Optimizing single-mode collection from pointlike sources of single photons with adaptive optics

ALEXANDER D. HILL,[1,*] DAVID HERVAS,[2] JOSEPH NASH,[1] MARTIN GRAHAM,[1] ALEXANDER BURGERS,[3] UTTAM PAUDEL,[4] DUNCAN STEEL,[4] CHRISTIAN SCHNEIDER,[5] MARTIN KAMP,[5] SVEN HÖFLING,[5,6] JIN WANG,[1] JIAHE LIN,[1] WANYING ZHAO,[1] AND PAUL G. KWIAT[1]

[1]*University of Illinois at Urbana-Champaign, Department of Physics, 1110 West Green Street, Urbana, IL 61801-3080, United States of America*
[2]*Universidad San Francisco de Quito, Av. Diego de Robles y Vìa Interoceánica, Riobamba, Ecuador*
[3]*Norman Bridge Laboratory of Physics MC12-33, California Institute of Technology, Pasadena, California 91125, USA*
[4]*University of Michigan, Department of Physics, 450 Church Street, Ann Arbor, MI 48109-1040, United States of America*
[5]*Technische Physik and Wilhelm-Conrad-Röntgen Research Center for Complex Material Sytems, University of Würzburg, 97074 Würzburg, Germany*
[6]*SUPA, School of Physics and Astronomy, University of St. Andrews, St. Andrews, KY16 9SS, United Kingdom*
[*]adhill4@illinois.edu

Abstract: The collection efficiency of light from a point-like emitter may be extremely poor due to aberrations induced by collection optics and the emission distribution of the source. Analyzing the aberrant wavefront (e.g., with a Shack-Hartmann sensor) and correcting accordingly can be infeasible on the single-photon level. We present a technique that uses a genetic algorithm to control a deformable mirror for correcting wavefront aberrations in single-photon signals from point emitters. We apply our technique to both a simulated point source and a real InAs quantum dot, achieving coupling increases of up to 50% and automatic reduction of system drift.

References and links

1. J. D. Sterk, L. Luo, T. a. Manning, P. Maunz, and C. Monroe, "Photon collection from a trapped ion-cavity system," Phys. Rev. A **85**, 062308 (2012).
2. R. Noek, G. Vrijsen, D. Gaultney, E. Mount, T. Kim, P. Maunz, and J. Kim, "High Speed, High Fidelity Detection of an Atomic Hyperfine Qubit," Opt. Lett. **38**, 4 (2013).
3. W. B. Gao, a. Imamoglu, H. Bernien, and R. Hanson, "Coherent manipulation, measurement and entanglement of individual solid-state spins using optical fields," Nat. Photonics **9**, 363–373 (2015).
4. K. H. Madsen, S. Ates, J. Liu, A. Javadi, S. M. Albrecht, I. Yeo, S. Stobbe, and P. Lodahl, "Efficient out-coupling of high-purity single photons from a coherent quantum dot in a photonic-crystal cavity," Phys. Rev. B **90**, 155303 (2014).
5. M. Arcari, I. Söllner, A. Javadi, S. Lindskov Hansen, S. Mahmoodian, J. Liu, H. Thyrrestrup, E. H. Lee, J. D. Song, S. Stobbe, and P. Lodahl, "Near-Unity Coupling Efficiency of a Quantum Emitter to a Photonic Crystal Waveguide," Phys. Rev. Lett. **113**, 1–5 (2014).
6. G. M. Akselrod, C. Argyropoulos, T. B. Hoang, C. Ciracì, C. Fang, J. Huang, D. R. Smith, and M. H. Mikkelsen, "Probing the mechanisms of large Purcell enhancement in plasmonic nanoantennas," Nat. Photonics **8**, 835–840 (2014).
7. N. Somaschi, V. Giesz, L. De Santis, J. C. Loredo, M. P. Almeida, G. Hornecker, S. L. Portalupi, T. Grange, C. Anton, J. Demory, C. Gomez, I. Sagnes, N. D. L. Kimura, A. Lemaitre, A. Auffeves, A. G. White, L. Lanco, and P. Senellart, "Near optimal single photon sources in the solid state," Nat. Photonics **10**, 1–6 (2015).
8. J. D. Wong-Campos, K. G. Johnson, B. Neyenhuis, J. Mizrahi, and C. Monroe, "High-resolution adaptive imaging of a single atom," Nat. Photonics **10**, 606–610 (2016).

