

AI4AO: artificial intelligence for robust and autonomous adaptive optics in space applications

Benoit Neichel¹, Francisco Oyarzun¹, Matteo Pasinetti¹, Edvardas Vineraskas¹, Thierry Fusco³, Nicolas Vedrenne³, Julien Charton⁴, Roméo Roudeix⁴, Axel Vincent⁵, Perrine Lognoné², Jonathan Dray¹, Morgan Gray¹, Pauline Trouvé³, Frederic Champagnat³, Baptiste Sinquin⁴, Jean-François Sauvage³, Sylvain Cetre⁶

¹Aix Marseille Univ., CNRS, CNES, LAM, Marseille, France; ²Durham University, Durham, UK; ³ONERA, DOTA, Salon-de-Provence, France; ⁴Bertin-ALPAO, Grenoble, France; ⁵ONERA, Toulouse, France; ⁶Wakea Consulting, France

ABSTRACT

Artificial intelligence is creating new opportunities for adaptive optics (AO) systems operating in rapidly changing and poorly calibrated environments. This is particularly relevant to ground-based space surveillance, satellite imaging, and free-space optical communication, where low-elevation propagation, fast line-of-sight motion, scintillation, variable target structure, and stringent latency requirements can invalidate conventional AO assumptions. AI4AO develops an integrated framework in which wavefront sensing, reconstruction, control, calibration, and real-time implementation are optimized jointly. The program combines deep-learning reconstruction from detector images, model-based and data-driven predictive control, self-calibration, and wavefront-sensor concepts designed for strong turbulence and scintillation. Algorithms are developed through end-to-end simulation and progressively validated on experimental platforms. PAPYRUS provided an on-sky environment for reinforcement-learning and frequency-domain data-driven control. RAMA, integrated with ONERA's 60 cm FEELINGS telescope, extends this approach to a visible non-modulated pyramid sensor, high-rate cameras, two pupil-conjugated deformable mirrors, and a modular real-time computer. This paper presents the scientific rationale, system architecture, representative algorithmic developments, validation strategy, and first integration lessons. Together, these developments define a path toward AO systems that identify their dynamics, adapt to changing propagation conditions, and maintain performance with reduced manual calibration.

Keywords: adaptive optics, artificial intelligence, space situational awareness, free-space optical communication, wavefront sensing, reinforcement learning, predictive control, scintillation

1. INTRODUCTION

Ground-based optical systems are increasingly required to observe, track, and communicate with objects in near-Earth space. High-resolution imaging supports space-domain awareness and the characterization of satellites, debris, and natural bodies, while optical feeder links offer a route to high-capacity communications. In both cases, atmospheric turbulence perturbs the phase and amplitude of the optical field. The resulting image blur, beam wander, speckle, scintillation, and loss of coupling efficiency can dominate the end-to-end performance of an otherwise diffraction-limited terminal. [1,2]

Adaptive optics corrects a measured wavefront in real time by combining a wavefront sensor (WFS), a deformable mirror (DM), and a real-time controller (RTC). Conventional AO has reached a high level of maturity in astronomy, but space-related observing geometries expose several limitations simultaneously. Targets may have high angular velocities and structured, time-dependent intensity distributions; observations at low elevation increase the integrated turbulence; the signal-to-noise ratio can change rapidly; and amplitude fluctuations violate the phase-only assumptions used in many reconstructors. A controller optimized for stationary statistics and fixed calibration can therefore lose sensitivity or stability precisely when the observation is most valuable. [1,2]

AI4AO addresses these limitations by treating AO as an adaptive, data-rich dynamical system rather than a sequence of independently designed subsystems. The objective is not to replace every classical method with a neural network. Instead, artificial intelligence is introduced where it can provide measurable value: learning nonlinear measurement models, predicting disturbances, identifying changing dynamics, combining heterogeneous measurements, and automating calibration. Physics-based models, conventional estimators, and safety constraints remain central to the architecture. [9]

The project is motivated by two complementary application families. For space surveillance and tracking, the principal metric may be image quality, detectability, astrometric precision, or the stability of an extended-object image. For optical communication, the relevant output may be received power, fibre-coupling efficiency, fade probability, or link availability. Both require high-rate correction, robustness to changing conditions, and traceable behavior. They also provide different supervisory signals that can be exploited for self-optimization.

