

PAPER • OPEN ACCESS

Visual assessment of multi-photon interference

To cite this article: Fulvio Flamini *et al* 2019 *Quantum Sci. Technol.* **4** 024008

View the [article online](#) for updates and enhancements.

You may also like

- [Signatures of many-particle interference](#)
Mattia Walschaers

- [Experimental generalized quantum suppression law in Sylvester interferometers](#)

Niko Viggianiello, Fulvio Flamini, Luca Innocenti et al.

- [Validating multi-photon quantum interference with finite data](#)

Fulvio Flamini, Mattia Walschaers, Nicolò Spagnolo et al.



Easy-to-use and Helium-3 free
cryogenics solutions

LEARN MORE

Quantum Science and Technology



PAPER

Visual assessment of multi-photon interference

OPEN ACCESS

RECEIVED
18 September 2018

REVISED
30 January 2019

ACCEPTED FOR PUBLICATION
7 February 2019

PUBLISHED
5 March 2019

Original content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Fulvio Flamini , Nicolò Spagnolo and Fabio Sciarrino

Dipartimento di Fisica, Sapienza Università di Roma, P.le Aldo Moro 5, I-00185 Roma, Italy

E-mail: fabio.sciarrino@uniroma1.it

Keywords: machine learning, boson sampling, quantum information, validation, multi-photon interference

Supplementary material for this article is available [online](#)

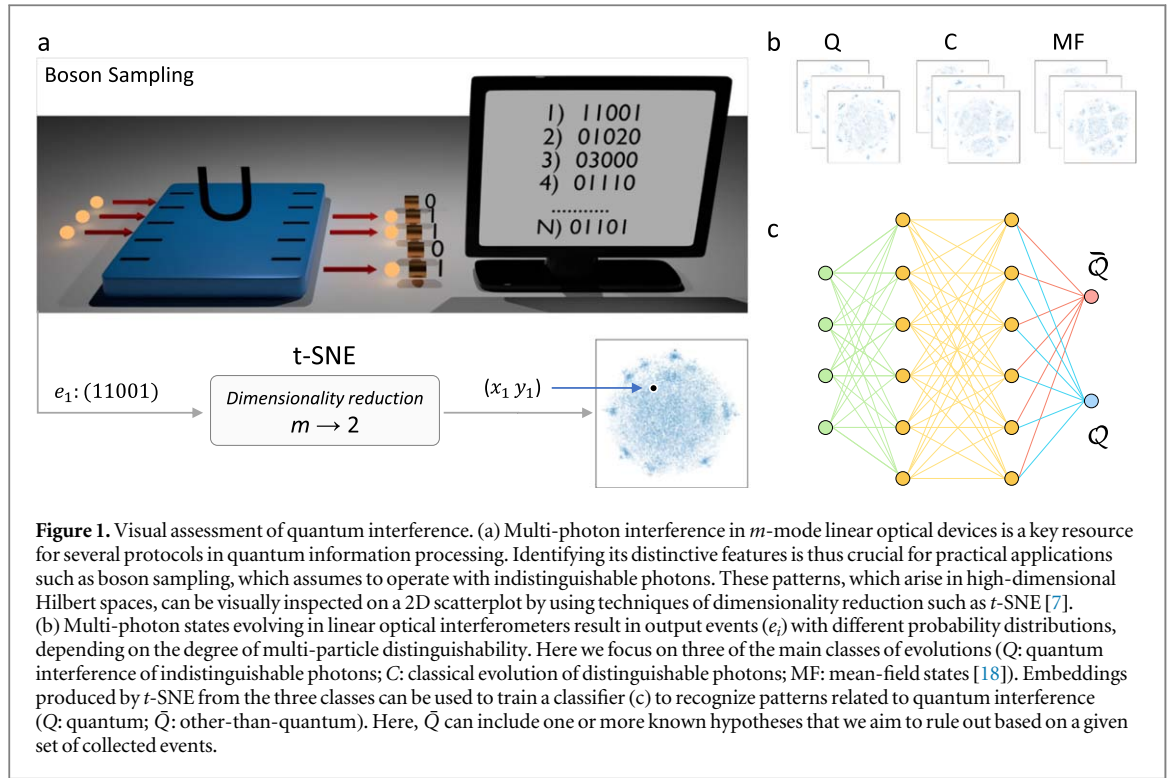
Abstract

Classical machine learning algorithms can provide insights on high-dimensional processes that are hardly accessible with conventional approaches. As a notable example, t -distributed Stochastic Neighbor Embedding (t -SNE) represents the state of the art for visualization of data sets of large dimensionality. An interesting question is then if this algorithm can provide useful information also in quantum experiments with very large Hilbert spaces. Leveraging these considerations, in this work we apply t -SNE to probe the spatial distribution of n -photon events in m -dimensional Hilbert spaces, showing that its findings can be beneficial for validating genuine quantum interference in boson sampling experiments. In particular, we find that nonlinear dimensionality reduction is capable to capture distinctive features in the spatial distribution of data related to multi-photon states with different evolutions. We envisage that this approach will inspire further theoretical investigations, for instance for a reliable assessment of quantum computational advantage.

Introduction

Motivated by the recent achievements enabled by machine learning, the last decades have seen a flowering of novel approaches developed to tackle hard problems. The strength of these techniques relies on their capability to address diverse and complex processes over very large data sets, for which a thorough understanding may be not available or accessible. The price to pay, a possible obscure interpretation of their findings [1], is often largely compensated for by the accuracy and simplicity of their predictions. Machine learning algorithms are conveniently divided in three classes [2]: (i) supervised, when they are trained on labeled examples, (ii) unsupervised, when they infer patterns in data with no labels, and (iii) reinforcement learning, when training aims at maximizing a reward. Beyond this classification, other statistical tools are employed to prepare [3] or to explore [4] the input data, usually spread in very large domains, or to test a statistical model [5]. Within these techniques, visualizing high-dimensional points would naturally represent a critical advantage for data analysis. To this aim, several manifold algorithms have been developed to reduce the dimensionality of the problem while preserving interesting features of the data [6]. Ultimately, t -distributed Stochastic Neighbor Embedding algorithm (t -SNE) [7] has now established itself as the new state of the art for this task in several fields [8–12].

