

COLUMN

Worldwide AI

TITLE

Melomics: A case-study of AI in Spain

AUTHORS

*Carlos Sánchez Quintana,
Francisco Moreno Arcas,
David Albarracín Molina,
Jose David Fernández Rodríguez,
Francisco J. Vico*

ABSTRACT

Traditionally focused on good old-fashioned AI and robotics, the Spanish AI community holds a vigorous computational intelligence substrate. Neuromorphic, evolutionary, or fuzzy-like systems have been developed by many research groups in the Spanish computer sciences. It is no surprise, then, that these nature-grounded efforts start to emerge, enriching the AI catalogue of research projects and publications and, eventually, leading to new directions of basic or applied research. In this article, we review the contribution of Melomics in computational creativity.

In Spain there are 74 universities, many of which have computer science departments that host AI-related research groups. AEPIA, the Spanish society for AI research, was founded in 1983 and has been vigorously promoting the advancement of AI since then. Along with several other societies and communities of interest, it promotes various periodic conferences and workshops. The Artificial Intelligence Research Institute (IIIA) of the Spanish National Research Council constitutes one of the flagships of local AI research. Ramón López de Mántaras, IIIA's renowned director, was one of the pioneers of AI in Spain, and he also was the recipient of the prestigious AAAI Englemore Award in 2011. Other researchers that have reached an outstanding position, and lead important research groups in Spain, include Antonio Bahamonde (University of Oviedo), Federico Barber (Polytechnic University of Madrid), Vicent Botti (Polytechnic University of Valencia), and Amparo Vila (University of Granada). In this column, we describe a new class of computer composer (Ball 2012), which is being considered as a milestone in AI research, currently developed at the Computer Science Department of the University of Málaga (UMA). This department, with more than one hundred faculty members, is organized in several research groups, three of which maintain active AI research lines.

Melomics is a new approach in artificial creativity (for a perspective on this discipline, see the 2009 fall issue of AI Magazine). More specifically, it focuses on algorithmic composition and aims at the full automation of the composition process of professional music. Before going into the details, and to better understand what is new in Melomics, it is worth mentioning that a wide range of AI techniques have been used for algorithmic composition in the past (like grammatical and knowledge-based systems, artificial neural networks, statistical machine learning, and evolutionary algorithms), as well as a wide collection of mappings from raw data to music notation. The first result goes back to the origins of AI: the Illiac Suite, a composition generated by computer (programmed by Hiller and Isaacson, late in 1956) as an experiment on the formal aspects of music composition. Since then, many researchers and artists have got notable results, as David Cope's Emily Howell algorithm or Kemal Ebcioglu's CHORAL expert system. Most strategies for computer composing have focused on imitating preexisting human styles, but Melomics' computer composers shy away from this trend, providing the system with knowledge about music composition (as a human learner is taught), which allows them to create their own styles.

Genetics, Embryology, and Evolution

To achieve these results, instead of traditional AI techniques, an approach based on evolutionary algorithms and indirect encoding has been followed. Since AI Magazine does not usually feature articles using these concepts, we will devote a few paragraphs to present them before going into further detail. See also Stanley and

Miikkulainen's (2003) paper for a more extensive presentation of the concept of indirect encoding.

Biological evolution is one of the main mechanisms that has contributed to the diversity and complexity of living forms. In computer science, evolutionary algorithms represent a kind of heuristic methodology inspired by evolutionary biology. In these algorithms, a changing set of candidate solutions (a population of individuals) undergoes a repeated cycle of evaluation (by means of a fitness function), selection, and reproduction with variation (mutation and crossover). Evolutionary algorithms have been studied in depth over the last few decades, and they have been applied to many different problem domains.

However, classical evolutionary algorithms tend to show problems of scalability (the performance degrades significantly as the size of the problem increases) and solution structure (the solutions generated by the algorithm tend to be unstructured, hard to adapt, and fragile). One of the most important factors in these problems is the use of direct encodings: classical evolutionary algorithms use genotypes (representations of solutions) that map to phenotypes (the solutions themselves) in a straightforward way. When an algorithm uses direct encoding, each part of the solution is mapped to a part of the representation. As a result, genotypes can grow too large for evolutionary optimization to be practical, and different parts of the genotype can evolve uncoordinatedly, thus inducing the previously mentioned problems of scalability and fragility in the solutions.

