

On the Use of an Effective Boltzmann Machine for Musical Style Recognition and Harmonisation

Matthew I Bellgard
Computer Science
Murdoch University
Murdoch, 6150, Australia
matthew@cs.murdoch.edu.au

Chi-Ping Tsang
Computer Science
The University of Western Australia
Nedlands, 6907, Australia
tsang@cs.uwa.edu.au

Abstract

This paper describes the application of the Effective Boltzmann Machine (EBM) to musical style recognition. After it has been trained with examples of four-part chorales, the EBM is able to distinguish different styles of chorales in unseen pieces of music. Our earlier studies showed that the EBM has many desirable properties. In particular, it is able to complete arbitrary length contextual sequences from learned local context, and synthesise musical harmony for polyphonic music. In this paper we show that the EBM is able to distinguish chorales from the Baroque period from Norwegian chorales. In summary, the EBM is able to distinguish, firstly, between chorales present in the training set and those that are not. Secondly, the EBM is also able to distinguish chorales of different styles, where one style was used in the training set. We also show that two different EBMs (one trained from examples from Choralbuch and the other trained from Norwegian chorales) are able to distinguish between two different harmonised styles of the same chorale melody. The EBMX system is a PC-based graphical interface, interactive system implementing the EBM for recognition and harmonisation of short musical pieces.

1 Introduction

To a human it is relatively easy to distinguish between different music styles even to a novice. However, it is difficult for a computer to perform this function. Although there are many rule based expert systems for music composition (for example [Eb88]), it is hard to use them for the purpose of musical style recognition. This is because there is no correctness evaluation measure and styles are not readily expressed in absolute rules. Furthermore, all these systems rely on the use of human perceived rules, and they are only as good as the human expert. In recent years there have also been some inductive learning systems, but they are far from being able to recognise musical styles. We present here a self learning music style recogniser which is able to distinguish musical styles.

Our system is based on our earlier work on the Effective Boltzmann Machine (EBM) which is an extension of the Boltzmann Machine (BM) neural network. Initially, we use a BM to learn local contexts from a set of music score examples and then an EBM is constructed to synthesise arbitrary length contextual sequences (using the trained BM) in a non-deterministic manner. This EBM system has many desirable properties. Firstly, it can perform pattern completion rather than pattern matching. Secondly, it can provide an evaluation of the synthesised result. A version of the EBM has

been applied to four-part chorale harmonisation. In this particular system, the EBM can find bass, tenor and alto parts to a given soprano part of any length. It also has the potential to find, say, the bass part for a given soprano, alto and tenor score. In fact, it is even more general than this. It can either complete a pitch for a given context and provide an evaluation of the completion or provide an evaluation for a completed piece. It is the latter property of the EBM that we will utilise in this paper.

In this paper, we show, by experiment, how the EBM is able to both distinguish between four-part chorales of different styles as well as recognise a given chorale as belonging to the style. It is also able to distinguish Baroque chorales from Norwegian chorales. The choice of chorale was made because different harmonised styles can be found for the same melody such that objective comparison can be made. In summary, the EBM is able to distinguish between chorales present in the training set and those that are not, as well as chorales of different styles. In addition we also show that by using two different EBMs (one trained from examples from Choralbuch and the other trained from Norwegian chorales) one can distinguish between two different harmonised styles of the same chorale melody. The EBMX PC-based harmonisation/recognition system will also be demonstrated.

The paper is structured as follows: section 2 briefly describes the EBM; section 3 describes the application of the EBM to music style recognition; and section 4 discusses the results and provides some conclusions.

synthesised solution. On one hand, it can behave as compositional/harmonisation system and on the other hand, it can behave as an music recogniser. Figure 1b shows the construction of an EBM to harmonise the given soprano melody in Figure 1a. The EBM in this instance is constructed from a BM that has a window size of 3 chords. That is, the BM

