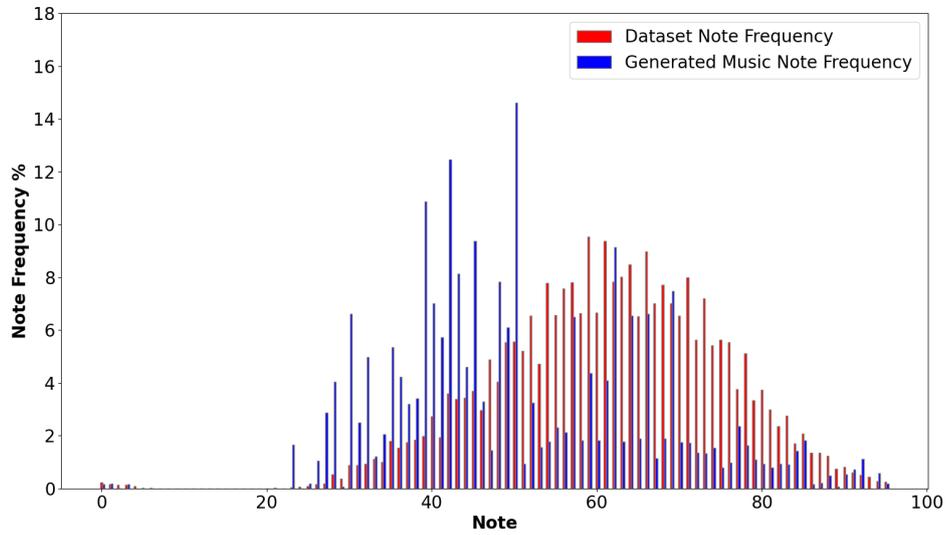


(a) Note Frequency (Undertale)

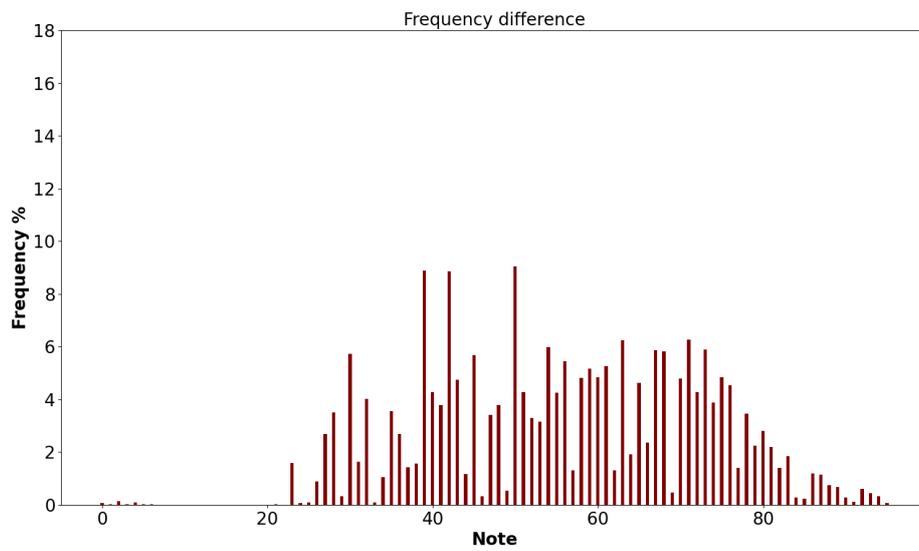


(b) Frequency Difference (Undertale)

Figure 4.6: Note frequency bar graphs (for the Undertale dataset) between real and (1000) generated samples by the VAE



(a) Note Frequency (Classic)



(b) Frequency Difference (Classic)

Figure 4.7: Note frequency bar graphs (for the Classic dataset) between real and (1000) generated samples by the VAE

## 4.2 GAN training results

In this subsection, the GAN modeling results are presented, replicating the VAE format. Firstly, training examples are shown along with loss plots. Afterwards, the generated examples are analysed with the same statistical process, as used for the VAE model. Each subsection evaluates the model based on all three datasets.

### 4.2.1 Training the GAN model

The GAN model, as it was built in the Methods section, consists of around 2 million number of trainable parameters (Generator and Discriminator combined). Also the number of epochs required to reach a desired stage was higher for all three datasets compared to the VAE.

Just like the VAE iteration, the GAN model is trained upon batches, splitting the dataset into groups of 64, executing back propagation after each batch, rather than updating the trainable parameters after every input sample is passed through, thus reducing the training process, while still maintaining effectiveness. Compared to the VAE, the GAN does not utilize the validation method, since loss is observed based on the two sub-models (generator, discriminator) rather than based on the complete model. For practical reasons, the script that is responsible for training also saves the model as a checkpoint, along with the generator, offering the ability to retrain or generate samples outside of the training process.

Figure 4.8, presents the training history of the GAN based on some generated samples. Those samples are the result of utilizing the generator model, feeding it with random noise input vector of 100 values ranging from  $-1$  to  $+1$ . Compared to the VAE, there is no latent space representation or clustering of the random input that can be visualised.

