

A Novel Human-Computer Collaboration: Combining Novelty Search with Interactive Evolution

In: *Proceedings of the 16th annual conference on Genetic and evolutionary computation, GECCO '14*, New York, NY. ACM

Winner of the Best Paper Award in Artificial Life / Robotics / Evolvable Hardware

Brian G. Woolley
Department of ECE
Air Force Institute of Technology
WPAFB, OH 45433, USA
brian.woolley@ieee.org

Kenneth O. Stanley
Department of EECS
University of Central Florida
Orlando, FL 32816, USA
kstanley@eecs.ucf.edu

ABSTRACT

Recent work on novelty and behavioral diversity in evolutionary computation has highlighted the potential disadvantage of driving search purely through objective means. This paper suggests that leveraging human insight during search can complement such novelty-driven approaches. In particular, a new approach called *novelty-assisted interactive evolutionary computation* (NA-IEC) combines human intuition with novelty search to facilitate the serendipitous discovery of agent behaviors in a deceptive maze. In this approach, the human user directs evolution by selecting what is interesting from the on-screen population of behaviors. However, unlike in typical IEC, the user can now request that the next generation be filled with *novel* descendants. The experimental results demonstrate that combining human insight with novelty search not only finds solutions significantly faster and at lower genomic complexities than fully-automated processes guided purely by fitness or novelty, but it also finds solutions faster than the traditional IEC approach. Such results add to the evidence that combining human users and automated processes creates a synergistic effect in the search for solutions.

Categories and Subject Descriptors:

I.2.6 [Artificial Intelligence]: Learning

General Terms: Algorithms

Keywords: Evolutionary computation, interactive evolutionary computation, human-led search, fitness, deception, non-objective search, novelty search, stepping stones, serendipitous discovery

1. INTRODUCTION

Several results in recent years have hinted at the limitations of traditional objective functions, wherein the more a candidate resembles the objective, the higher its fitness. An early hint that such an approach to fitness may be flawed was from experiments with the *novelty search* algorithm [13], which rewards novel behaviors instead of rewarding objective performance. Interestingly, novelty search significantly outperformed objective-based fitness in a deceptive maze-navigation domain [13, 16], showing counter-intuitively that in some deceptive cases it is possible that having no

specific objective may work better than rewarding progress toward the objective. Other results with novelty search in recent years have echoed this outcome [10, 17, 20].

At the same time, interactive evolutionary computation (IEC) [32], wherein humans guide evolution through interactive breeding, has also shown promise at finding needles in the haystack of possibilities. For example, Woolley and Stanley [35] found that images evolved by human users on the Picbreeder online service [26] could not be re-evolved by the very same evolutionary algorithm when they are made the automated objective, suggesting that humans are uniquely adept at identifying promising stepping stones.

Yet both novelty search and IEC face their own challenges. Work with novelty search has shown that it may become lost in especially large spaces [12, 15], and Takagi [32] warns that IEC is limited by human fatigue. However, interestingly, the advantages of both approaches can potentially address each other's disadvantages. That is, the main insight in this paper is that the ability of humans to identify promising stepping stones is naturally complemented by the ability of novelty search to generate *candidate sets* of potential stepping stones. In other words, novelty search can mitigate the main weakness of IEC (i.e. that humans grow tired quickly [32]) by offloading most of the exploratory work, while IEC can mitigate the tendency of novelty search to get lost in vast behavior spaces.

In this paper, a hybrid approach, called *novelty-assisted interactive evolutionary computation* (NA-IEC), is compared via the evolution of neurocontrollers for robots in the deceptive mazes of Lehman and Stanley [16] to pure novelty search, objective-based search, and a human-led search without the novelty assistance. Interestingly, while novelty search was previously shown significantly more effective than objective-based search in this domain [16], NA-IEC outperforms novelty search by a multiple of three to four times, yielding by far the fastest solution on these deceptive problems. This work is complementary to recent results from Bongard and Hornby [1] and Hornby and Bongard [11] that showed that humans and more conventional objective-driven search can work together synergistically to solve practical control problems rather than just to explore creative domains. This paper builds on this idea by showing that novelty also has an important role to play in such a human-machine collaboration, in part to push the human out of his or her comfort zone to explore new ideas.

2. BACKGROUND

This section reviews deception in EC and the non-objective methods that are the basis for the approach introduced in this paper.

2.1 Deceptive Task Domains

The key question in research on deception is what causes evolutionary algorithms (EAs) to fail and how to mitigate such fail-

ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the United States government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

GECCO'14, July 12–16, 2014, Vancouver, BC, Canada.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2662-9/14/07 ...\$15.00.

<http://dx.doi.org/10.1145/2576768.2598353>.




Figure 1: Maze Navigation Maps (Lehman and Stanley [13, 16]). The deceptive maze domain is a metaphor for search and is *not* a path-planning problem. Rather, the aim is to evolve a neural network that drives a robot through the maze. The walls represent barriers to search and the cul-de-sacs represent local-optima that can deceive objective-based search.




Figure 2: Maze Navigation Robot (Lehman and Stanley [13, 16]). The sensor package (a) includes six rangefinders that detect walls and four pie-slice sensors that signal the general direction to the objective. The navigation behavior, encoded as an ANN (b), maps sensor inputs to actions, i.e. turn rate (*left/right*) and velocity (*forward/backward*). Under this construction, navigators cannot see the whole maze and must evolve a control policy that traverses the maze based on sensory input.

ures [8, 34]. For the purpose of this work, we are interested in the case in which pursuing what appears to be a reasonable *objective* produces an unreasonable *objective function*. In this context, an intuitive definition of deception, as stated by Lehman and Stanley [16], is: “A deceptive objective function will *deceive* search by actively pointing the wrong way.” That is, the fitness function can point the wrong way because not only must it reward the objective, but it must also reward the intermediate solutions (i.e. the *stepping stones*) that lead to the objective. Often these stepping stones do not improve performance on the objective function (and may even decrease it), causing search algorithms to forsake the most promising candidates.

