

Bayesian Optimization for Robust State Preparation in Quantum Many-Body Systems

Tizian Blatz^{1,2}, Joyce Kwan³, Julian Léonard⁴, and Annabelle Bohrdt^{2,5}

¹Department of Physics and Arnold Sommerfeld Center for Theoretical Physics, Ludwig-Maximilians-Universität München, Munich D-80333, Germany

²Munich Center for Quantum Science and Technology (MCQST), Munich D-80799, Germany

³Department of Physics, Harvard University, Cambridge, MA 02138, USA

⁴Vienna Center for Quantum Science and Technology, Atominstitut, TU Wien, Vienna 1020, Austria

⁵Institute of Theoretical Physics, University of Regensburg, Regensburg D-93053, Germany

New generations of ultracold-atom experiments are continually raising the demand for efficient solutions to optimal control problems. Here, we apply Bayesian optimization to improve a state-preparation protocol recently implemented in an ultracold-atom system to realize a two-particle fractional quantum Hall state. Compared to manual ramp design, we demonstrate the superior performance of our optimization approach in a numerical simulation – resulting in a protocol that is 10× faster at the same fidelity, even when taking into account experimentally realistic levels of disorder in the system. We extensively analyze and discuss questions of robustness and the relationship between numerical simulation and experimental realization, and how to make the best use of the surrogate model trained during optimization. We find that numerical simulation can be expected to substantially reduce the number of experiments that need to be performed with even the most basic transfer learning techniques. The proposed protocol and workflow will pave the way toward the realization of more complex many-body quantum states in experiments.

1 Introduction

Recent advances in quantum simulation platforms based on ultracold atoms in optical lattices [1, 2] have allowed the creation of increasingly complex quantum many-body states in larger systems, necessitating improvements in state-preparation quality and speed. Well-established approaches like adiabatic ramping [3], manual mapping of the quantum many-body gap [4], and optimal control techniques [5–8] have allowed experimental access to exciting new phenomena, such as magnetic and

topological states [4, 9–11].

In recent years, machine learning methods, such as reinforcement learning [12, 13], have emerged as a powerful tool for quantum optimal control (QOC) applications. Bayesian optimization (BO) [14, 15], an alternative machine-learning approach, has been used to generate efficient protocols for fast creation of high-particle number Bose-Einstein condensates [16, 17], cooling in magneto-optical traps [18], and state preparation [19, 20]. In these works, an online optimization loop is formed between the learner, and either an experimental setup or numerical simulation. The resulting optimized protocols have been demonstrated to compare favorably to manual ramp design [17] and other optimization techniques, such as differential evolution and Nelder-Mead [20–22]. BO is well suited to these tasks featuring experimental disorder or low repetition rates, as it typically requires fewer iterations, than e.g., methods based on gradient descent, and training is robust to moderate noise.

In this work, we use an exact-diagonalization (ED) simulation and apply BO to prepare a few-body fractional quantum Hall (FQH) state which has recently been realized in an ultracold-atom experiment [4]. This allows us to compare the optimized result to a protocol that is manually designed and has been implemented in practice. We extensively investigate how taking disorder in the experimental system into account changes the optimal strategy and how to engineer effective control paths in a noisy environment. Based on this analysis, we explore how exploiting classical simulation resources can boost efficiency by reducing experimental costs. From this, we motivate a general optimization strategy, which we believe to be applicable to any experimental system in which classical simulation is available in some capacity. These include not only the preparation of the Laughlin state [23], which acts as our benchmark, but also other current issues including challenging state preparation, such as the realization of quantum spin liquid states [24], and cooling in Fermi-Hubbard

Tizian Blatz: blatz.tizian@physik.uni-muenchen.de

Annabelle Bohrdt: annabelle.bohrdt@physik.uni-regensburg.de

type systems [25, 26].

In contrast to other optimal-control applications realizing a phase transition [20], e.g., the superfluid to Mott-insulator transition [10, 19, 27, 28], the FQH problem features a topologically non-trivial target state prepared by tuning a larger number of control fields (4 vs 1-2) to be optimized over time. This more demanding problem can shed light on the prospects of machine-learning-optimized state preparation, entering a regime where simple linear ramping is impossible and an optimized protocol's robustness and generality become primary concerns. To determine whether BO is a tool fit to tackle these challenges, our work proceeds as follows: We discuss our optimization strategy and investigate the experimental setting motivating our work by comparing experimental data to ED simulation. In the core part of our work, we apply BO to the simulated system and thoroughly investigate the performance of optimized ramping protocols regarding fidelity, speed, and robustness to disorder. We conclude with a proposal that we believe will facilitate the creation of more complex states in the system under investigation and also act as a general tool to boost efficiency in QOC applications going beyond this specific system.

2 Optimization Strategy

In this work, we implement a BO training loop as illustrated in Fig. 1. Gaussian processes (GP) [29] are used as a surrogate model of the cost function or figure of merit of ramping protocols over the entire parameter space. For details about BO and our implementation, the reader is referred to Appendix A. Fig. 1 also shows the proposed optimization strategy to be used in experiments. The strategy is focused on making use of the trained surrogate model and numerical simulation to reduce experimental costs. We develop this strategy in the following sections.

3 Experimental Setting

Our application of BO is focused on the state-preparation problem of a bosonic two-particle FQH state recently realized in experiment [4]. The experiment features a ramp protocol, adiabatically¹ transforming an initial state of two pinned particles into the delocalized FQH state. To achieve this, four global lattice parameters in the interacting Harper-

¹In the context of optimization, we consider an adiabatic protocol to be any strategy aiming to minimize losses of the target-state population by limiting the ramp speed to the maximum set by the inverse of the many-body gap.

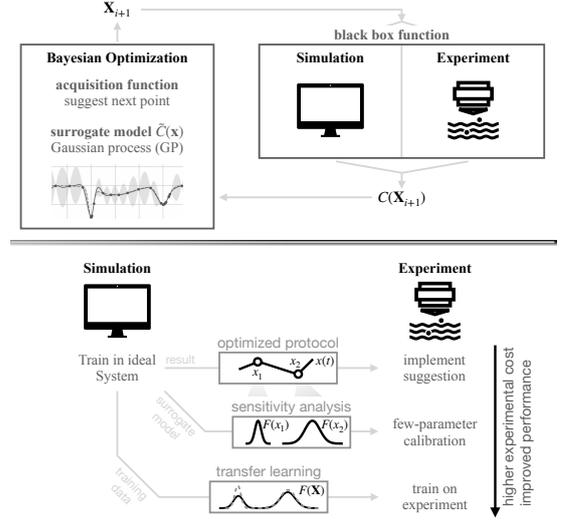


Figure 1: Optimization Strategy: **(top)** Illustration of the Bayesian optimization training loop. The experiment or numerical simulation acts as a black-box function, providing a quality measure for a given ramping protocol. Based on evaluations of this function, the BO learner builds a surrogate model of the cost landscape and suggests the next protocol to implement. **(bottom)** Proposed optimization strategy exploiting numerical resources to boost optimization efficiency in the experiment. Depending on the task at hand, several avenues with varying experimental costs may be explored. The most straightforward way is to directly implement the optimization result obtained from simulation in the experiment, requiring only a single evaluation of the protocol's quality. On top of this, the BO surrogate model can provide a sensitivity analysis that can be exploited to manually calibrate only the most sensitive protocol parameters in ~ 10 experimental evaluations. Finally, when substantial adaptation to the experimental setting is necessary, online optimization can be performed in ~ 100 to ~ 1000 iterations directly with the experiment. Transfer-learning approaches can leverage numerical resources to accelerate training significantly, e.g., by initializing the surrogate model used for training in the experiment with training data from the simulation. The prospective reduction of the number of cost-function evaluations in the experiment is crucial in the context of optical lattice systems where typical snapshot-based detection schemes require many experimental runs to determine an accurate figure of merit. In this work, we highlight and motivate these three scenarios by comparing the numerical simulation of an idealized system to a more realistic description of the experimental setting that takes experimental disorder into account.

