

287,872 Supermassive Black Holes Masses: Deep Learning Approaching Reverberation Mapping Accuracy

Yuhao Lu^{1a,b}, HengJian SiTu^{2c}, Jie Li^{3c}, Yixuan Li^{4c,b}, Yang Liu^{5a,b,d}, Wenbin Lin^{6a,c,e} and Yu Wang^{7f,b,e,g,*}

^a*School of Computer Science, University of South China, Hengyang, 421001, China*

^b*ICRANet-AI, Brickell Avenue 701, Miami, FL 33131, USA*

^c*School of Mathematics and Physics, University of South China, Hengyang, 421001, China*

^d*Department of Physics E. Pancini, University Federico II, Naples, 80126, Italy*

^e*ICRANet, Piazza della Repubblica 10, Pescara, 65122, Italy*

^f*ICRA and Dipartimento di Fisica, Sapienza Università di Roma, P.le Aldo Moro 5, Rome, 00185, Italy*

^g*INAF – Osservatorio Astronomico d'Abruzzo, Via M. Maggini snc, Teramo, I-64100, Italy*

ARTICLE INFO

Keywords:

supermassive black holes
quasars
machine learning
black hole mass estimation
SDSS-RM

ABSTRACT

We present a population-scale catalogue of 287,872 supermassive black hole masses with high accuracy. Using a deep encoder-decoder network trained on optical spectra with reverberation-mapping (RM) based labels of 849 quasars and applied to all SDSS quasars up to $z = 4$, our method achieves a root-mean-square error of 0.058 dex, a relative uncertainty of $\approx 14\%$, and coefficient of determination $R^2 \approx 0.91$ with respect to RM-based masses, far surpassing traditional single-line virial estimators. Notably, the high accuracy is maintained for both low ($< 10^{7.5} M_\odot$) and high ($> 10^9 M_\odot$) mass quasars, where empirical relations are unreliable.

1. Introduction

Supermassive black holes (SMBHs) with masses spanning from roughly $10^5 M_\odot$ to $10^{10} M_\odot$ are commonly observed at the centers of most massive galaxies (Kormendy and Richstone, 1995; Ferrarese and Ford, 2005; Kormendy and Ho, 2013). Recent breakthroughs, particularly the Event Horizon Telescope's imaging of the SMBH at the core of the elliptical galaxy M 87, have provided unprecedented direct observational evidence (Event Horizon Telescope Collaboration et al., 2019). It is now firmly established that SMBH masses are strongly correlated with the characteristics of their host galaxies, including bulge mass, stellar velocity dispersion, surface brightness, and luminosity (Ferrarese and Merritt, 2000; Merritt and Ferrarese, 2001; Häring and Rix, 2004; Saglia et al., 2016). These correlations appear to persist across both local and high-redshift galaxies (Graham and Scott, 2013; Schramm and Silverman, 2013; Izumi et al., 2019), suggesting a fundamental co-evolutionary link despite the vast difference in physical scales between SMBHs and their hosts (Hopkins et al., 2008; Schawinski et al., 2010; Izumi et al., 2019).

Nevertheless, significant challenges remain, particularly in understanding how such massive black holes could have formed within the universe's first billion years (Wu et al., 2015; Inayoshi et al., 2020). Current models suggest that SMBHs grow predominantly through gas accretion and galaxy mergers, releasing substantial energy that profoundly affects host galaxy evolution (Alexander and Hickox, 2012; Ciotti and Ostriker, 2007; Sijacki et al., 2007).

Observationally, SMBHs manifest as active galactic nuclei (AGNs) or quasars, whose extreme luminosities offer key insights into accretion processes and black hole growth (Soltan, 1982). However, directly measuring SMBH masses remains challenging. Spectroscopic techniques, though widely used, are labor-intensive and have yielded only about one million estimates over the past two decades (Shen et al., 2011; Kelly and Shen, 2013). The advent of large-scale surveys such as the Vera C. Rubin Observatory's Legacy Survey of Space and Time (LSST), expected to detect nearly 10^8 quasars (Ivezić et al., 2019), will require far more efficient and scalable mass estimation methods. Alternative approaches based on variability measurements in optical and X-ray bands show promise (McHardy et al., 2006; Burke et al., 2021), but are complicated by nonlinear physical dependencies and the massive data volumes involved. Meanwhile, direct dynamical measurements of SMBH masses remain limited to only a small sample of nearby galaxies (Kuo et al., 2011; Kormendy and Ho, 2013; McConnell and Ma, 2013; Shankar et al., 2016, 2019).

While classical machine learning techniques, such as symbolic regression, random forests, and photometric regressions, have contributed to early progress in black hole mass estimation by extending traditional scaling relations (Jin and Davis, 2023; He et al., 2022), their performance remains fundamentally constrained by shallow architectures and limited feature representations. In contrast, deep learning approaches have shown greater promise in capturing the non-linear dependencies inherent in high-dimensional astrophysical data. For instance, variability-based models such as AGNet achieve a scatter of approximately 0.37 dex

*Corresponding author

✉ lwb@usc.edu.cn (W. Lin); yu.wang@icranet.org (Y. Wang)

ORCID(s):

with respect to virial mass estimates derived from single-epoch spectra (Lin et al., 2023). On the other hand, models trained on X-ray reverberation-mapped AGNs, where more physically grounded mass measurements are available, have demonstrated predictive accuracies within 2-5% of the reference values (Chainakun et al., 2022). Although these performance metrics are not directly comparable due to differing calibration standards, they collectively highlight the advantages of deep learning over traditional approaches.

To further advance black hole mass estimation, we propose a deep learning-based framework, QuasarSpecNet, which integrates a convolutional neural network (CNN) and a transposed convolutional network (DeconvNet). The model utilizes an encoding-decoding architecture to process spectral data, and performs mass prediction through representations learned in a latent space. Through joint training, the model not only enables efficient reconstruction of the input spectra but also allows accurate prediction of black hole masses based on latent features. The architecture is designed to leverage the hierarchical feature extraction capabilities of deep convolutional layers, while incorporating skip connections to retain fine-grained spectral details during decoding. This design improves predictive precision and more effectively preserves physical information. Additionally, it introduces a novel approach for estimating black hole masses and demonstrates superior performance in processing and analyzing spectroscopic data. The method achieves a root-mean-square error of 0.05 dex in $\log(M_{\text{BH}})$.

A central aspect of this work is that the network is trained on SDSS-RM spectra with reverberation-mapped black hole masses as labels, which provide a relatively small (849 quasars) but physically accurate sample. After training, the model is applied to an extended SDSS sample of 287,872 quasars spanning $0 \lesssim z \lesssim 4$, producing mass estimates across $\log_{10}(M_{\text{BH}}/M_{\odot}) \approx 6-10$. The combination of RM-based training and large-scale application is the main advance of this study. It provides a uniform mass catalog large enough to trace black hole demographics over cosmic time, while combining the accuracy of RM calibration with the statistical power of survey-scale quasar populations.

This article is structured as follows: Section 2 describes the dataset and preprocessing steps; Section 3 details the proposed methodology; Section 4 presents model evaluation metrics; Section 5 applies the trained model to 287,872 SDSS quasars and presents the resulting mass catalog; Section 6 discusses the findings; finally, Section 7 provides conclusions and acknowledgments.

2. Dataset

2.1. Extraction of Observational Parameters

For the SDSS-RM sample, each quasar was observed repeatedly over multiple epochs, producing time-series spectra that capture the variability of both continuum and broad emission lines. Reverberation mapping relies on measuring the delayed response of lines such as $\text{H}\alpha$, $\text{H}\beta$, Mg II , and CIV to continuum changes. Each emission line is recorded

across many epochs, enabling variability studies and lag determinations. As a result, the dataset contains both single-epoch observational parameters and multi-epoch spectral sequences.

