Active Architecture Optical Systems with Kaleid for unified internal prescription control

YUEZE WANG[1], CHE-HANG YU[1], AND SPENCER LAVERE SMITH[1]

[1] University of California Santa Barbara

Active Architecture Optical Systems (AAOS) is a platform technology in which lenses or lens groups in a multi-lens assembly can be precisely re-positioned along multiple degrees-of-freedom (DOFs) during operation, to optimize performance at application time. These degrees of freedom can make the optimization process complex. Here we show that the high dimensional solution space can be efficiently navigated using a machine learning approach. The approach we developed, called Kaleid, combines learned, physically grouped motion proposals with measured corrections across dozens of DOFs. We use detailed modeling, including realistic noise, to show that Kaleid can recover diffraction-limited performance across a range of systems and metrics: a multiphoton dry objective under a nine field location wavefront-error criterion, a high numerical aperture immersion objective under a central wavefront-error criterion, and a deep ultraviolet semiconductor inspection lens system under a spatial-frequency contrast criterion. We also assessed five comparison procedures including the current state-of-the-art, and none were as successful as Kaleid across the test cases. Together, this analysis demonstrates a feasible, robust, multi-application machine learning approach for AAOS, enabling application-time, sub-micron lens position optimization in an array of imaging modalities.

1. INTRODUCTION

High performance imaging components, including high-numerical aperture (NA) objectives, are built around a prescription that becomes fixed at assembly. Tolerance analysis, precision manufacturing, active alignment, and quality control reduce manufacturing and assembly errors [13, 14], but at application time, the elements are not necessarily in the same configuration. Mechanical shocks during shipping and use, temperature changes, and gravity can all lead to position changes that reduce performance. Moreover, at application time, there can be differences in conditions that alter the optimal element positions. Correction collars for spherical aberration are a familiar 1D adjustment. In addition, imaging depth, sample properties, immersion media refractive index, temperature, excitation light wavelengths, and other factors can degrade image quality [1–3]. These factors result in changes to the wavefront and suboptimal performance metrics, e.g., lower resolution, more severe vignetting, and weaker signals. These real world challenges influence optical designers, in that lens designs that require high precision (submicron) lens positioning suffer from low yield in manufacturing, and are not feasible to produce commercially. Even with active alignment, lenses cannot be reliably positioned to submicron precision, and normal shipping and use can result in component movement that degrades performance.

Conventional adaptive optics (AO) is an effective mitigation when the main problem can be expressed as a phase error at one pupil or aberrator-conjugate plane [4–6]. A deformable mirror can update quickly and without mechanically moving an objective. Its correction, however, is limited by stroke, actuator sampling, clear aperture, and the optical relay that places the device at the desired conjugate [7]. Many applications of AO are limited to relatively minor corrections and offer only small fields of view due to anisoplanatism. Outside this correction regime, AO can fail to restore image quality and may even degrade it by adding high-order phase ripples.

With Active Architecture Optical Systems (AAOS) multiple movement DOFs are retained for lenses within a multi-element module. The positions of internal optical elements can be adjusted during operation, down to submicron precision, to optimize performance at application time. This both relieves manufacturing challenges, making tighter tolerance designs feasible and maintaining their performance despite shocks or drift, and it unlocks more extensive optimization potential at application time.

Previous work has developed machine learning methods for AO and alignment [8, 9, 12]. Hu et al. demonstrated one embedded network design across two-photon, three-photon, and structured-illumination microscopes [11]. In those experiments the network estimates 5-9 Zernike coefficients for one pupil location, for AO correction with a deformable mirror (DM) or spatial light modulator (SLM). Slor et al. used machine learning to diagnose misalignments in multi-lens systems, estimating 20

mechanical coordinates from ray tracing or 16 coordinates from camera images in an image-based application [10]. These results established machine learning approaches for pupil correction and mechanical diagnosis. The next step to enable AAOS is to implement control and scale the DOFs. Here, we developed methods we call Kaleid to control 35-40 coupled coordinates distributed through an optical system, guided by multi-field measurements and fine tuned at application time using in-task metrics.

This work evaluates Kaleid across three optical systems: A large field-of-view two-photon microscope objective (Trepan2p) with wavefront recovery across the field-of-view, a high-NA microscope objective (Nikon) with on-axis wavefront recovery, and a deep-ultraviolet inspection objective (KLA) with recovery of contrast and its spatial uniformity. We compare the complete procedure with several alternative approaches.

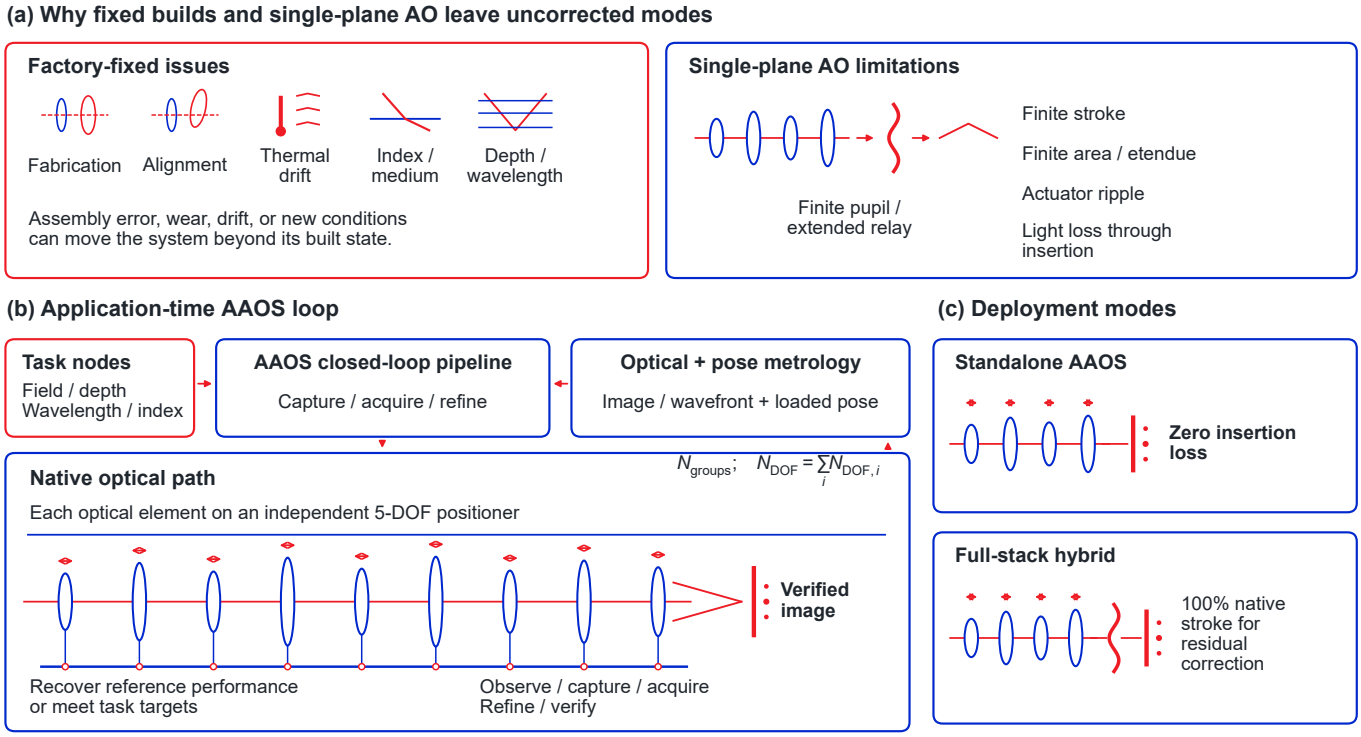


Fig. 1. AAOS rationale. (a) Limits of factory-fixed optics and single-plane AO. (b) Closed-loop control with each optical element on an independent 5-DOF positioner. (c) Standalone and hybrid deployment.

2. OPTICAL OVERVIEW

AAOS specifies the movable optical elements (lenses or lens groups, or mirrors), their controllable motions (e.g., 5-6 degrees of freedom, DOF), sensors (type and placement), and control system. Optionally, a high DOF AO element can be included in the AAOS stack (similar to woofer-tweeter style AO) [37]. The AAOS control system we present here, called Kaleid, controls the AAOS system through sensing, learned estimation, and measured correction.

