

Novelty Wins, “Straight Toward Objective” Loses! or Book Review: Why Greatness Cannot Be Planned: The Myth of the Objective, by Kenneth O. Stanley and Joel Lehman

Reuben Hersh
University of New Mexico

Follow this and additional works at: <https://scholarship.claremont.edu/jhm>



Part of the [Robotics Commons](#), and the [Social and Behavioral Sciences Commons](#)

Recommended Citation

Reuben Hersh, "Novelty Wins, “Straight Toward Objective” Loses! or Book Review: Why Greatness Cannot Be Planned: The Myth of the Objective, by Kenneth O. Stanley and Joel Lehman," *Journal of Humanistic Mathematics*, Volume 5 Issue 2 (July 2015), pages 161-165. DOI: 10.5642/jhummath.201502.15.
Available at: <https://scholarship.claremont.edu/jhm/vol5/iss2/15>

©2015 by the authors. This work is licensed under a Creative Commons License.

JHM is an open access bi-annual journal sponsored by the Claremont Center for the Mathematical Sciences and published by the Claremont Colleges Library | ISSN 2159-8118 | <http://scholarship.claremont.edu/jhm/>

The editorial staff of JHM works hard to make sure the scholarship disseminated in JHM is accurate and upholds professional ethical guidelines. However the views and opinions expressed in each published manuscript belong exclusively to the individual contributor(s). The publisher and the editors do not endorse or accept responsibility for them. See <https://scholarship.claremont.edu/jhm/policies.html> for more information.

Novelty Wins, “Straight Toward Objective” Loses!
OR
Book Review: *Why Greatness Cannot Be Planned:
The Myth of the Objective*,
by Kenneth O. Stanley and Joel Lehman

Reuben Hersh

Department of Mathematics and Statistics, University of New Mexico, New Mexico, USA
rhersh@gmail.com

Synopsis

Experiments in evolutionary artificial intelligence demonstrate that progress toward an important, difficult goal is not best achieved by attempting to go directly toward that goal, but rather, by rewarding novelty.

Why Greatness Cannot Be Planned: The Myth of the Objective. By Kenneth O. Stanley and Joel Lehman, Springer, Cham, 2015. (Softcover US\$29.95, ISBN: 978-3-319-15523-4. 135 pages.)

This astonishing book starts out as a straightforward report of very specific, concrete results in the field of evolutionary artificial intelligence. In other words, getting programs to be smart by running them over and over and over, each time with some more-or-less random variations, each time selecting “the best” according to some definite criterion of superiority. Stanley is known as the creator of “NEAT”, an evolutionary algorithm for improving neural network artificial intelligence.

After a few dozen pages, the authors arrive at an empirical, quantitative justification of a long-standing *cri de coeur* of scientists: “Let me do what I find interesting! It might yield a bigger payoff than goal-directed, bottom-line-oriented compulsion!”

We all are accustomed to cliché oratorical tributes to “pure science”. But when it comes down to pounds and pence, it’s goal-directed “practical” proposals that get the most. Us pointy-headed nerds yearn to do our own thing, but the practical decision-makers know better.

Now Stanley and Lehman are here, backing us up, with evidence, with empirical, quantitative results! This book shows that when you’ve got a really hard problem to solve (cancer, say, or poverty, or inequality, or obesity) it’s more practical to search for something new and interesting, not stick to what seems to move you closer to your practical goal.

“Objectives are well and good when they are sufficiently modest,” they write “but things get a lot more complicated when they’re more ambitious. In fact, objectives actually become obstacles towards more exciting achievements, like those involving discovery, creativity, invention, or innovation—or even achieving true happiness. In other words (and here is the paradox), the greatest achievements become less likely when they are made objectives. Not only that, but this paradox leads to a very strange conclusion—if the paradox is really true then the best way to achieve greatness, the truest path to “blue sky” discovery or to fulfill boundless ambition, is to have no objective at all” (pages 7–8).

“It often turns out that the measure of success—which tells us whether we are moving in the right direction—is deceptive because it’s blind to the true stepping stones that must be crossed” (page 9).

Two experiments are reported in detail, to back up these claims.

The first experiment models rats in a maze. The rats try to get out. After each failure, their search is modified in several different ways, and the “best” new search is selected. After it fails, it is used to produce the next generation of variants. And so on. Evolution!

The question is, which do you choose as “the best”? Stanley and Lehman compare two different criteria. In one criterion, “best” is the one that gets the rat closest to the goal: escape. In the second criterion, “best” is the one that is most novel, most different from what has been already tried, regardless of whether it gets the rat closer to escape! Which criterion for choosing a search strategy works best? The search for novelty wins, by a huge margin. “To be specific, we repeated the experiment with novelty search 40 times and in 39 of these a robot behavior was discovered that solved the maze. The result with objective-based search: three times out of 40” (page 52).

