

# GIANTSTEPS – PROGRESS TOWARDS DEVELOPING INTELLIGENT AND COLLABORATIVE INTERFACES FOR MUSIC PRODUCTION AND PERFORMANCE

P. Knees,<sup>1</sup> K. Andersen,<sup>2</sup> S. Jordà,<sup>3</sup> M. Hlatky,<sup>4</sup> G. Geiger,<sup>5</sup> W. Gaebele,<sup>6</sup> R. Kaurson<sup>7</sup>

<sup>1</sup> Dept. of Computational Perception, Johannes Kepler University Linz, Austria

<sup>2</sup> Studio for Electro-Instrumental Music (STEIM), Amsterdam, Netherlands

<sup>3</sup> Music Technology Group, Universitat Pompeu Fabra, Barcelona, Spain

<sup>4</sup> Native Instruments GmbH, Berlin, Germany <sup>5</sup> Reactable Systems S.L., Barcelona, Spain

<sup>6</sup> Yadastar GmbH, Cologne, Germany <sup>7</sup> JCP-Connect SAS, Rennes, France

info@giantsteps-project.eu

## ABSTRACT

We present the GiantSteps project, an EU-funded project involving institutions from academia, practitioners, and industrial partners with the goal of developing new concepts for intelligent and collaborative interfaces for music production and performance. At the core of the project is an iterative, user-centric research approach to music information retrieval (MIR) and human computer interaction (HCI) that is designed to allow us to accomplish three main targets, namely (1) the development of intelligent musical expert agents to support and inspire music makers, (2) more intuitive and collaborative interfaces, and (3) low-complexity methods addressing low-cost devices to enable affordable and accessible production tools and apps. In this paper, we report on the main findings and achievements of the project’s first 18 months.

**Index Terms**— music production, live performance, music information retrieval, intelligent agents, user interfaces, user-centric design

## 1. INTRODUCTION

The stated goal of the GiantSteps project is to create the so-called “seven-league boots” for future music production.<sup>1</sup> Built upon an iterative and user-centric research approach to music information retrieval (MIR) and human computer interaction (HCI), the project is developing digital musical tools and music analysis components that provide more intuitive and meaningful interfaces to musical data and knowledge in order to empower music practitioners to use their creative potential. In particular, we want to achieve this by targeting three directions:

1. Developing and integrating **musical expert agents**, i.e., supportive systems for melody, harmony, rhythm, structure, or style to guide users when they lack inspiration or technical or musical knowledge.

<sup>1</sup><http://www.giantsteps-project.eu>



**Fig. 1.** Music performance during an expert user session.

2. Developing **improved user interfaces** and paradigms for (collaborative) musical human-computer interaction that are easily graspable by novices and lead to unbroken workflows for professionals.
3. Developing low-complexity algorithms for music analysis addressing **low-cost devices** to enable affordable and accessible production tools and apps.

In order to meet these goals, GiantSteps is set up as a trans-disciplinary research project, carried out by a strong and balanced consortium including leading music information research institutions [*UPF-MTG, JKU*], leading industry partners [*Native Instruments (NI), Reactable, JCP-Connect*] and leading music practitioners [*STEIM* (cf. Fig. 1), *Red Bull Music Academy/Yadastar*].<sup>2</sup> With this consortium, the project aims at combining techniques and technologies in new and

<sup>2</sup><http://redbullmusicacademy.com>; From the webpage: “The Red Bull Music Academy is a world-travelling series of music workshops and festivals [in which] selected participants – producers, vocalists, DJs, instrumentalists and all-round musical mavericks from around the world – come together in a different city each year. For two weeks, each group will hear lectures by musical luminaries, work together on tracks and perform in the city’s best clubs and music halls.”



**Fig. 2.** Impressions from the user sessions: children’s workshop, noise-making prototype, interface mock-up, interview on workflow, and practitioners at the Red Bull Music Academy in Tokyo (from left to right).

unprecedented ways, all driven by users’ practical needs and aligned with existing market requirements, allowing for a smooth integration of outcomes into real-world systems.

