

Conversion Rate Optimization through Evolutionary Computation

Risto Miikkulainen^{1,2}, Neil Iscoe¹, Aaron Shagrin¹, Ron Cordell¹, Sam Nazari¹, Cory Schoolland¹,
Myles Brundage¹, Jonathan Epstein¹, Randy Dean¹, Gurmeet Lamba¹

¹Sentient Technologies, Inc.

²The University of Texas at Austin

ABSTRACT

Conversion optimization means designing a web interface so that as many users as possible take a desired action on it, such as register or purchase. Such design is usually done by hand, testing one change at a time through A/B testing, or a limited number of combinations through multivariate testing, making it possible to evaluate only a small fraction of designs in a vast design space. This paper describes Sentient Ascend, an automatic conversion optimization system that uses evolutionary optimization to create effective web interface designs. Ascend makes it possible to discover and utilize interactions between the design elements that are difficult to identify otherwise. Moreover, evaluation of design candidates is done in parallel online, i.e. with a large number of real users interacting with the system. A case study on an existing media site shows that significant improvements (i.e. over 43%) are possible beyond human design. Ascend can therefore be seen as an approach to massively multivariate conversion optimization, based on a massively parallel interactive evolution.

CCS CONCEPTS

•Computing methodologies → Genetic algorithms; Online learning settings; Sequential decision making; •Applied computing → E-commerce infrastructure;

KEYWORDS

Conversion optimization, e-commerce, interactive evolution, online evolution, design

ACM Reference format:

Risto Miikkulainen^{1,2}, Neil Iscoe¹, Aaron Shagrin¹, Ron Cordell¹, Sam Nazari¹, Cory Schoolland¹, Myles Brundage¹, Jonathan Epstein¹, Randy Dean¹, Gurmeet Lamba¹. 2017. Conversion Rate Optimization through Evolutionary Computation. In *Proceedings of the Genetic and Evolutionary Computation Conference 2017, Berlin, Germany, July 15–19, 2017 (GECCO '17)*, 7 pages.
DOI: 10.475/123.4

1 INTRODUCTION

In e-commerce, designing web interfaces (i.e. web pages and interactions) that convert as many users as possible from casual browsers to paying customers is an important goal [1, 11]. While there

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

GECCO '17, Berlin, Germany

© 2017 Copyright held by the owner/author(s). 123-4567-24-567/08/06...\$15.00

DOI: 10.475/123.4

are some well-known design principles, including simplicity and consistency, there are often also unexpected interactions between elements of the page that determine how well it converts. The same element, such as a headline, image, or testimonial, may work well in one context but not in others—it is often hard to predict the result, and even harder to decide how to improve a given page.

An entire subfield of information technology has emerged in this area, called conversion rate optimization, or conversion science. The standard method is A/B testing, i.e. designing two different versions of the same page, showing them to different users, and collecting statistics on how well they each convert [8]. This process allows incorporating human knowledge about the domain and conversion optimization into the design, and then testing their effect. After observing the results, new designs can be compared and gradually improved. The A/B testing process is difficult and time-consuming: Only a very small fraction of page designs can be tested in this way, and subtle interactions in the design are likely to go unnoticed and unutilized. An alternative to A/B is multivariate testing, where all value combinations of a few elements are tested at once. While this process captures interactions between these elements, only a very small number of elements is usually included (e.g. 2-3); the rest of the design space remains unexplored.

This paper describes a new technology for conversion optimization based on evolutionary computation. This technology is implemented in Ascend, a conversion optimization product by Sentient Technologies, deployed in numerous e-commerce websites of paying customers since September 2016 [13]. Ascend uses a customer-designed search space as a starting point. It consists of a list of elements on the web page that can be changed, and their possible alternative values, such as a header text, font, and color, background image, testimonial text, and content order. Ascend then automatically generates web-page candidates to be tested, and improves those candidates through evolutionary optimization.