9. A. Tiranov, P. C. Strassmann, J. Lavoie, N. Brunner, M. Huber, V. B. Verma, S. W. Nam, R. P. Mirin, A. E. Lita, F. Marsili, M. Afzelius, F. Bussières, and N. Gisin, "Temporal multimode storage of entangled photon pairs," Phys. Rev. Lett. **117**, 240506 (2016).
10. G. N. M. Tabia, "Recursive multiport schemes for implementing quantum algorithms with photonic integrated circuits," Phys. Rev. A **93**, 012323 (2016).
11. M. L. Plett, P. R. Barbier, and D. W. Rush, "Compact adaptive optical system based on blind optimization and a micromachined membrane deformable mirror." Appl. Optics **40**, 327–330 (2001).
12. A. Courteville, "Optimization of single-mode fiber coupling efficiency with an adaptive membrane mirror," Opt. Eng. **41**, 1073 (2002).
13. P. Villoresi, S. Bonora, M. Pascolini, L. Poletto, G. Tondello, C. Vozzi, M. Nisoli, G. Sansone, S. Stagira, and S. De Silvestri, "Optimization of high-order harmonic generation by adaptive control of a sub-10-fs pulse wave front." Opt. Lett. **29**, 207–209 (2004).
14. C. Bonato, S. Bonora, A. Chiuri, P. Mataloni, G. Milani, G. Vallone, and P. Villoresi, "Phase control of a path-entangled photon state by a deformable membrane mirror," J. Opt. Soc. Am. B **27** (2010).
15. M. Minozzi, S. Bonora, a. V. Sergienko, G. Vallone, and P. Villoresi, "Optimization of two-photon wave function in parametric down conversion by adaptive optics control of the pump radiation." Opt. Lett. **38**, 489–491 (2013).
16. V. N. Mahajan, "Strehl ratio of a Gaussian beam." J. Opt. Soc. Am. A **22**, 1824–1833 (2005).
17. A. P. Burgers, J. R. Schaibley, and D. G. Steel, "Entanglement and Quantum Optics with Quantum Dots," in "From Atomic to Mesoscale," (World Scientific, 2015), pp. 103–120.
18. W. B. Gao, P. Fallahi, E. Togan, J. Miguel-Sanchez, and A. Imamoglu, "Observation of entanglement between a quantum dot spin and a single photon," Nature **491**, 426–430 (2012).
19. K. De Greve, L. Yu, P. L. McMahon, J. S. Pelc, C. M. Natarajan, N. Y. Kim, E. Abe, S. Maier, C. Schneider, M. Kamp, S. Höfling, R. H. Hadfield, A. Forchel, M. M. Fejer, and Y. Yamamoto, "Quantum-dot spinâĂŞphoton entanglement via frequency downconversion to telecom wavelength," Nature **491**, 421–425 (2012).
20. J. R. Schaibley, A. P. Burgers, G. A. McCracken, L.-M. Duan, P. R. Berman, D. G. Steel, A. S. Bracker, D. Gammon, and L. J. Sham, "Demonstration of quantum entanglement between a single electron spin confined to an inas quantum dot and a photon," Phys. Rev. Lett. **110**, 167401 (2013).
21. B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. L. Vermeulen, R. N. Schouten, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminiau, and R. Hanson, "Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres," Nature **526**, 682–686 (2015).
22. J. Enderlein, "Theoretical study of detection of a dipole emitter through an objective with high numerical aperture," Opt. Lett. **25**, 634–636 (2000).
23. R. Maiwald, A. Golla, M. Fischer, M. Bader, S. Heugel, B. Chalopin, M. Sondermann, and G. Leuchs, "Collecting more than half the fluorescence photons from a single ion," Phys. Rev. A **86**, 043431 (2012).

1. Introduction

Efficient collection of photons emitted from point-like sources (e.g., trapped ions [1, 2], nitrogen-vacancy centers [3], quantum dots [3, 4], etc.) is often critical for the interrogation of single quantum systems and the efficient realization of quantum computing and communication protocols. Collection optics often introduce significant wavefront aberrations due to the clipping and focusing of collected light from an effective point source. Due to these aberrations and a small collection solid angle, coupling of the light emitted by an interrogated quantum system into a single-mode optical fiber (SMF) – the most natural way to connect separate quantum systems – may be extremely inefficient, despite the fact that the light is emitted into a single spatial mode (e.g., a dipole radiation mode). Recent work in producing radiation from pointlike objects in the desired mode has been quite successful using microcavities and antennae matched to the source [5–7]; however, they may be difficult to implement in existing experimental setups in general. Other purely optical approaches have proven successful [8], but are designed for isotropic sources. Note that as photon emitters are contemplated for use in multi-photon quantum information processing applications, e.g., multi-mode quantum repeater networks [9] or demonstrations of photonic integrated circuit based quantum algorithms [10], even modest improvements in coupling efficiency can lead to large net enhancements – especially in multiphoton experiments.

Adaptive optics (AO) allows experimental apparatus to correct for the aberrations created by large numerical aperture lenses and other collection system optics, and thus to optimize coupling to collection fibers (which can then transport the light to detectors, other quantum

systems, photonic circuitry, etc.). AO is effective for optimizing spatial modes for coupling applications [11, 12] and nonlinear optics [13]; however, AO can be difficult in the low-intensity limit, where wavefront measurements are unreliable and photon count fluctuations can thwart traditional gradient methods of system optimization. In this regime, AO must be performed entirely at the single-photon level [14, 15]. We have developed a "drop-in" AO-enhanced collection system that incorporates a novel noise-resistant genetic algorithm to optimize the shape of a deformable mirror for collection of single photons from a point-like emitter into a single-mode optical fiber (SMF) (although we use a SMF as the most common photonic transfer element, our techniques should also work in a large variety of other applications, e.g., enhanced coupling to/from waveguides, plasmonic devices, nano-antennas, etc.). We have tested our system extensively with a simulated point source (directing light through a sub-wavelength aperture) as well as with actual quantum dot photon emitters, and observed significant enhancements in both cases. After describing the genetic algorithm strategy we will present the results of optimizing collection from a sub-wavelength pinhole and from an InAs quantum dot.

2.   Genetic Algorithm

We control the 69 independent electromechanical actuators of a deformable mirror (Alpao DM69) using an in-house designed genetic algorithm specially suited for AO applications with single photons [Fig. 1]. Our algorithm generates random mirror shapes based on previous configurations and weighted by the performance of the generating families; the algorithm first estimates the effect of deforming the mirror on the measured signal (e.g., photon counts coupled into single-mode fiber per second) by computing the count-rate variance observed while randomly permuting subsets of the basis elements. Basis elements may be the full 69 actuators or the 30 lowest-order Zernike polynomials (e.g., tilt, focus, coma, etc.) created over the 69-actuator space of the mirror surface, the choice of which leads to changes in optimization behavior (discussed below). "Child" shapes are constructed by randomly weighting all basis elements to form new generations of mirror configurations. After initialization, randomized deformations are then applied to the mirror surface, creating the "children" of the subsequent generation. The count rate for each of these children is compared; the best test deformations in each generation (the new parents) are weighted by performance and then probabilistically combined to form a new generation of mirror deformations. For our tests, generations were comprised of 20 children from 10 parents selected from the previous generation. Over time (of order 100 generations in our tests) the variance in the generated children is lowered, which allows convergence to an optimal mirror shape.

