SIGEVOlution

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

Volume 10
Issue 3

in this issue

GECCO 2017 Report & Statistics

Genetic Improvement Workshop @ GECCO 17

Turing Award - Celebrating 50 years

EDITORIAL

Welcome to the 3rd issue of the SIGEVO newsletter this year!

Many of you will have enjoyed the wonderfully organised GECCO 2017 in Berlin earlier this summer. As ever the standard and variety of the work presented was fantastic - we certainly found that there were so many interesting papers it was difficult to choose where to go for each session! Although the hard work put in Peter Bosman and Gabriela Ochoa as General Chair and EiC was acknowledged many times during the conference, they definitely deserve one last shout out for their efforts! You can find a summary report from Gabriela inside which contains a fascinating statistical analysis of the submission data. Still on the subject of GECCO, there is also a report on the Genetic Improvement work- shop from Brendan Cody-Kenny. The final feature is a reflection on the Turing Award Celebrations from the 5 lucky students who were sponsored by SIGEVO to attend this event.

We hope everyone's experiments for GECCO 2018 papers are already running (only 4 months to the submission date!) – stay tuned for more details about GECCO 2018 in Kyoto in later newsletters. As ever, please get in touch if you would like to contribute an article for a future issue or have ideas for the newsletter.

Emma Hart and **Gabriela Ochoa**, Editors

The front cover artwork is based on methods developed in the paper "Evolutionary Image Transition Using Random Walks" by **Aneta Neumann**, **Bradley Alexander** and **Frank Neumann** which has been presented at EvoMusArt 2017. The authors use evolutionary computation methods for image transition and design different mutation operators based on uniform and biased random walks. The Evolutionary Image Transition (EIT) creates, in a surprising manner, a wide range of images with vibrant visual effects. The methods are used to inspire the creating of aesthetic artworks such as paintings.

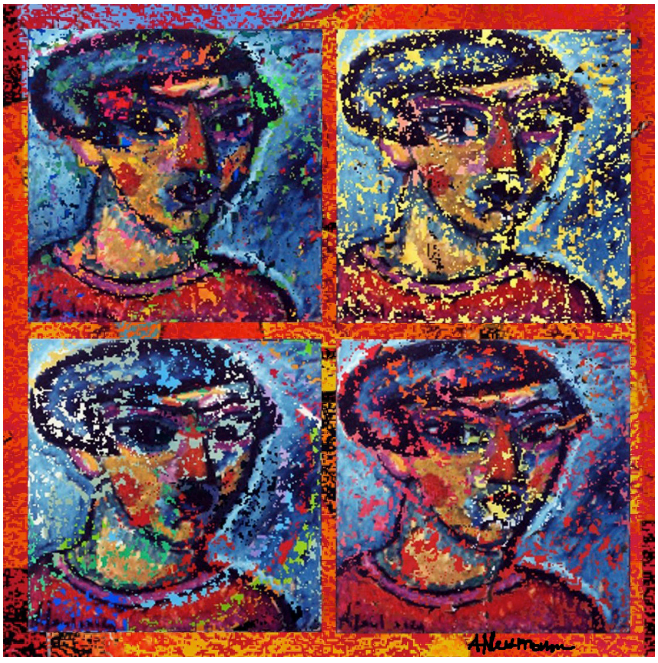

Interested in learning more about this? You can find more information about this and other evolutionary artworks at https://researchers.adelaide.edu.au/profile/aneta.neumann.

Please contact Aneta Neumann if you would like to discuss this work.

