

Efficient Characterization of Coherent and Correlated Low-Degree Noise in Layers of Gates

Marianna Crupi^{1,2}, J. Ignacio Cirac^{1,2}, and Flavio Baccari³

¹*Max Planck Institute of Quantum Optics, 85748 Garching, Germany*

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

³*Dipartimento di Fisica e Astronomia “G. Galilei” & Padua Quantum Technologies Research Center, Università degli Studi di Padova, I-35131 Padova, Italy*

Summary Today’s quantum devices are affected by the problem of noise, which is still a non-negligible part of quantum computation. Accurate noise characterization is therefore a crucial task that must be implemented for further development of quantum technologies. Existing methods, though widely used, are lacking in at least one of three crucial aspects of an ideal noise characterization protocol: scalability, ability to capture the full structure of realistic noise and experimental feasibility. In this work, we address the problem of reconstructing the noise affecting a *layer of gates*, and present a protocol that combines all three features by relying on the physically motivated ansatz of *low-degree quantum channels*—a structure that arises naturally in the regime of weak noise. We present two main contributions that make this possible. First we introduce a single-qubit randomized-measurement protocol that enables the efficient and experimentally-accessible reconstruction of low-degree quantum channels, with recovery guarantees in ℓ_2 and diamond-norm. This is achieved through a tailored choice of the shadow-tomography reconstruction function, which we exploit to obtain the first exponential-to-polynomial improvement in shadow sample complexity for a practically relevant scenario. Second, we extend the protocol to the case of noise arising from a layer of gates. To do that we show how to isolate the noise component of the evolution (the truly low-degree one) by implementing gate inversion in classical post-processing.

Setting. Quantum process tomography, i.e. the complete reconstruction of an unknown quantum channel, has long been a central challenge in quantum information [1]. Although it provides full knowledge of system dynamics, its standard formulations require resources that scale exponentially with system size [1, 2], severely limiting practical applicability. This limitation is particularly relevant for noise characterization in quantum devices, where understanding operational imperfections is crucial for error mitigation, correction, and hardware benchmarking [3–8]. Existing approaches face a fundamental trade-off: full tomographic methods, such as gate-set tomography, offer complete information but are prohibitively costly [1, 2, 9–12], while scalable techniques like randomized benchmarking and shadow-based methods efficiently estimate global figures of merit but cannot reconstruct the full noise process [8, 13–21].

An intermediate strategy is to assume an ansatz for the noise structure and attempt to reconstruct only its parameters, rather than the full exponential description. This strategy can dramatically reduce complexity, but so far it has come with important drawbacks. Common ansätze, such as restricting to Pauli-diagonal channels [22, 23], cannot capture coherent errors, which are known to be relevant in many physical platforms [24]. More expressive ansätze, on the other hand, have either relied on experimentally unrealistic requirements, such as multi-qubit measurements, or lacked rigorous recovery guarantees [25–29]. Among the proposed models,

the low-degree ansatz [29] has emerged as a physically motivated and expressive description of realistic noise, capturing both correlated and coherent errors while reducing the parameter space to a polynomial size. This ansatz arises naturally in the weak noise regime, and captures the intuition that errors arise as a stochastic mixture of local events: while the global channel may affect all qubits, each elementary error acts non-trivially only on a bounded number d of qubits. Formally, in the n -qubit Pauli basis $\{P_{\vec{\alpha}}\}_{\vec{\alpha} \in \{0,1,2,3\}^n}$ (with $P_0 = \mathbb{1}$), the channel can be written as: $\mathcal{E}(\rho) = \sum_{\vec{\alpha}, \vec{\beta}} \chi_{\vec{\alpha}\vec{\beta}} P_{\vec{\alpha}} \rho P_{\vec{\beta}}$, with $\chi_{\vec{\alpha}\vec{\beta}} = 0$ whenever $\text{wt}(P_{\vec{\alpha}}) > d$ or $\text{wt}(P_{\vec{\beta}}) > d$. This reduces the number of independent parameters from exponential in n to $\mathcal{O}(n^d)$, while still capturing important features such as coherent errors and crosstalk.

In this work, we adopt this ansatz and focus on reconstructing the noise generated by a layer of two-qubit gates. This setting raises two central challenges: (i) developing an efficient and experimentally friendly protocol for learning low-degree quantum channels, and (ii) isolating the noise process (the truly low-degree one) from the ideal implementation of the gate layer. In the following, we show how to address these challenges and highlight our main contributions.

