

Crowdsourcing human common sense for quantum control

Jesper Hasseriis Mohr Jensen , Miroslav Gajdacz, Shaeema Zaman Ahmed , Jakub Herman Czarkowski, Carrie Weidner , Janet Rafner, Jens Jakob Sørensen, Klaus Mølmer, and Jacob Friis Sherson*

Department of Physics and Astronomy, Aarhus University, Ny Munkegade 120, 8000 Aarhus C, Denmark



(Received 7 September 2020; accepted 7 December 2020; published 19 January 2021)

Citizen science methodologies have over the past decade been applied with great success to help solve highly complex numerical challenges. Here, we take early steps in the quantum physics arena by introducing a citizen science game, *Quantum Moves 2*, and compare the performance of different optimization methods across three different quantum optimal control problems of varying difficulty. Inside the game, players can apply a gradient-based algorithm (running locally on their device) to optimize their solutions and we find that these results perform roughly on par with the best of the tested standard optimization methods performed on a computer cluster. In addition, cluster-optimized player seeds was the only method to exhibit roughly optimal performance across all three challenges. Finally, player seeds show significant statistical advantages over random seeds in the limit of sparse sampling. This highlights the potential for crowdsourcing the solution of future quantum research problems.

DOI: [10.1103/PhysRevResearch.3.013057](https://doi.org/10.1103/PhysRevResearch.3.013057)

I. INTRODUCTION

Despite amazing advances in the past years, it has become increasingly clear that pattern-matching results from deep learning algorithms alone can be surprisingly brittle [1,2]. This failure has been attributed to, among other reasons, a lack of hierarchical learning, transfer between problems, and common sense comprehension of real world phenomena [3]. Common sense has been defined in many different ways, but here we refer to it as implicit and shared fundamental assumptions that people have about the world [4]. Additionally, there are indications that humans may sometimes solve computationally hard problems quickly and near-optimally [5,6]. However, humans just as often fail miserably [7]. Thus, many argue that optimized synergetic systems integrating individuals or collectives of humans and machines offer a promising, human-in-the-loop, approach to tackle complex problems [8–12]. One key challenge in this approach is that it requires large-scale studies of human capacities, for example common sense and the development of rich cognitive models [3]. Initial steps in this direction can be taken by exploring problems in research-relevant contexts, such as in the related fields of citizen science and collective intelligence [9,11], and detailed comparisons between human [6] and AI [13] performance are becoming feasible for problems such as protein folding.

In this paper, we take steps toward exploring how citizen input can be applied to the solution of relevant problems in quantum mechanics. In particular, we consider a dynamic quantum state transfer task (Sec. II) where the goal is to maximize the *fidelity*, $0 \leq F \leq 1$, by finding appropriate control functions $\{u_1(t), u_2(t)\}$ (solutions) as illustrated in Fig. 1. These are typically identified by locally ascending *seeds* (starting points) to the nearest maximum in the associated optimization landscape. Possible seeding strategies range from uniform random guessing (least imposed structure) to parametrizations based on highly domain-specific expert knowledge or heuristics (most imposed structure) as depicted by the colored regions in Fig. 1, each with associated probabilities of obtaining $F \approx 1$ upon optimization. Discovering the “good” regions of the optimization landscape can often be very challenging. In this work, we investigate whether nonexpert citizen scientists may be useful in efficiently identifying good regions by gathering their input through our game, *Quantum Moves 2* (Sec. III and Appendix A), which allows the player to both create seeds and subsequently engage with a local optimization algorithm embedded in the game and see its action in real time. Each player can be considered as an independent, adaptive seeding strategy that incorporates high-level heuristics and complex decision-making processes into an optimization loop which are otherwise difficult to capture and implement programmatically. This methodology may be used as a general means to extract features and heuristics for a given problem which could then aid experts in further analysis and guide the development of seeding strategies. We analyze three distinct scientific challenges—i.e., optimization landscapes—where the performance of the hybrid approach (human-computer) is compared against standard methodologies from quantum optimal control and computer science. In this sense, this work does not aim to present or promote an inimical competition between players and computer algorithms,

*sherson@mgmt.au.dk

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

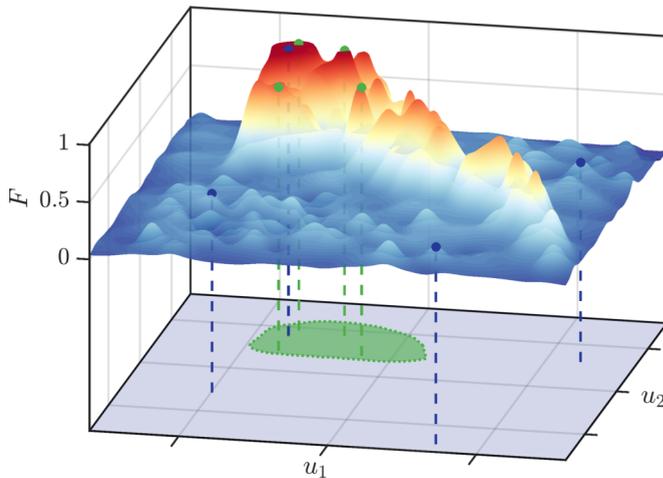


FIG. 1. Abstract illustration of an optimization (control) landscape with two control parameters. Points in the plane correspond to a particular set of functions $\{u_1(t), u_2(t)\}$ and the height of the landscape is given by the associated fidelity. Only a small region of the control space is optimal (with respect to $F \approx 1$) and otherwise contains many local “traps”—the relative efficiency of a local “hill-climbing” optimizer is thus correlated to where it is seeded. Purple full region: Uniform random seeding. Obtaining $F \approx 1$ is possible, but with low probability. Green smaller region: Specialized seeding. Obtaining $F \approx 1$ is possible with high probability, but not guaranteed. Such a region is conventionally targeted by experts through their accumulated knowledge, problem insights, heuristics, and programming proficiency. Nonexpert citizen scientist players may be useful in uncovering this region without extensive prior training.

but rather, explore their possible interplay in terms of solution strategy and the usefulness of player seeding in such problems. Our optimization code is available at Ref. [14].

Note that a simpler game interface, *Quantum Moves 1*, was previously accessible, and studies of the user experience [15,16] and numerics of the *Bring Home Water* level [17,18] were published. However, due to an error in the optimization code [19], the quantitative comparisons presented in the publication Ref. [18] are not valid and the article has been retracted. The research presented in the present work, which is broader both in terms of scope and analyses, was conducted with a wholly new codebase, including algorithms explored in Refs. [20,21], and free of the error that hampered the results in Ref. [18].

To address the tenability of game-based exploration of quantum research problems, we pose two specific questions about the current and potential scientific contribution to quantum physics and related fields:

(Q1) Can a suitable gamified interface allow citizen scientists to solve quantum control problems entirely on their own, using their own hardware to run the required computation (e.g., using an in-game optimizer), with a quality on par with traditional expert-driven optimization?

(Q2) Can the player-generated solutions, in combination with concrete algorithms, provide an edge against fully algorithmic solutions, and how does that depend on the type and mathematical complexity of the problem?

If Q1 can be answered in the affirmative, then such an interface would allow for a novel form of online, quantum citizen science combining human problem solving with crowd computing. In such a framework, one could imagine quantum researchers continually feeding in optimization challenges that are then solved efficiently by the community at no computational cost to the researchers.

Before moving on to a detailed description, we provide an overview geared toward the broader audience.

A. Overview of results and discussion

To address the questions Q1 and Q2 in a systematic manner and to establish a baseline, we compare optimization of player seeding to uniform random seeding (Sec. IV) across a range of distinct problems. Our results show that these problems differ vastly in terms of both landscape- and numerical complexity which impacts the difficulty of locating and identifying the optimal strategy (or strategies) corresponding to the green region(s) of Fig. 1. Notably, in the first problem (Sec. V), the structure in the citizen player seeds allows—while the lack of structure in the random seeds prevents—discovery of the exponentially gainful optimal strategy when applying a local optimization algorithm. The second problem (Sec. VI), however, turns out to have a very benign landscape and all seeding strategies can successfully identify the optimal solutions. The third problem (Sec. VII) exhibits, on the contrary, a very complex landscape with an interleaved plurality of strategies that are optimal in different regimes. Player and random seeds probe different parts of the landscape, and, in this case, this complementarity leads to enhanced collective results. Outside of their respective problem-dependent capabilities of discovering the optimal strategies, we find in a separate analysis (Sec. VIII) that across all problems, player seeds are statistically much more efficient in obtaining high-fidelity results than random seeds when restricting the sample size. (Performance in the limit of restricted sample size is relevant when extrapolating to problems of increased numerical complexity or when available computational resources are restricted and thus effects a less dense probing of the landscape.) Additionally, we find further significant enhancements when sampling only from the best raw player solutions (as measured by the fidelity) and even more advantages arise when sampling only from the best player-optimized solutions. Such enhancements are not found for the best random seeds. This indicates that both the structure of the player seeds and their donated computing power in “pre-optimizing” these are contributing factors to this statistical advantage. Thus, it is suggested that it is indeed beneficial to crowdsource the solution of, at least, this particular suite of quantum control problems.

This broad contribution of the players stands in stark contrast to, e.g., citizen science projects like Foldit [22], where only a small fraction of players provide a scientific contribution after an extensive training process. To distinguish these two types of citizen science challenges, we assert in this work that the *Quantum Moves 2* game, to some extent, taps into certain aspects of common sense (shared tacit knowledge of reasonable behavior) of the player population at large. Although it is not within the scope of this work to analyze the explicit nature of this common sense, it seems reasonable

to conjecture that the liquid analogy of the wave function dynamics taps into the classical intuition for sloshing water, which was also argued by Sels in his numerical analysis of one of the problems [20]. Along similar lines, we believe that successful generalizations of this citizen science methodology are correlated with the interfaces devised to represent the underlying optimization tasks. That is, the design of interfaces that are well-suited for facilitating the strengths of—and interplay between—both components of the human-computer hybrid approach.

Based on these answers to Q1 and Q2, we suggest posing additional questions for this line of research:

(Q3) Could citizen science games be first steps toward playful, explorative tools for domain experts? This approach is currently being pursued in microbiological research settings [23] but not yet in quantum physics. As a related question, can the citizen science experience be expanded to include citizen contribution in more aspects of the scientific method such as data analysis and, ultimately, hypothesis and problem formulation? If such so-called extreme citizen science [24] could be understood and systematically implemented, then this would constitute a major advance in complex problem solving, one of the most demanded skills as per the World Economic Forum [25].

(Q4) If the games are sufficiently challenging for humans, then can they be used for systematic studies of human problem solving? In our group we have started to investigate this within the setting of quantum experimental optimization [26] and by developing cognitive science variants of quantum challenges [27].

(Q5) If larger portions of the player base can make non-trivial contributions to several classes of research problems using the gamified interface, then could this contribute to the solution of one of the major roadblocks in the path to domain-general AI according to, e.g., the author of Ref. [3]—that is—the crowdsourcing and algorithmization of human common sense?

The latter two questions move well beyond the realm of quantum physics and the scope of this paper. In terms of generalizability outside the citizen science context, it is remarked that Q3 typically reframes the design goals of a given interface as it caters to the inherently different target audience of domain (semi)experts. In this instance, for example, classical analogies and interpretations become much less important with focus instead on an expanded list of more advanced concepts and flexible visualizations. We currently pursue elements of Q3 in other work [28] by developing an intuitive and visual quantum programming environment which was originally inspired by the *Quantum Moves* games.

The extent to which our work may be extended to gain broader implications for the fields of quantum research (Q3), social science (Q4), and computer science (Q5) constitutes interesting topics for future studies.

II. QUANTUM OPTIMAL CONTROL

In this section, we describe the context and define key goals of *Quantum Moves 2*.

The hallmarks of the second quantum revolution [29] are the exploitation and engineering of fragile, isolated quantum

objects. Quantum computing with any platform predicates precise control of the constituent qubits and associated gate operations. Additionally, the control must also be expeditiously carried out such as to avoid decoherence and other detrimental effects to the overall goal. Controls meeting these criteria and more can be obtained within the well-established theory of quantum optimal control.

A common class of quantum optimal control problems deals with facilitating a particular initial-to-target state transfer, $|\psi_0\rangle \rightarrow |\psi_{\text{tgt}}\rangle$, for some fixed process duration T . The manipulatory access to the state evolution is through a set of *control* parameters $\{u_1(t), u_2(t), \dots\}$ where each solution [specific choice of functions $u_i(t)$] uniquely maps to a final state $|\psi_0\rangle \rightarrow |\psi(T)\rangle$ (see Appendices B and C). In this context, the transfer fidelity

$$0 \leq F[\{u_1(t), u_2(t), \dots\}; T] = |\langle \psi_{\text{tgt}} | \psi(T) \rangle|^2 \leq 1, \quad (1)$$

for each fixed T can be interpreted as a high-dimensional optimization (or control) landscape as illustrated in Fig. 1. *Optimal* controls (or solutions) can then be associated with points in the landscape that are globally or locally maximal. Additionally, for a given fidelity requirement, we associate a fundamental minimal duration, T_{min}^F , below which no maximum exceeds F . Thus, T_{min}^F is defined as the shortest duration at which at least one control corresponding to a maximum can obtain the given F . Depending on the context, common choices for threshold values are $F = 0.99, 0.999, 0.9999, \dots$, characterized by a tradeoff between F and T_{min}^F (increased precision leads to longer durations).

In the limit $T \rightarrow \infty$, most problems become easy in the sense that many global maxima with $F \approx 1$ exist. As $T \rightarrow 0$, however, the control problem becomes increasingly difficult as previously global maxima gradually become only locally maximal and the control landscape becomes more rugged. The usually unfavorable topography of the control landscape in the $T \approx T_{\text{min}}^F$ regime therefore makes uncovering global maxima especially difficult for the aforementioned high-fidelity requirements.

At its core, any iterative optimization algorithm attempting to locate global maxima must prescribe a way to traverse the optimization landscape in a meaningful way. It must thus strike a balance between local (exploitation) and global (exploration) search methodologies. A common optimization paradigm initializes an algorithm, e.g., one excelling in finding the nearest local optimum, from many different seeds, thereby introducing a simple global component [30]. In principle, a local optimizer maps each seed to its nearest *attractor* [31]. Nonglobal local optima are often called *local traps*, referring to the propensity for said optimizers to locate these and terminate (since they are “stuck”). The effectiveness of this paradigm is then necessarily strongly correlated to a combination of the seeding strategy (the mechanism with which seeds are generated) and the choice of optimization algorithm.

Effective seeding strategies, e.g., those targeting the green region of optimality depicted in Fig. 1, and algorithms naturally become increasingly important with growing problem complexity and this can broadly be characterized by two axes: the computational (or numerical) complexity of the underlying simulations and the inherent topographical landscape

complexity. The computational difficulty can also be interpreted as the amount of required resources. Below, we discuss problems with two different degrees of computational difficulty drawing upon the single particle Schrödinger equation and the nonlinear dynamics of Bose-Einstein condensates.

III. OVERVIEW OF QUANTUM MOVES 2

In *Quantum Moves 2*, the player's goal is to solve various state transfer problems, referred to in-game as *levels*. Each level concerns 1D transfers of either single particle or Bose-Einstein condensate (BEC) wave functions $\psi(x, t) = \langle x | \psi(t) \rangle$, both describable by the Hamiltonian

$$\hat{H} = -\frac{\hbar^2}{2m} \frac{\partial^2}{\partial x^2} + V + g|\psi|^2, \quad (2)$$

where taking the nonlinear coupling parameter $g = 0$ corresponds to the single-particle case. While there is no in-game distinction in their representation, the numerics involving BEC ($g \neq 0$) are more intricate, and this has consequences for the efficiency of some optimization algorithms as discussed later. The potential has up to two controllable parameters, $V = V(u_1(t), u_2(t))$, depending on the level.

The main features of *Quantum Moves 2* are

- (1) A device-embedded algorithm that enables players optimize the seeds they produce.
- (2) In-game tools for analyzing previous solutions.
- (3) Constrained exploration to the regime near T_{\min}^F and below with $F = 0.99, 0.999$.
- (4) Scientific problems of various complexity.

Here we briefly describe and motivate the three main levels in the game. Figure 2 displays their in-game representation. See Appendix A for a more complete description of the game interface.

Bring Home Water. A single atom resides in the ground state of a static tweezer and must be picked up and shuttled back into the ground state at the original location of the movable tweezer. This type of transfer is necessary for implementing quantum computations in neutral atoms based on collision gates [32].

Splitting. A BEC initially resides in the ground state of a single-well configuration on an atom chip and must be transferred into the ground state of a double-well configuration by deforming the potential. The split condensate can then be used for matter-wave interferometry [33–36].

Shake Up. A BEC initially resides in the ground state of a single-well configuration on an atom chip and must be transferred into the first excited state by shaking the potential. The excited state of the BEC acts as a source for twin-atom beams [35–39].

