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] 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 measuring the coherent distortions of galaxy images caused by foreground matter, WL is sensitive not only to geometry but also to the structure growth, providing exactly the information needed to deliver an independent test of the DESI-motivated hint of dynamical dark energy. Delivering this 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 the DOE's 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. Joint pixel-level detection increases the effective galaxy number density of the LSST WL sample by at least $\sim 50\%$ [10] (catalog-level forecast; image-level joint detection is expected to deliver larger gains [11]); AI-based shape estimation [12, 13] further reduces the per-galaxy shape noise by 10–20%; and Euclid near-infrared photometry reduces the catastrophic photo-z outlier rate by 25–40% at $z \gtrsim 1$ [14, 15]. Together, these advances conservatively deliver 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) [16]—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, together with the emergence of **field-level** weak-lensing inference that forward-models the survey footprint directly, further amplifies what a joint analysis of these surveys can deliver. Realizing this opportunity, however, hinges on a technical bottleneck in **shear calibration**: the complicated 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 and calibrate 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 along the line of sight [17, 18, 19, 20], and is measured statistically across very large galaxy populations because the per-galaxy signal is more than an order of magnitude smaller than the intrinsic shape variation. Dark energy affects WL both by altering the distance–redshift relation that sets lensing efficiency and by suppressing the late-time growth of cosmic structure that sets the amplitude of WL correlations, so WL surveys probe cosmic geometry and structure growth simultaneously, complementing the BAO distance ladder of DESI. Equally important, the dominant WL systematics—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 test the DESI-motivated hint of dynamical dark energy. The pioneering Stage-III imaging surveys—DES [21], HSC [22], and KiDS [23]—have established the modern cosmic-shear methodology and produced precise measurements of the structure-growth amplitude

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

finding values 0.5–2.5σ below the *Planck* CMB [24] value under ΛCDM. With substantially larger area, greater depth, improved redshift information, and tighter control of shear systematics, the Stage-IV imaging surveys—now operating or coming online over the next few years (Figure 1)—will transform this emerging "S_8 tension" into a high-precision measurement of structure growth and an independent test of the DESI-motivated hint of dynamical dark energy. Among the LSST cosmological probes, WL 3×2pt is forecast to deliver the tightest single-probe constraint on the dark-energy equation-of-state parameters w_0 and w_a (Figure 2).

The scientific case for joint LSST–Euclid weak-lensing analysis is well established. Catalog-level forecasts show that combining LSST and Euclid shape measurements of the same galaxies increases the effective galaxy number density $n_{\rm eff}$ of the LSST WL sample by $\sim 50\%$ [10], and pixel-level joint detection and shape measurement—the approach pursued here—is expected to deliver still larger gains, as identified by the 2020 Joint Survey Processing Study Group [11]. Physics-informed AI shape estimators (D_4-equivariant CNNs trained with denoising score matching) [12, 13] further reduce the per-galaxy shape noise by an additional 10–20% on top of this joint-detection $n_{\rm eff}$ gain. Joint processing also tightens the photometric redshifts that underpin tomographic WL: adding Euclid near-infrared (NIR) photometry to LSST *ugrizy* imaging reduces the catastrophic-outlier rate by 25–40%, especially at $z \gtrsim 1$, where 4000 Å-break/Lyman-break confusion makes optical colors degenerate [14, 15]. Both gains arise from the same complementarity—LSST contributes optical depth and multi-band sampling, Euclid contributes space-based resolution and

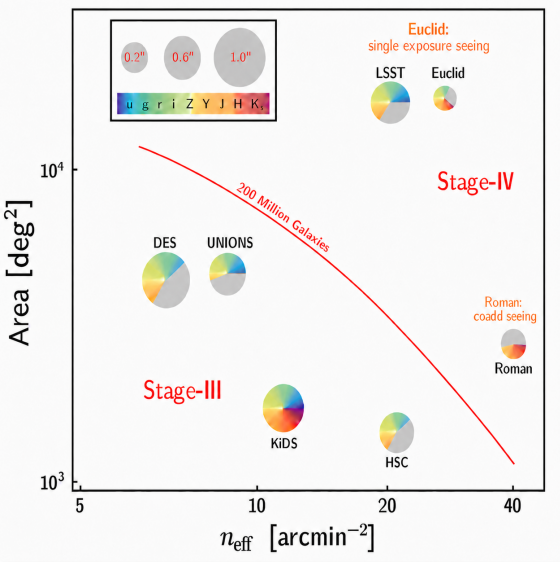

Figure 1: The Stage-III and Stage-IV imaging surveys. The diameter of each disk represents the size of the point-spread function (PSF) of the corresponding survey, and the color encodes its wavelength coverage from the ultraviolet (u band) to the near-infrared (K band).

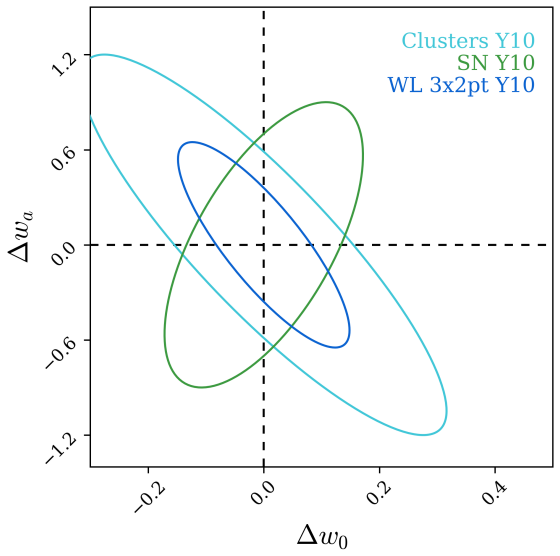

Figure 2: Forecast LSST ten-year 68% constraints on Δw_0, Δw_a from WL 3×2pt (blue), Type Ia SNe (green), and clusters (cyan). WL gives the tightest individual constraint and is the focus of this proposal. Adapted from Fig. G2 of the LSST DESC Science Requirements Document [25].