A good example of a deceptive domain, which is also the experimental domain in this paper, is the deceptive maze domain introduced by Lehman and Stanley [13, 16], in which a simulated robot must navigate through a maze with deceptive cul-de-sacs (figure 1). The maze-navigation agents that act within the maze have a sensor package with six rangefinders that detect the walls and four pie-slice sensors that signal the direction to the goal (figure 2a). Each robot’s navigation behavior, encoded as an artificial neural network (ANN), maps sensor inputs to actions, i.e. turn rate (*left/right*) and velocity (*forward/backward*), as shown in figure 2b. Under this construction, navigators must evolve a control policy that traverses the maze based on sensory input.

The medium (293 x 134 units) and hard (195 x 202 units) maps in figure 1 are deceptive by design because the maps contain cul-de-sacs that represent local optima in the search space. If fitness is assigned based on reducing the distance to the goal, then the objective function prunes out of the search the deceptive intermediate solutions (i.e. those that move away from the goal location) needed to reach the global objective. While an alternative objective function that rewards specific intermediate solutions is conceivable

(and in fact will be explored later in this paper as well), the original point of this domain was to explore the effect of objective fitness when the precise stepping stones are *not known*, which is the typical predicament in most domains. In such cases, as in the standard objective function here, performance is generally rewarded for its proximity to the target behavior. Thus evolution driven by proximity to the goal often converges to a cul-de-sac from which the goal is inaccessible.

Despite extensive research, deception remains a significant problem in the field of EC [9, 18, 22]. The problem is that evolutionary algorithms (EAs) ultimately respond to the selection pressures created by the fitness function, which is often misleading. The challenge is to determine how to reward the intermediate steps that are required to reach the goal. In effect, what appears to be a reasonable heuristic may actually *prevent* the objective from being reached. Therefore, any similarity metric that guides the search toward an a priori objective is potentially a false compass to the optimal solution [28]. In this spirit, Lehman and Stanley [13, 16] introduced the idea of abandoning objectives as a search heuristic in deceptive domains, electing instead to reward individuals only for novel behaviors, as described next.

2.2 Novelty Search

Because pursuing objectives can be deceptive, Lehman and Stanley [13, 16] proposed that searching without regard to the objective, i.e. searching only for novel behavior, can be more effective at discovering solutions in some deceptive domains than rewarding objective performance.

Novelty search works with EAs by replacing the fitness function with a *novelty metric*. The novelty metric is a measure of the uniqueness of an individual’s behavior at a given task. Instead of rewarding performance, novelty search rewards individuals in the population for finding new ways to complete the evaluation task, thus creating a constant pressure to do something new [16].

Because novelty search operates in *behavior space*, it is important first to characterize the space of unique behaviors in a way that is meaningful to the domain. The novelty search algorithm then computes the *sparseness* in the behavior space as the average distance to the k -nearest neighbors [2] around that behavior. The sparseness ρ of behavior x is given by

$$\rho(x) = \frac{1}{k} \sum_{i=0}^k \text{dist}(x, \mu_i), \quad (1)$$

where μ_i is the i th-nearest neighbor of x with respect to the distance function $\text{dist}(x, \mu)$. In this way, if the average distance is large, then the candidate solution is considered to be in a sparse area of the behavior space, thus making it more likely to be selected by the EA. Optionally, as in coevolution [4], an archive of past behaviors may serve to avoid backtracking through the behavior space. If the novelty metric is sufficiently high for a new individual (i.e. above some minimal threshold $\rho_{\min}$), then the individual may be recorded in the permanent archive to provide a comprehensive sample of where the search has been, thereby increasing the pressure to discover new ways of behaving in the domain [13, 16].

Characterizing behaviors so that they can be compared is the most challenging aspect of novelty search. For the deceptive maze domain [13, 16] (figure 1), the behavior of a maze navigation robot is usually defined as its *final position*. In this way, the novelty metric rewards controllers that end at new locations in the maze. At first, the collection of behaviors may include robots that do nothing, get stuck in corners, run in circles, and so forth. However, at some point, the collection of simple behaviors becomes saturated and the pressure to do something new increases, i.e. evolution favors mutations that take the navigator to new places in the maze.

While the idea of selecting anything novel may sound potentially similar to exhaustive search, searching in the space of *behaviors* is often tractable because many points in the space of possible *genomes* collapse to a single behavior. Furthermore, when applied in conjunction with complexifying algorithms like NEAT [13, 15] and GP [14], simple behaviors become associated with minimal representations, and only mutations that increase the size of the genome and lead to novel behaviors are explored further. Therefore, this approach, operating without regard to an objective, moves into complex spaces in a meaningful way because new behaviors are those that could not be expressed at lower levels of complexity, i.e. complexity is rewarded when it is warranted.

However, experience has also shown that novelty search becomes lost in unrestricted domains [15]. In such domains there is an opportunity to leverage human knowledge rather than exhaustively exploring the space of all possible solutions. Thus the next section provides relevant background on the field of human-led evolution.

