Adaptive Procedures for Discriminating Between Arbitrary Tensor-Product Quantum States

Sarah Brandsen*, Mengke Lian†, Kevin D. Stubbs‡, Narayanan Rengaswamy† and Henry D. Pfister†‡
*Department of Physics, Duke University, Durham, North Carolina 27708, USA
†Department of Electrical and Computer Engineering, Duke University, Durham, North Carolina 27708, USA
‡Department of Mathematics, Duke University, Durham, North Carolina 27708, USA
Email: {sarah.brandsen, mengke.lian, kevin.stubbs, narayanan.rengaswamy, henry.pfister}@duke.edu

***Abstract*—Discriminating between quantum states is a fundamental task in quantum information theory. Given two quantum states, ρ_+ and ρ_-, the Helstrom measurement distinguishes between them with minimal probability of error. However, finding and experimentally implementing the Helstrom measurement can be challenging for quantum states on many qubits. Due to this difficulty, there is a great interest in identifying local measurement schemes which are close to optimal. In the first part of this work, we generalize previous work by Acin et al. (Phys. Rev. A 71, 032338) and show that a locally greedy (LG) scheme using Bayesian updating can optimally distinguish between any two states that can be written as a tensor product of arbitrary pure states. We then show that the same algorithm cannot distinguish tensor products of mixed states with vanishing error probability (even in a large subsystem limit), and introduce a modified locally-greedy (MLG) scheme with strictly better performance. In the second part of this work, we compare these simple local schemes with a general dynamic programming (DP) approach. The DP approach finds the optimal series of local measurements and optimal order of subsystem measurement to distinguish between the two tensor-product states.** [1]

***Index Terms*—quantum state discrimination, LOCC, Helstrom measurement, dynamic programming, quantum hypothesis testing**

I. Introduction

Measurement lies at the heart of quantum mechanics. Since the exact state of a quantum system cannot be directly observed, measurement is the primary means of understanding real world quantum systems [1]–[6]. Due to the inherent uncertainty in quantum systems it is impossible to design a quantum measurement capable of perfectly discriminating between two non-orthogonal quantum states [7], [8]. The optimal measurement for state discrimination was described by Helstrom [9]. However, for large composite quantum systems, the Helstrom measurement can be computationally expensive to compute and impractical to implement experimentally because it may require simultaneously measuring all subsystems.

Several works in the literature have investigated techniques that use only local operations to distinguish between two possible qubit states. Given N copies of the state, the aim is to achieve or approximate the Helstrom probability of success. Such algorithms perform N rounds of measurement where one local subsystem is measured in each round. The next measurement then is chosen as a function of the past measurement results. The simplest strategy, a "majority vote", has been shown to have probability of error which approaches zero exponentially fast in N [10], [11]. For N copies of mixed qubit states, tight bounds on the error rate of the best locally adaptive protocol can be found [12]. Dynamic programming has also been used to recursively minimize the expected future error over all possible allowed measurements and thus compute the optimal adaptive strategy for any given family of measurements [11]. Finally, for the special case where the states are tensor powers of pure states, it has been shown that a greedy adaptive strategy, involving Bayesian updates of the prior after each measurement result, is optimal and achieves the same success probability as the collective Helstrom measurement [10].

In this paper, we generalize known results and consider the problem of discriminating between two arbitrary tensor product quantum states (TPQS) with a focus on qubit and qutrit subsystems. More specifically, we suppose that we are given either ρ_+ or ρ_- with prior probability q and $1-q$ respectively, where $\rho_\pm = \rho_\pm^{(1)} \otimes \cdots \otimes \rho_\pm^{(N)}$ and $\rho_\pm^{(j)}$ is potentially different for each $j \in \{1, \cdots, N\}$. Thus, we relax the constraint that all subsystems are identical copies, and instead also consider cases where distinct subsystems are different from each other. This problem is of practical interest in quantum communications, where we might modulate a classical binary codeword into a TPQS in order to transmit information through multiple uses of the channel, and each subsystem could experience a (slightly) different channel noise parameter. The received codeword set then consists of non-orthogonal noisy TPQS, which may be distinguished by locally adaptive methods. We consider three locally adaptive algorithms: a locally-greedy (LG) algorithm, a modified locally-greedy (MLG) algorithm, and a more general dynamic programming (DP) algorithm. The general DP algorithm we introduce extends previous works to additionally optimize over the order in which subsystems are measured.