Because e-commerce sites often have high volume of traffic, fitness evaluations can be done live with a large number of real users in parallel. The evolutionary process in Ascend can thus be seen as a massively parallel version of interactive evolution, making it possible to optimize web designs in a few weeks. From the application point of view, Ascend is a novel method for massively multivariate optimization of web-page designs. Depending on the application, improvements of 20-200% over human design have been observed through this approach [13].

This paper describes the technology underlying Ascend, presents an example use case, and outlines future opportunities for evolutionary computation in optimizing e-commerce.

2 BACKGROUND

With the explosive growth of e-commerce in recent years, entirely new areas of study have emerged. One of the main ones is conversion rate optimization, i.e. the study of how web interfaces should be designed so that they are as effective as possible in converting users from casual browsers to actual customers. Conversion means taking a desired action on the web interface such as making a purchase, registering for a marketing list, or clicking on other desired link in an email, website, or desktop, mobile, or social media application [1, 11]. Conversions are usually measured in number of clicks, but also in metrics such as resulting revenue or time spent on the site and rate of return to the site.

Conversions are currently optimized in a labor-intensive manual process that requires significant expertise. The web design expert or marketer first creates designs that s/he believes to be effective. These designs are then tested in an A/B testing process, by directing user traffic to them, and measuring how well they convert. If the conversion rates are statistically significantly different, the better design is adopted. This design can then be improved further, using domain expertise to change it, in another few rounds of creation and testing.

Conversion optimization is a fast-emerging component of e-commerce. In 2016, companies spent over \$72 billion to drive customers to their websites [5]. Much of that investment does not result in sales: conversion rates are typically 2-4% (i.e. 2-4% of the users that come to the site convert within 30 days). In 2014, only 18% of the top 10,000 e-commerce sites did any conversion optimization; in January 2017, 30% of them did so [4]. The growth is largely due to available conversion optimization tools, such as Optimizely, Visual Website Optimizer, Mixpanel, and Adobe Target [4]. These tools make it possible to configure the designs easily, allocate users to them, record the results, and measure significance.

This process has several limitations. First, while the tools make the task of designing effective web interfaces easier, the design is still done by human experts. The tools thus provide support for confirming the experts' ideas, not helping them explore and discover novel designs. Second, since each step in the process requires statistical significance, only a few designs can be tested. Third, each improvement step amounts to one step in hillclimbing; such a process can get stuck in local maxima. Fourth, the process is aimed at reducing false positives and therefore increases false negatives, i.e. designs with good ideas may be overlooked. Fifth, while the tools provide support for multivariate testing, in practice only a few combinations can be tested (e.g. five possible values for two elements, or three possible values for three elements). As a result, it is difficult to discover and utilize interactions between design elements.

Evolutionary optimization is well suited to address these limitations. Evolution is an efficient method for exploration; only weak statistical evidence is needed for progress; its stochastic nature avoids getting stuck in local maxima; good ideas will gradually become more prevalent. Most importantly, evolution searches for effective interactions. For instance, Ascend may find that the button needs to be green, but *only* when it is transparent, *and* the header is in small font, *and* the header text is aligned. Such interactions are very difficult to find using A/B testing, requiring human

insight into the results. Evolution makes this discovery process automatic. With Ascend, it is thus possible to optimize conversions better and at a larger scale than before.

Technically, Ascend is related to approaches to interactive evolution [12, 15] and crowdsourcing [2, 9] in that evaluations of candidates are done online by human users. The usual interactive evolution paradigm, however, employs a relatively small number of human evaluators, and their task is to select good candidates or evaluate the fitness of a pool of candidates explicitly. In contrast in Ascend, a massive number of human users are interacting with the candidates, and fitness is derived from their actions (i.e. convert or not) implicitly.

3 THE ASCEND METHOD

Ascend consists of defining the space of possible web interfaces, initializing the population with a good coverage of that space, allocating traffic to candidates intelligently so that bad designs can be eliminated early, and testing candidates online in parallel. Each of these steps is described in more detail in this section.

