Table of Contents

PROJECT NARRATIVE

1.0 Background/Introduction

Dark Energy and Cosmic Acceleration

Understanding the physical origin of **cosmic acceleration** is one of the central open questions in modern cosmology. In the standard ΛCDM model, this acceleration is attributed to **dark energy** in the form of a cosmological constant, Λ, whose energy density remains constant as the Universe expands. A more general possibility is that dark energy is dynamical, with an equation of state that evolves with cosmic time. This evolution is commonly parameterized as [1]

$$w(a) = w_0 + w_a(1-a), \tag{1}$$

where $w(a)$ is the ratio of dark-energy pressure to density, w_0 is its present-day value, and w_a describes its time evolution. A cosmological constant corresponds to $w_0 = -1$ and $w_a = 0$. Detecting a statistically significant departure from this point would be a profound indication of physics beyond the standard cosmological model. Recognizing the fundamental importance of this question, the U.S. Department of Energy (DOE) has made major investments in the Vera C. Rubin Observatory's Legacy Survey of Space and Time (**LSST**) [2]—including the LSST Camera (LSSTCam)—and in the Dark Energy Spectroscopic Instrument (DESI) [3], as part of a broad suite of dark-energy experiments designed to determine whether $w(a)$ deviates from a cosmological constant.

DESI began survey operations in 2021 and is measuring tens of millions of galaxy and quasar redshifts over roughly 14,000 deg^2 to map baryon acoustic oscillations (BAO) [4] and the growth of structure across cosmic time. The DESI DR2 BAO analysis used more than 14 million galaxies and quasars from three years of observations and provides the most precise BAO measurements to date [5, 6]. When combined with cosmic microwave background and Type Ia supernova datasets, these BAO measurements show 3–4σ evidence that dark energy may evolve with time [7]. The present hint is driven primarily by BAO distance information from DESI in combination with these external CMB and supernova datasets, so establishing whether dark energy truly evolves with time demands an independent cosmological probe with different parameter degeneracies and different dominant systematics.

Weak gravitational lensing (WL) [8] is precisely such a probe: by directly measuring the coherent distortions of galaxy images caused by foreground matter, WL is sensitive not only to cosmic geometry but also to the growth of structure, providing exactly the information needed to deliver an independent test of the DESI-motivated hint of dynamical dark energy. Delivering this independent WL test is one of the central dark-energy science goals of Rubin LSST.

A robust WL confirmation of a time-evolving dark-energy equation of state with Rubin LSST would represent one of the most important breakthroughs in cosmology in more than a decade, opening a new window on the physics driving cosmic acceleration. This proposal directly addresses this opportunity by combining LSST and Euclid [9] imaging at the pixel level for the first time at survey scale, reducing the effective shape noise by $\sim 25\%$ through joint ground–space shape measurement while improving photometric-redshift precision by $\sim 30\%$ through the addition of complementary space-based near-infrared imaging, leading to at least a 30% improvement in the LSST WL constraint on dark energy.

Weak Gravitational Lensing with Stage-IV Surveys

The next five years will mark a golden age for weak-lensing cosmology. Data from the Stage-IV imaging surveys—Rubin LSST, Euclid, and the Nancy Grace Roman Space Telescope (Roman) [10]—are arriving in close succession, and **joint pixel-level processing** of Rubin LSST and Euclid is already recognized as an immediate priority for Stage-IV dark-energy science. The simultaneous rapid development of **artificial intelligence (AI)** technologies further amplifies what a joint analysis of these surveys can deliver. Realizing this opportunity, however, hinges on a single technical bottleneck: the complicated shear-calibration procedure required to meet the LSST DESC sub-percent multiplicative-bias requirement, together with the large computational cost of the joint multi-survey image processing and the overlapped simulation campaigns needed to validate it. To remove this bottleneck, in the next subsection we propose a fast and simple analytical shear-estimation framework that is compatible with AI by design.

WL is the small, coherent distortion ("shear") of distant galaxy images by the intervening matter distribution—both visible matter and invisible dark matter—along the line of sight [11, 12, 13, 14]. For any individual background galaxy, the shear signal is more than an order of magnitude smaller than the intrinsic variation in galaxy shapes, so WL must be measured statistically across very large galaxy populations. Dark energy affects WL through two coupled effects: it changes the expansion history of the Universe, altering the distance–redshift relation and the lensing efficiency; and it suppresses the late-time growth of cosmic structure, reducing the amplitude of WL correlations. By correlating the shapes of large numbers of galaxies across the sky and redshift, WL surveys therefore probe cosmic geometry and structure growth simultaneously, providing information complementary to the BAO distance ladder measured by DESI. Equally important, the dominant systematics of WL—shear estimation bias, point-spread-function (PSF) modeling, galaxy blending, intrinsic alignments, and photometric-redshift errors—are largely orthogonal to those of spectroscopic BAO, making WL the decisive independent probe with which to provide an independent test of the DESI-motivated hint of dynamical dark energy.

A coordinated international program of Stage-III and Stage-IV imaging surveys has been built to realize this scientific potential: the pioneering Stage-III surveys are now nearing completion, while several Stage-IV surveys are already operating on sky and the rest will come online over the next few years (Figure 1). The Stage-III imaging surveys—the Dark Energy Survey (DES) [16], the Hyper Suprime-Cam Subaru Strategic Program (HSC) [17], and the Kilo-Degree Survey (KiDS) [18]—have established the modern cosmic-shear methodology and produced some of the most precise measurements to date of the structure-growth amplitude,

$$S_8 \equiv \sigma_8 \sqrt{\Omega_{\rm m}/0.3}. \tag{2}$$

The three Stage-III cosmic-shear analyses provide a compelling and largely independent view of late-time structure growth, finding values of S_8 roughly 0.5–2.5σ below the value inferred from *Planck* CMB measurements [19] under ΛCDM. This emerging "S_8 tension" highlights the power of weak lensing as a complementary probe of cosmology and motivates a decisive Stage-IV test. With substantially larger area, greater depth, improved redshift information, and tighter control of shear systematics, Stage-IV imaging will transform this suggestive Stage-III result into a high-precision measurement of structure growth and an independent test of the DESI-motivated hint of dynamical dark energy.

The scientific case for joint LSST–Euclid weak-lensing analysis is well established and quantitative. Catalog-level forecasts show that combining LSST and Euclid shape measurements of the same galaxies can already increase the effective galaxy number density $n_{\rm eff}$ of the LSST WL sample by $\sim 50\%$, equivalent to adding $\sim 50\%$ more LSST exposure time [20], and performing the detection and shape measurement jointly at the

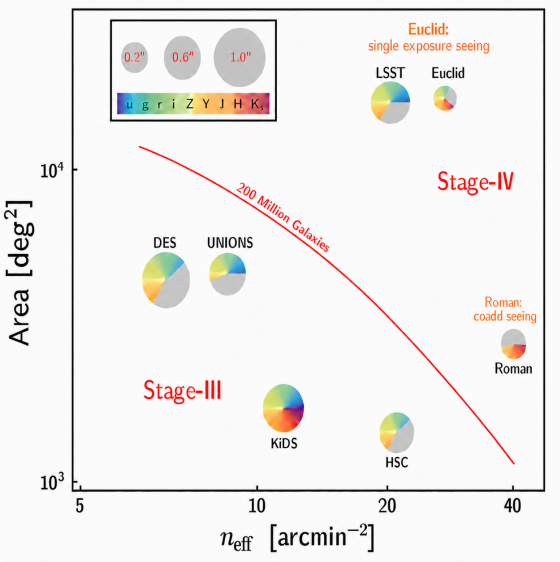

Figure 1: The Stage-III and Stage-IV imaging surveys. The diameter of each disk represents the size of the point-spread function (PSF) of the corresponding survey, and the color encodes its wavelength coverage from the ultraviolet (u band) to the near-infrared (K band).