It is observed that initially, epoch 0, the model generates images that cannot be decoded to sound, as they appear either gray (PMD), or black (Undertale). The Classic dataset appears to output only some gray spots that are still unable to be converted to sound. Afterwards, for PMD, at epoch 40, and for Undertale, epoch 3, the models are observed to generate some non-black spots, although conversion

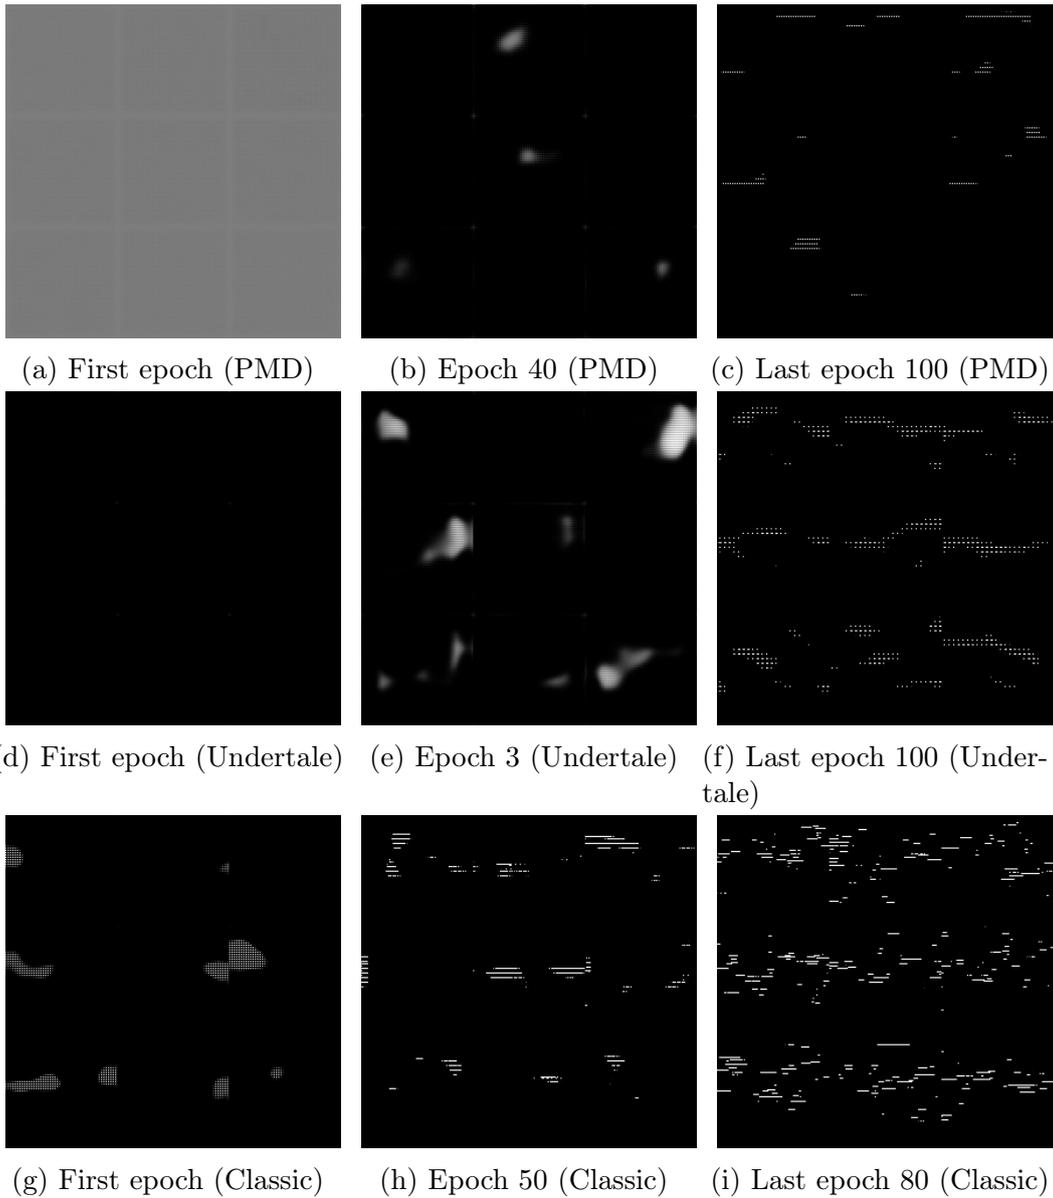


Figure 4.8: Generated samples for GAN during training.

to music is still impossible. At last, the training is stopped at epoch 100 for PMD, epoch 100 for Undertale, and epoch 80 for the Classic dataset, as the first two iterations seem unable to generate consistent images compared to the Classic iteration, which is closer to what an ideal result would be.

Figure 4.9, presents the losses in a graph based on each epoch. The loss with color

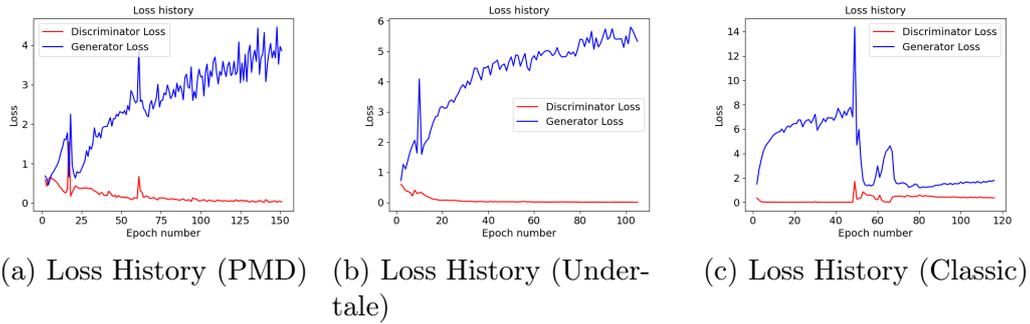


Figure 4.9: GAN loss history for each dataset

blue corresponds to the average generator loss for all the batches in each epoch, and the loss with color red, to the average discriminator loss. The loss values of generator and discriminator are expected to behave in an opposite manner, meaning that if the generator is performing well, low loss value, then the discriminator is fooled, thus outputting high loss value. More details and analysis on GAN performance based on losses are presented in the Discussion section.

