SIGEVOlution

newsletter of the ACM Special Interest Group on Genetic and Evolutionary Computation

Volume 11
Issue 2

in this issue

Welcome to the second 2018 newsletter. Our cover illustrates the beauty of nature inspired generative art. In their report, Leonardo Trujillo and Bill Langdon summarise the panelists' views and audience participation in the heated debate about Genetic Programming in the Era of Deep Neural Networks, which took place this April during EuroGP in Parma, Italy. We also report on the success of multi-objective evolutionary algorithms in software testing at Facebook. Regarding forthcoming events, we report on GECCO 2018 taking place this July in the historic and beautiful city of Kyoto in Japan. GECCO will host fascinating keynotes by Japanese researchers in topics ranging from AI for happiness, understanding human states from brain signals, to connecting humans and technology in space. The closing SIGEVO lecture features the influential David E. Goldberg (Dave) who share his views on the tensions between theory, experiment, and practice; and the need for practitioners to become reflective. As ever, please get in touch if you would like to contribute an article for a future issue or have ideas for the newsletter.

Gabriela Ochoa, Editor.

The front cover artwork is based on algorithms developed in the paper "Visual Art inspired by the collective feeding behavior of sand-bubbler crabs" by **Hendrik Richter** which has been presented at EvoMusArt 2018. Sand-bubblers are tiny crabs dwelling tropical beaches. Their feeding behavior involves creating patterns consisting of tiny sand balls that are placed in curves or spirals, straight or bent lines, which finally form overall structures, thus producing astonishing works of natural art. The algorithms presented by the author produce generative art by recreating these patterns. In nature, the patterns are monochromatic as the balls all have the color of the sand they are made from. The artistic interpretation of the patterns suggests using colors for making them visually more appealing.

Interested in learning more about this? You can find more information, images and videos at sand-bubbler art.

Please contact **Hendrik Richter** (hendrik.richter@htwk-leipzig.de) if you would like to discuss this work.

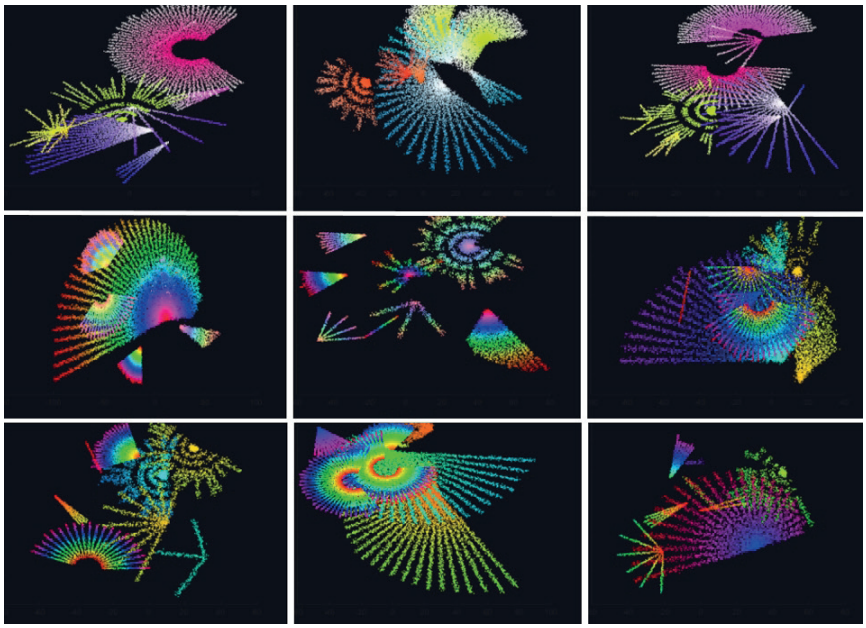

EuroGP 2018 Panel Debate: Genetic Programming in the Era of Deep Neural Networks

Wednesday 4 April 2018, University of Parma, Italy.