3.1 Defining the Search Space

The starting point for Ascend is a search space defined by the web designer. Ascend can be configured to optimize a design of a single web-page, or a funnel consisting of multiple pages such as the landing page, selections, and a shopping cart. For each such space, the designer specifies the elements on that page and values that they can take. For instance in the landing page example of Figure 1, logo size, header image, button color, content order are such elements, and they can each take on 2-4 values.

Ascend searches for good designs in the space of possible combinations of these values. This space is combinatorial, and can be very large, e.g. 1.1M in this example. Interestingly, it is exactly this combinatorial nature that makes web-page optimization a good application for evolution: Even though human designers have insight into what values to use, their combinations are difficult to predict, and need to be discovered by search process such as evolution.

3.2 Initializing Evolution

A typical setup is that there is already a current design for the web interface, and the goal for Ascend is to improve over its performance. That is, the current design of the web interface is designated as the Control, and improvement is measured compared to that particular design.

Because fitness is evaluated with real users, exploration incurs real cost to the customer. It is therefore important that the candidates perform reasonably well throughout evolution, and especially in the beginning.

If the initial population is generated randomly, many web interfaces would perform poorly. Instead, the initial population is created using the Control as a starting point: The candidates are created by changing the value of one element systematically. In a small search space, the initial population thus consists of all candidates with one difference from the control; in a large search space, the population is a sample of the set of such candidates. With such an initialization, most of the candidates perform similarly to the

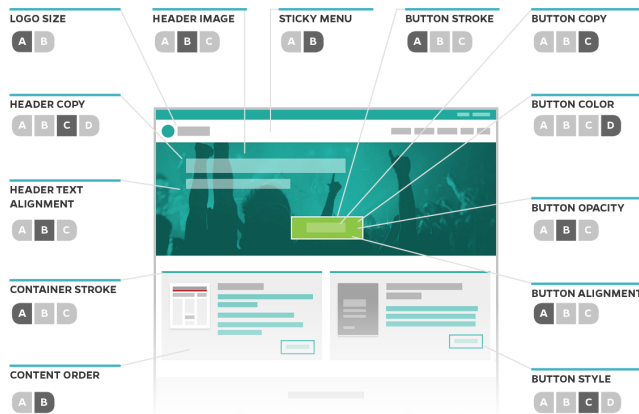


Figure 1: Elements and Values of an Example Web Page Design. In this example, 13 elements each have 2-4 possible values, resulting in 1.1M combinations.

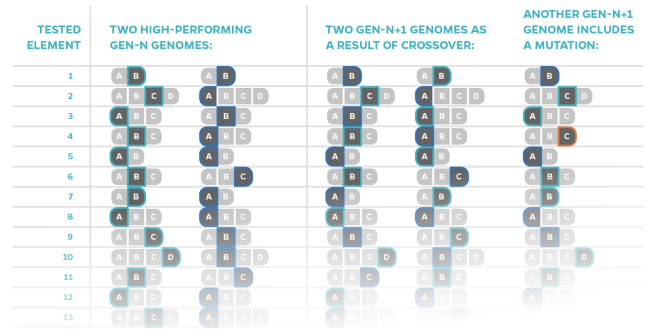


Figure 2: Genetic Encoding and Operations on Web Interface Candidates. The pages are represented as concatenations of their element values with one-hot encoding. Crossover and mutation operate on these vectors as usual, creating new combinations of values.

