

Task-Driven Lens Design

XINGE YANG,¹ QIANG FU,¹ YUNFENG NIE,² AND WOLFGANG HEIDRICH^{1,*}

¹King Abdullah University of Science and Technology, Thuwal, 23955-6900, Saudi Arabia

²Vrije Universiteit Brussel (VUB), Brussels Photonics Team, Department of Applied Physics and Photonics, Pleinlaan 2, B-1050 Brussels, Belgium

*wolfgang.heidrich@kaust.edu.sa

Abstract: Classical lens design minimizes optical aberrations to produce sharp images, but is typically decoupled from downstream computer vision tasks. Existing end-to-end optical design learns optical encoding through joint optimization, but often suffers from an unstable training process. We propose task-driven lens design, a new optimization philosophy for joint optics-network systems. We freeze the pretrained vision model and optimize only the lens so that the image formation better fits the model’s feature preferences. This network-frozen setting yields a low-dimensional and stable optimization process, enabling lens design from scratch without human intervention, thereby exploring a broader design space. Multiple computer vision experiments show that TaskLenses outperform classical ImagingLenses with the same or even fewer elements. Our analysis reveals that the learned optics exhibit long-tailed point spread functions, better preserving preferred structural cues when aberrations cannot be fully corrected. These results highlight task-driven design as a practical route for optical lenses that are compatible with modern vision models, and also inspire new optical design objectives beyond traditional aberration minimization.

© 2026 Optica Publishing Group

1. Introduction

Classical lens design is typically decoupled from any downstream image analysis task and focuses on maximizing image quality by minimizing optical aberrations. To achieve optimal performance for modern computer vision networks, such as those used for image classification [1, 2], object detection [3, 4], semantic segmentation [5, 6], and vision-language models (VLMs) [7–9], optical engineers design lenses that minimize aberrations to capture sharp, high-quality images. However, these high-end lenses often come at a prohibitively high cost and possess a bulky form factor with complex structures. For instance, modern smartphone cameras may include more than five highly aspheric elements [10, 11], while professional camera lenses can incorporate six or more precision-manufactured elements [12, 13] to achieve satisfactory optical performance. For deployment into mobile and robotic platforms, it is desirable to reduce this optical complexity without compromising the performance of any downstream task. However, when residual aberrations persist in captured images, downstream computer vision performance can suffer dramatically.

Recent studies in end-to-end lens design [14–17] have demonstrated the potential for achieving optimal system-level performance by jointly optimizing camera optics and computer vision algorithms. This approach enables the learning of optical encoding power during the joint optimization of lenses and neural networks, yielding promising results across a variety of tasks, including compact imaging systems [18–20], extended depth-of-field imaging [15, 17, 21–24], hyperspectral imaging [25–28], high-dynamic-range imaging [29, 30], object detection [16, 31], and depth estimation [32, 33]. These works suggest that producing conventionally sharp images is often suboptimal for downstream tasks, especially when using simple optics where aberrations cannot be fully corrected.

However, learning optical encoding power through joint optimization of optics and neural networks presents several significant challenges. Optical systems typically involve optimizing only tens of physical parameters, whereas modern neural networks can contain millions to billions of parameters; optimizing them together often leads to unstable training dynamics and optimization oscillations. To achieve promising convergence, existing studies typically rely on pre-optimized lenses as starting points [15–17, 31, 34], which may trap the design process in local minima and limit exploration of the optical design space. Moreover, modern vision foundation models have evolved to encapsulate rich prior knowledge of the physical world; retraining or fine-tuning these large-scale networks is not only prohibitively expensive but also risks disrupting their robust, pre-learned representations. Prior works attempt to explore the possibility of simplifying optical systems according to computer vision requirements [31, 35]; however, a comprehensive evaluation across diverse tasks and network architectures, as well as an explanation of the underlying optical mechanisms behind performance gains, remains missing.

In this paper, we introduce a new optimization philosophy for optics-network joint systems: freeze the network and optimize only the lens (i.e., **task-driven lens design**). While we use a differentiable image formation model (as in prior end-to-end works) to propagate gradients to the optical parameters, our network-frozen setting turns lens design into a low-dimensional and stable optimization problem. During the optimization, the lens learns to encode the image features preferred by the downstream computer vision models, which not only reduces optimization difficulty, also represents an explainable optical design objective.

We first demonstrate our task-driven lens design using image classification. We designed three distinct lenses (“TaskLenses”) fully automatically from scratch, using a frozen image classification network as the optimization objective. This automated task-driven lens design not only demonstrates the efficacy of stable network-provided gradients, but also enables a broader search space for lens design. For comparison, we designed three corresponding classical lenses (“ImagingLenses”) for each “TaskLens” by minimizing optical aberrations. We evaluated the image classification accuracies on the ImageNet benchmark [36] and found that our “TaskLenses” consistently outperform the “ImagingLenses” in classification accuracy while utilizing the same number of lens elements. Moreover, we discovered that “TaskLenses” can achieve comparable, or even superior, image classification accuracy with fewer elements.

We further investigate the generalizability of the task-driven design across different computer vision tasks, including object detection, semantic segmentation, and vision-language models. By evaluating the task-driven designed lenses across various applications, we found that different computer vision tasks can share similar preferable image features, making it feasible to generalize from simpler to more complex tasks. We also observed the potential for generalization from simpler network architectures to more complex ones. Analysis of the optical characteristics of the designed lenses reveals that the TaskLenses tend to converge on long-tailed point spread functions (PSFs) with a sharp central peak, which is a significant deviation from the compact PSFs with a relatively broad central peak, obtained with conventional lens design. While such PSF profiles may not appear sharp or meet traditional optical design standards, it is more effective at preserving crucial high-frequency features against optical aberrations, thereby benefiting various computer vision tasks. The contributions of this paper are summarized as follows:

- We introduce a new optimization philosophy for joint optics-network systems by freezing a pretrained vision model and optimizing only the lens. This task-driven design represents an explainable objective that is aligned with modern computer vision feature extraction.
- Task-driven lens design reduces optimization difficulty and enables broader exploration of the design space by starting from scratch, showing the capability of designing simplified structures with comparable computer vision performance using fewer lens elements.
- We provide a comprehensive evaluation of the task-driven designed lenses across various

