An Intelligent Musical Rhythm Variation Interface

Richard Vogl
Department of Computational Perception
Johannes Kepler University
Linz, Austria
richard.vogl@jku.at

Peter Knees
Department of Computational Perception
Johannes Kepler University
Linz, Austria
peter.knees@jku.at

IUI'16 Companion, March 7–10, 2016, Sonoma, CA, USA.
ACM 978-1-4503-4140-0/16/03.
http://dx.doi.org/10.1145/2876456.2879471

Abstract
The drum tracks of electronic dance music are a central and style-defining element. Yet, creating them can be a cumbersome task, mostly due to lack of appropriate tools and input devices. In this work we present an artificial-intelligence-powered software prototype, which supports musicians composing the rhythmic patterns for drum tracks. Starting with a basic pattern (seed pattern), which is provided by the user, a list of variations with varying degree of similarity to the seed pattern is generated. The variations are created using a generative stochastic neural network. The interface visualizes the patterns and provides an intuitive way to browse through them. A user study with ten experts in electronic music production was conducted to evaluate five aspects of the presented prototype. For four of these aspects the feedback was generally positive. Only regarding the use case in live environments some participants showed concerns and requested safety features.

Author Keywords
Rhythm pattern generation; restricted Boltzmann machines; machine learning; neural networks; generative stochastic models.

ACM Classification Keywords
H.5.2 [User Interfaces]: Graphical user interfaces, Input devices and strategies

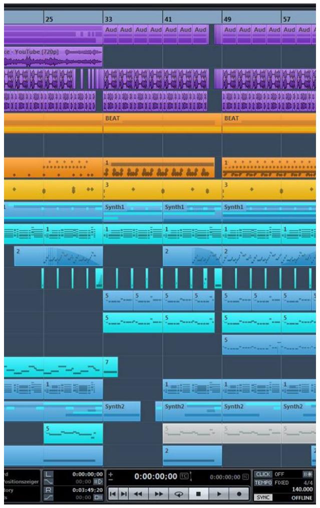

Figure 1: An EDM track being arranged in a DAW software. DAWs are programs used to produce music. The horizontal colored bars represent the tracks for instruments. The orange track contains the basic rhythmic pattern and its variations during build-ups, breaks, and fills.

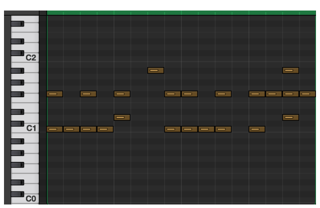

Figure 2: A piano roll editor showing a manually entered rhythm pattern.

Introduction

Nowadays, more than ever before, digital tools for music production play an important role in the workflow of music producers. Such tools cover applications like digital audio workstations (DAWs; see figure 1), integrated hardware/software solutions like grooveboxes, and software tools and plugins like synthesizers and audio effects. The *GiantSteps*[1] project focuses on simplifying the workflow of music producers by developing *intelligent agents* for the usage in electronic dance music (EDM) production and performance.

As yet, drum tracks are built by arranging rhythm patterns from a pattern library, or by creating patterns manually. Using predefined patterns bears the risk of sounding unoriginal, while creating them manually is a time consuming task and requires more musical knowledge. Entering rhythm patterns in a DAW is done using a mouse or MIDI controllers (keyboards and drum pads) to set notes in a piano roll (see figure 2) or similar editor. Step-sequencer-like interfaces are usually a feature of grooveboxes and drum machines and are typically found in setups for live performances.

When it comes to samplers and synthesizers for drums in EDM, a wide variety of commercial products as well as a lively research community exist. However, there are few works on automated drum rhythm variation and creation. In the works of Kaliakatsos–Papakostas et al. [2] and Ó Nuanáin et al. [4] genetic algorithms to generate rhythmic patterns are used. Genetic algorithms tend to produce random variations and the results strongly depend on the used fitness function.

Restricted Boltzmann machines (RBM, introduced in [5]) form a group of generative stochastic neural networks which are well suited for pattern generation. Battenberg et al. [1] use a variant, the conditional RBM, to classify the meter of drum patterns. They mention the capability of the learned model to generate drum patterns similar to the training data, given a seed pattern. Lattner et al. [3] use a similar method to predict the evolution of features for song segmentation.

