

Active galactic nuclei-heated dust revealed in “little red dots”

I. Delvecchio^{1,*}, E. Daddi², B. Magnelli², D. Elbaz², M. Giavalisco³, A. Traina¹, G. Lanzuisi¹,
H. B. Akins⁴, S. Belli⁵, C. M. Casey⁶, F. Gentile^{2,1}, C. Gruppioni¹, F. Pozzi⁵, and G. Zamorani¹

¹ INAF – Osservatorio di Astrofisica e Scienza dello Spazio di Bologna, Via Gobetti 93/3, I-40129 Bologna, Italy

² Université Paris-Saclay, Université Paris Cité, CEA, CNRS, AIM, 91191 Gif-sur-Yvette, France

³ Astronomy Department, University of Massachusetts, Amherst, MA 01003, USA

⁴ The University of Texas at Austin, 2515 Speedway Boulevard Stop C1400, Austin, TX 78712, USA

⁵ University of Bologna, Department of Physics and Astronomy (DIFA), Via Gobetti 93/2, I-40129 Bologna, Italy

⁶ Department of Physics, University of California, Santa Barbara, Santa Barbara, CA 93106, USA

Received 10 September 2025 / Accepted 27 October 2025

ABSTRACT

Little red dots (LRDs) are a puzzling population of extragalactic sources whose origin is highly debated. In this work, we performed a comprehensive stacking analysis of NIRC*am*, MIRI, and ALMA images of a large and homogeneously selected sample of LRDs from multiple JWST Legacy fields. We report clear evidence of hot-dust emission in the median stacked spectral energy distribution (SED) that features a rising near-infrared continuum up to rest-frame $\lambda_{\text{rest}} \sim 3 \mu\text{m}$, which is best explained by a standard dusty active galactic nucleus (AGN) structure. Although LRDs are likely to be a heterogeneous population, our findings suggest that most ($\geq 50\%$) LRDs show AGN-heated dust emission, regardless of whether the optical and ultraviolet (UV) continua are stellar or AGN-dominated. In either case, the best-fit dusty-AGN SED, combined with the lack of X-ray detection in the deep Chandra stacks, suggests that Compton-thick ($N_{\text{H}} > 3 \times 10^{24} \text{ cm}^{-2}$) gas obscuration is common, and likely confined within the dust sublimation radius ($R_{\text{sub}} \sim 0.1 \text{ pc}$). Therefore, we argue that AGN-heated dust does not directly obscure either the optical-UV continuum or the broad-line region emission, in order to explain the observed blue UV slopes and prominent Balmer features. While a gas-dust displacement is in line with several models, the formation scenario (in-situ or ex-situ) of this pre-enriched hot dust remains unclear.

Key words. galaxies: active – galaxies: evolution – galaxies: high-redshift – quasars: general – quasars: supermassive black holes

1. Introduction

The James Webb Space Telescope (JWST) has ushered extragalactic astronomy into a transformative era, offering unprecedented access to the high-redshift ($z > 4$) universe with the discovery of previously missed active galactic nuclei (AGN; e.g. Matthee et al. 2024). These AGNs are mostly identified via broad emission lines (full width at half maximum $> 1000\text{--}2000 \text{ km s}^{-1}$), corresponding to bolometric luminosities of $L_{\text{bol,AGN}} \sim 10^{44\text{--}46} \text{ erg s}^{-1}$ and harboring supermassive black holes (SMBHs, with mass $M_{\text{BH}} \sim 10^6\text{--}10^8 M_{\odot}$). In contrast to their lower- z analogs, they do not show significant X-ray or radio emission (Ananna et al. 2024; Mazzolari et al. 2025). A non-negligible fraction of these AGN ($\sim 10\text{--}30\%$; Hainline et al. 2025) feature unique and puzzling properties, and are dubbed “little red dots” (LRDs, e.g., Matthee et al. 2024; Greene et al. 2024; Labbe et al. 2025). These LRDs are extremely compact (often unresolved at $\lesssim 100 \text{ pc}$ scales) and show red optical colors combined with blue ultraviolet (UV) slopes, which result in a characteristic “V-shaped” spectral energy distribution (SED), with a sharp turnover near the Balmer break at rest 3645 \AA (e.g., Setton et al. 2024; Hainline et al. 2025). The vast majority of them are seen at $4 < z < 8$, and about $\sim 60\%$ of the spectroscopically confirmed LRDs display broad emission lines (e.g., Greene et al. 2024; Hviding et al. 2025). Notably, LRDs appear to be up to two orders of magnitude more abundant than luminous quasars at $z > 4$ (Kokorev et al. 2024;

Akins et al. 2025), which makes them accessible even in narrow-field JWST surveys.

However, an observational consensus on the physical nature of LRDs is still lacking. On the one hand, if LRDs are compact star-forming galaxies (e.g., Labbé et al. 2023), their stellar masses ($M_{\star} \sim 10^{8\text{--}11} M_{\odot}$) would imply extreme stellar densities (Pacucci & Narayan 2024). On the other hand, the red optical colors and broad emission lines hint at an AGN origin (Inayoshi et al. 2020; Volonteri et al. 2021; Ji et al. 2025), from which the (e.g., H α -based) M_{BH} estimates would imply overmassive black holes (BHs) ($M_{\text{BH}}/M_{\star} \geq 1:100$; see e.g., Pacucci et al. 2023; Juodžbalis et al. 2024a, 2025; Maiolino et al. 2024) compared to local scaling relations (Reines & Volonteri 2015), even without accounting for dust attenuation (e.g., Greene et al. 2025).

In the AGN scenario, the optical and/or UV continuum would be moderately reddened by dust ($A_V \sim 2\text{--}3 \text{ mag}$; Casey et al. 2024; Akins et al. 2025; Wang et al. 2025). However, this interpretation is apparently at odds with their extreme X-ray and radio weakness implied by the lack of stacked X-ray (Ananna et al. 2024; Yue et al. 2024; Maiolino et al. 2025) and radio emission (Akins et al. 2025; Gludemans et al. 2025). Recent studies attribute the X-ray weakness to super-Eddington accretion (Madau & Haardt 2024; Lambrides et al. 2024; Pacucci & Narayan 2024; King 2025; Madau 2025; but see e.g., Sacchi & Bogdán 2025).

To also alleviate the tension between the prominent optical and/or UV Balmer features and the apparent X-ray weakness, several models postulate the presence of a turbulent,

* Corresponding author: ivan.delvecchio@inaf.it

dust-free, and highly dense gas envelope (e.g., Liu et al. 2025; Cenci & Habouzit 2025; Inayoshi & Maiolino 2025; Naidu et al. 2025; Rusakov et al. 2025; de Graaff et al. 2025). These conditions explain the strong Balmer breaks and the red continua through pure gas opacity, since the gas temperature is too high for the dust to survive. In addition, the broad Balmer lines and other absorption features can be explained via electron scattering (e.g., Rusakov et al. 2025; Inayoshi et al. 2025) or resonant Balmer scattering (e.g., Naidu et al. 2025). These models mostly differ on the nature of the central object within the gas envelope, which could be either a direct-collapse BH (e.g., Cenci & Habouzit 2025), a Super-Eddington BH (Liu et al. 2025), a late-stage “quasi star” (Begelman & Dexter 2025), a “BH-Star” (e.g. Naidu et al. 2025), or a metal-free object such as a super-massive star (Nandal & Loeb 2025).

