

AUTOMATIC CHORD EXTRACTION AND MUSICAL  
STRUCTURE PREDICTION THROUGH  
SEMI-SUPERVISED LEARNING

Application to human-computer improvisation

TRISTAN CARSAULT

Supervised by Philippe Esling and Jérôme Nika

IRCAM, UMR STMS 9912  
Paris 75004, France  
Musical Representations Team



Master's degree  
SAR - ATIAM  
Faculty of Computer Science  
Université Pierre-et-Marie Curie / IRCAM / Télécom ParisTech  
February-August 2017



## ABSTRACT

---

Human computer co-improvisation aims to rely on a computer in order to produce a musical accompaniment to a musician's improvisation. Recently, the notion of guidance has been introduced to enhance the process of human computer co-improvisation. Although this concept has already been studied with a step-by-step guidance or by guiding with a formal temporal structure, it is usually only based on a past memory of events. This memory is derived from an annotated corpus which limits the possibility to infer the potential future improvisation structure. Nevertheless, most improvisations are based on long-term structures or grids. Our study intends to target these aspects and provide short term predictions of the musical structures to improve the quality of the computer co-improvisation.

Our aim is to develop a software that interacts in real-time with a musician by inferring expected structures. In order to achieve this goal, we divide the project into two main tasks: a *listening module* and a *symbolic generation module*. The listening module extracts the musical structure played by the musician whereas the generative module predicts musical sequences based on these extractions.

In this report, we present a first approach towards this goal by introducing an automatic chord extraction module and a chord label sequence generator. Regarding the structure extraction, as the current state-of-the-art results in automatic chord extraction are obtained with Convolutional Neural Networks (CNN), we first study new architectures derived from the CNNs applied to this task. However, as we underline in our study, the low quantity of audio labeled dataset could limit the use of machine learning algorithms. Hence, we also propose the use of Ladder Networks (LN) which can be trained in a semi-supervised way. This allows us to evaluate the use of unlabeled music data to improve labeled chord extraction. Regarding the chord label generator, many recent works showed the success of Recurrent Neural Networks (RNN) for generative temporal applications. Thus, we use a family of recurrent networks, the Long Short-Term Memory (LSTM) unit, for our generative task.

Here, we present our implementations and the results of our models by comparing to the current state-of-the-art and show that we obtain comparable results on the seminal evaluation datasets. Finally, we introduce the overall architecture of the software linking both modules and propose some directions of future work.

## RÉSUMÉ

---

La co-improvisation homme-machine a comme but l'utilisation d'un ordinateur dans l'accompagnement musical d'un musicien lors d'une improvisation. Les technologies existantes en improvisation homme-machine permettent un guidage pas-à-pas ou un guidage basé sur des formalismes structuraux pré-enregistrés. Ces structures calculées sont issues de corpus annotés, ce qui limite grandement la génération de structures improvisées construites sur un potentiel futur. Ainsi, notre logiciel aspire à combiner ces deux aspects en proposant à chaque pas des séquences de symboles en cohérence avec la progression des structures renseignées par le musicien.

Notre but est de concevoir un programme pouvant interagir en temps-réel avec un musicien en inférant des structures musicales futures. Cette conception peut se décomposer en deux étapes : un *module d'écoute* et un *module d'improvisation*. Le module d'écoute permet d'extraire la structure musicale jouée par le musicien, alors que le module d'improvisation est un générateur de séquences musicales symboliques.

Dans ce rapport, nous présentons les premières étapes du projet : l'extraction automatique d'accords musicaux et la génération de séquences de labels d'accords. A ce jour, les meilleurs algorithmes d'extraction automatique d'accords intègrent un réseau de neurones convolutif. Ainsi, nous avons étudié ces modèles pour la tâche d'extraction automatique d'accords. Cependant, nous avons souligné durant notre étude la faible quantité de bases de données audio annotées. Nous avons donc étudié le ladder network, un modèle de réseau de neurones pouvant être entraîné de façon semi-supervisée. Cet entraînement semi-supervisé nous permet de paramétrer notre modèle aussi bien avec des données labellisées qu'avec des données non labellisées. D'autre part, plusieurs logiciels de génération musicale ont été récemment développés avec des réseaux de neurones récurrents. C'est pourquoi, nous utilisons une catégorie de réseau de neurones récurrents, les long short term memory.

Ici, nous présentons nos implémentations ainsi que les résultats de nos modèles en les comparant à l'état de l'art actuel et en montrant que nous obtenons des résultats comparables sur les bases de données de références. Finalement, nous introduisons l'architecture globale du logiciel reliant les deux modules et nous proposons des directions pour des études ultérieures.

## ACKNOWLEDGEMENTS

---

This internship at IRCAM within the Music Representations team was a priceless experience. First of all, I would like to express my gratitude to my tutors Philippe Esling and Jérôme Nika. Due to their wisdom on interdependent topics, these two persons were absolutely complementary for a great supervision.

Thanks to the Music Representations team, Gérard Assayag, Carlos Agon, Jean Bresson, Jean-Louis Giavitto, Léopold Crestel, Axel Chemla-Romeu-Santos and Karim Haddad, for their kind advices, their sharing of knowledge and their daily happiness.

Thanks to the ATIAM pedagogical team for their implication in the development of their students.

Thanks to the ATIAM students for this unforgettable moments spent together.

A special thanks to my family and friends for their indefectible support.

## CONTENTS

---

1	MUSICAL MOTIVATION AND RELATED WORKS	11
1.1	Guidance in human-computer co-improvisation . . . . .	11
1.1.1	“Guiding” step by step . . . . .	11
1.1.2	“Guiding” with a Formal Temporal Structure or Description . . . . .	12
1.1.3	Intelligent listening and inference of short-term scenarios . . . . .	12
2	STATE OF THE ART IN MACHINE LEARNING	14
2.1	Machine learning . . . . .	14
2.1.1	Function approximation . . . . .	14
2.1.2	Loss function . . . . .	14
2.1.3	Types of learning algorithms . . . . .	15
2.1.4	Numerical optimization and training . . . . .	18
2.2	Neural networks . . . . .	20
2.2.1	Artificial neuron . . . . .	20
2.2.2	Multi-Layer Perceptron . . . . .	22
2.2.3	Convolutional network . . . . .	24
2.2.4	Ladder Networks . . . . .	26
2.2.5	Reconstruction cost and loss function . . . . .	27
2.3	Temporal structure and memory . . . . .	27
2.3.1	Recurrent Neural Networks . . . . .	27
2.3.2	LSTM . . . . .	28
3	EXTRACTING CHORD LABELS FROM AUDIO SIGNAL	30
3.1	Chord Progression Extraction . . . . .	30
3.1.1	Definition of the workflow . . . . .	30
3.1.2	Related works in Automatic Chord Extraction . . . . .	30
3.2	Datasets and audio pre-processing . . . . .	32
3.2.1	Isophonics datasets . . . . .	32
3.2.2	Audio transformation . . . . .	33
3.2.3	Chord alphabet reduction . . . . .	34
3.2.4	Transforming the chords into pitch vectors . . . . .	35
3.3	Our models . . . . .	35
3.3.1	Convolutional networks . . . . .	36
4	INTRODUCING TEMPORAL CONTEXT AND MEMORY	37
4.1	Post-filtering methods for ACE . . . . .	37
4.1.1	Real Book . . . . .	37
4.2	Generating new data from memory . . . . .	37
4.2.1	Generation of a chord sequence . . . . .	38
4.3	Our models . . . . .	38
4.4	Towards real-time prediction . . . . .	39
5	CONCLUSION AND PERSPECTIVES	40
	BIBLIOGRAPHY	42

## LIST OF FIGURES

---

Figure 1	Extracting structures and prediction. . . . .	13
Figure 2	Example of elements inside a chord labeling dataset based on spectrogram windows. . . . .	15
Figure 3	An example of supervised training for an ACE task, where we compare the ground truth annotation $a(t)$ with the approximation of our model $\tilde{f}_{\Theta}(x(t)) = y(t)$ . . . . .	16
Figure 4	An example of unsupervised training for an ACE task, where we compare a frame of spectrogram $x(t)$ with a denoising corrupt version of it $\tilde{f}_{\Theta}(x^*(t)) = z(t)$ . . . . .	17
Figure 5	Example of a two-class cluster classification dataset where two sample are labeled and others are unlabeled. The addition of unlabeled data modifies the cluster sizes and shapes. [58]. . . . .	17
Figure 6	Gradient of a two parameters function with two local minima and example of the use of gradient descent based on two different random initializations, image taken from [35]. . . . .	19
Figure 7	Two different models showing either overfitting (in green) or generalization (in black) for separation in a two-dimensional feature space with two classes, image taken from [40]. . . . .	19
Figure 8	Evolution of train and valid errors versus training time, image taken from [40]. . . . .	20
Figure 9	Schema of an artificial neuron. . . . .	21
Figure 10	Space separation with one linear output layer. Image from [3]. . . . .	22
Figure 11	Space transform operated by successive hidden layer. Image from [3]. . . . .	22
Figure 12	Multi-Layer Perceptron. . . . .	23
Figure 13	Illustration of the convolution operation, image taken from [57]. . . . .	25
Figure 14	Example of a 2X2 max pooling. . . . .	25
Figure 15	Example of a convolutional neural network for a task of ACE. . . . .	26
Figure 16	Semi-supervised learning on ACE task, the supervised classification (in blue) is reinforced by the unsupervised denoising autoencoder task (in green) . . . . .	26
Figure 17	Loops of a RNN unfolded through time, image taken from [54]. . . . .	28
Figure 18	Internal structure of a classical RNN cell, image taken from [54]. . . . .	29

Figure 19	Internal structure of a LSTM cell, image taken from [54]. . . . .	29
Figure 20	Syntax of Chord Notation in Isophonics, image taken from [17]. . . . .	33
Figure 21	Complete ACE system using CNN/LN as feature extractor and LSTM as post-filtering method. . . . .	37
Figure 22	The chord sequence prediction is realized with adding element to the chord sequence at each time step. . . . .	38
Figure 23	Our overall system architecture. Each frame of the audio flux is processed by the ACE system. Then the chord sequence is sent to a LSTM in order to predict the next chord. . . . .	39

## LIST OF TABLES

---

Table 1	Contributors for MIREX 2016 in ACE, for more information see MIREX website [31] . . . . .	32
Table 2	ACE on Isophonics 2009 for MIREX 2016 . . . . .	32
Table 3	Harmonization of the C Major scale. . . . .	35
Table 4	5-Fold recognition accuracy for ACE on Major/Minor classification task on Isophonics dataset with a CNN, results taken from [23] . . . . .	36
Table 5	1-Fold recognition accuracy for ACE on three different alphabet with a classic CNN on the Isophonics dataset. . . . .	36
Table 6	1-Fold chord prediction task with a LSTM trained on 8 time-step sequence from the realbook dataset. . . . .	38

## ACRONYMS

---

MIR Music Information Retrieval

MIREX Music Information Retrieval Evaluation eXchange

ACE Automatic Chord Extraction

SGD Stochastic Gradient Descent

RBM Restricted Boltzmann Machine

DBN Dynamic Bayesian Networks

HMM Hidden Markov Model

NN Neural Network

MLP Multi-Layer Perceptron

CNN Convolutional Neural Network

LN Ladder Network

RNN Recurrent Neural Network

LSTM Long Short Term Memory

## INTRODUCTION

---

The concept of *musical structure* could be defined as the arrangement and relations of musical elements through time. Furthermore, a piece of music has different levels of structure depending on the temporal scale that is under scrutiny. Indeed, elements of music such as notes (defined by their pitch, duration and timbre) can be combined into groups like chords, motifs and phrases and these, in turn, combine into larger structures such as chord progressions or choruses and verses. Thus, there can be complex and multi-scaled hierarchical and temporal relationships between the different musical elements. Given these complexities, musical structure have been given different meanings throughout the literature.

**STRUCTURE:** In this work, we define a structure as a sequence of symbols in a chosen alphabet (e.g. chord progression, profile of audio features), which describes the temporal evolution of a musical sequence. Hence, here we will not focus on other music information retrieval notions such as the high-level segmentation of an entire piece of music [43, 45].

