



Vera C. Rubin Observatory
Rubin Observatory Document

Image Quality Improvement Plan

NSF-DOE Vera C. Rubin Observatory

SOTN-004

Latest Revision: 2026-06-25



Abstract

As of June 2026, the delivered image quality from the Vera C. Rubin Observatory is not yet meeting design specifications for the total system contribution to the PSF FWHM, ellipticity, and uniformity of the PSF size and shape across the field of view. A technical pathway exists to meet all of the SRD design specifications with continued work on the Active Optics System (AOS) and environment control. We present an image quality improvement plan based analyses and experience from the Early Operations system optimization period, including strategy, priorities, and milestones for image quality improvements. We summarize currently known issues, and describe the offline analyses, on-sky engineering needs, daytime engineering needs, and expected timeline to resolve each issue.

Change Record

Version	Date	Description	Owner name
1	2026-06-24	Initial version.	Keith Bechtol

Document source location: <https://github.com/lsst-so/sotn-004>

Contents

1 Introduction	1
2 Current Delivered Image Quality Budget	1
3 Image Quality Improvement Priorities, Milestones, and Overview of Current Issues	3
3.1 Progression of Technical Performance Milestones	3
3.2 Summary of Image Quality Known Issues and Mitigations	4
3.3 Roadmap	6
4 High-priority Upcoming AOS Deployments	7
4.1 Donut Selection and Fit Quality	8
4.2 Optimized Optical State (Fixed Pointing)	9
4.3 Optimized Optical State (Survey Mode)	11
4.4 AOS Closed-loop Correction Latency	12
4.5 AOS Closed-loop Control Algorithm	12
4.6 Tracking Drift	13
4.7 Angle of Incidence Throughput Effect	14
5 Ongoing Optimization of the Environmental Control	14
5.1 Dome Thermal Environment	14
5.2 M1M3 Thermal Control	15
5.3 Heat Sources Near Optical Beam	16
5.4 Temperature LUT	17
6 Longer Term AOS Optimization	17
6.1 Difference in Intra- versus Extra-focal Donuts	17
6.2 Wavefront Estimation Methods using Machine Learning	18
6.3 Guiding	18
6.4 Image Quality Optimization and Survey Strategy	19
6.5 Complete Optimization of the Open-loop Model	19

7 Methodology Notes and Resources	19
7.1 Work Management	20
7.2 Version Control	20
7.3 On-sky Test Planning	21
A Acknowledgements	22
B References	22
C Acronyms	22

Image Quality Improvement Plan

1 Introduction

This technote provides a snapshot of plans as of June 2026 to improve the delivered image quality at the Vera C. Rubin Observatory based on findings from the Early Operations system optimization period. The June 2026 timeframe marks a transition from intensive on-sky engineering activity in March-April 2026 towards a more steady-state development cycle, with sustained wide-area pre-LSST observations driven by the Feature Based Scheduler (FBS) on most nights. Rather than working through a steady backlog on-sky engineering needs, targeted tests are now being performed on designated nights as new capabilities are ready to be validated on sky. This is a natural moment to assess the current performance, and to consider strategies for making continued improvements largely *in parallel* with sustained survey-mode observations.

The technote Bechtol (RTN-122) reports the distribution of demonstrated performance during the Early Operations period from October 2025 through May 2026. The image quality from Rubin Observatory has improved substantially since the first on-sky observations with LSSTCam in April 2025, but is not yet consistently meeting design specifications in pre-LSST observations. While the progress during the Early Operations period is encouraging, and the current delivered image quality will enable many LSST science cases, the performance can still be improved, especially to reduce the ellipticity and to increase the uniformity of the PSF shape across the full field of view.

The following sections describe planned improvements for the delivered image quality over the next months considering the ensemble of subsystems at the summit facility as well as operational procedures. We place particular emphasis on the Active Optics System (AOS), as the needs for continued on-sky engineering during the next months are expected to largely be driven by testing to improve the AOS performance.

2 Current Delivered Image Quality Budget

Table 1 summarizes the current measured and estimated system contributions to the delivered image quality budget for PSF FWHM relative to design specifications. The total system

TABLE 1: Current delivered image quality budget for PSF FWHM

Source of Degradation	Budgeted (arcsec)	Measured (arcsec)
Static Optics		
Optical Figure (includes AOS and mirror thermal gradients)	0.202	0.4
Dynamic Optics		
Top end and M1M3 vibrations	0.091	small, in spec
Mount Motion		
Tracking drift	0.025	0.1
Elevation, azimuth, rotator jitter	0.083	0.02
Seeing		
Upper atmosphere	(not included)	(not included)
Dome seeing	0.084	0.3 – 0.4
Mirror seeing	0.027	small, in spec
Instrument		
Diffusion of photoelectrons	0.207	0.207
Total	0.4	0.55 – 0.6

contribution to delivered image quality is not yet meeting design specifications, with larger than allocated contributions from the optical figure (including AOS and mirror thermal gradients) and dome seeing, in particular. The quoted values correspond to estimated median values of the distribution of performance. The tail of the distribution for the total system contribution to the delivered image quality is likely to have a larger relative optical figure contribution, as evidenced by gradients of the PSF size and shape across the full field of view.