GECCO Statistics and Collaboration Network

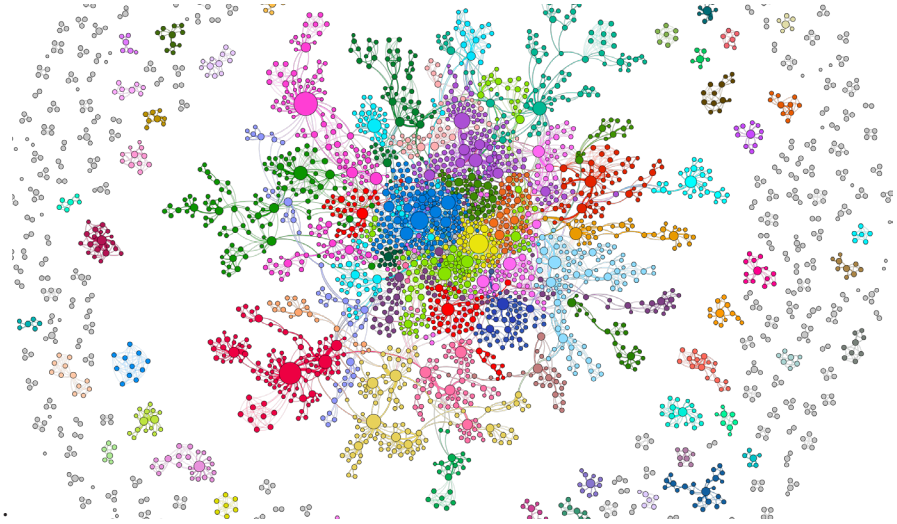

By **Gabriela Ochoa** and **Nadarajen Veerapen**

Computing Science and Mathematics, University of Stirling, Scotland, UK.

We conducted an analysis of the GECCO 2017 submission data. We also considered historical submission data since 2005 when the proceedings started to be handled by the ACM and thus accessible via the ACM digital library. Considering the full papers appearing in the proceedings since 2005, we extracted and visualised the co-authorship network.

2017 Data

Table 1 summarises the 2017 submissions. There was a first stage of abstract submissions, which did not fully translate to paper submissions. The poster acceptance groups submissions as full papers that were accepted as posters and accepted poster submissions.

	Submissions	Acceptances	%
Abstract	595		
Papers	462	178	39
Posters	32	19	59

Table 1: Overview of GECCO 2017 submissions and acceptances.

This year GECCO had 13 tracks, in alphabetical order: ACO-SI – Ant Colony Optimization and Swarm Intelligence, CS – Complex Systems, DETA – Digital Entertainment Technologies and Arts, ECOM – Evolutionary Combinatorial Optimization and Metaheuristics, EML – Evolutionary Machine Learning, EMO – Evolutionary Multiobjective Optimization, ENUM – Evolutionary Numerical Optimization, GA – Genetic Algorithms, GECH – General Evolutionary Computation and Hybrids, GP – Genetic Programming, RWA – Real World Applications, SBSE – Search-Based Software Engineering, THEORY – Theory. The Complex Systems track joins several previously separate tracks including: Artificial Life, Artificial Immune Systems, Generative and Developmental Systems, Evolutionary Robotics and Evolvable Hardware.

Figure 1 shows submissions and full-paper acceptances across the 13 tracks ordered by the number of submissions. The overall acceptance rate for the conference was 39% visualised as a purple line, with the acceptance rate per track visualised with a red line (for illustration purposes, no trend is implied) fluctuating between 32% (ACO-SI) and 58% (THEORY).

Figure 1: GECCO 2017 full-paper submissions and acceptances by track.

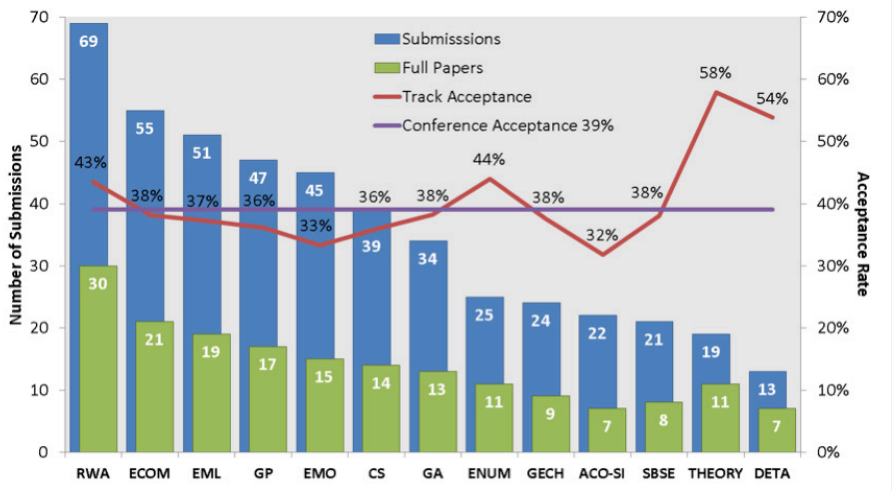

Where do GECCO 2017 authors come from?

