

Classifying Supermassive Black Hole Growth Regimes to Observables Across Cosmological Simulations with Forecasts for LSST

HITAISHI CHILLARA ¹

¹*Independent Researcher*

ABSTRACT

The possibility of over-massive black holes suggested by James Webb Space Telescope photometric discoveries of ‘little red dots’, **may disfavor light supermassive black hole (SMBH) seeds**. However, what should constitute the mass (range) of ‘heavy’ seeds remains relatively unconstrained. Moreover, Vera C Rubin Observatory’s Legacy Survey of Space and Time will photometrically characterize galaxies without direct black hole mass measurements. We forward-model the SIMBA, ILLUSTRISTNG, and EAGLE cosmological simulations into the photometric bands of LSST to train an ensemble machine learning classifier. Our framework achieves 91%–94% accuracy across SIMBA and ILLUSTRISTNG in distinguishing between over-massive and under-massive SMBH growth regimes under LSST magnitude limits, using only broadband photometry. Furthermore, cross-simulation transfer experiments (training on one cosmological simulation and evaluating on another using rank-normalized features) achieve 83%–89% accuracy. This suggests the relative photometric ordering of growth regimes is largely preserved even across fundamentally different sub-grid SMBH feedback prescriptions. Signal decomposition shows our classification is driven by host galaxy colors (82%–87% accuracy) and, relatedly, the accretion-state’s spectral energy distribution shape as opposed to an inversion of our forward model’s analytical luminosity prescription. Given that the evaluated simulations employ heavy seed prescriptions ($\geq 10^4 M_\odot$), our methodology establishes a validated baseline for classifying post-seeding growth regimes.

Keywords: supermassive black holes (1663) — galaxy evolution (594) — active galactic nuclei (16) — quasars (1319) — high-redshift galaxies (734) — computational methods (1965)

1. INTRODUCTION

The early growth and formation of supermassive black holes (SMBHs) in the first billion years of the universe poses key challenges for modern astrophysics. Luminous quasars observed at redshifts $z > 7$ reveal black holes with masses beyond $10^9 M_\odot$ (K. Inayoshi et al. 2020; M. Volonteri et al. 2021; X. Fan et al. 2023; P. Natarajan et al. 2024; A. Bogdán et al. 2024). Adding to this, recent observations from JWST have further complicated matters by revealing populations of “little red dots” at $z > 6$ where the inferred black hole masses surpass local scaling relations by orders of magnitude (Y. Harikane et al. 2023; R. Maiolino et al. 2024; D. D. Kocevski et al. 2023).

Understanding the manner in which these early objects grew relative to their hosts necessitates observable diagnostics that are capable of distinguishing between various SMBH growth regimes. While theoretical models do propose many seeding pathways, from “light seeds” ($\sim 100 M_\odot$ remnants of Population III

stars, which face severe growth challenges from radiative feedback and angular momentum transport; P. Madau & M. J. Rees 2001; V. Bromm & R. B. Larson 2004; S. Hirano et al. 2014; M. Volonteri 2010; T. Tanaka & Z. Haiman 2009; M. Milosavljević et al. 2009) to “heavy seeds” ($\sim 10^4$ – $10^6 M_\odot$ direct collapse or stellar collisions; V. Bromm & A. Loeb 2003; M. C. Begelman et al. 2006; G. Lodato & P. Natarajan 2006; J. Regan & M. Volonteri 2024; B. Devecchi & M. Volonteri 2009; A. Lupi et al. 2014), direct observational discrimination of the seeding epoch ($z > 10$ – 15) remains inaccessible. However, the subsequent history of accretion and black hole–host galaxy co-evolution leaves persistent imprints on the $M_{\text{BH}}\text{--}M_\star$ relation (M. Habouzit et al. 2021, 2022; F. Pacucci et al. 2023; F. Sassano et al. 2021; A. K. Bhowmick et al. 2025). The Trinity empirical model (H. Zhang et al. 2023; H. Zhang et al. 2024), for example, finds that luminosity-selected samples at high redshift are biased toward systematically over-massive systems by 0.35–1.0 dex. Through classification of post-seeding growth regimes from observables,

future attempts for purely-photometric seed-origin discrimination require that growth-regime signatures be robust and detectable across various theoretical models.

Imminently, the Vera C. Rubin Observatory’s Legacy Survey of Space and Time (LSST) will detect billions of galaxies that lack direct black hole mass measurements, and identifying broadband photometric signatures of these growth regimes is vital. Cosmological simulations such as SIMBA (R. Davé et al. 2019), ILLUSTRISTNG (V. Springel et al. 2018; A. Pillepich et al. 2018; D. Nelson et al. 2018), EAGLE (J. Schaye et al. 2015; R. A. Crain et al. 2015), and Horizon-AGN (Y. Dubois et al. 2014), utilize varying subgrid models of SMBH seeding, feeding, and feedback, enabling exploration of observable consequences of the BH-host connection (F. Shankar et al. 2025). Simultaneously, machine learning has emerged as a powerful tool for uncovering complex correlations for such datasets (A. F. L. Bluck et al. 2023).

In this paper, we present a machine learning framework that bridges theoretical simulations and observational predictions. We train various classifiers to distinguish between “over-massive” and “under-massive” black holes relative to a simulation-fitted $M_{\text{BH}}-M_*$ relation (Section 3.3), using the growth regime as a proxy for their accretion history. Since simulations analyzed here use heavy seed prescriptions ($\sim 10^4-10^6 M_\odot$), our methodology focuses on establishing the classification of post-seeding growth regimes. Furthermore, utilizing cross-simulation transfer experiments, we test whether the relative photometric ordering of growth regimes is physically preserved across simulation codes having fundamentally different accretion modes; these experiments demonstrate that our learned features truly reflect transferable physical signatures instead of simulation-specific artifacts.

2. DATA

2.1. Simulations and Data Selection

We draw upon data from the SIMBA (R. Davé et al. 2019) and ILLUSTRISTNG (A. Pillepich et al. 2018; D. Nelson et al. 2018; V. Springel et al. 2018; F. Marinacci et al. 2018; J. P. Naiman et al. 2018; D. Nelson et al. 2019a) cosmological simulations. SIMBA evolved a $(100 h^{-1} \text{ cMpc})^3$ comoving volume from $z = 249$ to $z = 0$ using the GIZMO code (P. F. Hopkins 2015), seeding SMBHs at $10^4 M_\odot/h$ and using a two-mode accretion model (P. F. Hopkins & E. Quataert 2011; D. Anglés-Alcázar et al. 2017). ILLUSTRISTNG 100-1 evolved a $(75 h^{-1} \text{ cMpc})^3$ volume using the AREPO code (V. Springel 2010), seeding SMBHs at $\sim 1.2 \times 10^6 M_\odot$ with Eddington-limited Bondi-Hoyle accretion

as well as dual-mode AGN feedback (R. Weinberger et al. 2018). To investigate resolution and volume trade-offs, we also analyze TNG300-1, as well as the higher-resolution TNG50-1 and TNG50-2 volumes (D. Nelson et al. 2019b; A. Pillepich et al. 2019).