Moving an optical element or rigid group changes aberrations and the ray footprint on downstream optics, which complicates control. For example, motions that look similar on axis can have divergent results at the edge of the field, and separately predicted corrections can interact when applied together. Kaleid predicts physically grouped motions, combines them into one system coordinate vector, and measures the combined optical response before refinement.

Three terms are used throughout: a *group* contains one or more optical elements (usually lenses) moved as one unit; a *subsystem* contains groups treated together for observation or update (e.g., a scan lens); and the *optical system* is the complete traced system, including the imaged sample. For example, the Nikon test contains seven actuated groups of lens elements, L1-L6 and L7′ [29]. The coverslip, immersion media, and tube lens remain in the optical model but are not positioned by the AAOS system.

In laser scanning imaging, we omit piston, tip, tilt, and defocus from the RMS wavefront error (WRMS) calculation. This is because piston typically has no effect on imaging, and tip, tilt, and defocus simply move the focus in space (repositioning and refocusing) and do not add aberrations, until extreme values, which are avoided in this study. The Trepan2p test uses a nine-point mean WRMS; the Nikon test uses the absolute center WRMS and reports edge WRMS separately. The KLA test uses nine-point Score_{2000} to reward contrast and uniformity.

Table 1. Three optical systems addressed by the same Kaleid proposal–measure–refine framework. Microscopy WRMS is in waves after piston, tip, tilt, and defocus removal.

System	Complete controlled state	Measurement supplied to Kaleid	Fixed-setting success criterion
Trepan2p	Seven lens elements and one coverglass, 40 internal coordinates	Diffuser-coded intensity over nine camera fields	Nine-field mean PTTD WRMS strictly below 0.07 waves
Nikon	Seven groups of lens elements, 35 internal coordinates	Differential pupil phase from three input fields	Center absolute PTTD WRMS below 0.07 waves; edge absolute WRMS is audited separately
KLA	Eight groups of optical elements, 40 internal coordinates	Pupil OPD over nine input fields	Nine-field $\text{Score}_{2000} \geq 0.45$; higher scores mean stronger, more uniform contrast

A. Incorporating DOFs in the optical prescription

In an ordinary imaging objective, lens spacings and centering errors become part of the as-built prescription. AAOS instead treats the element positions as measured coordinates. If group i contributes $N_{\text{DOF},i}$ useful degrees of freedom (DOFs), the complete coordinate vector is

$$\begin{aligned} \underbrace{\mathbf{u}}_{\substack{\text{all commanded and}\\ \text{encoder-read coordinates}}} &= \underbrace{\left[\mathbf{u}_1^{\mathsf{T}}, \ldots, \mathbf{u}_{N_{\text{groups}}}^{\mathsf{T}}\right]^{\mathsf{T}}}_{\text{stacked coordinates of groups } 1,\ldots,N_{\text{groups}}}, \\ \underbrace{\mathbf{u}_i}_{N_{\text{DOF},i}\text{ usable DOFs of group } i} &\in \underbrace{\text{Range}_i \subset \mathbb{R}^{N_{\text{DOF},i}}}_{\text{feasible motion set of group } i}, \\ \underbrace{N_{\text{DOF}}}_{\text{total controlled DOFs}} &= \sum_{i=1}^{\overbrace{N_{\text{groups}}}^{\text{number of active groups}}} N_{\text{DOF},i}. \end{aligned} \tag{1}$$

The vector $\mathbf{u}_i$ contains the $N_{\text{DOF},i}$ coordinate values of group i, while $N_{\text{DOF},i}$ specifies how many coordinates that group contributes. The set Range_i specifies the allowed combinations of these values under the group's motion limits. The group vectors are transposed into rows so that their entries can be joined end to end, and the outer transpose turns the combined row into a single column vector with N_{DOF} entries. Each group retains its own number of coordinates, so groups with different numbers of controlled DOFs can be combined without zero-padding.

Each $\mathbf{u}_i$ may contain values for axial spacing, lateral decenter, tip, tilt, or rotation. In the current implementation, each $\mathbf{u}_i$ contains five coordinate values, or DOFs, (the above list, without rotation) as offsets from the nominal design prescription. Selecting DOFs is system-dependent and part of the AAOS workflow. A coordinate is useful only when (i) it changes a task-relevant optical quantity, (ii) the sensor can distinguish that change from other coordinates, and (iii) the mechanism can move it repeatably without introducing vignetting, stress, or collision. Thus changing DOFs can change the correctable subspace, but can also make the inverse problem poorly constrained.

Let $\mathcal{U}_{\text{selected}}$ denote the chosen set of N_{DOF} actively controlled coordinates, whose values are collected in $\mathbf{u}$. Let $\mathbf{err}(\mathbf{u}) \in \mathbb{R}^{N_{\text{err}}}$ contain the differences between the measured and desired values of N_{err} optical quantities. The sensitivity matrix for these coordinates will be

$$\mathbf{J}_{\text{selected}} = \frac{\partial \mathbf{err}}{\partial \mathbf{u}} \in \mathbb{R}^{N_{\text{err}} \times N_{\text{DOF}}}, \tag{2}$$

evaluated at the current optical configuration. The subscript "selected" refers to the coordinate set $\mathcal{U}_{\text{selected}}$. Column j describes how all N_{err} optical error components change per unit motion of coordinate u_j. This matrix can be estimated by applying small known motions and measuring the resulting changes.

For a small coordinate adjustment $\Delta\mathbf{u}$, the optical error changes approximately by $\mathbf{J}_{\text{selected}}\Delta\mathbf{u}$. The columns of $\mathbf{J}_{\text{selected}}$ therefore describe the error changes that the selected motions can produce. Within this local linear approximation, and without motion constraints, the least-squares correction is $\Delta\mathbf{u}^* = -\mathbf{J}^+_{\text{selected}}\mathbf{err}$, where $\mathbf{J}^+_{\text{selected}}$ is the Moore-Penrose pseudoinverse. Substituting this adjustment gives the remaining error:

$$\begin{aligned} \underbrace{\mathbf{err}_{\text{res}}}_{\substack{\text{remaining}\\ \text{optical error}}} &\simeq \underbrace{\mathbf{err}}_{\substack{\text{measured}\\ \text{optical error}}} + \underbrace{\mathbf{J}_{\text{selected}}\Delta\mathbf{u}^*}_{\substack{\text{error change}\\ \text{from correction}}}, \\ &= \underbrace{\left(\mathbf{I} - \mathbf{J}_{\text{selected}}\mathbf{J}^+_{\text{selected}}\right)}_{\substack{\text{projector outside the}\\ \text{correctable subspace}}} \underbrace{\mathbf{err}}_{\substack{\text{measured}\\ \text{optical error}}}. \end{aligned} \tag{3}$$

Here $\mathbf{I}$ is the $N_{\text{err}} \times N_{\text{err}}$ identity matrix. The term $\mathbf{J}_{\text{selected}}\mathbf{J}^+_{\text{selected}}\mathbf{err}$ is the part of the measured error that the selected motions can cancel in this approximation. This calculation identifies the residual error that cannot be corrected through DOFs in the model. Ideally, this residual error will be negligible, and when it is not, tolerance analysis can guide design refinement to achieve a high performance AAOS.

B. Hardware realization and time scales

In this manuscript, we focus on controlling modeled AAOS devices, and the models are designed to be realistic and feasible to implement. In a physical AAOS device, an active lens group can be carried by an encoder-tracked motor stage, a piezo-flexure mount, or another precision motion mechanism. Position feedback from the manipulator or pose estimation from additional sensors is necessary but not sufficient, because precision positioning does not guarantee that the imaging performance is improved. After a move settles, the optical sensor must therefore measure the new state with imaging metrics. Hysteresis, creep, parasitic rotation, temperature-induced movement or sensing changes, gravity loading, and other influences can corrupt the calibration of an optomechanical setup. So, while precision pose estimation can inform optimization, an image-analysis assay should ultimately govern movement commands.

