

RESEARCH ARTICLE

Artificial intelligence in peer review: How can evolutionary computation support journal editors?

Maciej J. Mrowinski^{1*}, Piotr Fronczak¹, Agata Fronczak¹, Marcel Ausloos², Olgica Nedic^{3,4}

1 Faculty of Physics, Warsaw University of Technology, Koszykowa 75, PL-00-662, Warsaw, Poland, **2** School of Management, University of Leicester, University Road, Leicester LE1 7RH, United Kingdom, **3** Institute for the Application of Nuclear Energy (INEP), University of Belgrade, Banatska 31b, Belgrade-Zemun, Serbia, **4** Sub-editor in the Journal of the Serbian Chemical Society, Karnegijeva 4/III, 11120 Belgrade, Serbia

* mrow@if.pw.edu.pl



Abstract

With the volume of manuscripts submitted for publication growing every year, the deficiencies of peer review (e.g. long review times) are becoming more apparent. Editorial strategies, sets of guidelines designed to speed up the process and reduce editors' workloads, are treated as trade secrets by publishing houses and are not shared publicly. To improve the effectiveness of their strategies, editors in small publishing groups are faced with undertaking an iterative trial-and-error approach. We show that Cartesian Genetic Programming, a nature-inspired evolutionary algorithm, can dramatically improve editorial strategies. The artificially evolved strategy reduced the duration of the peer review process by 30%, without increasing the pool of reviewers (in comparison to a typical human-developed strategy). Evolutionary computation has typically been used in technological processes or biological ecosystems. Our results demonstrate that genetic programs can improve real-world social systems that are usually much harder to understand and control than physical systems.

OPEN ACCESS

Citation: Mrowinski MJ, Fronczak P, Fronczak A, Ausloos M, Nedic O (2017) Artificial intelligence in peer review: How can evolutionary computation support journal editors? PLoS ONE 12(9): e0184711. <https://doi.org/10.1371/journal.pone.0184711>

Editor: Sergi Lozano, Institut Català de Paleoeologia Humana i Evolució Social (IPHES), SPAIN

Received: January 2, 2017

Accepted: August 29, 2017

Published: September 20, 2017

Copyright: © 2017 Mrowinski et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All relevant data are within the paper and its Supporting Information files.

Funding: This publication is supported by the COST Action TD1306 "New frontiers of peer review" (http://www.cost.eu/COST_Actions/tdp/TD1306). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Introduction

Peer reviewed publications remain the cornerstone of the scientific world [1–5], despite the fact that the review process itself is far from perfect: The authors are frustrated by having to wait for the delaying reports [6, 7] and the editors are irritated by the time-consuming tasks of searching for and rounding up an excessive number of reviewers in the hope of receiving one or two reports. Surprisingly, despite the importance of peer review, it remains an understudied subject [8]. Existing papers concentrate on statistical or ethical issues relating to the process [9–11], with peer review only recently capturing the attention of complex systems scientists [12, 13]. The goal of our work, which is a result of a collaboration within COST action PEERE (New Frontiers of Peer Review; www.peere.org), is to provide quantitative improvements to the efficiency of the review process.

Competing interests: The authors have declared that no competing interests exist.

One of the biggest obstacles to research on peer review is the scarce availability of data. Although it is relatively simple to obtain information about citations or publication dates, the inner workings of the editorial process are usually kept secret. However, thanks to the cooperation with our partners at PEERE, we were able to acquire a very detailed dataset from one of the sub-editors of the Journal of the Serbian Chemical Society (JSCS). The dataset contained information about 58 articles submitted to the journal for publication and allowed us to study the review process from the viewpoint of a journal editor. We were able to separate the process into distinct phases (Fig 1a), determine the duration of each phase and the probability of transitions between phases [14]. Most importantly, we were able to create a simulation platform capable of generating artificial review threads (i.e. realisations of the review process). Formally, each review thread represents a path through the review process decision diagram (Fig 1b), with durations assigned to each transition. To our knowledge, most existing peer review simulations are based on abstract models. In contrast, our approach allows for data-driven simulations, which may lead to a quantifiable improvement in the efficiency of the peer review process.

Specifically, the main focus of the present study was how to improve the workflows of editors and decrease the review time. (i.e. the average number of days needed to acquire the