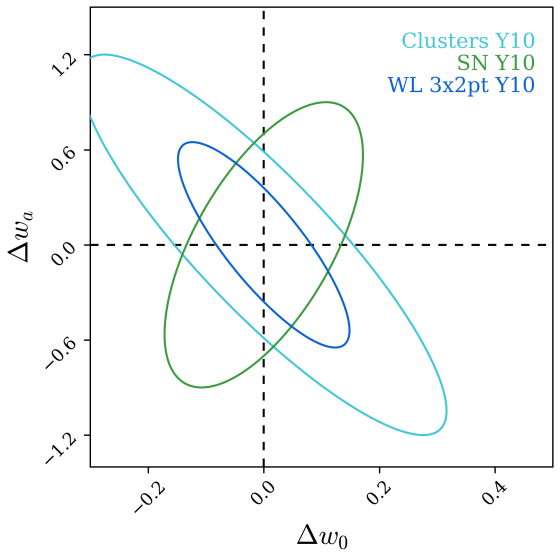

Figure 2: Forecast LSST ten-year 68% constraints on Δw_0, Δw_a from WL 3×2pt (blue), Type Ia SNe (green), and clusters (cyan). WL gives the tightest individual constraint and is the focus of this proposal. Adapted from Fig. G2 of the LSST DESC Science Requirements Document [15].

pixel level—the approach pursued in this proposal—is expected to deliver still larger gains, as identified by the 2020 Joint Survey Processing Study Group [21]. More importantly, joint processing also improves the redshift information that underpins tomographic weak-lensing cosmology: adding Euclid near-infrared photometry to LSST *ugrizy* imaging reduces the photometric-redshift catastrophic-outlier rate by $\sim$ 25–40%, especially at $z \gtrsim 1$, where the optical color sequence becomes degenerate because of 4000 Å-break/Lyman-break confusion [22, 23]. Both gains arise from the same underlying complementarity—LSST contributes optical depth and multi-band sampling, while Euclid contributes space-based resolution and NIR coverage—and they reinforce each other at the pixel level: pixel-level forced photometry across the two surveys delivers consistent u-through-H flux measurements for the same sources, avoiding the catalog-level matching ambiguities that limit existing multi-survey combinations and yielding narrower, less biased $n(z)$ for tomographic shear. Moreover, the proposed differentiable pipeline naturally propagates shear responses through both shape and flux measurements, allowing analytical correction of the magnitude-cut selection biases introduced when galaxies are binned into tomographic redshift slices.

These shape and redshift gains together make joint pixel-level processing of Rubin LSST and Euclid an immediate priority for Stage-IV dark-energy science, which is also recognized by the corresponding Astro2020 Decadal Survey white paper [24] listing joint shear analysis for weak-lensing cosmology among its priority science goals. This proposal focuses on the LSST–Euclid overlap, where both the area and the near-term data products enable a survey-scale demonstration. The resulting pipeline will be built to extend naturally to Roman once Roman coadd imaging products are released, with Roman's deeper, higher-resolution near-infrared imaging providing an additional anchor for blending, morphology, and photometric-redshift systematics.

However, realizing the joint-imaging opportunity laid out above hinges on a single technical bottleneck:

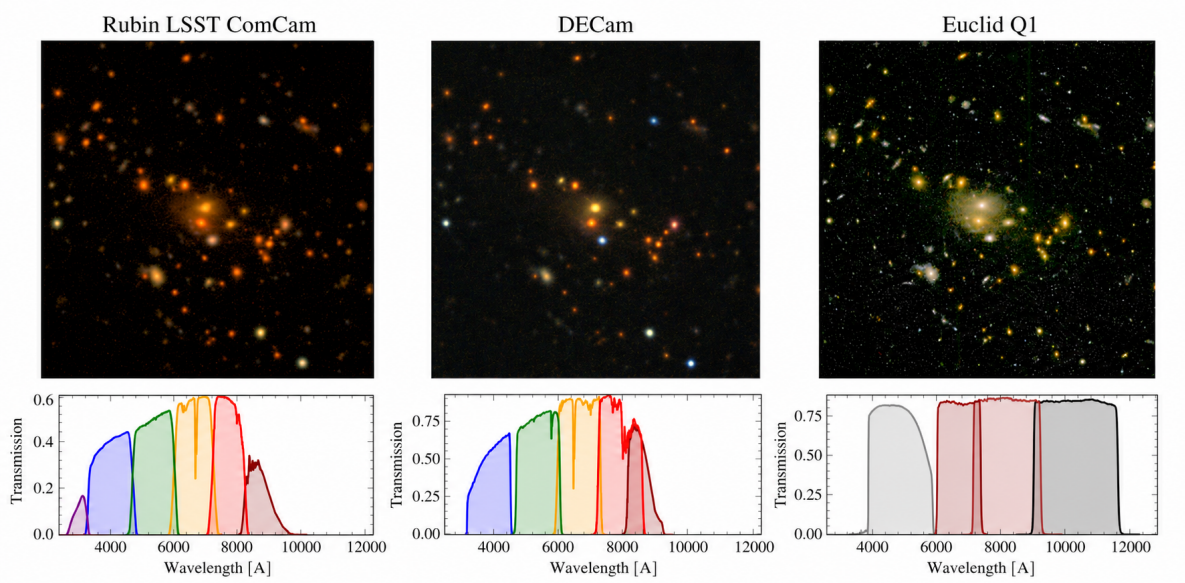

Figure 3: *Top:* Rubin Commissioning Camera Data Preview 1 (left), DES (middle), and Euclid Quick Data Release 1 (right) imaging of the same galaxy cluster, the X-ray cluster from the first eROSITA All-Sky Survey [25] (eRASS1; $\alpha = 59.4873°$, $\delta = -49.0003°$, $z = 0.6922$), located in the Euclid Deep Field South (EDF-S). *Bottom:* corresponding filter transmission curves for each survey, illustrating the wavelength coverage spanned by Rubin's optical *ugrizy* bands, DES's optical *grizy* bands, and Euclid's broad VIS and NIR (Y, J, H) bands. The three surveys differ in depth, angular resolution, and wavelength coverage, and joint pixel-level analysis across them is the basis of the proposed work.

meeting the LSST DESC subpercent-level shear-calibration requirement [15] at survey scale on *multi-resolution, multi-PSF, multi-band* data. Traditional high-accuracy approaches rely on applying small artificial shears to the observed images and rerunning key parts of the analysis many times. A full detection-bias correction requires not only remeasuring galaxy shapes, but also creating multiple sheared realizations of the imaging data, repeating detection on each realization, and propagating the response through deblending, measurement, and selection, while validating that the final residual bias meets the LSST DESC threshold. Furthermore, in a joint LSST–Euclid analysis, the same machinery must also supply the shear response of the measured *fluxes*, since these fluxes feed the photometric-redshift estimator that defines the tomographic redshift bins, and any bin-edge selection on sheared galaxies introduces an additional multiplicative bias that has to be corrected at the same sub-percent level. Implementing this in the existing re-rendering framework requires re-running detection, photometry, photo-z estimation, and tomographic binning on every sheared realization—a calibration chain that is technically intricate to validate and computationally demanding even at single-survey scale. These repeated image-level operations are technically complex to implement and computationally expensive for joint-survey image processing, which is why the long-recognized promise of joint ground–space weak-lensing analysis has not yet translated into an operational survey-scale shear pipeline.

AnaCal: Robust and Efficient Self Calibration for Shear Estimation

This proposal develops a survey-scale, multi-resolution joint-imaging shear estimator for combined ground- and space-based data, closing the missing-pipeline gap identified above: the absence of a method that can meet the LSST DESC sub-percent shear-calibration requirement with realistic resources. The proposed work builds on the PI's **AnaCal** [26, 27] framework, which is purpose-built to remove this bottleneck: AnaCal computes the shear response in closed form and propagates it by automatic differentiation, so that no counterfactual shearing of the images is required and the calibration is two orders of magnitude faster than the artificial-shearing approach. The same differentiable construction also makes AnaCal AI-compatible by design, supporting model fitting features and modern, symmetry-equivariant neural-network feature extractors within a single self-calibrated pipeline. The proposed program extends AnaCal into a joint ground–space shear-estimation pipeline for Stage-IV weak lensing.

The accuracy of any shear-estimation pipeline is commonly summarized by [28, 29]

$$g_{\rm obs} = (1+m)g_{\rm true} + c, \tag{3}$$

where $g_{\rm obs}$ is the observed shear, $g_{\rm true}$ is the true shear, m is the multiplicative calibration bias, and c is the additive bias. Controlling these bias terms is one of the most important systematics challenges for cosmic-shear analyses; the LSST DESC SRD requires $|m| < 3\times10^{-3}$ across different redshit bins [15]. Modern symmetry-informed, self-calibrating estimators have already demonstrated strong control of these biases in Stage-III analyses, achieving $|m| < 0.2\%$ with the Metadetection and AnaCal frameworks (see below). These state-of-the-art pipelines are *perturbative*: they calibrate the shear response by expanding the measured observables to first order in the applied shear, and exploit the spin-2 symmetry of the lensing distortion to suppress higher-order bias terms,

$$\langle e\rangle = \cancelto{0}{\langle e_0\rangle} + g\left\langle\frac{\partial e}{\partial g}\right\rangle + g^2\cancelto{0}{\left\langle\frac{\partial^2 e}{2\partial g^2}\right\rangle} + \mathcal{O}(g^3). \tag{4}$$

Metacalibration [30, 31] and its extension for detection-bias calibration, **Metadetection** [32], estimate the shear response *numerically*: small artificial shears are applied directly to the image pixels, the relevant analysis steps are rerun (including detection for Metadetection and shape measurement for Metacalibration), and the shear response is obtained by finite-differencing the measured shapes across the sheared realizations. This approach has demonstrated part-per-thousand calibration accuracy and underpins the DES Y6 shear catalog ($|m| < 3\times10^{-3}$), but it scales poorly with survey size: four to five sheared image realizations must be generated and processed for every object.