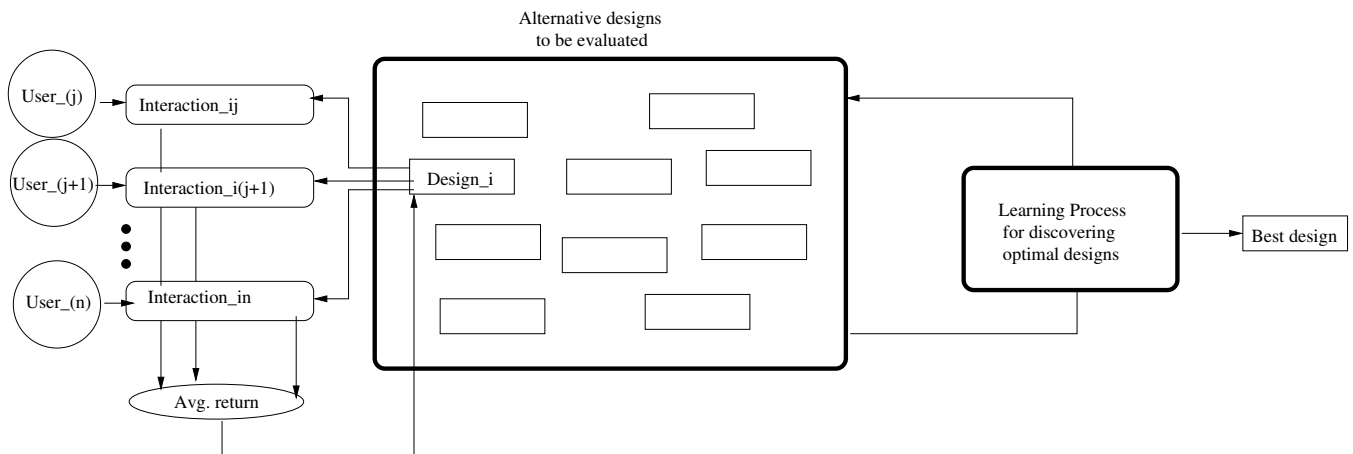


Figure 3: Overall Architecture of the Online Evolution System. The outcome of each interaction (i.e. whether the user converted or not) constitutes one evaluation of a design. Many such evaluations ij are run in parallel with different users j and averaged to estimate how good the design i is. After all designs have been evaluated, the adaptation process discards bad designs and generates more variations of the best designs. This process of generation, testing, and selection is repeated until a sufficiently good design has been found or the time allocated for the process has been spent. The best design found so far is output as the result of the learning process. The system thus discovers good designs for web interfaces through live online testing.

control. The candidates also cover the search dimensions well, thus forming a good starting point for evolution.

3.3 Evolutionary Process

Each page is represented as a genome, as shown for two example pages in Figure 2 (left side). The usual genetic operations such as crossover (re-combination of the elements in the two genomes; middle) and mutation (randomly changing one element in the offspring; right side) are then performed to create new candidates. In the current implementation, fitness-proportionate selection is used to generate offspring candidates from the current population. From the current population of n candidates, another n new candidates are generated in this way.

Because evaluations are expensive, consuming traffic for which most customers have to pay, it is useful to minimize them during evolution. Each page needs to be tested only to the extent that it is possible to decide whether it is promising, i.e. whether it should serve as a parent in the next generation, or should be discarded. A process similar to age-layering [7, 14] is therefore used to allocate fitness evaluations. At each generation, each new candidate and each old candidate is evaluated with a small number (a maturity age) of user interactions, such as 2000. The top n candidates are retained, and the bottom n discarded. In this manner, bad candidates are eliminated quickly. Good candidates receive progressively more evaluations, and the confidence in their fitness estimate increases.

In this process, Ascend learns which combinations of elements are effective, and gradually focuses the search around the most