Observational metadata are extracted from the headers of the raw spectral files. Key parameters including the spectroscopic plate number (PLATE), observation date (MJD), fiber identifier (FIBERID), mean observation epoch (Mean_MJD), redshift (ZFINAL), and observation epoch index (EPOCH) are extracted using regular expressions. These parameters were organized into structured records to facilitate sample selection, redshift dependent analyses, and cross-validation of spectral line measurements. Table 1 summarizes these observational parameters.

The SDSS-RM sample includes 849 unique sources with 75,915 spectra, spanning redshifts from $z = 0.12$ to $z = 4.33$ (mean $z = 1.64$, median $z = 1.61$). Only three sources (270 spectra) lie at $z \geq 4$, indicating that the sample is dominated by low- to intermediate-redshift AGNs.

2.2. Data Preprocessing and Standardization

To provide consistent and physically meaningful input spectra for the deep learning analysis, we preprocess quasar spectra obtained from the SDSS-RM project (Shen et al., 2011; Lyke et al., 2020; Shen et al., 2024). This preprocessing procedure involves extraction of key observational parameters, resampling of spectra to a uniform wavelength grid, exclusion of invalid or corrupted spectra, spectral normalization, and the compilation of a standardized dataset suitable for both training and evaluation of our model. Each preprocessing step is described in detail below.

2.3. Wavelength Resampling

The original spectral flux data points are sampled unevenly and lie on a non-uniform wavelength grid. To standardize the model input, the wavelengths were first transformed into logarithmic space and then interpolated onto a uniform grid of length 4618 (denoted as Wave1):

$$\lambda_i^* = \lambda_{\min} + i \cdot \Delta\lambda, \quad (1)$$

$$\Delta\lambda = \frac{\lambda_{\max} - \lambda_{\min}}{4617}, \quad (2)$$

where i ranges from 0 to 4617, and the wavelength range corresponds to 360–1032.5 nm (vacuum wavelength). Cubic spline interpolation was employed as the primary method, with linear extrapolation used at the boundaries. The resulting flux values were stored as Flux1 , with units preserved as $10^{-17} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ \AA}^{-1}$.

2.4. Invalid Spectrum Filtering

To remove spectra that contain no meaningful information or lack typical quasar spectral features, we applied the following criteria:

Table 1
Description of dataset fields

Field	Description
Defined_Mass	Adopted BH mass ($\log(M_{\odot})$), taking the first positive value from HA_MRM, HB_MRM, MG2_MRM, C4_MRM; set to -1 if none.
Flux	Raw observed flux on non-uniform grid ($10^{-17} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ \AA}^{-1}$).
Flux1	Flux interpolated onto uniform Wave1 grid (same units as Flux).
HA_MRM	RM-derived BH mass using H α line lag ($\log(M_{\odot})$).
HB_MRM	RM-derived BH mass using H β line lag ($\log(M_{\odot})$).
MG2_MRM	RM-derived BH mass using Mg II line lag ($\log(M_{\odot})$).
C4_MRM	RM-derived BH mass using C IV line lag ($\log(M_{\odot})$).
Mean_MJD	Mean Modified Julian Date of spectroscopic visits.
Norm_Flux	Min-max normalized Flux1 (unitless).
RMID	Unique ID in SDSS-RM catalog (0-999).
EPOCH	Index of spectroscopic observation epoch.
Rest_Flux	Rest-frame flux: $\text{Flux}/(1+z)$.
Rest_Flux_Normalized	Min-max normalized rest-frame flux.
Rest_Wave	Rest-frame wavelength: $\lambda_{\text{obs}}/(1+z)$ ($\text{\AA}$).
Wave	Observed vacuum wavelength ($\text{\AA}$).
Wave1	Uniform interpolated wavelength grid (model input).
ZFINAL	Adopted redshift of the object.
Name	SDSS-DR12 designation (J2000, hhmmss.ss+ddmmss.s).

1. **Flat Spectra:** If the flux vector $\mathbf{f} \in \mathbb{R}^{4618}$ exhibits no variation across the entire wavelength range:

$$\text{ptp}(\mathbf{f}) = \max(\mathbf{f}) - \min(\mathbf{f}) = 0, \quad (3)$$

the spectrum is considered static and physically uninformative, and is thus discarded.

2. **Quasi-linear Spectra:** After resampling, if the spectrum can be approximated by a first-order polynomial $f(i) = ai + b$ such that the maximum residual across all wavelengths satisfies:

$$\max_i |f_i - (ai + b)| < \epsilon, \quad (4)$$

it is deemed to lack spectral line structures and is removed accordingly. It is noted that $\epsilon = 10^{-5}$ that we take is a very small threshold, ensuring that spectra with even weak or narrow emission-line features are retained.

2.5. Spectral Normalization

To eliminate flux amplitude differences among different targets, all retained spectra were normalized to the $[0, 1]$ range using min-max scaling. Given the interpolated flux vector $\hat{\mathbf{f}} \in \mathbb{R}^{4618}$, the normalization is defined as:

$$\tilde{\mathbf{f}} = \frac{\hat{\mathbf{f}} - \min(\hat{\mathbf{f}})}{\max(\hat{\mathbf{f}}) - \min(\hat{\mathbf{f}})}. \quad (5)$$

To further improve model training stability, we also applied zero-mean normalization and unit-norm constraint. The final processed spectra are denoted as Norm_Flux.

2.6. Black Hole Mass Label Construction

To facilitate supervised learning applications, we adopt reverberation mapped black hole mass estimates from the SDSS-RM spectral analysis pipeline. These RM-based masses, obtained from the emission lines H α , H β , Mg II, and C IV, are represented as HA_MRM, HB_MRM, MG2_MRM, and C4_MRM, respectively. For each sample, a unified mass label, Defined_Mass, was determined by selecting the first valid estimate according to the priority sequence: C4_MRM, MG2_MRM, HB_MRM, and HA_MRM. If all lines mass estimates are invalid, no defined mass is assigned, and such quasars are excluded from supervised training but remain available for unsupervised spectral reconstruction. These labels are subsequently linked to their corresponding spectra using RMID identifiers and integrated into the dataset.

2.7. Dataset Partitioning

To promote robust generalization and reliable performance evaluation, we split the dataset by RMID into training (70%), validation (15%), and test (15%) subsets. The reference set, containing quasars with reverberation-mapped black hole masses, was included only in the training and validation subsets and excluded from the test set to avoid information leakage.

3. Methodology

3.1. Neural Network

In this study, we propose a deep autoencoder-based neural network that concurrently performs feature extraction, spectral reconstruction, and black hole mass prediction. The proposed architecture comprises three interconnected modules: an encoder, a decoder, and a dedicated mass prediction branch, as illustrated in Figure 1.




Figure 1: Architecture of the proposed deep autoencoder-based neural network.

The encoder compresses input spectral data into a compact latent representation that encapsulates crucial spectral characteristics. Specifically, the input spectral data, structured as single-channel one-dimensional tensors of dimensions $(batch_size, 1, sequence_length)$, undergo processing by three convolutional blocks. Each block comprises convolutional layers that progressively increase the number of feature channels, followed by batch normalization (BatchNorm1d) for training stabilization, a Rectified Linear Unit (ReLU) activation function for nonlinear transformation, and a dropout layer to mitigate overfitting and enable uncertainty estimation. Subsequently, the resulting feature maps are flattened and projected into a lower-dimensional latent vector of dimension $latent_dim$ via a fully connected layer (encoder_fc). Additionally, the encoder incorporates a specialized branch for directly predicting black hole mass from these latent representations. This prediction branch involves a fully connected layer that reduces dimensionality to 64, followed by a ReLU activation, a dropout layer (with a dropout probability of 0.5), and finally, a fully connected layer outputting a scalar corresponding to the predicted black hole mass.

The decoder reconstructs the original spectral input from the latent features produced by the encoder. Initially, a fully connected layer expands the latent vector to match the encoder's output dimensions (128×73) . The expanded vector is subsequently reshaped by an Unflatten operation into the dimensions $(batch_size, 128, 73)$ suitable for convolutional processing. Reconstruction then proceeds through three transpose convolutional layers, progressively reducing

the number of channels and increasing spatial resolution. Critically, these transpose convolutional layers utilize skip connections by concatenating their outputs with corresponding encoder feature maps, preserving detailed spatial information, facilitating gradient flow, and preventing vanishing gradients. The transpose convolutions sequentially reduce channels from 128 to 64, 64 to 32, and finally 32 to 1, with kernel sizes, strides, and padding chosen to ensure output sequence length matches the original spectral data. After concatenation, cropping operations (min_length) maintain consistent feature map dimensions throughout reconstruction.