This result is not so hard to understand. A maze has *culs-de-sac*. If you are a rat trapped in a *cul de sac*, it will do no good, to keep banging against it, in the direction of the exit. Better give up on that and try the opposite direction.

The next experiment is more surprising. Stanley and Lehman simulated a two-legged robot lacking a torso. The experimenters may not have known a method for this robot to learn to walk. But by letting it search for novelty, the robot did learn to balance and walk on its two legs.

“Novelty search isn’t trying to do anything in particular. It just looks at what the robot is doing and tries new versions of behaviors that were novel when they were discovered. So if the robot falls on its face, that’s good as long as the robot never fell down before in the same way. What do you think a biped robot looking for novelty would eventually end up doing? It turns out that the answer is that the robot learned to walk. In fact, it learned to walk better than when it tried to learn with the objective of walking. In other words, a robot that tried to walk farther and farther actually ended up walking less far than one that simply tried to do something novel again and again . . . Falling down and kicking your legs might be a better stepping stone [to walking] than trying to take a step, but if walking is the objective, falling down is considered one of the worst things you can do. So once again novelty search far out-performed the objective-based search” (pages 53–54).

These results have made a big impression in the field of evolutionary artificial intelligence.

Imagine a robot sent to travel down a hallway. If the hallway goes through a room with many closets, the robot can get trapped. A novelty-seeking robot will succeed in getting down to the end of the hall, much better than an objective-based robot. This picture of a hallway with closets can be a metaphor for any “landscape” where there are many suboptimal traps, inferior “hilltops” that interfere with finding the highest altitude.

A novelty-seeking search algorithm has to remember its past and avoid previous mistakes, so rewarding novelty means collecting information. Moreover, novelty-seeking goes past the simple algorithms and starts finding com-

plex ones. Being a kind of complexity-building device, it is more likely than an objective-based robot to solve a problem that requires a complex solution. Stanley and Lehman point out that the problem of improving education is a complex problem that demands a complex solution. A novelty-seeking strategy would be better than our present objective-based attempts to improve education.

Novelty-based search is similar to another evolutionary process that is not objective-based, namely, natural evolution! Natural selection is based on “fitness”—survivability and reproducibility in a given environment. It is different from selection based on novelty, but it also is one that is not objective-based. Complex organisms were an outcome of natural selection, which did not have *homo sapiens* as an objective. Complexity achieved by novelty-seeking, in evolutionary artificial intelligence, sheds light on complexity achieved in natural selection, another evolutionary experiment that is not objective-based.

Stanley and Lehman draw important consequences regarding policy in supporting scientific research. They argue that *as a general rule*, if the goal is not already “within sight”, but instead is best approached in some unknown direction, with sub-optimal *culs-de-sac* in between, then going “straight toward the objective” is delusory, and rewarding novelty is likely to succeed much faster. For a scientific researcher, there is a contrast between aiming at a well-defined objective, and following an interesting lead without knowing ahead of time where it will take him.

Novelty is not the only thing that makes a result interesting. Importance is also a major consideration, either practical importance—a direct payoff to human life—or theoretical importance—a major effect on how we see reality. This is the difference between mission-based research, and “pure” or open-ended research. The experience with evolutionary artificial intelligence strongly supports open-ended research, which is not always easy to reconcile with mission-oriented agencies. Stanley and Lehman bring many examples from science and technology, of advances made possible by unpredictable outcomes of open-ended research. This aspect of research makes it like the creative arts, like music or literature. A view of scientific research that many scientists would welcome, but a hard one to swallow, for most “funding agencies”.

“Serendipity” is a word for a chance occurrence or development of events in a happy or beneficial way. Novelty-based evolution is a way to “maximize serendipity”, to maximize the chance of finding or creating something interesting or valuable that you weren’t looking for. In fact, Stanley and Lehman’s advocacy of novelty-based evolution is itself an example of serendipity! They had created a program they called “Picbreeder.” It enabled anyone to evolve visual images by any criterion he wished to choose. They were intrigued that while users of Picbreeder did not do particularly well at creating whatever image they were seeking, they did surprisingly often come up with an attractive or useful image that was not what they had been looking for.

If you are yearning to do what’s interesting, rather than optimizing a “metric” of approach to a prescribed “objective”, you will love this book. Could it change the funding policies of governmental and non-governmental funders? I hope so. Maybe searching for novelty can yet save us.