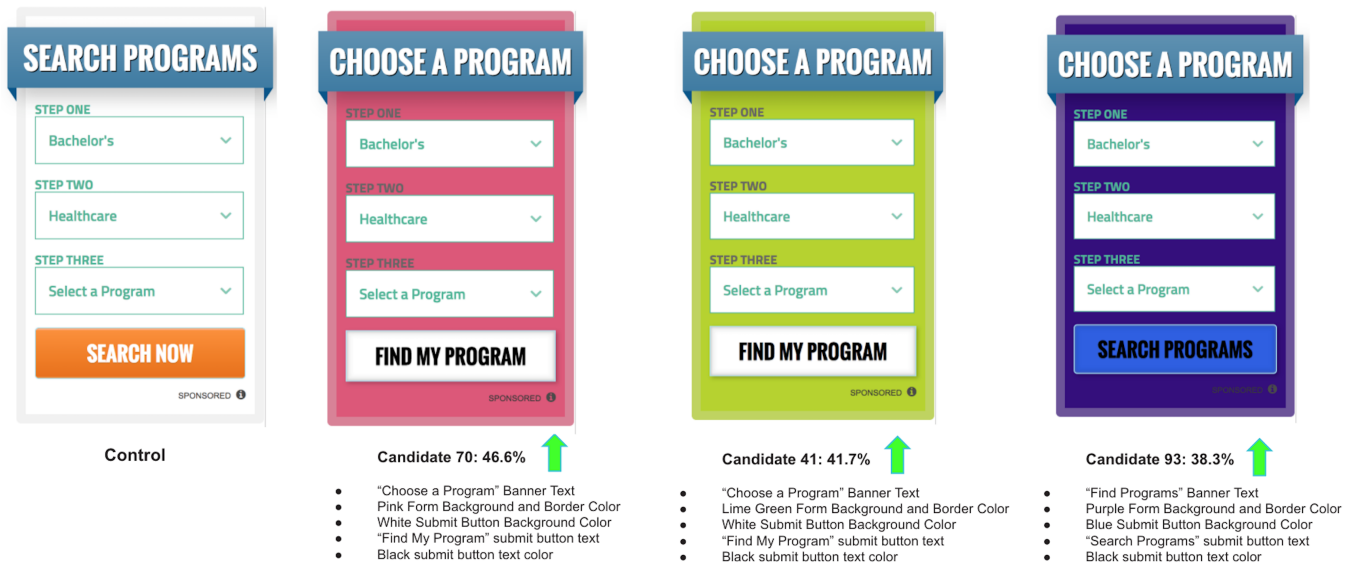


Figure 4: The control design and three best evolved designs. After 60 days of evolution with 599,008 user interactions, a design for the search widget was found that converted 46.6% better than the control (5.61% vs. 8.22%), as well as other good designs. Much of the improvement was based on discovering a combination of colors that draws attention to the widget and makes the call to action clear.

promising designs. It is thus sufficient to test only a tiny fraction of the search space to find the best ones, i.e. thousands of pages instead of millions or billions.

3.4 Online Evolution

While in simple cases (where the space of possible designs is small) such optimization can potentially be carried out by simpler mechanisms such as systematic search, hill-climbing, or reinforcement learning, the population-based approach is particularly effective because the evaluations can be done in parallel. The entire population can be tested at once, as different users interact with the site simultaneously. It is also unnecessary to test each design to statistical significance; only weak statistical evidence is sufficient to proceed in the search. In this process, thousands of page designs can be tested in a short time, which is impossible through A/B or multivariate testing.

Figure 3 shows the overall architecture of the system. A population of alternative designs (center) are adapted (right) based on evaluations with actual users (left). The population of designs (center) are evaluated with many users in parallel (left). The evolutionary process (right) generates new designs and outputs the best design in the end. The system also keeps track of which design has been show to which user, so that they get to see the same design if they return within a certain time limit (e.g. the same day).

4 CASE STUDY

As an example of how Ascend works, let us consider a case study on optimizing the web interface for a media site that connects users to online education programs. This experiment was run in September through November 2016 on the desktop traffic of the site.

The initial design for this page is shown in the left side of Figure 4. It had been hand designed using standard tools such as Optimizely. Its conversion rate during the time of the experiment was found to be 5.61%, which is typical of such web interfaces. Based on this page, the web designers came up with nine elements, with two to nine values each, resulting in 381,024 potential combinations (Figure 5). While much larger search spaces are possible, this example represents a mid-size space common with many current sites.

