Robot choreographies: artificial evolution between novelty and similarity

Mattia Barbaresi[a], Stefano Bernagozzi[a] and Andrea Roli[a]

[a]*Dipartimento di Informatica - Scienza e Ingegneria, Alma Mater Studiorum Università di Bologna, Campus of Cesena, Via dell'Università 50, Cesena, Italy*

Abstract

In this paper we introduce a novel fitness function for evolutionary art, which generates sequences of movements—i.e. robot choreographies—based on similarity to an inspiring repertoire. Similarity is counterbalanced by a novelty mechanisms, which makes it possible to sample unexplored areas of the choreography space. The approach is discussed together with preliminary results achieved in the context of Nō theatre. This work is a first step towards the development of a computational creativity system that can incorporate diverse generative mechanisms and can exploit information theoretic and complexity measures both for the generation and assessment of the choreographies produced.

Keywords

Computational, Creativity, Evolutionary Art, Robotics, Similarity, Novelty

1. Introduction

Computational creativity (CC) is an emerging field of research within AI that studies the ability of machines to both generate and evaluate novel outputs that would, if produced by a human, be deemed creative. Computers and creativity make interesting partners with respect to two different aims: one is understanding human creativity, the other is trying to produce "machine creativity", in which the computer at least appears to be creative, to some degree [1]. In this context there exists a wide spectrum of activities. Most researchers are focused on specific tasks, such as the generation of music [2], paintings [3], poetry [4], jokes [5] or games [6]. Some, such as Wiggins [7] and Ritchie [8], are interested in the generic frameworks within which such tasks take place, and how to evaluate them.

The majority of these works have been drawn from one of the most computationally tractable definitions of creativity, which has been proposed by Margaret Boden [9]: creativity is the ability to come up with ideas or artefacts that are new, surprising, and valuable [1]. So today it is widely accepted that something to be called creative needs to meet at least two main traits: *novelty* (or unexpectedness, unpredictability, surprise, originality) and *value* (or usefulness, quality, appropriateness, meaningfulness) [10].

These concepts—and novelty *in primis*—may find their definition in the context of Information Theory [11]; the theoretical tools it provides enable one to formalise the crucial concepts of

pattern and randomness that are central to learning and computer capability of generating, and relevant to evaluating the novelty of new generated products [12].

Our long term objective in CC aims at investigating these fundamental aspects with a two-fold perspective: exploring mechanisms for the generation of creative artefacts and exploiting information theory metrics for the assessment of such systems. Compared to the works mentioned above, our approach is different and tries to investigate CC on an empirical basis. With these assumptions, as an interdisciplinary research group, [1] we decided to explore the generation of robot choreoghraphies for the Nō theatre, a traditional form of classical Japanese musical drama.

In this work we present the first steps of this research. We conceived a genetic algorithm (GA) which, starting from a given inspiring repertoire and a set of unitary moves, generates symbolic sequences of movements (i.e., choreographies) exploiting similarity with the repertoire combined with the novelty search approach [13].

2. Proposed approach

The idea was to start from a simple, basic system, that includes primitive but essential mechanisms for creative generation (appropriateness/similarity and novelty) then increasingly add features that enhance the creative potential of the generative system, to evaluate the effects that these modifications have on the metrics used for the assessment. So we started encoding similarity and novelty in a genetic algorithm. To do this, we conceive a fitness function that evaluates the average similarity of an individual from the samples in the inspiring set. The entire generation process is guided by the novelty search approach described in [13]. These choreographies are performed by a (virtual) Nao robot [2] using CoppeliaSim [3] environment.

2.1. Encoding

We represent a choreography as a sequence of basic moves (i.e., figures or poses). Each figure is stored as a single keyframe (motion capture data) that specifies a particular position of each captured joint and is encoded as a symbol. Therefore, in our representation a choreography is a sequence of symbols. Hence, the GA generates strings (symbols sequences) that are mapped to keyframe series (sequences of joint angles sets). The actual movement of the robot is given by the transition from one keyframe to another. Only steady poses are encoded and we assume that the transition between stable moves is stable too, taking advantage of the typical gentle style of Nō movements. This, in general, represents a limitation for robot movements in real applications (e.g., some moves or transitions could make the robot fall) but runs well for our purpose.

---

[1]Performing Robots: https://site.unibo.it/performingrobots/en
[2]Softbank Robotics: http://doc.aldebaran.com/2-1/home_nao.html
[3]Coppelia Robotics: https://www.coppeliarobotics.com

2.2. Inspiring set and similarity

In general—but especially when dealing with artistic problem modelization—it is not trivial to conceive a (fitness) function that fully captures all the problem objectives. What is valuable? And how a choreography can be judged as more valuable than another? Note that in this case, we use the term *value* as a synonym of *appropriateness* or *membership*. The value, the quality, of an artefact will usually be assessed by other agents' (humans) judgement and will be based on cultural experience and knowledge, and hence are likely to reflect historical comparisons of the artefacts. These aspects refer more to the H-creativity [9] and will be addressed in future work. To model this aspect in this first attempt, the idea is to take a sample repertoire and to evaluate each individual through its similarity with respect to such a given set. The inspiring set represents choreographies that the robot, or the dancer, knows about a particular style or repertoire, which is the knowledge from which the artist pursues its style and creates new pieces. As a first step, we created the repertoire by manually encoding typical sequences of moves (*kata*) that have been performed by a Nō expert.[4]

Therefore, the fitness function evaluates sequences using string similarity between individuals and this repertoire. In formulas:

$$fitness(x) = \frac{1}{card} \sum_{n=1}^{card} similarity(repertoire[n], x) \tag{1}$$

$$similarity(x, y) = 1 - \frac{JW(x, y) + Jacc(x, y)}{2} \tag{2}$$

where $card = |repertoire|$ is the cardinality of the repertoire (i.e., the number of sequences forming the repertoire), $repertoire[n]$ is the nth element of that inspiring set and where the similarity is formulated as the combination of Jaro-Winkler and Jaccard [14] string distances, respectively $JW(x, y)$ and $Jacc(x, y)$ in equation 2. Similarity can be thought of as the convergent process in the generation that constrains the resemblance of generated artefact to given samples.