Structures are inherent in music, and their knowledge becomes crucial in the process of improvisation involving multiple musicians. Indeed, even though improvisation is often associated with spontaneity and randomness, it is largely based on rules and structures that allow these different musicians to play together correctly. If we focus on blues or jazz improvisation, we can see that it generally relies on a chord progression defining a guideline for the overall performance. This chord progression is the backbone structure that the musicians will follow to develop their improvisations, while performing with others. Therefore, in a collective improvisation, an understanding of the current musical structures is critical.

Therefore, in this internship, we will target the development of intelligent listening modules that could emulate this process of musical structure discovery, as performed in real-time by a musician. To that end, we separate this task as first the ability to perform automatic chord extraction and then to understand higher-level musical progressions by performing chord prediction.

The remainder of this work is organized as follows. First, we present existing human-computer co-improvisation systems (Section 1) and motivate our choice to use chord extraction as the first step of our musical inference system. Then, we present the state-of-the-art in the Automatic Chord Extraction (ACE) field (Section 2). In the following part, we introduce the basic concepts of machine learning, neural networks and deep learning models (Section 3). Finally, we present the results of our chord extraction models (Section 4) and the architecture of our complete inference and co-improvisation model (Section 5).

## MUSICAL MOTIVATION AND RELATED WORKS

---

### 1.1 GUIDANCE IN HUMAN-COMPUTER CO-IMPROVISATION

The motivation of this work comes from the field of interactive music generation, defined as "*music that is produced by a hardware and software system that importantly includes a computer*" [9]. Over the past years, many works have focused on the idea of relying on a computer algorithm to generate music. Here, we focus on the symbolic music without deep details on the sound synthesis. In these works, the definition of the musical structure underlying the generation is central to the success of these algorithms. Among them, we focus on human-computer co-improvisation processes. That is to say, as system that plays music with a musician in real time.

For instance, Omax [2] is a system that learns features specific to the musician's style in real-time, and then plays along with him interactively. This technology is based on Oracle's theory [1] and generates music with the help of an audio database that has been previously analyzed. Here, we focus on another family of software that retains a control and authoring on the musical output. Indeed, as introduced by Nika [36], human-computer co-improvisation processes can be enhanced through the notion of *guidance*. In this domain of *guided* human-computer co-improvisation, the notion of *guiding* music generation has two different meanings. On the one hand, *guiding* can be seen as a purely reactive and step-by-step process. This approach offers rich interaction possibilities but cannot take advantage of the prior knowledge of a temporal structure to introduce anticipatory behavior. On the other hand, *guiding* can mean defining the temporal structures or descriptions themselves, in order to drive the generation process of the whole musical sequence. These "scenario-based" systems are able to introduce anticipatory behaviours but require some prior knowledge about the musical context (a pre-defined scenario).

#### 1.1.1 "Guiding" step by step

The step-by-step process aims to produce an automatic accompaniment using purely reactive mechanisms without prior knowledge. The musician input signal is analyzed in real-time and the system compares it to the corpus and select the most relevant music in order to generate an accompaniment.

For instance, SoMax [6] uses a previously annotated corpus and extracts in real time multimodal observations of the musician's playing. Then it retrieves the most relevant music slices from the corpus to generate an accompaniment.

Other software as VirtualBand [34] or Reflexive Looper [41], also relies on *feature based interaction*. Given the extracted features (e.g. RMS, spectral centroid or chroma) of the musician’s audio signal, the software selects the audio musical accompaniment from the database .

### 1.1.2 “Guiding” with a Formal Temporal Structure or Description

Another approach is to see the guidance as constraints for the generation of complete musical sequences.

On the one hand, constraints can be used to preserve some structural patterns present in the database. In 2015 Herremans et al. built a system to generate bagana music [20], a traditional lyre from Ethiopia based on a first order Markov model. Another project using Markovian model is a model for corpus-based generative electronic dance music that has been proposed in 2013 by Eigenfeldt and Pasquier [15].

On the other hand, some research projects introduced a temporal specification to guide the music generation process. For instance, Donzé applies this concept in order to generate a monophonic solo similar to a given training melody upon a given chord progression [14]. Recently, the notion of *musical scenario* has been proposed [38] as the specification of the high-level musical structure that generated sequences should follow. This approach allows to define an overall direction to the music generation process. It introduces such motion or intention at each step of the generation.

SCENARIO: In our work, we define a scenario as a formalized temporal structure guiding the music generation.

An example of scenario-based system is ImproteK [37], which uses pattern-matching algorithms on symbolic sequences. Then, the symbolized inputs from the musician are compared with pre-defined scenarios. ImproteK is also reactive, as the scenario can be modified via a set of pre-coded rules or with parameters controlled in real-time via an external operator [39].

### 1.1.3 Intelligent listening and inference of short-term scenarios

In 1969, Schoenberg and Stein [49] formulated a fundamental distinction between progression and succession. From their point of view, a progression is aimed at a definite goal and is oriented towards its future, whereas a succession is a step-by-step conformity. The systems using a formal temporal structure or description consolidate this notion of progression developed by Schoenberg and Stein. Moreover, they introduce a notion of anticipatory behavior. However, this concept of scenario is limited to a pre-defined setup, which limits the capability of the algorithm to foresee any future movement or change inside the music. Moreover, Huron formalized the chain [25]

$$\text{“expectation} \rightarrow \text{prediction} \rightarrow \text{anticipation”} \quad (1)$$

Thus, the scenario is the temporal specification that takes the place of prediction and expectation. In other words, this chain becomes “specification -> anticipation” in scenario-based systems. Therefore, a promising direction would be the inference of short-term scenarios in order to obtain the entire chain (1).

The over-arching goal of interactive music systems is to emulate the behavior of a trained human musician. Hence, our aim is to design an intelligent listening module able to perform a real-time discovery of the structures existing inside a musical stream. At the end, an inference of short-term structures could be achieved from this discovery in order to feed scenario-based generative systems. (see figure 1)

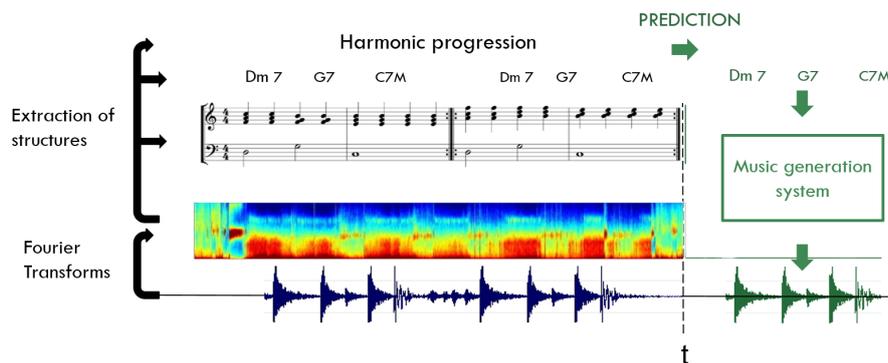


Figure 1: Extracting structures and prediction.

Multiple music generation applications could benefit from a real-time prediction of the musical structures, in particular, scenario-based generative softwares, such as ImproteK [37]. Based on this inference, music generation processes could combine the advantages of the two aforementioned forms of “guidance”. Hence, we aim for a system that could be responsive (in a real-time sense), while maintaining a long-term temporal vision, able to anticipate and generate improvisation from an inferred (and not only pre-defined) underlying structure.

**OUR MOTIVATION** is to infer future musical structures for improving human computer co-improvisation. In this report, we present our approaches for structure extraction in audio streams and for symbolic music prediction based on this information.

**OUR APPLICATION CASE** is musical chords. Indeed, chords are mid-level musical features which concisely describe the harmonic content of a piece. Moreover, chord sequences are often sufficient for musicians to play in an improvisation context. Thus, the models presented in this report principally concern the Automatic Chord Extraction (ACE) field and the generation of musical chord sequences.

## 2.1 MACHINE LEARNING

Machine learning aims to empower the computers with the capabilities to perform complex tasks that are innate to humans. The main idea is to develop algorithms able to learn by observing and modeling a set of examples. Thus, machine learning algorithms use computational methods to learn directly from the data by adapting the parameters of a pre-defined family of functions. The overarching goal of machine learning is to produce a model that could be able to generalize its understanding of a given (training) set to unseen data. In this section, we introduce the basic aspects and definitions underlying machine learning. Then, we introduce the different specific models that have been used along this internship.

2.1.1 *Function approximation*

In most machine learning problems, we start from a given dataset  $X = \{x_1, \dots, x_N\}; x_i \in \mathbb{R}^d$  and want to obtain a set of corresponding information in a given target space  $Y = \{y_1, \dots, y_N\}; y_i \in \mathbb{R}^t$ . This target space usually defines the goal of the learning system. However, in most cases, going from one space to the other can be seen as a transform  $F : \mathbb{R}^d \rightarrow \mathbb{R}^t$ , such that

$$F(x_i) = y_i, \forall i \in [1, N]$$

Hence, the aim of machine learning would be to find this function  $F \in \mathcal{F}$  that applies on the given data to obtain our desired solution. However, this family of functions  $\mathcal{F}$  is usually not defined in a straightforward manner and cannot be used directly. Therefore, across the set of all possible functions, we usually restrain the algorithm to consider only a given family of functions  $\mathcal{F}^*$ , that we define, hoping to be as close to  $\mathcal{F}$  as possible. Finally, to approximate the ideal solution, we select a parametric function  $f_\theta \in \mathcal{F}^*$ , in order to obtain an estimation

$$f_\theta(x_i) = \tilde{y}_i, \forall i \in [1, N]$$

by modifying the set of parameters  $\theta \in \Theta$  of this function, the algorithm learns to predict a set of solutions that should be as close as the real solutions as possible.

2.1.2 *Loss function*

Given this parametric function  $f_\theta$ , the goal of machine learning is to find the parameters that can provide the highest accuracy possible.

The parameters  $\theta$  are defined as a multi-dimensional vector containing the values of each parameter of the function. In order to find this vector, we need to define a loss function  $\mathcal{L}$ , that quantifies the difference between the correct solution and the current approximation

$$L(\mathbf{x} | \Theta) = \sum_{i=1}^N \mathcal{L}(f(x^i), \tilde{f}_{\Theta}(x^i)) \quad (2)$$

The goal will be, therefore, to minimize this function by adjusting iteratively the parameters of our function  $f_{\theta}$ . This error minimization process evaluates a given set of training examples and then performs the parameters update based on the derivative of the error.

### 2.1.3 Types of learning algorithms

The different approaches to machine learning can be broadly divided into three major types : supervised, unsupervised and reinforcement learning. The choice between these approaches depends principally on the information that we have on the data and the output that we seek for a given task. Here, we focus on detailing the supervised [2.1.3.1](#) and unsupervised [2.1.3.2](#) approaches as they will both be used in our subsequent methods, in the form of semi-supervised learning [2.1.3.3](#).

#### 2.1.3.1 Supervised learning

The supervised learning approach aims at classifying an unknown dataset with the help of a labeled dataset. The learning is done on a dataset  $\{x(i), a(i)\}_{1 \leq i \leq N}$  containing  $N$  samples. Here,  $x(i)$  is an input data and  $a(i)$  its associated label. We also define an alphabet  $C$  that contains all the different labels. In this way, we want the system to extract enough information from  $x(i)$  to classify it in the right class  $c \in C$ .

For instance, in the field of ACE, a potential supervised task would be to use a frame of the spectrogram  $x(t)$  as an input and develop a model that could find the associated chord label  $a(t)$ . Thus, we need a training set containing associations between frame and chord label (an example dataset is depicted in Figure 2).