2.1.   Detailed Description of Algorithm

The algorithm begins by estimating the initial size of the space to search during optimization, which results in the number of families chosen in each generation. The d basis elements (in our case, 69 actuators or 30 Zernike polynomials) are grouped into sets of five. In the case of the Zernike basis, the elements are grouped by polynomials of nearest order (using the Noll ordering, groupings are $\{Z_1, \ldots, Z_5\}, \{Z_6, \ldots, Z_{10}\}, \ldots$). In the actuator basis the elements are ordered by device index and grouped similarly, with thirteen groups containing 5 basis elements and one containing 4, to reach the full 69 actuators. For each grouping, each basis element is assigned a magnitude $\pm\beta$, where β is a parameter set before optimization; basis elements outside the current grouping are set to a neutral (flat) position. Each set (6 for Zernikes, 14 for actuators) of 5 basis elements is permuted through 32 combinations of voltages ($\pm\beta$ for each of the 5 basis elements in each group) so that every combination of elements within each grouping is applied to the deformable mirror. The collected output power or count rate is recorded for each permutation, and the variance in the count rate between all permutations is used to initialize the intergenerational variances a_i, discussed below. Assuming β is large enough, this ensures that the initial set of n children (n set by the user; we typically used 20 children per generation) adequately spans the
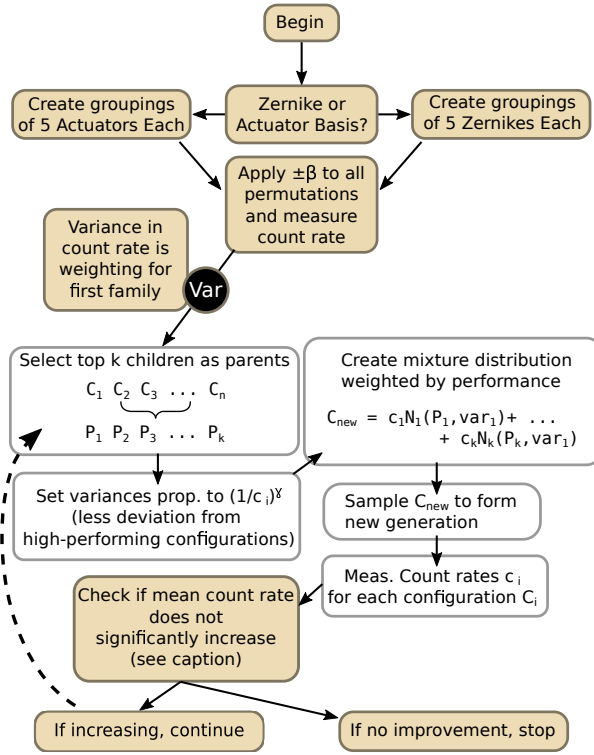
```
                          ┌─────────┐
                          │  Begin  │
                          └─────────┘
                               │
┌──────────────┐      ┌──────────────┐      ┌──────────────┐
│ Create groupings │◄──│ Zernike or   │──►│ Create groupings │
│ of 5 Actuators Each │ │ Actuator Basis? │ │ of 5 Zernikes Each │
└──────────────┘      └──────────────┘      └──────────────┘
                               │
                      ┌──────────────┐
                      │ Apply ±β to all │
                      │ permutations  │
                      │ and measure   │
                      │ count rate    │
┌──────────────┐      └──────────────┘
│ Variance in  │
│ count rate is │
│ weighting for │      ( Var )
│ first family │
└──────────────┘
```

Fig. 1. A schematic of the genetic algorithm. The first generation of mirror configurations is pseudo-randomly generated using a sum of gaussian distributions with variances established by the initialization procedure. Each new generation of mirror configurations is created by weighting gaussian random variables centered on each top-performing mirror configuration, according to count rate. An arbitrary number of children may be generated by sampling this mixture gaussian. The variance of the gaussian random variables is reduced in inverse proportion to configuration performance to allow the mirror generations to converge to an optimal solution. The algorithm completes when it is estimated that the mean has not statistically increased for at least 6 generations.

configuration space needed to optimize collection, since more effective (higher variance) basis elements are allowed to contribute heavily to the first generation of mirror configurations.

After initialization, in each generation we choose k future parents from the previous n children. The next n children are generated by breeding the chosen k parents using the following procedure. Let $\mathbf{Y}_i^t$ be a random vector of d elements (d being dimension of the space, as defined above) representing the test mirror configuration of the i-th parent at time step t. We select the k highest-performing configurations (those resulting in the most photon counts registered leaving the single-mode fiber in a fixed measurement interval) and form the new generation Y^{t+1} by sampling

$$\mathbf{Y}^{t+1} = \sum_{i=1}^{k} \omega_i N(\mathbf{Y}_i^t, a_i \mathbf{I}_d), \tag{1}$$

where $N(\mathbf{u}, \sigma)$ is a multivariate normal distribution with mean $\mathbf{u}$ and covariance σ, and $\omega_i = c_i / \sum_{i=1}^{k} c_i$ is a weighting given by the relative count rates c_i of each of the selected parents. The a_i are set to $a_i = (C/c_i)^\gamma \rho$, with γ and ρ set by the user; here C is an estimate of the maximum possible count rate in the absence of all aberrations, ρ sets the extent of the search

space, and γ sets the sensitivity of the search on the count rate relative to the estimated maximum. In our tests, we typically used $n = 20$, $k = 10$, $\rho = 3.6 \times 10^{-6}$, and $\gamma = 0.5$. These values were influenced by the simulations shown below; however, ρ and γ are strongly dependent on the experimental setup (very sensitive corrections require small ρ to avoid losing all coupling, for example).

The algorithm completes when the statistics of the count rate do not change significantly for several generations. First, we check if the mean count rate for the past 3 generations differs significantly from the mean for the prior 3 generations (comparing 6 generations total). If the mean has not significantly changed, we then take three new generations with twice the measurement time per mirror configuration to reduce fluctuations. If the mean does not change significantly between these three generations, we stop. In each case we consider the statistics of two generations to be significantly different if $|\bar{x}_n - \bar{y}_n| < 0.93\sqrt{2\bar{y}_n/n}$, where x_n is the set of the counts from generation x and y_n is the set of counts from the current generation y (the two-sample Z-test); 0.93 is a parameter that balances the probabilities of either accepting or rejecting convergence erroneously (for $p < 0.05$ rejection of both Type I and Type II errors, this would be 1.96; 0.93 is chosen to balance both errors).

The algorithm depends on a number of parameters that may be customized to the experimental implementation. For example, the size of the initial test mirror deformations should be close to the magnitude of aberrations encountered in the laboratory. Through numerical simulations of the algorithm we have indentified some general starting parameters (discussed in Appendix A).

