



# Single-step beam intensity and profile optimization using a $256 \times 256$ micromirror array and reinforcement learning

YUNHUI XIE,<sup>\*,†</sup>  MATTHEW PRAEGER,<sup>†</sup>  JAMES A. GRANT-JACOB,<sup>†</sup>  AND BEN MILLS 

*Optoelectronics Research Centre, University of Southampton, UK*

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

\*[yunhui.xie@soton.ac.uk](mailto:yunhui.xie@soton.ac.uk)

**Abstract:** Many optical applications require accurate control over a beam's spatial intensity profile, in particular, achieving uniform irradiance across a target area can be critically important for nonlinear optical processes such as laser machining. This paper introduces a novel control algorithm for Digital Micromirror Devices (DMDs) that simultaneously and adaptively modulates both the intensity and the spatial intensity profile of an incident beam with random and intricate intensity variations in a single step. The algorithm treats each micromirror within the DMD as an independent Bernoulli distribution characterized by a learnable parameter. By integrating reinforcement learning with fully convolutional neural networks, we demonstrate that the control of 65,536 ( $256 \times 256$ ) micromirrors in a DMD can be achieved with modest computational expense. Furthermore, we implement the Error Diffusion (ED) algorithm as a sampling method and show that an incident beam with random and intricate intensity variations can be modulated to a predefined shape with high uniformity in intensity, both in simulated and experimental environments.

Published by Optica Publishing Group under the terms of the [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

## 1. Introduction

Digital Micromirror Devices (DMDs) are extensively studied opto-electro-mechanical instruments, gaining significant interest in various optical applications due to their distinctive capabilities. DMDs function as programmatically configurable, binary light masks composed of arrays of micro-scale tilting mirrors [1]. Each individual mirror within the array can be electro-mechanically tilted to one of two fixed angles. This precise control over the micromirrors' tilt states allows for the modulation of light by altering the reflection angles, thereby determining the spatial presence (referred to as the “on” state henceforth) or absence (referred to as the “off” state henceforth) of light in the diffraction outputs. DMDs stand out among other light modulation hardware due to their advantageous characteristics [2], most importantly, rapid update rate ( $\sim 30$  kHz) [3], affording a greater control bandwidth.

Recent literature has documented the capability of DMDs to modulate the intensity of an incident beam [4–15]. In the typical use-case of a digital display, the optical system is designed so that each DMD micromirror can be individually resolved (as an image pixel) [16,17]. In this work, however, the DMD is incorporated in an optical system where spatial frequency bandwidth is limited, i.e., where the diffraction limit exceeds the physical dimensions of the ideal images of individual micromirrors. In this configuration (which could be termed “optical oversampling”), the diffraction limit forces light reflected from neighboring micromirrors to overlap at the same spatial positions. Consequently, altering the tilt states of neighboring micromirrors changes

the number of overlapping beams in their combined output, thereby affecting the intensities at corresponding spatial positions.

To simultaneously modulate both the overall intensity level and the spatial intensity profile of an incident beam using a DMD requires control over a substantial number of micromirrors. This necessity arises from the requirement for multiple DMD “pixels” to oversample a single output intensity “pixel” in a spatial frequency bandwidth limited system. Furthermore, the spatial intensity profile of the incident beam may exhibit random and intricate intensity variations due to aberrations in the optical system. These intensity variations can degrade the quality of the output beam if the incident beam is modulated naively. Adaptively correcting these unintended spatial intensity variations requires that the controlling algorithm observe the incident beam before determining an appropriate array of micromirror tilt states. In this work, we present a novel algorithm capable of adaptively determining the tilt states of individual micromirrors in a single step, enabling control of both the overall intensity and the spatial intensity profile of the output beam.