2.3 Interactive Evolutionary Computation

In *interactive evolutionary computation* (IEC) the traditional objective-function is replaced by a user who performs selection [32]. IEC is effective in creative domains [25] where the term *fitness* is subjective because what people experience as pleasing or interesting is based on individual preferences.

Like traditional EAs, IEC systems also typically begin from a random initial population that evolves over generations by selecting, mating and mutating members. However, IEC differs from traditional automated EAs in that a human user is now responsible for the evaluation and selection of promising candidate solutions. While this difference typically leads to smaller population sizes and higher mutation rates, the most profound implication is that evolution is no longer bound to a rigid expression of what is fit and unfit. While efforts have been applied to model human selection criteria [11], the human evaluator’s breadth of experience makes it likely that his or her selection criteria will change over the course of evolution. Such an ability to make *serendipitous discoveries*, i.e. to identify and pursue important artifacts as they emerge, is an important motivation for the NA-IEC approach introduced in this paper.

To interface with the human evaluator, the majority of IEC systems are modeled after the original Blind Watchmaker *Biomorphs* application by Dawkins [3]. In this approach the user is presented with a panel of individuals (e.g. 3×4) from which the parents of the next generation are selected. The IEC system then mates, recombines, and mutates the genetic material of the parents to create the next generation, which is then presented to the user. This process is repeated at the user’s direction until the user is satisfied.

Despite the benefits of having a human in the loop, such IEC systems are limited by user fatigue. According to Takagi [32], typical IEC only lasts 10–20 generations per session. The problem is that the vast majority of significant discoveries exist beyond the reach of a single-user session. One response, which has become known as *collaborative interactive evolution* (CIE) [31]), is to leverage the efforts of many users. That way, when one user fatigues, another can take over and continue the chain. One particular example is *Picbreeder* (<http://picbreeder.org> [26, 27]), where such CIE allows long chains of users to branch from each other’s discoveries, which indeed demonstrates that when the generational limits of a single user are extended significantly, remarkable discoveries begin to accumulate (figure 3).

However, chaining users is not the only way to augment the power of a single user. An alternative approach is to combine a human with a form of automated search, thereby amplifying the




Figure 3: Images Evolved on Picbreeder (Secretan et al. [26]). These images were interactively evolved by a community of human users. They demonstrate the system’s ability to discover interesting and meaningful images, including such seminal images as the *Skull* and the *Car*—which were evolved in 74 and 106 generations, respectively, by chains of users branching from each other.

human’s capabilities. While recent work by Bongard and Hornby [1] and Hornby and Bongard [11] shows that evolution driven by an objective can help to push search farther than the human alone, the novel approach in this paper adds the option to run novelty search into the mix. That way, users can request that the computer quickly gather a set of unique stepping stones that the user can then assess to decide which path is most promising, as explained next.

3. NA-IEC FRAMEWORK

The main idea in this section is to combine for the first time the intuitive ability of human users to identify what is interesting and important in a domain, i.e. interactive evolutionary computation (IEC), with a stepping stone generator based on a short-term novelty search and an objective optimizer to create a synergistic effect that expedites the evolution of controller solutions. Under this new approach, called *Novelty-Assisted Interactive Evolutionary Computation* (NA-IEC), a human user is asked to select individuals from a population of candidate behaviors and then apply one of three evolutionary operations: a traditional IEC step, a short-term novelty search, or a fitness-based optimization.

In this way, the user can apply the evolutionary operations where appropriate, even changing the mode of evolution during the course of the search, to reach a satisfactory (or just interesting) solution. The ability of a human user to apply powerful automated approaches like objective-based search [5–8] and novelty search [13, 16] in short bursts and when appropriate is a key contribution of the NA-IEC approach. The primary hypothesis is that letting the user make a relatively small number of critical selections during evolution, and leaving the remainder of search to automated approaches seeded by those user selections, can significantly augment the pace of evolution and the quality of its discoveries.

Figure 4 shows the main interface for the system, where the user can choose among the *Step*, *Novelty*, and *Optimize* operations. Choosing the *Step* operation creates a new generation of offspring through the recombination and mutation of the selected candidate behaviors. This classic approach to IEC is simple and computationally inexpensive, i.e. it only creates a handful of new candidates.

Choosing the *Novelty* operation causes evolution to explore the space of agent behaviors without regard to an objective and then present the human evaluator with a broad view of where the evolutionary search can go from its current position. To accomplish this aim, the next IEC population is generated by seeding a larger population with variations of the user-selected candidate behaviors and then running novelty search in the background to find novel individuals (in comparison to what has been encountered previously in the search) based on the sparseness measure $\rho(x)$ from equation 1 and the threshold $\rho_{\min}$. The underlying evolutionary algorithm is NEAT [29, 30], which is often the base algorithm under novelty search [13, 16]. Furthermore, to ensure that novelty is measured




Figure 4: Screenshot of the NA-IEC user interface. The user interface for the NA-IEC framework consists of the *Evolution Controls*, the *Evolution Options*, and the *Evaluation Population*. Candidate solutions are represented by a gradient trail that shows the robot’s behavior in a particular maze. Selected candidates are shown with a green border and solutions are highlighted with a white background. Unlike traditional IEC applications, the user can now select one of three evolution modes: *Step*, *Novelty*, and *Optimize*. The *Publish* button saves the results of a completed run for later analysis. Like Picbreeder, projects based on NA-IEC could additionally leverage crowdsourcing by implementing them an online format. [21].

with respect to the entire search completed so far, all individuals encountered during *both* traditional IEC steps and interleaved novelty searches throughout a session of NA-IEC are measured for their novelty and entered into the permanent archive if their novelty score is greater than the threshold $\rho_{\min}$.