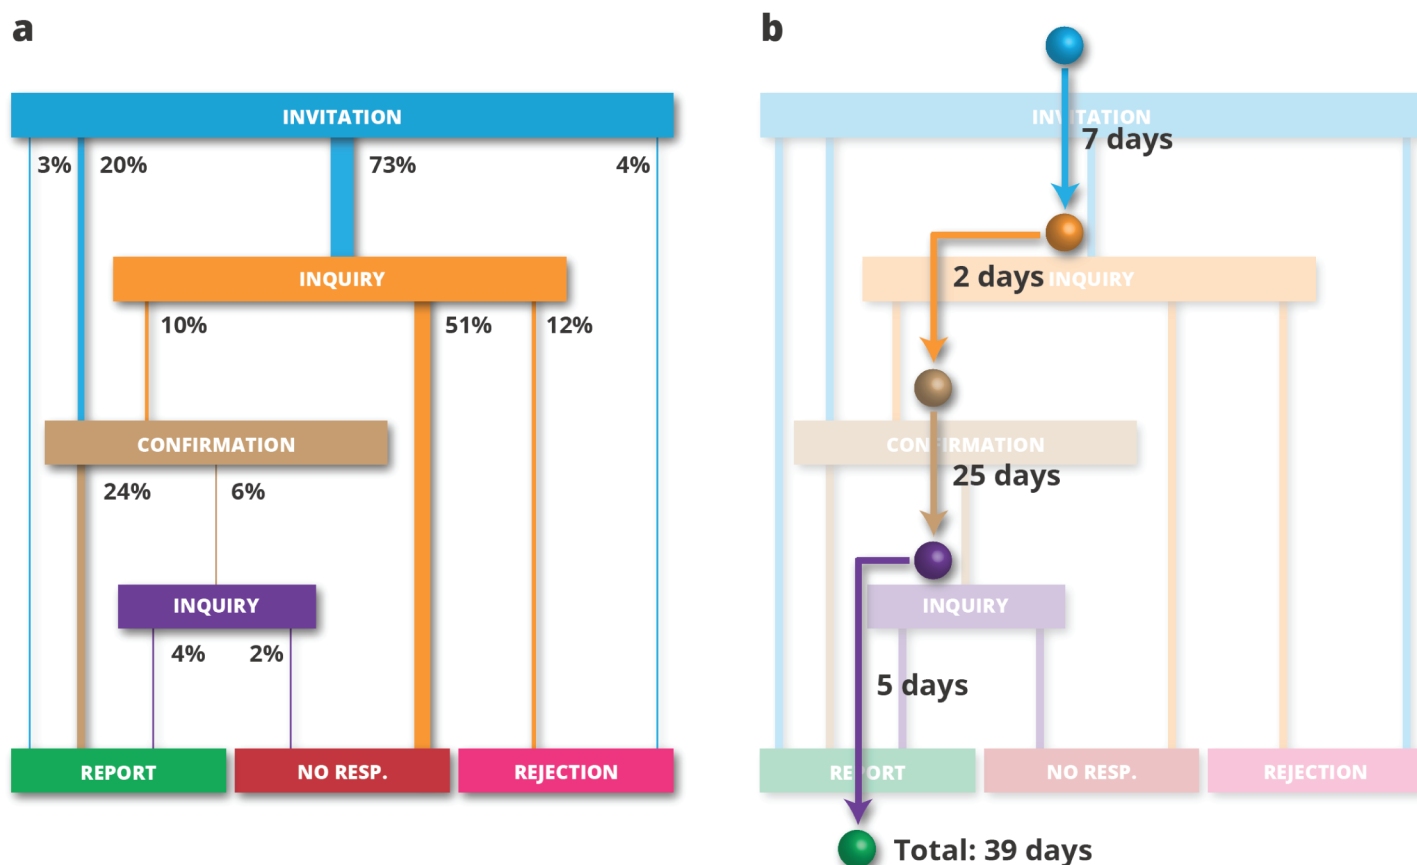


Fig 1. Decision diagram of the review process. a: Transition probabilities, showing the percent of review threads that passed through the edge. b: An example of a review thread, showing a path through the decision diagram, coupled with the duration of each phase.

<https://doi.org/10.1371/journal.pone.0184711.g001>

required number of reviews for a submitted manuscript). In the simulations, we assumed that two reviews were required, as is the case with the JSCS.

Materials and methods

Editorial workflows usually comprise many time-consuming tasks. First, editors need to find appropriate reviewers for each manuscript. Although this can be done with the help of various scientific databases, it is still a very laborious process. Second, they send invitations, handle the communication with reviewers and evaluate the received reports. Finally, they have to deal with situations where they do not receive enough reviews from the invited reviewers or the reviewers give conflicting recommendations.

Before we delve into a more in-depth explanation of how editorial workflows can be optimised, we would like to take a closer look at the artificial review threads. A realisation of the review process is initiated when the editor issues an invitation to a potential reviewer. Thus, to mirror the behaviour of real-life realisations, artificial review threads begin their life in the INVITATION phase (Fig 1). The next phase is determined by transition probabilities calculated from data. Unfortunately, the most probable route (73%) leads to the INQUIRY phase, which corresponds to the situation in which the invited reviewer does not answer. This prompts the editor, who waits for seven days after sending the initial invitation, to send another e-mail to the reviewer (the inquiry). When this second e-mail remains unanswered for ten days, then the editor assumes that the process is over with a negative outcome (without a review), which corresponds to the artificial review thread ending in the NO RESPONSE phase. However, if the reviewer does actually answer the inquiry (or the initial invitation) and agree to write the review, then the review thread enters the CONFIRMATION phase. In such a case, the editor waits for twenty-five days for the review and sends an inquiry if it is not received. To sum it up, each artificial review thread can end in one of three states – REPORT (which is the desired outcome), NO RESPONSE, and REJECT (when the invited reviewer answers but refuses to write the review). The probabilities of all outcomes, as well as probabilities of transitions between intermediate phases and durations of these transitions, were calculated using JSCS data.

To make the simulations as realistic as possible, all the review threads strictly followed the decision diagram presented in Fig 1a. After the initial invitations, the system handled all the interactions with the reviewers using real-world data. The remaining part of the review process was labelled ‘editorial strategies’. The editorial strategies are sets of rules that tell the editors how many new invitations they should issue at the beginning of the process and how many they should issue each time one review threads ends.

Informal discussions with editors revealed two editorial strategies, both of which are visualised in Fig 2. According to the first strategy (for simplicity called the strategy of Editor A), the editor should start a set number of review threads (called the batch size) at the beginning of the process and wait for all the threads to end before repeating this procedure if the required number of reviews is not received. In all of our simulations, the batch size is a parameter which, as in the strategy of Editor A, determines the upper limit of the number of simultaneously running review threads. The second strategy (for simplicity called the strategy of Editor B) begins in a similar way, with the editor sending invitations to a batch of reviewers. The difference between Editor A and Editor B is that the latter is always keeping the entire batch running. A finished review thread is immediately replaced with a new one, which ensures that the total number of running threads is equal to the batch size at all times.