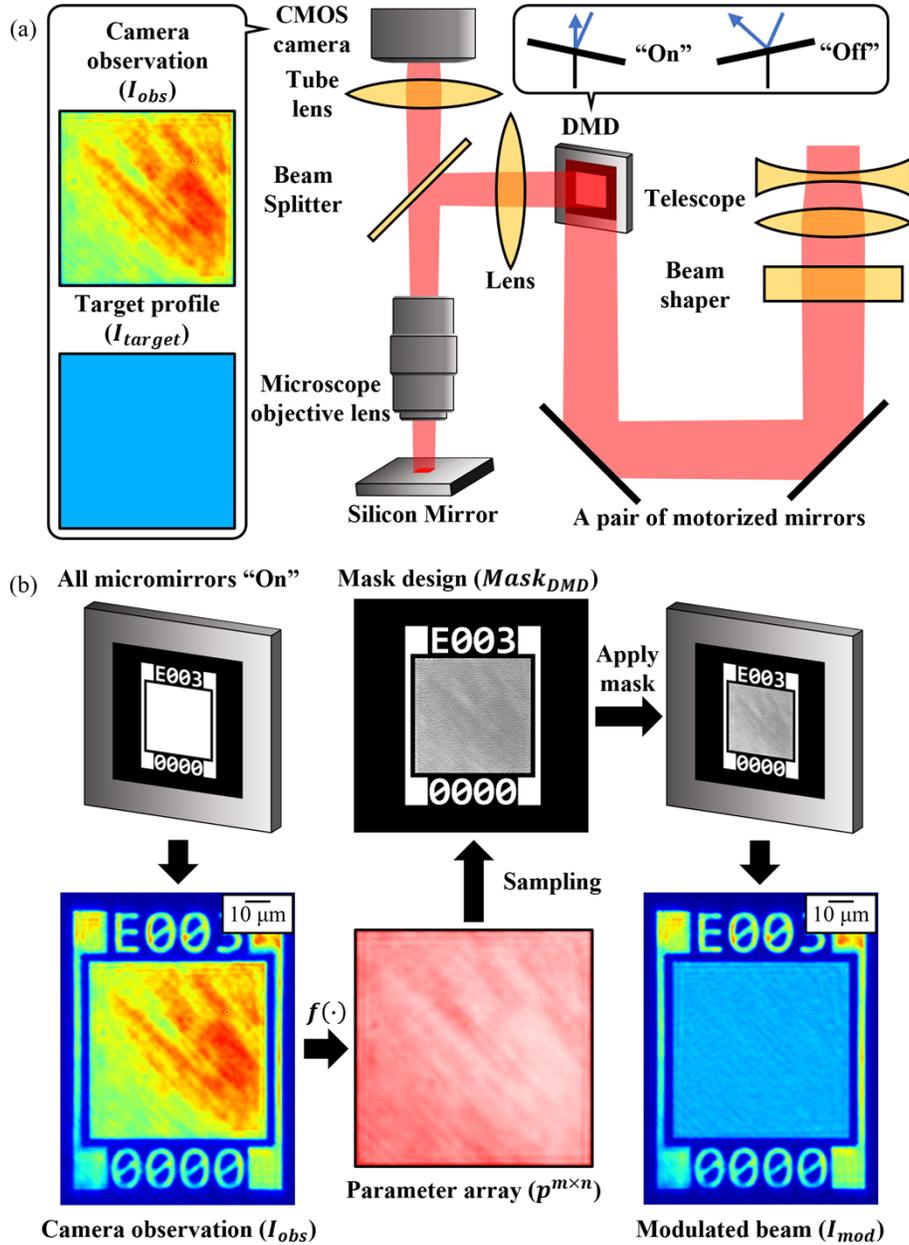
## 2. Method

We begin the introduction to our proposed algorithm with mathematical modeling of a DMD. As previously mentioned, each micromirror can only be in one of two fixed tilt states. This characteristic allows us to conveniently represent each micromirror as a Bernoulli distribution  $B(p)$  (i.e., a discrete probability distribution representing the occurrence of events with exactly two possible outcomes), parameterized by a single parameter  $p \in [0, 1]$ , where  $p$  denotes the probability of setting a micromirror to the “on” state (and naturally  $1 - p$  denotes the probability of setting a micromirror to the “off” state). Consequently, a mosaic comprising  $m \times n$  ( $m, n \in \mathbb{Z}^+$ , i.e., positive integers) micromirrors (i.e., a DMD) can be described as an  $m \times n$  array of independent Bernoulli distributions  $B(\cdot)^{m \times n}$ , parameterized by an  $m \times n$  array of  $p^{m \times n}$ .

Whilst it may seem at odds with the experimental arrangement, where light from neighboring micromirrors is combined, treating the individual micromirrors as independent, learnable probability distributions offers important practical advantages: It greatly simplifies the problem, and it allows the incorporation of stochasticity which could conceivably improve solution discovery by improving exploration [18]. The stochastic nature of sampling from probability distributions allows for the accidental discovery of optimal combinations of micromirror states through a sufficient number of trials, thereby facilitating a more effective search for solutions.

Upon establishing the mathematical model of a DMD, the problem of modulating the intensity and spatial intensity profile using a DMD can be formulated as follows: In a spatial frequency bandwidth limited system where the reflected beams from multiple neighboring micromirrors overlap, given an incident beam characterized by a spatial intensity profile  $I_{obs}$ , the objective is to optimize a function  $f(\cdot)$  that maps this intensity observation to an  $m \times n$  output array of parameters  $p^{m \times n} = f(I_{obs})$ . This output parameter array parameterizes an  $m \times n$  array of Bernoulli distributions, from which the state of each individual micromirror can be determined via sampling. Through sampling from each Bernoulli distribution, an  $m \times n$  array of binary values  $Mask_{DMD}$  is obtained, where each binary value (0 or 1) determines the state of the micromirror in the corresponding position (“off” or “on”). Upon displaying the DMD mask design  $Mask_{DMD}$  on the DMD, the incident beam  $I_{obs}$  is modulated, resulting in a beam with a spatial intensity profile  $I_{mod}$  that possesses desired, predetermined characteristics. This adaptive intensity modulation process is visualized in Fig. 1(b). It should be noted that  $I_{obs}$  represents the intensity distribution observed when all micromirrors are set to their “on” state. By switching a portion of these micromirrors to their “off” state to modulate  $I_{obs}$  to  $I_{mod}$ , some of the light is redistributed to other diffraction orders, thereby reducing the total intensity of  $I_{mod}$  compared to  $I_{obs}$ .

Our approach leverages the Proximal Policy Optimization (PPO) algorithm [19], derived from Deep Reinforcement Learning (DRL) [20], to optimize the mapping function  $f(\cdot)$ . In conjunction



**Fig. 1.** (a) Schematic of the experimental setup. (b) Flowchart of the proposed adaptive DMD intensity modulation process.

with the PPO algorithm, we employ a fully convolutional neural network [21] to represent the mapping function  $f(\cdot)$ , facilitating the control of a substantial number of micromirrors with modest computational expense.

Reinforcement learning optimizes the mapping function  $f(\cdot)$  by maximizing a performance metric known as a reward. The reward metric numerically evaluates the efficacy of the output produced by  $f(\cdot)$  in achieving a specified objective. In this work, the reward function is defined as the sum of the negative pixel-wise absolute difference between  $I_{mod}$  and  $I_{target}$ , mathematically represented as:

$$\mathcal{R}(I_{mod}, I_{target}) := - \sum_i \sum_j |I_{mod}^{(i,j)} - I_{target}^{(i,j)}|$$

Here,  $I_{target}$  denotes the target spatial intensity profile possessing the desired, predetermined characteristics. Maximizing  $\mathcal{R}(I_{mod}, I_{target})$  minimizes the discrepancy between  $I_{mod}$  and  $I_{target}$ . This optimization incentivizes  $f(\cdot)$  to produce  $p^{m \times n}$ , which parameterizes an array of Bernoulli distributions that, upon sampling, generate DMD mask designs. These DMD mask designs, when applied to the DMD, in turn, lead to modulated spatial intensity profiles  $I_{mod}$  that more closely resemble the target intensity  $I_{target}$ .

The mapping function  $f(\cdot)$ , optimized by the PPO algorithm, is represented by a specialized Convolutional Neural Network (CNN). Over the past decade, CNNs have emerged as the dominant paradigm in computer vision [22–26]. A pivotal feature underpinning their efficacy is “restricted connectivity” [27]. This feature mandates that only a localized receptive field of spatially contiguous pixels within a given layer contributes inputs to the computation of a corresponding output in the subsequent layer. This constraint is based on the premise that salient information within real-world imagery tends to be spatially concentrated. Consequently, restricting the input region in accordance with spatial locality reduces the computational demands inherent in processing imagery data. In this work, the DMD is purposefully placed in a spatial frequency bandwidth limited system, where the diffraction limit causes the reflected beams from neighboring micromirrors to overlap, creating spatial locality similar to that observed in real-world imagery data. The fully convolutional neural network employed to represent the mapping function  $f(\cdot)$  consists solely of convolutional layers throughout its architecture, ensuring that “restricted connectivity” is always enforced. This design allows us to demonstrate individual control of 65,536 ( $256 \times 256$ ) micromirrors in a DMD.

The models were trained on a desktop workstation equipped with an *NVIDIA Quadro P6000* GPU for one million steps (i.e.,  $I_{obs}$ ), which required approximately 48 hours. The model weights from the final training steps were saved for the DRL models. Inference time is hardware-dependent; for reference, the mapping  $f : I_{obs} \rightarrow p^{256 \times 256}$  takes  $0.0133 \pm 0.0011$  seconds on a high-performance GPU (*NVIDIA RTX A6000*) and  $0.2489 \pm 0.0001$  seconds on a high-performance CPU (*Intel Xeon w5-3423*). Currently, the sampling algorithm, which is discussed later, is implemented to run only on a CPU, with a computation time of  $0.1319 \pm 0.0174$  seconds on the same high-performance CPU (*Intel Xeon w5-3423*).

The implementation details for the software are provided in the [Supplement 1](#).

### 3. Physical and virtual environment

We first present our results derived from a simulated environment, henceforth referred to as the “gym”, which approximates our real-world experimental setup. This simulation specifically models the diffraction limit imposed by the aperture of the microscope objective. However, in our experimental setup, additional factors such as the apertures of input lenses, mirrors, and beam splitters could also contribute to spatial frequency clipping. We then demonstrate that the DRL models, trained within the gym, can be effectively and seamlessly applied to the real-world experimental setup despite these approximations.