Inspired by its successful applications, a natural question is whether manifold algorithms can provide useful solutions also for problems in quantum mechanics. Indeed, the observation that the size of Hilbert spaces scales exponentially fast in the number of qubits already encourages to look in this direction. An interesting test-bed is offered, for instance, by applications that need to probe the spatial distribution of systems living in large Hilbert spaces. A relevant example in this sense is represented by the problem of discerning multi-photon quantum interference in high-dimensional unitary evolutions [13]. This task has recently drawn attention for the validation of boson sampling [14], in order to assess the correct operation of single-photon sources and quantum devices. In a linear optical framework, boson sampling consists in sampling from the output probability distribution of n indistinguishable photons evolving in a m -mode interferometer. Assuming highly plausible conjectures, the problem was proved to be intractable for classical computers for a sufficiently large number of photons, paving the way for a first demonstration of quantum computational advantage [13, 15].



In this work we discuss a positive answer to this question, namely whether machine learning algorithms for dimensionality reduction can provide an effective approach also in quantum mechanics. We address this task by using t -SNE to visualize the result of different dynamics for states with n distinguishable and indistinguishable photons. Leveraging the approach of [16, 17], we consider their spatial distribution in m -dimensional Hilbert spaces and apply t -SNE to take a snapshot of these points (figure 1). As we will show, these images present interesting properties: (i) for a given unitary evolution and fixed degree of distinguishability, different sets of events yield similar images; (ii) data with the same (different) degree of distinguishability converge to images with similar (different) features; (iii) the number of events to sample to produce a reliable image appears to scale favorably with the size of the problem. We investigate this approach by numerically generating samples of increasing dimension, providing evidence for its efficacy in assessing multi-photon interference.

The manuscript is structured in two parts: first, after a short introduction on boson sampling and its validation, we introduce t -SNE as a promising tool to study multi-particle interference. Then, in the second part we discuss a possible improvement that takes advantage of convolutional neural networks (CNN) to automate and boost the process. We believe our analysis might inspire further theoretical researches, which could exploit t -SNE and related techniques thoroughly also in other scenarios.

Multi-photon interference and its assessment

Multi-photon interference occurring in a scattering process through a linear optical network lies at the basis of several quantum computational and simulation models. One of the reasons behind its relevance is given by the very complexity of its dynamics, which is related to the evaluation of complex physical quantities. This aspect has been recently clarified by the formulation of the boson sampling problem [14]. The latter is a computational task that corresponds to sampling from the output probability distribution obtained after the evolution of n indistinguishable bosons through a m -mode linear optical network, described by a unitary matrix U randomly drawn according to the Haar measure. More specifically, in this scenario the system configurations are given by ordered lists of m integer numbers (n_1, n_2, \dots, n_m) , where n_k is the number of particles occupying mode k and $\sum_k n_k = n$. Given n indistinguishable bosons, the transition amplitude \mathcal{A} from the input state configuration $S = (s_1, s_2, \dots, s_m)$ to the output $T = (t_1, t_2, \dots, t_m)$ is obtained as $\mathcal{A}(S \rightarrow T) = c_{S,T} \text{per}(U_{S,T})$. Here, $c_{S,T} = (\prod_i^m s_i! \prod_j^m t_j!)^{-1/2}$ is a prefactor taking into account multiple-occupancy degeneracies. The relevant term in the amplitude is $\text{per}(U_{S,T})$, that is, the permanent of the $n \times n$ submatrix of U retrieved by selecting rows and columns according to the occupied modes specified by S and T . Sampling from the output distribution of this system is classically hard, being the evaluation of matrix permanents with normally-distributed complex entries a $\#P$ -hard problem. Boson sampling thus represents a benchmark for the complexity of multi-photon

interference, and provides a well-established candidate for the first experimental observation of quantum computational advantage. These features have led to several experimental implementations [19–34], starting from small instances to progressively larger-size demonstrations.

Strongly related to its computational complexity, a crucial open problem is the challenge of discerning patterns in output data produced by quantum interference of indistinguishable photons from other scenarios, such as with distinguishable particles or other classically-simulatable models. Indeed, since the evaluation of a single input-output transition amplitude relies on the calculation of a matrix permanent, other non-trivial computationally efficient approaches have to be developed and tested [16–18, 24, 25, 27, 29, 30, 33, 35–49].

Furthermore, the space of possible configurations for n particle in m modes increases as $\binom{m+n-1}{n}$, thus leading to very sparse samples when the size of the system is large. Hence, no straightforward methods can be applied to find patterns hidden in such high-dimensional data. The identification of these patterns can be a helpful resource to assess the presence of genuine multi-photon interference.

t -SNE for validation of quantum interference

t -SNE is a technique for nonlinear dimensionality reduction, suitable to embed points from high-dimensional spaces into a plane for visualization [7]. The algorithm finds a high-to-low-dimensional map with the property that proximity, as measured with a given metric, is preserved by the map with high probability. More specifically, t -SNE consists in two stages. In the first stage it derives a probability distribution P , whose elements $p^{(i,j)}$ depend on the distance between points (i, j) in the high-dimensional space. In the second stage, t -SNE looks for a second distribution P' (the embedding) in the low-dimensional space that best represents P . The choice of the optimal P' is made by minimizing the Kullback–Leibler divergence [48] between P and P' with respect to the positions of the N points. The algorithm is efficient, in that the computational complexity scales as $O(N^2)$. The required resources can be further reduced with various approaches, most notably in the Barnes–Hut-SNE implementation [49], for which various open-source libraries are already available [51]. Very recently, a novel approach has been reported with linear computational complexity, enabling a quasi-real-time version for the browser that leverages the graphic processing unit [50].

