Remote optimization of an ultracold atoms experiment by experts and citizen scientists

Robert Heck[a], Oana Vuculescu[a,b], Jens Jakob Sørensen[a], Jonathan Zoller[c,d], Morten G. Andreasen[a], Mark G. Bason[e], Poul Ejlertsen[a], Ottó Elíasson[a], Pinja Haikka[a], Jens S. Laustsen[a], Lærke L. Nielsen[a], Andrew Mao[a], Romain Müller[a], Mario Napolitano[a], Mads K. Pedersen[a], Aske R. Thorsen[a], Carsten Bergenholtz[a,b], Tommaso Calarco[c,d,1], Simone Montangero[c,d,f,g], and Jacob F. Sherson[a,2]

[a]ScienceAtHome, Department of Physics and Astronomy, Aarhus University, 8000 Aarhus C, Denmark; [b]Department of Management, School of Business and Social Sciences, Aarhus University, 8000 Aarhus C, Denmark; [c]Institute for Complex Quantum Systems, Ulm University, 89081 Ulm, Germany; [d]Center for Integrated Quantum Science and Technology, Ulm University, 89081 Ulm, Germany; [e]Department of Physics and Astronomy, University of Sussex, Falmer, Brighton BN1 9QH, United Kingdom; [f]Theoretische Physik, Universität des Saarlandes, 66123 Saarbrücken, Germany; and [g]Dipartimento di Fisica e Astronomia "G. Galilei," Università degli Studi di Padova, 35131 Padova, Italy

We introduce a remote interface to control and optimize the experimental production of Bose–Einstein condensates (BECs) and find improved solutions using two distinct implementations. First, a team of theoreticians used a remote version of their dressed chopped random basis optimization algorithm (RedCRAB), and second, a gamified interface allowed 600 citizen scientists from around the world to participate in real-time optimization. Quantitative studies of player search behavior demonstrated that they collectively engage in a combination of local and global searches. This form of multiagent adaptive search prevents premature convergence by the explorative behavior of low-performing players while high-performing players locally refine their solutions. In addition, many successful citizen science games have relied on a problem representation that directly engaged the visual or experiential intuition of the players. Here we demonstrate that citizen scientists can also be successful in an entirely abstract problem visualization. This is encouraging because a much wider range of challenges could potentially be opened to gamification in the future.

citizen science | optimal control | ultracold atoms | human problem solving | closed-loop optimization

In modern scientific research, high-tech applications such as quantum computation (1) require exquisite levels of control while taking into account increasingly complex environmental interactions (2). This necessitates continual development of optimization methodologies. The "fitness landscape" (3) spanned by the possible controls and their corresponding solution quality forms a unifying mathematical framework for search problems in both natural (4–9) and social sciences (10, 11). Generally, search in the landscape can be approached with local or global optimization methods. Local solvers are efficient and analogous to greedy hill climbers. In nonconvex landscapes, however, they might get trapped locally and cannot reach the global optimum. The global methods attempt to escape these traps by taking larger, stochastic steps. That, however, typically increases the runtime dramatically compared with that of the local solvers. Achieving the proper balance between local and global methods is often referred to as the exploration/exploitation trade-off in both machine learning (ML) (12) and social sciences (13).

Much effort in computer science is therefore focused on developing algorithms that exploit the topology of the landscape to adapt search strategies and make better-informed jumps (6, 14). ML algorithms have achieved success across numerous domains. However, among researchers pursuing truly domain-general artificial intelligence, there is a growing call to rely on insights from human behavior and psychology (15, 16). Thus, emphasis is currently shifting toward the development of human–machine hybrid intelligence (17, 18).

At the same time quantum technology is starting to step out of university laboratories into the corporate world. For the realization of real-world applications, not only must hardware be improved but also proper interfaces and software need to be developed. Examples of such interfaces are the *IBM Quantum Experience* (19) and *Quantum in the Cloud* (20), both of which give access to their quantum computing facilities and have ushered in an era in which theoreticians can experimentally test and develop their error correction models and new algorithms directly (21). The optimal development of such interfaces, allowing the smooth transformation of human intuition or experience-based insights (heuristics) into algorithmic advances, necessitates understanding and explicit formulation of the search strategies introduced by the human expert.

The emerging field of citizen science provides a promising way to investigate and harness the unique problem-solving abilities humans possess (22). In recent years the creativity and intuition of nonexperts using gamified interfaces have enabled scientific contributions across different fields such as quantum physics (8,

Significance

The emerging field of gamified citizen science continually probes the fault line between human and artificial intelligence. A better understanding of citizen scientists' search strategies may lead to cognitive insights and provide inspiration for algorithmic improvements. Our project remotely engages both the general public and experts in the real-time optimization of an experimental laboratory setting. In this citizen science project the game and data acquisition are designed as a social science experiment aimed at extracting the collective search behavior of the players. A further understanding of these human skills will be a crucial challenge in the coming years, as hybrid intelligence solutions are pursued in corporate and research environments.

PHYSICS

SOCIAL SCIENCES

23), astrophysics (24), and computational biology (25–27). Here, citizen scientists often seem to jump across very rugged landscapes and solve nonconvex optimization problems efficiently using search methodologies that are difficult to quantify and encode in a computer algorithm.