IV. ALGORITHMS AND SEEDING STRATEGIES

In this section we specify the suite of different algorithms and seeding strategies under consideration. We define a *method* as a particular combination of algorithm and seeding strategy. For example, GRAPE Pr-RS is the method that uses the GRAPE algorithm to optimize preselected random seeds. Further details of numerics and each algorithm are included in Appendices B and C.

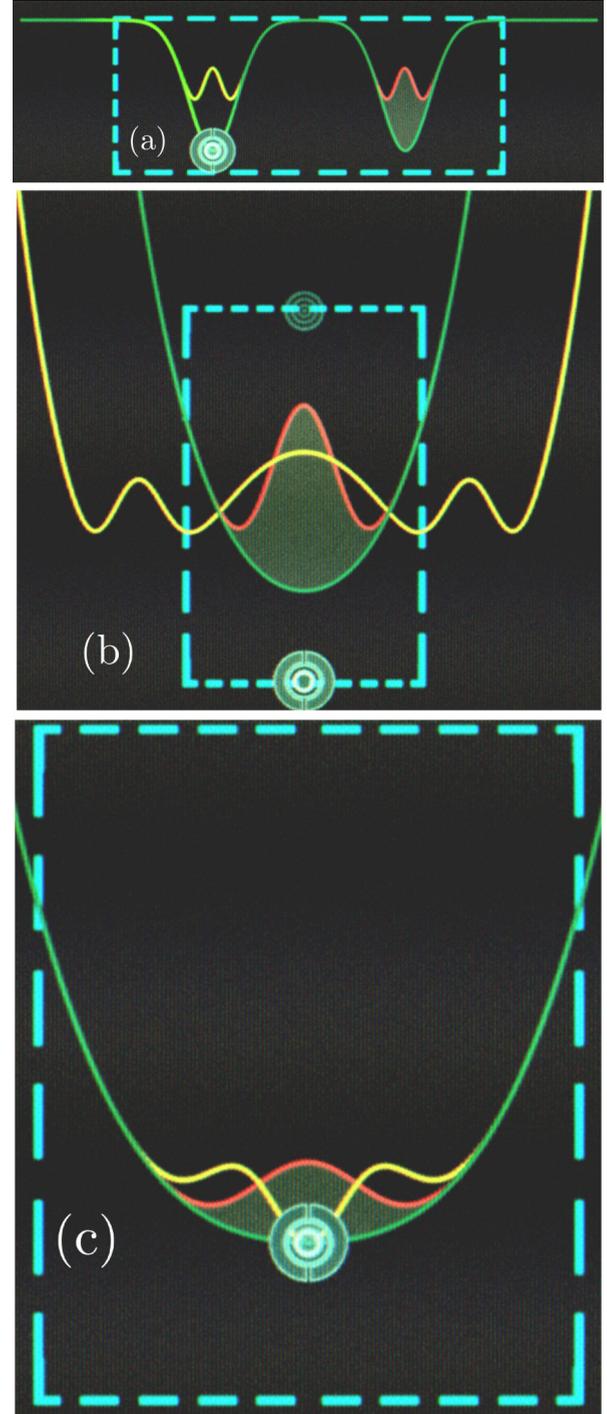


FIG. 2. The three central state transfer problems in *Quantum Moves 2*. (a) *Bring Home Water* (single-particle), (b) *Splitting* (BEC), and (c) *Shake Up* (BEC). The instantaneous density $|\psi(x, t)|^2$ (red line with shaded area) must be transferred into the target density $|\psi_{\text{tgt}}(x)|^2$ (yellow line) without residual excitation. The trapping potential (green line) is parametrized by the position of the round, draggable cursor, $\{u_1(t), u_2(t)\} = \{f_1(x_{\text{cursor}}(t)), f_2(y_{\text{cursor}}(t))\}$, which is unable to leave the turquoise bounding box (control boundaries). The functions f_1 and f_2 are linear. The wave function densities are offset by the potential, $|\psi(x)|^2 + V(x)$, for illustrative purposes (e.g., the *Splitting* target density represents two equally sized wave packets trapped in a double well—the large center bump is the barrier and two smaller bumps are the wave packets).

A. Algorithms

GRAPE. Standard gradient-based optimization using the L-BFGS quasi-Newton search direction with line search. Bandwidth limitation (smoothness) is included through a derivative-regularization cost term.

PGRAPE. Player GRAPE. The player can start and stop the optimization. Otherwise, it is identical to GRAPE. The algorithm is executed locally on the player's device as a part of the game.

Stochastic Ascent (SA). Gradient-free maximally greedy time-local search. The process duration T is segmented into n_b bins of equal width within which the control values are constant. The bin values are updated in a stochastic order. The bandwidth limitation (smoothness) is inversely proportional to n_b (there is no additional regularization cost term in the current implementation).

B. Seeding Strategies

As outlined in Sec. II, seeding strategies are a fundamental component of optimization. Drawing from random distributions provides the most generic way of seeding [40]:

Random Seed (RS). The control is assembled by independently sampling a uniform distribution within given boundaries (Appendix B) for each control parameter,

$$u_p(t) = \text{uniform}(u_{\min}, u_{\max}), \quad (3)$$

for $p = 1, 2$. Subsequent control values at $u_p(t + \delta t)$ are completely uncorrelated. To introduce correlations, one can also segment T into n_b bins of equal width w such that the control value is initially constant within each bin (see also Appendix C).

Quantum Moves 2 provides two novel seeding strategies:

Player Seed (PS). The control is assembled by mapping the players' cursor position during gameplay as a function of time,

$$\begin{pmatrix} u_1(t) \\ u_2(t) \end{pmatrix} = \begin{pmatrix} f_1[x_{\text{cursor}}(t)] \\ f_2[y_{\text{cursor}}(t)] \end{pmatrix}. \quad (4)$$

Player Optimized Seed (PO). This seeding strategy describes a PS seed that has been optimized by the player (i.e., PO \equiv PGRAPE PS),

$$\begin{pmatrix} u_1(t) \\ u_2(t) \end{pmatrix} = \text{PGRAPE} \begin{pmatrix} f_1[x_{\text{cursor}}(t)] \\ f_2[y_{\text{cursor}}(t)] \end{pmatrix}. \quad (5)$$

If the optimization was stopped before convergence, then the optimized control can be used as a seed for further optimization. If the optimization converged, then the seed itself is already a local optimum. From a resource perspective, these seeds are very valuable since they come partially or fully pre-optimized at no cost to the research team.

A heuristic extension of any seeding strategy is preselection:

Preselection (Pr). A naïve greedy heuristic to choose which seeds should be picked for optimization. Given a set of candidate seeds and their associated fidelities, optimize only the N seeds with highest initial fidelity.

V. BRING HOME WATER

In this section we present optimization results obtained in *Bring Home Water* for the different methods described in Sec. IV. The allotted resources for the optimizations are discussed in Appendix C. We denote by PGRAPE PS \cup PS the joint set of optimized and unoptimized solutions produced solely by the players in-game [41].

Figure 3 shows the density distribution of solution infidelities ($1 - F$) [42] as a function of process duration T in the high-fidelity regime for various methods. For a given method, the solution densities estimate the probability distribution $\mathcal{P}(F|T)$ of obtaining a particular fidelity for a given T , since each column is individually normalized. For reference, the solid line shows the best obtained GRAPE PS results from Fig. 4. Figure 4 shows the aggregate, monotonically best infidelity as a function of process duration T . The green crosses show all results produced only by players in-game, i.e., player seeds (PS) and player-optimized seeds (PGRAPE PS). The blue dots show the result of GRAPE PS, i.e., computer cluster optimized player seeds (PS) only (inducing a difference in number of seeds in Fig. 4). Lines with (without) dots show the best results obtained by methods based on player (random) seeds.

For the GRAPE PS density and PGRAPE PS \cup PS results shown in Fig. 4, we observe two clear bands of solutions that are each described by distinct exponential behavior, which hints at two corresponding solution strategies. This is verified and analyzed in Sec. VA using clustering techniques. The identification of the two exponentially-gapped solution strategies and the relative likelihoods of different methods identifying each strategy is the first of three main findings in this section. The duration-dependent globally optimal strategy changes from one to the other near $T = 0.092$ ms, explaining the kink in the best result reference curve seen in Fig. 3 and the departing line of blue dots in Fig. 4. The gap between the global and local optimal strategies increases exponentially, leading to significantly different minimal duration T_{\min}^F estimates (as defined in Sec. II). Outside of these strategies, the densities are mostly sparse but nonzero, indicating a topographically complex optimization landscape containing many isolated local traps or regions with near-vanishing gradient. This is also understood by the interspersed dot and cross distributions in Fig. 4.

We now examine GRAPE RS, where the only difference with respect to the former methods is the seeding strategy. In this case, we observe that only the inferior, locally optimal strategy is discovered, with only two (out of 20 734) solutions located in the strategy gap. Evidently, for the same optimization algorithm, the structure of player seeds is preferable. The deficiency of GRAPE RS is analyzed in Sec. VB.

Next, we turn our attention to SA RS variants with full resolution ($n_b = n_r$) and reduced resolution ($n_b = 40$). As discussed in detail in Appendix C, SA is expected to be efficient for linear problems ($g = 0$) with preferably a single control parameter. For *Bring Home Water* the former is satisfied by the problem definition and the latter can be satisfied by choosing the tweezer amplitude control such that it is maximally deep at all times, $u_2(t) = u_2^{\min}$, as is also done in Ref. [20] and SA thus only optimizes $u_1(t)$ [43]. Generally, in all the optimized GRAPE solutions, we indeed observe that the tweezer tends to

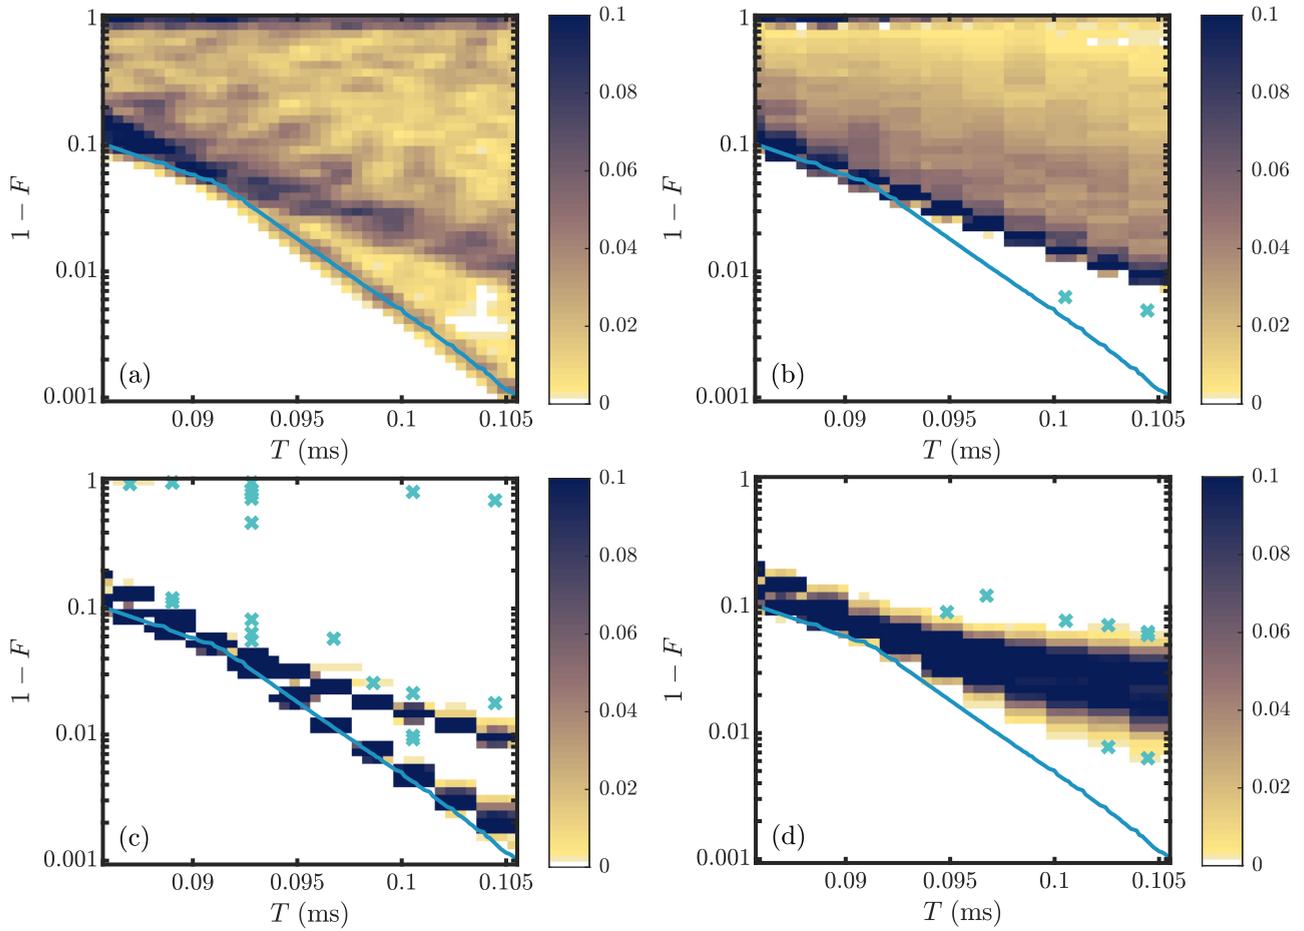


FIG. 3. *Bring Home Water* solution densities for different methods. (a) GRAPE PS, (b) GRAPE RS, (c) SA RS ($n_b = n_t$), and (d) SA RS ($n_b = 40$). Each density is normalized for every individual T and thus represents an estimate of the probability distribution for obtaining a particular F for a given T [$\mathcal{P}(F|T)$]. The reference curve shows the best obtained results for GRAPE PS. The crosses indicate solutions at densities lower than 0.002.

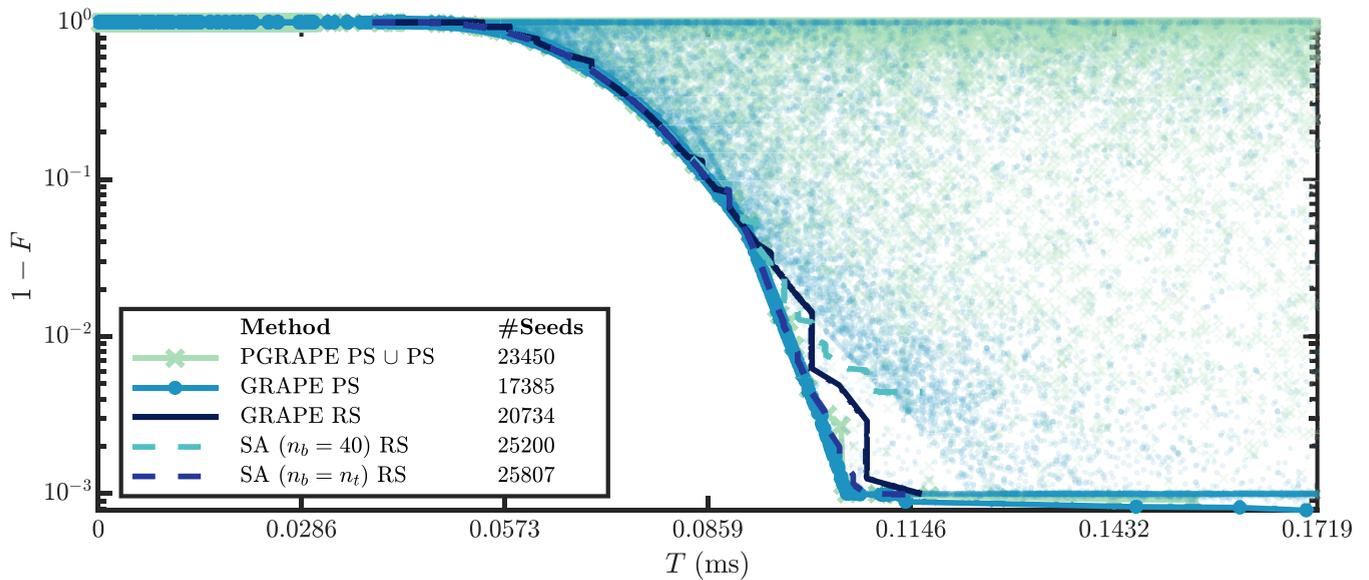


FIG. 4. Aggregate, monotonically best optimization results (lower is better) for several methods in *Bring Home Water*. Solid lines with (without) symbols show the best results obtained with (without) player influence. The scattered green crosses show all results produced by 536 players seeding on average ~ 32 solutions and optimizing approximately 1/4 of these for ~ 131 iterations on average. The scattered blue dots show the same, except the optimization is carried out on all player seeds with the computational resources described in the text (i.e., no player influence after seeding). The symbol translucency indicates the density distribution.

attain its maximal depth whenever the tweezer is in contact with the atom (overlap with nonzero probability density).