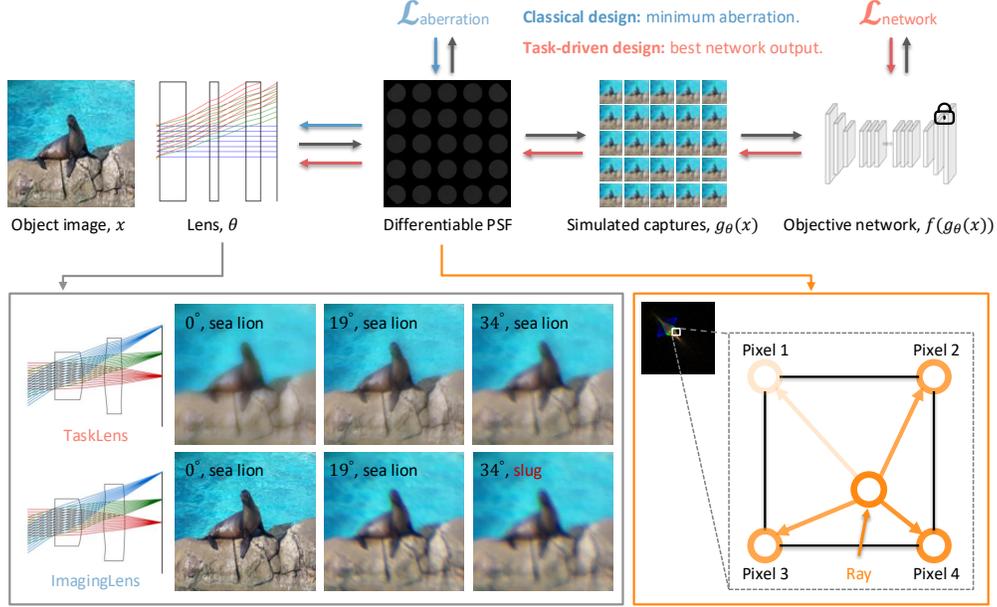


Fig. 1. The task-driven lens design pipeline begins by calculating the PSF via differentiable ray tracing. This PSF is then convolved with input images to simulate camera captures. A well-trained computer vision network (e.g., for image classification) evaluates the final output error and backpropagates the loss function to optimize lens parameters. Unlike classical lens design, which minimizes optical aberrations, task-driven lens design prioritizes capturing image features that are most effective for downstream computer vision networks. After optimization, the designed TaskLens better preserves image structure details even in the presence of optical aberrations, leading to improved performance on the target computer vision task. In contrast, an ImagingLens, designed to produce clear captures, is more susceptible to performance degradation from optical aberrations, potentially resulting in misclassification.

computer vision tasks and network architectures, demonstrating their generalizability and the emergence of a novel long-tailed PSF as a key characteristic of computer vision-preferred optical properties.

2. Methods

2.1. Task-driven lens design

Lens design for a computer vision task can be formulated as:

$$\theta^* = \operatorname{argmin}_{\theta} \|f_{\phi}(g_{\theta}(x)) - y\|, \quad (1)$$

In Equation (1), f_{ϕ} is a pre-trained neural network for the vision task:

$$f_{\phi} = \operatorname{argmin}_{\phi} \|f(x) - y\|. \quad (2)$$

Here, x represents the object image, θ and ϕ respectively denote the parameters of the lens design and the neural network, g represents the imaging process, y is the ground truth label for the computer vision task f_{ϕ} , and f_{ϕ} is a pre-trained computer vision model trained on a large dataset.

Classical lens design typically focuses on minimizing optical aberrations, such as root mean square (RMS) spot size or wavefront error. Under this design philosophy, optical engineers assume that minimizing optical aberrations to achieve optimal image quality will yield the best possible performance for downstream computer vision tasks. This classical lens design approach is illustrated by the blue arrow in Figure 1 and can be expressed as:

$$\theta^* = \underset{\theta}{\operatorname{argmin}} h_{\theta}, \quad (3)$$

where h_{θ} represents the optical aberration function, which depends on the lens parameters θ . Ideally, when $h_{\theta} = 0$, the camera-captured image $g_{\theta}(x)$ is identical to the object image x (i.e., $g_{\theta}(x) = x$). However, the classical lens design objective is decoupled from computer vision, and Equation (3) represents an *intermediate* optimization objective for the target computer vision task. The mismatch between minimizing optical aberrations and the actual requirements of diverse computer vision applications leads to a performance gap. Consequently, when optical aberrations cannot be fully corrected due to constraints on the optical form factor or manufacturing limitations, especially in edge devices such as robotics, the corresponding computer vision performance can be significantly affected.

A more straightforward and promising lens design objective is to directly optimize the objective function in Equation (1), which can be expressed as:

$$\theta^* = \underset{\theta}{\operatorname{argmin}} \|f(g_{\theta}(x)) - y\|. \quad (4)$$

This task-driven lens design objective directly optimizes the lens design for optimal computer vision performance on the target task, as illustrated by the red arrow in Figure 1. During optimization, this objective evaluates how well the captured image is suited for the downstream computer vision network. However, realizing this task-driven lens design objective has been challenging in classical lens design due to the lack of a differentiable lens simulator $g_{\theta}(x)$, which is necessary to enable direct gradient backpropagation from the network f outputs to the lens parameters θ . Recent advances in differentiable optics have provided a new approach for task-driven lens design, as demonstrated in [17, 22] and detailed in Section 2.2.

Modern vision foundation models trained on massive image datasets learn the rich information of image data distributions and can effectively extract useful latent features from image inputs. This inspires lens design to focus on capturing the most important image features from the objects, especially when the optical structure is limited by form factor and manufacturing constraints. Decomposing the original object image x as $x = x_f \oplus x_{bg}$, where x_f denotes the latent image features preferred by vision tasks and x_{bg} denotes background information, combined via the operator \oplus , the task-driven lens design objective can further transform the imaging problem into a feature encoding problem, with the preferable image features x_f defined by the computer vision models. Importantly, modern computer vision models are usually trained with massive image datasets that include various image degradations as data augmentation, thereby preventing the lens design from converging to that of a perfect imaging lens.