In our experimental setup, depicted in Fig. 1(a), a pulsed laser with a duration of 190 fs, repetition rate of 6 kHz, energy output of up to 1 mJ, and central wavelength of 1030 nm, emitted from a *Light Conversion Pharos SP*, was expanded, and collimated via a telescope before being directed onto a *TOPAG GTH-S-5.0-16.0* beam shaper. This beam shaper transformed the spatial intensity profile of the input beam from a Gaussian circular shape to a nominally top-hat square shape, however, additional random fringes were also observed. These fringing patterns arose from aberrations caused by intentionally deviating the incident beam from its ideal optical axis and from multiple reflections between optical elements. The top-hat, square-shaped beam was then directed onto a *Texas Instruments DLP 7000* DMD via a pair of motorized mirrors. The DMD modulated the incident beam in accordance with its displayed DMD mask design, which was controlled by a desktop workstation. The motorized mirrors adjusted the angle and position of the incident beam relative to the DMD, allowing the incident beam to exhibit varying fringing patterns on the DMD output. The modulated light exiting the DMD was imaged onto a silicon mirror target through a lens, a beam splitter, and a 50× microscopic objective (*Mitutoyo M Plan Apo NIR 50×*, NA 0.42). The resultant reflected light was then captured in real-time via the microscope's tube lens and a *Thorlabs CS165MU/M* CMOS camera. The experimental setup was controlled by a proprietary, multi-threaded program running on a desktop workstation, whilst the DRL models were hosted on a separate workstation. These two workstations were connected via Ethernet in a server-client architecture. The experimental workstation (i.e., client) transmitted input data (i.e., observations of the incident beam) to the DRL workstation (i.e., server), which then processed the input to generate a DMD mask design and sent the result back to the experimental workstation. Given the Ethernet-based communication and the multi-threaded design of the control software, accurately profiling the operational timing of individual steps is challenging. Nonetheless, the full process of beam profile optimization, including incident beam observation, data transfer between workstations, model inference, application of the DMD mask, and optimized beam observation, took approximately one second.

This experimental setup was designed for laser machining and thus was equipped with a pulsed femtosecond laser source. Ultrashort laser machining is, however, a highly complex process, where laser-induced material removal can exhibit variations due to non-linearities [28,29]. Consequently, in this proof-of-concept work we will concentrate on the modulated intensity profiles as observed on the camera, rather than on laser machining depth and uniformity, although this is certainly an area that we aim to explore in future research.

The gym approximates the experimental setup by simulating the spatial frequency clipping of the experimental optical system:

$$I_{mod} = IFFT(Mask_{pupil}(NA, \lambda) \cdot FFT(Mask_{DMD} \cdot I_{obs}))$$

where  $NA$  and  $\lambda$  denote the numerical aperture of the microscopic objective and the central wavelength of the incident beam, respectively. This mathematical expression can be conceptualized as follows: The pupil mask function (i.e.,  $Mask_{pupil}(\cdot)$ ) represents the restricted aperture of the optical system, which acts as a low-pass filter that eliminates high spatial frequency components from the modulated spatial intensity profile of the incident beam exiting the DMD (i.e.,  $Mask_{DMD} \cdot I_{obs}$ ). In both the gym and the experimental setup, the pitch size of an individual micromirror is approximately 13.68  $\mu\text{m}$ , and so the image size of the same micromirror under ideal 50× magnification would be approximately 0.27  $\mu\text{m}$ . However, the real microscope objective has a diffraction limit of approximately 1.23  $\mu\text{m}$ , which exceeds the size of the ideal image of the micromirror. This means each individual micromirror cannot be resolved after imaging via the real microscope objective. Nevertheless, by adaptively manipulating the states of neighboring micromirrors it is possible to modulate the intensity of an incident beam with greater control than would be possible if each binary micromirror exactly matched the diffraction limit. In this configuration, each output intensity “pixel” is approximately controlled by 16 neighboring

micromirrors (i.e.,  $4 \times 4$ ). In the gym environment, the mapping  $Mask_{DMD}, I_{obs} \mapsto I_{mod}$  does not include any hyperparameters derived from experimental observations. The calibration between the gym environment and the experimental setup is performed by optimally aligning the experimental apparatus to minimize discrepancies between the experimentally observed intensity profiles and their simulated counterparts, given identical DMD mask designs.

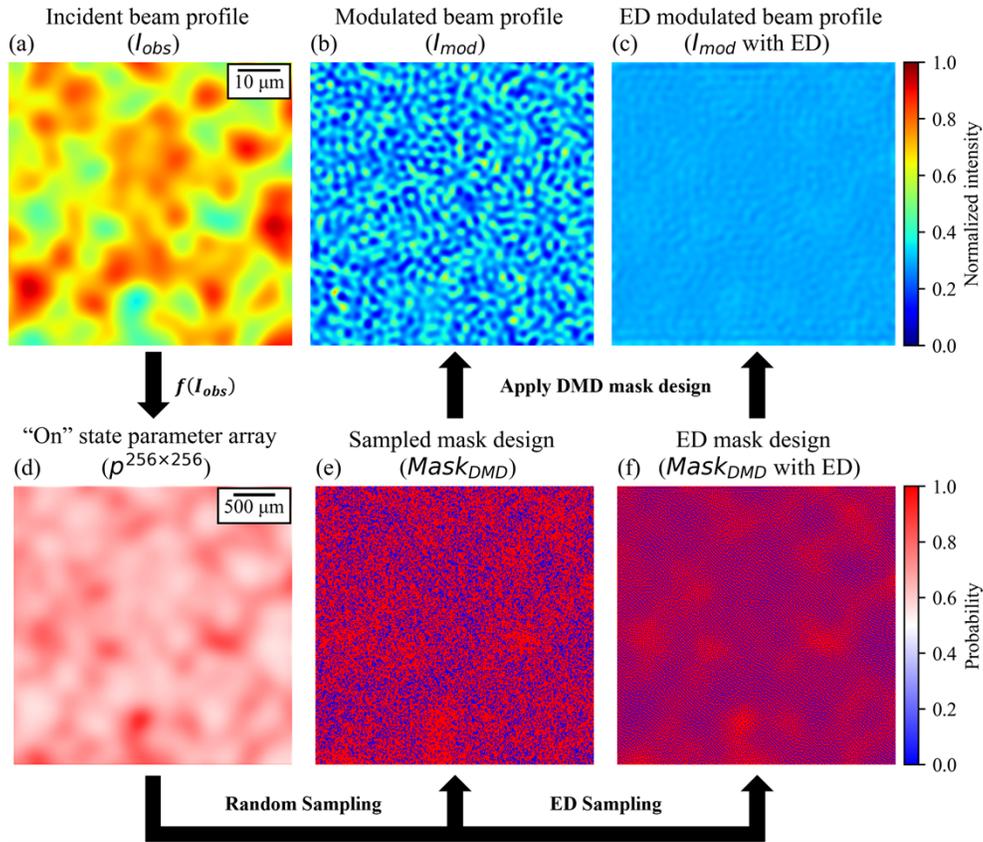
The gym serves as a “digital twin” to our experimental setup, meaning it closely approximates the experimental conditions. This allows a model trained within the gym to be directly applied to the experimental setup without requiring additional training. The simulated “experience” provided by the gym enables the DRL models to develop two capabilities: (1) modulating the intensity of an incident beam to a predetermined value, which is discussed in the first part of the following section, and (2) modulating the spatial intensity profile of an incident beam to a predefined shape, which is discussed in the second part of the following section.