Hofstadter model [30–32]

$$\begin{aligned}
 \hat{H} = & -t_x \sum_{x,y} (e^{i\phi_y} \hat{a}_{x,y}^\dagger \hat{a}_{x+1,y} + \text{H.c.}) + \Delta_x \sum_{x,y} x \hat{n}_{x,y} \\
 & -t_y \sum_{x,y} (\hat{a}_{x,y}^\dagger \hat{a}_{x,y+1} + \text{H.c.}) + \Delta_y \sum_{x,y} y \hat{n}_{x,y} \\
 & + \frac{U}{2} \sum_{x,y} \hat{n}_{x,y} (\hat{n}_{x,y} - 1)
 \end{aligned} \tag{1}$$

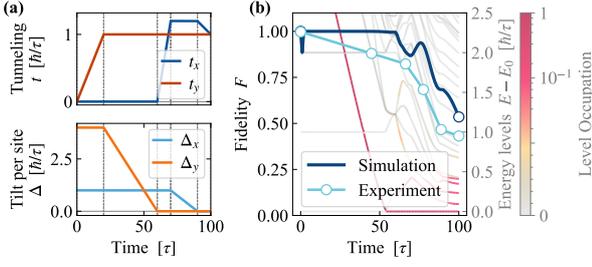


Figure 2: Experimental Ramp: **(a)** Experimental ramping protocol for the 4 lattice parameters adapted from [4]. **(b)** Fidelity w.r.t. the instantaneous ground state throughout the ramp. The ED simulation is compared to experimental data [4]. The color-coded lines in the background show the energy levels $E - E_0$ of $\hat{H}(t)$ and their respective occupation during the simulated ramping protocol. Here, E_0 is the instantaneous ground-state energy. Before t_x is turned on at $T = 60 \tau$, the system is effectively one-dimensional, since all particles are initialized in the 1st column of the 4×4 lattice and there is no coupling in x-direction. In this regime, we define the fidelity wrt. this 1D system to avoid spurious drops to $F \sim 0$ when $\Delta_x < \Delta_y$ caused by the initial-state geometry.

are varied independently over time: (i, ii) the tunneling strengths $t_{x,y}$ in x- and y-direction and (iii, iv) the chemical potential gradients $\Delta_{x,y}$. Here, $\hat{a}_{x,y}^{(\dagger)}$ is the bosonic annihilation (creation) operator at site (x, y) and $\hat{n}_{x,y} = \hat{a}_{x,y}^\dagger \hat{a}_{x,y}$ is the particle number operator. The tunneling phase $\phi/2\pi = 0.27$ is constant throughout this work, corresponding to the Laughlin-like state [23] at the fractional filling $\nu = 1/2$, where ν is the filling factor. This target state is realized in the flat lattice $\Delta_x = \Delta_y = 0$, at equal tunneling amplitudes $t_x = t_y = \hbar/\tau$. We fix the value of the tunneling amplitudes in the target state as our unit of energy while the tunneling time τ acts as the unit of time. In the experiment, $\tau = 4.3(1)$ ms [4]. Finally, the on-site repulsion is fixed to $U = 8 \hbar/\tau$.

The ramping protocol used in the experiment was devised by mapping out the many-body gap over the four-dimensional parameter space and manually identifying a path that avoids a gap closing. In addition, the tunneling t_x in x-direction is only turned on in the latter half of the experimental protocol to reduce the effects of Floquet heating [33] resulting from the synthetic magnetic field engineered in the experiment. The protocol is illustrated in Fig. 2. It serves as the primary benchmark in this work – representing the capabilities of sophisticated adiabatic ramp design.

The quality of a ramp is quantified by the fidelity

$$F = |\langle \psi_{\text{final}} | \psi_0 \rangle|^2, \quad (2)$$

where $|\psi_{\text{final}}\rangle$ is the state after the preparation protocol, and $|\psi_0\rangle$ is the ground state of \hat{H} at the target parameters $t_x = t_y = \hbar/\tau$ and $\Delta_x = \Delta_y = 0$. This is a good figure of merit since it is experimentally accessible by inverting the ramp and measuring the

fraction that has returned to the initial state using snapshots in the number basis. For the experimental protocol it was measured to be $F = 43(6)\%$ [4]. In Fig. 2, we compare the experimental data to an ED time-evolution simulation. Even though the simulation is performed at zero temperature, does not take into account any lattice defects, and does not model the Floquet sequence employed in the experiment – instead simulating the unperturbed effective Hamiltonian – the fidelity-loss over time (see Fig. 2b)) shows qualitative agreement with the experiment. The fidelity of $F = 53.5\%$ observed in the simulation is not much higher than the value reported in the experiment. This supports non-adiabatic losses to excited states during the state-preparation protocol as the leading source of fidelity loss, providing a promising foundation for ramp optimization.

4 Optimization

For each lattice parameter, we parametrize the time-dependent ramp by N_s linear segments of equal length. Fixing the initial and target parameters to be the same as in the experiment leads to a BO problem with $4 \times (N_s - 1)$ parameters. Motivated by the experimental conditions and concerns about Floquet heating, we limit the search space to

$$0 \leq t_{x/y} \leq 1.2 \hbar/\tau \text{ and } -4 \hbar/\tau \leq \Delta_{x/y} \leq 4 \hbar/\tau, \quad (3)$$

which corresponds to a Floquet regime that minimizes excitations to higher bands and achieves efficient dynamics. Specifically, the range for $t_{x/y}$ is chosen to limit the broadening of heating channels [33], as its strength is determined by the driving amplitude. We start ramping t_x and Δ_x only after passing the time $T_x^{(0)} = 0.6 T$ where T is the total ramp time. In summary, the optimizer may choose the values of t_y and Δ_y at times $T_y^{(1)} = T/3$ and $T_y^{(2)} = 2T/3$, and the values of t_x and Δ_x at times $T_x^{(1)} = T_x^{(0)} + (T - T_x^{(0)})/3$ and $T_x^{(2)} = T_x^{(0)} + 2(T - T_x^{(0)})/3$.

After training with 500 iterations, BO identifies an optimal ramp with $F = 94.5\%$ as illustrated in Fig. 3. Interestingly, the protocol is qualitatively similar to the one used in the experiment. As depicted in Fig. 3, it achieves near-perfect fidelity due to significant gains in the ground-state population (w.r.t. the instantaneous ground state) near the final protocol time. Such a resurgence of fidelity has also been reported in other simulation works on QOC applications [34]. While captured by optimal control theory, it is outside the scope of manual design principles, such as adiabatic ramp design, which aim to minimize fidelity losses, but do not take processes that can return population to the ground

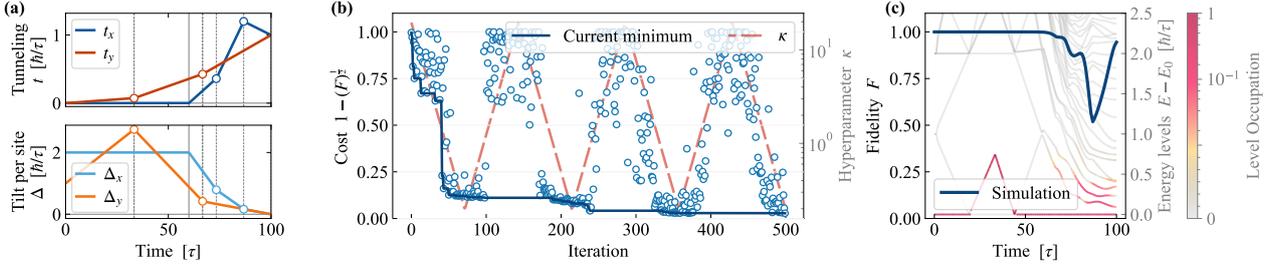


Figure 3: Optimized Ramp: **(a)** Optimized ramp protocol for the 4 lattice parameters. White dots show the parameter values controlled by the optimizer at fixed points in time (vertical dashed lines). **(b)** Evaluated cost function at each iteration of the BO training. The visible oscillations in the evaluated fidelity follow sweeps of the BO hyperparameter κ , which tunes the strategy between exploration and exploitation (see Appendix A for optimization details). **(c)** Fidelity w.r.t. the instantaneous ground state throughout the ramp. The color-coded lines in the background show the energy levels of $\hat{H}(t)$ and their respective occupation during the simulated ramping protocol. All energies are w.r.t. the ground state energy which sets $E = 0$. Before t_x is turned on at $T = 60 \tau$, the system is effectively one-dimensional, since all particles are initialized in the 1st column of the 4×4 lattice and there is no coupling in x-direction. In this regime, we define the fidelity wrt. this 1D system to avoid spurious drops to $F \sim 0$ when $\Delta_x < \Delta_y$ caused by the initial-state geometry.

state into account. We now turn to whether this feature, observable in simulations, can be expected to be robust in experimental conditions.