Distinguishing contributions from the dome seeing and upper atmosphere is challenging, and accordingly, we provide a range for the estimated dome seeing contribution in Table 1. The median value of the atmospheric contribution (at 500nm at zenith) from the baseline survey simulations is 0.72". An upper atmosphere seeing contribution of 0.72" at zenith corresponds to 0.8" at a typical median airmass of 1.2 for LSST observations. This atmosphere contribution added in quadrature with a total system contribution of 0.6", including a dome seeing contribution of 0.35", aligns well with the measured distribution of delivered image quality from pre-LSST observations with median PSF FWHM of ~1.0" in the riz bands. We have observed correlations of dome seeing indicators (e.g., donut blur) with wind speed and dome aperture orientation relative to the wind direction (Section 5.1).

3 Image Quality Improvement Priorities, Milestones, and Overview of Current Issues

This section describes strategy, priorities, milestones for image quality improvements, and a summary of specific known issues. Over the next months, as the currently known issues are addressed, and new issues emerge, the overall image quality effort will need continuous management and clear prioritization of tasks.

3.1 Progression of Technical Performance Milestones

At a high level, the current analyses, development, and dedicated on-sky tests center around two primary objectives to continue improving the delivered image quality:

1. characterize the “target” optical state that the system is converging to at fixed telescope pointing
2. characterize the stability / visit-to-visit variations of delivered image quality, initially at fixed telescope pointing, and then for wide-area survey observations

A sequential progression of technical performance milestones (with emphasis on optics control)

1. Achieve design performance for delivered image quality for FBS-driven stability tests with MTAOS at a single fixed telescope pointing (elevation, azimuth, rotator angle, filter)
 - validates that system has capability to converge to intended “target” optical state
 - validates stability of delivered image quality / visit-to-visit variation in most ideal observing configuration
 - until Rubin Observatory meets this performance benchmark, we will need regular dedicated on-sky engineering
2. Achieve design performance for delivered image quality for FBS-driven stability test with MTAOS at fixed pointing for any particular choice of (elevation, azimuth, rotator angle, filter) and for several environmental conditions

- validates that the system has capability to converge to the intended “target” optical state for the range of operational telescope configurations and environmental conditions.
 - foundational for developing and validating complete and robust open-loop model, also called look-up tables (LUTs)
3. Achieve design performance for delivered image quality in controlled experiments, such as alternating between two elevations, two rotator angles, or two filters
 - validates application of differential corrections in the open-loop model
 4. Achieve design performance for delivered image quality for standard LSST observations
 - validates the closed-loop feedback and open-loop model together in full range of nominal operating conditions

It is expected that progress will be made on multiple aspects of the delivered image quality budget at varying rates. Some of the issues are connected and require improvements in multiple related areas. For example, the performance of the active optics is somewhat correlated with the M1M3 thermal stability.

3.2 Summary of Image Quality Known Issues and Mitigations

Table 2 summarizes currently known image quality issues. The issues are grouped approximately by timescale for expected improvements.

The first category (see Section 4) includes high-priority AOS-related issues that we have clear plans to address within the next 3 months. For these issues, diagnostic data has already been collected that enables offline analysis and feature development, and feature deployments are already in progress. Several of these deployments will require some amount of dedicated on-sky engineering to validate performance in controlled tests, or to collect additional targeted diagnostic data.

The second category (see Section 5) includes ongoing optimization of the environment control over the next 6-12 months, primarily through daytime engineering work that does not directly compete for on-sky time. Several of the issues include multiple related items, or multiple hardware components that are expected to be installed over a period of time.

TABLE 2: Summary of currently known delivered image quality issues

Issue	Planned Developments / Mitigations	Resolution Timescale
<i>High-priority Upcoming AOS Deployments</i>		
Donut selection and fit quality	WCS-based donut selection; pupil model; unbinned fit; numerical stability and internal consistency investigations	Jul 2026
Optimized optical state (fixed pointing)	Implement measured intrinsic wavefront and optimized open loop model for fixed pointing	Aug 2026
Optimized optical state (survey mode)	Perform full array mode “bounce test” observing blocks to build elevation / rotator LUT	Jan 2027
AOS closed-loop correction latency	Achieve consistent N+2 latency >95% of the time through algorithm improvements	Nov 2026
AOS closed-loop control algorithm	Optimize PID controller, adding I term; use predictive control algorithm such as Smith Predictor or Kalman Filter	Dec 2026
Tracking drift	Apply synchronized AOS and pointing corrections	Jun 2026
u,y band image quality	Angle-of-incidence throughput correction	Jul 2026
<i>Ongoing Optimization of the Environment Control</i>		
Dome thermal environment	Dome Louvers; facility HVAC; increase insulation for thermal sources	May 2027
M1M3 thermal control	M1M3 nozzles; M1M3 glycol pump and valves; control loop optimization	Nov 2026
Heat sources near optical beam	Continue regular thermal surveys; hexapod idle feature; heat sources on Top End Assembly; M1M3 Mirror cover motors	Dec 2026
Temperature LUT	Additional temperature sensors; study the variation of focus as a function of temperature, after the M1M3 thermal control is optimized	Apr 2027
<i>Longer Term AOS Optimization</i>		
Difference in Intra/Extra Donuts	Explore pupil model; optical modeling; targeted diagnostic data to evaluate wavefront measurement	Oct 2026
ML-based Wavefront Estimation Methods	Tune training; brief on-sky tests; quantitative assessment of performance	Dec 2026
IQ Optimization and Survey Strategy	Study interplay of IQ conditions on LSST Survey	Feb 2027
Complete Optimization of the open-loop model	Optimize the open-loop model looking at trends in the corrections as a function of parameters such as elevation, rotation; perform analyses together with intrinsics	Jun 2027

