Interactive Generation of Musical Corpora for Piano Education: Opportunities and Open Challenges

Filippo Carnovalini[1] [a], Antonio Rodà[1] and Geraint A. Wiggins[2,3]

[1] *Centro di Sonologia Computazionale, University of Padova, via Gradenigo 6, Padova, Italy*

[2] *Vrije Universiteit Brussel, Pleinlaan 9, 1050 Brussel, Belgium*

[3] *Queen Mary University of London, Mile End Road, London E1 4NS, U.K.*

Keywords: Music Education, Integrated Teaching Systems, Music Generation, Computational Creativity.

Abstract: Learning to play a musical instrument such as a piano requires many hours of exercises, generally taken from a "method" book. These books are collections of progressive exercises intended to teach specific techniques and address the commonest mistakes and difficulties that players face while learning. One downside of these books is that the exercises are not personalized to the students and thus cannot address specific difficulties and characteristics of each learner. Given the many recent advances in the field of music generation, we propose that it should be possible to generate exercises automatically to form a personalized method for each student. The teacher would describe the characteristics of the student and their strengths and weaknesses to a software system, as well as the teaching goals that should be covered in the generated exercises, and the system would create exercises that are specific to the needs of the student and the concerns of the teacher, allowing for a more effective and engaging learning experience.

In this paper, we describe a project trying to design such a system, stating research questions, describing the tentative methodology, and outlining its potential impact for both research in music generation and in computer-supported education.

SCIENCE AND TECHNOLOGY PUBLICATIONS

1 INTRODUCTION

Computational Creativity is, as defined by the homonymous Association[1], "the art, science, philosophy, and engineering of computational systems which, by taking on particular responsibilities, exhibit behaviors that unbiased observers would deem to be creative." Having worked both on musical applications of Computational Creativity and on Computer-Supported Education in the past, the authors of this position paper wish to explore how the two disciplines can fit together. We believe that Computational Creativity, and more specifically Music Generation, can open novel avenues for music education, and that the exploration of these new applications can also give further insights into creativity and its computational formalization.

In the first part of this paper, we will describe some open problems of these two disciplines, searching for those contact points that we believe can lead to fascinating hybridization of approaches. Then, we will focus on some specific research questions that we wish to address in future work, and introduce the project *CALIOPE* (Co-creativity And Learning: Interactive Opus generation for Piano Education), which was recently awarded funding by the European Union under a Marie Skłodowska-Curie Actions Postdoctoral Fellowship and is set to start in early 2024. This project will specifically address these research questions, hopefully giving novel insights into Computer-Supported Music Education and Music Generation, by creating a prototype system for the automatic generation of piano exercises tailored around a description of the learner's characteristics and skill level.

2 STATE OF THE ART

2.1 Music Generation

While the interest in AI-generated music is growing recently, the task of algorithmically creating music is not novel and has examples that even predate com-

[a] https://orcid.org/0000-0002-2996-9486

[1] https://computationalcreativity.net

Carnovalini, F., Rodà, A. and Wiggins, G.
Interactive Generation of Musical Corpora for Piano Education: Opportunities and Open Challenges.
DOI: 10.5220/0012059200003470
In *Proceedings of the 15th International Conference on Computer Supported Education (CSEDU 2023) - Volume 1*, pages 407-412
ISBN: 978-989-758-641-5; ISSN: 2184-5026

Copyright © 2023 by SCITEPRESS – Science and Technology Publications, Lda. Under CC license (CC BY-NC-ND 4.0)

puters (Nierhaus, 2009). Among some of the more common approaches to music generation are the use of statistical methods (e.g., Markov Chains) to predict suitable notes to continue a musical sequence, the use of formal grammars to create sequences, genetic algorithms, and rule-based systems paired with optimization algorithms (Carnovalini and Rodà, 2020). However, in recent years the method that has become most common for the generation of musical sequences is the use of Neural Networks, commonly applying Deep Learning approaches such as Generative Adversarial Networks, Convolutional Networks, Autoencoders, and Long Short-Term Memory Networks (LSTM). The latter have proven to be particularly suitable for the generation of sequences in time, which can well represent melodies (Briot and Pachet, 2020). Although many advancements have been made in this field, there are still a variety of open challenges that need to be addressed. One of the most notable difficulties of music generation algorithms is the lack of long-term structure: generated pieces are credible for a couple of bars but seem to "wander off" as the piece progresses, lacking the formal aspects of repetition and reuse that characterize human-composed melodies (Carnovalini, 2019). Interestingly, even if Machine Learning approaches have become more common for this task, these systems offer little control over the generated results (Briot and Pachet, 2020) making it harder to solve this problem. Some researchers have started designing ad hoc architectures to obtain well-structured results (Roberts et al., 2018), while others have embedded multiple pre-existing architectures to delegate structural aspects to a specialized part of the system (Cho et al., 2014). Other researchers have directly addressed the problem of long-term structure. Morpheus (Herremans and Chew, 2017) is another EU-funded project that leveraged a specialized mathematical representation of harmonic tension to recreate the "tension profile" learnt from existing pieces in newly generated pieces. Hierarchical representations of melodic content can also be used to describe reuse of melodic content and obtain multi-layered descriptions of structure in music pieces (Carnovalini et al., 2021b).

