Quantum-accessible reinforcement learning beyond strictly epochal environments

A. Hamann¹ · V. Dunjko² · S. Wölk^{1,3}

Received: 27 August 2020 / Accepted: 7 May 2021 / Published online: 2 August 2021
© The Author(s) 2021

Abstract

In recent years, quantum-enhanced machine learning has emerged as a particularly fruitful application of quantum algorithms, covering aspects of supervised, unsupervised and reinforcement learning. Reinforcement learning offers numerous options of how quantum theory can be applied, and is arguably the least explored, from a quantum perspective. Here, an agent explores an environment and tries to find a behavior optimizing some figure of merit. Some of the first approaches investigated settings where this exploration can be sped-up, by considering quantum analogs of classical environments, which can then be queried in superposition. If the environments have a strict periodic structure in time (i.e. are strictly episodic), such environments can be effectively converted to conventional oracles encountered in quantum information. However, in general environments, we obtain scenarios that generalize standard oracle tasks. In this work, we consider one such generalization, where the environment is not strictly episodic, which is mapped to an oracle identification setting with a changing oracle. We analyze this case and show that standard amplitude-amplification techniques can, with minor modifications, still be applied to achieve quadratic speed-ups. In addition, we prove that an algorithm based on Grover iterations is optimal for oracle identification even if the oracle changes over time in a way that the “rewarded space” is monotonically increasing. This result constitutes one of the first generalizations of quantum-accessible reinforcement learning.

Keywords Reinforcement learning · Quantum-classical hybrid agent · Amplitude amplification

1 Introduction

In the last few years, there has been much interest in combining quantum computing and machine learning algorithms. In the domain of quantum-enhanced machine learning, the objective is to utilize quantum effects to speed-up or otherwise enhance the learning performance. The possibilities for this are numerous (Dunjko and Briegel 2018). E.g. variational circuits can be used as a type of “quantum neural network” (more precisely, using them as function approximators which cannot be evaluated efficiently on a conventional computer), which can be

trained as a supervised learning (classification) (Havlíček et al. 2019; Farhi and Neven 2018) or unsupervised learning model (generative models) (Aïmeur et al. 2013). There also exist various approaches where algorithmic bottlenecks of classical algorithms are sped-up, via annealing methods (Farhi and Neven 2018), quantum linear-algebraic methods (Harrow et al. 2009), or via sampling enhancements (Dunjko et al. 2016). If the data is assumed to be accessible in a quantum form (“quantum database”), then anything from polynomial to exponential speed-ups of classical algorithms may be possible (Biamonte et al. 2017; Dunjko and Briegel 2018; Chia et al. 2019; Gyurik et al. 2020)¹.

Modern reinforcement learning (RL), an interactive mode of learning, combines aspects of supervised and unsupervised learning, and consequently allows a broad spectrum of possibilities how quantum effects could help.

In RL (Sutton and Barto 1998; Russell and Norvig 2003; Briegel and De las Cuevas 2012), we talk about a learning

✉ S. Wölk
sabi.woelk@gmail.com

¹ Institut für Theoretische Physik, Universität Innsbruck, Technikerstraße 21a, 6020, Innsbruck, Austria

² LIACS, Leiden University, Niels Bohrweg 1, 2333, CA, Leiden, The Netherlands

³ Present address: Institute of Quantum Technologies, German Aerospace Center (DLR), D-89077 Ulm, Germany

¹In recent times, due to progress in quantum-inspired algorithms, the domain of algorithms where exponential speed-ups are to be expected has reduced, but many possibilities for classically intractable computations still exist.

agent which interacts with an environment, by performing actions, and perceiving the environmental states, and has to learn a “correct behavior”—the optimal policy—by means of a feedback rewarding signal. Unlike a stationary database, the environment has its own internal memory (a state), which the agent alters with its actions.

In quantum-enhanced RL, we can identify two basic scenarios: (i) where quantum effects can be used to speed up the internal processing (Paparo et al. 2014; Jerbi et al. 2019), and the interaction with the environment is classical, and (ii) where the interaction with the environment (and the environment itself) is quantum. The first framework for such “quantum-accessible” reinforcement learning modeled the environment as a sequence of quantum channels, acting on a communication register, and the internal environmental memory—this constitutes a direct generalization of an unknown environment as a map-with-memory (other options are discussed shortly). In this case, the action of the environment cannot be described as unitary mapping, without considering the entire memory of the environment. In general, this memory is inaccessible to the agent. However, as discussed in Dunjko et al. (2016), under the assumptions that the environmental memory can be purged or uncomputed in pre-defined periods, such blocks of interaction do become a (time-independent) unitary and amenable to oracle computation techniques. For instance, in Dunjko et al. (2016), it was shown that the task of identifying a sequence of actions which leads to a first reward (a necessary step before any true learning can commence) can be sped up using quantum search techniques, and in Dunjko et al. (2017), it was shown how certain environments encode more complex oracles—e.g. Simon’s oracle and Recursive Fourier Sampling oracles, leading to exponential speed-ups over classical methods.

For the above techniques to work, however, the purging of all of environmental memory is necessary to achieve time-independent unitary mappings. However, real task environments are typically not (strictly) episodic, motivating the question of what can be achieved in these more general cases. Here, we perform a first step towards generalization by considering environments where the length of the episode can change, but this is signaled and the estimate of the episode lengths is known. This RL scenario is well-motivated and fortunately maps to an oracle identification problem where the oracles change. While this generalizes standard oracular settings, it is still sufficiently simple such that we can employ standard techniques (essentially amplitude amplification) and prove the optimality of our strategies for oracle identification problems with changing oracles and increasing rewarded space.

The paper is organized as follows. We will first summarize the basics scenario of quantum-accessible

reinforcement learning in Section 2 and discuss the mappings from constrained (episodic) RL scenarios to oracle identification. We show how this must be generalized for more involved environments, prompting our definition of the “changing oracle” problem stemming from certain classes of RL environments. In Section 3, we focus on the changing oracle problem, analyze the main regimes, and provide an upper bound for the average success probability for the case of monotonically increasing rewarded space in Section 3.1. We prove in Section 3.2 that performing consecutive Grover iterations saturates this bound. We then discuss the more general case of only overlapping rewarded spaces in Section 3.3. In Section 3.4, we provide a numerical example demonstrating the possible advantages of consecutive Grover iterations with changing oracles. We conclude by summarizing our results, by discussing possible extensions, and by noting on the implications of our results of the changing oracle problem for QRL in Section 4.

2 Quantum-accessible reinforcement learning

RL can be described as an interaction of a learning agent A with a task environment E via the exchange of messages out of a discrete set which we call actions $\mathcal{A} = \{a_j\}$ (performed by the agent) and percepts $\mathcal{S} = \{s_j\}$ (issued by the environment). In addition, the environment also issues a scalar reward $\mathcal{R} = \{r_j\}$, which informs the agent about the quality of the previous actions and can be defined as being a part of the percepts. The goal of the agent is to receive as much reward as possible in the long term.

In theory of RL, the most studied environments are exactly describable by a Markov decision process (MDP). An MDP is specified by a transition mapping $T(s'|a, s) \in \mathbb{R}^{\geq 0}$ and a reward function $R(s, a) \in \mathbb{R}$. The transition mapping T specifies the probability of the environment transiting from state s to s' , provided the agent performed the action a , whereas the reward function assigns a reward value to a given action of an agent in a given environmental state.

Note that in standard RL, the agent does not have a direct access to the mapping T , but rather to learn it, it must explore, i.e. to act in the environment which is governed by T . On the other hand, in dynamical programming problems (intimately related to RL), one often assumes access to the functions T and R directly. This distinction leads to two different tasks on how agent-environment interaction can be quantized.

In recent works (Cornelissen 2018; Neukart et al. 2018; Levit et al. 2017), coherent access to the transition mapping T is assumed; in this case, lower quantum bounds for finding the optimal policy have been found (Ronagh 2019).

In this paper, we consider the other class of generalization, proposed first in Dunjko et al. (2016). Here, the agent-environment interaction is modeled as a communication between an agent (A) and the environment (E) over a joint communication channel (C), thus in a three-partite Hilbert space $\mathcal{H}_E \otimes \mathcal{H}_C \otimes \mathcal{H}_A$, denoting the memory of the environment, the communication channel, and the memory of the agent. The two parties A and E interact with each other by performing alternately completely positive trace preserving (CPTP) maps on their own memory and the communication channel. Different AE combinations are defined as equivalent in the classical sense, if their interactions are equivalent under constant measurements of C in the computational basis. For classical luck favoring AE settings with a deterministic strictly epochal environment E, it is possible to create a classical equivalent quantum version $A^q E^q$ which outperforms AE in terms of a given figure of merit as shown in Dunjko et al. (2016).

2.1 Strictly epochal environments

This can be achieved by slightly modifying the maps as to purge the environmental memory which couples to the overall interaction preventing a unitary time evolution of the agents memory. A detailed discussion of this procedure and necessary condition on the setting is outlined in Dunjko et al. (2016). However, for our setting, it is sufficient that the interaction of the agent with the environment can be effectively described as oracle queries. Specifically if environments are strictly episodic, meaning after some fixed number of steps the setting is re-set to an initial condition, then the environmental memory can be uncomputed or released to the agent at the end of an epoch. With this modification (called memory scavenging and hijacking in earlier works), blocks of interactions effectively act as one, time-independent unitary O , which can be queried using standard quantum techniques to obtain an advantage. We encode different actions the agent can perform, like $a \in \{0, 1\}$, into orthogonal quantum states, i.e. as $\{|0\rangle, |1\rangle\}$. As a result, the complete sequence of actions $\mathbf{a} = a_1, \dots, a_M$ the agent executes during one epoch of length M is encoded in the product state $|\mathbf{a}\rangle = |a_1\rangle \otimes |a_2\rangle \otimes \dots \otimes |a_M\rangle$. For strictly epochal environments, it is possible to re-express the effect of the environment by an unitary oracle

$$O|\mathbf{a}\rangle = \begin{cases} -|\mathbf{a}\rangle & \text{if } \mathbf{a} \in W \\ |\mathbf{a}\rangle & \text{else} \end{cases} \quad (1)$$

Here, W denotes the rewarded space containing all sequences of actions of length M which obtained a reward $r(\mathbf{a})$ larger than a predefined limit.

A learning agent can use such an oraculized environment to find rewarded action sequences faster. For this purpose, the agent prepares an equal superposition state of all

possible action sequences

$$|\psi\rangle = \frac{1}{\sqrt{N}} \sum_{i=1}^N |\mathbf{a}_i\rangle \quad (2)$$

with typically $N = |\mathcal{A}|^M$. Then, it interacts with the environment and thus effectively queries the oracle O . Afterwards, it performs a reflection over the initial state $|\psi\rangle$. In this way, it can perform amplitude amplification performing consecutive Grover iterations (Grover 1997; 1998; Brassard et al. 2000) $G_\psi |\psi\rangle$ with

$$G_\psi = (\mathbb{1} - 2|\psi\rangle\langle\psi|) O. \quad (3)$$

The agent can increase the probability to find a first rewarded sequence by performing several rounds of amplitude amplification.

The first found rewarded action sequence is in general not the optimal one. However, in general, it provides information which can help to find an optimal strategy. A quantum-enhanced learning agent can use the found rewarded sequence of actions to learn by using classical policy updates. Thus, quantum-enhanced reinforcement learning combines quantum search and classical reinforcement learning as, e.g. demonstrated experimentally in Saggio et al. (2021). An agent with access to such an oracle thus finds in average the first rewarded sequence faster which increases in so-called luck-favoring settings (Dunjko et al. 2016) also the probability to be rewarded in the future. This approach leading to a quadratic speed-up in exploration can be applied to many settings and even super-polynomial or exponential improvements can be generated for special RL settings (Dunjko et al. 2017).

2.2 Beyond strictly epochal environments

The simplest scenarios of task environments, which cannot be reformulated as an oracular problem, are arguably those which involve two oracles. We will consider this slight generalization in this work, as it still allows for a relatively simple treatment. This setting includes environments which simply change as a function of time such as reinforcement learning for managing power consumption or channel allocation in cellular telephone systems (Han 2018; Tesauro et al. 2008; da Silva et al. 2006; Singh and Bertsekas 1996). If the instances of change are known, again the blocking is possible, in which case we obtain the setting where we can realize access to an oracle but which changes as a function of time. Closely related to this is a more specific case of variable episode length. This setting, although more special, is in particular interest in RL. Episodic environments are usually constructed by taking an arbitrary environment, and establishing a cut-off after a certain number of steps. The resulting object is again an environment derived from the

initial setting. This construction is special in that given any sequence of actions $\mathbf{a}$ which is rewarding in a derived environment with cut-off after m steps, any sequence of actions in the environment which has a larger cut off $M > m$ which has $\mathbf{a}$ as a prefix is rewarded in the second. An example for such an environment is the Grid-world problem which consists in navigating a maze and the task is to find a specific location that is rewarded (Russell and Norvig 2003; Sutton and Barto 1998; Melnikov et al. 2018).

