SIGEVOlution

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

Volume 12
Issue 3

in this issue

Editorial

Welcome to the third 2019 newsletter! Our cover, a contribution from **Catarina Maçãs** and **Penousal Machado** from the University of Coimbra, Portugal, illustrates how Evolutionary Computation and Swarm Intelligence can be combined with Data Aesthetics to visualize retail consumption data. Our main article describes the winning entries of the sixteenth annual Human-Competitive Results Awards – "Humies" 2019, as reported by its host **Erik Goodman** the director of the BEACON Center for the Study of Evolution in Action at Michigan State University. The winning entries this year illustrate the power of evolutionary computation for designing wireless communication networks, cancer treatment and software data structures. We continue with a contribution from **Bill Langdon** at the University College London. **Bill** argues that evolutionary algorithms and other optimization techniques can be used for maintaining years of legacy software systems, not by automatically changing code, but instead the data associated to the programs. Finally, we overview the major evolutionary computation events taking place in 2020, be ready with your submissions and do not miss the deadlines.

As ever, please get in touch if you would like to contribute an article for a future issue or have suggestions for the newsletter.

Gabriela Ochoa, Editor.

About the Cover

The artwork of the front cover was created by **Catarina Maçãs** and **Penousal Machado** from the University of Coimbra, Portugal, through the combination of Data Aesthetics, Swarm intelligence, and Evolutionary Computation. This system is capable of generating artefacts that are meant to represent and distort the everyday retail consumption data. The visualisation rationale is based on a set of graphing bodies — swarming boids — that represent different types of consumption through colour, and are attracted or repulsed by the neighbours, depending on their type of consumption. Their size is influenced by the amount of consumption, which in turn influence the swarming forces. The visual artefacts are portraits of consumption as they represent the evolution of consumption patterns over time. All portraits are unique, as in the generation process the evolutionary system defines the parameters, such as the forces of attraction and repulsion of the graphing bodies, creating disturbances in the drawing of the visual artefacts. With this, the system achieves a diverse set of artefacts, although maintaining the visualisation rationale. Through the use of an Interactive Evolutionary Algorithm, the system allows the users to guide the evolution towards artefacts that are more meaningful and aesthetically pleasant to them. In the end, each visual artefact may not be truly representative of the data, but is influenced by it, which makes the artefacts a distorted portrayal of consumption.

The visualisation model was first presented at IJCAI 2015 in the paper "Swarm Systems in the Visualisation of Consumption Patterns", and the EA system in ArtsIT 2018 in the paper "Interactive Evolution of Swarms for the Visualisation of Consumptions"

For more details see: https://cdv.dei.uc.pt/ie-of-swarms-in-visualisation/

Human-Competitive Results Awards – “Humies” 2019 – Announces Winners at GECCO

Erik Goodman

Humies host for 2019, Michigan State University, USA

The Sixteenth Annual Humies Competition, which awards $10,000 in cash prizes for computational results deemed to be competitive with results produced by human beings, but are generated automatically by computer, was held as part of the GECCO-2019 program, the Genetic and Evolutionary Computation Conference, held July 13-17 in Prague, Czech Republic. The awards, sponsored by **John Koza** (who is widely acknowledged as the “Father of Genetic Programming”) annually solicit newly published papers that describe work fulfilling one or more of eight criteria, including such features as winning a regulated competition against humans or other programs, producing results that are publishable in their own right, not because they were created by a computer program, patentability, and several others, described fully on the competition’s website, www.human-competitive.org.

This year’s competition attracted 19 entries, and the complete entries are available at the website just above. They spanned very diverse domains, from design of medical treatments to urban traffic management to automatic optimization of space and runtime in computer programs to automatic discovery of new statistical tests, and many other amazing results. Eleven entries were selected as finalists, all judged to meet one or more of the criteria for being judged human-competitive. All eleven projects were presented for the judges and a public audience at GECCO on Monday, July 15, and the winners were announced and awarded certificates at the SIGEVO Annual Meeting and GECCO Award Ceremony on Wednesday, July 17, 2019.