The novelty search runs until *at least* n new individuals are added to the evaluation population (more than n such novel candidates may be found when the novelty search is first started by generating an initial pool of candidates based on the user-selected choices on the screen), where n is the size of an on-screen IEC population. At that point, the collected novel individuals become the next IEC generation and control is returned to the user. By convention, the n novel individuals are sorted by their novelty score before the NEAT-based speciation adjustment to place the most novel candidate behaviors on the first visible page of the on-screen IEC population (figure 4). While the *Novelty* operation is significantly more computationally expensive than the *Step* operation, it provides the human user with a breadth of stepping stones that would have been time-consuming or unlikely to discover on his or her own under the narrow view of a traditional single IEC step, which only presents the user with a handful of direct one-generation descendants. In a sense, the set of stepping stones returned to the user by novelty search is like the set of images evolved by *other* users from which a visitor to Picbreeder can branch: In both cases, someone or something else has put in effort to collect a set of interesting jumping-off points and present them to the user.

By augmenting the human-led interactive search with interleaved novelty searches, a small population can be constructed that contains a set of novel stepping stones around the currently-selected candidates. In the event that evolution cannot fill the next generation with a sufficient number of new archive members in a reasonable amount of time, or if the user cancels the operation, the search will return the most novel individuals yet discovered.

This approach does *not* imply that the set of novel agent behaviors presented to the evaluator will be good at a potential task. What is important is that they are *behaviorally* diverse; it is the human evaluator who will direct the search by recognizing what is promising for a given domain. The goal is to promote innovation through serendipitous discovery, and presenting the various directions that the search can take leverages the human evaluator’s inherent ability to recognize what is important or interesting in a particular domain.

Finally, because objective-based optimization is likely the best option for *perfecting* well-formed behaviors already discovered, the user is also given the option to request seeding a traditional objective-based search with currently-selected individuals. The objective-based search will run until a specified solution criterion is met or until the user requests it to terminate, at which point the most fit individuals discovered so far will update the on-screen IEC population. Providing this traditional option will allow users to optimize candidates that are near an objective attractor that the user would prefer to approach automatically once it is within striking distance (i.e. once the search is no longer deceptive and the primary discovery is already made).

In this multifaceted approach, the user is free to change the mode of evolution between generations, thus allowing evolution to proceed in the capacity best suited for the current context. In this way, the human user may begin NA-IEC by exploring the space of behaviors agnostically, and once an interesting behavior is established, the mode of evolution may be changed to optimize it. For more information on NA-IEC, or to download the framework, please visit <http://eplex.cs.ucf.edu/naiec>.

4. EXPERIMENT

To demonstrate the synergistic effects of augmenting a human-directed search with novelty search, the experiment is conducted in the deceptive maze domain introduced by Lehman and Stanley [13, 16, Section 2.1]. That way, the NA-IEC approach can be compared against pure novelty search, fitness-based search, and a novelty-free IEC-based search directly. In the deceptive maze domain, the goal is to evolve a navigation behavior that drives a robot from the start to the finish of the medium maze or the hard maze shown in figure 1, which are constructed with several cul-de-sacs that create local optima in the fitness landscape. Because these local optima are so deceptive, Lehman and Stanley [13, 16] found that novelty search significantly outperforms objective-based search in both mazes. The question here is, can NA-IEC do even better?

To compare performance, each approach is evaluated over 30 runs on the medium and hard maps. While novelty search and fitness-based search are both automated algorithms, the NA-IEC approach requires a human evaluator. To accomplish the NA-IEC portion of this experiment, six users (who are not the authors) were recruited who were familiar with novelty search and EAs. These users were introduced to the NA-IEC framework and each asked to evolve five solutions to the medium map and five solutions to the hard map. The aim is to characterize the performance that can be reasonably expected from a practitioner in EC when evolving with NA-IEC. Users were permitted to restart if they felt that evolution had become stuck. However, all evaluations before such restarts were recorded as a part of the same run.

Inevitably, some will argue that such human guided runs have an unfair advantage because the user can see the path through the maze. To address this concern, an additional fitness-based experiment, inspired by Risi and Stanley [24] is conducted. In this additional experiment the primary deceptive element of the maze navigation domain, i.e. the attraction of agents to cul-de-sacs, is removed. In this alternative reward scheme, candidates are rewarded




Figure 5: Maze navigation waypoints. To compare how much advantage is gained from knowing the path to the goal, waypoints (which are *not* seen by the agent) are provided for the medium and hard maps. In this way, deception is removed by allowing a traditional fitness-based search to reward solutions that discover stepping stones that are on the path to the goal.

for progressing along a path that actually leads to the goal. Figure 5 shows the waypoints (which are invisible to the agent) in the medium and hard maps. In this waypoint-directed version of the experiment, the fitness function f is defined such that agents are rewarded for each waypoint crossed (including the goal); they also receive a partial reward for approaching the next waypoint:

$$f = n + (1 - d), \quad (2)$$

where n is the number of waypoints reached and d is the distance to the next waypoint (proportional to distance between waypoint w_n and w_{n+1} , in the range $[0, 1]$).