NIR coverage—and pixel-level forced photometry across the two surveys yields consistent u-through-H flux measurements for the same sources, avoiding catalog-matching ambiguities that limit existing multi-survey combinations and enabling consistent detection bias correction in shear estimation. Figure 3 illustrates this multi-band, multi-resolution complementarity in joint imaging of an eRASS1 galaxy cluster across Rubin, DES, and Euclid. These gains make pixel-level LSST–Euclid joint processing an immediate priority for Stage-IV dark-energy science, also recognized by the Astro2020 Decadal Survey white paper [26], and the pipeline developed here will extend naturally to Roman once Roman coadd imaging products are released.

A further driver is the recent emergence of *field-level* weak-lensing inference [28, 29], which forward-models the observed shear and density maps directly rather than compressing them into two-point summary statistics. Field-level analysis is particularly natural for joint LSST×Euclid data: traditional cosmic-shear analyses implicitly assume statistical homogeneity across the survey, but the LSST–Euclid overlap, LSST-only, and Euclid-only regions differ markedly in depth, PSF, wavelength coverage, selection, and photo-z performance. Forward-modeling these spatially varying survey properties lets the calibrated overlap anchor coherent inference across the larger non-overlapping footprints, extracting the maximum cosmological information from the full Stage-IV mosaic.

Realizing this opportunity, however, hinges on meeting the LSST DESC sub-percent shear-calibration requirement [25] at survey scale on heterogeneous survey data. Classic high-accuracy shear calibration approaches apply small artificial shears to the observed images and rerun detection, shape and flux measurement, photo-z estimation, and tomographic binning on each sheared realization to capture detection and selection biases at the sub-percent level. This calibration chain is technically intricate to validate and computationally demanding for joint-survey analysis, and remains the principal obstacle to turning the long-recognized promise of joint ground–space weak-lensing analysis into an operational survey-scale pipeline.

AnaCal: Robust and Efficient Self Calibration for Shear Estimation

This proposal develops a survey-scale joint-imaging shear estimator for combined ground- and space-based data, closing the gap identified above: the absence of a method that can meet the LSST DESC sub-percent shear-calibration requirement at affordable cost. The work extends the PI's **AnaCal** [30, 31] framework into a joint ground–space pipeline for Stage-IV weak lensing. AnaCal computes the shear response in closed form and propagates it by automatic differentiation, eliminating counterfactual image shearing and running two orders of magnitude faster than artificial-shearing methods. The same differentiable construction makes AnaCal AI-compatible by design, supporting model-fitting features and symmetry-equivariant neural feature extractors within a single self-calibrated pipeline.

Shear-estimation accuracy is summarized by $g_{\rm obs} = (1+m)g_{\rm true} + c$ [32, 33], where m is the multiplicative bias and c the additive bias; the LSST DESC SRD requires $|m| < 0.3\%$ across redshift bins [25]. *Perturbative* self-calibrating estimators (Metadetection and AnaCal) already achieve this target even on heavily blended images by expanding the measured ellipticity e to first order in the applied shear g and exploiting the spin-2 symmetry of weak lensing:

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

Under isotropically distributed galaxy orientations, both $\langle e_0 \rangle$ (spin-2) and the second-order contribution (spin-2 plus spin-6) vanish, leaving only the spin-0 mean shear response $\langle \partial e / \partial g \rangle$. **Metacalibration** [34, 35] and **Metadetection** [36] estimate this response *numerically*, by finite-differencing measured shapes across image realizations with small artificial shears applied. This approach delivers part-per-thousand

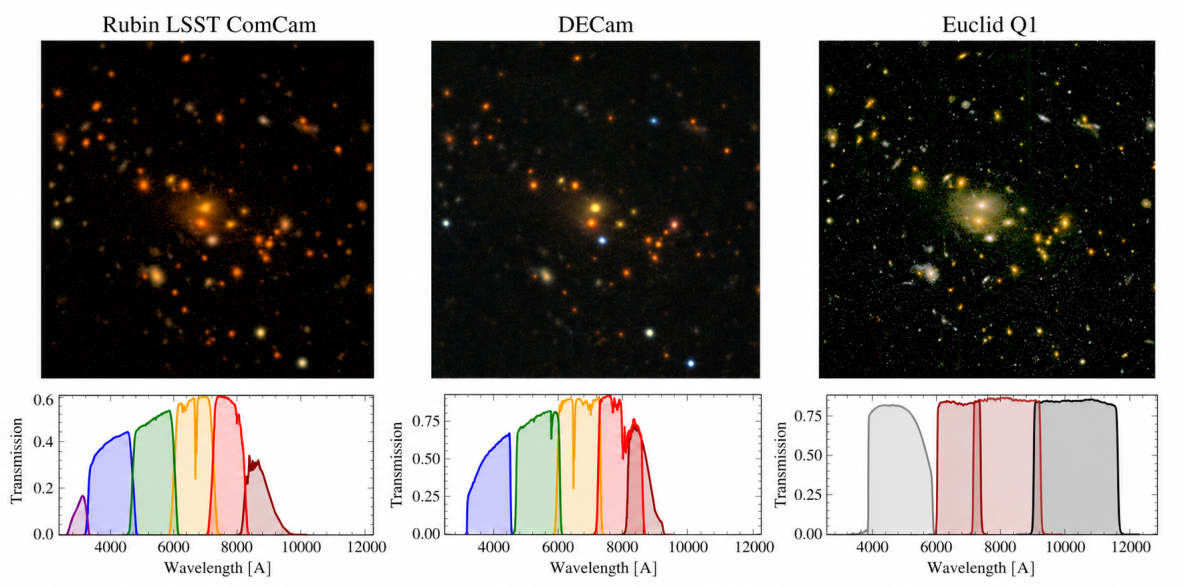

Figure 3: *Top:* Rubin Commissioning Camera DP1, DES, and Euclid Quick Release 1 imaging of the same eRASS1 galaxy cluster [27] ($z = 0.69$, EDF-S). *Bottom:* corresponding filter transmission curves: Rubin *ugrizy*, DES *grizy*, and Euclid Visible Imager (VIS) plus Near-Infrared Spectrometer and Photometer (NISP; Y, J, H). The three surveys differ in depth, resolution, and wavelength coverage; joint pixel-level analysis across them underpins this proposal.