The third catalog (see Section 6) includes outstanding open questions and exploratory efforts where detailed work plans are under development, and timelines for improvement are more uncertain. Several of these items require large statistical samples of survey observations sampling a range of environment conditions as input.

3.3 Roadmap

We describe an example (highly simplified) roadmap to help connect individual tasks as part of an overall action plan for image quality improvement.

The highest priority areas to achieve the first milestone of meeting design performance at a single fixed telescope pointing are

- Donut selection and fit quality – Section 4.1
- Optimized Optical State (Fixed Pointing) – Section 4.2
- M1M3 thermal control – Section 5.2

It is expected that substantial progress can be made in these areas by ~September 2026.

The next highest priority area to achieve design performance at arbitrary fixed telescope pointing, and in controlled experiments, is

- Optimized optical state (survey mode) – Section 4.3

It is expected that substantial progress can be achieved by late 2026.

The next highest priority areas to achieve design performance in survey-mode observations are

- AOS closed-loop correction latency – Section 4.4
- AOS closed-loop control algorithm – Section 4.5

It is expected that substantial progress can be achieved by late 2026

The environment control is high priority, but involves daytime engineering work and does not directly compete for on-sky time

- Dome thermal environment – Section 5.1
- Heat sources near optical beam – Section 5.3

Progress in these environment control areas is expected to be incremental over time.

The other areas are expected to improve the overall consistency of performance across the range of observing conditions. Several of these investigations require analyses of the accumulated telemetry from extended periods of survey operations, and accordingly, these optimization efforts will likely extend into calendar year 2027.

Figure 1 provides a visual summary of the image quality improvement plan roadmap.




FIGURE 1: Image quality improvement plan roadmap visual reference.

4 High-priority Upcoming AOS Deployments

We describe specific areas of active research and development with more details on current challenges and plans for each area.

4.1 Donut Selection and Fit Quality

Several improvements for reliable wavefront retrieval are being investigated. We discuss several of these investigations in turn below.

First, the current online wavefront estimation pipeline identifies donut images for fitting on each individual image from the corner wavefront sensors during runtime. Provided that the telescope pointing model is sufficiently accurate, improvements in both computational speed and robustness could be achieved by using a catalog of bright isolated stars to identify donut images for fitting, which is the baseline method of donut selection. The pointing model is connected to open-loop and closed-loop corrections applied by the AOS since the AOS applies corrections to the hexapods (Section 4.6). As a result, there have been multiple iterations to refine the pointing model as the AOS is developed in parallel.

Second, the corner wavefront sensors are located at the periphery of the field of view (~ 1.7 deg off boresight) where the pupil shape is affected by vignetting due to the Camera body, as well as the M2 edge, M3 edge, and Camera L1 entrance and filter entrance intruding on the M1 annulus. The most highly vignetted parts of the corner wavefront sensors are currently excluded from the online wavefront estimation pipeline. Measurements of optical baffles around M1M3 have been acquired at the summit to more accurately model the absolute pupil size. There are ongoing studies to validate and refine the pupil model used for wavefront estimation.

Third, for speed of the online wavefront estimation (Section 4.4), “donut” pixel images are binned to 2×2 before fitting. Direct comparison of the retrieved optical aberrations using binned (2×2) versus unbinned (1×1) donut images shows some significant differences, including spatial structure across the field of view in fits of full array mode data which might be attributed to variations in the pupil model. If the 1×1 binning is needed, further work to speed up the online wavefront estimation pipeline will be required.