In the special case where all subsystems are pure states, then the order of measurement does not matter and the Bayesian update-based strategy with locally-greedy measurements is optimal. This generalizes the result in [10] mentioned above. In the case where not all subsystems are pure, we show that in general the probability of success is affected by the order in

[1] An extended version of this paper is accessible at: https://arxiv.org/abs/1912.05087

which the subsystems are measured. Thus, we need to optimize over both the order in which the subsystems are measured and the measurement implemented on each subsystem. Additionally, for general TPQS, the optimal order depends on the measurement outcomes and must be determined round-by-round.

When the states are mixed, the locally-greedy algorithm is not optimal and, in fact, performs worse than most nonadaptive local strategies in the limit as $N \to \infty$. We show that this poor asymptotic performance arises from the local Helstrom measurement becoming noninformative for sufficiently imbalanced priors. To overcome this, we introduce a modified locally-greedy adaptive strategy with strictly better performance.

We also discuss a dynamic programming-based strategy that finds the optimal locally-adaptive strategy, generalizing the technique introduced in [11] to include optimizing over the order in which subsystems are measured. This dynamic programming approach is the optimal locally-adaptive technique subject to some simple constraints and includes the locally-greedy techniques as a special case of itself.

Finally, we consider the performance of ternary and binary projective measurements over qutrit states and show that, in general, multiple-outcome measurements are needed for optimality. This holds even for a binary state discrimination problem. Numerical results are provided for all these scenarios and the source code used to generate them is available at https://github.com/SarahBrandsen/AdaptiveStateDiscrimination.

II. Locally Greedy Algorithms

First, we describe a simple locally-greedy algorithm, which was called the "locally optimal locally adaptive" algorithm in [11]. Suppose we are given candidate states $\rho_\pm = \bigotimes_{j=1}^{N} \rho_\pm^{(j)}$ and prior $q = \Pr(\rho = \rho_+)$. By construction, the Helstrom measurement for ρ_+, ρ_-, and q is the optimal measurement for distinguishing between the two candidate states. If we are only given access to subsystem j and the prior $p = \Pr(\rho^{(j)} = \rho_+^{(j)} \big|$ past) conditioned on past measurement results, then the locally-optimal measurement is the (local) Helstrom measurement for $\rho_+^{(j)}$, $\rho_-^{(j)}$, defined as:

$$\Pi(p,j) \triangleq \sum_{|v\rangle \in \mathcal{V}(p,j)} |v\rangle\langle v| \text{ where} \tag{1}$$

$$\mathcal{V}(p,j) \triangleq \left\{ |v\rangle : M(p,j)|v\rangle = \lambda |v\rangle,\ \lambda \geq 0 \right\} \tag{2}$$

and where $M(p,j) \triangleq (1-p)\rho_-^{(j)} - p\rho_+^{(j)}$.

Round j of the locally-greedy algorithm consists of measuring subsystem j using the local Helstrom measurement for the current prior, then updating the prior via Bayes Theorem based on the measurement result. Hence, all information from previous measurements is compressed into the updated prior, and the adaptivity of the locally-greedy algorithm comes exclusively from these prior updates.

The set of allowed measurement outcomes in each round are $\{+,-\}$, and the outcome of the kth round is denoted by $d_k \in \{+,-\}$. All measurement results before the j-th round are denoted $\mathbf{d}_{1:j-1}$ (or more succinctly $\mathbf{d}_{[j-1]}$ with $[j] \triangleq \{1,\ldots,j\}$). Finally, the updated prior given these measurement is denoted by $p_j = P_j(q, \mathbf{d}_{[j-1]})$ and the locally-greedy algorithm is equivalent to implementing $\Pi(p_j, j)$ in round j. At the end of the algorithm, the state is decoded as ρ_+ if $p_{N+1} \geq \frac{1}{2}$ and ρ_- else, such that the total success probability is $\max(p_{N+1}, 1-p_{N+1})$.

In the special case where $\rho_\pm$ is a tensor product of arbitrary pure states, we prove analytically that the locally-greedy algorithm achieves the same success probability as the optimal Helstrom measurement.