2.2. Differentiable point spread function

The point spread function (PSF) characterizes the blurring effect of an optical system on a point light source. In image simulation, convolution of the PSF with the object image allows for simulating the image formation process of an optical lens. To compute the PSF, we trace optical rays emanating from point sources in the object space to the sensor plane. As illustrated in Figure 1, the PSF is then computed by integrating the energy deposited by the ray intersections over the sensor pixels. This process can be mathematically expressed as:

$$\operatorname{PSF}(\mathbf{o}_p) = \sum_{i=1}^N u_i \sigma\left(\frac{|(\mathbf{o}_p - \mathbf{o}_i) \cdot \hat{\mathbf{e}}_x|}{L}\right) \sigma\left(\frac{|(\mathbf{o}_p - \mathbf{o}_i) \cdot \hat{\mathbf{e}}_y|}{L}\right), \quad (5)$$

where \mathbf{o}_p represents the pixel coordinate, \mathbf{o}_i is the intersection point of the i -th ray with the sensor plane, and N denotes the number of rays traced from the point source. The scalar u_i signifies the radiant energy carried by the i -th ray, assigned a value of 1 in the experiments conducted. The vectors $\hat{\mathbf{e}}_x$ and $\hat{\mathbf{e}}_y$ are unit vectors situated within the sensor plane, and L denotes the physical width of a sensor pixel. The weighting function $\sigma(\cdot)$ is defined as:

$$\sigma(x) = \begin{cases} 1 - x, & 0 \leq x \leq 1 \\ 0, & \text{otherwise} \end{cases}, \quad (6)$$

which quantifies the contribution of a ray to its neighboring pixels. As depicted in Figure 1, the energy of each ray is distributed among its four nearest neighboring pixels. Thus, Equation (5) can be interpreted as an instance of inverse bilinear interpolation. By leveraging this sub-pixel information, we can approximate the continuous light distribution using a finite number of rays. Since Equation (5) is differentiable, we can compute the gradient of the PSF concerning the ray positions. During the backpropagation phase, these gradients direct the rays toward the target pixels, facilitating updates through the optical rays to the lens surfaces, which in turn adjusts the lens parameters.

3. Results

In this section, we first assess the effectiveness of the task-driven lens design objective for optimizing image classification lenses from scratch (Sections 3.1 and 3.2). We demonstrate that lenses designed with specific computer vision objectives can achieve competitive, or even superior, image classification accuracy with fewer optical elements compared to classical lens designs (Section 3.2). Subsequently, we extend this task-driven design approach to a variety of downstream computer vision tasks, including object detection, semantic segmentation, and vision-language models. We examine the underlying similarities across different computer vision tasks (Section 3.3) and their preferences regarding optical feature encoding, demonstrating the potential to design task-specific lenses for simpler tasks that generalize to more complex computer vision tasks.

3.1. Implementation details

We conduct differentiable ray tracing for optical simulation using the open-source simulator DEEPLENS [17, 22]. Each lens surface is modeled as an asphere, with curvature, axial position, and fourth- to tenth-order polynomial coefficients (α_4 to α_{10}) treated as optimizable parameters. For lens parameter optimization, we employ the Adam optimizer [37], setting the learning rate to 1×10^{-4} for curvature, position, and α_4 coefficients. An exponential decay factor of 0.02 is applied to the learning rates of the higher-order coefficients. To prevent self-intersections, we impose penalties on the distances between lens surfaces. The lens materials are selected from a library of common smartphone-grade polymers. The image sensor has a resolution of 1080×1920 pixels and a pixel pitch of $1.8 \mu\text{m}$. For the calculation of the PSF, we trace rays at three wavelengths: 656.3 nm, 589.3 nm, and 486.1 nm. More implementation settings are detailed in the supplementary material. Instead of simulating full sensor resolution images, we consider low-resolution image patches for different image regions covering the full field of view (FoV). This choice stems from the fact that modern computer vision models are typically trained on low-resolution images (e.g., 224×224), even though camera sensors often possess megapixel resolutions (e.g., 12 MP). This discrepancy in resolution necessitates simulating low-resolution images for task-driven lens design. During training, we randomly sample FoVs and, for each FoV, we trace 4096 rays and compute a 101×101 PSF kernel. Starting from a random initialization, the lens is trained for approximately 3,000 iterations, with the downstream network held constant. For evaluation, we select nine FoVs to represent the lens’s performance.

3.2. Task-driven lens design for image classification

We first demonstrate the effectiveness of task-driven lens design by designing lenses for image classification, a fundamental task in computer vision that is essential for image understanding and feature encoding. Utilizing a well-trained ResNet-50 network [1], we automatically design three distinct lenses comprising two, three, and four elements, entirely without human intervention (Figure 2). All three lens designs achieve a FoV of 68.8° and an F-number of 2.8. This success validates the effectiveness of employing gradients derived from network outputs and the robustness of the proposed task-driven lens design methodology. By initiating design from scratch and employing automated optimization, we circumvent the potential local minima associated with human expertise, thereby exploring a broader design space that extends beyond classical lens design paradigms.

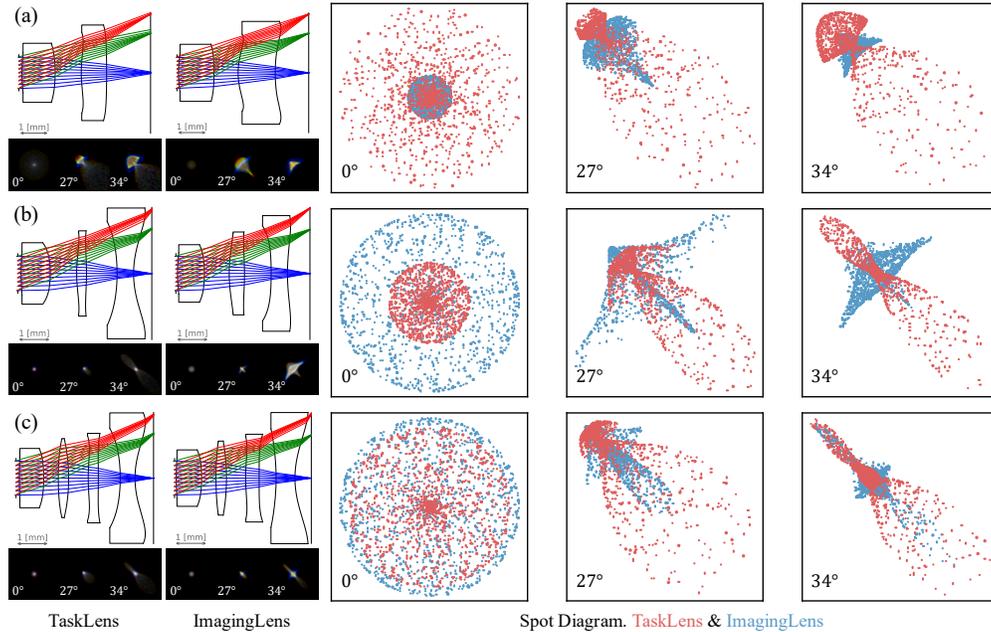


Fig. 2. The lens layout, PSFs and corresponding spot diagrams at different fields are shown for 2P (a), 3P (b), and 4P (c) lenses. Although RMS spot size of the TaskLens at some fields is larger than that of the ImagingLens, the majority of optical rays always converge within a small region. This spot diagram distribution results in a long-tailed PSF, which is characterized by a small, concentrated center and sparsely populated outer regions. In scenarios where the optical system cannot fully correct all aberrations, this long-tailed PSF with a sharp central peak is effective in preserving essential features of the object images, which benefits the performance of computer vision models.