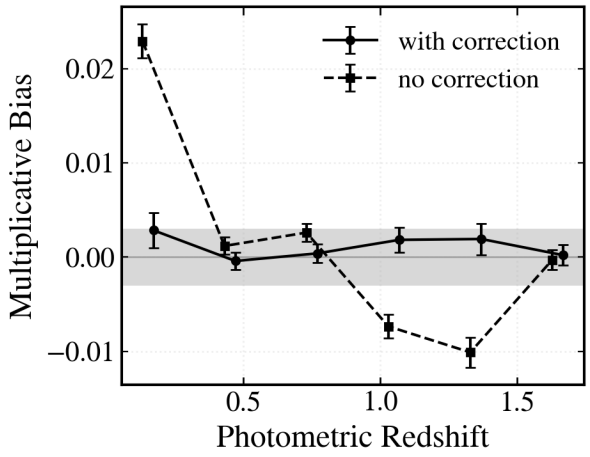

Figure 4: Multiplicative shear bias m as a function of `FlexZBoost` photometric redshift on LSST ten-year-like image simulations with blending. The *solid* line includes the analytical self-calibration for the selection bias induced by binning galaxies into tomographic redshift bins; the *dashed* line is without that correction. The *grey* band marks the LSST DESC requirement [25].

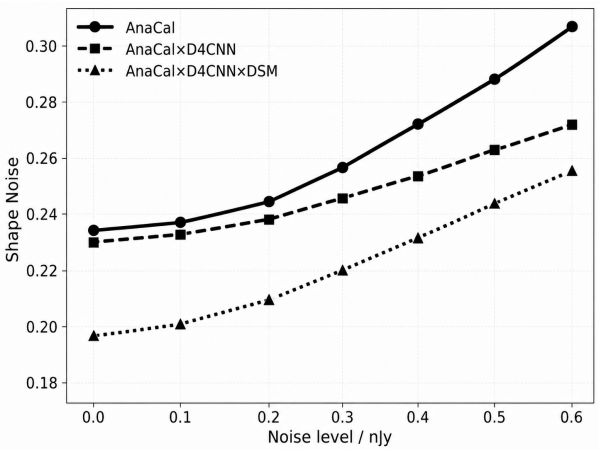

Figure 5: Per-component galaxy shape noise as a function of pixel-noise standard deviation $\sigma_{\rm noise}$ on LSST-like image simulations. *Solid:* AnaCal alone; *dashed:* AnaCal combined with a D_4-equivariant CNN [40]; *dotted:* AnaCal + D_4CNN further optimized with denoising-score matching. 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* unblinded and published yet); the corresponding cosmic-shear analysis is currently unpublished and blinded, and final published values may differ slightly from those quoted here.

accuracy and underpins the DES Y6 shear catalog ($|m| < 0.3\%$), but at the cost of additional computational complexity: four extra sheared image realizations must be generated and processed.

AnaCal, by contrast, computes the analytical shear response for each smoothed image pixel and propagates it forward through detection, deblending, and measurement to the final observables [37], while noise-induced bias is corrected analytically using the symmetry 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 bias terms (Equation (3)), and achieves $|m| < 0.3\%$ across all redshift bins. Figure 4 demonstrates this on LSST Y10 image simulations, with photo-z estimated using `FlexZBoost` [38] in the `RAIL` [39] framework: the solid curve includes AnaCal's analytical self-calibration for tomographic-binning selection effects (without external simulations) and lies safely within the LSST DESC SRD band.

Avoiding the artificial-shearing step makes AnaCal $\sim 100\times$ faster than Metadetection: processing the full LSST *ugrizy* coadd at survey scale requires only $\sim 10^4$ CPU hours (a few hundred node-hours on NERSC

Perlmutter). The same differentiable construction extends naturally to physics-informed AI shear estimation, propagating the shear response through machine-learned D_4-equivariant representations of galaxy images [40] while preserving analytical self-calibration.

AnaCal is already 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 Preview 1 (DP1) [41], is being extended to LSST Data Preview 2 (DP2), and has been validated for the chromatic-PSF effects [42] relevant to joint Rubin–Roman weak lensing [43, 44]. Table 1 compares the resulting HSC Y6 cosmic-shear analysis with DES Y6 in survey area, source density, redshift range, and S_8 uncertainty (Equation (2)).

Toward Physics-Informed AI Shear Estimation with AnaCal

The fast, AI-compatible AnaCal pipeline is the natural foundation on which to build the next generation of shear estimators. Deep neural networks learn flexible pixel-level representations relevant to shape measurement, deblending, and multi-band galaxy morphology, going beyond hand-crafted moment features. **Physics-informed AI shear estimators**—neural networks acting as the feature extractor while the shear response is analytically propagated and calibrated through AnaCal—are therefore a natural extension: they 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 ML shear estimators are not, on their own, calibratable to the LSST DESC sub-percent requirement [45, 46]: blindly trained networks introduce shape-dependent multiplicative bias at the 10^{-2} level, and simulation-based recalibration is prohibitive at survey scale. Two ingredients remove this obstacle:

(i) **AnaCal.** The analytical AnaCal calibration is applied to the network outputs by forward-propagating the analytical pixel-to-feature shear response through the trained network via automatic differentiation. This solves the shear-bias perturbation problem consistently through second order in g (Equation (3)).

(ii) **D_4-equivariant network.** The network [47] embeds the D_4 group of 90° rotations and reflections into the network weights, so the predicted ellipticity transforms exactly as a spin-2 quantity under image rotations. This eliminates orientation-dependent biases and cuts the parameter count by $\sim 8\times$, yielding a $\sim 10\times$ training-data-efficiency gain at no accuracy cost.

