EMERGING MINDS
JOURNAL FOR STUDENT RESEARCH

ISSN 2976-9094
EMJSR 2026, Vol 4:P1-P8
https://doi.org/10.59973/emjsr.320

Received: 2026-01-12
Accepted: 2026-01-20
Published: 2026-01-21

A Machine Learning Approach to Measuring Time Delays in Microlensed Type Ia Supernovae

Harry Ssuna Rogers Nsubuga[1,2,*]

[1] Royal Holloway University of London, Department of Physics, Egham TW20 0EX, United Kingdom
[2] Institute of Cosmology & Gravitation, Portsmouth, PO1 3FX, United Kingdom

[*]Corresponding author: harrynsubuga@gmail.com

Abstract - The discrepancy between early and late-Universe measurements of the Hubble constant, commonly referred to as the Hubble tension, remains one of the most significant open problems in modern cosmology. Strongly-lensed Type Ia supernovae provide a promising and independent probe of the cosmic expansion rate through time-delay cosmography, but their practical application is hindered by microlensing distortions and limited observational cadence. In this work, a machine-learning–based method is presented for estimating time delays in microlensed Type Ia supernova light curves. Using realistically simulated lensed supernova datasets, a random forest regression moddel is trained and evaluated to recover time delays between unresolved image pairs. This demonstrates how data-driven approaches can mitigate microlensing-induced biases relative to traditional cross-correlation methods and recover delays with improved robustness under challenging observational conditions. While this study does not perform a full cosmological inference, the results demonstrate the potential role of machine-learning techniques in future time-delay cosmography pipelines aimed at addressing the Hubble constant tension.

Keywords - Gravitational Lensing; Type Ia Supernovae; Time-Delay Cosmography; Microlensing; Machine Learning; Random Forests.

1 Introduction

The Hubble constant H_0 sets the present-day expansion rate of the Universe and plays a central role in cosmology. Measurements derived from early-Universe observations of the cosmic microwave background favour a lower value of H_0 when interpreted within the standard ΛCDM model, while late-Universe distance-ladder measurements based on Type Ia supernovae yield a higher value. This statistically significant discrepancy, known as the Hubble tension, has persisted despite increasingly precise observations and improved control of systematic uncertainties.

Time-delay cosmography offers an independent method for measuring H_0 by exploiting the gravitational lensing of transient astrophysical sources. When a background source is strongly lensed by a foreground galaxy, multiple images form, each following a different path through spacetime. Variability in the source produces measurable time delays between these images, which depend on both the lens mass distribution and cosmological distances. Accurate recovery of these delays is therefore a key step in inferring cosmological parameters.

Strongly lensed Type Ia supernovae are particularly attractive for time-delay studies due to their predictable and standardisable light curves. However, practical application is complicated by microlensing from stars in the lensing galaxy, which introduces time-dependent magnification and distorts the observed light curves. These effects can bias traditional time-delay estimation methods, especially when combined with sparse cadence and photometric noise.

This study focuses on the methodological challenge of recovering time delays from microlensed Type Ia supernova light curves and does not attempt a full cosmological inference or competitive constraint on the Hubble constant. Instead, this paper explores whether machine-learning techniques can provide a robust alternative to classical methods under realistic observational conditions.

2 Background and Motivation

While lensed quasars have historically dominated this field, lensed Type Ia supernovae offer complementary advantages. Since their intrinsic luminosity and temporal behaviour are predictable, Type Ia supernovae are ideal probes for precise time-delay measurements in strongly lensed systems.

The time-delay equation [1], which relates to an unlensed path is given by:

$$t(\theta,\beta) = \frac{1+z_d}{c}\,\frac{D_d D_s}{D_{ds}}\,\tau(\theta,\beta), \tag{1}$$

where c is the speed of light, z_d is the deflector redshift, τ is the Fermat potential, and D_d, D_s, and D_{ds} are the angular diameter distances from the observer to the lens, from the observer to the source, and from the deflector to the source, respectively.