From SIMBA, we obtain BH catalogs that contain ~ 20 million entries spanning $z = 0$ to $z \approx 10$, of which ~ 2 million lie at $z \geq 4$. All entries include: black hole mass (M_{BH}), mass accretion rate ($\dot{M}$), host galaxy stellar mass (M_*), star formation rate (SFR), scale factor (a), and a unique identifier. From the ILLUSTRISTNG simulation suite, we collect entries that span $z = 0$ to $z = 10$, including 17,254 entries from TNG100-1, 31,762 from TNG300-1, 292,308 from TNG50-1, and 289,483 from TNG50-2. Both simulation suites are analyzed across their full redshift range; high-redshift subsets ($z \geq 4$) are examined separately where noted (Section 4). The diverse seeding prescriptions and resolutions robustly test our methodology across varying black hole physics baselines.

2.2. LSST DP1 Observational Baseline

For establishing an observational baseline for our synthetic forecasts, we construct a comparison sample from LSST Data Preview 1 (DP1; NSF-DOE Vera C. Rubin Observatory 2025) observations of the Extended Chandra Deep Field South (ECDFS). Here, selecting extended sources that have signal-to-noise > 5 in the r -band yields 2,533 low- z AGN candidates ($g - r < 1.0$), 286 high- z g -dropout candidates, and 60 r -dropout candidates. Further cross-matching with SIMBAD recovers 439 unique objects with spectroscopic redshifts ($z = 0.04$ to 5.93 , median $z = 1.38$).

3. METHODS

3.1. Forward Modeling Pipeline

In order to generate synthetic LSST-compatible observations, the created forward modeling pipeline transforms simulated physical properties into predicted observed fluxes. We compute bolometric luminosities for AGN as $L_{\text{bol,AGN}} = \eta \dot{M} c^2$ ($\eta = 0.1$). We adopt a multi-break power-law SED from the D. E. Vanden Berk et al. (2001) composite quasar spectrum and G. T. Richards et al. (2006) templates:

$$L_\nu \propto \nu^\alpha, \quad \alpha = \begin{cases} -1.8 & \lambda_{\text{rest}} < 1000 \text{ \AA} \\ -0.7 & 1000 \text{ \AA} < \lambda_{\text{rest}} < 1400 \text{ \AA} \\ -0.44 & 1400 \text{ \AA} < \lambda_{\text{rest}} < 3000 \text{ \AA} \\ -0.30 & 3000 \text{ \AA} < \lambda_{\text{rest}} < 5000 \text{ \AA} \\ -0.80 & 5000 \text{ \AA} < \lambda_{\text{rest}} < 10000 \text{ \AA} \end{cases} \quad (1)$$

We utilize a luminosity-dependent UV slope correction (A. T. Steffen et al. 2006) on top of an Eddington ratio-dependent bolometric correction (E. Lusso et al. 2012).

For $z > 3$, we implement the P. Madau (1995) intergalactic medium (IGM) absorption prescription that captures the g -band dropout signature of $z > 4$ AGN. Then, we assign 50% of the AGN as obscured “Type 2” with dust reddening from an exponential distribution (scale parameter $E(B - V) = 0.3$), attenuated with the SMC extinction curve (K. D. Gordon et al. 2003).

Host galaxy r -band luminosities are computed using stellar mass-to-light ratios obtained from E. F. Bell & R. S. de Jong (2001): $(M/L)_r = 2.5 L_\odot/M_\odot$ for quiescent systems, and $1.0 L_\odot/M_\odot$ for star-forming galaxies. Thus, the total observed flux is the sum of AGN and host galaxy components. Lastly, we apply a Galactic (E. F. Schlafly & D. P. Finkbeiner 2011) and internal host extinction (D. Calzetti et al. 2000), introducing magnitude-scaled Gaussian scatter ($\sigma_m = 0.05 \times (m/23)$) to match uncertainties in DP1 observational data.

3.2. AGN Selection Criteria

Using standard quasar selection criteria, for low- z AGN ($z \lesssim 3$), we impose $g - r < 1.0$ and $15 < r < 24$ mag (G. T. Richards et al. 2002). High- z AGN ($z > 4$) are chosen through Lyman-break criteria (X. Fan et al. 2006; R. A. Overzier et al. 2006): g -dropouts ($z > 4.0$) require $g - r > 1.2$, $r - i < 0.6$, $i < 25.5$; r -dropouts ($z > 5.5$) require $r - i > 1.3$, $i - z < 0.8$, $z < 25.0$.

3.3. Machine Learning Framework

The classifier we developed distinguishes “over-massive” from “under-massive” black holes, determined by their variance from a linear fit to the simulation’s $M_{\text{BH}}-M_*$ population. Following W. Ma et al. (2025), we utilize the same relation with binned-median regression, yielding $\log_{10}(M_{\text{BH}}) = 1.49 \log_{10}(M_*) - 6.59$ for SIMBA. The steeper slope relative to the local relation (J. Kormendy & L. C. Ho 2013) demonstrates elevated M_{BH}/M_* normalizations at high redshift (F. Pacucci & A. Loeb 2024). We split objects at this median offset, ensuring balanced classes (Table 7 confirms an insensitivity to exact percentiles).

We use 16 purely observable features for the classifier: magnitudes ($griz$), colors ($g-r, r-i, i-z, g-i$), redshift, clipped $\log_{10} \lambda_{\text{Edd}}$, as well as six engineered features (e.g., SED color curvature $\kappa = (g-r) - 2(r-i) + (i-z)$, spectral range $m_g - m_z$, and a proxy for absolute magnitude, $m_r - 5 \log_{10}(d_L)$). We stack four models (Neural Network, FT-Transformer (Y. Gorishniy et al. 2021), Random Forest, XGBoost) via a logistic regression meta-learner (D. H. Wolpert 1992).

We employ 5-fold cross-validation, splitting by unique black hole ID to prevent data leakage across temporal snapshots. Evaluation is done in two primary modes: a “Magnitude-Limited” regime (LSST co-added depth limits from v. Ivezić et al. 2019 with conservative AGN cuts $g, r < 26.0$, $i < 25.5$, $z < 25.0$) and an “Intrinsic” regime (no magnitude cuts). Our training sets are then balanced through random undersampling.

To reduce potential label noise, we track the subhalo IDs across snapshots and filter objects whose growth regime classification fluctuates. Such transitioning black holes lie near the classification boundary, introducing ambiguity. The removal fraction scales inversely with simulation resolution ($\sim 2\%$ in SIMBA, $\sim 11\%$ in TNG50-1, and $\sim 27\%$ in TNG50-2), reflecting coarser mass discretization near this boundary as opposed to epoch-selective bias (Figure 1). Thus, removing them provides a cleaner training set without artificially creating the physical signal: evaluating even TNG50-2 without removal still yields $\sim 88\%$ magnitude-limited accuracy.