A first realization of this D_4CNN $\times$ AnaCal architecture [12], demonstrated on isolated galaxies in LSST-like single-band simulations, achieves $|m| < 0.1\%$ **across all tested noise levels, PSF sizes and ellipticities, and magnitude cuts (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, *denoising score matching* [13], trains the equivariant feature extractor to align with the analytical score function of the image likelihood (provably the minimum-variance unbiased shear estimator), reducing the per-component shape noise by another $\sim 10\%$ at LSST Y10 noise levels (Figure 5) while preserving $|m| < 0.1\%$ calibration.

Field-Level Inference for Joint Stage-IV Imaging

The second methodological pillar of this proposal is enabling *field-level* weak-lensing inference on joint LSST×Euclid data. Traditional two-point cosmic-shear analyses implicitly assume statistical homogeneity across the survey footprint, an assumption already violated within a single survey by depth and observing-condition variations [48], and broken much more severely when LSST-only, Euclid-only, and LSST–Euclid overlap regions are combined: these regions differ in depth, PSF, wavelength coverage, selection, and photo-z performance. Field-level inference forward-models these spatially varying survey properties directly, letting the calibrated overlap anchor coherent inference across the larger non-overlapping footprints [28, 29, 49, 50, 51].

Catalog-level cosmic-shear pipelines either restrict analyses to nominally homogeneous patches—losing the area-driven statistical power of the joint mosaic—or absorb survey-property variations into approximate mode-coupling matrices, leaving residual inhomogeneity systematics that threaten to become a dominant systematic floor at Stage-IV precision [48]. The field-level analysis pipeline `Miko`, developed by a student co-supervised by the PI [28], has demonstrated on HSC Year-3-like mocks that a forward-modeled, pixel-space treatment of analysis systematics (aliasing, boundary effects, mode coupling, density-induced shape noise) yields *unbiased* tomographic power-spectrum amplitudes when a Gaussian field-level prior is used, whereas a non-Gaussian log-normal prior introduces significant bias from prior misspecification. Building on this finding, the proposed program adopts the Gaussian-prior `Miko` backbone for unbiased field-level inference, which is the natural framework into which the LSST/Euclid survey inhomogeneities can be absorbed via the forward model. Extending this framework to the joint LSST×Euclid mosaic requires precisely the calibrated, co-registered, multi-resolution shear and photo-z catalogs developed here as its data layer: each sub-region (LSST-only and LSST-Euclid overlap) enters the field-level likelihood with its own shape noise, number density and photometric redshifts. By producing this calibrated anchor in the LSST–Euclid overlap, the proposed pipeline opens the door to field-level dark-energy analyses that coherently exploit the full Stage-IV footprint, rather than only its homogeneous subsets. Gradient-based posterior sampling over cosmological and nuisance parameters at field-level resolution is made tractable by coupling the pipeline to differentiable cosmological emulators: `CosmoPower`-like frameworks [52, 53] accelerate the mapping from cosmological parameters to linear and nonlinear power spectra inside the forward model, providing analytic gradients that integrate seamlessly with the AnaCal autodiff calibration layer. The same calibrated catalog will also feed downstream pipelines that target the *non-Gaussian* information in the cosmic shear field, which is inaccessible to two-point statistics: log-normal and N-body-prior field-level inference [49, 50, 51, 54], and simulation-based / deep-learning cosmology analyses based on higher-order summaries (scattering transforms, wavelet harmonics, mass-map moments, or learned features) [55, 56, 57]. This breaks the Gaussian-prior ceiling of two-point and Gaussian field-level analyses and broadens the scientific reach of the joint LSST×Euclid catalog beyond any single inference framework.

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 analysis 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 four coupled thrusts (see Section 3.0 for details): (i) joint pixel-level detection that raises the LSST WL effective galaxy number density; (ii) AI-based joint shear estimation that reduces the per-galaxy shape noise on top of the sub-percent-calibrated AnaCal baseline; (iii) joint photometric-redshift estimation that exploits LSST optical + Euclid NIR coverage; and (iv) Gaussian-prior field-level inference that forward-models

the spatially varying LSST×Euclid survey footprint. The resulting infrastructure is designed as a foundation for next-generation AI-based image-to-cosmology inference and is built to extend naturally to a joint LSST×Euclid×Roman analysis once Roman coadd data become available.

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 of the DESI-motivated hint of dynamical dark energy. The combined gain from joint $n_{\rm eff}$, AI shape noise, and joint photo-z improves the LSST WL constraint on w_0–w_a by at least $\sim 30\%$ relative to single-survey LSST, and a Gaussian-prior field-level analysis on the joint footprint provides a complementary, inhomogeneity-aware route to extract this signal. Together these advances accelerate the survey epoch at which Rubin LSST delivers an independent weak-lensing test of the 3–4σ hint for time-evolving $w(a)$ reported by DESI DR2 with external probes [7]. The key advance is independence: weak lensing probes geometry and late-time growth, with degeneracies and systematics distinct from BAO, supernovae, and CMB. A consistent result would provide independent evidence for dynamical dark energy; an inconsistent one would point to statistical fluctuation, probe-combination effects, or residual systematics in the current evidence.

(2) Sub-percent measurement of S_8 to resolve the S_8 tension. Stage-III cosmic-shear surveys (DES, HSC, KiDS) consistently report values of S_8 that lie 0.5–2.5σ below the *Planck* CMB value under ΛCDM (Equation (2)) [58], but lack the statistical reach to discriminate *statistical fluctuation*, *residual systematics*, and *new physics*. Stage-IV surveys provide that reach only if the accompanying systematics budget—shear calibration, cross-redshift blending, photo-z, intrinsic alignments, baryons, and small-scale modeling—is brought under matching sub-percent control. This proposal is built around that requirement: AnaCal delivers sub-percent shear bias on joint imaging, joint pixel-level detection plus AI shape estimation raise $n_{\rm eff}$ by $\sim 50\%$ and reduce per-galaxy shape noise by a further 10–20%, joint photo-z cuts catastrophic outliers by 30–40% at $z \gtrsim 1$, and the analytical magnitude-response correction removes selection-induced multiplicative bias from tomographic binning. Together with the LSST-DR1×Euclid-DR2 statistical reach, these advances deliver $\sigma(S_8) \lesssim 0.003$, sufficient by construction to classify the S_8 tension.