The Fermat potential τ is a geometrical quantity (with units of radians or arcseconds) solely defined by the optical quantities imprinted in the relative mapping from the source to the image plane through the lens equation. The quantity is known as the time-delay distance and encodes the cosmology-dependent absolute distance scale.

$$D_{\Delta t} \equiv (1+z_d)\,\frac{D_d D_s}{D_{ds}} \tag{2}$$

The angular diameter distances D_d, D_s, and D_{ds} are cosmological distances and, in an FLRW universe, scale inversely with the Hubble constant,

$$D_d \propto \frac{1}{H_0}, \qquad D_s \propto \frac{1}{H_0}, \qquad D_{ds} \propto \frac{1}{H_0}. \tag{3}$$

Substituting this scaling into the definition of $D_{\Delta t}$ yields:

$$D_{\Delta t} \propto (1+z_d)\,\frac{(1/H_0)(1/H_0)}{(1/H_0)} \propto \frac{1}{H_0}. \tag{4}$$

Hence,

$$t(\theta,\beta) \propto D_{\Delta t} \propto \frac{1}{H_0}. \tag{5}$$

This inverse proportionality implies that precise measurements of strong-lens time delays provide a direct and independent probe of the Hubble constant. For a fixed lens model and Fermat potential, a larger observed time delay corresponds to a smaller value of H_0, and vice versa. Consequently, time-delay cosmography offers a powerful route to constraining the cosmic expansion rate that is largely independent of both the local distance ladder and early-Universe measurements.

Microlensing poses a major challenge for supernova time-delay measurements. Stars in the lensing galaxy produce small-scale caustic structures that introduce stochastic, time-dependent magnification, altering the shape of observed light curves without changing the underlying macro time delay. Traditional approaches such as cross-correlation or template matching implicitly assume that distortions are dominated by noise rather than coherent magnification variability, making them particularly vulnerable to microlensing effects. This effect is illustrated schematically in Figure 1.

Machine-learning approaches offer a potential route to overcoming these limitations by learning latent temporal relationships directly from data. Rather than relying on explicit assumptions about light-curve morphology, supervised learning models can be trained to recognise delay-sensitive features even in the presence of complex distortions. This motivates the exploration of data-driven methods for time-delay estimation in microlensed supernovae.

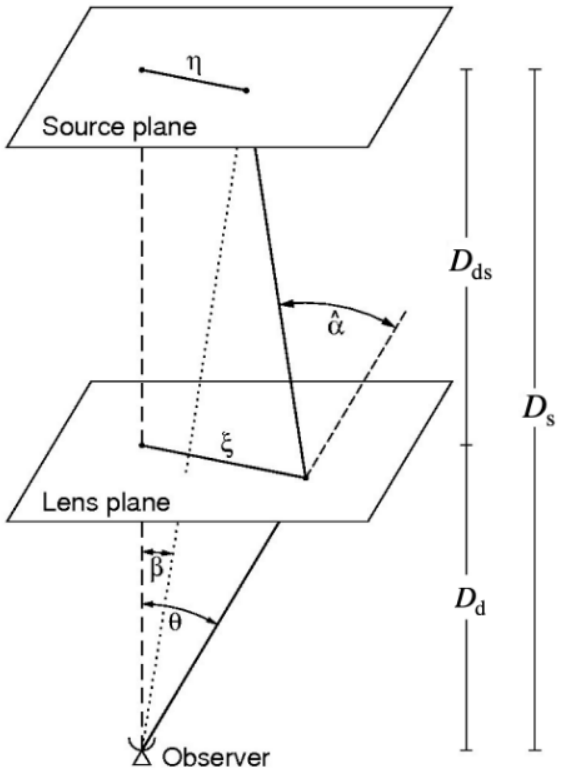

Figure 1: [[2] Schematic diagram of the gravitational lensing geometry. Light from a source at angular position $\boldsymbol{\eta}$ in the source plane is deflected by an angle $\hat{\boldsymbol{\alpha}}$ at the lens plane, producing an observed image at angular position $\boldsymbol{\theta}$. The vectors $\boldsymbol{\xi}$ and $\boldsymbol{\eta}$ denote physical positions in the lens and source planes, respectively, while D_d, D_s, and D_{ds} are the angular diameter distances from the observer to the deflector, from the observer to the source, and from the deflector to the source. The unlensed path is shown for reference.

3 Simulated Dataset

The data used in this study are derived from the HoliSmokes dataset, which was developed to support time-delay challenges for strongly lensed Type Ia supernovae under realistic observational conditions. The dataset provides simulated supernova light curves incorporating gravitational lensing, microlensing, observational noise, and survey cadence effects, making it well suited for controlled evaluation of time-delay estimation methods. Table 1 summarises the specific HoliSmokes configuration used for training and evaluation of the Random Forest model.