In determining which strategy is better, one should consider that two opposite forces affect the efficiency of editorial strategies. On the one hand, both authors of articles and editors

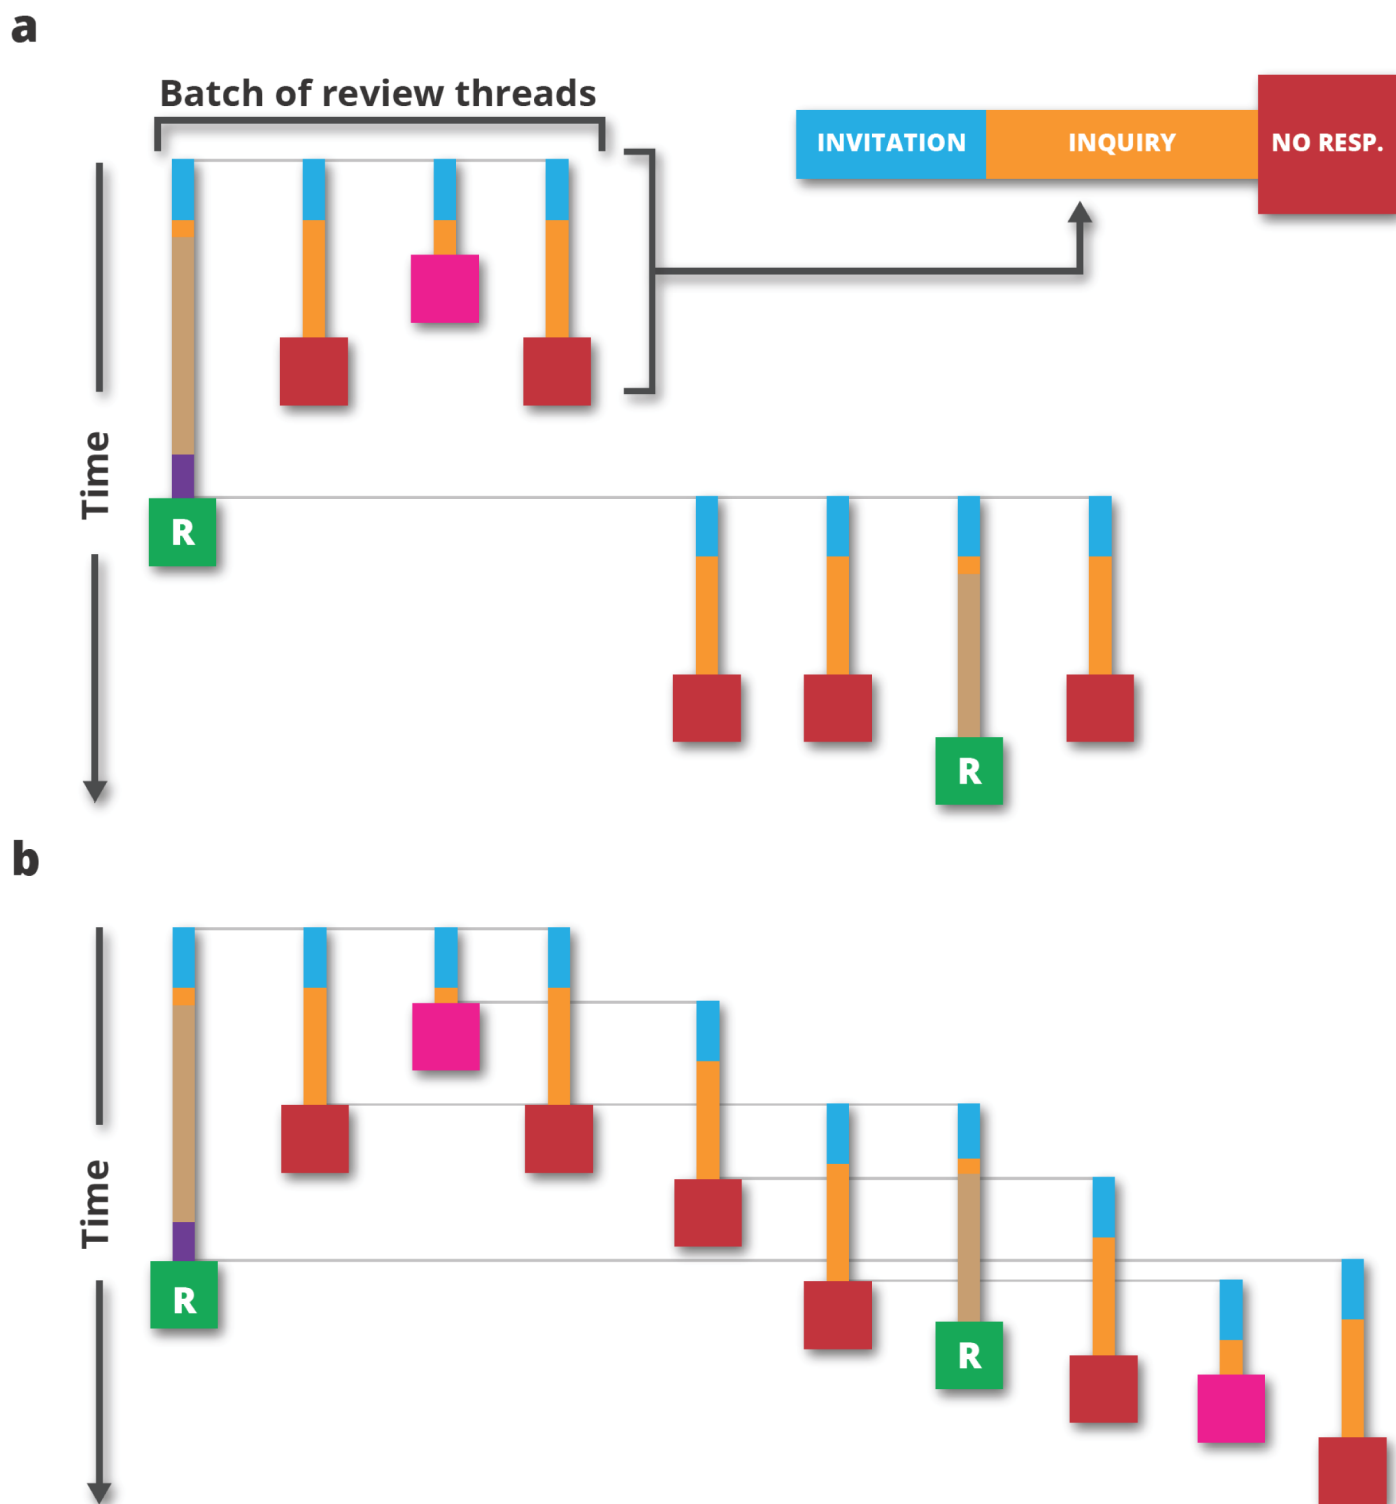


Fig 2. Two strategies used by editors. a: Strategy of Editor A, where the editor waits for all review threads to end before starting new ones. b: Strategy of Editor B, where the number of active review threads is always constant, and a new invitation is sent immediately after one of the active threads ends.

<https://doi.org/10.1371/journal.pone.0184711.g002>

would like to receive reviews as soon as possible. On the other hand, the editors prefer to minimise the number of invitations sent to reviewers. Studies show that authors become impatient if the review process takes too long and may withdraw the manuscript [15]. In theory, simply inviting many reviewers to assess the manuscript could shorten the review time. However, this is not a feasible solution for editors, as there are costs (e.g. searching time) associated with each reviewer. It follows that the best editorial strategy is the one that achieves the shortest review time using the smallest number of reviewers.

Taking the aforementioned factors into account, the strategies can be compared using the simple diagram presented in Fig 3a, which shows the review time as a function of the effective number of reviewers for both strategies. The effective number of reviewers is simply the average number of reviewers required by a particular strategy in order to acquire two reviews per article. In general, the effective number of reviewers will be higher than the batch size needed to achieve a given review time. As can be clearly seen, the strategy of Editor A is better than the strategy of Editor B. This may seem counter-intuitive at first, as the review time of Editor B for the same batch size is shorter than that of Editor A (Fig 3c). However, the effective number of reviewers that Editor B employs is larger (Fig 3d), which ultimately makes the strategy of Editor B less efficient.

The strategies of Editor A and Editor B are simple and intuitive. Many new editors would likely adopt one of these strategies as the first go-to solution in the peer review process. However, editors, especially in smaller journals, usually organise their workflows on their own without any point of reference. The development of an automated tool could provide more efficient strategies, thereby making the editor's job much easier, shortening review times and increasing the satisfaction of authors of manuscripts with the review process.