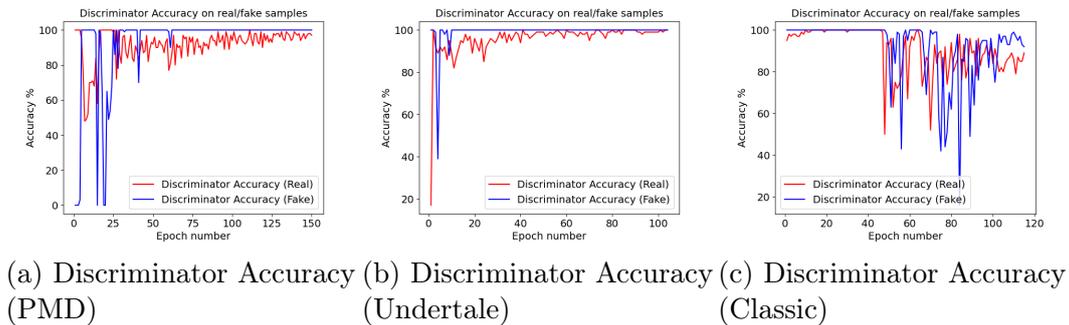


Figure 4.10: Discriminator Accuracy history for each dataset

In Figure 4.10 the accuracy of the discriminator is presented. This accuracy value refers to the ability of the discriminator to successfully identify whether its input is real (given by the dataset) or fake (generated by the generator).

Finally, just like in the VAE analysis, some iterations output results faster than others, in terms of epochs, due to the fact that each dataset contains different number of samples. More input samples lead to more calculations, but more training time for each epoch. The PMD and Undertale iterations are stopped at

epoch 100, as they are deemed unable to generate desired results, and the Classic iteration is stopped at epoch 80, as the model generates close to ideal results. More details are presented in the Discussion section about model performance and behavior.

## 4.2.2 GAN evaluation

In this subsection the GAN model is evaluated the same way as the VAE model. Generated samples are presented along with the probability of each note for every dataset. Once again a generative task is one of the hardest deep learning problems, as there is no clear way to determine whether the results are desired or not (subjectivity), whereas in classic deep learning tasks the performance is usually dictated by metrics.

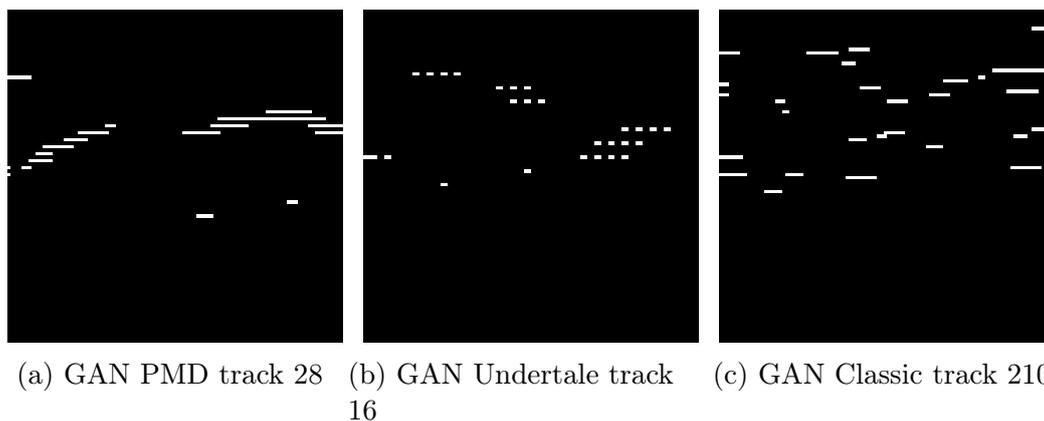
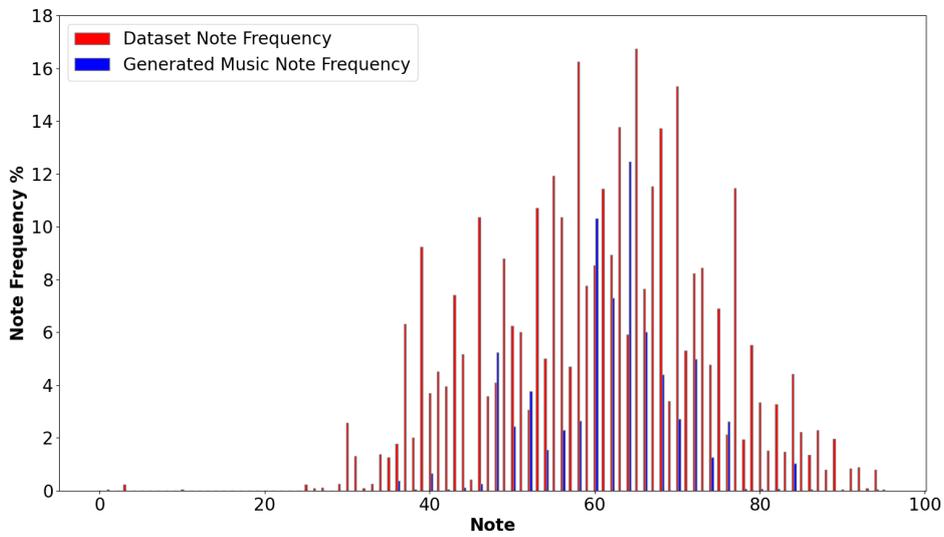


