


A First Look at Music Composition using LSTM Recurrent Neural Networks

Douglas Eck
doug@idsia.ch

Jürgen Schmidhuber
juergen@idsia.ch



Technical Report No. IDSIA-07-02

IDSIA / USI-SUPSI
Istituto Dalle Molle di studi sull' intelligenza artificiale
Galleria 2
CH-6900 Manno, Switzerland

A First Look at Music Composition using LSTM Recurrent Neural Networks

Douglas Eck
doug@idsia.ch

Jürgen Schmidhuber
juergen@idsia.ch

Abstract

In general music composed by recurrent neural networks (RNNs) suffers from a lack of global structure. Though networks can learn note-by-note transition probabilities and even reproduce phrases, attempts at learning an entire musical form and using that knowledge to guide composition have been unsuccessful. The reason for this failure seems to be that RNNs cannot keep track of temporally distant events that indicate global music structure. Long Short-Term Memory (LSTM) has succeeded in similar domains where other RNNs have failed, such as timing & counting and CSL learning. In the current study we show that LSTM is also a good mechanism for learning to compose music. We compare this approach to previous attempts, with particular focus on issues of data representation. We present experimental results showing that LSTM successfully learns a form of blues music and is able to compose novel (and we believe pleasing) melodies in that style. Remarkably, once the network has found the relevant structure it does not drift from it: LSTM is able to play the blues with good timing and proper structure as long as one is willing to listen.

1 Introduction

The most straight-forward way to compose¹ music with an RNN is to use the network as single-step predictor. The network learns to predict notes at time $t + 1$ using notes at time t as inputs. After learning has been stopped the network can be seeded with initial input values—perhaps from training data—and can then generate novel compositions by using its own outputs to generate subsequent inputs. This note-by-note approach was first examined by Todd (1989); Bharucha and Todd (1989) and later used by others (Stevens and Wiles, 1994; Mozer, 1994).

A feed-forward network would have no chance of composing music in this fashion. Lacking the ability to store any information about the past, such a network would be unable to keep track of where it is in a song. In principle an RNN does not suffer from this limitation. With recurrent connections it can use hidden layer activations as memory and thus is capable of exhibiting (seemingly arbitrary) temporal dynamics. In practice, however, RNNs do not perform very well at this task. As Mozer (1994) aptly wrote about his attempts to compose music with RNNs, “While the local contours made sense, the pieces were not musically coherent, lacking thematic structure and having minimal phrase structure and rhythmic organization”.

The reason for this failure is likely linked to the problem of *vanishing gradients* (Hochreiter et al., 2001) in RNNs. In gradient methods such as Back-Propagation Through Time (BPTT) (Williams and Zipser, 1995) and Real-Time Recurrent Learning (RTRL) (Robinson and Fallside, 1987) error flow either vanishes quickly or explodes exponentially, making it impossible for the networks to deal correctly with long-term dependencies. In the case of music, long-term dependencies are at the heart of what defines a particular style, with events spanning several notes or even many bars contributing to the formation of metrical and phrasal structure (Cooper and Meyer, 1960). The

¹In this paper we use the terms *composition* and *improvisation* loosely. It is probably more accurate to describe the behavior of the network as improvisation because it is inventing new melodies on top of a set form; however, the end goal is the creation of new melodies and new forms, thus the use of the term *composition*.

clearest example of these dependencies are chord changes. In a musical form like early rock-and-roll music for example, the same chord can be held for four bars or more. Even if melodies are constrained to contain notes no shorter than an eighth note, a network must regularly and reliably bridge time spans of 32 events or more.

The most relevant previous research is that of Mozer (1994), who did note-by-note composition of single-voice melodies accompanied by chords. In the “CONCERT” model, Mozer used sophisticated RNN procedures including BPTT, log-likelihood objective functions and probabilistic interpretation of the output values. In addition to these neural network methods, Mozer employed a psychologically-realistic distributed input encoding (Shepard, 1982) that gave the network an inductive bias towards chromatically and harmonically related notes. He used a second encoding method (Laden and Keefe, 1989) to generate distributed representations of chords.