Such a dust paucity is further supported by the non-detection from deep Atacama Large Millimeter Array (ALMA) stacking, which poses a stringent upper limit to the cold dust budget, corresponding to $M_{\text{dust}} < 10^6 M_{\odot}$ at $z \sim 6$ (Casey et al. 2025; see also Chen et al. 2025; Setton et al. 2025). In the rest-frame near and mid-IR, a broad diversity of LRD SEDs has been reported in the recent literature. Stacking of Mid-Infrared Instrument (MIRI)-undetected LRDs reveals a flat SED beyond rest 1.6 μm (e.g., Williams et al. 2024; Carranza-Escudero et al. 2025), in contrast to the rising mid-infrared (MIR) continuum expected in obscured AGN, which conveys the idea that MIR emission in LRDs is typically dominated by dust-poor stellar populations and not AGN dust. However, based on individual MIRI detections, other studies found that some LRDs display a flat or rising MIR continuum (Pérez-González et al. 2024; Barro et al. 2024; Leung et al. 2025; de Graaff et al. 2025; Ronayne et al. 2025), which is best fitted via composite stellar and AGN templates, suggesting a heterogeneous origin of the IR emission of LRDs.

In this highly debated framework, we caution that previous MIRI-SED studies of LRDs might have been biased toward bright individual objects (e.g., Labbe et al. 2024; Barro et al. 2024; Chen et al. 2025; de Graaff et al. 2025; Wang et al. 2025), or possibly affected either by individual MIRI non-detections (e.g., Pérez-González et al. 2024) or by stacking peculiar LRD subsamples (e.g., *H*-dropouts; Williams et al. 2024), whose SED could intrinsically differ from that of “typical” LRDs. To explore this possibility, in this paper we present a comprehensive stacking analysis of a large and homogeneously selected LRD sample with broad Near Infrared Camera (NIRCam), MIRI, and ALMA coverage. Throughout, we assume a Chabrier (2003) initial mass function (IMF) and a flat Λ CDM cosmology with $H_0 = 70 \text{ km/s/Mpc}$, $\Omega_m = 0.30$, and $\Omega_{\Lambda} = 0.70$.

2. Photometric selection of LRDs

We selected LRDs from the compilation of Kocevski et al. (2025), which consists of 341 photometrically selected LRDs in the redshift range $z \sim 2\text{--}11$ that were retrieved from several JWST/NIRCam surveys. These LRDs were homogeneously selected to display compact morphologies and “V-shaped” continua from spectral slope fitting between blueward and redward of the Balmer break (rest-frame 3645 Å). These criteria enable a self-consistent search for LRD candidates with red optical and blue UV colors over a wide redshift interval (Hainline et al. 2025), although we acknowledge that it does not guarantee an homogeneous physical selection of LRDs.

We replaced photometric redshifts (z_{phot}) with NIRSpec-based spectroscopic redshifts (z_{spec}) taken from the NIRSpec

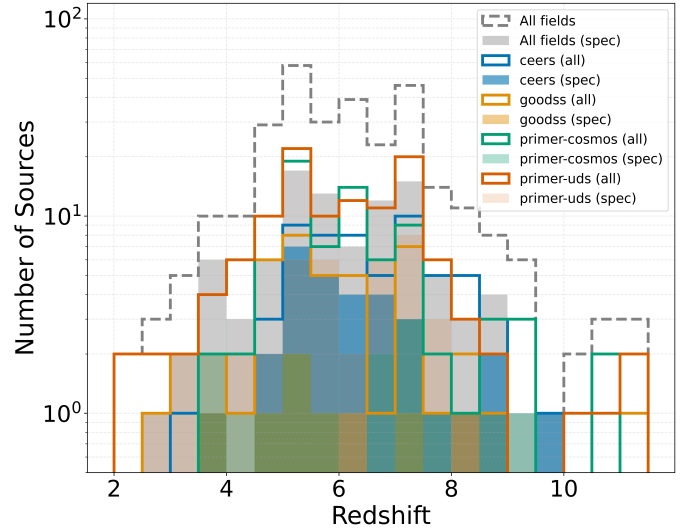


Fig. 1. Redshift distribution of the final LRD sample (dashed black histogram) used in this work. The full sample is split by field (CEERS, GOODS, PRIMER-COSMOS, PRIMER-UDS), as highlighted by the colored open histograms, while filled histograms mark the corresponding subset with spectroscopic redshifts.

Merged Table (version 4¹; Valentino et al. 2025) of the public Dawn JWST Archive (DJA; using MSAEXP; Brammer 2023a; de Graaff et al. 2024; Heintz et al. 2025), whenever it was of high quality (grade = 3, see archive). We augmented this list by adding z_{spec} measurements from the JADES public database² (for the GOODS-S). This enhanced the spectroscopic fraction from 11% (39/341) to 31% (107/341). The mutual agreement between z_{phot} and z_{spec} estimates is very good, with a median absolute deviation (MAD) $\langle |z_{\text{spec}} - z_{\text{phot}}| / (1 + z_{\text{spec}}) \rangle \approx 1.6\%$, and only five catastrophic outliers with $|z_{\text{spec}} - z_{\text{phot}}| / (1 + z_{\text{spec}}) > 0.2$. Hence, we do expect a negligible impact of photo- z uncertainties on the median stacks. Fig. 1 displays the redshift distribution of our final LRD sample (302 sources, grey dashed), split by field (colored open histograms). Filled histograms highlight the spectroscopic subsets. Corresponding numbers are listed in Table 1.

Finally, we restricted ourselves to the NIRCam fields with (at least partial) MIRI coverage in two or more bands (including 18 μm and/or 21 μm ; see Table 1), which yield the strongest IR-SED constraints. Our final sample counts 302 LRDs taken from the following surveys: the Cosmic Evolution Early Release Science Survey (CEERS; Finkelstein et al. 2023) in the EGS; the JWST Advanced Deep Extragalactic Survey (JADES; Eisenstein et al. 2023) in the Great Observatories Origins Deep Survey-South (GOODS-S) field, complemented by MIRI imaging from the Systematic Mid-infrared Instrument Legacy Extragalactic Survey (SMILES; Rieke et al. 2024; Alberts et al. 2024); the Public Release IMAGING for Extragalactic Research (PRIMER³) survey in the COSMOS and UDS fields. Out of 302 LRDs, 163 (54%) are covered by MIRI 7.7 μm , 126 (42%) also by MIRI-18 μm , and 96 (32%) have a robust z_{spec} . Moreover, we exploited archival ALMA/B6 (1.1 mm) data from the combined Automated Mining of the Public ALMA Archive (A3) A3COSMOS+A3GOODSS surveys (Liu et al. 2019; Adscheid et al. 2024; Magnelli et al. 2024), which are

¹ <https://zenodo.org/records/15472354>

² <https://jades.herts.ac.uk/search/>

³ <https://primer-jwst.github.io/>

Table 1. List of the survey fields considered in this study.