3.2. Physically Motivated Architecture

We construct the network from the physics of quasar spectra rather than from a purely accuracy-driven objective. Single epoch mass estimators follow the virial form $M_{BH} \propto V^2 R(L)$, where V is a velocity scale of the broad-line region and $R(L)$ is a luminosity dependent size. This relation is a useful first-order guide, but its accuracy is limited when $R(L)$ is forced to be a simple power law and when line-dependent effects such as profile asymmetry, outflows, and orientation are ignored. Our design goal is to encode the information that enters $V^2 R(L)$ and, at the same time, allow the model to use additional spectral indicators.

The architecture in Fig. 1 is an encoder-decoder with skip connections. The encoder compresses each spectrum into a low dimensional latent vector intended to carry global, slowly varying information. The decoder reconstructs the

spectrum using both the latent representation and intermediate encoder features passed through the skips. The mass regression branch reads only the latent vector, ensuring that the predicted M_{BH} depends on physically interpretable global features rather than on local variations.

Input spectra are min-max scaled and standardized to zero mean and unit variance. Normalization removes the absolute flux scale, so the network cannot recover luminosity from amplitude alone. It must infer luminosity from dimensionless features such as continuum shape, line ratios, and equivalent widths, which are known to correlate with ionizing power (for example, through the radius–luminosity relation and Baldwin-type trends). This choice is deliberate: it encourages the encoder to represent physically invariant relations between continuum and lines rather than memorize survey-specific flux levels.

We limit the latent space to be as small as practical to preserve interpretability and control. Quasar spectra are set by a few effective degrees of freedom: a luminosity or radius scale, a velocity scale, the Eddington ratio, broad spectral slope or curvature, and metal-line strength or ionization state. A compact latent vector makes these factors explicit and limits the tendency to spread information across many entangled components. With skip connections, the reconstruction loss no longer forces the latent vector to carry high frequency line details, because the shortest gradient path for reconstruction runs through shallow layers and the skips. The latent representation is therefore shaped mainly by the mass loss and concentrates on slowly varying features. Since the mass head reads only the latent vector, it naturally learns a combination of a luminosity or radius scale with a coarse velocity summary rather than a catalogue of profile minutiae. This aligns with the virial structure of the problem, in which $R(L)$ is a global quantity and V can be summarized by low-order width measures, and it keeps the inferred M_{BH} tied to a physically interpretable global representation rather than to local variability driven by variabilities.

3.3. Model Training and Optimization

We adopt a joint optimization approach, simultaneously minimizing two mean squared error (MSE) loss functions: one for spectral reconstruction and the other for black hole mass estimation. These loss functions are combined with equal weights (1:1) and optimized using the Adam algorithm with an initial learning rate of 1×10^{-4} and a batch size of 64. The maximum number of training epochs is set to 500. At each epoch, model performance is evaluated on both the training and validation sets. An early stopping mechanism, with a patience of 50 epochs based on the validation loss, is applied to mitigate overfitting, and the model weights corresponding to the best validation performance are retained.

4. Evaluation

4.1. Effectiveness of neural network for feature extraction

To evaluate the robustness and interpretability of the learned spectral representations, we trained the autoencoder

using different latent space dimensionalities. Varying the latent dimension provides a way to quantify how much spectral-mass information can be retained with minimal model complexity. A lower-dimensional latent space favors physical interpretability, since each latent variable tends to encode broad and physically meaningful spectral variations. Increasing the dimensionality can capture finer spectral details but may also fit stochastic noise and obscure clear physical meaning. Therefore, we aim to achieve an accurate black hole mass estimation with the smallest possible number of latent variables, balancing predictive performance and interpretability. Technically, we employed Uniform Manifold Approximation and Projection (UMAP), a dimensionality reduction technique that maps high-dimensional data to a two-dimensional plane (McInnes et al., 2018). Specifically, we applied UMAP to project latent representations of dimensions 1, 2, 4, 5, 10, and 50, and visualized the resulting projections by coloring data points according to their normalized logarithmic virial masses ($\log M_{\text{vir}}$), as presented in Figure 2.

The results shown in Fig. 2 visualize how the learned latent representation organizes spectra as the latent dimensionality varies. Points are colored by $\log M_{\text{BH}}$, and a smooth ordering with $\log M_{\text{BH}}$ is visible at all settings.

At one latent dimension, the embedding collapses to a single arc with a monotonic $\log M_{\text{BH}}$ gradient, showing that the dominant variation relevant to mass is already captured. At two dimensions, the map retains a principal ridge with a clear mass gradient and exhibits transverse structure, indicating that additional modes of spectral variation are represented in the 2D projection. At four and five dimensions, the manifold appears thicker in the UMAP plane, the mass ordering remains continuous along the main ridge. This suggests that added capacity captures secondary structure, for example broad continuum shape or cross line correlations, without disrupting the primary ordering. When the latent dimensionality increases to ten, the global pattern and mass gradient are preserved and the cloud becomes slightly more compact in the UMAP plane, consistent with extra capacity absorbing secondary variance rather than changing the core geometry. At fifty dimensions, the central structure and mass ordering persist, but a small fraction of points is displaced toward the periphery as scattered outliers, which likely correspond to unusual line shapes or lower signal to noise. The scattering reflects overcapacity that allows peripheral directions in latent space to be populated, not a fragmentation of the main manifold. The above analysis supports that five latent representations provide a practical balance between representation power and interpretability.

For evaluate the model's fitting performance, we adopt two metrics: the root mean squared error (RMSE) and the Pearson correlation coefficient r . The RMSE is given by

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (\hat{y}_i - y_i)^2}, \quad (6)$$




Figure 2: UMAP projections of the learned latent representation for latent dimensionality $d = 1, 2, 4, 5, 10$, and 50 . Points are colored by $\log M_{\text{BH}}$. A smooth mass gradient is visible in all panels. For $d = 1$, the embedding collapses to a single arc with a monotonic gradient; for $d = 2$, a principal ridge with transverse structure appears; for $d = 4-5$, the manifold thickens while the gradient remains continuous; for $d = 10$, the global pattern and gradient are preserved with a slightly more compact cloud; for $d = 50$, the central manifold persists and a small fraction of points is displaced to the periphery. UMAP preserves local neighborhoods rather than global geometry, so axis orientation is not physically meaningful.

where $\hat{y}_i$ and y_i are the predicted and true values for sample i , respectively, and N is the total number of samples. A lower RMSE indicates better predictive performance.

The correlation coefficient r measures the strength of the linear relationship between predictions and ground truth values:

$$r = \frac{\sum_{i=1}^N (\hat{y}_i - \bar{\hat{y}})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^N (\hat{y}_i - \bar{\hat{y}})^2} \sqrt{\sum_{i=1}^N (y_i - \bar{y})^2}}, \quad (7)$$

where $\bar{\hat{y}}$ and $\bar{y}$ denote the means of the predicted and true values, respectively. A value of r closer to 1 indicates a stronger linear correlation.

The fitting performance results, as presented in Table 2, are characterized by an R^2_{train} of 0.9570 and an R^2_{val} of 0.9099. These values indicate strong predictive accuracy on the training data and robust generalization to unseen data. Regarding error metrics, the RMSE is 0.00310 for the training set and 0.05532 for the validation set, with an MAE of 0.03896 on the validation set. These metrics, which are among the most favorable in Table 2, demonstrate that the model achieves high accuracy in predicting black hole masses with the chosen five-dimensional latent space.