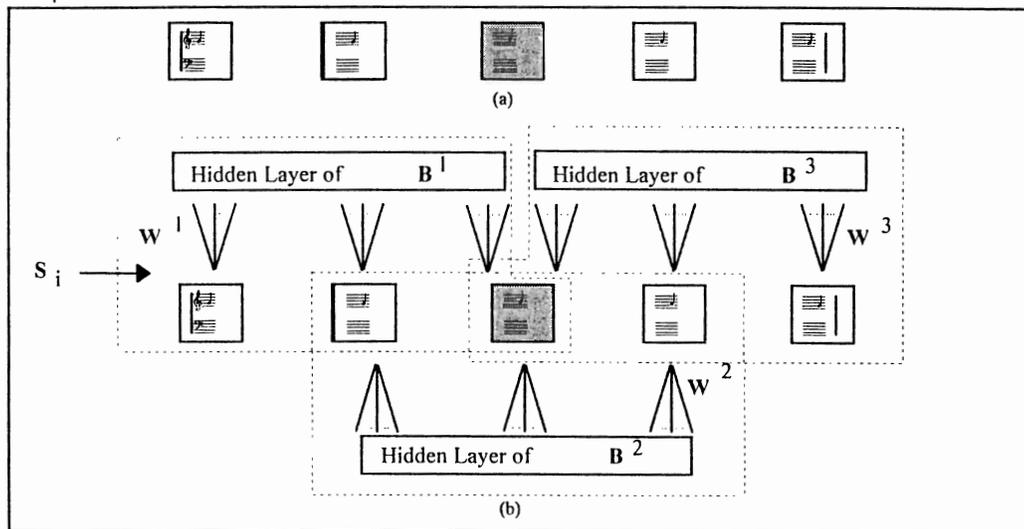


Figure 1. (a) a soprano melody to be harmonised. (b) construction of the EBM using 3 identical BMs. Note that, for example, the shaded pitch (to be harmonised) appears in each of the three BMs but in different positions.

2 The Effective Boltzmann Machine

The Effective Boltzmann Machine (EBM) [TB90, BT92, BT94, BT95] is a relaxation network. It is able to complete arbitrary length contextual sequences from learned local context. A version of the EBM has been applied to the synthesis of musical harmony for polyphonic music [BT92, BT94]. It has been further extended [BT95] so that it can be used for learning and not just completion. In addition, it has been shown that a hierarchical EBM (HEBM) can be constructed to learn longer-term contexts and not be limited by the window size. Like the Boltzmann Machine (BM), an Artificial Neural Network on which the EBM is based, the EBM is also a completion machine. Completions of partially specified sequences are produced in a non-deterministic manner. The completion process is not directional in terms of the temporal/spatial nature of the sequence. That is, it can capture implicational constraints in both directions. This is unlike other ANN systems that can only learn contextual information in one-direction (for example, [EI90]). Completions of partially specified sequences are produced in a non-deterministic manner. Another important property which is the emphasis of this paper, is that the EBM can provide an energy measure for a given

has been trained with all the local contexts (of 3 chords in length) in a set of chorales (the training set). Note that, for example, the shaded pitch (to be harmonised) appears in each of the three BMs but in different positions. It is overlapping property of contexts that enables long-range effects.

3 Using EBM to distinguish different styles of music

Experiment 1 : Distinguishing musical styles

In this experiment, we show that an EBM may be used to recognise styles of music. Different pieces of music were clamped to the input layer of the EBM which was then relaxed to determine an evaluation for the piece. The BM used was the one trained on a set of five chorales taken from Choralbuch [Do50]. We shall refer to this training set as *training set A*. Another BM was trained with 3 Norwegian Chorales [Wo48]. We shall refer to this training set as *training set B*. Six different chorales were chosen for evaluation. Only the first ten and the last ten chords of each chorale were used in the recognition process. This was done to ensure that the evaluation was unbiased in the number of chords for each evaluation. The same number of chords means that the evaluation will only distinguish the content of the different pieces