There are also ongoing studies to validate the numerical stability and robust convergence of fitting routines, to optimize the donut quality selection criteria, and to improve the fidelity of the optical model to approach fit residuals that are consistent with statistical noise. A variety of internal consistency checks can be performed, e.g., comparing fits from neighboring donuts on the focal plane, comparing the repeatability of fits from consecutive visits with the same optical state, and comparing fits for pupil images of different sizes corresponding to pistoning

the focal plane across a range of intra- and extra-focal positions. Most recently, an alternative algorithmic approach has been developed that projects the telescope pupil model to the focal plane during each step of the fitting, rather than solving the inverse problem of mapping the observed “donut” image back to the telescope pupil. Offline tests show that this new algorithm yields more numerically stable and repeatable wavefront estimates without the occasional discontinuities in the χ^2 surface observed with previous algorithm. This approach might also more naturally accommodate higher fidelity modeling of the telescope pupil, e.g., including features such as obscuration by the spider struts that support the telescope Top End Assembly.

Improvements for the wavefront retrieval will be primarily developed through offline analyses. Once specific algorithms and configuration settings for the wavefront estimation pipeline have been tested in offline analyses and are ready to be validated on sky, they will be first evaluated using FBS-driven closed loop stability tests (1-2 hrs per configuration change) before progressing to standard survey observations.

4.2 Optimized Optical State (Fixed Pointing)

After running many iterations of the AOS closed loop at fixed telescope pointing (elevation, azimuth, Camera rotator angle, filter), we find that there are often systematic differences between the measured wavefront across the field of view and the optical model for the intrinsic wavefront (or reference wavefront) that defines the target optical state that the AOS is trying to converge towards.

Possible explanations for these differences include the following:

1. wavefront retrieval is systematically inaccurate (Sections 4.1 and 6.1)
2. optical model for the reference wavefront across the field of view, including corner wavefront sensors, is an inaccurate description of the as-built system and thus does not minimize the wavefront error across the field view
3. enabled linear combinations of degrees of freedom are insufficient to drive the system towards to the optimized optical state (corresponding to the reference wavefront)
4. information from the corner wavefront sensors alone is insufficient to fully specify the set of linear combinations of degrees of freedom that are needed to drive the system

towards to optimized optical state

Current analyses imply that multiple of these issues are present at some level, and we are developing approaches to address each issue.

As a first step towards addressing the accuracy of the intrinsic wavefront model, the team has deployed a set of calibrated offsets for each of the corner wavefront sensors using the measured residual optical aberrations (defocus in particular) from a selection of recent best-quality images. Including these target residuals has noticeably reduced the residual defocus tilt across the focal plane and produced more stable on-sky results. The target residuals likely provide a more accurate model of differences in the heights of each corner wavefront sensor relative to the array of science detectors.

A more comprehensive approach is to use full array mode (FAM) observations with the entire focal plane array of science detectors pistoned to intra- and extra-focal positions to study how the optical wavefront varies across the full focal plane, and thereby empirically determine the reference wavefront. Aberrations in the measured wavefront can be described using a Double Zernike (k, j) basis for the wavefront (pupil Zernikes indexed by j), and its variation across the field of view (focal Zernikes indexed by k). The pupil Zernikes can be measured up to at least $j = 22$. Measurements with the four corner wavefront sensors located at the periphery of the field of view can fully constrain the lowest order focal Zernikes — common mode ($k=1$) and dipole ($k=2, k=3$) variations — but only partially constrain higher order focal Zernikes — in particular, “bowl” ($k = 4$) and quadrupole ($k=5, k=6$) variations across the field of view. FAM measurements are needed to fully constrain the $k \geq 4$ focal Zernikes.

In order to minimize degeneracy, and to control the degrees of freedom that are most constrained by information from the corner wavefront sensors, the AOS closed loop is currently using 22 of the 50 effective degrees of freedom available. Optical modeling suggests that additional degrees of freedom will be need to be used to reach the optimized optical state.

The path forward is to use FAM data to directly measure the intrinsic wavefront and derive open-loop corrections for the optical aberrations that are not well constrained by measurements with the corner wavefront sensors alone. Initially, we expect this to be applied as a fixed “average” correction that is only optimal for a single fixed telescope pointing, likely using a larger number of degrees of freedom than used for the AOS closed loop. Initial versions of this analysis are already in progress, and will require on-sky validation using AOS closed loop

stability tests (1-2 hours for each configuration update). The team is developing the capability to run AOS closed loop with interspersed FAM measurements to provide additional diagnostic information on the optical state that the system converges towards.

4.3 Optimized Optical State (Survey Mode)

To extend the work above (Section 4.2) and achieve an optimal optical state across the field of view for any telescope pointing, we intend to build a refined open loop model or look-up table (LUT) of the calibrated intrinsics as a function of telescope elevation and Camera rotator angle using FAM measurements.

Initial sets of FAM measurements have already been acquired for a grid of telescope elevation and Camera rotator angles, and show that there are repeatable (i.e., predictable) changes of optical aberrations across the field of view between telescope pointings. Already, these FAM data show that there is one component of the measured intrinsic wavefront that is fixed in the Camera focal plane coordinate system (at least partially attributed to micron-scale variations in the detector heights across the focal plane) and a second component of the measured intrinsic wavefront that is fixed in the observatory coordinate system related to the telescope optics, as expected. Accordingly, the measured intrinsic wavefront includes two separate components fixed to the Camera Coordinate System and Observatory Coordinate System, respectively, to accurately model the optical state as a function of telescope elevation and Camera rotator angle.