Therefore, considering both interpretability and predictive performance, we adopt five latent dimensions as the optimal configuration. The five-dimensional representation preserves a coherent mass gradient while effectively suppressing noise and non-physical spectral variations. Quantitative regression tests further support this choice, showing that increasing the dimensionality beyond 5 does not bring significant improvement in the mass fitting accuracy.

The physical analysis and the cosmological studies based on these results will be the main subjects of the subsequent papers.

4.2. Redshift-mass coverage and consistency with SDSS-RM

Figure 4 visualizes our predictions on the SDSS-RM sample as a mass-redshift heatmap overlaid on the RM-based masses. The two layers show close agreement across the full dynamic range, with both distributions peaking near $z \sim 1-2$ and spanning $\log_{10}(M_{\text{BH}}/M_{\odot}) \sim 7-10$. This pattern is consistent with the coverage and mass range reported in the SDSS-RM Key Results release, which summarizes the RM sample over $0 \lesssim z \lesssim 4$ and typical black-hole masses of $10^7-10^{9.5} M_{\odot}$ (Shen et al., 2011; Lyke et al., 2020;

Table 2

Model performance for different numbers of latent features. The table lists the coefficient of determination (R^2), root-mean-square error (RMSE), and mean absolute error (MAE) of training and validation datasets for different latent space dimensionalities. The results show that the predictive accuracy improves rapidly up to ~ 5 latent dimensions and then grows slowly, indicating that five dimensions provide the optimal balance between model complexity and performance.

Latent Dim	R^2_{train}	R^2_{val}	RMSE _{train}	RMSE _{val}	MAE _{val}
1	0.9265	0.8549	0.00512	0.07154	0.05329
2	0.9432	0.9075	0.00326	0.05712	0.03920
4	0.9495	0.9081	0.00324	0.05693	0.03911
5	0.9570	0.9099	0.00310	0.05532	0.03896
10	0.9579	0.9132	0.00306	0.05534	0.03768
50	0.9568	0.9142	0.00303	0.05503	0.03617




Figure 3: Comparison between black hole mass estimates from our autoencoder-based model and those obtained using traditional single-epoch virial methods for the test dataset. The blue points show the predictions from the neural-network model, while the purple, yellow, and green points correspond to single-line estimates based on $H\beta$, $Mg II$, and $C IV$, respectively. The dashed line marks the one-to-one relation expected for perfect agreement with the reverberation-mapping (RM) masses. The scatter of the blue points mainly reflects variations among individual spectra from the multi-epoch SDSS-RM observations, as each spectrum is treated independently in the estimation. Traditional estimators exhibit larger scatter and systematic deviations, particularly at the low- and high-mass ends, and many sources lack reliable single-line measurements. In contrast, our model maintains a tight correlation with the RM masses ($R^2 = 0.909$) and achieves a low overall RMSE of 0.058 dex with respect to the RM-based mass, and also providing predictions for objects where the conventional methods fail.

Supermassive Black Holes Mass




Figure 4: Heatmap of predicted versus RM-based black hole masses across different redshifts, with two overlaid layers: The blue colormap shows RM-based masses; the semi-transparent red colormap shows network-predicted masses.

Table 3

Summary of Black Hole Mass Prediction Methods. The table lists the RMSE (in dex) and R^2 evaluated against RM-based masses as the reference, reported in the respective studies.

Method	RMSE (dex)	R^2	Data
Single Epoch Virial(Vestergaard and Peterson, 2006)	0.35	–	Single epoch spectra
AGNet Deep Learning(Lin et al., 2023)	0.371	0.428	Photometric time series
X-ray Reverberation NN(Chainakun et al., 2022)	0.1077	0.9124	X-ray light curves
Lasso Regression(He et al., 2022)	0.50	–	Photometry + morphology
CQR(Yong and Ong, 2023)	0.198 (H β); 0.222 (Mg II)	0.7 (H β); 0.85 (Mg II)	SDSS single epoch spectra
QuasarSpecNet (This Work)	0.058	0.909	SDSS spectra

Shen et al., 2024). The agreement indicates that our network reproduces the established population-level trends in the SDSS-RM dataset without introducing apparent selection-driven distortions.

4.3. Comparison with Single-Line Virial Estimators

The model's predictions on the validation set are highly consistent with the RM based black hole masses, with the scatter plot (figure 3) tightly clustered around the ideal consistency line. In other words, the model does not introduce significant systematic biases and shows no evidence of overfitting-related biases, which strengthens our confidence in its generalizability.

Table 3 compares our method with traditional single-line virial estimators based on H β , Mg II, and C IV scaling relations, as well as with other machine learning approaches.

Conventional single-epoch relations typically generate scatterers of ~ 0.2 – 0.5 dex and are prone to systematic biases (Coatman et al., 2017), particularly for C IV-based estimates in high-accretion quasars. In contrast, our model achieves an overall RMSE of only 0.058 dex, representing almost an order-of-magnitude improvement in precision. Our model also achieves a high $R^2 > 0.9$, indicating reduced scatter and improved accuracy across the full sample.

The comparison between our autoencoder-based predictions and the traditional single-line virial estimators is shown in Figure 3. Our predictions (blue) closely follow the one-to-one relation across the full dynamic range in mass, demonstrating excellent agreement with the RM masses. In contrast, the single-line estimates (H β : purple; Mg II: yellow; C IV: green) exhibit larger systematic deviations at both the low- and high-mass ends, which is a known

problem in previous studies. At the high-mass end, C_{IV}-based estimates are strongly affected by non-virial outflows and blueshifts that distort the line width and centroid, leading to overestimation of the black hole mass (e.g., Denney et al., 2013; Coatman et al., 2017). At the low-mass end, the weaker and noisier broad lines, combined with uncertainties in continuum subtraction and line fitting, increase measurement scatter and bias (Denney et al., 2013; Shen et al., 2008). Additional differences arise from variations in the virial factor and broad-line region geometry across luminosity and accretion states (Mejía-Restrepo et al., 2017). Moreover, sample selection and luminosity-dependent Malmquist bias can further amplify these edge effects in flux-limited surveys (Shen et al., 2008).

A further advantage of our approach is its ability to recover reliable black hole masses even when prominent emission lines are weak, noisy, or fall outside the observed spectral range, where traditional virial estimators fail. For instance, H β -based estimates are unavailable for high-redshift quasars where the line is redshifted beyond the optical coverage of SDSS, and Mg_{II}-based estimates become inaccessible beyond $z \sim 2$ when the line shifts into the near-infrared (e.g., Shen et al., 2008). In contrast, C_{IV}-based estimates remain available at higher redshift but are strongly affected by line blueshifts and non-virial outflow components, leading to substantial scatter and systematic bias in the inferred masses (e.g., Denney et al., 2013; Coatman et al., 2017). These limitations highlight the advantage of our model, which leverages the full spectral information to obtain consistent mass estimates even in regimes where single-line methods are unreliable or inapplicable.

5. 287,872 Black Hole Mass

After training, the network can be applied to the large SDSS quasar sample. The SDSS-RM program is a dedicated subproject of SDSS that uses the same spectroscopic instrumentation and reduction pipeline as the main SDSS survey. The difference lies in the observing strategy: while the main SDSS typically provides a single-epoch spectrum for each quasar, SDSS-RM obtained repeated spectra over many epochs, enabling the measurement of emission-line lags and accurate reverberation-mapped black hole masses. Our model is trained on these RM-based labels, but once trained it requires only a single spectrum as input. This allows the network to achieve RM-level accuracy when applied to single-epoch spectra from the main SDSS quasar survey. In this work, we apply the trained model to 287,872 quasars of 377,549 spectra spanning $0 \lesssim z \lesssim 4$, providing a very large catalog of massive black hole mass.

The result is shown in Fig. 5, that the two dimensional distribution of black hole mass as a function of redshift. The inferred black hole masses cover a broad range from $10^6 M_\odot$ to $10^{10} M_\odot$. The overall trend agrees with the statistical behavior found in large optical quasar samples (e.g., Shen et al., 2011; Lyke et al., 2020), where the most luminous and massive quasars cluster around $z \sim 2-3$. The persistence of




Figure 5: SDSS DR16 quasar sample: hexbin map of predicted black-hole mass versus redshift, with a logarithmic color scale for counts. Two LOWESS curves are overplotted: solid for $\log M$ as a function of z , dashed for z as a function of $\log M$, tracing the central trend. Top and right panels show marginal histograms. The redshift is limited to $0 \leq z \leq 4$.