Cartesian Genetic Programming (CGP; implementations in various programming languages are readily available at www.cartesiangp.co.uk) [16–18] was originally conceived as a way of designing digital circuits. However, the actual scope of its applications is much larger, and it may be ideal for use in the optimisation peer review process. At the core of CGP is an evolutionary algorithm, which creates computer programs (functions) using concepts from the theory of evolution (i.e. natural selection and mutation) by optimising some fitness function. In the present case, these programs corresponded to editorial strategies, which told the editor how many new review threads should be started at any given time.

The editorial strategies in CGP can be encoded as a grid of nodes (Fig 4). Leftmost nodes represent inputs to the strategy – this is where the editor provides the integer numbers that describe the system (the required number of reviews, the number of received reviews, the number of running review threads, and the size of the batch) at a particular time. Input nodes are followed by intermediate nodes—a grid of n rows by m columns filled with nodes that correspond to basic mathematical functions. Each node has two inputs, which are used to calculate the value of the function realised by the node, and one output. It should be emphasised that inputs of nodes in a given column can be connected only to outputs of nodes in preceding columns. Forward (recurrent) connections are not allowed in order to avoid loops. In our simulations, we used only a single row and two thousand columns—that is, there were two thousand intermediate nodes between input and output arranged on a single, one-dimensional line. From our experience, this number of nodes is a good starting point as having too few nodes effectively prevents optimisation (CGP is not able to find a working solution to the problem in a reasonable time). All functions operated only on integers and were protected, which means that for inputs that are not allowed (e.g. division by zero) they would always return the number one. The rightmost node in the strategy is the output node, which by connecting to one of the preceding nodes tells the user of the strategy where to look for the final value.