To decide how to move its lenses, an AAOS device needs some form of optical feedback to assess its current performance. This feedback can come from a variety of sensors, including a camera, a Shack-Hartmann wavefront sensor (SHWFS), an interferometer, a task image (an image recorded during the actual application), or a combination of these [6]. Whatever the sensor, it records light that has passed through every optical element between the moved lenses and the detector, that is, the full downstream prescription. This matters because moving one lens changes not only its own aberrations but also where

rays strike the lenses that follow it, so the elements interact. A measurement taken through the complete optical path captures these interactions directly, whereas encoder readings of individual lens positions cannot. The three test systems differ in how this measurement is organized. In the Trepan2p case, the forty coordinates are split between a front subsystem (including the coverglass and the first three lenses at the tip) and a rear subsystem (the remaining four lenses), each with its own predictor. While one subsystem is being observed, the other is held at its reference position, which its motor encoders define. This keeps the measurement taken through the whole system while leaving the unobserved subsystem in a known state. The two subsystem corrections are then combined into a single correction for the complete system, which is evaluated at the center and edges of the field. In the Nikon and KLA systems, the sensor simply observes the full optical system in its current state, with no subsystem splitting.

3. MULTI-FIELD SENSING AND THE UNIVERSAL PDM-NET

The deep learning system we built for Kaleid is called the Universal Physically-Disentangled Multi-Scale Wavefront Network (Universal PDM-Net) and it maps an N_{fields}-point measurement stack to grouped mechanical coordinates. It is universal in the sense that the same network architecture is effective across diverse optical systems. Each optical system has its own sensor channels, group heads, calibration, and fitted weights. Multiple observation positions help distinguish motions with similar on-axis effects but different responses at the edge.

The measurements describe the optical system at the coordinates $\mathbf{u}$ defined in Section 2.A. Trepan2p encodes pupil optical path difference (OPD) into camera intensity; Nikon and KLA supply differential or direct pupil-OPD measurements. All three simulated inputs include noise matched to the sensing modality. The following subsections describe these measurements and their common network interface.

A. Trepan2p diffuser-coded camera measurements

For the Trepan2p system, pupil phase is converted into camera intensity before being supplied to the network. A fixed phase diffuser produces intensity patterns that change with the optical configuration. For field node $k = 1, \ldots, N_{\text{fields}}$, with $N_{\text{fields}} = 9$, let $\text{OPD}_k(\boldsymbol{\rho}; \mathbf{u})$ be the pupil OPD at pupil position $\boldsymbol{\rho}$. The simulated frame at detector position $\mathbf{r}$ is

$$\begin{aligned} \Phi_k(\boldsymbol{\rho}; \mathbf{u}) &= \frac{2\pi}{\lambda}\text{OPD}_k(\boldsymbol{\rho}; \mathbf{u}) + \phi_{\text{diffuser}}(\boldsymbol{\rho}), \\ \text{Frame}_k(\mathbf{r}; \mathbf{u}) &= \text{Cam}\left\{ \left| \text{Prop}_{z_{\text{prop}}}\left[\text{Amp}_k(\boldsymbol{\rho}) e^{i\Phi_k(\boldsymbol{\rho};\mathbf{u})} \right] (\mathbf{r}) \right|^2 \right\}. \end{aligned} \tag{4}$$

Here Amp_k is the valid-pupil amplitude, ϕ_{diffuser} is the fixed diffuser phase, and $\text{Prop}_{z_{\text{prop}}}$ denotes angular-spectrum propagation over $z_{\text{prop}} = 2$ mm. The camera model Cam applies photon counting, pixel-response variation, read noise, dark signal, full-well clipping, and digitization to the propagated intensity. A 64×64 OPD pupil is centered in a 128×128 complex array. The fixed diffuser spans 0–4π phase, and the propagated intensity is recorded on a 128×128 sensor grid. The network receives raw digital frames on the training-data scale, without per-channel maximum normalization. Appendix A gives the detector settings.

The nine Trepan2p input positions in degrees are $(-2, 0)$, $(-1, 0)$, $(1, 0)$, $(2, 0)$, $(0, -1.8)$, $(0, -1)$, $(0, 0)$, $(0, 1)$, and $(0, 2.2)$. The separate terminal grid contains the center, four cardinal points at $2°$, and four diagonal points $(\pm 1.414°, \pm 1.414°)$; terminal WRMS is their mean.

B. Pupil-OPD measurements for Nikon and KLA

Using the same field notation, Nikon and KLA measure pupil OPD in waves, $\text{OPD}_k(\boldsymbol{\rho}; \mathbf{u})/\lambda$, equivalent to pupil phase divided by 2π. Nikon samples 49 pupil nodes at each of $N_{\text{fields}} = 3$ control fields, $0, \pm 0.05$ mm. At each field k, these values are differenced against the same field at the fixed operating reference, and piston, tip, tilt, and defocus are removed. Gaussian noise with standard deviation 0.015 waves is added to these samples before fixed linear interpolation onto three 32×32 maps. A separate noisy absolute center WRMS measurement selects the gain. Terminal absolute WRMS is evaluated without feedback noise at $0, \pm 0.1$ mm.

KLA samples $N_{\text{fields}} = 9$ field positions from -0.132 to 0.132 mm, producing 64×64 maps of $\text{OPD}_k(\boldsymbol{\rho}; \mathbf{u})/\lambda$. Its sensor model applies sub-aperture slope perturbations, multiplicative pixel-response variation, additive read noise, and quantization to the valid pupil samples. Gain selection and joint phase/MTF correction use three 48×48 pupils at the center and two field edges; final evaluation uses the nine-field grid. Appendix A gives the noise settings.

C. Universal PDM-Net

The Kaleid proposal network is trained separately on each optical system's subsystem data. Its shared image encoder reads the field stack, and its small output heads estimate the misalignment of individual lens groups. The input stacks the measured maps over $k = 1, \ldots, N_{\text{fields}}$: camera frames Frame_k for Trepan2p and the pupil-OPD maps described above for Nikon and KLA. The stack dimensions are $9 \times 128 \times 128$, $3 \times 32 \times 32$, and $9 \times 64 \times 64$, respectively.

A multi-scale stem produces 64 channels, followed by three depthwise-separable convolution stages with 128, 256, and 256 channels [19]. Global average pooling produces a shared 256-value description. Each group head uses a $256 \to 64$ linear layer, a Gaussian Error Linear Unit (GELU), and a $64 \to N_{\text{DOF},i}$ output layer to predict the normalized misalignment of group i. Here $N_{\text{DOF},i} = 5$ for the coordinates $(x, y, z, \theta_x, \theta_y)$, using the same group index i as in Eq. (1). Table 4 gives the layer sequence.

The predictions are converted back from the training-label normalization to physical coordinate units and reversed to give the proposed group corrections $\Delta\mathbf{u}_i$. Appendix A gives the units and training configuration.

The proposal networks contain 691k parameters across the Trepan2p subsystems, 928k across Nikon, and 969k across KLA subsystems. Each optical system remains below 1M parameters.

The local correction $-\mathbf{J}^{+}_{\text{selected}}\mathbf{err}$ in Section 2.A relates a measured optical error to small coordinate changes. Universal PDM-Net instead learns a mapping from the configured measurements to grouped motion proposals over the training range, including nonlinear intensity measurements. Kaleid then tests and refines these proposals on the coupled optical system.

4. MEASURED OPTIMIZATION AND VERIFICATION

The Universal PDM-Net estimates element positions based on metrology (e.g., diffuser images, OPD). However, real-world noise, interacting elements, and imaging conditions can make the optimal position different from the nominal prescription.

Kaleid therefore uses optical feedback in four stages: it records the configured field measurements and *proposes* grouped motions; *measures gain* by testing their amplitudes on the optical system; jointly *refines* all N_{DOF} coordinates to account for coupling through downstream ray paths; and *verifies* the final prescription at center and edge fields with the coordinates held fixed. The local sensitivity analysis in Section 2.A explains how coordinate changes affect optical error; the measured loop applies this principle as the coupled configuration changes.