(3) A validated joint shear catalog for LSST DESC strong-lensing dark-energy and H_0 cosmology. The joint LSST×Euclid shear catalog produced by this program will anchor the LSST DESC strong-lensing program by tightening the constraint on the *mass-sheet degeneracy* (MSD) [59], which currently limits both time-delay H_0 measurements [60, 61] and tomographic strong-lensing constraints on dark energy [62]. Breaking the MSD requires extending the galaxy–galaxy WL shear profile around strong-lens deflectors down to small angular scales—a projected radius of ~ 0.04 Mpc, corresponding to only a few arcseconds at typical lens redshifts $z \sim 0.5$, just beyond the Einstein radius [63]. The proposed program is well matched to this requirement: joint LSST×Euclid pixel-level detection plus AI-based shape estimation together raise the small-scale signal-to-noise on the deflector shear profile by combining a higher source density with a lower per-galaxy shape noise. The resulting tightening of the stacked galaxy–galaxy WL constraint on the MSD reaches $\sim 30\%$ relative to LSST-only WL [63], supporting a percent-level H_0 from time-delay cosmography and a tomographic strong-lensing measurement of $w(z)$ to better than 0.1 out to $z \sim 3$, both independent of the CMB and the local distance ladder [62].

(4) A late-time growth anchor for the cosmological neutrino-mass measurement. The proposed joint LSST×Euclid WL constraint also provides a precise, low-redshift growth anchor for the cosmological sum of neutrino masses $\sum m_\nu$, which suppresses small-scale structure growth in a scale- and redshift-dependent

way. Joint analyses of LSST WL+clustering with future CMB-Stage-IV data forecast $\sigma(\sum m_\nu) \sim 30$ meV even in beyond-ΛCDM extensions that vary curvature and the dark-energy equation of state [64], breaking the degeneracies that limit CMB-only constraints and supporting a robust 3σ detection of the minimal ~ 60 meV normal-hierarchy mass once improved CMB optical-depth information is available. The sub-percent-calibrated, joint-imaging shear catalog produced by this program directly enables this neutrino science by ensuring that the neutrino-mass signal in the late-time growth measurement is not absorbed by shear-calibration, photo-z, or survey-inhomogeneity nuisance parameters.

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.

(1) A joint Rubin×Euclid image-to-shear pipeline with sub-percent bias control and a drop-in interface for AI shear estimators. Building on AnaCal [65, 30, 31], the pipeline replaces image re-rendering in Metadetection with analytical, autodiff-based shear responses propagated consistently across Rubin and Euclid pixel data, controlling $|m| < 0.3\%$ at survey scale at $\sim 100\times$ lower calibration cost than Metadetection. Joint pixel-level detection across the deep multi-band Rubin and the high-resolution stable-PSF Euclid imaging raises the LSST WL $n_{\rm eff}$ by at least $\sim$50% at the catalog level [10], with image-level processing expected to deliver more [11]. The pipeline's modular, differentiable structure also serves as a drop-in backbone for AI shear estimators: the D_4-equivariant-CNN [12] and denoising-score-matching extensions [13] inherit AnaCal's analytical calibration through automatic differentiation and cut shape noise by another 10–20%, and the same interface extends to **DeepDISC** [66, 67, 68] and to foundation models such as **AION-1** [69]. Combined with the photo-z improvement from Euclid NIR, this yields $>30\%$ improvement in the LSST WL constraint on dark energy.

(2) A joint LSST+Euclid photometric-redshift pipeline. Pixel-level forced photometry on LSST *ugrizy* + Euclid NIR delivers consistent *u*-through-*H* fluxes for the same sources, avoiding catalog matching ambiguities. Ingested into the LSST DESC RAIL [39, 38] estimators, the joint photometry cuts the catastrophic-outlier rate by 30–40% at $z \gtrsim 1$ [14, 15] and de-biases the tomographic $n(z)$. The same differentiable framework also propagates analytical *magnitude* responses to shear through the tomographic binning step, correcting the selection-induced multiplicative bias that leaks into $n(z)$—to our knowledge the first end-to-end propagation of the magnitude–shear response in a Stage-IV WL pipeline, and essential for keeping the total selection-induced bias sub-percent.

(3) End-to-end validation and reusable LSST DESC data products. $|m| < 0.3\%$ validation at survey scale requires controlled simulations of blending, anisotropic and chromatic PSFs [42, 43, 44], and detection systematics. The pipeline will be exercised on a continuous-integration basis against the PI-codeveloped `descwl-shear-sims` (the established LSST DESC shear-bias-validation tool) wired into the LSST `imSim` [70, 71, 72] framework, sharing identical inputs with Rubin Operations and DESC validation runs. The validated outputs—coadded multi-survey image stacks of the Rubin DDFs and overlapping Euclid/Roman deep fields, plus object catalogs, photo-z training sets, and shear responses—are packaged as a reusable LSST DESC data product leveraged by the cosmic-shear, cluster, photo-z, and strong-lensing working groups.

(4) A differentiable, emulator-accelerated field-level inference backbone. The calibrated, co-registered shear and photo-z catalog feeds the Gaussian-prior `Miko` [28] engine, with each LSST-only / Euclid-only / overlap region entering the forward model with its own number density, shape noise, and $n(z)$ so that survey inhomogeneity is absorbed by construction. Gradient-based posterior sampling at field-level resolution is

made tractable by coupling `Miko` to the differentiable `CosmoPower` [52, 53] emulators, which provide analytic gradients of the linear and nonlinear power spectra and integrate seamlessly with the AnaCal autodiff calibration layer—yielding an end-to-end-differentiable stack from raw pixels to the cosmological posterior.

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, and LSST DP1; and PI-codeveloped `descwl-shear-sims`—reducing the technical risk of the proposal.

3.0 Proposed Research and Methods

The proposed methodology extends AnaCal into an end-to-end, differentiable joint LSST×Euclid pipeline organized into four coupled thrusts: (#1) joint detection, (#2) joint shape measurement, (#3) joint flux and photo-z, and (#4) field-level cosmology analysis on the resulting calibrated catalog. For each thrust we describe the method, validation on `descwl-shear-sims/imSim` [70, 71, 72], and why it is appropriate and timely for the LSST-DR1 / Euclid-DR2 schedule.