This paper gives a program-level view of AI4AO and reports progress toward telescope validation. Section 2 describes the environmental and operational drivers. Section 3 presents the common architecture. Sections 4 and 5 discuss AI-enhanced sensing and control. Section 6 addresses real-time implementation and validation, including PAPYRUS and RAMA. Section 7 summarizes current evidence and the remaining maturation steps toward operational systems.




Figure 1. AI4AO programme architecture. Space-surveillance and optical-link requirements share a dynamic atmospheric channel and drive three coupled developments: AI-enhanced sensing and reconstruction, predictive control, and online calibration with deterministic real-time computing.

2. SPACE-RELEVANT OPERATING CONDITIONS

2.1 Non-stationary turbulence and propagation geometry

For an astronomical guide star observed near zenith, turbulence is often approximated as a locally stationary random process over the duration used to estimate a reconstructor or tune a controller. Satellite passes are less forgiving. Elevation, apparent angular velocity, illumination, and the sampled atmospheric volume evolve continuously. The effective wind profile and turbulence strength seen by the control loop may change within a pass, and telescope tracking can introduce additional vibrations or pointing-dependent dynamics. Robust operation therefore requires either frequent retuning or a controller that identifies and follows these changes online.

The servo-lag error is especially important at high correction order. Even when sensing and reconstruction are accurate, the command applied to the DM corresponds to a wavefront measured several fractions of a coherence time earlier. Frozen-flow turbulence creates exploitable spatio-temporal correlations, while structural vibrations produce narrow-band components. Predictive control can use both, provided that the model is updated as the line of sight and wind change.

2.2 Scintillation, low flux, and extended objects

Long slant paths produce amplitude as well as phase fluctuations. In a pupil-plane sensor, scintillation can be confused with phase-induced intensity redistribution and may bias conventional slope estimates. Flux can also vary across the pupil, between frames, and across an extended target. For pyramid sensors, the response depends on the focal-plane

illumination and residual wavefront, so changing target morphology and optical gain become coupled. The resulting inverse problem is nonlinear and condition-dependent.

The AI4AO sensing program therefore evaluates reconstruction directly from detector data, with explicit exposure to photon noise, detector noise, amplitude variations, calibration errors, and target diversity. Training data are generated from propagation models and experimental telemetry, while validation is performed on conditions and targets not used for training. This separation is essential: a small average error on familiar simulations is not evidence of robustness during an operational pass.

2.3 Operational constraints

An operational AO system must be more than an accurate offline estimator. It must produce commands at deterministic cadence, recognize when its inputs depart from the training domain, respect DM stroke and slew limits, and transition safely between acquisition, learning, optimized control, and fallback states. These requirements shape the project from the start. Latency, jitter, memory transfer, model uncertainty, and observability are treated as first-order design parameters rather than implementation details added after algorithm development.

3. AI4AO SYSTEM ARCHITECTURE

The AI4AO architecture is organized around a layered control concept. The fastest path contains the WFS camera, calibrated preprocessing, reconstruction, command computation, and DM interface. A supervisory layer estimates operating conditions and performance, updates models or gains at a slower rate, and manages safety. Science- or link-quality measurements provide additional information when available, but the AO loop remains capable of stable operation without them. This separation allows experimental algorithms to be evaluated without compromising the deterministic core of the system.

Table 1. Functional decomposition of the AI4AO architecture. The rates are deliberately separated so that learning and model updates do not compromise the deterministic command path.

Function	Inputs	AI4AO contribution	Update rate
Sensing	Raw WFS frames, detector state	Nonlinear reconstruction; phase/amplitude estimation; uncertainty	Frame rate
Control	Reconstructed state, commands, history	Prediction; residual control; vibration rejection	Frame rate
Calibration	Telemetry, probes, focal-plane metrics	Online identification; optical-gain and registration tracking	Slow/episodic
Supervision	Quality metrics, constraints, health flags	Model selection; safe exploration; fallback management	Slow/episodic

A common telemetry interface is used across simulation and experimental platforms. Timestamped WFS frames, reconstructed modes, DM commands, loop state, focal-plane images, and environmental metadata are recorded with sufficient precision to reconstruct the causal sequence of a control cycle. This enables reproducible replay, direct comparison of algorithms on identical sequences, and estimation of the incremental latency introduced by each component.