The ensemble of FAM measurements also show that optical aberrations across the field of view (e.g., expressed in a double Zernike basis, see Section 4.2) can vary on the timescale of minutes even at fixed telescope pointing with no commanded changes in the optical state due to latent variables in the system (e.g., temperature changes of the telescope structure and temperature gradients across the primary-tertiary mirror). To control optical aberrations across the field of view while in wide-area survey mode, we have developed an approach to measured intrinsic open-loop model by taking FAM measurements while alternating from two telescope elevations or two Camera rotator angles in rapid succession to mitigate the effects of temporal drifts in the optical state. By taking such “FAM bounce test” observations at a range of telescope pointings, we can build the measure intrinsic open-loop model for arbitrary pointing. This LUT will then be applied differentially as the telescope scans the sky.

Longer term, this open-loop model might be further improved by including additional vari-

ables for system telemetry, as described in Section 6.5.

4.4 AOS Closed-loop Correction Latency

The AOS closed loop design specifies that optical corrections derived from visit index N be applied to visit index $N+2$. Given the exposure time and typical gap time between successive visits, achieving $N+2$ latency requires that the total time delay between shutter close and delivery of optical wavefront measurements must be less than 34 seconds, including readout, image ingest, instrumental signature removal, and wavefront estimation.

Currently, the system is capable of achieving $N+2$ most of the time, and $N+3$ in most remaining instances, with occasional outliers to $N+4$ and larger latency. Most of possible timing improvements in the Camera DAQ, image ingest, and instrument signature removal have been realized, and further gains will involve optimization of the wavefront estimation pipeline, including robust selection of good quality pupil images (Section 4.1), optimizing fit convergence, and investigation of alternative faster online wavefront estimation approaches (see Section 6.2). Nearly all of this work involves offline analysis and development, and the latency impacts on the control loop can be evaluated in offline analysis (Section 4.5) prior to deploying at the summit.

Longer term, further optimization efforts could be pursued to reduce latency in crowded fields attributed to fitting blended donuts.

4.5 AOS Closed-loop Control Algorithm

The AOS closed loop currently uses a PID control algorithm. To mitigate the impact of the typical 2-visit latency between measurement and applied correction (Section 4.4), corrections derived from intermediate visits between a measurement and application of the corresponding closed loop corrections are currently being discarded to avoid overcorrection. The Integrated term of the PID controller is currently set to zero, since there were other issues with the wavefront retrieval and the reference wavefront. Data to include noise in the control loop has been taken for about 1 year to build a covariance term to minimize the noise impact but is not implemented yet.

While some improvements could be achieved by tuning parameters of the current PID con-

troller, alternative control algorithms better matched to address latency (e.g., Smith Predictor) and uncertainties in the wavefront measurement are likely needed to achieve faster and more robust convergence.

This area also includes investigation and handling of edge cases, such as logic around large slews and filter changes.

We have developed a method to reproduce the open loop behavior of the optical system (i.e., backing out the accumulated effect of applied “tweaks”) for the full set of archival Rubin observations in order to accurately simulate and compare the expected performance of PID controllers with different parameters, as well as alternative control algorithms, in offline analyses. Accordingly, control algorithms can be developed within minimal required on-sky engineering time, with targeted on-line validation of specific proposed updates. The anticipated needs for dedicated on-sky engineering are at the level of 1-2 hours for FBS-driven closed loop stability after a configuration change, followed by standard survey observations.

4.6 Tracking Drift

In general, corrections applied by the AOS to improve the optical focus and alignment also induce telescope pointing offsets. This is most apparent when the AOS applies hexapod tip/tilt or translational corrections. Large pointing model errors result in tracking drift during the standard 30 second exposures that contributes to the PSF size and shape (ellipticity).

Pointing offsets related to the AOS optics corrections are predictable based on the telescope optical model, and thus can be pre-emptively compensated by the pointing control loop. A sensitivity matrix relates changes in each optical degree of freedom to two-dimensional offsets in the camera coordinate system. Recent work on the closed loop operations enables synchronization of the AOS and pointing corrections within the standard slew-and-expose sequence, and these AOS pointing corrections have now been validated on sky. Next steps include an update of the core pointing model, potentially adjusting zero point offsets and other terms of the model.

Further work on the AOS might necessitate additional pointing model work. On-sky validation of the pointing accuracy and synchronized AOS and pointing corrections can be accomplished through standard sequences of the AOS closed loop initial focus and alignment at the start of night, and during regular wide-area survey observations, and thus minimal dedicated on-sky

engineering is expected.

4.7 Angle of Incidence Throughput Effect

The Rubin/LSSTCam filters are interference filters with photometric throughput that depend on the Angle of Incidence. The Angle of Incidence varies from the center to the edge of each “donut” pupil image, so the throughput (and hence resulting intensity) varies with radius inside the donut. This Angle of Incidence throughput effect is the leading candidate explanation for the AOS difficulty in the u and y bands, where the angle of incidence correction is roughly four times larger than in the mid-bands, potentially impacting the pupil and therefore focus.