Thrust #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. [10], since catalog-level merging by construction throws away the joint information at the likelihood level. Validation is performed in `descwl-shear-sims/imSim` [70, 71, 72] 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 [30], 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.

Thrust #2: Joint shape measurement

Method. For each detected galaxy candidate from thrust #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 [65], (b) galaxy model fitting [37], and (c) physics-informed AI shape estimators based on D_4-equivariant CNNs [12] optionally augmented with denoising score-matching extensions [13]. 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. The

same shear-response chain ingests *chromatic-PSF corrections* for the joint Rubin–Euclid–Roman analysis: the Euclid VIS PSF varies coherently with galaxy SED across the broad VIS bandpass, and the NISP $Y/J/H$ PSFs share a smaller but still significant chromatic component [43, 44]. We use the overlapping LSST *ugrizy* imaging to constrain the color-dependent VIS-band PSF correction, and overlapping Roman observations—which provide finer NIR wavelength sampling—to refine the NISP chromatic-PSF corrections in the Y, J, and H bands; the corrected PSFs are fed forward through the AnaCal pixel-level shear-response chain.

Appropriate. AnaCal has been demonstrated to deliver sub-percent multiplicative-bias control on each of these three feature families [31, 37, 12], 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 [10, 11]. Chromatic-PSF effects, in turn, produce shear biases at or above the LSST DESC sub-percent threshold if uncorrected [42], and the multi-band, multi-survey overlap exploited here is precisely what mitigates them [44]. We validate the measurement step in `descwl-shear-sims/imSim` [70, 71, 72] 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. The required LSST-DR1 / Euclid-DR2 / Roman overlap for the chromatic-PSF correction becomes available in the FY28–FY29 window covered by Year 3 of this proposal, the natural moment to deploy it at LSST-Y1 systematics-budget precision.

Thrust #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 [39, 38] 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, adding 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% [14, 15]. Validation uses `descwl-shear-sims/imSim` [70, 71, 72] 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.

Thrust #4: Field-level cosmology analysis on the joint LSST×Euclid mosaic

Method. The calibrated, co-registered shear and photo-z catalog from Thrusts #1–#3 is fed into the Gaussian-prior `Miko` [28] field-level inference engine. The forward model partitions the survey footprint into LSST-only, Euclid-only, and LSST–Euclid overlap subregions, and assigns each subregion its own depth, shape-

noise level, source number density, PSF, blending, and photo-z model, so that survey inhomogeneity is absorbed into the likelihood by construction. Posterior sampling over cosmological and nuisance parameters at field-level resolution is made tractable by coupling `Miko` to the differentiable cosmological emulators `CosmoPower` [52] and `CosmoPower-JAX` [53], which provide analytical gradients of the linear and nonlinear matter power spectra with respect to the cosmological parameters. These gradients integrate seamlessly with the AnaCal autodiff calibration layer of Thrusts #2–#3, yielding an end-to-end-differentiable stack from raw pixels to the cosmological posterior.

Appropriate. A Gaussian-prior field-level analysis has been demonstrated to recover *unbiased* tomographic power-spectrum amplitudes on HSC Year-3-like mocks [28], in contrast to log-normal field-level priors that introduce significant prior misspecification bias [49, 50]. The forward-model treatment of survey inhomogeneity is the natural way to combine LSST-only, Euclid-only, and LSST–Euclid overlap regions coherently, avoiding the residual inhomogeneity systematics that limit catalog-level cosmic-shear analyses at Stage-IV precision [48].

Timely. The `Miko` pipeline and the `CosmoPower-JAX` emulators are publicly released and already validated, so the field-level analysis can be deployed immediately on the joint LSST-DR1 / Euclid-DR2 catalog produced by Thrusts #1–#3 in the FY28–FY29 window. Downstream, the same calibrated catalog will also feed log-normal and N-body-prior field-level inference [49, 50, 51, 54] and simulation-based / deep-learning cosmology analyses based on higher-order summaries [55, 56, 57], broadening the scientific reach beyond the Gaussian-prior baseline.

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, as illustrated in Figure 6. **Each year requires 50% FTE of the PI and 100% FTE of two postdoctoral researchers.**

Pathfinder analysis 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 the joint image-processing pipeline (Thrusts #1/#2). Extend the AnaCal framework [31] from single-survey to multi-resolution, multi-survey joint imaging by performing joint source detection and correctly propagating the 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

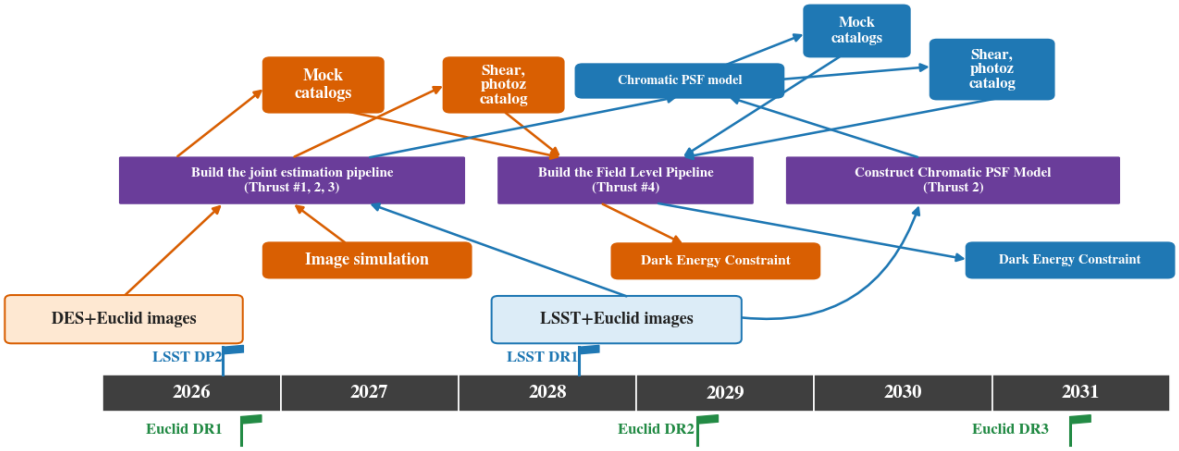

Figure 6: Five-year timeline of the proposed program, aligned with the Stage-IV imaging-survey data-release schedule (Euclid-DR1, LSST-DR1, Euclid-DR2, etc.). Each year couples a methodology deliverable to a real-data deliverable on the corresponding data.

simulation tests.