Survey/Field	Instrument	Filters	Area (arcmin ²)	N_{stack} range (spec)
CEERS/EGS	NIRCam	[f090w, f115w, f150w, f200w, f277w, f356w, f444w]	88.6	63 (31)
	MIRI	[f560w, f770w, f1000w, f1280w, f1500w, f1800w, f2100w]	55.0	2–39 (2–19)
JADES/GOODSS SMILES/GOODSS	NIRCam	[f090w, f115w, f150w, f200w, f277w, f356w, f444w]	62.1	44 (8)
	MIRI	[f560w, f770w, f1000w, f1280w, f1500w, f1800w, f2100w, f2550w]	34.0	22 (2)
A3-GOODSS	ALMA	Band-6 (1080 μm)	86.8	34 (4)
PRIMER/COSMOS	NIRCam	[f090w, f115w, f150w, f200w, f277w, f356w, f444w]	138.9	79 (15)
	MIRI	[f770w, f1800w]	73.0	41 (7)
A3-COSMOS	ALMA	Band-6 (1080 μm)	127.9	36 (3)
PRIMER/UDS	NIRCam	[f090w, f115w, f150w, f200w, f277w, f356w, f444w]	243.0	116 (42)
	MIRI	[f770w, f1800w]	130.0	61 (23)
All fields	NIRCam	–	531.7	≤ 302 (≤ 96)
	MIRI	–	292.0	≤ 163 (≤ 51)

Notes. For each survey, the sky area, the available filters, and the corresponding number of LRDs (including in brackets those with a spectroscopic redshift) are reported. The range of N_{stack} reflects the varying area across different filters of the same instrument. The global sample is reported at the bottom of this table. Field acronyms: EGS = Extended Groth Strip; GOODS-S = Great Observatories Origins Deep Survey–South; UDS = UKIDSS Ultra Deep Survey; COSMOS = Cosmic Evolution Survey.

available for a total of 70 LRDs in the COSMOS and GOODS-S fields. The number of sources split up by field, photometric coverage, and spectroscopic redshifts is summarized in Table 1.

3. Stacking method and results

Here we outline the input images used for stacking and the stacking method adopted in this work. We performed median stacking in all available NIRCam bands (seven in total) and MIRI bands (two in PRIMER, seven in CEERS, and eight in SMILES), at the optical position of all LRDs. We used a fixed pixel scale of $0.06''$, and we extracted fluxes from a circular aperture with fixed radius $r = 0.25''$. Image and noise maps were retrieved from the corresponding mosaics of each field (see below). At this step, we discarded narrow or medium-band filters (whenever available) to mitigate potential flux boosting in the case of strong emission lines.

We note that stacking was performed in the observed frame, since stacking in the rest-frame would require us to K-correct fluxes by assuming an intrinsic SED, which is instead what we would like to assess, particularly in the rest-frame near-IR (NIR, 1–3 μm), where an AGN-like or a galaxy-like SED alters the interpretation. However, a rest-frame stacking analysis was carried out in Casey et al. (2024) that delivered consistent results with ours. We discuss this comparison in Sect. 4.3.

3.1. Input image maps

We retrieved PRIMER-COSMOS images as part of the COSMOS-Web team, where the full PRIMER-COSMOS mosaic has been reduced following the same procedure as with the COSMOS-Web data (Franco et al. 2025; Harish et al. 2025). All NIRCam and MIRI mosaics have a pixel size of $0.06''$.

For the GOODS-S, we retrieved NIRCam and MIRI images from the JADES and SMILES surveys publicly available on the MAST⁴ archive. While SMILES maps are already at

$0.06''/\text{px}$, JADES maps were rescaled from their native resolution ($0.03''/\text{px}$) to $0.06''/\text{px}$ for consistency with the other fields.

For CEERS and PRIMER-UDS, which are unavailable on MAST, we leveraged the DJA (Valentino et al. 2023; using GRIZLI, Brammer 2023b), where the full reduced mosaics are publicly available in each NIRCam and MIRI band⁵, all with a pixel scale of $0.04''$ and in units of [10 nJy]. We did a resampling of all mosaics to $0.06''/\text{px}$ for consistency with SMILES and PRIMER-COSMOS maps, while preserving astrometry and total flux over a given aperture. For PRIMER-UDS, we also merged the east and west mosaics (listed separately in the DJA) into a full mosaic.

We only considered sources whose position in the map is at least 20 pixels inside the edge of the mosaic. For each object and in each band, we created a $N \times N$ pixel image (with $N = 51$, $0.06''/\text{px}$) from the corresponding map, each centred on the NIRCam position of the LRD. We then stacked all N_{stack} sources and retrieved the median flux at each pixel.

For ALMA, instead, we performed mean stacking in the uv -plane, using all publicly available Band-6 observations covering our LRDs in the COSMOS and GOODS-S fields. Mean stacking is needed in order to homogeneously combine archival data at a different frequency, angular resolution, and sensitivity. This analysis is carried out with CASA, following the method described in Wang et al. (2022) (see also Wang et al. 2024; Magnelli et al. 2024), to which we refer for further details of the procedure.

3.2. Stacked fluxes

In Fig. 2 we show the stacked cutouts ($3'' \times 3''$) in each band, sorted by increasing wavelength. For comparison, the dashed white circle marks the corresponding PSF FWHM. As MIRI images are strongly oversampled, to aid visualization we smoothed them with a Gaussian kernel, adopting a larger radius for larger PSF FWHM (see caption of Fig. 2). The stacked

⁴ SMILES: <https://archive.stsci.edu/hlsp/smiles>; JADES: <https://archive.stsci.edu/hlsp/jades>

⁵ DJA full reduced mosaics in each band are taken from <https://s3.amazonaws.com/grizli-v2/JwstMosaics/v7/index.html>