With these choices, SA RS ($n_b = n_t$) is capable of discovering the two distinct strategies with remarkable efficiency, despite the structureless nature of the $u_1(t)$ seeding mechanism. Only a few low fidelity solutions occur. SA RS ($n_b = 40$) does not identify the globally optimal strategy (lower branch) and instead concentrates on a broad band of solutions with a center shifted to above the inferior strategy (upper branch). This can be attributed to the reduced resolution of the algorithm preventing it from resolving the dynamics finely enough [44].

Both GRAPE PS and SA RS ($n_b = n_t$) find the same estimates $T_{\min}^{F=0.99} \approx 0.0973$ ms and $T_{\min}^{F=0.999} \approx 0.1057$ ms, whereas the estimate from PGRAPE PS \cup PS is off by less than 1% with respect to $F = 0.99$. The fact that the in-game player-optimized curve roughly matches the best results optimized on the computer cluster represents the second main finding of this section. This, coupled with similar findings for the two remaining challenges, represents our quantitative confirmation of Q1 (cf. Sec. I) for this range of control problems.

GRAPE PS, PGRAPE PS \cup PS, and SA RS ($n_b = n_t$) are all able to find both solution strategies, but SA RS ($n_b = n_t$) is the most efficient at optimizing low-fidelity solutions into high-fidelity solutions. This cannot be attributed to the different seeding mechanisms, since GRAPE RS fails to find the globally optimal solution strategy. Instead, the difference is due to how the two optimization algorithms traverse the landscape (Appendix C) and how they respond to being near a local trap [45].

Later in Sec. VIII we provide an alternative statistical characterization of each method, as well as the effect of the preselection heuristic, for each of the three problems.

A. Optimal strategies—Control clustering

To extract the identified solution strategies, we apply DBSCAN clustering [46] to the GRAPE PS method. Based on the results presented in the previous section, we expect the existence of distinct solution strategies, i.e., families of solutions that have a similar functional shape and characteristics but possibly different durations [47].

The clustering was performed only on the duration-normalized tweezer position, $u_1(t/T)$, justified by the fact that the optimized $u_2(t)$ control is, in general, maximally deep. To simplify the clustering, all solutions were given the same number of points by linearly interpolating on a 1000 point grid (the original number of points was defined by $T/\delta t$). For this analysis, we selected high-fidelity solutions $F \in [0.95, 0.999]$ with $T \in [0.093, 0.124]$ ms (i.e., near $T_{\min}^{F=0.99}$ and $T_{\min}^{F=0.999}$). We used the `sklearn` implementation of DBSCAN with Euclidean metric, $\epsilon = 3$, and a minimum number of neighbors $\text{min}_{\text{samples}} = 5$.

After filtering out physically insignificant delays [48] we find in Fig. 5 two major clusters (labeled 0 and 1) with increased populations and a “cluster” (labeled -1) in which local defects cause irreparable deviations from the two strate-

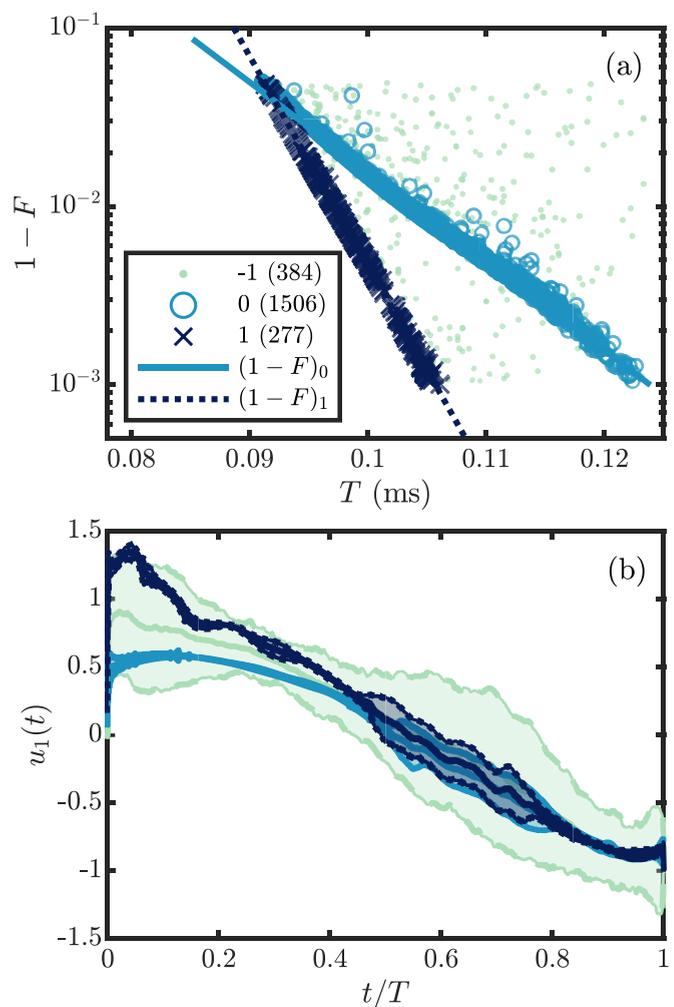


FIG. 5. *Bring Home Water*: Clustering of controls. (a) Each cluster (0 and 1) corresponds to a strategy and the gap between them exhibits an exponential $1 - F(T)$ behavior, leading to different estimates of minimal durations for a given threshold. The unclassified points (-1) are mostly populated by controls similar to either strategy, except for a few local defects that makes them appear distant to the cluster with respect to the Euclidean metric. (b) Lines correspond to cluster means and shaded areas to the standard deviation (cluster -1 has a dashed outline). The most populous cluster (0) corresponds to a front-swing strategy with the tweezer immediately being placed in front of the atom whereas the less populous cluster (1) corresponds to a back-swing strategy with the tweezer immediately being placed behind the atom.

gies. Each cluster exhibits a strikingly exponential tradeoff between fidelity and duration well-described by the fits

$$[1 - F(T)]_0 = 10^{-1.45 - 50.11 \cdot (T/\text{ms} - 0.0929)}, \quad (6)$$

$$[1 - F(T)]_1 = 10^{-1.50 - 117.27 \cdot (T/\text{ms} - 0.0929)}, \quad (7)$$

and it is clearly seen that the strategy gap widens exponentially. Assuming these trends can be extrapolated,

this yields $T_{\text{QSL},0}^F/T_{\text{QSL},1}^F = 1.17, 1.26, 1.33$ and infidelity ratios of $1 \times 10^2, 3 \times 10^3, 6 \times 10^4$ at $T_{\text{QSL},0}^F$ for $F = 0.999, 0.9999, 0.99999$, respectively.

Physically, the strategy corresponding to the cluster 0 begins by placing the tweezer in front of the atom, providing immediate acceleration toward the target position. We name this the *front-swing* strategy. Conversely, the strategy corresponding to the cluster 1 begins by placing the tweezer behind the atom, providing immediate acceleration away from the target position. We name this the *back-swing* strategy. In both instances, there is very little deviation from the cluster mean except during the shuttling of the atom where small deviations are allowed.

Intuitively one might expect that the back-swing strategy would be slower because the atom must travel an overall longer distance compared to the front-swing strategy, but this is evidently not the case. Instead, initially displacing the atom onto the static well's right side can serve as an additional accelerating force to that of the movable tweezer. With this analysis, we are in a position to explore in the following section why GRAPE RS fails to locate the back-swing strategy.

B. GRAPE RS—Efficiency versus optimality

It is clear from Fig. 4 that the fine-grained GRAPE RS (with $n_b = n_t$) fails to find the back-swing solution where the same optimization algorithm with PS succeeds. As shown in Sec. V A, this strategy requires the control to be immediately placed on the right-hand side of the static tweezer ($u_1 > x_0$) for a time interval sufficient to displace the atom. To gain a sense of timescales, the harmonic approximation to the fully overlapped tweezers ($u_1 = x_0$) yields an oscillation period of $\tau = 2\pi/\omega \approx 0.1$ (in simulation units, see Appendix B). Assuming $\tau/10$ is a sufficient response time for meaningful dynamics, the probability for consecutively randomly sampling the corresponding number of points $n = (\tau/10)/\delta t = 10^{-2}/(3.5 \times 10^{-4}) \approx 29$ such that they all have ($u_1 > x_0$) is vanishingly small $P(u_1 > x_0)^n = (0.25)^n = 6 \times 10^{-18}$ when successive control values are uncorrelated. Even at a much less conservative estimate of $\tau/40$, the probability 5×10^{-5} remains strongly suppressed. Since GRAPE traverses the landscape locally, it is therefore almost guaranteed that subsequent optimization leads to the front-swing strategy. The same is not observed for SA, since the one-dimensional exhaustive search allows for the individual adjustment of points ($u_1 < x_0$) into ($u_1 > x_0$).

One possible explanation for the observed difference between randomly and player-seeded behavior is thus that with a very high time resolution the random seeds oscillate rapidly whereas the player seeds are typically comparatively rather smooth due to the physical limitations in the speed of the players' cursor movements. One might then hypothesize that sufficiently coarse-grained piecewise constant random seeds would also yield similarly good behavior. If this was the case, then the player superiority would indeed be a trivial artifact of the choice of discretization. To test this, we therefore study the behavior of GRAPE RS optimizations when heuristically introducing correlations in the random seed by dividing the seed into $n_b \leq n_t$ piecewise constant control segments (like with SA, but in this case the subsequent optimization is not

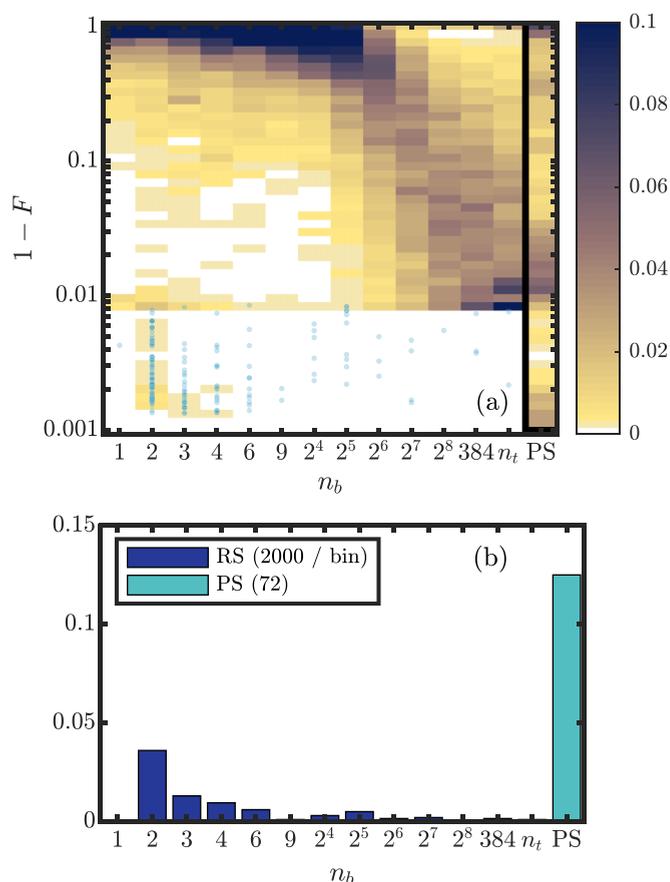


FIG. 6. (a) GRAPE RS solution densities as a function of number of bins n_b at $T = 0.1045$. Each dot denotes a solution in the back-swing strategy. The density in the right black box shows the GRAPE PS solution density from Fig. 3 at the same duration denoted by PS on the x axis (the duration is on the edge of one of the bin limits and the density is therefore taken as the mean of the two neighboring bins). (b) Histogram of the number of back-swing solutions, counted as those with fidelity higher than the best front-swing result at $n_b = n_t$, for GRAPE RS with variable n_b (blue) and GRAPE PS (turquoise) from Fig. 4. The total number of seeds per bin is indicated in the parentheses.

constrained to this parametrization). This makes it increasingly likely for the seed to initially begin and remain on the right side of the atom for a sufficient amount of time. Figure 6 shows the results of optimizing 2000 of these seeds near $T_{\text{min}}^{F=0.999}$ as a function of n_b , as well as the GRAPE PS results from Figs. 3 and 4. In the limit $n_b = n_t$, there is a high probability of finding solutions belonging to the front-swing strategy with associated fidelities around 0.99. This is consistent with the findings in Fig. 3. Only for $n_b \lesssim 4$ is there an appreciable albeit low probability (about 1 to 3.5%) of identifying back-swing solutions with much higher associated fidelities above 0.99, whereas more than 91% of the density resides below $F = 0.9$. The increase in low fidelity solutions can be attributed to the fact for, e.g., $n_b = 2$ (that has the highest empiric probability of finding the back-swing), there is an increased probability of placing the control tweezer far away from the atom and never touching the atom at all, leaving it in the initial state and resulting in a vanishing gradient (the

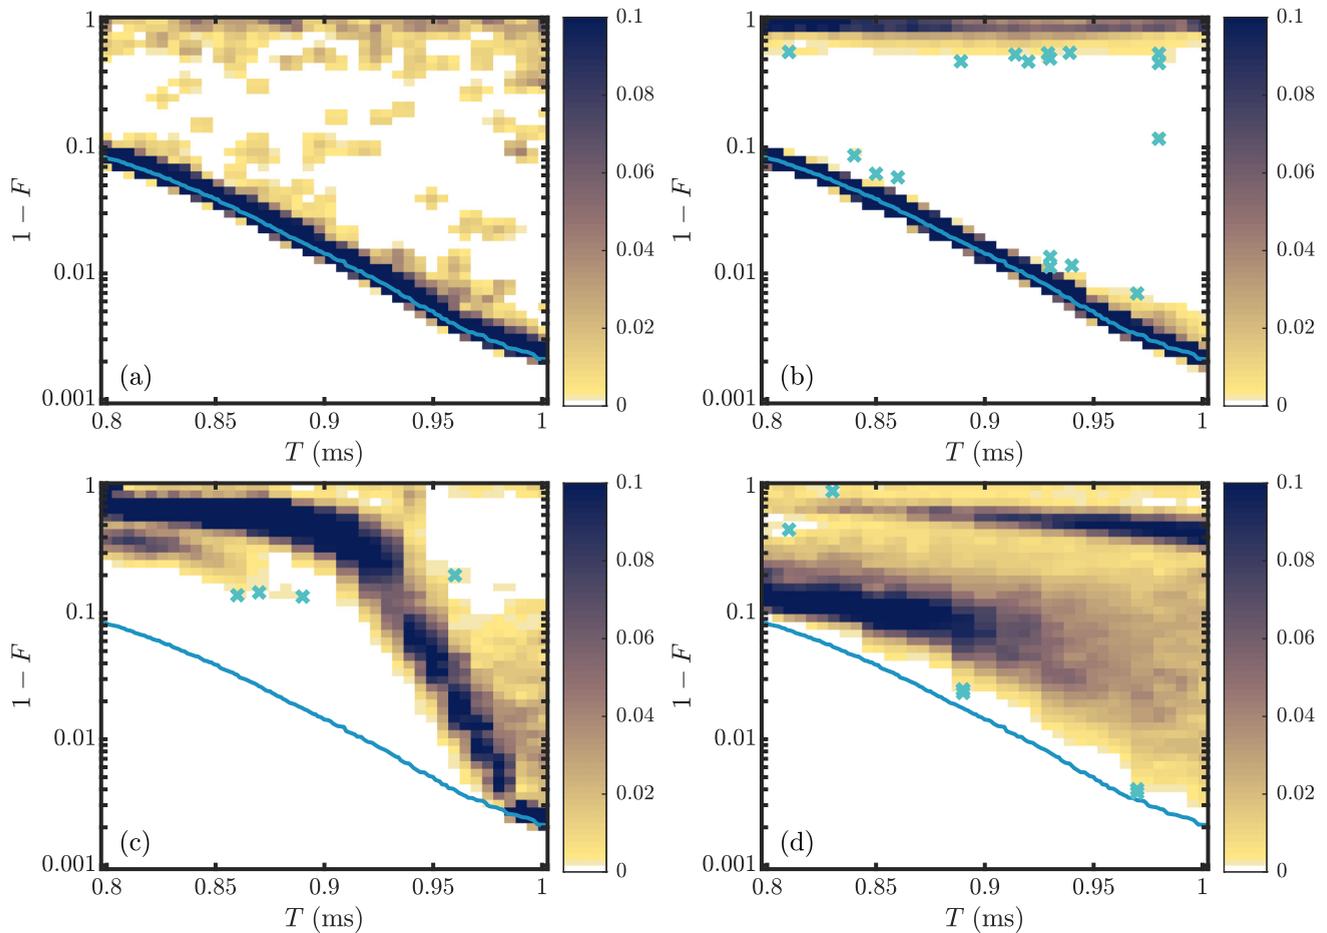


FIG. 7. *Splitting* solution densities for different methods. (a) GRAPE PS, (b) GRAPE RS, (c) SA RS ($n_b = n_t$), and (d) SA RS ($n_b = 40$). Each density is normalized for every individual T and thus represents an estimate of the probability distribution for obtaining a particular F for a given T ($\mathcal{P}(F|T)$). The blue reference curve shows the best obtained results for GRAPE PS. Crosses indicate solutions at densities lower than 0.002.

gradient is proportional to $\langle \psi_{\text{tgt}} | \psi(T) \rangle$). However, the probability of GRAPE PS finding the back-swing strategy is much larger (about 13%) with only 36% of the solutions below $F = 0.90$. The nontrivial observation that player seeds, for the same optimizer, outperform randomly generated piecewise constant seeds at all coarseness scales represents the third main contribution of this section.

