

Out-of-distribution generalisation for learning quantum channels with low energy coherent states

Jason L. Pereira,^{1,2,3} Quntao Zhuang,^{2,4} and Leonardo Banchi^{5,1,*}

¹*INFN Sezione di Firenze, via G. Sansone 1, Sesto Fiorentino (FI) I-50019, Italy*

²*Ming Hsieh Department of Electrical and Computer Engineering,
University of Southern California, Los Angeles, California 90089, USA*

³*Sorbonne Université, CNRS, LIP6, 4 Place Jussieu, Paris F-75005, France*

⁴*Department of Physics and Astronomy, University of Southern California, Los Angeles, California 90089, USA*

⁵*Department of Physics and Astronomy, University of Florence,
via G. Sansone 1, Sesto Fiorentino (FI) I-50019, Italy*

(Dated: July 8, 2026)

In our work [1], we present a general formalism for understanding out-of-distribution generalisation in Continuous Variable (CV) channels. We show that if any pair of CV channels has sufficiently close outputs for low energy coherent state inputs, it is possible to construct a bound on the output distance for any input state. Hence, we prove that out-of-distribution generalisation is always possible for learning quantum channels using low energy coherent states.

Machine learning models are typically trained and evaluated under the assumption that training and test data are drawn from the same underlying distribution. In practice, however, this assumption is rarely satisfied. Real-world systems are routinely deployed in environments that differ from those observed during training, where changes in operating conditions, measurement procedures, or the nature of the inputs can significantly degrade performance. The challenge of out-of-distribution (OOD) generalization is therefore to infer rules that remain valid beyond the support of the training data, rather than merely interpolating within it. This problem has attracted considerable attention across machine learning because it lies at the heart of building reliable, robust, and scientifically meaningful predictive models. At the same time, OOD generalization is fundamentally impossible without additional assumptions about the underlying data-generating process, making it essential to identify the structural properties that permit reliable extrapolation.

In quantum information processing, OOD generalization arises naturally in the characterization of quantum devices, hence when processing “quantum data”. Experimental constraints often limit the set of probe states that can be prepared, while one ultimately seeks to predict the behavior of the device on a much broader class of inputs. A particularly important example is learning quantum channels from low-energy coherent states, where accurate predictions are required for arbitrary, potentially high-energy or highly non-classical states. Our work [1] studies the conditions under which such extrapolation is possible and establishes rigorous guarantees relating performance on experimentally accessible training data to performance over the full input space. Our results demonstrate that, despite the apparent mismatch between training and deployment distributions, meaningful out-of-distribution generalization can be achieved under physically motivated assumptions, providing a theoretical foundation for experimentally efficient quantum process learning.

Our OOD generalization framework is constructed in three stages [1], as also illustrated in Fig. 1.

1. Given N samples of low energy coherent states sent through a channel Ψ , bound the error in learning the action of this channel on the same distribution of low energy coherent states. As sketched in Fig. 1 (c), we have N samples of the form $\Psi[|re^{i\phi}\rangle\langle re^{i\phi}|_{\text{coh}}]$, for $r \leq \tau$. Find a protocol that uses them to learn a channel Φ , such that (with high probability) $\|(\Psi - \Phi)[|re^{i\phi}\rangle\langle re^{i\phi}|_{\text{coh}}]\| \leq \epsilon_0$ for all $r \leq \tau$, where ϵ_0 is a decreasing function of N (so that more samples let us learn the channel better). In other words, we bound the “in-distribution” error for learning the action of channel Ψ on coherent states with an average photon number of $\leq \tau^2$. Note that the worst case error here follows $\epsilon_0 = \mathcal{O}(N^{-1/2})$ [2]. Decreasing ϵ_0 means using more samples to find a new channel Φ , such that the output distance (for our subset of coherent states) is smaller. We require that $\epsilon_0 \rightarrow 0$ as $N \rightarrow \infty$.
2. Assume stage 1 was successful. In stage 2, depicted in Fig. 1 (d), we aim to solve the restricted case of OOD generalisation in which we only consider coherent states beyond the energy of the actual probe states from stage 1. More precisely, we need to find a function $\epsilon(\epsilon_0, r^2)$ such that $\|(\Psi - \Phi)[|re^{i\phi}\rangle\langle re^{i\phi}|_{\text{coh}}]\| \leq \epsilon(\epsilon_0, r^2)$ for all r . We also require that $\epsilon(\epsilon_0, r^2)$ is a decreasing function of ϵ_0 that reaches 0 for $\epsilon_0 = 0$, ensuring that our OOD error reaches 0 if our in-distribution error does.

