



PROBABILISTIC MODELS FOR MELODIC PREDICTION

Jean-Francois Paiement ^a Samy Bengio ^b
Douglas Eck ^c

IDIAP-RR 08-50

JUNE 2008

SUBMITTED FOR PUBLICATION

^a Idiap Research Institute, Rue Marconi 19, C.P 592, CH-1920, Martigny, Switzerland.

^b Google, 1600 Amphitheatre Pkwy, Mountain View, CA 94043, USA.

^c Université de Montréal, Department of Computer Science, CP 6128, Succ. Centre-Ville, Montréal, Québec, Canada, H3C 3J7.

PROBABILISTIC MODELS FOR MELODIC PREDICTION

Jean-Francois Paiement

Samy Bengio

Douglas Eck

JUNE 2008

SUBMITTED FOR PUBLICATION

Abstract. Chord progressions are the building blocks from which tonal music is constructed. The choice of a particular representation for chords has a strong impact on statistical modeling of the dependence between chord symbols and the actual sequences of notes in polyphonic music. Melodic prediction is used in this paper as a benchmark task to evaluate the quality of four chord representations using two probabilistic model architectures derived from Input/Output Hidden Markov Models (IOHMMs).

1 Introduction

Probabilistic models for analysis and generation of polyphonic music would be useful in a broad range of applications, from contextual music generation to on-line music recommendation and retrieval. However, modeling music involves capturing long term dependencies in time series. This has proved very difficult to achieve with traditional statistical methods. Note that the problem of long-term dependencies is not limited to music, nor to one particular probabilistic model [5]. This difficulty motivates our exploration of chord progressions and their interaction with melodies. A chord is a group of three or more notes. A chord progression is simply a sequence of chords. In probabilistic terms, the current chord can be seen as a latent variable (local in time) that conditions the probabilities of choosing particular notes in other music components, such as melodies or accompaniments. Chord changes occur at fixed time intervals in most of the musical genres, which makes them much simpler to detect than beginnings and endings of musical notes, which can happen almost everywhere in music signal. Thus, knowing the relations between such chords and actual notes would certainly help to discover long-term musical structures in tonal music. For instance, an interesting challenge arising in the music information retrieval context is transcription, i.e. converting audio data into any kind of symbolic representation such as MIDI or traditional music notation. However, because of fundamental difficulties inherent to the nature of sound, state-of-the-art techniques are not able to accomplish this task with a sufficient level of precision for most practical applications. An intermediate goal is to try to infer chord symbols from audio data [3, 18]. This task is simpler than complete transcription of polyphonic audio signal. Combining reliable chord transcription with a model of the conditional distribution of other music components (e.g. melodies) given chords could help to improve transcription error rates of existing algorithms. Following the same idea, such models could even be included in genre classifiers or automatic composition systems [9] to increase their performance.

In most tonal music theories, chord names are defined by a root note that can either be expressed by its absolute pitch-class¹ or by its relation with the current key. The key of a song is designated by a note name (the tonic), and is the base of a musical scale from which most of the notes of the piece are drawn. Most commonly, that scale can be either in major or minor mode. Despite the simple and relatively universal chord building principles, many different notations have been used through music history to represent chord progressions [7]. We face the same problem in the computer science literature, where each author uses a notation corresponding to his musical background [1, 16, 18, 19, 22]. All these papers describe arbitrary chord representations embedded in probabilistic models, which are mostly variants of Hidden Markov Models (HMMs). However, to the best of our knowledge, there is no available quantitative comparative study of the effect of the choice of particular chord representations to solve practical applications. Each chord representation carries specific information that could be more adapted to certain tasks (or musical styles) than others. At the same time, all of these notations encapsulate basic information about a chord, such as its root. What are the usual probabilistic relations between these chord representations and the actual sequences of notes in polyphonic music? This paper explores this issue by using melodic prediction [21] given chords as a benchmark task to evaluate the quality of many chord representations. Likelihoods and conditional and unconditional prediction error rates are used as complementary measures of the quality of each representations.