A. Optical task objective

To specify the task, we collect optical errors at the selected field positions, sample depths or focus settings, wavelengths λ, and refractive indices into the error vector $\mathbf{err}(\mathbf{u})$ from Section 2.A. Each component represents one optical quantity at one selected operating condition. A general AAOS cost is

$$\begin{aligned}\mathrm{Cost}(\mathbf{u}) = & \underbrace{\sum_q \mathrm{weight}_q \left[\frac{\mathrm{err}_q(\mathbf{u})}{\mathrm{scale}_q}\right]^2}_{\substack{\text{dimensionless optical error: WFE, MTF, distortion,}\\ \text{telecentricity, chromatic/focus error, throughput, etc.}}} \\ & + \underbrace{\|\mathbf{u}-\mathbf{u}_{\mathrm{ref}}\|^2_{\mathbf{Weight}_{\mathrm{motion}}}}_{\substack{\text{motion away from the preferred}\\ \text{mechanically central prescription}}} \\ & + \underbrace{\mathrm{Penalty}(\mathbf{u})}_{\substack{\text{travel, clearance, stress, vignetting,}\\ \text{and ray-footprint constraints}}}.\end{aligned} \tag{5}$$

Here q indexes the error components, scale_q places unlike quantities on comparable scales, and weight_q sets scientific priority. The matrix $\mathbf{Weight}_{\mathrm{motion}}$ sets both the overall strength of the motion penalty about the preferred reference $\mathbf{u}_{\mathrm{ref}}$ and the relative coordinate weights according to their usable range or risk. The experiments use the WRMS and MTF-based scalar scores below.

B. From grouped proposals to a verified prescription

Starting from the current coordinates $\mathbf{u}_t$ at proposal-and-gain update t, Kaleid records the configured measurement stack for each scheduled predictor and obtains physical group corrections $\Delta\mathbf{u}_{i,t}$, including the configured proposal scaling (Appendix A). Unscheduled groups have $\Delta\mathbf{u}_{i,t} = \mathbf{0}$. Stacking all N_{groups} groups in the order of Eq. (1) forms the system trial direction:

$$\Delta\mathbf{u}_t^{\mathrm{prop}} = \left[\Delta\mathbf{u}_{1,t}^{\mathsf{T}}, \ldots, \Delta\mathbf{u}_{N_{\mathrm{groups}},t}^{\mathsf{T}}\right]^{\mathsf{T}} \in \mathbb{R}^{N_{\mathrm{DOF}}}. \tag{6}$$

The measured optical response selects the step amplitude:

$$\begin{aligned}\mathrm{gain}_t^* &\in \underset{\mathrm{gain}\in\mathrm{Gains}_t}{\arg\min}\ \mathrm{Cost}_{\mathrm{gain},t}(\mathbf{u}_t + \mathrm{gain}\Delta\mathbf{u}_t^{\mathrm{prop}}), \\ \mathbf{u}_{t+1} &= \mathbf{u}_t + \mathrm{gain}_t^*\Delta\mathbf{u}_t^{\mathrm{prop}}.\end{aligned} \tag{7}$$

Here Gains_t contains the tested gains, including zero motion. For the optical task described by Eq. (5), $\mathrm{Cost}_{\mathrm{gain},t}$ is the measured feedback cost used to select a step: center WRMS for Trepan2p and Nikon, or negative mean MTF over 500–4000 lp/mm for KLA (Table 2). The tests use discrete candidates or a bounded scalar search. Trial commands respect the configured motion bounds, and zero motion allows a proposal to be rejected on the measured score.

After the scheduled proposal-and-gain updates, joint correction adjusts all N_{DOF} coordinates together using the system's optical objectives. The final prescription $\mathbf{u}_{\mathrm{final}}$ is then held fixed for center-and-edge evaluation. Table 2 gives the implemented operations and terminal criteria.

Table 2. Implementation of Kaleid across the three optical systems. Grouped proposals and measured gains follow Eqs. (6)–(7); joint refinement and fixed-setting verification use the optical measurements and objectives shown below.

System	Proposal and gain measurement	Joint refinement	Fixed-setting verification
Trepan2p	Nine diffuser-coded intensity fields for the active front or rear subsystem; center WRMS of the complete optical system selects bounded subsystem gains	40-coordinate L-BFGS-B on center WRMS with the remaining budget; best measured system state retained	Mean PTTD-removed WRMS over nine terminal points; success is strictly below 0.07 waves
Nikon	Phase differences from the fixed operating reference at three fields; a separate noisy absolute center WRMS selects the gain	Damped calibrated-response increments with secant updates and measured local five-axis responses	Center absolute WRMS below 0.07 waves; absolute edge WRMS and ray coverage also recorded
KLA	Multi-field pupil phase from the current optical system; mean MTF over 500–4000 lp/mm selects the lens-wise gain	Calibrated three-field phase correction followed by broad-band and Score_{2000} contrast searches over all 40 coordinates	Nine-field Score_{2000}; success is at least 0.45

C. Optical implementation in each system

Trepan2p supplies measured score changes from small coordinate moves to L-BFGS-B. Nikon uses a calibrated 147×35 phase response, damped inverse increments, and secant updates; stalled steps can trigger a measured local five-axis response. KLA first uses a calibrated three-field phase response, then applies local Powell searches to improve contrast over 500–4000 lp/mm and subsequently Score_{2000}. The Nikon and KLA phase responses instantiate the sensitivity matrix $\mathbf{J}_{\mathrm{selected}}$ in Eq. (2) for their respective residuals and coordinate scales. The residual $\mathbf{err}_{\mathrm{res}}$ in Eq. (3) describes the error outside the correctable subspace of a local linearization; final performance is determined by the fixed-setting optical evaluation. Appendix A summarizes the response measurements and search sequence.

5. OPTICAL RECOVERY AND METHOD COMPARISON

Table 3 compares six procedures on eleven test cases per system. The requirement is recovery of every case in all three optical systems under their respective criteria. Each mean includes all eleven cases.

Standard SVD uses a truncated singular value decomposition of a fixed nominal interaction matrix, with relative singular-value cutoff 10^{-2}, and applies one bounded unit-gain command. Appendix A gives the Standard SVD preprocessing and calibration settings.

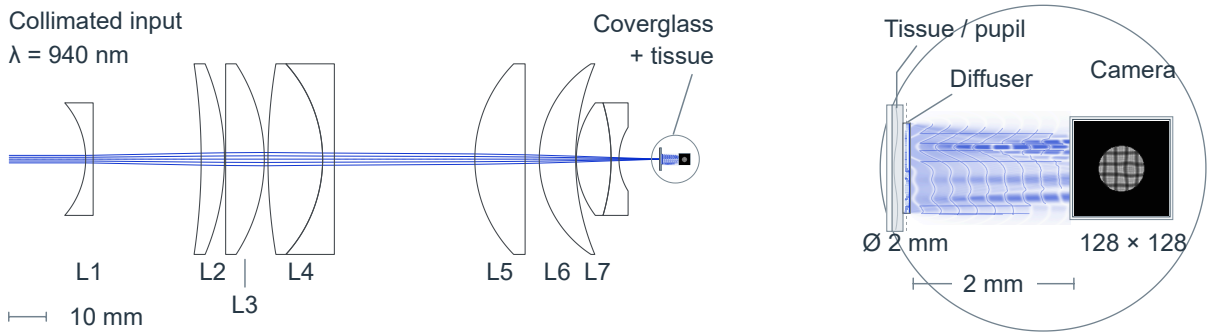


Fig. 2. Trepan2p optical layout and diffuser-camera readout schematic.

A. Trepan2p: two-photon imaging and coupled correction

The Trepan2p reconstruction uses the published 940-nm objective and scan engine, with 27.5-mm focal length, NA 0.43, a 3.5-mm imaging region, and a 2-mm object-space entrance pupil [25]. This dry objective focuses scanned excitation light into tissue, with the sample located after the final lens and coverslip. For calibration, the diffuser-camera readout will read directly or with 1:1 relay from the location normally occupied by the sample. Seven lens groups and a coverglass supply forty coordinates, partitioned into front and rear 20-coordinate predictors. The two predictions are combined before gain selection and correction of the complete optical system.

The eleven evaluation cases are selected from a deterministic seed stream separate from the training and validation sets using only initial optical metrics: initial nine-point total WRMS must exceed one wave and PTTD-removed high-order WRMS must exceed 0.3 waves. The mean initial total and high-order WRMS are 51.091 and 4.239 waves, respectively. Defocus contributes to the total value and is removed from the recovery metric to account for autofocus.

