

# Dynamical learning of a photonics quantum-state engineering process

Alessia Suprano,<sup>a,†</sup> Danilo Zia,<sup>a,†</sup> Emanuele Polino,<sup>a</sup> Taira Giordani,<sup>a</sup> Luca Innocenti,<sup>b,c,d</sup> Alessandro Ferraro,<sup>b,c</sup> Mauro Paternostro,<sup>c</sup> Nicolò Spagnolo,<sup>a</sup> and Fabio Sciarrino<sup>a,\*</sup>

<sup>a</sup>Sapienza Università di Roma, Dipartimento di Fisica, Roma, Italy

<sup>b</sup>Palacký University, Department of Optics, Olomouc, Czech Republic

<sup>c</sup>Queen's University Belfast, School of Mathematics and Physics, Centre for Theoretical Atomic, Molecular, and Optical Physics, Belfast, United Kingdom

<sup>d</sup>Università degli Studi di Palermo, Dipartimento di Fisica e Chimica-Emilio Segrè, Palermo, Italy

**Abstract.** Experimental engineering of high-dimensional quantum states is a crucial task for several quantum information protocols. However, a high degree of precision in the characterization of the noisy experimental apparatus is required to apply existing quantum-state engineering protocols. This is often lacking in practical scenarios, affecting the quality of the engineered states. We implement, experimentally, an automated adaptive optimization protocol to engineer photonic orbital angular momentum (OAM) states. The protocol, given a target output state, performs an online estimation of the quality of the currently produced states, relying on output measurement statistics, and determines how to tune the experimental parameters to optimize the state generation. To achieve this, the algorithm does not need to be imbued with a description of the generation apparatus itself. Rather, it operates in a fully black-box scenario, making the scheme applicable in a wide variety of circumstances. The handles controlled by the algorithm are the rotation angles of a series of waveplates and can be used to probabilistically generate arbitrary four-dimensional OAM states. We showcase our scheme on different target states both in classical and quantum regimes and prove its robustness to external perturbations on the control parameters. This approach represents a powerful tool for automated optimizations of noisy experimental tasks for quantum information protocols and technologies.

Keywords: orbital angular momentum; state engineering; black-box optimization; algorithm; quantum walk.

Received Sep. 1, 2021; revised manuscript received Nov. 14, 2021; accepted for publication Nov. 18, 2021; published online Dec. 13, 2021.

© The Authors. Published by SPIE and CLP under a Creative Commons Attribution 4.0 International License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI.

[DOI: [10.1117/1.AP.3.6.066002](https://doi.org/10.1117/1.AP.3.6.066002)]

## 1 Introduction

Quantum-state engineering of high-dimensional states is a pivotal task in quantum information science.<sup>1–4</sup> However, many existing protocols are platform-dependent and lack universality.<sup>5–10</sup> Conversely, a scheme to engineer arbitrary quantum states, relying on quantum walk (QW) dynamics, was showcased in Ref. 11. QWs are a particularly simple class of quantum dynamics that can be considered to generalize classical random walks.<sup>12</sup> QWs have been implemented in experimental platforms ranging from trapped ions<sup>13,14</sup> and atoms<sup>15</sup> to photonics circuits.<sup>16–23</sup>

In particular, engineering of arbitrary qudit states has been experimentally demonstrated with QWs in orbital angular momentum (OAM) and polarization degrees of freedom of light.<sup>11,24,25</sup>

In the paraxial approximation, the angular momentum of light can be decomposed in spin angular momentum, also referred to as polarization in this context, and OAM, which is related to the spatial transverse structure of the electromagnetic field.<sup>26–28</sup> In the classical regime, OAM finds application in particle trapping,<sup>29</sup> microscopy,<sup>30,31</sup> metrology,<sup>32</sup> imaging,<sup>33–35</sup> and communication.<sup>36–40</sup> On the other hand, in the quantum regime, OAM provides a high-dimensional degree of freedom, useful, for example, to encode large amounts of information in single-photon states. Applications include quantum communication,<sup>41–45</sup> computing,<sup>3,4,46</sup> simulation,<sup>47,48</sup> and cryptography.<sup>49,50</sup>

\*Address all correspondence to Fabio Sciarrino, [fabio.sciarrino@uniroma1.it](mailto:fabio.sciarrino@uniroma1.it)

<sup>†</sup>These authors equally contributed to this work.

Optimization algorithms have been proven to be useful tools in tasks such as detection of qudit states<sup>51</sup> and quantum state engineering.<sup>52,53</sup> Machine learning and genetic algorithms have also found many uses in photonics,<sup>54,55</sup> including the use of generative models,<sup>56</sup> quantum state reconstruction,<sup>57,58</sup> automated design of experimental platforms,<sup>59–61</sup> quantum-state and gate engineering,<sup>52,53,62–65</sup> and the study of Bell nonlocality.<sup>66–68</sup> Moreover, genetic algorithms have been employed to design adaptive spatial mode sorters using random scattering processes.<sup>69</sup> Between these types of algorithms, those based on a black box approach have the advantage that they do not rely on specific knowledge of the underlying experimental apparatus. These algorithms are built to tune a set of control parameters based on the information acquired from their environment, without having a notion of what the parameters represent in the experimental platform. This makes such approaches flexible and easier to apply in several scenarios.

In this paper, we showcase the use of RBFOpt,<sup>70,71</sup> a gradient-free global optimization algorithm based on radial basis functions (RBFs),<sup>72–74</sup> to learn how to engineer specific quantum states by efficiently tuning the parameters of a given experimental apparatus. The algorithm seeks to optimize cost function  $C(\Theta)$ , with  $\Theta$  a set of real parameters determining the state produced by the apparatus. With cost function  $C(\Theta)$ , we use the quantum state fidelity between the target and current states, as estimated from a given finite number of measurement samples. This makes the cost function inherently stochastic and thus its optimization potentially more complex. As shown in Refs. 70 and 71, RBFOpt is particularly suited to optimize problems with few parameters, with a focus on operating regimes where only a small number of function evaluations are allowed. This is fully appropriate for our scenario, where functions evaluations involve the generation and measurement of photonic states and are thus relatively costly.

We apply the proposed protocol to an experimental apparatus implementing discrete-time one-dimensional QWs in the OAM and polarization degrees of freedom of light, using a platform based on polarization-controlling waveplates and q-plates (QPs):<sup>75</sup> devices able to couple polarization and OAM degrees of freedom. This platform was shown to be able to engineer arbitrary target quantum states.<sup>11,24</sup> Such an approach, however, requires full knowledge of the inner workings of the underlying experimental apparatus. This feature makes it harder to flexibly adapt a protocol to the perturbations arising in realistic noisy conditions. On the other hand, an adaptive algorithm operating in a black-box scenario, capable of finding the ideal control parameters independently of the physical substratum it operates in, is intrinsically more resilient to varying environmental and experimental circumstances. To ensure that the performance of our protocol is mostly independent of the specific task to which we apply it here, we avoided fine-tuning of the associated hyperparameters, using the default values presented in Refs. 76 and 77. To further verify the resilience of the learning process, we also performed numerical simulations introducing some noise.

In Sec. 2, we introduce the general optimization framework and the QW model underlying our experimental architecture and showcase the performance of the RBFOpt algorithm in numerical simulations with noise that mimics the experimental conditions. In Sec. 3, we describe the experimental platform and report how our optimization pipeline fares when operating

directly on the experimental data. In Sec. 4, we analyze the performance of the protocol when applied to recover the optimal control parameters following sudden changes due to possible external perturbations in order to probe its flexibility under different scenarios. Finally, in Sec. 5, we summarize the results and relate our conclusions.