In this paper, we describe the findings and achievements of the project within the first 18 months. Section 2 outlines the user-centric design approach that generates the requirements for the technical developments. Section 3 deals with the advances in MIR that are necessary in order to enable music expert agents, novel visualisations, and new interfaces as discussed in section 4, as well as more effective algorithms for low-resource devices. Sections 5 and 6 describes the outcomes of the project in terms of market-released products and academic impact and outreach, respectively. Section 7 concludes with an outlook to the second half of the project.

## 2. USER-CENTRIC DESIGN APPROACH

The project’s constellation as a collaboration between music research institutions, manufacturers of software and hardware for music production, R&D companies and music practitioners allows us to engage in an ongoing conversation with the professional musicians that we consider our key users, and to tailor musical tools to their needs. This takes the form of generating user requirements and testing prototypes with end-users in an iterative process throughout the project.

The overall goals of this user involvement are to establish a range of current creative practices for musical expression, explore mental models of musical qualities, produce user-generated ideas through explorative making, and inspire design and non-design related tasks within the project, cf. [1, 2]. To this end, we conduct a series of different workshop and interview sessions involving expert and non-expert users as well as children (cf. Fig. 2).

The user sessions comprise interface-specific and work-practice-related interviews and cognitive walkthroughs, e.g., to identify breaks in workflows, as well as ad-hoc, open-ended interviews, carried out on location at STEIM, NI, the Music Hack Day in Barcelona and the Red Bull Music Academy, resulting in interactions with 94 individual users so far.

From our open-ended interview sessions, a number of ideas and requirements have emerged addressing the areas

of structuring audio repositories and describing and finding sounds, embodiment and physical devices, and the role of the (collaborating) machine in music creation. We will exemplarily put some statements by users in the following sections to demonstrate how these findings are informing the technical research in the project. More details on the studies can be found in [3, 4, 5, 6].

## 3. MIR FOR MUSIC PRODUCTION

Music information retrieval plays a central role in the project. The goal of the research is to develop high-performance and low-complexity methods for music and audio analysis that allow for extraction of musical knowledge in order to drive intelligent composition agents and visualisations (cf. section 4) according to the specific user needs in the music production and performance domain. The following quotes from interviews we conducted demonstrate a need for accurate music analysis methods.

*“Onset detection, beat detection, tempo detection and harmony detection is pretty much all we need. [...] Being able to pick out a small musical phrase of any kind in a big bunch of noises could be extremely helpful for people like me. Instead of spending nights equalizing something to get out a small musical idea.”* [RBMA\_Tokyo\_003]

*“...if you had technology that could tag all your drum samples, that this one is like dirty or distorted, 43Hz is the dominant frequency...”* [RBMA\_Tokyo\_007]

In accordance with these and other statements, new and improved MIR algorithms have been developed. These regard the areas of onset detection, beat and downbeat tracking [7, 8, 9], tempo estimation, key detection, chord extraction, melody extraction, style description and classification, instrument classification, drum sample description and recommendation, and drum transcription [10] for electronic dance music (EDM). Additionally, further semantic categories such as the “loopability” of a segment have been investigated. The developed methods for onset detection, beat tracking, and

tempo estimation<sup>3</sup> have successfully competed in the scientific MIREX evaluation campaign and yielded the top ranks in their respective tasks.<sup>4</sup>

Furthermore, steps towards the optimisation of algorithms for mobile platforms have been undertaken by establishing an audio analysis and benchmarking framework for the iOS mobile platform and a real-time-capable analysis library<sup>5</sup> for use in Pure Data and Max/MSP environments, both based on the Essentia audio analysis library.<sup>6</sup> The released libraries are not only of interest to researchers but also address music hackers and developers who often are practitioners themselves. In addition to signal-based approaches to music analysis, we also investigate the potential of online resources to provide semantic information on music and music styles [11].