Kaleid passes 11/11 nine-point means within twenty acquisitions, at mean WRMS 0.034685 waves. It passes 98/99 individual positions and the center plus every edge in 10/11 cases. Prediction only passes all eleven means, while no-initializer search passes 3/11 at mean 0.103182 waves. Learned initialization therefore provides effective recovery within this budget. Standard SVD also passes 11/11 means at 0.036291 waves, but only 8/11 cases pass at every position, while Kaleid improves this count to 10/11.

The bounded shared-pupil DM passes 0/11 at mean WRMS 4.234778 waves. It fits one command across the center and eight edges, whereas motion of optical elements can adjust the internal ray paths that produce these aberrations.

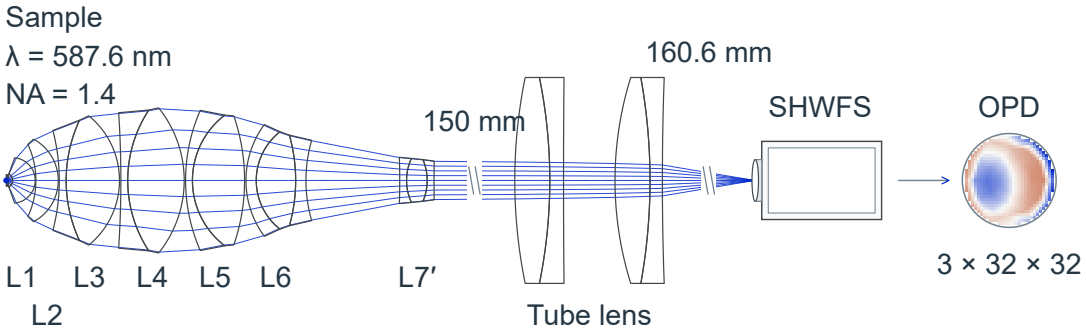


Fig. 3. Nikon optical layout and SHWFS readout schematic.

B. Nikon: high-NA center recovery and edge verification

The Nikon 100×, 1.4-NA objective follows embodiment 2 of U.S. Patent 6,519,092 B2 at 587.5618 nm [29]. This oil-immersion objective collects specimen light through the coverslip and immersion layer, with a tube lens forming the final image. A calibration reading at the specimen plane could provide a reference wavefront with an SHWFS placed at the relayed objective pupil. Three predictors (G1–G3) estimate 15/15/5 coordinates from the complete observation. Each of the eleven cases perturbs 2 to 4 of the seven lens groups, within ±20 μm and ±0.04° per coordinate. The controller can adjust all 35 coordinates. Kaleid passes 11/11 center at mean WRMS 0.034659 waves; no-initializer correction passes 5/11 and prediction only 0/11. Standard SVD passes 0/11 at mean WRMS 177.271330 waves.

Figure 6 reports absolute PTTD WRMS at the frozen center and both edges, ±0.1 mm. The 11/11 success count uses the center threshold; edge quality and ray coverage are recorded separately.

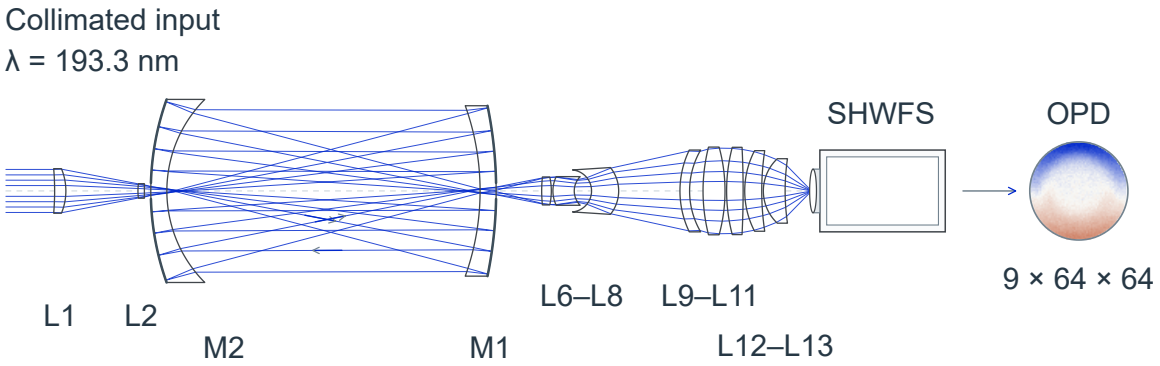


Fig. 4. KLA optical layout and SHWFS readout schematic.

C. KLA: frequency-specific contrast and task selection

The 193.3-nm, 0.8-NA catadioptric prescription is derived from KLA-Tencor-assigned U.S. Patent 6,842,298 B1 [30]. The KLA objective supports both bright-field and dark-field inspection; our alignment configuration follows the bright-field illumination direction, with collimated light entering from the instrument side and passing through the objective toward the sample plane. A sample-side wavefront-sensing head samples the delivered field at the location normally occupied by the sample. Three active subsystems supply 40 coordinates. 5 cases use limits of ± 80 μm and ± 2.5 arcmin; the other 6 use ± 40 μm and ± 1.2 arcmin. Instead of a microscopy WRMS criterion, the task rewards contrast transfer and its consistency over nine positions:

$$\overline{\mathrm{MTF}} = \frac{1}{9}\sum_{k=1}^{9} \mathrm{MTF}_k,$$
$$\mathrm{Score}_{2000} = \overline{\mathrm{MTF}} - 0.3\sqrt{\frac{1}{9}\sum_{k=1}^{9}(\mathrm{MTF}_k - \overline{\mathrm{MTF}})^2}. \quad (8)$$

Here MTF_k is tangential MTF near 2000 lp/mm at position k. A case passes at $\mathrm{Score}_{2000} \geq 0.45$. Kaleid reaches mean $\mathrm{Score}_{2000} = 0.624635$ and passes 11/11. Standard SVD reaches 0.437318 and passes 6/11, while the bounded DM reaches 0.233658 and passes 0/11. The Kaleid controller combines group-level gains with joint phase correction, followed by broad-band and target-frequency contrast optimization [24].

KLA illustrates task-specific prescription control: the final coordinates maximize the contrast target rather than reproduce the reference coordinates. Equation (5) extends this choice to wavelength, depth, refractive index and center-to-edge performance.

Task-specific validation. An additional KLA subsystem-2 validation asks whether task-specific control can exceed the nominal $\mathrm{Score}_{2000} = 0.5962$. The 11 training-excluded cases control L6–L8 (15 DOFs) with a 40-move cap; all finish above nominal, spanning 0.597878–0.617530 (median 0.601069; mean 0.602770).

Table 3. Method comparison. Each entry is passes/11 followed by mean terminal metric. Trepan2p uses nine-point mean WRMS, Nikon center WRMS (waves; lower is better), and KLA Score_{2000} (higher is better).

Procedure	Trepan2p	Nikon	KLA
Kaleid	11/11; 0.0347	11/11; 0.0347	11/11; 0.6246
Prediction only	11/11; 0.0387	0/11; 0.3568	0/11; 0.0097
No initializer	3/11; 0.1032	5/11; 3.5298	0/11; 0.0116
Standard SVD	11/11; 0.0363	0/11; 177.2713	6/11; 0.4373
Shared pupil DM	0/11; 4.2348	0/11; 25.8817	0/11; 0.2337
Single-pass SLOR	9/11; 0.0461	0/11; 0.1643	0/11; 0.0113