Year 2 (FY27–FY28): From images to catalogs (Thrusts #2/#3). 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` [39] 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 overlapping $\sim$ 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 (Thrust #4). Complete the Euclid-DR1×DES pathfinder program with a Gaussian-prior field-level cosmology analysis on the joint shear and photo-z catalogs above, using the `Miko` [28] field-level inference engine coupled to the differentiable `CosmoPower` [52, 53] emulators for gradient-based posterior sampling over cosmological and nuisance parameters, with a self-consistent forward-model treatment of shear bias, intrinsic alignments, photo-z uncertainty, and survey inhomogeneity. The result will be benchmarked against the DES Y6 cosmic-shear baseline.

Deliverables: cosmological parameter posteriors 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 is expected to grow 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): Chromatic PSF modeling (Thrust #2). The larger survey overlap and the more stringent LSST-Y1 systematic-error requirements make this the stage at which chromatic-PSF effects [42, 43, 44] 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): Image to catalog (Thrusts #1/#2/#3). 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 [70, 71, 72], 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 field-level dark-energy analysis (Thrust #4). Use the validated joint catalog to perform a Gaussian-prior field-level cosmology analysis on the LSST-DR1×Euclid-DR2 overlap with the `Miko` [28] engine coupled to the `CosmoPower` [52, 53] emulators, 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 field-level 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.

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 [30], and the analytical noise-bias correction [31] that together constitute the AnaCal formalism—and wrote the entire AnaCal code base from scratch [37]. 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 [73] and the subsequent HSC Y3 cosmic-shear analysis [74], developed the HSC Y6 image-simulation suite for shear calibration, is leading the ongoing HSC Y6 shear catalog, produced the LSST Commissioning Camera DP1 AnaCal shear catalog [41], and is currently leading the LSST DP2 AnaCal shear catalog. The PI is also supervising graduate students applying AnaCal to calibrate AI-based shear estimators [12, 13], and has co-supervised the graduate student who developed the `Miko` field-level weak-lensing inference pipeline [28]—giving the PI direct training-the-trainer experience that maps onto both the AI-extension and field-level-inference thrusts of this proposal.

Host institution. The PI is employed at Brookhaven National Laboratory (BNL), which regularly attracts top-tier postdoctoral scholars to its Cosmology and Astrophysics group. BNL has a long tradition of leader-

ship 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 [73] and the corresponding HSC Y3 cosmic-shear analysis [74], produced the LSST DP1 AnaCal shear catalog [41], 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 [12, 13], 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. *International Journal of Modern Physics D*, 10(2):213–223, January 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, March 2019.

[3] DESI Collaboration. The DESI Experiment Part I: Science, Targeting, and Survey Design. *arXiv e-prints*, page arXiv:1611.00036, October 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, November 2005.

[5] A. G. Adame et al. DESI 2024 VI: cosmological constraints from the measurements of baryon acoustic oscillations. JCAP, 2025(2):021, February 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(8):083515, October 2025.

[7] DESI Collaboration, K. Lodha, et al. Extended dark energy analysis using DESI DR2 BAO measurements. Phys. Rev. D, 112(8):083511, October 2025.

[8] Judit Prat and David Bacon. Weak gravitational lensing. In *Encyclopedia of Astrophysics*, volume 5, pages 508–537. Elsevier, January 2026.

[9] Euclid Collaboration. Euclid: I. Overview of the Euclid mission. A&A, 697:A1, May 2025.

[10] R. L. Schuhmann, C. Heymans, and J. Zuntz. Galaxy shape measurement synergies between LSST and Euclid. *arXiv e-prints*, page arXiv:1901.08586, January 2019.

[11] R. Chary et al. Joint Survey Processing of Euclid, Rubin and Roman: Final Report. *arXiv e-prints*, page arXiv:2008.10663, August 2020.

[12] 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.

[13] 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.

[14] 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 Near-ultraviolet Photometry. AJ, 159(6):258, June 2020.

[15] 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*, page arXiv:1501.07897, January 2015.

[16] D. Spergel et al. Wide-Field InfrarRed Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA 2015 Report. *arXiv e-prints*, page arXiv:1503.03757, March 2015.

[17] Matthias Bartelmann and Peter Schneider. Weak gravitational lensing. Phys. Rep., 340(4-5):291–472, January 2001.

[18] Henk Hoekstra and Bhuvnesh Jain. Weak Gravitational Lensing and Its Cosmological Applications. *Annual Review of Nuclear and Particle Science*, 58(1):99–123, November 2008.

[19] Martin Kilbinger. Cosmology with cosmic shear observations: a review. *Reports on Progress in Physics*, 78(8):086901, July 2015.

[20] Rachel Mandelbaum. Weak Lensing for Precision Cosmology. ARA&A, 56:393–433, September 2018.

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

[22] H. Aihara, N. Arimoto, R. Armstrong, et al. The Hyper Suprime-Cam SSP Survey: Overview and survey design. PASJ, 70:S4, January 2018.

[23] J. T. A. de Jong, G. A. Verdoes Kleijn, K. H. Kuijken, and E. A. Valentijn. The Kilo-Degree Survey. *Experimental Astronomy*, 35(1-2):25–44, January 2013.

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

[25] 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*, page arXiv:1809.01669, September 2018.

[26] P. Capak et al. Enhancing LSST Science with Euclid Synergy. *arXiv e-prints*, page arXiv:1904.10439, April 2019.

[27] 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, February 2024.

[28] 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.

[29] 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.