5 Training with Noise

To characterize robustness, and estimate the real-world fidelity we would expect from optimized protocols, we now model experimental disorder in our simulation by the sum of random and harmonic chemical potential offsets (see Appendix C). The disorder μ is taken to be Gaussian-distributed with mean 0 and standard deviation $\sigma = 0.1 \hbar/\tau$ which is an experimentally realistic disorder level. For each numerical evaluation of a ramp's fidelity, we perform time evolution with a newly sampled realization of the disorder.

Optimization is now performed in the presence of this disorder to represent training in an experimental setting. To mitigate fluctuations, we take the average fidelity of 5 disorder realizations as the cost function, which is a typical value used in the literature [17, 20]. After training for 1000 iterations, we compare the optimized protocol to the manually designed ramp and the optimum found from training without disorder in Fig. 4. The comparison is performed over a range of disorder strengths, covering both the clean as well as the experimentally realistic case. Firstly, this analysis highlights the competitiveness of the manually designed protocol and the complexity of the state preparation problem at hand. The experimentally implemented ramp is remarkably robust and evaluates to $F = 49(2)\%$ at $\sigma = 0.1 \hbar/\tau$, even closer to the experimentally measured ramp fidelity and more than $10\times$ the fidelity expected from simple linear ramping.

The fidelity of the ramp optimized in noisy training

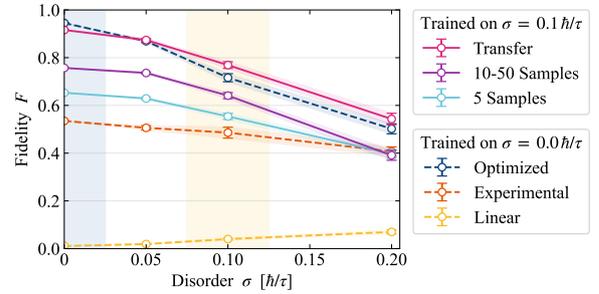


Figure 4: Experimental Noise: Expected fidelity F of different ramping protocols at varying standard deviation σ of the experimental disorder. The shaded areas show the low-disorder (blue) and the experimentally realistic (orange) regime. Solid lines show ramps trained in the presence of disorder, dashed lines show ramps designed in the ideal system. The error bars indicate the standard error of the mean fidelity obtained from averaging over 100 disorder realizations. The highest fidelity at $\sigma = 0$ is achieved by optimization in the clean system. The highest fidelity at $\sigma \geq 0.1 \hbar/\tau$ is achieved by training in the disordered system, initialized from training in the ideal system ('Transfer').

beats the fidelity of the manually designed protocol at all investigated disorder levels. However, it is inferior to the optimum found in training without disorder. We conclude that the noisy evaluation of the cost function makes training significantly more challenging.

We confirm this by significantly increasing the number of disorder realizations averaged over for an evaluation of the cost function to an adaptive number between 10 and 50 – taking more samples when the fidelity is high (see Appendix D for more details). Doing so boosts BO performance but still does not reach the fidelity achieved in training without disorder.

To reach optimum performance in the noisy environ-

ment, we turn to our knowledge of the ideal system. Usually, the surrogate model in BO is initialized by picking $\mathcal{O}(10)$ points in parameter space at random and evaluating the fidelity at these points. As a rudimentary implementation of transfer learning, we replace this random selection with a subset of the points visited during training in the ideal system. The most relevant points for training are chosen by a simple heuristic algorithm according to their fidelity (which is known for the clean system) and Euclidean distance in parameter space. For details about the algorithm, see Appendix E. The cost function at the chosen points is subsequently reevaluated in the presence of disorder to initialize the surrogate model. With this most basic transfer-learning approach, we reach the highest fidelity achieved in the noisy environment with a fraction of the necessary number of runs in the disordered system. Comparing this protocol with the optimized result from the clean system tells us something about the changes in the cost landscape when disorder is included. The new optimum in the disordered system is very close in parameter space to that in the clean system, so we find disorder to only lead to a slight shift in the ideal strategy. This scenario is distinct from one realized in a previous study [34] where the optimum location drastically changed with the onset of disorder. For a more detailed discussion of these two scenarios, see Appendix F which confirms our findings as a property of the system in question – not an artifact of optimization – via a brute-force search and examines a special case in our system that mirrors the scenario in the literature.

6 Short Ramps

Thus far, we have aimed to produce ramping protocols that achieve high fidelity both in an ideal as well as a noisy setting. The total ramp time T has been kept fixed to that of the adiabatic experimental protocol. However, a major advantage of QOC techniques is that they can significantly reduce the total protocol time while retaining high fidelities [10, 16]. To leverage this advantage, we run optimization with reduced final times $T = 10\tau$, and $T = 20\tau$, i.e., $10\times$, or $5\times$ shorter than the protocols previously considered. Based on our findings in the previous sections, we expect the global optimum of the optimization landscape to move only a little when introducing experimental disorder, so we chose to work in the ideal system.

Running the optimization for 500 iterations, we discover ramping protocols with $F = 53\%$, and $F = 78\%$, respectively. The shorter protocols are compared to the ones with $T = 100\tau$ in Fig. 5. We find the shortest ramp to be remarkably close to the experimental protocol in terms of the resulting

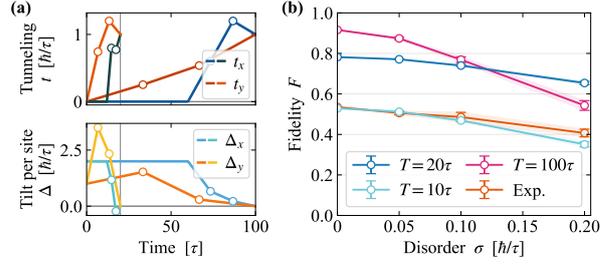


Figure 5: Short Ramps: **(a)** Ramp protocols for the optimized ramps with final times $T = 20\tau$ and $T = 100\tau$. The ramp shown for the final time $T = 100\tau$ is the most robust one considered in the previous section, i.e., the ‘Transfer’ ramp from Fig. 4. White dots show the parameter values controlled by the optimizer at fixed points in time. **(b)** Fidelities of the ramps optimized for various final times, as well as the experimental ramping protocol (which has final time $T = 100\tau$) at varying disorder level σ . The error bars indicate the standard error of the mean fidelity obtained from averaging over 100 disorder realizations.

fidelity – both in the clean as well as the disordered system. That is, at this fixed level of fidelity, optimization reaches a speedup by a factor of 10. We consider the ramp resulting from the ramping time $T = 20\tau$ to be optimal for an experimental implementation. While the fidelity in the disorderless scenario is still significantly lower than for optimized ramps with longer ramp times, we find the protocol to be extremely robust under the introduction of disorder. In fact, the fidelity at the strongest considered level of disorder $\sigma = 0.2\hbar/\tau$ is significantly higher than for any other protocol considered in this work. Similar to previous studies [10, 34], we attribute this result to the shorter protocol time presenting less of an opportunity for the accumulation of errors. This interpretation is supported by the significantly lower standard deviation of the fidelity in the noisy setting. At the experimental disorder level $\sigma = 0.1\hbar/\tau$, while the mean fidelity of the $T = 20\tau$ ramp is similar to that of the ramp optimized for $T = 100\tau$, the standard deviation is lower by more than a factor of 4.

7 Experimental Proposal

Based on the findings of the previous sections, we propose the following strategy and best practices for the realization of the BO approach in an experiment: (i) Implement suggestion: Due to the advantages in terms of speed and robustness of shorter ramping protocols, we suggest the $T = 20\tau$ ramp from the previous section for an experimental realization. Even though the training was performed on the ideal system, we believe the scheme to be directly implementable, as we found the global optimum location to be relatively insensitive to disorder.

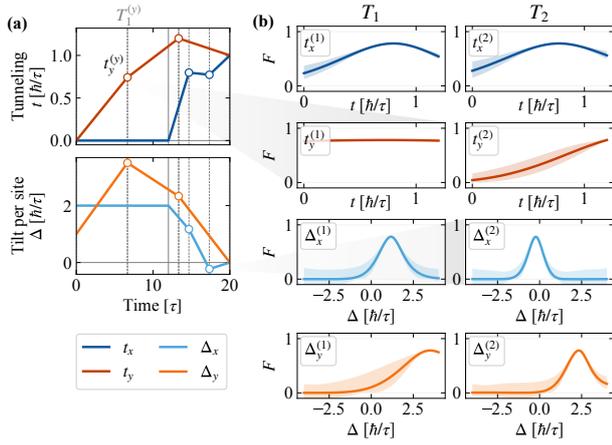


Figure 6: GP Model: Gaussian-process model of the cost landscape around the global optimum trained during optimization of the $T = 20 \tau$ ramp. **(a)** The optimum ramping protocol identified by optimization – characterized by 4 control fields at 2 time points (white dots). **(b)** Single-parameter variations around the global optimum of the modeled cost landscape with the remaining parameters fixed to their optimum value. The shaded regions depict the model uncertainty (one standard deviation), which is generally the lowest around the global optimum. The sensitivity of the model to changes in different parameters varies considerably. Grey, shaded regions link the least (most) sensitive protocol parameter $t_y^{(1)}$ ($\Delta_x^{(2)}$) to the corresponding modeled optimum shape.