The central purpose of this paper is to combine remote experimentation and citizen science with the aim of studying quantitatively how humans search while navigating the complex control landscape of Bose–Einstein condensate (BEC) production (Fig. 1A). Before presenting the results of this *Alice Challenge* (https://alice.scienceathome.org) we characterize the corresponding landscape using a nontrivial heuristic that we derive from analysis of our previous citizen science work (8). In addition, a team of experts using a state-of-the-art optimization algorithm explores the landscape of BEC production. This allows for a comparison of algorithmic and human search strategies.

Initial Landscape Investigations

On www.scienceathome.org, our online citizen science platform, more than 250,000 people have so far contributed to the search for solutions to fast, 1D single-atom transport in the computer game *Quantum Moves* (8). The solutions are investigated in terms of the fitness landscape: a spatial representation of a function $J(\vec{u}) \in \mathbf{R}$ describing the quality of a solution given by the control variables $\vec{u} \in \mathcal{S}$, where $\mathcal{S}$ is the set of possible solutions. (It is important to note that the topology of the fitness landscape critically depends on the choice of parameterization or representation of $\vec{u} \in \mathcal{S}$.) Clustering analyses of player solutions were

found to bunch into distinct groups with clear underlying physical interpretations. We refer to such a group of related solutions as a strategy.

A complete mapping of the fitness landscape topology would involve extensive numerical work such as the reconstruction of the full Hessian (28) or, as an approximation, random sampling of solutions around these strategies combined with methods like principal component analysis. Due to the high dimensionality of the problem, these approaches are difficult to scale. Instead, we introduce here the heuristics to explore the space between the identified strategies along the high-dimensional vector connecting the two as the "strategy-connecting heuristic" (SCH). Perhaps surprisingly, we identify a narrow path of monotonically increasing high-fidelity solutions between the two strategies, which we denote a "bridge." This demonstrates that for this problem a continuum of mixed-strategy solutions with no clear physical interpretation can be mapped out if all hundreds of control variables are changed synchronously in the appropriate way. (Note that if the identified bridge is narrow in high-dimensional space, the random sampling methods will likely fail to identify it.) Moreover, due to its monotonicity, a correctly set up greedy algorithm would be able to realize (at least partially) global exploration of the landscape along such a bridge (see *SI Appendix*, section A for details).

This leads us to ask whether other established strategies in physics are truly distinct or whether they are simply labels we attach to different points in a continuum of possible solutions due to our inability to probe the entire solution space. In the latter case, this, coupled with the human desire to create identifiable patterns, might cause us to terminate our search before discovering the true global optimum. This premature termination of the search nicely illustrates the "stopping problem" (29, 30) considered in both computer and social sciences: determining criteria to stop searching when the best solution found so far is deemed "good enough."

We now apply this methodology to the high-dimensional problem of experimental BEC production (31). In our case an increased BEC atom number, N_{BEC}, will provide significantly improved initial conditions for subsequent quantum simulation experiments using optical lattices (32). Although extensive optimization has been applied to the BEC creation problem over the past decade by using global closed-loop optimization strategies using genetic algorithms (33–36), little effort has been devoted to the characterization of the underlying landscape topology and thereby the fundamental difficulty level of the optimization problem. In the global landscape spanned by all possible controls, it is thus unknown whether there is a convex optimization landscape with a single optimal strategy for BEC creation as opposed to individual distinct locally optimal strategies of varying quality (illustrated in Fig. 2A) or a plethora of (possibly) connected solutions (Fig. 2B). Recent experiments (37) indicated a convex and thereby simple underlying landscape; however, that study did not explicitly optimize N_{BEC} and operated within a severely restricted subspace.

In our experiment (38), we capture ^{87}Rb atoms in a trap made of two orthogonal, focused 1,064-nm laser beams and a superimposed quadrupolar magnetic field which creates a magnetic field gradient at the position of the atoms and thereby forms a magnetic trap (illustrated in Fig. 1B). We evaporatively cool the atoms past the phase transition to a BEC by lowering the intensity of the laser beams as well as the magnetic field gradient. Then, the traps are turned off, and the atoms are imaged with resonant light. Image analysis yields the total and condensed atom numbers N_{tot} and N_{BEC}.

This setting allows for evaporative cooling in two widely used trap configurations. First, making use of only the laser beams, a purely optical trap can be created; this is commonly known as a "crossed dipole trap" (CDT) (39). Second, a single laser
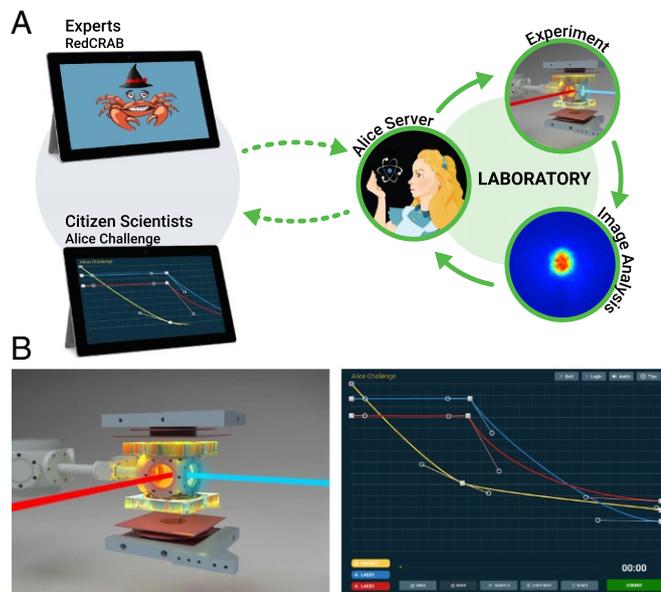
Fig. 1. (A) Real-time remote scheme for connecting experts and citizen scientists with the laboratory. The respective remote clients (RedCRAB for the experts, *Alice Challenge* game client for the citizen scientists) send experimental parameters through an online cloud interface. These parameters are turned into experimental sequences and executed by the Alice control program. The number of atoms in the BEC (N_{BEC}) serves as a fitness value and is extracted from images of the atom cloud taken at the end of each sequence. The Alice control program closes the loop by sending the resulting N_{BEC} back to the remote clients through the same cloud interface. (*B, Left*) Illustration of the experimental setup. The RedCRAB algorithm and the *Alice Challenge* players can control the magnetic field gradient depicted by the yellow shaded coils and the intensity of the two dipole beams drawn in red and blue. (*B, Right*) Screenshot of the *Alice Challenge* (https://alice.scienceathome.org). The game client features a spline editor for creating and shaping the experimental ramps.

