Research Prospects in the Design and Evaluation of Interactive Evolutionary Systems for Art and Science

Nadia Boukhelifa
Télécom ParisTech & CNRS LTCI
France
nadia.boukhelifa@telecom-paristech.fr

Evelyne Lutton
INRA, Grignon
France
evelyne.lutton@grignon.inra.fr

Anastasia Bezerianos
Univ Paris-Sud & CNRS (LRI), INRIA
France
anastasia.bezerianos@lri.fr

Alberto Tonda
INRA, Grignon
France
alberto.tonda@grignon.inra.fr

Paste the appropriate copyright statement here. ACM now supports three different copyright statements:

- ACM copyright: ACM holds the copyright on the work. This is the historical approach.
- License: The author(s) retain copyright, but ACM receives an exclusive publication license.
- Open Access: The author(s) wish to pay for the work to be open access. The additional fee must be paid to ACM.

This text field is large enough to hold the appropriate release statement assuming it is single spaced in a sans-serif 7 point font.
Every submission will be assigned their own unique DOI string to be included here.

Abstract
We report on our experience in designing and evaluating *seven* applications from *seven* different domains using an interactive evolutionary approach. We conducted extensive evaluations for some of these applications, both quantitative and qualitative, and collected rich feedback from our ongoing collaborations with end-user scientists and artists. To ground our discussion, we refer to two applications, from art and science, as exemplars of our work in order to identify strengths and weaknesses in our approach. We argue that human-centered design could play an important role in addressing some of the identified issues such as the “black box” and the “user-bottleneck” effects. We discuss research opportunities requiring human-computer interaction methodologies in order to support both the visible and hidden roles that humans play in interactive evolutionary computation and machine learning.

Author Keywords
interactive evolutionary computation, machine learning, human-computer interaction, visualization.

ACM Classification Keywords
H.5.2 [User Interfaces]: User-centered design.

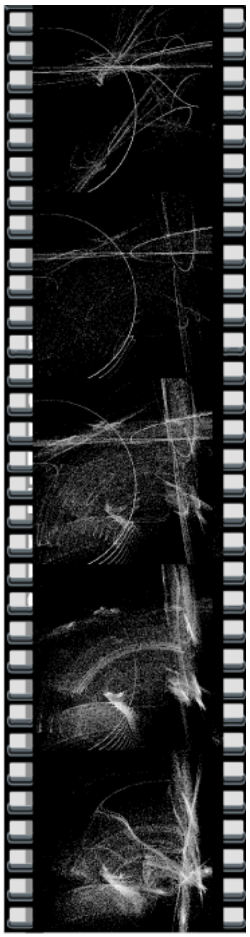

Figure 1: Art work created using *ArtiE-Fract* [10] where a population of images are interactively evolved then morphed (top to bottom) to create a movie clip.

Introduction

Machine learning is a subfield of artificial intelligence concerned with algorithms that discover patterns and knowledge from training data in order to make predictions of future events [7]. There are many approaches to machine learning including artificial neural networks, support vector machines and Bayesian networks. However, many machine learning problems can be modelled as optimisation problems where the aim is to find a trade-off between an adequate representation of the training set and a generalisation capability on unknown samples. In contrast to traditional local optimisation methods, Evolutionary Algorithms (EAs) have been widely used as a successful stochatic optimisation tool in the field of machine learning in the recent years [13]. In this sense, machine learning and the field of Evolutionary Computation (EC) that encompasses EAs are tightly coupled.

Evolutionary Algorithms (EAs) are stochastic optimisation heuristics that copy, in a very abstract manner, the principles of natural evolution that let a population of individuals be adapted to its environment [6]. An EA considers populations of potential solutions exactly like a natural population of individuals that live, fight, and reproduce, but the natural environment pressure is replaced by an "optimisation" pressure. Reproduction (see Fig. 2) consists of generating new solutions via variation schemes (the genetic operators), that, by analogy with nature, are called mutation if they involve one individual, or crossover if they involve two parent solutions. A *fitness function*, computed for each individual, is used to drive the selection process, and is thus optimised by the EA. More specifically, Interactive Evolutionary Computation (IEC) is focused on the optimisation of subjective quantities captured via a user interface (as described in the next section).

In this paper, we report on our experience in building *seven* interactive evolutionary systems from *seven* different domains in art and science (see Table. 1). In all these applications, humans play an important role in many areas, whether in the design of the core of the algorithm, the exploration of the data and the solution space, or in providing feedback and evaluating the output of the evolution or the IEC system itself. Conversely, the underlying computation assists user exploration in converging quickly towards "interesting" views or insight, and supporting creative work. Whereas current IEC research has focused on improving the robustness of the underlying algorithms, much work is needed to tackle human-factors in systems where adaptation between users and systems is likely to occur [11]. Based on our experience, we highlight *three* main research opportunities for HCI and human-centered design in IEC that are also applicable to machine learning in general.