AnaCal, by contrast, computes the analytical shear response for each smoothed image pixel, and the shear response of each image pixel is propagated forward through the detection, deblending, and measurement procedure to the final observables [33], while noise-induced bias is corrected analytically by exploiting the symmetry properties of the noise field. Automatic differentiation through this construction yields a fully differentiable estimator that enforces the spin-2 symmetry of weak lensing, suppresses the leading second-order shear estimation bias terms (see Equation (4)), and achieves $|m| < 0.3\%$ across all the redshift bins, comfortably below the LSST DESC shear-calibration requirement [15]. Figure 4 demonstrates this calibration accuracy on image simulations under LSST ten-year observing conditions: the residual multiplicative bias m is shown as a function of photometric redshift, with the photo-z estimated using `FlexZBoost` [34] within the `RAIL` [35] framework. The dashed line shows the bias from AnaCal alone, and the solid line shows the bias after the analytical self-calibration (without reling on any external image simulations) for

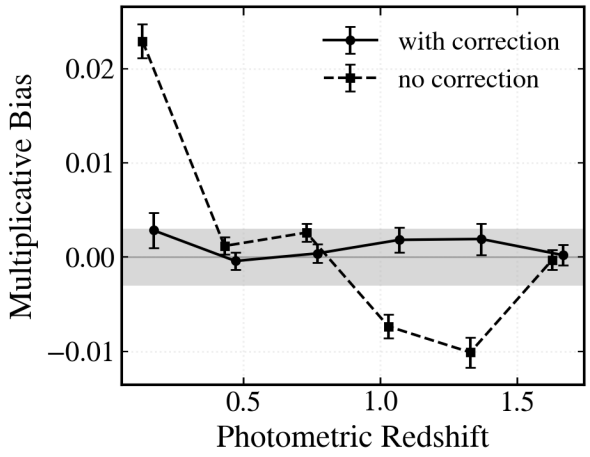

Figure 4: Multiplicative shear bias m as a function of `FlexZBoost` photometric redshift on image simulations under LSST ten-year observing conditions. The *solid* line includes the analytical self-calibration for the selection bias induced by binning galaxies into tomographic redshift bins; the *dashed* line is without that correction. The *grey* band marks the LSST DESC requirement [15].

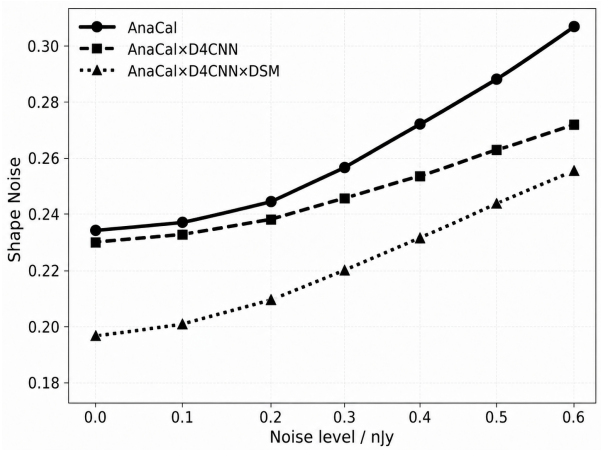

Figure 5: Per-component galaxy shape noise as a function of pixel-noise standard deviation $\sigma_{\rm noise}$ on LSST-like image simulations. *Solid:* AnaCal alone; *dashed:* AnaCal combined with a D_4-equivariant CNN [36]; *dotted:* AnaCal + D_4CNN further optimized with denoising-score matching. The expected LSST Year-10 i-band noise level is $\sigma_{\rm noise} = 0.59$.

Survey	**Method**	**Area [deg^2]**	$n_{\rm eff}$ **[arcmin^{-2}]**	z **range**	$\sigma(S_8)$
DES Y6	Metadetection	$\sim$ 4,200	$\sim$ 8.2	0.3–1.5	$\sim$ 0.015
HSC Y6	AnaCal	$\sim$ 800	$\sim$ 15	0.3–1.5	$\sim$ 0.019

Table 1: Footprint, effective source density, source-redshift range, and measured uncertainty on S_8 for the Stage-III cosmic-shear analyses. The DES Y6 entry is the forecast for the final DES analysis. HSC Y6 refers to the HSC-SSP collaboration's internal Year-6 dataset (*not* unbinded and published yet); the corresponding cosmic-shear analysis is currently unpublished and blinded, and final published values may differ slightly from those quoted here.

selection effects induced by the tomographic redshift binning. With the selection-bias correction included, AnaCal meets the LSST DESC SRD requirement (grey band) across all redshift bins of the LSST Y10 cosmic-shear analysis.

Because the pixel-level shear response is computed directly rather than by rerunning the analysis on artificially sheared images, AnaCal is roughly two orders of magnitude faster than Metadetection: processing the full LSST coadd image set in all six *ugrizy* bands at survey scale requires only $\sim 10^4$ CPU hours, equivalent to a few hundred node-hours on the NERSC Perlmutter system. The same differentiable construction extends naturally to physics-informed AI shear estimation: the automatic-differentiation framework that propagates the shear response through hand-crafted moment measurements and galaxy model fitting [33] can equally propagate it through machine-learned, D_4 symmetry-equivariant representations of galaxy images [36], enabling more flexible feature extraction while preserving analytical self-calibration.

AnaCal has already been deployed at survey scale: it is the shear pipeline for the ongoing HSC Year-6 cosmology analysis (PI is a core contributor), has been run on LSST Commissioning Camera Data Pre-

view 1 [37], is currently being extended to LSST Data Preview 2, and has been validated for the chromatic-PSF effects relevant to joint Rubin–Roman weak lensing [38, 39, 40].

Table 1 compares the resulting HSC Y6 cosmic-shear analysis with DES Y6 in terms of survey area, source density, redshift range, and measured uncertainty on S_8 (Equation (2)).

Toward Physics-Informed AI Shear Estimation with AnaCal

The fast, AI-compatible analytical pipeline introduced in the previous subsection is the natural foundation on which to build the next generation of weak-lensing shear estimators. Beyond the perturbative pipelines, deep neural networks can learn flexible pixel-level representations that capture information relevant to shape measurement, deblending, and multi-band galaxy morphology, going beyond a fixed set of hand-crafted moment features. **Physics-informed AI shear estimators**, in which neural networks act as the feature extractor while the shear response is still analytically propagated and calibrated through AnaCal, are therefore a natural extension of the framework: they aim to reduce the per-galaxy shape noise by an additional 10–20% on top of the perturbative AnaCal baseline, while preserving the self-calibration needed for sub-percent shear-bias control.

The central technical challenge is that black-box machine-learning shear estimators are not, on their own, calibratable to the LSST DESC sub-percent requirement [41, 42]: blindly trained networks readily introduce shape-dependent multiplicative bias at the 10^{-2} level, and recovering sub-percent calibration through traditional simulation-based recalibration is computationally prohibitive at survey scale. This obstacle can be removed by combining two ingredients. First, the network architecture is made D_4-*equivariant* [43]: the D_4 point group of 90° rotations and reflections is embedded directly into the network weights, so the predicted galaxy ellipticity transforms exactly as a spin-2 quantity under any rotation of the input image. This eliminates the orientation-dependent biases that plague non-equivariant ML shear estimators and reduces the network's parameter count by a factor of ~ 8, yielding a $\sim 10\times$ training-data-efficiency gain at no accuracy cost. Second, the analytical AnaCal calibration framework described above is applied to the network's outputs through automatic differentiation: backpropagating the analytical pixel-to-feature shear response through the trained network provides a closed-form shear response of every output, suppressing the leading second-order bias terms by construction and removing the need for expensive simulation-based recalibration entirely.