We compare each of our task-driven lenses, referred to as “TaskLens”, against three lenses designed with classical objectives focused on minimizing aberrations, labeled as “ImagingLens”. The first ImagingLens (#1) is optimized using the open-source DEEPLENS code [22] from scratch by minimizing the root-mean-square (RMS) spot size across 11×11 FoVs. The remaining two ImagingLens (#2 and #3) are designed in Zemax by experienced optical engineers who exerted their best efforts to optimize the performance, also starting from scratch. We use the default merit function (RMS spot radius, x+y relative to the centroid) with 16 rings and 12 arms, with thickness constraints to prevent negative values. The entrance pupil diameter is fixed, and the F-number controls first-order properties. All radii, aspheric coefficients, and thicknesses are set

Table 1. Performance metrics grouped by the number of lens elements. The “TaskLens” design is compared against three different “ImagingLens” designs within each group. The various evaluation metrics do not consistently yield the same results; specifically, an optical lens that achieves the highest image classification accuracy may not necessarily demonstrate the best RMS spot size performance. Through task-driven lens design, it is feasible to develop optical lenses with simpler structures (fewer lens elements) that achieve comparable or even superior image classification accuracy when compared to classical lens designs that prioritize minimizing optical aberrations.

Lens Structure	Lens Design	Accuracy (%)	PSNR (dB)	Avg RMS Spot Size (μm)
2-Element	TaskLens	70.08	19.46	27.27
	ImagingLens #1	65.63	22.43	13.61
	ImagingLens #2	68.54	23.65	9.85
	ImagingLens #3	67.01	22.90	10.99
3-Element	TaskLens	73.40	23.85	10.41
	ImagingLens #1	70.04	22.77	9.32
	ImagingLens #2	69.92	23.14	14.75
	ImagingLens #3	68.52	23.51	8.53
4-Element	TaskLens	73.61	23.67	9.97
	ImagingLens #1	72.27	25.00	7.32
	ImagingLens #2	68.88	23.63	10.11
	ImagingLens #3	68.56	23.98	11.37

as variables. Detailed lens data are provided in the supplementary material.

Table 1 presents the average top-1 classification accuracy across nine evaluation FoVs for each lens. The TaskLens consistently outperforms the corresponding ImagingLens in terms of image classification accuracy when the number of lens elements is the same. Due to optical aberrations, real-world image classification performance falls short of the upper bound of 75.63%, which is calculated on high-quality testing images. The ImagingLens experiences a more pronounced performance drop, attributed to the misalignment between the design objectives and the evaluation metrics. Notably, the two-element TaskLens surpasses all three-element ImagingLenses, while the three-element TaskLens outshines all four-element ImagingLenses. This suggests that task-driven lens design can identify simpler optical structures with competitive, or even superior, computer vision performance.

To elucidate the reason for the performance improvement, we analyze the spot diagrams and PSFs of the TaskLens and the best-performing ImagingLens. As illustrated in Figure 2, both designs exhibit noticeable blur due to aberrations, which results from the limited number of optical elements that can only partially correct optical aberrations. The distinction lies in their aberration-correction strategies: the TaskLens focuses the majority of light into a compact core while allowing a small fraction to form extended, low-energy tails. This results in a distinctive long-tailed PSF, particularly evident in the three- and four-element TaskLenses. While this long-tailed distribution redistributes optical energy and leads to reduced contrast (manifesting as a “haze” in images), it offers a critical advantage: a sharp central peak. In scenarios with strict hardware constraints, perfectly correcting all aberrations is infeasible. Traditional designs (minimizing RMS spot size) often result in a broadened central spot to minimize outlier rays, which suppresses high-frequency spatial information. In contrast, the TaskLens optimization allows for a long-tailed distribution (larger RMS spot size) to maintain a highly concentrated core. This sharp peak effectively preserves high-frequency structural details (such as edges) that are vital for feature extraction in computer vision models. Although these aberrations degrade standard image quality metrics, our experiments suggest that deep networks are robust to the

Table 2. Cross-task performance evaluation matrix. Each row represents a different computer vision task for design, while the columns indicate the computer vision tasks used for evaluation. Bold values indicate the best performance for each evaluation task.

Design Task \ Evaluation Task	RMS spot size μm	Image Classification Acc. (%)	Object Detection mAP.	Semantic Segmentation mAP	Image-Text Retrieval Recall@1
RMS spot size	9.32	70.04	35.45	36.08	54.4
Image Classification	10.41	73.40	40.80	41.44	68.2
Object Detection	15.27	72.18	42.56	40.21	67.4
Semantic Segmentation	14.86	72.60	41.75	43.41	65.5
Image-Text Retrieval	12.90	71.21	42.25	42.32	70.4

global contrast reduction caused by the tail but are highly sensitive to the loss of high-frequency information, allowing the network to effectively learn to extract useful feature information during the task-driven design process.

In contrast, the ImagingLens, which is designed according to classical objectives and decoupled from the computer vision task, strives to focus all rays as tightly as possible. When residual optical aberrations are present, the classical lens design objectives become ineffective, lacking guidance regarding which image features are preferred by computer vision tasks. These findings suggest that for resource-constrained systems, such as camera lenses for robotics, a task-driven lens design objective that accommodates certain task-agnostic aberrations in favor of preserving task-critical features can result in superior end-to-end performance.

3.3. Various downstream computer vision tasks

To assess the versatility of our design approach, we evaluate the lenses across a broader range of downstream applications, including object detection, semantic segmentation, and image-text retrieval in VLMs. These tasks effectively represent real-world challenges in computer vision. For object detection, we employ the Faster R-CNN model with a ResNet-50-FPN backbone from Detectron2 [38, 39], trained and tested on the COCO 2017 dataset [40]. Performance is quantified using the mean Average Precision (mAP) metric, which evaluates a model’s ability to accurately detect and classify objects. In semantic segmentation, we use the Mask2Former model with a Swin-Small backbone from HuggingFace [41], also trained and tested on the COCO 2017 dataset. This model’s performance is similarly quantified using the mAP metric, which assesses its ability to accurately segment objects. For image-text retrieval, we utilize the CLIP model with a ViT-Large-Patch14 backbone from HuggingFace [7], trained and tested on the Flickr30k dataset [42]. In this context, optical lenses are trained to maximize the similarity between image encodings and text encodings via a contrastive loss. Performance is evaluated using the top-1 recall metric.