In this work, we present an intuitive interface for drum pattern generation. Underlying the interface, an RBM, trained on a database of drum patterns is used to create variations of a seed pattern. In addition to presenting the implemented prototype, we report on user feedback we gathered from interviews conducted with experts during hands-on sessions to evaluate the prototype.

[1] http://www.giantsteps-project.eu/

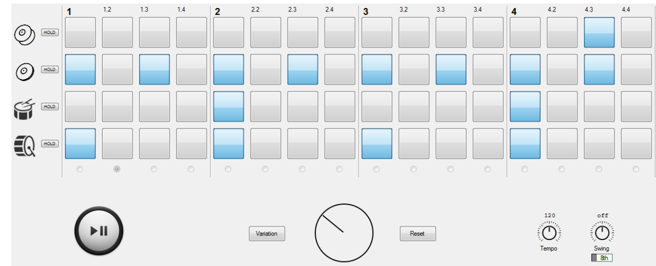

Figure 3: Screenshot of the prototype. The 4 by 16 step sequencer array poses as visualization and input for the drum rhythm patterns. Beneath the array are controls for playback, the pattern variation control to determine the degree of variation, and controls for tempo and swing (all also controllable through a hardware interface).

UI and Method

The developed prototype aims at supporting the producer of an EDM track creating variations of drum patterns, as well as providing creative input for creating new drum patterns. The visualization and input interface for these pat-

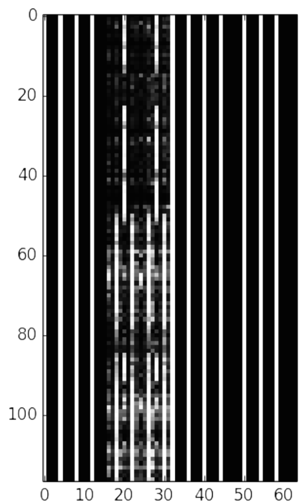

Figure 4: The evolution of the visible nodes of the RBM while creating pattern variations for the snare drum. The x-axis represent the index of the visible node of the RBM. The y-axis represents the number of Gibbs sampling step, starting at the top with the original input pattern and progressing downwards. Active nodes are represented by white, inactive nodes by black pixels.

terns within the prototype employ the well established *step sequencer* user interface (UI) paradigm. The controls for pattern variation are implemented as a dial on which the variations are placed ordered by sparsity and similarity to the seed pattern. Figure 3 shows a screenshot of the prototype's UI.

The output of the prototype is sent via MIDI (Musical Instrument Digital Interface),[2] making the integration into existing setups easy. All UI components, with the exception of the step sequencer array, can be controlled by an external MIDI hardware controller. Musicians and producers are familiar with controlling production software with MIDI controllers – especially in the context of live performances.

To generate meaningful, yet creative patterns, a process which combines obedience to musical rules with elements of surprise and unpredictability is needed. In order to fulfill this requirement, Gibbs sampling of an RBM was chosen as variation method. Apart from being well researched, RBMs feature a technique called *clamping*, which improves the quality of the generated patterns greatly. For details on Gibbs sampling, clamping, and RBM training, the reader is referred to the work by Hinton et al. [6]. RBMs are neural networks and have to be trained on representative training data. As training data a set of 16,513 one-bar drum patterns was used. The patterns were extracted from the sample drum loop library of Native Instrument's *Maschine*[3] software. The library consists of drum patterns for EDM, Hip Hop, and RnB. Since the main focus of this work is EDM, this library was well suited.

To generate variations of the seed pattern, first the seed pattern is entered into the visible layer of the RBM. Then variations for every instrument are generated individually by clamping all other instruments and performing several Gibbs sampling steps. Figure 4 shows the evolution of the visible layer of the RBM performing Gibbs sampling steps. It can be observed how the snare pattern (nodes 16-31) evolves while the other instruments (nodes 0-15 and 32-63) are clamped to their original values. The sorted single-instrument-pattern lists are then combined to full rhythm patterns by using bass drum, snare drum, open, and closed hi-hat patterns at the same indices.

