Musical Metacreation: Papers from the 2012 AIIDE Workshop
AAAI Technical Report WS-12-16

Embracing the Bias of the Machine: Exploring Non-Human Fitness Functions

Arne Eigenfeldt
School for the Contemporary Arts, Simon Fraser University, Vancouver, Canada
arne_e@sfu.ca

Abstract

Autonomous aesthetic evaluation is the Holy Grail of generative music, and one of the great challenges of computational creativity. Unlike most other computational activities, there is no notion of optimality in evaluating creative output: there are subjective impressions involved, and framing obviously plays a big role. When developing metacreative systems, a purely objective fitness function is not available: the designer is thus faced with how much of their own aesthetic to include. Can a generative system be free of the designer's bias? This paper presents a system that incorporates an aesthetic selection process that allows for both human-designed and non-human fitness functions.

Introduction

My recent computationally creative music systems have explored the potential of autonomous musical multi-agents exploring a controlled environment, in which they achieve a sense of musical meaningfulness through social interaction [Eigenfeldt 2010, Eigenfeldt and Pasquier, 2011a]. Such self-organisation for a composer is both exciting and frustrating: exciting in that the agents can end up producing music that is totally unexpected; and frustrating in that the agents can end up producing music that is totally banal.

The dilemma for a composer – who is naturally inclined to order things, and happy to find situations in which it is done for him – is whether to let a complex system run on its own, or to attempt to impart more order on the system, so as to guarantee certain results. Previous systems of mine have chosen the latter path: I can turn them on, and they will produce music that I find enjoyable to listen to. Although I still believe that artistic success – as judged by the system creator – is a necessary evaluation criteria for a metacreative system, I have come to acknowledge that these systems have tended to surprise me less and less (which is another criteria for a creative system) [Bruner, 1962].

The Curator Agent

Several of my metacreative systems have attempted to balance exploration – in which agents develop their own conceptual space – and exploitation, in which agents remain within a predefined space; for example, *Coming Together: Freesound* autonomously generates only soundscape compositions [Eigenfeldt and Pasquier, 2011b]. Sometimes these systems are presented as installations, where users can listen to their output over extended periods of time, and perhaps return hours, or even days, later, to hear very different results. Sometimes, these systems have been presented in concert, which presents its own set of problems. Self-organising systems have their own sense of evolutionary time, and tend to resist the "ten minute" compositional arc desirable within contemporary concert music.

As a compromise, my latest system works offline, and generates interactions which are not performed live, but written to a database. This database is analyzed for salient features – as the files are MIDI information, the analysis operates at a symbolic, rather than audio, level. Immediately prior to performance, a curator agent browses through the database, and selects movements that it thinks are suitable as a single, multi-movement composition [Eigenfeldt and Pasquier, 2012].

The premiere performance featured selections using a human-derived fitness function: choose a first movement that is bright and lively, then subsequently choose movements that are contrasting to immediately proceeding movements, ending with another bright, lively movement. This produced what I considered to be a successful result – in this case, a series of movements whose relationships were interesting when performed consecutively; however, their selection, while unpredictable in their specifics were predictable in the larger formal scheme.

One could argue that any generative system will display bias; (human) creative artists have aesthetic preferences, so shouldn't a metacreative system reflect the bias of its designer? Colton suggests that such systems are, in fact, mere assistants to the designing artist, and fully autonomous creative systems must strive for greater independence [Colton et al, 2012].

Colton argues for non-human fitness functions very eloquently, specifically with the FACE model. However, Colton's domain, in this case, is poetry, and his aesthetic selections rely heavily on external databases that provide meaning to his material: for example, selecting phrases from an online article, and then using tools such as WordNet and Afinn to ascertain meaning, and the DISCO API for ascertaining word similarities. In music, not only do we lack such databases, many of us composers will argue whether music can mean anything at all.

Contextual Fitness

For a recent performance, the curator agent in my system was extended to allow for greater variation in selecting movements, and determining relationships between movements. The selection of any movement can now be based on any analysed criteria (i.e. length of phrase, tempo, density, duration, etc.), and these criteria can be weighted separately. Subsequent movements can be selected through weighted relationships as well, and each feature rated for similarity or dissimilarity in any combination. Features can also be ignored entirely.