Chorales used for evaluation			EBM Using Training Set B			
Chorale Name	In Training Set A	Taken from	Evaluation of First 10 Chords		Evaluation of Last 10 Chords	
			Mean	Std	Mean	Std
Es ist genug, so nimm, Herr	Yes	ChoralBuch	-4220	14	-4165	10
Lobe den Herren, den machtigen	No	ChoralBuch	-3610	18	-3322	18
Nun freut euch, liebe Chrsteng'mein	No	ChoralBuch	-3037	20	-2971	21
Alene Gud I Himmerik	No	Norwegian Chorale Collection	-2640	23	-1908	13
Akk, Herre Gud	No	Norwegian Chorale Collection	-2251	30	-1993	15
Det Hev Ei Rosa Sprunge	No	Norwegian Chorale Collection	-2599	21	-2533	21

Table 1. Comparing different styles of chorale harmonisation. A variety of chorales are chosen, clamped to the 10 layer of the EBM and an evaluation is found for each clamping. The lower the value, the closer the clamped chorale to the chorales used in the training set. Only the first 10 and the last 10 chords of each chorale were clamped.

tested. The first chorale clamped was a chorale taken from training set A. The second and third chorales were also taken from Choralbuch but were not in training set A. The fourth, fifth and sixth chorales were taken from the set of Norwegian chorales [Wo48]. The results are tabulated in Table 1. The evaluations for each chorale segment were averaged over ten independent trials. The standard deviations which are nominal, are shown along with averages. The lower the value, the better the evaluation in terms of what the EBM learnt in the training set.

As is evident from Table 1, the evaluation of the chorale that was used in the training set (the first ten and last ten chords) has a better evaluation than the other chorales tested (ie. a lower value). The second and third chorales, which were taken from Choralbuch but were not present in the training set A, have a relatively lower evaluation (higher value) than the first chorale. However, they also have a relatively better evaluation than the Norwegian chorales.

The results in Table 1 confirm the judgement of a musical expert that the chorales taken from Choralbuch are a different style from the Norwegian chorales. The EBM is able to distinguish, firstly, between chorales present in the training set and those that are not. Secondly, the EBM is also able to distinguish chorales of different styles, where one style was used in the training set.

Experiment 2: Comparing two different harmonised styles of the same chorale melody

The chorale, *Lobe den Herren, den machtigen*, is harmonised by human composers in both the Choralbuch as well as the Norwegian style. The

Choralbuch version is not contained in Training set A where as the Norwegian version is contained in training set B. This next experiment was designed to determine if the EBM can distinguish between styles when 1) the chorale of either style was not contained in the training set (training set A) and 2) when the chorale of one style was contained in the training set (training set B).

Two EBMs, one using the training set A and the other using training set B, were clamped with the two different harmonised styles of the chorale, *Lobe den Herren, den machtigen* (the first 10 chords and the last 10 chords of the chorale). The evaluations were averaged from ten independent trials. The results are shown in Table 2.

As shown in Table 2, the EBM is able to distinguish between the two styles. Using training set A (chorales from Choralbuch), where the chorale of either style is not present in the training set, the EBM is able to distinguish between the two harmonised styles of the chorale. In other words, a better evaluation results for the Choralbuch style of the chorale than for the Norwegian version. As shown in Table 3, using training set B (chorales from Norwegian collection), which contained the Norwegian version of the chorale in the training set, the EBM was also able to distinguish between the two styles. The evaluation of the Norwegian style harmony is better than the Choralbuch version.

What is also evident from this experiment is that there is a difference in evaluation between the first 10 and the last 10 chords for the Choralbuch version of the chorale in Tables 2 and 3. For example, in the first row of Table 2, the evaluation for the first 10 chords (from Choralbuch) is better than the evaluation for the last 10 chords using

Chorales used for evaluation			EBM Using Training Set B			
Chorale Name	In Training Set A	Taken from	Evaluation of First 10 Chords		Evaluation of Last 10 Chords	
			Mean	Std	Mean	Std
Lobe den Herren, den machtigen	No	ChoralBuch	-3610	18	-3322	18
Lobe den Herren, den machtigen	No	Norwegian Choral Collection	-1708	34	-1745	26

Table 2. Comparison of a particular chorale in two different styles. Two different harmonisations of the chorale, *Lobe den Herren, den machtigen*, is clamped to the IO layer of an EBM with a BM trained with training set A.

