

Algorithm discovery by protein folding game players

Firas Khatib^a, Seth Cooper^b, Michael D. Tyka^a, Kefan Xu^b, Ilya Makedon^b, Zoran Popović^b, David Baker^{a,c,1}, and Foldit Players

^aDepartment of Biochemistry; ^bDepartment of Computer Science and Engineering; and ^cHoward Hughes Medical Institute, University of Washington, Box 357370, Seattle, WA 98195

Contributed by David Baker, October 5, 2011 (sent for review June 29, 2011)

Foldit is a multiplayer online game in which players collaborate and compete to create accurate protein structure models. For specific hard problems, Foldit player solutions can in some cases outperform state-of-the-art computational methods. However, very little is known about how collaborative gameplay produces these results and whether Foldit player strategies can be formalized and structured so that they can be used by computers. To determine whether high performing player strategies could be collectively codified, we augmented the Foldit gameplay mechanics with tools for players to encode their folding strategies as “recipes” and to share their recipes with other players, who are able to further modify and redistribute them. Here we describe the rapid social evolution of player-developed folding algorithms that took place in the year following the introduction of these tools. Players developed over 5,400 different recipes, both by creating new algorithms and by modifying and recombining successful recipes developed by other players. The most successful recipes rapidly spread through the Foldit player population, and two of the recipes became particularly dominant. Examination of the algorithms encoded in these two recipes revealed a striking similarity to an unpublished algorithm developed by scientists over the same period. Benchmark calculations show that the new algorithm independently discovered by scientists and by Foldit players outperforms previously published methods. Thus, online scientific game frameworks have the potential not only to solve hard scientific problems, but also to discover and formalize effective new strategies and algorithms.

citizen science | crowd-sourcing | optimization | structure prediction | strategy

Citizen science is an approach to leveraging natural human abilities for scientific purposes. Most such efforts involve visual tasks such as tagging images or locating image features (1–3). In contrast, Foldit is a multiplayer online scientific discovery game, in which players become highly skilled at creating accurate protein structure models through extended game play (4, 5). Foldit recruits online gamers to optimize the computed Rosetta energy using human spatial problem-solving skills. Players manipulate protein structures with a palette of interactive tools and manipulations. Through their interactive exploration Foldit players also utilize user-friendly versions of algorithms from the Rosetta structure prediction methodology (6) such as *wiggle* (gradient-based energy minimization) and *shake* (combinatorial side chain rotamer packing). The potential of gamers to solve more complex scientific problems was recently highlighted by the solution of a long-standing protein structure determination problem by Foldit players (7).

One of the key strengths of game-based human problem exploration is the human ability to search over the space of possible strategies and adapt those strategies to the type of problem and stage of problem solving (5). The variability of tactics and strategies stems from the individuality of each player as well as multiple methods of sharing and evolution within the game (group play, game chat), and outside of the game [wiki pages (8)]. One way to arrive at algorithmic methods underlying successful human Foldit play would be to apply machine learning techniques to the detailed logs of expert Foldit players (9). We chose instead to rely on a superior learning machine: Foldit players themselves.

As the players themselves understand their strategies better than anyone, we decided to allow them to codify their algorithms directly, rather than attempting to automatically learn approximations. We augmented standard Foldit play with the ability to create, edit, share, and rate gameplay macros, referred to as “recipes” within the Foldit game (10). In the game each player has their own “cookbook” of such recipes, from which they can invoke a variety of interactive automated strategies. Players can share recipes they write with the rest of the Foldit community or they can choose to keep their creations to themselves.

In this paper we describe the quite unexpected evolution of recipes in the year after they were released, and the striking convergence of this very short evolution on an algorithm very similar to an unpublished algorithm recently developed independently by scientific experts that improves over previous methods.

Results

In the social development environment provided by Foldit, players evolved a wide variety of recipes to codify their diverse strategies to problem solving. During the three and a half month study period (see *Materials and Methods*), 721 Foldit players ran 5,488 unique recipes 158,682 times and 568 players wrote 5,202 recipes. We studied these algorithms and found that they fell into four main categories: (i) *perturb and minimize*, (ii) *aggressive rebuilding*, (iii) *local optimize*, and (iv) *set constraints*. The first category goes beyond the deterministic minimize function provided to Foldit players, which has the disadvantage of readily being trapped in local minima, by adding in perturbations to lead the minimizer in different directions (11). The second category uses the *rebuild* tool, which performs fragment insertion with loop closure, to search different areas of conformation space; these recipes are often run for long periods of time as they are designed to rebuild entire regions of a protein rather than just refining them (Fig. S1). The third category of recipes performs local minimizations along the protein backbone in order to improve the Rosetta energy for every segment of a protein. The final category of recipes assigns constraints between beta strands or pairs of residues (*rubber bands*), or changes the secondary structure assignment to guide subsequent optimization.