We employ the three-element “TaskLens” and “ImagingLens” (#2) from the previous image classification section and design new lenses for the other three computer vision tasks, maintaining consistent experimental settings. The experimental results indicate that task-driven lens design is effective across all three tasks when designing lenses from scratch. We subsequently evaluate the performance of all these lenses across all metrics, with results presented in Table 2. Our findings reveal that training for a specific computer vision task invariably produces superior performance during evaluation. Moreover, the four computer vision tasks exhibit notable similarities. Lenses designed for one task perform well on others, significantly outperforming those designed with classical objectives. We believe this success stems from the need for accurate preservation of image textures and high-frequency features across these tasks, highlighting their generalizability. This finding further inspires the design of computer vision lenses by starting with simpler tasks (e.g., image classification) that can generalize to more complex tasks (e.g., VLM), which often present challenges for stable convergence. Example qualitative visualization examples for object

detection and segmentation are provided in the supplementary material, demonstrating the network outputs of the TaskLens and ImagingLens.

3.4. *Optical simulation validation*

To validate the fidelity of our optical simulation framework, we conducted experimental comparisons using a commercial camera system: a Canon EOS R6 paired with a Canon RF50mm f/1.8 lens. Experiments were performed in a controlled laboratory setting. We set the aperture to f/1.8 to maximize optical aberrations, thereby providing a more rigorous test of the simulator’s capabilities. Both the focus distance and the object plane were set to 1.0 m. We employed two standard metrics for optical characterization: the Modulation Transfer Function (MTF) and the Point Spread Function (PSF). MTF quantifies the system’s resolution and contrast transfer, while the PSF characterizes the impulse response and aberration profile. Data acquisition was performed in RAW format to avoid non-linear ISP processing. We used the “rawpy” library [43] to linearize the images, ensuring that the pixel values remained proportional to the scene illuminance.

For MTF validation, we imaged a standard test chart. We extracted a region of interest (ROI) and computed the MTF following the methodology described in [44]. We then simulated the same scene and optical parameters in our simulator. As shown in Fig. 3, the measured MTF curves exhibit a strong correlation with the simulated results across the spatial frequency spectrum. Minor deviations are attributed to manufacturing tolerances of the physical lens and sensor noise; however, the overall trend validates the frequency response modeling of our simulator.

To validate the PSF, we measured the system’s impulse response at three distinct fields of view (FoVs) to capture field-dependent variations. We generated a pseudo-point source by displaying a 2×2 pixel bright spot on a high-resolution iPad display. While not an ideal point source, this approximation allows for practical empirical measurement. To ensure a fair comparison, the simulated PSFs were convolved with a 2×2 pixel kernel matching the physical source dimensions. Figure 3 presents the comparison: the simulated PSFs closely reproduce the morphology and energy distribution of the captured PSFs. This strong agreement confirms the simulator’s accuracy in modeling complex optical aberrations.

4. **Extra validation and robustness analyses**

We present additional validation experiments that assess the robustness and practical utility of the proposed task-driven lens design. Specifically, we analyze manufacturing tolerance, examine generalization across different network architectures, compare against conventional end-to-end optical co-design, and evaluate performance with and without post-capture image restoration.

4.1. *Lens manufacturing tolerance analysis*

Manufacturing and assembly imperfections are inevitable in practical applications and can significantly degrade the performance of optical systems, subsequently impacting downstream applications. Consequently, evaluating a lens design’s tolerance to such errors is a critical component of the design process. To assess the robustness of our designs, we simulated the effects of manufacturing and assembly errors by generating three randomly “perturbed” versions of each lens. The classification performance of these perturbed lenses was then evaluated using a neural network model trained on the original, unperturbed design.

The results, as presented in Table 3, illustrate the performance of each lens before and after the introduction of random errors. A comparative analysis of the average classification accuracy reveals that our TaskLens consistently outperforms the conventional ImagingLens, even in the presence of these perturbations. Notably, TaskLens designs exhibit superior robustness to manufacturing tolerances, demonstrating a smaller decrease in performance relative to their ImagingLens counterparts. A striking example is observed in the three-element comparison, where the ImagingLens experiences an average performance degradation of 3.77%, whereas

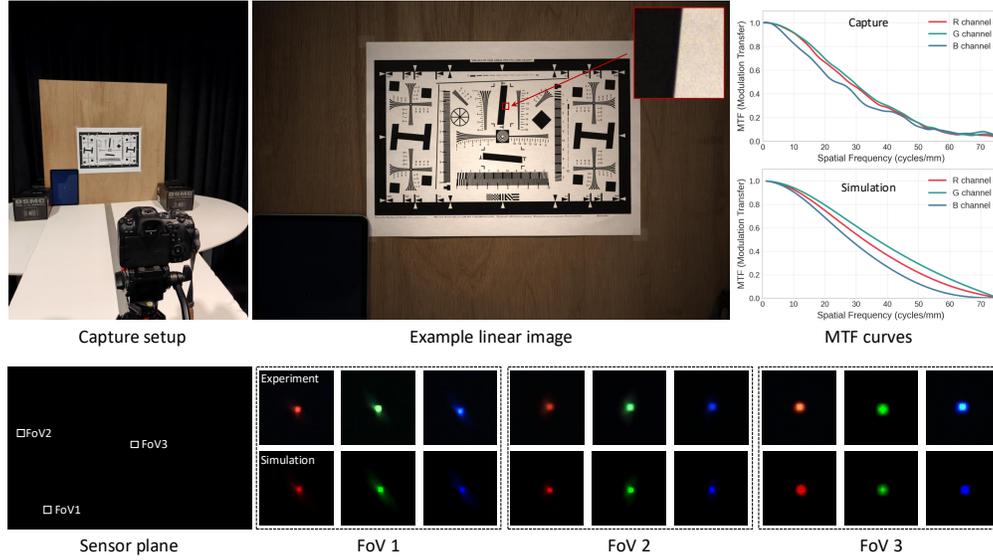


Fig. 3. Validation of the optical simulation pipeline using a Canon EOS R6 and RF50mm f/1.8 lens. The experimental setup utilizes a test chart for MTF calculation and a monitor-displayed point source for PSF measurement. The measured MTF curves (top) track the simulated values closely, confirming accurate frequency response modeling. The captured PSFs (bottom) at three FoV exhibit strong morphological agreement with the simulation, validating the modeling of field-dependent aberrations.

the three-element TaskLens only shows a minor decrease of 0.56%. This enhanced tolerance to fabrication errors constitutes a significant practical advantage. We attribute this robustness to our task-driven design philosophy, which does not necessitate perfect optical aberration correction, thereby making it more tolerant of the minor, unintended aberrations introduced during manufacturing and assembly processes.