GECCO
GECCO 2019 @ Prague
David
Lynch
University College Dublin
HUMAN COMPETITIVE RESULTS AWARD -- "HUMIES" 2019
Presented at the
Genetic and Evolutionary Computation Conference
July 13-17, 2019 Prague
Czech Republic
Winners of the
2019 Humies Gold Award
David Lynch, Michael Fenton, David Fagan, Stepan Kucera, Holger Claussen, Michael O'Neill
Automated Self-Optimisation in Heterogeneous Wireless Communications Networks
GECCO
Prague
sigevo
ThinkPad

Gold Award

The Gold and a $5,000 prize was awarded to a team including **David Lynch**, **Michael Fenton**, **David Fagan** and **Michael O'Neill**, of University College, Dublin, and **Stepan Kucera** and **Holger Claussen**, of Nokia Bell Labs in Dublin. Their paper was entitled, "Automated Self-Optimization in Heterogeneous Wireless Communications Networks," and appeared in IEEE/ACM Transactions on Networking, 27.1 (2019): 419-432. Their evolved, dynamic, self-optimizing algorithm produces, on-line, significantly better cell-edge throughput and less interference in a complex cellular telephone network than the algorithms written by humans or by other computer programs, and shows promise for deployment in next-generation networks. The presentation was made by **David Lynch**, a Ph.D. student in the O'Neill laboratory, who is shown above accepting the Gold Award at GECCO from **Erik Goodman**.

Silver Award

The Silver Award and $3,000 were earned by a team from the Department of Radiation Oncology at the University of Amsterdam Medical College, including **Stef Maree**, **Ernst S. Kooreman**, **Niek van Wieringen**, **Arjan Bel**, **Karel A. Hinnen**, **Henrike Westerveld**, **Bradley R. Pieters**, **Tanja Alderliesten**, and from the Life Science and Health Research Group of the Centrum Wiskunde & Mathematica (Center for Computer Science and Mathematics) in Amsterdam, Netherlands, including **Ngoc Hoang Luong** and **Peter A. N. Bosman**. Their paper, entitled "Evaluation of bi-objective treatment planning for high-dose-rate prostate brachytherapy - A retrospective observer study," appeared in the medical journal Brachytherapy, 18 (3), 396-403, 2019. The team devised an algorithm that designs a treatment plan for radioactive sources inserted into pre-implanted catheters for treatment of prostate cancer, such that the tumor receives a maximal dose while surrounding tissues receive minimal doses. The algorithm runs in real-time, after the catheters have been implanted but before the sources are inserted, dramatically reducing the waiting time from the 1-hour-plus time required when doctors design the plan. In trials, panels of doctors showed a strong preference for the machine-designed plans over those designed by human physicians, and the program is going into clinical use. Left-to-right, **Peter Bosnan**, **Tanja Alderliesten** and **Stef Maree** are shown accepting the Silver Award from **Goodman**.

Bronze Award

The Bronze Award and a $2,000 prize were given to a team from University College, London, including **Michail Basios**, **Lingbo Li**, **Fan Wu**, **Leslie Kanthan**, and **Earl T. Barr** for the work presented in their paper "Darwinian Data Structure Selection" that appeared in the Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE 2018), ACM, New York, NY, USA, 118-128. In their paper, they describe the capability of their genetic programming (or genetic improvement) algorithm to examine software from a third party and systematically to replace data structures in that software in order to reduce both size and runtime of the resulting software, in a completely automatic fashion. It can improve software written in a variety of languages, such as Java and C++. It has been demonstrated to work effectively on real-world software packages routinely used by millions of users. Pictured is **Michael Basios** accepting the Bronze Award from **Erik Goodman**.

Entries for the 2020 Seventeenth Annual Human-Competitive Results Awards will be accepted beginning in May, 2020. If you are publishing human-competitive results during the next year, you are cordially invited by **John Koza** and the other Humies organizers to visit the website www.human-competitive.org and to submit your work for this recognition, and if you are selected as a finalist, to present your work at GECCO-2020 in Cancun, Mexico, in July, 2020! We hope to see you there!

Big Data Driven Genetic Improvement for Maintenance of Legacy Software Systems

VC
XL
TL
AS
Part of gp–bibliography 2144 942 Revision:1.1088 25 Apr 2006

W. B. Langdon

Abstract

Software is vital to modern life, yet much of it is old and suffers from bit-rot. There are not and never will be enough software experts to keep it all up to date by hand. Instead we suggest combining data driven learning with evolutionary search to maintain computer systems.