As told by Leonardo Trujillo and WB Langdon

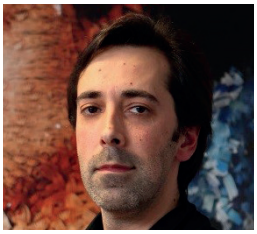

Penousal Machado, University of Coimbra, Portugal

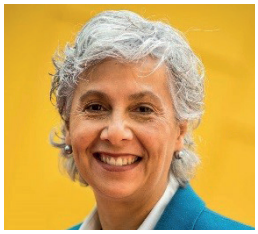

Una-May O'Reilly, MIT, USA.

Marco Gori, University of Siena, Italy.

Sebastian Risi, IT University of Copenhagen, Denmark.

In April this year the programme chairs of EuroGP, the largest conference dedicated to genetic programming, Mauro Castelli and Lukas Sekanina, organised a panel discussion and debate about genetic programming with respect to the undoubted recent successes of other Artificial Intelligence techniques, particularly Deep Learning. The discussion was moderated by Prof. Wolfgang Banzhaf who holds the John R. Koza chair in genetic programming at Michigan State University in the USA. He started by inviting each member of the panel to give their position:

Sebastian Risi:

- Evolutionary computation (EC) can be a "more creative" tool than Deep Learning
- It may also be useful to train Artificial Neural Nets with millions of parameters

Marco Gori:

- It seems that EC and Deep Learning may be related, but how?
- EC can be used to select or construct Artificial Neural Nets (such as some of the examples presented in EuroGP this year)
- EC may also be used to select the proper features of an Artificial Neural Net.
- Another topic is that most machine learning is supervised learning, deriving models based primarily on the training data. Another approach is to build machines that are consistent with specific rules related to learning, which can be converted to constraints, adopting tools from multivalued logic like fuzzy systems (his arguments can be clearly seen in his recent book Machine Learning: A Constraint-Based Approach)

Una-May O'Reilly:

- I have not yet combined Deep Learning and EC
- It seems that the main challenge in Deep Learning is the training phase
- EC can, in principle, replace gradient descent. But EC, such as evolution strategies, e.g. CMA-ES, can be efficient.
- EC can also be used to evolve structure and topology, such as NEAT and many of its variants.
- EC and GP can strive for comprehensibility. GP can be used to perform feature extraction in Deep Learning, since current Deep Learning produces black box Artificial Neural Nets. (LT: these arguments, mentioned also above, seem to be similar to those echoed by the GP community for several years now, Deep Learning is just an extreme case of incomprehensible models.)

- It can be possible to do adversarial learning for Deep Learning with EC techniques (using concepts and techniques from co-evolution), but how would examples differ for white box or black box models.

Penousal Machado:

- Penousal started by saying that he wants to be "critical" of current work.
- He has mixed feelings about the current Deep Learning hype, which is no doubt exciting, but everything is now done with Deep Learning, even if some work is only re-inventing the wheel (LT: this seems to happen frequently, also in GP), some of the work in Google Deep mind is very similar to Blondie24 by David Fogel published in 2001.
- There are many opportunities, but for many in EC they can only afford to do poor-man's machine learning (big money projects are not in EC or GP) and EC is computationally costly (on top of Deep Learning).
- Some workarounds are possible, such as EDAs or the use of surrogate fitness functions.
- Can we put "nature" back into Artificial Neural Nets, using insights regarding how the brain is constructed, to exploit modularity and transfer learning. (LT: this is related to a recent technique by Olague et al., called Brain Programming; but these ideas also run counter intuitive to the state-of-the-art in Deep Learning, see for instance the book by Goodfellow et al. who seem to suggest that much of the work in Deep Neural Nets has moved on from the biological inspiration.)

Marco Gori:

- Mostly agrees with previous points (particularly Penousal Machado), but there is a reason for the emphasis and hype in Deep Learning, and it is not because it is new, there has not been a "paradigm shift", as defined by Kuhn,
- Something is happening, it is simply because companies push machine learning / Deep Learning, and the reason is funding. This is probably the reason we are here today, can you "compete with the Dragon".
- Deep Learning experiments are expensive and time consuming. But there are many unused resources, idle computing hardware that is not exploited sufficiently. (LT: it brings to mind work done by Paco Fernandez on volunteer computing for EC, such as BOINC.)
- To get funding you need to show the community that you can exploit the computing power. Can we construct agents that can continuously learn and exploit the computational power that we do have before acquiring more?
- Another issue is that many in EC do not consider time. Not in the engineering sense, but in the scientific sense. How does time effect the way in which learning should take place.