3.4. Cross-Simulation Generalization

To evaluate whether learned features capture transferable physical signals instead of simulation-specific artifacts, we utilize a cross-simulation transfer experiment (training on one simulation, evaluating on another). As raw photometric features span different absolute ranges across simulation codes due to distinct accretion models, we construct domain-invariant relative features. Within adaptive redshift bins, we replace each photometric property with a corresponding percentile rank, quantifying how bright the AGN is relative to its peers at the same redshift. For tests across domains, we use tree-based models (XGBoost, Random Forest), as transformer architectures and their attention features tend to encode simulation-specific pairwise interactions.

4. RESULTS

4.1. Simulation Population Properties

Inspecting SIMBA’s $z \geq 4$ black holes reveals a population that exhibits downsizing (Figure 2). The Eddington ratio median is $\log_{10} \lambda_{\text{Edd}} = -1.07$, with a super-Eddington fraction of 1.5% (Figure 2). The $M_{\text{BH}}-M_*$ relation slope is 2.26 ± 0.07 (Figure 3), steeper than the ILLUSTRISTNG slopes overlaid in the same figure. This difference arises from the distinct accretion prescriptions: SIMBA’s torque-limited model channels cold gas more efficiently onto BHs in gas-rich hosts, producing steeper $M_{\text{BH}}-M_*$ scaling, while ILLUSTRISTNG’s Bondi-Hoyle accretion with strong kinetic feedback at low Eddington ratios suppresses growth in massive hosts




Figure 1. Label-flip removal diagnostics. *Top row:* A distribution of offsets from the simulation-fitted reference relation for retained (filled) and removed (dashed) populations that are split by class: (blue for under-massive, red for over-massive). In each simulation the removed population is denser towards the nearer classification boundary than the retained population: SIMBA ($|\Delta|_{\text{rem}} \approx 0.3$ vs. $|\Delta|_{\text{ret}} \approx 0.6$ dex), TNG50-1 (0.3 vs. 0.4 dex), and TNG50-2 (0.3 vs. 0.4 dex). This is consistent with boundary oscillation as opposed to systematic misclassification. This split is clearest in SIMBA and progressively tighter in TNG50-1 and TNG50-2, reflecting coarser mass resolution. *Bottom row:* Redshift distributions of removed and retained populations overlap, confirming that removal is not epoch-selective. Columns contain SIMBA (left), TNG50-1 (center), and TNG50-2 (right).

(M. Habouzit et al. 2021; R. Weinberger et al. 2018). The SIMBA median relation tracks above the Trinity empirical model (H. Zhang et al. 2023; H. Zhang et al. 2024), which expects high-redshift, luminosity-selected samples to be biased towards systematically over-massive systems (Figure 4). ILLUSTRISTNG produces a shallower median $M_{\text{BH}}-M_*$ relation that sits between SIMBA and the local J. Kormendy & L. C. Ho (2013) scaling at $z = 4-5$, consistent with its more conservative Bondi-limited accretion model.

4.2. Validation against LSST DP1

Our forward modeling pipeline accurately reproduces the broadband color locus and magnitude range of observed DP1 AGN (Figure 5). For objects where $z < 2$, the median magnitude offsets range from -0.34 in g -band to $+0.07$ in z -band (Figure 6). Though simplified UV SED treatments result in minor systematic offsets in bluer bands, the critical color-color separations used by the classifier remain robust as they rely on *differential* signatures between growth regimes at fixed redshift. Mock catalogs contain the same mass range observed in high- z quasars (K. Inayoshi et al. 2020), although the finite $(100 h^{-1} \text{ cMpc})^3$ simulation volume limits statistics for the rarest objects where $M_{\text{BH}} > 10^9 M_{\odot}$.

4.3. Classification Performance

Utilizing solely observable features, our stacked ensemble demonstrates robust discriminative power across growth regimes (Figures 9 and 10). For SIMBA, the model achieves 92.3% accuracy ($\text{AUC} = 0.978$) in magnitude-limited mode and 93.5% accuracy ($\text{AUC} = 0.983$) in intrinsic mode. The minimal performance drop of $\approx 1.2\%$ between modes demonstrates that discriminative power is preserved under realistic LSST observational constraints.

Through an ablation study applying each LSST magnitude limit separately (Table 1), we find that the r -band cut incurs the largest single penalty ($\Delta = -0.9\%$), as it appears in each primary optical color, but the combined effect of each of the 4 cuts reduces accuracy by only 0.2%. Evaluating across redshifts, intrinsic accuracy also remains highly stable ($\sim 92-96\%$) from $z < 0.5$ to $z > 5$ (Table 2), confirming that the classifier isolates persistent physical signatures instead of redshift-dependent artifacts.

With the FT-Transformer, attention heatmaps (Figure 8) reveal the network’s prioritization of the absolute magnitude proxy and spectral range ($m_g - m_z$). This demonstrates the network’s reliance on SED *shape*. Over-massive systems accreting near the Eddington




Figure 2. Analysis of SIMBA’s black hole population at $z \geq 4$. *Top Left:* Mass function showing downsizing behavior. *Top Right:* $M_{\text{BH}}-z$ relation colored by the Eddington ratio. *Bottom:* Number density evolution.

limit tend to produce flatter, bluer SEDs that are readily distinguished from sub-Eddington populations.

4.4. Signal Decomposition (Circularity Tests)

As our forward model sequentially links $M_{\text{BH}} \rightarrow \dot{M} \rightarrow L_{\text{bol}} \rightarrow$ magnitudes, the classifier could simply invert this analytic mapping instead of detecting genuine physical signatures. We conducted three decomposition experiments (Tests A, B, and C) (Table 3).

In Test A, we randomly permute $\dot{M}$ for black holes at fixed redshift, breaking the $M_{\text{BH}} \rightarrow \dot{M}$ correlation. Accuracy drops by only 2–4%, showing that the direct $L_{\text{bol}}(M_{\text{BH}})$ channel represents only a marginal fraction of the signal. In Test B, classifying using solely host galaxy photometry ($L_{\text{bol,AGN}} = 0$) yields 82–87% accuracy, further confirming that BH-host co-evolution is strongly imprinted on galaxy colors via mass-to-light correlations that are standard in photometric surveys. In Test C, we restrict to narrow 0.3 dex M_{BH} bins, and classifying yields high accuracy (82–91%); since mass is functionally constant within each bin, the model leverages genuine accretion-state SED differences rather than global mass recovery.

