

# End-to-end Optimization of Fluidic Lenses

MULUN NA\*, HÉCTOR A. JIMÉNEZ-ROMERO\*, XINGE YANG, JONATHAN KLEIN, DOMINIK L. MICHELS, and WOLFGANG HEIDRICH, KAUST, KSA

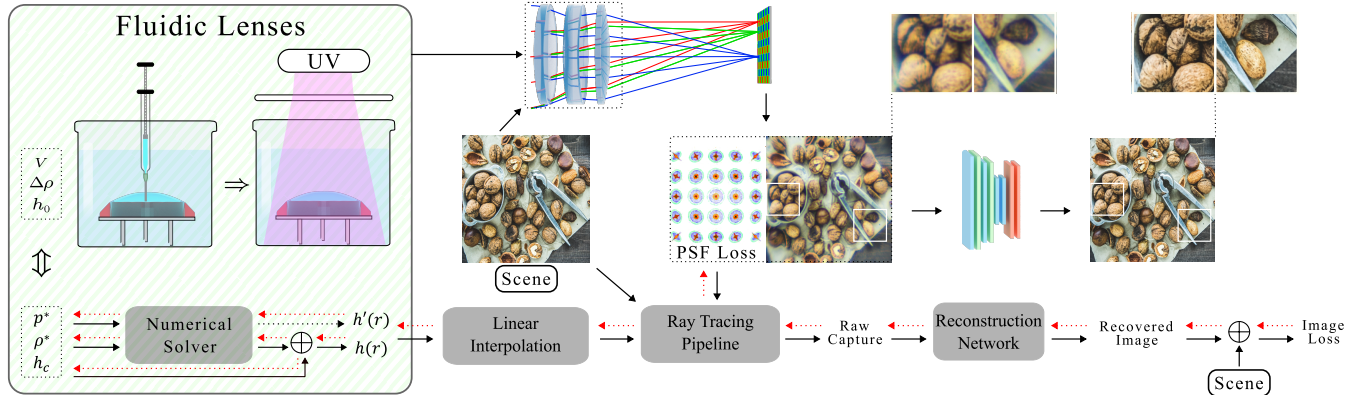


Fig. 1. Our end-to-end design process for cheap and customizable fluidic lenses. Left to right: Lenses are formed from liquid resin that is injected into a ring mold and hardened with UV light. A scene is imaged through the lens via differentiable ray tracing, and subsequently reconstructed through a neural network. Different loss functions are evaluated on the point spread function and finale image and back propagated through the whole pipeline to jointly learn the network parameters and lens design parameters.

Prototyping and small volume production of custom imaging-grade lenses is difficult and expensive, especially for more complex aspherical shapes. Fluidic shaping has recently been proposed as a potential solution: It makes use of the atomic level smoothness of interfaces between liquids, where the shape of the interface can be carefully controlled by boundary conditions, buoyancy control and other physical parameters. If one of the liquids is a resin, its shape can be “frozen” by curing, thus creating a solid optical element. While fluidic shaping is a promising avenue, the shape space generated by this method is currently only described in the form of partial differential equations, which are incompatible with existing lens design processes. Moreover, we show that the existing PDEs are inaccurate for larger curvatures. In this work, we develop a new formulation of the shape space lenses generated by the fluidic shaping technique. It overcomes the inaccuracies of previous models, and, through a differentiable implementation, can be integrated into recent end-to-end optical design pipelines based on differentiable ray tracing. We extensively evaluate the model and the design pipeline with simulations, as well as initial physical prototypes.

CCS Concepts: • **Computing methodologies** → **Computational photography**.

\* Authors contributed equally to this work.

Authors’ address: Mulun Na, mulun.na@kaust.edu.sa; Héctor A. Jiménez-Romero, hector.jimenezromero@kaust.edu.sa; Xinge Yang, xinge.yang@kaust.edu.sa; Jonathan Klein, jonathan.klein@kaust.edu.sa; Dominik L. Michels, dominik.michels@kaust.edu.sa; Wolfgang Heidrich, wolfgang.heidrich@kaust.edu.sa, KAUST, Thuwal, KSA.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SA Conference Papers ’24, December 3–6, 2024, Tokyo, Japan

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1131-2/24/12

<https://doi.org/10.1145/3680528.3687584>

Additional Key Words and Phrases: Lens Design, End-to-End Optimization, Fluid Dynamics, Reconstruction Network

## ACM Reference Format:

Mulun Na, Héctor A. Jiménez-Romero, Xinge Yang, Jonathan Klein, Dominik L. Michels, and Wolfgang Heidrich. 2024. End-to-end Optimization of Fluidic Lenses. In *SIGGRAPH Asia 2024 Conference Papers (SA Conference Papers ’24)*, December 3–6, 2024, Tokyo, Japan. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3680528.3687584>

## 1 INTRODUCTION AND RELATED WORK

Prototyping and small-volume production of custom lenses for imaging applications are both time-consuming and costly, as they necessitate specialized equipment and expertise. Despite the grinding and polishing of spherical lenses having a long-standing history and relatively low complexity, the fabrication of aspherical glass lenses necessitates advanced methods, such as single-point diamond turning, followed by magneto-rheological polishing [Brinksmeier et al. 2007; Kumar et al. 2022], which is not readily accessible. Small plastic lenses like those used in mobile phones can be economically mass-produced by injection molding [Spina et al. 2012], however the mold creation again requires methods like single-point diamond turning.

For this reason, there is a large and growing body of literature trying to develop more accessible techniques for fast optical prototyping. In graphics, Pereira et al. [2014] first proposed to fabricate fiberoptic bundles with a high-end but off-the-shelf polyjet 3D printer. Since then, 3D printing techniques have been refined substantially. On the high end, two-photon polymerization printers have been used to print anything from fiber optics [Berglund et al. 2022] to complete micro-scale lens assemblies [Gissibl et al. 2016]. Unfortunately, this technology is neither fully accessible,

nor does it scale to optical elements with larger form factors. On the low end, researchers have considered the use of resin-based 3D printers or selective laser sintering (SLS) for larger (cm-scale) lenses [Gawedzinski et al. 2017; Berglund et al. 2019; Berglund and Tkaczyk 2022; Aguirre-Aguirre et al. 2023; Christopher et al. 2023; Gonzalez-Utrera et al. 2024]. Unfortunately, the resolution of such additive processes is insufficient to provide optical grade surface finishes without either manual and laborious polishing, or some form of molding for transferring the final shape and surface quality from an existing reference surface. Notably, Akşit et al. [Akşit et al. 2019] have proposed to solve this issue by vacuum-forming a thin layer of optical resin onto a rough lens blank created by 3D printing, thus effectively transferring the smoothness of the vacuum forming sheet onto the lens. However, while improving the surface finish markedly, vacuum forming sheets are themselves not fabricated to optical surface quality, and also control of the final surface shape is limited.