#### 4. EXPERT AGENTS AND NEW INTERFACES

The extracted musical knowledge is used to inform supportive and inspirational music expert agents as well as enable new visualisations and interfaces. While users generally welcome the possibility of compositional support by intelligent systems, we found that this is a sensitive matter as it can not only disturb the creative process but also challenge the artistic identity of the user.

*“It can turn to be pretty invasive pretty fast, and, like, annoying, if it doesn’t work properly.”*  
[NIBLN002]

*“I’m sceptical about introducing, you know, stuff like melody into it, like, here’s a suggested kind of thing which fits nicely the two or three patterns you already got in there, then you are really kind of like creating melodies for them, then it’s like (laughs), then it’s really like, you know, who is the composer of this? (laughs)”* [NIBLN003]

Thus, we have to ensure that these systems are present when needed, but do not take over or inhibit the creative process. So far, the expert agents developed are concerned with rhythmic variations, concatenative sound generation based on timbre, tonality-aware scale restrictions, and arpeggiations. The integration of these agents in the workflow is inherently tied to the development of suitable interfaces for either existing desktop-based production and performance suites (such as NI’s MASCHINE), tangible and/or tabletop user interfaces like the Reactable [12], or smaller multi-touch interfaces of affordable portable devices such as tablets and smartphones. For instance, the automatic tonalizer expert agent integrates with the Reactable by displaying a virtual keyboard that is restricted to notes that match the scale of sample objects positioned on the table. The impact of the intelligent arpeggiator,

scaler, and chorder agent can be controlled by hardware dials on new a hardware interface (cf. section 5). Other interface developments relate to the collaborative control of multidimensional parameter spaces, leading to intuitive, expressive and tangible input modalities [13, 14].

#### 5. PRODUCT INTEGRATION

Apart from inclusion into publicly accessible developer libraries (cf. section 3), the maturity of the technical developments in the project allowed us to integrate some of the project’s outcomes into market-ready commercial products already. For instance, the Intelligent Arpeggiator, Scale, and Chord Engine has been integrated and released by NI into the KOMPLETE Kontrol Plugin, a plugin shell that is shipped with the KOMPLETE Keyboard (see Fig. 3) for seamlessly browsing through KOMPLETE instruments. The same features were also released as a free update to existing MASCHINE customers with the free MASCHINE 2.2 “Melody” update in Nov. 2014, reaching +100k users. The developed Automatic Tonalizer has been integrated by Reactable Systems and will be contained in a future release (see Fig. 4).



**Fig. 3.** The Native Instruments Komplete Kontrol-S keyboard released in Oct. 2014 containing GiantSteps technology.



**Fig. 4.** The Reactable Automatic Tonalizer being showcased at Musikmesse 2015.

<sup>3</sup>made available via <https://github.com/CPJKU>

<sup>4</sup>[http://www.music-ir.org/mirex/wiki/2014:MIREX2014\\_Results](http://www.music-ir.org/mirex/wiki/2014:MIREX2014_Results)

<sup>5</sup><http://mtg.upf.edu/technologies/EssentiaRT~>

<sup>6</sup><http://essentia.upf.edu>

## 6. PROJECT IMPACT AND OUTREACH

Apart from the early integration of our research into NI's and Reactable's interfaces, the consortium is actively pursuing multiple avenues for dissemination and exploitation in the form of sessions, demonstrations, and scientific workshops. GiantSteps was co-organizer (in collaboration and exchange with the EU projects PHENICX and CrowdRec) of a demo and hands-on session on NEM Summit 2014. Furthermore, the SoMeRA workshop at SIGIR 2014 was co-organized by GiantSteps partner JKU in collaboration with the PHENICX project [15] and the workshop "Collaborating with Intelligent Machines: Interfaces for Creative Sound" was held by three GiantSteps partners at CHI 2015 [4]. On top of that, we aim at making our results known to a wider audience by participating not only in academic or industrial forums, but also by taking active roles in events such as the Red Bull Music Academy, Music Hack Days, Sónar+D, and Musikmesse.