## 2 Quantum-State Engineering Process as a Black-Box and Simulated Optimization

In order to study the effects of noise on the RBFOpt algorithm and its feasibility to engineer target quantum states, we apply it to numerically simulated data, reproducing the most likely sources of noise in our experimental apparatus. We study, in particular, the effects of binomial and Poissonian fluctuations on the cost function used by the algorithm.

Generating arbitrary qudit states is a pivotal and ubiquitous task in quantum information science and quantum technologies, with applications ranging from quantum communications<sup>1,78–82</sup> to quantum computation.<sup>3,4,46,83</sup> The general quantum state engineering scenario we consider can be modeled with a parameterized unitary operation  $\mathcal{U}(\Theta)$  for some set of real parameters  $\Theta \in \mathbb{R}^N$ . Given a pair of initial and target states  $|\phi_{\text{in}}\rangle$  and  $|\phi_{\text{target}}\rangle$ , the state engineering task consists of finding values  $\Theta^* \in \mathbb{R}^N$  such that  $\mathcal{U}(\Theta^*)|\phi_{\text{in}}\rangle = |\phi_{\text{target}}\rangle$ .

To achieve this, we employ a numerical optimization algorithm to minimize the cost function  $C(\Theta) \equiv 1 - F(\Theta)$ , where  $F(\Theta) \equiv |\langle \phi_{\text{target}} | \mathcal{U}(\Theta) | \phi_{\text{in}} \rangle|^2$  is the fidelity between current and target states. The optimization is performed in a fully black-box scenario, meaning we want the optimization procedure to be independent of the specifics of the particular optimization task. In particular, the optimization algorithm can control and optimize only the generation parameters  $\Theta$ , even if it has no knowledge about both generation of the output state  $\mathcal{U}(\Theta)|\phi_{\text{in}}\rangle$  and computation of the cost function  $C(\Theta)$ . More specifically, we use RBFOpt,<sup>70,71</sup> which works by building an approximated model of the objective function—referred to as a surrogate model in this context—using a set of RBFs. RBFs are real-valued functions  $\phi_{\mathbf{p}}$  that depend only on the distance from some fixed point:  $\phi_{\mathbf{p}}(\mathbf{x}) = \phi(\|\mathbf{x} - \mathbf{p}\|)$  for some  $\phi$ . The goal of the surrogate model used in RBFOpt is to optimally exploit the information collected on the objective function from a limited number of function evaluations. Based on the current estimation of the surrogate model, the algorithm selects new values of the control parameters to improve its current estimation of the model (see Appendix A for further details). This algorithm is an extension of RBF algorithms<sup>72–74,84,85</sup> whose performances are enhanced by providing an improved procedure to find an optimal surrogate model. A comparison of its performances with basic gradient-free algorithms is proposed in Appendix C.

In our case,  $\mathcal{U}(\Theta)$  is the evolution corresponding to a one-dimensional discrete-time QW with time-dependent coin operations. In this model, one considers states in a bipartite space  $\mathcal{H}_W \otimes \mathcal{H}_C$ , where  $\mathcal{H}_W$  is a high-dimensional Hilbert space encoding the possible states of the walker degree of freedom, and  $\mathcal{H}_C$  is a two-dimensional space accommodating the coin degree of freedom. The dynamics are defined as a sequence of iterations, where each iteration is composed of a coin operation  $\hat{C}(\theta)$  followed by a controlled-shift operation  $\hat{S}$ . To simulate the experimental conditions, the operators are defined as



**Fig. 1** Experimental apparatus. (a) The engineering protocol has been tested experimentally in a three-step discrete-time QW encoded in the OAM of light with both single-photon inputs and classical continuous wave laser light (CNI laser PSU-III-FDA) with a wavelength of 808 nm. The single-photon states are generated through a type-II spontaneous parametric down-conversion process in a periodically poled KTP crystal. The input state is characterized by a horizontal polarization and OAM eigenvalue  $m = 0$ . Each step of the QW is made by a coin operator, implemented through a set of waveplates (QWP–HWP–QWP), and the shift operator, realized by a QP. To obtain the desired state in the OAM space, a suitable projection in the polarization space is performed through a quarter-waveplate, a half-waveplate, and a polarizing beam-splitter. The measurement station of the OAM-state is composed by an SLM followed by a single-mode fiber, and the coupled signal is measured through a power meter, in the classical regime, or an avalanche-photodiode detector, in the quantum one. In particular, in quantum optimizations, pairs of photons are generated, and heralded detection is performed, computing the two-fold coincidences between the detectors clicks from the QW evolved photon and the trigger one. The RBFOpt ignores the features of the experimental implementation that is seen as a black box. The algorithm has access only to the  $\Theta$  parameters of the coin operators and to the computed fidelity. (b) During the iterations of the algorithm, the RBFOpt samples the black-box function to construct a surrogate model that is employed in the optimization. In the  $k$ 'th iteration, the algorithm receives as input the fidelity computed in the previous iteration and uses it to improve the surrogate modeling. Moreover, the new parameters  $\Theta_k$  are computed based on the optimization process. This procedure is repeated for each iteration of the algorithm.

$$\hat{C}(\boldsymbol{\theta}) = \begin{pmatrix} e^{-i\beta} \cos \eta & (\cos \mu + i \sin \mu) \sin \eta \\ (-\cos \mu + i \sin \mu) \sin \eta & e^{i\beta} \cos \eta \end{pmatrix},$$

$$\hat{S} = \sum_k |k-1\rangle \langle k|_w \otimes |\downarrow\rangle \langle \uparrow|_c + |k+1\rangle \langle k|_w \otimes |\uparrow\rangle \langle \downarrow|_c, \quad (1)$$

where  $\beta \equiv \theta_1 - \theta_3$ ,  $\eta \equiv \theta_1 - 2\theta_2 + \theta_3$ ,  $\mu \equiv \theta_1 + \theta_3$ , and  $\boldsymbol{\theta} \equiv (\theta_1, \theta_2, \theta_3)$  are the control parameters tuned by the algorithm. This parameterization arises from the sequence of three polarization waveplates used to implement each coin operation. The case in which there are only two waveplates, as in the first step (Fig. 1), is simply obtained from having  $\theta_1 = 0$  and optimizing the values of  $\theta_2$  and  $\theta_3$ . Denoting with  $\boldsymbol{\theta}^{(i)} \equiv (\theta_1^{(i)}, \theta_2^{(i)}, \theta_3^{(i)})$  the free parameters characterizing the coin operation at the  $i$ 'th step, the full set of parameters characterizing an  $n$ -step QW dynamics is then  $\Theta = (\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(n)}) \in \mathbb{R}^{3n}$ . The overall evolution operator corresponding to  $n$  steps is then  $\mathcal{U}(\Theta) \equiv \prod_{i=1}^n \hat{S} \hat{C}(\boldsymbol{\theta}^{(i)})$ . Following the engineering protocol presented in Refs. 11 and 24, we project the coin degree of freedom at the end of the evolution so that our target state is  $|\phi_{\text{target}}\rangle \in \mathcal{H}_w$ .


We apply RBFOpt to optimize a three-step QW, where in the first iteration only two free parameters are used. This

corresponds to a total of eight control parameters:  $\Theta = (\boldsymbol{\theta}^{(i)})_{i=1}^3$  with  $\boldsymbol{\theta}^{(1)} \equiv (0, \theta_2^{(1)}, \theta_3^{(1)})$ . Importantly, the algorithm does not use the information of the correct model  $\mathcal{U}(\Theta)$  of the evolution. This feature permits us to use the present approach in conditions where a model of the experimental setup and noise processes is lacking.