The classical scenarios described above, under oracularization techniques, map onto the changing oracle problem (described in detail in the following section) where at a given time an oracle $\tilde{O}$ is exchanged by a different oracle O . This generalization especially captures the scenario of a single increment of an epoch length from m to $M > m$ for search in QRL. In this special case, the rewarded space $\tilde{W}$ of $\tilde{O}$ is a subspace of W of O . We will prove that the optimal algorithm in this case is given by a Grover search with a continuous coherent time evolution using both oracles consecutively. However, continuing the coherent time evolution of a Grover search can be suboptimal when $\tilde{W} \not\subseteq W$. The arguments following in the next section can be used iteratively to describe multiple changes/increments of the rewarded space.

3 The changing oracle problem

The situation above can be abstracted as a “changing oracle” problem which we specify here. As suggested, we consider an “oracle” to be a standard phase-flip oracle, such that $O|x\rangle = (-1)^{f(x)}|x\rangle$, where $f : X \rightarrow \{0, 1\}$ is a characteristic function on a set of elements X , with $|X| = N$; in our case X denotes sequences of actions of some prescribed length. The rewarded set is denoted by $W = \{x \in X | f(x) = 1\}$, and the states $|x\rangle$ denote a (known) orthonormal basis.

In the changing oracle problem, we consider two oracles $\tilde{O}$ and O , with respective rewarded sets $\tilde{W}$ and W . The problem specifies two time intervals, phases, in which only one of the two oracles is available: time-steps $1 \leq k \leq K$ during which only access to $\tilde{O}$ is available, and time-steps $K+1 \leq k < K+J$ during which only access to the second oracle O is available.

For simplicity, we assume that the values of K, J, N as well as the sizes of the rewarded sets $|\tilde{W}| = \tilde{n}$ and $|W| = n$ are known in advance, and in general, the objective is to either output an $x \in \tilde{W}$ before K , or, to output $x \in W$ in the remainder of the time. We will refer to both x as the solution. However, the exact time when the oracle changes, and does K and J , is not important and can be unknown as we show later. Unless K is in $\Omega(\sqrt{N/\tilde{n}})$, in general attempts to find a solution in the first phase will have a

very low success probability no matter what we do due to the optimality of Grover’s search. However, even in this case, having access to $\tilde{O}$ in the first phase may improve our chances to succeed in the second. This is the setting we consider.

The optimal strategies vitally depend on the known relationship between W and $\tilde{W}$. We will first briefly discuss all possible setting before focusing on the most interesting cases. Note, in this paper, we are not looking for a strategy which uses a minimal number of queries until a solution is found, but rather, a strategy which maximizes the success probability for a fixed number of queries. However, it is also known that Grover’s search achieves the fastest increase of success probability (Zalka 1999). Note, the here described algorithms can be also used to optimize the number of queries. However, the corresponding figure of merit, which needs to be optimized, has to be defined precisely for such tasks.

In the worst case, there may be no known correlation between W and $\tilde{W}$. In this case, we have no advantage from having access to $\tilde{O}$, and the optimal strategy is a simple Grover’s search in the second phase.

Another case with limited interest is when W and $\tilde{W}$ are known to be disjoint. In this case, the first oracle might be used to constrain the search space to the complement $\tilde{W}^c$, which contains W . The lower bounds for this setting are easy to find: we can assume that at K the set $\tilde{W}$ is made known (any state we could have generated using $\tilde{O}$ can be generated with this information). However, in this case, the optimal strategy is still to simply apply quantum search over the restricted space $\tilde{W}^c$ if it can be fully specified. But since we most often encounter cases where $\tilde{n} = |\tilde{W}|$ is (very) small compared to N , the improvement that could be obtained is also minor.

Similar reasoning follows also when the sets are not disjoint, but the intersection is small compared not just to N , but to $|W|$ and $|\tilde{W}|$. In this case, again we can find lower bounds by assuming that the non-overlapping complement becomes known. In addition, we assume that we can prepare any quantum state, which has an upper bound on the overlap with any state corresponding to the intersection, $x \in W \cap \tilde{W}$. Then, the optimal strategy is again governed by the optimality of Grover-based amplitude amplification²

This brings us to the situations which are more interesting, specifically, when the overlap $W_a = W \cap \tilde{W}$ is large (see Appendix A for exact definition).

²More generally, we can allow only states for which, under no quantum channel, allow us to determine such x with probability better than given by Grover iterations. This setting is a bit more involved, but it should be clear that as long as this probability is very small, whatever we do in the next phase, cannot be much better than starting from scratch.

Due to our motivation stemming from aforementioned RL settings, we are particularly interested in the case when $\tilde{W} \subseteq W$, for which we give the optimal strategy, which turns out to be essentially Grover's amplification where we "pretend" that the oracles hadn't changed.

The other cases, $W \subseteq \tilde{W}$, and the more generic case where the overlap is large, but no containment hold are less interesting for our purpose, so we briefly discuss the possible strategies without proofs of optimality.

3.1 Increasing rewarded spaces: upper bound on average final success probabilities

In the following, we consider the above-described changing oracle problem with monotonically increasing rewarded spaces $\tilde{W} \subseteq W$ and derive upper bounds for the maximal average success probability p_{K+J} of finding an element $x \in W$ at the end of the second phase. The changing oracle problem is outside the standard settings for which various lower bounding techniques have been developed (Arunachalam et al. 2019; Ambainis 2002; 2006), but the setting is simple enough to be treatable by modifying and extending techniques introduced to lower bound unstructured search problems (Zalka 1999).

To find upper bounds of the success probability, we first prove that we can restrict our search for optimal strategies to averaged strategies as defined in Appendix B. This induces certain symmetries which restrict the optimization to an optimization of two angles α and Δ , one for each phase. Finally, we derive bounds $\alpha(K)$ and $\Delta(J)$ for these angles depending on K, J which in turn restrict the optimal success probability p_{K+J} .

The search for an optimal strategy can be limited to strategies based on pure states and unitary time evolutions since it is possible to purify any search strategy by going from a Hilbert space $\mathcal{H}_A$ spanned by $\{|x\rangle\}$ into a larger Hilbert space $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$. As a consequence, every search strategy $T = (\{U_k\}, |\psi(0)\rangle)$ based on $K + J$ oracle queries can be described by a set of $K + J$ unitaries U_k and initial state $|\psi(0)\rangle$. Our knowledge about possible rewarded items after k oracles queries is then encoded in the quantum state

$$|\psi(k)\rangle = U_k O_k \cdots U_1 O_1 |\psi(0)\rangle \quad (4)$$

with $O_k = \tilde{O}$ for $1 \leq k \leq K$ and $O_k = O$ for $K + 1 \leq k \leq J$. The success probability at the end of the second phase is then given by

$$p_{K+J} = \text{Tr}[P_W |\psi(K+J)\rangle \langle \psi(K+J)|] \quad (5)$$

with

$$P_W = \left(\sum_{x \in W} |x\rangle_A \langle x| \right) \otimes \mathbb{1}_B. \quad (6)$$

Our goal is to maximize the success probability p_{K+J} average over all possible functions $\tilde{f}(x)$ and $f(x)$ with fixed sizes of the rewarded spaces $|\tilde{W}| = \tilde{n}$ and $|W| = n \geq \tilde{n}$. Different realization of $\tilde{f}(x)$ and $f(x)$ can be generated by substituting all oracle queries O_k by $\sigma O_k \sigma^\dagger$ and the projector P_W by $\sigma P_W \sigma^\dagger$ where σ denote a permutation operator acting on $\mathcal{H}_A$. As a consequence, an optimal strategy is a strategy T which maximizes

$$\bar{p}_T = \frac{1}{N!} \sum_{\sigma \in \Sigma_A} p_T(\sigma) \quad (7)$$

with

$$p_T(\sigma) = \text{Tr} \left[\sigma P_W \sigma^\dagger |\psi(k, \sigma)\rangle \langle \psi(k, \sigma)| \right] \quad (8)$$

$$|\psi(k, \sigma)\rangle_{AB} = U_k \sigma O_k \sigma^\dagger \cdots U_1 \sigma O_1 \sigma^\dagger |\psi(0)\rangle_{AB} \quad (9)$$

at the end of the second phase such that $k = K + J$. Here, Σ_A denotes the set of all possible permutations in $\mathcal{H}_A$.

We can further limit the search for optimal strategies to averaged strategies $\bar{T}$ as defined Appendix B because

Lemma 1 *The success probability $p_{\bar{T}}(\sigma)$ of the averaged strategy $\bar{T}$ is equal to the average success probability $\bar{p}_T$ of the strategy T for every permutation $\sigma \in \Sigma_A$.*

as proven in Appendix B. In the following, we consider only average strategies such that $p = \bar{p}$ and therefore omit the "bar" denoting an average value.

In addition, these strategies lead to symmetry properties of the unitaries U_k and resulting states $|\psi(k)\rangle$ under permutations σ as outlined in detail in Appendix B. Therefore, we can restrict the initial states $|\psi(0)\rangle$ to states with equal probability $q(x) = \text{Tr}[(|\psi\rangle_{AB} \langle \psi|) \cdot (|x\rangle_A \langle x| \otimes \mathbb{1}_B)]$ for all elements x . An example for such a symmetric state is the initial state of the Grover search algorithm given by the equal superposition $\sum_x |x\rangle_A / \sqrt{N}$. Yet, many other symmetric initial states are possible due to the additional degrees of freedom resulting from the additional Hilbert space $\mathcal{H}_B$. The unitaries U_k cannot break the symmetry between elements $|x\rangle$. Only the oracles $\tilde{O}$ and O can break the symmetry between rewarded and not-rewarded elements.

These symmetry properties will limit the optimization overall strategies to an optimization of a few parameters or angles as we will outline below. These parameters are then again upper bounded by the optimality of Grover search.

We can decompose the state $|\psi(K)\rangle$ at the end of the first phase into a rewarded and a not-rewarded component with respect to the eigenstates of the second oracle O . The not-rewarded or losing component $|\ell\rangle = |\ell_s\rangle$ is symmetric with respect to elements $x \in \mathcal{L}$. However, the rewarded component $|w\rangle$ is not completely symmetric because $\tilde{O}$ breaks the symmetry between elements $x \in \tilde{W} \cap W$

and $x \in W \setminus \tilde{W}$. Thus, we can further decompose the rewarded component into a symmetric component $|w_s\rangle$ and a component $|w_\perp\rangle$ orthogonal to it (see Appendix C). As a result, the state $|\psi(K)\rangle$ is given by

$$|\psi(K)\rangle = \cos \varepsilon |\phi_s\rangle + \sin \varepsilon |\phi_\perp\rangle \quad (10)$$

with the symmetric component

$$|\phi_s\rangle = \sin \phi |w_s\rangle + \cos \phi |\ell_s\rangle \quad (11)$$

and the orthogonal rewarded component

$$|\phi_\perp\rangle = |w_\perp\rangle. \quad (12)$$

The angles ε and ϕ are parameters depending on the strategy performed during the first phase. Their values are bounded by the success probability at the end of the first phase given by

$$p_K = \cos^2 \varepsilon \sin^2 \phi + \sin^2 \varepsilon. \quad (13)$$

The time evolution during the second phase described by $V = U_{K+J} O \cdots U_{K+1} O$ is also symmetric and thus transforms the symmetric component $|\phi_s\rangle$ into a symmetric component and $|w_\perp\rangle$ into a component orthogonal to $V|\phi_s\rangle$. As a consequence, the final success probability p_{K+J} can be divided into

$$p_{K+J} = \cos^2(\varepsilon) p_s + \sin^2(\varepsilon) p_\perp \quad (14)$$

with (see Appendix C)

$$p_s = \text{Tr} [P_{\mathcal{W}} V |\phi_s\rangle \langle \phi_s| V^\dagger] \quad (15)$$

$$p_\perp = \text{Tr} [P_{\mathcal{W}} V |w_\perp\rangle \langle w_\perp| V^\dagger]. \quad (16)$$

The reward probability $p_\perp$ of the orthogonal part is maximal if $p_\perp = 1$ which can be achieved if, e.g. V acts on $|w_\perp\rangle$ as identity. We parametrize the reward probability of the symmetric part via $p_s = \sin^2(\phi + \Delta)$, where the parameter Δ quantifies the increase of p_s during the second phase. Thus, Δ depends on the strategy performed during the second phase. We can quantify the final success probability via

$$p_{K+J} \leq \cos^2(\varepsilon) \sin^2(\phi + \Delta) + \sin^2(\varepsilon) \quad (17)$$

$$\leq 1 - \cos^2(\varepsilon) \cos^2(\phi + \Delta). \quad (18)$$

With the help of Eq. 13, we can rewrite $\cos^2 \varepsilon$ via $\cos^2 \varepsilon = (1 - p_K) / \cos^2 \phi$ leading to

$$p_{K+J} \leq 1 - (1 - p_K) \frac{\cos^2(\phi + \Delta)}{\cos^2 \phi}. \quad (19)$$

As a consequence, p_{K+J} is monotonically increasing with p_K , ϕ , Δ provided $0 \leq \phi \leq \pi/2$ and $0 \leq \phi + \Delta \leq \pi/2$. Thus, an optimal strategy optimizes p_K and ϕ during the first phase and Δ during the second phase.