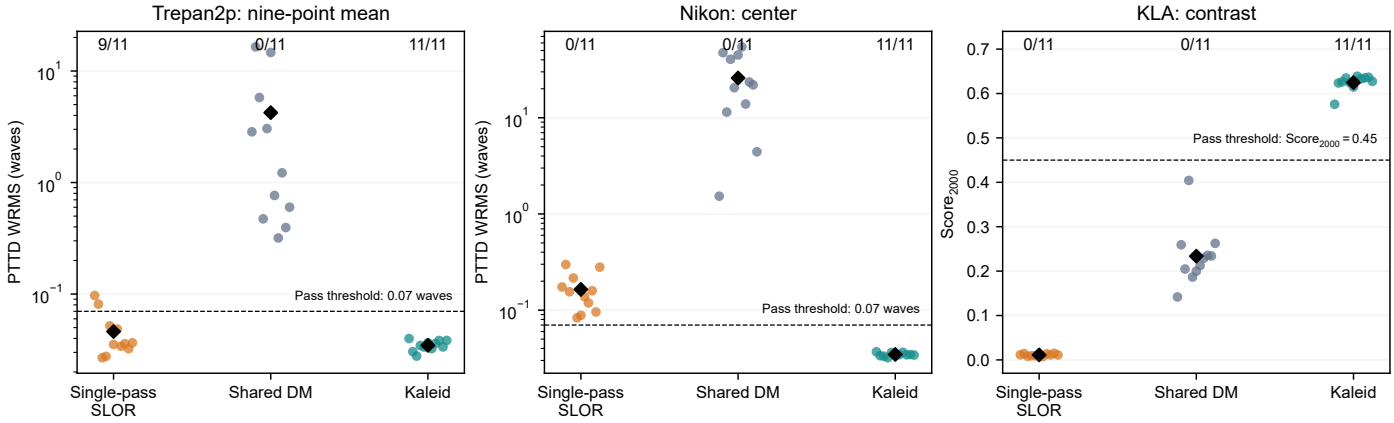


Fig. 5. Optical recovery on eleven cases. All three optical systems compare single-pass SLOR, a shared-pupil DM and Kaleid. Points are individual cases; diamonds are means. Dashed lines mark the success thresholds: WRMS $<$ 0.07 waves for Trepan2p and Nikon, and $\text{Score}_{2000} \geq 0.45$ for KLA.

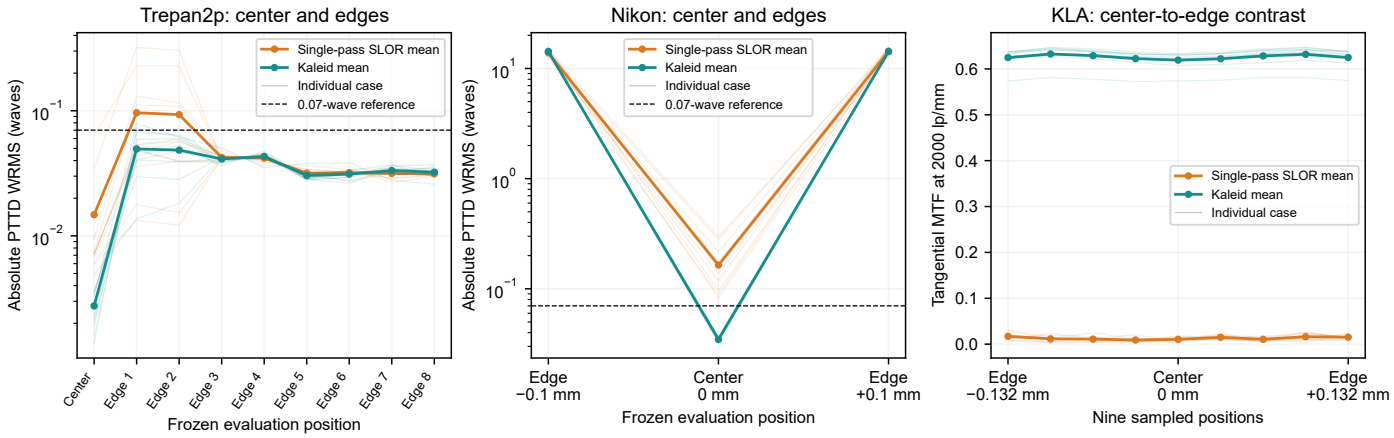


Fig. 6. Center and edge performance of single-pass SLOR and Kaleid at frozen terminal settings. Trepan2p uses the center and eight edges. Nikon uses the center and two edges at ± 0.1 mm. KLA samples nine positions from the negative edge through the center to the positive edge. Faint lines are individual cases; heavy lines with circles are eleven-case means. Dashed lines mark the 0.07-wave reference at each position in Trepan2p and Nikon.

D. From optical recovery to image formation

For a near-diffraction-limited terminal state, WFE can be translated into a familiar image-quality proxy with the extended Maréchal approximation [26],

$$\underbrace{\mathrm{Strehl}_{\mathrm{est}}}_{\text{estimated one-photon Strehl}} \approx \exp\left[-\left(2\pi \underbrace{\frac{\mathrm{WFE}_{\mathrm{rms}}}{\lambda}}_{\text{RMS WFE in waves}}\right)^2\right], \tag{9}$$

$$\underbrace{\mathrm{Excitation}_{2\mathrm{P,peak}}}_{\substack{\text{approximate normalized}\\ \text{two-photon peak-excitation proxy}}} \propto \underbrace{\mathrm{Strehl}_{\mathrm{est}}^2}_{\substack{\text{square-law response under}\\ \text{fixed delivered pulse energy}}}.$$

Two-photon fluorescence is quadratic in excitation intensity in the unsaturated regime [27, 28]. Figure 7 shows the full-prescription NAOMi/POPPY simulation at two sample depths [31, 32]. AAOS recovers sample structure and localized bright features blurred by the uncorrected and DM-relay responses.

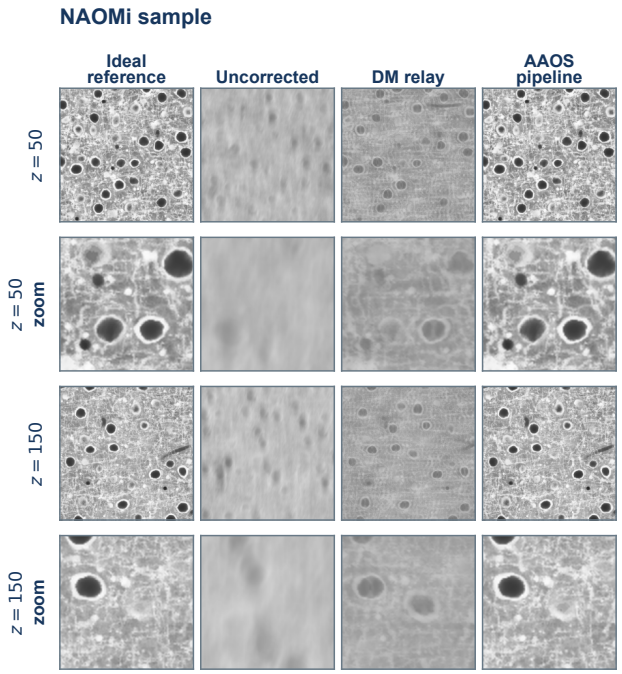


Fig. 7. Full-prescription NAOMi/POPPY images at two sample depths, with full views and zoomed regions. Columns compare ideal reference, uncorrected, 12 × 12 DM-relay and AAOS responses [31, 32].

6. CONCLUSION AND OUTLOOK

Kaleid combines compact learned motion proposals with measured correction of the optical system across diffuser-intensity, direct-OPD and differential-phase sensing, and across wavefront-error and contrast objectives.

Kaleid is the only procedure in Table 3 to pass all eleven cases in each of the three systems. Standard SVD passes 11/11 Trepan2p cases but only 0/11 Nikon and 6/11 KLA cases, so it does not meet the cross-system requirement. Kaleid also raises the Trepan2p center-and-all-edge pass count from 8/11 to 10/11. The no-initializer comparisons show the value of learned proposals within the available measurement budgets; combining those proposals with measured correction delivers recovery across all three optical systems.

A. Operating uses and complementary correction time scales

An adjustable prescription can compensate assembly drift and wear, or select a different operating task. Microscopy can change the target depth, wavelength or sample refractive index; inspection can prioritize contrast at selected spatial frequencies. Equation (5) combines these targets with mechanical constraints, and the NAOMi illustration shows their connection to image formation.

Standalone AAOS adjusts existing optical elements within the native optical path, with zero insertion loss because no extra optical elements or relay lenses are introduced. In a hybrid optical system, slow prescription correction removes static errors and drift, while faster AO handles sample dynamics and residual phase. Exposure, mechanical settling and hysteresis determine the physical update rate. Relatively slow 5-6 DOF mounts can optimize pose, while a deformable mirror or spatial light modulator remains available for faster sample- or atmosphere-induced residuals, such as in dual AO [15, 37].

B. Differentiable optics and measured adaptation

Differentiable ray and wave models could calculate prescription updates directly from an image-quality objective [33–36]. Combining refraction, propagation and image formation would provide gradients over the movable coordinates. Accurate gradients could reduce trial moves, while measured outcomes could calibrate the model to each optical system.