A first realization of this D_4CNN $\times$ AnaCal architecture [44] has been demonstrated on isolated galaxies in LSST-like single-band image simulations, achieving multiplicative biases consistent with zero across a wide range of noise levels, PSF sizes and ellipticities, and magnitude cuts (**all measurements satisfy $|m| < 10^{-3}$, most at the $\sim 10^{-4}$ level**), while delivering $\sim$10% lower shape noise than the moment-based estimator in the high-noise regime—equivalent to a $\sim$20% gain in effective galaxy number density. A complementary advance comes from *denoising score matching* [45], which trains the equivariant feature extractor to align with the analytical score function of the image likelihood—provably the minimum-variance unbiased shear estimator. Applying denoising score matching to the D_4CNN $\times$ AnaCal features further reduces the per-component shape noise by another 10% at the noise levels relevant for LSST Year 10, as shown in Figure 5, while preserving $|m| < 10^{-3}$ analytical calibration.

The joint-imaging infrastructure developed in this proposal (image simulations, validation suite, AnaCal calibration layer) is designed to host this evolving family of equivariant AI shear estimators as a drop-in replacement of the perturbative AnaCal feature extractor. Two parallel ecosystems already provide the natural backbone for this drop-in: the **DeepDISC** [46, 47, 48] deep-learning pipeline, which jointly performs detection, segmentation, classification, and probabilistic photometric-redshift estimation in a single forward pass;

and large pretrained foundation models such as **AION-1** [49], trained at scale on multi-survey astronomical imaging and capable of providing transferable, generic representations on which a D_4-equivariant AnaCal head can be fine-tuned. Deploying the calibrated analytical pipeline first and then dropping in an AI feature extractor that inherits its calibration through automatic differentiation is the practical path from today's perturbative methods to the AI-native, image-to-cosmology inference required for full Stage-IV exploitation.

2.0 Project Objectives

The overarching goal of this project is to deliver, for the first time at survey scale, a joint pixel-level weak-lensing shear measurement framework for ground- and space-based imaging from Rubin LSST and Euclid, and to use it to sharpen Rubin LSST's constraints on dynamical dark energy. The program is organized around two coupled thrusts—joint shear estimation and joint photometric-redshift estimation—both built on a fully differentiable, analytical shear-calibration pipeline (AnaCal) that meets the LSST DESC SRD sub-percent shear-bias requirement while improving both the precision of shear estimation and the performance of photometric-redshift estimation. The combined effect of these advances will improve the LSST WL constraint on dark energy by at least 30% using the Rubin LSST×Euclid joint analysis, and the pipeline is designed to extend naturally to a joint LSST×Euclid×Roman shear and photometric-redshift analysis once Roman coadd data become available. The resulting infrastructure is designed both as a drop-in component of the LSST DESC 3×2pt pipeline and as a foundation for next-generation AI-based image-to-cosmology inference.

Scientific Merit

The scientific case for this program rests on four cosmological deliverables, each of which directly addresses one of the outstanding open questions in late-Universe cosmology and feeds into the LSST DESC dark-energy science program.

(1) An independent weak-lensing test capable of confirming or refuting dynamical dark energy. The proposed joint LSST×Euclid analysis will improve the LSST weak-lensing 3×2pt constraint on the dark-energy equation of state, parameterized by w_0–w_a, by at least $\sim 30\%$ relative to a single-survey LSST analysis. This gain accelerates the survey epoch at which Rubin LSST can deliver an independent weak-lensing test of the current 3–4σ hint for a time-evolving $w(a)$ reported by DESI DR2 in combination with other cosmological probes [7]. The key advance is not simply higher precision, but independence: weak lensing probes both cosmic geometry and the late-time growth of structure, with parameter degeneracies and dominant systematics that are distinct from those of spectroscopic BAO, supernova distances, and CMB anisotropies. A consistent weak-lensing result would provide independent evidence for dynamical dark energy, while an inconsistent result would point to statistical fluctuation, probe-combination effects, or residual systematics in the current evidence.

(2) A validated joint shear catalog as the foundation for field-level LSST DESC analyses. The joint shear and photometric-redshift catalogs produced by this proposal will provide the calibrated data layer needed for future field-level weak-lensing inference. Field-level analyses aim to model the observed shear, density, and survey fields directly, rather than compressing the data only into two-point summary statistics. This is especially important for joint Stage-IV surveys, because the LSST–Euclid overlap, LSST-only regions, and Euclid-only regions have different depths, PSFs, wavelength coverage, noise properties, blending behavior, selection functions, and photo-z performance. In traditional catalog-level analyses, these discontinuities make it difficult to combine overlapping and non-overlapping regions optimally without introducing survey-inhomogeneity systematics. Field-level inference provides a natural path beyond this limitation: spa-

tially varying survey properties can be included directly in the forward model, allowing the LSST–Euclid overlap to serve as the highest-information, best-calibrated region while still coherently combining cosmological information from the larger LSST-only and Euclid-only footprints [50, 51]. By producing co-registered, multi-resolution, multi-PSF, multi-band shear and photo-z measurements in the overlap region, this proposal provides the calibrated anchor required for future field-level dark-energy analyses that can exploit the full Stage-IV imaging footprint.

(3) Sub-percent measurement of S_8 to resolve the S_8 tension. The Stage-III cosmic-shear surveys (DES, HSC, KiDS) consistently report values of S_8 that lie 0.5–2.5σ below the value inferred from the *Planck* CMB under ΛCDM (Equation (2))—an emerging tension recently reviewed in detail by [52]. Stage-III surveys lack the statistical power to discriminate between the three possible explanations—*statistical fluctuation*, *residual analysis systematics*, and *new physics beyond ΛCDM*. Stage-IV surveys provide that statistical reach, but only if the accompanying systematics budget—shear calibration, blending, photometric redshifts, intrinsic alignments, baryonic feedback, and small-scale modeling—is brought under matching, sub-percent-level control.

This proposal is built precisely around that requirement: AnaCal delivers sub-percent shear-bias control on multi-resolution joint imaging, the joint photometric-redshift pipeline reduces catastrophic outliers by 30–40% at $z \gtrsim 1$, the analytical magnitude-response correction removes selection-induced multiplicative bias from tomographic binning, and the joint pixel-level shape measurement reduces per-galaxy shape noise on the LSST WL sample. Together with the LSST-DR1×Euclid-DR2 statistical reach, these systematic-control advances will deliver $\sigma(S_8) \lesssim 0.003$—a sub-percent measurement that is sufficient, by construction, to classify the S_8 tension into one of the three categories above.

(4) A validated joint shear catalog to improve the precision of LSST DESC strong-lensing H_0 measurements. The same joint LSST×Euclid shear catalog produced by this program will provide a foundational calibration dataset for the LSST DESC strong-lensing program. Time-delay cosmography from strong lenses already delivers competitive H_0 constraints at the $\sim$ 1–2% level [53, 54], but further improvement is limited by the *mass-sheet degeneracy* (MSD) [55]. Breaking this degeneracy requires anchoring the deflector mass profile with high-precision galaxy–galaxy weak-lensing measurements at sub-arcminute scales [54]. The sub-percent-calibrated joint shear catalog produced here will enable stacked galaxy–galaxy WL profiles around strong-lens deflectors with the precision needed to reduce the MSD contribution to the H_0 error budget. In this way, the proposed joint shear catalog directly improves the precision of LSST DESC time-delay cosmography, supporting a percent-level H_0 measurement that is independent of both the CMB and the local distance ladder.

Technical Merit

To deliver the scientific outputs above, the program develops four coupled technical contributions that together form an end-to-end, differentiable, analytically calibrated pipeline running from raw multi-survey pixels to LSST DESC-ready cosmology likelihoods, all at realistic computational and implementation cost.

- **A joint Rubin×Euclid image-to-shear pipeline with sub-percent bias control.** Building on AnaCal [56, 26, 27], the pipeline replaces image re-rendering in Metadetection with analytical, autodiff-based shear responses that propagate consistently through PSF homogenization across Rubin and Euclid pixel data, controlling the multiplicative shear bias to $|m| < 3 \times 10^{-3}$ at survey scale at a calibration cost of $\sim 10^4$ CPU-hours—a few hundred node-hours on NERSC Perlmutter—per full LSST coadd, roughly two orders of magnitude below the cost of the corresponding Metadetection-based calibration. By combining shape estimates of the same galaxies at different resolutions—the deep, multi-band ground-based LSST imaging and the high-resolution, stable-PSF Euclid imaging—the

joint pipeline reduces the effective per-galaxy shape noise of the LSST WL sample by 20–25% relative to LSST-only shape measurement [20, 21].

- **A drop-in interface for next-generation physics-informed AI shear estimators.** The pipeline's modular, differentiable structure makes it directly compatible with the D_4-equivariant-CNN AI shear estimators [44] and with response-guided denoising score-matching extensions that align the learned features with the analytical score function of the image likelihood [45]. These AI upgrades reduce the per-galaxy shape noise by an additional 10–20% on top of the perturbative AnaCal baseline. Combined with the joint-imaging gain above, the overall reduction in per-galaxy shape noise reaches at least 30%, equivalent to adding roughly 60% more LSST exposure time at zero additional observing cost.