Table 3. Image classification accuracy of different lenses with random manufacturing and assembly errors. From left to right: 2-element TaskLens, 2-element ImagingLens (#2), 3-element TaskLens, 3-element ImagingLens (#1), 4-element TaskLens, 4-element ImagingLens (#1).

Lens	2-Element		3-Element		4-Element	
	TaskLens	ImagingLens (#2)	TaskLens	ImagingLens (#1)	TaskLens	ImagingLens (#1)
Designed	70.08%	68.54%	73.40%	70.04%	73.61%	72.27%
Lens #1	69.73%	67.43%	72.82%	66.24%	72.58%	71.07%
Lens #2	69.65%	67.27%	72.91%	66.36%	72.66%	70.99%
Lens #3	69.83%	67.23%	72.78%	66.21%	72.65%	70.98%
Avg (#1 ~ #3)	69.74%	67.31%	72.84%	66.27%	72.63%	71.10%
Avg decrease	-0.34%	-1.23%	-0.56%	-3.77%	-0.98%	-1.17%

Designed: The ideal lens design without manufacturing or assembly errors.

Lens #1-3: Three instances of the designed lens with random manufacturing and assembly errors.

4.2. Performance generalization across different network architectures

In our task-driven lens design, the lens is optimized specifically for a particular network architecture, namely ResNet-50. However, in practical applications, the downstream network may vary. For instance, smaller models might be deployed on edge devices with limited computational resources, while larger models may be employed in cloud environments to enhance performance. Since optical lenses cannot be modified post-manufacture, it is essential to ascertain whether the performance advantages of a TaskLens can be preserved when integrated with different network architectures.

To investigate this, we evaluated our lenses using three additional network models that vary significantly in size and architecture: MobileNetV3-Large [45] (“MobileNetV3-L”), a 5.5M-parameter CNN; Swin Transformer-Base [2] (“Swin-B”), an 88M-parameter Transformer; and Vision Transformer-Large [46] (“ViT-L/16”), a 304M-parameter Transformer. These are compared against the original 26M-parameter ResNet-50, which is also a CNN architecture. The image classification results are summarized in Table 4. Across all tested network models, our TaskLens consistently achieves the highest accuracy for each specified number of lens elements. This finding demonstrates that the TaskLens is compatible with various network architectures and remains robustly advantageous in terms of performance. Moreover, this result bears practical implications: one can utilize a smaller, computationally efficient network during the lens optimization phase, thereby reducing design time and resource requirements, and subsequently deploy the optimized lens with a larger, more powerful network to maximize overall performance.

Table 4. Classification accuracy with different network models. The designed TaskLenses are compatible with different network models while maintaining enhanced performance compared to the corresponding ImagingLenses. The results for the original ResNet50 are reported in Table 1. Red/blue highlighting indicates best/second-best performance within each configuration. These results support the claim that task-driven lens design can identify optical characteristics that are compatible with different network architectures.

Lens Elements	Method	MobileNetV3-L [47] (5.4M params)	Swin-B [2] (88M params)	ViT-L/16 [46] (304M params)
2-Element	TaskLens (Ours)	68.22%	81.19%	81.76%
	ImagingLens #1	64.60%	79.03%	79.39%
	ImagingLens #2	67.73%	80.87%	81.19%
	ImagingLens #3	66.05%	80.23%	80.79%
3-Element	TaskLens (Ours)	71.82%	82.65%	83.46%
	ImagingLens #1	68.36%	81.19%	81.62%
	ImagingLens #2	67.94%	81.08%	81.46%
	ImagingLens #3	68.01%	80.19%	81.61%
4-Element	TaskLens (Ours)	72.06%	82.82%	83.62%
	ImagingLens #1	71.00%	82.43%	82.52%
	ImagingLens #2	66.78%	80.75%	81.60%
	ImagingLens #3	67.19%	80.46%	81.50%

4.3. Comparison with end-to-end optical design from successful starting points

In this section, we investigate whether an end-to-end optical design approach can identify solutions comparable to our TaskLens. To test this hypothesis, we conducted a joint optimization of both the optical system and a pre-trained ResNet-50 network, which is a standard methodology for end-to-end design. We evaluated two distinct starting points for the lens: (1) a randomly initialized optical lens, representing a design from scratch, and (2) our best-performing, pre-

optimized ImagingLens. The final classification accuracies are summarized in Table 5. Our results indicate that the end-to-end training process fails to converge when initiated from an all-flat optics configuration. This phenomenon can be attributed to the highly non-convex optimization landscape inherent to optical design, which makes convergence from a random or null starting point particularly challenging.

When the training is initialized with a well-designed ImagingLens, the end-to-end optimization achieves only marginal performance improvements and fails to reach the accuracy of our TaskLens. We hypothesize that this occurs because the optimization becomes trapped in a strong local minimum. The initial ImagingLens configuration is already highly optimized for traditional objectives. Consequently, the backpropagated gradients are typically insufficient to move the design out of this local minimum towards a more globally optimal solution discovered by our method. As a result, the end-to-end process primarily fine-tunes the initial lens design, yielding classification performance that remains inferior to that of our TaskLens. This experiment underscores the limitations of conventional end-to-end fine-tuning and highlights the effectiveness of our design-space exploration in uncovering novel, task-optimal optical configurations.

Table 5. End-to-end design from scratch fails to converge. When starting with a well-designed ImagingLens, the end-to-end process fails to discover the optimal classification lens.

	TaskLens	End-to-End Training	
		From ImagingLens	From Scratch
2-Element	70.05%	69.55%	✗
3-Element	73.40%	71.94%	✗
4-Element	73.61%	73.44%	✗

4.4. Effect of post-capture image restoration

Image restoration is a widely employed technique for mitigating optical aberrations in captured images. Modern algorithms, particularly those based on neural networks [44, 48], have demonstrated effectiveness in recovering fine image details, thereby enhancing the performance of downstream vision tasks. This raises a critical question: Can post-capture image restoration bridge the performance gap between our TaskLenses and conventional ImagingLenses?