If we denote by

$$\sin^2 \alpha = \text{Tr} [P_{\tilde{W}} |\psi(K)\rangle \langle \psi(K)|] \quad (20)$$

the reward probability at the end of the first phase according to the first oracle $\tilde{O}$, then the success probability according to the second oracle O at this point is given by

$$p_K = \sin^2 \alpha + \cos^2 \alpha \frac{n_+}{n_+ + n_\ell} \quad (21)$$

following Eqs. 79 and 80 in Appendix C. Here, $n_+ = |\mathcal{W}_+|$ with $\mathcal{W}_+ = \tilde{\mathcal{L}} \cap W$ denotes the number of items x marked only by the second oracle O as rewarded and $n_\ell = |\mathcal{L}|$ the number of losing items according to O . Thus, p_K increases monotonically with α for $0 \leq \alpha \leq \pi/2$.

The angle ϕ is also upper bounded by α via (see Appendix C, Eq. 93)

$$\tan \phi \leq \tan \alpha \sqrt{\frac{\tilde{n}(n_+ + n_\ell)}{(\tilde{n} + n_+)n_\ell}} + \sqrt{\frac{n_+^2}{(\tilde{n} + n_+)n_\ell}}. \quad (22)$$

This bound also increases monotonically with α for $0 \leq \alpha \leq \pi/2$. As a result, the final success probability is upper bounded by the maximal achievable angles α (defined via the strategy during the first phase) and Δ (during the second phase) within the range $0 \leq \alpha \leq \pi/2$ and $0 \leq \phi(\alpha) + \Delta \leq \pi/2$.

The angles α and Δ can be upper bound with the help of a generalization of the optimality proof of Grover's algorithm from Zalka (1999) which can be stated in the following way

Lemma 2 *Given an oracle O which marks exactly n out of N items as rewarded, then performing Grover's quantum search algorithm gives the maximal possible average success probability $p_K = \sin^2[(2K + 1)\nu]$ for up to $0 < K < \pi/(4\nu) - 1/2$ with $\sin^2 \nu = n/N$.*

The proof of this lemma follows the optimality proof from Zalka for $n = 1$ given in Zalka (1999). We outline the difference in the proof for $n > 1$ in Appendix E. In general, the angle $2K\nu$ does not only limit the maximal success probability via $p \leq \sin^2[(2k + 1)\nu]$ when starting from a random guess, equal to $p_0 = \sin^2 \nu = n/N$, but to $p \leq \sin^2[2k\nu + \phi]$ when starting from any fixed initial success probability $p_0 = \sin^2 \phi$, as we also outline in Appendix E.

As a consequence, the maximal angle α is bounded by $\alpha \leq (2K + 1)\tilde{\nu}$ with $\sin^2 \tilde{\nu} = \tilde{n}/N$ which follows directly from Lemma 2 provided $(2K + 1)\tilde{\nu} \leq \pi/2$. And the reward probability of p_s is limited by $\sin^2(\phi + \Delta)$ with $\Delta < 2J\nu$ provided $2J\nu + \phi \leq \pi/2$.

3.2 Grover search is optimal for monotonically increasing rewarded spaces

In this section, we determine the (average) success probability p_{K+J} for the here defined changing oracle problem obtained via a generalized Grover algorithm and

show that it saturates the in Section 3.1 derived bound. Grover's algorithm starts in a equal superposition state given by

$$|\psi(0)\rangle = \frac{1}{\sqrt{N}} \sum_{x=1}^N |x\rangle \quad (23)$$

$$= \sin \tilde{v} |\tilde{w}\rangle + \cos \tilde{v} |\tilde{\ell}\rangle \quad (24)$$

with

$$|\tilde{w}\rangle = \frac{1}{\sqrt{\tilde{n}}} \sum_{x \in \tilde{\mathcal{W}}} |x\rangle \quad (25)$$

$$|\tilde{\ell}\rangle = \frac{1}{\sqrt{N - \tilde{n}}} \sum_{x \in \tilde{\mathcal{L}}} |x\rangle. \quad (26)$$

All unitaries U_k for $1 \leq k \leq K + J$ are given by

$$U_k = \mathbb{1} - |\psi(0)\rangle\langle\psi(0)|. \quad (27)$$

The time evolution during the first phase with oracle $\tilde{O}$ leads to a rotation of $|\psi(0)\rangle$ by an angle $2K\tilde{v}$ in the plane spanned by $|\tilde{w}\rangle$ and $|\psi(0)\rangle$ as depicted in Fig. 1. The state

Fig. 1 Visualization of the time evolution of $|\psi(0)\rangle$ under Grover iterations with changing oracles $O_k = \tilde{O}$ for $1 \leq k \leq K$ and $O_k = O$ for $K + 1 \leq k \leq J$ and $\tilde{\mathcal{W}} \subseteq \mathcal{W}$. The rewarded space (red plane) of O is spanned by $\{|w_s\rangle, |w_\perp\rangle\}$. The equal superposition state $|\psi(0)\rangle$ is rotated along the blue circle by an angle $2K\tilde{v}$ during the first phase leading to the state $|\psi(K)\rangle$. Consecutively, this state is rotated along the green circle changing only its component $|\phi_s\rangle$ but not $|w_\perp\rangle$

at the end of the first phase is given by

$$|\psi(K)\rangle = \sin[(2K + 1)\tilde{v}]|\tilde{w}\rangle + \cos[(2K + 1)\tilde{v}]|\tilde{\ell}\rangle \quad (28)$$

and thus saturates the upper limit $\alpha = (2K + 1)\tilde{v}$ leading to a maximal p_K and α . To describe the time evolution during the second phase, we perform a basis transformation into the new basis

$$|\ell_s\rangle = \frac{1}{\sqrt{n_\ell}} \sum_{x \in \mathcal{L}} |x\rangle \quad (29)$$

$$|w_s\rangle = \frac{1}{\sqrt{n}} \sum_{x \in \mathcal{W}} |x\rangle \quad (30)$$

$$|w_\perp\rangle = \sqrt{\frac{n_+}{\tilde{n} + n_+}} \sum_{x \in \tilde{\mathcal{W}}} |x\rangle - \sqrt{\frac{n}{\tilde{n} + n_+}} \sum_{x \in \mathcal{W}_+} |x\rangle \quad (31)$$

with $\mathcal{W}_+ = \tilde{\mathcal{L}} \cap \mathcal{W}$ and $n_+ = |\mathcal{W}_+|$. The states $|w_s\rangle$ and $|\ell_s\rangle$ are symmetric under permutations permuting only rewarded states with rewarded states and losing states with losing states similar to the symmetry properties of averaged strategies discussed in Appendix C. The state $|\psi(K)\rangle$ is given in this new basis by

$$|\psi(k)\rangle = \cos \varepsilon (\sin \phi |w_s\rangle + \cos \phi |\ell_s\rangle) + \sin \varepsilon |w_\perp\rangle \quad (32)$$

with the angle ϕ defined via

$$\tan \phi = \frac{\langle w_s | \psi(K) \rangle}{\langle \ell_s | \psi(K) \rangle} \quad (33)$$

$$= \tan \alpha \sqrt{\frac{\tilde{n}(n_+ + n_\ell)}{(\tilde{n} + n_+)n_\ell}} + \sqrt{\frac{n_+^2}{(\tilde{n} + n_+)n_\ell}} \quad (34)$$

saturating (93). The angle ε is given by

$$\sin \varepsilon = \sqrt{\frac{n_+}{n_+ + \tilde{n}}} \left[\sin(\alpha) - \sqrt{\frac{\tilde{n}}{n_+ + n_\ell}} \cos(\alpha) \right]. \quad (35)$$

The time evolution during the second phase, given by oracle O and U_k as given in Eq. 27, leads to a rotation of $|\psi(K)\rangle$ by an angle $2Jv$ in a plane parallel to the one spanned by $|\psi(0)\rangle$ and $|w\rangle$ as depicted in Fig. 1. As a consequence, the final state is given by

$$|\psi(K + J)\rangle = \cos \varepsilon [\sin(\phi + 2Jv)|w_s\rangle + \cos(\phi + 2Jv)|\ell_s\rangle] + (-1)^J \sin \varepsilon |w_\perp\rangle \quad (36)$$

leading to the maximal possible angle $\Delta = 2Jv$ and maximal $p_\perp = \sin^2 \varepsilon$ and thus to the maximal possible (average) success probability p_{K+J} .

As a result, performing consecutive Grover iterations in the first and second phase with in total $K + J$ oracle queries leads to the maximal possible average success probability p_{K+J} provided $\alpha = (2K + 1)\tilde{v} \leq \pi/2$ and $\phi(\alpha) + 2Jv \leq \pi/2$.

If more queries are available such that $(2K + 1)\tilde{v} > \pi/2$ or $\phi + 2Jv > \pi/2$, then it is possible to over-rotate the state $|\psi\rangle$ such that applying $\tilde{O}$ or O less often or performing

another algorithm like, e.g. fixed-point search (Yoder et al. 2014) leads to a higher success probability.

In general, the change of $|\psi(k)\rangle$ which can be created with a single oracle query O ($\tilde{O}$) is limited by $|\langle\psi(k+1)|\psi(k)\rangle| \geq \cos 2\nu$ ($\cos 2\tilde{\nu}$). The maximal possible difference between $|\psi(0)\rangle$ and $|\psi(K+J)\rangle$ achievable under this constrains would require that all states $|\psi(k)\rangle$ ly within a single plane (see discussion in Zalka (1999)). However, changing the oracle in Grover's algorithm leads to a change or tilt of the rotation plane/ axis as visualized in Fig. 1. Nevertheless, performing Grover iterations is the optimal strategy as we have proven. In addition, changing the oracle creates a component $|\phi_{\perp}\rangle$ which stays invariant under consecutive Grover iterations with the new oracle. Luckily, this component contains only rewarded items such that it does not prevent us from further increasing the success probability with Grover iterations if $\tilde{\mathcal{W}} \subseteq \mathcal{W}$. As a consequence, the optimality of Grover's algorithm in the case of a changing oracle might be not surprising but is also not obvious. Especially because performing Grover's algorithm with the maximal number of available oracle queries is not necessarily optimal if $\tilde{\mathcal{W}}$ and $\mathcal{W}$ only share a large overlap but $\tilde{\mathcal{W}} \not\subseteq \mathcal{W}$.

3.3 Grover iterations for $\tilde{\mathcal{W}} \not\subseteq \mathcal{W}$

In the following, we investigate the performance of Grover's algorithm if $\tilde{\mathcal{W}}$ and $\mathcal{W}$ share a large overlap (see Appendix A) but $\tilde{\mathcal{W}} \not\subseteq \mathcal{W}$. We will show that performing the maximal number K of oracle queries during the first phase is not always optimal depending on the number of available queries J in the second phase.