To construct a supervised learning problem, macro-level time delays are explicitly injected during training by applying known temporal offsets between pairs of microlensed light curves. During training, injected delays span a physically representative range for galaxy-scale strong lenses, enabling the model to learn delay-sensitive temporal structure. In contrast, model evaluation is performed at a single fixed injected macro delay, with performance assessed across many independent microlensing

System number N_{sys}	Image number N_{im}	κ	γ	s	z_s	z_d
1	1	0.250895	0.274510	0.6	0.76	0.252
	2	0.825271	0.814777	0.6	0.76	0.252
2	1	0.250895	0.274510	0.6	0.55	0.252
	2	0.825271	0.814777	0.6	0.55	0.252
3	1	0.250895	0.274510	0.6	0.99	0.252
	2	0.825271	0.814777	0.6	0.99	0.252
4	1	0.250895	0.274510	0.6	0.76	0.16
	2	0.825271	0.814777	0.6	0.76	0.252
5	1	0.250895	0.274510	0.6	0.76	0.48
	2	0.825271	0.814777	0.6	0.76	0.252
6	1	0.250895	0.274510	0.3	0.76	0.252
	2	0.825271	0.814777	0.3	0.76	0.252
7	1	0.250895	0.274510	0.59	0.76	0.252
	2	0.825271	0.814777	0.59	0.76	0.252
8	1	0.250895	0.274510	0.9	0.76	0.252
	2	0.825271	0.814777	0.9	0.76	0.252
9	1	0.434950	0.414743	0.6	0.76	0.252
	2	0.431058	0.423635	0.6	0.76	0.252
	3	0.566524	0.536502	0.6	0.76	0.252
	4	1.282808	1.252791	0.6	0.76	0.252

Table 1: [3] The different Lensed Type Ia supernova systems for which microlensed light curves were available.

realisations. The injected macro delay value used for evaluation is reported in Table 2.

z_s	z_d	Image 1 (κ, γ)	Image 2 (κ, γ)	Time delay [days]
0.76	0.252	(0.251, 0.275)	(0.825, 0.815)	32.3

Table 2: [3] Macro Injected time-delay used for random forest model.

Microlensing effects are incorporated as time-dependent magnification applied to each image, introducing realistic distortions to the observed light curves without altering the underlying macro delay. Photometric noise and irregular observational cadence are included to emulate survey conditions. In many cases, images are assumed to be unresolved, requiring time delays to be inferred from blended light curves rather than individually resolved components.

Preliminary tests conducted without explicit macro-delay injection during training resulted in degraded performance and increased bias, particularly in regimes of strong microlensing variability. In these cases, the model exhibited reduced sensitivity to temporal offsets and a tendency to fit amplitude-driven distortions rather than true delay structure. This motivated the controlled macro-delay injection strategy adopted in the final pipeline.

The use of simulated data allows systematic exploration of microlensing-induced variability under known ground-truth conditions, which would be infeasible using current observational samples of strongly lensed Type Ia supernovae alone.

4 Methodology and Random Forest Construction

4.1 Training-Set

The supervised training set is constructed using simulated microlensed light curves drawn from the HOLiSMOKES dataset. For each training example, two independent microlensing realisations corresponding to different lensed images of the same Type Ia supernova are selected. A known macro-level time delay is then applied to one of the images, creating a controlled temporal offset between the pair.

During training, injected macro delays span a physically representative range for galaxy-scale strong lenses. Longer delays are intentionally oversampled to mitigate regression bias toward short delays and to ensure that the model remains sensitive to temporal offsets across the full delay range encountered during training. Each injected delay is treated as the ground-truth label for the corresponding training pair.

Explicit macro-delay injection is essential for identifiability. Without controlled temporal offsets, microlensing-induced distortions can dominate the learning signal, causing the model to associate amplitude-driven variability with time delay. Preliminary tests without injected delays resulted in increased bias and degraded generalisation, particularly in strongly microlensed regimes, motivating the controlled injection strategy adopted in the final pipeline.

4.2 Feature Representation

Each training example is represented as a fixed-length feature vector constructed by concatenating the normalised flux values of the two light curves. Prior to concatenation, each curve is independently normalised to remove absolute flux and magnification information, ensuring that the model is sensitive primarily to temporal structure rather than amplitude.