Wolfgang Banzhaf:

- Time is very important!
- It is necessary to get creativity and novelty.
- The requirement of stochastic events may be hindrance for EC.
- Can we build tools that can go on and on, that can evolve continuously. (WBL: Wolfgang and I discussed long-term evolution at the recent GI Dagstuhl seminar.)

Sebastian Risi:

- For instance, can we branch out from previous runs of our EC algorithms?
- This seems possible in EC but not in Deep Learning.

Una-May O'Reilly:

- Deep Learning uses transfer learning quite often, which I relate to a sourdough starter in baking!

- There are systems that build up a corpus of text and model trained on this corpus. It is possible to train your system on this corpus to get a head-start on your model, and then contribute what you learn to the corpus (this is common in chat-bot design such as wit.ai)
- Branching out from previous runs seems like a very promising approach.

Leonardo Vanneschi:
- Image analysis was very rare in GP, but not anymore! We should not narrow our application domain.

Wolfgang Banzhaf:
- The success of Deep Learning is powered by new technology, exploiting massive parallelism, and this is very useful in image and video analysis.

Una-May O'Reilly:
- Images are just there, they are easy to generate and get now, in other domains it is more difficult to produce the required data.

Sebastian Risi:
- A lot of emphasis is in image recognition. Images seem simpler than much of the domains studied in EC and GP. (LT: some would disagree.)
- Animals cannot learn images without temporal information. A video with shuffled frames is meaningless to an animal but a learning algorithm can still learn something from unordered video frames.
- Other problems are also important.

Una-May O'Reilly:
- A child learns with much less information and effort than a Deep Neural Nets. There is innate structure in a child that is not present in Deep Learning models. (LT: this echoes the seminal work by Noam Chomsky).
- The Baldwin and Lamarckian models could be relevant again, they seem important in this context.

Sebastian Risi:
- We need to consider not just evolutionary learning but lifetime learning as well.
- Also, it is easier to measure success in Deep Learning because they have standardized benchmarks and tools, this has many benefits but it can also have some drawbacks (maybe limiting creativity and novel approaches)

Bill Langdon/ Una-May O'Reilly:
- Frederic Gruau worked on how there are different things you can learn at different time-scales.
- Cultural Evolution occurs across generations. (LT: there is interesting work by Daniel Dennett on this subject.)

Ernesto Costa:
- Deep Learning has a lot of hype in engineering problems, but results are not understandable. Can EC help? Comprehensible results are important.
- There is also a problem with time, we cannot wait 3 months for a solution, and some problems require solutions much faster than this.

Una-May O'Reilly:
- We can wait 3 months or more, because we can re-use a result. How long would a handmade model take to construct?

- One-shot learning should be studied more, learning from few examples.

Marco Gori:
- Time is a crucial question, it needs to be exploited further.
- Adversarial learning in Deep Learning has shown that many results are not robust, they are easily fooled.
- It is time to reconsider what is machine learning, which is currently based on the supervised learning approach. It is necessary to move to a constraint based approach. Not simply function approximation, but learning that is consistent with specific rules, rules than can be formulated as constraints on the learning process.

Sebastian Risi:
- Deep Learning needs millions of examples, humans do not learn this way, and they require much less examples (for most tasks).

Mengjie Zhang:
- DL does not depend on NNets, it is the stacked approach that seems the most important, and this is not something new from DL, only the scale is new. GP can do this, EC can be used to evolve both the architecture and the weights of a DL model.

Una-May O'Reilly:
- Deep Learning work is not always published, it is at the industrial and commercial level.
- How can we re-use datasets, what are the right datasets.
- Stressed the need for EC to develop shared open source tools and datasets but recognized the dangers of a whole field focusing on a single set of benchmarks.

Marco Gori:
- It is possible that in the next few years the importance and emphasis on data will be reduced.
- The Deep Learning approach is to throw all the data to a learning problem (some may say that EC also does this). It may be a better strategy to mine things that are already there.

Una-May O'Reilly:
- The importance of self-organization and emergence.