This approach requires efficient propagation and differentiation across many observation points, together with an accurate model of the built optical system. Surface and material errors, high-NA wave effects, sensor bias, temperature, sample refractive index and mechanical hysteresis can all change the measured response. Faster optical operators and system-specific calibration address these two requirements.

Within Kaleid, a differentiable model could propose a move, a learned residual model could compensate repeatable modeling errors, and measured gain selection could set the step. Logged observations, commands and scores would support adaptation during use, potentially reducing correction to one or a few steps with a final optical check.

C. A path toward dedicated optical-control hardware

An application-specific integrated circuit (ASIC) could combine sensor preprocessing, the sub-million-parameter proposal network and calibrated response calculations near the sensor and stage interfaces. A programmable prototype would establish the required precision, memory and data flow before a dedicated implementation.

A later optical-computation ASIC could accelerate ray–surface intersection, refraction, pupil propagation, image-quality evaluation and derivatives. Fixed geometry and calibration could remain in local memory, with unaffected upstream states reused after downstream moves. Event-driven or spiking implementations could exploit selective updates alongside con-

ventional digital pipelines. Hardware choices would depend on the changing optical states and required numerical precision; reduced computation time and energy would complement improvements in exposure and mechanical settling.

These directions form one development path: compact learned proposals and measured correction now, system-calibrated differentiable updates next, and dedicated computation as the repeated optical workload becomes well defined. AAOS supplies the physical freedom, while Kaleid connects that freedom to an optical task through measurement. More capable models and hardware can therefore extend the speed and range of adaptation while preserving the central operating principle: task-specific recovery by adjusting the internal prescription.

A. IMPLEMENTATION DETAILS

Network and training

Table 4. Universal PDM-Net layer sequence used by Kaleid.

Part	Operation	Output dim.
Input	N_{fields} measured two-dimensional maps	N_{fields}
Multi-scale stem	Parallel 3×3, 5×5, 7×7, and pool+1×1 paths	64
Stage 1	Two depthwise-separable convolution units; first has stride 2	128
Stage 2	Two depthwise-separable convolution units; first has stride 2	256
Stage 3	Two depthwise-separable convolution units; first has stride 2	256
Pooling	Global average pooling	256
Group head i	Linear $256\to64$, GELU, linear $64\to N_{\text{DOF},i}$	$N_{\text{DOF},i}$

Trepan2p and KLA labels use millimeters for translation and degrees/radians for rotation, respectively. Nikon uses normalized commands with scales of 20 μm and 0.04°.

Trepan2p and Nikon use standardized motion labels, Huber loss with threshold 0.25 [21], and AdamW [22] for 35 epochs. The learning rate decays from 2×10^{-4} to 2×10^{-6} with a cosine schedule; weight decay is 10^{-4}. Each Trepan2p subsystem uses 5,100 training and 1,300 validation states. Each Nikon predictor uses 4,000 training and 1,000 validation states, including simultaneous perturbations outside its predicted lens groups. KLA uses 5,000 states per subsystem.

SLOR uses mean-squared error and AdamW for 500 epochs, with batch size 250, learning rate 10^{-5} and weight decay 0.01. Model selection uses validation data for Trepan2p, Nikon and SLOR.

Trepan2p training includes translations within ±100 μm and tilts within ±2.5 arcmin. Its eleven test cases are generated separately from the training data and selected using initial optical quality: nine-point total WRMS above one wave, PTTD-removed WRMS above 0.3 waves, and at least 95% of nominal valid rays at every sampled position.

Measurement settings

Trepan2p camera measurements include Poisson shot noise at a 4000-electron peak exposure, 1.2-electron RMS read noise, 0.5% pixel-response nonuniformity, and dark-signal nonuniformity with a 2-electron mean and 0.5-electron standard deviation. The signal is clipped at 10,000 electrons and digitized to 16 bits at 4 ADU/electron with a 100-ADU offset. Terminal WRMS is averaged over the center, four cardinal field positions at 2°, and four diagonal positions $(\pm1.414°, \pm1.414°)$.

Nikon's control response uses 49 pupil samples at each of three field positions, $0, \pm0.05$ mm. Feedback noise has a standard deviation of 0.015 waves. Terminal absolute WRMS is evaluated without feedback noise at $0, \pm0.1$ mm.

KLA uses nine 64×64 pupil maps for network input and final evaluation. Gain selection and phase/MTF correction use three 48×48 pupils at the center and two field edges. On valid pupil samples, the noise model adds independent Gaussian slope coefficients with standard deviation 0.0075 waves per normalized pupil coordinate within an 8×8 sub-aperture partition, applies multiplicative pixel gains with mean one and standard deviation 0.01, and adds Gaussian read noise with standard deviation 0.015 waves. The resulting OPD range is quantized into 4096 intervals.

Correction settings

Trepan2p combines both subsystem proposals before selecting their gains on the complete optical system. Joint L-BFGS-B refinement uses the remaining portion of the twenty-measurement budget [23]. The returned coordinates have the lowest measured center WRMS; the nine-point terminal evaluation uses these fixed coordinates.

Nikon uses damping 10^{-4} and clips normalized commands to ±1.5. Correction stops when the measured center WRMS reaches 0.04 waves or below, the 80-reading limit is reached, or the solver terminates.

KLA scales its network proposals by 1, 1/2, and 1/4 in successive rounds before lens-wise gain selection. It then reduces the three-field phase difference from the operating reference using a calibrated response. A large remaining residual triggers a locally measured response. Joint Powell searches then optimize mean MTF over 500–4000 lp/mm, followed by Score_{2000} [24].

Standard SVD calibration

The interaction matrices are calibrated by signed coordinate perturbations about the nominal prescription. Trepan2p peak-normalizes each camera frame and averages nonoverlapping 8×8 blocks to obtain 16×16 samples. The nine frames form 2,304 features for each 20-coordinate subsystem. Calibration uses translations of ±15 μm and tilts of ±0.375 arcmin.

Nikon uses a 197×35 center-phase matrix with noiseless probes at ±0.001 normalized coordinate units. KLA projects three-pupil phase measurements onto orthonormal polynomials through degree six after removing piston and tilt, producing a 75×40 matrix. Its calibration probes are ±18 μm or ±0.525 arcmin.

The relative singular-value cutoff of 10^{-2} retains 20 modes per Trepan2p subsystem, eight Nikon modes and six KLA modes. Each inverse and its nominal reference remain fixed during testing.

DISCLOSURES

A provisional patent application has been filed covering the Active Architecture Optical Systems (AAOS) architecture described in this work.

DATA AVAILABILITY

Data are available upon request.