The shapes of both optimal strategies in Fig. 5 contain a major linear component, and this suggests prospects for improvement by employing piecewise linear rather than piecewise constant seeding. That is, a figure similar to Fig. 6 with results from GRAPE employing such seeds would results would therefore likely yield a dramatically different picture. In particular, we note that the exclusive identification of the inferior front-swing strategy from the presented GRAPE RS results by itself points to the piecewise linear seeding, which, if applied in subsequent optimization, would likely lead to the discovery of the globally optimal back-swing strategy. This relies on the fact that both solution strategies are captured within the parametrization of few piecewise linear segments. In general, such a similarity between distinct solution strategies is not guaranteed. As an example, we find in Sec. VII that the *Shake Up* problem also contains a plurality of strategies that are

related to one another in a more subtle way. Bulk analysis of solutions from one strategy would thus not necessarily yield a problem parametrization that also encapsulates the other strategies. This underscores the potential usefulness of having data sets available that have been generated with sets of basic underlying assumptions that are as different as possible.

Our conclusion is therefore not that a standard, heuristic-free numerical approach *will* fail on the problem but that it *can* fail, and this result should not be taken as proof that players can guarantee computational improvements to this or other problems. It does, however, constitute a necessary first demonstration of the value of examining the gamified approach further and comparing to or integrating it with more sophisticated expert heuristics.

VI. SPLITTING

In this section, we present optimization results obtained for the *Splitting* level using the same suite of methods, convergence criteria, and computational resources as in Sec. V.

The solution densities are shown in Fig. 7 and the aggregate results are shown in Fig. 8. Looking at the solution density for GRAPE PS we see a single, dominant band of solutions with

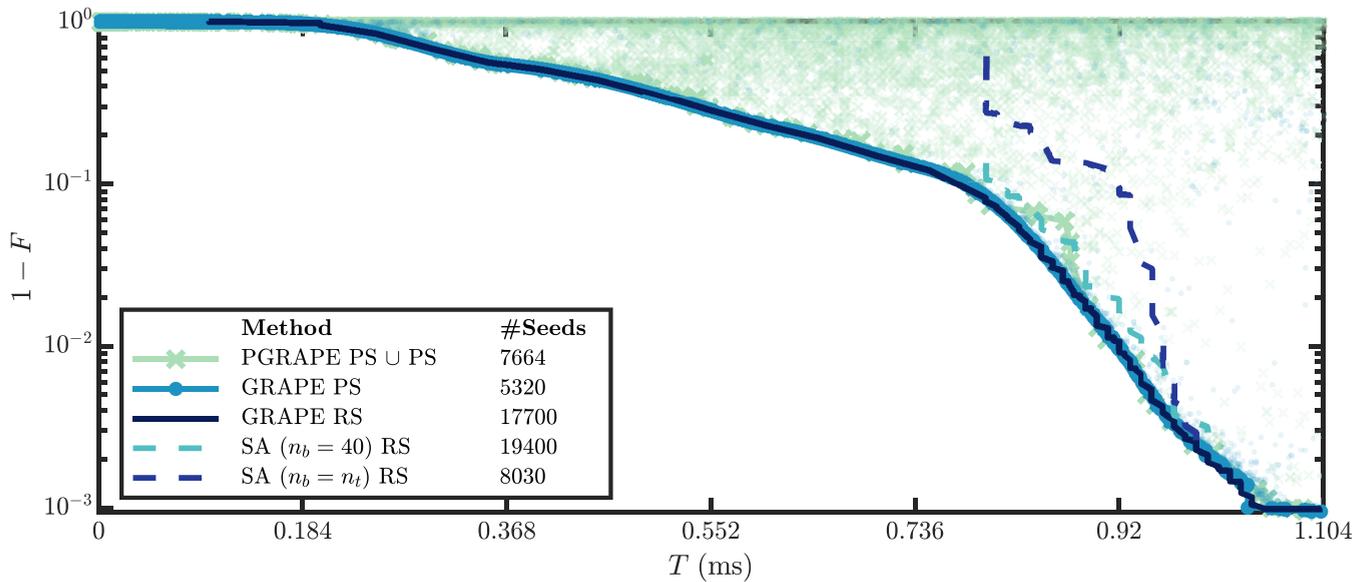


FIG. 8. Aggregate, monotonically best optimization results (lower is better) for several methods in *Splitting*. Solid lines with (without) symbols show the best results obtained with (without) player influence. The scattered green crosses show all results produced by 193 players seeding on average ~ 28 solutions and optimizing approximately $1/3$ of these for ~ 127 iterations on average. The scattered blue dots show the same, except the optimization is carried out on all player seeds with the computational resources described in the text (i.e., no player influence after seeding). The symbol translucency indicates the density distribution.

only a sparse population of solutions away from it. This hints at a simple, almost trap free, easily-navigated control landscape containing a very broad attractor for a single optimal strategy. Indeed, this is also understood by observing the dot density in Fig. 8: almost all the blue dots coincide with the best curve. A similar situation can be observed from GRAPE RS, except for a modest increase in low fidelity solutions but with virtually no population between these and the optimal band. Specifically, the average total density per bin below $F \approx 0.9$ is $(18 \pm 4)\%$ for RS and $(11 \pm 5)\%$ for PS. Player seeds are thus slightly more likely to be within reach of the optimal attractor. Both methods obtain estimates $T_{\min}^{F=0.99} \approx 0.92$ ms and $T_{\min}^{F=0.999} \approx 0.105$ ms. The estimates from the fully player-generated PGRAPE PS \cup PS are off by less than 2% in both instances.

The SA methods perform significantly worse than the GRAPE methods on this problem because $g \neq 0$. Even the best SA ($n_b = 40$) results do not reach the same fidelity, possibly due to a combination of reduced controllability (low resolution) and the computational penalties associated with $g \neq 0$ as described in Appendix C. The full resolution SA ($n_b = n_t$) also fails to converge almost everywhere, except at T sufficiently larger than $T_{\min}^{F=0.99}$. This reaffirms that the control landscape topography is benign enough that even an inefficient algorithm can find the optimal strategy with appreciable probability at these durations.

A. Optimal strategies—Control clustering

Even without a clustering analysis, the optimal strategy was apparent. Figure 9 shows the mean of all controls within 0.02 of the globally optimal fidelity as a function of $T > 0.2$ ms. Performing clustering (not shown) with $\epsilon = 5$ and $\text{min}_{\text{samples}} = 5$ on the GRAPE PS or GRAPE RS results verifies

that this problem has a simple optimization landscape: 3569 of the controls were associated with a single cluster, whereas 56 controls were unclassified.

For all T , the mean takes an initially (near) maximal control value for an extended period. This physically corresponds to raising the center barrier as much as possible (where the wave function is initially localized, see Fig. 2) and thus providing maximal acceleration to split the condensate into two equal wave packets. At low T the mean control exhibits a bang-bang structure that tapers off as T increases. Near $T = 0.4$ ms a new bang with smoothed edges appears roughly centered on $t/T \sim 0.5$, which is then bimodally split around $T = 0.65$ ms. The mean control then becomes increas-

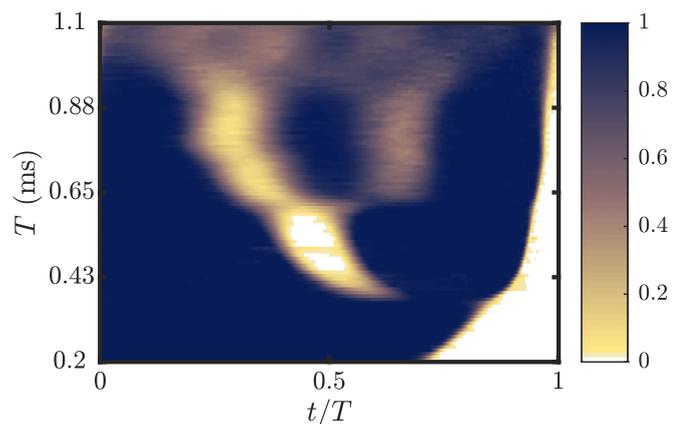


FIG. 9. *Splitting*: mean of (near) optimal solutions $(u_2(t/T))_{\text{opt}}$, corresponding to the height of the potential barrier. For a given T (y axis), the color indicates the mean optimal control value at a given t/T (x axis). Clustering analysis identified only a single cluster.

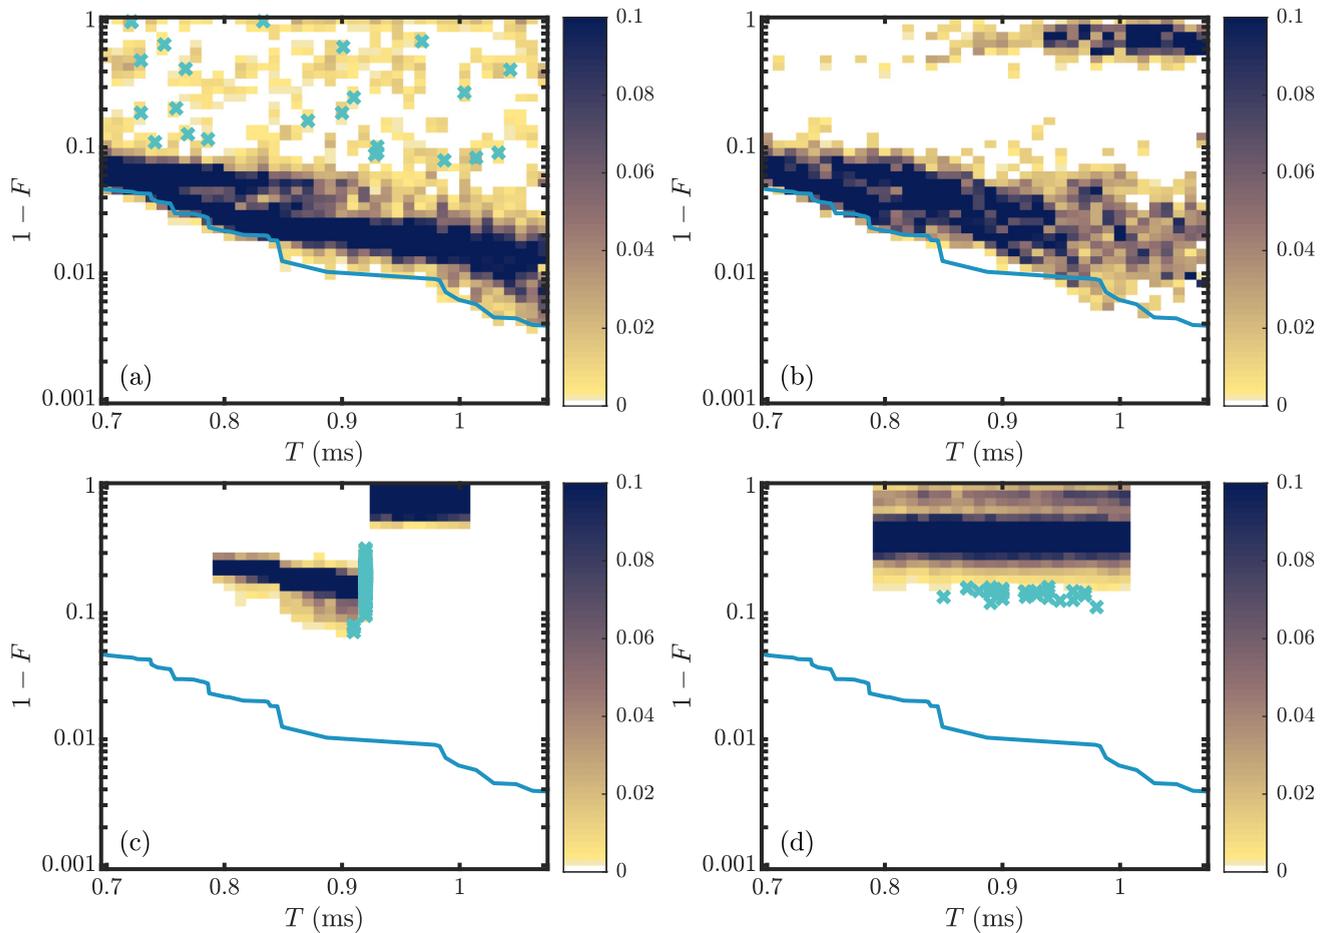


FIG. 10. *Shake Up* solution densities for different methods. (a) GRAPE PS, (b) GRAPE RS, (c) SA RS ($n_b = n_t$), and (d) SA RS ($n_b = 40$). Each density is normalized for every individual T and thus represents an estimate of the probability distribution for obtaining a particular F for a given T ($\mathcal{P}(F|T)$). The reference curve shows the best obtained results for GRAPE PS. Crosses indicate solutions at densities lower than 0.002.

ingly “blurry” indicating a growing departure from bang-bang structures.

VII. SHAKE UP

In this section, we present optimization results obtained in *Shake Up* for the same suite of methods, convergence criteria, and computational resources as in Sec. V.

The solution densities are shown in Fig. 10 and the aggregate results are shown in Fig. 11. The kinks in the reference curve in Fig. 10 can be more apparently understood by looking at the blue dots in Fig. 11 (corresponding to GRAPE PS solutions). The several pronounced, staircase-like plateaus suggest the existence of multiple strategies that are relevant at different duration intervals. These are examined more closely in Sec. VII A and are found to be associated with solution strategies defined by elements of periodic modulation. Each plateau extends over the next, meaning that the now inferior, locally optimal strategy remains a prominent attractor for quite some interval of duration. From looking at the GRAPE solution densities, this evidently makes the new globally optimal strategy much harder to find. For GRAPE PS the density splits when crossing the kinks and one plateau in

particular remains the main attractor, obfuscating the globally optimal strategy that coincidentally exists in the wing of the distribution. However, the density of very low fidelity solutions is sparse and independent of T . This occurs because partial transfers are not difficult to achieve as the initial state density can very quickly and easily be overlapped with one of the lobes of the double-peaked target state density (see Fig. 2) by a small constant displacement. Looking at the aggregate results, we indeed observe that the blue dots are separated from the upper, thick sea of green crosses dots. The abundance of the latter at low-fidelity is due to players terminating their optimization prematurely.

No single method is found to be the best for all T . This is contrary to both *Bring Home Water* and *Splitting*, where GRAPE PS, PGRAPE PS \cup PS and either SA RS or GRAPE RS [49], respectively, found the globally optimal solutions independently. In *Shake Up*, however, PS seeds seem to have the upper hand around $T = 0.85$ ms and after $T = 0.98$ ms, while the RS seeds seem to be better between those durations. The GRAPE RS density provides some nuance to this observation. It shows a broader, less dense distribution of high-fidelity solutions and the addition of many very low fidelity solutions as T increases. In fact, near $T = 1.05$ ms the monotonically best results are

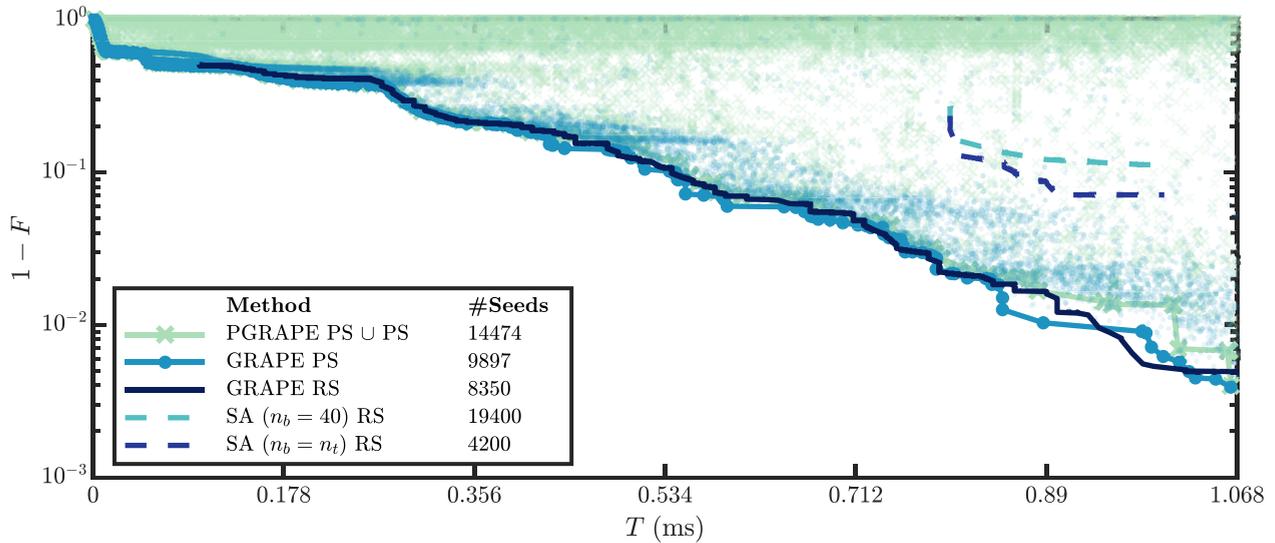


FIG. 11. Aggregate, monotonically best optimization results (lower is better) for several methods in *Shake Up*. Solid lines with (without) symbols show the best results obtained with (without) player influence. The scattered green crosses show all results produced by 266 players seeding on average ~ 37 solutions and optimizing approximately $1/3$ of these for ~ 116 iterations on average. The scattered blue dots show the same, except the optimization is carried out on all player seeds with the computational resources described in the text (i.e., no player influence after seeding). The symbol translucency indicates the density distribution.

due to only a few points in an otherwise empty density region. On the contrary, for GRAPE PS, the same region has a fairly high density and is thus the statistically superior method in the $T = 1.05$ ms regime. (We provide another perspective on these statistical performances in Sec. VIII.) This gives rise to slightly different minimal duration estimates: $T_{\min}^{F=0.99} \approx 0.939$ ms for RS and $T_{\min}^{F=0.99} \approx 0.969$ ms for PS, although the point at $T = 0.887$ ms with $F = 0.9897$ comes very close to tipping this (somewhat arbitrary) balance. Neither method obtained an estimate for $T_{\min}^{F=0.999}$ within the span of durations present in the game. The overall behavior of PGRAPE PS \cup PS shows that players terminated the optimization prematurely in this problem and therefore yields an approximately 7% worse estimate.