Our work is based on fluidic lens shaping pioneered by Frumkin and Bercovici [2021] and Elgarisi et al [2021], which forms lens surfaces through fluid dynamics. Specifically, a UV-curable resin is immersed in a surrounding liquid, forming an atomically smooth interface, whose shape can be controlled by the relative density of the resin and the surrounding liquid, the amount of resin used, and the lens boundary. After curing, this yields a solid lens with an optical grade finish.

Unfortunately, the original shape model for fluidic shaping [Elgarisi et al. 2021] contains several approximations that lead to inaccuracies intolerable for lens design. As a demonstrating example, for two liquids with the same density (neutral buoyancy), the resulting interface has a spherical shape, but their model can only express parabolic surfaces - thus an accurate fit is impossible. Moreover, because their shape space is only described in the form of a PDE, it has so far not been possible to integrate fluidic shaping into existing optical design pipelines due to computational constraints. We overcome both issues by developing a new shape model that is more accurate, but limited to circular boundaries that are common in lens design. We also apply this model in a differentiable ray tracer that simulates the camera-captured images using the fluidic lens. These simulated images are then input into an image reconstruction network, allowing for the joint design of optics and the network in an end-to-end fashion with gradients back-propagated from the network output. This approach enables a rapid prototyping method for computational lens design with outstanding imaging quality.

Finally, we also show preliminary work on the actual fluidic shaping fabrication process, describing how to control fabrication parameters, as well as achieve independent surface geometries on both sides of the lens.

In summary, our contributions are: (1) we introduce a new, more accurate shape model for fluidic shaping that is amenable to lens design tasks and analyse it in detail; (2) we prove some of its mathematical properties that are important for practical applications; (3) we introduce a new efficient numerical representation of lens surfaces to a differentiable ray tracing pipeline and use it for end-to-end design; (4) we demonstrate simulation results and initial prototype fabricated lenses and their optical performance. While more work remains on automating the fabrication process and controlling the

fabrication conditions to improve repeatability, we believe our work is a major step forward in establishing fluidic shaping as a viable option for rapid and inexpensive prototyping of optical components.

## 2 OVERVIEW

Our lenses are built from liquid, optically clear resin that can be hardened through UV light curing, once it assumes the desired shape. Shown in Fig. 1, the resin is injected into a ring shape (with customized height and diameter), that is closed at the bottom and resides inside a small tank filled with an immersion liquid. The immersion liquid is a mixture of water and glycerol, where the ratio is engineered to give the mixture a precise density. The lens resin inside the ring forms a single, compact shape, whose surface is defined by the interplay of two counteracting physical forces: interfacial tension between the UV resin and the immersion liquid and the buoyancy imbalance caused by their density difference. The shape of this lens depends on 1) the ring geometry, 2) the amount of injected volume, and 3) the density difference. By adjusting these quantities, different lens shapes are realized.

In this fabrication process, the lens surface will be perfectly smooth (much beyond the scale of the optical wavelength [Elgarisi et al. 2021]), without requiring any post processing such as polishing. This is an important advantage of fluidic lens fabrication.

We find optimal values for the 3 parameters for a specified imaging task through an end-to-end optimization pipeline. Our system consists of 3 main components, closer described in the following sections: 1) an accurate, physically-based model for the equilibrium state of fluidic lenses, 2) an analytical end-to-end optimization pipeline for optical lens systems that support non-analytical lens models and reconstruction networks, and 3) the manufacturing process to implement the designed lenses.

A part of the optimization pipeline is the reconstruction network, as shown in Fig. 1. It is distinct from an image post-processing network in that it is trained jointly during the lens parameter optimization stage. The goal of the lens optimization is then no longer to produce the sharpest possible image on the camera sensor but to produce the image from which the maximum amount of information about the scene can be decoded by the network. This enables higher image quality than simply training and applying the network as a post processing step.

While this general approach allows for a wide area of applications, we showcase it here by recreating optimized versions of three different classical lens systems with increasing complexity (single plano-convex lens, single biconvex lens, triplet lens). Single lens systems suffer from artifacts such as chromatic aberration and poor off-axis performance and by optimizing the lens shape, we reduce these.

While this type of fluidic lens fabrication has been demonstrated before, the applied physical model still had some limitations, and no imaging tasks were demonstrated [Elgarisi et al. 2021]. Here, we build upon and improve the mathematical model to reach the accuracy required for imaging. While the implementation of the lens manufacturing itself comes with its own set of problems, we show initial imaging results and outline future research directions on the practical side.

### 3 FLUIDIC LENS FORMATION MODEL

At the heart of the optimization is the mapping between the design parameters and the resulting lens shapes. While this is in principle described by classical fluid dynamics, we are only interested in the equilibrium state that the system converges towards after the injection. The ring has a radius of  $R_0$  and a height of  $h_0$ . At its center, the lens has a thickness of  $h_c$ . The immersion liquid has a density of  $\rho_{im}$ , the resin has a density of  $\rho_r$ , and the effective density difference is  $\Delta\rho := \rho_{im} - \rho_r$ . The radial symmetric lens surface is generated through revolving a generatrix function  $h(r)$ ,  $r \in [0, R_0]$  around the height axis. The total resin volume injected is  $V$ . An overview table of all physical units is given in the supplemental material.

#### 3.1 Surface from System Energy

Following Elgarisi et al. [2021], the energy of the above described system is given by

$$\Pi = \Pi_S + \Pi_G + \Pi_C, \quad (1)$$

where  $\Pi_S$  denotes the surface energy,  $\Pi_G$  denotes the energy from the buoyancy imbalance, and  $\Pi_C$  is a Lagrange multiplier for the volume constraint.

**3.1.1 Surface energy.** The energy  $\Pi_S$  is a measurement of the energy required to break the intermolecular bonds when the surface between the liquids changes. It can be obtained by integrating the (relative) surface energy  $\gamma$  (defined as work per area) over the surface  $S$  in polar coordinates, where  $h_r$  is the first derivative of  $h(r)$ :

$$\Pi_S = \iint_S \gamma \sqrt{1 + h_r^2} r \, dr \, d\theta. \quad (2)$$

**3.1.2 Buoyancy energy.** The gravitational pull acting on the optical resin is a function of the Earth's gravity  $g$  and the mass of the resin, scaling with its density. The surrounding liquid exerts an upward force equal to the weight of the liquid displaced. Depending on the density difference  $\Delta\rho$ , this resultant force will either cause the resin to rise or remain at the bottom of the container. Neutral buoyancy can be achieved by matching the density of the surrounding liquid with that of the optical resin. The volume  $V$  is a function of  $h(r)$ , where  $V(h) = \iint_S h r \, dr \, d\theta$ . Therefore, we have:

$$\Pi_G = \int \Delta\rho g V(h) dh = \iint_S \frac{1}{2} \Delta\rho g h^2 r \, dr \, d\theta. \quad (3)$$