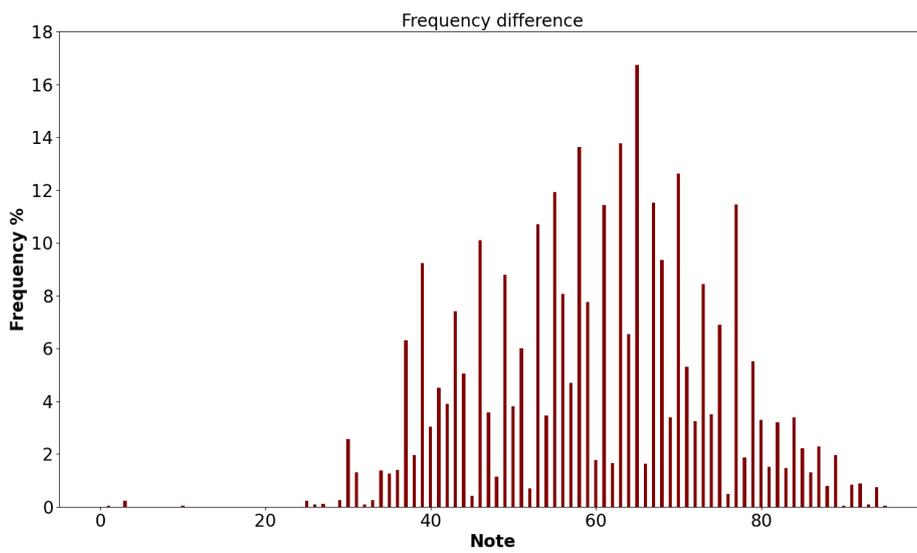
Figure 4.11: Samples generated by the GAN

The generated samples produced by the GAN are shown in Figure 4.11. Just like the VAE results, the GAN samples are normalized to output images with black or white pixels for the MIDI conversion, although, compared to the VAE, the GAN samples appear cleaner, with little noise, especially in the Classic dataset’s case.

In Figures 4.12, 4.13, 4.14, we present a statistical analysis on note frequency between tracks from the dataset and generated tracks, just like the VAE, along with difference graphs. Each iteration again generates 1000 samples and calculates a frequency for each note presented. This feature will be further analyzed in the



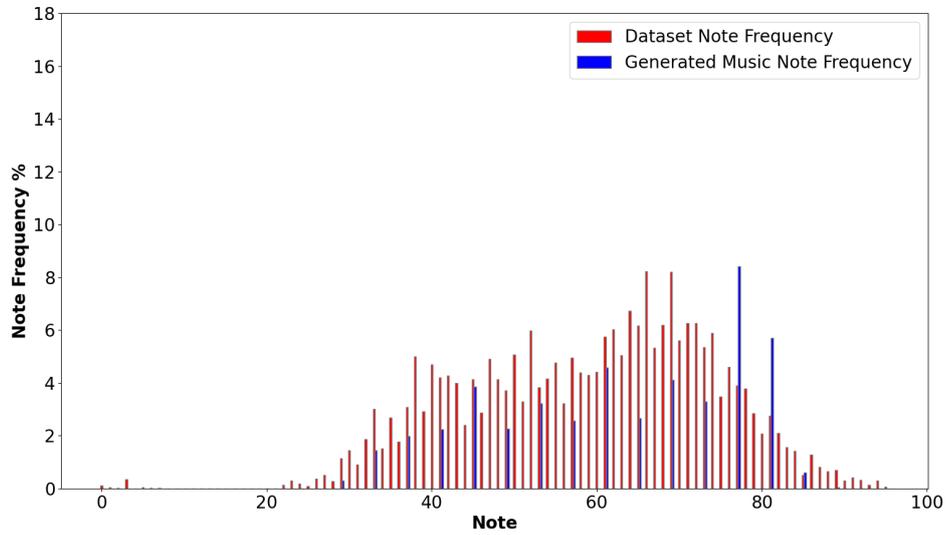
(a) Note Frequency (PMD)



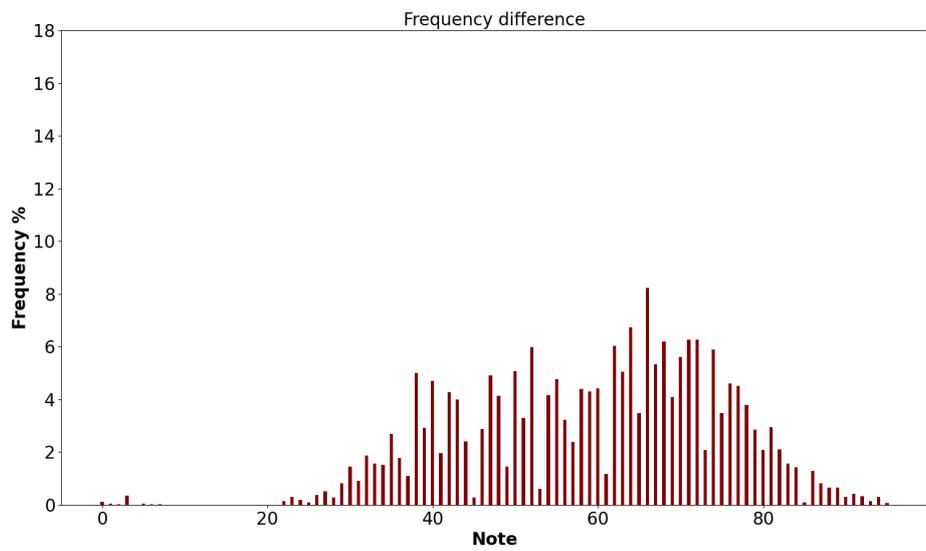
(b) Frequency Difference (PMD)

Figure 4.12: Note frequency bar graphs (for the PMD dataset) between real and (1000) generated samples by the GAN

Discussion section.