Theorem 2.1: Let $P_{\mathrm{s,h}}(q,\rho_\pm)$ and $P_{\mathrm{s,lg}}(q,\rho_\pm)$ denote the probabilities of successful state discrimination, given initial prior $\mathbb{P}(\rho = \rho_+) = q$, using the joint N-system Helstrom measurement and the locally greedy measurement technique, respectively. If ρ_+ and ρ_- are pure states, i.e., $\rho_\pm^{(j)} = |\pm\theta_j\rangle\langle\pm\theta_j|$ where $|\theta\rangle \triangleq \cos\frac{\theta}{2}|0\rangle + \sin\frac{\theta}{2}|1\rangle$, for some $\theta_j \in (0, 2\pi)$ for every $j \in [N]$, then

$$P_{\mathrm{s,h}}(q,\rho_\pm) = P_{\mathrm{s,lg}}(q,\rho_\pm) \tag{3}$$

$$= \frac{1}{2}\left(1 + \sqrt{1 - 4q(1-q)\Pi_{j=1}^{N}\cos^2(\theta_j)}\right). \tag{4}$$

Sketch of Proof: The strategy is to prove the result for $N = 2$ and then extend via induction for arbitrary N. ■

Plateau with locally greedy algorithm: When $\rho_\pm$ are tensor products of depolarized pure states, for some depolarizing parameter γ, we observe that the average probability of success (asymptotically) approaches a value strictly less than 1 when the depolarizing parameter is nonzero. Thus, despite the optimal performance of the locally-greedy algorithm for pure states, it is no longer even asymptotically optimal for mixed states (and can be outperformed by nonadaptive strategies as $N \to \infty$).

The experimental setup is as follows. A set of 1000 candidate pure states is generated by sampling $\theta_{t,\pm}$ from the continuous uniform distribution on the interval $[0,\pi]$ for $t \in [1,..,1000]$ and forming the corresponding quantum states where $|\theta\rangle \triangleq \cos\frac{\theta}{2}|0\rangle + \sin\frac{\theta}{2}|1\rangle$. Then, for every γ in the set of allowed depolarizing parameters, $\{0.01, 0.05, 0.1, 0.3\}$ and for every $N \in [1,2,...,12]$, the candidate states for the t-th trial are given by

$$\rho_\pm(\gamma,t,N) \triangleq \left((1-\gamma)|\theta_{\pm,t}\rangle\langle\theta_{\pm,t}| + \frac{\gamma}{2}I\right)^{\otimes N}.$$

Since each candidate state consists of N identical subsystems, the locally greedy method amounts to performing Helstrom measurements based on the updated prior over N copies of the same state. We plot the Monte Carlo average performance, $P_{\mathrm{succ}}(N,\gamma)$, for fixed N and γ. The results of this computational experiment are shown in Fig. 1. In the case where $\gamma = 0$, the probability of success must approach 1 with increasing N because the locally greedy approach recovers the optimal Helstrom performance (see Theorem 2.1).

Next, we provide a result that explains the performance plateaus in Fig. 1 and then we define a modified locally greedy approach that overcomes this sub-optimality.

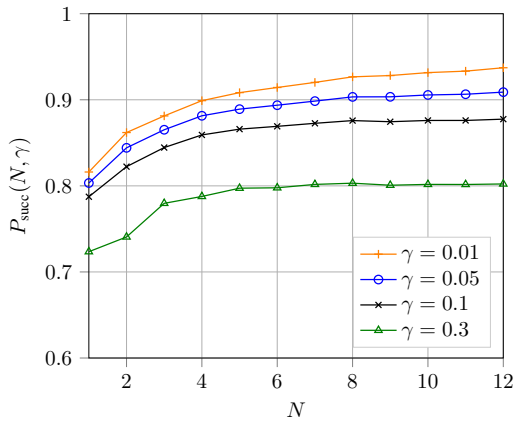

Fig. 1. Comparison of probability of success for varying γ in the case of identical copies, as a function of the number of available systems. Based on the computational results, we observe that as the depolarizing parameter increases, the probability of success levels off for large N.