As was mentioned above, a BPTT-trained RNN does a poor job of learning long-term dependencies. To offset this, Mozer used a distributed encoding of duration that allowed him to process a note of any duration in a single network timestep. By representing in a single timestep a *note* rather than a slice of *time*, the number of time steps to be bridged by the network in learning global structure is greatly reduced. For example, to allow sixteenth notes in a network which encodes slices of time directly requires that a whole note span at minimum 16 time steps. One CONCERT composition (Figure 8 in the paper) contained only 76 notes but would require 172 time steps presuming the sixteenth notes found in the piece are to be allowed.

Even with sophisticated RNN techniques and psychologically-realistic distributed representation of melody, chords and duration the CONCERT architecture failed to capture global musical structure². Though networks regularly outperformed third-order transition table approaches, they failed in all cases to find global structure. In analyzing this performance Mozer suggests that, for the note-by-note method to work it is necessary that the network can induce structure at multiple levels. We agree and offer the following architecture as one possible solution.

2 An LSTM Music Composer

2.1 LSTM Architecture

Due to space constraints it is impossible to describe LSTM in depth. See Gers and Schmidhuber (2001); Gers et al. (2000) for details. In summary, LSTM is designed to obtain constant error flow through time and to protect this error flow from undesired perturbations. LSTM uses linear units called *Constant Error Carousels* (CECs) to overcome error decay problems plaguing previous RNNs (Hochreiter, 1991; Hochreiter et al., 2001). Each CEC has a fixed self-connection and is surrounded by a cloud of nonlinear units responsible for controlling the flow of information in and out of the CEC. Typically, a multiplicative *input gate unit* learns to protect the flow from perturbation by irrelevant inputs. Likewise, a multiplicative *output gate unit* learns to protect other units from perturbation by currently irrelevant memory contents. A *forget gate* learns to reset a memory cell when its content is obsolete. Learning is done by gradient descent method that uses a slightly modified, truncated BPTT and a customized version of RTRL. Output units use BPTT; output gates use the truncated version of BPTT; while weights to cells, input gates and forget gates use truncated RTRL. LSTM performance is improved in online learning situations by using a Kalman filter to control weight updates (Pérez-Ortiz et al., 2002).

2.2 Data Representation

We avoid psychologically realistic distributed encodings and instead represent the data in a simple local form (similar to (Todd, 1989)). We use one input/target unit per note, with 1.0 representing on and 0.0 representing off. (In later experiments we used the common technique of adjusting input units to have a mean of 0 and a standard deviation of 1.0.) Unlike CONCERT this representation

²Mozer cannot be faulted for overstating the achievements of the model. He suggest that the “ER” in CONCERT is meant to stand for “ERudite” but that it could also mean “ERSatz” or ERratic’.

leaves it to the network to learn an inductive bias towards chromatically and harmonically related notes. Despite this, we prefer a localist representation for several reasons. First it is implicitly multi-voice and makes no artificial distinction between chords and melodies. (In fact, we implement chords by simply turning on the appropriate notes in the input vector.) Second it is an easy task to generate probability distributions over the set of possible notes, with the flexibility to treat single note probabilities as independent or dependent from one another.

The representation of time is straightforward, with one input vector representing a slice of real time. The stepsize of quantization can of course vary; if the quantization is set at the eighth note level (as it is for all experiments in this study) then a eight network time steps are required to process a whole note. This method is preferable for LSTM because it forces the network to learn the relative durations of notes, making it easier for the counting and timing mechanisms to work (Gers and Schmidhuber, 2000).

Two representational issues are ignored in this representation. First, there is no explicit way to determine when a note ends. This means that eight eighth notes of the same pitch are represented exactly the same way as, say, four quarter notes of the same pitch. One way to implement this without changing input and target data structures is to decrease the stepsize of quantization and always mark note endings with a zero. With this method, a quantization level of sixteen steps per whole note would generate unique codes for eight eighth notes and four quarter notes of the same pitch. A second method is to have special unit(s) in the network to indicate the beginning of a note. This method was employed by Todd (1989) and is certainly viable. However, it is not clear how such a method would scale to data sets with multi-voice melodies.