Previous papers describe probabilistic models to solve music related problems. Using graphical models, Cemgil [6] introduces a somewhat complex probabilistic model that generates a mapping from audio to a piano-roll using a simple model for representing note transitions based on Markovian assumptions. This model takes as input audio data, without any form of preprocessing. While being computationally costly, this approach has the advantage of being completely data-dependent. However, strong Markovian assumptions are necessary in order to model the temporal dependencies between notes. Hence, a proper chord transition model could be appended to such a transcription model in order to improve polyphonic transcription performance. Raphael and Stoddard [18] use graphical models for labeling MIDI data with traditional Western chord symbols. Lavrenko and Pick-

¹A pitch-class is a note name, like C or D. In this paper, we consider enharmonics (e.g. Eb and D#) to be completely equivalent.




Figure 1: A simple probabilistic model where the influence of each chord is direct on melodic observations. Circles represent random variables and arrows represent conditional probability distributions. When all chord variables in level 1 are observed, the dashed lines are irrelevant, thus making each observation completely local in time. Variables in level 2 correspond to melodic observations.

ens [14] propose a generative model of polyphonic music that employs Markov random fields. While being very general, this model would benefit from having access to more specific musical knowledge. For instance, we go a step further in this paper by using abstract chord representations as a smoothing technique towards better generalization. Begleiter et al. [2] provide an interesting comparison of variable order Markov models used for polyphonic music prediction, but without any use of chord information to overcome long term dependencies. Harmonization is the generation of a chord progression for a given melody. Allan and Williams [1] designed a harmonization model for Bach chorales using HMMs. While generating excellent musical results, this model has to be provided polyphonic music with specific 4 voice structure as input, restricting its applicability to very specific settings. Our proposed models are more general in the sense that it is possible to extract the appropriate chord representation from any polyphonic music, whatever the specific labeling, harmonic structure, or musical style.

2 Melodic Prediction Models

In order to assess the effect of using particular chord representations for melodic prediction, we propose two kinds of probabilistic models.

2.1 A Local Model

The first proposed strategy is to look at the direct effect of particular chord representations without any influence from past observations. In the remainder of this paper, we refer to this model as the *local* model, presented in Figure 1. Variables in level 1 are associated with chord observations. Their particular form is described in Section 2.3. Such variables are observed during training *and* testing. Thus, the dashed horizontal links are not relevant in this context, making each time-step independent of the others. While these variables are discrete, their number of possible values depend on the chosen chord representation, as described in Section 2.3. Variables in level 2 correspond to melodic observations. In this paper, we assume octave invariance for the melodic observations. In other words, all notes belonging to the same pitch-class are considered to be the same (e.g. all C notes regardless of octave are associated to the same random variable value). In level 2, we assign one value to each pitch-class, plus one extra value for silence, leading to a total of 13 possible melodic values. This melodic representation does not account for note similarities. In the proposed models, the probability of the melodic observations given appropriate other random variables is modeled by multinomials. Such a distribution does not embed any notion of similarity between its possible outcomes. However, we chose




Figure 2: Variant of an IOHMM model. The variables in level 1 are observed and correspond to chord observations. Variables in level 2 are hidden, while variables in level 3 correspond to melodic observations.

this melodic representation for its simplicity and also to avoid introducing bias while measuring the quality of the chord representations for melodic prediction. While this model is overly simplistic for practical purposes, it has the advantage of isolating the direct effect of particular chord representations on the choice of melodic notes. Since all the variables are observed, parameters for this model can be easily computed over a training set by standard maximum likelihood techniques.

2.2 IOHMM Model

A more realistic model can be designed by adding extra hidden variables in the previous model to consider influences from the past when trying to predict a melody note. The model presented in Figure 2 is very similar to an input/output hidden Markov model (IOHMM) [4]. In our implementation, all the variables in the model are discrete. Variables in level 1 are always observed and again correspond to chord observations. Variables in level 2 are always hidden and are used to introduce dependencies between time frames in the model. Variables in level 3 correspond to melodic observations and have 13 possible values as in the local model presented in Section 2.1. There is no link between level 1 and level 3 variables on Figure 2, contrary to standard IOHMMs [4]. The number of possible values is highly variable from one chord representation to another. Considering that, we chose to remove the usual links between inputs and outputs in IOHMMs in order to limit the impact of the particular choice of a chord representation on the capacity of the model. This way, the number of possible values of the chosen chord representation has an impact on the parameterization of the conditional distribution of the hidden variables, but not on the conditional distributions of the predicted melodic variables. Learning in the model is done with the standard EM algorithm [8]. Marginalization must be carried out in the proposed model both for learning (during the expectation step of the EM algorithm) and for evaluation. Exact marginalization with the standard Junction Tree Algorithm [13] is tractable in IOHMMs because of their limited complexity.