Figure 6: Distribution of the relative standard deviation ($\Delta M/M$) of black hole mass estimates for sources with multiple spectroscopic observations. The majority of sources show a relative uncertainty below $\sim 10\%$, indicating high internal consistency of the trained network.

high inferred M_{BH} beyond $z \approx 3$ is physically reasonable, since only the most massive and luminous accreting black holes can remain above the survey flux limit at those distances. The absence of a turnover suggests that the neural network estimator does not inherit the redshift dependent bias that affects virial masses based solely on C_{IV}, which often underpredict masses for high blueshift quasars (e.g., Coatman et al., 2017; Mejía-Restrepo et al., 2018). Instead, the model appears to combine continuum and multiband




Figure 7: Comparison between original and reconstructed spectra for four randomly selected sources. The reconstructed spectra (green) nearly coincide with the original observed spectra (orange) over the entire wavelength range, showing the network's high reconstruction accuracy. The residuals (gray) remain small except near the spectral edges, where the instrumental uncertainties dominate.

spectral information to maintain a consistent scaling of M_{BH} with redshift.

To evaluate the internal consistency of the model predictions, we analyzed sources with multiple independent spectroscopic observations. A total of 49,404 sources were repeatedly observed. For each source, the standard deviation and relative standard deviation of the predicted black hole masses were calculated, see Fig. 6. The model shows a median dispersion of 0.046 dex in black hole mass estimates, corresponding to a median relative uncertainty of 8.6%. This small internal scatter indicates that the model provides stable and consistent mass estimates across repeated observations, supporting the reliability of its predictions.

The four randomly selected spectra shown in Fig. 7 compare the original SDSS spectra with those reconstructed by the neural network. The reconstructed spectra almost completely overlap with the originals across most of the wavelength range, indicating that the model accurately reproduces both the continuum shape and emission-line structures. The small and structureless residuals confirm the high reconstruction fidelity of the model. The slightly larger discrepancies seen at the shortest and longest wavelengths are expected from the observational characteristics of the SDSS spectrographs. At the spectral edges, the instrumental response and CCD quantum efficiency decrease, the flux calibration becomes less reliable, and the SNR ratio drops sharply. These instrumental effects dominate the uncertainty




Figure 8: Differential black hole mass function for sources with $z < 4$. Blue steps show the binned estimate of dN/dM ; the green curve is the posterior median of SBPL fit obtained with a Poisson likelihood and MCMC sampling. The fit gives a break at $\log_{10}(M_b/M_{\odot}) = 8.52$ (vertical gray line), a low-mass logarithmic slope $\alpha_{\text{low}} = +1.1$ and a high-mass slope $\alpha_{\text{high}} = -2.7$.

in those regions, rather than the limitation of the neural network itself.

6. Discussion

6.1. Differential Mass function

We model the differential black hole mass function, dN/dM , see Fig. 8, for all sources with $z < 4$, using a smoothed broken power-law (SBPL) of the form

$$\frac{dN}{dM} = A \left[\left(\frac{M}{M_b} \right)^{p_1 s} + \left(\frac{M}{M_b} \right)^{p_2 s} \right]^{1/s}, \quad (8)$$

where A is the normalization, M_b is the break mass, p_1 and p_2 are the slopes below and above the break, and s controls the smoothness of the transition. The fitting is performed using a Poisson-binned maximum-likelihood estimator followed by an MCMC analysis to sample the posterior parameter distributions. While the total black hole mass function in inactive galaxies follows a Schechter-like form with an exponential cutoff (Schechter, 1976), the active population is shaped by the convolution of this total distribution with a mass-dependent duty cycle and a broad Eddington-ratio distribution. This convolution transforms the exponential cutoff into a smoother turnover, producing a double power-law form (Kollmeier et al., 2006; Greene and Ho, 2007; Schulze and Wisotzki, 2010).

The best-fit model gives a normalization of $\log_{10} A = -2.84$, a low-mass slope of $p_1 = 2.70$, a high-mass slope of $p_2 = -1.11$, a break mass of $\log_{10}(M_b/M_{\odot}) = 8.52$, and a smoothness parameter $s = 0.58$. In logarithmic space, the corresponding differential slopes are $\alpha_{\text{low}} = +1.1$ and $\alpha_{\text{high}} = -2.7$. Indicating roughly $\sim 10\%$ of the sample has mass more than $10^9 M_{\odot}$, and only $\sim 0.1\%$ exceeds $10^{10} M_{\odot}$.

The SBPL fit reveals a clear turnover at $\log_{10}(M_b/M_\odot) = 8.5$ ($M_b = 3 \times 10^8 M_\odot$). Below this mass, the rising slope reflects the increasing number of moderate-mass black holes, consistent with the expected population of typical quasars. At higher masses, the steep decline reflects the relative scarcity of the most massive quasars. A mild excess of sources below $M_{\text{BH}} \lesssim 10^{7.5} M_\odot$ likely originates from low-redshift, low-luminosity quasars that dominate flux-limited samples. Conversely, the observed deficit at $M_{\text{BH}} \gtrsim 10^{10} M_\odot$ is plausibly primarily due to a feedback-limited upper bound on black hole growth set by gas supply and radiation pressure, or/and selection incompleteness for extremely luminous objects near the bright end of the quasar luminosity function.

6.2. Broad and Narrow Line Quasars

RM requires measurable time delays between the variable continuum and the response of the broad emission lines that originate from the broad-line region (BLR). The BLR gas lies within light-days to light-weeks of the central ionizing source and exhibits Doppler-broadened velocities of several thousand km s^{-1} , producing the broad line profiles seen in Type 1 AGNs and quasars. In contrast, the narrow-line region (NLR) is located at much larger radii (tens to thousands of parsecs) and is photoionized by the central source but does not respond measurably to continuum variations on humanly observable timescales. Its line widths are only a few hundred km s^{-1} , reflecting the gravitational potential of the host galaxy rather than that of the black hole. As a result, narrow emission lines show little or no correlated variability with the continuum, making them unsuitable for RM-based black hole mass measurements. Therefore, the RM-based masses, including those in the SDSS-RM project, are derived from broad-line AGNs, for which the BLR variability can be directly observed and modeled (e.g., Blandford and McKee, 1982; Peterson, 1993; Shen et al., 2015, 2011; Shen et al., 2024).

In our framework, the RM-based black hole masses serve as the ground-truth labels for supervised training, while the model input consists of single-epoch spectra. Hence, the neural network does not rely on the temporal information required by reverberation mapping and, in principle, is not restricted to sources showing variable broad lines. However, since all RM measurements used for training originate from broad-line AGNs, the learned spectral representations inherently reflect the some physical conditions of the broad-line region. To assess possible extrapolation beyond this regime, we further divide all quasars into broad-line and narrow-line classes based on their emission-line widths. For broad-line quasars, the model performance on the previous validation set should directly reflect its accuracy. For narrow-line quasars, there is no prior benchmark for true RM-equivalent masses, and the model's reliability cannot be guaranteed a priori since the network is almost a black box. Nevertheless, given that the neural network captures, not only the emission lines features, but also the spectral details and global correlations, it is plausible that its predictions

for narrow-line quasars remain physically meaningful. Comparing the inferred masses between the broad- and narrow-line populations may therefore provide an empirical estimate of the model's generalization accuracy outside the RM-calibrated domain.

Fig. 9 shows the distribution of the predicted black hole mass as a function of redshift for the broad line and narrow line subsamples. Both panels share the same dynamic range in mass and redshift ($7 \leq \log(M_{\text{BH}}/M_\odot) \leq 10$ and $0 \leq z \leq 4$). The overall envelope of the two distributions is remarkably similar, indicating that the network assigns consistent mass scales to quasars regardless of their spectral classification. For both subsamples, the source density peaks at $z \approx 2.5$, corresponding to the epoch of maximum quasar activity, and the median mass increases from $\log M_{\text{BH}} \approx 7.5$ at low redshift to about $\log M_{\text{BH}} \approx 9.0$ near the peak.

