Self-learning optical system based on a genetic algorithm driven spatial light modulator

Scott D. Carpenter and Peter M. Weber
Chemistry Department, Box H, Brown University, Providence RI 02912

János Péter and Gábor Szabó
Department of Optics and Quantum Electronics, JATE Szeged, 9 Dom Square, Szeged, Hungary, H-6720

Tamás Szakács and András Lőrincz
Department of Photophysics, Institute of Isotopes, Budapest, P.O.B. 77, Hungary, H-1525

ABSTRACT

We demonstrate the applicability of a genetic algorithm (GA) to control the focus of an adaptive optical system using a liquid crystal spatial light modulator (SLM). The optical setup and the algorithm applied are set to fitness type reinforcement for learning. The particular GA developed optimizes the phase shifts in 32 independent pixels, and is biased towards approximating continuous functions that suit the focusing problem. The learning process is demonstrated to work reliably even in the presence of experimental noise.
Key words : adaptive optics, spatial light modulators, genetic algorithms

1. INTRODUCTION

The recent development of liquid crystal (LC) based SLM's has opened new prospects for adaptive optics (AO) [1-4]. The LC techniques, though not quite mature yet, offer important advantages over conventional adaptive mirrors such as low cost, large number of controllable elements, and a lack of moving parts associated with low power consumption. These advantages, when fully utilized, may bring AO methods to everyday practice. The same SLM's when introduced into systems with appropriate group velocity dispersion characteristics can be used for ultrashort pulse shaping [5], having important applications in communications, and in the control of molecular dynamics on femtosecond time scales [6-9].

Combination of adaptivity with learning algorithms may extend the potentials of AO techniques since in this case the feedback signal can be an arbitrary parameter of the system to be optimized. Moreover, as it has been shown in Ref. [9] for practical quantum control it is even a vital prerequisite to find a successful combination of a global search algorithm with an optics setup capable of shaping the temporal profile of the pulses.

In the present paper we demonstrate that a GA based learning routine combined with a SLM can successfully solve phase front control problems even in the presence of experimental noise. The potential of the method for AO in general is also illustrated.

The particular search algorithm we have developed is a version of the canonical GA [10] that was modified to fit the needs of the present problem of AO. GA's operate in analogy to biological evolution. The starting point is a random population of individuals. In the context of GA's, an individual is a bit string that represents a trial function for a solution. In our case the first 8 bits contain the voltage setting of the first SLM pixel, the second 8 bits that of the second pixel, etc. The collection of all individuals is called a population and the populations evolve over a large number of generations. The individuals of each generation are tested for fitness which is a monotonic function of some appropriate figure of merit measuring the quality of the solution. In the present experiment the fitness was tested by applying the voltage setting encoded in a given individual on the SLM and measuring the corresponding intensity distribution with CCD camera. The fitness was proportional to the intensity inside a (predefined) narrow stripe of the CCD target. The individuals are then ranked based on this fitness, and a series of genetic operations is performed on the population in order to create a new population. First a mating pool is created from individuals from the previous generation. The operator that fills up the pool is called selection. The main effect of this operator is to give individuals of better quality (fitness) more weight in the next generation. It is implemented by choosing individuals into the pool in proportion to their fitness. In standard GA's two other operators called crossover and mutation are used that operate on the mating pool. Mutation relates to random bit flips, with a constant probability for every bit. Mutation is important to keep the population diverse and to reduce the likelihood of premature

SPIE Vol. 3423 • 0277-786X/98/$10.00

convergence to a local maximum. Crossover consists of cutting two bit strings at the same location, and exchanging cut parts between the two strings. The role of crossover is to produce fitter individuals than either of the „parents". Generally crossover is not applied to all individuals: pairs of individuals undergo crossover with a certain probability. We have introduced a new operator, that pushes the population toward approximating a continuous solution, while simultaneously allowing discontinuities. Copy allowed one bit string unit that represents the voltage setting of a given SLM pixel to copy in overwrite mode onto a neighboring unit. (The effect of the copy operator will be discussed later.)

The genetic operators alter the composition of the individuals between generations. The program therefore features an increasingly fit population of individuals, which ultimately may converge to the global fitness maximum. Comparisons have been made between the theoretically well based simulated annealing method and the GA [11]. For the finite trial experiments corrupted by noise it was found that GA has better convergence properties than simulated annealing. This was attributed to the feature that GA is not making sharp decisions, and repeatedly inspects and selects parts of the full control series.