In order to simulate the experimental calculation of the fidelity of a given target state, an orthonormal basis  $\{|\psi_j\rangle\}_{j=1}^d$ , where  $d$  is the dimension of the target state and  $|\psi_1\rangle = |\phi_{\text{target}}\rangle$ , is built through the Gram–Schmidt algorithm. This approach to estimate the cost function is used to simulate the experimental statistics collection process. Furthermore, we consider both Poissonian  $[\mathcal{P}(\lambda)]$  and binomial  $[\mathcal{B}(N, p)]$  fluctuations. Poissonian fluctuations are introduced to take into account laser oscillations, whereas binomial fluctuations reflect the probabilistic nature of the measurement setup.

The number of events of the binomial distribution  $N$  is extracted from a Poissonian distribution with a parameter  $\lambda = 10^4$ , while the probability  $p$  is equal to the fidelity between the state proposed by the algorithm in the  $k$ 'th iteration and the specific element of the basis. Therefore, for each element of the orthonormal basis, the number of detected events is extracted from





**Fig. 2** Simulated optimization: infidelity  $1 - F$  obtained at different stages of the optimization. We test the algorithm on 10 random target states, repeating the optimization 10 times for each. The reported results are obtained as the mean over the average behavior for each of the 10 states. The highest average fidelity obtained is  $0.994 \pm 0.002$ . The shaded area represents the standard deviation of the mean.

the binomial distribution. The noisy fidelity between the proposed state and the target state is then calculated as the ratio between the counts for the element  $|\psi_1\rangle$  and the total number of counts.

We apply the optimization protocol to 10 random four-dimensional target states, repeating the optimization 10 times for each state. In Fig. 2, we show the value of the cost function—i.e., the infidelity between current and target states—obtained at different stages of the algorithm, up to the fixed maximum number of 1000 iterations. For each iteration number, we report the infidelity obtained as the mean over the average behavior of each of the 10 states. The obtained trend demonstrates that, also in noisy conditions, the algorithm manages to minimize the function, and promising results are obtained. Moreover, we also investigate the scalability of the proposed approach when the number of parameters increases. In particular, we simulated QWs of up to 17 steps (50 parameters) and observed in the investigated regime a linear increase in the mean number of iterations needed to achieve a fidelity value of at least 98%. Further details are reported in [Appendix B](#).

### 3 Experimental Dynamical Learning

The capability of manipulating the OAM of light enabled effective experimental implementations of high-dimensional discrete-time QWs. Therefore, to test experimentally the optimization procedure, we exploit a setup based on the scheme proposed in Ref. 11. In particular, we implemented three steps of a discrete-time QW encoding the coin state in the photon polarization and the walker in the OAM degree of freedom. At each iteration, the coin operation is implemented as a set of polarization waveplates, while the controlled-shift occurs via a QP, a device that acts on the OAM conditionally on the polarization state of light:<sup>75</sup>

$$\hat{Q} = \sum_m |m-1\rangle\langle m| \otimes |L\rangle\langle R| + |m+1\rangle\langle m| \otimes |R\rangle\langle L|, \quad (2)$$

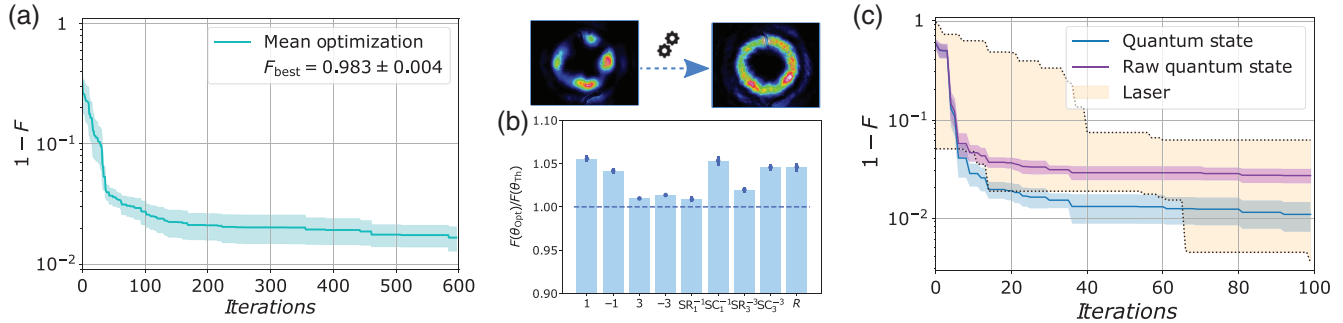
where  $m$  is the OAM value, and  $R$  and  $L$  are the right and left circular polarizations, respectively. We implement arbitrary coin operations using two quarter-waveplates (QWPs) interspaced

with a half-waveplate (HWP). The output OAM state is then obtained performing a suitable projection on the polarization. This is implemented with a set of waveplates followed by a polarizing beam-splitter [cf., Fig. 1(a)].

To measure the fidelity of the output states, we use a measurement apparatus composed of a spatial light modulator (SLM)<sup>86,87</sup> and a single-mode fiber. Since the SLM modulates the beam shape through computer-generated holograms, the operation of this measurement station is equivalent to a projective measurement on the state encoded in the employed hologram. To characterize an incident beam, we thus display on the SLM the hologram corresponding to each element of an orthonormal basis, obtaining the corresponding fidelities. The optimization speed is mainly limited by the measurement process, since significant statistics have to be collected for each projected hologram. Therefore, the use of algorithms able to limit the objective function evaluations, such as those based on the building of a surrogate model, is preferable.

The computed fidelities are then fed to the RBFOpt algorithm to tune the waveplate parameters  $\Theta$ . To achieve this, the algorithm does not require knowledge on the final target state or on the generation and measurement functioning, as shown in Fig. 1(b). However, since the algorithm has no control over the measurement station, the parameters of the latter have been fine-tuned *a priori*, and we are confident of the correctness of this step. Therefore, through a dynamic control of the waveplates' orientation, the algorithm is able to optimize the fidelity value in real time.

To showcase the efficiency of the protocol on our experimental platform, we applied it to engineer different kinds of target states in both the classical and quantum regimes. In Fig. 3, we show the results of running the optimization algorithm on nine different classical states. In particular, we focus our analysis on the elements of the computational basis  $|m\rangle$  with  $m \in \{-1, 1, -3, 3\}$  and on the balanced superposition of two OAM values. We considered both real  $\text{SR}_{m_1}^{m_2} = \frac{|m_1\rangle - |m_2\rangle}{\sqrt{2}}$  and complex superpositions  $\text{SC}_{m_1}^{m_2} = \frac{|m_1\rangle - i|m_2\rangle}{\sqrt{2}}$ , where  $m_1, m_2 \in \{-1, 1, -3, 3\}$  with  $|m_1| = |m_2|$ . Moreover, to verify the efficiency of the protocol, we optimize the engineering of a randomly extracted state ( $R$ ) in the four-dimensional Hilbert space with no zero coefficients corresponding to each basis element. As shown in Fig. 3(a), optimal average values are obtained in 600 algorithm iterations. In particular, the reported infidelity  $1 - F$  is computed, averaging over all the experimentally engineered states, and the minimization is compatible with the numerical results reported in Fig. 2. In Fig. 3(b), we report, for each engineered state, the ratio between the fidelities found by the RBFOpt algorithm and those found using the method presented in Ref. 24 to find the optimal values of the parameters. Indeed, as demonstrated in Ref. 24, it is possible to find coin parameters resulting in an arbitrary target state—albeit possibly with different projection probabilities—regardless of the experimental conditions. We find the fidelities reached by RBFOpt to always be higher than the ones computed using the direct method presented in Ref. 24. This is due to the dynamical learning algorithm we employ, which shows higher performances in compensating experimental imperfections. This showcases the advantages of real-time optimization algorithms for quantum state engineering in realistic scenarios. Notably, we extended the experimental demonstration of the protocol, also in the quantum regime of single photon states. We showcased the engineering of