The development workflow follows a progressive validation ladder: numerical simulation; hardware-in-the-loop tests; internal-source closed loop; on-sky observations on bright point sources; then more demanding objects and propagation conditions. At every stage, learned methods are compared with transparent baselines such as an integrator, optimized modal gain control, linear predictive control, or minimum-variance reconstruction. The goal is to quantify where adaptation is useful, not merely to demonstrate that a model can close a loop.

4. AI-ENHANCED WAVEFRONT SENSING

4.1 Learning nonlinear reconstruction

Classical reconstruction commonly linearizes the WFS response around a reference and applies a matrix calibrated with small DM perturbations. This approach is efficient and interpretable, but its accuracy decreases when residual aberrations become large, the sensor response saturates, optical gain changes, or amplitude variations enter the

measurement. A learned reconstructor can approximate the nonlinear inverse mapping from an image or intensity vector to a modal phase estimate. Convolutional architectures are natural for pupil images because they exploit local structure and share parameters spatially. [4,6]

In AI4AO, the learned estimate is not treated as an unconstrained image-to-command shortcut. The preferred formulation produces a physical state, modal coefficients, phase samples, or a correction to a linear estimate, together with confidence information. Residual learning is particularly attractive: a conventional reconstructor provides a stable first estimate, and the network learns only the condition-dependent bias or nonlinear remainder. This reduces the dynamic range required from the model and provides a defined fallback if its confidence is low.

Training sets must span atmospheric phase, scintillation strength, guide-object morphology, flux, detector noise, pupil shifts, WFS/DM misregistration, and non-common-path aberrations. Domain randomization is used to prevent the network from learning a single idealized instrument. Experimental calibration data then anchor the simulator to measured optical behaviour. Evaluation includes closed-loop residual error and task metrics, because a low open-loop phase error does not necessarily yield a stable or optimal closed loop.

4.2 Co-design of sensors and reconstruction

AI4AO also investigates wavefront sensors whose optical encoding is chosen jointly with the reconstruction algorithm. In a conventional design, the optical element is optimized for a linear analytic estimator. A learned inverse model relaxes this constraint and permits encodings that preserve information over a larger dynamic range or separate phase and amplitude more effectively. Candidate concepts are assessed using information content, photon efficiency, chromaticity, manufacturability, calibration burden, and robustness. [4,6]

This co-design is especially relevant under scintillation. If two aberration states create similar detector intensities in the chosen configuration, no reconstruction algorithm can separate them reliably. The optical design must therefore retain observability before learning is applied. Physics-informed training, multiple diversity states, or multi-sensor fusion can then exploit this information. The project studies combinations of WFS images, focal-plane images, and auxiliary flux measurements while controlling the added latency and hardware complexity.

4.3 The non-modulated pyramid as a demanding test case

RAMA uses a visible non-modulated pyramid wavefront sensor (PWFS). Removing mechanical modulation simplifies the optical chain and retains high small-signal sensitivity, but increases nonlinearity and sensitivity to residual aberration, centring, and imperfections of the pyramid apex. The physical pyramid inherited from PAPYRUS has a finite rooftop region. Numerical analysis for the RAMA configuration showed that stable operation strongly depends on the ratio between this feature and the focal-plane point-spread function, motivating an F/125 beam at the apex. This illustrates why optical design and reconstruction cannot be separated. [3–5,12]

The non-modulated PWFS also provides an ideal benchmark for hybrid methods. Linear reconstruction can be used near the reference point, while a neural correction addresses gain variations and nonlinear coupling at larger residuals. A focal-plane science camera can supply an independent estimate of optical gain or image quality. The relative performance of these strategies will be evaluated under controlled amplitude perturbations and then on sky. [5,6]



Figure 2. RAMA on the FEELINGS 60 cm telescope.

5. PREDICTIVE AND SELF-OPTIMIZING CONTROL

5.1 From integral control to prediction

The integrator remains a valuable baseline because it is computationally light, robust when carefully tuned, and well understood. Its principal weakness is that it reacts to an error after measurement and computation delays. Modal gain optimization improves the trade-off between rejection and noise propagation but does not fully exploit temporal correlations. Predictive controllers instead estimate the future disturbance at the time when a command will be effective. Frozen-flow layers, vibrations, and repeated tracking disturbances all provide predictable structure.