A total of 837 authors coming from 50 countries contributed to accepted GECCO 2017 full papers and posters. Figure 2 visualises their countries of affiliation. The area of the circles is proportional to the number of authors, and the colours identify continents as indicated in the legend (also including percentages). The country with the largest number of authors was the United States with 131 authors, followed by the UK with 65. France and Germany share the 3rd place with 59 authors. Interestingly, Brazil is located in the 4th position, the only South American country featuring a significant number of authors (51).

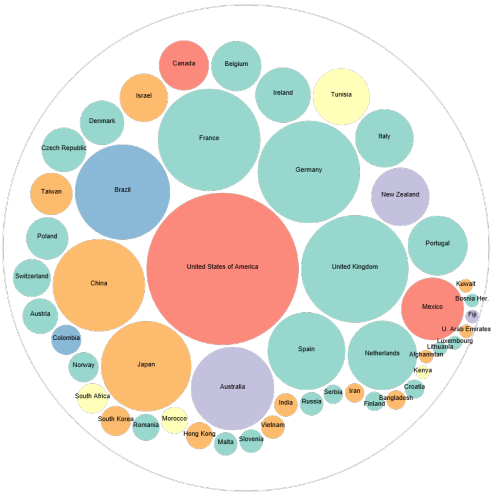

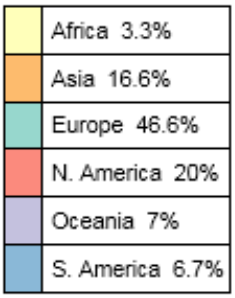

Figure 2: Countries of affiliation of 2017 papers and posters authors.

If we group the countries by continent (Figure 3), we can clearly observe that Europe has the largest contingent of GECCO 2017 authors. With North America and Asia following.

Figure 3: Countries of affiliation of 2017 papers and posters authors grouped by continent.

2005 – 2017 Data

Figure 4 shows submissions and full-paper acceptances by year since GECCO 2005 in Washington D.C. The submission numbers show an oscillating dynamic, where lower figures are observed when GECCO takes place in the United States. The acceptance rates are also shown with the red line helping to see a trend. With the exception of Portland in 2010, there has been a decreasing acceptance rate from 2005 until 2014. Since 2014 this trend has reversed and acceptance rates have increased.