Some might still argue that the waypoint-directed search above is at a disadvantage. Thus an additional human-led experiment is also conducted to further characterize the specific advantage provided by the human evaluator. In fact, this experiment is identical to the NA-IEC experiment with two exceptions: (1) the short-term novelty operation is not available to the human user, and (2) the mutation rate of the step operation is matched to the rates of the other experiments. In this way, the human user can evaluate individual generations of candidate behaviors just as in IEC, but can then tune (without waypoints) their solutions with the optimize operation to achieve the final solution. In this setup, five human subjects performed six runs each on both mazes.

4.1 ANN Representation

In this experiment, as in Lehman and Stanley [13, 16], the ANN controllers in all variants of the deceptive maze experiment are evolved by the NeuroEvolution of Augmenting Topologies (NEAT) approach [29, 30]. More specifically, the NEAT algorithm starts with a population of simple ANNs and *complexifies* them over generations by adding new nodes and connections through structural mutations. By evolving networks in this way, the topology of the network does not need to be known a priori. As applied to maze navigation policies, this process begins with an initial population of simple behaviors that are represented by fully-connected networks with 22 connections, no hidden nodes, and the inputs/outputs in figure 2b. As the underlying networks add complexity (i.e. new nodes and connections), features and nuances emerge in the resulting behaviors that could not be expressed by the simpler ANNs.

4.2 Experimental Parameters

The evolutionary parameters in this experiment are based on the deceptive maze navigation experiment by Lehman and Stanley [13, 16] and on the established parameters for NEAT [29]. All experiments were run with a version of the public domain ANJI package augmented to support steady-state evolution, interactive evolution, and novelty search. The IEC population size was 12, while the novelty search and fitness-based search population sizes were 250, with each run limited to 250,000 total evaluations. Note that

when the user initiates novelty search or optimization from within NA-IEC, a starting pool of 250 candidates are first generated from the user-selected candidates on the screen. The speciation threshold, δ_s , was 0.2 and the compatibility modifier was 0.3. Offspring in all automated runs, short-term novelty operations, and IEC step operations had a 5% chance of adding a node, a 10% chance of adding a link, a 1% chance of losing a link, and the weight mutation power was 0.8. During optimization operations and Step operations in the NA-IEC experiment the connection weight mutation power was 0.1 to support behavioral fine-tuning. The ANN employed unsigned sigmoidal activation with recurrent connections, resulting in a shifted output range of $[-0.5, 0.5]$.

The parameters specific to novelty were also based on the original deceptive maze navigation experiment [13, 16]. They include the nearest neighbors value ($k = 15$) and the novelty threshold, ρ_{min} , which begins at 3.0 and is adjusted after every 2,500 evaluations. Each navigation robot was given 400 time-steps to reach the goal, which only allows behaviors that proceed directly to the goal. It is important to note that the experiments of Lehman and Stanley [13, 16] were *re-run* with this setup to ensure a fair comparison and to validate our implementation.

5. RESULTS

As with the original experiment by Lehman and Stanley [13, 16], a navigation behavior that finishes within five units of the goal location is considered successful. The main result is that NEAT with NA-IEC discovers such solutions in significantly fewer evaluations than both NEAT with novelty search and fitness-based NEAT on the medium and hard maps. Furthermore, despite the expense of waiting on the human to evaluate a panel of candidate solutions, NA-IEC also consumes less wall-clock time in search, suggesting that the value of the user’s direction easily offsets the delay of waiting for human input. Another result is that NA-IEC produces solutions with significantly fewer hidden nodes than both novelty search and fitness-based search, further suggesting the importance of allowing a human evaluator to make key decisions about the direction of evolution. While some may dismiss such improvements based on the human evaluator’s ability to see the path through the maze, results from the waypoint-directed search, a non-deceptive fitness-based experiment, are on par with NEAT with novelty search, which is still well below the performance of NEAT with NA-IEC. Furthermore, NA-IEC not only succeeds in discovering solutions in significantly fewer evaluations than IEC without novelty, but also reduces fatigue by requiring many fewer user operations. The implication is that NEAT with NA-IEC not only exposes key stepping stones, but also provides evolution with subtle insights about the domain not easily incorporated into a traditional fitness function a priori.

On the medium map, users directing NEAT with NA-IEC found 30 solutions in an average of 6,729 (sd = 8,068) evaluations. These results are significantly ($p < 10^{-6}$; Student’s t-test) faster than NEAT with novelty search (22,116 evaluations, sd = 10,157), fitness-based NEAT (55,066 evaluations sd = 47,339), waypoint-directed NEAT (22,594 evaluations sd = 11,982), and human-led NEAT without novelty (12,741 evaluations sd = 12,792), each averaged over 30 runs (figure 6a). Furthermore, users solved the medium map in an average of 294 (sd = 359) seconds, which is 2.8 times faster than novelty search, 9.1 times faster than fitness-based search, and 2.0 times faster than the waypoint-directed search.¹ While solutions from the novelty-based, fitness-based, waypoint-

¹Note that clock times are not reported for IEC without novelty because these runs took so long that users generally could not complete them in one sitting, which would confound the results.

directed, and human-led searches have on average 3.2 (sd = 1.9) hidden nodes, 2.9 (sd = 1.65) hidden nodes, 3.0 (sd = 1.8) hidden nodes, and 1.0 (sd = 3.1) hidden nodes respectively, solutions produced by NA-IEC are significantly simpler, averaging just 0.23 (sd = 0.5) hidden nodes per solution ($p < 10^{-10}$; Student's t-test).