Direct implementation of a numerically optimized (local adiabatic) control path has previously been used successfully to realize few-particle crystallization in an experiment [35].

(ii) Sensitivity analysis: Going beyond this innate robustness, we argue that a sensitivity analysis based on the BO surrogate model can reduce manual calibration in the experiment to a few parameters. Since BO takes an increased number of samples around the discovered optimum, the model is very accurate in this region and can faithfully predict the optimum’s shape in parameter space. Thus, the surrogate model contains accurate information about the optimum’s sensitivity to the different parameters. When the optimum is only slightly shifted in parameter space by the onset of disorder, good performance can be recovered at moderate experimental cost by manually tuning only the most sensitive parameters.

The sensitivity analysis drawn from training the proposed $T = 20 \tau$ ramp is depicted in Fig. 6. By varying single parameters around the optimum, we identify the fidelity to be much more sensitive to tuning the global tilts than to changes in hopping amplitudes. In particular, the GP model of the cost function is almost completely insensitive to the value $t_y^{(1)}$ at the first time point $T = 20/3 \tau$ but sharply peaked around the optimum value of $\Delta_x^{(2)}$ at the second fixed time point. By taking advantage of the trained learner in this way, manual calibration in the experiment can be reduced from full functional

dependencies to only a few parameters.

(iii) Transfer learning: In case sensitivity analysis fails to recover the optimum, or the system is intractable numerically, it becomes favorable to train on the experiment directly. In Section 5, we have demonstrated that making use of simulation data in a simple transfer-learning approach can significantly boost performance while reducing experimental costs (represented by the number of simulation runs in the disordered system). In Section H we also find that finding high-fidelity state-preparation protocols for a 4-particle Pfaffian-like state [36] on a larger $L_x = L_y = 6$ lattice is made much easier by transferring the results for the smaller state obtained in Section 4. This presents an avenue that could extend efficiency gains to experimental systems beyond the reach of classical simulation.

8 Discussion and Outlook

We applied a Bayesian optimization approach to optimize parameter ramps for state preparation. The optimization scheme was demonstrated to be efficient at finding protocols that outperform manual design not only in an ideal system but also in the presence of significant disorder – identifying a ramp that is both $5 \times$ faster as well as twice as close to ideal fidelity as the manual reference. The superior performance holds when the training data itself is noisy, e.g., when training directly on an experiment, but at an increased cost of training.

We propose several approaches to decrease experimental costs in cases where numerical simulation of the system is available. Namely, we demonstrated that optimized solutions for the ideal system can be intrinsically robust due to a significant reduction in protocol duration and the surrogate model of the fidelity-landscape naturally provides information about the most sensitive parameters to tune in an experiment. The reductions in experimental cost achievable by leveraging numerical resources are especially valuable to systems of ultracold atoms in optical lattices, where estimating the figure of merit usually requires many experimental snapshots.

Even when learning directly on the disordered system representing an experiment, prior training in a simulated environment was demonstrated to significantly reduce the total cost in Section 5.

The central challenge ahead lies in extending efficient optimization to system sizes that are not tractable using numerical simulation. Promising recent results in this regard were obtained by Y.-J. Xie et al. [20], who report efficient transfer of BO to iteratively larger system sizes in a Heisenberg chain. For the FQH system discussed here, we demonstrate good transferability to a larger 4-particle state with different model param-

ters in Section H. In particular, this marks an example, where the substantial speedup offered by protocols optimized using machine learning can allow the preparation of a more complex quantum state within the already available coherence times of contemporary ultracold-atom experiments.

By using either such a transfer-learning approach or alternative schemes merging several FQH plaquettes [37], we expect the optimized protocols developed here to be instrumental in realizing larger FQH states, such as the Pfaffian state, and different filling factors. The more general optimization strategy, we believe to be widely applicable to boost experimental efficiency and improve the interface between classical and quantum simulation.

Data and Code Availability – The data presented in the figures is available at <https://github.com/TizianBlatz/bo-state-preparation>, which also contains a minimal working example of the optimization routine.

Acknowledgments – We are very grateful to F. Grusdt, F. Palm, and C. Weitenberg for fruitful discussions, and we wish to thank M. Greiner, B. Bakkali-Hassani, P. Segura, and Y. Li for valuable insights regarding the experimental system. This research was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy—EXC-2111—390814868. J.L. acknowledges support by the FWF grant no. Y-1436 and the FFG grant no. F0999896041.

A Bayesian Optimization

Bayesian optimization (BO) [14, 15] is an optimization method designed to find the optimum of a – possibly noisy and expensive to evaluate – black box function. In our case, the task is finding the optimal ramping protocol $\hat{H}(t)$ for the time-dependent Hamiltonian that transforms a given initial state $|\psi_{\text{initial}}\rangle$ into the target state $|\psi_{\text{target}}\rangle$. The quality of a given ramping protocol is quantified by the fidelity

$$F = |\langle \psi(T) | \psi_{\text{target}} \rangle|^2, \quad (4)$$

where $|\psi(T)\rangle$ is the initial state evolved according to $\hat{H}(t)$ up to the final ramp time T .

To phrase this as a standard optimization problem, we consider Hamiltonians $\hat{H}(t)$ characterized by a set of parameters $\{h_i(t)\}_i$. Then, we take each Hamiltonian parameter $h_i(t)$ as a piece-wise linear function characterized by a set of time-value tuples $\{x_j = (t_j, v_j)\}_j$. The optimization task then becomes minimizing the cost function

$$C(\{x_j\}_j) = 1 - \sqrt{F}. \quad (5)$$

To find optimal values for the parameter set $\mathbf{x} = \{x_j\}_j$, a BO loop, based on the open-source

`scikit-optimize` package, is run in the following way: First, the search is initialized by randomly choosing some \mathbf{x}_i – typically $\mathcal{O}(10)$ points – and evaluating the corresponding cost function values $C(\mathbf{x}_i)$ – either by running an experiment or, in our case, via numerical simulation. We then use Gaussian processes (GP) [29] to build a surrogate model $\tilde{C}(\mathbf{x})$ of the cost function over the entire parameter space, based on the set of evaluated points. A GP is a generalization of the Gaussian distribution, defining a distribution over functions by replacing the mean vector and covariance matrix by the mean function $m(\mathbf{x})$, and the kernel $k(\mathbf{x}, \mathbf{x}')$ [29]. Given a set of previously evaluated points, the GP model predicts the value of the cost function over the entire parameter space and gives a measure of the model’s uncertainty in less-explored areas. Key properties of the GP model, such as the smoothness of and variation between sampled functions are controlled via a set of hyperparameters, which are continually adjusted during training via maximum-likelihood estimation. These hyperparameters include the noise level σ_F of the data as well as the length scales l which give the characteristic scale on which functions sampled from the GP change their value. Control over these parameters can be a valuable tool to achieve optimum BO performance. For example, manually setting $\sigma_F = 0$ for training in the clean system can boost convergence and contribute to the low number of iterations needed compared to training in the noisy setting. We also demonstrate how limiting the lower bound of l – constraining the maximum complexity of the constructed model – can improve robustness while speeding up convergence in Appendix G.

After the surrogate model is initialized, a new point \mathbf{x}_{i+1} is suggested based on the existing cost-function model, and the cost function is subsequently evaluated at this new point. The new tuple

$$(\mathbf{x}_i, C(\mathbf{x}_i)) \quad (6)$$

is then used to update $\tilde{C}(\mathbf{x})$, completing the loop.

The suggestion \mathbf{x}_i is made by the acquisition function, which can prioritize either exploitation – choosing an \mathbf{x}_i which minimizes $\tilde{C}(\mathbf{x})$, or exploration – choosing an \mathbf{x}_i where $\tilde{C}(\mathbf{x})$ has high uncertainty. In the lower confidence bound (LCB) [38] acquisition function

$$\text{LCB}(x) = m(x) + \kappa \sigma_{\text{GP}}(x) \quad (7)$$

used in this work, this tradeoff is tuned via the hyperparameter κ , which sets the weight of the model uncertainty σ_{GP} . In practice, we sweep between these two strategies, starting with a focus on exploration ($\kappa \sim 20$) and prioritizing exploitation ($\kappa \sim 0.2$) during the final iterations of the optimization. Fig. 3b) of the main text showcases 4 sweeps of κ and the

corresponding evaluations of the cost function for a training run with a total of 500 iterations.