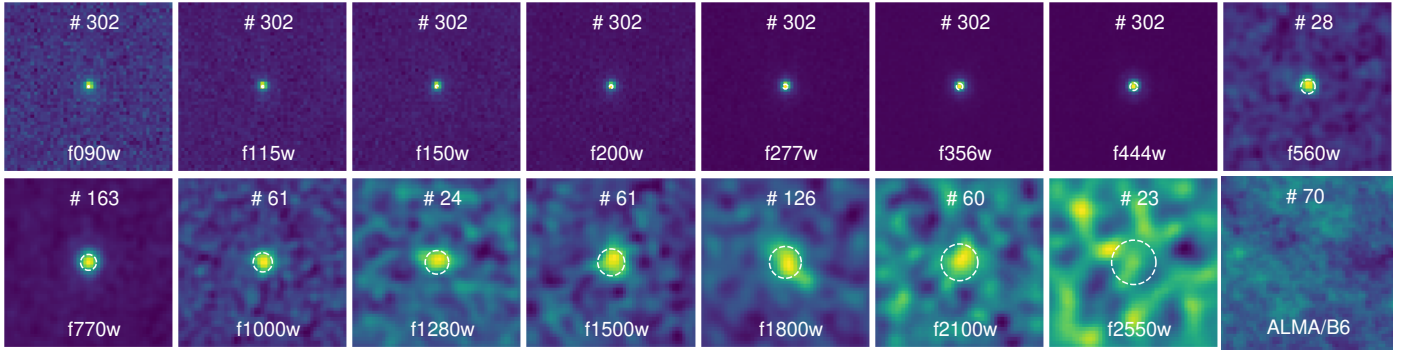


Fig. 2. Median stacked cutouts ($3'' \times 3''$) of LRDs in NIRCcam (7; from f090w to f444w), MIRI (8; from f560w to f2550w), and ALMA/B6 (1.1 mm) bands. All NIRCcam and MIRI images have the same pixel size ($0.06''$). Each cutout reports the corresponding PSF's FWHM (white circle) and the number of stacked sources. To ease visualization, a smoothing with a Gaussian kernel was applied to all MIRI images, with a larger radius for larger FWHM: two pixels for f560w, f770w, and f1000w; three pixels for f1280w and f1500w; four pixels for f1800w and f2100w, and five pixels for f2550w. The stacked ALMA/B6 image was obtained in the uv -plane and imaged at $0.1''/\text{px}$ scale. Except for MIRI/f2550w and ALMA/B6, all other median stacks lead to a $S/N > 3$ detection.

ALMA/B6 image was obtained in the uv -plane and imaged at $0.1''/\text{px}$ scale. Except for MIRI/f2550w and ALMA/B6, all other median stacks lead to a detection with a signal-to-noise ratio (S/N) greater than 3. We stress that all detections are consistent with being unresolved, as their extent is smaller or comparable with the PSF FWHM of the corresponding band (white circles in Fig. 2).

The uncertainty on the stacked flux density was quantified through a bootstrapping technique: if N_{stack} is the number of input targets in a given band, in each random realization we reshuffled the input sample, preserving the same N_{stack} by allowing source duplication. We repeated this 100 times and took the median of the resulting flux distribution as our formal stacked flux. The 1σ dispersion around this value is interpreted as the flux error. The rms of the stacked image was computed by stacking the inverse-variance weighted noise map of each target. Our median stacks yield a detection in each band up to $21 \mu\text{m}$ with signal-to-noise⁶ $S/N > 3$. All fluxes, errors, and rms of the stacked images are listed in Table 2 along with the effective number of sources stacked in each band.

We refer the reader to Appendix A, where we test the robustness of our stacking results against possible biases. In particular, it is worth noting that (rms-weighted) mean stacking delivers systematically higher fluxes than median stacking, by a factor of $\times 2\text{--}3$ in all bands. In Appendix A.2, we demonstrate that, at least in MIRI bands, these fluctuations are entirely due to individual MIRI detections (dropping from 76% at f770w to $\approx 0\%$ at f2100w), without which the mean and median stacked fluxes are matched one another. Therefore, we reiterate that mean stacking is highly sensitive to a few brighter detections, whereas median stacking is more representative of the bulk LRD population.

4. Discussion

4.1. SED fitting with AGN versus without AGN

In order to interpret the stacked photometry of LRDs, we performed SED fitting using the Code Investigating GALaxy Emission (CIGALE; Boquien et al. 2019). We fixed the input redshift to the median value of the LRD sample ($\langle z \rangle \sim 6.2$). In order to account for the broad redshift distribution of the sample

⁶ We consider as “noise” the nominal rms of the stacked image; though using the flux error derived via bootstrapping does not change our conclusions.

Table 2. Median stacked fluxes and errors (open circles in Fig. 3).

Wavelength (μm)	N_{stack}	rms (mJy)	Flux (mJy)	Error flux (mJy)
0.90	302	$3.47\text{e-}09$	$8.49\text{e-}06$	$9.31\text{e-}07$
1.15	302	$4.64\text{e-}09$	$1.31\text{e-}05$	$1.06\text{e-}06$
1.50	302	$8.55\text{e-}09$	$1.47\text{e-}05$	$1.07\text{e-}06$
2.00	302	$1.32\text{e-}08$	$1.78\text{e-}05$	$1.39\text{e-}06$
2.77	302	$1.10\text{e-}08$	$3.03\text{e-}05$	$2.13\text{e-}06$
3.56	302	$1.87\text{e-}08$	$5.89\text{e-}05$	$4.41\text{e-}06$
4.44	302	$3.84\text{e-}08$	$1.13\text{e-}04$	$7.17\text{e-}06$
5.60	28	$3.20\text{e-}06$	$1.23\text{e-}04$	$3.06\text{e-}05$
7.70	163	$2.07\text{e-}06$	$2.13\text{e-}04$	$1.60\text{e-}05$
10.00	61	$1.07\text{e-}05$	$1.94\text{e-}04$	$3.35\text{e-}05$
12.80	24	$3.93\text{e-}05$	$2.68\text{e-}04$	$5.67\text{e-}05$
15.00	61	$3.35\text{e-}05$	$3.42\text{e-}04$	$5.73\text{e-}05$
18.00	126	$4.99\text{e-}05$	$5.33\text{e-}04$	$9.97\text{e-}05$
21.00	60	$1.30\text{e-}04$	$6.08\text{e-}04$	$1.36\text{e-}04$
25.50	23	$1.69\text{e-}03$	$<5.07\text{e-}03$	–
1080.00	70	$5.89\text{e-}03$	$<1.77\text{e-}02$	–

Notes. For each filter, the table reports the number of stacked sources (N_{stack}), the rms of the stacked image, the median stacked flux, and its 1σ uncertainty from bootstrapping, respectively. In the case of non-detection, the 3σ upper limit is provided.

(mostly at $4.5 < z < 8$; see Fig. 1) and the correspondingly wide rest-frame range sampled by each filter, we convolved the filter response curves at $\langle z \rangle \sim 6.2$ with the band-specific redshift distribution of the underlying LRD sample. Then, we ran CIGALE both with and without AGN templates to test the relative significance of the AGN component. In both SED runs, the library of stellar templates was taken from Bruzual & Charlot (2003), including a delayed star formation history with random bursts. Energy balance is assumed between the dust-absorbed UV radiation and the reprocessed IR emission, modeled via a dust attenuation curve from Charlot & Fall (2000), and dust templates from Dale et al. (2014). The AGN templates come from the SKIRTOR module (Stalevski et al. 2012, 2016), which includes emission from an accretion disk, a dusty torus, and polar dust. A comprehensive list of the SED-fitting parameters is given in Appendices B and C.