For the SA RS methods, the ($n_b = n_t$) and ($n_b = 40$) variants both fail to find any meaningful results. Whereas the *Splitting* control landscape was benign enough to compensate for the computational difficulties associated with $g \neq 0$, this is clearly not the case in *Shake Up*.

A. Optimal strategies—Control clustering

For *Shake Up*, the optimized controls did not possess any readily apparent structure. When individually plotted alongside the corresponding position expectation value of the wave function, however, an oscillatory structure begins to emerge. Subtracting the expectation values, $u(t) - \langle x(t) \rangle$, we observe that these relative controls are dominated by low frequency cosine components. Decomposing the relative controls as

$$c_k = \frac{1}{T} \int_0^T (u(t) - \langle x(t) \rangle) \cos(\pi kt/T) dt, \quad (8)$$

thus yields low-dimensional vectors $\vec{c} = (c_0, \dots, c_5)$ in frequency space (with corresponding number of oscillations $N_k = k/2$) on which we apply clustering. We join the

GRAPE PS and GRAPE RS result sets, selected by the criteria $0.267 \text{ ms} < T < 1.068 \text{ ms}$ and $F > 0.6$, since the player and random seeds were dominant in different regions. We use $\epsilon = 0.1$ and $\min_{\text{samples}} = 250$. These parameters were chosen such that a single cluster is identified per cosine component.

Figure 12 shows the clustering results. Each cluster is labeled by its dominant coefficient corresponding to a half-integer or integer number of oscillations, which is evident from the cluster means when transforming back into real space. When solutions are colored according to their cluster membership in the aggregate plot, we uncover a clear hierarchy of the solution strategies corresponding to the clusters: within the full set of time intervals considered, each strategy is sequentially globally optimal in ascending order of oscillation number with approximately equal interval lengths. As the globally optimal strategy transitions, the now inferior locally optimal strategy remains a relatively broad attractor for an appreciable interval of duration T , leading to the plateaus observed in Fig. 11.

Based on the presented analysis, we do not believe we have found the true best results and associated minimal durations since (i) neither PGRAPE PS \cup PS, GRAPE PS, or GRAPE RS alone produce the best results, and (ii) the solution densities are shifted away from the optimal strategies with high variances. In particular, the GRAPE PS solution at 0.887 ms in Fig. 11 strongly suggests the existence of optimal solutions in this vicinity from the $k = 4$ solution strategy shown in Fig. 11. Their discovery, however, is evidently obfuscated due to the other active, locally optimal solution strategies ($k = 3, 5$) in this region. Based on this analysis one could imagine alleviating this issue by developing a seeding mechanism parametrized to specifically target the $k = 4$ strategy.

This demonstrates that this is the most challenging of the three examined problems in terms of landscape complexity. We have, however, successfully identified the hierarchy of

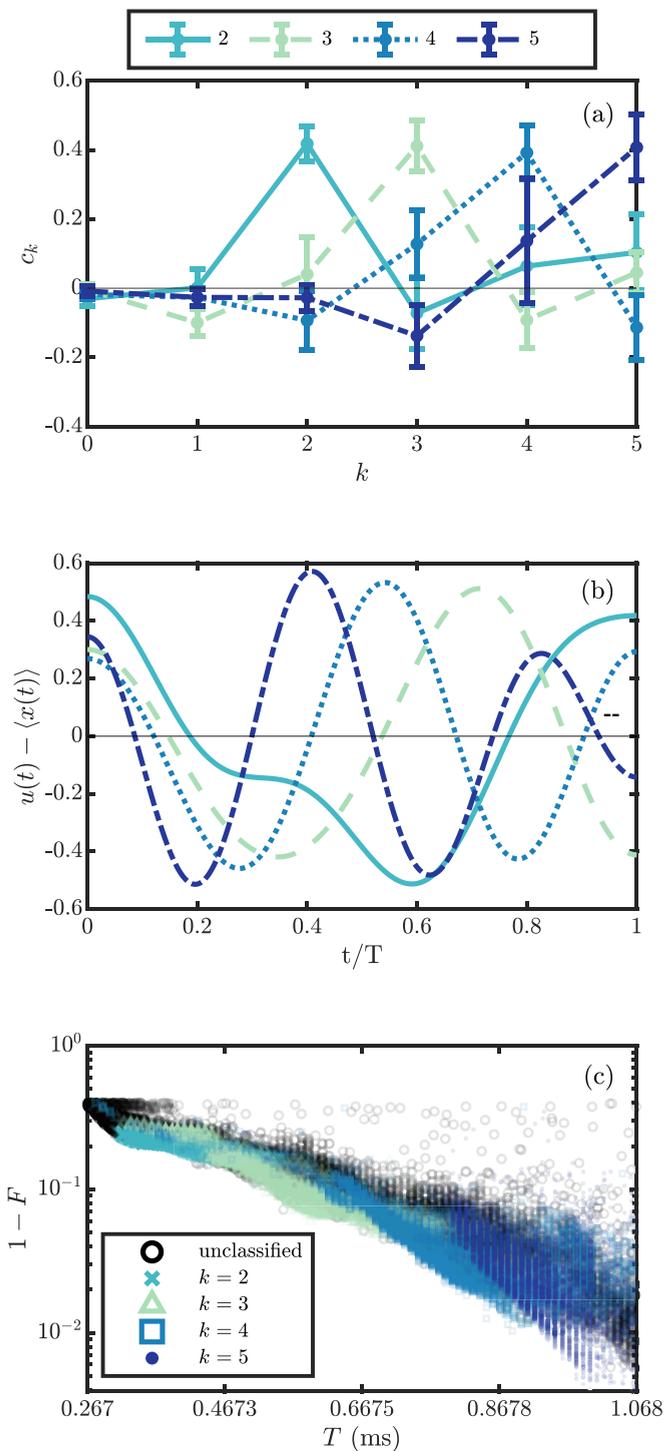


FIG. 12. *Shake Up* clustering results. (a) Component weights c_k with colors corresponding to the different clusters. The bars denote the mean and standard deviation whereas the different lines only help to guide the eye. (b) Cluster means in real space color-coded as above. (c) Fidelity distribution color-coded as above.

solution strategies that could guide further attempts at locating the ultimate minimal durations.

VIII. STATISTICAL PERFORMANCE

Here, we provide an alternative way of characterizing the statistical method performance, which is not based on solution densities, and in addition assess the relative usefulness of the preselection heuristic introduced in Sec. IV. Instead, we compare statistics when only a restricted number of seeds are allowed to be optimized. This emulates either the restriction of computational resources (e.g., no cluster for parallel computation is available) or an increased numerical difficulty of the problem (i.e., operations are more expensive in wall time, allowing fewer seeds to be optimized within the same time frame).

Regardless of either interpretation, we estimate a mean minimal duration $\langle T_{\min}^{\text{fit}} \rangle$ for $F = 0.99$ as a function of the sample size, N_{samples} , on the aggregate results in Figs. 4, 8, and 11. The procedure is based on fitting randomly sampled solutions in a specified interval \mathcal{T}_{\cup} . Concretely, the procedure runs as follows:

- (1) Divide the interval $\mathcal{T}_{\cup} = [0.8, 1.2] \cdot T_{\min}^{F=0.99}$ into 15 subintervals of equal width.
- (2) Sample N_{samples} solutions within \mathcal{T}_{\cup} .
- (3) Within each subinterval, select the solution with the highest fidelity and discard the rest.
- (4) Linearly fit the subset from step 3 with $\log_{10}[1 - F(T)]$, i.e., assume exponential behavior in T , and denote by T_{\min}^{fit} the duration where the fit value corresponds to $F = 0.99$.
- (5) If $T_{\min}^{\text{fit}} \in \mathcal{T}_{\cup}$ (interpolation), then it counts as a success. If it lies outside \mathcal{T}_{\cup} (extrapolation), then it counts as a failure.
- (6) Repeat 2–5 $N_{\text{trials}} = 1000$ times.
- (7) Compute the mean value $\langle T_{\min}^{\text{fit}} \rangle$ over successful trials.
- (8) Compute the empirical success rate (estimated probability of success) $\mathcal{P}(T_{\text{QSL}}^{\text{fit}} \in \mathcal{T}_{\cup}) = N_{\text{successes}}/N_{\text{trials}}$.

Including only the successful trials in $\langle T_{\min}^{\text{fit}} \rangle$ avoids skewing the mean due to extreme outliers (i.e., due to too small negative slope), and the information about these is instead captured in $\mathcal{P}(T_{\text{QSL}}^{\text{fit}} \in \mathcal{T}_{\cup})$. Thus, for a randomly sampled set of solutions, their T_{\min}^{fit} will in the mean be $\langle T_{\min}^{\text{fit}} \rangle$ with probability \mathcal{P} , or fail with probability $1 - \mathcal{P}$. In the following we denote triples of these quantities as $\{N_{\text{samples}}, \langle T_{\min}^{\text{fit}} \rangle, \mathcal{P}\}$.

Figure 13 shows the results for different methods in all three levels. Taking the generic GRAPE RS method as a baseline comparison, one finds that the average behavior of PGRAPE PS \cup PS (in-game player seeding and player optimization) tends to be worse in all levels: it has a comparatively high $\langle T_{\min}^{\text{fit}} \rangle$ and a low \mathcal{P} for small sample sizes, for instance $\{30, 1.15, 0.2\}$ in *Bring Home Water*. An exception occurs beyond $N_{\text{samples}} = 200$ where the players in-game perform better on average, but only in this level.

Thus player methods without *additional* optimization did not yield a superior approach on average. Upon optimization, however, GRAPE PS displays significant statistical performance increases over GRAPE RS across all three levels (implying also that players terminated their optimization prematurely and thus did not themselves realize the full potential of their seeds). This is particularly illuminating for *Shake Up* as neither of the two methods was clearly shown to be better in Sec. VII. There, $\mathcal{O}(10^3)$ seeds were optimized, and at those sample sizes we indeed find coincidence of the two methods' statistical performance capacities in Fig. 13.

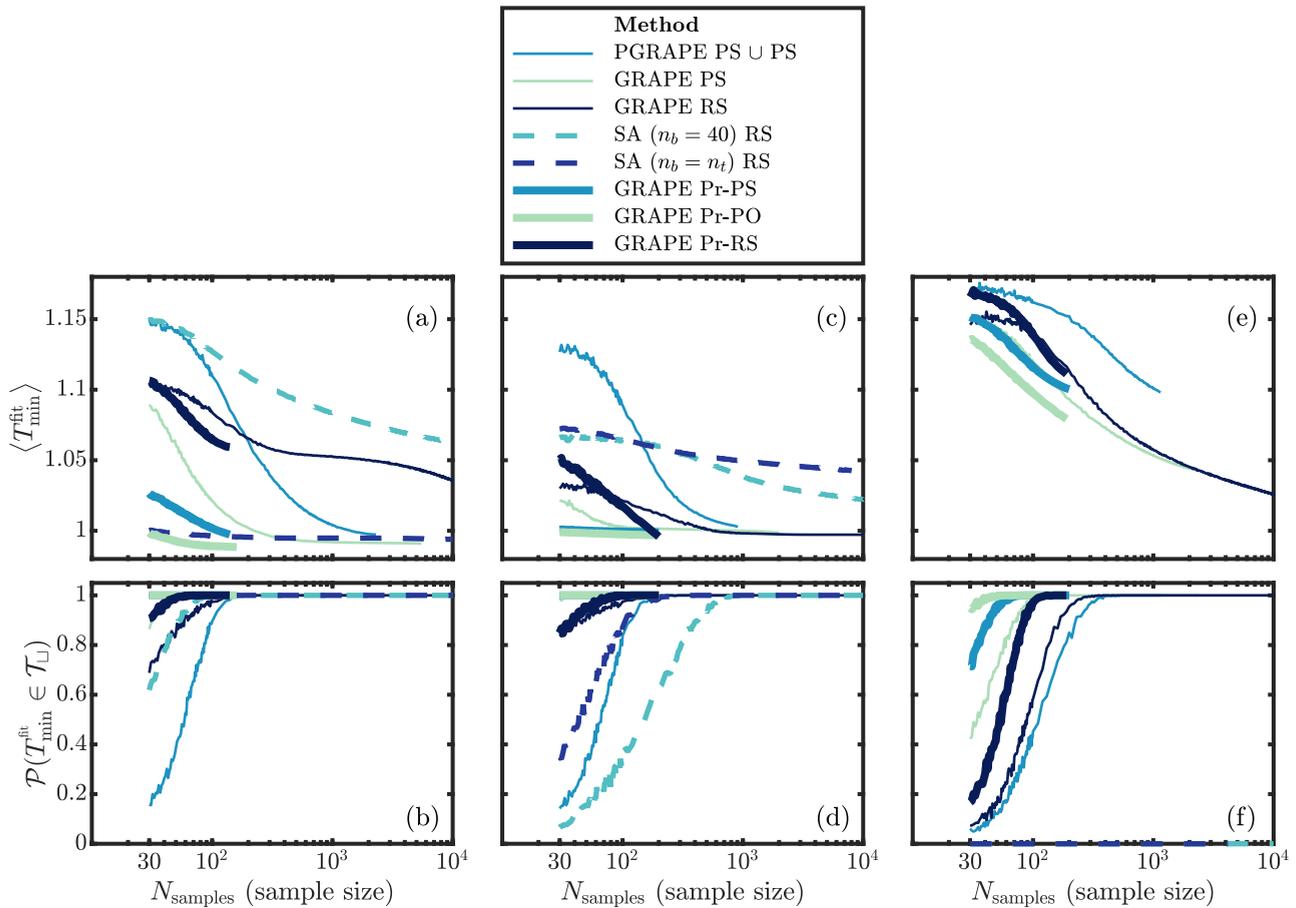


FIG. 13. Estimated $\langle T_{\min}^{\text{fit}} \rangle$ in units of $T_{\min}^{F=0.99}$ and empirical success rate (estimated probability of success) $\mathcal{P}(T_{\min}^{\text{fit}} \in \mathcal{T}_{\cup})$ as a function of sample size. (a), (b) *Bring Home Water*, (c), (d) *Splitting*, and (e), (f) *Shake Up*. Strong statistical performance is signified by low N_{samples} having simultaneously low $\langle T_{\min}^{\text{fit}} \rangle$ and high \mathcal{P} .

However, as N_{samples} is decreased, GRAPE PS becomes statistically superior.

We now turn to the effectiveness of the preselection heuristic. Preselecting (Pr) the best individual 600 PS and PO (i.e., the output from PGRAPE PS) seeds and optimizing these we see a significant shift in statistical performance. The GRAPE Pr-PO and GRAPE Pr-PS methods are observed to require up to several orders of magnitude fewer samples to produce the same $\langle T_{\min}^{\text{fit}} \rangle$ across all three levels compared to the baseline GRAPE RS (without preselection). Moreover, their success rates are significantly higher than any other method for small sample sizes (rivaled only by SA in *Bring Home Water*) and only dips below unit \mathcal{P} in *Shake Up*. For the most extreme case, compare $\{30, 1.135, 0.95\}$ for GRAPE Pr-PO and $\{30, 1.146, 0.07\}$ for GRAPE RS. Obtaining the same $\langle T_{\min}^{\text{fit}} \rangle$ and \mathcal{P} in GRAPE RS occurs at much higher sample sizes ($\{120, 1.135, 0.70\}$, and $\{260, 1.109, 0.95\}$, respectively). Similar trends are seen for the other levels, e.g., in *Bring Home Water* at $N_{\text{samples}} = 30$, where GRAPE Pr-PO obtains $\{30, 0.999, 1\}$ and GRAPE RS obtains $\{30, 1.104, 0.688\}$.