- **A joint LSST+Euclid photometric-redshift pipeline.** Pixel-level forced photometry combining LSST *ugrizy* and Euclid NIR imaging delivers consistent *u*-through-*H* flux measurements for the same sources, avoiding the catalog-level matching ambiguities that limit existing multi-survey combinations. The resulting joint photometry, ingested into the LSST DESC RAIL/FlexZBoost [35, 34] photo-z estimators, achieves a 30–40% reduction in the catastrophic-outlier rate at $z \gtrsim 1$ [22, 23], narrows and de-biases the tomographic $n(z)$, and shares the same differentiable forward model as the shear pipeline, so $n(z)$ uncertainties propagate analytically into the w_0–w_a posterior. Crucially, the same differentiable framework also delivers analytical *magnitude* responses to shear, which are used inside the tomographic binning step to correct the selection-induced multiplicative bias that otherwise leaks into $n(z)$—to our knowledge, the first time the magnitude–shear response is propagated end-to-end in a Stage-IV WL pipeline, and an essential ingredient for keeping the total selection-induced bias sub-percent in joint Rubin+Euclid analyses.

- **End-to-end validation and reusable LSST DESC data products.** Validation of $|m| < 3 \times 10^{-3}$ at survey scale requires controlled simulations covering galaxy blending, anisotropic PSFs, chromatic effects [38, 39, 40], and detection systematics. `descwl-shear-sims`, co-developed by the PI, provides this testbed and is the established LSST DESC tool for shear-bias validation; the proposed pipeline will be exercised against it on a continuous-integration basis and wired into the LSST imSim framework [57, 58, 59] so that Rubin Operations and DESC validation runs share identical inputs. The resulting validated outputs—coadded, PSF-homogenized multi-survey image stacks of the Rubin DDFs and the overlapping Euclid and Roman deep fields, together with object catalogs, photo-z training sets, and shear responses—are packaged as a reusable LSST DESC data product, broadly leveraged across DESC working groups (cosmic shear, galaxy clusters, photo-z, strong lensing) and amplifying the scientific return of the methodological advances above well beyond the 3×2pt analyses that motivate them.

The proposed work will be carried out within the LSST DESC, primarily in the *Pixel-to-Object (PO)* and *Photometric Redshift (PZ)* working groups, with the resulting joint catalogs and simulations delivered to the large-scale-structure, cluster, and strong-lensing working groups for downstream cosmological analysis. The staged approach is built on PI-developed, validated foundations—AnaCal in HSC-Y6, LSST DP1, and `descwl-shear-sims`—reducing the technical risk of the proposal.

3.0 Proposed Research and Methods

The overarching methodology of this proposal is to extend AnaCal from its established single-survey form to an end-to-end, differentiable, joint ground–space pixel-level pipeline running from raw LSST and Euclid

images to LSST DESC-ready cosmology likelihoods. The pipeline is organized into four work packages, each a coupled pair of a methodological development and a validation campaign on `descwl-shear-sims`: (1) joint detection on the combined LSST–Euclid pixel data, (2) joint shape measurement with moments, model-fitting, and AI feature extractors, (3) joint flux measurement with analytical magnitude-response propagation to photo-z tomography, and (4) chromatic-PSF correction. For each work package we summarize the proposed method, validation strategy, and specifically argue why the approach is *appropriate* and *timely* for delivery on the LSST-DR1 / Euclid-DR2 timescale. For the downstream cosmology analysis built on the resulting joint shear and photo-z catalogs, we will use the LSST DESC analysis software stack directly: the joint shape and photo-z catalog products are formatted identically to those already delivered for LSST DP1 and DP2, so they plug into the LSST DESC `TXPipe` [60] two-point analysis pipeline and the `firecrown` [61] cosmology likelihood without any additional adaptation, and inherit the LSST DESC-validated treatment of intrinsic alignments, baryonic feedback, scale cuts, and covariances.

Work package #1: Joint detection from LSST and Euclid

Method. For each pair of overlapping fields we construct a log-likelihood detection image for both Euclid and LSST imaging—LSST oversampled onto the Euclid pixel grid—together with the corresponding analytical shear response of every pixel derived from the AnaCal pixel-level shear-response formalism. The two log-likelihood images are combined into a single joint log-likelihood map, and galaxy candidates are identified as peaks in this joint map. The detection-bias correction of AnaCal then propagates analytically through the joint detection because the shear response of each pixel in the joint map is known by construction.

Appropriate. This procedure combines the information from ground- and space-based imaging at different angular resolutions and wavelengths at the detection step itself, so that the resulting catalog already takes advantage of Euclid's space-based PSF and LSST's depth. The expected gain in the number of usable galaxies is therefore strictly larger than the catalog-level merging forecast in Schuhmann et al. [20], since catalog-level merging by construction throws away the joint information at the likelihood level. Validation is performed in `descwl-shear-sims` with controlled truth shear distortions, which allows us to verify directly that the analytical shear response of the detection weight is computed correctly.

Timely. This is a direct extension of the analytical detection-bias correction already validated for AnaCal at single-survey scale [26], and the implementation is a small delta on top of the existing AnaCal code base. The overlapping LSST and Euclid imaging needed to test the joint detection becomes available with Euclid-DR1 and the DES–Euclid-DR1 pathfinder fields beginning in FY26–FY27, exactly aligned with the start of this proposal.

Work package #2: Joint shape measurement

Method. For each detected galaxy candidate from work package #1, we perform shape measurement on the joint LSST–Euclid pixel data using three complementary estimators within the same AnaCal-calibrated framework: (a) Fourier moment-based shape estimators [56], (b) galaxy model fitting [33], and (c) physics-informed AI shape estimators based on D_4-equivariant CNNs [44] optionally augmented with denoising score-matching extensions [45]. In all three cases the analytical pixel-to-feature shear response is propagated by automatic differentiation, so the multiplicative-bias correction is shared across estimators.

Appropriate. AnaCal has been demonstrated to deliver sub-percent multiplicative-bias control on each of these three feature families [27, 33, 44], and the joint LSST–Euclid pixel-level measurement reduces the per-galaxy shape noise on the LSST WL sample relative to LSST-only shape measurement [20, 21]. We validate the measurement step in `descwl-shear-sims` with controlled truth shear distortions to verify directly that

the analytical shear response of the measured shape is computed correctly at the joint-imaging level.

Timely. The training and inference cost for the D_4-equivariant CNN feature extractor is small (see the FastLens narrative for details), so the AI estimator can be exercised at survey scale within the proposed compute budget. The moment-based and model-fitting estimators are already deployed at survey scale in the HSC-Y6 and LSST DP1/DP2 AnaCal analyses, providing a validated baseline on which the AI extension is built.

Work package #3: Joint flux measurement and photometric redshift

Method. For each joint-detected candidate we measure fluxes in all LSST *ugrizy* and Euclid VIS + NISP bands using a fixed-kernel aperture after PSF homogenization, and propagate the same analytical pixel-level shear response forward to obtain the shear response of the measured fluxes. These fluxes feed the LSST DESC RAIL/FlexZBoost [35, 34] photometric-redshift estimator, and the analytical flux-shear response is propagated through the photo-z estimator to deliver the analytical shear response of the photometric redshift itself.

Appropriate. Propagating the shear response analytically all the way to the photo-z estimate is the missing ingredient that allows the magnitude-cut selection bias arising at tomographic redshift binning to be corrected in closed form, a step that has not previously been implemented end-to-end in any Stage-IV WL pipeline. In addition, the addition of Euclid near-infrared photometry to LSST *ugrizy* breaks the 4000 Å/Lyman-break degeneracy at $z \gtrsim 1$ and reduces the photometric-redshift catastrophic-outlier rate by 30–40% [22, 23]. Validation uses `descwl-shear-sims` with controlled truth shear distortions to verify the analytical shear response of measured fluxes at the multi-band, multi-resolution joint level.

Timely. The RAIL-based photo-z pipeline has already been validated by the LSST DESC photometric-redshift team in HSC and on LSST DP1, providing a turnkey backend onto which the joint LSST–Euclid photometry produced here can be ingested.