Fig. 2. (A) Illustration of the apparent global landscape topology for BEC production after performing 1D parameter scans. It seems to contain distinct local optima. However, as B illustrates, connecting bridges were found both between some 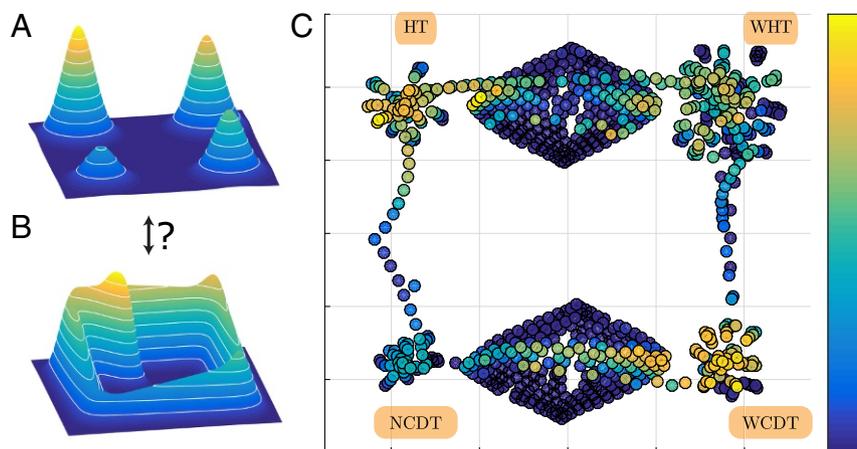of the conventional strategies and to nontrivial high-yield solutions in the high-dimensional search space. (C) A 2D T-distributed stochastic neighbor embedding (t-SNE) (42) representation of the landscape showing the variety of different trap configurations that are accessible in our experiment. (Note that the displayed data points stem from a different set of measurements, where high N_{BEC} are underestimated due to saturation effects in the imaging, which was alleviated for the main experiments of this paper. Therefore, the labeling of the color scale was omitted.) The plot contains data of the four main configurations which were scanned and optimized by 1D and 2D parameter scans. For more details, see main text.

beam can be combined with a weak magnetic gradient to form a "hybrid trap" (40). In both cases, the traps are initially loaded from a pure tight magnetic trap (see *SI Appendix*, section B for details). Conventionally, two types of geometrically differing loading schemes are pursued: loading into a large-volume trap that exhibits a nearly spatially mode-matched type of loading from the initial magnetic trap into the final trap configuration (31) or loading into a small-volume trap with only a small spatial overlap with the initial trap. The latter leads to a "dimple" type of loading (41) in which a smaller but colder atom cloud is produced. We can directly control the effective volume of the trap by translating the focus position of one of the dipole trap beams. This inspired us to identify four initial "conventional" trap configurations (BEC creation strategies): a small volume, "narrow" crossed dipole trap (NCDT); a large volume, "wide" counterpart (WCDT); and similarly a hybrid trap (HT) and a wide hybrid trap (WHT).

We first optimize the system by applying a simple standard experimental approach (*SI Appendix*, section B): Starting from the set of control variables associated with a known strategy, we iteratively perform 1D scans of single variables until a specified level of convergence is reached. The 1D scans yield four distinct strategies, and we find that the HT is the best-performing strategy. This hints at the landscape topology sketched in Fig. 2A. Further systematic studies would then proceed to scans of two or more parameters simultaneously. However, allowing for scans of combined parameters enables prohibitively many different 2D parameter scans. Therefore, we proceed by applying the SCH derived from the *Quantum Moves* investigations.

Both the low-yield NCDT configuration and the WCDT are types of CDTs but with different effective volumes and thus represent ideal candidates for first exploration. However, a simple linear interpolation of all of the available parameters between the NCDT and the WCDT fails to locate a bridge. Treating the effective trap volume as an independent second parameter realizes an extended 2D interpolation and leads to the emergence of a bridge (*SI Appendix*, section B). In this case, the change of the trap depth induced by changing the trap volume has to be counterbalanced by a quadratic increase of the laser intensities involved. Thus, changing to a different representation (i.e., a particular combination of parameters) efficiently encapsulating the underlying physics yields a bridge and disproves the initial

assumption that each strategy was distinct as illustrated in Fig. 2A. To illustrate these data, we created a dimensionality-reduced visualization (42) of the parameter scans (Fig. 2C). The four initial strategies are represented by the four clusters in the corners. The data points forming the bridge between NCDT and WCDT lie in the diamond shape at the bottom. A few other 1D and 2D interpolations between other pairs of strategies are shown, but none form a bridge. In an attempt to locate a bridge between NCDT and HT, extended 3D scans are performed (*SI Appendix*, section B, not shown in Fig. 2C). These scans identify a previously undiscovered optimum away from the four initially defined experience-based trap configurations. This demonstrates that the HT, our initial candidate for a global optimum, is not even a local optimum when appropriate parameter sets are investigated. One is therefore inclined to view the topology of the landscape as closer to what is depicted in Fig. 2B, where the four conventional strategies are now connected with bridges and at least one other higher-yield solution exists in the full landscape.