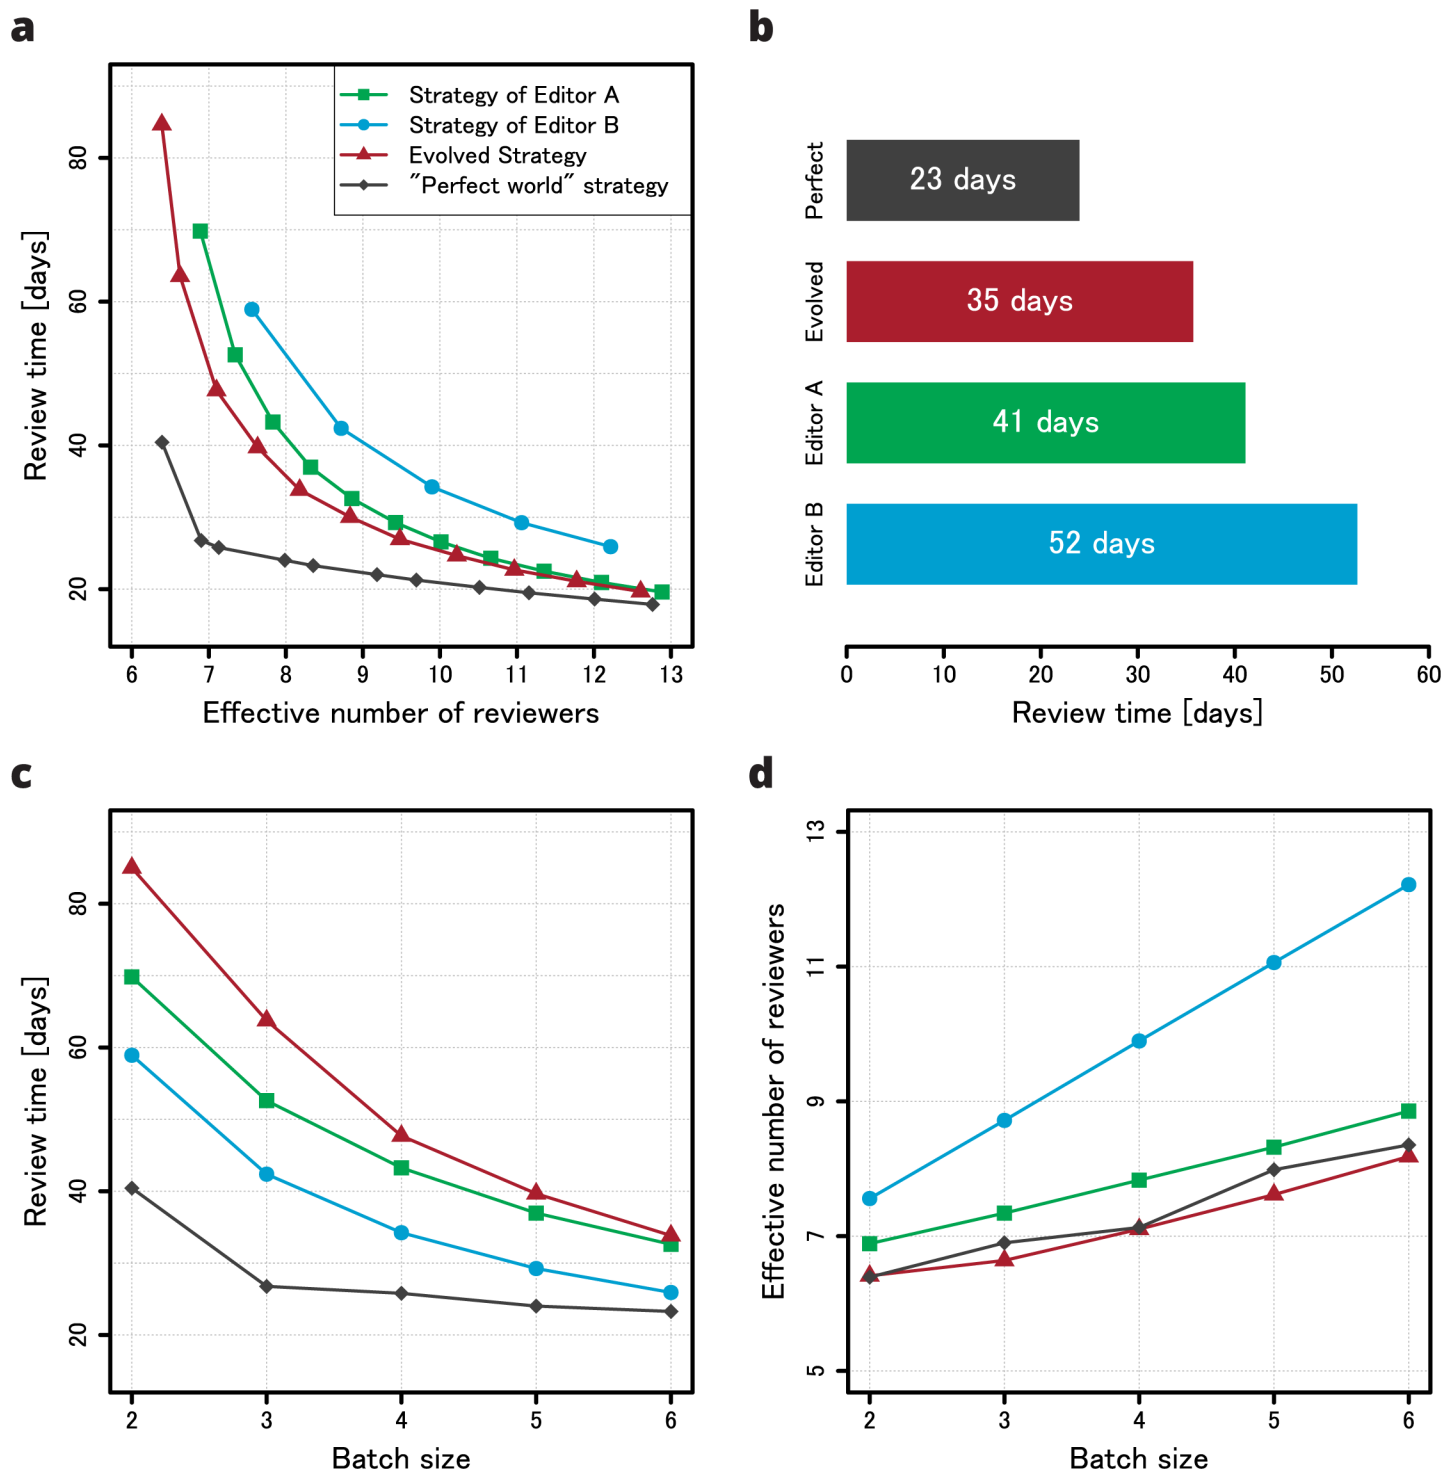


Fig 3. Efficiency of strategies. Strategies were evaluated by simulating the review process of a large number of articles (using the same procedure that was employed in the fitness function) and while the process itself is stochastic, deviations from these efficiency curves are negligible. (a) Review time as a function of the effective number of reviewers (the average number of reviewers needed to achieve a given review time). Each point on the plot corresponds to a single batch size—the first point on each curve represents the batch of two reviewers and for subsequent points the batch size increases by 1. (b) Comparison of the review time of various strategies for eight effective reviewers. (c) Review time as a function of the batch size. (d) The effective number of reviewers as a function of the batch size.