(i) Efficient and experimentally-friendly tomography of low-degree quantum channels. Until now, efficient tomography of low-degree quantum channels from experimentally accessible data was lacking. We address this by presenting a protocol that enables complete and scalable reconstruction of low-degree channels using only single-qubit state preparation and measurement. Unlike previous learning-based approaches [29], which assume access to multi-qubit operations, our restriction to single-qubit control is a key conceptual and practical advance, making the protocol directly applicable to current quantum devices.

Our approach operates within the same randomized-measurement framework as shadow tomography [30]. In this setting, properties of quantum states or processes are inferred from data obtained via randomized single-qubit measurements and post-processed using a reconstruction function. Standard shadow tomography employs a fixed, canonical choice of this function, which is highly effective for estimating expectation values of quantum states. However, when applied directly to quantum process tomography, this choice leads to exponential sample complexity [20]. It has long been observed [31–33]—though often overlooked in practice—that the reconstruction function is not uniquely defined. In this work, we exploit this freedom in a concrete and practically relevant way: by tailoring the reconstruction function to the low-degree structure of the noise channel, we reduce the sample complexity from exponential to polynomial. To our knowledge, this represents the first example where customizing the reconstruction function yields a rigorous exponential-to-polynomial improvement in a practically relevant scenario. More precisely the sample complexity M to estimate a d -degree noise channel with probability $1 - \delta$ and precision ϵ in ℓ_2 and diamond-distance is respectively:

$$M_2 = e^{\mathcal{O}(d \log(1/\epsilon))} \cdot \mathcal{O}\left(\log\left(\frac{n}{\delta}\right)\right), \quad M_{\diamond} = n^{\mathcal{O}(d)} e^{\mathcal{O}(\log(1/\epsilon))} \cdot \mathcal{O}\left(\log\left(\frac{1}{\delta}\right)\right). \quad (1)$$

We believe this result to be of independent technical interest beyond the scope presented here. Moreover, we highlight that the experimental procedure remains unchanged: the measurements are identical to those of standard shadow tomography, involving only single-qubit preparations and measurements. This has two important implications. First, such operations are universally available and relatively reliable on today’s platforms, unlike multi-qubit preparations and measurements which require entangling gates and thus introduce significant noise. Second, since single-qubit operations are much less noisy, it is reasonable to treat them as effectively error-free,

which allows us to isolate and characterize the noise arising specifically from multi-qubit gates. These aspects make our tomography protocol not only experimentally feasible with current technology but also self-consistent in its reconstruction of noise processes.

(ii) Isolating the noise process. The second challenge concerns gate inversion. In fact, the discussion so far assumed that the experimental data is taken directly from the noise channel. In practice, however, we never have access to the isolated noise process. What we observe is the combined effect of the intended gate and the accompanying noise. This distinction is crucial: even if the noise itself is low-degree, the composition of gate and noise will in general no longer have that structure—for instance, when considering a full layer of entangling gates. The need to separate the gate action from the noise has long been recognized and has been addressed in various ways. The most common approaches are: (i) introducing randomness through gate sequences [8], which averages out the action of the gate but sacrifices all coherent information about the noise, or (ii) performing the inversion directly on hardware, which introduces further errors due to the additional gate applications.

Our approach is different: we propose to invert the ideal gates entirely in classical post-processing. This strategy avoids the extra noise introduced by hardware inversion while still giving experimental access to the full noise channel. Such post-processing, however, could undermine the efficiency of the protocol. We analyze this issue and identify broad classes of circuits where the classical inversion can be carried out with polynomial overhead. We find that this includes layers of two-qubit gates that contain only a polylogarithmic number of non-Clifford operations. Sample complexity for the reconstruction of the noise arising from such a layer is:

$$M_2 = e^{\mathcal{O}(d \log(1/\epsilon))} \cdot \mathcal{O}\left(n^{p+1} \log\left(\frac{n}{\delta}\right)\right), \quad M_\diamond = n^{\mathcal{O}(d)+p} e^{\mathcal{O}(\log(1/\epsilon))} \cdot \mathcal{O}\left(\log\left(\frac{1}{\delta}\right)\right), \quad (2)$$

where p is the product of the exponents appearing in the polynomial dependence on $\log$ of the support of non-Clifford gates. By embedding this classical inversion within the shadow-tomography framework, we obtain an efficient and experimentally feasible protocol for the *complete characterization* of low-degree noise channels arising from a broad class of gate layers.