To address this question, we applied an image restoration process to the outputs of all lenses and compared the resulting classification accuracies. For image restoration, we utilized NAFNet [49], a state-of-the-art restoration model, and we trained a separate NAFNet for each lens on its simulated images. The results are summarized in Table 6. While image restoration consistently enhances classification accuracy across all lenses, the TaskLens continues to outperform the corresponding ImagingLens. This finding suggests that the performance advantage of the TaskLens is not solely attributable to a type of blur that can be easily corrected by an image restoration algorithm. Instead, the task-specific optical encoding provides a persistent advantage, indicating that the gap in classification performance cannot be bridged by image restoration alone.

5. Discussion and conclusion

Application scenarios In applications where high-quality, perfectly imaging lenses are infeasible due to form-factor and cost constraints (e.g., robotics), residual aberrations can significantly affect computer vision performance. In these scenarios, simple lenses designed with our task-driven objective bridge the gap between optical design and downstream tasks. This approach better preserves features preferred by computer vision models, improving performance despite

Table 6. Classification performance after image restoration. Improvements are observed in both PSNR and classification accuracy for all designs. Despite the TaskLens do not demonstrate the best PSNR scores, they maintain the best image classification performance. This indicates that the image classification advantages stem from fundamental optical properties rather than aberration tolerance. Red/blue highlighting indicates best/second-best performance within each configuration.

Lens Elements	Method	PSNR [dB]	Classification Accuracy
2-Element	TaskLens (Ours)	27.24	72.03%
	ImagingLens #1	27.54	68.19%
	ImagingLens #2	30.30	71.24%
	ImagingLens #3	30.87	71.08%
3-Element	TaskLens (Ours)	32.31	74.43%
	ImagingLens #1	29.44	72.42%
	ImagingLens #2	30.47	73.35%
	ImagingLens #3	29.95	70.90%
4-Element	TaskLens (Ours)	33.58	74.61%
	ImagingLens #1	34.29	73.98%
	ImagingLens #2	29.69	72.16%
	ImagingLens #3	29.38	71.52%

aberrations. Furthermore, while fine-tuning vision models on corrupted images is a common approach to enhance performance, it poses extra computational cost and may be unaffordable in some cases as modern network architectures become more complex and training datasets grow larger. In scenarios where users directly deploy pre-trained models, our task-driven lens design objective offers a valuable alternative for designing lenses that are inherently more compatible with computer vision models than those designed with traditional objectives.

Challenges and future work Our experiments reveal that as the complexity of computer vision models increases, the gradients provided by the networks become more unstable, leading to oscillatory optimization behavior. For instance, when training with the BLIP model [8] for image-text retrieval, we encountered significant convergence difficulties with our task-driven lens design. We believe this struggle arises from the noise in gradients produced by the complex network architecture during backpropagation. Moreover, propagating gradients through the entire network poses challenges related to GPU memory consumption. These issues not only complicate direct optimization for certain modern computer vision tasks and network architectures but also motivate us to explore more general evaluative objectives for the feature encoding capabilities of optical lenses. Consequently, we are currently undertaking follow-up work to simplify the task-driven lens design objective while elucidating key underlying principles for task-driven lens design.

Conclusion In this work, we introduce task-driven lens design. By leveraging a well-trained neural network to inform optical optimization, we demonstrate the feasibility of designing simpler yet more effective lenses from scratch for computer vision applications. When lens structures alone are insufficient to correct all optical aberrations, our resulting “TaskLenses” consistently outperform those created using classical design objectives. The unique optical characteristics and PSF effectively preserve critical image features identified by vision foundation models, and the TaskLenses exhibit greater tolerance for manufacturing errors. The generalizability across various computer vision tasks and network architectures encourages optimization using simpler tasks and networks, with the potential to scale to more complex tasks and networks. We believe

that task-driven lens design presents a new paradigm for the next generation of computational camera lenses, particularly for optical lenses constrained by practical considerations such as form factor.

Funding. King Abdullah University of Science and Technology (Individual Baseline Funding); King Abdullah University of Science and Technology (Center of Excellence for Generative AI); Fonds Wetenschappelijk Onderzoek (G0A3O24N); Vrije Universiteit Brussel (Hercules, Methusalem, OZR).

Disclosures. The authors declare no conflicts of interest.

Data availability. Data underlying the results presented in this paper are not publicly available at this time but may be obtained from the authors upon reasonable request.

Supplemental document. See Supplement 1 for supporting content.

References

1. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, (2016), pp. 770–778.
2. Z. Liu, H. Hu, Y. Lin, *et al.*, "Swin transformer v2: Scaling up capacity and resolution," in *CVPR*, (2022), pp. 12009–12019.
3. J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *CVPR*, (2016), pp. 779–788.
4. H. Zhang, F. Li, S. Liu, *et al.*, "DINO: DETR with improved denoising anchor boxes for end-to-end object detection," arXiv preprint arXiv:2203.03605 (2022).
5. Z. Liu, Y. Lin, Y. Cao, *et al.*, "Swin transformer: Hierarchical vision transformer using shifted windows," in *ICCV*, (2021), pp. 10012–10022.
6. X. Zhai, A. Kolesnikov, N. Houlsby, and L. Beyer, "Scaling vision transformers," in *CVPR*, (2022), pp. 12104–12113.
7. A. Radford, J. W. Kim, C. Hallacy, *et al.*, "Learning transferable visual models from natural language supervision," in *International conference on machine learning*, (PmLR, 2021), pp. 8748–8763.
8. J. Li, D. Li, C. Xiong, and S. Hoi, "Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation," in *International conference on machine learning*, (PMLR, 2022), pp. 12888–12900.
9. J. Zhang, J. Huang, S. Jin, and S. Lu, "Vision-language models for vision tasks: A survey," *IEEE transactions on pattern analysis machine intelligence* **46**, 5625–5644 (2024).
10. G. Shabtay, E. Goldenberg, M. Dror, *et al.*, "Folded camera lens designs," (2020). US Patent 10,571,644.
11. X. Zhou, "Camera optical lens," (2023). US Patent 11,567,301.
12. H. Abe, "Zoom lens and image pickup apparatus including the same," (2018). US Patent 9,952,446.
13. S. Sugita, "Zoom lens and image pickup apparatus including same," (2015). US Patent 9,110,278.
14. V. Sitzmann, S. Diamond, Y. Peng, *et al.*, "End-to-end optimization of optics and image processing for achromatic extended depth of field and super-resolution imaging," *ACM Trans. Graph.* **37**, 1–13 (2018).
15. Q. Sun, C. Wang, F. Qiang, *et al.*, "End-to-end complex lens design with differentiable ray tracing," *ACM Trans. Graph.* **40**, 1–13 (2021).
16. E. Tseng, A. Mosleh, F. Mannan, *et al.*, "Differentiable compound optics and processing pipeline optimization for end-to-end camera design," *ACM Trans. Graph.* **40**, 1–19 (2021).
17. C. Wang, N. Chen, and W. Heidrich, "do: A differentiable engine for deep lens design of computational imaging systems," *IEEE Trans. on Comput. Imaging* **8**, 905–916 (2022).
18. Q. Jiang, S. Gao, Y. Gao, *et al.*, "Minimalist and high-quality panoramic imaging with psf-aware transformers," *IEEE Trans. on Image Process.* **33**, 4568–4583 (2024).
19. M. Na, H. A. Jimenez Romero, X. Yang, *et al.*, "End-to-end optimization of fluidic lenses," in *SIGGRAPH Asia 2024 Conference Papers*, (ACM, 2024), SA '24, p. 1–10.
20. E. Tseng, S. Colburn, J. Whitehead, *et al.*, "Neural nano-optics for high-quality thin lens imaging," *Nat. Commun.* **12**, 6493 (2021).
21. Z. Li, Q. Hou, Z. Wang, *et al.*, "End-to-end learned single lens design using fast differentiable ray tracing," *Opt. Lett.* **46**, 5453–5456 (2021).
22. X. Yang, Q. Fu, and W. Heidrich, "Curriculum learning for ab initio deep learned refractive optics," *Nat. Commun.* **15**, 50835 (2024).
23. Y. Liu, C. Zhang, T. Kou, *et al.*, "End-to-end computational optics with a singlet lens for large depth-of-field imaging," *Opt. Express* **29**, 28530–28548 (2021).
24. X. Yang, M. Souza, K. Wang, *et al.*, "End-to-end hybrid refractive-diffractive lens design with differentiable ray-wave model," in *SIGGRAPH Asia 2024 Conference Papers*, (ACM, 2024), SA '24, p. 1–11.
25. D. S. Jeon, S.-H. Baek, S. Yi, *et al.*, "Compact snapshot hyperspectral imaging with diffracted rotation," *ACM Trans. Graph.* **38**, 1–13 (2019).
26. X. Dun, H. Ikoma, G. Wetzstein, *et al.*, "Learned rotationally symmetric diffractive achromat for full-spectrum computational imaging," *Optica* **7**, 913–922 (2020).