<https://doi.org/10.1371/journal.pone.0184711.g003>

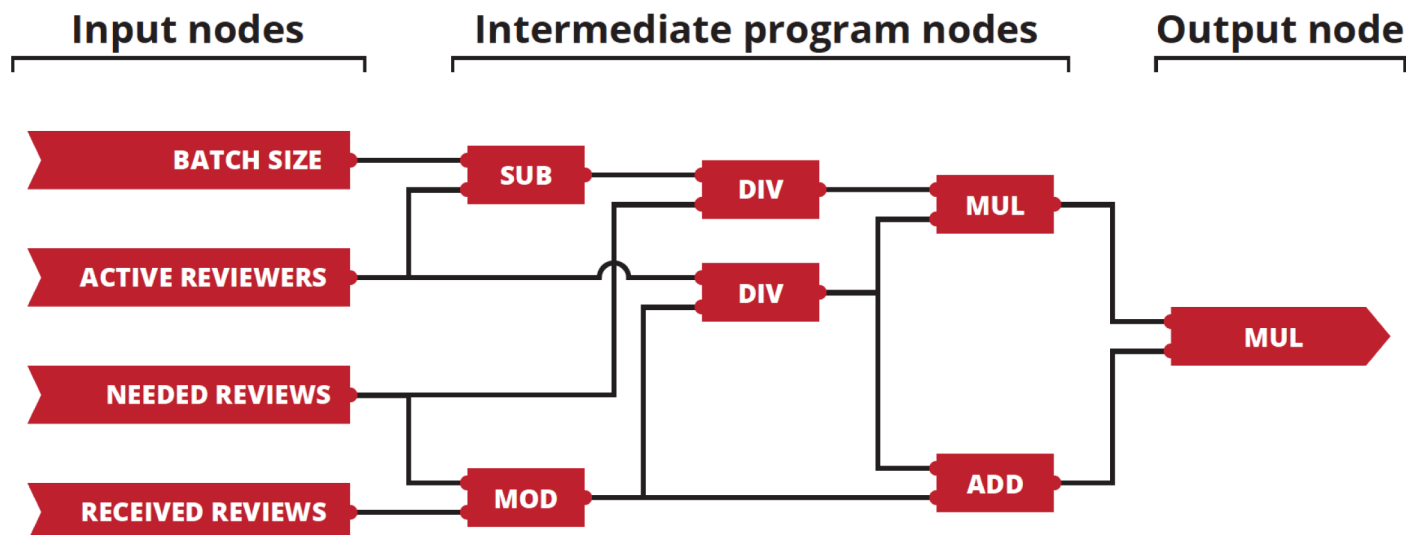


Fig 4. Evolved editorial strategy. BATCH SIZE, the size of the batch; ACTIVE REVIEWERS, the number of active review threads (reviewers for whom the review process has not ended yet); NEEDED REVIEWS, the number of reviews required per article; RECEIVED REVIEWS, the number of reviews received thus far; SUB, subtraction; DIV, division; MUL, multiplication; ADD, addition; MOD, modulo.

<https://doi.org/10.1371/journal.pone.0184711.g004>

Editorial strategies were evolved using a very simple algorithm, called the ‘4 + 1 evolutionary algorithm.’ [18]

1. At the beginning of the simulation five strategies are created by randomly assigning functions and links (connections) to nodes. Each of these random strategies is evaluated using a fitness function and the best strategy is chosen as the parent – the seed of the simulation.
2. At each step of the simulation the current parent is copied four times and the mutation operator is applied to these copies. The mutation operator traverses all inputs of all nodes in a strategy and changes them randomly with a given probability (called the mutation probability). It also traverses all function realised by nodes and changes them randomly with the same probability. It means that both the function and inputs can be changed by the mutation operator at the same time. Then, each of the mutated copies is evaluated using the fitness function. If one of the copies is at least as good (fit) as the parent, then it becomes the new parent. Otherwise, if all copies are worse strategies than the parent, then they are discarded and the parent carries over into the next iteration. We tested mutation probability values in the range $p_{mut} \in [0.01, 0.10]$. Since we were able to reach optimised solutions for all probabilities in this range, we decided to use $p_{mut} = 0.03$ in all simulations. It is worth mentioning at this point that while a graph representing an editorial strategy may be comprised of a large number of nodes, usually only a small fraction of these nodes (<10%) are actually used to calculate the output value. That is, most of the nodes in the graph are not parts of paths that connect inputs with the output. We call such nodes *inactive*. When the mutation operator is applied to an inactive node, the genotype may change but the phenotype remains the same. On the other hand, applying the mutation operator to an active node may change some of the previously inactive nodes into active as a result of changing the paths that lead from inputs to the output. This gradual build-up of changes in the inactive parts of graphs is one of the most important features of CGP.

3. Step 2 is repeated until a strategy characterised by a sufficiently low fitness function value is found (when compared to real strategies developed by human editors). There is no natural stopping condition for this algorithm – in theory it may try to optimise the strategy indefinitely.