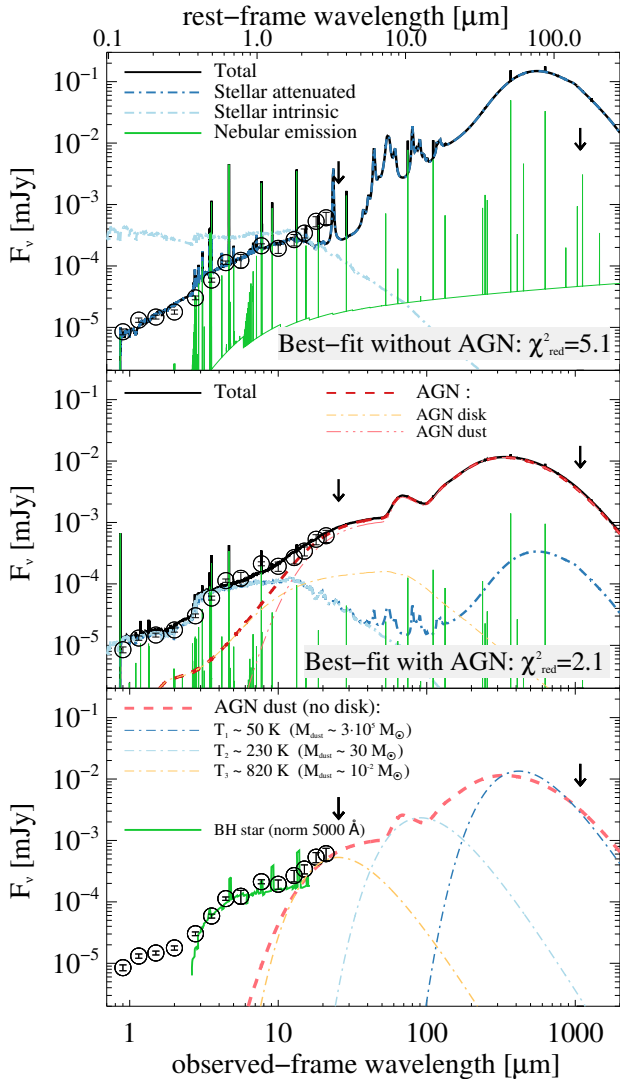


Fig. 3. Best SED fits obtained with CIGALE on the stacked photometry (open circles) at $z = 6.2$. Top panel: Fit including only galaxy templates. Middle panel: Fit including both galaxy and AGN templates. The fitting from the AGN run is strongly favored based on χ^2 arguments and ALMA constraints. Bottom panel: Decomposition of the AGN-dust template (dashed red) in three multi-temperature graybody models (dot-dashed lines). The normalized BH-star SED (Naidu et al. 2025) is overlaid for comparison (green) and not included in the SED fitting. Details are given in Section 4.1.

In Fig. 3, we display all stacked fluxes (or 3σ upper limits) overlaid to the total best-fit SED (black line) obtained from the run without AGN (top panel) and with AGN (middle panel). In both runs, the total best fit (solid black line) is obtained by adding up the attenuated stellar template (dot-dashed blue line), the nebular emission (green lines), and the AGN (dashed red lines). The AGN template shown in Fig. 3 breaks down into an accretion disk component (dot-dashed yellow line) peaking in the optical NIR, and a dusty component (clumpy torus and polar dust; dot-dashed red line) peaking at longer wavelengths. We further note that the best-fit stellar template chosen in the AGN run has very little dust attenuation ($A_V \sim 0.01$) and the far-IR emission is therefore much fainter than the ALMA upper limit and entirely attributed to AGN (polar) dust.

Regardless of the origin of optical and/or UV emission in LRDs (see Sect. 4.2 for more details), we argue that AGN-

heated dust is needed to reproduce the monothonic increase in the rest 1–3 μm , which pure dust-attenuated stellar light fails to match. This is further supported by our ALMA upper limit, which is very similar to that derived by Casey et al. (2025) (see also Xiao et al. 2025; Setton et al. 2025), and that translates into $M_{\text{dust}} \lesssim 10^6 M_{\odot}$ for $T \sim 100$ K. Indeed, the best-fit SED without AGN strongly overshoots the ALMA upper limit when attempting to reproduce the power-law-like SED at rest 1–3 μm . While this could be resolved by lifting the energy balance implemented in CIGALE, even ignoring the upper limits (at f2550w and ALMA/B6), a simple F-test comparing the two reduced χ^2 values favors the AGN solution at $>95\%$ significance.

4.2. Caveats and limitations of broadband SED fitting

We caution the reader that our SED analysis does not necessarily endorse a pure stellar origin of the optical and/or UV continuum, despite the stellar template chosen by CIGALE in the SED fitting. We highlight some caveats and limitations inherent to our photometric SED analysis. For instance, CIGALE assumes a single-zone obscuration between optical, UV, and IR light, which might be oversimplistic in the case of no co-spatiality (see Sect. 4.4). In addition, broad emission lines – such as those observed in most LRD spectra – are not properly plugged into the CIGALE library. More importantly, despite the fact that our median stacking mitigates broad emission line contamination, it relies on broadband photometry, which most likely irons out the strongest Balmer breaks, with a potential impact on the output stellar ages and M_{\star} .

Regarding limitations on the AGN templates, instead, while a BH-Star model (Naidu et al. 2025), or a quasi-star model (Santarelli et al. 2025), could be better suited for reproducing the strong Balmer breaks, Balmer line absorption, and broad emission lines (e.g., de Graaff et al. 2025), a satisfactory fit of the full SED (as shown in Fig. 3, bottom panel) further requires a dusty torus and (in most cases) an additional UV component. In this framework, accommodating a joint BH star (or quasi star) and a dusty-torus template demands radiative transfer modeling (following e.g., SKIRTOR; Stalevski et al. 2012, 2016).

Therefore, a more faithful modeling of the spectrophotometric LRD features would likely require an ad-hoc grid of (AGN) models with those built-in features. Besides the uncertainty related to the representativeness of such realistic models, this task is beyond the scope of this paper, which focuses on hot-dust emission from stacking photometric data.

4.3. Comparison with literature LRD SEDs

In the literature, LRDs are inevitably selected with distinct criteria, depending on the availability of multi-band NIRCcam and MIRI imaging, and rest-frame optical spectra. While our LRD sample meets the same compactness and V-shape criteria across all fields, we compare our stacked SED against other LRD SEDs based on different selections. Fig. 4 displays a number of stacked and single LRD SEDs for comparison, all in the rest frame and normalized to the rest 1 μm flux density to emphasize the difference in NIR slope.