**Fig. 3** Experimental results: (a) minimization of the quantity  $1 - F$  averaged over the algorithm performances for different experimental states. The mean maximum value reached is  $0.983 \pm 0.004$ . (b) Ratio between the maximum experimental values of the fidelities resulted after the optimization  $F(\Theta_{\text{opt}})$  and the fidelities measured with the theoretical parameters  $F(\Theta_{\text{Th}})$ . For each engineered state, the ratio is higher or compatible with the value 1 highlighted by the dashed line. This confirms that the adopted algorithm can reach performances compatible or even superior with respect to the one obtained with the direct method presented in Ref. 24 that considers ideal experimental platforms. In this sense, the algorithm can take into account and compensate for the experimental imperfections. All of the error bars reported are due to laser fluctuations affecting each measurement and are estimated through a Monte Carlo approach. (c) Comparison between the performances reached in 100 iterations using classical or single-photon input states. In yellow is reported the area between the best and worst optimization performed in the classical case. The blue and violet curves are associated with the minimization of the quantity  $1 - F$  averaged over five different optimizations for the state  $\text{SR}_1^{-1}$  engineered in the quantum domain. In particular, the raw data are shown in violet, whereas the data after accidental counts subtraction are in blue.

the superposition state  $\text{SR}_1^{-1}$ , and we repeated the optimization 5 times, considering only 100 iterations. No differences are expected between the employment of the laser and single-photon states. In Fig. 3(c), we compared the two performances and observed a good agreement between the approaches. In particular, we reported the optimization curves obtained in the quantum regime plotting the raw data, corresponding to  $\sim 4000$  Hz coincidences, and by subtracting the accidental counts. This allows us to distinguish the contribution to the cost function given by either the engineering or the measurement system. The corresponding maximum mean fidelities are  $F = 0.972 \pm 0.003$  and  $F = 0.989 \pm 0.003$ , respectively. In conclusion, since very high fidelities are reached in only 100 steps, the proposed approach can be efficiently applied to quantum situations.

#### 4 Dynamical Learning Protocol with External Perturbations

In realistic conditions, noise is unavoidable, which makes the capability of an algorithm to adapt to real-world perturbations pivotal. To test the robustness of RBFOpt, we have thus added external perturbations to the experimental setup. In particular, we consider a scenario where a sudden perturbation on the parameters is introduced. The algorithm is then tasked with finding again the optimal parameters required to engineer the target state. We assess the performances of the algorithm throughout the optimization, to determine whether a perturbation occurred, and thus the control parameters need to be reoptimized. More specifically:

1. Every 10 iterations, we used the optimal parameters found by the algorithm up to that time  $\Theta_{\text{best}}$  to obtain a new estimate of the cost function  $C_{\text{new}}(\Theta_{\text{best}})$ .

2. To spot if a perturbation occurred, we compared the new value with the one obtained during the algorithm evolution  $C_{\text{sampled}}(\Theta_{\text{best}})$ . So choosing a threshold  $t$ , we proceed as follows.

- (a) If  $C_{\text{new}}(\Theta_{\text{best}}) \leq C_{\text{sampled}}(\Theta_{\text{best}}) + t$ , the optimization is continued.
- (b) If  $C_{\text{new}}(\Theta_{\text{best}}) > C_{\text{sampled}}(\Theta_{\text{best}}) + t$ , the algorithm is restarted.

Therefore, within this approach, the surrogate model is discarded and rebuilt from scratch every time the quantity of interest deteriorates. We performed this check every 10 algorithm iterations in order to have a quick response to perturbations without excessively increasing the optimization time. Indeed, each function evaluation consists of a time-consuming projective measurement with the SLM. For each engineered state, the value of the threshold was fixed by analyzing the fluctuations in the value of the measured fidelity  $F$ , and these values are reported in Table 1.

The considered perturbations act on the HWP of the second step and on the first QWP of the third step. This disturbance consists of a permanent offset in the waveplates rotation of a quantity  $\delta$ . In particular, at each iteration and with probability  $q$ , the orientation of the waveplates optical axis is changed by the addition of an angle sampled from a normal distribution with mean  $\mu = -30$  deg and standard deviation  $\sigma = 5$  deg  $[\mathcal{N}(-30 \text{ deg}, 5 \text{ deg})]$ . We investigated the algorithm response using elements of the computational basis, balanced superpositions of such elements, and a random state. In these cases, several values for the parameter  $q$  are used. The engineered states and the probability  $q$  used for them are reported in Table 1.

An example of the dynamics under perturbations is reported in Fig. 4(a); here, the iteration in which a disturbance is introduced is highlighted by a vertical red or green line, respectively,

**Table 1** The parameters used in the study of the optimization under perturbations for the engineered states. In the second column, we report the values of the perturbation occurrence probability  $q$ , whereas in the third column, we report the threshold values  $t$  used for deciding the algorithm restart.

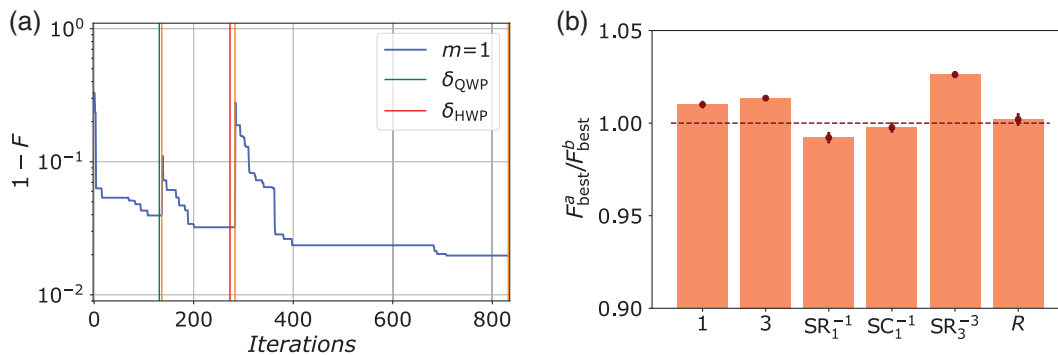
| Target state                                  | Perturbation probability | Restart threshold |
|---|--------------------------|-------------------|
| $ 1\rangle$                                   | 0.0015                   | 0.02              |
| $ 3\rangle$                                   | 0.0015                   | 0.02              |
| $\frac{1}{\sqrt{2}}( -1\rangle +  1\rangle)$  | 0.008                    | 0.02              |
| $\frac{1}{\sqrt{2}}( -1\rangle + i 1\rangle)$ | 0.004                    | 0.02              |
| $\frac{1}{\sqrt{2}}( -3\rangle +  3\rangle)$  | 0.0015                   | 0.05              |
| Random  | 0.0015                   | 0.02              |

for a shift on the HWP or on the QWP. Instead, the restart of the algorithm is indicated with a vertical orange line. As shown, after the perturbation, the minimum found by the algorithm is no longer the optimal solution, thus triggering a restart. The latter allows the algorithm to reach a new optimal solution in a different environmental condition. Moreover, in Fig. 4(b), the mean ratio between the best fidelity found before and after perturbation is reported for each analyzed state. Knowing that for each state more than one perturbation could be performed, the mean ratio is computed, averaging over all of them. Here, values close to or greater than 1 point out how, thanks to the restart, the algorithm is able to readapt its optimal solution and eventually improve the previously obtained fidelity.

## 5 Conclusions

The black-box optimization paradigm we discussed is highly flexible, thus promising to be a powerful tool with the potential to be applicable to problems ranging from optimizations of quantum information platforms to the study of nonclassicality.