3. Laboratory Simulation of a Point Source

We physically simulate the collection of light from a point-like emitter using a 600-nm pinhole etched through a 200-nm thick aluminum film on the tip of a SMF [Fig. 2, inset]; the result is a highly-divergent source (NA > 0.8) of 675-nm light. A 4-mm pane of glass representing, e.g., the vacuum window of a simulated emitter's setup, separates the pinhole from either a 10X, 0.25-NA microscope objective or a 0.5-NA aspheric lens; these approximately collimate ~ 500 μW (16%) of the collected light and also serve to define the effective numerical aperture of the collection system by clipping the mode emitted from the fiber. Using two-lens imaging we expand the resulting beam to cover the 10.5-mm-diameter surface of our deformable mirror (9.6° incidence). After reflecting from the mirror, the light is collected by a microscope objective and focused onto a single-mode fiber (Thorlabs 460HP). In order to obtain the optimal coupling without the AO mirror, the fiber position is first optimized using a precision piezo-controlled 6-axis stage (APT Nanotrak), though in practice most of the alignment is optimized using only the positional (X/Y/Z) degrees of freedom. The fiber output is then projected through a beamsplitter; the transmitted light is attenuated and sent to a free-space single-photon detector, while light from the reflected port is recorded by a photodiode to monitor the coupling efficiency while the genetic algorithm optimizes single-photon counts.

Figures 3(a)-3(b) demonstrate the results of typical optimization runs to maximize the coupling efficiency into the final, single-mode fiber. It is important to note, however, that because the algorithm begins optimization with a randomized mirror configuration, the initial couplings recorded in these plots are typically one-third to one-half the coupling for the pre-algorithm, flat-mirror configuration. Thus, the actual overall improvement was somewhat more modest, as the first set of mirror configurations necessarily deviated from the initial optimized alignment with a "flat mirror" surface (a mirror configuration calibrated for near-flatness). Nevertheless, improvements of 100-200% in coupling efficiency over the flat-mirror configuration were typical, though the initial and final couplings are dependent on initial beam misalignment and beam-waist mismatch into the final fiber. Given the finite number of actuators in our deformable mirror, correction of wavefront aberrations is accurate to 14th-15th order Zernike polynomials (Noll index); higher-order aberrations from lens distortion and the resulting Airy pattern will not be

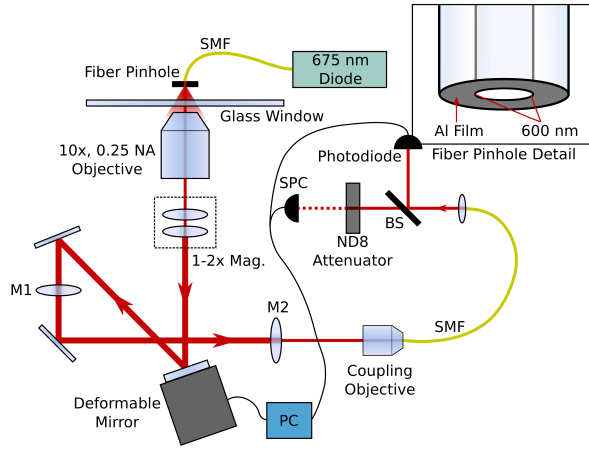completely corrected, which limits the maximum coupling achievable in our simulation.

Fig. 2. Schematic of the experimental setup. Laser light at 675 nm is collected from a pinhole, collimated, and manipulated using a deformable mirror before being contracted (minification lenses M1-2) and focused into a single-mode fiber. The coupling efficiency is monitored using a photodiode, but optimized using a single-photon counter (SPC) after attenuation.

In Fig. 3(a) we show a comparison of the system's performance for various initial count rates. For lower count rates (10 kcps), the count rate is more significantly affected by shot noise and direct comparison of mirror configuration performance is difficult. Nevertheless, our system is able to optimize collection almost as well as for higher count rates (75 kcps or greater) with a small decrease in optimization speed and an increase in the variability of the final count rate. As shown in Fig. 3(b), use of the actuator basis is better able to correct aberrations in the long term given the higher number of basis elements (all 69 actuators); in contrast, the Zernike basis allows for a faster initial optimization due to a shorter initialization time. The actuator basis was used in all tests (decribed below) for consistency.