At $z \lesssim 3$, the LOWESS trends of the two subsamples nearly coincide, showing a common rise of the typical black hole mass with increasing redshift. Beyond $z \approx 3$, a mild divergence appears: the mean mass of the narrow line subsample increases toward high redshift, while the mean of the broad line subsample exhibits a slower increase. This difference is modest compared to the overall scatter and does not represent a discontinuity in the mass distribution. Instead, it likely reflects observational selection effects. The selection of broad line samples requires a $\text{SNR} > 5$ and $\text{FWHM} > 1000 \text{ km s}^{-1}$. At $z \gtrsim 3$, the Balmer and Mg II lines move out of the optical band, while C IV and the UV continuum dominate the observable region, and the SNR degrades. Very broad lines with shallow wings are preferentially lost from the measurements. This truncates the high-FWHM, high-mass tail in the broad line subsample and pulls down its mean mass. At the same time, C IV is affected by outflows that shift the line centroid blueward and distort the profile. These effects can bias the measured FWHM to smaller values and push some intrinsically Type 1 objects below the broad line threshold, moving them into the narrow line bin and raising its mean mass.

The broad-line subsample also shows a noticeably narrower mass distribution compared to the narrow-line subsample, and it lacks the low-mass tail present in the latter. This asymmetry is likewise a consequence of selection effects. Low-mass black holes tend to have lower luminosities, which reduce the equivalent width and signal-to-noise ratio of their emission lines. As a result, their broad components are more easily missed or fall below the adopted selection thresholds.

To conclude, the similarity of the two distributions demonstrates that the model produces internally consistent mass estimates for different quasar types.

6.3. Redshift Distribution

Fig. 5 shows that the overall black hole mass distribution approximately follows a log-normal shape as expected. In contrast, the redshift distribution exhibits several pronounced peaks. The clustering of quasars around $z \approx 2 - 2.5$ arises because strong ultraviolet emission lines such as




Figure 9: Distribution of broad-line (left) and narrow-line (right) quasars in the redshift-mass plane. Each panel shows a hexbin density map with logarithmic color scaling. Two LOWESS trends are plotted (solid: $\log M$ vs. z ; dashed: z vs. $\log M$), with marginal histograms shown on the top and right. The redshift is limited to $0 \leq z \leq 4$, and the source number is indicated as N .

Mg II $\lambda 2798$ and C IV $\lambda 1549$ enter the SDSS optical wavelength range (3800-9200 Å) in this redshift interval, making quasars in this range easier to identify and classify spectroscopically. The SDSS quasar target selection algorithm, which relies on broadband colors, is also most efficient in this regime since these emission features significantly affect the u, g, r, i, z colors, increasing selection completeness. The secondary rise at lower redshift ($z \lesssim 0.8$) is produced by quasars with strong optical Balmer emission lines, particularly H β $\lambda 4861$ and H α $\lambda 6563$, which fall within the SDSS spectral window and produces high signal-to-noise detections. Together, these effects produce peaks in the observed redshift distribution that reflect both instrumental sensitivity and the redshift dependent efficiency of the quasar selection function.

6.4. High Redshift Limitations

The RM dataset contains only 3 quasars at $z > 4$ and does not contain any quasars with $z > 4.5$ or $\log M_{\text{BH}} > 10.5$. Since our network is supervised by RM labels, its accuracy depends on the size and coverage of the training sample. With so few high- z sources, it is difficult to guarantee predictive accuracy after training.

In addition, spectra at high redshift differ substantially from those at low redshift. Rest-frame optical lines such as H β and Mg II move out of the observed wavelength window, while the far-UV dominates, including lines such as Ly α , C IV, and Si IV+O IV]. These features exhibit different systematics, including stronger winds, larger blueshifts, and generally lower signal-to-noise ratios (Richards et al., 2002; Matthews et al., 2023; Denney et al., 2013). This domain shift means that the mapping learned from low- z RM spectra does not transfer cleanly to $z > 4$.

Therefore, although the network can produce mass estimates for high-redshift AGN, we do not include these predictions in the main results due to the limited training support at high z .

7. Conclusion

We trained a single epoch spectral model using RM masses as labels, then applied it to a uniform sample of 287,872 SDSS quasars. This combines RM level calibration with survey scale application, maintaining accuracy while reaching a massive sample size. The predictions track the RM masses with an RMSE scatter of 0.058 dex, which corresponds to a relative uncertainty of about 14%. From the full catalog we measured the redshift evolution of the black hole mass distribution and derived the differential mass function dN/dM for $z < 4$. A smoothed broken power-law fit obtains a break at $M_b \simeq 3 \times 10^8 M_{\odot}$, a low-mass index $\alpha_{\text{low}} \approx +1.1$, and a high-mass index $\alpha_{\text{high}} \approx -2.7$. The cumulative form implies that roughly 10% of objects have $M_{\text{BH}} > 10^9 M_{\odot}$, while fewer than $10^{-3} - 10^{-4}$ exceed $10^{10} M_{\odot}$.

From the machine-learning perspective, the network is constructed to refine the virial structure $M_{\text{BH}} \propto V^2 R(L)$. An encoder-decoder with skip connections allows high frequency line details to be used for spectral reconstruction, while a compact latent vector carries the global information that is most predictive for mass. We adopt a five-dimensional latent representation to preserve interpretability and to align with physically meaningful factors, for example continuum shape, broad line width statistics, and cross line correlations. This physically motivated design achieves accurate predictions without sacrificing transparency, and the latent coordinates will be analyzed in follow-up work.

This mass catalog enables population analyses that were previously limited by sample size or calibration and will be pursued in forthcoming papers, including tests of continuity equation growth models, duty cycle estimates from the active mass function, consistency checks with the quasar luminosity function via the Eddington ratio distribution, and joint constraints on black hole and galaxy coevolution using host scaling relations.

Due to the limited number of high redshift RM samples and the sparse SDSS spectroscopic coverage at $z > 4$, the accuracy of our model is not guaranteed in this regime. As near infrared spectroscopy, such as James Webb telescope, delivers more high z quasars with rest frame optical lines (e.g., $H\beta$, $Mg\ II$), we can finetune the network on these new data to improve performance at high redshift, which will enable more reliable mass estimates for $z > 4$ quasars.