We are all familiar with the idea that big data is transforming the way we use computers, e.g. automatically translating Chinese web pages into English. However we also know that computers are often held back by the vast volume of existing old, outdated software that legacy systems are obliged to continue to use. What if we could in some way combine the potential of big data to bring legacy software from the 1960s into the third millennium? We have recently begun investigating the power of artificial evolution in the form of genetic improvement of software to bring about radical change in how we maintain and update computer software.

Genetic Improvement (GI) [20, 21] uses search based software engineering (SBSE) [6] techniques, principally genetic programming [9, 2, 22], to improve existing human written code. For example, to automatically fix bugs [16] or speed up programs [11]. Although genetic improvement can be applied to program binaries [24], assembler [23] and byte code [26, 17] it has mostly updated program's source code [10], e.g. C and Java. However, recently we have started applying search to optimise *data* to improve existing software.

In Data Driven Genetic Improvement we apply SBSE to the maintenance of legacy systems via the data often embedded in programs. In some cases this can be done without code changes. In others, minimal code changes may be needed. Whilst there are examples of GI modified code being maintained by software engineers in the traditional way [12, 13, 5, 4, 1], one of the potential advantages of data driven GI is avoiding code changes so potentially making it easier to test [19] and reducing fears that mixing automatically generated code with human written code will lead to more difficult to understand systems.

Recently the USA government[1] has predicted that continued growth in the software industry will lead to an annual increase of 24% in the number of people employed as software developers. Such exponential growth is clearly not sustainable in the long term. (If continued, in twenty years time most of the US population would be Software Developers!) Thus there will be an huge demand for the available software developers to be more productive and we suggest that this will be achieved by increased automation of software production, particularly of software maintenance.

As with any learning approach, one of the bottle necks to progress in automatically searching for better software is obtaining sufficient training data to drive evolution. It is here that big data might step in. For example, in [11] we statistically sub-sampled from millions of DNA sequences and showed by continuously changing from where the sample was drawn, evolution could learn from a huge data source at a reasonable rate.

So far genetic improvement has concentrated upon genetic programming and other evolutionary approaches [25, 4, 8]. However, if we are primarily dealing with numbers (rather than programming languages code), other search or optimisation techniques, such as IRace [18] and ParamILS [7], might also be applied to improving code via big data.

In [15] we showed data driven GI could automatically improve the predictions of a state-of-the art C program which predicts RNA folding by updating data only. The new data give on average an improvement: some predictions are worse and some better but overall they are more accurate. The new data have been shipped in new releases since 2018.

Other experiments, using CMA-ES [3], have shown (albeit including limited manual code changes) that in some cases fitness directed evolution of data can result in new functionality [14].

There are potentially many data driven programs where maintenance is delayed either because people with experience of them have left or retired or where the in-house expertise is desperately overloaded so that they do not have large units of time to devote to the tedious and error prone update of the program's data. In addition to scientific and medical research programs lagging years behind publication of new results we may consider automatic updating of data driven programs to keep pace with new social and technical regulations.

References

[1] N. Alshahwan. Industrial experience of genetic improvement in Facebook. In J. Petke, S. H. Tan, W. B. Langdon, and W. Weimer, editors, GI-2019, ICSE workshops proceedings, page 1, Montreal, 28 May 2019. IEEE. Invited Keynote.

[2] W. Banzhaf, P. Nordin, R. E. Keller, and F. D. Francone. Genetic Programming – An Introduction; On the Automatic Evolution of Computer Programs and its Applications. Morgan Kaufmann, San Francisco, CA, USA, Jan. 1998.

[3] N. Hansen and A. Ostermeier. Completely derandomized self-adaptation in evolution

[1] US Bureau of Labor Statistics occupational outlook handbook Software Developers, 2016-26.

strategies. Evolutionary Computation, 9(2):159–195, Summer 2001.

[4] S. Haraldsson, A. Brownlee, and J. R. Woodward. Computers will soon be able to fix themselves - are IT departments for the chop? The Conversation, page 3.29pm BST, Oct. 12 2017.