4.5. Cross-Validation and Resolution Effects

Our developed framework performs robustly across varying simulation physics and resolutions (Table 4). TNG50-1 and TNG50-2 achieve comparable magnitude-limited accuracies (90.7% and 93.6%, respectively). The lower-resolution but larger-volume TNG300-1 simulation yields 94.3% accuracy in intrinsic mode, but provides insufficient samples in magnitude-limited mode, underscoring the need for cosmological simulations spanning a wide range of resolutions. Conversely, extremely low-resolution volumes (TNG100-2, TNG300-2) artificially inflate accuracy ($> 99\%$) as a result of significantly coarser mass discretization leading to trivial separability of populations; these results are presented in Appendix B. The consistency in performance between SIMBA and the high-resolution ILLUSTRISTNG runs confirms the model captures intrinsic physical scatter rather than numerical artifacts. The EAGLE simulation suite is detailed in Appendix A. While intrinsic classification achieves 96–98% across all nine EAGLE configurations, the magnitude-limited mode fails because EAGLE’s conservative accretion physics produces AGN too faint to pass LSST detection thresholds, reducing samples below the training minimum.




Figure 3. SIMBA black hole scaling relations at $z \geq 4$ that follow [M. Habouzit et al. \(2021\)](#), with fit slopes of ILLUSTRITNG overlaid and dashed lines of local scaling relations.

Table 1. Ablation Study: LSST Magnitude Cut Impact

Selection	N	Accuracy	Δ
No cuts	100,000	$83.0\% \pm 0.2\%$	...
$g < 26.0$	69,070	$82.7\% \pm 0.6\%$	-0.3%
$r < 26.0$	78,338	$82.1\% \pm 0.9\%$	-0.9%
$i < 25.5$	74,678	$82.6\% \pm 0.2\%$	-0.4%
$z < 25.0$	69,134	$82.4\% \pm 0.2\%$	-0.6%
All 4	65,594	$82.8\% \pm 0.2\%$	-0.2%

Note. Simplified pipeline (single Random Forest, no label noise removal) to focus on cut impacts.

4.6. Cross-Simulation Transfer

When training on one simulation code and evaluating on another, we find models trained on raw absolute photometry to perform poorly (~ 66 – 69% intrinsic accuracy) due to differing baseline luminosities. However, by training on domain-invariant rank features we suc-

cessfully recover model performance, yielding 83–89% intrinsic transfer accuracy (Table 5).

The *relative* physical ordering of growth regimes (how a specific AGN ranks against similar redshift peers) is preserved despite highly distinct subgrid accretion prescriptions. Under LSST magnitude limits, transfer accuracy drops to 74–82%, reflecting the difficulty of maintaining rank normalization under truncated, flux-




Figure 4. Comparison of median $M_{\text{BH}}-M_*$ relations from SIMBA (red) and ILLUSTRISTNG (green) against the Trinity empirical model (blue; H. Zhang et al. 2023) at $z = 4-7$. ILLUSTRISTNG is shown at $z = 4-5$, where sufficient black hole growth has occurred for a reliable median relation. Gold stars mark JWST-discovered AGN (H. Zhang et al. 2024).

limited sample distributions. The $\sim 4-11\%$ accuracy gap between cross-simulation transfer and within-simulation performance reflects genuine differences in subgrid feedback physics: SIMBA employs torque-limited accretion (P. F. Hopkins & E. Quataert 2011; D. Anglés-Alcázar et al. 2017) while ILLUSTRISTNG uses Bondi-Hoyle accretion with kinetic-mode feedback at low accretion rates (R. Weinberger et al. 2018). These distinct prescriptions produce different $M_{\text{BH}}-M_*$ scatter patterns and Eddington ratio distributions. The classifier learns *relative* photometric signatures of growth regimes, i.e., how a black hole’s accretion luminosity alters its host’s composite broadband colors, rather than the absolute $M_{\text{BH}}-M_*$ normalization, which differs substantially between codes (Figure 3). Since the radiation physics mapping accretion rate to photometric signature is simulation-independent, the differential color imprint of over-massive versus under-massive growth is preserved even when the underlying scaling relations di-

verge. That 83–89% accuracy still transfers indicates the growth-regime signature is broadly consistent across simulation codes, while the residual gap quantifies the degree to which simulation-specific feedback imprints limit generalization.

5. DISCUSSION

5.1. Framework Interpretation

The robustness of our pipeline demonstrates that broadband photometric observables reliably encode black hole growth history. The physical mechanism is as follows: at fixed redshift, over-massive black holes produce higher L_{bol} relative to their host stellar luminosity, shifting the AGN-to-host flux ratio and thus the composite broadband colors. Simultaneously, the host galaxy itself encodes M_* through its mass-to-light ratio, providing an independent constraint on the M_{BH}/M_* offset. The circularity tests (Section 4.4) confirm that both channels contribute: host-only photometry achieves 82–




Figure 5. Comparison of forward model predictions with LSST DP1. Top: color-color diagrams (SIMBA predictions in red, DP1 in blue). Bottom: cumulative magnitude distributions.

87% (Test B), while shuffling accretion rates drops accuracy by only 2–4% (Test A), indicating the signal is

distributed across the composite SED rather than concentrated in a single feature.




Figure 6. Band-by-band comparison of SIMBA forward model predictions against DP1 observations restricted to $z < 2$, with median offsets and KS statistics annotated. The DP1 sample consists of 1,737 AGN candidates at $z < 2$ (58 with spectroscopic redshifts from SIMBAD cross-matching, 1,679 with photometric redshifts).

The signal decomposition and attention heatmap studies confirm that the classification captures physical signals rather than a structural circularity in the forward model. That these signatures transfer effectively across SIMBA and the ILLUSTRISTNG suite implies that they represent genuine consequences of BH-host co-evolution. Rank computation provides an entirely unsupervised bridge; calculating percentile positions relative to a survey population allows simulation-trained models to adapt to real data without a strict luminosity calibration. The color-space separation between growth regimes is also robustly preserved across both magnitude-limited and intrinsic modes (Figure 11).

5.2. Photometry-Only Validation

Because M_{BH} enters the Eddington ratio $\lambda_{\text{Edd}} = L_{\text{bol}}/L_{\text{Edd}}$, the features derived from it could in principle introduce partial data leakage. To definitively

rule this out, we construct a 19-feature photometry-only pipeline entirely excluding λ_{Edd} and replacing it with purely observable engineered features (e.g., empirical color $\times$ magnitude cross terms and redshift interactions).

This strictly photometry-only pipeline attains 91.3% accuracy (AUC = 0.971) for SIMBA, 90.5% (AUC = 0.968) for TNG50-1, and 93.2% (AUC = 0.978) for TNG50-2 in magnitude-limited mode. The results remain within 0.2–1.0% of baseline 16-feature performance. This demonstrates that the classification accuracy is not an artifact of leakage via λ_{Edd} : the physical signal is fully accessible from photometric colors and magnitudes alone. The inclusion of λ_{Edd} (which is observationally accessible via single-epoch virial mass estimators (M. Vestergaard & B. M. Peterson 2006; Y. Shen et al. 2011)) provides a genuine 0.2–1.0% improve-




Figure 7. Redshift distribution of LSST DP1 AGN candidates from SIMBAD cross-match ($N = 439$, median $z = 1.38$).