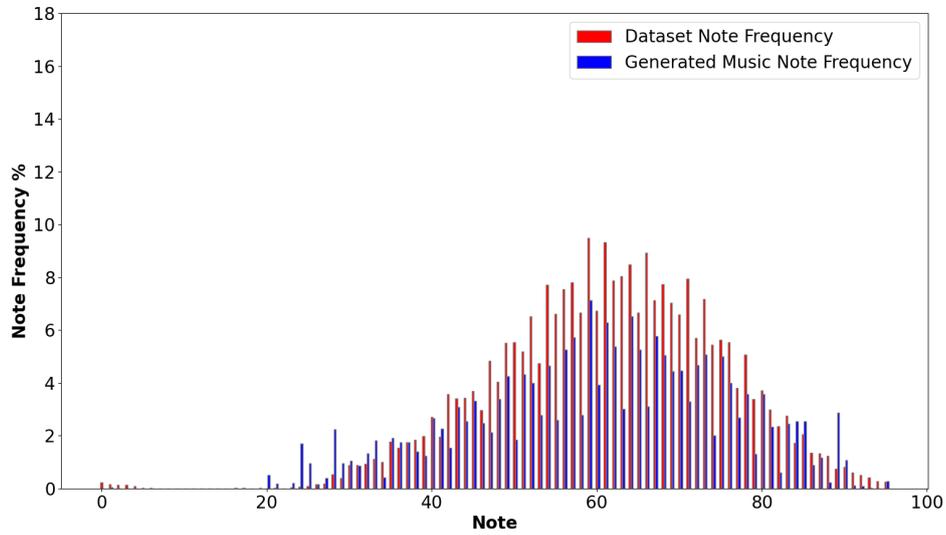


(a) Note Frequency (Undertale)

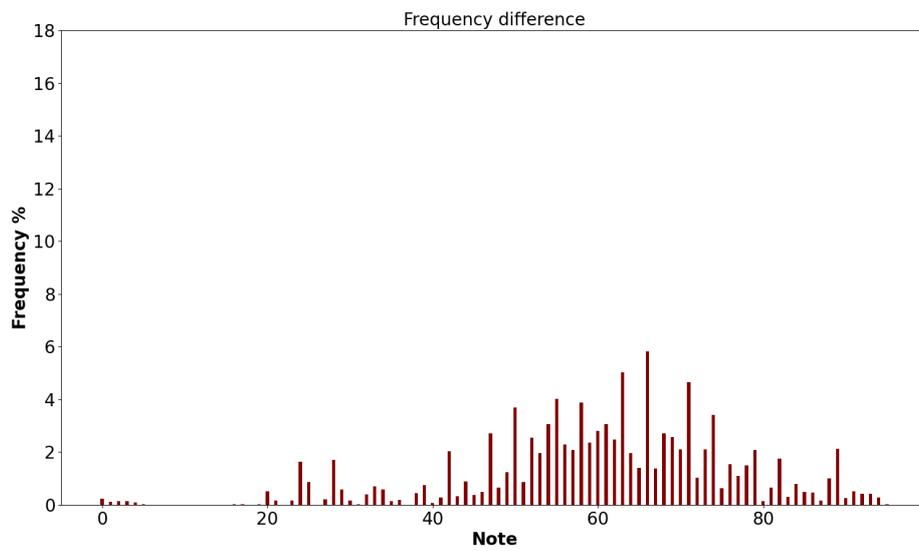


(b) Frequency Difference (Undertale)

Figure 4.13: Note frequency bar graphs (for the Undertale dataset) between real and (1000) generated samples by the GAN



(a) Note Frequency (Classic)



(b) Frequency Difference (Classic)

Figure 4.14: Note frequency bar graphs (for the Classic dataset) between real and (1000) generated samples by the GAN

# Chapter 5

## Discussion

This chapter's role is to comment, discuss and observe upon the results given by the training of the VAE and GAN models presented previously. The main goal is to identify patterns among observations, give reason to possible errors, as well as present possible solutions.

In general, based on Figures 4.1 and 4.8, the models behave differently for each dataset. The GAN in general generates cleaner images with little noise compared to the VAE. Although, in PMD there is little note selection, in Undertale the notes appear choppy and non-continuous, and in Classic the tracks seem complex and somewhat random (Figure 4.8). In VAE's case, noise is more present, PMD and Undertale iterations output some fair results, and Classic iteration is the noisiest and seems to not produce clear note sequences (Figure 4.1). It should be noted that those results in said figures are the raw unprocessed outputs of the networks, but for listening purposes then get filtered to end up either with a white or a black pixel (not grayscale).

### 5.1 Note probability graphs

Both models, VAE and GAN, are able to be trained upon all three datasets and generate desired images that resemble piano tablatures. The VAE iteration generates more noisy images, the Classic's case being the most noisy, and the GAN

outputs cleaner images. The main goal of this work, which is the extraction of images representing music and generating new ones based on that behavior, is achieved.

This offers the ability to then comment on those results and evaluate them based on melody and musical potential. Based on the probability graphs in the Results section (Figures 4.5, 4.6, 4.7, 4.12, 4.13, 4.14), we observe that the notes selected from the generated samples resemble the ones from the datasets, with some offsets as seen in the difference of probability graphs. Those offsets can be present due to various reasons:

- The models are unable to replicate the datasets' images, meaning a big value of difference is considered a non-ideal result.
- The models are defined in a way to achieve exploration or generate samples that are original compared to the inputs, meaning some value of difference is considered an ideal result.

In order to distinguish whether the models lack the ability to replicate the datasets or offer exploration and originality we observe the graphs considering spatial difference. We infer that big value of difference can be acceptable in spaces that appear frequently, but not in spaces where note probability from the datasets is small. For example, in both Figures 4.5 and 4.6, differences are presented in a continuous way, yet in the first case, PMD, the average difference is greater than the Undertale dataset, which makes the latter case's differences acceptable.