[5] S. O. Haraldsson, J. R. Woodward, A. E. I. Brownlee, and K. Siggeirsdottir. Fixing bugs in your sleep: How genetic improvement became an overnight success. In J. Petke, D. R. White, W. B. Langdon, and W. Weimer, editors, GI-2017, pages 1513–1520, Berlin, 15-19 July 2017. ACM. Best paper.

[6] M. Harman and B. F. Jones. Search based software engineering. Information and Software Technology, 43(14):833–839, Dec. 2001.

[7] F. Hutter, H. H. Hoos, K. Leyton-Brown, and T. Stuetzle. ParamILS: An automatic algorithm configuration framework. JAIR, 36:267–306, 2009.

[8] G. Kendall. Evolutionary computation has been promising self-programming machines for 60 years - so where are they? The Conversation, page 8.54am BST, Mar. 27 2018.

[9] J. R. Koza. Genetic Programming: On the Programming of Computers by Natural Selection. MIT press, 1992.

[10] W. B. Langdon. Genetic improvement of programs. In R. Matousek, editor, 18th International Con- ference on Soft Computing, MENDEL 2012, Brno, Czech Republic, 27-29 June 2012. Brno University of Technology. Invited keynote.

[11] W. B. Langdon and M. Harman. Optimising existing software with genetic programming. IEEE Transactions on Evolutionary Computation, 19(1):118–135, Feb. 2015.

[12] W. B. Langdon, B. Y. H. Lam, J. Petke, and M. Harman. Improving CUDA DNA analysis soft- ware with genetic programming. In S. Silva, A. I. Esparcia-Alcazar, M. Lopez-Ibanez, S. Mostaghim, J. Timmis, C. Zarges, L. Correia, T. Soule, M. Giacobini, R. Urbanowicz, Y. Akimoto, T. Glasmach- ers, F. Fernandez de Vega, A. Hoover, P. Larranaga, M. Soto, C. Cotta, F. B. Pereira, J. Handl, J. Koutnik, A. Gaspar-Cunha, H. Trautmann, J.-B. Mouret, S. Risi, E. Costa, O. Schuetze, K. Kraw- iec, A. Moraglio, J. F. Miller, P. Widera, S. Cagnoni, J. Merelo, E. Hart, L. Trujillo, M. Kessentini, G. Ochoa, F. Chicano, and C. Doerr, editors, GECCO ’15: Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation, pages 1063–1070, Madrid, 11-15 July 2015. ACM.

[13] W. B. Langdon and R. Lorenz. Improving SSE parallel code with grow and graft genetic programming. In J. Petke, D. R. White, W. B. Langdon, and W. Weimer, editors, GI-2017, pages 1537–1538, Berlin, 15-19 July 2017. ACM.

[14] W. B. Langdon and J. Petke. Evolving better software parameters. In T. E. Colanzi and P. McMinn, editors, SSBSE 2018 Hot off the Press Track, volume 11036 of LNCS, pages 363–369, Montpellier, France, 8-9 Sept. 2018. Springer.

[15] W. B. Langdon, J. Petke, and R. Lorenz. Evolving better RNAfold structure prediction. In M. Castelli, L. Sekanina, and M. Zhang, editors, EuroGP 2018: Proceedings of the 21st European Conference on Genetic Programming, volume 10781 of LNCS, pages 220–236, Parma, Italy, 4-6 Apr. 2018. Springer Verlag.

[16] C. Le Goues, M. Pradel, and A. Roychoudhury. Automated program repair. Communications of the ACM. To appear.

[17] M. Orlov. Evolving software building blocks with FINCH. In J. Petke, D. R. White, W. B. Langdon, and W. Weimer, editors, GI-2017, pages 1539–1540, Berlin, 15-19 July 2017. ACM.

[18] L. Perez Caceres, M. Lopez-Ibanez, H. Hoos, and T. Stuetzle. An experimental study of adaptive capping in irace. In R. Battiti, D. E. Kvasov, and Y. D. Sergeyev, editors, Learning and Intelligent Optimization - 11th International Conference, LION 11, Nizhny Novgorod, Russia, June 19-21, 2017, Revised Selected Papers, volume 10556 of Lecture Notes in Computer Science, pages 235–250. Springer, 2017.