Table 2. Classification Accuracy by Redshift Bin (SIMBA)

Redshift Bin	N	Intrinsic Acc.	Mag-Lim Acc.
<i>Magnitude-limited</i>			
$z < 0.5$	41,834	...	91.7%
$0.5 \leq z < 1$	8,723	...	94.7%
$1 \leq z < 2$	2,073	...	94.1%
<i>Intrinsic</i>			
$z < 0.5$	72,888	92.2%	...
$0.5 \leq z < 1$	28,547	94.7%	...
$1 \leq z < 2$	22,106	95.3%	...
$2 \leq z < 3$	5,978	96.2%	...
$3 \leq z < 5$	2,737	95.1%	...
$z \geq 5$	478	95.3%	...

ment for classification, confirming it contributes additional physical information rather than artificial leakage.

5.3. Systematic Error Budget

To quantify the sensitivity of our classification model results with respect to assumptions in the forward model, we performed a perturbation analysis varying

between baseline and perturbed values of five physical parameters independently (Table 6).

Total systematic uncertainty sums to $\pm 1.91\%$ in magnitude-limited mode. The dominant parameter is the star-forming mass-to-light ratio $(M/L)_{\text{SF}}$ ($\Delta = +1.53\%$ when doubled), reflecting the model's physical reliance upon signatures of host galaxies. Our reliance on an analytic forward model serves as a conservative




Figure 8. FT-Transformer attention heatmaps for SIMBA. The absolute magnitude proxy and spectral range receive the highest weights in both modes, indicating a consistent reliance on SED shape regardless of selection effects.

Table 3. Signal Decomposition and Circularity Tests

Test Configuration	Magnitude-Limited		Intrinsic	
	Accuracy	AUC	Accuracy	AUC
Baseline	92.3%	0.978	93.5%	0.983
A: Shuffled $\dot{M}$	87.9%	0.950	91.4%	0.971
B: Host-only	82.3%	0.892	87.2%	0.931
<i>C: Fixed M_{BH} bins</i>				
$\log_{10} M_{\text{BH}} = 6.9\text{--}7.2$	88.4%	0.952	90.1%	0.963
$\log_{10} M_{\text{BH}} = 7.2\text{--}7.5$	89.7%	0.960	91.2%	0.968
$\log_{10} M_{\text{BH}} = 7.5\text{--}7.8$	87.3%	0.944	88.9%	0.951
$\log_{10} M_{\text{BH}} = 7.8\text{--}8.1$	85.8%	0.932	87.4%	0.943
$\log_{10} M_{\text{BH}} = 8.1\text{--}8.4$	82.1%	0.908	83.5%	0.917

baseline; the incorporation of advanced spectral popu-

lation synthesis (e.g., FSPS, CIGALE) would increase

Table 4. Cross-Simulation Classification Results

Simulation Volume	Mode	N	Best Model	Accuracy	AUC
<i>Primary results (high-resolution, no caveats)</i>					
SIMBA	Magnitude-Limited	52,668	Ensemble	$92.3\% \pm 0.6\%$	0.978
SIMBA	Intrinsic	132,734	Ensemble	$93.5\% \pm 0.2\%$	0.983
TNG50-1	Magnitude-Limited	58,602	Ensemble	$90.7\% \pm 0.4\%$	0.970
TNG50-1	Intrinsic	231,846	Ensemble	$90.1\% \pm 0.0\%$	0.970
TNG50-2	Magnitude-Limited	30,628	Ensemble	$93.6\% \pm 0.4\%$	0.981
TNG50-2	Intrinsic	107,148	Ensemble	$93.9\% \pm 0.1\%$	0.987
TNG100-1	Intrinsic	7,228	NN	$88.4\% \pm 0.7\%$	0.962
TNG300-1	Intrinsic	7,418	NN	$94.3\% \pm 0.4\%$	0.988
EAGLE and resolution-limited runs are deferred to Appendices A and B					

Note. TNG100-1 and TNG300-1 appear in intrinsic mode only as their magnitude-limited samples fall below the 500-sample threshold.

Table 5. Cross-Simulation Transfer Accuracy

Transfer Direction	Mode	Raw Photometry	Rank Only	Both Combined
SIMBA $\rightarrow$ TNG50-1	Mag-Lim	72.5%	77.1%	78.1%
SIMBA $\rightarrow$ TNG50-1	Intrinsic	66.3%	89.1%	72.6%
TNG50-1 $\rightarrow$ SIMBA	Mag-Lim	72.2%	70.8%	73.6%
TNG50-1 $\rightarrow$ SIMBA	Intrinsic	68.7%	82.9%	81.9%
SIMBA $\rightarrow$ TNG50-2	Mag-Lim	75.2%	72.3%	82.1%
SIMBA $\rightarrow$ TNG50-2	Intrinsic	65.8%	85.0%	72.3%
TNG50-2 $\rightarrow$ SIMBA	Mag-Lim	72.8%	70.9%	73.8%
TNG50-2 $\rightarrow$ SIMBA	Intrinsic	69.0%	83.2%	72.3%

Table 6. Forward Model Systematic Error Budget

Parameter Perturbation	Baseline $\rightarrow$ Pert.	Δ Mag-Lim Acc. (%)	Δ Intrinsic Acc. (%)
η	0.1 $\rightarrow$ 0.06	−0.01	−0.16
η	0.1 $\rightarrow$ 0.2	+0.65	−0.28
$f_{\text{Type 2}}$	0.5 $\rightarrow$ 0.3	+0.63	+0.50
$f_{\text{Type 2}}$	0.5 $\rightarrow$ 0.7	−0.15	−0.41
$E(B - V)$ scale	0.3 $\rightarrow$ 0.1	+0.65	+0.62
$E(B - V)$ scale	0.3 $\rightarrow$ 0.5	−0.04	−0.35
$(M/L)_{\text{SF}}$	1.0 $\rightarrow$ 0.5	−0.06	+0.21
$(M/L)_{\text{SF}}$	1.0 $\rightarrow$ 2.0	+1.53	+0.64
Extinction	Calzetti $\rightarrow$ SMC	+0.05	−0.09
Total (quadrature)	...	± 1.91	± 1.22

Note. ML = magnitude-limited mode; Int = intrinsic mode. Baseline accuracies are 87.7% (ML) and 89.8% (Int) using a simplified 2-model ensemble (XGBoost + RF, 3-fold CV). η : radiative efficiency; $f_{\text{Type 2}}$: obscured AGN fraction; $E(B - V)$ scale: mean dust reddening for Type 2 AGN; $(M/L)_{\text{SF}}$: stellar mass-to-light ratio for star-forming hosts; Extinction: dust law applied to host galaxy. The quadrature sum of the maximum absolute deviation per parameter quantifies the systematic uncertainty.




Figure 9. SIMBA classification in magnitude-limited mode. *Top Left:* ROC curves. *Bottom Left:* Color-color distribution by class. *Bottom Right:* $M_{\text{BH}}-M_*$ relation mapping.