As before, an initial movement is selected based upon the selection fitness function. And, as before, the database is rated on the variable relationship fitness measures; however, these ratings are then passed back to the selection function, which ranks the top 5 movements based upon the selection criteria. A subsequent movement is then selected from these ratings, using a roulette-wheel section. This is typical of the author's systems, so as to ensure that the same conditions do not always produce the same results.

Figure 1 shows how the previous human-informed fitness function would look. The selection function, shown above, weights certain features higher (tempi and density, for example), and ignores others (pitchrange, pattern, and modulations, for example). The comparison function, shown below, weights tala (phrase length) and tempo dissimilarity as most important in selection, while ignoring key, modulation ratio, pitch and onset motive relationships.

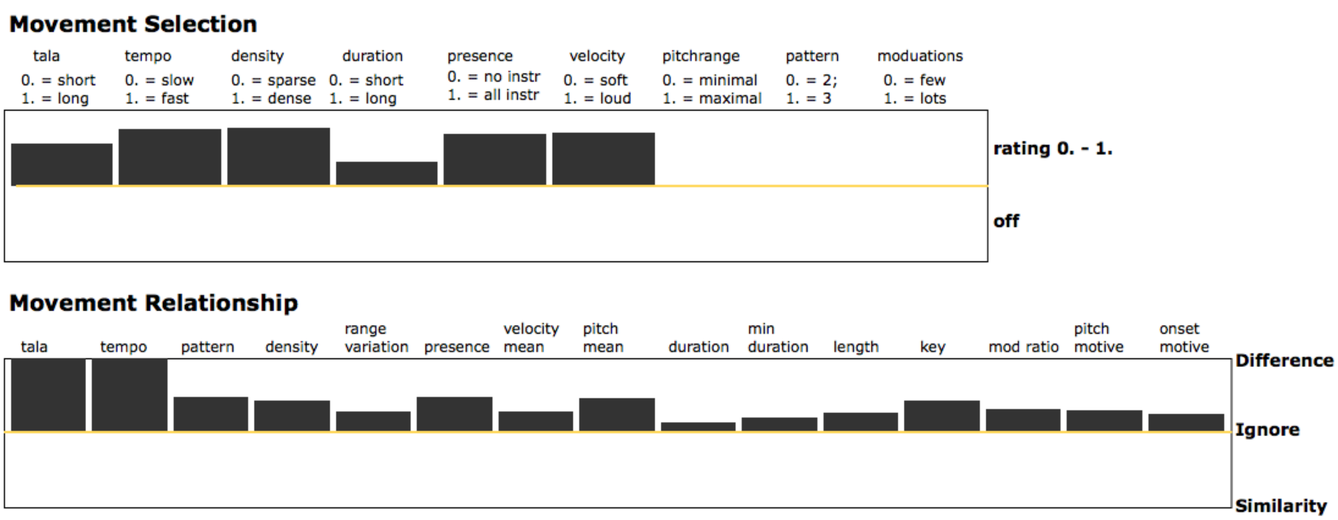

Figure 1. *Selection and Relationship Functions for a traditional multi-movement work, displaying weightings for selection (above) and relationships (below) between movements.*