#### 4. Simulation results and discussion

In each training iteration, the gym generated 16 spatial intensity profiles  $I_{obs}$  with dimensions  $512 \times 512$ , each exhibiting randomized intensity variations. These intensity variations were generated by applying Gaussian smoothing to a randomly generated array of values sampled from a uniform distribution. These spatial intensity profiles were then processed through the DRL model (i.e.,  $f(\cdot)$ ) to derive 16 corresponding parameter arrays  $p^{256 \times 256}$  with dimensions  $256 \times 256$ . Next, 16 arrays of Bernoulli distributions  $B(\cdot)^{256 \times 256}$ , parameterized by the derived 16 arrays of  $p^{256 \times 256}$ , were sampled to generate 16 binary DMD mask designs  $Mask_{DMD}$ , each with dimensions  $256 \times 256$ . It is important to note that the probability arrays  $B(\cdot)$  were unique to each micromirror and were independent of the states of other neighboring micromirrors. These binary mask designs were then fed back into the gym to modulate the spatial intensity profiles of their corresponding incident beams  $I_{obs}$ , resulting in modulated spatial intensity profiles  $I_{mod}$ . The reward function  $\mathcal{R}(\cdot)$  was employed to evaluate the resemblance between each modulated spatial intensity profile  $I_{mod}$  and a predefined target spatial intensity profile  $I_{target}$ , which was, in this case, a uniform intensity profile having the desired value. This process optimized the DRL model to more accurately map parameter arrays that produce modulated spatial intensity profiles closely matching the target. Figure 3(a) presents a flowchart visualizing the RL training steps.

Figure 2 illustrates a simulation of DMD intensity modulation facilitated by a trained DRL model. Figure 2(a) shows the spatial intensity profile of an incident beam with random intensity variations (i.e.,  $I_{obs}$ ). Figure 2(d) shows the parameter array (i.e.,  $p^{256 \times 256}$ ) mapped from Fig. 2(a) by the DRL model (i.e.,  $f(I_{obs})$ ). Figure 2(e) shows a binary DMD mask design sampled from Fig. 2(d) (i.e.,  $Mask_{DMD}$ ). Figure 2(b) shows the spatial intensity profile (i.e.,  $I_{mod}$ ) that results from modulating the incident beam in Fig. 2(a) with the DMD mask shown in Fig. 2(e). The target spatial intensity profile  $I_{target}$  is a square shape, with the intensity uniformly set to 30% of the maximum normalized intensity across the entire spatial profile.

From Fig. 2(b), it is evident that intensity variations persisted in the modulated intensity profiles. These variations could be largely attributed to the independence of each Bernoulli distribution in the previously described  $B(\cdot)^{256 \times 256}$ . Let  $q$  neighboring micromirrors, which collectively determine the intensity of an output pixel, be represented by  $q$  Bernoulli distributions  $B(p)^q$ . Sampling from these independent and identical distributions once is equivalent to performing  $q$  trials on  $B(p)$ . Given the limited number of these micromirrors (determined by the experimental configuration,  $q \sim 16$ ), the law of large numbers cannot ensure, for  $q$  trials on  $B(p)$ , the convergence to the summed expectation. Here, the “summed expectation” can be interpreted as the expected number of neighboring micromirrors set to the “on” state (i.e.,  $\sum_q \mathbb{E}[B(p)] = \sum_q (1 \cdot p + 0 \cdot (1 - p)) = q \cdot p$ ), which in turn determines the intensity of the output “pixel”, as previously discussed.



**Fig. 2.** DMD intensity modulation simulation. (a) The spatial intensity profile of the incident beam. (b) The spatial intensity profile after DMD modulation using the sampled DMD mask shown in (e). (c) The spatial intensity profile after DMD modulation using the Error Diffusion (ED) DMD mask shown in (f). (d) The array of learned probabilities that set micromirrors to the “on” state. (e) The DMD mask design sampled from (d). (f) The DMD mask design sampled from (d) using ED.

To address this sampling issue, we adopt an alternative sampling method, Error Diffusion (ED), for binary quantization of the parameter array [13–15,30–32]. ED sampling determines the state of a micromirror based on the corresponding parameter in the parameter array, whilst also considering the states of neighbouring pixels. This approach ensures better adherence to the learned local statistics, resulting in a more uniform intensity distribution.

Sampling from a Bernoulli distribution binarizes a continuous parameter  $p \in [0, 1]$  into a binary state  $x \in \{0, 1\}$ , where  $P(x = 1) = p$  and  $P(x = 0) = 1 - p$ . For a parameter array  $p^{m \times n}$ , where each parameter  $p^{(i,j)}$  belongs to the set  $\{p^{(i,j)} | i, j \in \mathbb{Z}^+, 0 < i \leq m, 0 < j \leq n\}$ , sampling directly from the Bernoulli distribution  $B(p^{(i,j)})$  requires that the expected value  $B(p^{(i,j)})$  matches the probability  $p^{(i,j)}$  of the micromirror at position  $(i, j)$  being in the “on” state (i.e.,  $\mathbb{E}[B(p^{(i,j)})] = 1 \cdot p^{(i,j)} + 0 \cdot (1 - p^{(i,j)}) = p^{(i,j)} \in [0, 1]$ ). However, since micromirrors can only be either fully “on” or “off”, this is not physically possible. To overcome this, we consider a group of neighboring micromirrors that together determine the intensity of an output pixel due to diffraction effects. The objective then becomes assigning binary states  $x \in \{0, 1\}$  to these micromirrors such that  $\sum_i \sum_j x^{(i,j)} \approx \sum_i \sum_j \mathbb{E}[B(p^{(i,j)})]$  is satisfied, which serves as the basis of the ED algorithm.