27. S.-H. Baek, H. Ikoma, D. S. Jeon, *et al.*, “Single-shot hyperspectral-depth imaging with learned diffractive optics,” in *ICCV*, (2021), pp. 2651–2660.
28. L. Li, L. Wang, W. Song, *et al.*, “Quantization-aware deep optics for diffractive snapshot hyperspectral imaging,” in *CVPR*, (2022), pp. 19780–19789.
29. C. A. Metzler, H. Ikoma, Y. Peng, and G. Wetzstein, “Deep optics for single-shot high-dynamic-range imaging,” in *CVPR*, (2020), pp. 1375–1385.
30. Q. Sun, E. Tseng, Q. Fu, *et al.*, “Learning Rank-1 diffractive optics for single-shot high dynamic range imaging,” in *CVPR*, (2020), pp. 1386–1396.
31. G. Côté, F. Mannan, S. Thibault, *et al.*, “The differentiable lens: Compound lens search over glass surfaces and materials for object detection,” in *CVPR*, (2023), pp. 20803–20812.
32. J. Chang and G. Wetzstein, “Deep optics for monocular depth estimation and 3D object detection,” in *ICCV*, (2019), pp. 10193–10202.
33. H. Ikoma, C. M. Nguyen, C. A. Metzler, *et al.*, “Depth from defocus with learned optics for imaging and occlusion-aware depth estimation,” in *Int. Conf. Comput. Photog.*, (IEEE, 2021), pp. 1–12.
34. G. Côté, J.-F. Lalonde, and S. Thibault, “Deep learning-enabled framework for automatic lens design starting point generation,” *Opt. Express* **29**, 3841–3854 (2021).
35. A. Teh, D. Vicini, B. Bickel, *et al.*, “Automated design of compound lenses with discrete-continuous optimization,” arXiv preprint arXiv:2509.23572 (2025).
36. J. Deng, W. Dong, R. Socher, *et al.*, “ImageNet: A large-scale hierarchical image database,” in *CVPR*, (Ieee, 2009), pp. 248–255.
37. I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” arXiv preprint arXiv:1711.05101 (2017).
38. R. Girshick, “Fast r-cnn,” in *Proceedings of the IEEE international conference on computer vision*, (2015), pp. 1440–1448.
39. Y. Wu, A. Kirillov, F. Massa, *et al.*, “Detectron2,” <https://github.com/facebookresearch/detectron2> (2019).
40. T.-Y. Lin, M. Maire, S. Belongie, *et al.*, “Microsoft coco: Common objects in context,” in *Computer vision—ECCV 2014: 13th European conference, zurich, Switzerland, September 6–12, 2014, proceedings, part v 13*, (Springer, 2014), pp. 740–755.
41. B. Cheng, I. Misra, A. G. Schwing, *et al.*, “Masked-attention mask transformer for universal image segmentation,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, (2022), pp. 1290–1299.
42. P. Young, A. Lai, M. Hodosh, and J. Hockenmaier, “From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions,” *Trans. association for computational linguistics* **2**, 67–78 (2014).
43. rawpy contributors, “rawpy,” [Software] <https://github.com/letmaik/rawpy> (accessed 22 Jan. 2026) (2026).
44. S. Chen, T. Lin, H. Feng, *et al.*, “Computational optics for mobile terminals in mass production,” *IEEE Trans. on Pattern Anal. Mach. Intell.* **45**, 4245–4259 (2023).
45. A. G. Howard, M. Zhu, B. Chen, *et al.*, “MobileNets: Efficient convolutional neural networks for mobile vision applications,” arXiv preprint arXiv:1704.04861 (2017).
46. A. Dosovitskiy, L. Beyer, A. Kolesnikov, *et al.*, “An image is worth 16x16 words: Transformers for image recognition at scale,” arXiv preprint arXiv:2010.11929 (2020).
47. A. Howard, M. Sandler, G. Chu, *et al.*, “Searching for MobileNetv3,” in *ICCV*, (2019), pp. 1314–1324.
48. X. Yang, C. Nguyen, W. Wang, *et al.*, “Efficient depth-and spatially-varying image simulation for defocus deblur,” arXiv preprint arXiv:2507.00372 (2025).
49. L. Chen, X. Chu, X. Zhang, and J. Sun, “Simple baselines for image restoration,” in *Computer Vision – ECCV 2022*, (Springer Nature Switzerland, 2022), pp. 17–33.