We also evaluated our technique's adaptibility to changes in both experimental parameters, such as the coupling enhancement for different collection optics, and various algorithmic parameters (e.g., number of children per generation, number of parents selected per generation, etc). In the former case, we tested both a 10x, 0.25-NA microscope objective and a 1", 0.5-NA aspheric lens to collect light from the fiber pinhole. We anticipated that the ability of the mirror and algorithm to correct for aberrations should depend heavily on the order of those aberrations, which in turn depend on the quality and numerical aperture of the collection lens. A comparison of the algorithm performance for both lenses is shown in Fig. 3(c). The microscope objective produces predominantly higher-order aberrations, while the aspheric lens also introduces significant lower-order aberrations; this is reflected in the coupling optimization time required for both cases: the system is able to correct for low-order aberrations (spherical aberration, coma, etc.) much faster than high-order aberrations due to errors introduced in higher-order corrections from the finite number of actuators. The overall coupling efficiency for the aspheric case, however, was comparatively low ($< 5\%$, versus 50% with the microscope objective), as the aspheric lens produces larger aberrations of all orders. Our particular deformable mirror, the Alpao DM-69, accurately corrects most low orders other than severe spherical aberration and secondary astigmatism. The maximum coupling efficiency into a single mode goes as $\eta = 1 - (2\pi/\lambda)^2\sigma^2$, where σ is the RMS wavefront error [16]. For example, for our mirror, addressing a peak-to-valley spherical aberration of 7.26 μm results in an RMS residual error of only 107 nm, but even this corresponds to a maximum theoretical coupling efficiency of only 5% *after* optimization. For
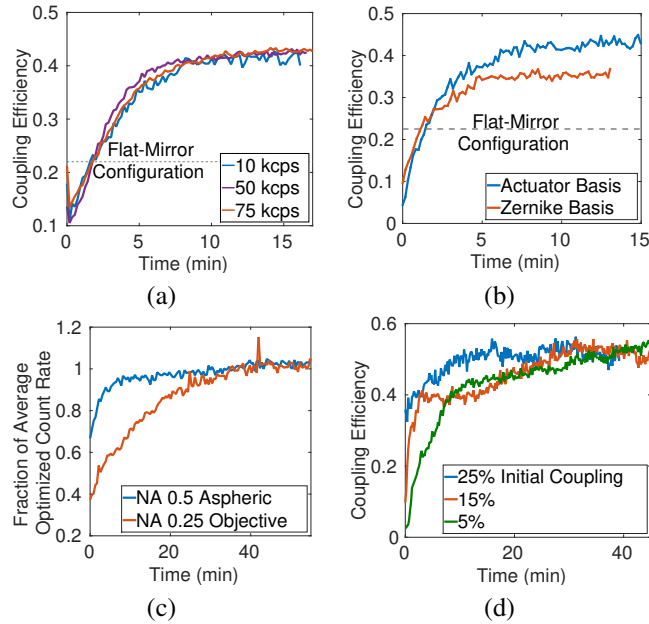
Fig. 3. Performance of the algorithm (a) for various initial count rates (in the actuator basis after first optimizing the flat-mirror coupling), showing a robustness against signals affected heavily by shot noise and (b) for optimization in the basis of the first 30 Zernike polynomials or the basis of all 69 mirror actuators. (c) A comparison of AO correction for 1", 0.5-NA aspheric and 0.25-NA objective collection lenses using the actuator basis, normalized to the average count rate after optimization, in order to emphasize the convergence speed difference. The unnormalized final coupling efficiency with the aspheric lens (< 5%) was actually much lower with the objective (50%). (d) Performance for system misalignment where the assumed 'best' alignment results in a coupling of 50%. The algorithm is able to correct even severely misaligned systems, with a small cost in optimization time. Each line plots the maximum counts per generation, with the classical coupling given by the monitoring photodiode.

other orders this effect is much less pronounced. This is a limitation of the device itself, and should be reduced significantly for higher-resolution models.

We investigated the effect of the initial system alignment (initial coupling) on the algorithm's performance. For these tests, the algorithm achieved a 50% coupling efficiency after a manual alignment; the system was then intentionally misaligned varying amounts by tilting the final mirror before the single-mode fiber. Finally, the algorithm-controlled deformable mirror was used to try to compensate for this misalignment. The system was able to correct these issues (Fig. 3d), though the optimization time increased for poorer alignments.

The parameters controlling the genetic algorithm itself significantly alter its performance. We identified the total time to optimization (defined as the time for the algorithm to reach a steady state with respect to coupling efficiency), and the ratio of AO-corrected coupling efficiency to uncorrected coupling efficiency as metrics for evaluating overall system performance. We tested several strategies: holding the number of children constant in each generation, or linearly or exponentially increasing the number of generated children in proportion to the iteration number. The amount of time spent per generation increased in all cases, though the number of children per generation (and therefore the time per generation) was significantly larger at the end of optimization for the exponential case. Because our algorithm automatically reduces generational variance as the coupling approaches the user-specified theoretical maximum, this additional time

required for these final runs does not necessarily improve the overall optimization efficiency. Monte Carlo simulations of the algorithm indicate that the fastest strategy on average (comparing using a fixed number of children for each iteration, or linearly or exponentially increasing the number of children per iteration) is to fix the number of children at 20 mirror configurations per generation. These simulations also suggest that, for our experimental setup, the algorithm should reach completion after approximately 100 generations (see Appendix A).

4.   Application to Quantum Dots

4.1.   Stationary Collection

The optimization of light collection is crucial for performing experiments using quantum optics techniques in QD systems [17]. Recent work on spin-photon entanglement using single charged QDs are largely limited by collection efficiency [18–20]. Adaptive optics could have a significant impact on experiments such as entanglement swapping via intermediate entangled photons, and were, in fact, used in recent work using NV centers to perform a loophole-free bell test [21].

   We applied our technique to collecting photons from self-assembled InAs quantum dots (QD) grown using molecular beam epitaxy and embedded in a distributed Bragg-reflector (DBR) cavity for enhanced light collection. As shown in Fig. 4(a), the sample is cooled by a liquid-helium optical cryostat; excitation and collection are performed by a high numerical aperture lens (NA=0.68). A Ti:Saph continuous-wave laser is tuned above the ~890-nm band-gap, exciting carriers into the conduction band; these carriers then radiatively recombine at the exciton resonance. The excitation laser (<890 nm) is filtered out by a 925-nm long-pass filter while the QD luminescence (~950 nm) is sent to a single-grating spectrometer where single dot signatures are seen on a liquid nitrogen-cooled CCD. In order to isolate a single QD for optimization, we use an etalon (10-nm free spectral range, finesse of 100) placed after the long-pass filter. The beam path is sent to a single-mode fiber via the deformable mirror; the fiber output is sent to the spectrometer to verify that only QD luminescence is coupled into the fiber. Finally, a fiber-coupled single-photon avalanche photodiode (SPAD) monitors the single photon counts from the dot and serves as the input for the genetic algorithm software.