Consider a one-dimensional parameter array  $p^{m \times 1}$ . The binarization of its first element  $p^{(1,1)}$  is performed by comparing it to a universal threshold variable  $\vartheta \in [0, 1]$ , defined as:

$$ED(p^{(1,1)}) := \begin{cases} 1, & p^{(1,1)} \geq \vartheta \\ 0, & p^{(1,1)} < \vartheta \end{cases}$$

This process introduces a residual error  $\epsilon^{(1,1)}$  given by:

$$\epsilon^{(1,1)} := p^{(1,1)} - ED(p^{(1,1)})$$

To minimize local discrepancies, this residual error is propagated to the neighboring micromirror  $p^{(2,1)}$ , and its binarization is adjusted as:

$$ED(p^{(2,1)}) := \begin{cases} 1, & p^{(2,1)} + \epsilon^{(1,1)} \geq \vartheta \\ 0, & p^{(2,1)} + \epsilon^{(1,1)} < \vartheta \end{cases}$$

Generalizing this recurrence, the binarization of any subsequent parameter  $p^{(i,1)}$  follows:

$$ED(p^{(i,1)}) := \begin{cases} 1, & p^{(i,1)} + \epsilon^{(i-1,1)} \geq \vartheta \\ 0, & p^{(i,1)} + \epsilon^{(i-1,1)} < \vartheta \end{cases}$$

where the residual error  $\epsilon^{(i-1,1)}$  is given by:

$$\epsilon^{(i-1,1)} := p^{(i-1,1)} + \epsilon^{(i-2,1)} - ED(p^{(i-1,1)})$$

At the end of each row, the final residual error  $\epsilon^{(m,1)}$  is discarded. The general residual error is thus:

$$\epsilon^{(i,1)} := p^{(i,1)} + \epsilon^{(i-1,1)} - ED(p^{(i,1)})$$

or equivalently,

$$\epsilon^{(i,1)} - \epsilon^{(i-1,1)} = p^{(i,1)} - ED(p^{(i,1)})$$

Summing the stepwise residual error difference  $\epsilon^{(i,1)} - \epsilon^{(i-1,1)}$  over the entire sequence with  $m$  steps, starting from  $\epsilon^{(1,1)}$ , yields:

$$\sum_{i=1}^m \epsilon^{(i,1)} - \epsilon^{(i-1,1)} = \sum_{i=1}^m p^{(i,1)} - ED(p^{(i,1)})$$

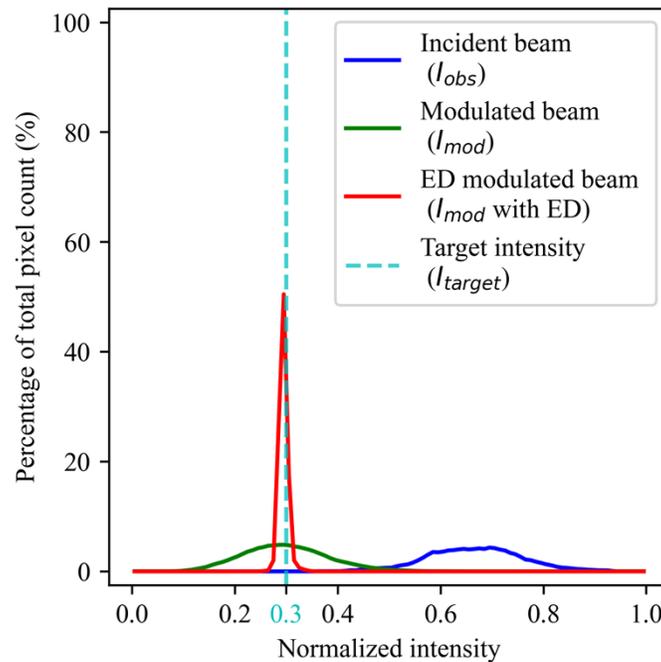
which simplifies to:

$$\epsilon^{(m,1)} - \epsilon^{(0,1)} = \sum_{i=1}^m p^{(i,1)} - \sum_{i=1}^m ED(p^{(i,1)})$$

It can be easily noticed that  $\epsilon^{(i,1)} \in [-1, 1]$ , and  $\epsilon^{(0,1)} = 0$  (i.e.,  $\epsilon^{(1,1)} = p^{(1,1)} - ED(p^{(1,1)}) = p^{(1,1)} + \epsilon^{(0,1)} - ED(p^{(1,1)})$ ), thus for sufficiently large  $\sum_{i=1}^m ED(p^{(i,1)}) \gg |\epsilon^{(m,1)} - \epsilon^{(0,1)}| = 1$ , it follows that:  $\sum_{i=1}^m p^{(i,1)} = \sum_{i=1}^m \mathbb{E}[B(p^{(i,1)})] \approx \sum_{i=1}^m ED(p^{(i,1)})$ . Similarly, summing the stepwise residual error difference over a local sequence with  $k$  steps, starting from  $\epsilon^{(j,1)}$  (i.e.,  $2 \leq j < j+k \leq m$ ), yields  $\sum_{i=j}^k \mathbb{E}[B(p^{(i,1)})] \approx \sum_{i=j}^k ED(p^{(i,1)})$ , provided that  $\sum_{i=j}^k ED(p^{(i,1)}) \gg |\epsilon^{(j+k,1)} - \epsilon^{(j-1,1)}| = 2$ .

The core principle of this ED algorithm is that the residual error arising from the binarization of a continuous variable  $p^{(i,j)} \in [0, 1]$  (i.e., a probability) into a binary value of 0 or 1 is not simply discarded. Instead, it is propagated to the binarization of the neighboring probability  $p^{(i+1,j)}$ . This “error diffusion” helps alleviate the accumulation of errors, thereby preventing the number of neighboring micromirrors set to the “on” state in a local area (which collectively

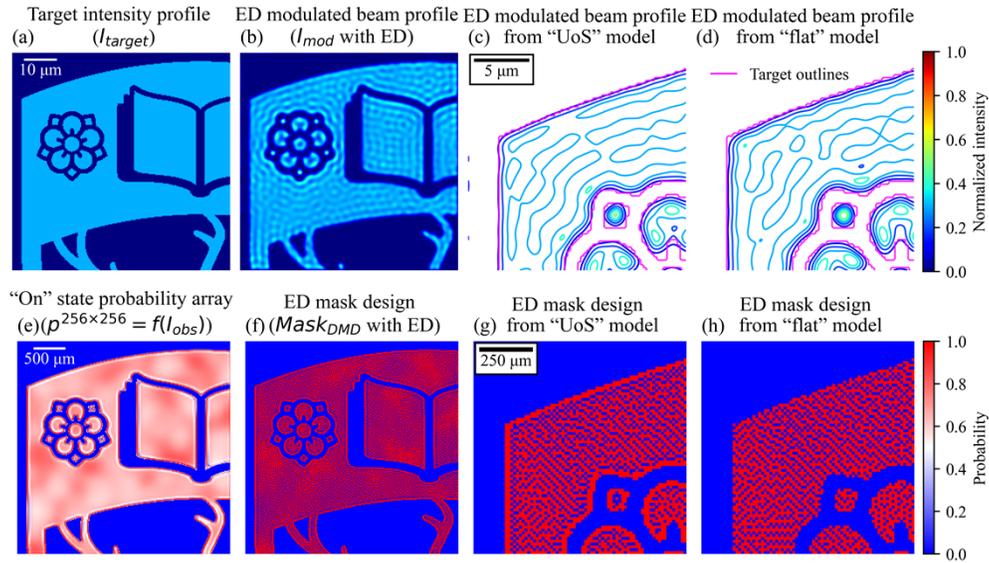




**Fig. 4.** The histograms of the distributions of intensity values for the simulated spatial intensity profile of the incident beam (blue line), the simulated spatial intensity profile modulated using the sampled DMD mask design (green line), and the simulated spatial intensity profile modulated using the ED DMD mask design (red line). The target intensity value is indicated by the cyan dashed line.

intensity for the background is set to 0%. The same spatial intensity profile of the incident beam, illustrated in Fig. 2(a), is applied to the trained “UoS” model, and the resultant probability array is illustrated in Fig. 5(e). Similarly, a DMD mask design can be sampled from this learned probability array via ED sampling, which is presented in Fig. 5(f). Lastly, the spatial intensity profile produced when the incident beam is modulated using the ED sampled DMD mask design is illustrated in Fig. 5(b).