**3.1.3 Free system energy.** Substitution of Eq. (2) and Eq. (3) into Eq. (1), we obtain the free energy of the system

$$\Pi = \iint_S \left( \gamma \sqrt{1 + \left(\frac{dh}{dr}\right)^2} + \frac{1}{2} \Delta\rho g h^2 + \lambda h \right) r \, dr \, d\theta, \quad (4)$$

under a fixed volume constraint

$$\Pi_C = \lambda \iint_S h r \, dr \, d\theta, \quad (5)$$

which is enforced with a Lagrange multiplier  $\lambda$ . We denote the integrand of Eq. (4) with

$$F(r) = \left( \gamma \sqrt{1 + h_r^2} + \frac{1}{2} \Delta\rho g h^2 + \lambda h \right) r \quad (6)$$

to formulate the Euler-Lagrange equation  $\frac{\partial F}{\partial h} = \frac{d}{dr} \frac{\partial F}{\partial h_r}$ , by which we can obtain a second-order ordinary differential equation (ODE) of our system:

$$\lambda r + \Delta\rho g h r = \frac{\gamma(1 + h_r^2)(h_r + r h_{rr}) - \gamma r h_{rr} h_r^2}{(1 + h_r^2)^{\frac{3}{2}}}, \quad (7)$$

where  $h_{rr}$  represents the second-order derivative.

#### 3.2 Solving the Surface ODE

Solving the ODE in Eq. (7) will give us the generatrix  $h(r)$  of the fluidic-freeform surface. We introduce two parameters:

$$\begin{cases} \rho^* = \frac{\Delta\rho g}{\gamma}, \\ p^* = \frac{\lambda}{\gamma}. \end{cases} \quad (8)$$

Since the lens surface should be continuously differentiable, we have  $h_r(0) = 0$  and, from the assumption that the resin volume is sufficient to touch the ring edge, but not too large to overflow, we have  $h(R_0) = h_0$ .

$$\begin{cases} h_{rr} = (\rho^* \cdot h + p^*)(1 + h_r^2)^{\frac{3}{2}} - \frac{h_r(1 + h_r^2)}{r}, \\ h(R_0) = h_0, \\ h_r(0) = 0. \end{cases} \quad (9)$$

**3.2.1 General solution.** In Eq. (9), the initial values for  $h$  and  $h_r$  are given for two different positions, which does not match the structure required for standard numerical solvers. We show, however, that the solution for Eq. (9) is *shift-invariant*, i.e., the shape of the surface does not depend on the value of  $h_c$  if the injected volume is adjusted appropriately. Please, see the supplementary material for the derivation of the shift-invariance.

Thus, we formulate the shifted problem  $h^*$  with the initial values  $h^*(0) = 0$  and  $h_r^*(0) = 0$ , that we can solve using any numerical solver:

$$h^*(\rho^*, p^* | r) = \text{SOLVER}(\rho^*, p^*). \quad (10)$$

The solution for  $h(r)$  is then obtained through  $h(r) = h^*(r) + h_c$ .

To enable fluidic lens design in an end-to-end fashion, it's crucial to know the derivatives of  $h$  with respect to the design parameters  $\rho^*$ ,  $p^*$ , and  $h_c$ . To address this challenge, we utilize a differentiable numerical solver from `torchdiffeq` [Chen 2018], which is a PyTorch library of ODE solvers. We use the default solver `dopri5`, Dormand–Prince 5(4) [Dormand and Prince 1980], which is an embedded Runge-Kutta numerical method with adaptive step size control and uses the adjoint sensitivity method to compute gradients for memory efficiency. This method solves an additional ODE in a reverse direction to compute the gradients, which avoids storing intermediate states and thus keeps memory usage low.

**3.2.2 Close form solution when  $\rho^* = 0$ .** For neutral buoyancy (i.e. when  $\Delta\rho = 0$ , leading to  $\rho^* = 0$ ), ODE in Eq. (9) has an analytical solution in the shape of a spherical surface. We find that

$$h(r) = \frac{\frac{p^*}{2} r^2}{1 + \sqrt{1 - \left(\frac{p^*}{2}\right)^2}}. \quad (11)$$

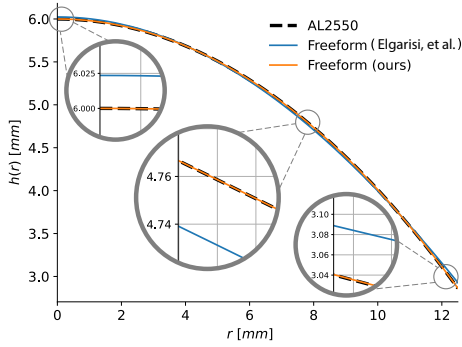


Fig. 2. Comparison of geometrical fittings of aspherical lens *AL2550*.




Fig. 3. Parameter space exploration of Eq. 10. The stable parameter region is highlighted in blue, parameters outside this region don't have a solution on the whole cross section. Values are shown for  $R_0 = 25\text{mm}$ , but results generalize through scale-invariance. ODE solutions corresponding to selected parameters from the left are shown on the right. The dashed lines correspond to negated  $\rho^*$  values and form convex lenses, while positive  $\rho^*$  values form concave lenses.

### 3.3 Parameter Space Analysis

With the above formulation, we introduced a new design space for lenses, that we now analyze in more detail. Overall, there are three independent parameters:  $\rho^*$  related to the density difference,  $p^*$  related to the injected volume; and  $h_c$ , the center thickness that directly relates to the ring height through  $h_c = h_0 - h^*(R_0)$ . Due to the above mentioned shift-invariance, we mostly give general results for  $h^*(r)$  and refer to the shift  $h_c$  only when needed.

**3.3.1 Improved accuracy.** In Elgarisi et al. [Elgarisi et al. 2021] the assumption, that the maximal height difference on the lens surface is significantly smaller than the lens radius, is used to derive an efficient, analytical solution of the system of ODE. However, this constrains the permitted parameters significantly and limits possible lens design. If the assumption is not fulfilled, lenses can still be designed, but the mismatch between model and reality then hinders optical performance. To illustrate, for neutral buoyancy ( $\Delta\rho = 0$ ),

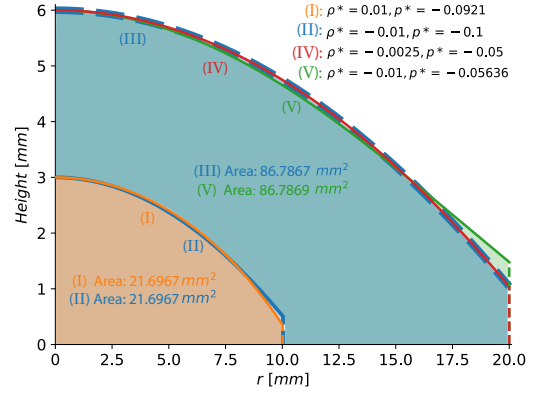


Fig. 4. Scale-invariance and impact of the density relationships on the lens shape. 4 shapes with specified parameters are shown. 1) Flipping the sign of  $\rho^*$  while keeping the volume constant (through adjusting  $p^*$ ), changes the lens shape ((I) & (II)). 2) Geometric scaling of shape (II) yields shape (III). Attempting to achieve the scaled shape by increasing the injected volume yields (V) as the best fit. Only when the density of the immersion liquid is changed as well (changing  $\rho^*$ ), the scaled shape can be fit perfectly, (IV). The scaled lens parameters are computed according to Eq. 16.

the model of Elgarisi et al. computes a parabolic lens shape when the correct solution is a spherical one, as shown above.