In simulations for this study, a range of 12 notes were possible for chords and 13 notes were possible for melodies Figure 1. Though we divided chord notes from melody notes for these simulations, this division is artificial: Chord notes are represented no differently than melody notes and in future experiments we intend to blend the two in a more realistic manner.




Figure 1: Possible note values for these simulations

2.3 Training Data

For the experiments in this study, a form of 12-bar blues popular among bebop jazz musicians is used. With a quantization stepsize of 8 notes per bar, this yields a single song length of 96 network time steps. The chords used did not vary from song to song and are shown in Figure 2. Chords inversions were chosen so that the chords would fit into the allocated range of notes. For Experiment 1, only these chords were presented. For Experiment 2, a single melody line was



Figure 2: Chords for training data (transposed up one octave).

presented along with the chords. The melody line was built using the pentatonic scale (Figure 3) commonly used in this style of music. Training melodies were constructed by concatenating bar-long segments of music written by the first author to fit musically with each chord. Datasets were



Figure 3: Pentatonic scale used for training data melodies.

constructed by choosing randomly from the space of unique complete pieces ($n = 2^{12} = 4096$). Only quarter notes were used. No rests were used. Space constraints make it impossible to include examples. However several of these training pieces are provided as sheet music (Acrobat .pdf) and audio (MIDI, MP3 and wav) at <http://www.idsia.ch/~doug/blues/index.html>.

In Figure 8.1 at the end of this paper is an example of the melodies and chords presented to the network.

3 Experiment 1 — Learning Chords

In this experiment we show that LSTM can learn to reproduce a musical chord structure. My motivation is to ensure that the chord structure of the song can in fact be induced in absence of its melody. Otherwise it is unclear whether LSTM is taking advantage of the local structure in melody to predict global chord structure. This is especially a risk when input data is generated using a random recombination of a relatively small number of musical examples, as was done here.

3.1 Network Topology and Experimental Parameters

The chords used are the ones described in Section 2.3. No melodies are presented. The quantization timestep is eight events per whole note. In the network four cell blocks containing 2 cells each are fully connected to each other and to the input layer. The output layer is fully connected to all cells and to the input layer. Forget gate, input gate and output gate biases for the four blocks are set at -0.5, -1.0, -1.5 and -2.0. This allows the blocks to come online one by one. Output biases were set at 0.5. Learning rate was set at .00001. Momentum rate (Plaut et al., 1986) was set at .9. Weights being burned after every timestep. Experiments showed that learning was faster if the network was reset after making one (or a small number) of gross errors. Resetting went as follows: on error, burn existing weights, reset the input pattern and clear partial derivatives, activations and cell states. Gers and Schmidhuber (2000) use a similar strategy. The squashing function at the output layer was the logistic sigmoid with range [0,1].

3.2 Training and Testing

The goal was to predict at the output the probability for a given note to be on or off. For predicting probabilities root mean squared error (RMSE) is not appropriate. Instead the network was trained using cross-entropy as the objective function. The error function E_i for output activation y_i and target value t_i is

$$E_i = -t_i \ln(y_i) - (1 - t_i) \ln(1 - y_i) \quad (1)$$

This yields a δ term at the output layer of $(t_i - y_i)$. See, e.g., Joost and Schiffmann (1998) for details. Note that by a series of binomial formulations rather than a single multinomial formulation (softmax) we are treating outputs as statistically independent of one another. Though this assumption is untrue, it allows the network to predict chords and melodies in parallel and also allows for multi-voice melodies.

The network was tested by starting it with the inputs from the first timestep and then using network predictions for ensuing time steps. Chord notes were predicted using a decision threshold of 0.5. Training was stopped after the network successfully predicted the entire chord sequence.