If $\tilde{\mathcal{W}} \not\subseteq \mathcal{W}$, then the perpendicular component $|\phi_{\perp}\rangle$, Eq. 10 also includes a losing component $|\ell_{\perp}\rangle$ such that the state $|\psi(K)\rangle$ can be written via

$$|\psi(K)\rangle = \cos \varepsilon |\phi_s\rangle + \sin \varepsilon |\phi_{\perp}\rangle \quad (37)$$

$$|\phi_s\rangle = \sin \phi |w_s\rangle + \cos \phi |\ell_s\rangle \quad (38)$$

$$|\phi_{\perp}\rangle = \sin \chi |w_{\perp}\rangle + \cos \chi |\ell_{\perp}\rangle. \quad (39)$$

Applying Grover iterations with unitaries U_k as defined in Eq. 27 does not change the success probability of the component $|\phi_{\perp}\rangle$. It only changed the success probability of the component $|\phi_s\rangle$ leading to

$$p_{K+J} = \cos^2 \varepsilon \sin^2(\phi + \Delta) + \sin^2 \varepsilon \sin^2 \chi \quad (40)$$

$$\leq 1 - \sin^2 \varepsilon \cos^2 \chi \quad (41)$$

with $\Delta = 2J\nu$. As a consequence, the success probability at the end of second phase is limited by $1 - |\langle\ell_{\perp}|\psi(K)\rangle|^2$ and thus by the weight of the orthogonal losing component created during the first phase. The contribution of this component is increasing with K for $K < \frac{\pi}{4\nu} - 1/2$ as shown in Fig. 2.

Fig. 2 Comparison of the success probabilities $p_{K'+J}$ for different numbers K' , J of Grover iterations during the first and second phase with $K' = 0$ (blue), $K' = 5$ (green) and $K' = 10$ (red) for $n_{\ell} = 5000$, $\tilde{n} = 15$, $n = 10$, $n_{+} = 5$ and thus $n_{-} = 10 = |\tilde{\mathcal{W}} \cap \mathcal{L}|$

In this case, the success probability p_{K+J} is still monotonically increasing with Δ . Therefore, performing the maximal possible number (J) of Grover iterations during the second phase is still a good idea provided $\phi + 2J\nu \leq \pi/2$. However, performing the maximal number (K) of Grover iterations during the first phase is not optimal if it leads to phases $\phi = \phi(K)$ and $\chi = \chi(K)$ such that

$$\sin^2 \chi(K) < \sin^2[2J\nu + \phi(K)]. \quad (42)$$

In this situation, performing less Grover iterations $K' < K$ during the first phase can lead to a higher final success probability $p_{K'+J} > p_{K+J}$ as shown in Fig. 2. Here, performing only $K' = 5$ instead of $K' = 10$ Grover iterations leads to higher final success probabilities for $6 \leq J \leq 11$. In general, it is optimal to perform the maximal number K of Grover iterations during the first phase if $J = 0$ (provided $(2K+1)\tilde{\nu} < \pi/2$). However, less and less effective queries to the first oracle $\tilde{O}$ should be used the more queries to the second oracle are available as demonstrated in Fig. 2.

3.4 Minimizing the cost to find a rewarded item

In the previous section, we have proven that it is possible to pursue a Grover-based search strategy and maintaining optimality even when the oracle is changing provided the rewarded space is monotonically increasing. We annotate that changing oracle settings have not been considered in literature before from the present perspective. Hence, it was unclear whether continuing Grover iterations in settings which changing oracles would lead to suboptimal results (compare Section 3.3). Note that the standard results for Grover search (Grover 1997; 1998; Brassard et al. 2000) can be only applied to the changing oracle problem if the two

phases are treated separately. That is, a Grover search is performed during the first phase. At the end of this phase, a measurement is performed. If no rewarded item is found, another Grover search is started with queries to the second oracle.

In the following, we discuss a numerical example demonstrating the possible improvements we can gain by proceeding with a coherent Grover-based strategy in a setting where the oracle changes rather than stopping it and starting a new one.

In our example, we assume that K Grover iterations $\tilde{G}_\psi$ with $\tilde{O}$ have been performed during a first phase. Then, the oracle is changed to O with $\tilde{W} \subseteq W$ and an arbitrary number of queries to the second oracle are allowed. We compare in our example the following two procedures: (a) a measurement is performed immediately after the change of the oracle and (b) the coherent time evolution is continued with the new oracle before performing a measurement. In both cases, we continue with a standard Grover search based solely on the second oracle O if the first measurement did not reveal a rewarded item.

In many scenarios, the goal is to minimize a given cost function, such as the number of oracle queries, rather than optimizing the reward probability. Therefore, we compare in this section the cost to find a rewarded item if we either proceed or interrupt the coherent search when the oracle changes. Here, we use $C(j) = 2j + 1$ as a typical cost function for a Grover algorithm with j steps (see e.g. Dunjko et al. (2016) and Sriarunothai et al. (2019)).

In the following, we neglect the cost produced by queries to the first oracle $\tilde{O}$ because this cost is identical for both cases and optimize the number of Grover iterations during the second phase (see Appendix F). In general, it is advantageous to continue the coherent time evolution if

$$4\nu \leq \pi - \arcsin\left(\frac{1}{1.38(1 - \sin^2 \varepsilon_K)} - 2\phi_K\right) \quad (43)$$

with $\sin^2 \nu = |W|/N = 1/N'$. Here, ϕ_K quantifies the reward probability of the symmetric part after the first phase (compare Eq. 35) and ε_K the ratio between the symmetric component and the orthogonal component (compare (33)). The cost functions for both procedures scale with $\sqrt{N'}$ just like in typical amplitude amplification algorithms. The cost difference between stopping or continuing Grover's algorithm also scales with $\sqrt{N'}$ for fixed angles ϕ_k and ε_k as we have evaluated in Appendix F and visualized in Fig. 3.

4 Conclusion

Research in quantum enhanced reinforcement learning has motivated quantum computation scenarios involving two

Fig. 3 Minimal expected average cost C of a continuous Grover algorithm with changing oracle with optimized j (red) for $\psi = \sqrt{|W|/N} = 1/\sqrt{N'}$ and fixed $\phi_K = 0.6$ and $\varepsilon_K \approx 0.32$. The expected average cost for an interrupted Grover algorithm with fixed oracles is denoted by a blue dashed line for comparison

systems, the agent and its environment, with restricted access to each other. In special cases, the interaction of the agent with its environment can be reduced to unitary oracle queries. However, general settings do not allow such a treatment due to memory effects induced by the environment.

In this paper, we generalized the basic case, where the environment acts effectively as a single fixed oracle, to settings where the oracle changes in time. This was motivated by standard grid-world type problems, where the number of consecutive actions within a single epoch can grow or shrink. We have demonstrated that the search for a rewarded action sequence of increasing length can be described as a search in a data based with fixed sequence length (equal to the maximal sequence length) but changing oracle leading to an increase of the rewarded space. We analyzed this setting and identified Grover-type amplitude amplification as optimal strategy for monotonically increasing rewarded spaces.

However, continuing coherent Grover iterations when the target space decreases will partially trap the resulting state within the losing subspace. As a consequence, the reward probability will be limited, with a limit clearly below unity, if we continue with Grover iterations after the oracle has changed.

It is easy to conceive a cascade of ever more general problems. For example, in slightly more general settings the agent might be allowed to choose if and when to change the effective oracle. In this way, the agent might combine breadth-first and depth-first search in a single coherent search for RL. Often, shorter rewarded action sequences are preferred but longer rewarded action sequences are more likely. Increasing the sequence length during a

coherent quantum search will amplify the probability for shorter rewarded sequences more than for longer sequences. Combining different oracles, corresponding to different sequence length, within a single Grover search might therefore help to balance the tradeoff between the desire for short rewarded sequences on the one side and high reward probabilities on the other. Further research investigating these possible benefits is needed. Quantum-enhanced reinforcement learning agent using different oracles might be developed for many different scenarios such as grid world, navigation or routing. In addition, many problems where we are searching for the shortest correct sequence of actions to achieve a goal, like optimal elements to create entangled photons (Melnikov et al. 2018), are equivalent to the grid-world problem.

We envision that the results presented here may become even more important in the future when large-scale quantum networks (Kimble 2008; Cacciapuoti et al. 2020) become available where the interactions are naturally quantum.

The goal in RL is in general to minimize a given cost function C instead of maximizing solely the success probability. In general, performing consecutive Grover iterations can be also used to minimize the average number of oracle queries necessary until a rewarded item is found. In general, we expect for our algorithm a quadratic improvement of the cost $C_g \propto \sqrt{C_{\text{class}}}$ compared to the cost C_{class} of a classical algorithm. In Section 3.4, we discussed a cost function depending solely on the number of oracle queries. We compared quantum algorithms based on a single oracle or multiple oracles leading to the costs C_s and C_g , respectively. Here, we found that the difference $C_s - C_g \propto C_s$ scales with the cost.

An optimal algorithm will depend on the exact cost function we want to minimize. For example, the search algorithm described in Boyer et al. (1998) is only optimal in terms of oracle queries. However, the number of elementary qubit gates necessary to perform a Grover search can be reduced by using a recursive Grover search (Arunachalam and de Wolf 2015) which separates the database into several subgroups. In RL, queries to different oracles might be connected to different cost. In such setting, an optimal algorithm might use different oracles in a recursive way for a quantum search. In this way, improvements in terms of cost might go beyond the quadratic improvements achievable in quantum exploration.

Finally, possibly the most interesting extensions would avoid reductions of environments to unitary oracles, and identify new schemes to obtain improvements in settings which may be more applicable in real-world RL settings. We leave these more general considerations for follow-up investigations.

Appendix A: Large overlap of $\tilde{\mathcal{W}}$ and $\mathcal{W}$

We say that the rewarded spaces $\tilde{\mathcal{W}}$ and $\mathcal{W}$ have a large overlap if increasing the probability $\tilde{p}$ for $x \in \tilde{\mathcal{W}}$ uniformly also increases the probability p to find $x \in \mathcal{W}$.

In general, optimal search strategies can be always constructed in such a way that the probability for all rewarded states $p(x|f(x) = 1)$ are equal as outlined in Appendix B. The same holds for losing states $|x\rangle$ with $f(x) = 0$. Let $n_a = |\tilde{\mathcal{W}} \cap \mathcal{W}|$ (n_ℓ) be the number of states which are marked as rewarded (not-rewarded/losing) by both oracles and $n_- = |\tilde{\mathcal{W}} \cap \mathcal{L}|$ (n_+) the number of states which are rewarded only according to the first (second) oracle. Thus, the total number of items is given by $N = n_a + n_\ell + n_- + n_+$. We denote the probabilities to find any state which is always rewarded, always not rewarded, rewarded only in the first phase and rewarded only during the second phase by p_a, p_ℓ, p_-, p_+ . Increasing the initial probability $\tilde{p} = p_a + p_- = (n_a + n_-)/N$ during the first phase in a symmetric way as outlined in Appendix B by a factor α leads to

$$p_a \rightarrow \alpha p_a, \quad p_- \rightarrow \alpha p_-, \quad p_\ell \rightarrow \beta p_\ell, \quad p_+ \rightarrow \beta p_+ \quad (44)$$

with

$$\beta = \frac{N - \alpha(n_a + n_-)}{n_\ell + n_+} \quad (45)$$

due to normalization. This leads to a change of p given by

$$p = p_a + p_+ \rightarrow \frac{n_+}{n_\ell + n_+} + \alpha \frac{n_a n_\ell - n_+ n_-}{n_\ell + n_+}. \quad (46)$$

As a result, we can increase p by increasing $\tilde{p}$ in a symmetric way whenever

$$n_a n_\ell > n_+ n_-. \quad (47)$$

As a result, we say $\tilde{\mathcal{W}}$ and $\mathcal{W}$ share a large overlap if they fulfill (47).

Appendix B: Averaged search strategies

In the following we consider search problems defined via some set of N orthonormal states $\{|n\rangle_A\}$ forming the basis of the Hilbert space $\mathcal{H}_A$ which can be separated into two subsets $\mathcal{H}_A = \mathcal{W} \cup \mathcal{L}$, the set of rewarded states $\mathcal{W}$ and the set of losing states $\mathcal{L}$ with $\mathcal{W} \cap \mathcal{L} = \emptyset$. Information about rewarded states can be obtained by querying phase-flip oracles

$$O_k = P_{\mathcal{W}_k} - P_{\mathcal{L}_k}. \quad (48)$$

where $P_{\mathcal{W}_k}$ and $P_{\mathcal{L}_k}$ denote projectors on some subspaces $\mathcal{W}_k$ and $\mathcal{L}_k$ forming together again the complete Hilbert space $\mathcal{H}_A = \mathcal{W}_k \cup \mathcal{L}_k$ with $\mathcal{W}_k \cap \mathcal{L}_k = \emptyset$. For standard

search problems we have $\mathcal{W}_k = \mathcal{W} \forall k$ and $\mathcal{L}_k = \mathcal{L} \forall k$. However, for more general search problems such as the here considered changing oracle problem, the subspaces $\mathcal{W}_k$ and $\mathcal{L}_k$ might differ from query to query.