Work package #4: Chromatic-PSF correction for joint Rubin–Euclid–Roman analysis

Method. The Euclid VIS PSF varies coherently with galaxy spectral energy distribution across the broad VIS bandpass, and the Euclid NISP $Y/J/H$ PSFs share a smaller but still significant chromatic component [39, 40]. We will use the overlapping multi-band LSST imaging to constrain the color-dependent PSF correction in the Euclid VIS band, and use overlapping Roman observations—which provide finer near-infrared wavelength sampling—to refine the Euclid NISP chromatic-PSF corrections in the Y, J, and H bands. The corrected PSFs are then fed forward through the AnaCal pixel-level shear-response chain established in work packages #1–#3.

Appropriate. Atmospheric and instrumental chromatic-PSF effects are known to produce shear biases at or above the LSST DESC sub-percent threshold if uncorrected [38], and a recent optimization of the Roman survey strategy has shown that mitigating these effects requires precisely the multi-band, multi-survey overlap exploited here [40]. The combination of LSST *ugrizy* for the optical bands and Roman for the NIR bands provides the densest wavelength sampling currently feasible for an external chromatic calibration of Euclid PSFs.

Timely. The required overlap between LSST-DR1, Euclid-DR2, and Roman wide-field imaging becomes available in the FY28–FY29 window covered by Year 3 of this proposal, which is the natural moment to deploy the chromatic-PSF correction at LSST-Y1 systematics-budget precision.

4.0 Timetable of Activities

The proposed program is organized into a five-year plan that is tightly synchronized with the Stage-IV imaging-survey data-release schedule: Euclid-DR1 (21 October 2026), Roman launch (by May 2027), LSST-DR1 (June 2028), Euclid-DR2 (March 2029), and LSST-DR2 (anticipated 2030). Each year couples a methodology deliverable to a real-data deliverable, with built-in risk-buffer fallbacks that keep the program productive if any single survey schedule slips. **Each year requires 65% FTE of the PI and 100% FTE of one postdoctoral researcher.**

Pathfinder analyses on Euclid-DR1×DES

The LSST-DP2 footprint does not provide sufficient overlap with Euclid to support a joint shear analysis, so no intermediate LSST-DP2 deployment is attempted; the validated joint pipeline is instead brought into production on LSST-DR1 (released June 2028). Before LSST-DR1, the $\sim$ 1,000 deg^2 overlap between Euclid-DR1 and DES—comparable in depth to LSST-Y1—provides the natural pathfinder for joint ground–space shear measurement, and its analysis defines the first two and a half years of the program.

Year 1 (FY26–FY27): Develop and Test Joint Image-processing Pipeline. Extend the AnaCal framework [27] from single-survey to multi-resolution, multi-survey joint imaging by performing joint source detection and correctly propagate shear response of the detection weight. In addition to moment-based and model-fitting shape estimators, the pipeline will support physics-informed, D_4-equivariant machine-learned feature extractors. Build the overlapped ground–space image-simulation suite within `descwl-shear-sims`, with controlled truth shear, realistic PSFs and pixel scales, and survey-specific noise.

Deliverables: a public release of the joint-imaging AnaCal pipeline; the overlapped ground–space image-simulation suite, used to quantify the magnitude measurement error and to calibrate and validate the shear pipeline, with the simulation code released alongside; and a journal paper documenting the pipeline and its simulation tests.

Year 2 (FY27–FY28): From images to catalogs. Apply the pipeline to the Euclid-DR1×DES overlap. Carry out null tests of PSF modeling, shear estimation, and star–galaxy shape correlation, and compare WL cluster mass estimates against the DES Y6 baseline. Calibrate the joint photometric-redshift estimates against the rich spectroscopic sample in the Extended Chandra Deep Field–South (ECDFS), processed end-to-end in the LSST DESC `RAIL` [35] framework, to anchor the joint $n(z)$ tomography with a spectroscopic-quality redshift truth set.

Deliverables: a public DES×Euclid-DR1 joint shear catalog with joint photometric-redshift estimates covering the overlaped $\sim$ 1,000 deg^2; one journal paper presenting the shear catalog and its null tests, and a companion paper presenting the photo-z catalog with its standard validation.

Year 3 (FY28–FY29): Cosmology from the pathfinder catalogs. Complete the Euclid-DR1×DES pathfinder program with a full cosmic-shear cosmology analysis on the joint shear and photo-z catalogs above. The two-point statistics, scale cuts, blinding, and covariances will be computed within the LSST DESC `TXPipe` [60] two-point analysis framework, and the cosmological inference will be performed using the LSST DESC `firecrown` [61] likelihood framework, which provides a self-consistent treatment of the WL, intrinsic-alignment, and photo-z-error models within `cobaya` and `cosmosis` samplers. The result will be benchmarked against the DES Y6 cosmic-shear baseline.

Deliverables: cosmological parameter posteriors on cosmological parameters from the Euclid-DR1×DES joint analysis, and a journal paper reporting the cosmic-shear cosmology analysis.

Joint LSST-DR1×Euclid-DR2 analysis

The LSST-DR1×Euclid overlap exceeds 1,000 deg^2 once LSST-DR1 (June 2028) is released, and grows beyond 5,000 deg^2 once Euclid-DR2 (March 2029) is available. The remainder of the program runs the validated pipeline on this overlap fully coordinated with LSST DESC's DR1 analysis.

Year 3 (FY28–FY29): Joint shear estimation with LSST-DR1 and Euclid-DR2. The larger survey overlap and the more stringent LSST-Y1 systematic-error requirements make this the stage at which chromatic-PSF effects [38, 39, 40] must be treated in detail. We will use the overlapping multi-band LSST imaging to constrain color-dependent PSF corrections for the Euclid VIS band, and will use overlapping Roman observations, with their finer near-infrared wavelength sampling, to refine the Euclid NISP chromatic-PSF corrections in the Y, J, and H bands. We will then perform dedicated PSF-systematics tests to quantify the improvement in PSF modeling and its impact on shear calibration.

Deliverables: updated chromatic-PSF models for Euclid VIS and NISP; a paper presenting the chromatic-PSF model, the associated PSF-systematics tests, and the resulting impact on weak-lensing shear calibration.

Year 4 (FY29–FY30): LSST-Y1-scale joint shear catalog and image-simulation calibration. Run the full joint pipeline on the LSST-DR1×Euclid-DR2 overlap as Euclid-DR2 becomes available (March 2029), perform null tests and cross-survey validation, and quantify residual shear and selection biases. In parallel, build the image-simulation suite at LSST-Y1 scale on top of the LSST imSim infrastructure [57, 58, 59], and use it to validate and calibrate the joint shear catalog against the LSST DESC SRD shear-bias requirement.

Deliverables: a public LSST-Y1 joint image-simulation suite; a calibrated LSST-DR1×Euclid-DR2 joint shear catalog; one journal paper documenting the image-simulation tests and shear-bias calibration; and a companion paper presenting the joint shear catalog and its null-test validation.

Year 5 (FY30–FY31): Joint LSST×Euclid 3×2pt dark-energy analysis. Use the validated joint catalog to perform a 3×2pt cosmological analysis on the LSST-DR1×Euclid-DR2 overlap, delivering the first joint-imaging WL constraint on w_0–w_a. Combine with DESI BAO and *Planck* CMB priors to deliver an independent Stage-IV cross-check of the DESI dynamical-dark-energy hint, and establish the joint-pipeline foundation for the LSST-DR2 era by delivering tools and catalogs to the LSST DESC large-scale-structure (LSS) and cluster working groups for downstream WL science.

Deliverables: cosmological parameter posteriors from the joint-imaging WL 3×2pt analysis; a peer-reviewed publication of the dark-energy result; and a DESC-supported public release of the joint LSST×Euclid analysis pipeline and simulation infrastructure.

Each year couples a methodological deliverable (pipeline, simulations, photo-z, calibration) to a data-release deliverable, and each is risk-buffered by an analysis-ready fallback dataset (HSC PDR3, DES Y6, the Rubin Deep Drilling Fields, and the EDF-N/EDF-S deep fields), so the program produces high-impact LSST DESC science even if any single survey schedule slips.

5.0 Competency of Applicant's Personnel and Adequacy of Proposed Resources

PI track record. The PI developed the AnaCal analytical shear-calibration framework from first principles—deriving the analytical pixel-level shear response, the analytical detection-bias correction [26], and the analytical noise-bias correction [27] that together constitute the AnaCal formalism—and wrote the entire AnaCal code base from scratch [33]. Building on this foundation, the PI has consistently delivered survey-scale weak-lensing shear catalogs and cosmic-shear analyses for the most demanding cosmology surveys of the past decade: led the HSC Y3 image-simulation campaign and shear catalog production [62] and the subsequent HSC Y3 cosmic-shear analysis [63], developed the HSC Y6 image-simulation suite for shear

calibration, is leading the ongoing HSC Y6 shear catalog, produced the LSST Commissioning Camera Data Preview 1 AnaCal shear catalog [37], and is currently leading the LSST Data Preview 2 AnaCal shear catalog. The PI is also supervising graduate students applying AnaCal to calibrate AI-based shear estimators [44, 45], providing direct training-the-trainer experience that maps onto the AI-extension thrust of this proposal.