[2]https://en.wikipedia.org/wiki/MIDI
[3]http://www.native-instruments.com/en/products/maschine/production-systems/maschine-studio/

Early Prototype Evaluation

To evaluate the quality of the generated patterns, as well as the interaction with the UI, a user study was conducted of which we report first findings. To this end we interviewed ten experts in EDM creation in a guided, informal way while they were exploring the prototype – see figure 5. Table 1 summarizes the number of positive responses for five evaluated aspects we deemed crucial for the success of such an interface.

Over two thirds of the users (seven of ten) considered the variations to maintain the basic rhythmic idea they entered. Only participants who entered patterns untypical for EDM, complained that the variations did not conserve their basic rhythmic idea. This can easily be explained by the fact that the RBM was trained on EDM patterns and therefore tried to converge on these kind of patterns.

Nine out of ten participants considered the variations produced by the prototype to be musically meaningful. Eight participants commented positively on the way they interact with the prototype (step sequencer, variation dial and hardware controller), as exemplified by these quotes:

> "It works like it should, so I think it is quite user friendly. [...] I also think the scrolling [through the variations] is cool because it is fast and practical."
> JKU-15-05

Figure 5: A study participant using the prototype to explore pattern variations. A MIDI hardware controller is used to enable a more direct interaction.

> "I have tested quite a lot of hardware sequencer things and I think a feature like that would be pretty cool, actually. Especially if it has lots of variations like we had right there." JKU-15-08

Regarding the use of the prototype in live performances, the participants presented themselves cautious. Six out of ten participants stated that they could imagine to use the prototype in a live environment. Some of the participants would use this kind of tool only with the addition of features like a preview function (visually or audible) or the option to limit the degree of variation. The idea of using the prototype in a studio environment was met with enthusiasm. Participants were eager to use the prototype to create variations and get inspiration from it in a production context.

Conclusion

We presented a prototype for an intelligent rhythm agent to assist musicians and producers in the context of EDM production and live performances. A user study was conducted to evaluate both the pattern variation algorithm as well as the UI of the prototype. The study shows that the interaction concept of the prototype is something most participants can imagine working with. It also implies that the acceptance of such a tool in a studio environment would be high, while concerns were raised about precision and reliability when it comes to live performance scenarios. The created patterns were mostly considered musical and in many cases perceived to reflect the basic rhythmic idea of the seed pattern.

aspect	positive comments (out of 10)
seed rhythm is preserved	7
patterns are meaningful	9
prototype interaction	8
would use live	6
would use in studio	9

Table 1: Number of participants giving positive responses wrt. the topics of interest of the user study. The total number of participants (N) was ten.

Acknowledgments

We thank Michael Hlatky for sharing the source code of a simple step sequencer which was used to build the UI and Matthias Leimeister for extracting the *Maschine* drum loop library. This work is supported by the European Union Seventh Framework Programme FP7 / 2007-2013 through the GiantSteps project (grant agreement no. 610591).

REFERENCES

1. E. Battenberg and D. Wessel. 2012. Analyzing Drum Patterns Using Conditional Deep Belief Networks. In *Proc 13th International Society for Music Information Retrieval Conference*.
2. M.A. Kaliakatsos–Papakostas, A. Floros, and M.N. Vrahatis. 2013. evoDrummer: Deriving Rhythmic Patterns through Interactive Genetic Algorithms. In *Evolutionary and Biologically Inspired Music, Sound, Art and Design*. Lecture Notes in Computer Science, Vol. 7834. Springer, 25–36.
3. S. Lattner, M. Grachten, K. Agres, and C.E. Cancino Chacón. 2015. Probabilistic Segmentation of Musical Sequences using Restricted Boltzmann Machines. In *Proc 5th Mathematics and Computation in Music Conference*.
4. C. Ó Nuanáin, P. Herrera, and S. Jordà. 2015. Target-Based Rhythmic Pattern Generation and Variation with Genetic Algorithms. In *Proc 12th Sound and Music Computing Conference*.
5. P. Smolensky. 1986. Information processing in dynamical systems: Foundations of harmony theory. *Parallel Dist. Proc.* (1986), 194–281.
6. G. E. Hinton, S. Osindero, and Y. Teh. 2006. A Fast Learning Algorithm for Deep Belief Nets. In *Neural Comput.* Vol. 18. MIT Press, 1527–1554.