AI4AO compares several families on common telemetry: empirical orthogonal-function predictors, linear state-space methods, frequency-domain data-driven control, and reinforcement learning. The comparison separates the benefit of prediction from other effects such as different regularization, command filtering, or effective delay. Performance is reported versus flux, wind, turbulence strength, delay, calibration error, and model age, using multiple random realizations where simulation is involved.

5.2 Model-based reinforcement learning

Model-based reinforcement learning (MBRL) learns a compact dynamics model from closed-loop telemetry and optimizes a policy using that model. For AO, this is appealing because useful training data are produced continuously during operation and because the controller can adapt to turbulence, vibrations, and misregistration. The Policy Optimization for Adaptive Optics (PO4AO) approach has demonstrated predictive and self-calibrating behaviour in simulation and on laboratory benches. Its implementation can train in parallel with inference, allowing the active policy to remain fixed while a candidate update is prepared and validated. [7,8]

Within the PAPYRUS programme, PO4AO was interfaced to the DAO RTC through shared-memory telemetry and tested on sky at the Observatoire de Haute-Provence. The published 2026 results report stable operation across different targets, flux levels, and atmospheric conditions, with performance above the reference integrator in the tested configurations and successful rejection of vibration patterns. They also expose an engineering lesson: a research-oriented Python implementation added roughly 0.75 ms of latency and introduced jitter and occasional frame loss. This motivates the optimized inference and scheduling work described in Section 6. [10]

5.3 Frequency-domain data-driven control

Reinforcement learning is not the only path to adaptation. DD4AO identifies disturbance spectra from telemetry and formulates controller optimization with explicit stability and command constraints. On-sky PAPYRUS experiments reported in 2026 demonstrated stable hour-scale operation, online adaptation, and direct switching between DD4AO, an integrator, and optimized modal gain control. On the bright star Arcturus, the reported Strehl ratio at 1310 nm increased from 28.8% with the integrator to 33.9% with DD4AO. On a fainter target, the performance difference was smaller, but the adaptive controller used less DM stroke and reduced vibration peaks. [11]

These results are important for AI4AO because they provide a complementary, interpretable adaptive-control baseline. A learned controller should be compared not only with a fixed integrator but also with modern data-driven linear control. Conversely, neural dynamics models may help when the system is nonlinear or when operating conditions change faster than a spectral estimate can converge. Hybrid strategies can use identified linear dynamics for guaranteed stability and a bounded learned residual for additional performance.

5.4 Safe adaptation and evaluation

Online learning creates specific risks: an exploratory command can degrade the image or link, a model update can be trained on corrupted telemetry, and an apparently small prediction error can destabilize the feedback loop. AI4AO therefore uses guarded deployment. Candidate policies are evaluated offline or in shadow mode; command amplitude and slew are constrained; health monitors detect saturation, flux loss, abnormal residual growth, and timing faults; and the supervisor can revert immediately to a verified controller.

The evaluation protocol includes correction quality, stability margin, convergence time, sensitivity to hyperparameters, computational latency, jitter, DM effort, and recovery after a distribution shift. For surveillance, image sharpness and target-dependent observables are added. For optical links, power or fibre-coupling statistics and fade duration are more relevant than a mean Strehl ratio alone. Reporting these task-level metrics avoids optimizing a surrogate that does not improve the operational objective.

6. REAL-TIME IMPLEMENTATION AND EXPERIMENTAL VALIDATION

6.1 Deterministic real-time pipeline

The DAO RTC provides the integration backbone for the experimental platforms. High-rate camera streams, DM commands, and algorithm processes exchange data through low-overhead buffers. This modular design enables rapid comparison of controllers while retaining a deterministic acquisition and actuation path. In the present RAMA optical configuration, the target loop rate is up to 2 kHz, with the practical ceiling set by camera readout and region-of-interest choices. [13]

For learned components, average execution time is insufficient: worst-case latency and jitter must fit within the frame budget. The implementation strategy therefore separates data collection and training from real-time inference. Models are exported to optimized runtimes, memory transfers are minimized, and fixed tensor dimensions are used in the fast path. GPU implementations are appropriate during development and for convolutional workloads; FPGA or other deterministic accelerators are considered when power, latency, or deployment constraints justify the additional engineering effort.

Timing is measured end to end, from detector exposure and transfer to the instant at which the DM command becomes effective. Timestamp consistency is critical when training predictors: an unnoticed one-frame offset can make a sophisticated model learn the wrong temporal relationship. Replay tests, synthetic delay injection, and telemetry integrity checks are consequently part of the verification plan.