However, the use of direct encodings can be regarded as a historical artifact. In contemporary evolutionary biology, development (the process that transforms a zygote into a full-fledged multicellular organism) is known to play a key role in evolutionary processes. Unfortunately, when evolutionary algorithms were first proposed, this role had been downplayed for decades. This, together with the limited capacity of early computers, resulted in the widespread, implicit adoption of direct encoding in evolutionary algorithms.

As time passed, the importance of development in the context of evolutionary biology became better and better understood, and this eventually gave rise to a specific branch of evolutionary thought: evolutionary developmental biology (evo-devo for short). Developmental processes can be described as self-organized choreographies of precisely timed events, with cells dividing and arranging themselves into layers of tissues that fold in complex shapes, resulting in the formation of a multicellular organism from a single zygote. In evo-devo, evolutionary changes are interpreted as small mutations in the genome of organisms that modulate their developmental processes in complex and orchestrated ways, resulting in altered forms and novel features.

As the field of evo-devo matured in the context of evolutionary biology, it inspired an analogous field in the

context of evolutionary computation: artificial development. In opposition to the direct encodings of classical evolutionary algorithms, artificial development uses indirect encodings, that is, formal abstractions of developmental processes that define complex mappings between genotype and phenotype. Using an effective indirect encoding, a small genotype can potentially specify a large and complex phenotype, accounting for the scalability problem previously mentioned. Additionally, a small change in the genotype can potentially provoke a variety of coordinated changes in the phenotype, thus helping to mitigate the problem of solution fragility.

Because of these characteristics, evolutionary algorithms with well-engineered indirect encodings can obtain complex solutions, and potentially generate complex variations of these solutions. In some concrete fields traditionally left to human expertise alone (like industrial design, or the arts), these algorithms can perform in a truly disruptive way. They are currently being used to a certain extent for automating tasks that demand creativity, proposing different variations to existing solutions, which evolve toward desired design targets, resembling an automated form of brainstorming.

Remarkable examples of real-life applications of artificial development have appeared in recent years, such as the set of antennas designed to fit the technical requirements for the satellites of a NASA space mission (Hornby, Lohn., and Linden 2011); the design of microstructured optical fibers (Manos, Large, and Poladian 2007); and the automatic generation of board games (Browne 2008). Our group has successfully developed and applied this paradigm of artificial development to two different domains in computational creativity: new techniques for automatic character animation (Lobo, Fernández, and Vico 2012), and new systems for algorithmic composition.

Computer Composers

Iamus is a computer composer specialized in contemporary classical music. What makes its strategy unique is that, instead of learning by examples, it is taught to compose in a way that resembles how humans learn music composition rules. At its core, Iamus follows an evo-devo approach, as it combines an evolutionary algorithm with a sophisticated indirect encoding scheme to map genotypes into phenotypes. In essence, Iamus implements the evolution of complex musical structures, encoded into artificial genomes (resembling multicellular living organisms, which develop from a genome, and also evolve in time). These genomes represent the musical information in an indirect and very compact way: each genome encodes the specifications to generate a music piece following a complex developmental process. As a population of compositions evolves, mutations performed on their genomes will provoke the resulting musical structures to change. A fitness function evaluates the result, determining if the new structures are better or worse fitted to musical formal constraints and basic aesthetic principles. As evolution proceeds, genomes will undergo



Figure 1. Examples of scores generated by Iamus, considered by professional musicians good enough as to have been written by an avant-garde composer (Berger 2013). (Name and dedication were not decided by Iamus.)

transformations, making the corresponding music pieces more and more complex, and better fitted to the requirements.

Artificial development provides very powerful ways of encoding music pieces. It means a very high computational cost, but it also opens the way to a groundbreaking automation of creative tasks. New fusion genres could be discovered by recombination operators that merge musical genomes of different styles, so the offspring might show combined features of the parental genomes, speeding up the appearance of new hybrid music styles. Also, the genes encoding a particular jingle could be incorporated into another genome, resulting in compositions with musical fragments that evoke the original melody that has been inserted in the genome. Finally, a theme could adapt by evolution to certain constraints (based, for example, on a video script). Applications of this kind could go disruptive, if the music community is provided with efficient tools to implement them. As we will see, doing this is far from simple and constitutes one of the main problems for the diffusion of AI-based solutions.