We believe that the way forward is to expand on hierarchical representations of melodic content, as these representations can capture cognitive aspects of musical structure (Temperley, 2011). Applying such structures to the context of musical exercises could also open new directions into the description of stylistic coherence and reuse of melodic content between multiple pieces, considering long-term structure at a corpus/book level rather than at the song/piece level.

2.2 Computer Assisted Music Education

Music Education in a traditional Western context includes a variety of aspects, all of which can benefit from the use of computational systems to some extent. It is possible to divide these aspects into four main categories: fundamentals of music (music theory), performance skills, music analysis, and composition (Brandao et al., 1999). In this paper, we wish to focus on performance skills, i.e., those skills related to playing one instrument (the piano). Research shows that the use of a Computer-Assisted Instruction (CAI) framework can effectively increase the effectiveness of piano lessons compared to a control group (Kaleli, 2020). In this framework, the computational system does not substitute the role of the teacher, but rather complements the teaching by offering a learning environment (Abdullah and Mustafa, 2019). For example, a CAI system can embed score following to help a student visualize (through a display or through augmented reality) what is being played and the mistakes they make (Smoliar et al., 1995; Rigby et al., 2020). Another approach is to use an Integrated Teaching System (ITS), that is, a system that has greater control over the teaching activity by proactively helping the teacher by leveraging knowledge about the subject being taught, the curriculum, and the student. In this context, there are systems that can select exercises from a given corpus based on the needs of the student and the teacher's concerns (Wiggins and Trewin, 2000; Sébastien et al., 2012). To the best of our knowledge, there is no system that generates novel piano exercises to propose to students in a similar fashion. One similar project is the one by Takidaira et al (Takidaira et al., 2022), that generates rhythmic scores for the popular rhythm game Taiko no Tatsujin to improve the player's ability in the game by analyzing their game logs. While this is not directly meant for musical education, performance in a rhythm game is strongly related to musical skills, and thus this can be considered a music education tool. However, that system only focuses on rhythmic aspects, making the system quite limited compared to the one we propose. Other researchers dealing with music generation have considered applications in the context of learning an instrument, and there is some research towards controlling the difficulty of a generated music piece. Ariga et al (Ariga et al., 2017) proposed a system to generate guitar solos using fingering as a means of controlling difficulty. Nakamura and Yoshii (Nakamura and Yoshii, 2018) describe a system that generated piano reduction of ensemble scores, allowing different difficulty settings based on tempo and fingering, and similar techniques have been applied to score difficulty analysis (Ramoneda et al., 2022). However, none of these systems really try to build an ITS that uses music generation for the goal of piano education.

The current level of customization that is possible with ITS is limited by the selected corpora and therefore cannot offer a fully personalized learning experience. We posit that the use of computational creativity can improve the personalized learning experience by creating a specific music generation system that takes into account the characteristics of the student.

2.3 Research Questions

Wishing to explore further levels of personalization for musical education, while also deepening our understanding of music cognition and representation for Computational Creativity, we formulate the following research question:

Q1: How can Computational Creativity be applied to improve the quality and personalization of musical exercises used in a training course?

Two sub-questions arise from this:

Q2: How can hierarchical knowledge representation improve the longer-term structure of computationally created sequences of exercises?

Q3: How can a model of the learner and the educational context improve the computational creation of sequences of exercises?

We decided to address these questions through the development of a prototype system for teaching piano. The objective would be to create a system capable of generating musical exercises in a progressive manner, considering the playing skills of the user as input. This goal can be further divided into three subobjectives:

SO1: Create a system for the generation of a set of music pieces with progressive difficulty.

SO2: Design a knowledge representation to describe the user's weaknesses, strengths, learning needs, and personality (Learner Model).

SO3: Integrate the Learner Model into the generation system, so that the generated exercises are personalized to the student.

The next section will describe in more detail the design choices of the project and the features we desire from the final prototype, as well as the methodology we wish to follow both for the implementation of the prototype and for finding answers to the research questions outlined above.