6.2 PAPYRUS as an on-sky algorithm pathfinder

PAPYRUS was developed as an open AO platform on the 1.52 m telescope at the Observatoire de Haute-Provence. It enabled external algorithms to be connected to an operational RTC and evaluated against standard controllers on the same night. Beyond the individual PO4AO and DD4AO demonstrations, PAPYRUS established procedures for safe controller switching, telemetry capture, repeatable calibration, and collaboration between optics, control, and machine-learning teams. These procedures are directly inherited by RAMA. [10,11]

The PAPYRUS results also show the need to distinguish algorithmic gain from implementation artefacts. Additional latency can erase part of the predictive advantage, while asynchronous processes can create control jitter. A fair experiment must therefore report both the residual performance and the actual timing of each controller. This principle is embedded in the RAMA test plan.

6.3 RAMA on the FEELINGS telescope

RAMA is installed on ONERA's 60 cm FEELINGS telescope and is designed as an agile pathfinder for space-oriented AO. The current hardware baseline contains two pupil-conjugated ALPAO mirrors (DM97 and DM192, with the latter initially used as a flat relay mirror), an Emergent IMX426LLJ camera for the PWFS, an Emergent IMX421LLJ science and gain-scheduling camera, and the DAO RTC. The common path uses off-axis parabolic relays to conjugate the telescope pupil to the mirrors. A long-focus branch produces the required F/125 beam at the pyramid apex, while the science branch samples the focal-plane image. [12,13]

The bench was designed from available components to shorten integration time and remain modifiable. An internal calibration source supports alignment, interaction-matrix acquisition, and repeatable closed-loop tests independently of telescope availability. The two-camera architecture enables simultaneous WFS and image-quality measurements. It also makes RAMA suitable for experiments in which a focal-plane metric supervises optical-gain estimation or controller adaptation.

Laboratory work has established internal-source closed-loop operation and characterized the sensitivity of the non-modulated pyramid configuration to apex geometry and alignment. Telescope integration has established optical coupling to FEELINGS and provides the basis for staged on-sky commissioning. The immediate program is to consolidate point-source control, quantify chromatic and scintillation effects, and then compare classical, hybrid neural, and predictive controllers under matched conditions. [12]




Figure 3. Example of the RAMA GUI with the non-modulated signal and the corresponding corrected PSF.

7. DISCUSSION AND ROADMAP

The AI4AO program results support three conclusions. First, prediction and online identification can improve AO performance in realistic conditions, as demonstrated by the recent PAPYRUS on-sky experiments. Second, implementation details are part of the control law: latency and jitter must be included in training and evaluation. Third, difficult sensors such as a non-modulated PWFS cannot be treated as a purely algorithmic problem; optical encoding, calibration, reconstruction, and control must be developed together. [10,11]

Several challenges remain before autonomous AI-enhanced AO can be deployed operationally. Learned reconstructors must demonstrate robustness to conditions absent from training and provide reliable uncertainty indicators. Online control must maintain explicit safety margins during rapid transitions. Scintillation and extended-object morphology must be represented faithfully in both simulation and experiments. Finally, performance should be linked to mission metrics rather than judged only through residual phase variance.

The next RAMA milestones are staged to reduce these risks. The first phase consolidates the bench calibration and reference integrator on bright stars. The second introduces controlled comparisons of linear and nonlinear reconstruction, including optical-gain tracking. The third evaluates predictive controllers under measured delays and vibrations. The fourth targets more representative objects and stronger amplitude fluctuations. Lessons from these campaigns will inform larger future systems, including the 2.5 m PROVIDENCE telescope and optical ground-station concepts.

A longer-term objective is a self-configuring AO system that can identify its interaction matrix and temporal dynamics, select an appropriate reconstruction/control model, monitor its own confidence, and retune as conditions evolve. Such autonomy is not equivalent to unrestricted learning. It is a structured combination of physical modelling, bounded adaptation, validated baselines, and explicit supervisory logic. This is the architecture most likely to transfer from research benches to dependable space operations.

8. CONCLUSION

AI4AO develops adaptive optics systems that learn where conventional calibration and stationary control assumptions become limiting. The project combines nonlinear and hybrid wavefront reconstruction, co-designed sensing, predictive and data-driven control, self-calibration, and real-time implementation. Its validation strategy links end-to-end simulation to laboratory and telescope experiments, using open and modular platforms.