Among the median stacked SEDs, we show: Akins et al. (2025; red starred symbols); the best-fit galaxy and AGN SED fits from Casey et al. (2024; dot-dashed lines); and the median stacked SED obtained from the LRD sample of Kokorev et al. (2024; blue squares). Akins et al. (2025) selected 434 LRDs in the COSMOS-Web area (0.5 deg^2 ; Casey et al. 2023) that simultaneously meet a compactness cut (in f444w), a very red color

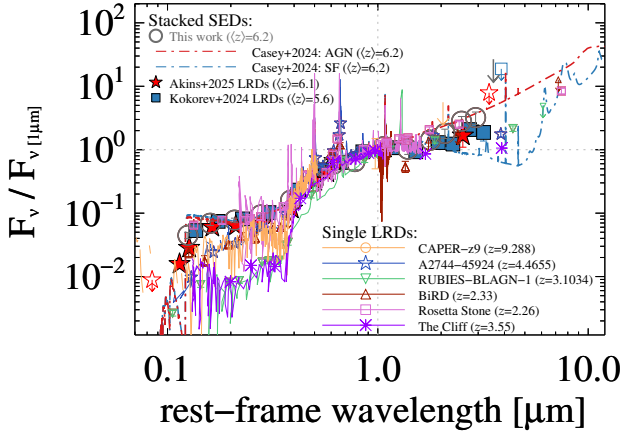


Fig. 4. Comparison between our rest-frame median stacked SED (open gray circles) and other literature spectra or SEDs of LRDs. All fluxes are normalized to the rest-frame $1 \mu\text{m}$ flux density. We show the stacked SEDs from Kokorev et al. (2024), Akins et al. (2025) and Casey et al. (2024; both AGN and galaxy fits as red and blue dot-dashed lines, respectively). Other single LRD SEDs include: CAPER-z9 ($z = 9.288$; Taylor et al. 2025); A2744-45924 ($z = 4.4655$; Labbe et al. 2024); RUBIES-BLAGN-1 ($z = 3.1034$; Wang et al. 2025); the BiRD ($z = 2.33$; Loiacono et al. 2025); the Rosetta Stone (or GN 280754; $z = 2.26$; Juodžbalis et al. 2024b) and The Cliff ($z = 3.55$; de Graaff et al. 2025).

cut ($f_{277\text{w}} - f_{444\text{w}} > 1.5$), and a bad brown dwarf fit, without explicitly requiring a V-shaped SED due to limitations in NIR-Cam imaging. Their median SED agrees remarkably well with ours, though their statistics in MIRI is limited to $7.7 \mu\text{m}$ (over 1/3 of COSMOS-Web) and $18 \mu\text{m}$ (from the inner PRIMER-COSMOS footprint, $\approx 4\%$ of the full sample), this latter leading to tentative detection ($0.43 \pm 0.15 \text{ mJy}$). We show that their $18 \mu\text{m}$ flux is fully consistent with the rising MIR slope seen longward $10 \mu\text{m}$. Their upper limits are here re-scaled to 3σ . Despite the different LRD selection and survey field ($\lesssim 25\%$ of our LRDs are in Akins et al. 2025), this agreement further suggests that the compactness and the conservative color selection by Akins et al. (2025) broadly compensate for the need for a V-shaped SED, which leads to self-consistent LRD samples.

Kokorev et al. (2024) selected 260 LRDs over the same JWST Legacy fields as this study, based on a compactness (in $f_{444\text{w}}$) criterion and various color cuts that mimic a V-shape requirement (see Sect. 3.1 in Kokorev et al. 2024 for details). From their original sample, we further updated the redshift measurements with additional z_{spec} as done in Sect. 2, and repeated the same median stacking analysis carried out for our LRD sample; we found good consistency.

In Fig. 4 we also compare our stacks against the median best-fit SED by Casey et al. (2024), which was obtained from stacking a joint LRD sample (675 LRDs) from Akins et al. (2025) and Kokorev et al. (2024). In their work, Casey et al. (2024) performed stacking in rest-frame wavelength bins, applying a K-correction based on the observed SED of each LRD. The median stacked SED was then redshifted to $z = 6.2$ (as for our sample) and fitted with either a pure AGN (dashed red) or a pure galaxy template (dashed blue). We emphasize that our median stacks are fully consistent with their median AGN SED within their uncertainties (not reported here). Our ALMA upper limit at $\approx 1.1 \text{ mm}$ is also fully consistent with the ALMA 1.3 mm upper limit derived in Casey et al. (2025) via stacking 60 LRDs retrieved from the Akins et al. (2025) sample.

In addition, in Fig. 4 we display the spectra and photometric SEDs of some well-studied LRDs for comparison, namely: CAPER-z9 ($z = 9.288$; Taylor et al. 2025); A2744-45924 ($z = 4.4655$; Labbe et al. 2024); RUBIES-BLAGN-1 ($z = 3.1034$; Wang et al. 2025); the “Big Red Dot” (BiRD, $z = 2.33$; Loiacono et al. 2025); the “Rosetta Stone” (or “GN 280754”; $z = 2.26$; Juodžbalis et al. 2024b) and “The Cliff” ($z = 3.55$; de Graaff et al. 2025). This compilation showcases the variety of rest-frame $1-3 \mu\text{m}$ SEDs of bright LRDs, as compared to median stacked SEDs. As noted here and in other works (e.g., Barro et al. 2024; Loiacono et al. 2025; Setton et al. 2025; Ronayne et al. 2025), most of the brightest LRDs do not seem to show a rising NIR slope but flat, or even decreasing, rest-frame $1-3 \mu\text{m}$ SEDs. This is at odds with recent MIRI stacked SEDs, which are dominated by non-detections (only $\lesssim 10\%$ are detected longward $f_{1000\text{w}}$, see Appendix A.2). Such a comparison further highlights that LRDs are not a single homogeneous population, and that our median stacking in the longest MIRI bands is more representative of intrinsically fainter sources.

Another meaningful comparison concerns the work by Williams et al. (2024), who stacked LRDs over all MIRI bands within the GOODS-SMILES footprint. They found a flat NIR SED up to $18 \mu\text{m}$ (and upper limits longward), thus excluding a dominant contribution from AGN-heated dust. Besides the small numbers (only nine objects) and the slightly higher redshift compared to our sample ($\langle z \rangle \sim 7.7$ instead of $\langle z \rangle \sim 6.2$), their LRDs were pre-selected from a pool of H-band dropouts (i.e., Hubble Space Telescope; HST-dark galaxies), hence likely biased toward the most obscured sources. We repeated the stacking for the same nine sources and found the same median SED as in Williams et al. (2024). This check suggests that their sample might be sensitive to more obscured sources than conventional LRDs, especially considering their H-dropout selection, combined with the lack of compactness and V-shaped requirements. These differences might link to a different obscuration origin relative to typical LRDs. Hence we caution that these nine LRDs are not representative of our photometrically selected LRDs (from Kocovski et al. 2025), despite being an interesting subsample to follow up.

4.4. Are LRDs typically dust-obscured AGN?

By considering the best-fit solution with AGN, the rest-frame $6 \mu\text{m}$ luminosity of the AGN torus can be used as a proxy for the intrinsic AGN X-ray (rest $[2-10] \text{ keV}$) luminosity ($L_{\text{X, intr}}$ (e.g. Stern 2015)). From $L_{6\mu\text{m}} = (3.8 \pm 0.6) \times 10^{43} \text{ erg/s}$, we derive $L_{\text{X, intr}} = 2 \times 10^{43} \text{ erg/s}$, which nicely matches that expected from the best-fit AGN bolometric luminosity, if we assume standard X-ray bolometric corrections (e.g., Lusso et al. 2012). By stacking the deepest X-ray images of the 44 LRDs that fall inside the Chandra Deep Field South (Luo et al. 2017), we obtain a 3σ $[2-10] \text{ keV}$ luminosity limit $L_{\text{X}} \lesssim 10^{42} \text{ erg/s}$ at $z = 6.2$ (Comastri et al. 2025). In order to match it with the $L_{\text{X, intr}}$ expected from the AGN SED, the hydrogen column density has to be $N_{\text{H}} \gtrsim 3 \times 10^{24} \text{ cm}^{-2}$, which implies Compton-thick gas obscuration (Ananna et al. 2024; Yue et al. 2024; Maiolino et al. 2025; Sacchi & Bogdán 2025).