In order to apply t -SNE to study multi-particle quantum interference, we need to construct a suitable space where each output event can take place. To this aim, we choose to work in a space where each axis is associated to a different output mode [16, 17] (figure 1(a)). For each point, the coordinate on the k -th axis is given by the occupation number on mode k of the corresponding n -photon state. For instance, for $n = 4$ photons in $m = 7$ modes we could have the output event $(0, 1, 0, 0, 2, 1, 0)$, which corresponds to the point with coordinates $(1, 2, 1)$ on, respectively, the 2nd, 5th and 6th axis and 0 on the others. For a given set of n -photon events over m modes, t -SNE can then output a 2D embedding that aims to preserve the proximity between points. An example of this analysis with Euclidean distance is shown in figure 2, in the case of $n = 5$ indistinguishable photons injected in interferometers with $m = 25$ modes. Data for this test were classically simulated by employing the exact sampling algorithm by Clifford and Clifford [51] for indistinguishable photons. Here we considered four Haar-random unitary evolutions (figures 2(a)–(d)) and, for each unitary, we constructed the t -SNE projection of four data sets with $N = 10^4$ events each. One can recognize in the panels of figure 2 the above-mentioned property (i): when we fix the transformation and the degree of distinguishability, t -SNE yields similar embeddings for different sets of events despite the very large Hilbert space dimensions. Furthermore, patterns in these embeddings slightly change from unitary to unitary, indicating that t -SNE is indeed capturing features of multi-photon quantum interference in a consistent way.

One interesting aspect to investigate, then, is the dependence of these patterns on the degree of distinguishability, which affects the nature of the dynamics at a deeper level. To show that t -SNE embedding is effective to separate quantum and non-quantum sampling we require that it preserves the capability to discriminate between different types of interference. We provide evidence for this property in appendix A based on numerical simulations. For a first analysis along this direction, we applied t -SNE to data related to specific unitary evolutions: the quantum Fourier transform and the Hadamard transform [52]. The choice of these transformations is motivated by the fact that, for specific input states called ‘cyclic’, they are known to exhibit marked differences in the output probability distributions between the cases of fully indistinguishable and fully distinguishable input states [18, 38, 44]. Specifically, in the case of a fixed cyclic input with all indistinguishable photons, we observe the suppression of a large number of output configurations, which in turn implies that most events are concentrated in a subset of all possible output configurations. Conversely in the second case, when the input cyclic state consists of distinguishable photons, the output distribution is fully flat. The analysis was extended also to a third class of input states with different degree of distinguishability: besides states where all photons are fully indistinguishable, exhibiting quantum interference (Q: blue), fully distinguishable,