Host institution. The PI is employed at Brookhaven National Laboratory (BNL), which regularly attracts top-tier postdoctoral scholars and graduate students to its Cosmology and Astrophysics group. BNL has a long tradition of leadership in weak-lensing shear measurement, catalog production, and cosmological analysis, dating back to the founding contributions of Erin Sheldon to the SDSS, DES, and LSST WL programs and continued through the PI's work on the HSC and LSST WL pipelines—providing exactly the institutional expertise required to deliver the joint Rubin×Euclid pipeline proposed here.

Computational resources. The proposal also relies on world-class computational facilities through BNL's Scientific Data and Computing Center (SDCC), a DOE-recognized HPC and large-scale data-handling center supporting the LSST, ATLAS, DUNE, and DESI experiments, among others. SDCC provides petabyte-scale storage, GPU-equipped HPC nodes, and a direct, low-latency connection to NERSC and to the LSST/Rubin Data Management infrastructure. These resources are essential for the overlapped ground–space image simulations, joint-survey image processing, and end-to-end shear-pipeline validation that underpin the proposed work, and are already in routine use by the PI's group for LSST DP1/DP2 AnaCal shear analyses.

6.0 Potential For Leadership Within the Scientific Community

The PI is a recognized leader in weak-lensing shear measurement and catalog production for the leading optical imaging surveys of the past decade. The PI developed the AnaCal analytical shear-calibration framework from first principles—using a perturbative pixel-level formulation to derive the analytical shear response of the full image processing chain and validating it against image simulations at the LSST DESC sub-percent multiplicative-bias requirement.

Beyond methodology, the PI has consistently led the survey-scale catalog production efforts that turn shear estimators into cosmology deliverables: the PI led the HSC Year-3 image-simulation campaign and shear catalog [62] and the corresponding HSC Y3 cosmic-shear analysis [63], produced the LSST Commissioning Camera Data Preview 1 AnaCal shear catalog [37], is currently leading the LSST DP2 AnaCal shear catalog, and is the convener of the HSC weak-lensing working group. The PI is also co-supervising graduate students applying AnaCal to AI-based shear estimators [44, 45], positioning the analytical-calibration framework for the AI-driven era of WL methodology that is now rapidly emerging.

This proposal will consolidate and extend that leadership position along three axes. First, by bringing AnaCal into joint LSST×Euclid pixel-level processing, the program positions the PI as the methodological bridge between the U.S. DOE/NSF Rubin LSST effort and the ESA Euclid and NASA Roman missions, ensuring that the LSST-Y1 shear-and-photo-z cosmology analysis benefits directly from AnaCal-anchored systematic control. Second, by extending AnaCal to AI-based shear estimators within the LSST DESC analysis path, the program keeps the PI at the front of WL methodology through the next decade, in which physics-informed AI is expected to become the dominant analysis paradigm. Third, the proposal further strengthens BNL's role within LSST DESC by anchoring its WL pipeline contributions in the joint-imaging pixel-level layer, complementing BNL's long-standing leadership in shear-measurement methodology established by Erin Sheldon and ensuring that BNL realizes a strong return on its sustained investment in Rubin Observatory cosmology infrastructure.

APPENDIX 1: Bibliography & References Cited

References

[1] Michel Chevallier and David Polarski. Accelerating Universes with Scaling Dark Matter. *Int. J. Mod. Phys. D*, 10(2):213–223, 2001.

[2] Ž. Ivezić, S. M. Kahn, J. A. Tyson, et al. LSST: From Science Drivers to Reference Design and Anticipated Data Products. *ApJ*, 873(2):111, 2019.

[3] DESI Collaboration. The DESI Experiment Part I: Science, Targeting, and Survey Design. *arXiv e-prints*, 2016.

[4] Daniel J. Eisenstein, Idit Zehavi, David W. Hogg, et al. Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies. *ApJ*, 633(2):560–574, 2005.

[5] A. G. Adame et al. DESI 2024 VI: Cosmological Constraints from the Measurements of Baryon Acoustic Oscillations. *JCAP*, 2025.

[6] DESI Collaboration, M. Abdul-Karim, et al. DESI DR2 Results II: Measurements of Baryon Acoustic Oscillations and Cosmological Constraints. *Phys. Rev. D*, 112:083515, 2025.

[7] DESI Collaboration, K. Lodha, et al. Extended Dark Energy analysis using DESI DR2 BAO measurements. *Phys. Rev. D*, 112:083511, 2025.

[8] Judit Prat and David Bacon. Weak gravitational lensing. In *Encyclopedia of Astrophysics*, volume 5, page 508. Elsevier, 2026.

[9] Euclid Collaboration. Euclid: I. Overview of the Euclid mission. *A&A*, 2025.

[10] D. Spergel et al. Wide-Field InfrarRed Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA 2015 Report. *arXiv e-prints*, 2015.

[11] Matthias Bartelmann and Peter Schneider. Weak Gravitational Lensing. *Phys. Rep.*, 340(4-5):291–472, 2001.

[12] Henk Hoekstra and Bhuvnesh Jain. Weak Gravitational Lensing and Its Cosmological Applications. *Annu. Rev. Nucl. Part. Sci.*, 58:99–123, 2008.

[13] Martin Kilbinger. Cosmology with cosmic shear observations: a review. *Rep. Prog. Phys.*, 78(8):086901, 2015.

[14] Rachel Mandelbaum. Weak Lensing for Precision Cosmology. *Annu. Rev. Astron. Astrophys.*, 56:393–433, 2018.

[15] The LSST Dark Energy Science Collaboration, R. Mandelbaum, T. Eifler, et al. The LSST Dark Energy Science Collaboration (DESC) Science Requirements Document. *arXiv e-prints*, 2018.

[16] The Dark Energy Survey Collaboration. The Dark Energy Survey: more than dark energy – an overview. *MNRAS*, 460(2):1270–1299, 2016.

[17] H. Aihara, N. Arimoto, R. Armstrong, et al. The Hyper Suprime-Cam SSP Survey: Overview and Survey Design. *PASJ*, 70:S4, 2018.

[18] J. T. A. de Jong, G. A. Verdoes Kleijn, K. H. Kuijken, and E. A. Valentijn. The Kilo-Degree Survey. *Experimental Astronomy*, 35:25–44, 2013.

[19] Planck Collaboration, N. Aghanim, Y. Akrami, et al. Planck 2018 results. VI. Cosmological parameters. *A&A*, 641:A6, 2020.

[20] R. L. Schuhmann, C. Heymans, and J. Zuntz. Galaxy shape measurement synergies between LSST and Euclid. *arXiv e-prints*, 2019.

[21] R. Chary et al. Joint Survey Processing of Euclid, Rubin and Roman: Final Report. *arXiv e-prints*, 2020.

[22] Melissa L. Graham, Andrew J. Connolly, Željko Ivezić, Samuel J. Schmidt, R. Lynne Jones, et al. Photometric Redshifts with the LSST. II. The Impact of Near-infrared and *u*-band Photometry. *AJ*, 159(6):258, 2020.

[23] B. Jain, D. Spergel, R. Bean, A. Connolly, et al. The Whole is Greater than the Sum of the Parts: Optimizing the Joint Science Return from LSST, Euclid and WFIRST. *arXiv e-prints*, 2015.

[24] P. Capak et al. Enhancing LSST Science with Euclid Synergy (Astro2020 Science White Paper). *Bull. AAS*, 51(7):199, 2019.

[25] A. Merloni, G. Lamer, T. Liu, M. E. Ramos-Ceja, et al. The SRG/eROSITA all-sky survey: First X-ray catalogues and data release of the western Galactic hemisphere. *A&A*, 682:A34, 2024.

[26] Xiangchong Li and Rachel Mandelbaum. Analytical weak-lensing shear responses of galaxy properties and galaxy detection. *MNRAS*, 521(4):4904–4926, 2023.

[27] Xiangchong Li, Rachel Mandelbaum, and The LSST Dark Energy Science Collaboration. Analytical noise bias correction for precise weak lensing shear inference. *MNRAS*, 536(4):3663–3676, 2025.

[28] Richard Massey, Catherine Heymans, Joel Bergé, Gary Bernstein, Sarah Bridle, et al. The Shear Testing Programme 2: Factors affecting high-precision weak-lensing analyses. MNRAS, 376(1):13–38, March 2007.