Marco Gori:
- Explaining AI is important, models need to be interpretable. It is necessary to convert sub-symbolic knowledge into logical constraints.

Final comments by the attendees:
- **James Foster**: The state of theory is much better in EC than in machine learning. Life took a long time to evolve intelligence, we should be prepared to run our artificial evolution for more generations. Let them run for 500 years. He also said in EC it is possible to analyze what has been evolved: what has it learnt. Finally the (human) self is not just one identity but a community of many different types of cell and organisms working together, albeit with some level of competition.
- **Leonardo Trujillo**: GP often promises interpretable models and solutions, but we rarely demand it from work inside our community. Another interesting feature of current Deep Learning approaches is that they rely on relatively simple and old techniques, and many "good heuristics", and we in GP and EC are good at such *tricks*.

Evolutionary Algorithms for Software Testing in Facebook

Mark Harman is widely known in our community for his work on source code analy-sis and testing and was instrumental in the founding of the field of Search Based Software Engineering (SBSE). Since its inception in 2001, SBSE has rapidly grown to include over 800 authors, from 270 institutions spread over 40 countries.

Since February 6th, 2017, **Mark** has been a full-time Engineering Manager at Face-book London and a part-time professor of Software Engineering in CREST (the Centre for Research on Evolution Search and Testing), and the Computing Science Department at UCL (University College London).

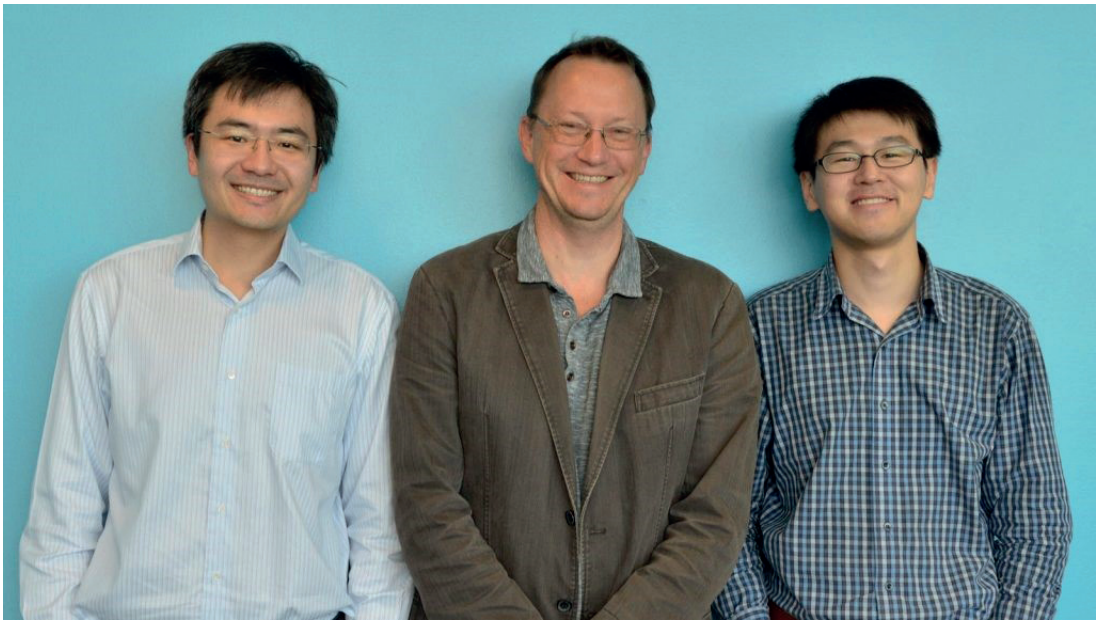

Yue Jia (left), Mark Harman (center) & Ke Mao (right).

About a year ago, the UCL team behind spinout software testing company MaJiCKe, moved to work with Facebook in London. The company has three co-founders **Mark Harman** (Scientific Advisor), **Yue Jia** (CEO), and **Ke Mao** (CTO). Here is the link to the announcement in Facebook.

The key technology this team has deployed at Facebook is called Sapienz –Multiobjective Automated Android Testing. A short video of Sapienz can be found here, and a paper describing it was published in 2016 at ISSTA, the leading research symposium on software testing and analysis.