The initial population of 37 candidates was formed by systematically replacing each of the values in the control page with one of the alternative values, as described in section 3.2. Evolution was then run for 60 days, or four generations, altogether testing 111 candidates with 599,008 user interactions total. The estimated conversion rates of the candidates over this time are shown in Figure 6. The conversion rates of the top 20 candidates are shown in Figure 7. These figures show that evolution was successful in discovering significantly better candidates than control.

As an independent verification, the three top candidates in Figure 4 were then subjected to an A/B test using Optimizely. In about 6500 user interactions, the best candidate was confirmed to increase the conversion rate by 43.5% with greater than 99% significance (and the other two by 37.1% and 28.2%)—which is an excellent result given that the control was a candidate that was already hand-optimized using state-of-the-art tools.

Unlike Control, the top candidates utilize bright background colors to draw attention to the widget. There is an important interaction between the background and the blue banner (whose color was fixed)—in the best two designs (in the middle) the background is distinct from the banner but not competing with it. Moreover,

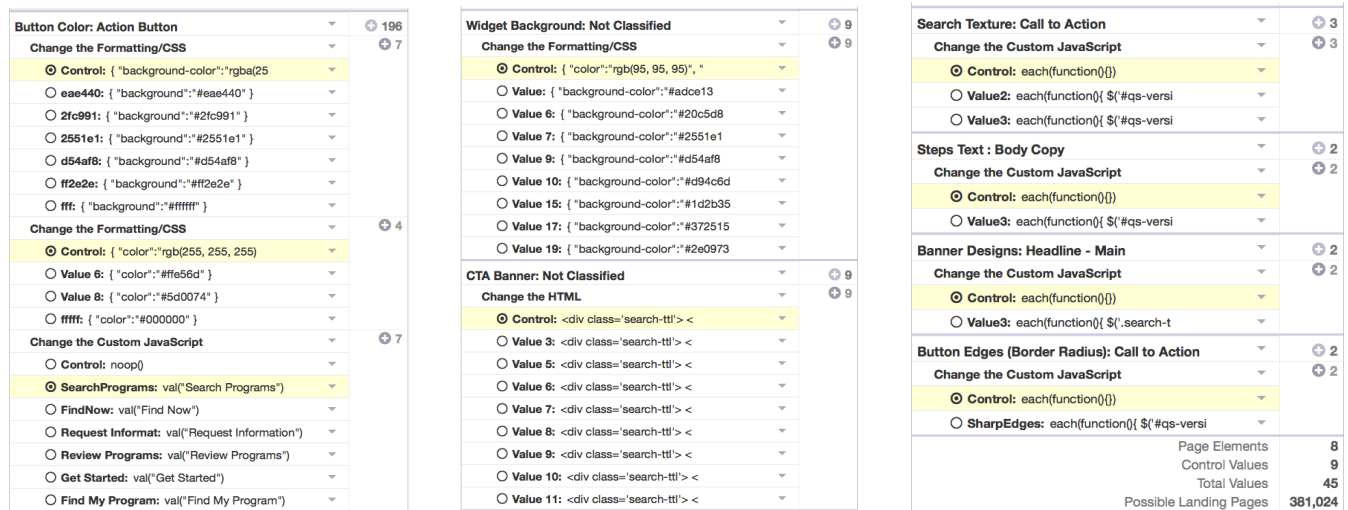


Figure 5: A screenshot of the user interface for designing Ascend experiments, showing the elements and values in the learn-howtobecome.com case study. Nine elements with two to nine different values each result in 381,024 potential web page designs; the first value in each element is designated as the control. This is a mid-size problem typical of current web interface designs.

given the colored background, a white button with black text provided the most clear call for action. It is difficult to recognize such interactions ahead of time, yet evolution discovered them early on, and many of the later candidates built on them. Other factors such as an active call to action (i.e. “Get Started” and “Find my Program” rather than “Request Info”) amplified it further. At the time evolution was turned off, better designs were still being discovered, suggesting that a more prolonged evolution and a larger search space (e.g. including banner color and other choices) could have improved the results further.