Different algorithms were used with very different frequencies during the experiment. Some are designated by the authors as public and are available for use by all Foldit players, whereas others are private and available only to their creator or their Foldit team. The distribution of recipe usage among different players is shown in Fig. 1 for the 26 recipes that were run over 1,000 times. Some recipes, such as the one represented by the leftmost bar, were used many times by many different players, while others, such as the one represented by the pink bar in the

Author contributions: F.K., S.C., Z.P., and D.B. designed research; F.K., S.C., M.D.T., and F.P. performed research; F.K., S.C., M.D.T., K.X., and I.M. analyzed data; and F.K., S.C., Z.P., and D.B. wrote the paper.

The authors declare no conflict of interest.

Freely available online through the PNAS open access option.

¹To whom correspondence should be addressed. E-mail: dabaker@u.washington.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1115898108/-DCSupplemental.




Fig. 1. Foldit recipes are used with very different frequencies. Each bar represents a different Foldit recipe; the height of the bar shows the number of times the recipe was used, and the colors denote different Foldit players. Blue Fuse, at the very left, was run by many players a total of 24,273 times; the creator of Blue Fuse alone (large pink region) used it over 2,000 times. Most recipes (those represented by many colors) were run by different players; in contrast, *0 bounce 00.080* (all pink bar) is a private unshared recipe.

middle, were used by only one player who chose not to share it with other players. Not surprisingly, the frequency of usage, indicated by the height of the bars, was significantly higher for the publicly shared recipes than the private ones.

Context Dependence. Observing the breadth of created recipes, it is clear that Foldit players use these algorithms to augment rather than to substitute for human strategizing. Players in essence perform a problem-informed search over the space of strategies, and use recipes to encode specific lower-level strategies. Different algorithms are employed at different stages in gameplay. Fig. 2 shows the relative frequency of recipe use in the opening, midgame, and endgame. Expert players exhibit different patterns of recipe use than the player population as a whole (compare Fig. 2A and B). The top Foldit players use recipes *tlaloc Contract 3.00* and *Aotearoas_Romance* in the endgame, while most players use them almost equally at every stage in a puzzle (Fig. S2A and B). Most Foldit players run *after-rebuild test* exclusively in the opening, but the top players often use it in the midgame as well. Local optimize recipes are heavily used in the endgame by top Foldit players (Fig. 2B and Fig. S2F).

Human strategic judgment plays an important role in choosing when and how to employ different recipes. Most recipes heavily rely on the interactive aspects of the game, and are used less as fully automated tools and more as power tools that serve as an extension of the player's strategy. We have not found any recipes that generate top models without human intervention; instead,

Foldit players employ recipes to perform specific tasks at different stages of the folding process. For example, the local optimize recipes that walk along the protein backbone performing local minimizations are useless on the initial state of a Foldit puzzle that starts in an extended chain configuration, because a successful prediction will no longer have that backbone in an extended conformation and will require a new round of local optimization. Many of the aggressive rebuilding recipes require specifying which region of the protein to rebuild, as it is rarely useful to entirely rebuild a protein chain from beginning to end. These recipes are launched by players once they have converged to a low-energy solution and want to search nearby regions of conformational space before local optimization (Fig. S2C and D); they are most useful in the midgame and often run for long periods without any human intervention. By contrast, all recipes in the set constraints category are designed to be run before launching other recipes or manually manipulating the protein (Fig. S2G and H).

Overall, the context dependent use of different recipes is key to successful gameplay, but also makes it difficult to quantitatively evaluate the effectiveness of individual recipes because they are optimally used on quite different input protein states. This context dependence also complicates efforts to incorporate the rich new algorithms developed by Foldit players back into Rosetta or other automated methods.

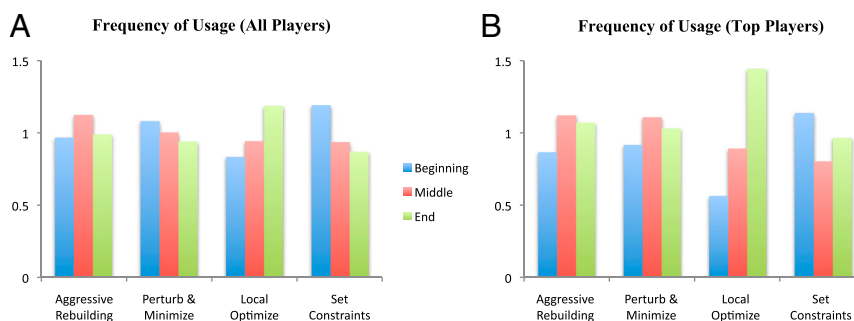


Fig. 2. Foldit players employ different recipes at different stages of gameplay. The relative frequency of usage for each recipe category is shown for three different stages of gameplay (blue, first third of gameplay; red, second third; green, final third) for all players (A) and for top ranked players (B). All players use local optimize recipes more frequently as gameplay progresses (A), and this trend is even stronger among top players (B). The usage frequency for each individual recipe is shown in detail in Fig. S2. The actual context dependence of the use of different recipes is likely greater than that shown in the figure; the division based on elapsed game time used here is relatively crude because different players in different puzzles may spend very different amounts of time on the opening, middle game, and end game.