2.3. Novelty search

Convergent/divergent thinking is a characteristic of creativity; to model these dynamics, we include a novelty search mechanism that guides the generation in the opposite way with respect to similarity. The novelty mechanism is inspired by the work of Vinhas et al. [13]. The main goal of this algorithm is to generate a more diverse set of individuals than the set that would be created by a traditional fitness based evolutionary algorithm. The method evolves individuals according to two criteria: *(i)* look for the best individuals according to the fitness function and *(ii)* take novelty and fitness as two different objectives to be maximised using a Pareto optimization. This bi-objective optimization is performed considering the number of individuals of the current generation that have a fitness above a given threshold. It uses also an archive to store the most novel individuals. The novelty score is then calculated as

$$novelty(x) = \frac{1}{k} \sum_{j=1}^{k} dissim(x, j) \tag{3}$$

---

[4]https://site.unibo.it/performingrobots/en/project/activities

where $dissim(x, j)$ is a dissimilarity metric, between the individual being evaluated and a set of k neighbours chosen from the population and the archive. We implemented the dissimilarity metric by complementing the similarity function.

2.4. Results

We tested the algorithm by varying several parameters, among which the size of the repertoire, and we compared the results of the combined use of similarity and novelty with simple fitness-based GA results and random search—which may represent a naive way of producing a divergent process. For lack of space, here we just briefly summarize the main results. As expected, the overall similarity slightly decreases with the size of the repertoire, meaning that the search process achieves a similarity balance between the generated individuals and the ones in the repertoire. The drift imposed by novelty turns out to be effective, as the best choreographies generated are usually sensibly different among themselves, considerably more diverse than those generated at random. Some samples of the choreographies can be found here: https://github.com/ste93/NaoNohVideos.

3. Ongoing and future work

Currently we are performing a quantitative assessment of the choreographies by applying information theory and complexity measures. These metrics make it possible to capture relevant features of the choreographies that can be used both for identifying their peculiarities and differences (and possibly compare the fingerprint provided by the metrics with human evaluation), and derive some heuristics for improving the generation process. This approach is analogous to that recently proposed for swarm robotics [15]. Among the metrics that we are currently applying we mention the Normalized Compression Distance [16]—which is an alternative to similarity and it is based on Kolmogorov complexity—and the Set-based Complexity [17], which is aimed at reckoning the heterogeneity of an ensemble of non-random strings.

References

[1] M. A. Boden, Creativity in a nutshell, Think 5 (2007) 83–96.
[2] J. Biles, et al., Genjam: A genetic algorithm for generating jazz solos, in: ICMC, volume 94, 1994, pp. 131–137.
[3] P. McCorduck, Aaron's code: meta-art, artificial intelligence, and the work of Harold Cohen, Macmillan, 1991.
[4] P. Gervás, Computational approaches to storytelling and creativity, AI Magazine 30 (2009) 49–49.
[5] G. Ritchie, R. Manurung, H. Pain, A. Waller, R. Black, D. O'Mara, A practical application of computational humour, in: Proceedings of the 4th International Joint Conference on Computational Creativity, 2007, pp. 91–98.
[6] A. Liapis, G. N. Yannakakis, J. Togelius, Computational game creativity, in: ICCC, 2014.

[7] G. A. Wiggins, A preliminary framework for description, analysis and comparison of creative systems, Knowledge-Based Systems 19 (2006) 449–458.

[8] G. Ritchie, Some empirical criteria for attributing creativity to a computer program, Minds and Machines 17 (2007) 67–99.

[9] M. A. Boden, et al., The creative mind: Myths and mechanisms, Psychology Press, 2004.

[10] P. Sarkar, A. Chakrabarti, J. Gero, Studying engineering design creativity, in: Workshop on Studying Design Creativity, 2008.

[11] C. Shannon, A mathematical theory of communication, The Bell System Technical Journal 27 (1948) 379–423,623–656.

[12] S. McGregor, Algorithmic information theory and novelty generation, Programme Committee and Reviewers (2007) 109.

[13] A. Vinhas, F. Assunção, J. Correia, A. Ekárt, P. Machado, Fitness and novelty in evolutionary art, in: International Conference on Computational Intelligence in Music, Sound, Art and Design, Springer, 2016, pp. 225–240.

[14] I. Koumarelas, A. Kroschk, C. Mosley, F. Naumann, Experience: Enhancing address matching with geocoding and similarity measure selection, Journal of Data and Information Quality 10 (2018) 1–16.

[15] A. Roli, A. Ligot, M. Birattari, Complexity measures: open questions and novel opportunities in the automatic design and analysis of robot swarms, Frontiers in Robotics and AI 6 (2019) 130.

[16] M. Li, X. Chen, X. Li, B. Ma, P. M. Vitányi, The similarity metric, IEEE transactions on Information Theory 50 (2004) 3250–3264.

[17] D. Galas, M. Nykter, G. Carter, N. Price, I. Shmulevich, Biological information as set-based complexity, IEEE Transactions on Information Theory 56 (2010) 667–677.