5 FUTURE WORK

Ascend has been applied to numerous web interfaces, and it has consistently improved conversion rates by 20-200% compared to hand designed controls [13]. The main limitation is often the human element: web designers, who are used to A/B and multivariate testing, often try to minimize the search space, i.e. the number of elements and values, as much as possible, thereby not giving evolution much space to explore and discover powerful solutions. Often the evolution discovers significant improvement in a couple of generations, and the designers are eager to adopt them right away, instead of letting evolution optimize the designs fully. Population-based optimization requires different thinking; as designers become more comfortable with it, we believe they will let evolution take its course, reaching more refined results.

Currently Ascend delivers one best design, or a small number of good ones, in the end as the result, again in keeping with the A/B testing tradition. In many cases there are seasonal variations and other long-term changing trends, making the performance of good designs gradually decay. It is possible to counter this problem by running the optimization again every few months. However, a new paradigm of “always-on” would be more appropriate: Evolutionary

optimization can be run continuously at a low volume, keeping up with changing trends (i.e. through dynamic evolutionary optimization; [3]). New designs can then be adopted periodically when their performance exceeds old designs significantly.

Furthermore, currently Ascend optimizes a single design to be used with all future users of a mobile or desktop site. An interesting extension would be to take user segmentation [16] into account, and evolve different pages for different kinds of users. Moreover, such a mapping from user characterizations to page designs can be automatic: A mapping system such as a neural network can take user variables such as location, time, device, any past history with the site as inputs, and generate the vector of elements and their values as outputs. Neuroevolution [6, 10] can discover optimal such mappings, in effect evolve to discover a dynamic, continuous segmentation of the user space. Users will be shown designs that are likely to convert well based on experience with other users with similar characteristics, continuously and automatically. It will be possible to analyze such evolved neural networks and discover what variables are most predictive, characterize the main user segments, and thereby develop an in-depth understanding of the opportunity.

Finally, the Ascend approach is not limited to optimizing conversions. Any outcome that can be measured, such as revenue or user retention, can be optimized. The approach can also be used in a different role, such as optimizing the amount of resources spent on attracting users, such as ad placement and selection, adword bidding, and email marketing. The approach can be seen as a fundamental step in bringing machine optimization into e-commerce, and demonstrating the value of evolutionary computation in real-world problems.




Figure 6: Estimated Conversion Rates through the 60-day Online Evolution Run. Days are in the x -axis and the conversion rate on the y axis. The dark blue dots (on top) indicate the current best candidate, the light blue dots (in the middle) an average of all currently active candidates, and the orange dots (at the bottom) the estimated performance of the control design. The shaded areas display the 95% confidence intervals (from the binomial distribution with the observed mean). The dark blue peaks indicate the start of each new generation. Such peaks emerge because during the first few days, the new candidates have been evaluated only a small number of times, and some of them have very high estimated rates through random chance. Eventually they will be evaluated in a maturity age of 2000 user interactions, and the estimates become lower and the confidence intervals narrower. The elite candidates are tested across several generations as described in section 3.3, resulting in very narrow intervals towards the end. Estimated conversion rates of the best candidates in later generations are significantly higher than control, suggesting that evolution is effective in discovering better candidates. Interestingly, the active population average is also higher than control, indicating that the experiment did not incur any cost in performance.

6 CONCLUSION

Sentient Ascend demonstrates how interactive evolution can be scaled up by testing a large number of candidates in parallel on real users. It includes technology for keeping the cost of exploration reasonable, and for minimizing the number of evaluations needed. From the application point of view, Ascend is the first automated system for massively multivariate conversion optimization—replacing A/B with AI. It finds the subtle combinations of variables that lead to conversion increases. The web designer can spend more time trying ideas and less time doing statistics, giving them the freedom they need to make a difference.