* leonardo.banchi@unifi.it




FIG. 1. Parts (a) and (b) show two different ways we could learn a channel, Φ , that mimics the action of a target process, Ψ , for low energy coherent states. (a) is “classical” learning of a quantum channel, whilst (b) is a form of quantum machine learning. In this illustration, we depict probing the optical properties of an unknown (green) sample, although the process could be a magnetic field, a non-linear medium, a reflective cavity, or any other transformation that could be applied to a state of light. We search through a set of possible (yellow) optical mediums to find the one that best imitates the green substance. In (a), we use a (finite energy) laser to send low energy coherent states through the target and characterise the outputs with measurements. The results are post-processed (e.g., by feeding them into a classical computer, per the diagram) to obtain classical knowledge of the input-output relations. We can then “simulate” the action of Ψ on an unknown state, using our classical computer and the transformation Φ . In (b), we have a tuneable quantum device (such as a different substance with tuneable optical properties, a tuneable magnetic field, or some parametrised optical circuit) that we want to use to imitate the target process. We probe both processes with the same coherent states (we show two lasers, but a real implementation could use a single laser and a balanced beamsplitter, to ensure the probing states are identical). Instead of characterising both outputs, we measure them jointly, to determine their trace distance. We could then use the measurement results to tune the quantum device and so reduce the distance. Once Φ is tuned, we hope that an unknown state sent through the quantum device will be transformed similarly to if it were sent through Ψ . Parts (c-e) illustrate the three types of input for which we can bound the closeness of two channels. (c) corresponds to the input states for parts (a) and (b); the probes used in our physical measurements are low energy coherent states. The dotted line shows the maximum average energy of the inputs. We can construct a bound on the output distance between our target and learned processes, Ψ and Φ , based on our measurement results. In (d), we consider higher energy coherent state inputs. These are outside of the class of inputs for which we have actual measurement results; rather, we want to be able to trust that our simulation, Φ , of Ψ is still accurate for these inputs (in the scenario of part (a)) or that our tuned quantum device still mimics our target well (in the scenario of part (b)). We use our bound for the inputs shown in (c) to construct a bound on the inputs shown in (d). The scaling of this bound depends on the class of channels to which the target and learned channel belong. In (e), we consider other types of input states, including non-classical states such as squeezed vacuums and Fock states. We again want to trust that Φ is a good simulation of Ψ for these types of input, and so we extend our bounds for the inputs in (d) to the inputs in (e).

3. Assume stage 2 was successful. In the final stage, depicted in Fig. 1 (e), we generalise to arbitrary states: for an arbitrary input ρ , we need to bound $\|(\Psi - \Phi)[\rho]\|$ in terms of $\bar{n}(\rho)$, $\epsilon(\epsilon_0, r)$, and possibly $\mathcal{N}(\rho)$ (if finite). This gives us a complete bound for OOD generalisation in the general case, and we require that it reaches 0 for $\epsilon = 0$.

Following such procedure, our work [1] provides two main formal results.