Based on these findings, it might be reasonable to expect that preselection of RS seeds would lead to similar improvements. In this case, however, much smaller relative improvements are observed. This shows that, across all levels, the structure of the best player seeds places them much more

prominently in the (abstract) green region of optimality of Fig. 1 than the best random seeds.

The $\langle T_{\min}^{\text{fit}} \rangle$ reduction gained by increasing the sample size for GRAPE Pr-PO is minimal and quickly saturated in *Bring Home Water* and *Splitting*, while this is not true for *Shake Up*. This reinforces the conclusion that *Shake Up* is the most difficult problem overall considered in this work. The only instance where GRAPE Pr-PO is matched in performance at the same N_{samples} occurs in *Splitting*. Here GRAPE Pr-RS intersects near $N_{\text{samples}} = 200$, the maximum sample size for the preselection-based methods. However, practically the same $\langle T_{\min}^{\text{fit}} \rangle$ can be achieved using just $N_{\text{samples}} = 30$ with GRAPE Pr-PO.

Additionally, the relative difference between Pr-PO and Pr-PS signifies roughly how valuable the in-game player optimizations were in terms of absolute results for a given level (cf. statistics in captions of Figs. 4, 8, and 11). *Bring Home Water* and *Shake Up* show a clear gain while *Splitting* is nearly unaffected, lending itself again to the interpretation that it has the least difficult landscape topography.

The SA ($n_b = n_t$) method performs very well in *Bring Home Water* (with the caveats described in Sec. V also discussed shortly). The SA ($n_b = 40$) method is unsurprisingly seen to have the unequivocally worst overall statistical performance, even in the linear case, $\langle T_{\min}^{\text{fit}} \rangle$ scales comparatively

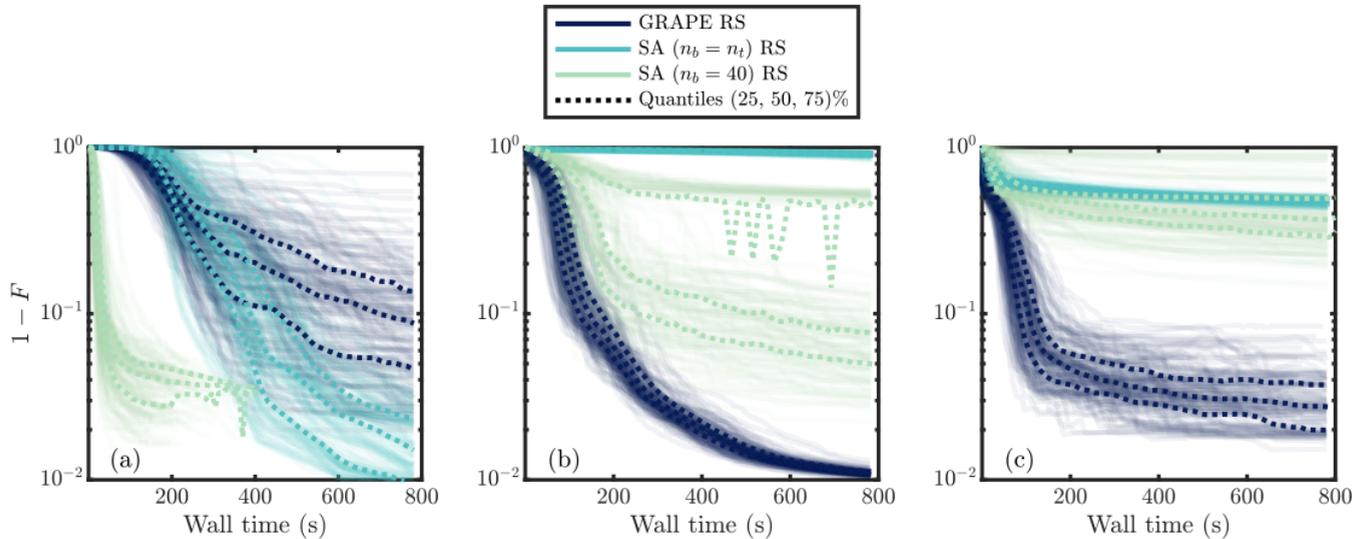


FIG. 14. Results (lower is better) for GRAPE (dark blue), SA ($n_b = n_t$) (turquoise), and SA ($n_b = 40$) (mint green) using the same 100 RS seeds at $T_{\min}^{F=0.99} \approx 0.0973, 0.92, 0.89$ ms for the three levels, respectively. (a) *Bring Home Water*, (b) *Splitting*, and (c) *Shake Up*. Each line denotes an optimization trajectory as a function of wall time. The translucency is related to the density distribution and the three dotted lines indicate the 25%, 50%, and 75% quantiles for all active optimizations.

poorly with N_{samples} . For *Shake Up* both these methods fail ($\mathcal{P} = 0$) for all N_{samples} . The failure is, again, less severe for *Splitting*.

A. Algorithmic run time

Observations on end-of-optimization results in Figs. 4, 8, and 11 do not account for the associated computation times, except in the time-out termination condition.

Figure 14 shows the optimization trajectories for RS methods as a function of wall time for the same 100 seeds at or near $T_{\min}^{F=0.99}$ in each level [50]. These are separate from the previous results, and each seed had in this case a maximum optimization time of roughly 13 min. The quantile statistics consider only ongoing optimizations, i.e., they do not include solutions converged at earlier times.

In *Bring Home Water*, we clearly see that SA ($n_b = 40$) rapidly improves the fidelity initially, lending merit to its usefulness (as discussed in Appendix C). This is owed to the speed at which a full iteration can be carried out, i.e., where all parts of the control domain are adjusted once. The mean time per complete iteration is 14.2 ± 1.4 s. Further speed up can be achieved as pointed out in Appendix C. However, progress stagnates and terminates before $F = 0.99$ due to the reduced resolution. A better choice for the heuristic resolution parameter $40 < n_b < n_t$ would address this issue. The full resolution SA ($n_b = n_t$) does not suffer from stagnation, but progress is initially orders of magnitudes slower since control values at successive times are completely uncorrelated and a full iteration takes much longer. The mean time per complete iteration is 219 ± 31 s. GRAPE completes an iteration in 0.27 ± 0.03 s and performs about the same as SA ($n_b = n_t$) until around 400s. Beyond this point SA ($n_b = n_t$) dominates in the mean, reflecting that GRAPE RS finds only the inferior, locally optimal strategy with less efficiency. Recall that, as opposed to GRAPE, SA optimizes only $u_1(t)$ and is not subject to derivative

regularization (smoothness criterion) [51] of the control, and the algorithm therefore effectively solves an easier problem in the current implementation. Inclusion of both points would likely lead to a doubling in computation time (i.e., stretching the turquoise and green lines by a factor 2 on the x axis in Fig. 14) to obtain similar results to the ones presented, making them more or less coincidental with GRAPE RS near 800 s. However, we do not believe this would drastically change the conclusions in terms of the overall performance in Sec. V.

In *Shake Up*, SA ($n_b = 40$) and SA ($n_b = n_t$) complete an iteration in 159 ± 13 s and >800 s, respectively, but never reach high fidelities. GRAPE completes an iteration in 0.76 ± 0.08 s and fares comparatively much better. Similar numbers are found in *Splitting*.

IX. DISCUSSION

We have presented *Quantum Moves 2*, a citizen-science game in which players act as seeding mechanisms and initial optimizers for quantum optimal control problems. Selecting three distinct problems in the game for analyses, we applied different optimization methods (combinations of algorithms and seeding strategies) to these. For each problem, we examined the (a) respective methods' results and efficiencies, (b) optimal solution strategies, (c) overall problem structures, and (d) statistical performance capacities through a separate random sampling procedure. Here we summarize the findings.

In *Bring Home Water*, a single-particle problem, we identified two solution strategies (front- and back-swing) characterized by an exponentially widening gap between them. Using the same resources and a gradient-based algorithm (GRAPE), the player-infused seeds uncover both strategies efficiently whereas the random seeds only find the inferior, locally optimal strategy at durations relevant for high-fidelity transfers. Imposing increasingly more structure *a posteriori* from physical insight on the random seeds allows

discovery of the globally optimal strategy, albeit with significantly reduced overall efficiency. Employing a gradient-free stochastic ascent algorithm (SA) with random seeds and a strong enforced behavior on one of the controls (tweezer maximally deep at all times), we find both strategies with high efficiency. The success of SA is explained by a combination of the optimization landscape's structure (many, small local traps and a few, sufficiently broad optimal attractors), the algorithm's search methodology (allowing it to escape the abundant local traps), and the linearity of the equations of motion (reducing algorithmic complexity). On an optimization algorithmic level, the comparatively reduced efficiency of the gradient-based algorithm is due to its inability to escape these traps. Thus, the efficiency is due to the algorithm and not the seeding strategy; using SA with player seeds is expected to yield results at least as good as SA with random seeds.

In *Splitting*, a BEC problem, we identified a single solution strategy. This strategy was found efficiently by both the player-infused and randomly seeded gradient-based methods. The gradient-free algorithms with random seeds fail due to the nonlinearity of the equations of motion (resulting in increased algorithmic complexity), with the exception that one variant is somewhat successful in the limit of large durations—the optimization landscape is thus so simple that even a numerically inefficient algorithm can possibly discover the globally optimal strategy in this regime with the allotted resources.

In *Shake Up*, another BEC problem, we identify four solution strategies involving low frequency oscillations around the BEC center-of-mass. Each strategy, individually characterized by a dominant half-integer number of oscillation periods, is globally optimal at different durations in order of lowest to highest number of oscillations. Outside of their respective regions, each strategy remains a broad locally optimal attractor, leading to plateaus in the optimization results. Based on these results, neither of the gradient-based methods are found to be the best on their own, but each exhibits exclusive regions of dominance. In this instance, the gradient-free algorithms completely fail. From these observations, we conclude that the optimization landscape is very complex and we do not believe we have found the ultimately best results. Further efforts could, however, be guided by the identified strategies.

Through random sampling, we then found that the statistical performance of player seeds was always better than random seeds upon cluster optimization in all the studied levels when restricting the sample size (equivalent to increased numerical difficulty or restricted available computational resources). This could be considered a tiebreaker for the methods' similar absolute performances in the *Shake Up* problem. Additionally, preselecting the player solutions based on best initial fidelity significantly increases their relative performance when optimized, whereas the same preselection procedure improves the random seeds to a lesser comparative extent. This can be interpreted as the best player solutions being more likely to be located in the green region of optimality in Fig. 1. The benefit diminishes gradually as the sample size is increased and the overall optimization landscape more densely explored. This echoes previously drawn conclusions—given enough resources (correspondingly a sufficiently large sample size) most combinations of algorithms and seeding strategies achieve similar absolute results,

given they at least partially cover the green region of optimality.

As outlined in the introduction, we attribute the observed advantage of player seeds over the random seeds to the use of human common sense which presumably draws upon analogies with classical fluid behavior. It remains to be seen if this advantage could be maintained in optimization of quantum processes involving decoherence (mixed states) or entanglement or even simultaneous evolution of different possible input states, such as in quantum gates. In the citizen science context, we believe that the answer depends on how strong these effects are (their effect on the state evolution) and, crucially, the degree to which relevant information about the state can be represented in an “intuitive” and meaningful graphical manner. For example, entanglement between motional degrees of freedom, as in, e.g., two interacting particles in 1D, could be visualized as a two-dimensional density distribution (one dimension per particle) represented as an image with shade/color coding of the values or even as a three dimensional surface. Further, the multiple orthogonal states representing unitary evolution operators, in, e.g., quantum gates, could be displayed in parallel (on top or above each other) while a common potential affecting them is moved by the player. Nevertheless, ample citizen science questions about single pure state control remain unanswered and are prospective candidates for immediate follow up works.

We now return to the two initial citizen science related questions Q1 and Q2 discussed in the introduction. A main feature of *Quantum Moves 2* was an in-game optimization button enabling players to store and optimize candidate solutions on their local device. This clearly underscored the player's role in the search for overall, global features in solutions (that is, locating the green region of optimality in Fig. 1), whereas fine-grained, local optimization could be left to the optimization algorithm. This certainly supports a sequential, “one-off,” player-computer interaction. However, with supporting tools (such as replay and the *ghost* feature, see Appendix A) the game also enables more intertwined, hybrid human-computer interactions in which players gain insight by examining the output of the computer optimization and can thereby improve their search for promising features and heuristics. Further study of this will be left for future work.

Figures 4, 8, and 11 demonstrate that the method of player-seeding with player-invoked local-device optimization (PGRAPe PS \cup PS) performs roughly on par (in terms of best achieved results) with the best performing fully algorithmic approaches under consideration in each problem. For the two hardest challenges, *Bring Home Water* and *Shake Up*, this method outperforms the randomly seeded GRAPe and SA, respectively. Thus, we suggest, as proposed in Q1, that it could indeed make sense to develop a framework, like Foldit has done for protein folding, in which the solution of quantum problems are outsourced to the general population. In Q2, we ask if the game-based approach could actually yield a computational advantage. A full answer to this question should entail comprehensive comparison to the best possible expert-driven optimization. Here, we take first steps in that direction with a baseline benchmark comparison to off-the-shelf optimization with initialization that is as heuristic-free as possible. We find that GRAPe with player seeds (GRAPe PS) is the only method

that is roughly optimal across all three levels, the exception being a small window in *Shake Up*. These results should only be understood as a necessary baseline study and a first demonstration for further exploration, and they should not be taken as a guarantee that player-based seeding is advantageous when comparing to increasingly complex algorithmic strategies. However, outside of immediate performance capabilities, player-generated data may show additional potential because, similar to machine learning-generated data [52–54], they constitute a means for researchers to address problems that is not influenced by any expert biases. This can inform the extraction of heuristics and insights that can subsequently be understood, utilized, and expanded upon by the domain expert. An indication of this was seen in the most complex challenge, *Shake Up*, in which player and randomly seeded methods were globally optimal at different durations, as they probed different parts of the interleaved optimal strategies.

ACKNOWLEDGMENTS

We acknowledge all the players of *Quantum Moves 2* and support from European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie QuSCo Grant Agreement No. 765267, the ERC, H2020 Grant No. 639560 (MECTRL), and the John Templeton and Carlsberg Foundations. The numerical computer cluster optimization results presented in this work were obtained at the Centre for Scientific Computing, Aarhus [55].

APPENDIX A: INTERFACE

The player is faced with two main views within a level: the dynamic view and the graph view. A simplified version of these is shown in Fig. 15. In the dynamic play view, the player generates a seed by dragging the round cursor (lower solid turquoise dot) to the round target point (upper transparent turquoise dot). The position of the cursor is linearly mapped by f_1 and f_2 [Eqs. (4)] to the instantaneous control function values, $\{u_1(t), u_2(t)\} = \{f_1(x_{\text{cursor}}(t)), f_2(y_{\text{cursor}}(t))\}$, which dynamically alter the potential (green). The instantaneous wave function density (dark green solid with red line) propagates in the current potential (green line) and the goal is to match the target state density (yellow line) without excess excitation at the end, as the target states are stationary. The path traced by the cursor can have any shape, but it is clamped to remain within the control boundaries (turquoise dashed bounding box). In the graph view, the solution is indicated by a dot on an $F(T)$ graph corresponding to its final fidelity and transfer duration. Note that the fidelity axis is nonlinear and the T axis is normalized by an approximate reference $T_{\min}^{F=0.99}$ bound found by a conventional optimization before the launch of the game. The T axis is divided into 12 blocks of equal size. Each block contains three green lines of increasing fidelity. Points are accrued within each block by placing a solution with the highest fidelity possible relative to these. The uppermost green line is the *challenge curve*, which are reference results found by conventional GRAPE methods. The remaining lower green lines are motivational game elements without any scientific significance.