We have showcased how the RBFOpt global optimization algorithm allows us to dynamically learn the quantum state generation process. In particular, such an approach enables the optimization of target states engineering without having to devise *ad hoc* platform-dependent protocols. First of all, we dynamically tune the QW parameters in order to optimize the engineering of nine different experimental states in the classical domain. The obtained results turned out to be comparable to the preliminary ones achieved in our numerical simulations. Moreover, the RBFOpt results in higher fidelities than those computed using the direct method of Ref. 24. Therefore, the real-time optimization allows us to take into account and compensate for experimental imperfections. Moreover, we optimized an experimental state using a single-photon source as input to prove the equivalence between the performances reached in the classical and quantum regimes and extend the proposed approach. In order to carry out a complete analysis, and as the adaptation capability of an algorithm is pivotal in realistic conditions, we simulated the effect of real-world perturbations. We have thus applied the optimization algorithm to different states while adding permanent offsets to the orientation of two waveplates in a probabilistic manner. The algorithm manages to adapt itself so as to reach fidelities comparable to those obtained before the perturbation. Our results prove the advantages of adopting real-time optimization algorithms for experimental quantum state engineering protocols. Therefore, practical experimental quantum information experiments can benefit from our work, increasing the engineering performances and employing a real-time fine-tuning of the parameters. The proposed approach can be extended to many different tasks; for example, by suitably modifying the cost function, it is possible to optimize not only the fidelity but also the success probability to generate a target state after the coin projection (see Refs. 11 and 24). Moreover, since the algorithm does not require information on the function to be optimized and on the employed experimental platform, our scheme can find applications in different engineering protocols and quantum information tasks that make use of controllable devices parameters employing, in principle, arbitrary degrees



**Fig. 4** Experimental perturbation results. (a) Optimization under external perturbation of the quantity  $1 - F$  for the state  $|1\rangle$ . The iterations in which a perturbation  $\delta$  occurs are highlighted by a vertical red line (second step HWP) or by a vertical green line (third step QWP), and a vertical orange line highlights the iteration in which the algorithm is restarted. (b) Mean ratio between the best value obtained for the fidelity after ( $F_{\text{best}}^a$ ) and before ( $F_{\text{best}}^b$ ) the perturbation for the different engineered states. The ratio is close to or higher than 1 for all of them, which showcases that the algorithm is able to reobtain and eventually improve the best value sampled before the perturbation. All of the error bars reported are due to laser fluctuations affecting each measurement and are estimated through a Monte Carlo approach.

of freedom. Furthermore, going beyond the fully black-box paradigm, in principle, the approach can be exploited also for different protocols. For instance in the theoretical design of experiments, it could be used to optimize the number of quantum gates needed for a specific desired task. Moreover, it could also be used in the calibration of complex optical circuits that find applications in tasks like boson sampling<sup>88–91</sup> and in the engineering of multiphoton quantum states.<sup>92</sup> In this case, it would be crucial to tailor a suitable cost function.

## 6 Appendix A: Description of the RBFOpt Algorithm

The RBFOpt optimization algorithm is based on the exploitation of a radial basis interpolant, called a surrogate model.<sup>70,71,84,85</sup> Given  $k$  distinct parameter points  $\Theta_1, \dots, \Theta_k \in \Omega$ , where  $\Omega$  is a compact subset of  $\mathbb{R}^N$ , with corresponding cost function values  $C(\Theta_1), \dots, C(\Theta_k)$ , the associated surrogate model  $s_k(\Theta)$  is defined as

$$s_k(\Theta) = \sum_{i=1}^k \lambda_i \phi(\|\Theta - \Theta_i\|) + p(\Theta), \quad (3)$$

where  $\phi(\cdot)$  is an RBF,  $\lambda_1, \dots, \lambda_k \in \mathbb{R}$ , and  $p(\cdot)$  is a polynomial of degree  $d$ . This degree is selected based on the type of the RBF function used in the surrogate model. The possible RBF function choices and the degree of their associated polynomial are reported in Table 2. The hyperparameter  $\gamma$  present in the expression of the RBFs is set to 0.1 by default.<sup>76,77</sup> Moreover, the RBFOpt algorithm automatically selects the RBF that appears to be the most accurate in the description of the problem. This selection is made using a cross-validation procedure, in which the performance of a surrogate model constructed with points  $(\Theta_i, C(\Theta_i))$  for  $i = 1, \dots, k$  is tested at  $(\Theta_j, C(\Theta_j))$  with  $j \neq i$ .<sup>70,71</sup>

The value of the parameters  $\lambda_i$  with  $i = 1, \dots, k$  and the coefficients of the polynomial can be determined solving the following linear system:<sup>70,71,84,85</sup>

$$\begin{cases} s_k(\Theta_i) = C(\Theta_i), & i = 1, \dots, k \\ \sum_{i=1}^k \lambda_i \hat{p}_j(\Theta_i) = 0, & j = 1, \dots, \tilde{d} \end{cases} \quad (4)$$

where  $\Pi_d$  is the space of polynomials of degree less than or equal to  $d$ ,  $\tilde{d}$  is the dimension of  $\Pi_d$ , and  $\hat{p}_1, \dots, \hat{p}_{\tilde{d}}$  are a basis of the space.

**Table 2** The RBFs exploited by the RBF algorithm and the degree of the polynomial used in the construction of the surrogate model.<sup>70,71,84,85</sup> When  $d = -1$ , the polynomial is removed from Eq. (3).

| RBF $\phi(x)$           | Polynomial degree $d$ |
|-------------------------|-----------------------|
| $x$                     | 0                     |
| $x^3$                   | 1                     |
| $\sqrt{x^2 + \gamma^2}$ | 0                     |
| $x^2 \log x$            | 1                     |
| $e^{-\gamma x^2}$       | -1                    |

At the beginning of the optimization procedure, the surrogate model is constructed from a set of parameter points tunable in number and sampled using a Latin hypercube design.<sup>76,77</sup> After that, the interpolant is used to choose the next point, in which the cost function is computed. So, the evolution of the RBFOpt algorithm is composed by the repetition of following steps (say  $k$ 'th step).

1. Compute the surrogate model  $s_k(\Theta)$  from the data points  $(\Theta_i, C(\Theta_i))$ , with  $i = 1, \dots, k$ , solving the linear system of Eq. (4).
2. Use the surrogate model to choose the next point  $\Theta_{k+1}$ . In particular, the metric stochastic response surface method is applied.<sup>70,71,85</sup> Within this framework, the algorithm does a number of global steps controlled by the hyperparameter `num_global_searches` (default value 5<sup>76,77</sup>) and a local step. The latter gives as the next point the one that minimizes the surrogate model.
3. Evaluate the cost function at  $\Theta_{k+1}$  and add  $(\Theta_{k+1}, C(\Theta_{k+1}))$  to the data points.
4. Decide whether to restart the model for lack of improvement. Specifically, if the algorithm does not find a new optimal solution after a number of evaluations defined by the hyperparameter `max_stalled_iterations` (default value 100<sup>76,77</sup>), the actual surrogate model is discarded, and the optimization procedure is restarted from scratch.

Moreover, during the optimization, the algorithm executes a refinement step, the purpose of which is to improve the optimal solution doing a local search around it through variation of a trust region method.<sup>70,71</sup> The refinement step is triggered at the end of point (3) with a frequency controlled by the hyperparameter `refinement_frequency`, with default value equal to 3.<sup>76,77</sup>

Furthermore, in the study concerning the evolution under external perturbation, we add, as explained in Sec. 4, a new condition for triggering a restart. Beyond the default one, we analyzed the deterioration of the optimal value founded for the cost function and decided whether to restart the optimization. This further check was done every 10 iterations in order to have a faster response to perturbations without increasing excessively the number of function evaluations that experimentally are expensive in time.