Numerical results. We complement our theoretical analysis with numerical simulations that validate the performance of the protocol. In a first scenario, we consider solely the noise channel consisting of low-degree but entangling errors, and apply our tomography method directly to reconstruct it. In a second, more realistic setting, we study a layer of *iSWAP* gates followed by single-qubit noise, which can be well approximated by a low-degree model. In both cases, the numerical results confirm that the sample complexity scales in accordance with our theoretical predictions, supporting the claim that our approach is both efficient and robust in practice.

Outlook. An important direction for future work is to understand and mitigate the impact of state preparation and measurement (SPAM) errors, which remain a dominant source of inaccuracy in current quantum platforms [24]. While our protocol assumes ideal single-qubit preparation and measurement, relaxing this assumption will be crucial for its practical implementation. Promising directions include adapting mitigation strategies developed within the shadow tomography framework, such as robust shadow estimation [34] and probabilistic error cancellation [35, 36]. Integrating these techniques could enable the suppression of SPAM-induced errors without fundamentally altering the structure of the protocol. A systematic treatment of SPAM errors in this setting, however, remains an open problem and is left for future investigation.

References

- [1] Isaac L. Chuang and M. A. Nielsen. Prescription for experimental determination of the dynamics of a quantum black box. *Journal of Modern Optics*, 44(11–12):2455–2467, 1997.
- [2] J. F. Poyatos, J. I. Cirac, and P. Zoller. Complete characterization of a quantum process: The two-bit quantum gate. *Physical Review Letters*, 78(2):390–393, 1997.
- [3] Youngseok Kim, Andrew Eddins, Shantanu Anand, Daniel J. Egger, Sergey Bravyi, Andrew W. Cross, Abhinav Kandala, Jerry M. Chow, and Jay M. Gambetta. Evidence for the utility of quantum computing before fault tolerance. *Nature*, 618:500–505, 2023.
- [4] Suguru Endo, Zhenyu Cai, Simon C. Benjamin, and Xiao Yuan. Hybrid quantum-classical algorithms and quantum error mitigation. *Journal of the Physical Society of Japan*, 90(3):032001, 2021.
- [5] Alireza Seif, Ze-Pei Cui, Sisi Zhou, Senrui Chen, and Liang Jiang. Shadow distillation: Quantum error mitigation with classical shadows for near-term quantum processors. *PRX Quantum*, 4(1):010303, 2023.
- [6] Barbara M. Terhal. Quantum error correction for quantum memories. *Rev. Mod. Phys.*, 87:307–346, Apr 2015.
- [7] Robin Blume-Kohout, John Gamble, Erik Nielsen, Kenneth Rudinger, Jonathan Mizrahi, Kevin Fortier, and Peter Maunz. Demonstration of qubit operations below a rigorous fault tolerance threshold with gate set tomography. *Nature Communications*, 8:14485, 2017.
- [8] Easwar Magesan, J. M. Gambetta, and Joseph Emerson. Scalable and robust randomized benchmarking of quantum processes. *Phys. Rev. Lett.*, 106:180504, May 2011.
- [9] Kenneth Rudinger, Craig W. Hogle, Ravi K. Naik, Akel Hashim, Daniel Lobser, David I. Santiago, Matthew D. Grace, Erik Nielsen, Timothy Proctor, Stefan Seritan, Susan M. Clark, Robin Blume-Kohout, Irfan Siddiqi, and Kevin C. Young. Experimental characterization of crosstalk errors with simultaneous gate set tomography. *PRX Quantum*, 2:040338, Nov 2021.
- [10] Erik Nielsen, John King Gamble, Kenneth Rudinger, Travis Scholten, Kevin Young, and Robin Blume-Kohout. Gate set tomography. *Quantum*, 5:557, Oct 2021.
- [11] Seth T. Merkel, Jay M. Gambetta, John A. Smolin, Stefano Poletto, Antonio D. Córcoles, Blake R. Johnson, Colm A. Ryan, and Matthias Steffen. Self-consistent quantum process tomography. *Physical Review A*, 87(6):062119, Jun 2013.
- [12] Trystan Surawy-Stepney, Jonas Kahn, Richard Kueng, and Madalin Guta. Projected least-squares quantum process tomography. *Quantum*, 6:844, oct 2022.
- [13] Jay M. Gambetta, A. D. Córcoles, S. T. Merkel, B. R. Johnson, John A. Smolin, Jerry M. Chow, Colm A. Ryan, Chad Rigetti, Stefano Poletto, Thomas A. Ohki, Mark B. Ketchen, and Matthias Steffen. Characterization of addressability by simultaneous randomized benchmarking. *Physical Review Letters*, 109(24):240504, 2012.
- [14] Ryan Levy, Di Luo, and Bryan K. Clark. Classical shadows for quantum process tomography on near-term quantum computers. *Physical Review Research*, 6(1):013029, 2024.