On the hard map, the NA-IEC approach evolved 30 successful navigators in an average of 7,481 (sd = 6,610) evaluations, which is a significant ($p < 10^{-5}$; Student's t-test) improvement over not only NEAT with novelty search alone (33,320 evaluations, sd = 20,949), but also over the non-deceptive (i.e. waypoint-directed) version of fitness-based NEAT (26,954 evaluations, sd = 18,464) and the human-led NEAT without novelty (24,447 evaluations, sd = 26,691), each averaged over 30 runs. In the case of fitness-based NEAT, as in Lehman and Stanley [13, 16], no comparison could be made because only four of 30 runs evolved solutions for the hard map. Evaluations required to solve the hard map are shown in figure 6b. In addition to evolving successful navigators for the hard map in fewer evaluations, NA-IEC did so on average in just 402 (sd = 374) seconds, which is 3.5 times faster than NEAT with novelty search and 2.5 times faster than the waypoint-directed search. Regarding complexity, solutions from novelty search have on average 3.3 (sd = 1.8) hidden nodes, solutions from the non-deceptive waypoint-directed search have an average of 3.5 (sd = 2.0) hidden nodes, and solutions from the human-led search averaged 1.2 (sd = 1.4) hidden nodes, while those evolved by NA-IEC are significantly smaller with 0.5 (sd = 1.01) hidden nodes ($p < 10^{-8}$; Student's t-test).

Typical patterns of exploration for each approach in the medium and hard maps are shown in figure 7, which compares the distribution of all ending points visited during a typical run. As Lehman and Stanley [13, 16] discovered previously, the traditional fitness-based approach is attracted to the cul-de-sacs in the maze (figures 7a and 7b), while selecting for behavioral novelty allows NEAT to explore the space of possible behaviors more evenly (figures 7c and 7d). Such search distributions are the result of selection pressure; thus when the objective-function rewards agents for following the solution path (figures 7e and 7f) the cul-de-sacs no longer deceive evolution. Interestingly, when the points visited during IEC (figures 7g and 7h) and NA-IEC (figures 7i and 7j) are plotted in this way the signatures of the human selector becomes evident. As expected, the first of these is that there are far fewer points in the cul-de-sacs than in both novelty search and even the waypoint-directed search, demonstrating the intolerance of the human user for behaviors that explore these spaces. The second signature is that there are frequently tight groupings of points at key junctions in the map, indicating that the user is probing these areas of the search space for a behavior that can turn a corner and enter a new chamber of the maze. Such observations demonstrate how the human evaluator is contributing his or her insights to the search. Furthermore, it is interesting how these human effects are so readily visible in the points plotted. However, the key difference between the IEC runs (i.e. without the short-term novelty operator) and NA-IEC runs is the regularity with which solutions can be found in a small number of evaluations, highlighting the advantage of having the novelty search option.

Finally, it is also important to analyze the behavior of the human users, especially in light of the human susceptibility to fatigue in IEC systems [32]. In the IEC experiment without novelty, users found solutions to the medium map in an average of 148.8 (sd = 128.5) operations, 79.1% of which were *Step* functions and 20.9% were *Optimize* functions. Similarly, solutions found for the hard map required an average of 1,179.5 (sd = 1,837.9) user operations, of which 95.6% were *Step* functions and 4.4% were *Optimize*




Figure 6: Evaluations required to find solutions. The number of evaluations required by NEAT to find solutions are shown for the medium (a) and hard (b) maps with NA-IEC, IEC augmented with optimization only, novelty search, fitness-based search (pure), and waypoint-directed search. The average number of evaluations to reach a solution is marked by a line, boxed regions extend out to one and two standard-deviations, and the distribution of the individual data points is shown. As in work by Lehman and Stanley [13, 16], fitness-based search is generally deceived in the hard map and is unlikely to produce solutions. The main result is that the NA-IEC approach consistently finds solutions for the medium and hard maps in significantly fewer evaluations than any other approach.

functions. In contrast, users working with NA-IEC found solutions for the medium map in an average of just 30.1 (sd = 40.5) choices, applying the *Step* function 29.8% of the time, the *Novelty* function 47.8% of the time, and the *Optimize* function 22.4% of the time. On the hard map, solutions required an average of 32.0 (sd = 23.5) human choices, of which 29.2% were *Step* functions, 58.9% were *Novelty* functions, and 11.9% were *Optimize* functions. Furthermore, results show that users tended to employ novelty early in evolution and applied optimization at the end of evolution, i.e. once a solution has been established. Combined with the overall performance results, these statistics demonstrate the important role played by the *Novelty* operator in dramatically reducing both the overall cost of a run and the likelihood for user fatigue [32]. As these results also show, the human's ability to see through the maze in the IEC runs could not make up for the lack of novelty search, demonstrating again that search (even interactive search augmented with objective optimization) benefits significantly from novelty.

6. DISCUSSION

In the deceptive maze domain, humans in NA-IEC make a good team with novelty search and objective optimization, which helps to finish the job. In both mazes, users choose *Novelty* to generate the next set of choices significantly more frequently than the other options. The stepping-stone generator of novelty search provides




Figure 7: Distribution of final points visited. Each maze shows the final position for all candidates in a typical run. The density of points shows how fitness-based NEAT, NEAT with novelty search, waypoint-directed NEAT, IEC-based NEAT, and NEAT with NA-IEC behave in the deceptive maze domain. As in Lehman and Stanley [13, 16], fitness-based search is attracted to the cul-de-sacs, while the points visited by novelty search are more evenly distributed throughout the maze. When rewarded for progressing along the solution path, the waypoint-directed runs are less deceived. For the IEC and NA-IEC runs, the human user’s influence is clearly visible, i.e. there are significantly fewer points in the major cul-de-sacs and tight groupings of points around key junctions. Such characteristics reveal how human selections are impacting evolution.

a desirable menu of possibilities to the human user, ultimately exceeding the performance of novelty search alone, as well as the reported performance of novelty–fitness multi-objective hybrids [19], by several times. Nevertheless, a natural question is whether such results are somehow specific to the maze domain. Perhaps humans harbor a particularly keen insight into the most promising robot behavior in mazes, but would lack such insight in other domains.