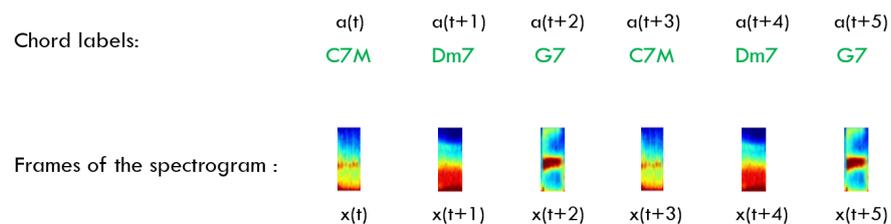


Figure 2: Example of elements inside a chord labeling dataset based on spectrogram windows.

For training an ACE model that could perform chord labeling, a potential choice of loss function is the distance between the predicted chord labels and the annotated chord labels. Then, training the model can be done as shown in Figure 3. Here, by following our previously defined notation (2), we have the desired output of the classification task given by  $f(x(t)) = a(t)$  and the approximation of our model  $\tilde{f}_{\Theta}(x(t)) = y(t)$ , which is here composed of several layers of transformations. By comparing these two values through the loss function  $\mathcal{L}(f(x(t)), \tilde{f}_{\Theta}(x(t)))$ , we can assess the errors made by the model in order to improve it by changing its parameters.

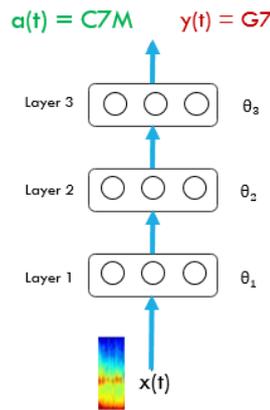


Figure 3: An example of supervised training for an ACE task, where we compare the ground truth annotation  $a(t)$  with the approximation of our model  $\tilde{f}_{\Theta}(x(t)) = y(t)$ .

### 2.1.3.2 Unsupervised learning

Most of the time the data that are available for a task are not labeled. Thus, we only have a dataset composed by  $M$  elements  $\{x(j)\}_{1 \leq j \leq M}$  and want to discover the structure of those data. The most well-known type of unsupervised learning is clustering algorithms. It is used for exploratory data analysis to find patterns or groupings inside a dataset. In other words, we extract the most salient features from the inputs. Then, we cluster the data by applying a distance between these features. Since we don't inform the model on which features we are focusing on, the unsupervised learning works without any assumptions.

Once again, in the context of ACE, an example task would be to present a noisy spectrogram frame  $x^*(t)$  as input in order to denoise it. This task is known as a *denoising* operation, as depicted in Figure 4. Here, one part serves to transform the input into a representation, which could then be used to reconstruct a noiseless version of the frame  $x^*(t)$ . The cost function is defined by the distance between the output  $z(t)$  and the initial frame  $x(t)$ .

Despite the fact that supervised learning has demonstrated impressive results in the past years, unsupervised learning would be an approach widely studied in the next years.

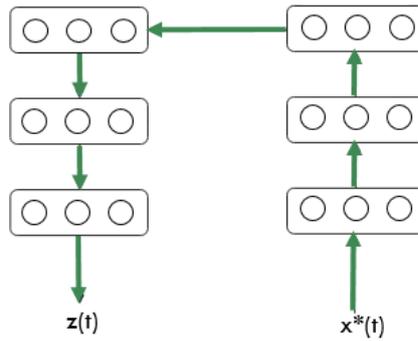


Figure 4: An example of unsupervised training for an ACE task, where we compare a frame of spectrogram  $x(t)$  with a denoising corrupt version of it  $\hat{f}_\Theta(x^*(t)) = z(t)$ .

“We expect unsupervised learning to become far more important in the longer term. Human and animal learning is largely unsupervised: we discover the structure of the world by observing it, not by being told the name of every object.”

– LeCun, Bengio, Hinton, Nature 2015 [29]

### 2.1.3.3 Semi-supervised learning

In most applications, we usually have access to few labeled data  $\{x_t, y_t\}_{1 \leq t \leq N}$  and, comparatively, to a lot of unlabeled data  $\{x_t\}_{N+1 \leq t \leq M}$  (i.e. music tracks without annotations). Consequently, the labeled data is scarce whereas the unlabeled data is plentiful ( $N \ll M$ ). The main idea is to assign probabilistic label to unlabeled data in order to use them in the training. By modeling  $P(x|y)$  as clusters, unlabeled data affects the shape and size of clusters [58]. Technically, we use the clustering effect of the unsupervised learning and assign a label at each cluster with the help of the labeled dataset.

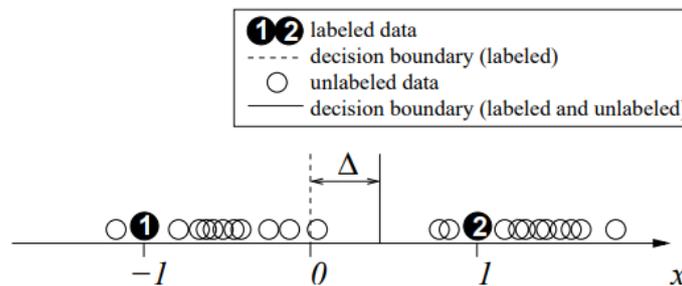


Figure 5: Example of a two-class cluster classification dataset where two sample are labeled and others are unlabeled. The addition of unlabeled data modifies the cluster sizes and shapes. [58].

In Figure 5, we represent a two-class cluster classification dataset where two sample are labeled and others are unlabeled. Thus, the addition of unlabeled data shifts the decision boundary and gives a most accurate classification on unseen data.

Nevertheless, the semi-supervised learning does not always give better results and the error convergence depends on the learning task [52].

#### 2.1.4 Numerical optimization and training

As discussed previously, we need to find an efficient way to adjust the parameters of a model, in order to increase its accuracy. The accuracy is calculated by the error function. One solution for an efficient training is to use its derivative depending on the parameters. Hence, at each training iteration we find which direction to explore, leading to the family of *gradient descent* algorithms.

##### 2.1.4.1 Gradient descent and learning rate

THE GRADIENT DESCENT is an iterative optimization algorithm aimed at finding a local minimum of an error function. Each step consists of moving the set of parameters in a direction that maximally lowers the value of this loss function. The gradient descent is therefore defined as starting with a random set of parameters  $\Theta_0$  and then repeatedly picking the direction of steepest descent (with respect to the loss function) in order to update the parameters

$$\Theta_{t+1} \leftarrow \Theta_t - \alpha * \nabla_{\Theta} L(\Theta_t) \quad (3)$$

where  $\nabla_{\Theta} L(\Theta_t)$  represents the gradient of the loss function with respect to the parameters and  $\alpha$  is an *hyperparameter* called the *learning rate*. This hyperparameter allows to control the magnitude of the step that we are taking at each parameter update. A bigger learning rate will accelerate the convergence but can cause to "jump" over the minima as the update steps are too large. These ideas are exemplified in Figure 6. This figure represents the shape of a two parameter loss function  $J(\theta_0, \theta_1)$ . On the figure we see two paths (in black). The starting points of the paths (in the red area) are set by a random initialization of the model parameters. Then, each path is constructed by choosing the best gradient descent at each step. They naturally end in a local minima (in blue).

Over the past decades, the cardinality and dimensionality of datasets have outpaced the expansion of the speed of processors. As argued by Bottou [7], the capabilities of machine learning has been limited by the computing time rather than the sample size. Hence, stochastic approximation of the gradient optimization method as Stochastic Gradient Descent (SGD) have been proposed [7], which update the parameters of the model after observing each example, rather than relying on the whole dataset.

##### 2.1.4.2 Overfitting, Generalization

The goal of learning methods is to obtain a model which provides a good *generalization* over unseen data. The generalization performance

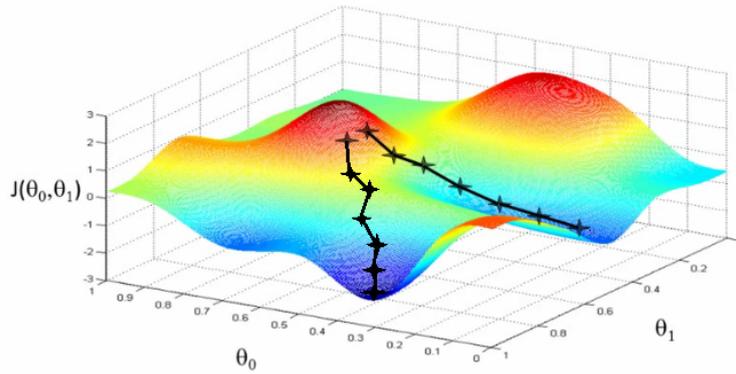


Figure 6: Gradient of a two parameters function with two local minima and example of the use of gradient descent based on two different random initializations, image taken from [35].

of a method refers to its capacity to make good predictions for independent (unseen) test data. Thus, we hope that the training set will reflect the characteristics of the entire sample space.

The overfitting appears when a model attached too many importance to the details of the training set. For instance, we show in Figure 7 the effect of overfitting for a task of classification between two classes inside a two-dimensional feature space. The black line represents a model with a good generalization performance whereas the green line shows a model displaying overfitting.

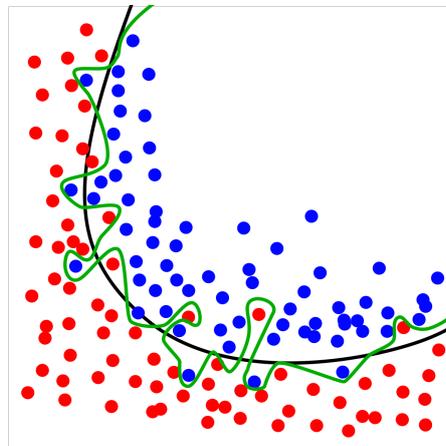


Figure 7: Two different models showing either overfitting (in green) or generalization (in black) for separation in a two-dimensional feature space with two classes, image taken from [40].

The overfitting can come from the complexity of the model. Indeed, if a model is excessively complex, it has too many parameters comparing to the number of observations. Among other solutions we can avoid overfitting and obtain better generalization by corrupting slightly the inputs. The overfitting also depends of the time of training. Thus, when we use learning methods, we have to stop the optimization process in order to avoid this effect. This optimization step can be seen as finding a tradeoff between the data that we use

in our training set and the general behavior of future unseen data. One widely used method to avoid overfitting is to validate the model with a dataset which is independent from the training and the test set. Therefore, we divide the data into three subsets.

- **THE TRAINING SET** is used to train the model. It contains the sample from which we adjust the parameters.
- **THE VALIDATION SET** is used to avoid overfitting. As seen in Figure 8, the validation error decreases and then increases during the training process. We must stop the training when the validation error is the lower in order to have the better generalization.
- **THE TEST SET** is the data set on which the performance of the model is calculated.

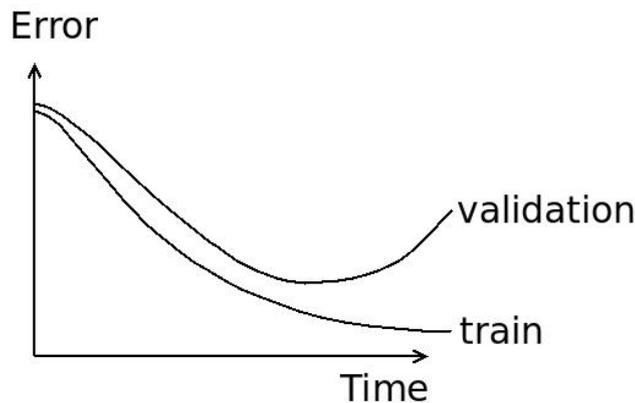


Figure 8: Evolution of train and valid errors versus training time, image taken from [40].

## 2.2 NEURAL NETWORKS

### 2.2.1 Artificial neuron

An artificial neuron can be compared to a biological neuron. In biological neurons, the dendrites are connected to a cell body. Each dendrite transmits electrical impulsion from the outside to the body. When the electrical charge exceeds a threshold the cell body sends an electrical impulsion through the axon. In an artificial neuron such as Figure 9, the dendrites are modeled by the inputs  $\{x_i\}_{1 \leq i \leq m}$ , the activation is a non-linear function  $f$  with a threshold  $b$  and the axon is the output  $h$ .