The fitness function used to evaluate strategies is one of the most important components of the optimisation process. The numerical value it assigns to each program is a measure of effectiveness and can be used to compare two strategies. Our approach is based entirely on real-world data. With the help of artificial review threads, the fitness function we designed performs numerical simulations of the review process for batch sizes between two and twenty. The effective number of reviewers and the review time are calculated by averaging the results of simulations of the review process of ten thousand articles per one batch size. The pseudo-code for the fitness function and the function responsible for the simulation of the review process can be found in the Supporting Information ([S1 Algorithm](#)).

During simulations we assume that time is discrete and that one time step corresponds to one day. However, we also assume editors send new invitations (and consult editorial strategies) only when one of the running review threads finishes its execution (that is, when it enters the NO RESPONSE, REPORT or REJECTION phase). A valid strategy must obey certain constraint. First, it is assumed that at the beginning of the review process (i.e. the zeroth day) invitations are sent to the entire batch of reviewers. Second, the batch size must never be exceeded. Third, at least one invitation must be sent when there are no active (running) review threads. If at least one of these constraints is not met by the strategy being evaluated (the first condition is imposed automatically during simulations), then the fitness function returns a very large value (proportional to the number of batch sizes for which the strategy is not valid), unobtainable by a proper strategy. Otherwise, when the aforementioned conditions are fulfilled, the fitness function calculates the area under the strategy efficiency plot ([Fig 3](#); the integral is calculated over the interval [7, 15]) and returns that value to the CGP algorithm. Despite the apparent simplicity of this definition of the fitness function, it is entirely data driven and leads to improvements in real-world strategies. Lastly, it should be noted that we did not use any well defined stop condition. We would wait until the value of the fitness function indicated that the evolved strategy was better than the strategy of Editor A and we would terminate the optimisation process manually once it reached a stable solution (the value of the fitness function stopped changing).

Although the simulations performed to evaluate the values of the fitness function were based on ‘full’ review threads (i.e. with full information about phases of the process that constitute each thread), the results of these simulations can be recreated using data presented in the Supporting Information – [S1](#) and [S2](#) Tables. Both tables contain empirical distributions of the durations of the review process in JSCS. The duration, expressed in days, can be found in the ‘Days’ column. The ‘Freq.’ column contains the number of review threads in the JSCS sample that correspond to each duration. We should add at this point that the journal agreed to the publication of the data in this form. Also, the data we received was completely anonymised and we did not have access to any personal information.

Results and discussion

With the help of CGP equipped with the aforementioned fitness function, we searched for evolved strategies that would be better than strategies conceived by humans. The evolved strategy can be visualised in a couple of ways. The most natural representation is presented in [Fig 4](#), which depicts the optimal strategy we found as a graph. It is worth mentioning that this is not the full program that was initially evolved by CGP. In order to simplify the graph, we

employed CGP once again—this time to look for a new strategy which would be functionally equivalent to the evolved strategy but implemented using a smaller number of active nodes (that is, a strategy with the same phenotype but simpler genotype). Also, to make it more readable, we removed all inactive nodes from Fig 4 – that is nodes that are not directly or indirectly connected to the output and thus do not take part in the calculation of the value returned by the strategy. For the sake of comparison, S1 Fig from Supporting Information shows an extract of the full strategy graph with inactive nodes intact.

Although the program in Fig 4 looks complicated, the strategy it represents is not. For example, if the batch size is four, then the editor should initially invite two reviewers. The editor should then wait until only one active thread remains. If only one review has been received up to that point, the editor should continue to wait and begin two new review threads only if the last remaining thread does not yield a review. Otherwise, the editor should send invitations to two new reviewers immediately, without waiting for that last thread.

Since editorial strategies are essentially deterministic functions that assign integer numbers to states of the system (to each possible input), one can recreate the mathematical formula that corresponds to a strategy by following nodes from output to inputs and writing down the nested functions encountered in nodes. It is important to keep in mind that all functions realised by nodes are protected and normal rules of simplifying mathematical expressions often do not apply. The function equivalent to the evolved strategy can be expressed in the following way:

$$N_{NRT} = \left[N_R \bmod N_{RR} + \frac{N_{EQ} \bmod N_{RR}}{N_{ART}} \right] \times \frac{N_B - N_{ART}}{N_{RQ}} \times \frac{N_{RQ} \bmod N_{RR}}{N_{ART}} \quad (1)$$

where N_{NRT} is the number of new review threads that should be started, N_{RQ} is the required number of reviews, N_{RR} is the number of received reviews, N_{ART} is the number of active review threads, and N_B is the batch size.

The evolved strategy performed well when compared to that of Editor B and Editor A (Fig 3). For eight effective reviewers, the review time was 17 days shorter using the evolved strategy as compared to the strategy of Editor B and six days shorter as compared to the strategy of Editor A. Considering that the review time of the JSCS is relatively short, editors of other journals can potentially shorten their peer review times even more. We do not claim that the strategy presented in Fig 4 is universal. It works in the ecosystem of the JSCS but may be less efficient in a different context. However, the method used is universal. As long as editors are able to perform a statistical analysis of the peer review process of their journal and use the results to create simulations of the process with artificial review threads, they can employ CGP to improve their workflow. We believe the work presented herein is an interesting application of an evolutionary technique to what, in part, is a sociological problem and that such a data-driven approach is the key to realistic models of peer review that can provide more insight into the process and make it more efficient.

It should be noted that our method of optimisation has no bearing on the quality of peer review. Editorial strategies determine when and how many invitations must be sent to reviewers but they do not interfere with the actual process of reviewing the manuscript by said reviewers. It means that it is still up to the editor to uphold the standard of the journal and the quality of the process.