## 7 Appendix B: Scalability of the Optimization Approach

In this section, we study the RBFOpt behavior as the number of parameters of the objective function increases. In particular, we simulated different experimental configurations with QW steps ranging from 3 to 17 and thus considered up to 50 parameters. In fact, with  $N_{\text{steps}}$  as the number of steps and considering only two waveplates in the first coin, the number of parameters  $N_{\text{par}}$  follows the relation:

$$N_{\text{par}} = 3N_{\text{steps}} - 1. \quad (5)$$

For each case, we generated at random 50 target states and investigated the optimization procedure stopping the process when a fidelity of at least 98% was reached. In all of the evolutions, we added the same Poissonian and binomial noises described in the main text to the fidelity between the target state and the one proposed by the algorithm.






**Fig. 5** Scalability: the plot shows the mean number of RBFOpt algorithm iterations as a function of the black-box problem parameters. Here, the optimization process is interrupted when a value of the fidelity between the target state and the one proposed by the algorithm of at least 98% is reached. For each configuration, the iteration values are obtained by averaging more than 50 random target states and simulating experimental noise using binomial and Poissonian distributions. The uncertainty associated with each point is provided by the standard deviation of the mean.

The computational cost of performing a black-box optimization in high-dimensional spaces can be extracted analyzing how the mean number of iterations changes in relation to the number of parameters. The values obtained averaging over the 50 states considered in our study are reported in Fig. 5 for each simulated configuration. As can be seen from the plot, the RBFOpt algorithm appears to have linear scaling over the parameters number when applied to our implementation. This theoretically showcases the effectiveness of the proposed approach for the engineering of higher dimensional OAM states, and similar behaviors are expected experimentally taking into account the devices response time and adapting properly the related implementation. Finally, while similar behaviors are expected in the regime of a few parameters for higher orders of magnitude, the time needed to perform an iteration step increases drastically. In such regimes, a more refined version of the algorithm might be useful to improve its efficiency.

## 8 Appendix C: Comparison Between RBFOpt and Basic Algorithms

In this section, we perform simulations to compare the RBFOpt algorithm with two basic gradient-free methods suitable to multi-parameter black-box optimization. In particular, we consider both nonadaptive and adaptive approaches.

Regarding the first class, among the simplest is the random search method. As suggested by the name, in each iteration of the optimization processes, the parameters are randomly extracted with a uniform distribution in the parameter space and independently from values assumed in previous steps. The second comparative algorithm is based upon the simplest gradient-free adaptive method known as the Powell method.<sup>93</sup> It attempts to find the local minimum nearest to the starting point. Initially, a set of directions is defined, and the algorithm moves along one of them until a minimum is reached. This minimum becomes the uploaded starting point for the following minimization performed on the second direction. After repeating this procedure



**Fig. 6** Comparison between different optimization algorithms: the plot reports the simulated performances of three different algorithms averaged over the optimization of 10 different states, each of which is repeated 10 times. Dotted blue, dashed green, and continuous orange lines report the trends corresponding to Powell, random search, and RBFOpt, respectively. RBFOpt is found to perform significantly better than the alternatives in most cases. All curves are generated simulating experimental noise with both Poissonian ( $\lambda = 10^4$ ) and binomial fluctuations.

for each direction, a new direction is defined, and the algorithm proceeds to upload the set of directions.

Figure 6, shows the reported trends corresponding to each compared algorithm obtained from averaging the optimizations of 10 distinct states, each of which is repeated 10 times. The experimental conditions are simulated adding both Poissonian ( $\lambda = 10^4$ ) and binomial fluctuations. As expected, both of the adaptive approaches results are advantageous with respect to the random approach for a considerable number of function evaluations. Moreover, since the RBFOpt spans the whole parameter space through the global steps, its performances are substantially better.

## Acknowledgments

We would like to acknowledge the support from the European Union's Horizon 2020 Research and Innovation Program (Future and Emerging Technologies) through project TEQ (Grant No. 766900), QU-BOSS-ERC Advanced Grant (Grant No. 884676), the QUSHIP PRIN 2017 (Grant No. 2017SRNBRK), the DfE-SFI Investigator Program (Grant No. 15/IA/2864), COST Action CA15220, the Royal Society Wolfson Research Fellowship (No. RSWFR3\183013), the Leverhulme Trust Research Project Grant (Grant No. RGP-2018-266), and the UK EPSRC (Grant No. EP/T028106/1).

## References

1. H. Bechmann-Pasquinucci and A. Peres, "Quantum cryptography with 3-state systems," *Phys. Rev. Lett.* **85**(15), 3313–3316 (2000).
2. T. Vértesi, S. Pironio, and N. Brunner, "Closing the detection loophole in Bell experiments using qudits," *Phys. Rev. Lett.* **104**(6), 060401 (2010).
3. B. P. Lanyon et al., "Simplifying quantum logic using higher-dimensional Hilbert spaces," *Nat. Phys.* **5**(2), 134–140 (2009).
4. T. C. Ralph, K. J. Resch, and A. Gilchrist, "Efficient Toffoli gates using qudits," *Phys. Rev. A* **75**(2), 022313 (2007).
5. B. Anderson et al., "Accurate and robust unitary transformations of a high-dimensional quantum system," *Phys. Rev. Lett.* **114**(24), 240401 (2015).