2.3 Chord Representations

The chord representations that we introduce in this section consider chord symbols as they are represented in musical analysis instead of actual instantiated chords. In other words, we observe chord symbols such as they appear in music sheets [20] instead of observing the notes that would be played by a musician reading these chord symbols. As we noted in Section 1, chords can be seen as a latent

variable (local in time) that conditions the probabilities of choosing particular notes in other music components, such as melodies or accompaniments. The chord symbol “C Maj7” is usually constructed using the pitch classes C, E, G, and B. However, it really defines a conditional probability over *all* pitch classes. For instance, the pitch class D would normally be much more likely over this particular chord than the pitch-class Bb. Note that it is easy to infer valid chord symbols from the actual notes in most datasets with deterministic methods, which is done by most of the MIDI sequencers today. Hence, a model observing chord symbols instead of actual notes could still be used over traditional MIDI data with minimal preprocessing effort.

Four chord representations have been used in the experiments described in Section 3. First, what we call a **Naive** representation is to consider every chord (including the choice of the root) as a distinct observation. This representation has the disadvantage of excluding any notion of chord similarity. Moreover, this representation leads to a large number of states for the associated random variables (e.g. 152 in the current experiments, corresponding to each different chord found in the dataset described in Section 3). This can be harmful when learning over small datasets because of the high number of parameters. Despite all these drawbacks, such a representation can be useful if the notions of chord similarities are included in others parts of the models, such as in the conditional probabilities between variables [15].

Another possible chord representation is to discard any information except the root, yielding random variables with 12 possible values. While having a reasonable number of possible values, such a representation introduces a lot of smoothing in the models. It is possible to automatically detect the key and the mode (major or minor) of a song. Given that information, the root is very often sufficient to predict the whole structure of the rest of the current chord. For instance, given a song in C major and observing a root C, it is very likely that the complete chord associated to this root is a variant of C major.

We can also restrict ourselves to a subset of all possible chords [3, 19] by mapping more complex observed chord symbols to a subset of simpler ones. Such a representation is also used in the experiments described in this paper. We define only whether a chord is minor, major or dominant, leading to chord random variables with only 3 possible values. In this case, if we observe for instance the chord C7#5b9, we map this chord to the value corresponding to a dominant (7) chord. In the remaining of the paper, this representation is referred to as the **mM7** (**minor-Major-dominant 7th**) representation.

Finally, we can combine the **Root** representation and the **mM7** representation. This leads to discrete random variables with 36 possible values (12 roots times 3 chord qualities).

3 Comparing Chord Representations

53 jazz standards melodies [20] were recorded by the first author in MIDI format. Corresponding chord labels were also manually added to this corpus. The complexity of the chord sequences and melodies found in the corpus is representative of the complexity of common jazz and pop music. The songs were transposed to the key of C. This simplification has no impact on the generality of the experiments since automatic key detection is relatively reliable. Every song was 16 bars long, with four beats per bar, and with one chord change every two beats. The shortest notes in the dataset are half notes. Hence, two melodic observations were made for each beat, yielding observed sequences of length 128. Since chords span multiple melodic observations, their symbols were repeated multiple times (e.g. a 6 beat chord is represented as 12 distinct observations). This has no impact on the quality of the model since chords are acting as latent variables. Hence, the fact that chords are repeated in input sequences does not imply actual repetitions in the corresponding melodic signal.

3.1 Local Model

Out-of-sample classification errors for the local model described in Section 2.1 and for each of the chord representations given in Section 2.3 are presented in Table 1. 5-fold cross-validation [10] has been used

Table 1: Out-of-sample classification error and average negative log-likelihood for local models