Figure 10. SIMBA classification in intrinsic mode. Similar to Fig. 9, the ensemble achieves 93.5% accuracy ($\text{AUC} = 0.983$).

the available spectral features and thus the information content available to the classifier.

Furthermore, perturbing the $M_{\text{BH}}-M_{*}$ label boundary from the 40th to 60th percentile (Table 7) impacts intrinsic mode accuracy by less than 2%, confirming the signal is robust to a specified boundary. As physically expected, magnitude-limited accuracy drops at higher percentiles due to limits preferentially removing faint under-massive objects, significantly shrinking the minority class sample size (N).




Figure 11. Color distribution comparison between magnitude-limited (solid) and intrinsic (dashed) modes. While magnitude limits truncate the faint population, the color-space separation between under-massive (blue) and over-massive (red) classes is robustly preserved.

5.4. Limitations and Future Directions

Because SIMBA, ILLUSTRISTNG, and EAGLE employ heavy seed prescriptions, our classifiers distinguish growth paths within heavy-seed paradigms. Future work extending this framework to simulations featuring explicit light-seed channels (e.g., BRAHMA; A. K. Bhowmick et al. 2025) is needed to test direct early-universe seed-origin discrimination. Additionally, incorporating spectral population synthesis codes and realistic photometric redshift uncertainties will further refine the pipeline for immediate LSST data deployment.

6. CONCLUSION

We developed a machine learning framework to connect synthetic AGN observations from cosmological simulations with underlying black hole growth histories. Our primary conclusions are:

- A purely photometric pipeline successfully distinguishes over-massive from under-massive growth regimes, achieving $\sim 91\%$ – 94% accuracy under realistic LSST magnitude limits across both SIMBA and ILLUSTRISTNG.

- Signal decomposition definitively rules out trivial mass-to-luminosity mapping. The classification is driven by genuine physical signatures: host galaxy color correlations (82%–87% accuracy from host photometry alone) and accretion-state SED shape variations at fixed black hole mass.
- The relative photometric ordering of growth regimes is physically robust. Classifiers trained on one simulation and evaluated on another using rank-normalized features achieve 83%–89% accuracy in intrinsic mode, confirming the universality of the signal across diverse subgrid physics.
- The framework maintains high fidelity across redshift bins ($z < 0.5$ to $z > 5$). While intrinsic physical ordering transfers well, LSST flux limits induce distribution shifts that degrade cross-simulation transfer (to 74%–82%), indicating that future applications to real survey data will benefit from joint training on multiple simulations or explicit domain adaptation techniques.

This work provides a foundational framework for interpreting upcoming LSST observations through the lens of simulation-informed machine learning.

APPENDIX

A. EAGLE CROSS-VALIDATION

We applied our pipeline to nine EAGLE configurations (J. Schaye et al. 2015; R. A. Crain et al. 2015; S. McAlpine et al. 2016). Intrinsic classification achieves 96.2–98.2% accuracy across all variants, confirming that photometric growth-regime signatures persist across different subgrid physics implementations. However, the magnitude-limited mode fails for EAGLE: its conservative Bondi-limited accretion produces systematically lower accretion rates ($\dot{M}/M_{\text{BH}} \sim 10^{-14}$, compared to $\sim 10^{-10}$ in ILLUSTRISTNG), yielding AGN too faint to survive LSST magnitude cuts and reducing the observable sample below the training threshold. We note that the universal $M_{\text{BH}}-M_*$ relation applied here (fitted to SIMBA) produces a split in EAGLE that largely tracks whether black holes grew beyond the seed mass ($1.475 \times 10^5 M_\odot$), a less nuanced distinction than in SIMBA or ILLUSTRISTNG where BHs span a wider dynamic range above their respective seed masses. Fitting an EAGLE-specific relation (cf. J. Schaye et al. 2015, Fig. 10) could yield more physically meaningful class boundaries and is deferred to future work.

B. RESOLUTION-ARTIFACT RUNS (TNG100-2, TNG300-2)

The lower-resolution TNG variants TNG100-2 and TNG300-2 achieve near-perfect classification accuracy ($> 99\%$) due to coarse mass discretization rather than genuine physical separability. At these resolutions, black hole masses are assigned in large discrete steps, producing artificially well-separated populations that are trivially distinguishable by any classifier. These results are therefore not representative of the framework’s scientific performance and are presented here for completeness only.

ACKNOWLEDGMENTS

We thank the SIMBA team for making their simulation data publicly available. We gratefully acknowledge the IllustrisTNG collaboration for providing public access to simulation data via the TNG API and the EAGLE team for

Table 7. Label Boundary Sensitivity

Mode	Percentile	Split (dex)	N	Accuracy (%)	AUC
Mag-limited	p40	+1.29	74,374	87.6 ± 1.5	0.928
	p45	+1.34	73,582	86.0 ± 2.0	0.911
	p50 (Fiducial)	+1.39	50,424	85.0 ± 2.0	0.903
	p55	+1.43	35,760	83.3 ± 2.1	0.890
	p60	+1.48	26,520	82.8 ± 2.5	0.887
Intrinsic	p40	+1.29	110,266	90.0 ± 0.7	0.951
	p45	+1.34	132,168	90.0 ± 1.0	0.950
	p50 (Fiducial)	+1.39	129,172	90.3 ± 0.7	0.952
	p55	+1.43	108,654	90.7 ± 0.9	0.953
	p60	+1.48	94,358	91.7 ± 0.9	0.957

Note. Accuracies reflect a single XGBoost classifier on 12 photometric features used as a lightweight diagnostic; the relative trend across percentiles, not the absolute values, is the relevant result. The $\sim 7\%$ offset from the headline ensemble accuracies (Table 4) arises from using one model instead of four and fewer engineered features. The split column gives the $M_{\text{BH}}-M_*$ offset value (dex from the simulation-fitted relation; Section 3.3) at which objects are divided into over- and under-massive classes. The fiducial p50 split corresponds to the median offset used throughout this work.

Table 8. Classification Results for Resolution-Limited TNG Runs

Simulation Volume	Mode	N	Best Model	Accuracy	AUC
TNG100-2	Magnitude-Limited	6,658	Ensemble	$99.8\% \pm 0.1\%$	1.000
TNG100-2	Intrinsic	20,146	Ensemble	$99.8\% \pm 0.0\%$	1.000
TNG300-2	Intrinsic	98,664	Ensemble	$99.3\% \pm 0.1\%$	1.000

Note. Near-perfect accuracy is a resolution artifact: coarse BH mass discretization produces trivially separable populations. These runs are excluded from the primary analysis.

their public data releases. The IllustrisTNG simulations were undertaken with compute time awarded by the Gauss Centre for Supercomputing (GCS) under GCS Large-Scale Projects GCS-ILLU and GCS-DWAR on the GCS share of the supercomputer Hazel Hen at the High Performance Computing Center Stuttgart (HLRS), as well as on the machines of the Max Planck Computing and Data Facility (MPCDF) in Garching, Germany. We acknowledge the Virgo Consortium for making the EAGLE simulation data available. The EAGLE simulations were performed using the DiRAC-2 facility at Durham, managed by the ICC, and the PRACE facility Curie based in France at TGCC, CEA, Bruyères-le-Châtel. This work used data from LSST Data Preview 1 (NSF-DOE Vera C. Rubin Observatory 2025), obtained during Rubin Observatory commissioning. This material is based upon work supported in part by the National Science Foundation through Cooperative Agreement AST-1258333 and Cooperative Support Agreement AST-1202910 managed by the Association of Universities for Research in Astronomy (AURA), and the Department of Energy under Contract No. DE-AC02-76SF00515 with the SLAC National Accelerator Laboratory managed by Stanford University. We thank Ryan Farber for mentorship and guidance throughout this work.