To transfer results from a completed optimization run to a new one, e.g., to save cost when training on a noisy system, we replace random initialization with selected parameter-space points evaluated in the completed run. The points are selected by an algorithm detailed in a separate section. The corresponding values of the cost function are reevaluated in the new (e.g., noisy) environment, and the resulting (point, cost) tuples are used to initialize the surrogate model for the new optimization run.

B Exact Diagonalization Simulation

Exact-diagonalization (ED) simulations are used to perform ground-state search and time evolution for the Harper-Hofstadter Hamiltonian [31, 32] on a 4×4 lattice with $N = 2$ bosons. The parameters $t_{x/y}$ and $\Delta_{x/y}$ are the control fields tunable in the experiment – representing tunneling amplitudes and global lattice tilts respectively. Site-dependent chemical potentials $\mu_{x,y}$ may be introduced to model disorder in the experiment.

For efficiency, we exploit the U(1) symmetry [39], resulting from particle number conservation. As in the experiment, the initial state is prepared by manually placing the 2 bosonic particles in the bottom-left corner of the 4×4 lattice, i.e.,

$$|\psi_{\text{initial}}\rangle = \hat{a}_{0,0}^\dagger \hat{a}_{0,1}^\dagger |0\rangle, \quad (8)$$

where $|0\rangle$ is the 0-particle state.

Time evolution is performed using the Runge-Kutta method RK5(4) [40] with a maximum step-size $\delta t_{\text{max}} = 0.1$.

All computations for this work were run on a commercial laptop (Apple M1 Pro processor), where a single time evolution to $T = 100 \tau$ takes ~ 3 s. When running BO, fitting the surrogate model becomes more expensive than time-evolution simulation after about 200 iterations. A typical optimization run to 500 iterations takes about 1 h.

C Experimental Disorder Model

For a more realistic description of the experimental system, we consider disorder in the shape of static chemical-potential offsets that can be divided into two components: One as random offsets, and another as a harmonic confinement.

The random offsets arise due to light scattering on surfaces, introducing disorder in the optical lattice. We model these offsets as uncorrelated with a maximum strength of $\sim 10 - 15$ Hz, or about a quarter of the tunneling amplitude $\hbar/\tau \sim 40$ Hz in the final FQH state which sets our unit of energy. The harmonic confinement is introduced by the optical potentials,

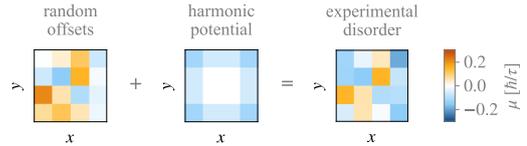


Figure 7: Experimental Disorder: Sample of a disorder realization with $\sigma = 0.1 \hbar/\tau$ composed of a random and a harmonic component.

or walls, that defines the region of interest. The walls are designed to have a Gaussian shape and are holographically projected by a digital micromirror device (DMD) [41]. The DMD is located in the Fourier plane, and the finite aperture of the DMD limits the sharpness of the walls, resulting in edge offsets that may be modeled as a harmonic confinement. Thus, on the 4×4 lattice, the offset is modeled as having strength $2\mu_{\text{DMD}}$ in the corners of the lattice and μ_{DMD} everywhere else on the edges. The strength μ_{DMD} of this type of disorder is taken to follow the same distribution as the random offsets, that is μ_{DMD} may have either sign and the two types of disorder are similarly strong.

In the simulation, we model the disorder $\mu [\hbar/\tau]$ as being Gaussian-distributed with mean 0 and consider a standard deviation of $\sigma = 0.1 \hbar/\tau$ to be the experimentally realistic disorder level. For each numerical time evolution, we sample $4 \times 4 + 1$ independent values from this distribution – one for each single-site offset and one for the global harmonic one.

We choose this disorder model as opposed to some other common experimental imperfections such as elevated temperatures and particle losses, as these can be well controlled in the small system by postselecting runs with the correct low-energy initial particle configuration and final particle number via density snapshots.

D Adaptive Sampling

To keep the total number of samples used for training low, while reducing the noise of cost-function evaluations near the optima, we use a simple adaptive sampling technique – increasing the number of samples based on the fidelity. Starting from a minimum number of 5 samples, more samples are added if the averaged fidelity is above a threshold value. The procedure continues with higher and higher fidelity thresholds until a predefined maximum number of 50 samples. We believe this technique to be beneficial as it allows us to more accurately map out the shape of discovered optima while saving cost by sacrificing

accuracy in regions where the fidelity is low.

E Initialization Algorithm

To select the most valuable points in parameter space to transfer to a new optimization run, we use a simple heuristic algorithm. Iterating through all evaluated points, a point \mathbf{x} is accepted, if it is more than a predefined distance r away from all previously accepted points (in terms of the Euclidean distance in parameter space). If the closest accepted point \mathbf{x}_{acc} is closer than r , between \mathbf{x}_{acc} and \mathbf{x} , the point that yielded a higher fidelity is accepted. The radius r is chosen such that ~ 100 points are accepted in total. The 10 – 100 points with the highest fidelity are then used to initialize the new optimization run.

This basic algorithm ensures that the point that yielded the highest fidelity in the BO run used for initialization is accepted and reevaluated for the new optimization. Apart from that, the algorithm prioritizes keeping points spread out over the parameter space (due to the division in ‘bubbles’ of size r) while keeping points with high fidelity. We therefore expect this algorithm to succeed in selecting the most important local optima alongside the global optimum. We expect it to work well, not only in our situation, where the location of the global optimum does not change much when introducing disorder to the system but also in more challenging situations, where the location changes significantly. This is the case since, in general, we believe the global optimum in a disordered system to very likely have been a local optimum in the ideal system. We discuss this in more detail in the following sections.

F Brute-Force Search

In the main text, using ramps with 3 linear segments, we find the best performance in the disordered system with a ramp very similar to the optimum obtained in the clean system. To confirm that this is a feature of the problem setting and not a result of poor optimization performance, we perform a brute-force search in the simpler four-dimensional parameter space characterizing ramps with only 2 linear segments. The fidelity with and without experimental disorder is performed within the search area

$$t_{x/y} \in [0.0, 1.2], \Delta_{x/y} \in [0.0, 4.0], \quad (9)$$

stepping through a grid with step sizes $\delta t_{x/y} = \delta \Delta_{x/y} = 0.2$. We motivate this uniform grid size, taking more than 3 times more points for the tilts Δ than for the hopping amplitudes t by the increased sensitivity to tilts suggested by the GP models trained during BO (see main text). Even if optimization performance were to be poor, we trust this general trend of sensitivities, as it is consistent in all BO training

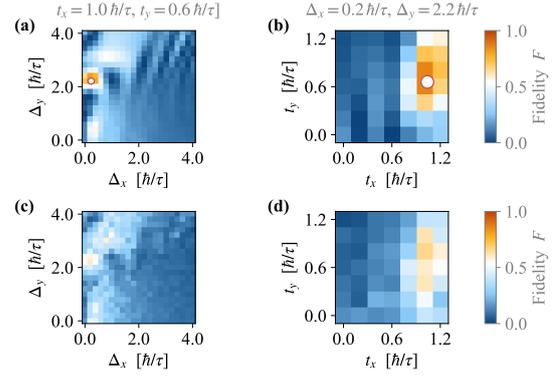


Figure 8: Brute-Force Search: The optimization landscape around the global optimum mapped out in the brute force search. **(a, b)** Ramp fidelity for variations of the tilts **(a)**, and tunneling amplitudes **(b)** with the 2 remaining parameters fixed to their optimum values (grey text). The value of each pixel is obtained from a time-evolution simulation of the corresponding ramping protocol. The white dots indicate the location of the optimum identified by BO. **(c, d)** The same as **(a, b)**, but for the brute-force search carried out in the disordered system with disorder strength $\sigma = 0.1 \hbar/\tau$. Each pixel represents the mean fidelity obtained from evaluating at least 10 independent disorder realizations.

performed in this work. To evaluate the fidelity in the presence of disorder, we use an adaptive sampling technique similar to the one used during noisy training – ensuring high accuracy in high-fidelity regions while cutting down on the total number of evaluations. Still, at least 10 independent disorder realizations are evaluated for each point in the search grid, so more than 200 000 time-evolution simulations are carried out for the full brute-force search in the disordered setting. Optimal fidelity is achieved at

$$[t_x = 1.0, t_y = 0.6, \Delta_x = 0.2, \Delta_y = 2.2], \quad (10)$$

determined to be $F = 85.8\%$ in the ideal system, and $F = 66(2)\%$ in the system including disorder. Crucially, optimum fidelity with- and without disorder is reached at the same point in parameter space, so – within the finite resolution provided by the grid – the two optima coincide.