6. A. Rossi et al., "Multipath entanglement of two photons," *Phys. Rev. Lett.* **102**(15), 153902 (2009).
7. M. Hofheinz et al., "Synthesizing arbitrary quantum states in a superconducting resonator," *Nature* **459**(7246), 546–549 (2009).
8. A. C. Dada et al., "Experimental high-dimensional two-photon entanglement and violations of generalized Bell inequalities," *Nat. Phys.* **7**(9), 677–680 (2011).
9. S. Rosenblum et al., "A CNOT gate between multiphoton qubits encoded in two cavities," *Nat. Commun.* **9**(1), 652 (2018).
10. R. W. Heeres et al., "Implementing a universal gate set on a logical qubit encoded in an oscillator," *Nat. Commun.* **8**(1), 94 (2017).
11. T. Giordani et al., "Experimental engineering of arbitrary qudit states with discrete-time quantum walks," *Phys. Rev. Lett.* **122**(2), 020503 (2019).
12. S. E. Venegas-Andraca, "Quantum walks: a comprehensive review," *Quantum Inf. Process.* **11**(5), 1015–1106 (2012).
13. H. Schmitz et al., "Quantum walk of a trapped ion in phase space," *Phys. Rev. Lett.* **103**(9), 090504 (2009).
14. F. Zähringer et al., "Realization of a quantum walk with one and two trapped ions," *Phys. Rev. Lett.* **104**(10), 100503 (2010).
15. M. Karski et al., "Quantum walk in position space with single optically trapped atoms," *Science* **325**(5937), 174–177 (2009).
16. L. Sansoni et al., "Two-particle bosonic-fermionic quantum walk via integrated photonics," *Phys. Rev. Lett.* **108**(1), 010502 (2012).
17. A. Crespi et al., "Anderson localization of entangled photons in an integrated quantum walk," *Nat. Photonics* **7**(4), 322–328 (2013).
18. F. Cardano et al., "Quantum walks and wavepacket dynamics on a lattice with twisted photons," *Sci. Adv.* **1**(2), e1500087 (2015).
19. F. Caruso et al., "Fast escape of a quantum walker from an integrated photonic maze," *Nat. Commun.* **7**(1), 11682 (2016).
20. T. Kitagawa et al., "Observation of topologically protected bound states in photonic quantum walks," *Nat. Commun.* **3**(1), 882 (2012).
21. X. Qiang et al., "Efficient quantum walk on a quantum processor," *Nat. Commun.* **7**(1), 11511 (2016).
22. J. O. Owens et al., "Two-photon quantum walks in an elliptical direct-write waveguide array," *New J. Phys.* **13**(7), 075003 (2011).
23. J. Boutari et al., "Large scale quantum walks by means of optical fiber cavities," *J. Opt.* **18**(9), 094007 (2016).
24. L. Innocenti et al., "Quantum state engineering using one-dimensional discrete-time quantum walks," *Phys. Rev. A* **96**(6), 062326 (2017).
25. A. Suprano et al., "Enhanced detection techniques of orbital angular momentum states in the classical and quantum regimes," *New J. Phys.* **23**(7), 073014 (2021).
26. L. Allen et al., "Orbital angular momentum of light and the transformation of Laguerre-Gaussian laser modes," *Phys. Rev. A* **45**(11), 8185–8189 (1992).
27. A. M. Yao and M. J. Padgett, "Orbital angular momentum: origins, behavior and applications," *Adv. Opt. Photonics* **3**(2), 161–204 (2011).
28. B. Piccirillo et al., "The orbital angular momentum of light: genesis and evolution of the concept and of the associated photonic technology," *Riv. Nuovo Cimento* **36**(11), 501–555 (2013).
29. Q. Zhan, "Trapping metallic Rayleigh particles with radial polarization," *Opt. Express* **12**(15), 3377–3382 (2004).
30. S. Fühapter et al., "Spiral phase contrast imaging in microscopy," *Opt. Express* **13**(3), 689–694 (2005).
31. F. Tamburini et al., "Overcoming the Rayleigh criterion limit with optical vortices," *Phys. Rev. Lett.* **97**(16), 163903 (2006).
32. M. P. Lavery et al., "Detection of a spinning object using light's orbital angular momentum," *Science* **341**(6145), 537–540 (2013).
33. L. Torner, J. P. Torres, and S. Carrasco, "Digital spiral imaging," *Opt. Express* **13**(3), 873–881 (2005).
34. D. S. Simon and A. V. Sergienko, "Two-photon spiral imaging with correlated orbital angular momentum states," *Phys. Rev. A* **85**(4), 043825 (2012).
35. N. Uribe-Patarroyo et al., "Object identification using correlated orbital angular momentum states," *Phys. Rev. Lett.* **110**(4), 043601 (2013).
36. A. E. Willner et al., "Optical communications using orbital angular momentum beams," *Adv. Opt. Photonics* **7**(1), 66–106 (2015).
37. N. Bozinovic et al., "Terabit-scale orbital angular momentum mode division multiplexing in fibers," *Science* **340**(6140), 1545–1548 (2013).
38. M. Malik et al., "Influence of atmospheric turbulence on optical communications using orbital angular momentum for encoding," *Opt. Express* **20**(12), 13195–13200 (2012).
39. J. Baghdady et al., "Multi-gigabit/s underwater optical communication link using orbital angular momentum multiplexing," *Opt. Express* **24**(9), 9794–9805 (2016).
40. J. Wang, "Advances in communications using optical vortices," *Photonics Res.* **4**(5), B14–B28 (2016).
41. D. Cozzolino et al., "High-dimensional quantum communication: benefits, progress, and future challenges," *Adv. Quantum Technol.* **2**(12), 1900038 (2019).
42. X.-L. Wang et al., "Quantum teleportation of multiple degrees of freedom of a single photon," *Nature* **518**(7540), 516–519 (2015).
43. M. Krenn et al., "Twisted photon entanglement through turbulent air across Vienna," *Proc. Natl. Acad. Sci. U. S. A.* **112**(46), 14197–14201 (2015).
44. M. Malik et al., "Multi-photon entanglement in high dimensions," *Nat. Photonics* **10**(4), 248–252 (2016).
45. A. Sit et al., "High-dimensional intracity quantum cryptography with structured photons," *Optica* **4**(9), 1006–1010 (2017).
46. S. D. Bartlett, H. deGuise, and B. C. Sanders, "Quantum encodings in spin systems and harmonic oscillators," *Phys. Rev. A* **65**(5), 052316 (2002).
47. F. Cardano et al., "Statistical moments of quantum-walk dynamics reveal topological quantum transitions," *Nat. Commun.* **7**(1), 11439 (2016).
48. I. Buluta and F. Nori, "Quantum simulators," *Science* **326**(5949), 108–111 (2009).
49. M. Mirhosseini et al., "High-dimensional quantum cryptography with twisted light," *New J. Phys.* **17**(3), 033033 (2015).
50. F. Bouchard et al., "Quantum cryptography with twisted photons through an outdoor underwater channel," *Opt. Express* **26**(17), 22563–22573 (2018).
51. J. Li, M. Zhang, and D. Wang, "Adaptive demodulator using machine learning for orbital angular momentum shift keying," *IEEE Photonics Technol. Lett.* **29**(17), 1455–1458 (2017).
52. J. M. Arrazola et al., "Machine learning method for state preparation and gate synthesis on photonic quantum computers," *Quantum Sci. Technol.* **4**(2), 024004 (2019).
53. J. Mackeprang, D. B. R. Dasari, and J. Wrachtrup, "A reinforcement learning approach for quantum state engineering," *Quantum Mach. Intell.* **2**(1), 5 (2020).
54. W. Ma et al., "Deep learning for the design of photonic structures," *Nat. Photonics* **15**(2), 77–90 (2021).
55. P. R. Wiecha et al., "Deep learning in nano-photonics: inverse design and beyond," *Photonics Res.* **9**(5), B182–B200 (2021).
56. M. Benedetti et al., "A generative modeling approach for benchmarking and training shallow quantum circuits," *npj Quantum Inf.* **5**(1), 45 (2019).
57. S. Yu et al., "Reconstruction of a photonic qubit state with reinforcement learning," *Adv. Quantum Technol.* **2**(7–8), 1800074 (2019).
58. T. Giordani et al., "Machine learning-based classification of vector vortex beams," *Phys. Rev. Lett.* **124**(16), 160401 (2020).
59. A. A. Melnikov et al., "Active learning machine learns to create new quantum experiments," *Proc. Natl. Acad. Sci. U. S. A.* **115**(6), 1221–1226 (2018).
60. Y. Ren et al., "Genetic-algorithm-based deep neural networks for highly efficient photonic device design," *Photonics Res.* **9**(6), B247–B252 (2021).