- [15] David C. McKay, Andrew W. Cross, Christopher J. Wood, and Jay M. Gambetta. Correlated randomized benchmarking. *arXiv preprint arXiv:2003.02354*, 2020.
- [16] J. Helsen, I. Roth, E. Onorati, A.H. Werner, and J. Eisert. General framework for randomized benchmarking. *PRX Quantum*, 3(2):020357, Jun 2022.
- [17] Hsin-Yuan Huang, Sitan Chen, and John Preskill. Learning to predict arbitrary quantum processes. *PRX Quantum*, 4(4):040337, Dec 2023.
- [18] Yuqing Wang, Guoding Liu, Zhenhuan Liu, Yifan Tang, Xiongfeng Ma, and Hao Dai. Robust estimation of nonlinear properties of quantum processes. *Phys. Rev. A*, 110:032415, Sep 2024.
- [19] J. Helsen, M. Ioannou, J. Kitzinger, E. Onorati, A. H. Werner, J. Eisert, and I. Roth. Shadow estimation of gate-set properties from random sequences. *Nature Communications*, 14(1):5039, 2023.
- [20] Jonathan Kunjummen, Minh C. Tran, Daniel Carney, and Jacob M. Taylor. Shadow process tomography of quantum channels. *Phys. Rev. A*, 107:042403, Apr 2023.
- [21] Mohan Sarovar, Timothy Proctor, Kenneth Rudinger, Kevin Young, Erik Nielsen, and Robin Blume-Kohout. Detecting crosstalk errors in quantum information processors. *Quantum*, 4:321, sep 2020.
- [22] Robin Harper, Steven T. Flammia, and Joel J. Wallman. Efficient learning of quantum noise. *Nature Physics*, 16(12):1184–1188, 2020.
- [23] Cambyse Rouzé and Daniel Stilck Franca. Efficient learning of the structure and parameters of local pauli noise channels, 2023.
- [24] Hongye Yu and Tzu-Chieh Wei. Efficient separate quantification of state preparation errors and measurement errors on quantum computers and their mitigation. *Quantum*, 9:1724, 2025.
- [25] Steven T. Flammia and Ryan O’Donnell. Pauli error estimation via Population Recovery. *Quantum*, 5:549, September 2021.
- [26] L. C. G. Govia, G. J. Ribeill, D. Ristè, M. Ware, and H. Krovi. Bootstrapping quantum process tomography via a perturbative ansatz. *Nature Communications*, 11(1):1084, 2020.
- [27] Giacomo Torlai, Christopher J. Wood, Atithi Acharya, Giuseppe Carleo, Juan Carrasquilla, and Leandro Aolita. Quantum process tomography with unsupervised learning and tensor networks. *Nature Communications*, 14(1):2858, 2023.
- [28] Chirag Wadhwa, Laura Lewis, Elham Kashefi, and Mina Doosti. Agnostic process tomography. *arXiv preprint arXiv:2410.11957*, 2024.
- [29] Srinivasan Arunachalam, Arkopal Dutt, Francisco Escudero Gutiérrez, and Carlos Palazuelos. Learning low-degree quantum objects, 2024.
- [30] Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16:1050–1057, 2020.

- [31] Jaakko Malmi, Kimmo Korhonen, Daniel Cavalcanti, and Guillermo García-Pérez. Enhanced observable estimation through classical optimization of informationally overcomplete measurement data: Beyond classical shadows. *Physical Review A*, 109(6):062412, 2024.
- [32] Andrea Caprotti, Joshua Morris, and Borivoje Dakić. Optimizing quantum tomography via shadow inversion. *Phys. Rev. Res.*, 6:033301, Sep 2024.
- [33] Laurin E. Fischer, Timothée Dao, Ivano Tavernelli, and Francesco Tacchino. Dual-frame optimization for informationally complete quantum measurements. *Phys. Rev. A*, 109:062415, Jun 2024.
- [34] Senrui Chen, Wenjun Yu, Pei Zeng, and Steven T. Flammia. Robust shadow estimation. *PRX Quantum*, 2:030348, Sep 2021.
- [35] Kristan Temme, Sergey Bravyi, and Jay M. Gambetta. Error mitigation for short-depth quantum circuits. *Phys. Rev. Lett.*, 119:180509, Nov 2017.
- [36] Hamza Jnane, Jonathan Steinberg, Zhenyu Cai, H. Chau Nguyen, and Bálint Koczor. Quantum error mitigated classical shadows. *PRX Quantum*, 5:010324, Feb 2024.