Model	% of Error	Neg. Log-Likelihood
Naive	77.39	357.48
Roots + mM7	80.93	316.33
Roots	81.40	293.95
mM7	82.41	295.29
Freq	83.34	301.40

to calculate the prediction error rates. Let $x_{n,i}^t$ and $c_{n,i}^t$ be respectively the melodic observation and the chord observation in sequence n of the test set associated to the i -th fold of the cross-validation process at time t . Assume also that the i -th fold in the cross-validation process contains N_i test sequences, that there is a total of I folds in the cross-validation process and that all the sequences have length T . The rate of error is given by

$$\frac{1}{I} \sum_{i=1}^I \frac{1}{N_i} \sum_{n=1}^{N_i} \frac{1}{T} \sum_{t=1}^T d_{n,i,t} \quad (1)$$

where

$$d_{n,i,t} = \begin{cases} 1 & \text{if } (\operatorname{argmax}_{x \in X} p_{n,i,t}(x | c_{n,i}^1, \dots, c_{n,i}^t)) \neq x_{n,i}^t \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

with X being the set of possible melodic observations and $p_{n,i,t}(x | c_{n,i}^1, \dots, c_{n,i}^t)$ being the probability of observing melodic value x in sequence n of the test set associated to the i -th fold of the cross-validation process at time t , estimated by the evaluated model. Note that we condition over all previous chords in Equation (2), while a melodic prediction in the local model only depends on the current chord observation. This notation will also be useful for evaluating the IOHMM models presented in Section 2.2. In words, the out-of-sample error is just the average number of times the algorithm makes a mistake when trying to predict melodic observations over songs *unseen* during training. It should be pointed out that this criterion is really the one we want to minimize when developing models for practical applications. In such a context, the system *must* make a decision at each time step. As an example, a melodic model could be appended to a transcription algorithm. In such a context, the model would have to guess what is the most likely next note, given the previous notes and the audio measurements.

On each iteration of cross-validation, one fold of the dataset is not used for training. This subset of the dataset (referred to as the *test set* in machine learning literature) can be used to evaluate the model. Hence, it is always possible to observe a chord symbol during evaluation (or testing) that has not been observed previously when training the model. Dirichlet priors [11] have been used on all the chord variables in the algorithms described in this paper in order to avoid propagating infinite negative log-probabilities in the models in this case.

Parameters are tied over time in the **Local** model presented in Figure 1. In other words, the arrows between level 1 and level 2 always correspond to the same probability table in one fold of the cross-validation. As a benchmark, we also introduce the **Freq** model in Table 1, which is simply an algorithm that always selects the most frequent melodic observation in the training set. It corresponds to removing level 1 random variables in Figure 1, just leaving independent observations.

The *capacity* is the complexity of a function chosen from a parameterized family of functions. For good generalization, we normally want to find in this family the function with the minimal capacity that is able to model the training set. With the **Local** model, it is impossible to optimize capacity because no parameter is available which relates to the number of degrees of freedom in the model. Thus in the **Local** model, the only way to increase capacity is to use a chord representation with more possible values. Looking at the obtained results, we see that the rate of error increases when using

Table 2: Out-of-sample classification error and average negative log-likelihood for IOHMM models

Model	% of Error	Neg. Log-Likelihood
Naive	79.05	281.84
Roots	82.09	223.67
Roots + mM7	83.17	247.48
mM7	84.71	212.47
HMM	86.56	196.27

chord representations with *fewer* possible values. It is then possible that all these models somehow underfit, making the **Naive** representation (with 152 possible chord values in the current experiments) the best choice in this context. Not surprisingly, the **Freq** model, which always chooses the most frequent melodic observation, is the worst model in terms of classification error rate. Generalization capabilities of a model can benefit from the smoothing produced by a simplified chord representation, since this simplification is done by clustering perceptual properties of chords (e.g. all chords with the same root are clustered together in the **Root** representation). This is obviously not the case when using the **Freq** model.

Table 1 also shows the average negative out-of-sample log-likelihoods obtained with the same models again using cross-validation. Using the same notation as in Equation (1) and (2), this measure is given by

$$-\frac{1}{I} \sum_{i=1}^I \log \left(\prod_{n=1}^{N_i} \prod_{t=1}^T p_{n,i,t}(x_{n,i}^t | c_{n,i}^1, \dots, c_{n,i}^t) \right) \quad (3)$$

which is equal to

$$-\frac{1}{I} \sum_{i=1}^I \sum_{n=1}^{N_i} \sum_{t=1}^T \log(p_{n,i,t}(x_{n,i}^t | c_{n,i}^1, \dots, c_{n,i}^t)). \quad (4)$$

This performance measure is the one that was optimized over the training set when learning the parameters of the models with the EM algorithm.

As can be observed, the negative log-likelihood results are not coherent with the classification error. For instance, the **Freq** model has a lower negative log-likelihood than the **Roots + mM7** and **Naive** models! This result is counterintuitive since one would expect that adding current chord information would help the model to guess what would be the current melody note.