DATA AVAILABILITY

The SIMBA simulation data are publicly available at <http://simba.roe.ac.uk/> (R. Davé et al. 2019). The IllustrisTNG simulation data are publicly available at <https://www.tng-project.org/data/> (D. Nelson et al. 2019a). The EAGLE simulation data are publicly available through the EAGLE database at <http://icc.dur.ac.uk/Eagle/database.php> (S. McAlpine et al. 2016). Rubin Observatory Data Preview 1 is available to data rights holders via the Rubin Science Platform (NSF-DOE Vera C. Rubin Observatory 2025). The classification pipeline code and derived data products generated in this work will be shared on reasonable request to the corresponding author.

Facilities: Rubin (LSSTCam)

Software: Python, NumPy, Pandas, Matplotlib, scikit-learn, PyTorch, XGBoost, Astropy ([The Astropy Collaboration et al. 2022](#)), h5py

REFERENCES

- Anglés-Alcázar, D., Faucher-Giguère, C.-A., Quataert, E., et al. 2017, *MNRAS*, 472, L109, doi: [10.1093/mnras/slx161](#)
- Begelman, M. C., Volonteri, M., & Rees, M. J. 2006, *MNRAS*, 370, 289, doi: [10.1111/j.1365-2966.2006.10467.x](#)
- Bell, E. F., & de Jong, R. S. 2001, *ApJ*, 550, 212, doi: [10.1086/319728](#)
- Bhowmick, A. K., Blecha, L., Torrey, P., et al. 2025, *MNRAS*, 538, 518, doi: [10.1093/mnras/staf269](#)
- Bluck, A. F. L., Piotrowska, J. M., & Maiolino, R. 2023, *ApJ*, 944, 108, doi: [10.3847/1538-4357/acac7c](#)
- Bogdán, A., Goulding, A. D., Natarajan, P., et al. 2024, *Nature Astronomy*, 8, 126, doi: [10.1038/s41550-023-02111-9](#)
- Bromm, V., & Larson, R. B. 2004, *ARA&A*, 42, 79, doi: [10.1146/annurev.astro.42.053102.134034](#)
- Bromm, V., & Loeb, A. 2003, *ApJ*, 596, 34, doi: [10.1086/377529](#)
- Calzetti, D., Armus, L., Bohlin, R. C., et al. 2000, *ApJ*, 533, 682, doi: [10.1086/308692](#)
- Crain, R. A., Schaye, J., Bower, R. G., et al. 2015, *MNRAS*, 450, 1937, doi: [10.1093/mnras/stv725](#)
- Davé, R., Anglés-Alcázar, D., Narayanan, D., et al. 2019, *MNRAS*, 486, 2827, doi: [10.1093/mnras/stz937](#)
- Devecchi, B., & Volonteri, M. 2009, *ApJ*, 694, 302, doi: [10.1088/0004-637X/694/1/302](#)
- Dubois, Y., Pichon, C., Welker, C., et al. 2014, *MNRAS*, 444, 1453, doi: [10.1093/mnras/stu1227](#)
- Fan, X., Bañados, E., & Simcoe, R. A. 2023, *ARA&A*, 61, 373, doi: [10.1146/annurev-astro-052920-102455](#)
- Fan, X., Strauss, M. A., Becker, R. H., et al. 2006, *AJ*, 132, 117, doi: [10.1086/504836](#)
- Gordon, K. D., Clayton, G. C., Misselt, K. A., Landolt, A. U., & Wolff, M. J. 2003, *ApJ*, 594, 279, doi: [10.1086/376774](#)
- Gorishniy, Y., Rubachev, I., Khrulkov, V., & Babenko, A. 2021, in *Advances in Neural Information Processing Systems*, Vol. 34, 18932–18943
- Habouzit, M., Li, Y., Somerville, R. S., et al. 2021, *MNRAS*, 503, 1940, doi: [10.1093/mnras/stab496](#)
- Habouzit, M., Somerville, R. S., Li, Y., et al. 2022, *MNRAS*, 509, 3015, doi: [10.1093/mnras/stab3147](#)
- Harikane, Y., Zhang, Y., Nakajima, K., et al. 2023, *ApJ*, 959, 39, doi: [10.3847/1538-4357/ad029e](#)
- Hirano, S., Hosokawa, T., Yoshida, N., et al. 2014, *ApJ*, 781, 60, doi: [10.1088/0004-637X/781/2/60](#)
- Hopkins, P. F. 2015, *MNRAS*, 450, 53, doi: [10.1093/mnras/stv195](#)
- Hopkins, P. F., & Quataert, E. 2011, *MNRAS*, 415, 1027, doi: [10.1111/j.1365-2966.2011.18542.x](#)
- Inayoshi, K., Visbal, E., & Haiman, Z. 2020, *ARA&A*, 58, 27, doi: [10.1146/annurev-astro-120419-014455](#)
- Ivezić, v., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, 873, 111, doi: [10.3847/1538-4357/ab042c](#)
- Kocevski, D. D., Onoue, M., Inayoshi, K., et al. 2023, *ApJL*, 954, L4, doi: [10.3847/2041-8213/ace5a0](#)
- Kormendy, J., & Ho, L. C. 2013, *ARA&A*, 51, 511, doi: [10.1146/annurev-astro-082708-101811](#)
- Lodato, G., & Natarajan, P. 2006, *MNRAS*, 371, 1813, doi: [10.1111/j.1365-2966.2006.10801.x](#)
- Lupi, A., Colpi, M., Devecchi, B., Galanti, G., & Volonteri, M. 2014, *MNRAS*, 442, 3616, doi: [10.1093/mnras/stu1120](#)
- Lusso, E., Comastri, A., Simmons, B. D., et al. 2012, *MNRAS*, 425, 623, doi: [10.1111/j.1365-2966.2012.21513.x](#)
- Ma, W., Cui, W., Davé, R., Anglés-Alcázar, D., & Guo, H. 2025, *A&A*, 702, A221, doi: [10.1051/0004-6361/202555571](#)
- Madau, P. 1995, *ApJ*, 441, 18, doi: [10.1086/175332](#)
- Madau, P., & Rees, M. J. 2001, *ApJ*, 551, L27, doi: [10.1086/319848](#)
- Maiolino, R., Scholtz, J., Curtis-Lake, E., et al. 2024, *A&A*, 691, A145, doi: [10.1051/0004-6361/202347640](#)
- Marinacci, F., Vogelsberger, M., Pakmor, R., et al. 2018, *MNRAS*, 480, 5113, doi: [10.1093/mnras/sty2206](#)
- McAlpine, S., Helly, J., Schaller, M., et al. 2016, *Astronomy and Computing*, 15, 72, doi: [10.1016/j.ascom.2016.02.004](#)
- Milosavljević, M., Couch, S. M., & Bromm, V. 2009, *ApJL*, 696, L146, doi: [10.1088/0004-637X/696/2/L146](#)
- Naiman, J. P., Pillepich, A., Springel, V., et al. 2018, *MNRAS*, 477, 1206, doi: [10.1093/mnras/sty618](#)
- Natarajan, P., Pacucci, F., Ricarte, A., et al. 2024, *ApJL*, 960, L1, doi: [10.3847/2041-8213/ad0e76](#)
- Nelson, D., Pillepich, A., Springel, V., et al. 2018, *MNRAS*, 475, 624, doi: [10.1093/mnras/stx3040](#)
- Nelson, D., Springel, V., Pillepich, A., et al. 2019a, *Computational Astrophysics and Cosmology*, 6, 2, doi: [10.1186/s40668-019-0028-x](#)
- Nelson, D., Pillepich, A., Springel, V., et al. 2019b, *MNRAS*, 490, 3234, doi: [10.1093/mnras/stz2306](#)