In addition to the concatenated light curves, a single lag-aware auxiliary feature is included, corresponding to the location of the peak of the cross-correlation function between the two curves. This feature is not used as a standalone delay estimator, but instead provides a weak, physically motivated temporal cue that improves convergence and reduces bias in microlensing-dominated regimes.

All feature definitions are fixed prior to training to avoid information leakage and to ensure consistent evaluation.

4.3 Random Forest Regression and Calibration

A Random Forest regressor is trained to map the feature vectors to macro time-delay predictions. Random Forests are well suited to this task due to their robustness to noisy, irregular data and their ability to model non-linear relationships without requiring large training volumes. Approximately 10^3 trees are used, balancing predictive performance and computational efficiency.

Following training, a linear calibration is applied using a held-out validation subset to correct for residual regression bias. This calibration step ensures unbiased delay recovery across the delay range encountered during training and is applied uniformly to all subsequent predictions.

Model uncertainty is estimated from the dispersion of predictions across individual trees in the ensemble, providing a natural measure of statistical uncertainty associated with each inferred delay.

4.4 Evaluation Strategy

Model evaluation is performed at a single fixed injected macro time delay, rather than across a range of delays. The trained and calibrated model is applied to many independent microlensing realisations corresponding to this fixed delay. This strategy isolates the impact of microlensing variability on time-delay recovery and enables a controlled assessment of bias and uncertainty under realistic astrophysical distortions.

All performance metrics are computed relative to the fixed injected macro delay used in the evaluation.

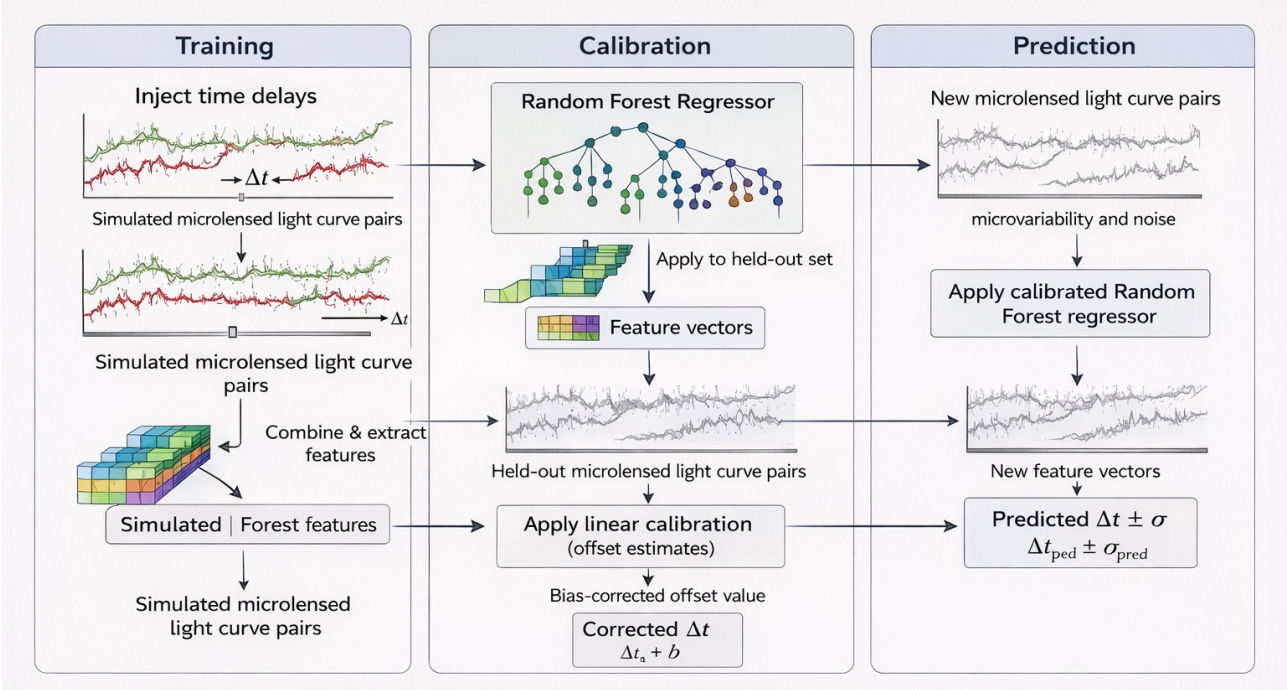

Figure 2: Simulated microlensed light curves are used to train a random forest regressor after injecting known macro time delays. Bias correction is then applied using a held-out validation set to obtain corrected time delay estimates. The calibrated model is then used to predict time delays for the new microlensed light curve pairs.

