Algorithmic composition: An overview of the field, inspired by a criticism of its methods

MUS-15

Sven Deserno

July 10, 2015

1 Introduction

Algorithmic composition, the "partial or total automation of the process of music composition by using computers" [2, p. 513] or the "technique of using algorithms to create music" [8], is a popular field of research and a particular strand of modern musical practice. These definitions are rather general, and their true scope will be mapped out in the course of this report. While some historical precedents and contemporary composers, e.g. Arvo Pärt [11, p. 48], employ algorithmic methods without using computers, these applications are – save for the purpose of imitation – of limited interest to the perspective of computer science. There exist a number of different terms such as *computer music*, *automatic composition*, and *computer-aided algorithmic composition* (CAAC), which are used in a similar or identical way [4, p. 1; 11, p. 48]. All statements in this report should apply to these areas as well, regardless of minor distinctions made by some authors.

The aim of this report is to present the criticism that was levelled by Pearce, Meredith, and Wiggins [1] against the field of algorithmic composition. Pearce, Meredith, and Wiggins (hereafter PMW) attest a stagnation in the body of research, and identify the cause in a failure to develop and adopt an appropriate methodology. Embedded in the presentation of this criticism is an overview of the variety and complexity within the field, in order to a) give context to the criticism, b) outline possible roots of the described problems, and c) motivate the improvements proposed by PMW.

This report is structured as follows: In Section 2, PMW's criticism is presented and supplemented with examples as well as with notes on the requirements for scientific progress in a field of research. Section 3 contains a brief history of algorithmic composition and a discussion of the immense variety of problems and approaches that fall under the umbrella of algorithmic composition. Once the criticism and the relevant field are thus mapped out, section 4 presents PMW's proposal of differentiating 4 distinct motivations for developing algorithms that compose music and discusses the benefits and implications. Finally, section 5 provides some concluding thoughts and a tentative outlook on the development since the publication of the presented paper.

2 PMW's criticism

Echoing a similar criticism, that was levelled 12 years earlier against the field of artificial intelligence, PMW speak of a "methodological malaise" [1, pp. 121ff] that resulted in a stagnation in the body of research concerning algorithmic composition. The accusation against the field is threefold: PMW attest

> 1. a failure to specify the precise practical or theoretical aims of research;
> 2. a failure to adopt an appropriate methodology for achieving the stated aims;
> 3. a failure to adopt a means of evaluation appropriate for judging the degree to which the aims have been satisfied. [1, p. 121]

According to PMW, these failures stem from "an implicit assumption that simply describing a computer program that composes music counts as a useful contribution to research" [1, p. 121]. This is an indictment of the common practice of describing algorithms (that compose music) with a focus on implementation details, but on a shaky theoretical foundation, and with a marked deficit in formulating goals and evaluating results.

The effects of this practice can be observed, for example, in the work of David Cope. Cope's program *Experiments in Musical Intelligence* (EMI) is able to learn from a given corpus of musical examples and produces pieces in a similar style. EMI works well enough to have gained some popularity, however the presentation of the algorithms in books and research articles has been deemed insufficient, vague, and ultimately unscientific [1, p. 125]. This criticism was so pronounced that Wiggins' [6] opens his review of one of Cope's works with a discussion of pseudoscience.

The severity of PMW's accusations can be underscored by taking note of some requirements for progress in any field of research, particularly those dealing with the development and analysis of algorithms. To make progress measurable it is paramount to a) deal with well-defined problems, b) propose solutions to these problems, and c) be able to meaningfully compare different solutions. This, of course, requires a solid methodology. New research might contribute to progress by describing new and quantifiably better solutions to existing problems, or by giving a valid definition of some new problem. The latter case should naturally also include a justification for the relevance of this particular problem. If these requirements are not met, the possibility of scientific progress becomes highly questionable. For example, in the field of algorithmic composition, the goals of research are often formulated so vaguely that the real problem is only defined retroactively by the described algorithm. This, of course, invalidates any question as to how well the algorithm solves the given problem. In this light, PMW's criticism of the failure to specify the precise aims of research and adopt an appropriate corresponding methodology becomes especially pertinent. PMW also cite cases where the produced output was deemed "reasonable", "quite pleasing musical forms", or "pleasant to listen to" [1, pp. 127, 129]. These subjective impressions reveal very little about the algorithm and cannot form a basis to define any kind of progress. In terms of the above criticism, such assessments are not an appropriate means of evaluation. Therefore, these examples underscore the validity of the accusations.

3 On algorithmic composition

Accepting this criticism as valid and as an accurate description of the current situation in algorithmic composition begs the question of why the field developed in this direction. A cautious answer might be derived from a look at the origins of algorithmic music composition on one hand, and a discussion of the factors that may impede the development of appropriate methodologies and means of evaluation on the other hand. To this end, this section presents a historical perspective, followed by a discussion of a selection of variables that can account for a vast number of problems and approaches in the field of algorithmic composition.