In VAE's case the total notes selected for all three datasets seem to resemble the notes selected from the dataset, although the differences vary for each dataset. The Classic iteration seems to generate lower notes more frequently than higher ones, compared to the dataset which favors high notes, so the difference is rather great (Figure 4.7). The PMD iteration presents some great variances between some notes, indicating that the model does not replicate the dataset (Figure 4.7). At last, in the Undertale iteration the frequency difference is smaller, even in notes that the variance is the greatest e.g. around 6% (Figure 4.6). In summary, the VAE generates images that contain a total number of notes close to the number

of the datasets, nevertheless the frequency differences vary for each dataset. The Undertale iteration seems to be performing better than the other two datasets in this case.

The GAN model's total note selection varies from the VAE iteration. Training upon the Undertale dataset, the GAN generates images that contain less than 20 different notes in total. The differences appear small, although that is the result of the model not selecting enough different notes (Figure 4.13). The PMD iteration outputs more total notes, but still not as many as the dataset, and the frequency differences are great (Figure 4.12). At last, in the Classic iteration the GAN generates notes close to the total number of different notes in the dataset, as well as the average frequency differences seem low, around 4% (Figure 4.14). In summary, the PMD and Undertale iterations do not resemble the input dataset, as they select fewer total different notes, and the Classic iteration is what an ideal result should look like, selecting a great number of notes (close to equal to the dataset) and having small frequency differences.

Regardless, there are two problems that are presented with the case of spatial frequency differences: firstly, in practice the result of altering a note to another, which is close spatially can be completely different musically in a non-ideal way, and, secondly, this analysis observes solely the selection of notes without considering the temporal attribute in the nature of music. For example, if we alter a melodic piano track temporally, keeping the same notes, but changing the order of the notes or the time offsets between them, the resulted track will sound completely different and in most cases random. Furthermore, we infer that this analysis is not sufficient enough to evaluate model performance, but it offers some insights about note selection and exploration.

## 5.2 Metric plots

Models in Deep Learning tasks are usually evaluated based on metrics, such as loss or accuracy. In a classic task, such as classification, loss is a way to determine when to stop training as well as observe the training history. Typically, in generative tasks, such as the one of the current work, metric evaluation is more complicated,

due to the nature of subjectivity of the results. Regardless, even generative models are expected to behave in a certain way considering metric plots.

### 5.2.1 VAE metric plots

The training of the VAE model is evaluated based on metrics, such as the validation loss, which is the loss computed based on a certain portion of the dataset reserved for prediction at the end of each epoch. The average training loss is also included for more insight. This loss is computed based on reconstruction loss and KL divergence, as stated in the Methods section (3.7), (3.4). For the three datasets (Figure 4.3) the VAE converges to around 0.1 on validation loss. The average training loss seems to be lower than the validation loss and in the cases of PMD and Undertale is not yet converged. In general, lower training loss compared to validation loss is an indication of overfitting, meaning the model is able to perform well on the training set, but struggles to replicate the same performance on data it has not yet processed. Ideally the difference of those two values should be close to 0, as well as the values themselves should be as low as possible. While there is a difference present, the values remain fairly low (close to 0.1). The PMD iteration starts with high losses at the first epoch and converge after a couple of iterations, also the variance between training loss and validation loss is greater due to the fact that the model overfits. Both of those reasons are the result of few input samples on the initial dataset (PMD). The greater number of samples a dataset contains, the more training it receives after one epoch, in comparison. The Undertale dataset contains more items, so the initial losses are smaller and the model does not overfit at the same rate. The same can be seen in the Classic dataset containing more items than both the other two datasets. In summary, the loss behavior of the VAE on all three datasets is what should be expected and can be considered as ideal results. With the loss histories, we could select models from earlier epochs for generation purposes, in hindsight.

## 5.2.2 GAN metric plots

In GAN's case, the model evaluation based on loss is completely different compared to traditional models, due to the fact that we observe two losses of two competing sub-models, the generator and the discriminator (Figure 4.9). Also, we get more information on the discriminator's decision-making based on the accuracy history (Figure 4.10).

Based on the Figure 4.9, we observe in the first two cases (PMD and Undertale) the discriminator outperforms the generator as the epochs increase, and, in the case of the Classic dataset, that is also present up to some point, around epoch 50, when the loss difference decreases.

The training of a GAN model is difficult for the fact that the two sub-models are trained at the same time in a zero sum game. This means that improvements to one model come at the expense of the other. GANs have plenty of problems that require solution during training, such as Mode Collapse, which is the result of the generator producing one or a small subset of different outcomes, and the Convergence Failure, which is the struggle of the whole model to reach loss convergence. The latter becomes present in this work's case, as in the first two datasets the model is unable to converge to a point of equilibrium.

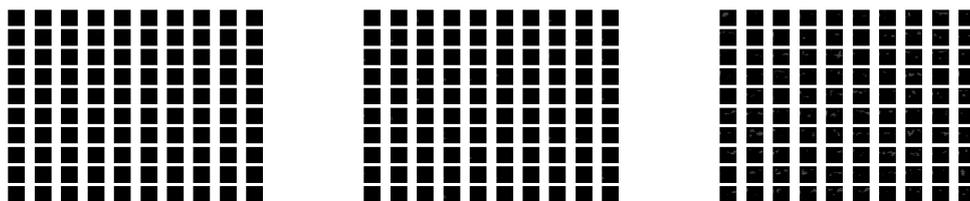
In an ideal scenario, a stable GAN will have a discriminator loss around 0.5 and a generator loss around 1.0. Also the accuracy of the discriminator on both real and generated (fake) samples should be around 70% and 80% (Brownlee 2019c). It should be noted that initially, in the early epochs, the loss variance is expected to be great, as the behavior of the sub-models are expected to start off erratic and move around a lot before convergence.