A neuron is then defined by its parameters, namely the bias  $b$  and the weights  $w_i$ , and by an activation function  $f$ . We have then a mathematic formalization that links the inputs  $x_i$  to the output  $h$  :

$$h = f\left(\sum_{j=1}^m w_j x_j + b\right) \quad (4)$$

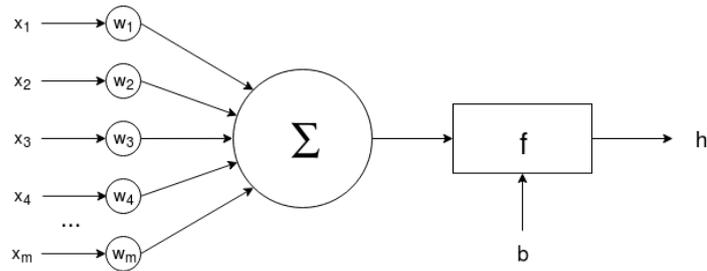
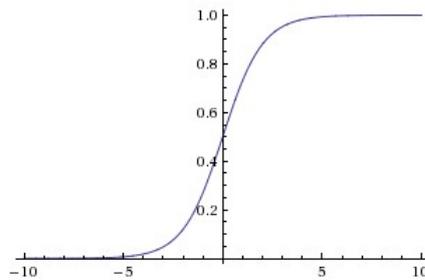


Figure 9: Schema of an artificial neuron.

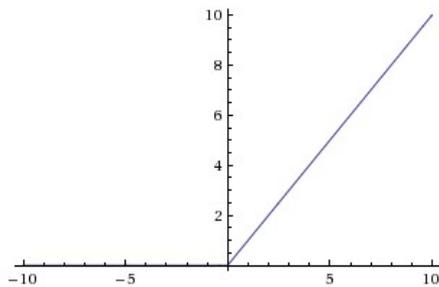
In the literature we find many activation function  $f$ . The activation should be a non-linear transformation. Here we present two examples of widely used activation function.

**THE SIGMOÏD FUNCTION** is a function that has been widely used in the field of machine learning. His derivative is very simple, making easier the chain-rule computations of the back propagation process (see back propagation 2.2.2.1).



$$f(x) = \frac{1}{1 + \exp(-x)}$$

**RELU** that stands for rectified linear unit, has been introduced in 2000 with strong biological and mathematical motivation. As of 2015 this is the most popular activation function used in deep neural architecture.



$$f(x) = \max(0, x)$$

### 2.2.1.1 Interpretation

Here we provide two ways of interpreting the Neural Networks (NN). First, by looking at Equation 4, we can see that if the transfer function is a threshold function, this defines an hyperplane. On the one hand, network with a single processing layer can be interpreted as separating the input space in two regions with a straight line (as seen in Figure 10). Here, the blue and red lines represents two data distributions, while the colored region represent the value of the output

neuron. Therefore, we can note that a single-layered network can not separate properly these two distributions (as this dataset is *non linearly separable*). Hence, we must use a non-linear transfer function for this classification task. On the other hand, a way of interpreting these networks is to see the successive layers as performing *space transformation* operations. As we can see in Figure 11, the space is gradually transformed by the successive neuron layers, as the coordinates of the next space is defined by the output values of the different neurons. Hence, for complex data we must use neural networks with multiple layers, to exploit this property of *compositionality*.

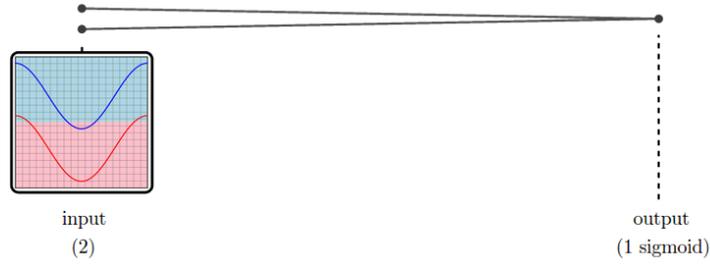


Figure 10: Space separation with one linear output layer. Image from [3].

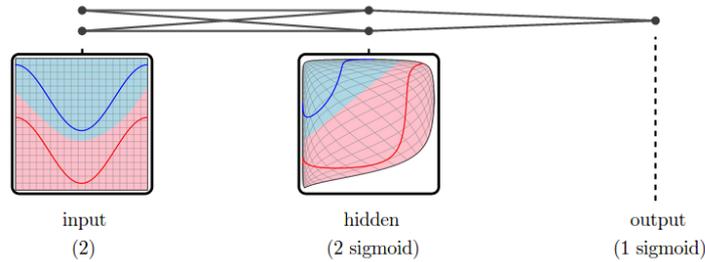


Figure 11: Space transform operated by successive hidden layer. Image from [3].

### 2.2.2 Multi-Layer Perceptron

The *Multi-Layer Perceptron* (MLP) is a feed-forward model organized as a succession of layers of neurons defined previously. Each layer receives the output of neurons in the previous layer as input and then applies a (linear) affine transform and a non-linear transfer function to these values. Therefore, a MLP is a fully-connected network of depth  $L$ . Here, we denote the output value of the  $l$ -th layer by vector  $h^l \in \mathbb{R}^{M_l}$ , where  $M_l$  is the number of neurons contained in this layer (see Figure 12). For  $1 \leq l \leq L$ , the parameters of a layer are defined by a weight matrix  $W^l \in \mathbb{R}^{M_l \times M_{l-1}}$  and a bias vector  $b^l \in \mathbb{R}^{M_l}$ . Therefore, the activation of neuron  $m$  in layer  $l$  is computed with the following equation

$$h_m^l = f\left(\sum_{j=1}^{M_{l-1}} (W_{m,j}^l \cdot h_j^{l-1}) + b_m^l\right) \quad (5)$$

Let us rewrite the equation 5 such as  $b_m^l = W_{m,0}^l \cdot h_0^{l-1}$ .

$$h_m^l = f\left(\sum_{j=0}^{M_{l-1}} (W_{m,j}^l \cdot h_j^{l-1})\right) = f(y_m^l) \quad (6)$$

Thus,  $y_m^l$  is the net input for the neuron  $m$  in the layer  $l$ .

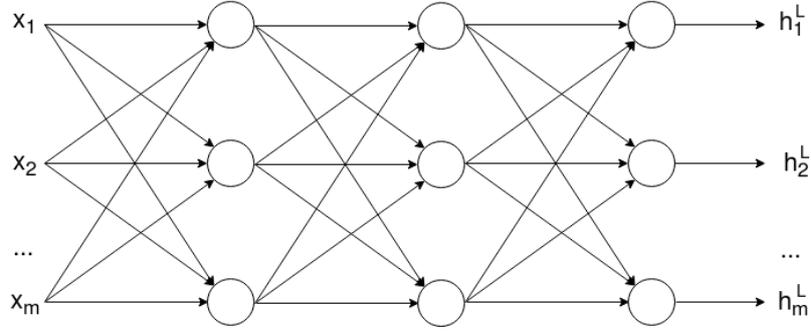


Figure 12: Multi-Layer Perceptron.

### 2.2.2.1 Back propagation

In order to train the MLP, it would seem that we should define a procedure that takes all layers into account, leading to a very complicated formulation. However, the introduction of the *back-propagation* method by Rumelhart and Hinton [56] allows to decompose this into a set of simple operations. The main idea is to see that adding layers amounts to add a function to the previous output. Therefore, in the *forward pass*, this output is updated by relying on the weight and bias of the neurons in this layer. Therefore, when trying to obtain the contribution of a given neuron to the final error value, we can use the *chain rule* of derivations, to separate its contribution inside the network. The error values obtained after a forward pass can be simply propagated backwards, starting from the output, and computing the derivative of each neuron output given its parameters. This procedure is repeated until all neurons in the network have been updated.

After the last layer  $L$ , we compute the loss of our network (see Section 2.1.2), that we take here to be the squared error between the desired output  $\delta_i^L$  and the output of the network  $h_i^L$  :

$$\mathcal{L}(\delta_i^L, h_i^L) = \frac{1}{2} \|\delta_i^L - h_i^L\|^2 \quad (7)$$

The derivative of this loss is then given by

$$\mathcal{L}(\delta_i^L, h_i^L)' = e_i^L = [\delta_i^L - h_i^L] \cdot f'(y_i^L) \quad (8)$$

Using the chain rule, the error is then back propagated by computing error for preceding layers :

$$e_i^l = f'(y_i^l) \sum_{j=0}^{M_{l+1}} W_{ij}^{l+1} \cdot e_j^{l+1} \quad ; \quad l \in [1 \dots L-1] \quad (9)$$

Finally, the weights are updated in the corresponding layer :

$$W_{ij}^l = W_{ij}^l + \lambda \cdot e_i^l h_j^{l-1} \quad (10)$$

Where  $\lambda$  is the learning rate.

### 2.2.2.2 Limitations and deep learning solutions

As seen in the previous part, a MLP is composed of layers of artificial neurons as a simple model that tries to mimic biological neurons. Hence, simple neural networks are composed by a few layer of neurons to reduce their computational costs. However, an architecture with such insufficient depth would require an infinite number of computational elements [5] to solve complex tasks. Conversely, deep architectures can exploit non-linear hierarchies to solve this issue. Thus, we use deep learning when we need higher level of abstractions for our problem. This approach allows to construct optimal abstractions by discriminating the important variations at each layer by focusing on information at different scales. Nevertheless, adding several layers leads to a gradient diffusion problem during the back-propagation. It can be solved by a layer pre-training [4].

Deep learning usually relies on an encoder/decoder (unsupervised) model that aims to deconstruct an object and reconstruct it from fewer parts. In other words we are learning its structure. This notion of encoder/decoder seems fit to our musical structure problem. Indeed, each layer can be seen as a different abstraction (notes, chords, keys).

### 2.2.3 Convolutional network

A Convolutional Neural Network (CNN) is a specific type of feed-forward neural network that is currently amongst the best performing systems for image processing tasks [13, 28]. The architecture of CNNs were inspired by the organization of the animal visual cortex [30]. The main idea of CNN is to introduce invariance properties (such as translation, rotation, perspective distortion) into the transform computed by neural networks. Compared to the MLP, the hidden layers of a CNN compute different operations, by relying on convolution and pooling operators. In the next parts, we describe these operators and the overall behavior of CNNs.

#### 2.2.3.1 Convolutional Layers

A convolutional layer is defined by a set of convolution kernels that are applied in parallel to the inputs and produce a set of output *feature maps*. These kernels are learned during the training and hopefully each of these should describe a recurrent characteristic of the training data. After applying each kernel to the input, we obtain a set of feature maps, that is defined by a three-dimensional tensor  $h \in \mathbb{R}^{M \times I \times J}$  where  $M$  is the number of kernel,  $I$  is the height and  $J$  the width of

each kernel. During the feed-forward pass we apply a convolution by sliding the kernel over every local area of the input.

For instance, as seen in Figure 13, we start from the top left region and move the kernel until it reaches the bottom border of the input matrix. If we define the input as the matrix  $X$ , then the output feature maps are defined by  $Y = X * h_m$  for every kernels. Here, the operator  $*$  is a 2D discrete convolution

$$(A * B)_{i,j} = \sum_{r=1}^T \sum_{s=1}^F A_{r,s} B_{r+i-1,s+j-1} \tag{11}$$

for  $A \in \mathbb{R}^{T \times F}$  and  $B \in \mathbb{R}^{I \times J}$  with  $1 \leq T \leq I - 1$  and  $1 \leq F \leq J - 1$

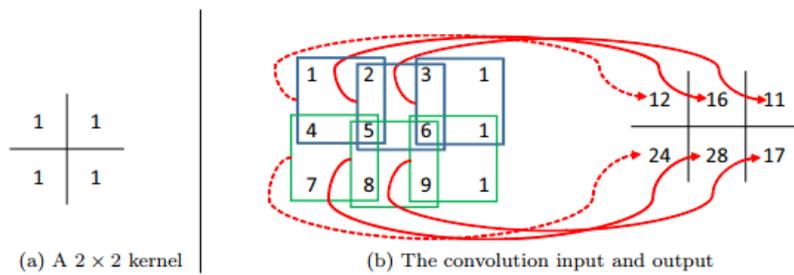


Figure 13: Illustration of the convolution operation, image taken from [57].