Our goal is to find any state $|n\rangle_A \in \mathcal{W}$ with the help of maximally K oracle queries. All possible search strategies can be represented via unitary operations and pure initial states since it is possible to purify any search strategy by going to a larger Hilbert space $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$ and defining the generalize operators

$$O_{AB} = O_A \otimes \mathbb{1}_B, P_{\mathcal{W},AB} = P_{\mathcal{W},A} \otimes \mathbb{1}_B, P_{\mathcal{L},AB} = P_{\mathcal{L},A} \otimes \mathbb{1}_B. \quad (49)$$

To avoid a notation with over boarding indices, we skip the labels indicating the different subspace the operators/unitaries are working on if they are not crucial. Operators with a subspace index, such as e.g. σ_A acting on a state from a larger Hilbert space, e.g. $|\psi\rangle_{AB}$ are meant as short forms of the generalized operators defined similar to Eq. 49.

Any search strategy T to find a state $|n\rangle \in \mathcal{W}$ can be described via $T = (\{U_k\}, |\psi(0)\rangle_{AB})$ with a pure initial state $|\psi(0)\rangle_{AB}$ and unitaries $\{U_k\}$ acting on the combined Hilbert space $\mathcal{H}_{AB}$ leading after K oracle queries to the final state

$$|\psi(K)\rangle_{AB} = U_K O_K \cdots U_2 O_2 U_1 O_1 |\psi(0)\rangle_{AB} \quad (50)$$

and a consecutive projective measurement. Without loss of generality, we apply first an oracle query since any unitary U_0 applied before can be subsumed into the initial state. The probability p_T to identify a rewarded state correctly for a given strategy T and set of oracles $\{O_k\}$ is then given by

$$p_T = \text{Tr} (P_{\mathcal{W}} |\psi(K)\rangle \langle \psi(K)|). \quad (51)$$

Let σ denote a permutation operator acting on $\mathcal{H}_A$ and Σ_A denoting the group of operators of all possible permutations. The average reward probability $\bar{p}_T$ of the strategy T is defined via

$$\bar{p}_T = \frac{1}{N!} \sum_{\sigma \in \Sigma_A} p_T(\sigma) = \frac{1}{N!} \sum_{\sigma \in \Sigma_A} \text{Tr} [\sigma P_{\mathcal{W}} \sigma^\dagger |\psi(K, \sigma)\rangle \langle \psi(K, \sigma)|] \quad (52)$$

with

$$|\psi(K, \sigma)\rangle_{AB} = U_K \sigma O_K \sigma^\dagger \cdots U_1 \sigma O_1 \sigma^\dagger |\psi(0)\rangle_{AB} \quad (53)$$

being the resulting state if we substitute every oracle O_k by $\sigma O_k \sigma^\dagger$.

For every search strategy $T_{AB} = (\{U_{k,AB}\}, |\psi(0)\rangle_{AB})$, we can define an averaged strategy $\bar{T}_{ABC}$ via

Definition 1 The averaged strategy $\bar{T}_{ABC} = (\{\bar{U}_{k,ABC}\}, |\bar{\psi}(0)\rangle_{ABC})$ of the strategy $T_{AB} = (\{U_{k,AB}\}, |\psi(0)\rangle_{AB})$ is defined via the averaged initial state

$$|\bar{\psi}(0)\rangle_{ABC} = \frac{1}{\sqrt{N!}} \sum_{\sigma_\gamma \in \Sigma_A} \sigma_\gamma^\dagger |\psi(0)\rangle_{AB} |\gamma\rangle_C. \quad (54)$$

and average unitaries

$$\bar{U}_{k,ABC} = \sum_{\sigma_\gamma \in \Sigma_A} \sigma_\gamma^\dagger U_{k,AB} \sigma_\gamma \otimes |\gamma\rangle_C \langle \gamma|. \quad (55)$$

Here, the states $\{|\gamma\rangle_C\}$ are given by an arbitrary orthonormal basis of a Hilbert space $\mathcal{H}_C$ with dimension $d_C = N!$ acting as labels for the applied permutation operator σ_γ acting on $\mathcal{H}_A$.

The averaged strategy $\bar{T}$ has the following properties:

Lemma 3 The success probability $p_{\bar{T}}(\sigma)$ of the averaged strategy $\bar{T}$ is equal to the average success probability $\bar{p}_T$ of the strategy T for every permutation $\sigma \in \Sigma_A$.

Proof The success probability $p_{\bar{T}}(\sigma)$ is given by

$$p_{\bar{T}}(\sigma) = \text{Tr}_{ABC} [\sigma P_{\mathcal{W},ABC} \sigma^\dagger |\bar{\psi}(K, \sigma)\rangle \langle \bar{\psi}(K, \sigma)|]. \quad (56)$$

The state $\sigma^\dagger |\bar{\psi}(K, \sigma)\rangle_{ABC}$ for $\sigma \in \Sigma_A$ is given by

$$\sigma^\dagger |\bar{\psi}(K, \sigma)\rangle_{ABC} = \frac{1}{\sqrt{N!}} \sum_{\sigma_\gamma \in \Sigma_A} \sigma^\dagger \sigma_\gamma^\dagger U_{K,AB} \sigma_\gamma \sigma O_K \sigma^\dagger \cdots \sigma_\gamma^\dagger U_1 \sigma_\gamma \sigma O_1 \sigma^\dagger \sigma_\gamma^\dagger |\psi(0)\rangle_{AB} |\gamma\rangle_C \quad (57)$$

$$= \frac{1}{\sqrt{N!}} \sum_{\tilde{\sigma}_\gamma \in \Sigma_A} \tilde{\sigma}_\gamma^\dagger U_{K,AB} \tilde{\sigma}_\gamma O_K \cdots \tilde{\sigma}_\gamma^\dagger U_1 \tilde{\sigma}_\gamma O_1 \tilde{\sigma}_\gamma^\dagger |\psi(0)\rangle_{AB} |\gamma\rangle_C \quad (58)$$

where we used

$$\sigma \sum_{\sigma_\gamma \in \Sigma_A} \sigma_\gamma = \sum_{\tilde{\sigma}_\gamma \in \Sigma_A} \tilde{\sigma}_\gamma \quad \forall \sigma \in \Sigma_A \quad (59)$$

because Σ_A is a symmetric group. As a consequence, the application of the permutation $\sigma^\dagger$ on $|\bar{\psi}(K)\rangle$ is equivalent to a relabeling of the permutations σ_γ such that we now apply the permutation $\tilde{\sigma}_\gamma^\dagger = \sigma^\dagger \sigma_\gamma^\dagger$ instead of $\sigma_\gamma^\dagger$ if subsystem C is in state $|\gamma\rangle_C$. However, these labels have been arbitrary and therefore we find for the success probability

$$p_{\bar{T}}(\sigma) = \text{Tr}_{ABC} [P_{\mathcal{W},ABC} \sigma^\dagger |\bar{\psi}(K, \sigma)\rangle \langle \bar{\psi}(K, \sigma)| \sigma] \quad (60)$$

$$= \frac{1}{N!} \sum_{\tilde{\sigma} \in \Sigma_{\mathcal{H}_A}} \text{Tr}_{AB} [\tilde{\sigma} P_{\mathcal{W},AB} \tilde{\sigma}^\dagger |\psi(K)\rangle \langle \psi(K)| \tilde{\sigma}_{AB} \langle \psi(K), \tilde{\sigma}|_{AB}] \quad (61)$$

$$= \bar{p}_T \quad (62)$$

□

The relabeling can be formalized in the following way. We define the index $\tilde{\gamma}$ via $\sigma \sigma_\gamma = \sigma_{\tilde{\gamma}}$. Then, we can define the permutation $\pi(\sigma)$ acting on $\mathcal{H}_C$ via

$$\pi(\sigma) |\gamma\rangle_C = |\tilde{\gamma}\rangle_C \quad (63)$$

which then leads to the following lemma:

Lemma 4 The averaged strategy $\bar{T}$ is permutation invariant under joined permutations $\sigma \otimes \pi(\sigma) \forall \sigma \in \Sigma_A$ such

that

$$[\bar{U}_k, \sigma \otimes \pi(\sigma)] = 0 \quad (64)$$

$$\sigma \otimes \pi(\sigma)|\bar{\psi}(0)\rangle = |\bar{\psi}(0)\rangle. \quad (65)$$

Proof For the symmetric initial state $|\bar{\psi}(0)\rangle$, we find

$$\sigma \otimes \pi(\sigma)|\bar{\psi}(0)\rangle_{ABC} = \sigma \otimes \pi(\sigma) \frac{1}{\sqrt{N!}} \sum_{\sigma_\gamma \in \Sigma_A} \sigma_\gamma |\psi(0)\rangle_{AB} |\gamma\rangle_C \quad (66)$$

$$= \frac{1}{\sqrt{N!}} \sum_{\sigma_\gamma \in \Sigma_A} \sigma \sigma_\gamma |\psi(0)\rangle_{AB} \pi(\sigma) |\gamma\rangle_C \quad (67)$$

$$= \frac{1}{\sqrt{N!}} \sum_{\sigma_\gamma \in \Sigma_A} \sigma_\gamma |\psi(0)\rangle_{AB} |\gamma\rangle_C = |\bar{\psi}(0)\rangle_{ABC}. \quad (68)$$

For the symmetric unitaries $\bar{U}_k$, we find

$$\sigma \otimes \pi(\sigma) \bar{U}_k \sigma^\dagger \otimes \pi^\dagger(\sigma) = \sigma \otimes \pi(\sigma) \left(\sum_{\sigma_\gamma \in \Sigma_A} \sigma_\gamma U_{k,AB} \sigma_\gamma^\dagger \otimes |\gamma\rangle_C \langle \gamma| \right) \sigma^\dagger \otimes \pi^\dagger(\sigma) \quad (69)$$

$$= \sum_{\sigma_\gamma \in \Sigma_A} \sigma_\gamma U_{k,AB} \sigma_\gamma^\dagger \otimes |\gamma\rangle_C \langle \gamma| \quad (70)$$

$$= \bar{U}_k. \quad (71)$$

$[\bar{U}_k, \sigma \otimes \pi(\sigma)] = 0$ follows immediately since permutation operators are unitary. $\square$

As a consequence of Lemmas 3 and 4, we can limit the search for the best strategy T , optimizing $\bar{p}_T$, to averaged strategies $\bar{T}$ which also optimize the worst case probability $\min_{\sigma} p_T(\sigma)$ and leads to certain symmetries as outlined in Appendix C.

Appendix C: Symmetry investigations for the changing oracle problem

In the following, we consider a search problem, where the oracle O_k changes at a certain time step. Thus, we can separate the search into two phases. The first phase contains K oracle queries to oracle $\tilde{O} = O_k$ for $1 \leq k \leq K$ with rewarded space $\tilde{\mathcal{W}}$ and losing space $\tilde{\mathcal{L}}$. Then, the oracle changes to $O = O_k$ for $K < k \leq J$ with the new rewarded space $\mathcal{W}$ and losing space $\mathcal{L}$ and the search is continued by another J queries to O . In addition, we restrict the problem to monotonically increasing rewarded spaces that is the rewarded space $\tilde{\mathcal{W}}$ of the first phase is a subset $\tilde{\mathcal{W}} \subseteq \mathcal{W}$ of the rewarded space $\mathcal{W}$ of the second oracle O . This automatically leads to $\tilde{\mathcal{L}} \supseteq \mathcal{L}$.

In the following, we investigate the appearing symmetries occurring during the first and second phase when applying averaged search strategies $\bar{T}$ to this problem. Since we only consider averaged strategies and thus averaged unitaries $\bar{U}_k$ and states $|\bar{\psi}(k)\rangle$, we omit the bar on all states and unitaries in this section to simplify the notation.

In the following, we investigate the symmetry properties of the states

$$|\psi(K)\rangle = U_K O_K \cdots U_1 O_1 |\psi(0)\rangle \quad (72)$$

$$|\psi(K+J)\rangle = U_{K+J} O_{K+J} \cdots U_{K+1} O_{K+1} |\psi(K)\rangle \quad (73)$$

at the end of the first and the second phase. This will allow us to determine an upper bound for the average success probability p .