An interesting feature observable from the learned probability array illustrated in Fig. 5(e) is that around the boundaries between the background and the target shape, the “UoS” model tends to increase the probabilities of setting the micromirrors to the “on” state (red). Conversely, just inside the boundary of the target shape the probabilities of setting the micromirrors to the “on” state are decreased (white). A comparison is presented in Fig. 5(g) and (h), where subfigure (g) shows a magnified upper-left portion of the ED DMD mask design from the “UoS” model shown in subfigure (f), and subfigure (h) shows a magnified view of the upper-left portion of the ED mask design from the “flat” model, shown in Fig. 2(e), masked by the shape of the UoS emblem. Figure 5(c) and (d) then show contour maps of the modulated spatial intensity distributions produced from the mask designs shown in Fig. 5(g) and (h), respectively. Comparing these contour maps reveals that the UoS model compensates for distortions in transferred mask shape induced by the diffraction limit and produces a smoother spatial intensity profile around the boundaries [34,35].



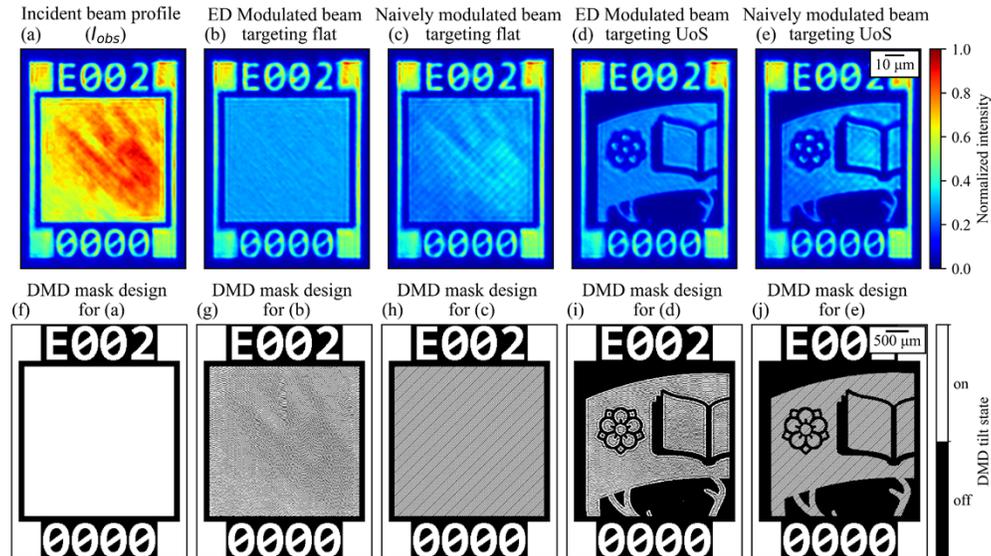
**Fig. 5.** DMD spatial intensity profile modulation simulation for University of Southampton (UoS) emblem target. (a) The upper-left portion of the UoS emblem used as the target spatial intensity profile. (b) The spatial intensity profile after DMD modulation using the ED DMD mask shown in (f). (c) The upper-left portion of (b) in contour plot, overlaid with the outlines of the target shape colorcoded in pink. (d) The upper-left portion of the modulated spatial intensity profile in contour plot, using the ED DMD mask shown in (h), overlaid with the outlines of the target shape colorcoded in pink. (e) The array of learned probabilities that set micromirrors to the “on” state. (f) The DMD mask design sampled from (e) using ED. (g) The upper-left portion of (f), which produces (c). (h) The upper-left portion of the DMD mask design illustrated in Fig. 2(f) masked by the shape of the UoS emblem, which produces (d).

## 5. Experimental results and discussion

The two trained DRL models, namely the “flat” and the “UoS” models, were employed in the experiments without necessitating further training. The experimental procedure mirrored the training steps described earlier, starting with the DRL model mapping the observed spatial intensity profile of an incident beam with random and intricate intensity variations to a parameter array. Subsequently, a binary DMD mask design was obtained by applying ED sampling to this parameter array and this mask design was applied to the DMD to modulate the incident beam. It is noteworthy that the experimental DMD mask design differed slightly from the simulation DMD mask design. The experimental DMD mask comprised two sections: An inner section consisting of  $256 \times 256$  micromirrors, whose tilt states were controlled by the DMD mask design. And an outer section, which served as a “frame” with identification codes and where the micromirrors were always in the “on” state.

Figure 6(a) illustrates the spatial intensity profile of an incident beam with random and intricate spatial intensity variations observed by the camera. Various random intensity patterns could be achieved by controlling the motorized mirrors that moved the beam relative to the DMD as previously detailed. This camera observation was captured by setting all  $256 \times 256$  controllable micromirrors to the “on” state, effectively turning the DMD into a mirror (shown in Fig. 6(f)). The spatial intensity profile of the incident beam, modulated by the  $256 \times 256$  controllable micromirrors in the inner section of the DMD, illuminated an area of approximately  $815 \times 815$  pixels in the camera observation. This illuminated area was cropped and then downsampled and

normalized to a  $512 \times 512$  input image, which was subsequently mapped to two parameter arrays by the two trained DRL models. Two DMD mask designs were generated by performing ED sampling on the obtained parameter arrays, shown in Fig. 6(g) and Fig. 6(i), targeting the square shape and the UoS emblem shape, respectively. The resultant modulated spatial intensity profiles captured by the camera whilst displaying these DMD masks are presented in Fig. 6, with subfigure (b) showcasing the square shape and subfigure (d) showcasing the UoS emblem shape.