BO trained on the clean system with 500 iterations finds the optimum to be at

$$[t_x = 1.04, t_y = 0.66, \Delta_x = 0.23, \Delta_y = 2.21], \quad (11)$$

with $F = 87.8\%$. These findings confirm that, in our problem setting, the global optimum of the cost landscape does indeed not move much with the introduction of disorder, and BO can efficiently and correctly identify said global optimum.

G Making Use of the Trained Model

Here, we perform a closer investigation of how the surrogate model trained during a BO run can be used to make predictions about robustness and how training simpler models can lead to protocols that are more robust against unknown types of noise.

Since the parameter space in the optimization directly translates to changes in Hamiltonian parameters over time – just like many types of noise or disorder do – we think it is reasonable to believe that the width of an optimum in parameter space can predict its performance against unknown types of noise.

If this relation between disorder in the experiment and noise in parameter space holds true, there is a simple intuitive picture of what is happening to the cost function when introducing noise to the system. In common with other noisy signals, peaks in the optimization landscape would be expected to become wider and flatter with increasing disorder. Depending on the type and strength of the noise, two distinct scenarios may occur:

- (i) The global optimum widens but remains the highest peak, thus the location of the global optimum does not change much.
- (ii) Due to the loss in fidelity, a wider – more robust – local optimum becomes the new global optimum, thus the location of the global optimum in parameter space may jump dramatically.

Based on the analysis presented in the main text, as well as the brute-force mapping of the cost landscape detailed above, we are confident that our system and disorder level realize scenario (i) – both for long ($T = 100 \tau$) and short ($T \leq 20 \tau$) ramping protocols. However, dramatic changes in the global optimum’s location with the introduction of noise have been reported in the literature [19] indicating scenario (ii). To look for (ii) in our system and test our conjecture, we turn to the intermediate ramp time $T = 50 \tau$. Here, in the clean system, BO identifies an optimum that it predicts to be extremely narrow in parameter space. The corresponding control path and GP model are illustrated in Fig. 9. The figure also shows the model’s uncertainty, which becomes high, even very close to the discovered optimum. To accommodate the narrow optimum, the length-scale hyperparameters of the GPs become very low during training, making the uncertainty away from evaluated points high.

If there is indeed a direct relation between different types of noise, because the optimum is so narrow, we would expect this control path to be sensitive to unknown noises, such as the experimental potential noise used previously. This is confirmed in Fig. 10. While the protocol evaluates to $F = 81\%$ in the noiseless system, the fidelity quickly drops with the onset of noise and falls below that of the manually

designed protocol for moderate to strong noise.

We can go even further, and ask whether the GP model can not only predict poor performance against an unknown source of noise but if training can be actively shaped to avoid unsuitably narrow optima. GPs provide natural approaches to tackle this issue, because they allow bounding the complexity of the model being constructed. In our example, by limiting the lower bound of the length scale hyperparameter, BO identifies an optimum that lies lower in terms of fidelity but is much wider in parameter space.

As hypothesized – despite only knowing about the ideal system – this new protocol is significantly more robust against the modeled experimental noise (see Fig. 10). It outperforms the more narrow optimum at experimentally realistic noise strengths in a meaningful way. The optimum strategy jumps from the narrow to the wider optimum between $\sigma = 0.05 \hbar/\tau$ and $\sigma = 0.1 \hbar/\tau$ which we interpret as a realization of (ii).

We conclude that the surrogate model has predictive power regarding an optimum’s robustness to unknown sources of noise reiterating our recommended strategy to exploit the information contained in this costly model. We also wish to highlight that limiting the model’s complexity – as has been done above using length scale bounds – can not only produce more robust protocols but also improve the speed of convergence significantly, thus reducing the cost of training. This is apparent in the uncertainties of the models compared in Fig 9 where the restricted model features substantially lower uncertainty after the same number of iterations.

We conclude this discussion with some technical details about the occurrence of (ii) in our system (only) at intermediate ramp times. We attribute this curious behavior to a qualitative transition in the optimal ramping protocol causing issues due to the simple parametrization using linear segments between fixed points in time. In the main text – for both the experimental and the optimized protocol – we observed a crossing in the control fields $t_x(t)$ and $t_y(t)$, where t_x is ramped up from $t_x = 0.0$ to more than unity and subsequently decreased back to $t_x = 1.0 \hbar/\tau$. This feature is absent in the $T = 20 \tau$ ramp – presumably due to a lack of time. Consequently, there must be a crossover at an intermediate total time $20 \tau \leq T \leq 100 \tau$ where both strategies are equally good. This complicates training and is likely exacerbated by the rigid parametrization which does not allow ‘doing nothing’. Similar to longer protocols, the optimized result for $T = 50 \tau$ shown in Fig. 9 features a crossing $t_x(t) = t_y(t)$ at $0 < t < T$ but has to speed up delocalization in the y-direction by ramping Δ_y to a value significantly below 0.0.

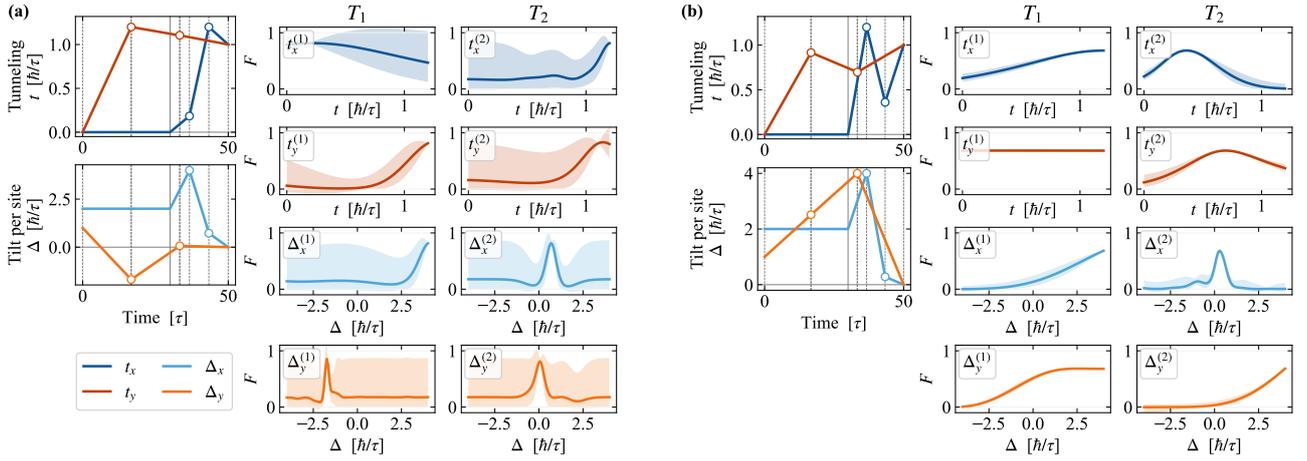


Figure 9: Different GP Models: Results of 2 independent BO runs with 500 iterations each for an intermediate ramp time $T = 50 \tau$. Based on different hyperparameter settings, the models identify different optima. Each subfigure ((a), and (b)) shows the optimized ramping protocol (left) and the modeled fidelity-landscape of the 8 control parameters around the optimum (right). **(a)** The model after training with the same parameter settings as in the main text. An optimum with fidelity $F = 81\%$ is identified. The optimum is very narrow in parameter space and is especially sensitive to the Δ_y control points. The model uncertainties (shaded areas) away from the optimum are large. **(b)** GP model with a $10\times$ lower bound for the length scale hyperparameter than in (a). The lower-complexity model identifies a different optimum in the optimization landscape that lies lower in terms of fidelity ($F = 70\%$) but is much wider. The model uncertainty around the optimum is very low.