The angle of incidence correction has now been implemented in the baseline wavefront estimation pipeline (Danish 1.1) and on-sky validation began in early June. A full lunation is likely needed to comprehensively validate the correction on sky, as only one of u or y filters is installed at a given time.

Minimal dedicated on-sky engineering time is needed for diagnostic measurements and on-sky validation, at the level of 1-2 hours for u and 1-2 hours for y. This work is already in progress as of June 2026.

5 Ongoing Optimization of the Environmental Control

5.1 Dome Thermal Environment

A major objective of the coming year is to reduce dome seeing through improved thermal conditioning and ventilation of the dome. 12 of the total 34 dome louvers are currently being used in regular nighttime operations, including 6 electromechanically operated louvers and 6 pneumatic louvers. The baseline plan is to have only electromechanically operated louvers and all 34 should be operable by the end of the calendar year. We could add some pneumatic louvers but they are not remote controlled meaning they are fully open during the entire night and they do not provide telemetry. The team is investigating the procurement of additional pneumatic louvers to accelerate the availability of louvers for nighttime operation, at some additional cost, especially as we are headed towards installing and operating the Light and Wind Screen panels that will close the dome and change the impact of the wind on the dome

temperature.

Note that with 10 additional pneumatic louvers for a total of 22 louvers, we will have opened all the large louvers and expect to gain almost all the benefits of the complete set of 34. Even with 12 louvers open, we see evidence that the dome seeing is correlated with wind speed and telescope pointing relative to the wind direction and wind speed. Typical dome seeing is likely at the level 0.3" -- 0.4", with relatively lower (higher) contributions pointing into (away from) the wind direction at moderate wind speeds. The dome seeing contribution is expected to be independent of the pointing direction once additional dome louvers are active.

Additional work related to the dome thermal environment involves increasing the cooling capacity load of the air handling units, increasing insulation for thermal sources/ingress points in the dome and pier, and optimizing operation of the downdraft airflow to inhibit convective heat transfer from lower areas of the Summit Facility into the dome environment.

Nearly all of the dome thermal environment work involves daytime engineering, with no dedicated on-sky engineering required.

5.2 M1M3 Thermal Control

The M1M3 mirror is a cast borosilicate monolith whose shape is extremely sensitive to temperature. Temperature-controlled air enters each honeycomb cell of the mirror through an air nozzle, before returning and being reconditioned by fan coil units. The mirror is cooled to stay ~ 1 C below ambient throughout the night to minimize mirror seeing. However, with its high thermal coefficient of expansion and large variation in thickness across the monolith, the optical figure of M1M3 is sensitive to temperature gradients that are induced by rapidly cooling the mirror. Some optical aberrations, spherical aberrations in particular, are excited almost entirely by M1M3 thermal gradients, and are not readily controlled by any of the M1M3 mirror bending modes.

The team continues to optimize the M1M3 thermal control as part of the Environmental Awareness System to minimize temperature gradients along the optical axis of M1M3 that affect the mirror figure. Two studies are currently in progress that reduce the maximum rate of change for the M1M3 temperature set point, or to keep the setpoint within a configurable range of the M1M3 bulk glass temperature, in order to avoid inducing large temperature gradients. The control optimization work also includes development of more accurate weather forecasting

tools for Cerro Pachón. In addition to internally developed forecasting tools, Rubin Observatory is working with an Italian collaborator who has a validated forecasting model used for other observatories such as the E-ELT.

Daytime engineering work to adjust the air flow rate through 1720 nozzles that circulate conditioned air throughout the M1M3 honeycomb cell structure has led to clear improvements for the control of radial temperature gradients in recent months. Adjustments to the air flow rate as a function of radial distance from the center of M1M3 account for radial variations in the glass thickness of M1M3.

Effectively all of the M1M3 thermal control work involves daytime engineering, offline analyses, and testing in parallel with on-sky observations.

5.3 Heat Sources Near Optical Beam

Rubin Observatory performs routine thermal surveys using an IR camera to identify individual heat sources near the telescope beam and develop mitigation plans for each source. Over the past months, this work has led to substantial reduction of heat sources around the telescope structure, Top End Assembly, and Camera body. A recent example mitigation is the installation of a neoprene thermal shroud covering the Camera hexapods legs where the optical beam crosses three times before reaching the focal plane.

For each individual thermal source, the mitigation approach is specific to the individual subsystem component, typically involving software controls (e.g., to reduce heat generation through more optimized duty cycle), increasing thermal insulation, or adjusting thermal control loops. Specific examples include enabling a feature to idle the hexapod motors between commanded movements to reduce heat generation, and turning off the M1M3 mirror cover motors while observing.

As with the other thermal environment work areas, the activity is primarily daytime engineering with no dedicated on-sky testing required.