Spatial light modulators can be assembled to either control the amplitude or the phase of an optical wave front. Sandwiching two SLM's allows for the independent control of both amplitude and phase [5]. The SLM (Cambridge Research & Instrumentation Inc. type SLM-256-NIR) we used in phase control mode was capable of delaying the optical phase of the wave front transmitted through a pixel element up to 10π.

2. EXPERIMENTS AND RESULTS

The aim of the proof-of-the-principle experiment reported here was to demonstrate the feasibility of the control of the position of the focus in a self-learning way. This choice was based on the following considerations. It is obvious that the phase delays across the SLM can be adjusted in such a way that the wave front emerging from the planar SLM surface has a certain curvature before entering the lens. If, for instance, this curvature is a parabolic function of the radial coordinate it results in a shift of the position of the focal point along the optical axis. Since the learning algorithm utilizes no a priori knowledge of the desired final shape of the phase front this problem is equivalent to any arbitrary shape as far as learning is concerned. Here, the advantage is that phase fronts extracted from the settings of the individual pixels of the SLM can easily be interpreted.

The experimental arrangement is shown in Fig. 1. The output of a 10 mW HeNe laser is passed through a spatial filter which expands the beam to approximately 5 cm diameter. The beam propagates as a plane wave up to the SLM, where the desired phase modulation is introduced. The beam is then focused by a plano-convex lens of 80 mm focal length. To facilitate the simulation of the expected intensity distributions (see later) the SLM is inserted in the Fourier plane (i.e. back focal plane) of the lens. The focal spot is imaged by a 25x microscope objective onto a CCD camera. The microscope objective provides aberration free imaging, ensuring a faithful image of the focal spot intensity distribution on the CCD. The small depth of focus of the microscope objective helps to identify the axial location of the target plane. The magnification of the microscope objective was chosen so that each CCD pixel corresponds to 1.4 microns on the image plane, which is much smaller than the 17 micron width of the cylindrical focal spot. For measuring the light intensity and the position of the focus we interfaced the CCD camera to a computer using an RS232 interface.

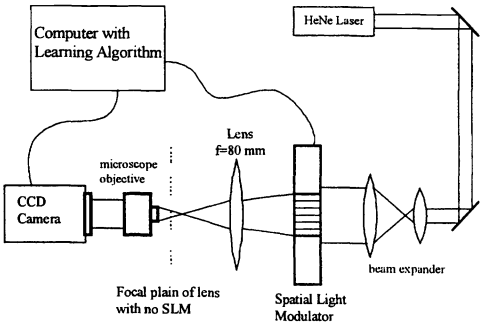

Fig. 1. Schematic of the experimental setup.

If the SLM transmits a planar phase front, the focal spot is nearly diffraction limited in both horizontal and vertical direction and coincides with the paraxial focal plane of the plano-convex lens. (The width of the focal spot is different in the horizontal and vertical directions because the housing of the SLM constitutes a rectangular aperture of 12 X 2 mm.) If a (horizontal) cylindrical wave front curvature is introduced by the SLM such that the phase delay is minimum on the optical axis the position of the (horizontal) focus moves closer to the lens. In the present optical arrangement the focus could also be moved by up to 10 mm in both directions.

With the help of the display software of the CCD camera a rectangular area was selected, representing the focal position of the laser beam for a certain wave front curvature provided by the SLM. The total intensity measured in this area was used as the fitness parameter for the GA.

In the present experiments we used the concatenated binary values of the pixels in the SLM as bit strings for GA encoding. The problem was defined as a 32 pixel phase only modulation with 8 bit phase resolution. This choice was based on

the following considerations: the theoretical shape of the phase front to be found is parabolic. The best approximation that the SLM can realize is a step function where the step-width corresponds to the width of the pixels and the phase setting to the average of the parabola taken over the spatial coordinate range of the actual pixel. It is clear that when the variation of theoretical phase setting over a single pixel is too large the approximation eventually looses its physical meaning. From this reason the region of optimization was limited to pixels not having theoretical phase setting variation values over $\pi/2$ within one pixel.