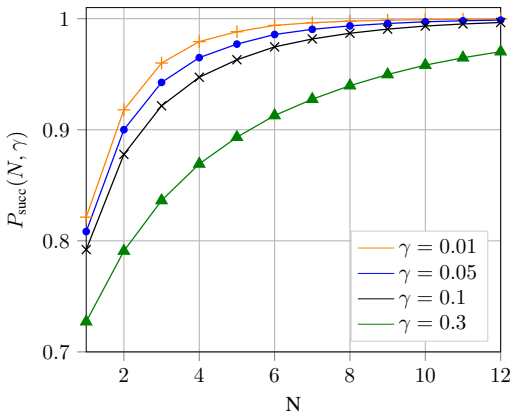

Fig. 2. Comparison of probability of success for varying γ in the case of identical copies, as a function of the number of available systems using the MLG algorithm. We observe that as the depolarizing parameter increases, the probability of success no longer levels off for large N.

Theorem 2.2: Consider the problem of distinguishing between two distinct single-qubit states ρ_+^{dep} and ρ_-^{dep}, with prior probabilities q and $1-q$ respectively. Assume that ρ_+^{dep} and ρ_-^{dep} are depolarized pure states[2] $|\psi_+\rangle\langle\psi_+|$, $|\psi_-\rangle\langle\psi_-|$ such that

$$\gamma_\pm \in [0,1] \quad \text{and} \quad \rho_\pm^{\text{dep}} \triangleq (1-\gamma_\pm)\,|\psi_\pm\rangle\langle\psi_\pm| + \frac{\gamma_\pm}{2} I.$$

For any choice of $\gamma_\pm$, $q \in [0,1]$ the probability of correctly distinguishing ρ_+^{dep} and ρ_-^{dep} is denoted by $P_{\text{succ}}^{\text{dep}}$ and satisfies

$$P_{\text{succ}}^{\text{dep}} \leq \max\left\{1-q, q, 1-\frac{\gamma_{min}}{2}\right\}, \tag{5}$$

where $\gamma_{\min} \triangleq \min(\gamma_+, \gamma_-)$.

This theorem implies that once the prior is updated so that either q or $1-q$ is greater than $1-\frac{\gamma}{2}$, the locally greedy algorithm is stuck making noninformative measurements on all subsequent subsystems. Therefore the error does not approach 0 as $N \to \infty$. Similarly, it can be shown that the locally greedy method also exhibits plateaus in more general scenarios. These arise from the fact that the Helstrom measurement becomes trivial, such that $\Pi(p,j) \in \{0, \mathbb{I}\}$, when $\max(q, 1-q) \geq 1-\frac{\gamma}{2}$.

III. Modified Locally Greedy Algorithms

The poor asymptotic behaviour of the locally-greedy algorithm provides motivation for our introduction of a modified locally greedy method (MLG method) with better asymptotic properties. In particular, the structure is identical to the locally-greedy algorithm except that a "modified Helstrom" measurement is implemented in each round. The modified Helstrom measurement is defined by:

$$\Pi^*(p,j) \triangleq \begin{cases} \Pi(p,j) & \text{if } \Pi(p,j) \notin \{\mathbb{I}, 0\} \\ |v_{\lambda_{\max}}\rangle\langle v_{\lambda_{\max}}| & \text{if } \Pi(p,j) = 0 \\ \mathbb{I} - |v_{\lambda_{\min}}\rangle\langle v_{\lambda_{\min}}| & \text{if } \Pi(p,j) = \mathbb{I}, \end{cases}$$

[2]Note that an arbitrary qubit state (density matrix) can always be expressed as a pure state passed through a depolarizing channel, because this procedure can define any state in the Bloch sphere [7].

where $\lambda_{\max} \triangleq \max_\lambda \left\{\lambda \;\middle|\; \left((1-p)\rho_-^{(j)} - p\rho_+^{(j)}\right)|v_\lambda\rangle = \lambda\,|v_\lambda\rangle\right\}$ and $\lambda_{\min} \triangleq \min_\lambda \left\{\lambda \;\middle|\; \left((1-p)\rho_-^{(j)} - p\rho_+^{(j)}\right)|v_\lambda\rangle = \lambda\,|v_\lambda\rangle\right\}$.

Whenever the Helstrom measurement is nontrivial, the modified Helstrom measurement is equivalent and thus locally-optimal by definition. In the case where the Helstrom measurement is trivial, it can be shown that any other measurement and outcome would lead to identical posterior-based decoding (i.e. any measurement is locally optimal). The modified Helstrom measurement takes advantage of this degeneracy to create a more informative measurement by separating out the projector that is most strongly predictive of the less-likely candidate state.