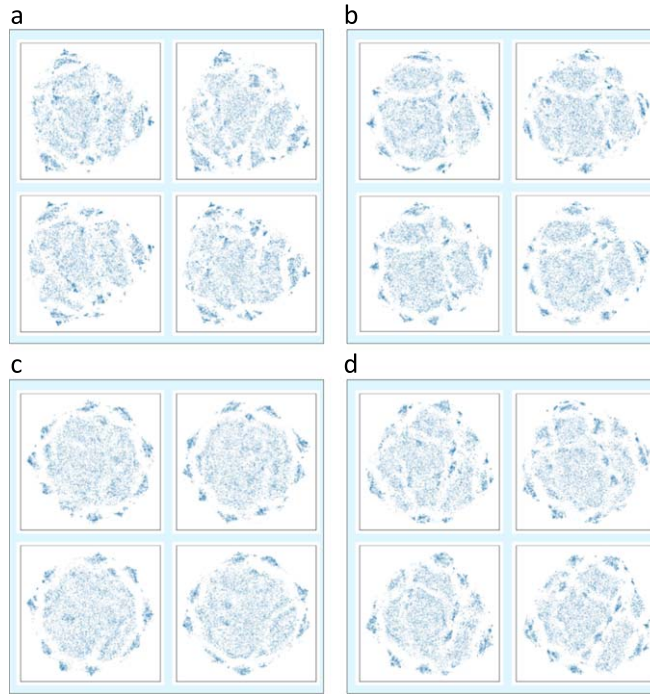
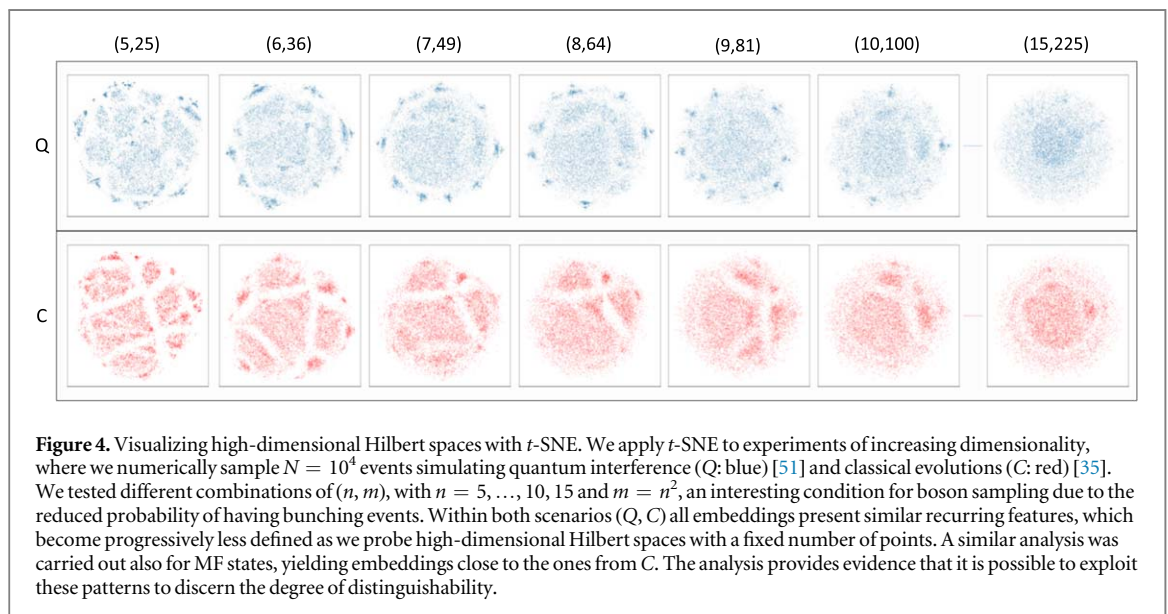
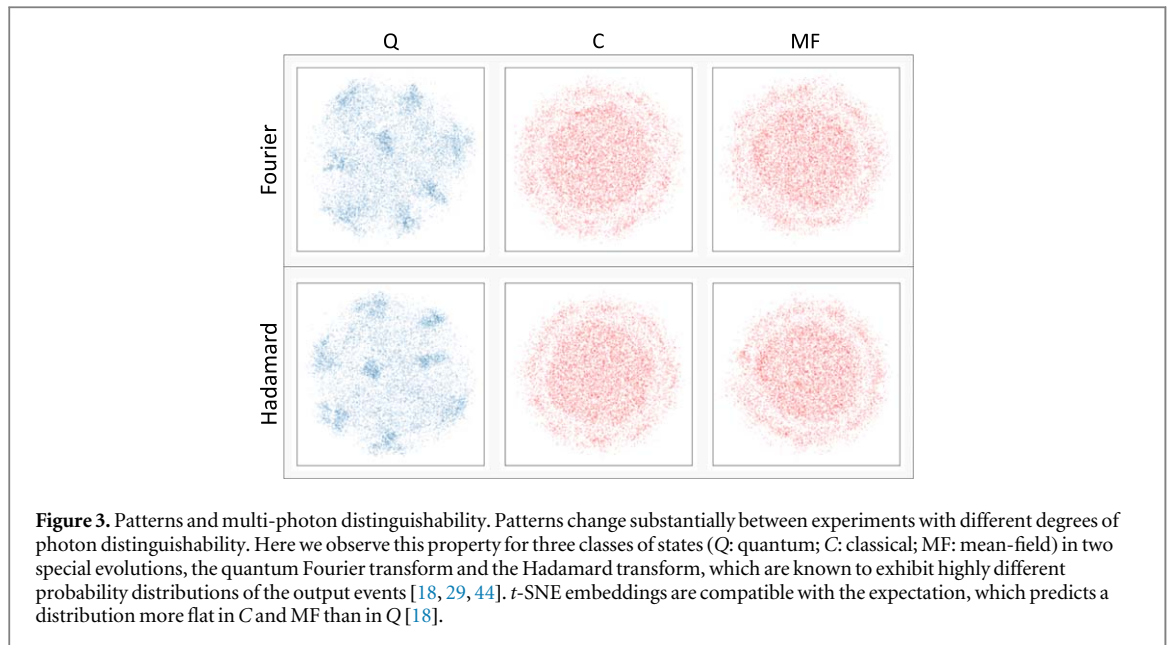


Figure 2. Reproducibility of t -SNE patterns. For all experiments we simulated, t -SNE embeddings show patterns that are reproduced in a clear way when we employ different output data sets from the same unitary evolution. Here, four panels (a)–(d) are associated to as many unitary evolutions with $n = 5$ indistinguishable photons in $m = 25$ -mode interferometers. For each unitary, four different sets of 10^4 events are projected via t -SNE to show reproducibility. See the supplemental material available online at stacks.iop.org/QST/4/024008/mmedia (section *Repeatability in t-SNE patterns*) for a comment on how the random initialization of the gradient descent affects the reproducibility.

corresponding to classical evolutions (C: red), we simulated also the so-called mean-field state (MF: red). The latter is a special state that has been studied in the context of boson sampling validation [18], in which photons are injected—one by one—in superposition over a set of n modes with random phases. The interest for these states is due to the observation that they reproduce some signatures of multi-photon interference, even though each photon evolves independently without interfering with the others. Events for the other two states have been numerically generated with two exact sampling algorithms: one by Aaronson and Arkhipov [35] for distinguishable photons, and the one by Clifford and Clifford [51] for indistinguishable photons. From these considerations, considering the different distribution of events in the high-dimensional space, we intuitively expected to find a trace of these patterns also in the 2D projection with t -SNE. Results for this analysis are shown in figure 3 for $n = 8$ photons in $m = 64$ -dimensional quantum Fourier and Hadamard transforms. The feature we observe is the one we anticipated in property (ii) data sets related to experiments with (Q) or without (C or MF) quantum interference yields embeddings that exhibit recurring features, which are different between the two scenarios.

So far, we have restricted the analysis in figure 3 to structured evolutions that, as such, might not reflect a general feature of multi-photon dynamics as seen through the lens of t -SNE. We address this question with a further analysis, reported in figure 4, discussing how the property (ii) arises in Haar-random unitary transformations with no special structure. Figure 4 reports seven images for increasing sizes of the problem ($n, m = n^2$) with indistinguishable photons, distinguishable photons and mean-field states, providing evidence that this behaviour is indeed not limited to low dimensions or to transformations with high symmetries. This observation suggests that the output distributions corresponding to different particle statistics present general dissimilarities that are sufficiently strong to exhibit different behaviours already with a relatively small number of sampled events (here $N = 10^4$). For instance, for $n = 15$ photons in $m = 225$ modes (rightmost images in figure 4) such a small sample size would be only a negligible fraction of the total number of configurations in the Hilbert space, namely $\sim 10^{-20}$. In the supplemental material (section ‘*Interference and its distinctive t-SNE patterns*’) we provide a more detailed description on the shape of these patterns for Haar-random transformations. The aspect of how many samples are required for a reliable assessment (iii) has been explored for the dimensions considered in these analyses, and ranges from $N \sim 2 \times 10^3$ to $N \sim 10^4$ output events, which is in line with the state of the art of validation protocols.