We define the set of permutations (operators) $\Sigma_{\tilde{O}} = \Sigma_{\tilde{\mathcal{W}}} \cup \Sigma_{\tilde{\mathcal{L}}}$ as the complete set of permutations operators which leave the rewarded space $\tilde{\mathcal{W}}$ and losing space $\tilde{\mathcal{L}}$ invariant. As a consequence, we find $[\tilde{O}, \sigma] = 0 \forall \sigma \in \Sigma_{\tilde{O}}$. The initial state $|\psi(0)\rangle$ and all unitaries U_k and $\tilde{O}_k$ during the first phase are permutation invariant under $\sigma \otimes \pi(\sigma) \forall \sigma \in \Sigma_{\tilde{O}}$ since $\Sigma_{\tilde{O}} \subseteq \Sigma_{\mathcal{H}_A}$. Thus, the state $|\psi(K)\rangle$ at the end of the first phase is also permutation invariant under $\sigma \otimes \pi(\sigma) \forall \sigma \in \Sigma_{\tilde{O}}$.

To determine the symmetry properties of $|\psi(K+J)\rangle$, we need to investigate how the rewarded and not-rewarded components of $|\psi(K)\rangle$ changes when we change the oracle. We define the normalized rewarded $|\tilde{w}\rangle$ and not-rewarded component $|\tilde{\ell}\rangle$ of $|\psi(K)\rangle$ via

$$\cos \alpha |\tilde{\ell}\rangle = P_{\tilde{\mathcal{L}}} |\psi(K)\rangle \quad (74)$$

$$\sin \alpha |\tilde{w}\rangle = P_{\tilde{\mathcal{W}}} |\psi(K)\rangle \quad (75)$$

with $\cos \alpha = |P_{\tilde{\mathcal{L}}} |\psi(K)\rangle|$. As a consequence, $|\psi(K)\rangle$ can be decomposed via

$$|\psi(K)\rangle = \cos \alpha |\tilde{\ell}\rangle + \sin \alpha |\tilde{w}\rangle. \quad (76)$$

The components $|\tilde{w}\rangle$ and $|\tilde{\ell}\rangle$ are permutation invariant under $\sigma \otimes \pi(\sigma) \forall \sigma \in \Sigma_{\tilde{O}}$ because the projectors $P_{\tilde{\mathcal{W}}}$ and $P_{\tilde{\mathcal{L}}}$ as well as $|\psi(K)\rangle$ are permutation invariant.

Let us now investigate the rewarded and not-rewarded components at the beginning of the second phase. The initial state of the second phase is given by $|\psi(K)\rangle$. Its component $|\tilde{w}\rangle$ is also a rewarded component according to the second oracle O such that $P_{\mathcal{W}} |\tilde{w}\rangle = |\tilde{w}\rangle$. However, $|\tilde{\ell}\rangle$ contains both rewarded and not-rewarded components

$$\cos \beta |\ell\rangle = P_{\mathcal{L}} |\tilde{\ell}\rangle \quad (77)$$

$$\sin \beta |w_+\rangle = P_{\mathcal{W}} |\tilde{\ell}\rangle \quad (78)$$

with $\cos \beta = |P_{\mathcal{L}} |\tilde{\ell}\rangle|$. Note, that $|w_+\rangle \in \mathcal{W}_+ = \tilde{\mathcal{L}} \cap \mathcal{W}$ and thus $|w_+\rangle \perp |\tilde{w}\rangle$. Therefore, we can divide the state $|\psi(K)\rangle$ into three orthogonal components via

$$|\psi(K)\rangle = \sin \alpha |\tilde{w}\rangle + \cos \alpha (\sin \beta |w_+\rangle + \cos \beta |\ell\rangle). \quad (79)$$

The angle β is given by

$$\sin \beta = \sqrt{\frac{n_+}{n_+ + n_\ell}} \quad (80)$$

where n_+ denotes the dimension of $\mathcal{W}_+$ and n_ℓ the dimension of $\mathcal{L}$ (see Appendix D).

Let us now invest the symmetries of $|\psi(K)\rangle$ with respect to permutations $\sigma \otimes \pi(\sigma) \forall \sigma \in \Sigma_O$ which leave the second oracle O invariant. Let P_S be the projector onto the symmetric subspace which can be written as

$$P_S = \sum_{|s\rangle} |s\rangle\langle s| \quad \text{with} \quad \sigma \otimes \pi(\sigma)|s\rangle = |s\rangle \quad \forall \sigma \in \Sigma_O \quad (81)$$

where $\{|s\rangle\}$ forms an orthonormal basis of the symmetric subspace. Then, we can define the symmetric component

$$\cos \varepsilon |\phi_s\rangle = P_S |\psi(K)\rangle \quad (82)$$

and its complement

$$\sin \varepsilon |\phi_\perp\rangle = (\mathbb{1} - P_S) |\psi(K)\rangle \quad (83)$$

with $\cos \varepsilon = |P_S |\psi(K)\rangle|$. The state $|\ell\rangle$ is permutation invariant under $\sigma \otimes \pi(\sigma) \forall \sigma_A \in \Sigma_O$ since $\mathcal{L} \subseteq \tilde{\mathcal{L}}$ such that $P_S |\ell\rangle = |\ell\rangle$. However, the (not normalized) rewarded component $\sin \alpha |\tilde{w}\rangle + \cos \alpha \sin \beta |w_+\rangle$ is not necessarily permutation invariant under $\sigma \otimes \pi(\sigma) \forall \sigma \in \Sigma_O$. As a consequence, there might exist a non-vanishing component $|\phi_\perp\rangle$; however, this component lies within the rewarded space $\mathcal{W}$ such that

$$\sin \varepsilon |\phi_\perp\rangle = (\mathbb{1} - P_S) |\psi(K)\rangle = P_{\mathcal{W}} (\mathbb{1} - P_S) |\psi(K)\rangle = \sin \varepsilon |w_\perp\rangle \quad (84)$$

The symmetric component $|\phi_s\rangle$ can be decomposed into a rewarded and a not-rewarded component

$$\cos \varepsilon \sin \phi |w_s\rangle = P_{\mathcal{W}} P_S |\psi(K)\rangle \quad (85)$$

$$\cos \varepsilon \cos \phi |\ell_s\rangle = P_{\mathcal{L}} P_S |\psi(K)\rangle = \cos \varepsilon \cos \phi |\ell\rangle \quad (86)$$

with $\cos \varepsilon \sin \phi = |P_{\mathcal{W}} P_S |\psi(K)\rangle|$. Thus, the state $|\psi(K)\rangle$ can be separated into the following three orthogonal components

$$|\psi(K)\rangle = \cos \varepsilon (\sin \phi |w_s\rangle + \cos \phi |\ell\rangle) + \sin \varepsilon |w_\perp\rangle. \quad (87)$$

A comparison with Eq. 79 leads to the following identities

$$\cos \varepsilon \cos \phi = \langle \ell | \psi(K) \rangle = \cos \alpha \cos \beta \quad (88)$$

$$\cos \varepsilon \sin \phi = \langle w_s | \psi(K) \rangle = \sin \alpha \langle w_s | \tilde{w} \rangle + \cos \alpha \sin \beta \langle w_s | w_+ \rangle \quad (89)$$

$$\sin \varepsilon = \langle w_\perp | \psi(K) \rangle = \sin \alpha \langle w_\perp | \tilde{w} \rangle + \cos \alpha \sin \beta \langle w_\perp | w_+ \rangle. \quad (90)$$

Note, all appearing scalar products are real due to the definition of $|w_s\rangle$ and $|w_\perp\rangle$ and they are upper bounded via

$$|\langle w_s | \tilde{w} \rangle| \leq \sqrt{\frac{\tilde{n}}{\tilde{n} + n_+}} \quad (91)$$

$$|\langle w_s | w_+ \rangle| \leq \sqrt{\frac{n_+}{n_\ell + n_+}}. \quad (92)$$

As a consequence, the angle ϕ is upper bounded by the angle α via

$$\tan \phi \leq \tan \alpha \sqrt{\frac{\tilde{n}(n_+ + n_\ell)}{(\tilde{n} + n_+)n_\ell}} + \sqrt{\frac{n_+^2}{(\tilde{n} + n_+)n_\ell}}. \quad (93)$$

Let us investigate the time evolution during the second phase. We denote with

$$V = U_{K+J} O \cdots U_{K+1} O \quad (94)$$

a unitary which described the complete time evolution during the second phase. The unitary V commutes with the projector P_S as the following considerations will prove. There exist a joined eigenbasis of V and $\sigma \otimes \pi(\sigma) \forall \sigma \in \Sigma_O$ since $[V, \sigma \otimes \pi(\sigma)] = 0$. Let $\{|v_x\rangle\}$ be an eigenbasis of V and wlog we assume that the first f states of this basis form the symmetric subspace such that

$$\sigma \otimes \pi(\sigma) |v_x\rangle = |v_x\rangle \quad \forall \sigma \in \Sigma_O \text{ and } 1 \leq x \leq f. \quad (95)$$

As a consequence, we find

$$P_S V = \sum_{x=1}^f |v_x\rangle\langle v_x| \sum_y \lambda_y |v_y\rangle\langle v_y| = \sum_{x=1}^f \lambda_x |v_x\rangle\langle v_x| = V P_S \quad (96)$$

where λ_y denote the eigenvalues of V . Thus, the time evolution of the symmetric component $V|\phi_s\rangle$ stays a symmetric state with

$$P_S V |\phi_s\rangle = V P_S |\phi_s\rangle = V |\phi_s\rangle \quad (97)$$

whereas $V|\phi_\perp\rangle$ stays orthogonal to this subspace since

$$P_S V |\phi_\perp\rangle = V P_S |\phi_\perp\rangle = 0 \quad (98)$$

and thus the symmetric part and the orthogonal part do not mix.

The reward probability of $|\psi(K+J)\rangle$ can be decomposed into a symmetric part and a part orthogonal to it via

$$\begin{aligned} \text{Tr} [P_{\mathcal{W}} |\psi(K+J)\rangle\langle\psi(K+J)|] &= \text{Tr} [P_{\mathcal{W}} P_S |\psi(K+J)\rangle\langle\psi(K+J)|] \\ &\quad + \text{Tr} [P_{\mathcal{W}} (\mathbb{1} - P_S) |\psi(K+J)\rangle\langle\psi(K+J)|] \\ &= \cos^2 \varepsilon \text{Tr} [P_{\mathcal{W}} V |\phi_s\rangle\langle\phi_s| V^\dagger] \\ &\quad + \sin^2 \varepsilon \text{Tr} [P_{\mathcal{W}} V |\phi_\perp\rangle\langle\phi_\perp| V^\dagger] \end{aligned} \quad (99)$$

where we used $[P_{\mathcal{W}}, P_S] = 0$ which follows directly from $[P_{\mathcal{W}}, \sigma \otimes \pi(\sigma)] = 0 \forall \sigma \in \Sigma_O$ and $P_S = P_S^2$.