Theorem (2 in [1]). *Suppose $\|(\Psi - \Phi)[|re^{i\phi}\rangle\langle re^{i\phi}|_{\text{coh}}]\| \leq \epsilon_0$ for all ϕ and all $r \leq \tau$, for some energy threshold $\tau > 0$. Then, there exists a concave (in r^2) function $\epsilon(\epsilon_0, r^2)$ such that $\|(\Psi - \Phi)[|re^{i\phi}\rangle\langle re^{i\phi}|_{\text{coh}}]\| \leq \epsilon(\epsilon_0, r^2)$ for all ϕ and r , with the following properties:*

1. For any r and for $\epsilon'_0 < \epsilon_0$, $\epsilon(\epsilon'_0, r^2) \leq \epsilon(\epsilon_0, r^2)$.

2. For any r , $\epsilon(\epsilon_0, r^2) \rightarrow 0$ as $\epsilon_0 \rightarrow 0$.

Theorem (3 + corollary 1 in [1]). Let Ψ and Φ be a pair of quantum channels for which $\|(\Psi - \Phi)[|re^{i\phi}\rangle\langle re^{i\phi}|_{\text{coh}}]\| \leq \epsilon(\epsilon_0, r^2)$ for all r and for some concave (in r^2) function ϵ . Let ρ be any input state with average photon number $\bar{n}$. For any $\kappa > 1$ and positive integer M ,

$$\|(\Psi - \Phi)[\rho]\| \leq \mathcal{O}[e^{M \log(M\kappa)} \kappa^{-1} \epsilon(\epsilon_0, \kappa^{-1})] + \mathcal{O}[\bar{n}^{\frac{1}{2}} M^{-\frac{1}{2}}] \quad (1)$$

This formal result allows us to make a few observations. Firstly, if we fix κ , then as ϵ_0 approaches 0, $\epsilon(\epsilon_0, \kappa^{-1})$ will also approach zero, and we will be able to send $M \rightarrow \infty$, so that Eq. (1) approaches 0 too. Hence, as long as we are able to keep decreasing ϵ_0 , by taking more samples, we are also able to achieve an arbitrarily small output trace distance for any input state. Secondly, fixing $M \sim \bar{n}$, we see that the leading order term scales at most polynomially in the input energy ($\sim \mathcal{O}[e^{\bar{n} \log(\bar{n})} \epsilon(\epsilon_0, \kappa^{-1})]$). A linear increase in the allowed $\bar{n}$ (for fixed error) can therefore always be achieved with an (at most) exponential decrease in $\epsilon(\epsilon_0, \kappa^{-1})$.

These results show that *out-of-distribution generalisation is possible for every input state and all target channels*, in the sense that a sufficiently small in-distribution error bound can always be translated to an out-of-distribution error bound. It is therefore possible to achieve an arbitrarily small error in reproducing the input-output relations for any input state, as long as sufficient samples are taken from the set of low energy coherent states.

The full proof of those theorems and details on how to construct the function $\epsilon(\epsilon_0, r^2)$ can be found in [1]. Here, on the other hand, we briefly comment on the physical original of such possibility, as we believe that it may span further developments in the theory of quantum machine learning.

Indeed, at first sight, it may seem surprising that information obtained from low-energy coherent states can determine the action of a quantum channel on arbitrarily energetic or highly non-classical states. The possibility of out-of-distribution generalization does not arise because quantum channels are linear, but because the chosen training distribution is informationally complete. Although the probe states are restricted to bounded energy, they do not span only a low-energy subspace. Instead, bounded-energy coherent states form an overcomplete family whose linear span is the entire Hilbert space [3]. Consequently, there is no hidden orthogonal subspace on which the channel could behave arbitrarily. High-energy states are therefore not genuinely outside the domain explored by the training data; they are encoded, albeit implicitly, in the interference structure of the coherent-state family.

-
- [1] J. L. Pereira, Q. Zhuang, and L. Banchi, Out-of-distribution generalization for learning quantum channels with low-energy coherent states, PRX Quantum **6**, 040306 (2025).
 - [2] L. Banchi, J. L. Pereira, S. T. Jose, and O. Simeone, Statistical Complexity of Quantum Learning, Adv. Quantum Technol. **n/a**, 2300311 (2024).
 - [3] J. Janszky, P. Domokos, and P. Adam, Coherent states on a circle and quantum interference, Physical Review A **48**, 2213 (1993).