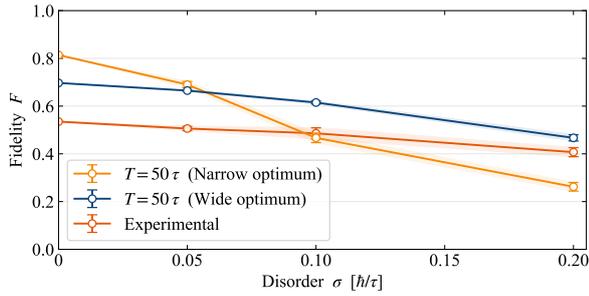


Figure 10: Robust Models: Fidelities of ramping protocols obtained from constrained and unconstrained training at varying disorder strength. The error bars indicate the standard error of the mean fidelity obtained from averaging over 100 disorder realizations. While the ramp trained without any restrictions features the highest fidelity in the ideal system, it rapidly loses fidelity with the onset of experimental noise – even falling below the manually designed protocol when crossing the experimentally realistic noise strength. The ramp trained with a simpler BO model with limited length scales is substantially more robust and retains its performance advantage over manual design in the presence of strong disorder.

Consequently, the control path becomes extremely sensitive to this Δ_y value which turns out to be the optimization parameter that is assigned the shortest length scale.

H Larger System

As an outlook beyond the minimal two-particle Laughlin-like state at filling factor $\nu = 1/2$ that has been realized experimentally, we conclude with first results on the $N = 4$, $L_x = L_y = 6$ Pfaffian-like state at $\nu = 1$ proposed by F. Palm et al. [36]. Following their work, we use the system parameters $U = 4 \hbar/\tau$ and $\phi/2\pi = 0.16$. We then optimize the ground-state preparation from an initial state of particles pinned in the corner of the system in a 1×4 configuration. While pinning the particles in a corner of the system becomes increasingly restrictive when going to larger systems, we choose this state to most naturally extend the optimization problem from the $N = 2$ case. Combined with the larger Hilbert-space dimension $\dim(\mathcal{H}^{N=4}) = 82251$, compared to $\dim(\mathcal{H}^{N=2}) = 136$, this makes for a significantly more challenging preparation task, where randomly chosen protocols are expected to yield almost vanishing fidelity. As the goal of this work is to allow the preparation of more complex states using contemporary experimental setups, we perform optimization with the same total protocol duration and restrictions imposed by Floquet heating as previously for the smaller system. That is, we again optimize protocols of total duration $T = 100 \tau$ with the coupling in x-direction only active for a shorter time $T_x = 40 \tau$. We wish to answer two main questions:

- (i) Can we still find good preparation protocols for the larger system? How do these compare to the adiabatic protocol for the smaller system?
- (ii) Can we use the knowledge gained in the small

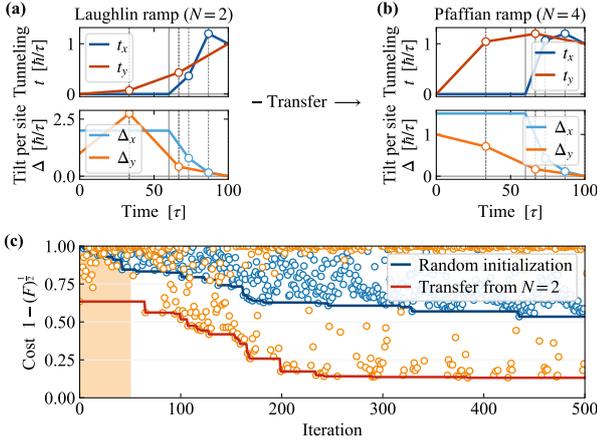


Figure 11: Transfer to Larger System: BO training for the $N = 4$, $L_x = L_y = 6$, $\nu = 1$ state. **(a)** The optimized protocol for the two-particle $\nu = 1/2$ state from Section 4. **(b)** Optimized ramping protocol for the larger $\nu = 1$ state obtained by transferring training data: 50 protocols from the training of (a) are used to initialize the surrogate model. **(c)** Comparison between training initialized with 50 transferred points and random initialization with 10 points. The transfer approach is much more efficient, reaching $F > 70\%$ in less than 200 iterations. The shaded orange region marks the evaluation of transferred data points. The optimized Laughlin ramp prepares the larger $\nu = 1$ state with $F = 13\%$ (Iteration = 0; orange) – significantly better than random protocols, but far from optimal.

system to boost optimization efficiency for the more complex state?

In order to answer these questions, we perform two separate BO runs with 500 iterations each. The first run follows the procedure described in Section 4, using 10 random initial points, while the second run uses initial points collected from training on the small systems following the initialization algorithm described in Section E. The performance of these two approaches is compared in Fig. 11. While the ‘cold-start’ approach that uses random initialization continues to find better and better protocols as training goes on, the rate of improvement is low, reaching only $F = 22\%$ after 500 iterations. The transfer approach performs much better, reaching fidelities as high as $F = 75\%$. While the optimized protocol for the $N = 2$ system is not directly applicable to the larger system ($F = 13\%$), the qualitative strategy does not change. Fig. 12 compares the evolution of the density profile for the optimized protocol in the two system sizes.

The optimization performance was boosted massively by utilizing knowledge of the smaller system. This extends a previous report of efficient scaling in the system size [20] to a case featuring different model parameters (U and α) and a different target state. In summary, the transfer technique allows us to positively answer both (i) and (ii). For our system, optimization allows preparing a significantly more com-

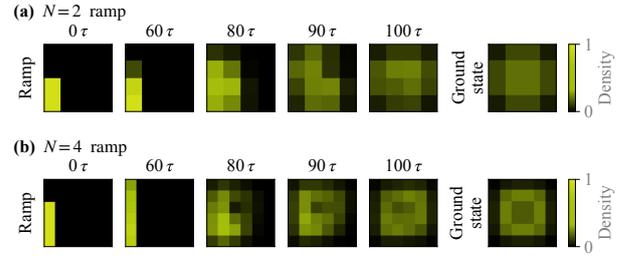


Figure 12: Particle Density: Evolution of the density profile from the initial state of pinned particles comparing different system sizes. The plots on the left show the density distribution at different times during the ramp. The rightmost plot shows the density profile of the (target) ground state. **(a)** Smaller system of 2 particles on 4×4 sites evolving according to the optimized protocol from Section 4. **(b)** Larger system of 4 particles on 6×6 sites evolving according to the optimized protocol obtained via transfer learning.

plex state at a higher fidelity than the manual benchmark on the small system with the same experimental constraints. For reference, applying the manual protocol to the $N = 4$ system results in $F = 3.4\%$ while a simple linear ramp for all control fields has $F = 9.7 \times 10^{-7}$. Additionally, optimization on a small system easily tractable using numerics was able to dramatically reduce the costs of optimization on a much more demanding system.

References

- [1] Immanuel Bloch, Jean Dalibard, and Sylvain Nascimbène. “Quantum simulations with ultracold quantum gases”. *Nature Physics* **8**, 267–276 (2012).
- [2] Christian Gross and Immanuel Bloch. “Quantum simulations with ultracold atoms in optical lattices”. *Science* **357**, 995–1001 (2017).
- [3] Martin Eckstein and Marcus Kollar. “Near-adiabatic parameter changes in correlated systems: Influence of the ramp protocol on the excitation energy”. *New Journal of Physics* **12**, 055012 (2010).
- [4] Julian Léonard, Sooshin Kim, Joyce Kwan, Perrin Segura, Fabian Grusdt, Cécile Repellin, Nathan Goldman, and Markus Greiner. “Realization of a fractional quantum Hall state with ultracold atoms”. *Nature* **619**, 495–499 (2023).
- [5] Patrick Doria, Tommaso Calarco, and Simone Montangero. “Optimal Control Technique for Many-Body Quantum Dynamics”. *Physical Review Letters* **106**, 190501 (2011).
- [6] Steffen J. Glaser, Ugo Boscain, Tommaso Calarco, Christiane P. Koch, Walter Köckenberger, Ronnie Kosloff, Ilya Kuprov, Burkhard Luy, Sophie Schirmer, Thomas Schulte-Herbrüggen, Dominique Sugny, and