3.3 Results

LSTM easily handled this task under a wide range of learning rates and momentum rates. Once a network could successfully generate one full cycle through the chord sequence, it could generate any number of continuing cycles. This indicates that there was no reason to continue learning for a longer time. As it is already well documented that LSTM excels at timing and counting tasks (Gers and Schmidhuber, 2000), success at this task is not surprising. Fast convergence was not a goal of this study, and learning times were not carefully collected. However, even with no attempt was made to optimize learning rates, this task was learned relatively quickly (on the order of 15 minutes to 45 minutes of processor time on a 1Ghz Pentium depending on parameter settings and initial random weights)

4 Experiment 2 — Learning Melody and Chords

In this experiment both melody and chords are learned. Learning continues until the chord structure is learned and cross-entropy error is relatively low. Note that there are far too many melodies for the network to learn them all. Once learning has been stopped, the network is started with a seed note or series of notes and then allowed to compose freely. The goal of the study was to see if LSTM could learn chord structure and melody structure and then use that structure to advantage when composing new songs.

4.1 Network Topology and Experimental Parameters

The network topology for this experiment differs from the previous task in that some cell blocks processed chord information while other cell blocks processed melody information. Eight cell blocks containing 2 cells each are used. Four of the cell blocks are fully connected to the input units for chords. The other four cell blocks are fully connected to the input units for melody. The chord cell blocks have recurrent connections to themselves and to the melody cell blocks. However, melody cell blocks are only recurrently connected to melody cell blocks. That is, melody information does not reach the cell blocks responsible for processing chords. At the output layer, output units for chords are fully connected to cell blocks for chords and to input units for chords. Output units for melody are fully connected to cell blocks for melody and to input units for melody. The implications of such a topology are discussed below in Section 5. Forget gate, input gate and output gate biases for the four blocks dedicated to processing chords are set at -0.5, -1.0, -1.5 and -2.0. Gates for processing melodies are biased in exactly the same way. All other parameters are identical to those described for Experiment 1 in Section 3.1.

4.2 Training and Testing

The goal was to predict at the output the probability for a given note to be on or off. For chords, the same method as Experiment 1 is used: the network applies a decision threshold of 0.5 for all chord notes. For melodies we restrict the network to choosing a single note at any given timestep. This is achieved by adjusting melody output activations so that they sum to 1.0 and then using a uniform random number in the range [0,1] to choose the appropriate next note. The implications of this decision are discussed below in Section 5.

The network was trained until it had learned the chord structure and until objective error had reached a plateau. Then the network was allowed to freely compose music. Music was composed by providing a single note or series of notes (up to 24) from the training set as input. After those were presented, network outputs were presented as inputs at the next timestep. No algorithmic or statistical method was used to evaluate the musical quality of the network. In all cases the network succeeded in reproducing chord structure while in parallel improvising new melodies.

4.3 Results

LSTM composed music in the form of blues. It learned the chord structure in training and used that structure to constrain its melody output in composition mode. Because it is difficult to evaluate the performance objectively—this point is commonly made in AI art research, e.g., (Mozer, 1994)—we urge the reader to visit <http://www.idsia.ch/~doug/blues/index.html>. On that page are examples of network blues composition in sheet music form (Acrobat .pdf) and audio (MIDI, MP3 and wav). It can be said that the network compositions are remarkably better sounding than a random walk across the pentatonic scale (this is also available for listening at the website). Unlike a random walk, the network compositions follow the structure of the musical form. They do diverge from the training set, sometimes significantly. But due to the influence of the chord structure, they never “drift” away from the form: the chord changes always bring it back. Also, an informal survey in our lab indicates that the compositions are at times quite pleasant. One jazz musician³ is struck by how much the compositions sound like real bebop jazz improvisation over this same chord structure. In particular, the network’s tendency to intermix snippets of known melodies with less-constrained passages is in keeping with this style.

5 Discussion

These experiments were successful: LSTM induced both global structure and local structure from a corpus of musical training data, and used that information to compose in the same form. This answers Mozer’s (Mozer, 1994) key criticism of RNN music composition, namely that an RNN is unable to compose music having global coherence. To my knowledge the model presented in this paper is the first to accomplish this.

That said, several parts of the experimental setup made the task easier for the model. More research is required to know whether the LSTM model can deal with more challenging composition tasks.