References

- Alexander, D.M., Hickox, R.C., 2012. What drives the growth of black holes? *New A Rev.* 56, 93–121. doi:10.1016/j.newar.2011.09.003.
- Blandford, R.D., McKee, C.F., 1982. Reverberation mapping of the emission line regions of Seyfert galaxies and quasars. *ApJ* 255, 419–439. doi:10.1086/159843.
- Burke, C.J., Suyu, S.H., Millon, M., Chen, G.C.F., Courbin, F., Refsdal, S., Treu, T., Agnello, A., Mao, S., Meyer, E.T., 2021. A characteristic optical variability time scale in astrophysical accretion disks. *Science* 373, 789–794. doi:10.1126/science.aba5344.
- Chainakun, P., Young, C., Stevens, A.L., Parker, M.L., Uttley, P., McHardy, I.M., 2022. Predicting the black hole mass and accretion rate in X-ray reverberating active galactic nuclei using neural networks. *MNRAS* 513, 648–661. doi:10.1093/mnras/stac1123.
- Ciotti, L., Ostriker, J.P., 2007. Radiative Feedback from Massive Black Holes in Elliptical Galaxies: AGN Flaring and Central Starburst Fueled by Recycled Gas. *ApJ* 665, 1038–1056. doi:10.1086/519529.
- Coatman, L., Hewett, P.C., Banerji, M., Richards, G.T., 2017. Correcting C IV-based virial black hole masses. *MNRAS* 465, 2120–2142. doi:10.1093/mnras/stw2797.
- Denney, K.D., Pogge, R.W., Assef, R.J., Barth, A.J., Bennert, V.N., Brewer, B.J., Cid Fernandes, R., Dalla Bontá, E., De Rosa, G., Dietrich, M., Filippenko, A.V., Grier, C.J., Ho, L.C., Jiang, L., Kaspi, S., Kelly, B.C., Kim, D.C., King, A.L., Kochanek, C.S., LaMassa, S., Netzer, H., Peterson, B.M., Rotter, A., Schimoia, J.S., Shappee, B., Shi, Y., Valenti, S., Vestergaard, M., Wang, J.M., Yoshii, Y., 2013. C IV line-width anomalies: the perils of low signal-to-noise spectra. *ApJ* 775, 60. doi:10.1088/0004-637X/775/1/60.
- Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al., 2019. First M87 Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole. *ApJ* 875, L1. doi:10.3847/2041-8213/ab0ec7, arXiv:1906.11238.
- Ferrarese, L., Ford, H., 2005. Supermassive Black Holes in Galactic Nuclei: Past, Present and Future Research. *Space Sci. Rev.* 116, 523–624. doi:10.1007/s11214-005-3947-6.
- Ferrarese, L., Merritt, D., 2000. A Fundamental Relation Between Supermassive Black Holes and Their Host Galaxies. *ApJ* 539, L9–L12. doi:10.1086/312838.
- Graham, A.W., Scott, N., 2013. The $M_{BH-L}^{spheroid}$ Relation at High and Low Masses, the Quadratic Growth of Black Holes, and Intermediate-mass Black Hole Candidates. *ApJ* 764, 151. doi:10.1088/0004-637X/764/2/151.
- Greene, J.E., Ho, L.C., 2007. The Mass Function of Active Black Holes in the Local Universe. *ApJ* 667, 131–148. doi:10.1086/520497, arXiv:0705.0020.
- Häring, N., Rix, H.W., 2004. On the Black Hole Mass-Bulge Mass Relation. *ApJ* 604, L89–L92. doi:10.1086/383567.
- He, D., Qin, Y., Wang, F.Y., Wu, X.F., 2022. Predicting Supermassive Black Hole Masses Using Machine Learning. *Research in Astronomy and Astrophysics* 22, 085014. doi:10.1088/1674-4527/ac8190.
- Hopkins, P.F., Hernquist, L., Cox, T.J., Kereš, D., Aravena, M., Bernardi, M., Blanton, M., Brinchmann, J., Bundy, K., Busha, M.T., Caldwell, R., Dalcanton, J., Davis, M., de Bruyn, A.G., Dekel, A., Egami, E., Finkbeiner, D., Giovanelli, R., Gnedin, N.Y., Gottl"ober, S., Greve, T.R., Grogin, N.A., Haiman, Z., Iono, D., Jarvis, M., Kravtsov, A.V., Lahav, O., Loeb, A., Lutz, D., McDonald, P., Murray, N., Norberg, P., Ostriker, J.P., Padmanabhan, T., Papovich, C., Pierce, C.M., Salvaterra, R., Scannapieco, E., Schmidt, B.P., Shapley, A., Sheth, R.K., Stewart, K.R., Strauss, M.A., Toft, S., Wetzel, A., Wyithe, J.S.B., Zabludoff, A.I., 2008. A Cosmological Framework for the Co-evolution of Quasars, Supermassive Black Holes, and Elliptical Galaxies. I. Galaxy Mergers and Quasar Activity. *ApJS* 175, 356–389. doi:10.1086/524362.
- Inayoshi, K., Visbal, E., Haiman, Z., 2020. The Assembly of the First Massive Black Holes. *ARA&A* 58, 27–97. doi:10.1146/annurev-astro-120419-014455.
- Ivezić, Ž., Kahn, S.M., Tyson, J.A., Abel, B., Acosta, E., Allsman, R., AlSayyad, Y., Anderson, S.F., Andrew, J., Angeli, G.Z.e.a., 2019. LSST: From Science Drivers to Reference Design and Anticipated Data Products. *ApJ* 873, 111. doi:10.3847/1538-4357/ab042c.
- Izumi, T., Onoue, M., Fujimoto, S., Ueda, Y., Matsuoka, Y., Nagao, T., Kohno, K., Toba, Y., Shirakata, H., Nakanishi, K., Yamaguchi, Y., Akiyama, M., Alexander, D.M., Asami, N., Bruno, G., Furusawa, H., Goto, T., Ishikawa, S., Kubo, M., Lee, K.S., Lu, T.M., Mawatari, K., Meyer, R.A., Ni, Y., Ono, Y., Shibuya, T., Taniguchi, Y., Tanaka, M., Wang, F., 2019. Subaru High- z Exploration of Low-Luminosity Quasars (SHELLQs) Project. VIII. Black Hole Masses, Accretion Rates, and Narrow Emission-Line Diagnostics for Quasars at $z \sim 6$. *PASJ* 71, 111. doi:10.1093/pasj/psz109.
- Jin, Z., Davis, B.L., 2023. Discovering black hole mass scaling relations from multi-wavelength AGN observations with symbolic regression. *arXiv e-prints arXiv:2310.19406*.
- Kelly, B.C., Shen, Y., 2013. The Demographics of Broad-line Quasars in the Mass-Luminosity Plane. II. Black Hole Mass Function and Eddington Ratio Distribution. *ApJ* 764, 45. doi:10.1088/0004-637X/764/1/45.
- Kollmeier, J.A., Onken, C.A., Kochanek, C.S., Gould, A., Weinberg, D.H., Dietrich, M., Cool, R., Dey, A., Eisenstein, D.J., Jannuzi, B.T., Le Flo'ch, E., Stern, D., 2006. Black Hole Masses and Eddington Ratios at $0.3 < z < 4$. *ApJ* 648, 128–139. doi:10.1086/505646, arXiv:astro-ph/0508657.
- Kormendy, J., Ho, L.C., 2013. Coevolution (or Not) of Supermassive Black Holes and Host Galaxies. *ARA&A* 51, 511–653. doi:10.1146/annurev-astro-082708-101811.
- Kormendy, J., Richstone, D., 1995. Inward Bound—The Search For Supermassive Black Holes In Galactic Nuclei. *ARA&A* 33, 581–624. doi:10.1146/annurev.aa.33.090195.003053.
- Kuo, C.Y., Braatz, J.A., Condon, J.J., Impellizzeri, C.M.V., Lo, K.Y., Zaw, I., Schenker, M., Henkel, C., Reid, M.J., Greene, J.E., 2011. The Megamaser Cosmology Project. III. Accurate Masses of Seven Supermassive Black Holes in Active Galaxies with Circumnuclear Megamaser Disks. *ApJ* 727, 20. doi:10.1088/0004-637X/727/1/20, arXiv:1008.2146.
- Lin, H.S., Leung, T.C., Shu, F., Song, H., Pei, L., Fu, H.Q., Peng, C., Ye, J.L., Liu, X., 2023. AGNet: An Attention-Based Neural Network to Predict Supermassive Black Hole Masses from Astronomical Light Curves. *MNRAS* 518, 4921–4933. doi:10.1093/mnras/stac3394.
- Lyke, B.W., Higley, A.N., McLane, J.N., Schurhammer, D.P., Myers, A.D., Ross, A.J., Dawson, K., Chabanier, S., Martini, P., Busca, N.G., du Mas des Bourboux, H., Salvato, M., Streblyanska, A., Zarrouk, P., Burtin, E., Anderson, S.F., Bautista, J., Bizyaev, D., Brandt, W.N., Brinkmann, J., Brownstein, J.R., Comparat, J., Green, P., de la Macorra, A., Muñoz Gutiérrez, A., Hou, J., Newman, J.A., Palanque-Delabrouille, N., Pâris, I., Percival, W.J., Petitjean, P., Rich, J., Rossi, G., Schneider, D.P., Smith, A., Vivek, M., Weaver, B.A., 2020. The Sloan Digital Sky Survey Quasar Catalog: Sixteenth Data Release. *ApJS* 250, 8. doi:10.3847/1538-4365/aba623, arXiv:2007.09001.
- Matthews, J.H., Strong-Wright, J., Knigge, C., Hewett, P., Temple, M.J.,