REFERENCES

- [1] Tim Ash, Rich Page, and Maura Ginty. 2012. *Landing Page Optimization: The Definitive Guide to Testing and Tuning for Conversions* (second ed.). Wiley, Hoboken, NJ.
- [2] Daren C. Brabham. 2013. *Crowdsourcing*. MIT Press, Cambridge, MA.
- [3] Jürgen Branke. 2002. *Evolutionary Optimization in Dynamic Environments*. Springer, Berlin.
- [4] Builtwith. 2017. A/B Testing Usage. (2017). <https://trends.builtwith.com/analytics/a-b-testing> Retrieved 1/9/2017.
- [5] eMarketer. 2016. US Digital Ad Spending to Surpass TV this Year. (2016). <https://www.emarketer.com/Article/US-Digital-Ad-Spending-Surpass-TV-this-Year/1014469> Retrieved 2/1/2017.
- [6] Dario Floreano, Peter Dürri, and Claudio Mattiussi. 2008. Neuroevolution: From Architectures to Learning. *Evolutionary Intelligence* 1 (2008), 47–62.
- [7] Babak Hodjat and Hormoz Shahrzad. 2013. Introducing an Age-Varying Fitness Estimation Function. In *Genetic Programming Theory and Practice X*, Rick Riolo, Ekaterina Vladislavleva, Marylyn D Ritchie, and Jason H. Moore (Eds.). Springer, New York, 59–71.




Figure 7: Estimated Conversion Rates of the Top 20 Candidates vs. Control. The candidates are identified in the x -axis with numbers corresponding to those in Figure 4. They are ordered according to their estimated conversion rate. All these candidates are significantly better than the control at the 95% level. The best one (on the right) is 46.6% better, as confirmed with an independent A/B test that showed 43.5% improvement.

- [8] Ron Kohavi and Roger Longbotham. 2016. Online Controlled Experiments and A/B Tests. In *Encyclopedia of Machine Learning and Data Mining*, Claude Sammut and Geoffrey I. Webb (Eds.). Springer, New York.
- [9] Joel Lehman and Risto Miikkulainen. 2013. Boosting Interactive Evolution using Human Computation Markets. In *Proceedings of the 2nd International Conference on the Theory and Practice of Natural Computation*. Springer, Berlin.
- [10] Joel Lehman and Risto Miikkulainen. 2013. Neuroevolution. *Scholarpedia* 8, 6 (2013), 30977. <http://nn.cs.utexas.edu/?lehman:scholarpedia13>
- [11] Khalid Salehd and Ayat Shukairy. 2011. *Conversion Optimization: The Art and Science of Converting Prospects to Customers*. O'Reilly Media, Inc., Sebastopol, CA.
- [12] Jimmy Secretan, Nicholas Beato, David B. D'Ambrosio, Adelein Rodriguez, Adam Campbell, J. T. Folsom-Kovarik, and Kenneth O. Stanley. 2011. Picbreeder: A Case Study in Collaborative Evolutionary Exploration of Design Space. *Evolutionary Computation* 19 (2011), 345–371.
- [13] Sentient Technologies 2017. It's not A/B, I's AI. (2017). <http://www.sentient.ai/ascend> Retrieved 1/9/2017.
- [14] Hormoz Shahrzad, Babak Hodjat, and Risto Miikkulainen. 2016. Estimating the Advantage of Age-Layering in Evolutionary Algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2016)*. ACM, New York, NY, USA.
- [15] H. Takagi. 2001. Interactive Evolutionary Computation: Fusion of the Capacities of EC Optimization and Human Evaluation. *Proc. IEEE* 89, 9 (2001), 1275–1296. <http://ieeexplore.ieee.org/iel5/5/20546/00949485.pdf?tp=&arnumber=949485&isnumber=20546>
- [16] Daniel Yankelovich and David Meer. 2006. Rediscovering Market Segmentation. *Harvard Business Review* 84, 2 (2006).