Appendix D: Determining the angle β

In the following, we give a more detailed derivation of Eq. 80 for determining β defined via

$$\sin^2 \beta = \langle \tilde{\ell} | P_{\mathcal{W}} | \tilde{\ell} \rangle. \quad (101)$$

Let wlog $\{|j\rangle_A\}$ with $1 \leq j \leq n_+ + n_\ell$ be a basis of the losing space $\tilde{\mathcal{L}}_A$. The state $|\tilde{\ell}\rangle_{ABC}$ can then be written as

$$|\tilde{\ell}\rangle = \sum_{j=1}^{n_++n_\ell} \xi_j |j\rangle_A |\gamma_j\rangle_{BC} \quad (102)$$

with some arbitrary normalized states $|\gamma_j\rangle_{BC}$. The probability for each state $|j\rangle_A$ is given by

$$p_j = || |j\rangle_A \langle j| \otimes \mathbb{1}_{BC} |\tilde{\ell}\rangle ||^2 = |\xi_j|^2. \quad (103)$$

However, the state $|\tilde{\ell}\rangle$ is permutation invariant $\forall \sigma \in \Sigma_{\tilde{\ell}}$ such that

$$\sigma \otimes \pi(\sigma) |\tilde{\ell}\rangle = \sum_{j=1}^{n_+ + n_\ell} \xi_j |j'(\sigma)\rangle_A |\gamma'_{j'}\rangle_{BC} = |\tilde{\ell}\rangle \quad (104)$$

with $|\gamma'_{j'}\rangle_{BC} = \mathbb{1}_B \otimes \pi_C |\gamma_j\rangle_{BC}$. As a consequence, we find for the probabilities

$$p_j = || (|j\rangle \langle j| \otimes \mathbb{1}_{BC}) (\sigma_A \otimes \pi_C(\sigma)) |\tilde{\ell}\rangle ||^2 = |\xi_{j'}|^2 \quad \forall \quad 1 \leq j, j' \leq n_+ + n_\ell \quad (105)$$

and due to normalization $p_j = 1/(n_+ + n_\ell)$. Since there exist n_+ orthonormal states within the subspace $\mathcal{W}_+ = \mathcal{W} \cap \tilde{\mathcal{L}}$ we find

$$\sin^2 \beta = \frac{n_+}{n_+ + n_\ell}. \quad (106)$$

Appendix E: Optimality proof of Grover's algorithm for multiple rewarded items

The optimality proof of Grover's algorithm for oracles with a single rewarded item by Zalka (1999) consist of two parts given by the inequality

$$2N - 2\sqrt{Np} - 2\sqrt{N(N-1)(1-p)} \leq \sum_{y=1}^N |||\phi_J\rangle - |\phi_J^y\rangle||^2 \leq 4N \sin^2(J\psi). \quad (107)$$

Here, N is the number of items, p the success probability to identify the single rewarded item y correctly, J the maximal number of oracle queries and the angle ψ is defined via $\sin^2 \psi = 1/N$. The two quantum states $|\phi_j\rangle$ and $|\phi_j^y\rangle$ are defined via

$$|\phi_j\rangle = V^j |\phi\rangle \quad (108)$$

$$|\phi_j^y\rangle = V_y^j |\phi\rangle \quad (109)$$

where $|\phi\rangle$ is some arbitrary state, V_y^j a unitary of the form Eq. 94 based on j queries to the oracle O_y and V^j is a unitary based on j queries to an empty oracle. The optimality of Grover's algorithm follows from the proof of both inequalities and the fact that Grover's algorithm saturates both.

We generalize the results from Zalka by going to oracles O_y which mark exactly n items out of N items as rewarded. In this case, y is now a label for the rewarded space $\mathcal{W}_y$ and there exist now $D = \binom{N}{n}$ different oracles. The success probability p now denotes the probability to identify any rewarded item $|z\rangle \in \mathcal{W}_y$ correctly. For a random guess, this

probability is given by $\sin^2 v = n/N$. As a consequence, Eq. 107 can be generalized to

$$2D - 2D\sqrt{p\frac{n}{N}} - 2D\sqrt{(1-p)\left(1 - \frac{n}{N}\right)} \leq \sum_{y=1}^D |||\phi_J\rangle - |\phi_J^y\rangle||^2 \leq 4D \sin^2(Jv) \quad (110)$$

which we will proof in the following and is equal to Eq. 107 for $n = 1$. Again, Grover's algorithm saturates these bounds.

We start with the right inequality and proof the following lemma

Lemma 5 *The maximal difference between $|\phi_J\rangle$ and $|\phi_J^y\rangle$ achievable with J oracle queries averaged over all possible oracles with n rewarded items is given by*

$$\frac{1}{D} \sum_{y=1}^D |||\phi_J\rangle - |\phi_J^y\rangle||^2 \leq 4 \sin^2(J\psi) = 2[1 - \cos(2J\psi)] \quad (111)$$

with $\sin^2 v = n/N$.

Proof This Lemma follows directly from the optimality proof of Grover's algorithm given in Zalka (1999) by generalizing the sum overall possible oracles which mark only one item y to all possible oracles which mark n items. In the following, we do not reproduce every step from Ref. (Zalka 1999) but concentrate only on steps where the generalization from one rewarded item to several rewarded items makes a difference. Following Ref. Zalka (1999) we find (Eq. 22)

$$\frac{1}{D} \sum_{y=1}^D |||\phi_J\rangle - |\phi_J^y\rangle||^2 \leq Df(x) \quad (112)$$

with the argument

$$x = \frac{4J}{D} \sum_{y=1}^D \sum_{j=1}^J ||P_{\mathcal{W}_y} |\phi_j\rangle||^2 \quad (113)$$

and $P_{\mathcal{W}_y}$ the projector onto the rewarded space of oracle y . The function $f(x)$ is defined in Zalka (1999) via

$$f(x = 4J^2 \sin^2 v) = 4 \sin^2(Jv). \quad (114)$$

Every state $|z\rangle \in \mathcal{H}_A$ is part of the rewarded space $\mathcal{W}_y$ for exactly $d = \binom{N-1}{n-1}$ different oracles. As a consequence, the argument x of the function f in Eq. 112 is given by

$$x = \frac{4J}{D} \sum_{y=1}^D \sum_{j=1}^J ||P_{\mathcal{W}_y} |\phi_j\rangle||^2 = \frac{4J}{D} \sum_{j=1}^J \sum_{y=1}^D \sum_{z \in \mathcal{W}_y} ||\langle z | \phi_j \rangle||^2 \quad (115)$$

$$= 4J \sum_{j=1}^J \frac{d}{D} \sum_{z=1}^N ||\langle z | \phi_j \rangle||^2. \quad (116)$$

The sum over all states $|z\rangle$ sums up to unity leading to

$$\frac{4J}{D} \sum_{y=1}^D \sum_{j=1}^J ||P_{\mathcal{W}_y}|\phi_j\rangle||^2 = 4J^2 \frac{n}{N} = 4J^2 \sin^2 v \quad (117)$$

where we used $d/D = n/N$. $\square$

Grover's algorithm saturates this inequality since we find for this algorithm

$$|\phi\rangle = \frac{1}{\sqrt{N}} \sum_{z=1}^N |z\rangle = \cos v |\ell\rangle + \sin v |w\rangle \quad (118)$$

$$|\phi_j^y\rangle = \cos[(2J+1)v]|\ell\rangle + \sin[(2J+1)v]|w\rangle \quad (119)$$

$$|\phi_J\rangle = \cos v |\ell\rangle + \sin v |w\rangle \quad (120)$$

leading to

$$\begin{aligned} \frac{1}{D} \sum_{y=1}^D |||\phi_J\rangle - |\phi_j^y\rangle||^2 &= 2 - 2 \sin[(2J+1)v] \sin v - 2 \cos[(2J+1)v] \cos v \quad (121) \\ &= 2[1 - \cos(2Jv)]. \end{aligned} \quad (122)$$

The right side of Eq. 110 is govern by the lemma

Lemma 6 *The average success probability p to identify any item $z \in \mathcal{W}_y$ out of the rewarded space $\mathcal{W}_y$ of the oracle O_y with $1 \leq y \leq D$ given the states ϕ_j^y average over all oracles is upper bounded by*

$$2 - 2\sqrt{p \frac{n}{N}} - 2\sqrt{(1-p) \left(1 - \frac{n}{N}\right)} \leq \frac{1}{D} \sum_{y=1}^D |||\phi_J\rangle - |\phi_j^y\rangle||^2 \quad (123)$$

Proof Again, in order to proof this lemma, we follow the proof in Zalka (1999) and only point out the generalizations we have to make when going from $n = 1$ rewarded state to $n > 1$ rewarded states. Similar to (Zalka 1999), we write the states

$$|\phi_j^y\rangle = \sum_{x=1}^X c_x^y |x\rangle \quad |\phi_J\rangle = \sum_{x=1}^X c_x |x\rangle \quad (124)$$

via some orthonormal basis $\{|x\rangle\}$ of some Hilbert space with dimension X . The optimal procedure to identify a rewarded item $|z\rangle$ is to perform projective measurements (see Ref. Zalka (1999)). Let $\{|x\rangle\}$ be the measurement basis and we denote with X_z the subspace containing all states $|x\rangle$ which correctly denote that $|z\rangle$ is a rewarded item. As a consequence, the success probability p_y if the unknown oracle is given by O_y is determined via

$$p_y = \sum_{z \in \mathcal{W}_y} \sum_{x \in X_z} |c_x^y|^2. \quad (125)$$

Similar, we can define a success probability a_y for the state $|\phi_J\rangle$ via (compare Eq. A7 in Zalka (1999))

$$a_y = \sum_{z \in \mathcal{W}_y} \sum_{x \in X_z} |c_x|^2. \quad (126)$$

In order to proof (6), Zalka determines the minimal distance an arbitrary state $|\phi_y\rangle$ with success probability p_y needs to have from a given state $|\zeta_y\rangle$ with success probability a_y . This minimal distance is given (compare Eq. A8 in Zalka (1999)) by

$$|||\phi_y\rangle - |\zeta_y\rangle||^2 \geq 2 - 2 \left(\sqrt{p_y a_y} + \sqrt{(1-p_y)(1-a_y)} \right). \quad (127)$$

The minimum of

$$\frac{1}{D} \sum_{y=1}^D |||\phi_y\rangle - |\zeta_y\rangle||^2 \quad (128)$$

for all possibly states $|\zeta_y\rangle$ and success probabilities p_y is reached if all $p_y = p$ and $a_y = a$ (see Zalka (1999)). Due to normalization we find

$$\sum_{y=1}^D a_y = d \sum_{x=1}^X |c_x| = d \quad (129)$$

where we have used that each item $|z\rangle$ belongs to the rewarded space of $d = \binom{N-1}{n-1}$ different oracles. As a consequence, the minimum is achieved for $a_y = d/D = n/N$ (see discussion before Eq. A10 in Zalka (1999)) leading finally to the modification of Eq. A10 Zalka (1999) to

$$\frac{1}{D} \sum_{y=1}^D |||\phi_J\rangle - |\phi_j^y\rangle||^2 \geq 2 - 2\sqrt{p \frac{n}{N}} - 2\sqrt{(1-p) \left(1 - \frac{n}{N}\right)} \quad (130)$$

which gives us directly Lemma 6. Also this bound is satisfied by Grover's algorithm. $\square$

The above-stated optimality proof of Grover's algorithm can be easily generalized to situation where we start in a state $|\zeta_y\rangle$ with success probability $a_y = a = \sin^2 \phi$ and try to optimize the success probability p_y of $V_y^J |\zeta_y\rangle$ with the help of maximal J oracle queries. Lemma 5 is independent from the initial state and can therefore directly be applied. From Eq. 127 we find

$$\frac{1}{D} \sum_{y=1}^D ||V_y^J |\zeta_y\rangle - |\zeta_y\rangle||^2 \geq 2 - 2 \sum_y \sqrt{p_y \sin^2 \phi} \sqrt{(1-p_y) \cos^2 \phi} \quad (131)$$

which is minimal if $p_y = p \forall y$. Thus, we find

$$\frac{1}{D} \sum_{y=1}^D ||V_y^J |\zeta_y\rangle - |\zeta_y\rangle||^2 \geq 2 - 2\sqrt{p \sin^2 \phi} \sqrt{(1-p) \cos^2 \phi}. \quad (132)$$

Lemma 5 and Eq. 132 can be simultaneously saturated by starting in a state

$$|\zeta_s\rangle = \sin \phi \frac{1}{\sqrt{|\mathcal{W}_y|}} \sum_{|z\rangle \in \mathcal{W}_y} |z\rangle + \cos \phi \frac{1}{\sqrt{|\mathcal{L}_y|}} \sum_{|z\rangle \in \mathcal{L}_y} |z\rangle \quad (133)$$

and performing Grover iterations via the unitary

$$V_y^J = [(\mathbb{I} - |\phi\rangle\langle\phi|)O_y]^J \quad (134)$$

$$|\phi\rangle = \frac{1}{\sqrt{N}} \sum_{z=1}^N |z\rangle. \quad (135)$$

Applying V_y^J with an empty oracle on $|\zeta_s\rangle$ does not change the success probability a_y leading to a maximal success probability $p = \sin^2(\phi + \nu)$ with $\sin^2 \nu = n/N$.

Appendix F: Minimal cost

In the following, we consider a data base search where the oracle $\tilde{O}$ is exchange at time t_0 by a new oracle O with rewarded space $W \supseteq \tilde{W}$. Before t_0 , several Grover iterations $\tilde{G}_\psi$ have been applied and created the state $|\psi(K)\rangle$ as given in Eq. 32. Then, the oracle changes. The probability to project $|\psi(K)\rangle$ onto a target state directly after t_0 is given by

$$p_K = \|P_{\mathcal{W}}|\psi(K)\rangle\|^2 = \sin^2 \varepsilon + \cos^2 \varepsilon \sin^2 \phi_K. \quad (136)$$

In the following, we want to minimize the cost to find a target state. A typical cost for Grover's algorithm with j steps is given by $C(j) = 2j + 1$ (see e.g. Dunjko et al. (2016) and Sriarunothai et al. (2019)).