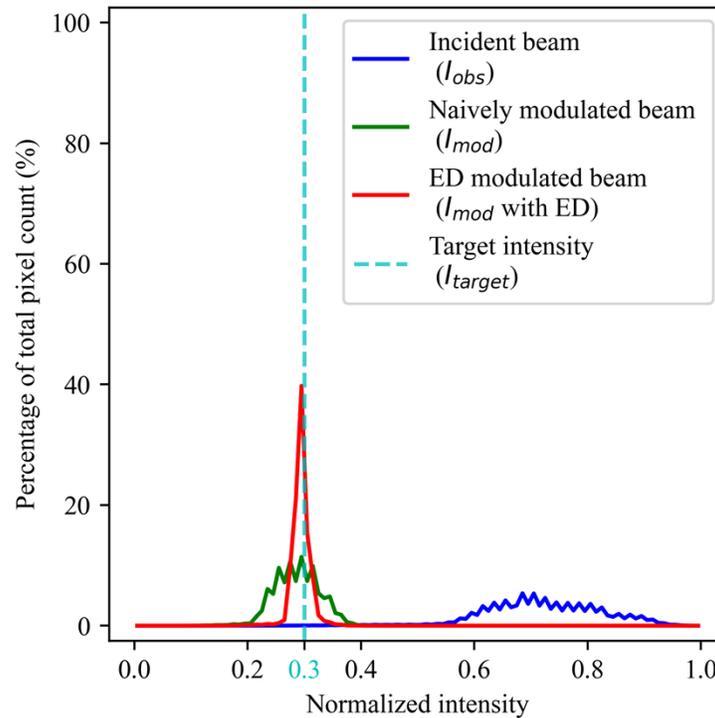


**Fig. 6.** Camera observations of the spatial intensity distributions are presented in the first row, and their corresponding displayed DMD masks are shown in the second row. (a) The observed spatial intensity profile of the incident beam. (b) The spatial intensity profile modulated with the “flat” model via ED sampling, targeting a uniformly distributed square shape. (c) The spatial intensity profile modulated with the naively designed periodic grating pattern shown in (h), targeting a uniformly distributed square shape. (d) The spatial intensity profile modulated with the “UoS” model via ED sampling, targeting a uniformly distributed UoS emblem shape. (e) The spatial intensity profile modulated with the naively designed periodic grating pattern shown in (j), targeting a uniformly distributed UoS emblem shape.

To address the concerns that the improved uniformity in intensity distributions shown in Fig. 6(b) and Fig. 6(d) might simply be due to the decreases in their mean intensities, two additional DMD mask designs were tested, as illustrated in Fig. 6(h) and Fig. 6(j) for the square shape and the UoS emblem shape, respectively. These DMD mask designs comprised periodic checkerboard patterns to equally reduce the intensity across the spatial intensity profile of the incident beam. Figure 6(c) and Fig. 6(e) illustrate the modulated spatial intensity profiles using these naively designed DMD masks for the square shape and the UoS shapes, respectively.

Comparatively, the intensity uniformities for the spatial intensity profiles modulated by the ED DMD masks are superior to those of the naively designed DMD masks. The difference in intensity uniformity is further illustrated in the intensity distribution histograms shown in Fig. 7. The blue curve represents the experimentally observed intensity histogram of the spatial intensity profile of the incident beam, as shown in Fig. 6(a), with a mean of 0.722 and a standard deviation of 0.101. The green curve represents the intensity histogram of the naively modulated spatial intensity profile whose intensities are naively and equally reduced across the spatial intensity profile, as shown in Fig. 6(c), with a mean of 0.287 and a standard deviation of 0.039. Lastly, the red curve represents the intensity histogram of the ED modulated spatial intensity profile, targeting a

uniformly distributed square shape, as shown in Fig. 6(b), with a mean and standard deviation of 0.295 and 0.016, respectively. In terms of the Root-Mean-Square (RMS) measurement of flatness, our single-step approach achieves approximately 1.29% after modulation and 9.00% before modulation in the ‘flat’ case, as shown in Fig. 6(b). To the best of our knowledge, the previously reported best performance is 1.11% after 17 iterative modulation steps and 1.50% before modulation, measured at a wavelength of 1064 nm with an array of  $230 \times 230$  micromirrors [14]. It is important to note that factors such as diffraction-limits-to-micromirror-pitch-size ratio (introducing inherent bit depth errors), camera configuration (affecting imaging errors), and target intensity level (determining the skewness of the modulated intensity distribution) contribute to the differences between these reported values. Therefore, these values are provided for reference only.



**Fig. 7.** Histograms of the distributions of intensity values for the experimental observed spatial intensity profile of the incident beam (blue line), the experimental observed spatial intensity profile modulated by the naively designed DMD mask (green line), targeting the uniform square shape, and the experimental observed spatial intensity profile modulated by the “flat” model and sampled via ED (red line), targeting the uniform square shape. The target intensity value is indicated by the cyan dashed line.

By comparing the experimental intensity distributions shown in Fig. 7 with the simulated intensity distribution shown in Fig. 4, it is evident that the experimentally observed ED modulated beam exhibits slightly less uniform intensity compared to the simulated ED modulated beam. This decrease in performance may stem from the discrepancies between the simulated environment, where the two DRL models were trained, and the experimental setup, to which the two DRL models were applied. Firstly, the intensity variations generated from the Gaussian-smoothed random values during training might not adequately generalize to the experimentally observed intensity variations. The differences in intensity variations can be observed by comparing the intensity distributions of the simulated and experimental incident beams, shown as the

blue curves in Fig. 4 and Fig. 7, respectively. The intensity distribution of the experimentally observed incident beam exhibits periodic spikes, likely resulting from interferences between optical elements, which are not presented in the intensity distribution of the simulated incident beam. Secondly, the gym environment used in simulations only accounts for spatial frequency clipping by the microscope objective, whilst in the experimental setup, other optical elements might also contribute to spatial frequency clipping. Lastly, the simulated environment implicitly assumes that the observed intensity profiles, both before and after modulation, are aberration-free and represented in single-precision floating-point format. However, in the experimental setup, slight misalignments of optical instruments, after the DMD, can introduce aberrations, potentially distorting camera observations and causing artifacts which cannot be corrected by the DMD (as these distortions and artifacts occur after DMD modulation). Furthermore, the camera used to capture the intensity profile has a bit depth of 10 (i.e., 1024 grayscale levels), which may introduce quantization errors when representing the continuous intensity profiles.

Although the use of a gym environment facilitates the training of DRL models without requiring time-consuming data collection in the experimental setup, it also introduces different complexities owing to the discrepancies between the simulated and experimental environments. The data-driven nature of the proposed algorithm necessitates a substantial number of high-quality data to be effective. Therefore, our future work will focus on large-scale automated collection of experimental data and methods of efficiently transferring learned experience from gym to real-world environments.

## 6. Conclusion

This paper introduces a novel control algorithm for DMD that can, in a single step, simultaneously and adaptively modulate both the overall intensity and the spatial intensity profile of an incident beam that has been aberrated with random and intricate spatial intensity variations. The paper begins with a unique perspective of treating each micromirror in a DMD as an independent Bernoulli distribution, characterized by a learnable parameter. This unique perspective leads to the development of a computationally efficient learning paradigm, enabling individual control of each micromirror. Critically, we apply error diffusion sampling to the learned parameter array which results in DMD mask designs that precisely and accurately modulate the intensity and spatial intensity profile to predefined targets. The effectiveness of our proposed control algorithm is demonstrated in both simulated and experimental environments.

**Funding.** Engineering and Physical Sciences Research Council (EP/T026197/1, EP/W028786/1).

**Disclosures.** The authors declare no conflicts of interest.

**Data availability.** Data underlying the results presented in this paper (including the DRL models) are available in Ref. [36].