Given the SED-based estimate of the AGN bolometric luminosity, $L_{\text{bol, AGN}} = (2.5 \pm 0.4) \times 10^{44} \text{ erg s}^{-1}$ and the best-fit stellar mass from the AGN run ($M_{\star} \sim 2.5 \times 10^9 M_{\odot}$; i.e., $\approx 2\times$ smaller than from a pure-galaxy fit), we emphasize that, if LRDs are highly accreting AGN (Eddington ratio $\gtrsim 10\%$), the BH-to-galaxy mass ratio $M_{\text{BH}}/M_{\star} \gtrsim 1:100$, implies strongly overmassive BHs compared to local scaling relations

(e.g. Reines & Volonteri 2015). Moreover, our $L_{\text{bol,AGN}}$ estimate corresponds to a dust sublimation radius of $R_{\text{sub}} \approx 0.1$ pc at $T_{\text{sub}} \approx 1500$ K (Barvainis 1987), meaning that any dusty AGN structure should begin at scales ≥ 0.1 pc from the central SMBH. In the bottom panel of Fig. 3, we try to decompose the AGN-dust template (dashed red curve) with three multi-temperature graybody curves (leaving T free to vary). Our rising MIRI SED up to rest-frame $2\text{--}3\mu\text{m}$ indicates hot dust at $T \approx 820$ K, though it could be even hotter if optically thick. Instead, the bulk dust mass ($M_{\text{dust}} \approx 3 \times 10^5 M_{\odot}$) comes from colder polar AGN dust at ≈ 50 K (though not constrained by our data). While a dusty structure could induce X-ray obscuration, it would also strongly suppress the optical and/or UV continuum, in tension with the blue UV slope and strong Balmer break features that imply relatively modest dust attenuation ($A_V \sim 2\text{--}3$; Akins et al. 2025; Casey et al. 2025).

This apparent contradiction can be resolved if X-ray absorption in LRDs occurs within R_{sub} (< 0.1 pc), in a gas-dense and dust-free region, implying no co-spatiality. We stress that, despite the fact that our stacking analysis suggests that most ($\geq 50\%$) LRDs contain hot-AGN dust, they do not need to be dust obscured. This gas-dust displacement implies that theoretical models of quasi stars (Begelman & Dexter 2025), BH stars (Naidu et al. 2025), and super-massive stars (Nandal & Loeb 2025), can in principle co-exist with outer hot dust, albeit its formation scenario (in-situ or ex-situ) remains open.

Our findings also fit with an alternative hypothesis (Setton et al. 2024; Maiolino et al. 2025; D'Eugenio et al. 2025a; Ji et al. 2025; Inayoshi & Maiolino 2025) that postulates that X-ray obscuration arises from Compton-thick gas clouds within the BLR (at $R < R_{\text{sub}}$, i.e., dust-free), which covers the compact X-ray corona without suppressing the broad Balmer lines. However, the origin of the optical and/or UV continuum remains open. In this debated framework, recent works report the presence of a forest of auroral and FeII lines in some luminous LRDs (e.g., Tripodi et al. 2025; Lin et al. 2025; Lambrides et al. 2025; Torralba et al. 2025; D'Eugenio et al. 2025b), which indicate that – at least in these bright objects – part of the UV continuum is dominated by a compact AGN embedded in a dense gas cocoon that (at least partly) explains the Balmer break and the broadening of the Balmer lines. Regardless of the AGN orientation, we argue that the hot dust revealed in this work should be placed most likely at outer radii ($R_{\text{sub}} > 0.1$ pc) than the BLR scales at which the (dust-free) gas absorption takes place. Given that $R_{\text{sub}} \propto L_{\text{bol,AGN}}^{1/2}$, we expect that even in the most luminous LRDs (e.g., Torralba et al. 2025; D'Eugenio et al. 2025b, i.e., $50\text{--}100\times$ brighter than our median stacked AGN), $R_{\text{sub}} \geq 1$ pc, which is $\approx 10^4$ times larger than the nominal thickness of the BLR (gas) cloud (D'Eugenio et al. 2025b).

Although we caution that LRDs might not consist of a single population, we speculate two simple scenarios in Fig. 5, depending on what powers the optical and/or UV continuum in LRDs. In either case, the observed optical and/or UV continuum comes from dust-unattenuated regions, despite the presence of hot AGN dust, and the outer BLR is directly visible along the line of sight.

(i) If the optical and/or UV comes from unobscured stellar light (Fig. 5, left), it must be placed at larger scales than the AGN torus. In this case the LRD resembles a Type-2 AGN configuration, in which the torus covers the AGN accretion disk, which would otherwise outshine the host. In the case of high torus clumpiness, mildly attenuated light from the outer accretion disk and the BLR could pierce through the torus and add to the optical and/or UV emission.

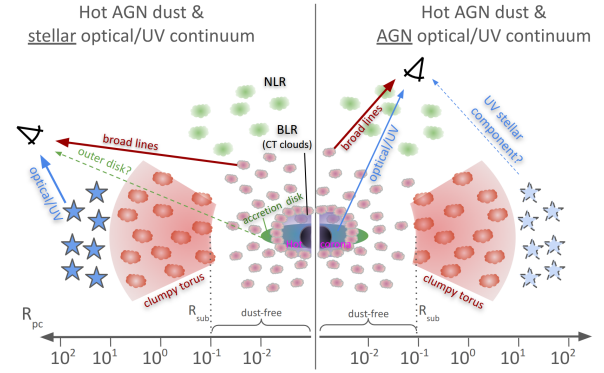


Fig. 5. Sketch of the proposed hot AGN-dust scenario, both for stellar-dominated (left) and AGN-dominated (right) optical and/or UV continuum.

(ii) Instead, if the optical and/or UV is dominated by AGN light (Fig. 5, right), the LRD resembles a Type-1 unobscured AGN configuration, though the origin of the sharp Balmer break remains unclear. Fainter optical and/or UV light arising from (kiloparsec-scale) star formation cannot be ruled out (e.g., Rinaldi et al. 2025).

5. Conclusions

In this manuscript, we investigate the nature of LRDs via an extensive NIRCcam, MIRI, and ALMA stacking analysis of a homogeneously selected sample of 302 sources across the CEERS, JADES, and PRIMER JWST fields. We fit the median stacked SED with either pure galaxy or galaxy+AGN templates to interpret the average IR emission properties of LRDs.