We have already seen, that our model produces spherical shapes when expected. For non-neutral buoyancy, solving the ODE through numerical solvers allows for arbitrary precision without relying on constraining assumptions, as shown in Fig. 2.

**3.3.2 Fitting common lens designs.** Different design spaces (i.e. parameterized formulas that describe lens surface) are used in lens design pipelines (e.g. ZEMAX/ CODE V). For example, the generatrix of the spherical surface is designed in the form of

$$\text{Spherical}(r) = \frac{cr^2}{1 + \sqrt{1 - (cr)^2}}, \quad (12)$$

where  $c$  is the curvature. Aspherical surfaces are described by

$$\text{Aspherical}(r) = \frac{cr^2}{1 + \sqrt{1 - (1 + \kappa)(cr)^2}} + \alpha_4 r^4 + \alpha_6 r^6 + \dots, \quad (13)$$

where  $\kappa$  is the conic constant, and  $\alpha_i$  denote coefficients of polynomials.

Through our surface formation model, we introduced a novel design space with the surface parameters  $p^*$ ,  $\rho^*$ , and the lens thickness  $h_c$ . Rather than being expressed through an analytical formula,  $h(r)$  is obtained by solving a differential equation obtained from a physical model.

We have already shown, that spherical lenses are a subset of our design space. We now show that classic aspherical lenses (Eq. 13) can be very well approximated in the fluidic design space. Fig. 2 shows geometric fits of the Thorlabs AL2550 aspherical plano-convex lens surface (described by 5 parameters in Eq. (13);  $c$ ,  $\kappa$ ,  $\alpha_4$ ,  $\alpha_6$  and  $\alpha_8$ ) by our method and Elgarisi et al.. We find that, in contrast to previous work, our method closely resembles the given shape. Note, however, that our design space has only 3 dimensions, thus not all possible shapes from Eq. (13) can fit equally well in theory. In

the supplemental material, we show, that all common lens designs tested have very close fits in our design space. Note, however, that the motivation of our work lies in the creation of new, end-to-end optimized, lens designs, thus fitting existing lens designs is considered a niche application of our work here. Additional fits are found in the supplemental material.

**3.3.3 Symmetry in volume parameter.**  $h^*(r)$  is symmetric with respect to the volume parameter  $p^*$ , meaning that

$$h^*(\rho^*, p^* | r) = -h^*(\rho^*, -p^* | r). \quad (14)$$

It immediately follows, that  $h^* = 0$  for  $p^* = 0$ . Otherwise, the surface is convex for  $p^* < 0$  and concave for  $p^* > 0$ . Examples of resulting curves and their symmetries, are shown in Fig. 3. Note, that  $p^*$  is in an abstract unit, rather than in a volume unit, as it depends on the Lagrange multiplier  $\lambda$ . Intuitively, it describes the deviation from the volume that gives a minimal surface area for any ring size.

**3.3.4 Impact of buoyancy imbalance.**  $\rho^*$  is proportional densities difference  $\Delta\rho$ . In previous work, only the absolute value of the force was considered, which does not allow the distinction between  $\rho_{im} > \rho_r$  and  $\rho_{im} < \rho_r$  and limits possible designs [Elgarisi et al. 2021], since both cases lead to different shapes.

To illustrate this, several resulting lens shapes are plotted in Fig. 4. There, different shapes with equal lens volume and flipped signs of  $\rho^*$  are shown, where  $p^*$  is adjusted to fulfill the previous two constraints accordingly. We find that an upward facing net force leads to taller and slimmer shapes, while a downward facing force leads to flatter lens surfaces.

**3.3.5 Valid parameter ranges.** During the flow of the resin, its surface tension has to be strong enough to keep all the resin in a compact shape. Since we model the equilibrium state of the surface, it is expected that for values beyond these limits, no solutions exist in our model.

We perform numerical experiments to explore these stability conditions. For a fixed ring size (with  $R_0 = 25$  mm and  $h_0$  undefined since we only analyze  $h^*$ ), we find all possible combinations of  $p^*$  and  $\rho^*$  (through brute-force sampling and checking for the existence of a solution), shown in Fig. 3. We find, that negative values of  $\rho^*$  (for which the net force faces downward) allow for significantly larger injected volumes. While these results appear specific for the initial ring size, the next section generalizes these results through scale remapping.

**3.3.6 Scale dependency.** Since surface tension scales with surface area while buoyancy scales with volume, the solution for Eq. (9) is scale-dependent. I.e., for a scaling factor  $s$  and unadjusted  $p^*, \rho^*$  we have

$$\begin{cases} R'_0 = s \cdot R_0, \\ h'_0 = s \cdot h_0, \\ h'(r) \neq s \cdot h(\frac{r}{s}). \end{cases} \quad (15)$$

However, we find that there exists a way to adjust  $p^*, \rho^*$  to perfectly compensate for the change in scale:

$$\begin{cases} \rho_s^* = \frac{\rho^*}{s^2}, \\ p_s^* = \frac{p^*}{s}. \end{cases} \quad (16)$$

and with this

$$s \cdot h\left(\rho^*, p^* \mid \frac{r}{s}\right) = h'\left(\rho_s^*, p_s^* \mid \frac{r}{s}\right). \quad (17)$$

In practice, this means, that in order to create a larger or smaller lens with the same shape, it is not sufficient to merely adjust the injected volume, but also the water/glycerol mixture must be adjusted to compensate for the different scaling between surface tension and buoyancy, see Fig. 4.

## 4 END-TO-END LENS OPTIMIZATION WITH DIFFERENTIABLE RAY TRACING

We integrate the differentiable physical lens formation model into an open-source differentiable ray tracer *DeepLens* [Wang et al. 2022; Yang et al. 2024] for end-to-end lens optimization, see Fig. 1.

### 4.1 Differentiable Ray Tracing

Ray tracing through an optical lens [Glassner 1989; Lee et al. 2010; Chen et al. 2021] can be utilized to calculate the point spread function (PSF) and for ray tracing-based rendering, both of which are effective in simulating the images captured by an optical lens. In recent advancements, this process has been developed in a differentiable manner, employing simulated images as objectives to optimize optical lenses [Sun et al. 2021; Wang et al. 2022; Côté et al. 2023; Yang et al. 2024]. The *DeepLens* ray tracer is capable of computing the intersection of optical rays with lens surfaces and the refraction of rays in a plug-and-play fashion, given the surface function  $h$  and its first-order derivative  $h_r$ . Optical rays emanating from single object points reach the sensor plane, and PSFs are calculated, which are subsequently used in image simulation. Operating on the PyTorch auto-diff framework, *DeepLens* is entirely differentiable, enabling gradient back-propagation from the final image output to optimize lens parameters. For further details on the pipeline, we refer to the aforementioned references.