**Supplemental document.** See [Supplement 1](#) for supporting content.

## References

1. L. J. Hornbeck, "Digital Light Processing(TM) for high-brightness, high-resolution applications," *P Soc Photo-Opt Ins* **3013**, 27–40 (1997).
2. S. Turtaev, I. T. Leite, K. J. Mitchell, *et al.*, "Comparison of nematic liquid-crystal and DMD based spatial light modulation in complex photonics," *Opt. Express* **25**(24), 29874–29884 (2017).
3. "Texas Instruments DLP7000 technical documentation," <https://www.ti.com/product/DLP7000>.
4. B. Mills, D. J. Heath, M. Feinaeugle, *et al.*, "Laser ablation via programmable image projection for submicron dimension machining in diamond," *J. Laser Appl.* **26**(4), 041501 (2014).
5. D. J. Heath, B. Mills, M. Feinaeugle, *et al.*, "Rapid bespoke laser ablation of variable period grating structures using a digital micromirror device for multi-colored surface images," *Appl. Opt.* **54**(16), 4984–4988 (2015).
6. D. J. Heath, J. A. Grant-Jacob, M. Feinaeugle, *et al.*, "Sub-diffraction limit laser ablation via multiple exposures using a digital micromirror device," *Appl. Opt.* **56**(22), 6398–6404 (2017).
7. D. J. Heath, J. A. Grant-Jacob, R. W. Eason, *et al.*, "Single-pulse ablation of multi-depth structures via spatially filtered binary intensity masks," *Appl. Opt.* **57**(8), 1904–1909 (2018).

8. C. Sun, N. Fang, D. M. Wu, *et al.*, "Projection micro-stereolithography using digital micro-mirror dynamic mask," *Sens. Actuators, A* **121**(1), 113–120 (2005).
9. R. C. Y. Auyeung, H. Kim, S. Mathews, *et al.*, "Spatially modulated laser pulses for printing electronics," *Appl. Opt.* **54**(31), F70–F77 (2015).
10. Y. Lu, G. Mapili, G. Suhali, *et al.*, "A digital micro-mirror device-based system for the microfabrication of complex, spatially patterned tissue engineering scaffolds," *J. Biomed. Mater. Res., Part A* **77a**(2), 396–405 (2006).
11. D. Akbulut, T. J. Huisman, E. G. van Putten, *et al.*, "Focusing light through random photonic media by binary amplitude modulation," *Opt. Express* **19**(5), 4017–4029 (2011).
12. S. A. Goorden, J. Bertolotti, A. P. Mosk, *et al.*, "Superpixel-based spatial amplitude and phase modulation using a digital micromirror device," *Opt. Express* **22**(15), 17999–18009 (2014).
13. J. Liang, R. N. Kohn Jr., M. F. Becker, *et al.*, "1.5% root-mean-square flat-intensity laser beam formed using a binary-amplitude spatial light modulator," *Appl. Opt.* **48**(10), 1955–1962 (2009).
14. J. Y. Liang, R. N. Kohn Jr., M. F. Becker, *et al.*, "High-precision laser beam shaping using a binary-amplitude spatial light modulator," *Appl. Opt.* **49**(8), 1323–1330 (2010).
15. J. Y. Liang, R. N. Kohn, M. F. Becker, *et al.*, "Homogeneous one-dimensional optical lattice generation using a digital micromirror device-based high-precision beam shaper," *J. Micro/Nanolith. MEMS MOEMS* **11**(2), 023002 (2012).
16. J. B. Sampell, "Digital Micromirror Device and Its Application to Projection Displays," *J. Vac. Sci. Technol. B* **12**(6), 3242–3246 (1994).
17. S. K. Nayar, V. Branzoi, T. E. Boulton, *et al.*, "Programmable imaging using a digital micromirror array," *Proc Cvpr Ieee*, 436–443 (2004).
18. R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction* (MIT press, 2018).
19. J. Schulman, F. Wolski, P. Dhariwal, *et al.*, "Proximal policy optimization algorithms," *arXiv*, (2017).
20. V. Mnih, K. Kavukcuoglu, D. Silver, *et al.*, "Human-level control through deep reinforcement learning," *Nature* **518**(7540), 529–533 (2015).
21. J. Long, E. Shelhamer, T. Darrell, *et al.*, "Fully convolutional networks for semantic segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, (2015), 3431–3440.
22. K. M. He, X. Zhang, S. Ren, *et al.*, "Deep Residual Learning for Image Recognition," *2016 Ieee Conference on Computer Vision and Pattern Recognition (Cvpr)*, 770–778 (2016).
23. J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," *arXiv*, (2018).
24. Z. Liu, H. Mao, C.-Y. Wu, *et al.*, "A ConvNet for the 2020s," *2022 Ieee/Cvf Conference on Computer Vision and Pattern Recognition (Cvpr)*, 11966–11976 (2022).
25. J. A. Grant-Jacob, M. Praeger, M. Loxham, *et al.*, "Lensless imaging of pollen grains at three-wavelengths using deep learning," *Environ. Res. Commun.* **2**(7), 075005 (2020).
26. J. A. Grant-Jacob, M. Praeger, R. W. Eason, *et al.*, "Single-frame 3D lensless microscopic imaging via deep learning," *Opt. Express* **30**(18), 32621–32632 (2022).
27. Y. Lecun, L. Bottou, Y. Bengio, *et al.*, "Gradient-based learning applied to document recognition," *Proc. IEEE* **86**(11), 2278–2324 (1998).
28. J. Bonse, S. Baudach, J. Krüger, *et al.*, "Femtosecond laser ablation of silicon-modification thresholds and morphology," *Appl. Phys. A* **74**(1), 19–25 (2002).
29. E. G. Gamaly, A. V. Rode, B. Luther-Davies, *et al.*, "Ablation of solids by femtosecond lasers: Ablation mechanism and ablation thresholds for metals and dielectrics," *Phys. Plasmas* **9**(3), 949–957 (2002).
30. C. Dorrer and J. D. Zuegel, "Design and analysis of binary beam shapers using error diffusion," *J. Opt. Soc. Am. B* **24**(6), 1268–1275 (2007).
31. C. Dorrer, "High-damage-threshold beam shaping using binary phase plates," *Opt. Lett.* **34**(15), 2330–2332 (2009).
32. M. Praeger, S. Papazoglou, A. Pesquera, *et al.*, "Laser-induced backward transfer of monolayer graphene," *Appl. Surf. Sci.* **533**, 147488 (2020).
33. R. Floyd, "An adaptive algorithm for spatial grey scale," *SID digest*, 1975 (1975).
34. J. P. Stürniman and M. L. Rieger, "Fast Proximity Correction with Zone Sampling," *Optical/Laser Microlithography Vii* **2197**, 294–301 (1994).
35. T. Matsunawa, B. Yu, D. Z. Pan, *et al.*, "Optical Proximity Correction with Hierarchical Bayes Model," *Proc. SPIE* **9426**, 94260X (2015).
36. Y. Xie, M. Praeger, J. A. Grant-Jacob, *et al.*, "Data supporting the publication 'Single-step beam intensity and profile optimization using a 256 × 256 micromirror array and reinforcement learning'," University of Southampton, 2024 <https://doi.org/10.5258/SOTON/D3259>.