Having established that the global optimum must be found by using unconventional strategies, we switch to the main topic of the paper: a remotely controlled strategy using closed-loop optimization performed by experts implementing a state-of-the-art optimization algorithm and citizen scientists operating through the *Alice Challenge* game interface. As detailed below, our particular implementation allows for a quantitative assessment of the citizen scientist search behavior, but only a qualitative assessment of their absolute performance. As a result, the search behavior of the two parties can only be compared qualitatively.

Expert Optimization

As mentioned above, closed-loop optimization has been explored extensively for BEC creation using random, global methods (33–37) and is also routinely used to tailor radio-frequency fields to control nuclear spins or shape ultrashort laser pulses to influence molecular dynamics (ref. 43 and references therein). In our remote-expert collaboration, we use the dressed chopped random basis (dCRAB) (44) algorithm, which is a basis-adaptive variant of the CRAB algorithm (45). The main idea of both algorithms is to perform local landscape explorations, using control fields consisting of a truncated expansion in a suitable random basis. This approach makes optimization tractable by limiting the number of optimization parameters and has, at the same

time, the advantage of obtaining information of the underlying landscape topology. It has been shown that the unconstrained dCRAB algorithm converges to the global maximum of an underlying convex landscape with probability one (44). That is, despite working in a truncated space, iterative random function basis changes allow the exploration of enough different directions in the functional space to escape traps induced by the reduced explored dimensionality (46, 47). CRAB was introduced for the theoretical optimization of complex systems in which traditional optimal control theory could not be applied (48). In a closed loop, CRAB was applied to optimize the superfluid to Mott insulator transition (49). Very recently, dCRAB (44) was used to realize autonomous calibration of single-spin qubits (50) and to optimize atomic beam-splitter sequences (51).

Unlike the closed-loop optimization performed in the past on other experiments in which the optimization libraries were installed directly in the laboratory control software, our Alice remote interface allows one to implement the dCRAB algorithm remotely [remote dCRAB (RedCRAB); Fig. 1A]. This allows the team of optimization experts to easily adjust algorithmic parameters based on the quality of preliminary optimization runs and thus exploit its full potential. Moreover, algorithmic improvements to the control suite can easily be transferred to future experiments. As in the case of the *IBM Quantum Experience*, we believe that this will enhance the efficiency of experimentation and lower the barrier for even wider adoption of automated optimization in different quantum science and technology aspects,

from fundamental science experiments to technological and industrial developments.

After initialization, the RedCRAB works unsupervised and controls the intensity ramps of the two dipole trap beams, the ramp duration T_{ramp} and a single parameter that represents the value of the magnetic field gradient during evaporation (*SI Appendix*, section C). A set of parameters is sent through the remote interface to the laboratory and realized in the experiment. The corresponding yield N_{BEC} is fed back through the same interface, which closes the loop. To prevent the algorithm from becoming misled due to noise-induced outliers, each RedCRAB iteration step is the result of an adaptive averaging scheme which repeats a set of parameters up to four times, depending on N_{BEC} reaching certain threshold values.