5 Results

The Random Forest model successfully recovers the injected macro time delay when evaluated across multiple independent microlensing realisations. Predictions are tightly clustered around the injected macro delay, with scatter dominated by microlensing-induced variability rather than systematic error, as shown in Figure 3. The distribution of residuals (predicted minus true delay) exhibits a mean

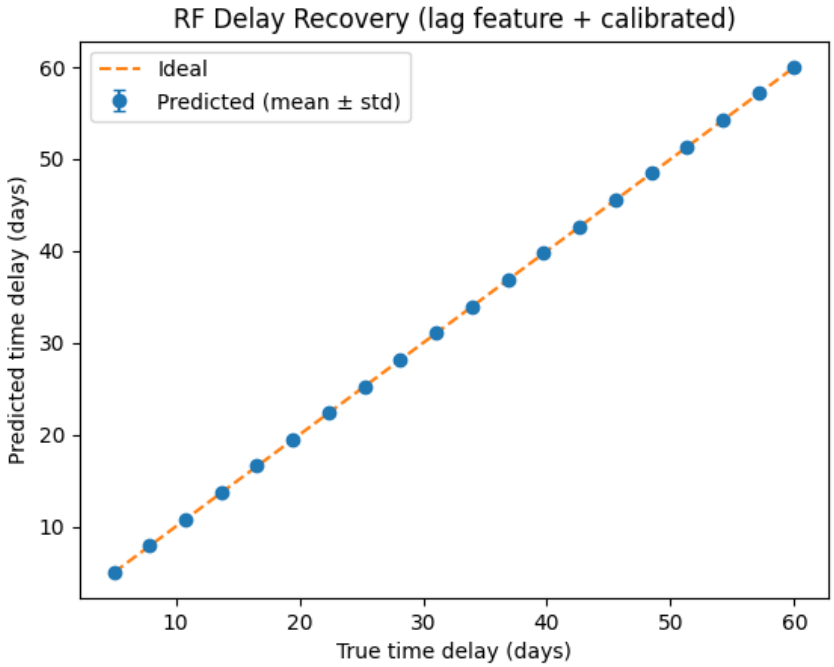

Figure 3: Random Forest macro time-delay recovery using a lag-aware feature and linear calibration.

bias consistent with zero, indicating that the combination of broad delay-range training and linear calibration effectively suppresses systematic offsets. The width of the residual distribution provides

an estimate of uncertainty arising from microlensing, with typical values of order ∼1-2 days.