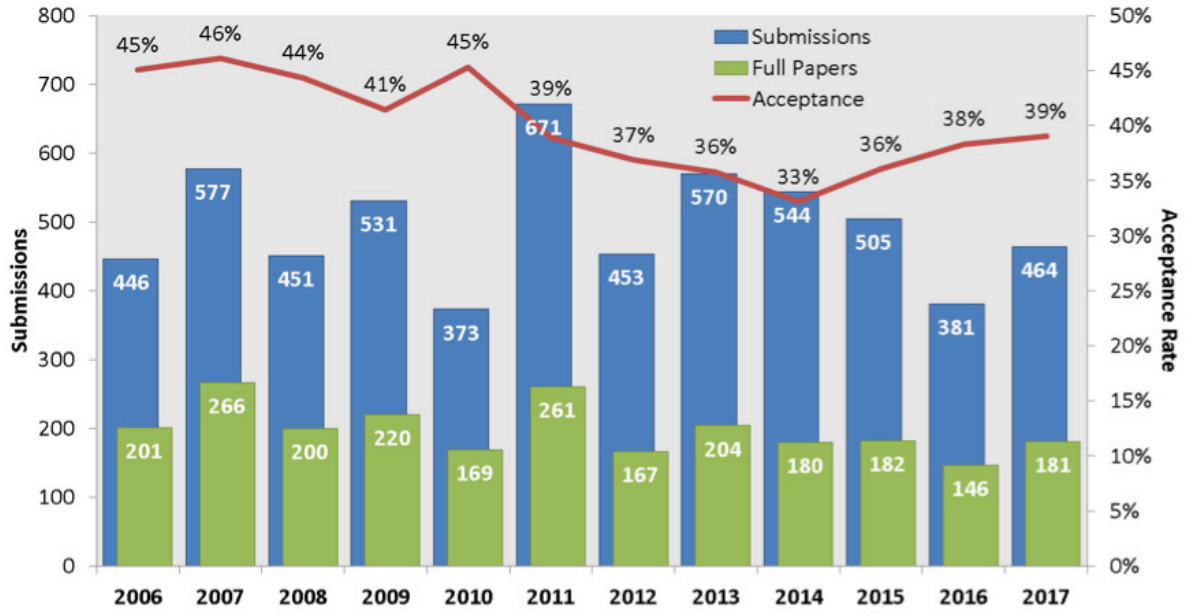

Figure 4: Full-paper submissions and acceptances by year since GECCO 2005

Figure 5 illustrates the dynamic of GECCO tracks' relative size (width) and rank (order from top to bottom) over the last 12 years using a so-called 'bump chart' or 'sorted stream graph'. By 'track size', we mean the number of published papers. Interestingly, the track ranks and sizes change over time. The plot reveals that the largest tracks have been GA (orange), RWA (Real World Applications, brown), and the now called CS (Complex Systems, blue), which combines previous related separate tracks. The GA track has recently clearly decreased, with ECOM (Evolutionary Combinatorial Optimization and Metaheuristics, green) and EML (Evolutionary Machine Learning, pink) in turn gaining in rank and size.

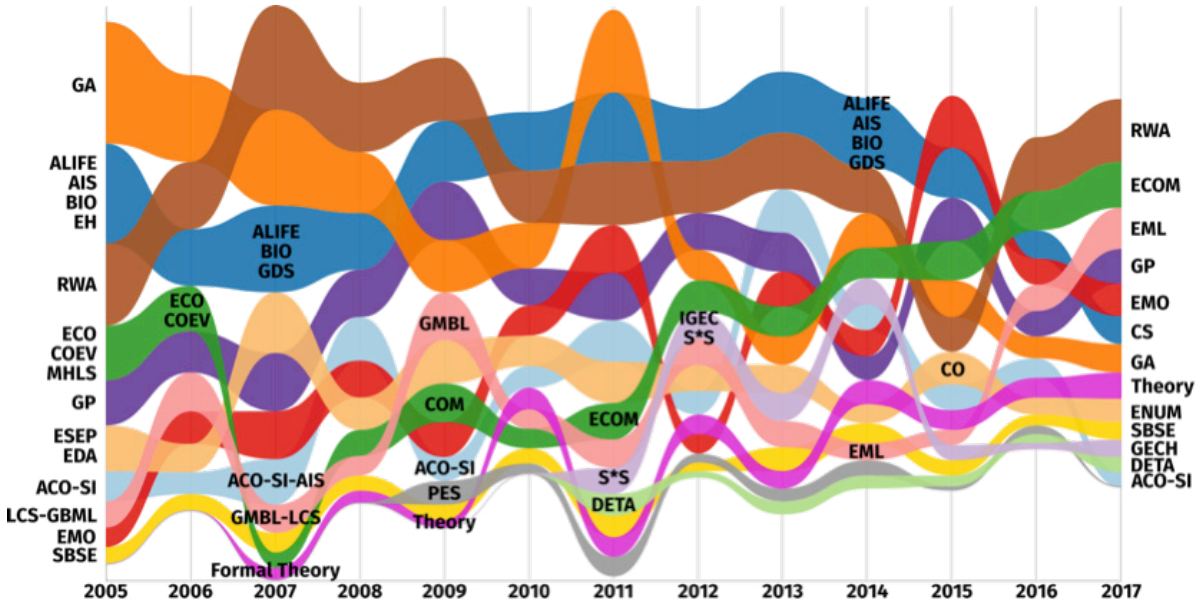

Figure 5: Size (number of papers) and rank of GECCO tracks 2005

Co-authorship Network

We extracted the GECCO co-authorship network[1] from 2005 until 2017. Each node is an author, and two authors are connected if they have co-authored one or more GECCO papers. The network does not include co-authorships in journals or any other conference. There is a single link between any pair of authors, even if they have more than one paper together. The strength of the connection is based on the number of papers co-authored by pairs of scientists, and the total number of authors of those papers. Table 2 overviews the main network metrics.