As illustrated in Fig. 3A, we achieve a new maximal solution in about 100 iterations that exceeds the performance of the HT by more than 10%. This solution is nontrivial in the sense that it can be seen as a type of CDT combined with the magnetic field gradient of the HT. The beam intensities are adjusted to lead to relatively similar trap depths as in the HT. However, especially in the beginning of the evaporation process, the trap is relaxed much faster, leading to an overall shorter ramp. By applying the SCH, a bridge connecting the HT to this solution could be identified (*SI Appendix*, section C). This illustrates that the RedCRAB algorithm is highly effective at both locating nontrivial optimal solutions and providing topological information of the underlying landscape. It substantiates the appearance of
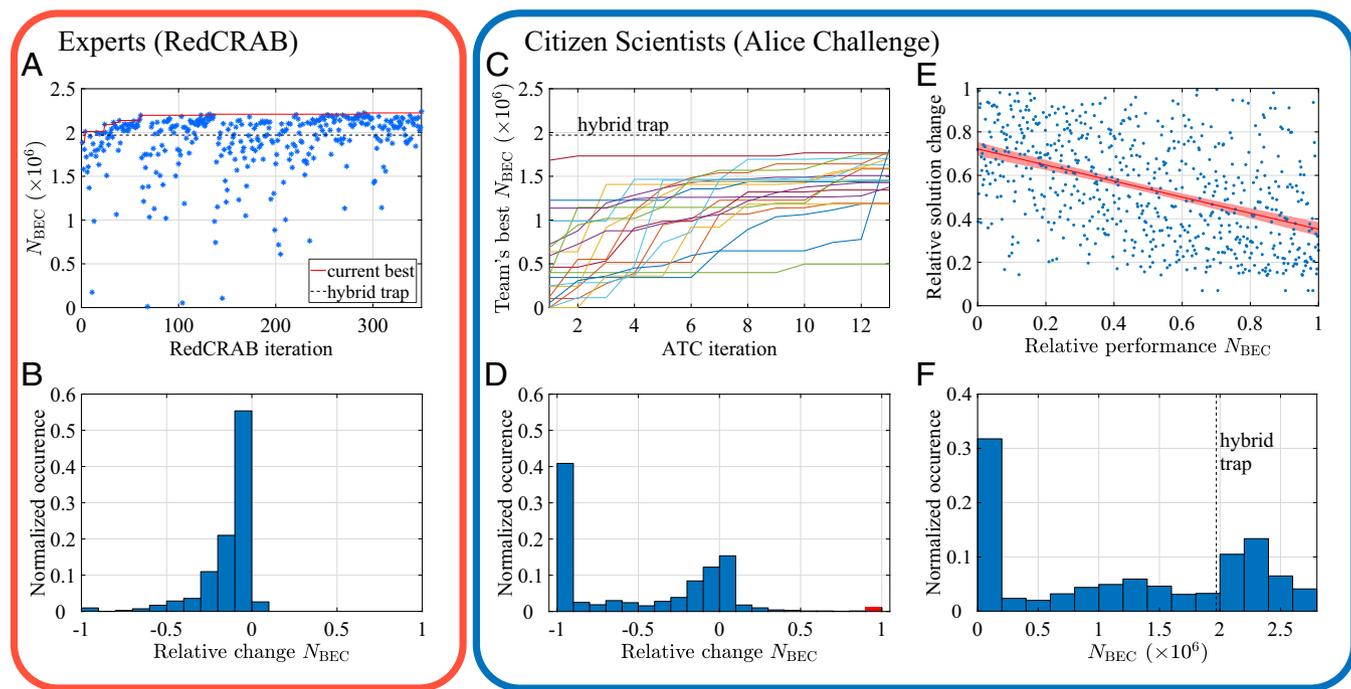

Fig. 3. (*A* and *B*) Experts' optimization with RedCRAB. (*A*) Single unsupervised optimization run. By applying an adaptive averaging scheme, the mean N_{BEC} is plotted in blue as a function of RedCRAB iteration steps (see main text for details). The red solid line denotes the current best N_{BEC}. Compared to the level of the previously best HT configuration (black dashed line), N_{BEC} was improved by 10%. (*B*) Histogram of relative N_{BEC} changes compared with the current best solution for the RedCRAB optimization. (*C*–*F*) Citizen scientists' optimization in the ATC (*C*–*E*) and the ASC (*F*). (*C*) Round-based performance in the ATC. The lines show the cumulated best achieved N_{BEC} for teams with three or more active players as a function of ATC iteration steps (see main text for definition). Although human players had only a very limited number of tries (13 iterations), they still achieve relatively good optimization scores. Overall, all teams but one achieve N_{BEC} above 1×10^6. (*D*) Histogram of N_{BEC} changes relative to the current best solution for the ATC. In contrast to the experts' RedCRAB searches (compare *B*), humans engage in many search attempts that lead to poor N_{BEC}. The red bar denotes all solutions which showed a relative change in $N_{BEC} > +0.9$. (*E*) The players' adaptive search behavior as a function of the relative performance with respect to the team's best N_{BEC}. A linear regression with a 95% confidence bound is shown in red and yields a correlation of −0.37(4). The distance measure captures the difference between the player's current and own previous solution. The N_{BEC} measure captures a player's performance relative to the team's best. Both measures are normalized across all ATC iterations and teams. (*F*) Histogram for the achieved N_{BEC} for all submitted solutions in the ASC. More than 73% of the submitted solutions successfully yielded a BEC.

a complex but much more connected landscape than initially anticipated.

Citizen Science Optimization

In our second approach to remote optimization, we involve citizen scientists by using a gamified remote user interface that we call the *Alice Challenge*. We face the difficulty of turning the adjustment of laser and magnetic field ramps into an interactive, engaging game. Therefore, we developed a client using the cross-platform engine Unity and promoted it through our online community www.scienceathome.org. The protocol was approved by a Human Subjects Committee at Aarhus University and all participants provided informed consent before participation. As depicted in Fig. 1*B*, the ramps are represented by three colored spline curves and are modified by adjustable control points. The total ramp duration, T_{ramp}, is fixed. After manipulating the splines, the player submits the solution, which is then executed on the experiment in the laboratory (Fig. 1*A*). The obtained N_{BEC} provides performance feedback to the players and is used to rank players in a high score list. Players have the option to see, copy, and adapt everyone's previous solutions. This setup generates a collective search setting where players emulate a multiagent genetic search algorithm. Note, that two different game modes were designed: the *Alice Team Challenge* (ATC) and the *Alice Swarm Challenge* (ASC), both of which are introduced later.

Citizen scientists have shown that they can solve highly complex natural science challenges (8, 25–27). However, data from previous projects suffer from the fact that they merely showed that humans solve the challenges but did not answer how a collective is able to balance a local vs. a global search while solving these complex problems. Social science studies in controlled laboratory settings have shown that individuals adapt their search based on performance feedback (52). Specifically, if performance is improving, humans tend to make smaller changes (i.e., local search), while if performance is worsening, humans tend to make larger changes (i.e., search with a global component). Therefore, experimental evidence suggests that human search strategies are neither purely local nor global (52, 53). Furthermore, studies have also established the importance of social learning and how humans tend to copy the best or most frequent solutions (54–56), which facilitates an improved collective search performance. However, these laboratory-based studies have been constrained by the low dimensionality and artificial nature of the tasks to be solved. This raises concerns with respect to the external validity of the results: Are these general human problem-solving patterns or are they merely behaviors elicited by the artificial task environment? Finally, previous citizen science results were based on intuitive game interfaces such as the close resemblance to sloshing water in *Quantum Moves*. In contrast, the *Alice Challenge* is not based on any obvious intuition. It is therefore interesting to investigate whether and how citizen scientists are able to efficiently balance local and global search when facing a real-world, rugged, nonintuitive landscape.