3 PROJECT CALIOPE

In this section, we wish to introduce project CALIOPE (Co-creativity And Learning: Interactive Opus generation for Piano Education), which stems from the desire to address how Computational Creativity can help improve personalization in music education. This project was recently awarded a Postdoctoral Fellowship under the Marie Skłodowska-Curie Actions funding scheme. Although the project has not started yet, this section will introduce the ideas and design choices that drove the formulation of this project, hoping to receive comments and suggestions from the scientific community before the project's official beginning.

3.1 Summary

When learning to play a musical instrument in a classical western setting, such as conservatory or music school, teachers employ exercise books meant to increase the difficulty of the provided studies progressively. Such a book is known as a "method", and it is easy to find numerous methods for any orchestral instrument. One of the major drawbacks of these books is that they do not provide a personalized learning experience as they do not focus on the specific needs of each student, but instead provide general exercises that can be useful to most. It is the role of the teacher to select exercises that are most fit for their students, potentially selecting exercises from different methods or even creating ad hoc exercises for each student. This activity requires abundant time and effort from teachers, who often resort to simply following the most popular methods instead of personalizing education, limiting the effectiveness of music education, and requiring more effort from students.

Given the considerable recent advancements in music generation, we propose that it should be possible to employ computers to automatically generate personalized methods that are tailored to the needs of each student. Therefore, the main output of the project will be a prototype Integrated Teaching System for piano education. A music teacher would use this system to generate exercises that are tailored to their students, by giving the system information about each student via specialized knowledge representations to describe the learner, called the Learner Model (Bull, 2020). The model describes the strengths, learning needs, and personality of the student. After the initial setup of the learner model, the prototype will give the teacher some starting exercises for the student. The student practices using the generated exercises and the teacher evaluates their performance.

Based on that, the teacher can describe to the system what kind of difficulties the student found and what they should focus on going forward, and what is the goal of the following practice session (i.e., what is the Teaching Concern (Wiggins and Trewin, 2000)). The system can use this information to generate further exercises, possibly also including annotations on things that the learner should pay attention to while exercising, allowing the learning cycle to progress.

In this way, the system would help provide a personalized exercise plan based on the characteristics of the student, addressing the weaknesses of the student while maintaining a satisfactory learning curve since the generation would consider their current level of expertise.

3.2 Methodology

To create the system described above and answer the research questions described in the previous section, the following methodological steps will be employed:

(i) Creation of the dataset. The AI system will require a dataset of relevant piano methods. Many exist, but at least 10 methods must be selected (estimating about 50 exercises per method for a total of 500 exercises). Methods in the public domain, available digitally, and that respect gender balance will be preferred. Syllabi for piano courses at music education institutes will be used to find well-established methods.

(ii) Model of Difficulty. A detailed study of the literature on automatic difficulty analysis (Ramoneda et al., 2022) will serve as a basis for the definition of a model of difficulty. Using a top-down approach to the matter, to define constraints that describe difficulty.

(iii) Generation of Progressive Exercises. Hierarchical representations of structure, jointly with machine learning (different algorithms will be tested, including LSTM networks), will be leveraged to generate musical pieces. The difficulty model will be embedded in the generation system as a self-reflection module that influences the output.

(iv) Evaluation (Difficulty Generation). A meta-evaluation procedure (Carnovalini et al., 2021a) will be used to evaluate the difficulty control system. Expert-based evaluation (e.g., CAT (Amabile, 1983), SPECS (Jordanous, 2012)) will be used to assess the quality of the produced exercises.

(v) Learner Model Design. A collaborative design (Chiu, 2002) approach will be used, creating a focus group of piano teachers that will help me describe the characteristics of the student in the model, as well as the teaching concerns in piano education.

(vi) Evaluation (Learner Model). A focus group will be used again to evaluate how well the designed model captures the requirements.

(vii) Personalization. The difficulty model will be expanded to consider the learner model.

(viii) Final Evaluation. The same as point iv, but also using a test set covering all the limit use cases of the learner model will be used, and the expert evaluation will consider the different applied settings.