Chorales used for evaluation			EBM Using Training Set C			
Chorale Name	In Training Set B	Taken from	Evaluation of First 10 Chords		Evaluation of Last 10 Chords	
			Mean	Std	Mean	Std
Lobe den Herren, den machtigen	No	ChoralBuch	-5419	11	-6010	11
Lobe den Herren, den machtigen	Yes	Norwegian Choral Collection	-6467	5	-6453	10

Table 3. Comparison of a particular chorale in two different styles. Two different harmonisations of the chorale, *Lobe den Herren, den machtigen*, is clamped to the IO layer of an EBM with a BM trained with training set B.

training set A in the EBM. That is, the first 10 chords collectively are "more like" the type of chords contained in the chorales in training set A and the last 10 chords are "less like" the chorales in training set A. This difference in the first row, is also reflected in the second row of Table 3. Using training set B (which contained Norwegian style chorales), the first 10 chords (from Choralbuch) are "less like" the chorales in training set B and the last 10 chords are "more like" the chorales in training set B. This confirms an expert opinion that the two chorales are of different styles.

4 Discussion

We have presented results of two experiments. The first experiment demonstrates that by using a single EBM one can recognise whether a given chorale segment is of the same style as those used in the training set. The discrimination is good. This is independent from the fact whether it is an unseen piece or not. Of course, the discrimination is even clearer for those pieces which are in the training set. In the second experiment, we employed two EBMs each trained with different styles. They are used to decide whether the chosen chorale "Lobe den Herren, den machtigen" is one or the other style. As this particular chorale actually has both style of harmonisation, we tested both of them. From our results, it is clear that EBM can distinguish the different styles.

We have demonstrated that EBM is useful for distinguishing musical styles. More importantly it can learn from a set of examples without human intervention. The discrimination is good. Further work is currently being carried out in improving the

design of EBM and experimenting with other styles of music.

References

- [AHS85] Ackley, D. H., Hinton, G. E. and Sejnowski, T. J. (1985) "A Learning Algorithm for Boltzmann Machines", *Cognitive Science*, 9, pp. 147-169.
- [AK88] Aarts E and Korst, J. (1988) *Simulated Annealing and Boltzmann Machines*, John Wiley & Sons, New York.
- [BT92] Bellgard, M. I. and Tsang, C. P. (1992) "Harmonizing Music Using a Network of Boltzmann Machines", in *Proceedings of Fifth International Conference of Artificial Neural Networks and their Application (Neuro-Nimes '92)*, EC2, pp. 321-332.
- [BT94] Bellgard, M. I. and Tsang, C. P. (1994) "Harmonising Music The Boltzmann Way", *Connection Science*, Vol. 6, Nos 2 & 3, pp. 281-297, 1994.
- [BT95] Bellgard, M. I. & Tsang, C. P. (1995), "Using an Effective Boltzmann Machine to Learn Context Dependencies of a Sequence", in the *Proceedings of the 1995 IEEE International Conference on Neural Networks (ICNN '95)*, Perth, WA., Vol. 5, pp.2841-2846.
- [Do50] Dorffel, A. 1950. *CHORALBUCH* C. F. Peters Corporation, New York.
- [Eb88] Ebcioğlu, K. 1988. "An Expert System for Harmonizing Four-part Chorals", *Computer Music Journal*, V.12, No.3, pp.43-51.
- [El90] Elman, J. L., (1990), "Finding Structure in Time", *Cognitive Science*, 14, pp. 179-211.
- [TB90] Tsang, C.P., and Bellgard, M.I. 1990. "Sequence Generation Using a Network of Boltzmann Machines", *Proceedings of the 4th Australian Joint Conference on Artificial Intelligence (AI '90)*, Tsang, C.P. (Ed), World Scientific, pp. 224-233.
- [Wo48] Wonderlich, E., (1948) *Chorale Collection*, 3rd edition, New York, Appleton.