To address this question, we created a controlled setting: the ATC. Unfortunately, due to the structure of the remote participation, sufficient data could not be gathered to quantitatively study both the initial search behavior and the convergence properties of the human players (*SI Appendix*, section D). Our previous work (8) demonstrated that the human contribution lay in roughly exploring the landscape and providing promising seeds for the subsequent, highly efficient numerical optimization. We therefore chose a design focusing on the initial explorative search of the players, knowing that this would preclude any firm statements about the absolute performance of the players in terms of final atom number. Concretely, teams of five players each were formed, with every team member being allowed one submission in each of the 13 rounds (ATC iterations). After the five solutions from the active team were collected, they were run on the

experiment and results provided to the players. All teams were provided initially with the same five low-performing solutions and T_{ramp} was fixed to about 4 s. Each ATC iteration lasted about 3 min, and a 13-ATC iteration game lasted ∼1 h. This was chosen as the best balance between keeping teams motivated over the whole experiment with minimal dropouts and gathering enough data for analysis.

As illustrated in Fig. 3*C* players showcase substantial, initial improvements across all game setups, even though the system and its response were completely unknown to them beforehand. This demonstrates that humans can indeed effectively search complex, nonintuitive solution spaces (*SI Appendix*, Fig. S7). To make sense of how citizen scientists do this, we asked some of the top players how they perceived their own gameplay. One of them explained that he tried to draw on his previous experience as a microwave engineer by applying a black-box optimization approach. Because he did not need a detailed understanding of the underlying principles of the search space, this suggests that humans might have domain generic search heuristics they rely on when solving such high-dimensional problems.

Differences in the player setup and the accessible controls preclude a direct comparison of the absolute performance of the RedCRAB and the citizen scientists. However, Fig. 3 *B* and *D* demonstrates the distribution of the relative changes in N_{BEC} and clearly reveals how fundamentally different the respective search behaviors are. The local nature of the RedCRAB algorithm leads to incremental changes in either positive or negative directions. When optimizing the system, 80% of RedCRAB's guesses correspond to changes of 20% or less in their current optimal value of N_{BEC}. In contrast, humans engage in many search attempts that lead to poor N_{BEC}. In that case, 60% of the solutions yield N_{BEC} which differs by more than 20% compared with the current best.

To investigate this quantitatively, we further analyzed observations from players in the ATC ($N = 110$). Supporting previous laboratory studies (52, 53), our results show that players engaged in adaptive search; i.e., if one identified a good solution compared with the other solutions visible to the player, the player tended to make small adjustments in the next attempt. In contrast, if the solution found by the player was far behind the best solution, the player tended to engage in more substantial adjustments to his or her current solution (Fig. 3*E* and *SI Appendix*, section I). Advancing previous studies, we were also able to identify that players engaged in the same type of adaptive search, when engaged in social learning, i.e., when copying a solution from someone else in the team and subsequently manipulating it before submitting their modified solution (*SI Appendix*, section I). The nature of adaptive search leads to a heterogeneous human search "algorithm" that combines local search with a global component. This search is prevented from stopping too early, since poorly performing individuals search more distantly, breaking free from exploitation boundaries, while individuals that are near the top perform exploitative, local search. This out of the laboratory quantitative characterization of citizen science search behavior represents the main result of this paper.

Finally, to explore the absolute performance of the citizen scientists, we created with the ASC an open "swarm" version within the *Alice Challenge*. The client was free to download for anyone and the number of submitted solutions was unrestricted. Participants could copy and modify other solutions freely. As this setting was uncontrolled, general statements about the search behavior are not possible. In the ASC, we had roughly 500 citizen scientists spanning many countries and levels of education. The submitted solutions were queued, and an estimated process time was displayed. In this way, players could join, submit one or a set of solutions, and come back at a later time to review the results. The game was open for participation for 1 wk, 24 h/d,

with brief interruptions to resolve experimental problems. As an additional challenge, the game was restarted two to three times per day while changing T_{ramp} as well as the suggested start solutions. In a total of 15 sessions, we covered a range of ramp times from 1.75 s to 8 s and a total of 7,577 solutions were submitted. Without the restriction of game rounds, players were able to improve the solutions further. Fig. 3F shows the distribution of the attained N_{BEC} across all sessions. For short ramp durations it became increasingly difficult to produce BECs with high N_{BEC} (*SI Appendix*, section F). Nonetheless, the players could adapt to these changing conditions and produce optimized solutions.

The largest BEC was found for $T_{ramp} = 4$ s and contained about $2.8 \cdot 10^6$ atoms, which set a new record in our experiment. The solutions found by the players were qualitatively different from those found by numerical optimization. Where the RedCRAB algorithm was limited by having control only over the evaporation process and being able to apply only a single specific value for the magnetic field gradient, the players had full control over all ramps throughout the whole sequence of loading and evaporation. This was used to create a smoother transition from loading to evaporation. The magnetic field gradient during evaporation was initially kept at a constant value but relaxed toward the end (*SI Appendix*, section F).