REFERENCES

1. Y. Zhang and H. Gross, "Systematic design of microscope objectives. Part I: System review and analysis," *Adv. Opt. Technol.*, vol. 8, no. 5, pp. 313–347, 2019, doi: 10.1515/aot-2019-0002.
2. M. J. Booth, "Adaptive optical microscopy: The ongoing quest for a perfect image," *Light Sci. Appl.*, vol. 3, Art. no. e165, 2014, doi: 10.1038/lsa.2014.46.
3. N. Ji, "Adaptive optical fluorescence microscopy," *Nat. Methods*, vol. 14, pp. 374–380, 2017, doi: 10.1038/nmeth.4218.
4. J. Mertz, H. Paudel, and T. G. Bifano, "Field of view advantage of conjugate adaptive optics in microscopy applications," *Appl. Opt.*, vol. 54, no. 11, pp. 3498–3506, 2015, doi: 10.1364/AO.54.003498.
5. J. Li, D. R. Beaulieu, H. Paudel, R. Barankov, T. G. Bifano, and J. Mertz, "Conjugate adaptive optics in widefield microscopy with an extended-source wavefront sensor," *Optica*, vol. 2, no. 8, pp. 682–688, 2015, doi: 10.1364/OPTICA.2.000682.
6. Q. Zhang, Q. Hu, C. Berlage, P. Kner, B. Judkewitz, M. Booth, and N. Ji, "Adaptive optics for optical microscopy [Invited]," *Biomed. Opt. Express*, vol. 14, no. 4, pp. 1732–1756, 2023, doi: 10.1364/BOE.479886.
7. N. Devaney, E. Dalimier, T. Farrell, D. Coburn, R. Mackey, D. Mackey, F. Laurent, E. Daly, and C. Dainty, "Correction of ocular and atmospheric wavefronts: A comparison of the performance of various deformable mirrors," *Appl. Opt.*, vol. 47, no. 35, pp. 6550–6562, 2008, doi: 10.1364/AO.47.006550.
8. L. Zhang, S. Zhou, J. Li, and B. Yu, "Deep neural network based calibration for freeform surface misalignments in general interferometer," *Opt. Express*, vol. 27, no. 23, pp. 33709–33723, 2019, doi: 10.1364/OE.27.033709.
9. M. Zhang, P. Jia, Z. Li, W. Xiang, J. Lv, and R. Sun, "Perception of misalignment states for sky survey telescopes with the digital twin and the deep neural networks," *Opt. Express*, vol. 31, no. 26, pp. 44054–44075, 2023, doi: 10.1364/OE.507254.
10. T. Slor, D. Oren, S. Baneth, T. Coen, and H. Suchowski, "Deep learning for optical misalignment diagnostics in multi-lens imaging systems," *Opt. Lett.*, vol. 51, no. 1, pp. 169–172, 2026, doi: 10.1364/OL.578126.
11. Q. Hu, M. Hailstone, J. Wang, M. Wincott, D. Stoychev, H. Atilgan, D. Gala, T. Chaiamarit, R. M. Parton, J. Antonello, A. M. Packer, I. Davis, and M. J. Booth, "Universal adaptive optics for microscopy through embedded neural network control," *Light Sci. Appl.*, vol. 12, Art. no. 270, 2023, doi: 10.1038/s41377-023-01297-x.
12. S. W. Paine and J. R. Fienup, "Machine learning for improved image-based wavefront sensing," *Opt. Lett.*, vol. 43, no. 6, pp. 1235–1238, 2018, doi: 10.1364/OL.43.001235.
13. H. Liu, W. Li, S. Gao, Q. Jiang, L. Sun, B. Zhang, L. Zhao, J. Zhang, and K. Wang, "Application of deep learning in active alignment leads to high-efficiency and accurate camera lens assembly," *Opt. Express*, vol. 32, no. 25, pp. 43834–43849, 2024, doi: 10.1364/OE.537241.
14. E. Hu, J. He, J. Zhou, Z. Ren, H. Sun, T. Jiang, H. Feng, and Y. Chen, "Fast and accurate active alignment of camera lenses with physics-informed deep learning," *Opt. Express*, vol. 33, no. 10, pp. 21256–21273, 2025, doi: 10.1364/OE.560123.
15. W. Zou, X. Qi, and S. A. Burns, "Wavefront-aberration sorting and correction for a dual-deformable-mirror adaptive-optics system," *Opt. Lett.*, vol. 33, no. 22, pp. 2602–2604, 2008, doi: 10.1364/OL.33.002602.
16. P. Berto, H. Rigneault, and M. Guillon, "Wavefront sensing with a thin diffuser," *Opt. Lett.*, vol. 42, no. 24, pp. 5117–5120, 2017, doi: 10.1364/OL.42.005117.
17. T. Wu, M. Guillon, G. Tessier, and P. Berto, "Multiplexed wavefront sensing with a thin diffuser," *Optica*, vol. 11, no. 2, pp. 297–304, 2024, doi: 10.1364/OPTICA.500780.
18. Y. Nishizaki, M. Valdivia, R. Horisaki, K. Kitaguchi, M. Saito, J. Tanida, and E. Vera, "Deep learning wavefront sensing," *Opt. Express*, vol. 27, no. 1, pp. 240–251, 2019, doi: 10.1364/OE.27.000240.
19. F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2017, pp. 1251–1258, doi: 10.1109/CVPR.2017.195.
20. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2016, pp. 770–778, doi: 10.1109/CVPR.2016.90.
21. P. J. Huber, "Robust estimation of a location parameter," *Ann. Math. Stat.*, vol. 35, no. 1, pp. 73–101, 1964, doi: 10.1214/aoms/1177703732.
22. I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, New Orleans, LA, USA, 2019.
23. C. Zhu, R. H. Byrd, P. Lu, and J. Nocedal, "Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization," *ACM Trans. Math. Softw.*, vol. 23, no. 4, pp. 550–560, 1997, doi: 10.1145/279232.279236.
24. M. J. D. Powell, "An efficient method for finding the minimum of a function of several variables without calculating derivatives," *Comput. J.*, vol. 7, no. 2, pp. 155–162, 1964, doi: 10.1093/comjnl/7.2.155.
25. J. N. Stirman, I. T. Smith, M. W. Kudenov, and S. L. Smith, "Wide field-of-view, multi-region, two-photon imaging of neuronal activity in the mammalian brain," *Nat. Biotechnol.*, vol. 34, pp. 857–862, 2016, doi: 10.1038/nbt.3594.
26. V. N. Mahajan, "Strehl ratio for primary aberrations in terms of their aberration variance," *J. Opt. Soc. Am.*, vol. 73, no. 6, pp. 860–861, 1983, doi: 10.1364/JOSA.73.000860.
27. W. Denk, J. H. Strickler, and W. W. Webb, "Two-photon laser scanning fluorescence microscopy," *Science*, vol. 248, no. 4951, pp. 73–76, 1990, doi: 10.1126/science.2321027.
28. P. Theer and W. Denk, "On the fundamental imaging-depth limit in two-photon microscopy," *J. Opt. Soc. Am. A*, vol. 23, no. 12, pp. 3139–3149, 2006, doi: 10.1364/JOSAA.23.003139.
29. K. Yamaguchi, "Immersion microscope objective lens," U.S. Patent 6 519 092 B2, Nikon Corp., Feb. 11, 2003.
30. D. R. Shafer and Y.-H. Chuang, "Broad band DUV/VUV catadioptric imaging system," U.S. Patent 6 842 298 B1, KLA-Tencor Corp., Jan. 11, 2005.
31. A. Song, J. L. Gauthier, J. W. Pillow, D. W. Tank, and A. S. Charles, "Neural anatomy and optical microscopy (NAOMi) simulation for evaluating calcium imaging methods," *J. Neurosci. Methods*, vol. 358, Art. no. 109173, 2021, doi: 10.1016/j.jneumeth.2021.109173.
32. M. D. Perrin, R. Soummer, E. M. Elliott, M. D. Lallo, and A. Sivaramakrishnan, "Simulating point spread functions for the James Webb Space Telescope with WebbPSF," in *Proc. SPIE Space Telescopes and Instrumentation 2012: Optical, Infrared, and Millimeter Wave*, vol. 8442, Art. no. 84423D, 2012, doi: 10.1117/12.925230.
33. C. Wang, N. Chen, and W. Heidrich, "dO: A differentiable engine for deep lens design of computational imaging systems," *IEEE Trans. Comput. Imaging*, vol. 8, pp. 905–916, 2022, doi: 10.1109/TCI.2022.3212837.
34. D. Deb, G.-J. Both, E. Bezzam, *et al.*, "Chromatix: A differentiable, GPU-accelerated wave-optics library," *Nat. Methods*, vol. 23, pp. 1388–1398, 2026, doi: 10.1038/s41592-026-03121-x.
35. Q. Sun, C. Wang, Q. Fu, X. Dun, and W. Heidrich, "End-to-end complex lens design with differentiable ray tracing," *ACM Trans. Graph.*, vol. 40, no. 4, Art. no. 71, pp. 1–13, 2021, doi: 10.1145/3450626.3459674.
36. L. Desdoigts, B. J. S. Pope, J. Dennis, and P. G. Tuthill, "Differentiable optics with ∂Lux: I—Deep calibration of flat field and phase retrieval with automatic differentiation," *J. Astron. Telesc. Instrum. Syst.*, vol. 9, no. 2, Art. no. 028007, 2023, doi: 10.1117/1.JATIS.9.2.028007.
37. C. Li, N. Sredar, K. M. Ivers, H. Queener, and J. Porter, "A correction algorithm to simultaneously control dual deformable mirrors in a woofer-tweeter adaptive optics system," *Opt. Express*, vol. 18, no. 16, pp. 16671–16684, 2010, doi: 10.1364/OE.18.016671.