Fig. 3. Rapid proliferation of Foldit recipes between June 2009 and August 2010. Each different shade of gray represents an individual recipe with the size of the gray bar denoting how many times that recipe was run that particular week. The most heavily used script recipe is Blue Fuse, shown in yellow. The most popular GUI recipe is Quake, shown in red, which has been used consistently by Foldit players weekly since its creation.

Recipe Evolution. In addition to providing tools for recipe creation, we enabled sharing, evaluation, and refinement of recipes. These tools spurred a rapid evolution of recipe use in the Foldit player population from June 2009 [when the graphical user interface (GUI) recipes were first introduced; see *Materials and Methods*] through the end of August 2010. Fig. 3 shows the number of distinct recipes run every week across this time period, with different shades of gray representing different recipes. Two recipes were particularly heavily used: *Quake* (in red) and *Blue Fuse* (in yellow). The essential features of the algorithms contained within these two most popular recipes are similar. *Quake* is a GUI recipe that repeatedly packs the side chains and minimizes the side chain and backbone torsion angles, increasing and decreasing the strength of an applied set of rubber band constraints to promote annealing of the structure. The *Blue Fuse* algorithm (a script recipe; see *Materials and Methods*) is conceptually similar to but simpler than *Quake*: rather than varying the strength of artificial constraint functions, the strength of atom-atom repulsive interactions is alternately increased and decreased in repeated cycles of side chain packing and complete torsion angle minimization (Fig. S3). In both cases, an initially poorly packed structure is alternately compressed (by minimization with strong constraints or weak repulsive interactions) and then relaxed (by minimization with weak constraints or strong repulsive interactions); this alternation appears to be effective in locating low energy, well packed conformations.

Foldit algorithms evolve in a social fashion, with popular recipes copied and then elaborated upon. An evolutionary tree depicting the evolution of the *Blue Fuse* algorithm family is shown in Fig. 4; size corresponds to the logarithm of the number of recipe uses and color represents the recipe author. Popular recipes spawn large numbers of descendants, and there are multiple independent lineages each spanning many generations.

While the evolutionary mechanisms by which new algorithms are created (copying and variation) are clear, the mechanisms through which individual recipes spread through the population are less so. The popularity of *Quake* and *Blue Fuse* could result from individual player experimentation, player communication through chat, or player communication within teams. Interviews with Foldit players suggest that the primary mechanism by which successful algorithms take over the population is word of mouth. Players test and then recommend new recipes to others and rate

recipes on the Foldit website, making decisions in part based on the reputation of the recipe author.

Similarity of Independently Developed Player and Scientist Algorithms. During this same time period, researchers in the Baker group were attempting to improve the core optimization methods within the Rosetta structure prediction and design program. A breakthrough was made with the *Fast Relax* algorithm, which achieves more effective energy optimization in less time than previous Rosetta methods. Benchmarks have shown *Fast Relax* to be considerably more efficient than an older algorithm, *Classic Relax* (12) used since 2002 (Fig. 5, compare green to red), and it is now used in all Rosetta de novo and homology modeling methods as well as enzyme design protocols.

The key innovation in the *Fast Relax* algorithm compared to *Classic Relax* and earlier methods is alternately increasing and decreasing the strength of the repulsive interactions. As diagrammed in Fig. 6, *Fast Relax* is comprised of interlaced cycles of full combinatorial repacking of the side chains of the protein (*repack*) and gradient-based minimization of the backbone and side chain degrees of freedom (*minimize*); within each round the repulsive weight is increased while from the end of one round to the start of the next (indicated by the arrow) the repulsive weight is decreased 50-fold. Typically 5–15 rounds of this repulsive




Fig. 4. Social evolution of Foldit recipes. Each circle represents a different Foldit recipe, each color denotes a different author, and the size is the logarithm of the number of times a recipe was used. Arrows represent the process of copying and subsequent modification: the recipe at the tip of an arrow was created from the recipe at the base of the arrow. *Acid Tweaker v0.5* is the parent of all of the recipes shown here with *Blue Fuse* its most popular offspring. The popularity of *Blue Fuse* led many Foldit players to create their own modified versions of the recipe.