In computational creativity, another drawback is the ill-defined nature of the problems under study. In the case of Melomics, a fitness function determines the conditions under which musical organisms evolve. These functions were designed under collaboration with professional musicians, and Melomics assesses music compositions according to various criteria of formal and basic aesthetic nature, enabling populations of music pieces to evolve toward formally correct and potentially beautiful musical forms. By way of illustration, piano chords of six notes (to be played with a single hand) would penalize the

composition, as they cannot be performed. Also, the polyphony of some instruments is rewarded, as they do combine well (the results are aesthetically valid). All in all, nearly a thousand rules have been coded to define the environment in which correct music is intended to evolve.

But it is the evo-devo nature of the algorithm behind Melomics that constitutes a milestone in the design of fully autonomous computer composers, as it provides the tools to access, manipulate, and turn into musical language an extraordinarily diverse world of complex structures. As it was described in the report dedicated to Melomics in Nature: "...unfolding complex structure from a mutable core has enabled the kind of dramatic invention found in biological evolution" (Ball 2012). What is remarkable in Melomics' music is that it allows the definition of new aesthetics, not imitating the style that any particular musician has previously developed. This distinguishes Melomics from previous attempts at algorithmic composition, making it a powerful tool for composers in any musical style.

This technology has been implemented in two computer systems to date. Iamus, on one hand, is a half-cabinet that has produced scores of complete contemporary classical works (figure 1), qualified by professional musicians and composers as indistinguishable from those written by avant-garde composers (Berger 2013). Melomics109, on the other hand, is a computer cluster operated by Melomics Media Inc., fully dedicated to composing and synthesizing popular music, a kind of music fundamentally different from Iamus's style, as it is very much constrained by cultural and instrumental restrictions.

From the Lab to the Market

In order to be truly disruptive, an AI-based solution requires more than engineering prowess: it has to interface effectively with the real world, solving hard technological problems by integration of standards, making a precise market analysis, and setting up cloud and hardware resources. Researchers and students at UMA are encouraged to promote and fund groundbreaking ideas in the scope of the University Spin-off competition, a call that has been published since 1998 to stimulate technology transfer from the academic laboratories toward the socioeconomic fabric. Melomics Media's mission is the commercialization of this generation of computer composers, its proposal was awarded in the 13th edition of this contest, and started to operate as a technology-based startup incubated at UMA's spin-off facilities.

Different strategies have been considered by Melomics Media in order to target the general public. In principle, client-based music generation could be implemented by hardware integration of the algorithms, or embedding them in a system on chip (SoC). These solutions would require a complex control of the parameter space by the user, costly composition and rendering times, and also the distribution of expensive music libraries for audio synthesis. Instead, Melomics Media adopted an offline approach: creating a vast repository of music (Infinitunes)² covering most genres, formats, and needs, and making it browsable in an efficient way, usable by a collection of applications, and opened to developers through an API.

Infinitunes contains pieces for direct consumption, to be interpreted by musicians, or used as a musical base for human composers. Consequently, the music pieces are presented in a variety of output formats, to cover these various purposes: music edition (MusicXML and MIDI), human performance (score in PDF), and playing (quality MP3). MIDI and MusicXML, standard formats for music representation, can be imported by most notation editors, which make the music freely modifiable by professional musicians. Scores in PDF format make the music ready for instrumental performance. A file in MP3 format is also obtained with custom processes for audio synthesis, using quality music libraries and track mixing to obtain high-quality results. Finally, the music is stored and made available to the general public with a sophisticated interface that allows users to efficiently browse, mark, and buy themes.

Infinitunes is paving the way for brand new applications and business models. For example, Melomics@life is a mobile application that provides streamed music for every scenario of daily life, in-store audiomarketing, and music therapy. The royalty-free character of these songs makes them available worldwide, allowing for a wide range of freemium options (in contrast to the limited options that constrained profit margins allow to companies like Spotify and Pandora Media). However, it also opens possibilities to implement

sound trademarking by matching catchy tunes to products or brands, and including them (by inserting their genes) into bigger music pieces. Finally, an API for developers will boost the implementation of a new class of applications that were not viable in practice until now.