In Classic's case, the GAN is able to escape from the great loss variance, at around epoch 50 (Figure 4.9), and the losses hover around 1.0 for the discriminator and 2.0 for the generator. This stage of convergence is close to the ideal outcome, so the case of convergence failure is not as present. Also, the discriminator accuracy in Figure 4.10 for the Classic dataset seems to behave close to the ideal scenario, as the discriminator seems to struggle to succeed in distinguishing fake and real

samples 100% of the time.

In the other two cases (PMD and Undertale), the models suffer from convergence failure, loss difference is high, as it can be seen in Figure 4.9, and discriminator accuracy is high, especially in fake samples in Figure 4.10. This is the cause of the generator outputting completely random images, making it fairly easy for the discriminator to distinguish fakes from reals.

In all cases the generator at some point generates full black images. In Classic’s case from epoch 5 to 50 (Figure 5.1), it generates only black images and then is able to recover and generate samples that replicate the initial dataset. Convergence failure is present due to the fact that the discriminator easily identifies the fake ones from the generator by being completely black images, not resembling at all the initial dataset.



(a) GAN Classic epoch 5 (b) GAN Classic epoch 30 (c) GAN Classic epoch 50

Figure 5.1: GAN generating fully black images up to epoch 50

There are ways to combat this type of failure in GANs, as it is a fairly common issue in training competing models. One solution would be to add noise to the discriminator input (Martin Arjovsky 2017). Another would be to penalize discriminator weights (Roth et al. 2017). Both of those solutions force the discriminator to fail more frequently, giving the generator a chance to recover.

### 5.3 Listening to the results

Due to the nature of music, it is not possible to describe music in words, or in this case dictate whether the generated images decode to melodic parts, without

considering music rules or principles (which can complicate this work). By decoding the generated images to MIDI files, we are able to listen to our results. Due to the complexity of input values to the generator sub-models (Decoder on VAE, Generator on GAN), it is not feasible to listen to every possible track that can be generated. The decoder receives two inputs with random values in a specific range (dictated by the space representation of clustering), and the generator receives a vector of 100 random values between  $-1.0$  and  $1.0$ .

By listening to many tracks generated based on all three datasets for each model, we infer that the models struggle to produce consumer-grade music. A future goal of the work would be to automate music production for certain situations, such as music for television or radio advertisements, but at this stage this is not feasible.

Nevertheless, we identified some notable images that we decoded to music in MIDI format that were interesting.

The best results were produced by the VAE, trained on the Undertale dataset. Although the VAE produces noisy images, after normalization, some tracks have a pleasant groove. Track 11 of the Undertale dataset, as shown in Figure 5.2, presents a close to ideal result, and when listened to, we understand the model has a tendency to create a sample with jazzy feeling, which is similar to the input dataset. By utilizing this track's encodings we can play around and produce samples that sound similar but are different. The other two datasets (PMD, Classic) produced fair results, although the sound is not as pleasant.

The GAN model generates cleaner images than the VAE. That being said, the sound those samples produce are not pleasant to listen to, as they do not resemble musical sequences.

We look for good overall results, rather than nitpicking some instances. In this case, we dictate the models are not ready for consumer-grade music generation.



Figure 5.2: VAE Track 11 (Undertale) sheet

## 5.4 Model comparison

Training two different models on three different datasets (6 iterations total) we get the option of comparing each combination for the same task we initially set to solve.

VAE pros and cons:

- + early loss convergence, less training time required
- + generated some notable samples on Undertale dataset
- + can manipulate inputs with encodings
- generates noisy images
- overall cannot produce consumer-grade music

GAN pros and cons:

- + sharper images
- + good performance on Classic dataset
- hard to train (convergence failure, many hours)

- output is harder to manipulate, based on input dimension (random vector of 100 items)
- overall cannot produce consumer-grade music

# Chapter 6

## Conclusions

The main goal of this work was to produce a model that could somewhat replicate images encoded from MIDI files, that could then be decoded to music.

In summary, we present the steps taken:

- A script converts MIDI files to 96 by 96 images, one axis for the note values and the other for the time ticks.
- The models, GAN and VAE, are trained upon the image data sets, saving model objects, loss values and space representation (for the VAE) at the end of each epoch. Training stops at a stage when desired samples are generated.
- Another script imports the generator sub-models (generator for GAN, decoder for VAE), generates and saves image samples for evaluation. Those images are also converted to music format and saved for listening.
- At last, an evaluation script is ran to create loss plots of training history, and bar graphs that indicate note probability of generated and input samples for comparison.

By extensive research and multiple iterations of deep learning models, we managed to create two solutions that were able to generate musical samples in a very efficient way regarding time and memory. Those models were trained upon three, different in length (number of samples) and genre, datasets each. Some iterations

outperformed others, but all successfully achieved what was initially desired.

The GAN model trained upon the Classic dataset escaped from convergence failure and based on loss plots and frequency difference bar graphs its training was successful and close to ideal, yet the decoding of images back to MIDI files was not that ideal. The VAE model in the Undertale case produced some samples that were notable musically, yet the samples produced were not as clean as the in the GAN's case. Both models were somewhat successful, although ideally we would want a model that is able to be trained upon different types of datasets and generate similar results.

The reason this work was undertaken was due to the fact that music has a huge role in human lives. With our exploratory proposal we are able to create a bridge between computer science, and more specifically deep learning, with music. Deep learning being a growing field of research in a very fast rate, could be a very useful tool for musicians and composers, eventually offering real-time generated and original music that the consumer has never listened to.

This work is a starting point, a simplistic solution for music generation, that other researchers may continue upon and perfect what was initiated. Furthermore, our results were not that bad, as there are very few instances of tracks that could be enjoyed and listened to, nevertheless, we prefer this work to be evaluated based on its efficiency and effectiveness of model training.