Interactive Evolutionary Computation

Interactive Evolutionary Computation (IEC) describes evolutionary computational models where humans, via suitable user interfaces, play an active role, implicitly or explicitly, in evaluating the outputs evolved by the evolutionary computation (Fig. 2). Applications of IEC range from artistic to scientific projects [14, 10]. IEC lends itself very well to art applications such as for melody or graphic art generation where creativity is essential, due to the subjective nature of the fitness evaluation function. For scientific and engineering applications, IEC is interesting when the exact form of a more generalised fitness function is not known or is difficult to compute, say for producing a visual pattern that would interest a particular user. Here, the human visual system, together with the emotional and psychological responses of the user in question are far superior than any pattern detection or learning algorithm.

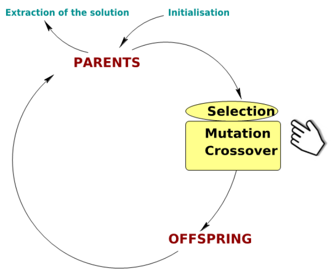

Figure 2: The evolutionary loop: user interactions can occur at any stage including selection/evaluation of individuals and the genetic operators.

system	domain	Fitness Function computed	interactive	learned
ArtiE-Fract [10]	art & design	✓	✓	
EvoGraphDice [4]	agronomy	✓	✓	✓
GraphCuisine [1]	graph theory	✓	✓	
ELISE [8]	databases	✓	✓	
BayesianVis [15]	networks	✓	✓	
E-Learning [16] (based on an Ant colony structure)	e-learning		✓	✓
HEVEA [9]	medecine		✓	

Table 1: Interactive evolutionary systems developed by the authors and components of their fitness functions. *Computed*: an automatically calculated value, *Interactive*: a subjective evaluation captured via a user interface, *Learned*: a value learned over past user interactions.

The Visible & Hidden Roles of Humans in IEC

The role of humans in IEC can be characterised by the evolutionary component at which they operate (Fig. 2), namely: initialisation, evolution, selection, genetic operators, constraints, local optimisation, genome structure variation and parameters tuning, which may or may not be desirable from a usability perspective especially for non-technical users. The general approach when humans are involved, especially for parameter tuning, is mostly by trial-and-error and reducing the number of parameters. Such tasks are often visible, in that they are facilitated by the user interface. However, there exists a hidden role of humans in IEC that has often been neglected. Algorithm and system designers play a central role in deciding the details of the fitness function to be optimised and in setting the default values of system parameters (contributing to the "black box" effect of IEC systems). Such tasks are influenced by the designers' previous experience and end-user task requirements.

Besides this hidden role in the design stage, there is a major impact of the "human in the loop" on the IEC. This problem is known as the "user bottleneck" [12], i.e. a human fatigue due to the fact that the user and machine do not live and react at the same speed. Various solutions have been considered in order to avoid systematic and repetitive or tedious interactions, and the authors themselves have considered several of them, such as: (i) reducing the size of the population and the number of generations; (ii) choosing specific models to constrain the exploration in a-priori "interesting" areas of the search space; and (iii) performing an automatic learning (based on a limited number of characteristic quantities) in order to assist the user and only present interesting individuals of the population, with respect to previous votes or feedback from the user. These solutions require considerable computational effort. A different approach and new ideas to tackle the same issue could come from HCI and usability research (as discussed in the last section of this paper).

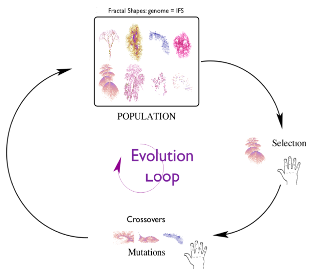

Figure 3: The interactive evolutionary loop of ArtiE-Fract.

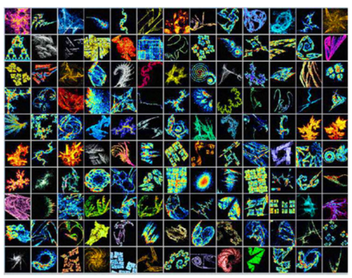

Figure 4: A population of fractals evolved by ArtiE-Fract.

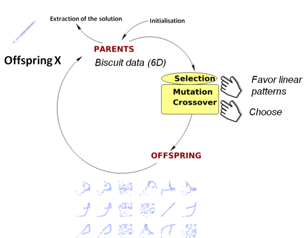

Figure 5: The interactive evolutionary loop of EvoGraphDice (using a biscuits dataset [4]).