Concluding Remarks

We hold that artificial intelligence can benefit much from cross-fertilization with other fields, and the systems described here are but an example. The field of artificial life currently leads the introduction of bioinspired models, and will continue to generate interesting applications of evo-devo systems. The results presented here have achieved an important milestone in computational creativity by connecting two disparate domains: evolutionary computation with artificial development and musical language. Iamus has been widely covered by the media, as the first computer that composes full works of professional contemporary classical music in its own style. Ten of these works have been recorded and compiled in a debut album, the very first one that has been fully composed by a nonhuman intelligence on its own, and interpreted by first-shelf musicians (including the London Symphony Orchestra).

As for industrial implementation of these evo-devo approaches to music composition, we envision three main areas of application. First, it enables novel and disruptive techniques for the music industry: (1) Infinitunes allows nonmusicians to produce original music, and (2) sentient computing on mobile devices, bringing up homeostatic³ music systems, that is, music that adapts online so as to keep constant the user's physiological state. As a direct consequence, it also opens new business opportunities: a marketplace where music is sold as a true commodity, and native advertisement strategies for musical branding, including evocative explicit jingles seamlessly inserted into otherwise normal compositions to be provided through streaming services. Finally, it also paves the way for music genetic engineering: tools to obtaining variations on a theme (by mutation of the genetic material), or automatic fusion of genres (by recombination of themes of different genres). In short: evolutionary computation might be the substrate for new forms of genetic engineering, applied to basically all forms of creative tasks, as we have seen in the industrial design and other artistic endeavors. Aligned for a new age of music (with emerging disciplines such as music medicine and music psychology), composers might soon be adopting the role of music genetic engineers.

Acknowledgments

The work on Iamus was partially supported by grants IPT-300000-2010-010 from the Spanish Ministerio de Ciencia e Innovación and TSI-090302-2011-8 from the Spanish Ministerio de Industria, Turismo y Comercio. The first and fourth authors were supported by grant P09-TIC-5123 from the Consejería de Innovación y Ciencia, Junta de Andalucía.

Notes

1. See selected press coverage at <http://melomics.com/news>
2. See <http://infinitunes.com/search>
3. Term proposed by Matt Peckham in an internal communication.

References

- Ball, P. 2012. Computer Science: Algorithmic Rapture. *Nature* 488(7412): 458.
- Berger, K. 2013. 100 Top Stories of 2012. Digital Composer Records with London Symphony Orchestra. *Discover Magazine* 24(1).
- Browne, C. 2008. Automatic Generation and Evaluation of Recombination Games. Ph.D. diss., Queensland University of Technology, Dept. of Computer Science, Brisbane, Australia.
- Hornby, G. S.; Lohn, J. D.; and Linden, D. S. 2011. Computer-Automated Evolution of an X-Band Antenna for NASA's Space Technology Mission. *Evolutionary Computation* 19(1):1–23.
- Lobo, D.; Fernández, J. D.; and Vico, F. J. 2012. Behavior-Finding: Morphogenetic Designs Shaped by Function. In *Morphogenetic Engineering: Toward Programmable Complex Systems*, ed. R. Doursat, H. Sayama, and O. Michel, 441–472. Berlin: Springer-Verlag.
- Manos, S.; Large, M. C. J.; and Poladian, L. 2007. Evolutionary Design of Single-Mode Microstructured Polymer Optical Fibres Using an Artificial Embryogeny

Representation. In *Proceedings of the 9th Annual Genetic and Evolutionary Computation Conference*, 2549–2556. New York: Association for Computing Machinery.

Stanley, K.; and Miikkulainen, R. 2003. A Taxonomy for Artificial Embryogeny. *Artificial Life* 9(2):93–130.

Biographic snippets

Carlos Sánchez Quintana is a Ph.D. student at the University of Málaga under the supervision of Francisco J. Vico. He is investigating new bioinspired techniques for behavior finding.

Francisco Moreno Arcas, BSc in Computer Science, is a senior developer under the supervision of Francisco Vico, specialized in AI techniques for automatic music composition and synthesis.

David Albarracín Molina, is a PhD student under the supervision of Francisco J. Vico, his work focuses on bio-inspired generative music methods.

Jose David Fernández Rodríguez is a postdoctoral research assistant at the UMA, working in bioinspired algorithmic techniques and biological modeling.

Francisco J. Vico is full professor of Computer Science at the UMA, leader of a research group in bioinspired algorithmic techniques, and a renowned authority in Artificial Creativity.