In conclusion, with the advent of machine-learning methods, the focus on cutting-edge algorithmic design is shifting to the subtle interplay between exploration and exploitation elements at which human inspiration is acknowledged to be important. In our opinion, the growing emphasis on human–computer interaction, as computer algorithms are integrated more deeply into the scientific research methodology, will challenge the clear divide between social and natural science. We see our work as an example of the growing usefulness of bridging this divide. Con-cretely, we have introduced an interface that allowed for remote closed-loop optimization of a BEC experiment, both with citizen scientists interacting through a gamified remote client and by connecting to numerical optimization experts. Both yielded solutions with improved performance compared with the previous best strategies. The obtained solutions were qualitatively different from those of well-known strategies conventionally pursued in the field. This hints at a possible continuum of efficient strategies for BEC creation. Although quantitative studies of player optimization performance were precluded by the design, it is striking that with regard to overall performance, the players seemed to be able to compete with the RedCRAB algorithm and also exhibited the ability to adapt to changes in the constraints (duration) and conditions (experimental drifts). The controlled design of the ATC yielded quantitative insight into the collective adaptive search performed by the players. This points toward a future in which the massive amounts of data on human problem solving from online citizen science games could be used as a resource for investigations of many ambitious questions in social science.

1. Ladd TD, et al. (2010) Quantum computers. *Nature* 464:45–53.
2. Devitt SJ, Munro WJ, Nemoto K (2013) Quantum error correction for beginners. *Rep Prog Phys* 76:076001.
3. Wright S (1932) The roles of mutation, inbreeding, crossbreeding, and selection in evolution. *Proceedings of the Sixth International Congress on Genetics*, ed Jones DF (Brooklyn Botanic Garden, Menasha, WI), Vol 1, pp 355–366.
4. Kryazhimskiy S, Tkacik G, Plotkin JB (2009) The dynamics of adaptation on correlated fitness landscapes. *Proc Natl Acad Sci USA* 106:18638–18643.
5. Rabitz HA, Hsieh M, Rosenthal CM (2004) Quantum optimally controlled transition landscapes. *Science* 303:1998–2001.
6. Malan KM, Engelbrecht AP (2014) Fitness landscape analysis for metaheuristic performance prediction. *Recent Advances in the Theory and Application of Fitness Landscapes: Emergence, Complexity and Computation*, eds Richter H, Engelbrecht A (Springer, Berlin), Vol 6, pp 103–132.
7. Rabitz H, Wu RB, Ho TS, Tibbetts KM, Feng X (2014) Fundamental principles of control landscapes with applications to quantum mechanics, chemistry and evolution. *Recent Advances in the Theory and Application of Fitness Landscapes: Emergence, Complexity and Computation*, eds Richter H, Engelbrecht A (Springer, Berlin), Vol 6, pp 33–70.
8. Sørensen JJWH, et al. (2016) Exploring the quantum speed limit with computer games. *Nature* 532:210–213.
9. Palittapongarnpim P, Wittek P, Zahedinejad E, Vedaie S, Sanders BC (2017) Learning in quantum control: High-dimensional global optimization for noisy quantum dynamics. *Neurocomputing* 268:116–126.
10. Levinthal DA (1997) Adaptation on rugged landscapes. *Manage Sci* 43:934–950.
11. Acerbi A, Tennie C, Mesoudi A (2016) Social learning solves the problem of narrow-peaked search landscapes: Experimental evidence in humans. *R Soc Open Sci* 3:160215.
12. Dueck G (1993) New optimization heuristics. *J Comput Phys* 104:86–92.
13. March JG (1991) Exploration and exploitation in organizational learning. *Organ Sci* 2:71–87.
14. Mersmann O, et al. (2011) Exploratory landscape analysis. *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation*, ed Krasnogor N (ACM, New York), pp 829–836.
15. Lake BM, Ullman TD, Tenenbaum JB, Gershman SJ (2017) Building machines that learn and think like people. *Behav Brain Sci* 40:e253.
16. Marcus G (2018) Deep learning: A critical appraisal. arXiv:1801.00631 [cs, stat]. Preprint, posted January 2, 2018.
17. Kamar E (2016) Directions in hybrid intelligence: Complementing AI systems with human intelligence. *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, ed Brewka G (AAAI Press, New York), pp 4070–4073.
18. Baltz EA, et al. (2017) Achievement of sustained net plasma heating in a fusion experiment with the optometrist algorithm. *Sci Rep* 7:6425.
19. IBM (2016) IBM Quantum Experience. Available at www-03.ibm.com/press/us/en/pressrelease/49661.wss, www.research.ibm.com/quantum/. Accessed October 25, 2018.
20. University of Bristol (2013) Quantum in the Cloud. Available at www.bris.ac.uk/news/2013/9720.html, www.bristol.ac.uk/physics/research/quantum/engagement/qcloud/. Accessed October 25, 2018.
21. Hebenstreit M, Alsina D, Latorre JI, Kraus B (2017) Compressed quantum computation using a remote five-qubit quantum computer. *Phys Rev A* 95:052339.
22. Bonney R, et al. (2014) Next steps for citizen science. *Science* 343:1436–1437.
23. The BIG Bell Test Collaboration (2018) Challenging local realism with human choices. *Nature* 557:212–216.
24. Lintott CJ, et al. (2008) Galaxy zoo: Morphologies derived from visual inspection of galaxies from the Sloan digital sky survey. *Mon Not R Astron Soc* 389:1179–1189.
25. Cooper S, et al. (2010) Predicting protein structures with a multiplayer online game. *Nature* 466:756–760.
26. Lee J, et al. (2014) RNA design rules from a massive open laboratory. *Proc Natl Acad Sci USA* 111:2122–2127.
27. Kim JS, et al. (2014) Space–time wiring specificity supports direction selectivity in the retina. *Nature* 509:331–336.
28. Shir OM, Roslund J, Whitley D, Rabitz H (2014) Efficient retrieval of landscape Hessian: Forced optimal covariance adaptive learning. *Phys Rev E* 89:063306.
29. Freeman PR (1983) The secretary problem and its extensions: A review. *Int Stat Rev Rev Int Stat* 51:189.
30. Blei DM, Smyth P (2017) Science and data science. *Proc Natl Acad Sci USA* 114:8689–8692.
31. Ketterle W, Durfee DS, Stamper-Kurn DM (1999) Making, probing and understanding Bose-Einstein condensates. *Bose-Einstein Condensation in Atomic Gases: Varenna on Lake Como, Villa Monastero, 7–17 July 1998, Proceedings of the International School of Physics "Enrico Fermi,"* eds Inguscio M, Stringari S, Wieman CE (IOS Press, Amsterdam), Vol 140, pp 67–176.
32. Gross C, Bloch I (2017) Quantum simulations with ultracold atoms in optical lattices. *Science* 357:995–1001.
33. Rohringer W, et al. (2008) Stochastic optimization of a cold atom experiment using a genetic algorithm. *Appl Phys Lett* 93:2006–2009.
34. Rohringer W, Fischer D, Trupke M, Schmiedmayer J, Schumm T (2011) Stochastic optimization of Bose-Einstein condensation using a genetic algorithm. *Stochastic Optimization - Seeing the Optimal for the Uncertain*, ed Dritsas I (InTech, Rijeka, Croatia), pp 3–28.
35. Geisel I, et al. (2013) Evolutionary optimization of an experimental apparatus. *Appl Phys Lett* 102:214105.