These discrepancies between classification error rate and negative log-likelihood could be explained by the fact that the terms of the sum on the right hand side of Equation (3) are not bounded negatively. Suppose that a model fits most of the data quite well but some of the out-of-sample examples have very low probabilities. Then, the terms associated to these examples in Equation (3) can take very large negative values that could dominate the average negative log-likelihood for all the examples. On the other hand, the cost of encountering a very unlikely out-of-sample sequence (with respect to the model being evaluated) in Equation (1) is only proportional to $\frac{1}{N}$, with N being the total number of examples in the dataset. This observation raises the following somewhat controversial question: Is the likelihood of the model over the observed data the best criterion to optimize when what we *really* want to do is to minimize the error of classification? We further discuss this question in Section 4.

3.2 IOHMM Model

Table 2 shows out-of-sample classification error and average negative log-likelihood for the IOHMM model (Section 2.2) These results are qualitatively similar to those in Table 1 for the **Local** model. This time, the number of possible values for hidden variables in level 2 of Figure 2 was optimized using 5-fold double cross-validation, which is a recursive application of cross-validation where both the optimization of the parameters of the model and the evaluation of the generalization performance

of the model are carried out simultaneously. Standard cross-validation was then applied to each subset of 4 folds with each hyper-parameter setting and tested with the best set of parameters (on average) on the remaining hold-out fold.

The same parameters are used over time to define the conditional probability distributions. For instance, all the vertical arrows between variables in level 1 and level 2 in Figure 2 represent the *same* probability table. The fact that it was possible in this context to optimize the capacity of the models by adjusting the number of states for the hidden variables makes the results in Table 2 more trustworthy than the ones found in Table 1, although they are similar. It should be pointed out that the capacity of the models was optimized with respect to the appropriate error measure. For instance, when reporting results about prediction error rates, capacity is optimized with respect to prediction error rate (while the models are trained by maximizing the likelihood with the EM algorithm).

The HMM model referred to in Table 2 is similar to the IOHMM model but removing the chord inputs layer. It can be represented by the model in Figure 1 with the horizontal dashed arrows being present. In this case, variables in level 1 are hidden and variables in level 2 are still melodic observations. As expected, the HMM model produces higher out-of-sample classification error than the IOHMM models, which can take advantage of the chord symbols given as inputs. Interestingly, the **Naive** representation for chords seems to be consistently efficient for melodic prediction. In Table 2, the **Naive** representation gives statistically significantly better results than the **Roots** representation with a confidence level of 99%². This is an indication that developing probabilistic models with this representation could be a viable approach, especially if perceived relations between chords are included in the models via the conditional probabilities related to these chords [15]. The representation including only the roots also performs well. Given these results, this representation appears to be a good compromise given its relative simplicity and the fact that it inherently embodies perceptually relevant smoothing. Using basic chord information (**mM7** representation) does not seem to help with respect to unconditional classification error.

Again, average negative log-likelihoods contradict average classification error rates. Even worse, the HMM model performs much better than the IOHMM models in terms of likelihood! This is an indication that such a measure favors models that are more uniform in essence, thus giving a relatively high probability to unseen sequences. However, such models are weaker when asked to make a decision, meaning that they define distributions with modes less precisely adapted to the training data.

3.3 Conditional Classification Error

The goal of the models presented here is to predict the melodies in the dataset. It is out of the scope of this work to evaluate the subjective artistic quality of the predicted melodies. A more interesting measure of melodic prediction is the out-of-sample *conditional* classification error, given by Equation (1), but using instead

$$d_{n,i,t} = \begin{cases} 1 & \text{if } (\operatorname{argmax}_{x \in X} p_{n,i,t}(x | c_{n,i}^1, \dots, c_{n,i}^t, x_{n,i}^1, \dots, x_{n,i}^{t-1})) \neq x_{n,i}^t \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

This measure is very similar to the unconditional error described in Section 3.2. However, the models have access to the true previous melodic observations when trying to guess the next one.