Here, we note that depending on the size of the kernel, the dimension of the output will be smaller than that of the input. Therefore, to keep the dimensions constant, we can use a zero-padding operation which consists to add zeros to the border of the original input. Furthermore, the norm is to apply an activation function after each convolutional layer.

### 2.2.3.2 Pooling Layer

Each convolutional layer significantly increases the dimensionality of the input data. Therefore, a pooling layer is often placed between convolutional layers in order to reduce the size of the feature maps. This downsampling layer can either perform an average pooling, L2-norm pooling or max pooling. For instance, the max pooling (depicted in Figure 15) operation only keeps the maximum value in each regions of a partition of the input. The pooling size refers to the factor to which the map is reduced.

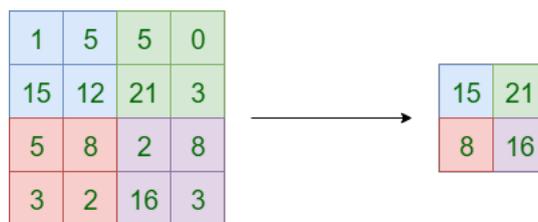


Figure 14: Example of a 2x2 max pooling.

2.2.3.3 Fully-connected layer

The overall structure of a convolutional network usually consists of a succession of convolution, activation and pooling layers. However, in order to then perform correlation between the different feature maps, this architecture is typically followed by one or many fully-connected layers (see Figure 15). This fully-connected layer is a standard MLP, where each neuron in layer  $l$  is connected to all neurons in layer  $l + 1$ .

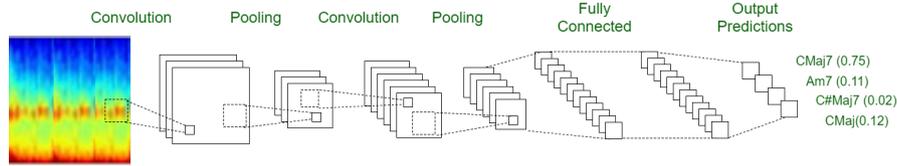


Figure 15: Example of a convolutional neural network for a task of ACE.

2.2.4 Ladder Networks

Recently, the *ladder network* (LN) has been proposed to allow a combination of supervised and unsupervised learning [47]. In this approach, we use an auxiliary unsupervised learning task (here denoising) to support supervised learning. This model has proved to give very satisfying results for the hand-written digit classification task [48]. However, as far as we know, the ladder network has not yet been applied to audio classification tasks. In this section we present the ladder network architecture.

2.2.4.1 Lateral connections and noise injection

One of the particularity of the LNs comes from the lateral connections of the encoder/decoder. As seen in Figure 16, each layer of the encoder is laterally connected to the decoder. These connections aim at relieving the pressure of higher layers. Thus, each layer concentrates more efficiently on their degree of abstraction.

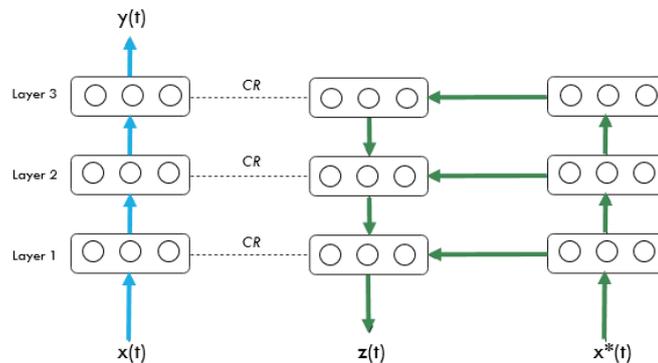


Figure 16: Semi-supervised learning on ACE task, the supervised classification (in blue) is reinforced by the unsupervised denoising autoencoder task (in green)

The lateral connection (from the layer  $l$  of the encoder to the layer  $l$  of the decoder) and the vertical connection (from the layer  $l + 1$  of the

decoder to the layer  $l$  of the decoder) are combined by a *combination function* before to enter in the layer  $l$  of the decoder. This combination function has a non-negligible impact on the overall performance of the LN. In [46], Pezeshki et al. improve gradually the combination function from a simple linear combination to a function enhanced by multiplicative terms and sigmoid non-linearity functions.

However, the choice of the combinator function seems less crucial than the type of noise introduced in the denoising autoencoder. Indeed, the introduction of noise into each layer of the encoder has a strong regularization effect that helps for the generalization. Nevertheless, the most important aspect of the LN lies in its semi-supervised training.

### 2.2.5 Reconstruction cost and loss function

The LN can be trained to minimize simultaneously the sum of supervised (Cross Entropy) and unsupervised ( $L_2$  reconstruction) cost functions. We consider the dataset with  $N$  labeled examples  $\{x(t), y(t)\}_{1 \leq t \leq N}$  and  $M$  unlabeled examples  $\{x(t)\}_{N+1 \leq t \leq M}$ . The output of  $x(t)$  performed by the supervised part of the network is written  $y(t)$ . The cost function is then defined by

$$\text{Cost} = - \sum_{t=1}^N \log P(y(t) = a(t)|x(t)) + \sum_{t=N+1}^M \sum_{l=1}^L \lambda_l \text{ReconsCost}(z^{(l)}(t), \hat{z}^{(l)}(t)) \quad (12)$$

Where  $L$  is the total number of layers,  $\hat{z}^{(l)}$  and  $z^{(l)}$  are respectively the outputs of the supervised and unsupervised path at the layer  $l$ .  $\lambda_l$  define the contribution of each layer by their reconstruction cost. Moreover, these reconstructions costs are the most important properties of the LD and provide the desired regularization from unlabeled data.

Finally, each layer of the LN, initially classic MLPs, can be replaced by more specific ones such as CNNs.

## 2.3 TEMPORAL STRUCTURE AND MEMORY

### 2.3.1 Recurrent Neural Networks

One of the largest issue in NN is its incapacity to store temporal information. However, in our setup, we need to discover the underlying structure of a music signal. Therefore, we need to design our model with an understanding of time. Recurrent Neural Networks (RNN) show promising results in many learning fields [11] including musical structure generation [16]. To define recurrent networks, we usually rely on loops inside the network, where the output of a neuron is directly linked to itself for the next sequence element. The RNN can also be interpreted as a neural networks with lateral connections as seen in Figure 17

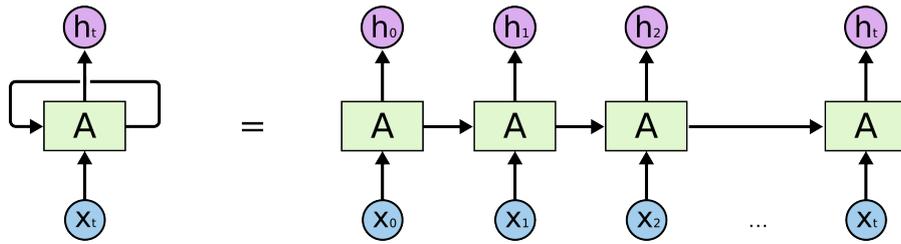


Figure 17: Loops of a RNN unfolded through time, image taken from [54].

The inputs  $x_i$  are processed by the RNN and it outputs a value  $h_i$ . The transformation is operated by a module  $A(x_i, h_{i-1})$  that takes in arguments the input  $x_i$  and the output of the previous step  $h_{i-1}$ .

When have then a mathematical formalism for the RNN :

$$h_i = A(x_i, h_{i-1}) \quad (13)$$

In a standard RNN, the function  $A$  has a very simple structure such as a linear layer followed by a tanh layer.

The training of a RNN is realized by the Backpropagation Through Time (BPTT). This algorithm works by fixing the amount of time steps, unrolling the RNN in order to obtain a classic NN, and training it with a standard backpropagation algorithm.

Nevertheless, standard RNNs tend to forget things quickly along the time steps. Gated memory mechanisms solved this problem by adding to standard RNNs few trainable gates that allow to select, stock and transport the most important information for the task at end. The two most used gated RNN structures are the Long Short-Term Memory (LSTM) and the Gated Recurrent Unit (GRU).

### 2.3.2 LSTM

Generic RNN are limited in there capacity to learn long-term dependencies and are hard to train [42]. An elegant variant, called LSTM, has been introduced by Hochreiter and Schmidhuber [21]. It provides good results in many use case [53].

Similarly to classic RNN, LSTM network performs a transformation of an input  $x_i$  by taking into account the transformations of the previous time steps. The difference between the two cells lies in the structure of the module  $A$ . In classic recurrent neuron, the output of the previous step  $h_{i-1}$  is send through a tanh layer (see Figure 18). The reminiscence of information between distant time steps is then relatively poor.

In LSTMs, an additional internal connection is present (see Figure 19). This connection carries information through time steps with only some minor linear interactions. Thus, the information between distant time steps could be unchanged.

In addition to this memory cell, the LSTM is composed by three gates which control how to update and used the memory cell. The gates are principally the combination of a sigmoid layer followed by a point-wise multiplication layer.

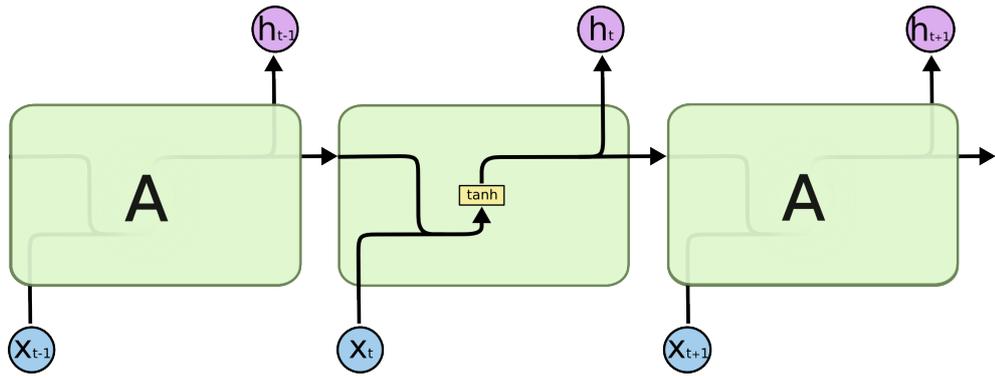


Figure 18: Internal structure of a classical RNN cell, image taken from [54].

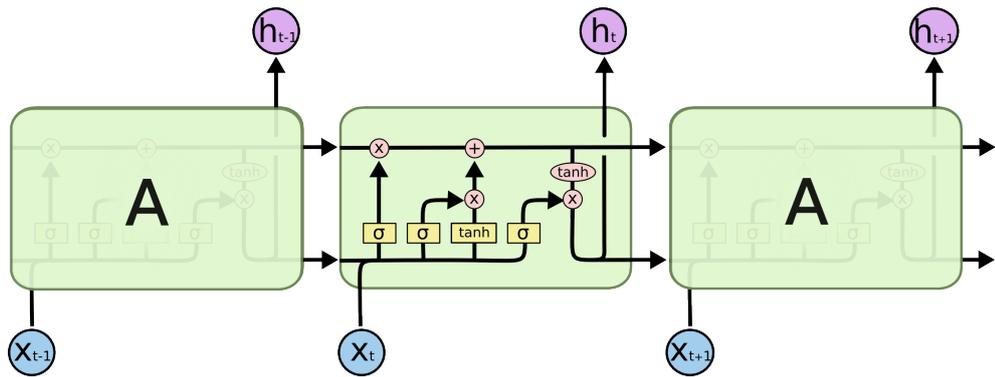


Figure 19: Internal structure of a LSTM cell, image taken from [54].