In either view the player can click a button to start or stop a GRAPE optimization of the currently selected solution. During

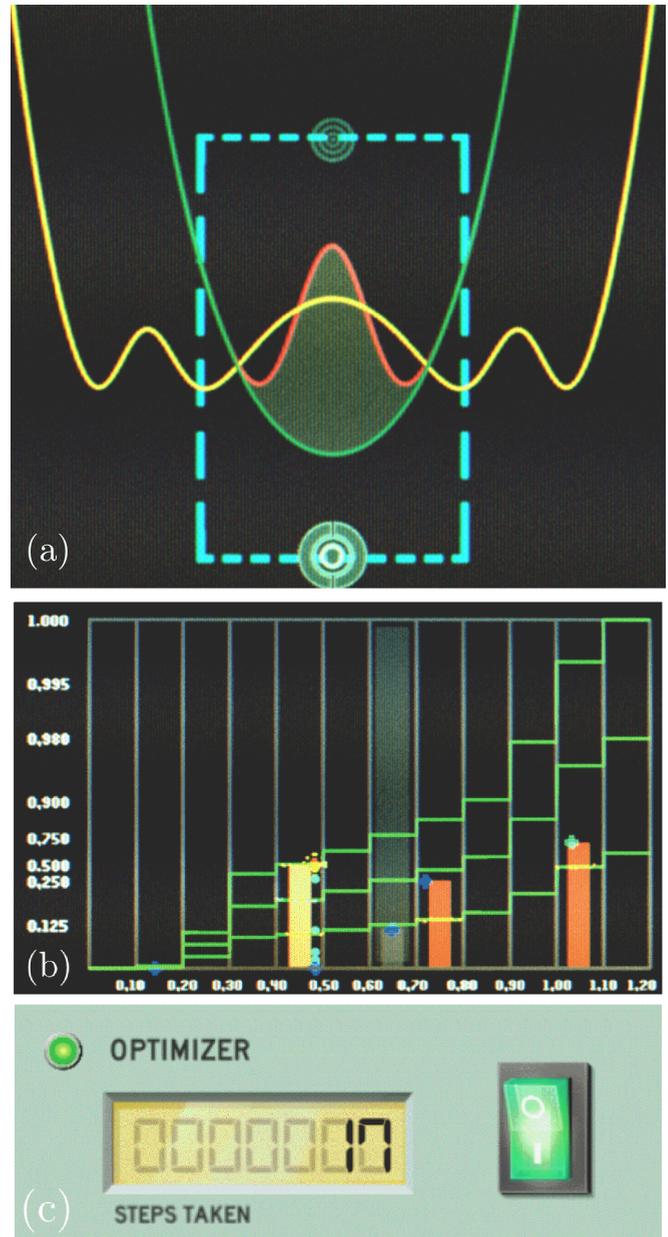


FIG. 15. Partial screen shots of the two main views in *Quantum Moves 2*. Much of the UI has been hidden for simplicity. (a) Dynamic play view with dynamic potential (green), instantaneous wave function density (red line with dark green area), and target state (yellow outline). (b) Graph view. Each solution corresponds to a point in the graph. The duration axis is normalized by an approximate $T_{\min}^{F=0.99}$ bound and is divided into 12 blocks. The uppermost green line (the challenge curve) corresponds to the best solution found by conventional methods prior to the launch of the game. Any point can be selected for optimization. (c) The player can invoke the embedded optimization algorithm (PGRAPE) by toggling the switch.

optimization the point will climb in the graph view and the final wave function density will become increasingly similar to the target outline in the dynamic view. The player also has other tools available (not shown) such as replaying the time evolution of a solution in real time or setting a solution as a *ghost* that plays along during dynamic play. This can

help the player determines strategies by recognizing certain high-level characteristics inspired by the solutions found by the optimizer, or players can try to mimic certain behaviors.

APPENDIX B: NUMERICS AND CONTROL PROBLEMS

The state transfers $\psi_0 \rightarrow \psi_{\text{tgt}}$ considered in this paper are simulated and solved in the QEngine [51], our C++ software package for quantum optimal control. On top of this library is an interfacing code layer exposing functionality and facilitating interoperability with the Unity C# game code.

The system dynamics is governed by the Schrödinger equation $i\hbar\dot{\psi} = \hat{H}\psi$, where the Hamiltonian is parametrized by up to two control parameters, $\hat{H} = \hat{H}(u_1(t), u_2(t))$. More conveniently we can write the state at final time T as $|\psi_T\rangle = \hat{U}|\psi_0\rangle$ where $\hat{U} = \mathcal{T} \exp(-i \int_0^T \hat{H}(u_1(t'), u_2(t')) dt')$ is the time-evolution operator and \mathcal{T} denotes time-ordering. The goal is thus to find a set of optimal controls $\{u_1^*(t), u_2^*(t)\}$ that implements $\psi_T = \psi_{\text{tgt}}$ (up to a global phase) or, equivalently, maximizes the fidelity $F = |\langle \psi_{\text{tgt}} | \psi_T \rangle|^2 = |\langle \psi_{\text{tgt}} | \hat{U} | \psi_0 \rangle|^2 \leq 1$.

The continuous time parameter is discretized on a regular grid $t \in \{t_1, t_2, \dots, t_j, \dots, t_n\}$ with sufficiently small spacing δt , and the spatial dimension is similarly discretized. We may then approximate $\hat{U} \approx \prod_{j=1}^{n-1} \hat{U}_{t_j} = \hat{U}_{t_{n-1}} \dots \hat{U}_{t_1}$ where $\hat{U}_{t_j} = \hat{U}[u_1(t_j), u_2(t_j)]$ is the local time evolution operator, $\psi_{t_{j+1}} = \hat{U}_{t_j} \psi_{t_j}$. The control functions are bounded

$$\{u_1^{\min}, u_2^{\min}\} \leq \{u_1(t), u_2(t)\} \leq \{u_1^{\max}, u_2^{\max}\}, \quad (\text{B1})$$

and the control values fixed at $t = 0$ and $t = T$, depending on the level. For optimizations only, controls and bounds are linearly transformed into $0 \leq u_i(t) \leq 1$ for $i = 1, 2$.

Units: For numerical purposes we obtain nondimensionalized [51] working equations such that effectively $\hbar = m = 1$ and

$$i\dot{\psi} = -\kappa \frac{\partial^2 \psi}{\partial x^2} + V\psi + g|\psi|^2\psi, \quad (\text{B2})$$

where κ is a constant that can be used to gauge the units. SI and simulation units are related by $\alpha_{\text{SI}} = \mu_{[\alpha]} \alpha_{\text{sim}}$ where $\mu_{[\alpha]}$ is the chosen unit for the dimension of quantity α and α_{sim} is the dimensionless number entering, e.g., Eq. (B2) (the subscript is henceforth omitted for brevity and quantities written without units imply simulation values). We take the atomic species to be rubidium-87 atoms such that the unit of mass is $\mu_{\text{mass}} = m_{\text{Rb}} = 87$ amu and we take the energy unit to be $\mu_{\text{energy}} = \hbar/\mu_{\text{time}}$. Fixing two elements of the triplet $\{\kappa, \mu_{\text{length}}, \mu_{\text{time}}\}$ determines the remaining element to produce Eq. (B2).

Bring Home Water: The units are fixed by

$$\kappa = 0.5, \quad \mu_{\text{length}} = 532 \text{ nm}, \quad (\text{B3a})$$

$$\Rightarrow \mu_{\text{time}} = 2\mu_{\text{mass}}\kappa\mu_{\text{length}}^2/\hbar = 0.38731 \text{ ms}. \quad (\text{B3b})$$

Time steps are of size $\delta t = 3.5 \times 10^{-4}$ and $x \in [-3, 3]$ with $n_x = 256$ grid points. The sum of the tweezer potentials, initial control values, and bounds are given by

$$V(u_1, u_2) = u_2(t) \exp\left(\frac{-2[x - u_1(t)]^2}{\sigma^2}\right) + A \exp\left(\frac{-2(x - x_0)^2}{\sigma^2}\right), \quad (\text{B4a})$$

$$\{u_1(0), u_2(0)\} = \{u_1(T), u_2(T)\} = \{-1, -130\}, \quad (\text{B4b})$$

$$\{-2, -150\} \leq \{u_1(t), u_2(t)\} \leq \{2, 0\}, \quad (\text{B4c})$$

where $\sigma = 0.5$, $x_0 = 1$, $A = -130$ (≈ -53.42 kHz h).

The frequency of the harmonic approximation when the control tweezer is maximally deep and centered on x_0 is $\omega = \sqrt{-4(A + u_2^{\max})/(m_{\text{Rb}}\sigma^2)} \approx 66.93$ (≈ 436.8 kHz).

Shake Up. The units are fixed by

$$\mu_{\text{length}} = 1 \mu\text{m}, \quad \mu_{\text{time}} = 1 \text{ ms}, \quad (\text{B5a})$$

$$\Rightarrow \kappa = \hbar\mu_{\text{time}}/(2\mu_{\text{mass}}\mu_{\text{length}}^2) = 0.36537. \quad (\text{B5b})$$

Time steps are of size $\delta t = 1 \times 10^{-3}$ and $x \in [-2, 2]$ with $n_x = 256$ grid points. The atom chip potential (see Ref. [51] and references therein), initial control values, and bounds are given by

$$V(u_1) = \sum_{r=2,4,6} p_r [x - u_1(t)]^r, \quad (\text{B6a})$$

$$u_1(0) = u_1(T) = 0, \quad -1 \leq u_1(t) \leq 1, \quad (\text{B6b})$$

with coefficients and nonlinear coupling strength

$$p_2 = 65.8392, \quad p_4 = 97.6349, \quad (\text{B7})$$

$$p_6 = -15.3850, \quad g_{\text{1D}} = 1.8299. \quad (\text{B8})$$

Splitting. The units are the same as in Eqs. (B5a)–(B5b) with the addition of $\mu_{\text{magnetic}} = 1 \text{ G} = 10^{-4} \text{ T}$. Time steps are of size $\delta t = 1 \times 10^{-3}$ and $x \in [-3.5, 3.5]$ with $n_x = 256$ grid points. The atom chip potential for $\omega = 2\pi 1.26 \text{ MHz}$ after nondimensionalization (see also Refs. [36,56]), initial control values, and bounds are given by

$$V(u_1) = p \sqrt{(B_S(x) - B_\omega)^2 + \left\{ \frac{[0.5 + 0.3u_2(t)]}{2B_S(x)} B_I \right\}^2}, \quad (\text{B9a})$$

$$u_2(0) = 0, \quad u_2(T) = 1, \quad 0 \leq u_2(t) \leq 1, \quad (\text{B9b})$$

In this form, $p = (\mu_{\text{magnetic}} \mu_{\text{Bohr}} m_F g_F) / \mu_{\text{energy}} = 8794.1$ is an overall factor (μ_{magnetic} has been factored out from under the square root), and $m_F = 2$, $g_F = 1/2$ are the internal hyperfine state and Landé factors, respectively. Additionally, $B_\omega = 0.9$ and $B_S(x) = \sqrt{(\text{Gr}x)^2 + B_I^2}$, where $\text{Gr} = 0.2$ is a magnetic field gradient and $B_I = 1$. We take $g_{\text{1D}} = 1.8299$.

APPENDIX C: ALGORITHMS AND RESOURCES

Here we expand on the algorithms described in Sec. IV and the optimization resources employed in Secs. V–VII.

Outside the game, optimizations were performed in large batches on a computer cluster. Each seed was optimized until it met any of the pertinent termination criteria listed in the

following subsections. The initial minimal allotted time was roughly 13 min per seed. To maximize resource use, excess time for seeds that terminated due to the other criteria was divided equally amongst the remaining seeds and added to their minimal times.

Inside the game, PGRAPE ran locally on the player devices until manually stopped by the player, convergence (step size below 10^{-7}), or $F = 0.999$, incurring no computational cost for us.

1. GRAPE

Our variant of the standard GRAPE methodology includes derivative regularization ($\gamma = 10^{-6}$) and boundary cost terms ($\sigma = 2 \times 10^3$, see Ref. [51] for details).

The pertinent stopping conditions for GRAPE are: exceeding the allotted optimization wall time (minimum ~ 13 min), exceeding the fidelity threshold ($F \geq 0.999$), or subceeding the line search threshold ($\alpha^k < 10^{-7}$). In the case of PGRAPE, the wall time condition is replaced by an in-game button as described in the text.

A single GRAPE iteration requires calculating the gradient ($2n_t$ time steps), calculating the L-BFGS quasi-Newton search direction (negligible cost), and performing a line search ($\sim 5n_t$ time steps on average in our implementation), for a total time step cost of $7n_t$. This algorithm is fully exploitative (local) with an exploratory (global) component induced through multistarting [30].

2. (Discrete) stochastic ascent

This section reviews and expands the analysis of the stochastic ascent algorithm from Ref. [20]. In the following, we assume a single control parameter for simplicity. The time axis is segmented into n_b bins of equal width w which have the same control value, such that $n_t = wn_b$. For example, for $w = 3$

$$\vec{u} = \underbrace{\{u_{t_1}, u_{t_2}, u_{t_3}\}}_{u_{b_1}}, \underbrace{\{u_{t_4}, u_{t_5}, u_{t_6}\}}_{u_{b_2}}, \dots, \underbrace{\{u_{t_{n_t-2}}, u_{t_{n_t-1}}, u_{n_t}\}}_{u_{b_{n_b}}}.$$

The propagator for the first bin is $\hat{U}_{b_1} = \hat{U}_{t_3}\hat{U}_{t_2}\hat{U}_{t_1}$ and the pattern continues for the other bins [57]. We then allow u_{b_k} to assume only values from a predefined discrete set $\Omega = \{u_d\}_{d=1}^{n_d}$ (this choice is discussed at the end of the section). Using $n_d = 128$, these values are linearly spaced from the lower to the upper control boundary, see Eq. (B1).

Updating u_{b_k} is done by exhaustively computing the fidelity for all possible values in Ω and setting u_{b_k} corresponding to the maximal value,

$$u_{b_k} \leftarrow \underset{u_{b_k} \in \Omega}{\operatorname{argmax}} F(u_{b_1}, \dots, u_{b_{k-1}}, u_{b_k}, u_{b_{k+1}}, \dots, u_{b_{n_b}}),$$

while keeping the other control values fixed (discrete coordinate ascent [58]). When u_{b_k} has been updated, it is not chosen for further updates until all the remaining points have also been updated. Updating all points once constitutes an iteration and the sequential control update order is stochastic within each of these. As is, bandwidth limitations are only imposed by the choice of n_b , but one could easily accommodate a derivative regularization term as in GRAPE. The pertinent stop-

ping conditions are: exceeding the allotted optimization wall time (minimum ~ 13 min), exceeding the fidelity threshold ($F \geq 0.999$), or when the algorithm achieves no gain in fidelity by changing any of the control values.

Exhaustive evaluation of the fidelity for bin k can be sped up for the linear Schrödinger equation ($g = 0$),

$$F(u_{b_k}) = |\langle \psi_{\text{tgt}} | \psi_{t_{n_t}} \rangle|^2 = |\langle \chi_{b_{k+1}} | \hat{U}_{b_k} | \psi_{b_k} \rangle|^2, \quad (\text{C1})$$

since the forward-propagated $|\psi_{b_k}\rangle = \prod_{j=1}^{k-1} \hat{U}_{b_j} |\psi_0\rangle$ and backward-propagated $|\chi_{b_{k+1}}\rangle = \prod_{j=n_b}^{k+1} \hat{U}_{b_j}^\dagger |\psi_{\text{tgt}}\rangle$ vectors for all times up to and after bin k , respectively, only need to be calculated and cached once per bin update. Upon finishing the evaluation, the update is applied to the control and forward-propagated state. Additionally, the matrix representation of the time-evolution operators \hat{U}_d corresponding to every element in Ω can be precomputed and cached in memory, changing the time-stepping method to a single matrix-vector multiplication [20] instead of the Fourier split-step method.

The first control value update requires ($n_t + n_d$) time steps after which the forward/backward vectors have been initialized and cached. Calculating the new forward/backward vectors and updating for a subsequent bin at k' requires only $w|k - k'|$ time steps when reusing the old vectors. If $k < k'$, then the ψ cache is updated and the χ cache otherwise.

We may write the average time step distance $w \langle |k - k'| \rangle = w\rho n_b = \rho n_t$ where $\rho \approx 1/3$ is found empirically. Performing subsequent updates thus costs $(n_t/3 + n_d)$ time steps when averaged over all bins. The average number of time steps required to complete a full iteration is thus $n_b(n_t/3 + n_d)$, except the first iteration which costs an additional $2n_t/3$ due to forward/backward vector cache initialization [59].

In the nonlinear case $g \neq 0$, the explicit state dependence has severe consequences for the algorithm's feasibility. First, the time evolution operators \hat{U}_j cannot be precomputed since they depend on ψ , which changes as the control changes. Second, it does not make sense to maintain a cache for backward-propagated vectors; altering the control at t_k changes the ψ state trajectory from k to n_t and the backward-propagated vectors depend on these in the nonlinear case. In this case one may just as well evaluate the fidelity at t_{n_t} using the first equality in Eq. (C1). The ψ cache only needs to be updated when $k < k'$, yielding $w \langle |k - k'| \rangle_{k < k'} = wn_b/6 = n_t/6$. From there, evaluating the fidelity of a single $u_{b_k} \in \Omega$ requires $w|n_b - k|$ time steps which must be done for all n_d elements. Averaging over a full iteration yields $w \langle |n_b - k| \rangle = wn_b/2 = n_t/2$. Consequently, the number of time steps needed to update a single point u_{b_k} changes as $n_t/3 \rightarrow n_t/6$ and $n_d \rightarrow n_d n_t/2$.