- NSF-DOE Vera C. Rubin Observatory. 2025, Legacy Survey of Space and Time Data Preview 1, doi: [10.71929/rubin/2570308](https://doi.org/10.71929/rubin/2570308)
- Overzier, R. A., Bouwens, R. J., Illingworth, G. D., & Franx, M. 2006, ApJ, 648, L5, doi: [10.1086/507678](https://doi.org/10.1086/507678)
- Pacucci, F., & Loeb, A. 2024, ApJ, 964, 154, doi: [10.3847/1538-4357/ad3044](https://doi.org/10.3847/1538-4357/ad3044)
- Pacucci, F., Nguyen, B., Carniani, S., Maiolino, R., & Fan, X. 2023, ApJL, 957, L3, doi: [10.3847/2041-8213/ad0158](https://doi.org/10.3847/2041-8213/ad0158)
- Pillepich, A., Springel, V., Nelson, D., et al. 2018, MNRAS, 473, 4077, doi: [10.1093/mnras/stx2656](https://doi.org/10.1093/mnras/stx2656)
- Pillepich, A., Nelson, D., Springel, V., et al. 2019, MNRAS, 490, 3196, doi: [10.1093/mnras/stz2338](https://doi.org/10.1093/mnras/stz2338)
- Regan, J., & Volonteri, M. 2024, The Open Journal of Astrophysics, 7, doi: [10.33232/001c.123239](https://doi.org/10.33232/001c.123239)
- Richards, G. T., Fan, X., Newberg, H. J., et al. 2002, AJ, 123, 2945, doi: [10.1086/340187](https://doi.org/10.1086/340187)
- Richards, G. T., Lacy, M., Storrie-Lombardi, L. J., et al. 2006, ApJS, 166, 470, doi: [10.1086/506525](https://doi.org/10.1086/506525)
- Sassano, F., Schneider, R., Valiante, R., et al. 2021, MNRAS, 506, 613, doi: [10.1093/mnras/stab1737](https://doi.org/10.1093/mnras/stab1737)
- Schaye, J., Crain, R. A., Bower, R. G., et al. 2015, MNRAS, 446, 521, doi: [10.1093/mnras/stu2058](https://doi.org/10.1093/mnras/stu2058)
- Schlaflly, E. F., & Finkbeiner, D. P. 2011, ApJ, 737, 103, doi: [10.1088/0004-637X/737/2/103](https://doi.org/10.1088/0004-637X/737/2/103)
- Shankar, F., Bernardi, M., Roberts, D., et al. 2025, MNRAS, 541, 2070, doi: [10.1093/mnras/staf747](https://doi.org/10.1093/mnras/staf747)
- Shen, Y., Richards, G. T., Strauss, M. A., et al. 2011, ApJS, 194, 45, doi: [10.1088/0067-0049/194/2/45](https://doi.org/10.1088/0067-0049/194/2/45)
- Springel, V. 2010, MNRAS, 401, 791, doi: [10.1111/j.1365-2966.2009.15715.x](https://doi.org/10.1111/j.1365-2966.2009.15715.x)
- Springel, V., Pakmor, R., Pillepich, A., et al. 2018, MNRAS, 475, 676, doi: [10.1093/mnras/stx3304](https://doi.org/10.1093/mnras/stx3304)
- Steffen, A. T., Strateva, I., Brandt, W. N., et al. 2006, AJ, 131, 2826, doi: [10.1086/503627](https://doi.org/10.1086/503627)
- Tanaka, T., & Haiman, Z. 2009, ApJ, 696, 1798, doi: [10.1088/0004-637X/696/2/1798](https://doi.org/10.1088/0004-637X/696/2/1798)
- The Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, ApJ, 935, 167, doi: [10.3847/1538-4357/ac7c74](https://doi.org/10.3847/1538-4357/ac7c74)
- Vanden Berk, D. E., Richards, G. T., Bauer, A., et al. 2001, AJ, 122, 549, doi: [10.1086/321167](https://doi.org/10.1086/321167)
- Vestergaard, M., & Peterson, B. M. 2006, ApJ, 641, 689, doi: [10.1086/500572](https://doi.org/10.1086/500572)
- Volonteri, M. 2010, Astronomy and Astrophysics Reviews, 18, 279, doi: [10.1007/s00159-010-0029-x](https://doi.org/10.1007/s00159-010-0029-x)
- Volonteri, M., Habouzit, M., & Colpi, M. 2021, Nature Reviews Physics, 3, 732, doi: [10.1038/s42254-021-00364-9](https://doi.org/10.1038/s42254-021-00364-9)
- Weinberger, R., Springel, V., Pakmor, R., et al. 2018, MNRAS, 479, 4056, doi: [10.1093/mnras/sty1733](https://doi.org/10.1093/mnras/sty1733)
- Wolpert, D. H. 1992, Neural Networks, 5, 241, doi: [10.1016/S0893-6080\(05\)80023-1](https://doi.org/10.1016/S0893-6080(05)80023-1)
- Zhang, H., Behroozi, P., Volonteri, M., et al. 2023, MNRAS, 518, 2123, doi: [10.1093/mnras/stac2633](https://doi.org/10.1093/mnras/stac2633)
- Zhang, H., Behroozi, P., Volonteri, M., et al. 2024, MNRAS, 531, 4974, doi: [10.1093/mnras/stae1447](https://doi.org/10.1093/mnras/stae1447)