1. The forget gate controls which part of the memory cell to forget. It looks at  $h_{i-1}$  and  $x_i$ , select which dimension we have to keep or forget and then update the memory cell.
2. The update gate decides which values to update, then compute a new memory vector and add it to the previous vector.
3. The output gate controls which memory dimension must be output to the next time step.

Furthermore, a LSTM network is composed by layer of LSTM units. This construction allows to detect patterns at different abstraction levels. This hierarchy vision is reminiscent of the structure hierarchy in an audio track. Thus, LSTM network could operate at different time scales.

## EXTRACTING CHORD LABELS FROM AUDIO SIGNAL

---

### 3.1 CHORD PROGRESSION EXTRACTION

#### 3.1.1 *Definition of the workflow*

Across the several tasks in the field of Music Information Retrieval (MIR), the Automatic Chord Estimation (ACE) from audio signal is a topic that has been widely studied over the past years [33]. Throughout the literature, chord recognition systems are usually designed by following the same workflow.

1. First of all, the audio signal is transformed into a time-frequency representation. In the chord classification task, the Short-Term Fourier Transform (STFT) or a logarithmic representation such as the Constant-Q Transform (CQT) are often used.
2. Secondly, features are extracted and the signal is represented by a *chromagram* (a matrix with 12 rows representing each pitch class, and N columns corresponding to the number of time frames of the audio track).
3. At the end, each time frame (or related chromagram frame) is assigned by the system to a chord label. A post-filtering step, which is also called *chord sequence decoding* is often performed to smooth out the sequence of outputs.

#### 3.1.2 *Related works in Automatic Chord Extraction*

In the seminal paper by Wakefield in 1999 [55], the first mathematical grounds for the definition of *chromagram* opened up its use for the ACE field. Consequently, Hidden Markov Models (HMM) have been amongst the first models investigated for chord extraction [51] as post-filtering methods. Following these works, multiple variants of this workflow were explored, either where other features have been exploited such as the tonal centroid [18], or replacing the classification methods with probabilistic N-Grams [50] or Dynamic Bayesian Networks (DBN) [32].

In the recent years, several works relied on neural networks for ACE. For instance, Humphrey and Belo [23] applied CNN to classify major and minor chords. This new approach has raised a different perspective on the problem, where a context (which takes several seconds of pitch spectra) is classified directly by a CNN, without relying on the explicit definition and computation of sound features beforehand.

Since then, the use of deep learning in the field of chord extraction has been explored in different parts of the traditional workflow. Boulanger-Lewandowsky et al. [8] trained a RNN as a post filtering method, over a feature extraction pipeline learned by a Restricted Boltzmann Machine (RBM), making this whole approach entirely based on learning methods.

The previously cited approaches were trained to distinguish between major and minor chords only. However, based on the excellent results provided by learning methods, following works considered other chord types (seventh, augmented, or suspended) mapped to major/minor as a novel classification problem. Hence, the Deep Chroma Extractor system was proposed [27] to target this problem, which aims to learn how to extract a chroma representation from the audio signal with deep networks. Following this work, the same authors proposed a method combining a CNN for feature extraction with a Conditional Random Field (CRF) for chord sequence decoding [26]. The CNN is trained to predict chord labels for each audio frame, which provides hidden representations that are used as features for the subsequent pattern matching and chord sequence decoding stage performed by the CRF.

#### 3.1.2.1 MIREX

In order to perform a systematic and objective evaluation of the robustness of the proposed approaches, ACE systems are compared in an annual evaluation called the Music Information Retrieval Evaluation eXchange (MIREX). For the ACE task, the MIREX chord label dictionary is composed of five different alphabets.

1. Chord root note only;
2. Major and minor N (no chord), maj, min;
3. Seventh chords N, maj, min, maj7, min7, 7;
4. Major and minor with inversions N, maj, min, maj/3, min/b3, maj/5, min/5;
5. Seventh chords with inversions N, maj, min, maj7, min7, 7, maj/3, min/b3, maj7/3, min7/b3, 7/3, maj/5, min/5, maj7/5, min7/5, 7/5, maj7/7, min7/b7, 7/b7.

In order to perform these comparisons, MIREX evaluates the quality of an automatic extraction by comparing the results to a ground truth created by one or more human annotators. The Chord Symbol Recall (CSR) measure is typically used to estimate how well the predicted chords match with the ground-truth.

$$\text{CSR} = \frac{\text{total duration of segments where (annotation = estimation)}}{\text{total duration of annotated segments}} \quad (14)$$

We present in Table 1 the methods proposed for the ACE task in MIREX 2016 as we will be using their results (detailed in Table 2) on the same datasets as a baseline in order to evaluate the success of our proposed methods (see Chapter 4).

Algorithm	Contributors
CM1 (Chordino)	Chris Cannam, Matthias Mauch
DK1-DK4	Junqi Deng, Yu-Kwong Kwok
FK2, FK4	Filip Korzeniowski
KO1 (shineChords)	Maksim Khadkevich, Maurizio Omologo

Table 1: Contributors for MIREX 2016 in ACE, for more information see MIREX website [31]

Algorithm	1	2	3	4	5
CM1	78.56	75.41	72.48	54.67	52.26
DK1	79.21	76.19	74.00	66.02	64.15
DK2	77.84	74.49	71.93	61.61	59.47
DK3	80.03	77.55	74.79	68.40	65.88
DK4	76.05	72.96	71.41	62.77	61.44
FK2	86.09	85.53	82.24	74.42	71.54
FK4	82.28	80.93	78.03	70.91	68.26
KO1	82.93	82.19	79.61	76.04	73.43

Table 2: ACE on Isophonics 2009 for MIREX 2016

### 3.2 DATASETS AND AUDIO PRE-PROCESSING

We introduce in this section the datasets and the audio pre-processing that will be used in the evaluation, based on the MIREX evaluation guidelines.

#### 3.2.1 *Isophonics datasets*

The Isophonics dataset is the major dataset for ACE over the past years. It has been collected by the Centre for Digital Music at Queen Mary of the University of London and is composed by 5 subsets. The subsets are from The Beatles, Carole King, Queen, Michael Jackson and Zweieck. The sum of all the subsets represents 300 MP3 files at a sampling rate of 44100Hz with annotations on chords, keys, beats and structural segmentations. Nevertheless, the different subsets that composed Isophonics have not the same level of confidence. Besides all of them the Beatles’s dataset (180 songs) created by Christopher Harte [17] is the one with the highest level of confidence. The annotations of the chords follows the syntax described on the Figure 20

<code>&lt;chord&gt;</code>	<code>::=</code>	<code>&lt;pitchname&gt; ":" &lt;shorthand&gt; [{"&lt;ilist&gt;"}]["/"&lt;interval&gt;]</code> <code>  &lt;pitchname&gt; ":" "("&lt;ilist&gt;)" ["/"&lt;interval&gt;]</code> <code>  &lt;pitchname&gt; ["/"&lt;interval&gt;]</code> <code>  "N"</code>
<code>&lt;pitchname&gt;</code>	<code>::=</code>	<code>&lt;natural&gt;   &lt;pitchname&gt; &lt;modifier&gt;</code>
<code>&lt;natural&gt;</code>	<code>::=</code>	<code>"A"   "B"   "C"   "D"   "E"   "F"   "G"</code>
<code>&lt;modifier&gt;</code>	<code>::=</code>	<code>"b"   "#"</code>
<code>&lt;ilist&gt;</code>	<code>::=</code>	<code>["*"] &lt;interval&gt; ["," &lt;ilist&gt;]</code>
<code>&lt;interval&gt;</code>	<code>::=</code>	<code>&lt;degree&gt;   &lt;modifier&gt; &lt;interval&gt;</code>
<code>&lt;degree&gt;</code>	<code>::=</code>	<code>&lt;digit&gt;   &lt;digit&gt; &lt;degree&gt;   &lt;degree&gt; "0"</code>
<code>&lt;digit&gt;</code>	<code>::=</code>	<code>"1"   "2"   "3"   "4"   "5"   "6"   "7"   "8"   "9"</code>
<code>&lt;shorthand&gt;</code>	<code>::=</code>	<code>"maj"   "min"   "dim"   "aug"   "maj7"   "min7"   "7"</code> <code>  "dim7"   "hdim7"   "minmaj7"   "maj6"   "min6"   "9"</code> <code>  "maj9"   "min9"   "sus2"   "sus4"</code>

Figure 20: Syntax of Chord Notation in Isophonics, image taken from [17].

### 3.2.2 Audio transformation

In order to obtain an efficient representation of input signals, we need to compute a time-frequency transform of the musical tracks. The analyzed data is then a two-dimensional matrix  $X[m, k]$  where  $k$  is a bin centered around a frequency and  $m$  is the temporal index of the frame. However, this representation might remain computationally intensive for learning. Therefore, a dimensionality reduction can be performed by reducing the frequency resolution or temporal quantification.

In this section, we discuss our choices regarding spectral transforms and also introduce the audio data transposition which allows us to augment the training dataset with transposed chords.

#### 3.2.2.1 Short Term Fourier Transform

The Short-Term Fourier Transform (STFT) is a spectral transform which is used to extract the sinusoidal frequency and phase content of local sections of a signal. In the STFT, we divide the pressure signal, into shorter segments of equal length (by performing a *windowing* operation) and then compute the Fourier transform separately on each segment.

$$X_{\text{STFT}}[m, k] = \sum_{n=0}^{N-1} x[n] \cdot w[n-m] \cdot e^{-i \frac{2\pi k}{N} n} \quad (15)$$

where  $x[m]$  is the  $m$ -th sample of the signal and  $w[n]$  is the window function, commonly a Hann or Gaussian window centered around zero. The variation of the window length  $N$  is a trade-off between time and frequency resolution.

### 3.2.2.2 Constant-Q Transform

The STFT gives a linear resolution of the frequency bands along a definite range. However, the pitches in western music are rather based on a logarithmic scale organization along the spectrum of audible frequency. Thus, the STFT might not be the most optimized representation for the study of musical chords. The Constant-Q Transform (CQT) is also a spectral transform of audio signals, akin to the Fourier transforms. However, this transform can be thought of as a series of logarithmically spaced filters. Thus, in the CQT, the Q value, which is the ratio between the central  $f_k$  frequency to the bandwidth  $\delta f_k$  is constant :

$$Q = \frac{f_k}{\delta f_k} \quad (16)$$

The windows length N becomes a function of the bin number  $N[k]$  and the windowing function becomes a function of the window length :

$$X_{\text{CQT}}[m, k] = \frac{1}{N[k]} \sum_{n=0}^{N[k]-1} x[n].W[k, n - m].e^{-i\frac{2\pi Q}{N[k]}n} \quad (17)$$

Therefore, an appropriate choice in the bank of filters allows us to have a correspondence between the calculated bins and the music notes. This property will be very helpful to simplify our implementation of the audio transposition for augmenting the training dataset.

### 3.2.2.3 Audio transposition

In the CQT we define a number of bin per octave. Hence, we can artificially transpose our signal with a translation on our CQT representation. We also transpose the metadataset by replacing the < natural > and the < modifier > accordingly (see Figure 20).

For our trainings, we choose to work with transpositions between -6 and +6 demi-tone.

### 3.2.3 Chord alphabet reduction

All ACE algorithms do not generate the same alphabets of chords and this choice influences the performance evaluation [44]. As mentioned in our motivation, we are interested in the functional aspect of the chord labels. Hence, we built our own alphabets that slightly differ from these presented in the subsection 3.1.2.1.

We use three different alphabets for our classification task :

1. Major and minor: N (which means no chord), maj, min;
2. Major and minor, seventh, diminished chords: N, maj, min, maj7, min7, 7, dim;

3. Major and minor, seventh, diminished, augmented chords: N, maj, min, maj7, min7, 7, dim, aug;

The most studied alphabet between them is the  $n^{\circ}2$ . It represents the harmonization of the major scale (Figure 3). With this alphabet we can describe nearly all the jazz, rock and pop songs.

I	II	III	IV	V	VI	VII
C Major	D Minor	E Minor	F Major	G Major	A Minor	B Diminished

Table 3: Harmonization of the C Major scale.