The methodology is dictated by the diverse objectives of the project. Computational creation of musical content requires the use of musical information, which can be provided either via data in a bottom-up manner or via description of rules and constraints in a top-down manner. For the generation of musical material, the bottom-up approach is chosen as literature shows that machine learning is effective in describing musical sequences, with added top-down information given by the hierarchical model to ensure long-term coherence. For the description of difficulty and of the learner model a top-down approach is used because machine learning in this context would require higher amount of musical data with difficulty annotations and might fail to capture the multidimensional concept of difficulty. To capture the real-life teaching requirements that piano teachers meet, collaborative design is used for SO2. This will help maintain a human-centric design and avoid bias introduced by AI in the most sensitive part of the project, which is the description of the learner model. Evaluation follows the principles established by the literature for computational creativity and music generation, requiring experts to gather insights on the capabilities of the system. The use of limit use cases for the final prototype serves as a proxy for the evaluation of the impact of this teaching system using actual students, which can only be considered in a longer-term study when the prototype created with this project has already been thoroughly tested in an offline setting.

3.3 Relevance and Open Directions

Being a project that stems from research questions coming from both Computational Creativity and Computer-Supported Education, the potential impact of the project is twofold.

From the viewpoint of music generation, this project offers a novel approach to the problem of long-term structure: multi-piece generation. By considering the entire method as the generated output, this project will need to generate a large amount of musical content that must be coherent and well-structured, both considering the usual criteria for musical well-formedness within each piece and considering higher structures joining the different exercises, maintaining a coherent style throughout. The fact that exercises are distinct music pieces gives some leeway, but the additional constraints regarding progressive difficulty and educational value make this an intriguing use case. Under the broader perspective of non-strictly musical Computational Creativity, the project offers a novel take on the exploration of a conceptual space (Boden, 2004). The Learner Model serves to define personalized conceptual spaces that are explored by the generation algorithm as per SO3: the space of the possible musical pieces is constrained and transformed dynamically by the Learner model, and the generation of novel pieces both explores the defined space by adapting the difficulty of pieces, and by transforming the space of useful exercises by updating the Learner model. This approach also deals with the problem of Explainability in AI: by having the system generate music that is specifically tailored to address a specific teaching concern, it is possible to assess how well the system understands the given task and it would be possible to leverage the knowledge in the system to describe what makes each exercise apt for the task at hand.

From the viewpoint of Computer-Supported Education, this project will offer a case study for the employment of Artificial Intelligence in the personalization of the educational path. This technology can offer novel insight on how an Integrated Teaching System can support both the Teacher and the Student through the use of the novel human-computer interaction routes offered by Artificial Intelligence and Computational Creativity.

Finally, the creation and public release of the data and technologies related to this project (exercise dataset, learner model, difficulty analyzer) will hopefully foster further research both in music generation and in Computer-Supported Education.

4 CONCLUSIONS

In this position paper we presented project *CALIOPE*, by giving an overview of the research questions it will try to answer, the methodology that we intend to follow to address those questions, and the potential impact of the project. In the latter part of the paper, we briefly reviewed relevant scientific literature.

The main goal of the project is to develop a prototype for an Integrated Teaching System that will leverage Music Generation techniques for the creation of personalized piano exercises. By doing so, we believe it will be possible to address some open problems of music generation and computational creativity, while advancing the knowledge in Computer-Supported education.

Although the project is still in its preliminary phase, the authors are open to suggestions and willing to find potential collaborations from the scientific community. With this paper, our desire was to state our intents and what we believe could result from this study, and let the project be known to the relevant community.

ACKNOWLEDGEMENTS

The first author wishes to thank his supervisors for all the help given in these years and for the sincere interest in the author's career development.
While this paper describes a project for which funding was approved by the European Union, the current work was not funded by the European Union and does not necessarily reflect the views of the European Union, and the authors are the sole responsible for the contents of this article.
This work was partially supported by the project Creative Recommendations to Avoid Unfair Bottlenecks, granted by the Department of Information Engineering, University of Padova.

REFERENCES

Abdullah, Z. D. and Mustafa, K. I. (2019). The Underlying Factors of Computer Self-Efficacy and the Relationship with Students' Academic Achievement. *International Journal of Research in Education and Science*, 5(1):346–354. Publisher: ERIC.

Amabile, T. M. (1983). A Consensual Technique for Creativity Assessment. In Amabile, T. M., editor, *The Social Psychology of Creativity*, Springer Series in Social Psychology, pages 37–63. Springer New York, New York, NY.

Ariga, S., Fukayama, S., and Goto, M. (2017). Song2Guitar: A Difficulty-Aware Arrangement System for Generating Guitar Solo Covers from Polyphonic Audio of Popular Music. In *ISMIR*, pages 568–574.

Boden, M. A. (2004). *The creative mind: Myths and mechanisms*. Routledge, London, United Kingdom.

Brandao, M., Wiggins, G., and Pain, H. (1999). Computers in music education. In *Proceedings of the AISB'99 Symposium on Musical Creativity*, pages 82–88. AISB

(Society for the Study of Artificial Intelligence and the Simulation of . . . .