# Chapter 7

## Future Work

The nature of this work is mostly technical, aside from the engineering of the data pre-processing. That being said, there are actions that can be taken to further improve training and results.

Some solutions are presented that could result in more desired generated samples. Most complex deep learning tasks usually require a lot of training. Such is the case for the GAN, being a cost-intensive model to train (many parallel calculations, many epochs), which could output better results, if trained more. Specifically, it is observed to output fully black images after a certain epoch (for each dataset), with high generator loss and low discriminator loss, but with further training it could stabilize and outperform past epochs, such as the case of the Classic training, where the model is able to escape from convergence failure.

Also, there is much engineering that could be done concerning the deep learning layers, meaning the addition of certain layers, the number of repetition of some layer sequences (number of Conv2Ds along with activations), and the tuning of hyperparameters or other attributes (such as batch number, filters, kernel sizes etc.) could prove more efficient. This process is typically a matter of trial and error.

In addition, the pre-processing of the datasets could be altered. One idea is to utilize three channels instead of one, meaning colored images instead of gray-scaled.

This extends the time ticks of each image by three times the initial range, or it could be utilized to indicate the power at which each note is being pressed, offering an extra layer of possible melodic results.

This work relied on training datasets that consist of piano tracks. Given that such a basic iteration of music generation worked in a desired way, there could be multiple instrument training, for parallel music composition, or inputs from multiple genres. This is feasible as VAE modeling can be represented in space, and clustering could be classified, meaning we could potentially train upon multiple instruments or genres (classes) at the same time and generating samples based on the values at which each class clusters. Another option would be to utilize Conditional GAN models as they offer the option to implement conditions, as their name suggests, based on musical principles allowing for more control in training.

While researching, we initially tried out a note prediction approach, utilizing Recurrent Neural Networks (RNNs) and Long-Short term Memory (LSTM) models, where the model would be given an input of a track and would predict the next action that should be taken, or note that should be played, acting as a classification task. The problem we came across was that the number of combinations of parallel notes, also known as chords, of a dataset is high, resulting in thousands of nodes at the final layer of classification. With better handling and engineering, this solution could be just as good, or even better than a typical generative task, offering composers the ability of assistance in completing musical parts.

Finally, it is the nature of deep learning to require heavy computing and machines equipped with complex graphical units, as well as memory. That being said, we recommend retraining the current models, with higher in cost values of hyperparameters and attributes, such as filters of each layer. With faster rate of learning, the models could be trained for more epochs, offering greater insight about each model's behavior and results.

# Bibliography

- Bjørndalen, Ole Martin (2013). *Mido library documentation*. URL: <https://mido.readthedocs.io/en/latest/index.html>.
- Briot, Jean-Pierre, Gaëtan Hadjeres, and François-David Pachet (2017). “Deep learning techniques for music generation—a survey”. In: *arXiv preprint arXiv:1709.01620*.
- Brownlee, Jason (2019a). *A Gentle Introduction to Generative Adversarial Networks (GANs)*. URL: <https://machinelearningmastery.com/what-are-generative-adversarial-networks-gans/>.
- (2019b). *How Do Convolutional Layers Work in Deep Learning Neural Networks?* URL: <https://machinelearningmastery.com/convolutional-layers-for-deep-learning-neural-networks/>.
- (2019c). *How to Identify and Diagnose GAN Failure Modes*. URL: <https://machinelearningmastery.com/practical-guide-to-gan-failure-modes/>.
- Cope, D. (2000). *The Algorithmic Composer*. Computer Music and Digital Audio Series. A-R Editions. ISBN: 9780895794543. URL: <https://books.google.gr/books?id=rFGH07I2KTcC>.
- Cuthbert, Michael (2006). *Music21 documentation*. URL: <https://web.mit.edu/music21/doc/>.
- Goodfellow, Ian, Yoshua Bengio, and Aaron Courville (2016). *Deep Learning*. <http://www.deeplearningbook.org>. MIT Press.
- Han, Jun and Claudio Moraga (1995). “The influence of the sigmoid function parameters on the speed of backpropagation learning”. In: *From Natural to Artificial Neural Computation*. Ed. by José Mira and Francisco Sandoval. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 195–201. ISBN: 978-3-540-49288-7.

- Kingma, Diederik P. and Jimmy Ba (2015). “Adam: A Method for Stochastic Optimization”. In: *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*. Ed. by Yoshua Bengio and Yann LeCun. URL: <http://arxiv.org/abs/1412.6980>.
- LeCun, Yann et al. (1988). “A theoretical framework for back-propagation”. In: *Proceedings of the 1988 connectionist models summer school*. Vol. 1, pp. 21–28.
- Martin Arjovsky, Leon Bottou (2017). *Towards Principled Methods for Training Generative Adversarial Networks*. URL: <https://arxiv.org/pdf/1701.04862.pdf>.
- Piyush Madan, Samaya Madhavan (2020). *An introduction to deep learning*. URL: <https://developer.ibm.com/articles/an-introduction-to-deep-learning/>.
- Roth, Kevin et al. (2017). “Stabilizing training of generative adversarial networks through regularization”. In: *arXiv preprint arXiv:1705.09367*.
- Schmidt-Jones, Catherine (2013). “Understanding basic music theory”. In: Shafkat, Irhum (2018). *Intuitively Understanding Variational Autoencoders*. URL: <https://towardsdatascience.com/intuitively-understanding-variational-autoencoders-1bfe67eb5daf>.
- Xu, Bing et al. (2015). “Empirical Evaluation of Rectified Activations in Convolutional Network”. In: *CoRR* abs/1505.00853. arXiv: 1505.00853. URL: <http://arxiv.org/abs/1505.00853>.