Isophonics’ dataset is annotated much richer than our desired alphabet, we then applied a chord alphabet reduction. This mapping is a surjection  $M : C_{M0} \rightarrow C_{M2}$  where  $C_{M0}$  is the original domain with the Harte’s notation [17] and  $C_{M2} = \{N, maj, min, maj7, min7, 7, dim\}$  is the domain of the alphabet  $n^{\circ}2$ . Our mapping associated all existing chords in  $C_{M0}$  with a chord in  $C_{M2}$ .

#### 3.2.4 Transforming the chords into pitch vectors

In a basic classification, a model does not take into account the distance between the output classes. Nonetheless, we want to include in our model some information from music theory. For instance, a CMaj7 is nearer to a Am7 than a C#Maj7. Our conception is based on harmonic network such as Tonnetz, that has been already exploited in many works on music information extraction [10, 24]. Thus, every chord of our alphabet is transformed into a 12-D vector where each dimension represents a pitch class.

### 3.3 OUR MODELS

In this section, we present the results obtained for our different models. We trained two families of models, namely CNN and LN. Each of these families are composed by different variants of these models. For the CNN, we have the *Vanilla* one which has been introduced earlier in this report (see Section 2.2.3). Here, we also tested the new proposals called Residual CNN and the Dense CNN. Both of these models are CNN with additional *skip-connections* between layers. These connections carry the untouched information from one layer to the next, in parallel to the transformation path. For more details, the interested reader can refer to the literature [19, 22]. We also implemented the LN, while performing the same types of improvements. We replaced the classic MLP layers with CNN layers and Residual CNN layers. For all the models we are working on the Beatles dataset introduced previously. We transform the input pressure signal through a CQT with 24 bins per octave and a frequency range of 27 Hz to 3.5kHz. Therefore, each frame contains 169 frequency bins, and we set the temporal context to 10 steps. Each temporal step is equivalent to 0.2s from an MP3 file, This implies that our overall context windows represent 2s

of sound. These parameter choices are based on previous works [23, 26].

To evaluate our results in an objective way, we provide in Table 4 the results of a CNN for Maj/Min chord extraction without any post-filtering method. These results have been obtained on the Isophonic dataset with multiple transpositions.

Train	Valid	Test
82.81	77.8	77.48

Table 4: 5-Fold recognition accuracy for ACE on Maj/Min classification task on Isophonics dataset with a CNN, results taken from [23]

A fold is a repartition of our dataset between train, valid and tests subsets. On the table 4 the authors present results that are the mean of 5 different folds. Nevertheless, our models were only tested on one fold.

We trained our models with different configurations : with and without transposition, with and without the chord to vector transformation and on the three different alphabets of the section 3.2.3.

### 3.3.1 Convolutional networks

The CNNs that we trained are composed by 3 layers, each layer has 500 neurons. The two first convolutional layers have 32 kernels (11 bin heights by 3 bin width), the last convolutional layer has also 32 kernels (5 bin heights by 2 bin width).

#### 3.3.1.1 Vanilla

The table 5 shows the results of the ACE task using a vanilla CNN without transposition on the dataset and without the chord to pitch vector space transformation.

	Train	Valid	Test
Alphabet n°1	85.51	73.16	73.34
Alphabet n°2	80.93	60.56	62.39
Alphabet n°3	81.71	55.27	59.80

Table 5: 1-Fold recognition accuracy for ACE on three different alphabet with a classic CNN on the Isophonics dataset.

*The other models are still under calculation, this report will be completed in an updated version.*

## INTRODUCING TEMPORAL CONTEXT AND MEMORY

The interest of adding a temporal context in our model is twofold. On one hand, this is a post-filtering method that completes our ACE models. On the other hand, the memory of the musical structures can be used for predictive models.

### 4.1 POST-FILTERING METHODS FOR ACE

The models presented on the previous part do not represent the whole chain of ACE described on section 3.1.1. Indeed, the post filtering step is missing. This step consists of smoothing the chord sequence obtained after the feature extractions. Among other methods, we decided to train LSTM layers to achieve this task.

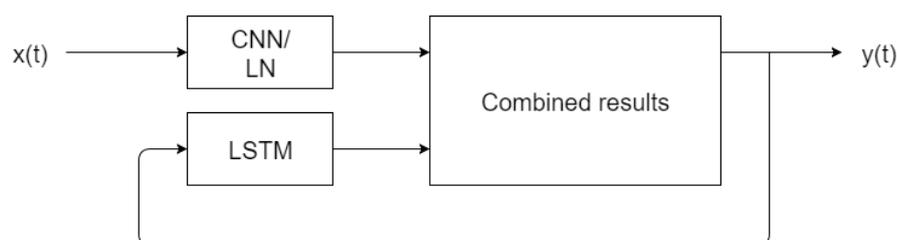


Figure 21: Complete ACE system using CNN/LN as feature extractor and LSTM as post-filtering method.

On the Figure 21 the CNN (or the LN) extracts the features from the signal  $x(t)$ . These feature are combined to the prediction of the LSTM realized on the last extracted features  $y(t-1), y(t-2), \dots, y(0)$ .

#### 4.1.1 Real Book

In order to augment the amount of data for the training of the LSTM, we used the dataset [12] composed of 2847 taken from the real book. The real book is a compilation of jazz classics that has been firstly collected by the students of the Berklee College of Music during the seventies. As of today we count a lot of existing books (e.g Realbook (1, 2, 3), New Realbook (1, 2, 3), the Fakebook, the real Latin book). The annotations of this dataset follow the ones of the Isophonics dataset.

### 4.2 GENERATING NEW DATA FROM MEMORY

In this part we present the generative aspect of the LSTM. We use the same training than for the ACE task. Nevertheless, instead of doing a prediction at only one step, we use the LSTM to generate chord sequences.

#### 4.2.1 Generation of a chord sequence

As shown in Figure 22, we have as input a chord sequence  $\{Gm, D7, Gm, D7, Bb, Gm\}$  where each element is contained in an alphabet  $C_M$ . The output of the LSTM is a vector with a dimension equals to the number of elements in  $C_M$ . This vector contains the prediction probabilities for the next chord. If we select the chord with the highest probability, we can add it to our initial sequence. Thus, at time  $t + 1$  we use this new sequence as input of the LSTM. At  $t + n$  we obtain a scenario of length  $n$  elements from the alphabet  $C_M$ .

Moreover, if we don't choose the chord with the highest probability but the second one, we obtain a slight variant in the predicted chord. Thus, we can generate multiple scenarios from the LSTM.

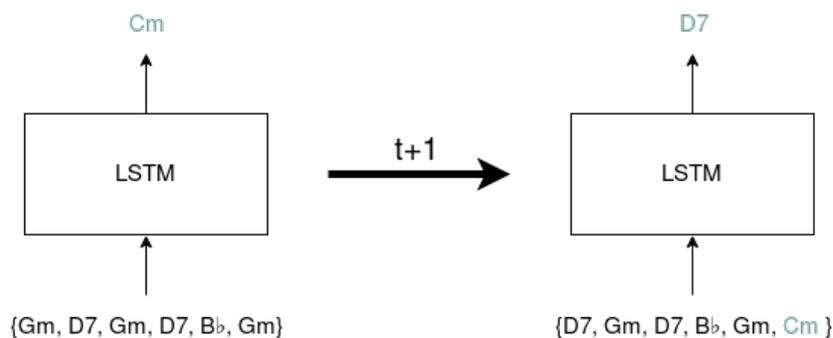


Figure 22: The chord sequence prediction is realized with adding element to the chord sequence at each time step.

### 4.3 OUR MODELS

Our LSTM model contains three layers. The number of LSTM units by layer not exceed 500. During the training we use a maximum of 8 time-steps for the back-propagation. The time-step is relative to each song and corresponds to its beat. The table 6 shows the results on the tree different alphabets (see section 3.2.3.)

	Train	Valid	Test
Alphabet n°1	38.55	30.05	31.44
Alphabet n°2	26.04	27.05	24.45
Alphabet n°3	14.62	15.88	16.62

Table 6: 1-Fold chord prediction task with a LSTM trained on 8 time-step sequence from the realbook dataset.

At first view these results seem perfectible. Nevertheless, deeper studies must be realized in order to analyze the output of the LSTM for a given chord sequence. Indeed, even if the predicted chord is not formally exact it could have a musical interest.

## 4.4 TOWARDS REAL-TIME PREDICTION

In our approach, we use machine learning methods for both the classification and the prediction task. We have then divided our system in two part : a detection system which aims to discover the chords playing and a predictive system.

The chord sequence sent to the LSTM is the one discovered by the ACE system in the musician's live record. Our overall software architecture is depicted on Figure 23.

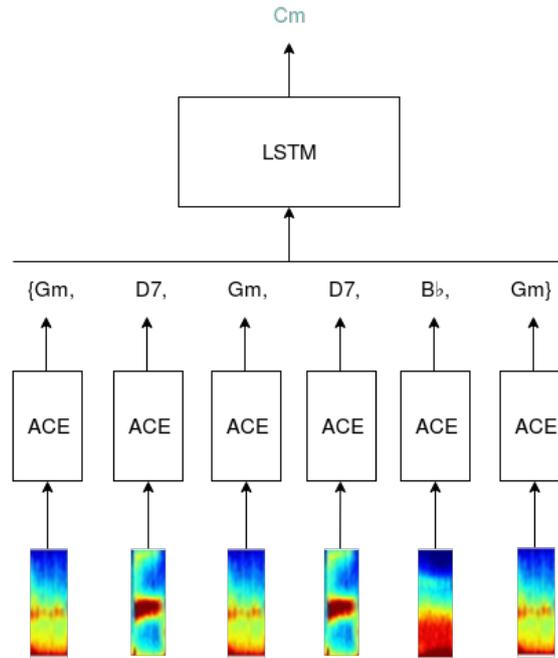


Figure 23: Our overall system architecture. Each frame of the audio flux is processed by the ACE system. Then the chord sequence is sent to a LSTM in order to predict the next chord.

We plan to use this software for real-time application. Indeed, the musician signal would be process by MAX/MSP in order to extract CQT frames. Then, the overall system should generate music chord sequences based on this audio stream. Finally, the predicted chord sequence would be sent at each time step to a scenario-based music generation software (e.g. ImproteK).

## CONCLUSION AND PERSPECTIVES

---

The goal of this internship was to enhance computer co-improvisation processes by introducing inference of short-term scenario. Thus, we separated this task between the ability to extract structures in an audio stream on one hand, and to generate symbolic music based on this information on the other hand. Our application case was the design of an ACE system and a subsequent chord sequence generator.

For the ACE task, we focused on the feature extraction part of the workflow. We pre-processed the data from Isophonics with CQT and STFT transforms. Then, we augmented the data by using sets of transpositions from the spectral transforms. We implemented various architectures to compare the results between CNN, Residual CNN, Dense CNN, LN, Convolutional LN and Residual LN models. We assessed different models as well as their parameters and their training and testing procedures. The CNN models are trained with a supervised learning algorithm whereas the LN models are trained in a semi-supervised way. Furthermore, we proposed another training process that takes as output a transformation from the chord label annotation to a pitch vector representation. This allowed us to define a more accurate distance criterion for chords. We obtained results comparable to the state of the art even without taking into account all our enhancements. Indeed, the global ACE system that includes LSTM post-filtering has not been tested yet. However, given its preliminary results, we can expect it to improve our results even further.

Secondly, we designed a LSTM network for chord sequence generation. We trained this LSTM network with the realbook dataset, which allows to cast our problem into a purely symbolic inference task. This approach allows to generate a scenario which is the continuation of a given chord sequence. Moreover, the scenarios are not pre-calculated but created from a probabilistic study on a large corpus. We showed the applicability of this system and its encouraging results.

Finally, we designed the overall architecture of the system and proposed an overarching workflow of our application that could be applicable to real-time setups. Nevertheless, we did not yet realize application case use with musicians.