For any TPQS $\rho_\pm$ and any q, denote by $P_{s,mlg}(q,\rho_\pm)$ (respectively, $P_{s,lg}(q,\rho_\pm)$) the probability of successfully distinguishing the states via the MLG algorithm (LG algorithm). It can be shown that, for any $\rho_\pm$ and any $q \in [0,1]$, we have $P_{s,mlg} \geq P_{s,lg}$. Thus, the MLG algorithm is always at least as good as the LG algorithm. Additionally, for any $\rho_\pm$, it follows that $P_{s,mlg}(\rho_\pm) \to 1$ as the number of subsystems j such that $\rho_+^{(j)} \neq \rho_-^{(j)}$ approaches infinity. Thus, the MLG algorithm no longer exhibits an asymptotic plateau in the success probability, and this is illustrated in Fig. 2.

IV. Dynamic-Programming Algorithms

The dynamic programming (DP) algorithm works backwards to recursively compute an expected future risk function $R_S\colon [0,1] \times S \to [0,1]$, where S denotes the set of subsystem indices that are yet to be measured. The domain corresponds to the current updated prior and unmeasured subsystems and the codomain is the expected probability of state discrimination error given that an optimal locally adaptive algorithm is applied in all remaining rounds. For each round, DP generates the optimal choice of the next subsystem to be measured and the optimal measurement to implement (given an allowed action set.)

Now, we introduce some notation for purposes of defining the expected risk function. Let σ be the ordering permutation

such that, at round $j \in \{1, \ldots, N\}$, the next subsystem measured is $\sigma(j)$. The algorithm must additionally determine the chosen measurement action $\mathbf{a}_{\sigma(j)} \in \mathcal{A}$ where $\mathcal{A}$ is the set of allowed measurement POVMs. The measurement result upon executing the action can take on values $d_{\sigma(j)} \in \mathcal{D}$, where $\mathcal{D}$ is the space containing possible outcomes for the chosen action set. For example, if $\mathcal{A}$ contains projective measurements on qubits, then $\mathcal{D} = \{\pm 1\}$. Then at round j, the past actions and results are recorded into the vectors $\mathbf{a}^{\sigma}_{[j-1]} = (\mathbf{a}_{\sigma(1)}, \ldots, \mathbf{a}_{\sigma(j-1)})$ and $\mathbf{d}^{\sigma}_{[j-1]} = (d_{\sigma(1)}, \ldots, d_{\sigma(j-1)})$ respectively.

Formally, at round j,

$$S_j \triangleq S_{j-1} \setminus \{\sigma(j-1)\} \tag{6}$$
$$= [N] \setminus \sigma([j-1]). \tag{7}$$

For the case $S = \emptyset$, one can make a hard decision on the updated prior $p_N \triangleq P^{\sigma}_N(q, \mathbf{a}^{\sigma}_{[N]}, \mathbf{d}^{\sigma}_{[N]})$, i.e., by comparing it to 0.5. Hence

$$R_{\emptyset}(p_N) = \min(p_N, 1 - p_N), \tag{8}$$

For the general case $S \neq \emptyset$ and $j = N - |S| + 1$, $N - |S|$ measurements have been performed. The goal is to choose the best subsystem $\sigma(j)$ to be measured during the j-th step in order to minimize the expected error probability assuming optimal future decisions. Thus, we have

$$R_S(p_{N-|S|}) = \min_{(k,\mathbf{a}_k) \in S \times \mathcal{A}} \sum_{d_k \in \mathcal{D}} \mathbb{P}\left(d_k \;\middle|\; p_{N-|S|},\ \mathbf{a}_k\right)$$
$$\times R_{S \setminus \{k\}}\left(P_{N-|S|+1}\left(p_{N-|S|}, \mathbf{a}_k, d_k\right)\right), \tag{9}$$

where the optimal subsystem and action pairing are thus those that achieve the minimum in the above equation. The DP algorithm optimizes over all other locally-adaptive algorithms with the same allowed measurement set, and thus is guaranteed to perform at least as well as the LG and MLG algorithms when the action space is over all binary POVMs.