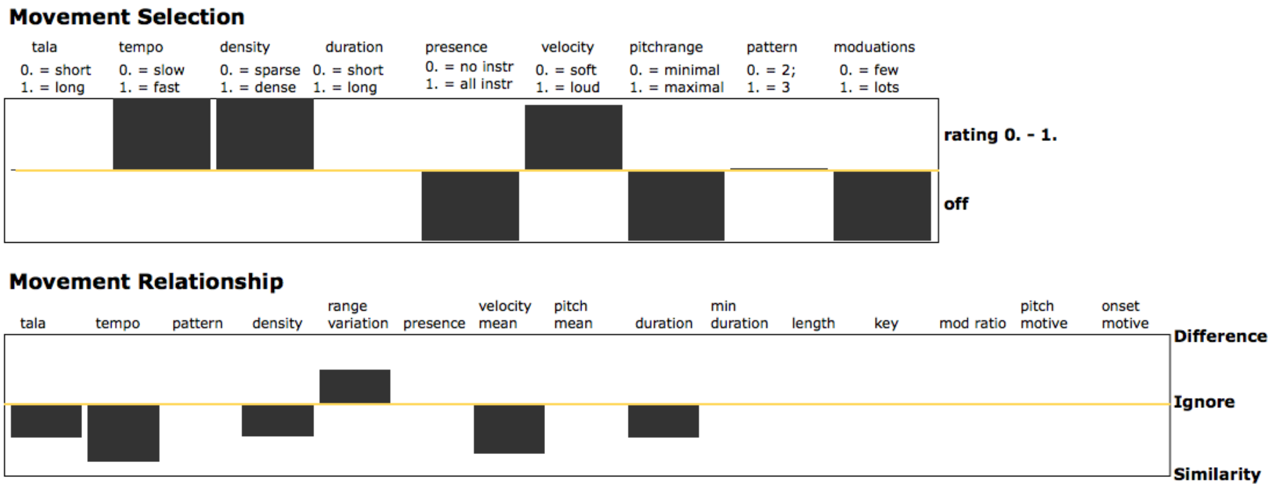

Figure 2. *Selection and Relationship Functions for a multi-movement work with a criteria that favours the selection of movements that contain fast tempi, high density, high velocity, short phrases, and short duration, and subsequent movements that are similar in phrase length, tempo, density, velocity, and duration, but dissimilar in pitch range.*

Figure 2. suggests a "thrash-metal" approach, in which movements that display short phrases, fast tempi, high density, short duration, and high volume are rated highly. Subsequent movements are selected based upon their similarity of phrase length, tempi, density, volume, and duration, and dissimilarity of pitch range.

What has proved interesting in performance has been how the eventual selections can evolve over time, due mainly to the roulette-wheel selection of rated movements. When a selection is made that is not at the top of the list, it may demonstrate certain aspects not found in those directly above it; as similarity (or dissimilarity) to this new object is made, aspects of the original criteria may no longer be present, and the system seems to wander towards its own preferences.

Towards a Machine Bias

Although as yet unexplored in performance, a completely machine derived fitness function can be produced, one that may not make what could be called "musical sense", but one that may display a bias not found in the designer. It is the author's intention to explore these possibilities in a new work in which continual compositions are generated, simultaneously printed out, and performed by mechanical instruments in an ongoing gallery installation.

The demonstration will present the system, including some initial experiments in generating autonomous machine fitness functions.

References

Bruner, J.S. 1962. The Conditions of Creativity. In H.E. Gruber, G. Terrell and M. Wertheimer, Contemporary Approaches to Creative Thinking, USA: Atherton Press, pp. 1–30.

Colton, S.: Goodwin, J.: Veale, T. 2012. Full-FACE Poetry Generation. In Proceedings of the International Conference on Computational Creativity, Dublin, pp. 95–102.

Eigenfeldt, A. 2010. Coming Together - Negotiated Content by Multi-agents. In Proceedings of the ACM Multimedia 2010 International Conference, Firenze, p.1433–1436.

Eigenfeldt, A.: Pasquier, P. 2011a A Sonic Eco-system of Self-Organising Musical Agents. In EvoApplications 2011, Part II, C. Di Chio et al. (Eds.), LNCS 6625, pp. 283–292.

Eigenfeldt, A.: Pasquier, P. 2011b. Negotiated Content: Generative Soundscape Composition by Autonomous Musical Agents in Coming Together: Freesound. In Proceedings of the Second International Conference on Computational Creativity, Mexico City, pp. 27–32.

Eigenfeldt, A.: Pasquier, P. 2012. Creative Agents, Curatorial Agents, and Human-Agent Interaction in Coming Together. In Proceedings of Sound and Music Computing, Copenhagen, pp. 181–186.