61. L. O'Driscoll, R. Nichols, and P. Knott, "A hybrid machine learning algorithm for designing quantum experiments," *Quantum Mach. Intell.* **1**(1), 5–15 (2019).
62. A. Lumino et al., "Experimental phase estimation enhanced by machine learning," *Phys. Rev. Appl.* **10**(4), 044033 (2018).
63. R. Santagati et al., "Witnessing eigenstates for quantum simulation of Hamiltonian spectra," *Sci. Adv.* **4**(1), eaap9646 (2018).
64. J. Wang et al., "Experimental quantum Hamiltonian learning," *Nat. Phys.* **13**(6), 551–555 (2017).
65. K. Rambhatla et al., "Adaptive phase estimation through a genetic algorithm," *Phys. Rev. Res.* **2**(3), 033078 (2020).
66. A. A. Melnikov, P. Sekatski, and N. Sangouard, "Setting up experimental bell tests with reinforcement learning," *Phys. Rev. Lett.* **125**(16), 160401 (2020).
67. D. Poderini et al., "Ab-initio experimental violation of bell inequalities," arXiv:2108.00574 (2021).
68. K. Bharti et al., "Machine learning meets quantum foundations: a brief survey," *AVS Quantum Sci.* **2**(3), 034101 (2020).
69. R. Fickler, M. Ginoya, and R. W. Boyd, "Custom-tailored spatial mode sorting by controlled random scattering," *Phys. Rev. B* **95**(16), 161108 (2017).
70. A. Costa and G. Nannicini, "RBFopt: an open-source library for black-box optimization with costly function evaluations," *Math. Programming Comput.* **10**(4), 597–629 (2018).
71. G. Nannicini, "On the implementation of a global optimization method for mixed-variable problems," arXiv:2009.02183 (2021).
72. M. J. D. Powell, "The theory of radial basis function approximation in 1990," in *Advances in Numerical Analysis, Vol. II: Wavelets, Subdivision Algorithms and Radial Functions*, W. Light, Ed., pp. 105–210, Oxford University Press, Oxford (1992).
73. M. J. D. Powell, "Recent research at Cambridge on radial basis functions," in *New Developments in Approximation Theory*, M. W. Müller et al., Eds., pp. 215–232, Birkhäuser, Basel (1999).
74. M. D. Buhmann, *Radial Basis Functions: Theory and Implementations, Cambridge Monographs on Applied and Computational Mathematics*, Cambridge University Press, Cambridge (2003).
75. L. Marrucci, C. Manzo, and D. Paparo, "Optical spin-to-orbital angular momentum conversion in inhomogeneous anisotropic media," *Phys. Rev. Lett.* **96**(16), 163905 (2006).
76. "The source code of the RBFopt algorithm," <https://github.com/coin-or/rbfopt>.
77. "RBFopt documentation," [https://rbfopt.readthedocs.io/\\_downloads/en/latest/pdf/](https://rbfopt.readthedocs.io/_downloads/en/latest/pdf/).
78. M. Fitz, N. Gisin, and U. Maurer, "Quantum solution to the byzantine agreement problem," *Phys. Rev. Lett.* **87**(21), 217901 (2001).
79. N. J. Cerf et al., "Security of quantum key distribution using  $d$ -level systems," *Phys. Rev. Lett.* **88**(12), 127902 (2002).
80. D. Bruß and C. Macchiavello, "Optimal eavesdropping in cryptography with three-dimensional quantum states," *Phys. Rev. Lett.* **88**(12), 127901 (2002).
81. A. Acin, N. Gisin, and V. Scarani, "Security bounds in quantum cryptography using  $d$ -level systems," arXiv:quant-ph/0303009 (2003).
82. N. K. Langford et al., "Measuring entangled qutrits and their use for quantum bit commitment," *Phys. Rev. Lett.* **93**(5), 053601 (2004).
83. E. T. Campbell, H. Anwar, and D. E. Browne, "Magic-state distillation in all prime dimensions using quantum Reed–Muller codes," *Phys. Rev. X* **2**(4), 041021 (2012).
84. H.-M. Gutmann, "A radial basis function method for global optimization," *J. Global Optim.* **19**(3), 201–227 (2001).
85. R. Regis and C. Shoemaker, "A stochastic radial basis function method for the global optimization of expensive functions," *INFORMS J. Comput.* **19**(4), 497–509 (2007).
86. E. Bolduc et al., "Exact solution to simultaneous intensity and phase encryption with a single phase-only hologram," *Opt. Lett.* **38**(18), 3546–3549 (2013).
87. A. Forbes, A. Dudley, and M. McLaren, "Creation and detection of optical modes with spatial light modulators," *Adv. Opt. Photonics* **8**(2), 200–227 (2016).
88. H. Zhong et al., "Quantum computational advantage using photons," *Science* **370**(6523), 1460–1463 (2020).
89. V. Cimini et al., "Calibration of multiparameter sensors via machine learning at the single-photon level," *Phys. Rev. Appl.* **15**(4), 044003 (2021).
90. F. Hoch et al., "Boson sampling in a reconfigurable continuously-coupled 3D photonic circuit," arXiv:2106.08260 (2021).
91. D. Brod et al., "Photonic implementation of boson sampling: a review," *Adv. Photonics* **1**(3), 034001 (2019).
92. J. Pan et al., "Multiphoton entanglement and interferometry," *Rev. Mod. Phys.* **84**(2), 777–838 (2012).
93. M. J. D. Powell, "An efficient method for finding the minimum of a function of several variables without calculating derivatives," *Comput. J.* **7**(2), 155–162 (1964).

**Alessia Suprano** is a PhD student in the Quantum Information Laboratory of Professor Fabio Sciarrino. Her current interests are focused on quantum optics for the implementation and exploitation of quantum walks in the orbital angular momentum degree of freedom of photons and for the study of nonlocality in networks. She graduated in October 2018 from the Sapienza Università di Roma.

**Danilo Zia** is a PhD student in the Quantum Information Laboratory of Professor Fabio Sciarrino. Currently, he is investigating the applications of quantum walk dynamics in the orbital angular momentum degree of freedom of the photons using machine learning techniques. He graduated in October 2020 from Sapienza Università di Roma.

**Emanuele Polino** graduated in November 2016 from the Sapienza Università di Roma and obtained his PhD in February 2020. Currently, he is working on photonic technologies for quantum foundations studies and quantum information tasks. In particular, he worked on fundamental tests on wave-particle duality of photons, study of nonlocality in different quantum networks, quantum metrology protocols, generation, transmission, and measurement of quantum states of light carrying orbital angular momentum, and realization of entangled photons sources.

**Taira Giordani** received her PhD in physics in 2020. She is currently a postdoc in the Physics Department of Sapienza University. Her research focuses on the experimental implementation of quantum walks in integrated photonic devices and the angular momentum of light. In this context, her efforts aim to develop machine learning and optimization methods for the certification and engineering of photonic quantum walk platforms.

**Luca Innocenti** received his PhD in physics in 2020 from Queen's University Belfast, with a thesis on applications of machine learning to quantum information science. He is currently a researcher in the Department of Physics at the University of Palermo. His research interests focus on devising ways to apply machine learning to quantum physics, with a focus on quantum optics and photonics.

**Alessandro Ferraro** received his PhD in physics from the University of Milan in 2005. Afterwards, he held fellowship and research positions at ICFO and UAB (Barcelona), UCL (London), and US (Sussex). He is now a senior lecturer and head of CTAMOP research Centre at QUB (Belfast), where he was appointed in 2013. His research focuses on the theory of quantum information with continuous variables and quantum optics, including quantum correlations and machine learning, thermodynamics, and many-body systems.

**Mauro Paternostro** received his PhD from Queen's University Belfast before securing fellowships from Leverhulme Trust and EPSRC. He was appointed as a lecturer at Queen's in 2008, reader in 2010, and

professor in 2013. He is the head of the School of Maths and Physics and co-leads the Quantum Technology group, where he works on quantum foundations and quantum technologies. His honors include a Royal Society Wolfson Fellowship. He has published more than 200 papers and attracted £20M+ in funding.

**Nicolò Spagnolo** received his PhD in 2012 in physical science of matter, with a thesis on experimental multiphoton quantum optical states. He is an assistant professor in the Department of Physics of Sapienza Università di Roma. His research interests have been focused on quantum computation, simulation, communication and metrology by

employing different photonic platforms, such as orbital angular momentum and integrated photonics.

**Fabio Sciarrino** received his PhD in 2004 with a thesis in experimental quantum optics. He is a full professor and head of the Quantum Information Lab in the Department of Physics of Sapienza Università di Roma. Since 2013, he has been a fellow of the Sapienza School for Advanced Studies. His main field of research is quantum information and quantum optics, with works on quantum teleportation, optimal quantum machines, fundamental tests, quantum communication, and orbital angular momentum.