The giant component groups more than half of GECCO authors, which is a common characteristic of social networks. The average distance (number of links) between two authors is 7.3, which indicates a small-world. The distribution of the number of collaborators per author (Figure 6) is also characteristic of social networks, featuring a long tail. The majority of authors have 3 collaborators, few prefer to work on their own, and there is a small number of 'hubs' with many collaborators.

total authors	3822
total papers	2614
papers per author	avg = 2.0 max = 40
authors per paper	avg = 2.9 max =12
collaborators per author	avg = 3.7 max = 51
no. of components	504
giant component	1963 (51%)
2nd largest component	26 (0.7%)
distance	avg = 7.3, max =15.1
clustering coefficient	0.4

Table 2: GECCO co-authorship network metrics.

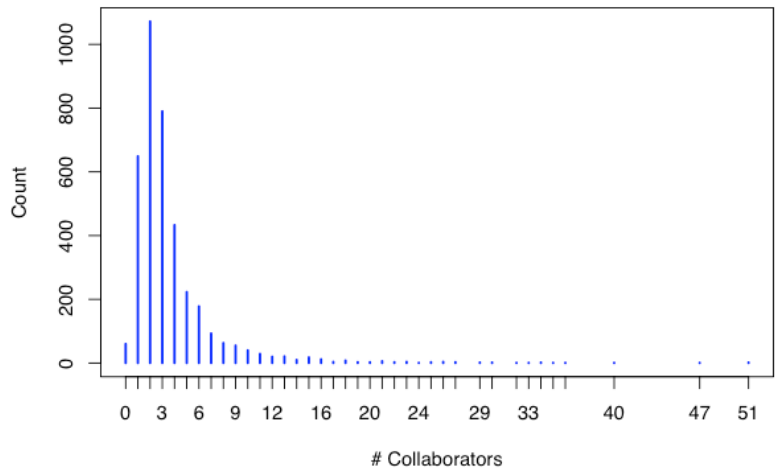

Figure 6: Distribution of the number of collaborators per author.

1 The GECCO co-authorship network is available online at http://www.cs.stir.ac.uk/~goc/gecco-network/index.html

Genetic Improvement Workshop at GECCO 2017

By **Brendan Cody-Kenny**[1]

The Genetic Improvement (GI) Workshop at GECCO this year was an event in itself. It had it's own keynote delivered by Wolfgang Banzhaf, followed by a mixture of long and short papers, and finished with best paper awards just before closing. From 10:40am to 6pm, one of the three largest workshops at GECCO covered bug-fixing, energy optimisation, new operators and landscape analysis.

An interesting trend at the workshop was the use of GP as part of a programmer's workflow. This is hardly surprising, as GI seeks to use Genetic Programming to improve existing programs and is closely related to Search-Based Software Engineering. In the workshop we saw the evolution of programs in C++, Java and Python in environments as diverse as mobile phones and GPUs. To evolve programs in these languages and environments, GI implementations are heavily influenced by concepts and techniques in Software Engineering. GI research at the workshop also drew heavily from GP. We saw landscape analysis of large programs and heard the need for new operators to traverse the sometimes jagged landscapes in this domain.

The conference was opened by David White who introduced Wolfgang Banzhaf as keynote speaker. Wolfgang presented on the origins of Evolutionary Computation and considered the goal of Automated Programming in the context of Software Engineering. Wolfgang highlighted that existing software appears to lack uniqueness, and questioned whether concepts such as regulatory networks could be used to reassemble code. Saemundur Haraldsson described how bugs in a live system were fixed overnight using GP. The result of nightly GP runs were presented to developers each morning, demonstrating the use of GP as part of a programmers daily workflow. Shin Yoo continued on this theme by proposing that GP be included in programming languages. This would allow developers to define at design time what parts of their program are amenable to evolution at run-time.