Ultimately, all these results encourage to consider *t*-SNE as an effective tool to study multi-particle scattering processes. To stress that its applicability can also extend beyond validation of quantum interference, in the supplemental material we have included further analyses on patterns that can be found in *t*-SNE embeddings (section ‘*t*-SNE beyond validation of quantum interference’). With the above considerations in mind, we are now ready to move further and try to exploit these features for a self-consistent off-the-shelf approach.

From *t*-SNE to classification

Embeddings produced via *t*-SNE can be inspected to discern patterns related to different multi-photon interference effects (figure 4). In particular, when we are to choose between only two hypotheses (e.g. quantum or classical evolutions) the task reduces to a binary classification problem, for which several algorithms are available [53]. In these cases, we want to train a suitable classifier to recognize characteristic patterns in data with different distinguishability. A sketch of the training stage would then proceed as follows. (a) First, we need to generate N_s labeled data sets, with N events each, for each of the two hypotheses: quantum (Q) or other-than-quantum (\bar{Q}) interference. A key observation here is that the classifier does not look for patterns within a subset

Table 1. Accuracy of a CNN in validating quantum interference for different numbers of photons n and modes m , trained with 2×10^3 images using 5×10^3 samples. Details on the simulation are reported in the supplemental material (section ‘Details on validation via CNNs’).

m	$n = 6$	$n = 8$	$n = 10$
36	0.95 (3)	1.00 (0)	0.94 (3)
64	0.99 (1)	1.00 (0)	1.00 (0)
100	0.98 (1)	1.00 (0)	0.99 (1)

of N events in each of the N_s sets, since single events are in general not indicative of the degree of interference. Rather, a classifier should explore both collections by looking at summary features of each set, in order to learn which characteristic best signals one scenario or the other. (b) Different choices could be made for the summary feature of N events. In [41, 47], for instance, the authors exploited statistical properties of the two-mode correlators to discriminate alternative evolutions. In the following, instead, we address this task by converting all N_s sets in as many 2D images by means of t -SNE, which ultimately will be fed into the classifier. As for the algorithm we choose to train deep neural networks, since they are known to be highly effective for computer vision [54]. More specifically, we employed CNN [55] built using the high-level Keras libraries [56] with TensorFlow backend, the open-source software library for numerical computation [57].

The analysis was carried out for various numbers of photons $n = 6, 8, 10$ and modes $m = 36, 64, 100$, generating $N_s = 2 \times 10^3$ data sets of $N = 5 \times 10^3$ output events for quantum and classical boson sampling (distinguishable photons and mean-field). Each such image takes approximately 25 s on a standard laptop, averaging over all pairs of (n, m) under investigation. The number of samples employed by the present protocol is comparable with other state-of-the-art validation protocols, such as in [16, 17, 41, 47]. For each combination of (n, m) , we split all sets in three subsets: 75% of all images is used to train the CNN (training set), 20% is used by the CNN for model selection (validation set) and 5% is used by the experimenter to test the prediction of trained CNNs on never-seen-before images, to exclude a bias due to training. A high accuracy on separate data sets suggests that the CNNs are properly trained, since they generalize the classification rules to new data. This choice of the split is typical in machine learning [53]. For the scope of this analysis, images related to distinguishable photons and mean-field states are treated as one class, from which the CNN is asked to distinguish images related to quantum interference. Results on the application of CNNs to validate boson sampling are reported in table 1. Trained CNNs prove effective in discriminating the two classes of images, achieving a high score in all combinations of (n, m) . Specifically, they show a higher accuracy whenever $m \sim n^2$ or $m < n^2$, which can be explained by the fact that in this regime the fraction of bunching effects (whose probability also depends on the type of distinguishability) is not negligible. Indeed, in these cases the two patterns are sufficiently different to be identified by simple human inspection (see for instance figure 4). Interestingly, the CNN is capable to witness distinctive patterns also in the majority of the hardest cases ($n, m > n^2$), where patterns tend to become much more ambiguous. The results in table 1 suggest also that the scaling of the number of samples is favorable with (n, m) , since the protocol is effective using a number of samples that ranges from 2×10^3 to 10^4 . While a thorough investigation on its potential goes naturally beyond the scope of the present work, we expect CNNs to be successful also for larger sizes of the problem when trained on more images and for a finer optimization of the hyper-parameters. Clearly, a stronger training requires also a fine tuning of the t -SNE hyper-parameters [58] and of the specific rendering of the plots. Notwithstanding, in some of the hardest instances our simple CNN was successful even though a trained experimenter could not identify any distinctive pattern to make a meaningful guess. Besides representing an effective way to automate the validation protocol after t -SNE, this feature suggests that neural networks can provide a useful tool for assessing quantum interference in boson sampling experiments.