[29] Dragan Huterer, Masahiro Takada, Gary Bernstein, and Bhuvnesh Jain. Systematic errors in future weak-lensing surveys: requirements and prospects for self-calibration. MNRAS, 366(1):101–114, February 2006.

[30] Eric Huff and Rachel Mandelbaum. Metacalibration: Direct Self-Calibration of Biases in Shear Measurement. *arXiv e-prints*, 2017.

[31] E. S. Sheldon and E. M. Huff. Practical Weak-lensing Shear Measurement with Metacalibration. *ApJ*, 841:24, 2017.

[32] E. S. Sheldon, M. R. Becker, N. MacCrann, and M. Jarvis. Metadetection: Calibrated Shear in the Presence of Object Detection. *ApJ*, 902:138, 2020.

[33] Xiangchong Li. Analytical weak-lensing shear response of galaxy model fitting. *arXiv e-prints*, page arXiv:2506.16607, June 2025.

[34] Rafael Izbicki and Ann B. Lee. Converting high-dimensional regression to high-dimensional conditional density estimation. *Electronic Journal of Statistics*, 11(2):2800–2831, 2017.

[35] The RAIL Team et al. Redshift Assessment Infrastructure Layers (RAIL): Rubin-era photometric redshift stress-testing and at-scale production. *arXiv e-prints*, 2025.

[36] Shurui Lin, Xiangchong Li, Ji Li, Shengcao Cao, Xin Liu, et al. D_4CNN$\times$AnaCal: Physics-Informed Machine Learning for Accurate and Precise Weak Lensing Shear Estimation. *arXiv e-prints*, page arXiv:2603.19046, March 2026.

[37] Xiangchong Li et al. AnaCal Shear Profile of Abell 360 in LSST ComCam Data Preview 1. SITCOMTN 164, LSST Project / Vera C. Rubin Observatory, 2025.

[38] Joshua E. Meyers and Patricia R. Burchat. Impact of Atmospheric Chromatic Effects on Weak Lensing Measurements. ApJ, 807(2):182, July 2015.

[39] Federico Berlfein, Rachel Mandelbaum, Xiangchong Li, et al. Chromatic Effects on the PSF and Shear Measurement for the Roman Space Telescope High-Latitude Wide Area Survey. *MNRAS*, 542(2):608–628, 2025.

[40] Federico Berlfein, Rachel Mandelbaum, Jiachuan Xu, and Tianqing Zhang. Optimizing the Roman Space Telescope High-Latitude Wide Area Survey for mitigating chromatic PSF effects on shear measurement. *arXiv e-prints*, page arXiv:2603.15763, March 2026.

[41] M. Tewes, T. Kuntzer, R. Nakajima, F. Courbin, H. Hildebrandt, et al. Weak-lensing shear measurement with machine learning. Teaching artificial neural networks about feature noise. *A&A*, 621:A36, 2019.

[42] Zekang Zhang, Huanyuan Shan, Nan Li, et al. FORKLENS: Accurate weak-lensing shear measurement with deep learning. *A&A*, 683:A209, 2024.

[43] Taco S. Cohen and Max Welling. Group Equivariant Convolutional Networks. *arXiv e-prints*, 2016.

[44] Shurui Lin, Xiangchong Li, Ji Li, Shengcao Cao, Xin Liu, et al. D_4CNN$\times$AnaCal: Physics-Informed Machine Learning for Accurate and Precise Weak Lensing Shear Estimation. *arXiv e-prints*, 2026.

[45] Shurui Lin, Xiangchong Li, Xin Liu, et al. D_4CNN$\times$AnaCal II: Shear Response as a Score Inner Product: A Unified Statistical Framework for Weak Gravitational Lensing Shear Estimation. *in preparation*, 2026.

[46] Colin J. Burke, Patrick D. Aleo, Yu-Ching Chen, Xin Liu, John R. Peterson, et al. Deblending and classifying astronomical sources with Mask R-CNN deep learning. *MNRAS*, 490(3):3952–3965, 2019.

[47] Grant Merz, Yichen Liu, Colin J. Burke, Patrick D. Aleo, Xin Liu, et al. Detection, instance segmentation, and classification for astronomical surveys with deep learning (DeepDISC): DETECTRON2 implementation and demonstration with HSC. *MNRAS*, 526:1122–1137, 2023.

[48] Grant Merz, Xin Liu, Samuel Schmidt, Alex I. Malz, Tianqing Zhang, et al. DeepDISC-photoz: Deep learning-based photometric redshift estimation for Rubin LSST. *arXiv e-prints*, 2025.

[49] The AION Collaboration et al. AION-1: A Multi-Survey Foundation Model for Astronomical Imaging. *arXiv e-prints*, 2025.

[50] Alan Junzhe Zhou, Xiangchong Li, Scott Dodelson, and Rachel Mandelbaum. Accurate field-level weak lensing inference for precision cosmology. *Physical Review D*, 110(2):023539, July 2024.

[51] A. Loureiro, L. Whiteway, E. Sellentin, J. S. Lafaurie, A. H. Jaffe, and A. F. Heavens. Almanac: Weak Lensing power spectra and map inference on the masked sphere. *The Open Journal of Astrophysics*, 6:6, February 2023.

[52] Ioannis Pantos and Leandros Perivolaropoulos. Status of the S_8 tension: A 2026 review of probe discrepancies. *Phys. Dark Universe*, 52:102286, June 2026.

[53] Kenneth C. Wong, Sherry H. Suyu, Geoff C.-F. Chen, Cristian E. Rusu, Martin Millon, et al. H0LiCOW - XIII. A 2.4 per cent measurement of H_0 from lensed quasars: 5.3σ tension between early- and late-Universe probes. MNRAS, 498(1):1420–1439, October 2020.

[54] S. Birrer, A. J. Shajib, A. Galan, M. Millon, T. Treu, et al. TDCOSMO. IV. Hierarchical time-delay cosmography - joint inference of the Hubble constant and galaxy density profiles. A&A, 643:A165, November 2020.

[55] E. E. Falco, M. V. Gorenstein, and I. I. Shapiro. On model-dependent bounds on H 0 from gravitational images : application to Q 0957+561 A, B. ApJ, 289:L1–L4, February 1985.

[56] Xiangchong Li, Nobuhiko Katayama, Masamune Oguri, and Surhud More. Fourier Power Function Shapelets (FPFS) shear estimator: performance on image simulations. *MNRAS*, 481(4):4445–4460, 2018.

[57] LSST Dark Energy Science Collaboration, Bela Abolfathi, et al. The LSST DESC DC2 Simulated Sky Survey. ApJS, 253(1):31, 2021.

[58] M. A. Troxel, C. Lin, A. Park, C. Hirata, R. Mandelbaum, et al. A Joint Roman Space Telescope and Rubin Observatory Synthetic Wide-Field Imaging Survey. MNRAS, 522(2):2801–2822, 2023.

[59] OpenUniverse, The LSST Dark Energy Science Collaboration, The Roman HLIS Project Infrastructure Team, The Roman RAPID Project Infrastructure Team, The Roman Supernova Cosmology Project Infrastructure Team, et al. OpenUniverse2024: A shared, simulated view of the sky for the next generation of cosmological surveys. MNRAS, 544(4):3799–3833, 2025.

[60] J. Prat, J. Zuntz, Y. Omori, C. Chang, T. Tröster, E. Pedersen, C. García-García, E. Phillips-Longley, J. Sanchez, D. Alonso, X. Fang, E. Gawiser, K. Heitmann, M. Ishak, M. Jarvis, E. Kovacs, P. Larsen, Y.-Y. Mao, L. Medina Varela, M. Paterno, S. D. Vitenti, Z. Zhang, and The LSST Dark Energy Science Collaboration. The catalog-to-cosmology framework for weak lensing and galaxy clustering for LSST. *The Open Journal of Astrophysics*, 6:13, May 2023.

[61] LSST DESC Firecrown Contributors. Firecrown: DESC Cosmology Likelihood Framework. https://github.com/LSSTDESC/firecrown, 2024. Python package for implementing LSST DESC likelihoods.

[62] Xiangchong Li, Hironao Miyatake, Wentao Luo, Surhud More, Masamune Oguri, et al. The three-year shear catalog of the Subaru Hyper Suprime-Cam SSP Survey. PASJ, 74(2):421–459, April 2022.

[63] Xiangchong Li, Tianqing Zhang, Sunao Sugiyama, Roohi Dalal, Ryo Terasawa, et al. Hyper Suprime-Cam Year 3 results: Cosmology from cosmic shear two-point correlation functions. Phys. Rev. D, 108(12):123518, December 2023.