### 4.2 Efficient Numerical Surfaces

During ray tracing, most of the computational time and power is dedicated to calculating interactions with lens surfaces, which involves multiple iterations of Newton's method. Consequently, the parameters  $h$  and  $h_r$  are frequently evaluated at various points. Traditionally, each evaluation requires a repeated ODE computation due to unknown intersection points, which is excessively costly in terms of both time and memory.

To address this challenge, we propose an efficient numerical solution to replace the naive per-ray calculation. After updating each lens parameter and prior to ray tracing, the function  $h$  is sampled once and stored in a table. Furthermore, solving the differential equation to compute  $h$  allows us to obtain and store function values of  $h_r$  at no extra cost. The nearest neighbour lookup proves inadequate, as the gradient may vanish entirely for nearby sample positions. Therefore, a differentiable interpolation step is necessary to approximate intermediate function values continuously. Our research indicates that typically, 1k samples of  $h$  and  $h_r$  along with linear interpolation suffice, given the inherent smoothness of lens surfaces. By adjusting the table size, the desired accuracy can be attained.

In a typical end-to-end optimization step, 100k rays are sampled, but the ODE is only evaluated once, in steps that modify the lens parameters. And since the lookup tables are small, there is virtually no memory or computation time impact, when using numerical surface descriptions. Contrary, this approach might even be used to speed up the evaluation of complex analytical surface descriptions. While the exact values vary between use cases and chosen resolutions, this method roughly realizes a speed-up factor of 500 in our case.

### 4.3 Lens Optimization

To maintain consistency between simulation and real-world experiments, a flat image is positioned at a set distance in front of the lens system to represent the actual scene. Image reconstruction is carried out using a lightweight U-Net [Ronneberger et al. 2015] shaped network, NAFNet [Chen et al. 2022], which is co-designed with the fluidic lens in an end-to-end optimization process. Three wavelengths (656 nm, 589 nm, 486 nm) are utilized to simulate RGB images and characterize the optical aberrations inherent in the lens design. During the end-to-end optimization, we first simulate the camera-captured images, followed by feeding them into the reconstruction network. In the backward pass, the loss function can be back-propagated through the network and the fluidic lens for joint design.

### 4.4 Loss Function

The loss function

$$\mathcal{L}_{\text{Total}} = \mathcal{L}_{\text{PSF}} + \mathcal{L}_{\text{Image}} \quad (18)$$

strives to improve both the optical performance of the lens and the final image quality simultaneously. The image loss determines the optimal combination of the lens and reconstruction network for enhancing image quality, while an additional PSF loss adjusts the lens to minimize the spot size of the ray distribution across all wavelengths, as shown in Fig. 1.

Specifically, the image loss itself consists of three components:

$$\mathcal{L}_{\text{Image}} = \mathcal{L}_{\text{MSE}} + \mathcal{L}_{\text{perceptual}} + \mathcal{L}_{\text{SSIM}}, \quad (19)$$

where  $\mathcal{L}_{\text{MSE}}$  represents Mean Squared Error,  $\mathcal{L}_{\text{perceptual}}$  denotes perceptual loss, and  $\mathcal{L}_{\text{SSIM}}$  signifies Structural Similarity Index Measure (SSIM) loss. All the loss terms are weighted by hyperparameters. For training, the lens parameters are initialized with parameters from a numerical fit to the geometrical shape of the comparison lenses. Then, for the first part, only the network is trained, until the results stabilize. After this, lens optimization is enabled and the network and lens are trained jointly until the end.

## 5 FLUIDIC LENS MANUFACTURING

We test our lens design by manufacturing some prototypes and evaluating their performance. A more detailed discussion about manufacturing difficulties and ways to overcome them is found in the discussion section.

### 5.1 Physical Parameters and Materials

In the lens design space, we optimize for  $\rho^*$  and  $p^*$  (Eq. 8), and  $h_c$ . From these, we compute the fabrication parameters  $\Delta\rho$ ,  $V_0$ , and  $h_0$ ,

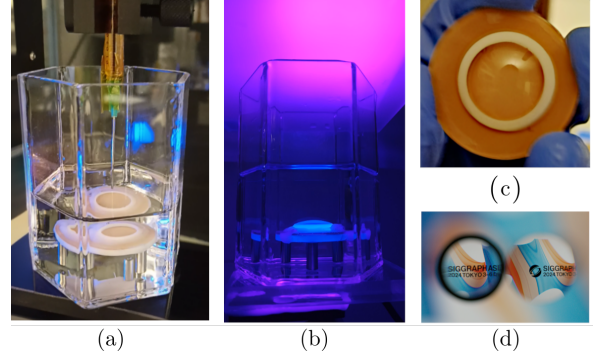


Fig. 5. Manufacturing overview. (a) UV resin injection, (b) curing process, (c) finished lens after curing, and (d) fluidic lens (left) qualitative comparison with a benchmark off the shelf lens (right).

the density difference, liquid lens volume, and ring height, respectively.

The relative density  $\Delta\rho$  is calculated assuming a constant value for the interfacial tension  $\gamma$ . This value is in fact an average over time since the molecular diffusion processes prevent it from having a fixed constant. However, according to our experimental measurements, as discussed later, this is still a suitable approximation.

From  $\Delta\rho$  the mixing ratio of deionized water and glycerol is determined. According to [Takamura et al. 2012], the resulting density of the mixture follows a linear relationship. In practice, we can accurately match the required mixing ratio by iteratively measuring the density using a density meter and adding water or glycerol accordingly. Following this process, we can retrieve a density error of less than 0.05 %.

We also need to compute the injection volume of the resin. This is not directly possible from  $p^*$ , since it is not given in physical units. However, knowing  $h(r)$ , numerical integration can be used to determine the lens volume with arbitrary precision.

For our prototypes, we used the UV resin brand Vida Rosa which has already been reported in the fluidic lenses fabrication literature [Elgarisi et al. 2021]. We experimentally determined its refractive index for the visible spectrum by using spectroscopic ellipsometry. This optical model is required for the ray-tracing simulation. We also measured its liquid density and characterized the average shrinkage caused by the curing process to reduce potential fabrication errors.

### 5.2 Manufacturing Steps

Refer to Fig. 5 for a summary of the key fabrication steps and a visual comparison of the optical performance between one of our produced lenses and a commercially available lens. While the ring diameter is chosen a-priori, the ring height  $h_0$  is an optimization parameter and will differ between lenses. We use a parametric 3D model to print the ring shape. A photograph of one of such rings can be seen in Fig. 5. The actual ring fabrication is challenging, since we require precision in the micro-meter scale, beyond what is visible to the naked eye. We use a stereolithography (SLA) Formlabs Form 3L 3D printer with standard resin that has an XY resolution of 25 microns and layer thickness of 50 microns.