For example, one hypothesis might be that humans in effect know the right path through the maze because they can see the whole maze. Yet this interpretation is not entirely accurate. The correct path *through the maze* is not equivalent to the correct path *through the search space*. While some behaviors seem clearly dead ends (such as being caught in the most obvious cul-de-sac in the hard maze), others are less obvious. It is not necessarily the case that just because one behavior drives the robot farther down the correct path that it must be a more promising stepping stone. Some such behaviors are themselves dead ends that cannot push farther. Also, humans perceive more subtle and nuanced indicators that are also important, such as path smoothness or unnecessary loops in the robot trajectory. A behavior in which the robot doubles back on itself and then turns back onto the correct path may be just as ominous as being stuck in a dead end. Humans intuitively understand *these* kinds of dangers, yet to articulate them in an objective function would be quite challenging, and would almost certainly take more time than simply guiding the search away from them, whether they are easy to formalize or not.

Interestingly, the NA-IEC approach found solutions in significantly fewer evaluations and with fewer user operations compared to the IEC runs conducted without the novelty operator. In *both* of these human-led experiments, the evaluators had full insight into the correct route through maze and the robot path behaviors for each candidate, as well as the option to run optimization, yet those advantages alone without novelty were still less effective. The insight is that the ability to request a set of novel candidate behaviors (whether good or bad) plays an important role by affording the user a broad view of where evolution could go next. This additional information appears to facilitate better decisions that lead to increased efficiency.

While only future empirical results can settle this issue, there is reason to believe that humans would carry similarly critical insights into other domains. For example, in a biped-walking task [16, 23, 33], humans can see that certain kinds of leg oscillations are promising even if the robot falls down. Yet to describe exactly what makes them promising in a fitness function is likely prohibitive. The human’s overhead view, and hence knowledge of the mazes, should be viewed metaphorically as like any intuitive understanding of the shape of a particular behavior space. Just as we can see in the maze that certain passageways must precede other passageways, so we can see in a biped robot that oscillations and balance must precede walking. While it is possible that the intuitive insight into some domains is less than in the maze domain, the highly significant advantage provided by such insight in the maze domain suggests that even if the advantage were less elsewhere, it could still be significant.

NA-IEC also may be important for more than just optimization. In some spaces, such as in morphological brain-body evolution or with sophisticated encodings, we may be more interested in *what is possible* than in achieving a particular end result. The apparent synergy that results from humans combined with novelty search could be leveraged in the future to show us more about such spaces than trying to solve specific problems. With all the limitations recently shown for objective-based search, NA-IEC provides an alternative without relinquishing our desire to have some say in the process.

7. CONCLUSION

This paper introduced the *novelty-assisted interactive evolutionary computation* (NA-IEC) approach, wherein the intuitive ability of human users to identify promising candidate behaviors is augmented by an agnostic stepping stone generator (i.e. novelty search) seeded with the behaviors selected by the user. In this way, the advantages from IEC and novelty search create a complementary effect that offsets each other's limitations. Humans in this study only spent up to ten minutes to make a few dozen selections among thousands of evaluations, much of which was automated by novelty search. This novel incorporation of single mutational steps, short-term novelty search, and fitness-based optimization into a single evolutionary tool resulted in an effective synergy that allowed human users to realize what was important for a given domain *during* evolution. Furthermore, such serendipitous exploration found solutions in fewer evaluations, at lower genomic complexities, and in significantly less time overall than the fully-automated processes (i.e. novelty search, fitness-based search, and the directed fitness search alone), while also reducing fatigue when compared to IEC without the novelty component. Thus the key contribution of the NA-IEC approach is that it can accelerate the rate and quality of evolution by leveraging human-level domain knowledge without burdening the user with the responsibility of evaluating every candidate created during evolution.