36. Lausch T, et al. (2016) Optimizing quantum gas production by an evolutionary algorithm. *Appl Phys B* 122:112.
37. Wigley PB, et al. (2016) Fast machine-learning online optimization of ultra-cold-atom experiments. *Sci Rep* 6:25890.
38. Bason MG, et al. (2018) Measurement-enhanced determination of BEC phase transitions. *J Phys B At Mol Opt Phys* 51:175301.
39. Grimm R, Weidemüller M, Ovchinnikov YB (2000) Optical dipole traps for neutral atoms. *Advances in Atomic, Molecular, and Optical Physics*, eds Bederson B, Walther H (Academic, New York) Vol 42, pp 95–170.
40. Lin YJ, Perry AR, Compton R, Spielman IB, Porto J (2009) Rapid production of Rb[87] Bose-Einstein condensates in a combined magnetic and optical potential. *Phys Rev A* 79:063631.
41. Pinkse PWH, et al. (1997) Adiabatically changing the phase-space density of a trapped Bose gas. *Phys Rev Lett* 78:990–993.
42. van der Maaten L, Hinton G (2008) Visualizing data using t-SNE. *J Mach Learn Res* 9:2579–2605.
43. Brif C, Chakrabarti R, Rabitz H (2010) Control of quantum phenomena: Past, present and future. *New J Phys* 12:075008.
44. Rach N, Müller MM, Calarco T, Montangero S (2015) Dressing the chopped-random-basis optimization: A bandwidth-limited access to the trap-free landscape. *Phys Rev A* 92:062343.
45. Doria P, Calarco T, Montangero S (2011) Optimal control technique for many-body quantum dynamics. *Phys Rev Lett* 106:190501.
46. Russell B, Rabitz H, Wu RB (2017) Control landscapes are almost always trap free: A geometric assessment. *J Phys A: Math Theor* 50:205302.
47. Caneva T, et al. (2014) Complexity of controlling quantum many-body dynamics. *Phys Rev A* 89:042322.
48. van Frank S, et al. (2016) Optimal control of complex atomic quantum systems. *Sci Rep* 6:34187.
49. Rosi S, et al. (2013) Fast closed-loop optimal control of ultracold atoms in an optical lattice. *Phys Rev A* 88:021601.
50. Frank F, et al. (2017) Autonomous calibration of single spin qubit operations. *Npj Quan Inf* 3:48.
51. Weidner CA, Anderson DZ (2018) Experimental demonstration of shaken-lattice interferometry. *Phys Rev Lett* 120:263201.
52. Billinger S, Stieglitz N, Schumacher TR (2013) Search on rugged landscapes: An experimental study. *Organ Sci* 25:93–108.
53. Vuculescu O (2017) Searching far away from the lamp-post: An agent-based model. *Strateg Organ* 15:242–263.
54. Mason W, Suri S (2012) Conducting behavioral research on Amazon's mechanical Turk. *Behav Res Methods* 44:1–23.
55. Morgan T, Rendell L, Ehn M, Hoppitt W, Laland K (2012) The evolutionary basis of human social learning. *Proc R Soc B* 279:653–662.
56. Muthukrishna M, Morgan TJ, Henrich J (2016) The when and who of social learning and conformist transmission. *Evol Hum Behav* 37:10–20.