Effect of ordering on success probability: One general question is whether subsystem ordering affects the probability of success when optimization is done over all "reasonable" adaptive protocols. We address the question of ordering by first demonstrating analytically that ordering can make a difference for a specific subset of candidate states when $N = 2$. Consider candidate states of the form:

$$\rho_+ = \begin{pmatrix} 1-x & 0 \\ 0 & x \end{pmatrix} \otimes |\theta\rangle\langle\theta|,$$
$$\rho_- = \begin{pmatrix} x & 0 \\ 0 & 1-x \end{pmatrix} \otimes |-\theta\rangle\langle\theta|.$$

Measuring the subsystems in the best order (diagonal matrices first followed by $|\pm\theta\rangle\langle\pm\theta|$) is equivalent to updating the prior from $q = \frac{1}{2}$ to x and then implementing a Helstrom measurement on the second subsystem with the updated prior. The resulting probability of success is optimal and performs as well as a composite Helstrom measurement on both subsystems, namely, $P_{\text{succ,best}} = \frac{1}{2}(1 + \sqrt{1 - 4(1-x)x\cos^2(2\theta)})$.

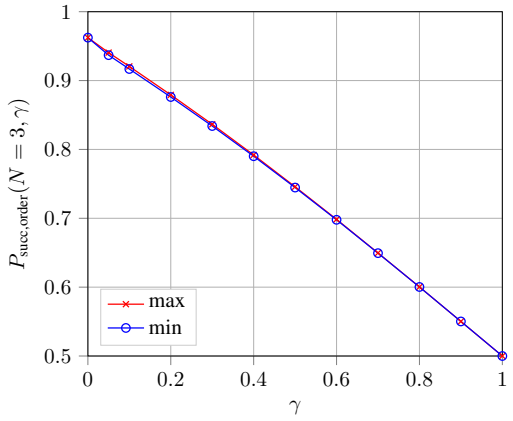

Fig. 3. Comparison of probabilities $P_{\text{succ,best}}(N = 3, \gamma)$ and $P_{\text{succ,worst}}(N = 3, \gamma)$ as a function of the depolarising parameter γ over 1000 trials. Although $P_{\text{succ,best}}(N = 3, \gamma) \neq P_{\text{succ,worst}}(N = 3, \gamma)$, the relative difference is small.

Measuring in the reverse order, the probability of success is $P_{\text{succ,worst}} = \max\{x, 1-x, \frac{1}{2}(1 + \sqrt{1 - \frac{1}{2}\cos^2(2\theta)})\}$ given that the diagonal subsystems are always optimally measured in the computational basis regardless of previous information. Thus, there is, in general, a difference between the best and worst ordering.

More generally, we show that a small but nontrivial difference persists for more general tensor product states under the DP algorithm. The experimental protocol involves taking 1000 random trials where $\hat{\rho}^{(j)}_{\pm}$ are all real qubit states. The set of measurements $\mathcal{A}_{\text{qubit}}$ is taken to be the standard action space of real orthogonal projectors [11]

$$\mathcal{A}_{\text{qubit}} \triangleq \left\{\{|\phi\rangle\langle\phi|, |\phi^{\perp}\rangle\langle\phi^{\perp}|\} : \phi \in \left[0, \frac{\pi}{2}\right]\right\}, \tag{10}$$

where we quantize ϕ into $Q_{\phi} = 128$ equally spaced points. For each N and γ considered, we then find the average success probability in two distinct cases– first, for the "best" ordering (the standard DP algorithm) and second for the "worst" ordering (where the risk function is maximized with respect to subsystem index k but still minimized with respect to the action $\mathbf{a}_k$).

We plot $P_{\text{succ,order}}(N = 3, \gamma)$ as a function of γ in Fig. 3 and we also compare the difference $P_{\text{succ,diff}}(N, \gamma) \triangleq P_{\text{succ,best}}(N, \gamma) - P_{\text{succ,worst}}(N, \gamma)$ for $N \in \{3, 4, 5, 6, 7\}$ in Fig. 4. From these results, we observe that the difference in probability of success with respect to ordering is quite small but persists even when using the DP algorithm.