5.4 Temperature LUT

This area of investigation broadly covers mitigations for image quality impacts attributed to the changing in-dome thermal environment at the Simonyi Survey Telescope. Specific known issues include

1. a rapidly changing optimal focus (100 microns dz in an hour compared to a depth of focus 10 microns wide in dz) that is at least partially correlated with temperature readings on the surface of the telescope mount assembly structure, and
2. changes in the optical figure of M1M3 (e.g., characterized using a basis of double Zernike for the optical wavefront aberrations in FAM data) are at least partially correlated with temperature gradients across M1M3

Currently, there is a focus offset LUT for each filter and a two-dimensional LUT for telescope elevation and rotator angle. As other aspects of the AOS stabilize, and telemetry correlation analyses are performed with additional survey data, further image quality gains might be realized by incorporating additional telemetry streams into open loop AOS corrections. For example, additional temperature sensors on the telescope structure and for M1M3 might be more predictive of the rapid focus drifts. M1M3 thermal gradients might be used to predict mirror figure changes that are not easily measured using the corner wavefront sensors.

6 Longer Term AOS Optimization

6.1 Difference in Intra- versus Extra-focal Donuts

The current baseline wavefront estimation pipeline performs a simultaneous fit of intra- and extra-focal pupil images on the two halves of each corner wavefront sensor, specifically, considering pairs of intra- and extra-focal pupil images. Visual comparisons of the intra- and extra-focal pupil images clear differences that are not yet fully understood, in particular, that the extra-focal pupil images have more sharply defined features (e.g., edges of the pupil including shadows from the spider struts that support the telescope top end assembly). Moreover, we find systematic differences between measured Zernikes of the optical wavefront when separately fitting the intra- and extra-focal pupil images. In practice, the extra-focal pupil fit results currently dominate the overall wavefront measurements.

Investigations into possible physical origins of apparent differences in the intra- and extra-focal images are under way, including studies of the pupil model (Section 4.1) and subtle diffraction effects. While some data has already been acquired to compare wavefront measurements from a range of pupil image sizes using focus sweeps, additional diagnostic data might be required, for example, targeting a bright star and gradually moving from in-focus to “giant donuts” to examine the behavior of donut features.

6.2 Wavefront Estimation Methods using Machine Learning

The team is currently pursuing two machine learning (ML) approaches for wavefront estimation.

The first approach involves an emulator for the current baseline wavefront estimation pipeline that is trained to reproduce the results of the conventional wavefront estimation pipeline. By construction, the emulator approach can only be as accurate as the conventional wavefront estimation pipeline, but still offers two important potential improvements: rapid inference to reduce correction latency and increased precision / reliability given a sufficiently large and carefully curated training dataset. This method is more robust to seeing conditions and blending donuts, making it valuable in dense fields.

The second ML-based approach is a fully independent model trained to infer the optical wavefront and associated optical state directly from simulated pupil images. In addition to the potential advantages noted above, this approach has the potential to serve as an independent validation and/or alternative to the conventional wavefront estimation pipeline. The accuracy of this approach could ultimately be limited by the accuracy of the optical model used to train the model.

6.3 Guiding

LSSTCam has a total of 8 guider sensors, with 2 guider sensors on each of the 4 corner rafts located on either side of the wavefront sensor for that raft. Currently, the guider images are being used only for image quality diagnostic purposes to evaluate tracking drift and to sample variations of the PSF on timescales much shorter than the standard 30/38-second visit exposure time (rates up to 9 Hz are possible). Tracking drifts could be further reduced by using feedback from the guiders for active control of the telescope pointing.

6.4 Image Quality Optimization and Survey Strategy

The distribution of delivered image quality is connected to the survey strategy through the distribution of telescope pointing elevation angles (i.e., airmass), distribution of telescope slew distances in elevation angle and Camera rotator angle, and effects related to the dome thermal environment and ventilation such as pointing direction relative to the wind vector (see Section 5.1). This work item is a placeholder for the general research topic of how image quality interfaces with the LSST survey strategy.

Joint optimization of the delivered image quality and survey strategy is likely an exercise better suited for after the first year of LSST, given that a large volume of survey observations will be needed as input to the study, other expected improvements for the delivered image quality, priorities for the first year of the LSST survey, and needs for statistical characterization of the distribution of atmosphere seeing. Any potential modifications to the survey strategy would need to be carefully considered, and technical recommendations passed to the Survey Cadence Optimization Committee (SCOC) for review.

6.5 Complete Optimization of the Open-loop Model

This item extends the open-loop modeling work described in Section 4.3 with further refinements that incorporate telemetry from an extended period of survey observations sampling a range of environment conditions. The work includes analyses of trends for the applied corrections (and wavefront residuals) as a function of parameters such as elevation and rotation. Optimization of the open-loop model will need to be considered together with the measured intrinsics.

7 Methodology Notes and Resources

We provide supplemental information on general principles for work management, version control, and the development of on-sky tests.