- Frank K. Wilhelm. “Training Schrödinger’s cat: Quantum optimal control”. *The European Physical Journal D* **69**, 279 (2015).
- [7] Christiane P. Koch. “Controlling open quantum systems: Tools, achievements, and limitations”. *Journal of Physics: Condensed Matter* **28**, 213001 (2016).
- [8] Christiane P. Koch, Ugo Boscain, Tommaso Calarco, Gunther Dirr, Stefan Filipp, Steffen J. Glaser, Ronnie Kosloff, Simone Montangero, Thomas Schulte-Herbrüggen, Dominique Sugny, and Frank K. Wilhelm. “Quantum optimal control in quantum technologies. Strategic report on current status, visions and goals for research in Europe”. *EPJ Quantum Technology* **9**, 1–60 (2022).
- [9] Jonathan Simon, Waseem S. Bakr, Ruichao Ma, M. Eric Tai, Philipp M. Preiss, and Markus Greiner. “Quantum simulation of antiferromagnetic spin chains in an optical lattice”. *Nature* **472**, 307–312 (2011).
- [10] S. van Frank, M. Bonneau, J. Schmiedmayer, S. Hild, C. Gross, M. Cheneau, I. Bloch, T. Pichler, A. Negretti, T. Calarco, and S. Montangero. “Optimal control of complex atomic quantum systems”. *Scientific Reports* **6**, 34187 (2016).
- [11] Sylvain de Léséleuc, Vincent Lienhard, Pascal Scholl, Daniel Barredo, Sebastian Weber, Nicolai Lang, Hans Peter Büchler, Thierry Lahaye, and Antoine Browaeys. “Observation of a symmetry-protected topological phase of interacting bosons with Rydberg atoms”. *Science* **365**, 775–780 (2019).
- [12] Murphy Yuezhen Niu, Sergio Boixo, Vadim N. Smelyanskiy, and Hartmut Neven. “Universal quantum control through deep reinforcement learning”. *npj Quantum Information* **5**, 1–8 (2019).
- [13] Iris Paparella, Lorenzo Moro, and Enrico Prati. “Digitally stimulated Raman passage by deep reinforcement learning”. *Physics Letters A* **384**, 126266 (2020).
- [14] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando de Freitas. “Taking the Human Out of the Loop: A Review of Bayesian Optimization”. *Proceedings of the IEEE* **104**, 148–175 (2016).
- [15] Roman Garnett. “Bayesian Optimization”. *Cambridge University Press*. Cambridge (2023).
- [16] P. B. Wigley, P. J. Everitt, A. van den Hengel, J. W. Bastian, M. A. Sooriyabandara, G. D. McDonald, K. S. Hardman, C. D. Quinlivan, P. Manju, C. C. N. Kuhn, I. R. Petersen, A. N. Luiten, J. J. Hope, N. P. Robins, and M. R. Hush. “Fast machine-learning online optimization of ultra-cold-atom experiments”. *Scientific Reports* **6**, 25890 (2016).
- [17] Zachary Vendeiro, Joshua Ramette, Alyssa Rudelis, Michelle Chong, Josiah Sinclair, Luke Stewart, Alban Urvoy, and Vladan Vuletić. “Machine-learning-accelerated Bose-Einstein condensation”. *Physical Review Research* **4**, 043216 (2022).
- [18] A. D. Tranter, H. J. Slatyer, M. R. Hush, A. C. Leung, J. L. Everett, K. V. Paul, P. Vernaz-Gris, P. K. Lam, B. C. Buchler, and G. T. Campbell. “Multiparameter optimisation of a magneto-optical trap using deep learning”. *Nature Communications* **9**, 4360 (2018).
- [19] Rick Mukherjee, Harry Xie, and Florian Mintert. “Bayesian optimal control of GHZ states in Rydberg lattices”. *Physical Review Letters* **125**, 203603 (2020).
- [20] Yan-Jun Xie, Han-Ning Dai, Zhen-Sheng Yuan, Youjin Deng, Xiaopeng Li, Yu-Ao Chen, and Jian-Wei Pan. “Bayesian learning for optimal control of quantum many-body states in optical lattices”. *Physical Review A* **106**, 013316 (2022).
- [21] Jeffrey C. Lagarias, James A. Reeds, Margaret H. Wright, and Paul E. Wright. “Convergence Properties of the Nelder–Mead Simplex Method in Low Dimensions”. *SIAM Journal on Optimization* **9**, 112–147 (1998).
- [22] Swagatam Das and Ponnuthurai Nagaratnam Suganthan. “Differential Evolution: A Survey of the State-of-the-Art”. *IEEE Transactions on Evolutionary Computation* **15**, 4–31 (2011).
- [23] R. B. Laughlin. “Anomalous Quantum Hall Effect: An Incompressible Quantum Fluid with Fractionally Charged Excitations”. *Physical Review Letters* **50**, 1395–1398 (1983).
- [24] C. Broholm, R. J. Cava, S. A. Kivelson, D. G. Nocera, M. R. Norman, and T. Senthil. “Quantum spin liquids”. *Science* **367**, eaay0668 (2020).
- [25] Anton Mazurenko, Christie S. Chiu, Geoffrey Ji, Maxwell F. Parsons, Márton Kanász-Nagy, Richard Schmidt, Fabian Grusdt, Eugene Demler, Daniel Greif, and Markus Greiner. “A cold-atom Fermi–Hubbard antiferromagnet”. *Nature* **545**, 462–466 (2017).
- [26] Christie S. Chiu, Geoffrey Ji, Anton Mazurenko, Daniel Greif, and Markus Greiner. “Quantum State Engineering of a Hubbard System with Ultracold Fermions”. *Physical Review Letters* **120**, 243201 (2018).
- [27] S. Rosi, A. Bernard, N. Fabbri, L. Fallani, C. Fort, M. Inguscio, T. Calarco, and S. Montangero. “Fast closed-loop optimal control of ultracold atoms in an optical lattice”. *Physical Review A* **88**, 021601 (2013).
- [28] J. J. Sørensen, J. H. M. Jensen, T. Heinzl, and J. F. Sherson. “QEngine: A C++ library for quantum optimal control of ultracold atoms”. *Computer Physics Communications* **243**, 135–150 (2019).
- [29] Carl Edward Rasmussen and Christopher K. I.

- Williams. “Gaussian Processes for Machine Learning”. *The MIT Press*. (2005).
- [30] M. Eric Tai, Alexander Lukin, Matthew Rispoli, Robert Schittko, Tim Menke, Dan Borgnia, Philipp M. Preiss, Fabian Grusdt, Adam M. Kaufman, and Markus Greiner. “Microscopy of the interacting Harper–Hofstadter model in the two-body limit”. *Nature* **546**, 519–523 (2017).
- [31] P. G. Harper. “The General Motion of Conduction Electrons in a Uniform Magnetic Field, with Application to the Diamagnetism of Metals”. *Proceedings of the Physical Society. Section A* **68**, 879 (1955).
- [32] Douglas R. Hofstadter. “Energy levels and wave functions of Bloch electrons in rational and irrational magnetic fields”. *Physical Review B* **14**, 2239–2249 (1976).
- [33] Christoph Sträter and André Eckardt. “Interband Heating Processes in a Periodically Driven Optical Lattice”. *Zeitschrift für Naturforschung A* **71**, 909–920 (2016).
- [34] Rick Mukherjee, Frédéric Sauvage, Harry Xie, Robert Löw, and Florian Mintert. “Preparation of ordered states in ultra-cold gases using Bayesian optimization”. *New Journal of Physics* **22**, 075001 (2020).
- [35] P. Schauß, J. Zeiher, T. Fukuhara, S. Hild, M. Cheneau, T. Macrì, T. Pohl, I. Bloch, and C. Gross. “Crystallization in Ising quantum magnets”. *Science* **347**, 1455–1458 (2015).
- [36] F. A. Palm, M. Buser, J. Léonard, M. Aidelsburger, U. Schollwöck, and F. Grusdt. “Bosonic Pfaffian state in the Hofstadter-Bose-Hubbard model”. *Physical Review B* **103**, L161101 (2021).
- [37] F. A. Palm, J. Kwan, B. Bakkali-Hassani, M. Greiner, U. Schollwöck, N. Goldman, and F. Grusdt. “Growing extended Laughlin states in a quantum gas microscope: A patchwork construction”. *Physical Review Research* **6**, 013198 (2024).
- [38] Niranjana Srinivas, Andreas Krause, Sham M. Kakade, and Matthias Seeger. “Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design”. *IEEE Transactions on Information Theory* **58**, 3250–3265 (2012).
- [39] J. M. Zhang and R. X. Dong. “Exact diagonalization: The Bose–Hubbard model as an example”. *European Journal of Physics* **31**, 591 (2010).
- [40] J. R. Dormand and P. J. Prince. “A family of embedded Runge-Kutta formulae”. *Journal of Computational and Applied Mathematics* **6**, 19–26 (1980).
- [41] Philip Zupancic, Philipp M. Preiss, Ruichao Ma, Alexander Lukin, M. Eric Tai, Matthew Rispoli, Rajibul Islam, and Markus Greiner. “Ultra-precise holographic beam shaping for microscopic quantum control”. *Optics Express* **24**, 13881–13893 (2016).