Fig. 6. Our Imaging Setup: (a) The scene being imaged, and (b) the front (top) and side (bottom) views of the camera setup, including the custom-manufactured fluidic lens.

Before starting the fabrication, the immersion liquid and the UV resin are placed inside a vacuum chamber for 10 minutes. The 3D printed ring is attached to a polycarbonate sheet that will form a flat face for our optical component. The base ring is placed inside the immersion liquid and carefully leveled using a high resolution bubble level.

For the resin injection, we use a high-end threaded plunger syringe brand Hamilton, which releases  $13.23 \mu\text{l}$  per turn. For the resin injection, we use a high-end threaded plunger syringe brand Hamilton that releases  $13.23 \mu\text{l}$  per turn, which determines our volume resolution. For our 20 mm in diameter lenses, we use around 80-120 syringe turns. The syringe is positioned vertically in the central zone of the ring, and the resin deposition begins at a slow enough rate to prevent the entrapment of immersion liquid bubbles. Subsequently, the resin volume spreads out and uniformly fills the ring. For upward facing buoyancy forces, special care must be taken, to keep the resin volume inside the ring. Surface wettability plays a crucial role here. While not experimentally characterized in our exploration, UV resin demonstrates better spreading on polycarbonate surfaces compared to glass.

For the photo-curing process, we utilize a 32 W UV lamp at 405 nm and apply an energy dose of  $2500 \text{ mJ}/\text{cm}^2$  at  $1.1 \text{ mW}/\text{cm}^2$ , resulting in an exposure time of approximately 38 minutes. Once the curing is finished, the lens and its base are taken out from the immersion liquid, followed by removal from the polycarbonate flat base and drying.

Lenses with two curved surfaces require two passes of this process. Two separate rings are fabricated, where the flat side of the first lens surface acts as bottom for the second ring. The liquid resin of the second surface attaches tightly to the already cured first surface and the final lens has no visible seam. We note, however, that precise leveling for the second half becomes more challenging and failure to do so results in asymmetric lens shapes.

## 6 COMPUTATIONAL IMAGING WITH END-TO-END FLUIDIC LENS DESIGN

To showcase the efficacy of the proposed method, we design and fabricate fluidic lenses for computational imaging in an end-to-end manner, and then compare them with off-the-shelf optical lenses in terms of imaging quality. Specifically, several lens prototypes

are demonstrated: (1) AL2550: plano-convex lens with one aspherical and one flat surface, (2) Best form: bi-convex lens with two spherical surfaces, and (3) Cooke: spherical triplet, consisting of a bi-concave lens between two bi-convex lenses (for a total of 6 optimized surfaces).

Note, that due to differences in the optical properties, it is insufficient to merely optimize for the same shape, we rather have to use end-to-end optimization to recreate or outperform the imaging capabilities.

### 6.1 Experimental Settings and Setup

The end-to-end optical design of the fluidic lens and an image reconstruction network use the DIV2K dataset consisting of 800 training images [Agustsson and Timofte 2017]. Data augmentation is used during optimization to improve the overall performance of both the optics and the network.

For real-world experiments, a virtual scene is set up with flat images positioned 1 m away from the first lens surface and various 3D objects included, as shown in Fig. 6. We use a camera Lucid Triton 5.4 MP, with an 4.5 mm aperture diameter, and the lens placed in the middle. A lens hood is added to prevent indirect reflections. All elements are fixed on tracks to ensure alignment and enable lens refocusing.

### 6.2 Results on Simulated Lenses

Our imaging results are shown in Fig. 7 (top). As the stock plano-convex lens already produces a comparatively sharp image on the sensor, our optimized fluidic lens shows a very similar performance. This carries over to the reconstruction results by the network, which are again very comparable.

For the bi-convex and triplet, the differences are clearer. In both cases, the optimized lens produces a sharper image on the sensor. We note however, that these are end-to-end optimizations, where a sharp sensor image is not the ultimate goal, but rather the overall reconstruction is optimized. Consequently, the performance gap is larger on the reconstructed images. These results show, that even on simple imaging systems with few adjustable surfaces, end-to-end optimization can improve results over these stock lenses.

The training time for all examples is around 24 hours. The plano-convex and bi-convex lenses were optimized on a single V100 GPU, while the triplet lens is computationally more demanding due to 6 surfaces being jointly optimized and runs on 2 A100 GPUs.

The scores on the right are averages over the whole images and over many test images of our dataset and show, that the general image quality improves.

### 6.3 Real Results

After lens fabrication, the reconstruction network should be fine-tuned to learn sensor characteristics such as noise and adjust to potential fabrication inaccuracies. This can help to mitigate artifacts, but more importantly, the network performance will likely be poor on unexpected input data. For real scenes, ground truth images are harder to obtain. We place a tablet in the scene to cover the whole field of view of the lens and display training images on it. These ground truth images must be carefully aligned with the sensor image