Briot, J.-P. and Pachet, F. (2020). Deep learning for music generation: challenges and directions. *Neural Computing and Applications*, 32(4):981–993.

Bull, S. (2020). There are Open Learner Models About! *IEEE Transactions on Learning Technologies*, 13(2):425–448.

Carnovalini, F. (2019). Open Challenges in Musical Metacreation. In *Proceedings of the 5th EAI International Conference on Smart Objects and Technologies for Social Good*, pages 124–125, Valencia Spain. ACM.

Carnovalini, F., Harley, N., Homer, S. T., Rodà, A., and Wiggins, G. A. (2021a). Meta-Evaluating Quantitative Internal Evaluation: a Practical Approach for Developers. In *Proceedings of the 12th International Conference on Computational Creativity*, page 5, México Virtual. Association for Computational Creativity.

Carnovalini, F. and Rodà, A. (2020). Computational Creativity and Music Generation Systems: An Introduction to the State of the Art. *Frontiers in Artificial Intelligence*, 3:14.

Carnovalini, F., Rodà, A., Harley, N., Homer, S. T., and Wiggins, G. A. (2021b). A New Corpus for Computational Music Research andA Novel Method for Musical Structure Analysis. In *Audio Mostly 2021 (AM '21)*, page 4, virtual/Trento Italy. ACM.

Chiu, M.-L. (2002). An organizational view of design communication in design collaboration. *Design Studies*, 23(2):187–210.

Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. (2014). Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.

Herremans, D. and Chew, E. (2017). MorpheuS: generating structured music with constrained patterns and tension. *IEEE Transactions on Affective Computing*, 10(4):16.

Jordanous, A. (2012). A standardised procedure for evaluating creative systems: Computational creativity evaluation based on what it is to be creative. *Cognitive Computation*, 4(3):246–279.

Kaleli, Y. S. (2020). The Effect of Computer-Assisted Instruction on Piano Education: An Experimental Study with Pre-Service Music Teachers. *International Journal of Technology in Education and Science*, 4(3):235–246. Publisher: ERIC.

Nakamura, E. and Yoshii, K. (2018). Statistical piano reduction controlling performance difficulty. *APSIPA Transactions on Signal and Information Processing*, 7.

Nierhaus, G. (2009). *Algorithmic Composition: Paradigms of Automated Music Generation*. Springer Science & Business Media. Google-Books-ID: jaowAtnXsDQC.

Ramoneda, P., Tamer, N. C., Eremenko, V., Serra, X., and Miron, M. (2022). Score Difficulty Analysis for Piano Performance Education based on Fingering. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 201–205, Singapore. IEEE.

Rigby, L., Wünsche, B. C., and Shaw, A. (2020). piARno - An Augmented Reality Piano Tutor. In *32nd Australian Conference on Human-Computer Interaction*, pages 481–491, Sydney NSW Australia. ACM.

Roberts, A., Engel, J., Raffel, C., Hawthorne, C., and Eck, D. (2018). A hierarchical latent vector model for learning long-term structure in music. In *International conference on machine learning*, pages 4364–4373. PMLR.

Smoliar, S. W., Waterworth, J. A., and Kellock, P. R. (1995). pianoFORTE: a system for piano education beyond notation literacy. In *Proceedings of the third ACM international conference on Multimedia - MULTIMEDIA '95*, pages 457–465, San Francisco, California, United States. ACM Press.

Sébastien, V., Ralambondrainy, H., Sébastien, O., and Conruyt, N. (2012). Score analyzer: Automatically determining scores difficulty level for instrumental e-learning. In *ISMIR 2012*, pages 571–576.

Takidaira, S., Shoji, Y., and Dürst, M. J. (2022). Slow, Repeat, Voice Guidance: Automatic Generation of Practice Charts to Improve Rhythm Action Games. In Fang, X., editor, *HCI in Games*, volume 13334, pages 119–138. Springer International Publishing, Cham. Series Title: Lecture Notes in Computer Science.

Temperley, D. (2011). Composition, Perception, and Schenkerian Theory. *Music Theory Spectrum*, 33(2):146–168.

Wiggins, G. A. and Trewin, S. (2000). A System for Concerned Teaching of Musical Aural Skills. In Goos, G., Hartmanis, J., van Leeuwen, J., Gauthier, G., Frasson, C., and VanLehn, K., editors, *Intelligent Tutoring Systems*, volume 1839, pages 494–503. Springer Berlin Heidelberg, Berlin, Heidelberg. Series Title: Lecture Notes in Computer Science.