7.1 Work Management

Analysis and development tasks to support image quality improvements are captured, ranked by priority, and have their progress and outcomes recorded using the Jira epic RSO-9.

7.2 Version Control

All of the planned AOS improvement tests will be interleaved with routine survey operations. From May 2026 onwards, there is increased emphasis on stable delivery of acceptable science-grade images, even as new AOS features and configuration updates are deployed to the summit. It is essential to ensure rigorous tracking of code versions, configurations of the MTAOS, open loop LUT version, intrinsic aberration model, values of any constant offsets applied, number of degrees of freedom enabled and truncation index for the basis of characteristic non-degenerate v-modes. Software versions deployed at the summit will be tagged following a standardized convention.

Process for deployment of new algorithm/versions/configuration at the summit:

- Maintain 2-3 day look-ahead for specific upcoming deployments to coordinate with on-sky test planning
- Proposed configuration changes must be approved by the Software Versioning Acceptance Board (SVAB) prior to be deployed at the summit
- Summit deployments that are not completed before the tailgate 4:30pm Chile will be postponed to the next opportunity
- Preference to schedule deployments Tuesday-Thursday; Monday + Friday as second choice; Deployment on Saturday and Sunday will only be for emergency fixes.
- Preference for atomic updates; avoid more than one new deployment per night and all deployments need to be tracked using the confluence page above.
- Before deployment, determine what is the current stable tagged configuration version that we would revert to if needed
- After deployment, use the following sequence of on-sky tests to evaluate performance

- FBS-driven closed loop stability for 1-2 hrs at fixed telescope (elevation, azimuth, rotator angle, filter)
- Pre-LSST to evaluate performance in wide-area survey mode

Resources used

- Software Versioning Acceptance Board Confluence page
- Tracking AOS Deployment Confluence page
- Spreadsheet w/ configuration status: AOS Configuration Changes

7.3 On-sky Test Planning

The Rubin Observatory Summit Facility Concept of Operations is described in (Kleinman et al., RTN-094).

During the current period with frequent AOS configuration updates and dedicated on-sky engineering tests interleaved with survey observations on a regular basis (e.g., multiple nights per week), we have used test planning on 1-2 week timescales to help connect strategic priorities to the development and execution of individual night test plans. The 1-2 week planning timescale facilitates development of new observing blocks, arranging the schedule of specific feature deployments, and coordinating with other summit activities, including daytime engineering and summit software updates. We outline draft day-by-day test plans with a 7 day lookahead, and detailed day-by-day plans with a 2-3 day lookahead.

Reminder of the sources used

- Week Test Plan (Confluence pages, Canvas)
- Day-by-day detail test planning (maintain 2-3 day look-ahead) Canvas

A Acknowledgements

This material is based upon work supported in part by the National Science Foundation through Cooperative Agreements AST-1258333 and AST-2241526 and Cooperative Support Agreements AST-1202910 and AST-2211468 managed by the Association of Universities for Research in Astronomy (AURA), and the Department of Energy under Contract No. DE-AC02-76SF00515 with the SLAC National Accelerator Laboratory managed by Stanford University. Additional Rubin Observatory funding comes from private donations, grants to universities, and in-kind support from LSST-DA Institutional Members.

B References

- [RTN-122], Bechtol, K., 2026, *Rubin Observatory Early Operations On-sky Performance*, Technical Note RTN-122, NSF-DOE Vera C. Rubin Observatory, URL <https://rtn-122.lsst.io/>
- [RTN-094], Kleinman, S.J., Dennihy, E., Christensen, E.J., et al., 2026, *Summit Facility ConOps v1.00*, Technical Note RTN-094, NSF-DOE Vera C. Rubin Observatory, URL <https://rtn-094.lsst.io/>

C Acronyms

Acronym	Description
AOS	Active Optics System
AST	NSF Division of Astronomical Sciences
AURA	Association of Universities for Research in Astronomy
DAQ	Data Acquisition System
DE-AC02	Department of Energy contract number prefix
DOE	Department of Energy
ELT	Extremely Large Telescope
FBS	Feature-Based Scheduler
FWHM	Full Width at Half-Maximum

HVAC	Heating, Ventilation, and Air Conditioning
IR	infrared
L1	Lens 1
LSST	Legacy Survey of Space and Time
LSST-DA	LSST Discovery Alliance
LSSTCam	LSST Science Camera
LTS	LSST Telescope and Site (Document Handle)
LUT	Look-Up Table
M1	Primary Mirror
M1M3	Single piece of glass for Primary Mirror/Tertiary Mirror
M2	Secondary Mirror
M3	Tertiary Mirror
ML	Machine Learning
MTAOS	Main Telescope Active Optics System
NSF	National Science Foundation
PID	Proportional-Integral-Derivative
PSF	Point Spread Function
RSO	Rubin Summit Operations
RTN	Rubin Technical Note
SCOC	Survey Cadence Optimization Committee
SLAC	SLAC National Accelerator Laboratory
SRD	LSST Science Requirements; LPM-17
WCS	World Coordinate System