The only objective performance measure we can provide about a melodic prediction given a chord progression is to tell if a predicted melody is similar or not to the one provided in the dataset with the same chord progression. However, while the space of plausible melodies is huge, we only have access to a very small number of sequences to measure the performance of the models given a chord progression. Moreover, given a particular sequence of chords, one can imagine that a very high number of melodies would be considered more or less musically similar to the ones in the dataset. Among all these melodies, some of them may not share a single note with the true melody associated with this

²We used a standard proportion test, assuming a binomial distribution for the errors and using a normal approximation.

Table 3: Out-of-sample conditional classification error rates for IOHMM models

Model	% of Error
Roots	57.41
Roots + mM7	58.21
mM7	58.32
Naive	69.77
HMM	85.27

sequence of chords in the test set. A good melodic prediction model would be likely to generate any one of these melodies, thus producing a very high unconditional error rate.

The conditional error rate alleviates this problem by measuring the prediction performance of a model in regions of the observation space where data is present, leading to a much more reliable performance measure in this context. Moreover, distributions that would generalize well according to such a measure could also be sampled to generate realistic melodies given chord progressions and initial melodic motives. The Junction Tree Algorithm for marginalization [13] allows us to fix the state of some variables in the graphical model while finding the most probable states for other variables. Out-of-sample conditional classification error rates for the IOHMM models presented in Section 2.2 are shown in Table 3. We do not provide conditional classification error rates for the `Local` model described in Section 2.1, because they would be identical to the unconditional classification error rates shown in Table 1. When making a prediction, such models are completely unaware of previous observations in time. Also, we do not provide average negative log-likelihoods, since this measure is not well adapted to prediction tasks, as we noted in Section 3.

Again, the HMM model produces higher out-of-sample conditional classification error rate than the IOHMM models which benefit from chord symbols given as inputs. In Table 3, the conditional error rates are much lower than the unconditional ones for the same models. For each chord representation, the prediction accuracy gained when observing previous melodic notes is much higher than the differences in error rates for each chord representations obtained in Table 2. This means that observing previous melodic notes gives more information about the likely choices for the current melody than any chord information.

In Table 2, the `Naive` representation was the best one in terms of unconditional prediction error rate. On the other hand, this representation has the highest conditional prediction error rate among all the IOHMMs. When knowing nothing about the previous melodic observations, the model performs better when provided with a more detailed chord representation. However, given previous melodic observations, smoothed chord representations lead to better generalization in terms of prediction error rate. The `Naive` representation overfits the training data because it leads to models with higher capacity. Finally, no representation is statistically significantly better than another with a confidence level of 90% among the three best representations in Table 3.

4 Conclusions

The main motivation behind this paper was to better understand the statistical relations between chord representations and the actual choice of notes in polyphonic music. To this end, we compared four chord representations using melodic prediction as a benchmark task. Surprisingly, the `Naive` representation where each chord is conceived as a discrete observation apparently performs well in terms of unconditional prediction error rates. Nevertheless, this representation overfits when past melodic observations are used to condition the predictions. In this case, smoothed chord representations seems more appropriate. Given the obtained results, representing chords only by their roots seems to be a good compromise, especially when all the songs to be analyzed are transposed to the same key. While being extremely simple, this representation inherently includes smoothing related to perceptual

relations between notes. The `Root+mM7` representation that is used in some important music information retrieval papers [3, 18] is not optimal in terms of out-of-sample classification error and average negative log-likelihood for both probabilistic models presented in this paper. However, in practice, the actual choice of a chord representation should always be made considering the application to be developed.

An interesting observation when looking at the results of the experiments done in Section 3 is that the behavior of the average out-of-sample likelihood does not follow the trends of the average prediction error rate (conditional or unconditional). On the one hand, the likelihood is a measure of the fit of a whole distribution to a dataset. However, the classification error seems to be a better descriptor of the fit of the *modes* of a distribution. Provided a nearly infinite amount of data, these two measures would lead to the same ranking of the models. Thus, likelihood and prediction error would probably be more comparable when measured with models trained and evaluated with much more data. How much data is needed in this particular framework to obtain such a behavior is still an open question that needs to be addressed. However, given realistic datasets, optimizing the likelihood of a model with respect to training data may not be the best strategy when one is only interested in the modes of the distribution, which can be the case when doing prediction. An alternative learning strategy would be to maximize the sum of the differences between the probabilities of the observed classes and the probabilities of the most probable wrong classes instead of just maximizing the sum of the log probabilities of the observed classes. This approach is referred to as the minimum classification error (MCE) algorithm [12].