## 7. CONCLUSIONS AND OUTLOOK

The consortium and orientation of GiantSteps allow for a genuinely target-user-focussed MIR and HCI research approach. This unusual combination of disciplines makes it possible for user's requests and desires to be present in the earliest stages of the MIR algorithm design.

Throughout our conversations with users, there is a strong desire for improved retrieval mechanisms as well as an ongoing acknowledgment of the difficulty of engaging in the collaborative aspects of music making across different interfaces and practices.

*"...what takes me really long time is organising my music library for DJing. [...] it could be something like Google image search for example."* [RBMA\_Tokyo\_011]

*"Because we usually have to browse really huge libraries [...] that most of the time are not really well organised."* [RBMA\_Tokyo\_003]

*"We have a boundary between us, that is, like, we defined that, like, where do we cross, where do our systems cross, and then they will eventually either run as own processes, or we will copy the code eventually, into one process..."* [BCN011]

These are two of the areas where GiantSteps will focus its energies in the following months, looking to build and create experimental interfaces for new methods of retrieval and collaboration. These interfaces will then in turn be tested and reviewed by our users, not only to determine if they work but maybe more importantly, if they are interesting and productive as interfaces for creative expression in digital sound.

## 8. ACKNOWLEDGMENTS

This work is supported by the European Union Seventh Framework Programme FP7 / 2007-2013 through the GiantSteps project (grant agreement no. 610591).

## 9. REFERENCES

- [1] K. Andersen and D. Gibson, "The Instrument as the Source of new in new Music," in *Proc 2nd Biennial Research Through Design Conference*, 2015.
- [2] K. Andersen, "Using Props to Explore Design Futures: Making New Instruments," in *CHI workshop on Alternate Endings: Using Fiction to Explore Design Futures*, 2014.
- [3] K. Andersen and F. Grote, "GiantSteps: Semi-structured conversations with musicians," in *Proc CHI-EA*, 2015.
- [4] F. Grote, K. Andersen, and P. Knees, "Collaborating with intelligent machines: Interfaces for creative sound," in *Proc CHI-EA*, 2015.
- [5] F. Grote, "Jamming with machines: Social technologies in musical creativity," in *Proc 8th midterm Conference of the European Research Network Sociology of the Arts*, 2014.
- [6] F. Grote, "The music of machines – investigating culture and technology in musical creativity," in *Proc XIII. Conference on Culture and Computer Science*, 2015.
- [7] S. Böck and G. Widmer, "A Multi-Model Approach to Beat Tracking Considering Heterogeneous Music Styles," in *Proc ISMIR*, 2014.
- [8] M. Davies and S. Böck, "Evaluating the Evaluation Measures for Beat Tracking," in *Proc ISMIR*, 2014.
- [9] F. Korzeniowski, S. Böck, and G. Widmer, "Probabilistic Extraction of Beat Positions From Neural Network Activations," in *Proc ISMIR*, 2014.
- [10] M. Leimeister, "Feature learning for classifying drum components from nonnegative matrix factorization," in *Proc 138th AES*, 2015.
- [11] P. Knees, "The use of social media for music analysis and creation within the GiantSteps project," in *Proc Workshop on Social Media Retrieval and Analysis*, 2014.
- [12] S. Jordà, M. Kaltenbrunner, G. Geiger, and R. Bencina, "The reacTable," in *Proceedings of the International Computer Music Conference (ICMC)*, 2005.
- [13] K. Gohlke, M. Hlatky, and B. de Jong, "Physical construction toys for rapid sketching of tangible user interfaces," in *Proc TEI*, 2015.
- [14] C. Ó Nuanáin and L. Ó Sullivan, "Real-time Algorithmic Composition with a Tabletop Musical Interface: A First Prototype and Performance," in *Proc Audio Mostly*, 2014.
- [15] M. Schedl, P. Knees, and J. Shen, "SoMeRA 2014: Social media retrieval and analysis workshop," in *Proc SIGIR*, 2014.