IEC Example 1: Fractals for Art and Design

ArtiE-Fract [10] is the result of joint work between engineers, artists and designers to provide a set of user-oriented tools for fractal image design. It is an interactive software that allows the user (artist or designer) to explore the space of fractal 2D shapes with the help of an interactive genetic programming scheme. The user can interfere in the evolution at five different levels: initialisation, fitness function, genome, parameter setting and strategy choices (e.g. how a fractal image is represented in the system).

The direct manipulation of the genome, for instance through the fractals' control points, was also intended to address the "user bottleneck" problem in providing diverse interaction means. We evaluated ArtiE-Fract with a dozen of fashion, art and design students and two professional artists who use it on a regular basis for their work since several years. It appears that direct manipulation, as well as having a variety of interaction techniques at hand (e.g. the geometric modification of fractals can be performed by manually moving the control points or in a semi-automatic fashion using random walks [10] or different morphing methods) seem to reduce user fatigue and allows for more focus on creative work. Thus, on one hand users drive the evolutionary computation in order to converge to an aesthetically-pleasing area of the design space; and on the other hand, the evolutionary computation offers a systematic way to explore this space *creatively* given the input criteria. Indeed, artists' feedback [10] suggests that the right balance between randomness and user-guided search can lead to augmented creativity. For instance, we were able to observe that artists found patterns they had never encountered before and used the system itself in surprising ways (e.g. utilising the underlying database of ArtiE-Fract as a creation tool rather than just a storage facility).

IEC Example 2: Guided Search for Agronomy

EvoGraphDice [4, 2], was designed to aid the exploration of multidimensional datasets where 2D projections of combined dimensions could be of interest to agronomists. Starting from dimensions whose values are automatically calculated by a Principle Component Analysis (PCA), an IEA progressively builds non-trivial viewpoints in the form of linear and non-linear dimension combinations, to help users discover new interesting views and relationships in their data. The criteria for evolving new dimensions is not known a-priori and is partially specified by the user via an interactive interface. Pertinence of views is modeled using a fitness function that plays the role of a predictor: (i) users select views with meaningful or interesting visual patterns and provide a satisfaction score; (ii) the system calibrates the fitness function optimised by the evolutionary algorithm to incorporate user's input, and then calculates new views. A learning algorithm was implemented to provide pertinent projections to the user based on their past interactions.

The evaluation of EvoGraphDice [3, 2] followed a mixed-approach where, on the one hand we observed the utility and effectiveness of the system for the end-user (user-centered approach); and on the other hand we analysed the computational behaviour of the system (algorithm-centered approach). Based on these evaluations, it appears that the interactive evolutionary algorithm, with the help of user feedback, was able to converge quickly to an interesting view when a clear task was specified [2]. In the other direction, the IEC allowed users to laterally explore different possibilities, better formulate their research questions and build new hypotheses for further investigation [4]. However, there were limitations to using the tool, such as the interpretation of the proposed solutions and the lack of feedback regarding the possible "convergence" of the algorithm.

Research Opportunities For Human-Centered IEC

User-driven optimisation processes rely on systems that adapt their behaviour based on user feedback, while users themselves adapt their goals and strategies based on the solutions proposed by the system. This two way communication and adaptation presents prospects to conduct future HCI research for EC. We discuss these below as research opportunities aiming to facilitate and support the different roles humans play in IEC, i.e. in the design, interaction and evaluation of IEC systems.

Human-Centered Design: during the design, development and evaluation of many of our tools, we worked with domain experts at different levels. For EvoGraphDice, for instance, we largely benefited from having a domain expert as part of the design and evaluation team. However, this was carried out in an informal way. Involving end-users in the design team is a long-time tradition in the field of HCI as part of the user-centered design methodology. Participatory design, for instance, could be conducted with IEC end-users to incorporate their expertise in the design of the fitness function. This is a recommendation we should consider in a more systematic way, both as a design and as a system validation approach.

Interaction and Visualization: often the solutions proposed by the IEC are puzzling to the end-user. This is because the inner workings of the evolutionary algorithm and user exploration strategies that led to the solution are often not available to the user. This “black box” effect is challenging to address as there is a fine balance to find between the richness of a transparent interface and the simplicity of a more obscure one. Finding the tipping point requires an understanding of evolving user expertise in manipulating the system, and the task requirements. Whereas HCI and user-centered design can help elicit these requirements and tailor tools to user needs over time, visualization techniques can make the provenance of views and the system status more accessible.

At the interaction level, HCI can contribute techniques to capture rich user feedback without straining the user, that are either implicit: e.g. using eye-tracking; or explicit such as using simple gestures or interactions mediated by tangible objects to indicate user subjective assessment of a given solution. Here, our recommendation is to investigate rich and varied interaction techniques to facilitate user feedback in parallel to developing robust user models that try to learn from the provided input.