   The algorithm consistently improved coupling of single photons from the QD [Fig. 4(b)]. Each mirror configuration was tested for one second, with an initial count rate of approximately 15-20 kcps. Though the approximate 1-ns lifetime of the QD would imply a maximum photon production rate of 1 GHz, poor detector efficiency (2%), relatively low-NA collection optics (0.6 NA), and light lost in the sample itself resulted in a low detection rate (<50 kcps). Furthermore, due to the variability in the brightness of individual quantum dots and uncertainty in the alignment of driving lasers onto the QD, establishing an absolute coupling efficiency is impossible (though in this case the absolute efficiency is likely to be quite low). Better index matching from GaAs to air using a solid immersion lens could improve this absolute efficiency. Nevertheless, the algorithm was able to improve the count rate $35 - 50\%$ over the non-optimized (flat-mirror) case across all runs. The discrepancies between the runs were due to variability in the initial dot coupling and overall system alignment between each run, which the system was not able to completely correct.
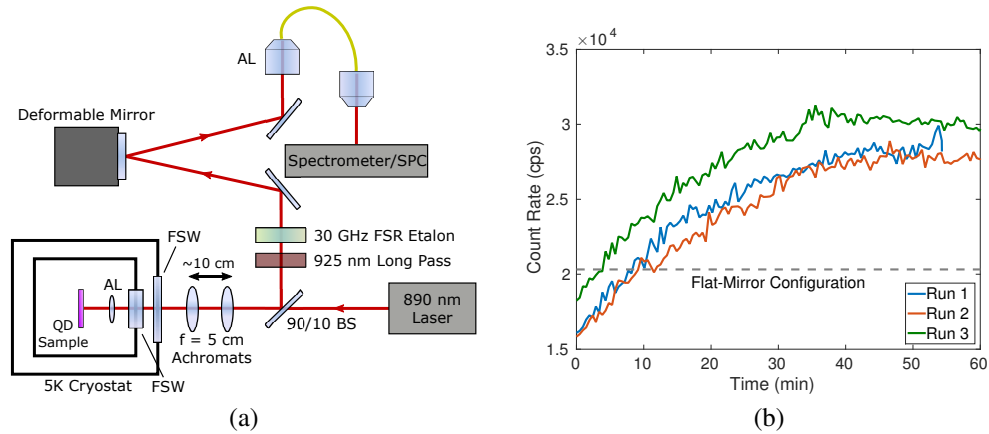
Fig. 4. Optimizing collection from an InAs quantum dot (QD). (a) Photoluminescence light is collected from the QD and collimated with a focusing asphere (AL, Thorlabs 352330-B) before passing through two fused-silica cryostat windows (FSW, 0.2" and 0.125", respectively). Two 5-cm focal length achromatic lenses (Thorlabs AC254-050-B-ML) are used for fine adjustment of the beam. The dot light is then filtered with a 925-nm long-pass and a 30-GHz free-spectral-range etalon to remove residual pump light. The filtered QD light is manipulated by the deformable mirror and coupled into a single-mode fiber (SM980) with an aspheric lens (Thorlabs C260TME-B). The light is then sent to a spectrometer to confirm the collection of QD photoluminescence. After tuning the system for a particular QD, the QD emission is sent to a single-photon counter for optimization using our algorithm. (b) Several optimizations of the collection from a single quantum dot, showing up to 50% improvements in coupling. The optimization rate and final count rate were dependent on the overall system alignment and temperature of the dot.

4.2. Simulation of Source Drift

In practice, almost all point emitters that emit into free space are likely to display some amount of drift, e.g., situations where effects such as thermal expansion or ion drift translate the photon source relative to the collection optics. Ideally, the adaptive collection system should compensate for such drifts. Here we evaluate how well our system achieved this. Because our experimental setup did not enable us to realize such a displacement of the actual source in a controllable way, we instead used a piezoelectric translation stage to apply a 3.95-μm peak-to-peak displacement in the horizontal position of the collection fiber with respect to the final coupling lens. The movement of the fiber results in an oscillation in the count rate as seen by the final single-photon detector, just as a lateral shift in the position of the quantum dot emitter would. In this way, we are able to emulate some of the aberrations caused by a slowly moving source.

We tested the behavior of our QD collection system in the presence of such simulated drifts of two frequencies. Figure 5 shows data with imposed sinusoidal "drifts" with periods of 1000 s (slow drift) and 200 s (moderate drift). For these drift speeds (3.95-μm full displacement, giving peak speeds of 0.12 μm/s peak speed for the 200s period case), we observed an improvement in the mean count rate and reduction in the oscillation magnitude when the AO stabilization was employed ($\sim$ 100% improvement in mean coupling in both cases, with oscillations reduced from $\sim$ 50% peak-to-peak variation to $\sim$ 25% for the 1000-s period case); this indicates that our AO system should be able to provide significant automatic stabilization against the types of drift found in many experiments with point sources.
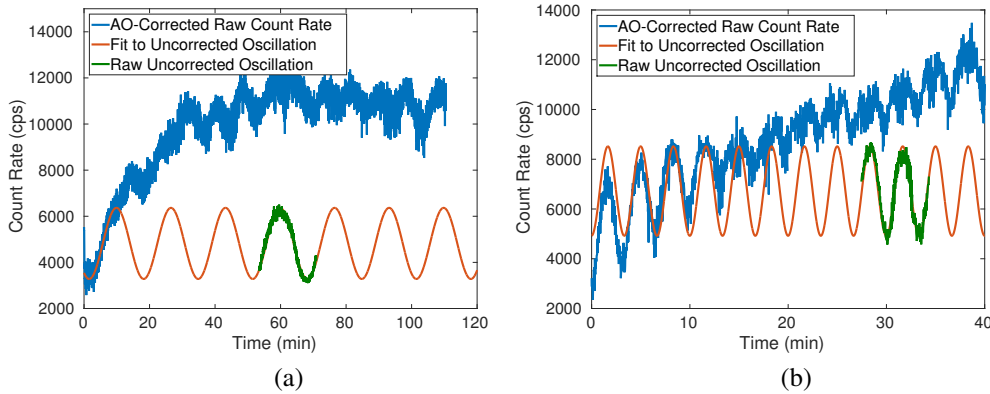
Fig. 5. Behavior of the adaptive optic system in the presence of source drift (oscillating count rates of (a) 1000 s period (1 mHz) and (b) 200-s period (5 mHz). Sinusoidal fits to measured oscillations (green data) before correction are shown in red as a guide to the eye. For slow oscillations (> 200-s period) the algorithm improved the mean photon count and was able to largely suppress the oscillation amplitude. For faster oscillations, the ability of the algorithm to reduce the oscillation amplitude and improve the mean count rate decreased approximately inversely proportionally to the "drift" rate.