Further development of the models proposed in this paper are possible. First, it would be interesting to append these models to existing music information retrieval algorithms to increase their performance. As a comparison, language models are regularly used in speech recognition algorithms to constrain their search space to reasonable solutions [17].

Acknowledgment

The authors would like to thank Sean Wood for useful discussions. This work was supported in part by the IST Program of the European Community, under the PASCAL Network of Excellence, IST-2002-506778, funded in part by the Swiss Federal Office for Education and Science (OFES) and the Swiss NSF through the NCCR on IM2.

References

- [1] M. Allan, C. K. I. Williams, Harmonising chorales by probabilistic inference, in: *Advances in Neural Information Processing Systems*, vol. 17, 2004.
- [2] R. Begleiter, R. El-Yaniv, G. Yona, On prediction using variable order Markov Models, *Journal of Artificial Intelligence Research* 22 (2004) 385–421.
- [3] J. P. Bello, J. Pickens, A robust mid-level representation for harmonic content in music signals, in: *Proceedings of the Sixth International Conference on Music Information Retrieval*, London, 2005.
- [4] Y. Bengio, P. Frasconi, Input/output HMMs for sequence processing, *IEEE Transactions on Neural Networks* 7 (5) (1996) 1231–1249.
- [5] Y. Bengio, P. Simard, P. Frasconi, Learning long-term dependencies with gradient descent is difficult, *IEEE Transactions on Neural Networks* 5 (2) (1994) 157–166.
- [6] A. T. Cemgil, Bayesian music transcription, Ph.D. thesis, Radboud University of Nijmegen (2004).

- [7] C. Dahlhaus, *Studies on the Origin of Harmonic Tonality*, Princeton University Press, 1990, XV-389 p.
- [8] A. P. Dempster, N. M. Laird, D. B. Rubin, Maximum likelihood from incomplete data via the EM algorithm, *Journal of the Royal Statistical Society* 39 (1977) 1–38.
- [9] D. Eck, J. Schmidhuber, Finding temporal structure in music: Blues improvisation with LSTM recurrent networks, in: H. Bourlard (ed.), *Neural Networks for Signal Processing XII, Proc. 2002 IEEE Workshop*, IEEE, New York, 2002.
- [10] T. Hastie, R. Tibshirani, J. Friedman, *The Elements of Statistical Learning*, Springer series in statistics, Springer-Verlag, 2001.
- [11] D. Heckerman, D. Geiger, M. Chickering, *Learning Bayesian networks: The combination of knowledge and statistical data*, Tech. rep., Microsoft Research (1994).
- [12] B.-H. Juang, S. Katagiri, Discriminative learning for minimum error classification, in: *IEEE Trans. on Signal Processing*, vol. 10 (12), 1992.
- [13] S. L. Lauritzen, *Graphical Models*, Oxford University Press, 1996.
- [14] V. Lavrenko, J. Pickens, Polyphonic music modeling with random fields, in: *Proceedings of ACM Multimedia*, Berkeley, CA, 2003.
- [15] J.-F. Paiement, D. Eck, S. Bengio, A probabilistic model for chord progressions, in: *Proceedings of the 6th International Conference on Music Information Retrieval*, 2005.
- [16] J.-F. Paiement, D. Eck, S. Bengio, Probabilistic melodic harmonization, in: *Proceedings of the 19th Canadian Conference on Artificial Intelligence*, Springer, 2006.
- [17] L. Rabiner, B.-H. Juang, *Fundamentals of speech recognition*, 1st ed., Prentice All, 1993.
- [18] C. Raphael, J. Stoddard, Harmonic analysis with probabilistic graphical models, *Computer Music Journal* 28 (3) (2004) 45–52.
- [19] A. Sheh, D. P. Ellis, Chord segmentation and recognition using EM-trained Hidden Markov Models, in: *Proceedings of the 4th ISMIR*, Baltimore, Maryland, 2003.
- [20] C. Sher (ed.), *The New Real Book*, vol. 1-3, Sher Music Co., 1988.
- [21] D. Temperley, A probabilistic model of music perception, in: *Proceedings of the 7th Conference on Music Information Retrieval*, Victoria, Canada, 2006.
- [22] B. Thom, Predicting chords in jazz: the good, the bad, and the ugly, in: *IJCAI-95, Music and AI Workshop*, Montreal, Canada, 1995.