Multifaceted Evaluation: the evaluation of an IEC system remains a difficult task as the system adapts to user preferences but also the user interprets and adapts to system feedback. Getting a clear understanding of the subtle mechanisms of this co-adaptation [11] is challenging and requires to consider evaluation criteria other than speed of algorithm convergence and the usability of the interface.

In the context of exploration, both for scientific and artistic applications, creativity is sought and can be characterised by lateral thinking, surprising findings, and the way users learn how to operate the interactive system and construct their own way to use it. Our observation is that augmented creativity can be achieved with the right balance between randomness and user-guided search. What is important to consider for evaluating an IE system, in the context of creativity, are the exploration components. Our recommendation with this respect is two-fold: first, to work towards creating tools that support creativity (something that the HCI community is already looking into [5]); and to investigate objective and subjective metrics to study creativity within IEC and identify impacting factors such as the optimisation constraints, user engagement and the presence or absence of direct manipulation. Some of these measures may only be identifiable through longitudinal observations of this co-adaptation process.

Acknowledgement: This work was supported by the Chair Big Data and Market Insights https://bdmi.wp.mines-telecom.fr/en/.

References

[1] Benjamin Bach, André Spritzer, Evelyne Lutton, and Jean-Daniel Fekete. 2012. Interactive Random Graph Generation with Evolutionary Algorithms. In *Graph Drawing (Lecture Notes in Computer Science)*, Springer (Ed.). Springer.

[2] Nadia Boukhelifa, Anastasia Bezerianos, Waldo Cancino, and Evelyne Lutton. 2015b. Evolutionary Visual Exploration: Evaluation of an IEC Framework for Guided Visual Search. *Evolutionary Computation* (2015).

[3] Nadia Boukhelifa, Anastasia Bezerianos, and Evelyne Lutton. 2015a. A Mixed Approach for the Evaluation of a Guided Exploratory Visualization System. In *EuroVis Workshop on Reproducibility, Verification and Validation in Visualization (EuroRV^3).*

[4] Nadia Boukhelifa, Waldo Cancino, Anastasia Bezerianos, and Evelyne Lutton. 2013. Evolutionary Visual Exploration: Evaluation With Expert Users. *Computer Graphics Forum* 32, 3 (2013), 31 – 40.

[5] Erin Cherry and Celine Latulipe. 2014. Quantifying the Creativity Support of Digital Tools Through the Creativity Support Index. *ACM Transactions on Computer-Human Interaction.* 21, 4 (2014), 21:1–21:25.

[6] David E. Goldberg. 1989. *Genetic Algorithms in Search, Optimization and Machine Learning* (1st ed.). Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA.

[7] Ron Kohavi and Foster Provost. 1998. Glossary of terms. *Machine Learning* 30, 2-3 (1998), 271–274.

[8] Yann Landrin-Schweitzer, Pierre Collet, and Evelyne Lutton. 2006. Introducing Lateral Thinking in Search Engines. *Genetic Programming an Evolvable Hardware Journal.* 1, 7 (2006), 9–31.

[9] P. Legrand, C. Bourgeois-Republique, V. Pean, E. Harboun-Cohen, J. Lévy Véhel, B. Frachet, E. Lutton, and P. Collet. 2007. Interactive evolution for cochlear implants fitting. *GPEM* 8, 4 (2007), 319–354.

[10] Evelyne Lutton. 2006. Evolution of Fractal Shapes for Artists and Designers. *IJAIT, International Journal of Artificial Intelligence Tools* 15, 4 (2006), 651–672. Special Issue on AI in Music and Art.

[11] Wendy Mackay. 2000. Responding to cognitive overhead: co-adaptation between users and technology. *Intellectica* 30, 1 (2000), 177–193.

[12] Riccardo Poli and Stefano Cagnoni. 1997. Genetic Programming with User-Driven Selection: Experiments on the Evolution of Algorithms for Image Enhancement. In *Genetic Programming Conference.* Morgan Kaufmann, 269–277.

[13] Kenneth Stanley and Risto Miikkulainen. 2002. Evolving Neural Networks Through Augmenting Topologies. *Evolutionary Computation* 10, 2 (2002), 99–127.

[14] Hideyuki Takagi. 1998. Interactive evolutionary computation : System optimisation based on human subjective evaluation. In *Proceedings of Intelligent Engineering Systems (INES'98).* IEEE.

[15] Alberto Tonda, Andre Spritzer, and Evelyne Lutton. 2013. Balancing User Interaction and Control in Bayesian Network Structure Learning. In *Artificial Evolution Conference (LNCS 8752).* Springer.

[16] Gregory Valigiani, Evelyne Lutton, Yannick Jamont, Raphael Biojout, and Pierre Collet. 2006. Automatic Rating Process to Audit a Man-Hill. *WSEAS Transactions on Advances in Engineering Education* 3, 1 (2006), 1–7.