5.1 Training Data

There was no variety in the underlying chord structure. For this reason it is perhaps better to talk about network performance as improvisation over a predefined (albeit learned) form rather than composition. This lack of variation made it easier for LSTM to generate appropriately-timed chord changes. Furthermore, quantization stepsize for these experiments was rather low, at 8 time steps per whole note. As LSTM is known to excel at datasets with long time lags, this does not pose a serious problem. However it remains to be seen how much more difficult the task will be at, say, 32 time steps per whole note, a stepsize which would allow two sixteenth notes to be uniquely identified from a single eighth note.

subsectionNetwork Architecture There network connections were divided between chords and melody, with chords influencing melody but not vice-versa. We believe this choice makes sense: in real music improvisation the person playing melody (the soloist) is for the most part following the chord structure supplied by the rhythm section. However this architectural choice presumes that we know ahead of time how to segment chords from melodies. When working with jazz sheet music, chord changes are almost always provided separately from melodies and so this does not pose a great problem. Classical music compositions on the other hand make no such explicit division. Furthermore in an audio signal (as opposed to sheet music) chords and melodies are mixed together.

These are preliminary experiments, and much more research is warranted. A comparison with BPTT and RTRL (and other candidate RNNs) would help verify the claim that LSTM performance is better. A more interesting training set would allow for more interesting compositions. Finally, recent evidence suggests (Gers et al., 2002) that LSTM works better in similar situations using a Kalman filter to control weight updates. This should be explored. Finally, the current architecture

³Admittedly, this musician isn’t particularly good, and also happens to be the first author.

is limited to working with symbolic representations (i.e. modified sheet notation) of music. If the architecture were to be extended to handle real-time performed music (i.e. MIDI or audio) it would have potential as a tool for interactive improvisation. This would require an ability to deal with temporal noise found in real performed music. One possibility is to apply research on oscillator beat tracking models (Eck, 2002, 2001) to LSTM in order to create an inductive bias towards coupling with sometimes-noisy rhythmical elements in the input.

6 Conclusion

A music composition model based on LSTM successfully learned the global structure of a musical form, and used that information to compose new pieces in the form. Two experiments were performed. The first verified that LSTM was not relying on regularities in the melody to learn the chord structure. The second experiment explored the ability for LSTM to generate new instances of a musical form, in this case a bebop-jazz variation of standard 12-bar blues. These experiments are preliminary and much more work is warranted. However by demonstrating that an RNN can capture both the local structure of melody and the long-term structure of a of a musical style, these experiments represent an advance in neural network music composition.

7 Acknowledgements

We would like to thank Mike Mozer for answering many questions about his model and simulation methods. The division of labor between authors is as follows: The first author devised and constructed the datasets, implemented the program code, ran the simulations and wrote the paper. The second author wrote the grant proposal and eventually obtained a grant (SNF 21-49144.96) for doing this type of work. He provided guidance on how to structure the task and how to best use LSTM; he also edited the paper. The research referred to in this paper has been submitted to the ICANN-2002 conference (Eck and Schmidhuber, 2002).

8 Musical Examples

The reader is referred to <http://www.idsia.ch/~doug/blues/index.html> for a full set of examples.

8.1 Sample Training Input to Network

The image displays a musical score for a blues piece in C major, consisting of ten staves of music. Each staff contains a sequence of notes and rests, with chord symbols written above. The chord symbols are: C, F7, C, Gm7 C, F7, Cdim (Staff 1); C, Em A7, Dm7, G7, C, A7, Dm G7 (Staff 2); C, F7, C, Gm7 C, F7, Cdim (Staff 3); C, Em A7, Dm7, G7, C, A7, Dm G7 (Staff 4); C, F7, C, Gm7 C, F7, Cdim (Staff 5); C, Em A7, Dm7, G7, C, A7, Dm G7 (Staff 6); C, F7, C, Gm7 C, F7, Cdim (Staff 7); C, Em A7, Dm7, G7, C, A7, Dm G7 (Staff 8); C, F7, C, Gm7 C, F7, Cdim (Staff 9); C, Em A7, Dm7, G7, C, A7, Dm G7 (Staff 10).