3.1 A history of algorithmic composition

The idea of using computers for the composition of music can be traced back to the very beginnings of modern computer science. In 1843, Ada Lovelace wrote concerning the Analytical Engine:

> Supposing, for instance, that the fundamental relations of pitched sounds in the science of harmony and of musical composition were susceptible of such expressions and adaptations, the engine might compose elaborate and scientific pieces of music of any degree of complexity or extent. [7]

It took another century before this idea became reality. In 1954 Lejaren Hiller and Leonard Isaacson produced the Illiac Suite, using the ILLIAC I computer at the University of Illinois, Urbana-Champaign, as a proof of concept. It consisted of 4 movements for string quartet and is widely considered the first composition by a computer program. Each movement was designed to test different principles of algorithmic composition, employing different sets of rules, and experimenting with probabilities and randomness [11, p. 49; 1, p. 120].

Besides this interest in algorithmic composition from the perspective of computer science, many composers were curious about the possibilities created by the increasing availability of computers. One well-known pioneer is Iannis Xenakis (1922 – 2001), who championed the use of mathematical models in music composition and used the output of algorithms as material in several compositions. He was also interested in developing algorithms as tools that could be used by other composers to produce musical material [1, p. 120].

Since these beginnings, the field of algorithmic composition has attracted a vast number of publications, which exhibit a great deal of variety. Researchers have aspired to different musical styles, from baroque-style counterpoint to improvised jazz and completely new genres, and adapted a huge array of technical approaches from other areas of computer science. Some of this variety is accounted for and explored in the next section.

It should be noted that the idea of using algorithms to create music, without reliance on a computer, is much older. In the eleventh century, Guido d'Arezzo, originator of modern musical notation, used algorithmic principles to derive melodies from the text of songs. For this purpose, he assigned a different pitch to each vowel sound, which effectively defined a deterministic mapping from text to melody. A completely different approach is the 18th century "Dice Music", attributed to W. A. Mozart [5, pp. 1, 38]. It consisted of a number of pre-composed parts that could be randomly combined to create countless similar-but-different compositions. The distinction between deterministic and randomised approaches is only one of the variables that will be discussed in the next section.

3.2 Scope and variety

The previous definitions of algorithmic composition have remained extremely general. To gain an understanding of its scope and the variety within, it is prudent to examine the different characteristics an algorithm that composes music might have. First, the *input* may take different forms. It may consist of musical examples in the form of scores or recordings, it may consist of rules and probabilities that encode some facet of musical theory, or it may be extra-musical data that is meant to be mapped to music by the algorithm. Second, the *output* is somewhat less variable in its form, it usually consists of a musical score or synthesised sound, but the issue of evaluation becomes prominent. In contrast to other areas of computer science, the question of what a "correct" or even good output is, is not straightforward and needs some serious considerations. PMW's criticism applies here, when these considerations are not taken or made appropriately clear. Third, there are general characteristics, such as the question of *determinism* and the *degree of human intervention*. Finally, the actual technical approach taken by the algorithm, governing how input and possible outputs are related, can vary greatly. The choices made on any of these considerations can define very different problems, which all fall in the area of algorithmic composition.

The variety of technical approaches can be observed in a comprehensive survey of papers in algorithmic composition. Fernández and Vico [2] briefly review several hundred papers, and their categorisation gives an insight into the breadth of approaches that have been used for algorithmic composition. On the coarsest level, they differentiate between 1. grammars, 2. symbolic, knowledge-based systems, 3. Markov chains, 4. artificial neural networks, 5. evolutionary and other population-based methods, and 6. self-similarity and cellular automata.

Each of these approaches can be realised in a number of ways and employed to different ends. This abundance of choices means that, in turn, the choice of any one approach should be sufficiently justified.

Turning to the evaluation of the output, at least two perspectives present themselves. First, one could try to define what constitutes "good" music. However, this is a question of music theory, philosophy, or plain subjective impression, and it is unclear how this could be turned into a computable measure for an algorithm. Second, there is the question of specific aims and expectations for an algorithm. Here, a common distinction is made between the imitation of a style or corpus of examples, and what has been called "genuine composition" [5] or "automation of compositional tasks" [2]. In the former case, meaningful measures could be devised to assess the similarity of the output and a predefined style. In the latter case, a more specific formulation of the expectations is needed to define such a measure.

Looking at the breadth of problems and approaches, as well as the difficulty of evaluating outputs, it becomes plausible that the field of algorithmic composition evolved in a way that warrants PMW's criticism. A compounding factor is that, in this area, artistic and scientific interests meet, and the outlined difficulties might lead researchers to substitute artistic curiosity for a difficult but sound methodology. A similar view is shared by Fernández and Vico:

> As machines became less expensive, more powerful and in some cases interactive, algorithmic composition slowly took off. However, aside from the researchers at Urbana (Hiller's university), there was little continuity in research, and reinventing the wheel in algorithmic composition techniques was common. This problem was compounded by the fact that initiatives in algorithmic composition often came from artists, who tended to develop ad hoc solutions, and the communication with computer scientists was difficult in many cases. [2, p. 518]

4 PMW's proposal of categorisation by motivation

It remains the question of how the criticised status quo can be improved. First of all, the criticised threefold failure, which was outline above, can be turned into a challenge to specify precise aims, adopt appropriate methodology, and adopt appropriate means of evaluation. However, PMW do not only criticise, but also offer suggestions to facilitate the adoption of sound methodologies [1].