### FUTURE WORK

Given the large amount of models and enhancements proposed in this work, we are still currently waiting for the complete results for all the models. As soon as all models are assessed, we plan to first investigate the reasons behind their relative successes and failures and to perform hyper-optimization of the parameters for the best models. Then, this model will be trained on different musical annotations

(chroma, chord progressions), in order to develop a multi-scale system able to scrutinize temporal relationships at different time scales.

Furthermore, we strongly believe that the post-filtering method would also largely improve our results. On the other hand, the LSTM outputs could also be interpreted in a musical way. We intend to define this evaluation procedure in order to gain some insights on the inference and generation process.

Regarding the overall system, a real-case use of our modules including real-time human improvisation will be performed to confirm the qualities and drawbacks of our proposals. Therefore, we plan to test the complete system in real case application in order to have feedback from musicians.

Finally, a study on multivariate signals that could allow structure inference and generation for multiple musicians could be performed by relying on our overall framework.

## BIBLIOGRAPHY

---

- [1] Gérard Assayag and Georges Bloch. “Navigating the oracle: A heuristic approach.” In: *International Computer Music Conference’07*. 2007, pp. 405–412.
- [2] Gérard Assayag, Georges Bloch, Marc Chemillier, Arshia Cont, and Shlomo Dubnov. “Omax brothers: a dynamic topology of agents for improvisation learning.” In: *Proceedings of the 1st ACM workshop on Audio and music computing multimedia*. ACM. 2006, pp. 125–132.
- [3] Yoshua Bengio, Ian J Goodfellow, and Aaron Courville. “Deep learning.” In: *Nature* 521 (2015), pp. 436–444.
- [4] Yoshua Bengio, Pascal Lamblin, Dan Popovici, and Hugo Larochelle. “Greedy layer-wise training of deep networks.” In: *Advances in neural information processing systems*. 2007, pp. 153–160.
- [5] Yoshua Bengio et al. “Learning deep architectures for AI.” In: *Foundations and trends® in Machine Learning* 2.1 (2009), pp. 1–127.
- [6] Laurent Bonnasse-Gahot. “An update on the SOMax project.” In: *Ircam-STMS, Internal report ANR project Sample Orchestrator 2* (2014).
- [7] Léon Bottou. “Large-scale machine learning with stochastic gradient descent.” In: *Proceedings of COMPSTAT’2010*. Springer, 2010, pp. 177–186.
- [8] Nicolas Boulanger-Lewandowski, Yoshua Bengio, and Pascal Vincent. “Audio Chord Recognition with Recurrent Neural Networks.” In: *ISMIR*. Citeseer. 2013, pp. 335–340.
- [9] Joel Chadabe. “Some reflections on the nature of the landscape within which computer music systems are designed.” In: *Computer Music Journal* (1977), pp. 5–11.
- [10] Elaine Chew. “Towards a mathematical model of tonality.” PhD thesis. Massachusetts Institute of Technology, 2000.
- [11] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. “Learning phrase representations using RNN encoder-decoder for statistical machine translation.” In: *arXiv preprint arXiv:1406.1078* (2014).
- [12] Keunwoo Choi, George Fazekas, and Mark Sandler. “Text-based LSTM networks for automatic music composition.” In: *arXiv preprint arXiv:1604.05358* (2016).

- [13] Dan Claudiu Cireșan, Ueli Meier, Luca Maria Gambardella, and Jürgen Schmidhuber. “Convolutional neural network committees for handwritten character classification.” In: *Document Analysis and Recognition (ICDAR), 2011 International Conference on*. IEEE. 2011, pp. 1135–1139.
- [14] Alexandre Donzé, Rafael Valle, Ilge Akkaya, Sophie Libkind, Sanjit A Seshia, and David Wessel. “Machine improvisation with formal specifications.” In: *ICMC*. 2014.
- [15] Arne Eigenfeldt and Philippe Pasquier. “Realtime generation of harmonic progressions using controlled Markov selection.” In: *Proceedings of ICCX-X-Computational Creativity Conference*. 2010, pp. 16–25.
- [16] Alex Graves. “Generating sequences with recurrent neural networks.” In: *arXiv preprint arXiv:1308.0850* (2013).
- [17] Christopher Harte. “Towards automatic extraction of harmony information from music signals.” PhD thesis. 2010.
- [18] Christopher Harte, Mark Sandler, and Martin Gasser. “Detecting harmonic change in musical audio.” In: *Proceedings of the 1st ACM workshop on Audio and music computing multimedia*. ACM. 2006, pp. 21–26.
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. “Deep residual learning for image recognition.” In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016, pp. 770–778.
- [20] Dorien Herremans, Stéphanie Weisser, Kenneth Sørensen, and Darrell Conklin. “Generating structured music for bagana using quality metrics based on Markov models.” In: *Expert Systems with Applications* 42.21 (2015), pp. 7424–7435.
- [21] Sepp Hochreiter and Jürgen Schmidhuber. “Long short-term memory.” In: *Neural computation* 9.8 (1997), pp. 1735–1780.
- [22] Gao Huang, Zhuang Liu, Kilian Q Weinberger, and Laurens van der Maaten. “Densely connected convolutional networks.” In: *arXiv preprint arXiv:1608.06993* (2016).
- [23] Eric J Humphrey and Juan Pablo Bello. “Rethinking automatic chord recognition with convolutional neural networks.” In: *Machine Learning and Applications (ICMLA), 2012 11th International Conference on*. Vol. 2. IEEE. 2012, pp. 357–362.
- [24] Eric J Humphrey, Taemin Cho, and Juan P Bello. “Learning a robust tonnetz-space transform for automatic chord recognition.” In: *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*. IEEE. 2012, pp. 453–456.
- [25] David Brian Huron. *Sweet anticipation: Music and the psychology of expectation*. MIT press, 2006.
- [26] Filip Korzeniowski and Gerhard Widmer. “A fully convolutional deep auditory model for musical chord recognition.” In: *Machine Learning for Signal Processing (MLSP), 2016 IEEE 26th International Workshop on*. IEEE. 2016, pp. 1–6.

- [27] Filip Korzeniowski and Gerhard Widmer. “Feature learning for chord recognition: the deep chroma extractor.” In: *arXiv preprint arXiv:1612.05065* (2016).
- [28] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. “Imagenet classification with deep convolutional neural networks.” In: *Advances in neural information processing systems*. 2012, pp. 1097–1105.
- [29] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. “Deep learning.” In: *Nature* 521.7553 (2015), pp. 436–444.
- [30] Yann LeCun, Yoshua Bengio, et al. “Convolutional networks for images, speech, and time series.” In: *The handbook of brain theory and neural networks* 3361.10 (1995), p. 1995.
- [31] MIREX ACE 2016. [http://www.music-ir.org/mirex/wiki/2016:Audio\\_Chord\\_Estimation\\_Results](http://www.music-ir.org/mirex/wiki/2016:Audio_Chord_Estimation_Results). Accessed: 2017-09-20.
- [32] Matthias Mauch. “Automatic chord transcription from audio using computational models of musical context.” In: (2010).
- [33] Matt McVicar, Raúl Santos-Rodríguez, Yizhao Ni, and Tijl De Bie. “Automatic chord estimation from audio: A review of the state of the art.” In: *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)* 22.2 (2014), pp. 556–575.
- [34] Julian Moreira, Pierre Roy, and François Pachet. “Virtualband: Interacting with Stylistically Consistent Agents.” In: *ISMIR*. 2013, pp. 341–346.
- [35] Andrew Ng. “Coursera Machine Learning, Lecture 2 Slides.” In: ().
- [36] Jérôme Nika. “Guiding human-computer music improvisation: introducing authoring and control with temporal scenarios.” PhD thesis. Paris 6, 2016.
- [37] Jérôme Nika, Marc Chemillier, and Gérard Assayag. “ImproteK: introducing scenarios into human-computer music improvisation.” In: *Computers in Entertainment (CIE)* 14.2 (2016), p. 4.
- [38] Jérôme Nika, José Echeveste, Marc Chemillier, and Jean-Louis Giavitto. “Planning human-computer improvisation.” In: *ICMC*. 2014.
- [39] Jérôme Nika, Dimitri Bouche, Jean Bresson, Marc Chemillier, and Gérard Assayag. “Guided improvisation as dynamic calls to an offline model.” In: *Sound and Music Computing (SMC)*. 2015.
- [40] *Overfitting*. <https://en.wikipedia.org/wiki/Overfitting>. Accessed: 2017-09-20.
- [41] François Pachet, Pierre Roy, Julian Moreira, and Mark d’Inverno. “Reflexive loopers for solo musical improvisation.” In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM. 2013, pp. 2205–2208.

- [42] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. "On the difficulty of training recurrent neural networks." In: *ICML (3)* 28 (2013), pp. 1310–1318.
- [43] Jouni Paulus, Meinard Müller, and Anssi Klapuri. "State of the Art Report: Audio-Based Music Structure Analysis." In: *ISMIR*. 2010, pp. 625–636.
- [44] Johan Pauwels and Geoffroy Peeters. "Evaluating automatically estimated chord sequences." In: *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*. IEEE. 2013, pp. 749–753.
- [45] Geoffroy Peeters. "Sequence Representation of Music Structure Using Higher-Order Similarity Matrix and Maximum-Likelihood Approach." In: *ISMIR*. 2007, pp. 35–40.
- [46] Mohammad Pezeshki, Linxi Fan, Philemon Brakel, Aaron Courville, and Yoshua Bengio. "Deconstructing the ladder network architecture." In: *arXiv preprint arXiv:1506.02351* (2016).
- [47] Antti Rasmus, Harri Valpola, and Tapani Raiko. "Lateral connections in denoising autoencoders support supervised learning." In: *arXiv preprint arXiv:1504.08215* (2015).
- [48] Antti Rasmus, Mathias Berglund, Mikko Honkala, Harri Valpola, and Tapani Raiko. "Semi-supervised learning with ladder networks." In: *Advances in Neural Information Processing Systems*. 2015, pp. 3546–3554.
- [49] Arnold Schoenberg and Leonard Stein. *Structural functions of harmony*. 478. WW Norton & Company, 1969.
- [50] Ricardo Scholz, Emmanuel Vincent, and Frédéric Bimbot. "Robust modeling of musical chord sequences using probabilistic N-grams." In: *Acoustics, Speech and Signal Processing, 2009. ICASSP 2009. IEEE International Conference on*. IEEE. 2009, pp. 53–56.
- [51] Alexander Sheh and Daniel PW Ellis. "Chord segmentation and recognition using EM-trained hidden Markov models." In: (2003).
- [52] Aarti Singh, Robert Nowak, and Xiaojin Zhu. "Unlabeled data: Now it helps, now it doesn't." In: *Advances in neural information processing systems*. 2009, pp. 1513–1520.
- [53] Bob L Sturm, João Felipe Santos, Oded Ben-Tal, and Iryna Korshunova. "Music transcription modelling and composition using deep learning." In: *arXiv preprint arXiv:1604.08723* (2016).
- [54] *Understanding LSTM Networks*. <http://colah.github.io/posts/2015-08-Understanding-LSTMs/>. Accessed: 2017-09-20.
- [55] Gregory H Wakefield. "Mathematical representation of joint time-chroma distributions." In: *International Symposium on Optical Science, Engineering, and Instrumentation, SPIE*. Vol. 99. 1999, pp. 18–23.

- [56] Paul John Werbos. *The roots of backpropagation: from ordered derivatives to neural networks and political forecasting*. Vol. 1. John Wiley & Sons, 1994.
- [57] Jianxin Wu. "Introduction to Convolutional Neural Networks." In: 2017.
- [58] Xiaojin Zhu and Andrew B Goldberg. "Introduction to semi-supervised learning." In: *Synthesis lectures on artificial intelligence and machine learning* 3.1 (2009), pp. 1–130.

#### COLOPHON

This document was typeset using the typographical look-and-feel `classicthesis` developed by André Miede. The style was inspired by Robert Bringhurst's seminal book on typography "*The Elements of Typographic Style*". `classicthesis` is available for both  $\text{\LaTeX}$  and  $\text{\LyX}$ :

<https://bitbucket.org/amiede/classicthesis/>

*Final Version* as of August 21, 2017 (`classicthesis` version 4.2).