The problem of finding the optimal phase front of the wave emerging from the SLM is far from trivial, since the global optimum to be found is contained in a space volume of approximately 10^{77} SLM settings. Furthermore, since phase jumps of integer times 2π are not relevant to the propagation of a plane wave there is a tendency in the learning process to break up the phase fronts into segments that are "offset" by these phase jumps. It is evident, however, that aside from the 2π phase jumps, the best solution must be a slow function of pixel position. The algorithm was modified in a way that it could take advantage of this preference of continuity. The introduction of the genetic operator "copy" in the present work addresses this problem by preferentially adjusting pixel elements to the voltage level of their neighbors. This accelerates the learning process by favoring continuity, provided that neighboring bit string units have similar values in most cases. At the same time this does not restrict the solution to be continuous; if a copying results in a gene of low fitness then it will not survive. The operator copy has another favorable property. It works effectively as an intelligent mutation, since it introduces new bits into a solution. In fact, considering typical mutation rates, the operator copy may be considered also as an extremely strong mutation since it may change several bits at a time. Since the optimal size of the population decreases with increasing mutation rate, the copy operator serves to decrease the required size of the population.

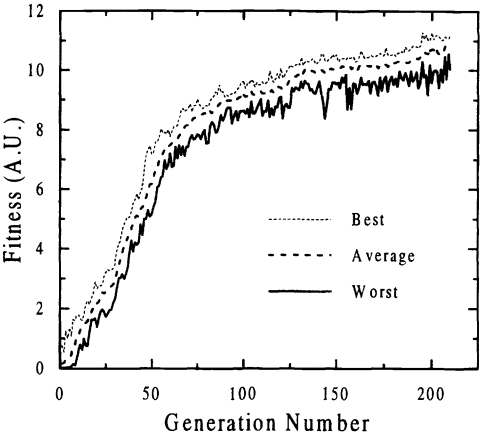

Fig. 2. Experimental learning curve for 20 member population. Curves correspond to best, average, and worst fitness for each generation

We performed experiments with population sizes of 20 and 100 individuals, respectively. The probability of crossover was set to 0.8, and the probability of a mutation was 10^{-4}. We found that the algorithm approximates the global optimum phase setting within about 200 generations. As illustrated in Fig. 2, the algorithm quickly improves the fitness of the population within the first 50 generations. The learning curve subsequently flattens, and improvements in the fitness become very slow. The run was terminated after 210 generations, corresponding to a total measurement time of 3.5 hours. The duration of the learning process was entirely determined by the rate (approximately 3 s) at which the SLM adjusts to a new voltage setting.

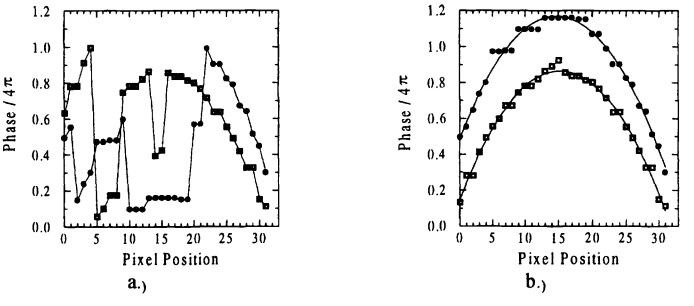

Fig. 3. Phase found by genetic algorithm vs. pixel position with parabolic fits. Circles correspond to 20 individuals per population and squares correspond to 100 individuals per population.
a) before correction b) after correction for 2π phase jumps.

The phase front found by the GA is shown in Fig. 3(a), for both the run with 20 individuals and the one with 100 individuals. (The phase data presented in the graphs were calculated from the pixel voltage settings by using the measured voltage-phase calibration curves of the SLM.) The discontinuities are due to the limited range of voltage settings: slight change in optimal phase settings may require very different voltage values. Having made the corrections at the discontinuity points the distribution of the phases follows nicely the theoretical parabolic shape Fig. 3(b).