Discussion

The task of validating multi-particle quantum interference is progressively drawing more attention, from a fundamental perspective as well as for a reliable assessment of quantum technologies. In this work we tackle this challenge by leveraging machine learning techniques, with an approach that naturally fits in photonic applications. Our strategy consists in highlighting hidden patterns within the statistics of multi-photon collected events, which turn out to be representative of the degree of interference that generated them. Such patterns are projected from high-dimensional Hilbert spaces onto 2D embeddings by means of t -SNE, which allows the

experimenter to directly inspect them. The final validation stage, where patterns are scrutinized to decide whether or not they could be due to quantum interference, can be conveniently recast in a problem of classification for which several other machine learning techniques are available. As an example, we trained a convolutional neural network to automate and boost this process, thus also making it independent of possible human biases. Our numerical investigations demonstrate that this approach is effective for a broad range of experiments up to a large number of photons and modes, which amply covers mid-term applications. Interesting directions opened by this work include the development of a theoretical description of the above-mentioned patterns, which could pave the way for the design of novel validation protocols.

Acknowledgments

This work was supported by the ERC-Advanced Grant CAPABLE (Composite integrated photonic platform by femtosecond laser micromachining; grant agreement no. 742745) and by QUANTERA HiPhoP (High dimensional quantum Photonic Platform; grant agreement no. 731473).

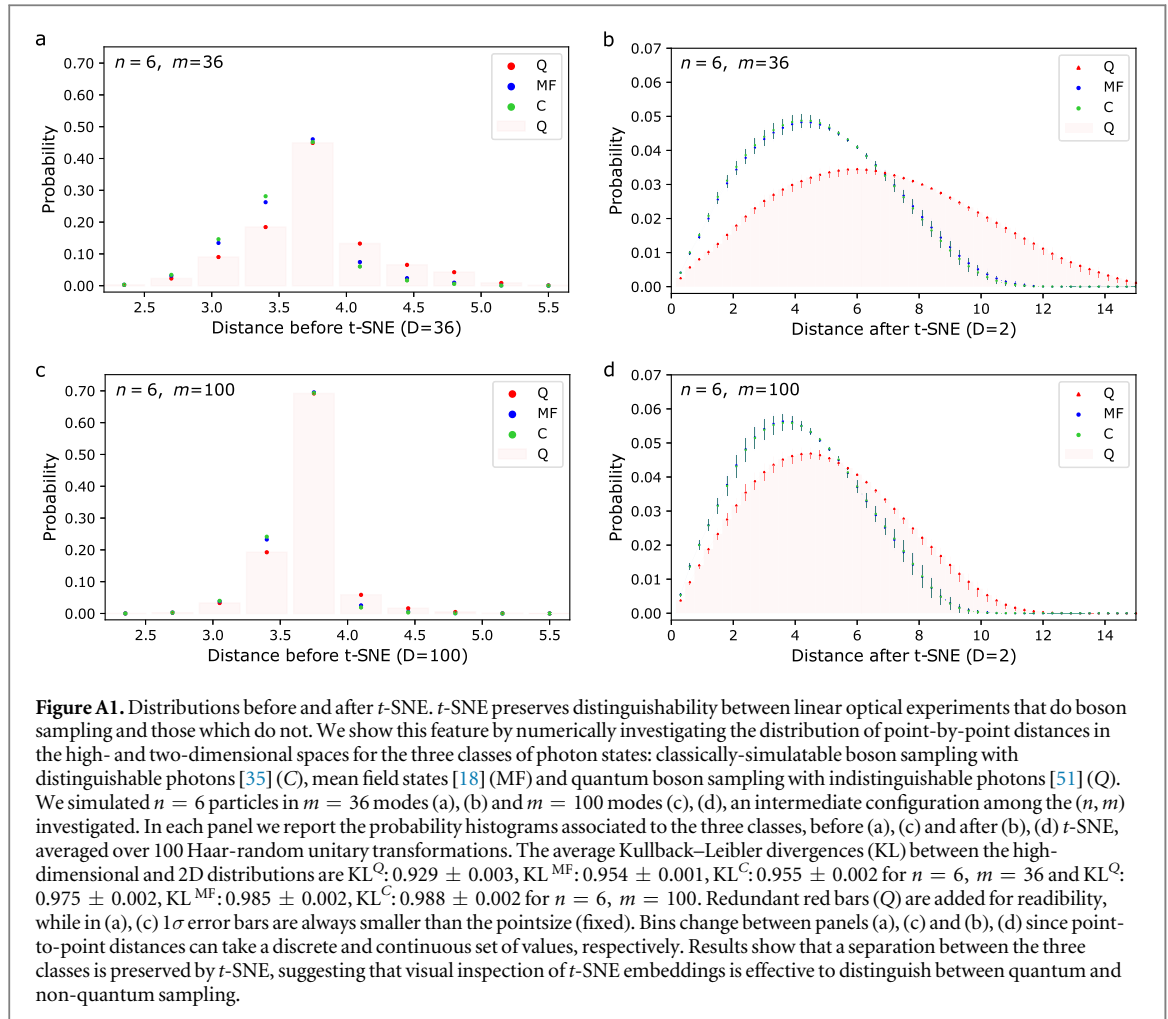
Appendix A. Feasibility of t -SNE for validation

To show that t -SNE is a well performing embedding algorithm for our purpose, namely to distinguish quantum and non-quantum sampling, we need evidence that it preserves the capability to discriminate between classes of multi-photon interference. To search for this evidence, we study the differences between the data in the high-dimensional space (before t -SNE) and their 2D embeddings (after t -SNE). To this end, we focus on the set of point-to-point distances \mathcal{D} , which contains information that is roughly preserved in the projection and that, hence, should be characteristic of each class. The reason for this choice is that the set of pairwise distances (similarities) is related to the conditional probability distribution $p_{ij} = \frac{1}{2}(p_{i|j} + p_{j|i})$ constructed by t -SNE over all pairs of points, where [7]

$$p_{j|i} = \frac{\exp\left(-\frac{\|x_j - x_i\|^2}{2\sigma_i^2}\right)}{\sum_{k \neq i} \exp\left(-\frac{\|x_k - x_i\|^2}{2\sigma_i^2}\right)}. \quad (\text{A.1})$$

The intuition is that, if this distribution has different properties for each class and this difference survives after t -SNE, then this would be a good signal that t -SNE embeddings do not distort the capability to discriminate them.