References

- [1] J. C. Bongard and G. S. Hornby. Combining fitness-based search and user modeling in evolutionary robotics. In *Proceedings of the 15th annual conference on Genetic and evolutionary computation*, GECCO '13, pages 159–166, New York, NY, USA, 2013. ACM.
- [2] T. M. Cover and J. Thomas. *Elements of Information Theory*. Wiley, Hoboken, NJ, 1991.
- [3] R. Dawkins. *The Blind Watchmaker*. Longman, Essex, U.K., 1986.
- [4] E. D. De Jong. The incremental pareto-coevolution archive. In *Proceedings of the 6th annual conference on Genetic and evolutionary computation*, GECCO '04, pages 525–536. Springer Berlin Heidelberg, 2004. ACM.
- [5] K. A. De Jong. *Evolutionary Computation: A Unified Perspective*. MIT Press, Cambridge, MA, 2002.
- [6] A. Eiben and J. Smith. *Introduction to evolutionary computing*. Springer-Verlag, 2003.
- [7] D. B. Fogel. *Evolutionary computation: toward a new philosophy of machine intelligence, third edition*. Wiley-IEEE Press, 2006.
- [8] D. E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Reading, MA, 1989.
- [9] D. E. Goldberg and J. Richardson. Genetic algorithms with sharing for multimodal function optimization. In *Proceedings of the Second International Conference on Genetic Algorithms*, pages 148–154. San Francisco: Kaufmann, 1987.
- [10] J. Gomes, P. Urbano, and A. Christensen. Evolution of swarm robotics systems with novelty search. *Swarm Intelligence*, pages 1–30, 2013.
- [11] G. S. Hornby and J. C. Bongard. Accelerating human-computer collaborative search through learning comparative and predictive user models. In *Proceedings of the 14th annual conference on Genetic and evolutionary computation*, GECCO '12, pages 225–232, New York, NY, USA, 2012. ACM.
- [12] S. Kistemaker and S. Whiteson. Critical factors in the performance of novelty search. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, GECCO '11, pages 965–972, New York, NY, USA, 2011. ACM.
- [13] J. Lehman and K. O. Stanley. Exploiting open-endedness to solve problems through the search for novelty. In *Proceedings of the Eleventh International Conference on Artificial Life (Alife XI)*, Cambridge, MA, 2008. MIT Press.
- [14] J. Lehman and K. O. Stanley. Efficiently evolving programs through the search for novelty. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation*, GECCO '10, pages 837–844, New York, NY, USA, 2010. ACM.
- [15] J. Lehman and K. O. Stanley. Revising the evolutionary computation abstraction: minimal criteria novelty search. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation*, GECCO '10, pages 103–110, New York, NY, USA, 2010. ACM.
- [16] J. Lehman and K. O. Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2):189–223, 2011.
- [17] J. Lehman and K. O. Stanley. Novelty search and the problem with objectives. In *Genetic Programming Theory and Practice IX (GPTP 2011)*, New York, NY, 2011. Springer.
- [18] G. E. Liepins and M. D. Vose. Deceptiveness and genetic algorithm dynamics. Technical Report CONF-9007175-1, Oak Ridge National Lab., TN (USA); Tennessee Univ., Knoxville, TN (USA), 1990.
- [19] J.-B. Mouret. Novelty-based multiobjectivization. In *New Horizons in Evolutionary Robotics*, volume 341 of *Studies in Computational Intelligence*, pages 139–154. Springer Berlin / Heidelberg, 2011.
- [20] E. Naredo and L. Trujillo. Searching for novel clustering programs. In *Proceedings of the 15th annual conference on Genetic and evolutionary computation*, GECCO '13, pages 1093–1100, 2013.
- [21] G. Newman, A. Wiggins, A. Crall, E. Graham, S. Newman, and K. Crowston. The future of citizen science: emerging technologies and shifting paradigms. *Frontiers in Ecology and the Environment*, 10(6):298–304, 2012.
- [22] M. Pelikan and D. Goldberg. Escaping hierarchical traps with competent genetic algorithms. Technical Report 2001003, Illinois Genetic Algorithms Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL, 2001.
- [23] T. Reil and P. Husbands. Evolution of central pattern generators for bipedal walking in a real-time physics environment. *IEEE Transactions on Evolutionary Computation*, 6(2):159–168, April 2002.
- [24] S. Risi and K. O. Stanley. Enhancing es-hyperneat to evolve more complex regular neural networks. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, GECCO '11, pages 1539–1546, New York, NY, USA, 2011. ACM.
- [25] J. Romero and P. Machado. *The art of artificial evolution: A handbook on evolutionary art and music*. Springer-Verlag, New York, NY, 2008.
- [26] J. Secretan, N. Beato, D. B. D'Ambrosio, A. Rodriguez, A. Campbell, J. T. Folsom-Kovarik, and K. O. Stanley. Picbreeder: A case study in collaborative evolutionary exploration of design space. *Evolutionary Computation*, 19(3):373–403, 2011.
- [27] J. Secretan, N. Beato, D. B. D'Ambrosio, A. Rodriguez, A. Campbell, and K. O. Stanley. Picbreeder: Evolving pictures collaboratively online. In *CHI '08: Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 1759–1768, New York, NY, USA, 2008. ACM.
- [28] K. O. Stanley. Searching without objectives. Speech presented at SPLASH 2010, Reno, Nevada. October 21, 2010. <http://www.infoq.com/presentations/Searching-Without-Objectives>.
- [29] K. O. Stanley and R. Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10:99–127, 2002.
- [30] K. O. Stanley and R. Miikkulainen. Competitive coevolution through evolutionary complexification. *Journal of Artificial Intelligence Research*, 21:63–100, 2004.
- [31] S. R. Szumlanski, A. S. Wu, and C. E. Hughes. Conflict resolution and a framework for collaborative interactive evolution. In *Proceedings of the 21st national conference on Artificial intelligence - Volume 1*, pages 512–517. AAAI Press, 2006.
- [32] H. Takagi. Interactive evolutionary computation: Fusion of the capacities of EC optimization and human evaluation. *Proceedings of the IEEE*, 89(9):1275–1296, 2001.
- [33] M. van de Panne and A. Lamouret. Guided optimization for balanced locomotion. In *Sixth Eurographics Workshop on Animation and Simulation*, volume 95, pages 165–177. Springer-Verlag, 1995.
- [34] L. D. Whitley. Fundamental principles of deception in genetic search. *Foundations of genetic algorithms*, 1(3):221–241, 1991.
- [35] B. G. Woolley and K. O. Stanley. On the deleterious effects of a priori objectives on evolution and representation. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, GECCO '11, New York, NY, USA, 2011. ACM.

The views expressed in this paper are those of the authors and do not reflect the official policy or position of the United States Air Force, the Department of Defense, or the United States Government.