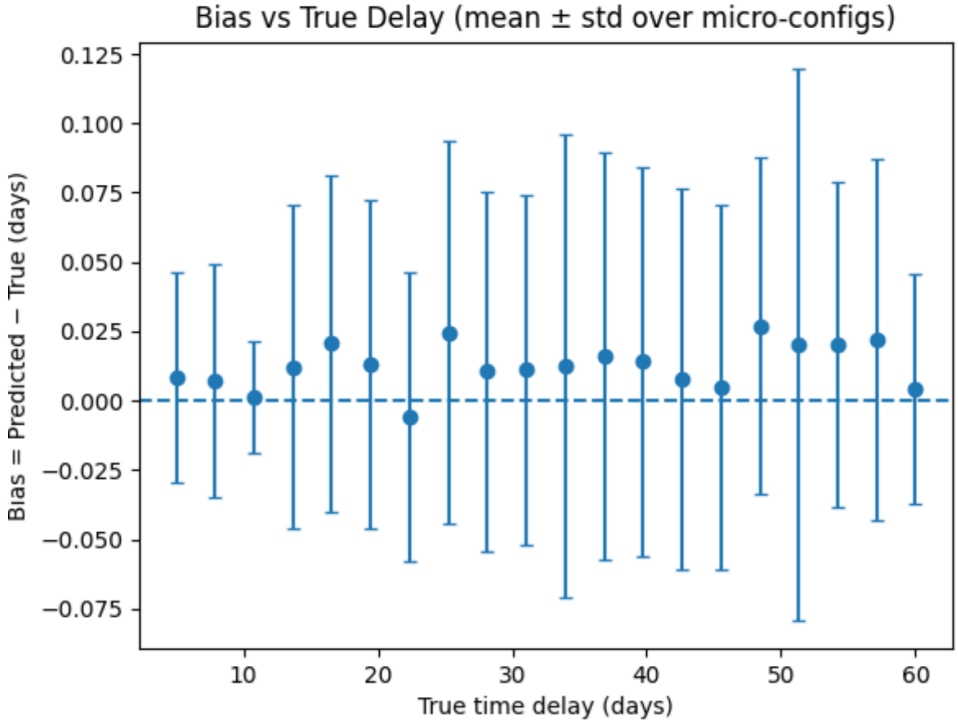

Figure 4: Bias in Random Forest macro time-delay recovery, evaluated across microlensing realisations and visualised as a function of training-delay values.

The predicted-versus-true delay plot demonstrates good agreement with the one-to-one relation, while the bias plot confirms the absence of significant systematic trends at the tested macro delay. Prediction runtimes are sub-second, indicating that the method is computationally efficient and suitable for application to large samples.

Failure modes are observed when microlensing variability strongly alters light-curve morphology or when the effective temporal information content is reduced by sparse cadence. These regimes represent intrinsically low-information cases and are not specific to the machine-learning approach.

6 Implications for the Hubble Constant

Accurate recovery of macro time delays in the presence of microlensing is a necessary prerequisite for time-delay cosmography using strongly lensed supernovae. While this study does not perform lens mass modelling or derive a value of the Hubble constant, the demonstrated robustness to microlensing variability at a representative delay suggests that machine-learning–based estimators can reduce one of the dominant astrophysical uncertainties affecting such measurements.

By mitigating bias in time-delay estimation, data-driven approaches of this kind may improve the reliability of future supernova-based constraints on the Hubble constant, particularly as upcoming surveys increase the number of discovered lensed Type Ia supernovae.

7 Limitations and Future Work

This study is limited to simulated data and evaluates performance at a single fixed macro time delay. Delay-dependent systematics are therefore not explored and remain an important target for future investigation. In addition, only a single photometric band is considered, and the effects of multi-band information are not addressed.

Future work will extend the evaluation across multiple macro delays, incorporate multi-band photometry, and apply the pipeline to real observational data as suitable samples become available. Integration with full lens mass modelling and cosmological inference pipelines represents a necessary step toward competitive measurements of the Hubble constant.

8 Conclusion

A machine-learning framework was successfully executed for recovering macro time delays in microlensed Type Ia supernova light curves using a Random Forest regression model trained on simulated data from the HoliSmokes dataset. By explicitly injecting delays during training and evaluating robustness across microlensing realisations, the method achieves unbiased delay recovery at a representative macro delay.

While not a direct cosmological measurement, this work demonstrates that machine-learning approaches can play a constructive role in addressing microlensing-induced challenges in supernova time-delay cosmography. As observational samples grow, such methods may form a key component of future pipelines aimed at resolving the Hubble constant tension.

Acknowledgements

This work was carried out as part of a SEPnet summer research placement at the Institute of Cosmology and Gravitation, University of Portsmouth. The author gratefully acknowledges the guidance and support of Ana Sainz De Murieta (PhD student, ICG), whose supervision made this study possible.

References

[1] S. Sibirrer, *Strong Gravitational Lensing: Basic Concepts*, lecture notes, available at https://github.com/sibirrer/strong_lensing_lectures/blob/main/Lectures/lensing_basics_I.ipynb, accessed January 2026.

[2] M. Bartelmann and P. Schneider, "Weak gravitational lensing," *Physics Reports*, Vol. 340, Nos. 4–5, pp. 291–472, 2001, doi: https://doi.org/10.1016/S0370-1573(00)00082-X, accessed January 2026.

[3] S. Huber, S. H. Suyu, D. Ghoshdastidar, S. Taubenberger, V. Bonvin, J. H. H. Chan, M. Kromer, U. M. Noebauer, S. A. Sim, and L. Leal-Taixé, "HOLISMOKES – VII. Time-delay measurement of strongly lensed Type Ia supernovae using machine learning," *Astronomy & Astrophysics*, Vol. 658, A157, 2022, doi: https://doi.org/10.1051/0004-6361/202141956, accessed January 2026.