To investigate this aspect, we numerically simulate 100 experiments for each of the three classes (classical interference C, quantum interference Q and mean-field MF) and we construct a histogram of \mathcal{D} before and after t -SNE. The analyses shown in figure A1 demonstrate how the type of interference is reflected in \mathcal{D} and, correspondingly, in the related $p_{i,j}$ within t -SNE. Interestingly, the behaviour of C and MF is similar and, at the same time, different from Q (in agreement with what already found in the t -SNE embeddings), and MF is always intermediate between C and Q, as expected [18]. These results support the adoption of t -SNE for validation of quantum interference.



Appendix B. Validation via CNN for $n = 4$

To investigate the efficacy of CNNs for validating multi-particle interference based on t -SNE embeddings, we have simulated several experiments for $n = 4$ and various m . The choice of this specific scenario is due to the possibility to observe a transition from high accuracy to lower accuracy at lower m . For other (n, m) , in fact, the behaviour would be more flat in the regime $m \leq 100$, where t -SNE embeddings are faster to compute. We ascribe this fact to the different statistics of bunching events when $m > n^2$, to which t -SNE is sensitive.

Results reported in table B1 show how the accuracy progressively decreases for increasing m , from nearly perfect predictions to a sudden drop at $m \sim 64$. Looking at the images, we attribute this fact to the—too similar— t -SNE embeddings rather than to a suboptimal design of the CNN. This notwithstanding, better performances might be probably achieved by increasing the sample size (which has been fixed in this analysis).

Data sets	Number of images		
	Q	C + MF	
Training	750	750	75%
Validation	200	200	20%
Test (A)	50	50	5%
	1000	1000	

Figure B1. Data sets used to train the CNNs. Hyper-parameters have been optimized using training sets and validation sets (for model selection), while accuracies are estimated on a separate test set using training sets plus validation sets.

Table B1. Validation via CNN for $n = 4$. We numerically simulated 2×10^3 experiments (Q [51], MF [18], C [35]) with 5×10^3 samples for various m up to 64. The CNN was trained to recognize the two classes (Q and \bar{Q}) with equal weights, i.e. without requiring a focus on Q. Accuracies are averaged over 3 runs of the CNN, and the standard deviation refers to the last digits. See figure B1 and the supplemental material (section ‘Details on validation via CNNs’) for details on the simulation. With these settings, the employed CNN was effective until $m \sim 49$.

n	m	A	A_C	A_{MF}	A_Q
4	16	0.99 (1)	1.00 (0)	0.99 (1)	0.99 (1)
	25	0.96 (1)	1.00 (0)	0.99 (1)	0.94 (2)
	36	0.95 (2)	0.97 (2)	0.95 (3)	0.99 (1)
	49	0.85 (3)	1.00 (0)	1.00 (0)	0.70 (6)
	64	—	—	—	—

ORCID iDs

Fulvio Flamini  <https://orcid.org/0000-0003-4999-2840>

Nicolò Spagnolo  <https://orcid.org/0000-0003-3471-2252>

Fabio Sciarrino  <https://orcid.org/0000-0003-1715-245X>

References

- [1] Doshi-Velez F and Kim B Towards a rigorous science of interpretable machine learning arXiv:1702.08608v2
- [2] Dunjko V and Briegel H J 2018 *Rep. Prog. Phys.* **81** 074001
- [3] Adèr H J, Mellenbergh G J and Hand D J 2008 *Advising on Research Methods: a consultant's companion* (Huizen: Johannes van Kessel Publishing) ch 14, pp 333–356
- [4] Andrienko N and Andrienko G 2005 *Exploratory Analysis of Spatial and Temporal Data: a Systematic Approach* (Berlin: Springer) (<https://doi.org/10.1007/3-540-31190-4>)
- [5] Lehmann E L and Romano J P 2005 *Testing Statistical Hypotheses* 3rd edn (New York: Springer) (<https://doi.org/10.1007/0-387-27605-X>)
- [6] Lee J A and Verleysen M 2007 *Nonlinear Dimensionality Reduction* 1st edn (Berlin: Springer) (<https://doi.org/10.1007/978-0-387-39351-3>)
- [7] van der Maaten L J P and Hinton G E 2008 *J. Mach. Learn. Res.* **9** 2579–605
- [8] Gashi I, Stankovic V, Leita C and Thonnard O 2009 An experimental study of diversity with off-the-shelf antivirus engines 2009 8th IEEE Int. Symp. on Network Computing and Applications pp 4–11
- [9] Hamel P and Eck D 2010 Learning features from music audio with deep belief networks 11th Int. Society for Music Information Retrieval Conf. (ISMIR 2010) (<http://ismir2010.ismir.net/proceedings/ismir2010-58.pdf>)
- [10] Jamieson A R, Giger M L, Drukker K, Li H, Yuan Y and Bhooshan N 2009 *Med. Phys.* **37** 339–51
- [11] Wallach I and Lilien R 2009 *Bioinformatics* **25** 615–20
- [12] Birjandtalab J, Pouyan M B and Nourani M 2016 Nonlinear dimension reduction for EEG-based epileptic seizure detection IEEE-EMBS Int. Conf. on Biomedical and Health Informatics (BHI) pp 595–8
- [13] Flamini F, Spagnolo N and Sciarrino F 2019 *Rep. Prog. Phys.* **82** 016001
- [14] Aaronson S and Arkhipov A 2011 The computational complexity of linear optics *Proc. 43rd Annual ACM Symp. on Theory of Computing* pp 333–42
- [15] Harrow A W and Montanaro A 2017 *Nature* **549** 203–9
- [16] Wang S T and Duan L M Certification of boson sampling devices with coarse-grained measurements arXiv:1601.02627
- [17] Agresti I, Viggianiello N, Flamini F, Spagnolo N, Crespi A, Osellame R, Wiebe N and Sciarrino F 2019 *Phys. Rev. X* **9** 011013
- [18] Tichy M C, Mayer K, Buchleitner A and Molmer K 2014 *Phys. Rev. Lett.* **113** 020502
- [19] Spagnolo N et al 2013 *Phys. Rev. Lett.* **111** 130503
- [20] Broome M A, Fedrizzi A, Rahimi-Keshari S, Dove J, Aaronson S, Ralph T C and White A G 2013 *Science* **339** 794–8
- [21] Crespi A, Osellame R, Ramponi R, Brod D J, Galvão E F, Spagnolo N, Vitelli C, Maiorino E, Mataloni P and Sciarrino F 2013 *Nat. Photon.* **7** 545–9
- [22] Spring J B et al 2013 *Science* **339** 798–801
- [23] Tillmann M, Dakic B, Heilmann R, Nolte S, Szameit A and Walther P 2013 *Nat. Photon.* **7** 540–4
- [24] Spagnolo N et al 2014 *Nat. Photon.* **8** 615–20
- [25] Carolan J et al 2014 *Nat. Photon.* **8** 621–6
- [26] Bentivegna M et al 2015 *Sci. Adv.* **1** e1400255
- [27] Carolan J et al 2015 *Science* **349** 711
- [28] Tillmann M, Tan S H, Stoeckl S E, Sanders B C, de Guise H, Heilmann R, Nolte S, Szameit A and Walther P 2015 *Phys. Rev. X* **5** 041015
- [29] Crespi A, Osellame R, Ramponi R, Bentivegna M, Flamini F, Spagnolo N, Viggianiello N, Innocenti L, Mataloni P and Sciarrino F 2016 *Nat. Commun.* **7** 10469
- [30] He Y et al 2017 *Phys. Rev. Lett.* **118** 190501
- [31] Lored J C, Broome M A, Hilaire P, Gazzano O, Sagnes I, Lemaitre A, Almeida M P, Senellart P and White A G 2017 *Phys. Rev. Lett.* **118** 130503
- [32] Wang H et al 2017 *Nat. Photon.* **11** 361–5