References

- Bharucha, J. J. and Todd, P. M. (1989). Modeling the perception of tonal structure with neural nets. *Computer Music Journal*, 13(4):44–53.
- Cooper, G. and Meyer, L. B. (1960). *The Rhythmic Structure of Music*. The University of Chicago Press.
- Eck, D. (2001). A network of relaxation oscillators that finds downbeats in rhythms. In Dorffner, G., editor, *Artificial Neural Networks – ICANN 2001 (Proceedings)*, pages 1239–1247, Berlin. Springer.
- Eck, D. (2002). Finding downbeats with a relaxation oscillator. *Psychological Research*. In press.
- Eck, D. and Schmidhuber, J. (2002). Learning the long-term structure of the blues. In *Artificial Neural Networks – ICANN 2002 (Proceedings)*. Submitted.
- Gers, F. A., Perez-Ortiz, J., Eck, D., and Schmidhuber, J. (2002). DEKF-LSTM. In *Proc. 10th European Symposium on Artificial Neural Networks, ESANN 2002*.
- Gers, F. A. and Schmidhuber, J. (2000). Recurrent nets that time and count. In *Proc. IJCNN'2000, Int. Joint Conf. on Neural Networks*, Como, Italy.
- Gers, F. A. and Schmidhuber, J. (2001). LSTM recurrent networks learn simple context free and context sensitive languages. *IEEE Transactions on Neural Networks*, 12(6):1333–1340.
- Gers, F. A., Schmidhuber, J., and Cummins, F. (2000). Learning to forget: Continual prediction with LSTM. *Neural Computation*, 12(10):2451–2471.
- Hochreiter, S. (1991). Untersuchungen zu dynamischen neuronalen Netzen. Diploma thesis, Institut für Informatik, Lehrstuhl Prof. Brauer, Technische Universität München. See www7.informatik.tu-muenchen.de/~hochreit.
- Hochreiter, S., Bengio, Y., Frasconi, P., and Schmidhuber, J. (2001). Gradient flow in recurrent nets: the difficulty of learning long-term dependencies. In Kremer, S. C. and Kolen, J. F., editors, *A Field Guide to Dynamical Recurrent Neural Networks*. IEEE Press.
- Joost, M. and Schiffmann, W. (1998). Speeding up backpropagation algorithms by using cross-entropy combined with pattern normalization. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 6(2):117–126.
- Laden, B. and Keefe, D. H. (1989). The representation of pitch in a neural net model of chord classification. *Computer Music Journal*, 13(4):44–53.
- Mozer, M. C. (1994). Neural network composition by prediction: Exploring the benefits of psychophysical constraints and multiscale processing. *Cognitive Science*, 6:247–280.
- Pérez-Ortiz, J. A., Gers, F. A., Eck, D., and Schmidhuber, J. (2002). Kalman filters improve lstm network performance in problems unsolvable by traditional recurrent nets. *Neural Networks*. Accepted pending minor revisions.
- Plaut, D. C., Nowlan, S. J., and Hinton, G. E. (1986). Experiments on learning back propagation. Technical Report CMU-CS-86-126, Carnegie-Mellon University, Pittsburgh, PA.
- Robinson, A. J. and Fallside, F. (1987). The utility driven dynamic error propagation network. Technical Report CUED/F-INFENG/TR.1, Cambridge University Engineering Department.
- Shepard, R. N. (1982). Geometrical approximations to the structure of pitch. *Psychological Review*, 89:305–333.

- Stevens, C. and Wiles, J. (1994). Representations of tonal music: A case study in the development of temporal relationship. In Mozer, M., Smolensky, P., Touretsky, D., Elman, J., and Weigend, A. S., editors, *Proceedings of the 1993 Connectionist Models Summer School*, pages 228–235. Erlbaum, Hillsdale, NJ.
- Todd, P. M. (1989). A connectionist approach to algorithmic composition. *Computer Music Journal*, 13(4):27–43.
- Williams, R. J. and Zipser, D. (1995). Gradient-based learning algorithms for recurrent networks and their computational complexity. In Chauvin, Y. and Rumelhart, D. E., editors, *Back-propagation: Theory, Architectures and Applications*, chapter 13, pages 433–486. Hillsdale, NJ: Erlbaum.