Insufficiency of binary projective measurements: Finally, we investigate whether adaptive binary projective measurements are always sufficient for general quantum states. This question is motivated in part by recalling that in the special case when $\rho_{\pm}$ are both pure, Theorem 2.1 shows that the optimal adaptive strategy consists of binary projective measurements. To this aim, we define an action space $\mathcal{A}$ to be *s*ufficient for state space $\mathcal{H}$ if and only if for all $\rho_{\pm} \in \mathcal{H}$ and $q \in [0, 1]$,

$$P_{\text{succ},\mathcal{A}}(q, \rho_{\pm}) = P_{\text{succ},\mathcal{A}_{\text{all}}}(q, \rho_{\pm}),$$

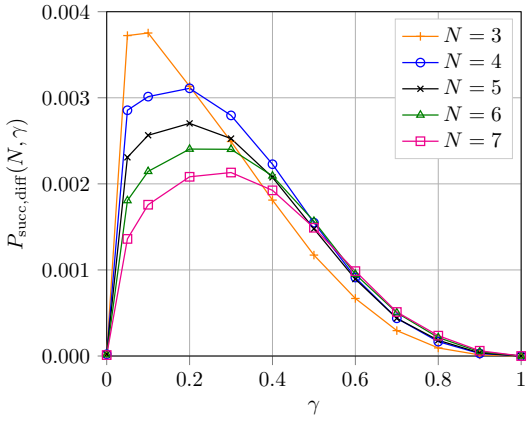

Fig. 4. Difference in maximum and minimum probability of success, $P_{\text{succ,diff}}(N,\gamma)$, as a function of the depolarizing parameter γ over 1000 trials for $N = 3, 4, 5, 6, 7$.

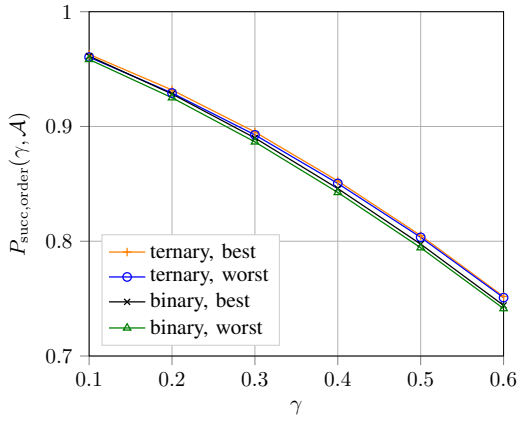

Fig. 5. The average probability of success for the best and worst ordering using both ternary and binary projective measurements for qutrit product states when $N = 3$. Results are averaged over 1000 trials.

where $\mathcal{A}_{\text{all}}$ is the set of all quantum measurements of dimension $\dim(\rho_{\pm})$ and $P_{\text{succ},\mathcal{A}}(q,\rho_{\pm})$ is the probability of success of the DP algorithm for a given action space $\mathcal{A}$.

We show by example that binary projective measurements are not sufficient for general state spaces. To this aim, we define $\mathcal{H}_{\text{qutrit}}$ to be the space of depolarized, real qutrit states and define the action space of real binary (ternary) measurements as $\mathcal{A}_{\text{b}}$ ($\mathcal{A}_{\text{t}}$).

We fix $N = 3$ and randomly generate 1000 sets of pure TPQS candidate states, where each subsystem is a qutrit. Then for each depolarizing parameter γ, we compute the average success probability of four separate methods: use both the best and worst ordering in combination with both binary ($\mathcal{A} = \mathcal{A}_b$) and ternary ($\mathcal{A} = \mathcal{A}_t$) action spaces. Results for all four methods are shown in Fig. 5. The difference between the remaining three methods and the ternary best-ordering method ($P_{\text{diff,order}}(\gamma,\mathcal{A}) = P_{\text{succ,best}}(\gamma,\mathcal{A}_t) - P_{\text{succ,order}}(\gamma,\mathcal{A})$) is shown in Fig. 6. We observe that the best ternary ordering is better than best binary ordering and again the ordering affects performance. From this, we conjecture that, for any action space and any adaptive approach, the order of subsystem measurement will affect the success probability. It remains an open question whether it is sufficient to consider d rank-1 orthogonal projectors for a state space $\mathcal{H}_d$ containing d-dimensional real quantum states.