- [33] Wang H *et al* 2018 *Phys. Rev. Lett.* **120** 230502
- [34] Wang X J, Jing B, Sun P F, Yang C W, Yu Y, Tamma V, Bao X H and Pan J W 2018 *Phys. Rev. Lett.* **121** 080501
- [35] Aaronson S and Arkhipov A 2014 *Quantum Inf. Comput.* **14** 1383–423
- [36] Gogolin C, Kliesch M, Aolita L and Eisert J 2013 Boson-sampling in the light of sample complexity arXiv:1306.3995
- [37] Bentivegna M *et al* 2014 *Int. J. Quantum Inf.* **12** 1560028
- [38] Crespi A 2015 *Phys. Rev.* **91** 013811
- [39] Liu K, Lund A P, Gu Y J and Ralph T C 2016 *J. Opt. Soc. Am. B* **33** 1835–41
- [40] Shchesnovich V S 2016 *Phys. Rev. Lett.* **116** 123601
- [41] Walschaers M, Kuipers J, Urbina J D, Mayer K, Tichy M C, Richter K and Buchleitner A 2016 *New J. Phys.* **18** 032001
- [42] Dittel C, Keil R and Weihs G 2017 *Quantum Sci. Technol.* **2** 1
- [43] Viggianiello N *et al* 2018 *New J. Phys.* **20** 033017
- [44] Viggianiello N, Flamini F, Bentivegna M, Spagnolo N, Crespi A, Brod D J, Galvão E F, Osellame R and Sciarrino F 2018 *Sci. Bull.* **63** 1470–8
- [45] Wang L J *et al* 2017 *Phys. Rev. A* **95** 012301
- [46] Dittel C, Dufour G, Walschaers M, Weihs G, Buchleitner A and Keil R 2018 *Phys. Rev. A* **97** 062116
- [47] Giordani T *et al* 2018 *Nat. Photon.* **12** 173–8
- [48] Kullback S and Leibler R A 1951 *Ann. Math. Stat.* **22** 79–86
- [49] van der Maaten L Barnes–Hut-SNE arXiv:1301.3342
- [50] Pezzotti N, Mordvintsev A, Holtt T, Lelieveldt B P F, Eisemann E and Vilanova A Linear tsne optimization for the web arXiv:1805.10817
- [51] Clifford P and Clifford R 2018 The classical complexity of boson sampling *Proc. 29th Annual ACM-SIAM Symp. on Discrete Algorithms SODA '18* (Philadelphia, PA: Society for Industrial and Applied Mathematics) pp 146–55
- [52] Nielsen M A and Chuang I L 2010 *Quantum Computation and Quantum Information* (Cambridge: Cambridge University Press) (<https://doi.org/10.1017/CBO9780511976667>)
- [53] Shalev-Shwartz S and Ben-David S 2014 *Understanding Machine Learning: from Theory to Algorithms* (Cambridge: Cambridge University Press) (<https://doi.org/10.1017/CBO9781107298019>)
- [54] Rawat W and Wang Z 2017 *Neural Comput.* **29** 2352–449
- [55] Lecun Y, Bottou L, Bengio Y and Haffner P 1998 Gradient-based learning applied to document recognition in *Proceedings of the IEEE* **86** 2278–324
- [56] Chollet F 2015 Keras
- [57] Martín A *et al* 2015 ‘TensorFlow: Large-scale machine learning on heterogeneous systems’ tensorflow.org
- [58] Wattenberg M, Viégas F and Johnson I 2016 ‘Distil’ How to use *t*-SNE effectively (<https://doi.org/10.23915/distill.00002>)