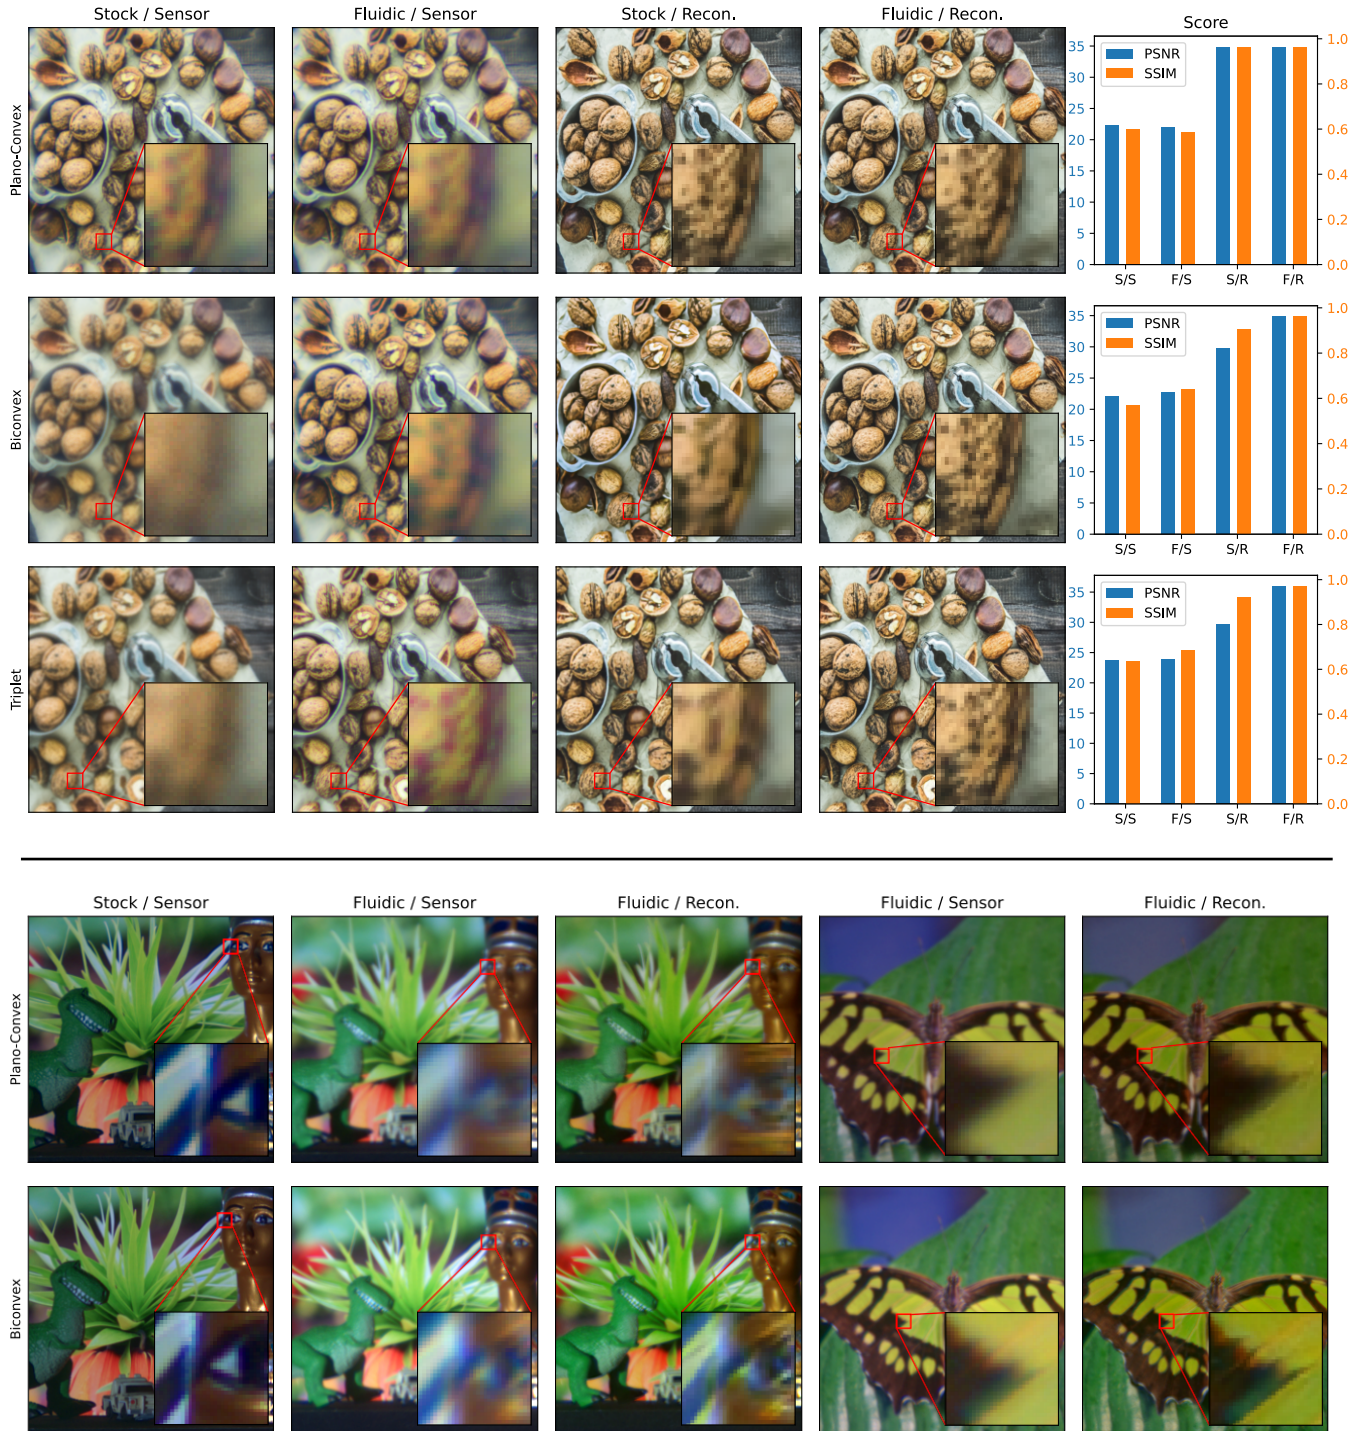


Fig. 7. Imaging results of our lens designs. (Top) Optical performance of simulated lenses. 3 types of setups are compared, where *Stock* refers to bought reference lenses, and *Fluidic* refers to our optimized lenses. Results on the sensor and after the reconstruction are shown. Images are compared against the ground truth image in the virtual scene by *peak signal to noise ratio* and *structural similarity metric*, for both, higher values are better. Scores are averaged over the whole dataset, exact values are found in the supplementary material. (Bottom) Optical performance of fabricated lenses. For these preliminary prototypes, an aperture with F-number 10 was used to enhance sharpness and show case the effect of the reconstruction network better.



and potentially color corrected. The fine-tuning requires around 6 additional hours for good training results.

Our measurements are shown in Fig. 7 (bottom). Our findings indicate that there is room for improvement in the performance of our lenses compared to the expectations set by the previous simulations. A detailed discussion of likely reasons is found in the discussion section.

To increase image quality, we add an aperture (F-number 11) behind the lens. This filters out rays passing through the edge of the lens, where the surface shape is most affected by ring inaccuracies. The image quality is improved by the fine-tuned reconstruction network, however, also the bi-convex lens does not outperform the stock lens.

The comparison images show the stock lenses with and without the aperture. While with a small aperture, the stock lenses are capable of producing sharp images that do not necessarily require network post-processing, these single lens systems show significant problems without the aperture.

In the future, we hope to improve fabrication quality through enhanced engineering. Designing lens systems means dealing with compromises. Our simulated results show for example, how off-axis performance can be improved when the focus on on-axis performance is reduced. Such kind of trade-offs mean, that it is not always necessary to have an overall better all-around performance to outperform in specialized usage scenarios. Then, our lens optimization would have the potential to improve over the stock lenses, as shown in the simulated results.

## 7 FUTURE WORK AND CONCLUSION

In this work, we have introduced a new shape model for fluidic shaped lenses, which is more accurate than earlier models, and for the first time includes spherical surfaces, which are the expected shape for neutral buoyancy conditions, since in this case surface tension is the only force. We have also shown that, through a differentiable implementation, this new shape model can be directly integrated into end-to-end optical design pipelines, and we have evaluated this approach through extensive simulations. For now, a discrepancy remains between the performance of the simulated and fabricated lenses for which we identify some fabrication difficulties that likely have the biggest impact on quality.