After the break, Jason Landsborough began the next session with mutation analysis of open source GNU core utilities binaries to find semantics-preserving program transforms beyond those considered by compilers.Jason presented findings showing roughly 50% of the binary code for a program can be modified without failing any tests. Brad Alexander presented the challenges involved in defining a fitness function for energy optimisation on Android

A large and attentive crowd at the Genetic Improvement Workshop, GECCO, Berlin 2017.

1 codykenny@gmail.com

Smartphones. Accurate measures of program energy or run-time can be difficult to gather due to variation in noisy environments.Michael Orlov suggested extracting building blocks from tree-based representations of Java bytecode, adding structure to otherwise linear bytecode so it can be more effectively evolved. David White presented GIN, a GI implementation for evolving patches targeting the evolution of Java source code. The focus of this work is providing a simple reference implementation and a set of benchmark programs ideal for getting started in GI.Aniko Ekart raised the possibility of using GP to improve road traffic models, with a focus on reusing existing symbolic regressors and time series models. Justyna Petke suggested that new operators could be derived by analysing program variants in software repositories, using more formal methods to guide mutation or by analysing program variants produced during the GP process.

In the last session, Nadarajen Veerapen showed results of landscape analysis for two C programs, showing these programs are relatively robust to mutations. Further to this point, Saemundur Haraldsson showed how a bioinformatics program is robust to mutation, but also that mutations which do not affect functionality can yield performance improvements. Bill Langdon gave the penultimate talk on performance improvement of C code, re-affirming software robustness and the ability to find performance improvements.I then posed questions on how software engineering might change if we could perform landscape analysis on all existing software in a language. After an audience vote, Justyna presented best paper awards to Saemundur and Nadarajen.

Overall the workshop had a heavy focus on analysing existing code. We saw that GP techniques are highly relevant to this analysis.Given the wide range of software development environments, the amount of existing source code and the number of requirements a software engineer has to consider, I was left with the sense that there is a deep well of questions for the GP researcher.

To get more detail on these ideas you can visit the GI website which lists all papers.

Excellent food and lively discussion at the social event following the Genetic Improvement Workshop, GECCO, Berlin 2017.

About the author

Brendan Cody-Kenny is a Post-Doctoral Researcher in the Natural Computing Research and Applications Group (NCRA) at University College Dublin. He holds a DPhil in Computer Science from Trinity College Dublin. His current research interests include Performance Improvement and Software Recomposition toward the goal of Automated Programming.

ACM Marks 50 Years of the ACM A.M. Turing Award and Computing's Greatest Achievements.

Since its inauguration in 1966, the ACM A. M. Turing Award has recognized major contributions of lasting importance in computing. Through the years, it has become the most prestigious technical award in the field, often referred to as the "Nobel Prize of computing."

During 2017, ACM celebrated 50 years of the Turing Award and the visionaries who have received it. The aim was to highlight the significant impact of the contributions of the Turing Laureates on computing and society, to look ahead to the future of technology and innovation, and to help inspire the next generation of computer scientists to invent and dream.

The celebrations culminated in a conference on the present and future in computing on June 23-24 in San Francisco, CA that was attended by numerous Turing laureates as well as other distinguished researchers and entrepreneurs.

ACM and its special interest groups offered scholarships to students so they could attend the conference and meet the laureates. SIGEVO was delighted to provide sponsorship to 5 students who share their reflections on the conference below.