5. Final Notes

In conjunction with a deformable mirror, our algorithm is capable of improving the single-mode fiber coupling of aberrant beams from a variety of sources, including from real-world stationary and drifting light sources. At least 1000 counts per test configuration is desirable for reliable performance; this number may be increased at the cost of overall optimization time (assuming the increase is simply due to longer accumulation times), leading to a reliably higher final coupling efficiency as long as the longer collection time does not approach the timescale of any system drifts. For some single-photon emitter applications, low count rates may increase the overall optimization time to unacceptable levels (e.g., molecules that emit only a finite number of photons before bleaching). One possible solution for some cases is to run the collection system backward, stimulating the emitter with high-intensity light reflected from the deformable mirror and optimizing the counts collected from the subsequent emitter fluorescence over a much larger solid angle, e.g., detected by a camera or large area photomultiplier tube. We have verified that our system performs identically in both the forward (e.g., collection from a pinhole or emitter) and backward (e.g., coupling into a pinhole, or driving an emitter) directions. This is clearly a strategy to be tested *in situ*, but one that could greatly improve the time for optimizing the coupling between the single-photon emitter and a fiber or other single-mode optical element.

Our technique is general and can be extended to other applications, such as side-on or oblique coupling of dipoles. The absolute theoretical coupling into *free space* under certain conditions may be near unity [22], though coupling performance into a single Gaussian mode (SMF) will be optics-dependent and therefore difficult to estimate generally. Indeed, though experiments in near-4π collection from dipole emitters with parabolic mirrors have shown possible experimental collection efficiencies of over 50% into free space [23], such experiments have additional optical limitations that can significantly reduce the measured efficiencies of the complete optical system (< 25%), even without the additional constraint of fiber coupling. At least in principle, the single (but non-Gaussian) free-space mode can be mapped with high efficiency onto a SMF mode using adaptive optics.

6. Appendix A: Simulations

The other parameters influencing our algorithm's behavior (the number of parents k, the step rate between generations γ, the scale of the step rate ρ, and the initialization voltage β) are also tunable by the user. Some of these parameters are situation-dependent; for example, in cases where coupling light from a point emitter is extremely sensitive to very small beam aberrations, using large initial deformations (large β) or allowing the children to differ significantly (large ρ) may cause algorithm performance to degrade. However, in some cases we are able to set general guidelines for choosing parameters as a result of numerical simulations of the algorithm. A summary of these results is presented in Fig. 6.

In general, simulations show that using $n = 30-50$ children per generation and selecting $\leq 50\%$ of the children as parents for the new generation is optimal. For our experiments, performed before these simulations were completed, we chose $n = 20$ and $k = 10$. The parameters ρ and γ must be chosen depending on the application. Simulations [Fig. 6(c)] show that ρ should not be too low (too little intergenerational variance) or too high (too much intergenerational variance); for our pinhole coupling application we chose 3.8×10^{-6}, which is related to the maximum actuator voltage allowed on the deformable mirror and the magnitude of the aberrations being corrected. Furthermore, simulations show that γ should be large for short-term optimization, though long-term optimization is unaffected by γ [Fig. 6(d)]; in practice, however, large γ values may result in unstable behavior. For this reason we chose $\gamma = 0.5$ for all experimental tests.
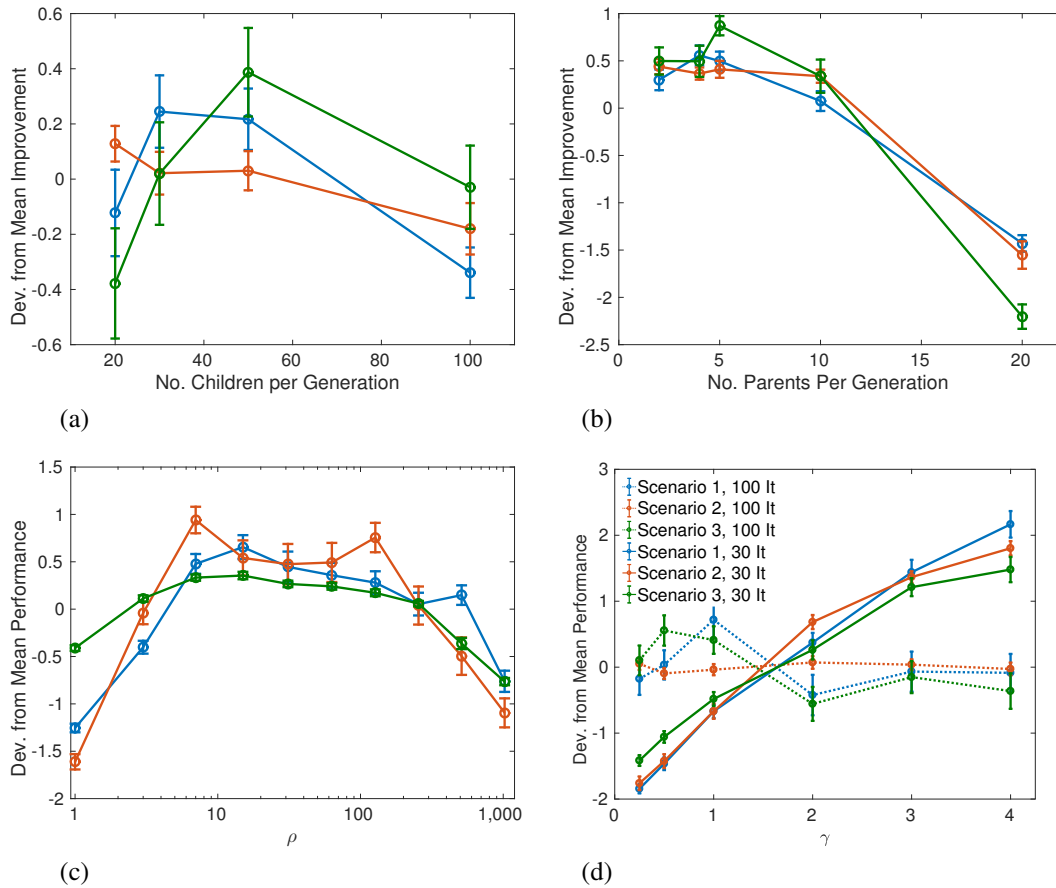
Fig. 6. The results of numerically simulating the genetic algorithm while adjusting various parameters for three aberration scenarios (red, blue, green), adjusted for the final mean performance across the parameter space (i.e., ±1 represents a ±100% difference in algorithm performance from the average case). Each plot represents the coupling improvement relative to the starting efficiency using 2000 test configurations (e.g., 100 generations of 20 children). In (a)-(c), the three colored curves represent three different sets of wavefront aberrations to correct. (a) Adjusting the number of child mirror configurations per generation n while keeping the total time constant (total number of test configurations across all generations, where the number of generations is set to $2000/n$) suggests that, given a constant 10 parents selected per generation, the optimal family size is 20-50 children. For all experimental tests we used 20 children per generation. (b) Selecting all 20 of the children to form the next generation (no culling) degrades performance. In order to balance speed and the final optimized coupling, in all other simulations we used 10 parents and 20 children per generation. (c) Simulations suggest ρ may be optimized for the experimental scenario (ρ as plotted has arbitrary units); for these simulations, an optimal value would be $\rho = 5 - 100$. (d) Finally, in the short term (solid lines, 30 test generations), large γ values (the power scaling used to determined the variance between each new generation) are better, though we have observed that the convergence may be unreliable for larger γ (not shown here). In the long term (dotted lines, 100 generations) there was no advantage for large γ.