- Long, K.S., Rankine, A.L., Stepney, M., Banerji, M., Richards, G.T., 2023. A disc wind model for blueshifts in quasar broad emission lines. *MNRAS* 526, 3967–3986. doi:10.1093/mnras/stad2895.
- McConnell, N.J., Ma, C.P., 2013. Revisiting the Scaling Relations of Black Hole Masses and Host Galaxy Properties. *ApJ* 764, 184. doi:10.1088/0004-637X/764/2/184.
- McHardy, I.M., Koending, E., Knigge, C., Uttley, P., Fender, R.P., 2006. Active galactic nuclei as scaled-up Galactic black holes. *Nature* 444, 730–732. doi:10.1038/nature05389.
- McInnes, L., Healy, J., Melville, J., 2018. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *arXiv e-prints* arXiv:1802.03426.
- Mejía-Restrepo, J., Lira, P., Netzer, H., Trakhtenbrot, B., Capellupo, D., 2017. The virial factor and biases in single epoch black hole mass determinations. *Frontiers in Astronomy and Space Sciences* 4, 70. doi:10.3389/fspas.2017.00070.
- Mejía-Restrepo, J.E., Trakhtenbrot, B., Lira, P., Netzer, H., Capellupo, D.M., 2018. The reliability of C IV λ 1549 broad-line as a virial black hole mass estimator for active galactic nuclei. *Monthly Notices of the Royal Astronomical Society* 478, 1929–1952. doi:10.1093/mnras/sty1153.
- Merritt, D., Ferrarese, L., 2001. The $M_{\text{bh}}\text{-}\sigma$ Relation for Supermassive Black Holes. *ApJ* 547, 140–145. doi:10.1086/318372.
- Peterson, B.M., 1993. Reverberation Mapping of Active Galactic Nuclei. *PASP* 105, 247. doi:10.1086/133140.
- Richards, G.T., Vanden Berk, D.E., Reichard, T.A., Hall, P.B., Schneider, D.P., SubbaRao, M., Thakar, A.R., York, D.G., Strateva, I.V., 2002. Broad Emission-Line Shifts in Quasars: An Orientation Measure for Quasars? *AJ* 124, 1–17. doi:10.1086/339056.
- Saglia, R.P., Opitsch, M., Erwin, P., Thomas, J., Beifiori, A., Fabricius, M., F"orster Schreiber, N.M., Mendel, J.T., Bender, R., 2016. The SINFONI Black Hole Survey: The Black Hole Fundamental Plane Revisited and the Paths of (Co)evolution of Supermassive Black Holes and Bulges. *ApJ* 818, 47. doi:10.3847/0004-637X/818/1/47.
- Schawinski, K., Dowlin, N., Thomas, D., Urry, C.M., Edmondson, E.I., 2010. The Fundamentally Different Co-Evolution of Supermassive Black Holes and Their Host Galaxies. *ApJ* 711, 284–302. doi:10.1088/0004-637X/711/1/284.
- Schechter, P., 1976. An analytic expression for the luminosity function for galaxies. *ApJ* 203, 297–306. doi:10.1086/154079.
- Schramm, M., Silverman, J.D., 2013. The Black Hole-Bulge Mass Relation of Active Galactic Nuclei in the Extended Chandra Deep Field-South Survey. *ApJ* 767, 13. doi:10.1088/0004-637X/767/1/13.
- Schulze, A., Wisotzki, L., 2010. Low redshift AGN in the Hamburg/ESO Survey . II. The active black hole mass function and the distribution function of Eddington ratios. *A&A* 516, A87. doi:10.1051/0004-6361/201014193, arXiv:1004.2671.
- Shankar, F., Bernardi, M., Richardson, S.N., Lang, P., Xie, M., Koss, M.J., Krishnan, C., Eracleous, M., Colpi, M., 2019. Black hole scaling relations of active and quiescent galaxies: addressing selection effects and constraining virial factors. *MNRAS* 485, 1278–1300. doi:10.1093/mnras/stz509.
- Shankar, F., Bernardi, M., Sheth, R.K., Ferrarese, L., Graham, A.W., Marconi, A., Savorgnan, G., La Barbera, F., Lasky, P.D., Haigh, I.P., 2016. Selection bias in dynamically measured supermassive black hole samples: its consequences and the quest for the most fundamental relation. *MNRAS* 460, 3119–3144. doi:10.1093/mnras/stw803.
- Shen, Y., Brandt, W.N., Dawson, K.S., Hall, P.B., MacLeod, C.L., Richards, G.T., Schneider, D.P., Trump, J.R., Byun, D.Y., Chen, Y., et al., 2015. The Sloan Digital Sky Survey Reverberation Mapping Project: Technical Implementation and Sample Characterization. *ApJS* 216, 4. doi:10.1088/0067-0049/216/1/4.
- Shen, Y., Greene, J.E., Strauss, M.A., Richards, G.T., Schneider, D.P., McDonald, P., Bailey, S., Fan, X., Hall, P.B., Jiang, L., Ross, N.P., Brandt, W.N., Kurtz, M.J., Vanden Berk, D.E., 2008. Biases in virial black hole masses: an SDSS perspective. *ApJ* 680, 169–190. doi:10.1086/587475.
- Shen, Y., Grier, C.J., Horne, K., Stone, Z., Li, J.I., Yang, Q., Homayouni, Y., Trump, J.R., Anderson, S.F., Brandt, W.N., Hall, P.B., Ho, L.C., Jiang, L., Petitjean, P., Schneider, D.P., Tao, C., Donnan, F.R., AlSayyad, Y., Bershad, M.A., ..., Zou, H., 2024. The sloan digital sky survey reverberation mapping project: Key results. *The Astrophysical Journal Supplement Series* 272, 26. doi:10.3847/1538-4365/ad3936.
- Shen, Y., Richards, G.T., Strauss, M.A., Hall, P.B., Schneider, D.P., Snedden, S., Bizyaev, D., Brewington, H., Brinkmann, J., Eisenstein, D.J., Frieman, J., Gray, J., Gunn, J.E., Hao, L., Harrington, R., Harvanek, M., Kent, S.M., Kleinman, S.J., Kniazev, A., Krzesinski, J., Long, D., Loveday, J., Malanushenko, E., Malanushenko, V., Oravetz, D., Pan, K., Pier, J.R., Simmons, A., de Haas, E., York, D.G., 2011. A Catalog of Quasar Properties from Sloan Digital Sky Survey Data Release 7. *ApJS* 194, 45. doi:10.1088/0067-0049/194/2/45.
- Sijacki, D., Springel, V., Di Matteo, T., Hernquist, L., 2007. A unified model for AGN feedback in cosmological simulations of structure formation. *MNRAS* 380, 877–900. doi:10.1111/j.1365-2966.2007.12153.x.
- Soltan, A., 1982. Masses of quasars. *MNRAS* 200, 115–122. doi:10.1093/mnras/200.1.115.
- Vestergaard, M., Peterson, B.M., 2006. Determining central black hole masses in distant active galaxies and quasars. ii. improved optical and uv scaling relationships. *The Astrophysical Journal* 641, 689.
- Wu, X.B., Wang, F., Fan, X., Yi, W.M., Zuo, W., Bian, F., Jiang, L., McGreer, I.D., Wang, R., Yang, J., Yang, Q., Thompson, D., Stern, D., Wang, J., Wang, Y., Zhang, X., Yang, J., Wang, X., Kang, W., 2015. An ultraluminous quasar with a twelve-billion-solar-mass black hole at redshift 6.30. *Nature* 518, 512–515. doi:10.1038/nature14241.
- Yong, S.Q., Ong, D., 2023. Conformal quantile regression for virial black hole mass estimation and uncertainty quantification. *MNRAS* 524, 3116–3127. doi:10.1093/mnras/stad2190.