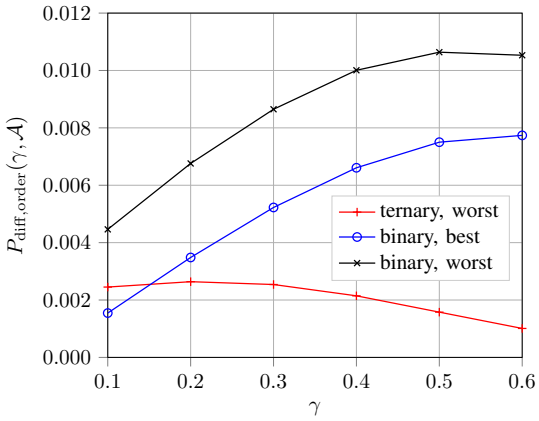

Fig. 6. Loss in average success probability for the various methods relative to ternary best, namely, $P_{\text{diff,order}}(\gamma,\mathcal{A})$ as a function of γ when $N = 3$. Results are averaged over 1000 trials.

V. Conclusion

In this work, we investigated simple locally-greedy and modified locally-greedy algorithms as well as more general dynamic programming algorithms for quantum state discrimination when the given states are tensor products of N arbitrary qubit or qutrit states. We prove analytically that, when the individual subsystems are pure states, the simple locally-greedy algorithm achieves the optimal performance of the joint N-system Helstrom measurement. For the scenario where each subsystem contains arbitrary qubit states, we demonstrate a plateau in the probability of success attained by the locally-greedy algorithm with increasing N. The reason for this plateau is found and an explicit bound is derived for the success probability as a function of the channel depolarizing parameter and initial prior. Based on these results, a modified locally-greedy algorithm is introduced with strictly better performance and its state discrimination becomes perfect in the large N limit.

For the general DP algorithm, we show that ordering of subsystems affects the performance when the individual subsystems have distinct states. For qutrit states, we show that binary projective measurements are inadequate to achieve optimal performance.

Acknowledgment

The authors would like to thank Iman Marvian for helpful discussions. The work of Brandsen, Rengaswamy, and Pfister was supported in part by the National Science Foundation (NSF) under Grant No. 1908730 and 1910571. Any opinions, findings, conclusions, and recommendations expressed in this material are those of the authors and do not necessarily reflect the views of these sponsors.

REFERENCES

[1] A. Peres, *Quantum theory: Concepts and methods*, vol. 57. Springer Science & Business Media, 2006.

[2] J. Renes, R. Blume-Kohout, A. Scott, and C. Caves, "Symmetric informationally complete quantum measurements," *Journal of Mathematical Physics*, vol. 45, 06 2004.

[3] A. Bisio, G. Chiribella, G. M. D'Ariano, S. Facchini, and P. Perinotti, "Optimal quantum tomography of states, measurements, and transformations," *Phys. Rev. Lett.*, vol. 102, p. 010404, 2009.

[4] M. Dall'Arno, S. Brandsen, F. Buscemi, and V. Vedral, "Device-independent tests of quantum measurements," *Phys. Rev. Lett.*, vol. 118, p. 250501, 2017.

[5] P. Busch, "Quantum states and generalized observables: A simple proof of gleason's theorem," *Phys. Rev. Lett.*, vol. 91, p. 120403, 2003.

[6] R. W. Spekkens, "Contextuality for preparations, transformations, and unsharp measurements," *Phys. Rev. A*, vol. 71, p. 052108, 2005.

[7] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition*. New York, NY, USA: Cambridge University Press, 10th ed., 2011.

[8] N. Gisin, "Quantum cloning without signaling," *Physics Letters A*, vol. 242, pp. 1–3, 1998.

[9] C. W. Helstrom, "Quantum detection and estimation theory," *Journal of Statistical Physics*, vol. 1, no. 2, pp. 231–252, 1969.

[10] A. Acín, E. Bagan, M. Baig, L. Masanes, and R. Muñoz Tapia, "Multiple-copy two-state discrimination with individual measurements," *Phys. Rev. A*, vol. 71, p. 032338, 2005.

[11] B. L. Higgins, A. C. Doherty, S. D. Bartlett, G. J. Pryde, and H. M. Wiseman, "Multiple-copy state discrimination: Thinking globally, acting locally," *Phys. Rev. A*, vol. 83, p. 052314, 2011.

[12] J. Calsamiglia, J. I. de Vicente, R. Muñoz Tapia, and E. Bagan, "Local discrimination of mixed states," *Phys. Rev. Lett.*, vol. 105, p. 080504, 2010.