7. Appendix B: Tests of Varying Family Sizes for Pinhole Coupling Optimization

Using our pinhole-simulated point source, we experimentally evaluated the algorithm's performance as the number of children per generation was varied [Fig. 7]. We compared holding the number constant at 20 children per generation, adding one child every generation (linearly increasing from 20 children), and exponentially increasing the number of children as $e^{0.04m} + 20$, where m is the generation number. The final optimized value was similar for all cases; however, the optimization speed of the linear case was significantly lower than for the constant or exponential cases. We chose to hold the number of children constant at 20 children per generation for all other tests.


Fig. 7. Experimental algorithm performance as the number of children per generation is varied (holding the generation size at 20 children per generation, increasing linearly from 20 children per generation, and exponentially increasing from 20 children per generation). Each run is normalized to the final optimized value, which was comparable across all cases. The speed of optimization was worst for the linearly-increasing case.

8. Appendix C: Repeated Runs for a Single Quantum Dot

Coupling into a SMF efficiently requires flattening the phase profile of the aberrated beam. The mirror configuration that optimally couples into a SMF will therefore be conjugate to the wavefront of the coupled light. Figure 8 shows the final optimized AO mirror phase profiles for three repeated optimizations of a single quantum dot as discussed in the main text, as well as the corresponding Zernike decompositions of the final mirror configurations. From the average of the three runs we see a tendency to correct significant defocus (index 4), spherical aberration (index 12), and some higher-order aberrations at indices 16 and 17 that may indicate correction of the typical "Maltese-cross" pattern emitted from the quantum dot. Qualitatively, there is some agreement in the general sign and magnitude of each Zernike order applied to the mirror in the optimal configuration; however, the variation between the runs may imply that either a global maximum was not achieved in each case, or the aberrations were not addressable by the mirror (e.g., aberrations were of too high an order). We anticipate that a higher-resolution mirror or longer run times may reduce the variation in final optimization conditions.
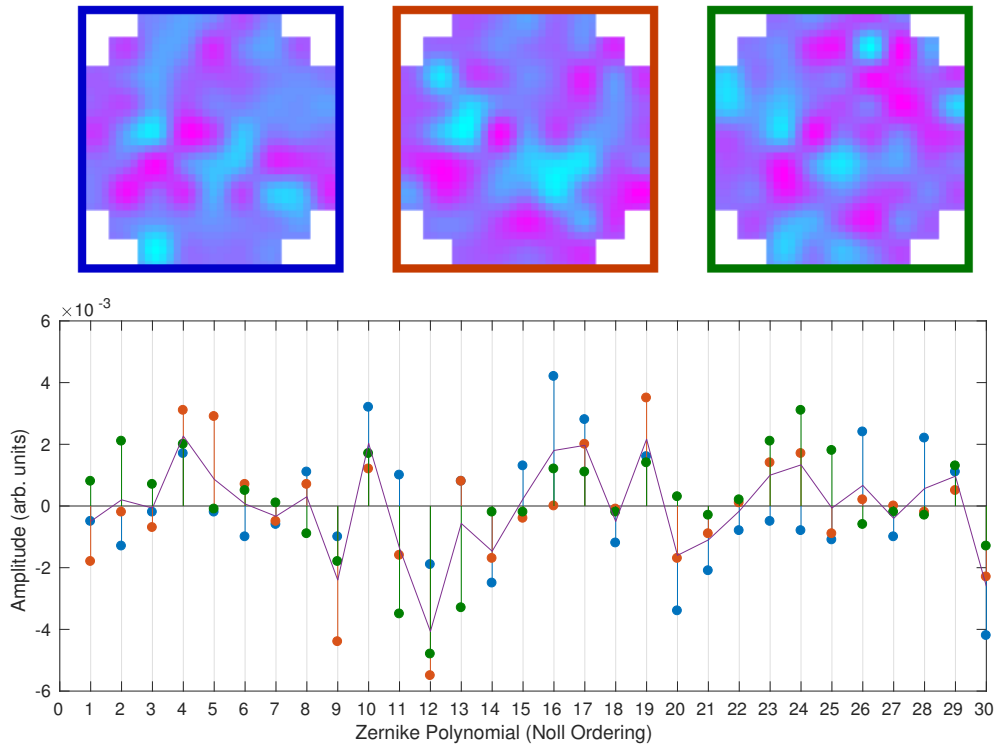
Fig. 8. Zernike amplitudes after optimizing collection from a single quantum dot several times in succession (optimization data presented in main text). The three runs are presented in blue, green, and red, and the solid line displays the average of the three runs. The corresponding final mirror configurations for each run are reproduced above the chart. Actuator voltages of ±0.02 V are magenta and cyan, respectively. The absolute throw per volt is dependent on the overall membrane tension and therefore cannot be estimated directly, but the peak-to-valley displacement is on the order of a few microns or less. There is significant variation in the optimal configuration for each run, which may suggest the presence of higher-order aberrations which are not fully addressable using our finite-resolution mirror.

Funding