The standard Grover algorithm considers only static oracles. Hence, the coherent evolution has to be terminated when the oracle changes. Thus in the standard case, we perform a measurement to finish the coherent time evolution based on $\tilde{G}_\psi$. If no target state has been found, we continue by performing Grover iterations with the new oracle described by G_ψ . The average cost C_s for the standard Grover algorithm is then given by

$$C_s = 1 + (1 - p_K) \frac{2j + 1}{\sin^2[(2j + 1)\nu]}, \quad (137)$$

for a fixed number of j steps and $\sin^2 \nu = |\mathcal{W}|/N$. Here, we neglected all costs accumulated before t_0 . The minimum cost is achieved for $\tan[(2j + 1)\nu] \approx 2(2j + 1)\nu$ which can be solved numerically. The numerical solution corresponds to apply Grover's algorithm until a probability to find a target is approximately $\sin^2[(2j + 1)\nu] \approx 0.844$ and leads

to the average cost of

$$C_s^{\text{opti}} \approx 1 + (1 - p_K) \frac{1.38}{\nu}. \quad (138)$$

We generalize Grover's algorithm by continuing the coherent time evolution after the oracle has changed. The cost for continuing the Grover search with j applications of G_ψ is given by $(2j + 1)$. A consecutive measurement will project the resulting state $|\psi(K + j)\rangle$ onto a target state with probability (see Eq. 36)

$$p_{K+j} = \sin^2 \varepsilon_K + \cos^2 \varepsilon_K \sin^2(2j\nu + \phi_K). \quad (139)$$

If we do not find a target, we continue with a standard Grover search based on the new oracle leading to the same expected average cost of $1.38/\nu$. Thus, we expect in total an average cost of

$$C_g \approx (2j + 1) + (1 - p_{K+j}) \frac{1.38}{\nu} \quad (140)$$

if we continue the coherent time evolution by another j Grover iterations G_ψ although the oracle has changed before we perform a measurement. By migrating j into a real valued variable, we can optimize C_g over j . Extremal points of C_g can be found for values of j satisfying

$$\sin[4j\nu + 2\phi_K] = \frac{1}{1.38 \cos^2 \varepsilon_K}. \quad (141)$$

For $\cos^2 \varepsilon_K > 1/1.38$, we find a maximum of C_g at

$$j \approx \frac{1}{4\nu} \left[\arcsin\left(\frac{1}{1.38 \cos^2 \varepsilon_K}\right) - 2\phi_K \right] \quad (142)$$

and a minimum at

$$j \approx \frac{1}{4\nu} \left[\pi - \arcsin\left(\frac{1}{1.38 \cos^2 \varepsilon_K}\right) - 2\phi_K \right]. \quad (143)$$

This minimum corresponds to a probability to find a target state at the first measurement given by

$$p_{\text{opti}}(\varepsilon_K) = \frac{1}{2} \left(1 + \sin^2 \varepsilon_K - \sqrt{\cos^4 \varepsilon_K - \frac{1}{1.38^2}} \right) \quad (144)$$

which is independent of ν . As a consequence, the average cost of a continuous Grover algorithm optimized over j can be written as

$$C_g^{\text{opti}} = 1 + \frac{f(\varepsilon_K, \phi_K)}{\nu} \quad (145)$$

with the function

$$f = \frac{1}{2} \left[\pi - \arcsin\left(\frac{1}{1.38 \cos^2 \varepsilon_K}\right) - 2\phi_K \right] + 1.38 \cdot (1 - p_{\text{opti}}(\varepsilon_K)). \quad (146)$$

Finally, we can determine the difference between the optimal cost using the continuous or interrupted Grover algorithm by

$$C_s^{\text{opti}} - C_g^{\text{opti}} = \frac{1.38(1 - p_K) - f(\varepsilon_K, \phi_K)}{\nu} \quad (147)$$

Thus, there exist j such that $C_g(j) \leq C_s$ if $1.38 \cos^2 \varepsilon_K > 1$. For $1.38 \cos^2 \varepsilon_K < 1$, we achieve a minimal cost at $j = 0$ such that we do not obtain any advantage continuing the Grover algorithm. In general, the maximal possible different scales as $C_s(\psi) - C_g(\psi) \sim 1/\nu \sim \sqrt{N'}$ if we assume that the classical probability to find a target state with the second oracle is given by $\sin^2 \nu = |W|/N = 1/N'$ (see Fig. 3).

Acknowledgements The authors thanks H.J. Briegel, F. Flamini, S. Jerbi, D. Orsucci and L. Trenkwalder for fruitful discussions.

Funding Open Access funding enabled and organized by Projekt DEAL. SW acknowledges support from the Austrian Science Fund (FWF) through the SFB BeyondC. AH acknowledges support from the Austrian Science Fund (FWF) through the project P 30937-N27. This work was also supported by the Dutch Research Council (NWO/OCW), as part of the Quantum Software Consortium programme (project number 024.003.037).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Aïmeur E, Brassard G, Gambs S (2013) Quantum speed-up for unsupervised learning. *Mach Learn* 90:261. <https://doi.org/10.1007/s10994-012-5316-5>
- Ambainis A (2002) Quantum lower bounds by quantum arguments. *J Comp Syst Sci* 64:750
- Ambainis A (2006) Polynomial degree vs. quantum query complexity. *J Comp Syst Sci* 72:220
- Arunachalam S, Briët J., Palazuelos C (2019) Quantum query algorithms are completely bounded forms. *SIAM J Comp* 48:903
- Arunachalam S, de Wolf R (2015) Optimizing the number of gates in quantum search. *Quantum Inform Comput* 17
- Biamonte J, Wittek P, Pancotti N, Rebentrost P, Wiebe N, Lloyd S (2017) Quantum machine learning. *Nature* 549:195
- Boyer M, Brassard G, Hoyer P, Tappa A (1998) Tight bounds on quantum searching. *Fortschr Phys* 46:493. <https://doi.org/10.1002/3527603093.ch10>
- Brassard G, Hoyer PF, Mosca M, de Montreuil ATDU, Aarhus BU, Waterloo CU (2000) Quantum amplitude amplification and estimation. *arXiv:quant-ph/0005055*
- Briegel HJ, De las Cuevas G (2012) Projective simulation for artificial intelligence. *Sci Rep* 2:400. <https://doi.org/10.1038/srep00400>
- Cacciapuoti AS, Caleffi M, Tafuri F, Cataliotti FS, Gherardini S, Bianchi G (2020) Quantum internet: networking challenges in distributed quantum computing. *IEEE Netw* 34(1):137. <https://doi.org/10.1109/MNET.001.1900092>
- Chia NH, Gilyén A, Li T, Lin HH, Tang E, Wang C (2019) Sampling-based sublinear low-rank matrix arithmetic framework for dequantizing quantum machine learning. *arXiv:1910.06151*
- Cornelissen A (2018) Quantum gradient estimation and its application to quantum reinforcement learning. Master's thesis, Delft University of Technology
- da Silva BC, Basso EW, Bazzan PM, Engel ALC (2006) Dealing with non-stationary environments using context detection. In: *Proceedings of the 23rd international conference on machine learning, ICML, vol 2006*, p 217
- Dunjko V, Taylor JM, Briegel HJ (2016) Quantum-enhanced machine learning. *Phys Rev Lett* 117:130501. <https://doi.org/10.1103/PhysRevLett.117.130501>
- Dunjko V, Liu YK, Wu X, Taylor JM (2017) Super-polynomial and exponential improvements for quantum-enhanced reinforcement learning. *arXiv:1710.11160*
- Dunjko V, Briegel H (2018) Machine learning & artificial intelligence in the quantum domain: a review of recent progress. *Rep Prog Phys* 81:074001. <https://doi.org/10.1088/1361-6633/aab406>
- Farhi E, Neven H (2018) Classification with quantum neural networks on near term processors. <https://arxiv.org/abs/1802.06002v2>. *arXiv:1802.06002*
- Grover LK (1997) Quantum mechanics helps in searching for a needle in a haystack. *Phys Rev Lett* 79:325. <https://doi.org/10.1103/PhysRevLett.79.325>
- Grover LK (1998) Quantum computers can search rapidly by using almost any transformation. *Phys Rev Lett* 80:4329. <https://doi.org/10.1103/PhysRevLett.80.4329>
- Gyurik C, Cade C, Dunjko V (2020) Towards quantum advantage for topological data analysis. *arXiv:2005.02607*
- Han M (2018) Reinforcement learning approaches in dynamic environments. *Databases [cs.DB].Télécom ParisTech*. English. tel-01891805
- Harrow AW, Hassidim A, Lloyd S (2009) Quantum algorithm for linear systems of equations. *Phys Rev Lett* 103:150502. <https://doi.org/10.1103/PhysRevLett.103.150502>
- Havlicek V, Córcoles AD, Temme K, Harrow AW, Kandala A, Chow JM, Gambetta JM (2019) Supervised learning with quantum-enhanced feature spaces. *Nature* 567(7747):209. <https://doi.org/10.1038/s41586-019-0980-2>
- Jerbi S, Poulsen Nautrup H, Trenkwalder LM, Dunjko BHJV (2019) A framework for deep energy-based reinforcement learning with quantum speed-up. *arXiv:1910.12760*
- Kimble HJ (2008) The quantum internet. *Nature* 453:1023. <https://www.nature.com/articles/nature07127>
- Levit A, Crawford D, Ghadermarzy N, Oberoi JS, Zahedinejad E, Ronagh P (2017) Free energy-based reinforcement learning using a quantum processor. *arXiv:1706.00074*
- Melnikov AA, Makmal A, Briegel HJ (2018) Benchmarking projective simulation in navigation problems. *IEEE Access* 6:64639. <https://doi.org/10.1109/ACCESS.2018.2876494>
- Melnikov AA, Poulsen Nautrup H, Krenn M, Dunjko V, Tiersch M, Zeilinger A, Briegel HJ (2018) Active learning machine learns to create new quantum experiments. *Proc Natl Acad Sci* 115(6):1221. <https://www.pnas.org/content/115/6/1221>
- Neukart F, Von Dollen D, Seidel C, Compostella G (2018) Quantum-enhanced reinforcement learning for finite-episode games with discrete state spaces. *Front Phys* 5:71. <https://doi.org/10.3389/fphy.2017.00071>
- Paparo GD, Dunjko V, Makmal A, Martín-Delgado MA, Briegel HJ (2014) Quantum speedup for active learning agents. *Phys Rev X* 4:031002. <https://doi.org/10.1103/PhysRevX.4.031002>
- Ronagh P (2019) Quantum algorithms for solving dynamic programming problems. *arXiv:1906.02229*
- Russell SJ, Norvig P (2003) *Artificial intelligence: a modern approach*, 2nd edn. Pearson Education, London

- Saggio V, Asenbeck B, Hamann A, Strömberg T, Schiansky P, Dunjko V, Friis N, Harris NC, Hochberg M, Englund D, Wölk S, Briegel HJ, Walther P (2021) Experimental quantum speed-up in reinforcement learning agents. *Nature* 591:229. <https://doi.org/10.1038/s41586-021-03242-7>
- Singh S, Bertsekas D (1996) Reinforcement learning for dynamic channel allocation in cellular telephone systems. In: *Proceedings of the 9th International Conference on Neural Information Processing Systems, NIPS 1996*, p 974
- Sriarunothai T, Wölk S., Giri GS, Friis N, Dunjko V, Briegel HJ, Wunderlich C (2019) Speeding-up the decision making of a learning agent using an iontrap quantum processor. *Quantum Sci Technol* 4:015014
- Sutton R, Barto A (1998) *Reinforcement learning*. The MIT Press, Cambridge
- Tesauro G, Das R, Chan H, Kephart J, Levine D, Rawson F, Lefurgy C (2008) Managing power consumption and performance of computing systems using reinforcement learning. In: *Advances in neural information processing systems*, vol 20, p 1497
- Yoder TJ, Low GH, Chuang IL (2014) Fixed-point quantum search with an optimal number of queries. *Phys Rev Lett* 113:210501. <https://doi.org/10.1103/PhysRevLett.113.210501>
- Zalka C (1999) Grover's quantum searching algorithm is optimal. *Phys. Rev. A* 60:2746. <https://doi.org/10.1103/PhysRevA.60.2746>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.