To fully understand the experimental results we simulated the intensity distributions in the target plane. We assume in accordance with experimental situation that the SLM is illuminated by a monochromatic plane wave traveling along the optical axis (z). The field amplitude at a given point (x_L,y_L) on the surface of the lens facing the SLM can be calculated by using the Huygens-Fresnel integral

$$E(x_L, y_L) = C \cdot \iint_{SLM} f(x_s) \cdot \frac{\exp(-i \cdot \mathbf{k} \cdot \mathbf{R}')}{R'} dx_s dy_s \tag{1}$$

where (x_s,y_s) and (x_L,y_L) are the spatial coordinates in the plane of the SLM and the lens respectively, $f(x_s)$ is the phase function of the SLM, **k** is the wave vector, **R'** is the vector connecting an arbitrary point from the plane of the SLM with the point where the E-field is evaluated in the plane of the lens, C is a constant. The integration goes over the surface of the SLM. Taking into account the geometry of the problem we can treat the lens in paraxial approximation. In this case the effect of the lens can simply be represented by adding a spherical phase factor with a radius of f, where f is the focal length of the lens. Thus, the field at the surface of the lens facing the target plane can be written as

$$\Psi(x_L, y_L) = E(x_L, y_L) \cdot \exp\left(i \cdot k \cdot \frac{(x_L^2 + y_L^2)}{2 \cdot f} \right) \tag{2}$$

Finally, the E-field in the target plane can be calculated from $\psi(x_L,y_L)$ by using again the Huygens-Fresnel integral

$$P(\xi,\eta) = C \cdot \iint_{Lens} \Psi(x_L, y_L) \cdot \frac{\exp(-i \cdot \mathbf{k} \cdot \mathbf{R})}{R} dx_L dy_L \tag{3}$$

where (ξ,η) are the spatial coordinates in the target plane, and **R** is a vector connecting an arbitrary point from the plane of the lens with the point in the target plane where the E-field is evaluated. By using relatively straightforward algebra and the approximation that the distances in planes perpendicular to the z axis are much smaller than the ones along the z axis (3) can be rewritten as

$$P(\xi,\eta) = C \cdot K \cdot G(\xi,\eta) \cdot \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x_s) \cdot H(x_s, y_s) \cdot \exp\left(i \cdot k \cdot b \cdot (x_s \cdot \xi + y_s \cdot \eta)\right) dx_s dy_s \tag{4}$$

where $b = \dfrac{f}{f \cdot (z + z') - z \cdot z'}$, $K = \dfrac{i \cdot b}{\lambda} \cdot \exp\left(-i \cdot k \cdot (z + z')\right)$,

$$G(\xi,\eta) = \exp\left(-i \cdot k \cdot \frac{\xi^2 + \eta^2}{2 \cdot z} \right) \cdot \exp\left(i \cdot k \cdot \frac{b \cdot z'}{2 \cdot z} \cdot (\xi^2 + \eta^2) \right), \quad H(x_s, y_s) = \exp\left(-i \cdot k \cdot \frac{x_s^2 + y_s^2}{2 \cdot z'} \right) \cdot \exp\left(i \cdot k \cdot \frac{b \cdot z}{2 \cdot z'} \cdot (x_s^2 + y_s^2) \right)$$

From the fact that we inserted the SLM into the back focal plane of the lens b = 1/f follows which, in turn, means that G=1 simplifying (4) substantially. Evaluating (4) and taking $|P(x,h)|^2$ the intensity distribution in the target plane (i.e. in the image plane of the CCD) can be obtained.

The measured and calculated intensity distributions are shown in Fig. 4a. and Fig. 4b respectively. For the calculations we assumed that the SLM introduces a parabolic phase front into the beam corresponding to 5 mm focus shift that is the expected best solution to the problem. From Fig. 4. it is clear that agreement between the measured and calculated intensity distributions is fairly good. Closer inspection of the vertical intensity distribution can provide further help to identify the absolute position along the z axis. In Fig. 5. we have depicted the calculated intensity distribution along the y axis for 4 (dashed line), 5 (solid line), and 6 mm (dotted line) focal shifts along with the measured data (squares). The figure clearly supports that the algorithm found the best solution since it shows that the vertical intensity would be significantly different even for errors in the focal shift well below 1 mm.