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

[31] 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, February 2025.

[32] 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.

[33] 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.

[34] Eric Huff and Rachel Mandelbaum. Metacalibration: Direct Self-Calibration of Biases in Shear Measurement. *arXiv e-prints*, page arXiv:1702.02600, February 2017.

[35] E. S. Sheldon and E. M. Huff. Practical Weak-lensing Shear Measurement with Metacalibration. ApJ, 841(1):24, May 2017.

[36] E. S. Sheldon, M. R. Becker, N. MacCrann, and M. Jarvis. Mitigating Shear-dependent Object Detection Biases with Metacalibration. ApJ, 902(2):138, October 2020.

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

[38] 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.

[39] The RAIL Team et al. Redshift Assessment Infrastructure Layers (RAIL): Rubin-era photometric redshift stress-testing and at-scale production. *The Open Journal of Astrophysics*, 9:58200, February 2026.

[40] 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.

[41] 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.

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

[43] 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, September 2025.

[44] 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.

[45] 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, January 2019.

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

[47] Taco S. Cohen and Max Welling. Group Equivariant Convolutional Networks. *arXiv e-prints*, page arXiv:1602.07576, February 2016.

[48] Husni Almoubayyed, Rachel Mandelbaum, Humna Awan, Eric Gawiser, R. Lynne Jones, Joshua Meyers, J. Anthony Tyson, Peter Yoachim, and The LSST Dark Energy Science Collaboration. Optimizing LSST observing strategy for weak lensing systematics. *Monthly Notices of the Royal Astronomical Society*, 499(1):1140–1158, November 2020.

[49] Supranta S. Boruah, Eduardo Rozo, and Pier Fiedorowicz. Map-based cosmology inference with lognormal cosmic shear maps. *Monthly Notices of the Royal Astronomical Society*, 516(3):4111–4122, October 2022.

[50] Pier Fiedorowicz, Eduardo Rozo, Supranta S. Boruah, Chihway Chang, and Marco Gatti. KaRMMa – Kappa Reconstruction for Mass Mapping. *Monthly Notices of the Royal Astronomical Society*, 512(1):73–85, March 2022.

[51] Natalia Porqueres, Alan Heavens, Daniel Mortlock, Guilhem Lavaux, and T. Lucas Makinen. Field-level inference of cosmic shear with intrinsic alignments and baryons. *arXiv e-prints*, April 2023.

[52] H. T. Jense, I. Harrison, E. Calabrese, A. Spurio Mancini, B. Bolliet, J. Dunkley, and J. C. Hill. A complete framework for cosmological emulation and inference with CosmoPower. *RAS Techniques and Instruments*, 4:rzaf002, 2025.

[53] Davide Piras and Alessio Spurio Mancini. CosmoPower-JAX: high-dimensional Bayesian inference with differentiable cosmological emulators. *The Open Journal of Astrophysics*, 6:20, 2023.

[54] Divij Sharma, Biwei Dai, Francisco Villaescusa-Navarro, and Uroš Seljak. A field-level emulator for modelling baryonic effects across hydrodynamic simulations. *Monthly Notices of the Royal Astronomical Society*, 538:1415, 2025.

[55] Tianhuan Lu, Zoltán Haiman, and Xiangchong Li. Cosmological constraints from HSC survey first-year data using deep learning. *Monthly Notices of the Royal Astronomical Society*, 521(2):2050, 2023.

[56] M. Gatti, G. Campailla, N. Jeffrey, L. Whiteway, A. Porredon, J. Prat, J. Williamson, M. Raveri, B. Jain, V. Ajani, et al. Dark Energy Survey Year 3 results: Simulation-based cosmological inference with wavelet harmonics, scattering transforms, and moments of weak lensing mass maps. II. Cosmological results. *Physical Review D*, 111(6):063504, 2025.

[57] Biwei Dai and Uros Seljak. Multiscale Flow for Robust and Optimal Cosmological Analysis. In *Machine Learning for Astrophysics*, page 10, 2023.

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

[59] 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.

[60] 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.

[61] 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.

[62] Anowar J. Shajib, Graham P. Smith, Simon Birrer, Aprajita Verma, Nikki Arendse, Thomas Collett, Tansu Daylan, Stephen Serjeant, and LSST Strong Lensing Science Collaboration. Strong gravitational lenses from the Vera C. Rubin Observatory. *Philosophical Transactions of the Royal Society of London Series A*, 383(2295):20240117, May 2025.

[63] Narayan Khadka, Simon Birrer, Alexie Leauthaud, and Holden Nix. Breaking the mass-sheet degeneracy in strong lensing mass modelling with weak lensing observations. MNRAS, 533(1):795–806, September 2024.

[64] Siddharth Mishra-Sharma, David Alonso, and Joanna Dunkley. Neutrino masses and beyond-ΛCDM cosmology with LSST and future CMB experiments. *Physical Review D*, 97(12):123544, 2018.

[65] 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, December 2018.

[66] 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, December 2019.

[67] 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 Hyper Suprime-Cam data. MNRAS, 526(1):1122–1137, November 2023.

[68] Grant Merz, Xin Liu, Samuel Schmidt, Alex I. Malz, Tianqing Zhang, et al. DeepDISC-photoz: Deep Learning-Based Photometric Redshift Estimation for Rubin LSST. *The Open Journal of Astrophysics*, 8:40, April 2025.

[69] The AION Collaboration et al. AION-1: Omnimodal Foundation Model for Astronomical Sciences. *arXiv e-prints*, page arXiv:2510.17960, October 2025.

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

[71] 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):2801–2820, June 2023.

[72] 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):3799–3823, December 2025.

[73] 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.

[74] 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.

APPENDIX 2: Facilities & Other Resources

Facilities and other resources go here.

APPENDIX 3: Equipment

Equipment information goes here.

APPENDIX 4: Data Management Plan

Data management plan goes here.

APPENDIX 5: Synergistic Activities (optional)

Synergistic activities go here.

APPENDIX 6: Transparency of Foreign Connections

Transparency of foreign connections information goes here.

APPENDIX 7: Other Attachments

Other attachments go here.