Thus roughly an additional $n_b n_d n_t/2$ time steps must be performed when $g \neq 0$, each of which has an increased computation time because \hat{U}_d cannot be cached. The differences between the two cases is summarized in Table I.

The speed with which the stochastic (coordinate) ascent operates comes at the cost of not being able to perform correlated, simultaneous control value updates between bins. This could easily be remedied, in principle, by updating n_p bins instead of just a single one. However, such an approach is untenable for the discrete version even for small $n_p > 1$ if one desires exhaustive search: the update cost would then depend

TABLE I. Summary for $g = 0$ and $g \neq 0$ for SA.

	Linear ($g = 0$)	Nonlinear ($g \neq 0$)
Caching	\hat{U}_d, ψ, χ	ψ
Time evolution	Matrix \times vector	Fourier split-step
Time steps/iter	$n_b(\frac{1}{3}n_t + n_d)$	$n_b(\frac{1}{6}n_t + \frac{1}{2}n_t n_d)$

on the largest index distance between the parameters and an exponential number of discrete combinations $n_d^{n_p}$.

The fast linear evaluation ($g = 0$) for the SA algorithm is independent of the choice to restrict control values to a discretized set of values. For example, it would be straightforward to use the same fast evaluation methodology with derivative-based methods and perform line searches. Such a change in update rule shifts the exploration versus exploitation tradeoff: the discrete version is in a sense fully exploratory (the search is globally exhaustive), but only along one axis at a time. Obviously, discretization and fixing the remaining axes produces a reduced representation of the underlying control landscape with respect to which the discrete version exhibits a mix of global- and local search properties. Abandoning

discretization and performing line searches turns the algorithm into a fully exploitative one, but again only along a single axis. This is the more standard version of coordinate ascent [60]. In our implementation, line searching usually requires about $5 \ll n_d$ objective evaluations [61], allowing potentially orders of magnitude fewer time steps ($n_d \rightarrow 5$ in Table I) when close to an optimum. Although such an approach does not allow caching of the unitary time evolution operator since the controls can take any value, the aforementioned benefits should more than compensate for this during the local adjustment phase. In this setting, however, the advantageous convergence rates associated with derivatives and adaptive step sizes is only with respect to the chosen axis—there are no theoretical guarantees for convergence in the full dimensional landscape [60]. In light of these observations, it would be interesting to combine the three methodologies with handover techniques, for example, starting with the most global algorithm and ending with the most local one,

discrete SA \rightarrow gradient SA \rightarrow GRAPE.

The performance of such a combination would be interesting to try on the different seeding strategies discussed in the main text and is a potential subject of future work.

- [1] D. Heaven, Why deep-learning AIs are so easy to fool, *Nature* **574**, 163 (2019).
- [2] S. Jiang, S. Lu, and D.-L. Deng, Vulnerability of machine learning phases of matter, [arXiv:1910.13453](https://arxiv.org/abs/1910.13453) (2019).
- [3] G. Marcus, The next decade in AI: Four steps towards robust artificial intelligence, [arXiv:2002.06177](https://arxiv.org/abs/2002.06177) (2020).
- [4] G. J. Fletcher, Psychology and common sense, *IEEE Eng. Manage. Rev.* **14**, 30 (1986).
- [5] S. Carruthers and U. Stege, On evaluating human problem solving of computationally hard problems, *J. Probl. Solving* **5**, 4 (2013).
- [6] B. Koepnick, J. Flatten, T. Husain, A. Ford, D.-A. Silva, M. J. Bick, A. Bauer, G. Liu, Y. Ishida, A. Boykov *et al.*, De novo protein design by citizen scientists, *Nature* **570**, 390 (2019).
- [7] D. Kahneman, *Thinking, Fast and Slow* (Macmillan, New York, 2011).
- [8] D. Dellermann, P. Ebel, M. Söllner, and J. M. Leimeister, Hybrid intelligence, *Bus. Inf. Syst. Eng.* **61**, 637 (2019).
- [9] P. Michelucci and J. L. Dickinson, The power of crowds, *Science* **351**, 32 (2016).
- [10] E. Kamar, Directions in hybrid intelligence: Complementing AI systems with human intelligence, in *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)* (2016), pp. 4070–4073.
- [11] A. Berditchevskaia and P. Baeck, Future of minds and machines, Nesta report (2020), https://media.nesta.org.uk/documents/FINAL_The_future_of_minds_and_machines.pdf.
- [12] E. Baltz, E. Trask, M. Binderbauer, M. Dikovsky, H. Gota, R. Mendoza, J. Platt, and P. Riley, Achievement of sustained net plasma heating in a fusion experiment with the optometrist algorithm, *Sci. Rep.* **7**, 1 (2017).
- [13] A. W. Senior, R. Evans, J. Jumper, J. Kirkpatrick, L. Sifre, T. Green, C. Qin, A. Žídek, A. W. R. Nelson, A. Bridgland, H. Penedones, S. Petersen, K. Simonyan, S. Crossan, P. Kohli, D. T. Jones, D. Silver, K. Kavukcuoglu, and D. Hassabis, Improved protein structure prediction using potentials from deep learning, *Nature* **577**, 706 (2020).
- [14] <https://gitlab.com/quatomic/quantum-moves2>.
- [15] A. Lieberoth, M. K. Pedersen, A. C. Marin, T. Planke, and J. F. Sherson, Getting humans to do quantum optimization—user acquisition, engagement and early results from the citizen cyberscience game quantum moves, *Hum. Comput.* **1** 219 (2014).
- [16] C. Díaz, M. Ponti, P. Haikka, R. Basaiawmoit, and J. Sherson, More than data gatherers: exploring player experience in a citizen science game, *Qual. User Exp.* **5**, 1 (2020).
- [17] J. J. Sørensen, M. K. Pedersen, M. Munch *et al.*, Retraction Note: Exploring the quantum speed limit with computer games, *Nature* **584**, 484 (2020).
- [18] J. J. Sørensen, M. K. Pedersen, M. Munch, P. Haikka, J. H. Jensen, T. Planke, M. G. Andreassen, M. Gajdacz, K. Mølmer, A. Lieberoth *et al.*, Exploring the quantum speed limit with computer games, *Nature* **532**, 210 (2016).
- [19] A. Grønlund, Explaining the poor performance of the Kass algorithm implementation, [arXiv:2003.05808](https://arxiv.org/abs/2003.05808) (2020).
- [20] D. Sels, Stochastic gradient ascent outperforms gamers in the quantum moves game, *Phys. Rev. A* **97**, 040302(R) (2018).
- [21] A. Grønlund, Algorithms clearly beat gamers at quantum moves: A verification, [arXiv:1904.01008](https://arxiv.org/abs/1904.01008) (2019).
- [22] S. Cooper, F. Khatib, A. Treuille, J. Barbero, J. Lee, M. Beenen, A. Leaver-Fay, D. Baker, Z. Popović *et al.*, Predicting protein structures with a multiplayer online game, *Nature* **466**, 756 (2010).
- [23] S. Cooper, A. L. Sterling, R. Kleffner, W. M. Silversmith, and J. B. Siegel, Repurposing citizen science games as software tools for professional scientists, in *Proceedings of the 13th International Conference on the Foundations of Digital Games* (Association for Computing Machinery, New York, NY, 2018), pp. 1–6.

- [24] M. Haklay, Citizen science and volunteered geographic information: Overview and typology of participation, in *Crowdsourcing Geographic Knowledge* (Springer, Berlin, 2013), pp. 105–122.
- [25] W. E. Forum, Future of jobs report (2018), http://www3.weforum.org/docs/WEF_Future_of_Jobs_2018.pdf.
- [26] R. Heck, O. Vuculescu, J. J. Sørensen, J. Zoller, M. G. Andreassen, M. G. Bason, P. Ejlertsen, O. Elfasson, P. Haikka, J. S. Laustsen, L. L. Nielsen, A. Mao, R. Müller, M. Napolitano, M. K. Pedersen, A. R. Thorsen, C. Bergenholtz, T. Calarco, S. Montangero, and J. F. Sherson, Remote optimization of an ultracold atoms experiment by experts and citizen scientists, *Proc. Natl. Acad. Sci. USA* **115**, E11231 (2018).
- [27] <https://www.scienceathome.org/games/quantum-minds/>.
- [28] S. Z. Ahmed, J. H. M. Jensen, C. A. Weidner, J. J. Sørensen, M. Mudrich, and J. F. Sherson, Quantum composer: A programmable quantum visualization and simulation tool for education and research, [arXiv:2006.07263](https://arxiv.org/abs/2006.07263) (2020).
- [29] J. P. Dowling and G. J. Milburn, Quantum technology: The second quantum revolution, *Philos. Trans. R. Soc. London, Ser. A* **361**, 1655 (2003).
- [30] Z. Ugray, L. Lasdon, J. Plummer, F. Glover, J. Kelly, and R. Martí, Scatter search and local nlp solvers: A multistart framework for global optimization, *INFORMS J. Comput.* **19**, 328 (2007).
- [31] Numerical and implementation details such as inexact line searches introduce the possibility of “skipping over” the true nearest attractor.
- [32] C. Weitenberg, S. Kuhr, K. Mølmer, and J. F. Sherson, Quantum computation architecture using optical tweezers, *Phys. Rev. A* **84**, 032322 (2011).
- [33] T. Schumm, S. Hofferberth, L. M. Andersson, S. Wildermuth, S. Groth, I. Bar-Joseph, J. Schmiedmayer, and P. Krüger, Matter-wave interferometry in a double well on an atom chip, *Nat. Phys.* **1**, 57 (2005).
- [34] U. Hohenester, P. K. Rekdal, A. Borzi, and J. Schmiedmayer, Optimal quantum control of Bose-Einstein condensates in magnetic microtraps, *Phys. Rev. A* **75**, 023602 (2007).
- [35] G. Jäger, D. M. Reich, M. H. Goerz, C. P. Koch, and U. Hohenester, Optimal quantum control of bose-einstein condensates in magnetic microtraps: Comparison of gradient-ascent-pulse-engineering and krotov optimization schemes, *Phys. Rev. A* **90**, 033628 (2014).
- [36] J. J. Sørensen, M. Aranburu, T. Heinzl, and J. Sherson, Approaching the quantum speed limit with global-local optimization, [arXiv:1802.07521](https://arxiv.org/abs/1802.07521) (2018).
- [37] R. Bücker, J. Grond, S. Manz, T. Berrada, T. Betz, C. Koller, U. Hohenester, T. Schumm, A. Perrin, and J. Schmiedmayer, Twin-atom beams, *Nat. Phys.* **7**, 608 (2011).
- [38] S. van Frank, M. Bonneau, J. Schmiedmayer, S. Hild, C. Gross, M. Cheneau, I. Bloch, T. Pichler, A. Negretti, T. Calarco *et al.*, Optimal control of complex atomic quantum systems, *Sci. Rep.* **6**, 34187 (2016).
- [39] J. J. Sørensen, M. O. Aranburu, T. Heinzl, and J. F. Sherson, Quantum optimal control in a chopped basis: Applications in control of Bose-Einstein condensates, *Phys. Rev. A* **98**, 022119 (2018).
- [40] Another common seeding strategy is random seeding in frequency space which has the advantage of enabling frequency cutoffs and thereby adjustable levels of smoothness of the seeds. Although we have also conducted the analysis below for this type of seed parametrization, it will not be presented here in order to preserve the clarity of the presentation, as it generally performed on par with or slightly worse than the uniform random seeding for the chosen parameters.
- [41] Due to technical limitations at the time, all optimizations of player solutions (inside and outside the game) in *Bring Home Water* were performed with the same numerical implementation of Eq. (2) as $g \neq 0$ but with $g = 0$. This (artificially) increases the operational costs compared to optimization of RS solutions. As we shall see in the following, however, the conclusions are robust toward this slight skew and fully leveling the playing field would serve to reinforce them.
- [42] The densities are obtained by kernel density estimation using the Epanechnikov (parabolic) kernel with bandwidth 0.08 in the infidelity dimension [evaluated as $\log_{10}(1 - F)$], which attributes a localized density peak to each sampled data point and yields a density map when adding all the points.
- [43] This *a priori* fixation heuristic is the strongest type of structure to impose on the problem, the motivation of which is rooted in insight regarding the particular problem (steeper potentials yield faster dynamical time scales). This type of heuristic may not always be readily available, e.g., if the control parameters must simultaneously change in a nontrivial way to facilitate the transfer (e.g., the $\theta(t)$ and $\beta(t)$ parameters in Ref. [62]). Nevertheless, here we choose the *a priori* heuristically constrained version to be consistent with previous studies [20,21]. Likewise, the number of bins for the reduced resolution variant ($n_b = 40$) was chosen based on these previous studies.
- [44] It is safe to conjecture that a more careful choice of the heuristic bin parameter $40 < n_b < n_t$ would yield results matching the full resolution ($n_b = n_t$) above some resolution threshold.
- [45] SA performs exhaustive search (within its discretization) along a single, stochastically chosen dimension at a time and it is therefore possible to escape from initial trapping regions likely to be found by structureless random seeding. Eventually, the optimal strategies are discoverable given that their attractors are sufficiently broad (which for an arbitrary problem is not generally guaranteed). However, GRAPE is principally a strict local optimizer and will rapidly converge to the local trap. Thus, when using GRAPE the structureless RS seeding is more likely to get stuck in local traps, whereas the more structured PS seeding is more likely to yield better results.
- [46] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, A density-based algorithm for discovering clusters in large spatial databases with noise, *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining (KDD)* (1996), 226.
- [47] The DBSCAN clustering algorithm allows controls that are far away from each other (as measured by, e.g., the Euclidean metric) to be in the same cluster, provided there are other controls bridging the gap between them. This property enables solutions from a given optimal strategy to be grouped within the same cluster.
- [48] Initial results revealed two major and three minor clusters. By inspection it was found that each smaller cluster was in fact identical to one of the major clusters, apart from an initial delay in placing the tweezer near the atom. In order to filter out this physically insignificant delay we apply a correction procedure

to all the controls, dropping all initial points before the first instance of $u_1(t/T) \geq 0$. Applying the clustering with the same clustering parameters returns a significantly simpler result as seen in Fig. 5. Examination of the unclassified controls suggest that those solutions do belong to one of the two strategies but have a few local defects (sudden displacements) that make them sub-optimal with respect to the nearest cluster and appear far away from other cluster members based on the Euclidean metric. Thus, effectively only two (high-fidelity) optimal strategies exist.

- [49] For the *Shake Up* analysis we had initially produced a factor of 6 more RS seeds than PS seeds. To make the comparison fair, we form subsets by discretizing the T axis into 24 bins and randomly drawing RS seeds equal to the number of PS seeds in each bin (the ensuing difference in total number of seeds in Fig. 11 is due to the lack of RS seeds at very short durations). The presented solution density and aggregate results are representative of the statistical outcome of this procedure.
- [50] Comparing iterations is not very instructive since the algorithms are so fundamentally different.
- [51] J. J. Sørensen, J. H. M. Jensen, T. Heinzl, and J. F. Sherson, QEngine: A C++ library for quantum optimal control of ultracold atoms, *Comput. Phys. Commun.* **243**, 135 (2019).
- [52] M. Dalgaard, F. Motzoi, J. J. Sørensen, and J. F. Sherson, Global optimization of quantum dynamics with AlphaZero deep exploration, *npj Quantum Inf.* **6**, 6 (2020).
- [53] R. Iten, T. Metger, H. Wilming, L. del Rio, and R. Renner, Discovering Physical Concepts with Neural Networks, *Phys. Rev. Lett.* **124**, 010508 (2020).
- [54] M. Krenn, M. Malik, R. Fickler, R. Lapkiewicz, and A. Zeilinger, Automated Search for New Quantum Experiments, *Phys. Rev. Lett.* **116**, 090405 (2016).
- [55] phys.au.dk/forskning/cscaa.
- [56] I. Lesanovsky, T. Schumm, S. Hofferberth, L. M. Andersson, P. Krüger, and J. Schmiedmayer, Adiabatic radio-frequency potentials for the coherent manipulation of matter waves, *Phys. Rev. A* **73**, 033619 (2006).
- [57] Calculations could in principle be sped up by subsuming the individual propagators into a single step of length $w\delta t$, although such approach would require small enough w to ensure the error associated with time evolution remains small.
- [58] S. J. Wright, Coordinate descent algorithms, *Math. Program.* **151**, 3 (2015).
- [59] $(n_t + n_d) + (n_b - 1)(n_t/3 + n_d) = \frac{2}{3}n_t + n_b(n_t/3 + n_d)$.
- [60] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed. (Springer Science & Business Media, New York, 2006).
- [61] The discrete SA algorithm could also be adapted to, e.g., only check an informed subset of the n_d possibilities in Ω , or let Ω itself be dynamically changing.
- [62] J. H. M. Jensen, J. J. Sørensen, K. Mølmer, and J. F. Sherson, Time-optimal control of collisional $\sqrt{\text{swap}}$ gates in ultracold atomic systems, *Phys. Rev. A* **100**, 052314 (2019).