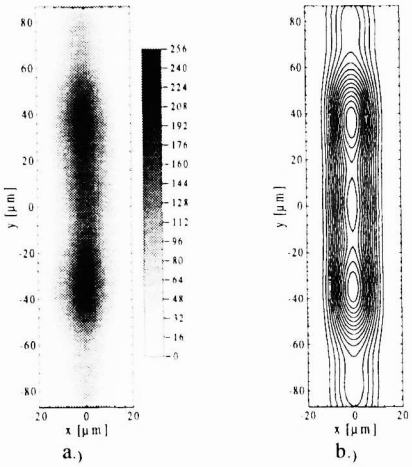

Fig. 4. Intensity distributions when the horizontal focal plane was shifted by 5mm
a) measured b) calculated

It is important to note here that the physical parameter to be learnt by the GA was the voltage setting of the individual pixels. This means that the algorithm solved the control problem without using the phase calibration curves of the SLM. This is a significant advantage of our method that may prove to be especially important in cases when the dependence of the control parameters is nonlinear (as in the case of LC SLM-s) or when calibration of the control device is difficult.

The overall phase of the transmitted phase front is irrelevant to the fitness as defined for the current experiment, and it can be seen that there is a constant phase difference between the two experimental curves. As expected, the curve with 100 individuals resulted in a closer approximation to the theoretical parabolic shape.

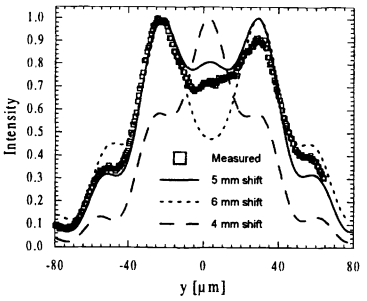

Fig. 5. Intensity distributions along the y axis: the measured data (squares) and the calculated intensity distributions for 4 (dashed line), 5 (solid line), and 6 mm (dotted line) shifts in the position of the horizontal focus

3. CONCLUSION

In conclusion, we have demonstrated the operation of a global search algorithm to steer AO in order to optimize an experimental quantity. The methodology is expected to find many applications, including the optimal control of wavepacket evolution and the operation of AO in problems where diffraction limited focusing through optically inhomogeneous media is necessary, like confocal microscopy. In our particular experiment, a GA steered the focus of a laser beam to a given target plane. The GA was chosen because of its tolerance towards experimental noise. The experimental learning curve shows the effective interplay of the global search algorithm with the AO. At the present time, the speed of the focusing experiment is entirely given by the update rate of the spatial light modulator and with a faster SLM, or an acousto-optic device [12] orders of magnitude could be gained. Using a fast 2-dimensional light modulator, possibly with a circular pixel structure, real time focusing of optical devices should become possible. Finally, we would like to stress that the presented self-learning scheme offers a fundamental advantage over servo systems in that it can be used even if the target is a nonlinear or even unknown function of a large number of control variables, provided that an appropriate signal measuring the quality of the different solutions can be given.

ACKNOWLEDGMENTS

This work was supported by the Army Research Office under contract numbers DAAL03-92-G-0312 and DAAH04-93-G-0356, and by an US-Hungarian NSF grant, INT-9307718, OTKA Grants: T017110 and T017364.

REFERENCES

1. J. Amako, H. Miura, and T. Sonehara, *Appl. Opt.* **32**, 4323 (1993).
2. G.D. Love, J.V. Major, and A. Purvis, *Opt. Lett.* **19**, 1170 (1994).
3. G.D. Love, and R. Bhandari, *Opt. Comm.* **110**, 475 (1994).
4. R. Dou, and M.K. Giles, *Opt. Lett.* **20**, 1583 (1995).
5. J.P. Heritage, A.M. Weiner, and R.N. Thurston, *Opt. Lett.* **10**, 609 (1985).
6. D.J. Tannor and S.A. Rice, *J. Chem. Phys.* **83**, 5013 (1985).
7. M. Shapiro and P. Brumer, *J. Chem. Phys.* **84**, 4103 (1986).
8. A.P. Pierce, M.A. Dahleh, and H. Rabitz, *Phys. Rev. A* **37**, 4950 (1988)
9. R.S. Judson and H. Rabitz, *Phys.Rev. Lett.* **68**, 1500 (1992).
10. D.E. Goldberg, *Genetic algorithm in search, optimization and machine learning* (Add.-Wesley, Reading, MA, (1989).
11. G.J. Tóth, A. Lörincz, and H. Rabitz, *J. Chem. Phys.* **101**, 3715 (1994).
12. C.W. Hillegas, J.X. Tull, D. Goswami, D. Strickland, and W.S. Warren, *Opt. Lett.* **19**, 737 (1994).