## ACKNOWLEDGMENTS

This work has been supported by KAUST under individual baseline funding. The authors are grateful to Dmitry A. Lyakhov and Helmut Pottmann for helpful discussions. The reviewers' valuable comments are gratefully acknowledged.

## REFERENCES

J.R. Dormand and P.J. Prince. 1980. A family of embedded Runge-Kutta formulae. *J. Comput. Appl. Math.* 6, 1 (1980), 19–26. [https://doi.org/10.1016/0771-050X\(80\)90013-3](https://doi.org/10.1016/0771-050X(80)90013-3)

Andrew S Glassner. 1989. *An introduction to ray tracing*. Morgan Kaufmann.

Ekkard Brinksmeier, Oltmann Riemer, and Ralf Gläbe. 2007. Merging Technologies for Optics. In *Towards Synthesis of Micro-/Nano-systems*. Springer London, London, 1–9.

Sungkil Lee, Elmar Eisemann, and Hans-Peter Seidel. 2010. Real-time lens blur effects and focus control. *ACM Transactions on Graphics (TOG)* 29, 4 (2010), 1–7.

Roberto Spina, Paul Walach, Julian Schild, and Christian Hopmann. 2012. Analysis of lens manufacturing with injection molding. *International Journal of Precision Engineering and Manufacturing* 13 (2012), 2087–2095.

Koichi Takamura, Herbert Fischer, and Norman R. Morrow. 2012. Physical properties of aqueous glycerol solutions. *Journal of Petroleum Science and Engineering* 98–99 (2012), 50–60. <https://doi.org/10.1016/j.petrol.2012.09.003>

Thiago Pereira, Szymon Rusinkiewicz, and Wojciech Matusik. 2014. Computational light routing: 3d printed optical fibers for sensing and display. *ACM Transactions on Graphics (TOG)* 33, 3 (2014), 1–13.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5–9, 2015, proceedings, part III 18*. Springer, 234–241.

Timo Gissibl, Simon Thiele, Alois Herkommer, and Harald Giessen. 2016. Two-photon direct laser writing of ultracompact multi-lens objectives. *Nature photonics* 10, 8 (2016), 554–560.

Eirikur Agustsson and Radu Timofte. 2017. NTIRE 2017 Challenge on Single Image Super-Resolution: Dataset and Study. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.

John Gawdzinski, Michal E Pawlowski, and Tomasz S Tkaczyk. 2017. Quantitative evaluation of performance of three-dimensional printed lenses. *Optical engineering* 56, 8 (2017), 084110–084110.

Ricky T. Q. Chen. 2018. torchdiffeq. <https://github.com/rtqichen/torchdiffeq>

Gregory D. Berglund, Tomasz S. Tkaczyk, Gregory D. Berglund, Tomasz S. Tkaczyk, and Tomasz S. Tkaczyk. 2019. Fabrication of optical components using a consumer-grade lithographic printer. *Optics Express*, Vol. 27, Issue 21, pp. 30405–30420 27, 21 (Oct. 2019). <https://doi.org/10.1364/OE.27.030405>

Kaan Akşit, Praneeth Chakravarthula, Kishore Rathinavel, Youngmo Jeong, Rachel Albert, Henry Fuchs, and David Luebke. 2019. Manufacturing application-driven foveated near-eye displays. *IEEE transactions on visualization and computer graphics* 25, 5 (2019), 1928–1939.

Qilin Sun, Congli Wang, Fu Qiang, Dun Xiong, and Heidrich Wolfgang. 2021. End-to-end complex lens design with differentiable ray tracing. *ACM Trans. Graph* 40, 4 (2021), 1–13.

Mor Elgarisi, Valerie Fumkin, Omer Luria, and Moran Bercovici. 2021. Fabrication of freeform optical components by fluidic shaping. *Optica* 8, 11 (Nov. 2021), 1501–1506.

Valeri Frumkin and Moran Bercovici. 2021. Fluidic shaping of optical components. *Flow* 1 (May 2021), E2. <https://doi.org/10.1017/flo.2021.1>

Shiqi Chen, Huajun Feng, Dexin Pan, Zhihai Xu, Qi Li, and Yueting Chen. 2021. Optical aberrations correction in postprocessing using imaging simulation. *ACM Transactions on Graphics (TOG)* 40, 5 (2021), 1–15.

Gregory Berglund, Anna Wisniewiecki, John Gawdzinski, Brian Applegate, and Tomasz S. Tkaczyk. 2022. Additive manufacturing for the development of optical/photonics systems and components. *Optica* 9, 6 (Jun 2022), 623–638. <https://doi.org/10.1364/OPTICA.451642>

Congli Wang, Ni Chen, and Wolfgang Heidrich. 2022. dO: A differentiable engine for deep lens design of computational imaging systems. *IEEE Transactions on Computational Imaging* 8 (2022), 905–916.

Gregory D. Berglund and Tomasz S. Tkaczyk. 2022. Enabling consumer-grade 3D-printed optical instruments &#x2013; a case study on design and fabrication of a spectrometer system using low-cost 3D printing technologies. *Opt. Continuum* 1, 3 (Mar 2022), 516–526. <https://doi.org/10.1364/OPTCON.447693>

Manjesh Kumar, Hari Narayan Singh Yadav, Abhinav Kumar, and Manas Das. 2022. An overview of magnetorheological polishing fluid applied in nano-finishing of components. *Journal of Micromanufacturing* 5, 2 (2022), 82–100.

Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. 2022. Simple baselines for image restoration. In *European conference on computer vision*. Springer, 17–33.

Geoffroi Côté, Fahim Mannan, Simon Thibault, Jean-François Lalonde, and Felix Heide. 2023. The differentiable lens: Compound lens search over glass surfaces and materials for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 20803–20812.

Daniel Aguirre-Aguirre, Dulce Gonzalez-Utrera, Brenda Villalobos-Mendoza, and Rufino Diaz-Urbe. 2023. Fabrication of biconvex spherical and aspherical lenses using 3D printing. *Applied Optics*, Vol. 62, Issue 8, pp. C14–C20 62, 8 (March 2023).

Jay Christopher, Liam M. Rooney, Mark Donnachie, Deepak Uttamchandani, Gail McConnell, and Ralf Bauer. 2023. Low-cost 3D printed lenses for brightfield and fluorescence microscopy. *bioRxiv* (2023), 2023.11.22.568227. <https://doi.org/10.1101/2023.11.22.568227>

Xinge Yang, Qiang Fu, and Wolfgang Heidrich. 2024. Curriculum learning for ab initio deep learned refractive optics. *Nature communications* 15, 1 (2024), 6572.

Dulce Gonzalez-Utrera, Brenda Villalobos-Mendoza, Rufino Diaz-Urbe, and Daniel Aguirre-Aguirre. 2024. Modeling, fabrication, and metrology of 3D printed Alvarez lenses prototypes. *Optics Express*, Vol. 32, Issue 3, pp. 3512–3527 32, 3 (Jan. 2024).