They propose the distinction between 4 disparate motivations underlying research in algorithmic composition:

1. "algorithmic composition" in a stricter sense,
2. design of compositional tools,
3. computational modelling of musical styles, and
4. computational modelling of music cognition.

For each motivation, PMW point to appropriate methods and important issues. Furthermore, they hold that, in the past, a failure to distinguish between these categories has led to bad methodology and ultimately bad science.

In the first category of *algorithmic composition* in a stricter sense, the objective is purely artistic. The algorithm is a tool in the compositional process and reflects the composer's idiosyncratic needs and vision. When such an algorithm is published, the theoretical or practical relevance must be demonstrated in order to create a valuable contribution to research [1, pp.125f].

In the second category, *design of compositional tools*, the problem becomes a software engineering task. Consequently, software engineering standards should be upheld, and a publication should adequately document the analysis, design, implementation, and testing stages [1, pp.126ff].

The third category, *computational modelling of musical styles*, is closely linked with music theory. The algorithmic modelling permits testing hypotheses about the properties of different styles. PMW discuss devising sound tests for over- and undergeneration, that constitute a measure of how well the algorithm emulates a style [1, pp.129ff].

The fourth category, *computational modelling of music cognition*, is related to cognitive science and has the goal to test hypotheses about the cognitive processes that are involved in and required for musical composition. In this case, the relations and differences between cognitive processes and features of the algorithmic model must be made clear and justified [1, pp. 134ff].

PMW state that they “consider it imperative that developers of such [algorithmic composition] systems be clear about their motivations”, that this “would have potential benefits for both the research itself and for others interested in their work”, and that a failure “hinders the communication of theories between different academic communities and the comparison of different theories and practical applications.” [1, p. 140] This last claim is particularly important, as it explains how failures of individual research projects can result in a stagnation in the entire field of research.

5 Conclusion and outlook

This report started from PMW’s assertion that the field of algorithmic composition suffers from a lack of appropriate methods. As a rational for the status quo, a view on the history of algorithmic composition and the scope of and variety within the field was delineated, followed by a presentation of PMW’s proposed solutions. In a nutshell, PMW proffer a categorisation of research interests due to 4 distinct motivations, dividing the general field of algorithmic composition into sub-fields that each require different approaches to a sound methodology.

Since the publication of the presented paper, there has been a limited impact. As of July 10, 2015 it has been cited 60 times (according to Google Scholar), and several researchers have put forward ideas for better frameworks in algorithmic composition, e.g. Ariza [4], who offers a model for a more fine-grained categorisation. However, no approach seems to have gained enough traction to bring about significant change in the field of research. Pearce and Wiggins have continued to work in the field of computational modelling of music cognition, the fourth motivational category discussed above, and computational creativity. They published papers proposing methods for the evaluation of such models [9, 10].

References

[1] Pearce, Marcus, Meredith, David, and Geraint Wiggins. "Motivations and methodologies for automation of the compositional process." *Musicae Scientiae* 6.2 (Fall 2002): 119–147.

[2] Fernández, Jose David, and Francisco Vico. "AI Methods in Algorithmic Composition: A Comprehensive Survey." *Journal of Artificial Intelligence Research* 48 (2013): 513–582.

[3] Papadopoulos, George, and Geraint Wiggins. "AI Methods for Algorithmic Composition: A Survey, a Critical View and Future Prospects." *AISB Symposium on Musical Creativity*, 1999.

[4] Ariza, Christopher. "Navigating the Landscape of Computer-Aided Algorithmic Composition Systems: A Definition, Seven Descriptors, and a Lexicon of Systems and Research." *Proceedings of the International Computer Music Conference.* San Francisco: International Computer Music Association, 2005. 765–772.

[5] Nierhaus, Gerhard. *Algorithmic Composition: Paradigms of Automated Music Generation.* Wien: Springer, 2009.

[6] Wiggins, Geraint. "Computer Models of Musical Creativity: A Review of Computer Models of Musical Creativity by David Cope." *Literary and Linguistic Computing* 23.1 (2008): 109–116.

[7] <http://blogs.bodleian.ox.ac.uk/adalovelace/about-ada-lovelace/> Last retrieved: July 10, 2015.

[8] Wikipedia. <https://en.wikipedia.org/wiki/Algorithmic_composition> Last retrieved: July 10, 2015.

[9] Wiggins, Geraint A., Marcus T. Pearce, and Daniel Müllensiefen. *Computational modeling of music cognition and musical creativity.* na, 2009. <http://www.researchgate.net/publication/253385688>

[10] Pearce, Marcus T., and Geraint A. Wiggins. "Evaluating cognitive models of musical composition." *Proceedings of the 4th international joint workshop on computational creativity.* Goldsmiths, University of London, 2007.

[11] Supper, Martin. "A Few Remarks on Algorithmic Composition." *Computer Music Journal* 25.1 (Spring 2001): 48-53.