Sapienz is an approach to Android testing that uses multi-objective evolutionary algorithms to automatically explore and optimise test sequences, minimising length, while simultaneously maximising coverage and fault revelation. It is in production now helping to improve the quality of Facebook software!

GECCO 2018

This year the Genetic and Evolutionary Computation Conference (GECCO) will take place in Kyoto, a major historical city and former imperial capital of Japan. Kyoto is home of 17 UNESCO World Heritage Sites and many more Japanese national treasures. GECCO presents the latest high-quality results in genetic and evolutionary computation since 1999. Topics include: genetic algorithms, genetic programming, ant colony optimization and swarm intelligence, complex systems (artificial life/robotics/evolvable hardware/generative and developmental systems/ artificial immune systems), digital entertainment technologies and arts, evolutionary combinatorial optimization and metaheuristics, evolutionary machine learning, evolutionary multiobjective optimization, evolutionary numerical optimization, real world applications, search-based software engineering, theory and more.
The GECCO 2018 program and schedule is now available.

Best Paper Candidates

Each of the GECCO tracks (or combinations for small tracks) will be giving a best-paper award. The nominees can be found here.

GECCO 2018 INVITED SPEAKERS

The event will host 3 enticing keynote talks and a SIGEVO lecture.

AI for Happiness of People

Kazuo Yano
Fellow, Corporate Officer, Hitachi, Ltd., Tokyo, Japan.

Exploitation of Bio Signal Data to Understand Human State

Tatsuya Okabe
AI R&I Division and Value Innovation Division, DENSO Co., Ltd.

©NASA

Connecting Human and Technologies in Space

Naoko Yamazaki
The University of Tokyo, former JAXA astronaut.

SIGEVO Lecture - On Becoming a Reflective Practitioner

David E. Goldberg
ThreeJoy Associates, Big Beacon, and University of Illinois (Emeritus).

About this newsletter

SIGEVOlution is the newsletter of SIGEVO, the ACM Special Interest Group on Genetic and Evolutionary Computation. To join SIGEVO, please follow this link: [WWW]

Contributing to SIGEVOlution

We solicit contributions in the following categories:

Art: Are you working with Evolutionary Art? We are always looking for nice evolutionary art for the cover page of the newsletter.

Short surveys and position papers: We invite short surveys and position papers in EC and EC related areas. We are also interested in applications of EC technologies that have solved interesting and important problems.

Software: Are you are a developer of an EC software and you wish to tell us about it? Then, send us a short summary or a short tutorial of your software.

Lost Gems: Did you read an interesting EC paper that, in your opinion, did not receive enough attention or should be rediscovered? Then send us a page about it.

Dissertations: We invite short summaries, around a page, of theses in EC-related areas that have been recently discussed and are available online.

Meetings Reports: Did you participate to an interesting EC-related event? Would you be willing to tell us about it? Then, send us a short summary, around half a page, about the event.

Forthcoming Events: If you have an EC event you wish to announce, this is the place.

News and Announcements: Is there anything you wish to announce, such as an employment vacancy? This is the place.

Letters: If you want to ask or to say something to SIGEVO members, please write us a letter!

Suggestions: If you have a suggestion about how to improve the newsletter, please send us an email.

Contributions will be reviewed by members of the newsletter board.

We accept contributions in LATEX, MS Word, and plain text.

Enquiries about submissions and contributions can be emailed to editor@sigevolution.org

All the issues of SIGEVOlution are also available online at: www.sigevolution.org

Notice to Contributing Authors to SIG Newsletters

By submitting your article for distribution in the Special Interest Group publication, you hereby grant to ACM the following non-exclusive, perpetual, worldwide rights:

- to publish in print on condition of acceptance by the editor
- to digitize and post your article in the electronic version of this publication
- to include the article in the ACM Digital Library
- to allow users to copy and distribute the article for noncommercial, educational or research purposes

However, as a contributing author, you retain copyright to your article and ACM will make every effort to refer requests for commercial use directly to you.

Editor: **Gabriela Ochoa**
Associate Editors: **Emma Hart**, **James McDermott**, **Una-May O-Reilly** & **Darrell Whitley.**
Design & Layout: **Kate Simpson**