[19] J. Petke. Constraints: The future of combinatorial interaction testing. In 2015 IEEE/ACM 8th International Workshop on Search-Based Software Testing, pages 17–18, Florence, May 2015.

[20] J. Petke, S. O. Haraldsson, M. Harman, W. B. Langdon, D. R. White, and J. R. Woodward. Genetic improvement of software: a comprehensive survey. IEEE Transactions on Evolutionary Computation, 22(3):415–432, June 2018.
[21] J. Petke, M. Harman, W. B. Langdon, and W. Weimer. Specialising software for different downstream applications using genetic improvement and code transplantation. IEEE Transactions on Software Engineering, 44(6):574–594, June 2018.
[22] R. Poli, W. B. Langdon, and N. F. McPhee. A field guide to genetic programming. Published via http://lulu.com and freely available at http://www.gp-field-guide.org.uk, 2008. (With contri- butions by J. R. Koza).
[23] E. Schulte, S. Forrest, and W. Weimer. Automated program repair through the evolution of assembly code. In Proceedings of the IEEE/ACM International Conference on Automated Software Engineering, pages 313–316, Antwerp, 20-24 Sept. 2010. ACM.
[24] E. Schulte, W. Weimer, and S. Forrest. Repairing COTS router firmware without access to source code or test suites: A case study in evolutionary software repair. In W. B. Langdon, J. Petke, and D. R. White, editors, Genetic Improvement 2015 Workshop, pages 847–854, Madrid, 11-15 July 2015. ACM. Best Paper.
[25] J. R. Woodward, J. Petke, and W. Langdon. How computers are learning to make human software work more efficiently. The Conversation, page 10.08am BST, June 25 2015.
[26] K. Yeboah-Antwi and B. Baudry. Online genetic improvement on the java virtual machine with ECSELR. Genetic Programming and Evolvable Machines, 18(1):83–109, Mar. 2017.

Forthcoming Conferences

evostar

Seville, Spain
15-17 April 2020

EvoStar comprises of four co-located conferences run each spring at different locations throughout Europe. These events arose out of workshops originally developed by EvoNet, the Network of Excellence in Evolutionary Computing, established by the Information Societies Technology Programme of the European Commission, and they represent a continuity of research collaboration stretching back over 20 years.

The four conferences include:
EuroGP 23rd European Conference on Genetic Programming
EvoApplications 23rd European Conference on the Applications of Evolutionary and bio-inspired Computation
EvoCOP 20th European Conference on Evolutionary Computation in Combinatorial Optimisation
EvoMUSART 9th International Conference (and 14th European event) on Evolutionary and Biologically Inspired Music, Sound, Art and Design

Important Dates

EvoApplications Special Session proposals: 10 September, 2019
Submission deadline: 1 November, 2019
EvoStar Conference: 15-17 April, 2019

GECCO 2020 @ Cancun

The Genetic and Evolutionary Computation Conference

July 8th-12th 2020

Important Dates

Workshop Papers, Hot off the Press, Late-breaking Abstracts and Competitions abstracts

Abstract submission deadline: January 30, 2020

Full paper submission deadline: February 6, 2020

The Genetic and Evolutionary Computation Conference (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 IEEE World Congress on Computational Intelligence (IEEE WCCI) is the world's largest technical event in the field of computational intelligence. WCCI 2020 features the flagship conference of the Computational Intelligence Society: The 2020 International Joint Conference on Neural Networks (IJCNN 2020), the 2020 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE 2020), and the 2020 IEEE Congress on Evolutionary Computation (IEEE CEC 2020) under one roof. It encourages cross-fertilisation of ideas among the three big areas and provides a forum for intellectuals from all over the world to discuss and present their research findings on computational intelligence.

Important Dates

15 Nov 2019	Special Session & Workshop Proposals Deadline
15 Dec 2019	Competition and Tutorial Proposals Deadline
15 Jan 2020	Paper Submission Deadline
15 Mar 2020	Paper Acceptance Notification Date
15 April 2020	Final Paper Submission and Early Registration Deadline

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 purpose

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** and **Darrell Whitley**
Design & Layout: **Kate Simpson**