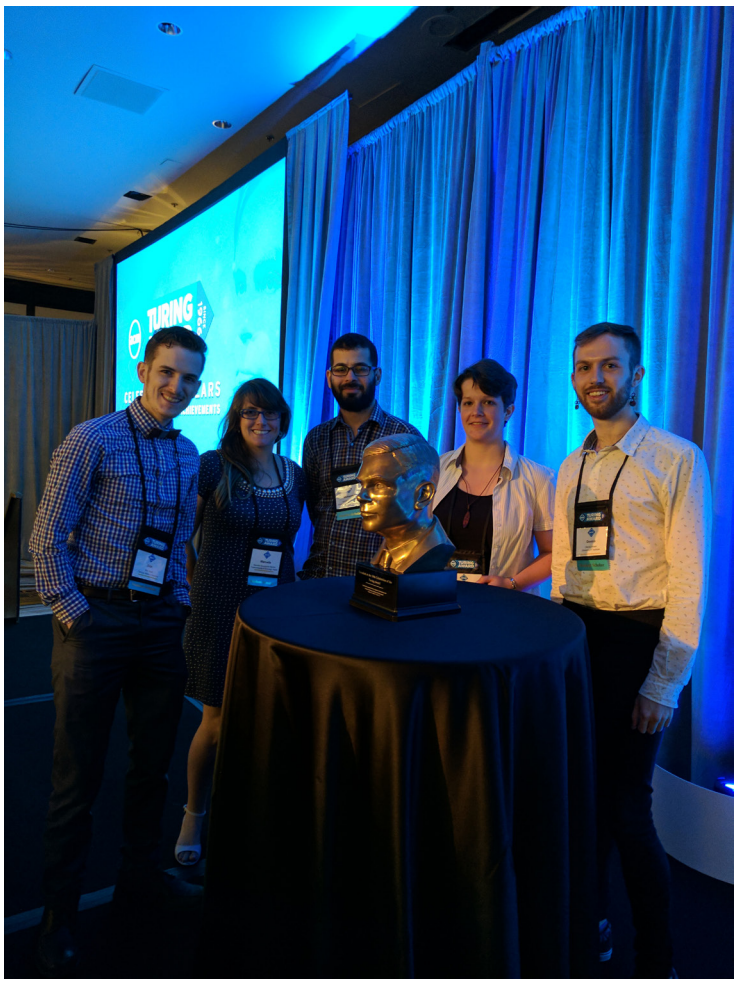

Eric Scott, Marcella Scoczynski Ribeiro Martins, Mohamed El Yafrani, Vanessa Volz and Dennis G Wilson next to the bust of Turing unveiled at the conference.

Overview

The conferences consisted of several very diverse panels, commencing with an overview of the first decade of the Turing award and the work of its recipients. The first panel discussed the popular topic of Deep Learning, with panelists from both sides of the aisle giving praise and criticism for recent applications of deep learning. The celebration continued to a panel discussing privacy and security, which involved a heated discussion over the responsibility of security researchers to assist government agencies.

The next panel was on digital preservation, with a focus on the openness of information and the importance of executability preservation. After that, there was a discussion on computing in a post-Moore's Law world. The day finished with a panel on ethics, especially focusing on biases in data mining.

The next day started with a brief talk and excellent Q&A session with Donald Knuth, and then continued to a theoretical view of quantum computing with speculation on the near future of the field. The last panel reviewed the history of augmented reality and highlighted current problems and directions.

Student Reflections

Vanessa Volz:

I found the fact that the audience was much more diverse in comparison to most conferences very refreshing and it resulted in many enlightening conversations for me. For instance, I was able to meet two researchers with similar interests as me, but with completely different approaches. Apart from my rediscovered appreciation for flat landscapes and sunscreen (thanks San Francisco!), my three main takeaways are the following:

- The panel on digital preservation renewed my motivation to preserve and make accessible my data and algorithms. For this reason, I have recently started to campaign for statistically sound benchmarks with verifiable results in my field of research.
- Even with deep learning, you still need to pay for your lunch. Properly characterising a problem and forgoing complete black-box end-to-end solution might result in better performance (see AlphaGo), despite the bias you introduce.
- I was delighted to see that people from very different backgrounds, including Donald Knuth, endorse using games as interesting research problems and benchmarks for AI and other fields of computing.

Eric Scott:

On a personal level, what I took away from the conference is just how much of a shared language and culture there is across computer science. Sometimes it can feel like we are "way over here" in evolutionary computation doing our own thing, with our own lingo and ideas that have nothing to do with systems, databases, or expert systems, etc.

But when you look at the Turing awards, you get a sense that there is this big, core computer science worldview that transcends sub-disciplines. This is also true of AI more locally: as we listened to the panelists debate the limits of deep learning and the future of AI, I couldn't help but notice how often they appealed to the same tension between general-purpose problem solving and specialized domain knowledge that we are always wrestling with in EC. Edward Feigenbaum, for instance, believes that the old slogan "knowledge is power" applies just as surely to modern forms of machine learning and AI as it does to expert systems.