PAPYRUS has already provided on-sky evidence that adaptive controllers can outperform a standard integrator and reject real disturbances, while also revealing the importance of optimized latency. RAMA transfers this methodology to a space-oriented 60 cm telescope and a challenging visible non-modulated pyramid sensor. The bench provides a direct route to studying scintillation, target structure, nonlinear response, predictive control, and operational safeguards under real atmospheric conditions. Together, these elements establish a scalable path toward robust and increasingly autonomous AO for space surveillance, tracking, and optical communications. [10–12]

ACKNOWLEDGEMENTS

The AI4AO project is funded by ANR-DGA-AID ASTRID program (ANR-25-ASTR-0015)

REFERENCES

- [1] Roddier, F., ed., *Adaptive Optics in Astronomy*, Cambridge University Press, Cambridge (1999).
- [2] Tyson, R. K., *Principles of Adaptive Optics*, 4th ed., CRC Press, Boca Raton (2015).
- [3] Ragazzoni, R., “Pupil plane wavefront sensing with an oscillating prism,” *Journal of Modern Optics* 43(2), 289–293 (1996). <https://doi.org/10.1080/09500349608232742>
- [4] Fauvarque, O., Neichel, B., Fusco, T., Sauvage, J.-F. and Girault, G., “General formalism for Fourier-based wave front sensing,” *Optica* 3(12), 1440–1452 (2016). <https://doi.org/10.1364/OPTICA.3.001440>
- [5] Chambouleyron, V., Fauvarque, O., Janin-Potiron, P., Neichel, B., Fusco, T., Sauvage, J.-F. and Levraud, N., “Pyramid wavefront sensor optical gains compensation using a convolutional model,” *Astronomy & Astrophysics* 644, A6 (2020). <https://doi.org/10.1051/0004-6361/202038034>
- [6] Landman, R. and Haffert, S. Y., “Nonlinear wavefront reconstruction with convolutional neural networks for Fourier-based wavefront sensors,” *Optics Express* 28(11), 16644–16657 (2020). <https://doi.org/10.1364/OE.391789>
- [7] Nousiainen, J., Rajani, C., Kasper, M. and Helin, T., “Adaptive optics control using model-based reinforcement learning,” *Optics Express* 29(10), 15327–15344 (2021). <https://doi.org/10.1364/OE.414200>
- [8] Nousiainen, J., Engler, B., Kasper, M., Rajani, C., Helin, T., Heritier, C. T., Quanz, S. P. and Glauser, A. M., “Laboratory experiments of model-based reinforcement learning for adaptive optics control,” *Journal of Astronomical Telescopes, Instruments, and Systems* 10(1), 019001 (2024). <https://doi.org/10.1117/1.JATIS.10.1.019001>
- [9] Fowler, J. and Landman, R., “Machine learning for adaptive optics: a review,” *Journal of Astronomical Telescopes, Instruments, and Systems* 9(1), 010901 (2023).
- [10] Nousiainen, J., Chambouleyron, V., Neichel, B., Cetre, S., Sauvage, J.-F., Alagao, A., Kasper, M., Dray, J., Fetick, R. and Engler, B., “On-sky demonstration of reinforcement learning for adaptive optics control,” *arXiv:2606.10771* (2026). <https://doi.org/10.48550/arXiv.2606.10771>
- [11] Dinis, I., Schinde, M., Blind, N., Chazelas, B., Restori, N., Lovis, C., Cetre, S., Alagao, A., Fetick, R., Héritier, T., Chambouleyron, V., Neichel, B., Fusco, T. and Sauvage, J.-F., “DD4AO control law for RISTRETTO: robustness, real-time performance, and on-sky validation with PAPYRUS,” *arXiv:2607.14925* (2026). <https://doi.org/10.48550/arXiv.2607.14925>
- [12] Pasinetti, M. et al., “RAMA: a new agile AO bench for the telescope FEELINGS,” *Proc. SPIE Adaptive Optics Systems X* (2026), publication details to be updated.
- [13] Barr, D. et al., “Durham Adaptive Optics real-time controller: modular high-performance control for adaptive optics,” *Proc. SPIE* (2024/2025), bibliographic details to be confirmed.