We would like to end this article by mentioning one more discovery. Editorial workflows are incredibly important, but what can scientists do to speed up the review process? Surely, not

everything is in the hands of editors. As it turns out, scientists can do quite a lot. Fig 3 shows what we call ‘the perfect world strategy.’ In order to create the efficiency curve for this strategy, the evolved strategy was evaluated under a very simple, yet unrealistic assumption, as shown by data: All potential reviewers answer their mails and always immediately respond to an initial invitation by either accepting or refusing to write the review. That is, the NO RESPONSE and REJECTION outcomes of the review process are merged into one without changing the probability of accepting to write a review. As can be clearly seen, this one assumption makes a striking difference. In the end, it seems that a shorter review time is easily obtainable. It is, quite literally, one mouse click away.

Supporting information

S1 Table. Distribution of the length of the review process (single review thread; process ended with a review) in JSCS.

(PDF)

S2 Table. Distribution of the length of the review process (single review thread; process ended without a review) in JSCS.

(PDF)

S1 Algorithm. Algorithms for the fitness function and the simulation of the review process.

(PDF)

S1 Fig. Editorial strategy with both active and inactive nodes visible.

(PDF)

Acknowledgments

This publication is supported by the COST Action TD1306 “New frontiers of peer review”.

Author Contributions

Conceptualization: Piotr Fronczak, Agata Fronczak.

Data curation: Olgica Nedic.

Formal analysis: Maciej J. Mrowinski.

Investigation: Maciej J. Mrowinski.

Methodology: Maciej J. Mrowinski, Piotr Fronczak, Agata Fronczak.

Project administration: Piotr Fronczak, Agata Fronczak.

Software: Maciej J. Mrowinski.

Validation: Maciej J. Mrowinski, Piotr Fronczak, Agata Fronczak.

Visualization: Maciej J. Mrowinski.

Writing – original draft: Maciej J. Mrowinski, Piotr Fronczak, Agata Fronczak, Marcel Ausloos, Olgica Nedic.

Writing – review & editing: Maciej J. Mrowinski, Piotr Fronczak, Agata Fronczak, Marcel Ausloos, Olgica Nedic.

References

1. Baker D. The peer review process in science education journals. *Research in science Education*. 2002; 32(2):171–180. <https://doi.org/10.1023/A:1016073927593>
2. Ware M, Monkman M. Peer review in scholarly journals: Perspective of the scholarly community—an international study. Mark Ware Consulting; 2008. Available from: <http://publishingresearchconsortium.com/>
3. Ware M, Mabe M. The STM report: An overview of scientific and scholarly journal publishing (4th edition). International Association of Scientific, Technical, and Medical Publishers; 2015.
4. Mulligan A, Hall L, Raphael E. Peer review in a changing world: An international study measuring the attitudes of researchers. *Journal of the American Society for Information Science and Technology*. 2013; 64(1):132–161. <https://doi.org/10.1002/asi.22798>
5. Nicholas D, Watkinson A, Jamali H, Herman E, Tenopir C, Volentine R, et al. Peer Review: still king in the digital age. *Learned Publishing*. 2015; 28(1):15–21. <https://doi.org/10.1087/20150104>
6. Bourke-Waite A. Innovations in scholarly peer review at Nature Publishing Group and Palgrave Macmillan. *Insights*. 2015; 28(2):93–99.
7. Powell K. Does it take too long to publish research? *Nature*. 2016; 530:148–151. <https://doi.org/10.1038/530148a> PMID: 26863966
8. Squazzoni F, Takács K. Social simulation that 'peers into peer review'. *Journal of Artificial Societies and Social Simulation*. 2011; 14(4):3. <https://doi.org/10.18564/jasss.1821>
9. Cawley V. An analysis of the ethics of peer review and other traditional academic publishing practices. *International Journal of Social Science and Humanity*. 2011; 1(3):205–213. <https://doi.org/10.7763/IJSSH.2011.V1.36>
10. Wager E. Ethics: What is it for. *Nature: Web Debate—Peer-review*. 2006;
11. Resnik DB, Gutierrez-Ford C, Peddada S. Perceptions of ethical problems with scientific journal peer review: an exploratory study. *Science and engineering ethics*. 2008; 14(3):305–310. <https://doi.org/10.1007/s11948-008-9059-4> PMID: 18311477
12. Ausloos M, Nedic O, Fronczak A, Fronczak P. Quantifying the quality of peer reviewers through Zipf's law. *Scientometrics*. 2016; 106(1):347–368. <https://doi.org/10.1007/s11192-015-1704-5>
13. Ausloos M, Nedic O, Dekanski A. Day of the week effect in paper submission/acceptance/rejection to/in/by peer review journals. *Physica A*. 2016; 456:197–203. <https://doi.org/10.1016/j.physa.2016.03.032>
14. Mrowinski MJ, Fronczak A, Fronczak P, Nedic O, Ausloos M. Review time in peer review: quantitative analysis and modelling of editorial workflows. *Scientometrics*. 2016; 107(1):271–286. <https://doi.org/10.1007/s11192-016-1871-z> PMID: 27073291
15. Nguyen VM, Haddaway NR, Gutowsky LFG, Wilson ADM, Gallagher AJ, Donaldson MR, et al. How Long Is Too Long in Contemporary Peer Review? Perspectives from Authors Publishing in Conservation Biology Journals. *PLoS ONE*. 2015; 10(8):1–20. <https://doi.org/10.1371/journal.pone.0132557>
16. Miller JF, Job D, Vassilev VK. Principles in the Evolutionary Design of Digital Circuits—Part I. Genetic Programming and Evolvable Machines. 2000; 1(1):7–35. <https://doi.org/10.1023/A:1010016313373>
17. Miller JF, Job D, Vassilev VK. Principles in the Evolutionary Design of Digital Circuits—Part II. Genetic Programming and Evolvable Machines. 2000; 1(3):259–288. <https://doi.org/10.1023/A:1010066330916>
18. Miller JF. Cartesian Genetic Programming. Springer; 2011.