Mohamed El Yafrani:

The experience I had in the Turing award celebration was quite different from my previous experiences in other conferences. I had the chance to meet world-renowned researchers and students from various fields of computing. Being part of such a diverse community was very interesting and insightful. I was pleased to see that many researchers from other areas of computing that I have met have an interest in evolutionary computation (e.g. machine learning, theorem proving, software testing).

I particularly appreciated the panel about deep neural networks, especially Pr. Russell's critique and comments about the limitations of deep neural networks. I also liked very much the discussions about Moore's law, quantum computing, and privacy.

Dennis G Wilson:

A theme in the ceremony that I noticed from the introductory talk through most of the panels was the extent to which the application of computer science defines the field. The first panel was split around this divide, with advocates for deep neural networks praising the application advances made with this technology and other panelists urging that this alone wasn't enough to prove deep learning's worth. In the security panel, application was paramount as Nadia Heninger aptly explained that virtually all security flaws lie in the application of different algorithms and not in the underlying cryptography. The quantum computing panel was in stark contrast to this. There was a large amount of discussion about the differences of quantum computing theory and the need for further work in this domain in tandem with the admission that application of quantum computing may still be too far off to predict its utility.

Computer science as a field is closely tied to its application. The extent to which application leads the direction of the field, and the responsibility of computer scientists to influence the application of their work, seemed variable amongst the ceremony presenters. In the small group session after the main ceremony that I attended, Vint Cerf spoke about his current work guiding the development of the Internet. Following his part in the design of TCP/IP, Vint Cerf has taken an active part in guiding the application of this and other technologies as the Internet has grown. His dedication to this unwieldy force that has undoubtedly shaped the world was a testament to both the power a computer scientist has to influence the application of their work and the limitations in doing so. I was very motivated by this ceremony to consider the possible applications, both positive and negative, that might result from my own work.

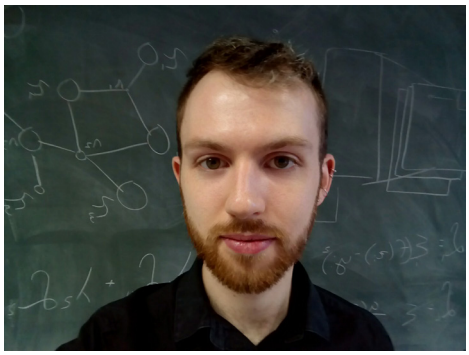

Marcella Scoczynski Ribeiro Martins:

Beyond the celebration giving us the possibility to share our experiences, gain new knowledge, and listen to and ask questions at the panel discussions, it has motivated and inspired us in our own research. The event has provided insights for us and for the future generations in order to use computation to contribute to solving the main issues of our society.

Proposals for FOGA 2019

As announced at the GECCO-2017 business meeting, SIGEVO is soliciting proposals for hosting FOGA-2019. The deadline for submitting hosting proposals is September 30, 2017. Submissions should be electronically sent to **Darrell Whitley** and **Kenneth De Jong**. Selection will be made by the SIGEVO board and notifications sent out by October 15, 2017.

Calls for papers

EvoCOP 2018 - The 18th European Conference on Evolutionary Computation in Combinatorial Optimisation

April 5 - April 6, 2018
Parma, Italy
Part of Evo* 2018

The 18th European Conference on Evolutionary Computation in Combinatorial Optimisation is a multidisciplinary conference that brings together researchers working on evolutionary computation methods and other metaheuristics for solving difficult combinatorial optimisation problems appearing in various industrial, economic, and scientific domains. Prominent examples of metaheuristics include: evolutionary algorithms, estimation of distribution algorithms, swarm intelligence methods such as ant colony and particle swarm optimisation, local search methods such as simulated annealing, tabu search, variable neighbourhood search, iterated local search, scatter search and path relinking, and their hybridisation, such as memetic algorithms. Automatic algorithm configuration and design, meta-optimisation, model-based methods, and hyperheuristics are also topics of interest.

Important dates

- Submission deadline (Regular papers): **November 1, 2017**
- Submission deadline (LBAs): **January 15, 2018**
- EvoStar: **April 4-6, 2018**

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