In summary, we find clear evidence for a rising MIRI continuum up to $\lambda_{\text{rest}} \sim 3\mu\text{m}$, which we interpret via SED fitting as hot (≈ 820 K) AGN dust, regardless of whether the optical and/or UV continuum is stellar- or AGN-dominated, or a combination thereof. Despite the presence of hot dust, these LRDs selected via V-shaped and compactness criteria must be not strongly dust-obscured. The non-detection from deep X-ray stacking suggests that the apparent X-ray weakness of LRDs is caused by Compton-thick ($N_{\text{H}} > 3 \times 10^{24} \text{ cm}^{-2}$) gas obscuration.

We acknowledge that our homogeneous LRD selection does not necessarily translate into a single LRD population with similar physical properties. However, we reiterate that our main empirical finding is unchanged, that is, the majority ($\geq 50\%$) of LRDs show AGN-heated dust emission. If this were not the case, the rising NIR slope would not appear in the median stacked data, and the mean and median MIRI fluxes would differ significantly after removing individual detections (contrary to our findings in Appendix A.2). The small fraction of individual MIRI detections ($\leq 10\%$ longward $f1000w$, i.e., longer than rest $1.3\mu\text{m}$) contributes negligibly to the median stacked SED at rest-frame $1\text{--}3\mu\text{m}$. Therefore, the rising NIR slope is predominantly driven by the MIRI non-detections, which suggests that AGN-heated dust is characteristic of the fainter, more typical LRD population.

That being said, our results are contingent upon the adopted LRD selection criteria, which may be inherently biased. For instance, if the rest-frame UV emission of LRDs arises from an AGN accretion disk, we speculate that photometric LRD selection criteria (especially the blue UV-slope requirement) would preferentially select Type-1 AGN, according to the Unified Model. The same criterion would therefore induce a bias

against the Type-2 LRD analogs, in which the disk is hidden by an obscuring dusty structure. With this caveat in mind, our results should be taken as an empirical benchmark for motivating deeper investigations of the full LRD population.

Acknowledgements. The authors thank the anonymous referee for their constructive report. ID thanks Federica Loiacono for sharing their SED datapoints, and also Giovanni Mazzolari, Yoshiki Matsuoka and Guillermo Barro for insightful discussions. ID acknowledges funding by the European Union – NextGenerationEU, RRF M4C2 1.1, Project 2022JZJBHM: “AGN-sCAN: zooming-in on the AGN-galaxy connection since the cosmic noon” – CUP C53D23001120006. SB is funded by ERC grant 101076080. MG wishes to acknowledge the generous support and hospitality, by the CEA – Université Paris-Saclay during his visit, when part of the research related to this work was carried out. (Some of) the data products presented here were retrieved from the DJA, an initiative of the Cosmic Dawn Center (DAWN), which is funded by the Danish National Research Foundation under grant DNR140. We acknowledge the contribution of the COSMOS collaboration, consisting of more than 200 scientists. More information about the COSMOS survey can be found at <https://cosmos.astro.caltech.edu/>. This work is based [in part] on observations made with the NASA/ESA/CSA James Webb Space Telescope. The data were obtained from the Mikulski Archive for Space Telescopes at the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-03127 for JWST. These observations are associated with program #1180, #1181, #1210, #1286, #1895, #1963, #3215, #1345, #1837 and #1207. The data described here may be obtained from the MAST archive at <https://doi.org/10.17909/et3f-zd57> (SMILES) and <https://dx.doi.org/10.17909/8tdj-8n28> (JADES) and on the DJA at <https://s3.amazonaws.com/grizli-v2/JwstMosaics/v7/index.html> (CEERS and PRIMER-UDS).

References

- Adscheid, S., Magnelli, B., Liu, D., et al. 2024, *A&A*, **685**, A1
- Akins, H. B., Casey, C. M., Lambrides, E., et al. 2025, *ApJ*, **991**, 37
- Alberts, S., Lyu, J., Shivaee, I., et al. 2024, *ApJ*, **976**, 224
- Ananna, T. T., Bogdán, Á., Kovács, O. E., Natarajan, P., & Hickox, R. C. 2024, *ApJ*, **969**, L18
- Barro, G., Perez-Gonzalez, P. G., Kocevski, D. D., et al. 2024, *ApJ*, submitted [arXiv:2412.01887]
- Barvainis, R. 1987, *ApJ*, **320**, 537
- Begelman, M. C., & Dexter, J. 2025, *AJ*, submitted [arXiv:2507.09085]
- Boquien, M., Burgarella, D., Roehlly, Y., et al. 2019, *A&A*, **622**, A103
- Brammer, G. 2023a, <https://doi.org/10.5281/zenodo.8319596>
- Brammer, G. 2023b, <https://doi.org/10.5281/zenodo.8370018>
- Bruzual, G., & Charlot, S. 2003, *MNRAS*, **344**, 1000
- Carranza-Escudero, M., Conselice, C. J., Adams, N., et al. 2025, *ApJ*, **989**, L50
- Casey, C. M., Kartaltepe, J. S., Drakos, N. E., et al. 2023, *ApJ*, **954**, 31
- Casey, C. M., Akins, H. B., Kokorev, V., et al. 2024, *ApJ*, **975**, L4
- Casey, C. M., Akins, H. B., Finkelstein, S. L., et al. 2025, *ApJ*, **990**, L61
- Cenci, E., & Habouzit, M. 2025, *MNRAS*, **542**, 2597
- Chabrier, G. 2003, *PASP*, **115**, 763
- Charlot, S., & Fall, S. M. 2000, *ApJ*, **539**, 718
- Chen, K., Li, Z., Inayoshi, K., & Ho, L. C. 2025, *ApJ*, submitted [arXiv:2505.22600]
- Comastri, A., Lanzuisi, G., Vito, F., et al. 2025, *A&A*, submitted [arXiv:2510.00112]
- Dale, D. A., Helou, G., Magdis, G. E., et al. 2014, *ApJ*, **784**, 83
- de Graaff, A., Rix, H.-W., Carniani, S., et al. 2024, *A&A*, **684**, A87
- de Graaff, A., Rix, H.-W., Naidu, R. P., et al. 2025, *A&A*, **701**, A168
- D’Eugenio, F., Maiolino, R., Perna, M., et al. 2025a, *MNRAS*, submitted [arXiv:2503.11752]
- D’Eugenio, F., Nelson, E., Ji, X., et al. 2025b, *ApJ*, submitted [arXiv:2510.00101]
- Eisenstein, D. J., Johnson, B. D., Robertson, B., et al. 2023, *ApJS*, accepted [arXiv:2310.12340]
- Finkelstein, S. L., Bagley, M. B., Ferguson, H. C., et al. 2023, *ApJ*, **946**, L13
- Franco, M., Casey, C. M., Koekemoer, A. M., et al. 2025, *ApJ*, submitted [arXiv:2506.03256]
- Gloudeans, A. J., Duncan, K. J., Eilers, A.-C., et al. 2025, *ApJ*, **986**, 130
- Greene, J. E., Labbe, I., Goulding, A. D., et al. 2024, *ApJ*, **964**, 39
- Greene, J. E., Setton, D. J., Furtak, L. J., et al. 2025, *ApJ*, submitted [arXiv:2509.05434]
- Hainline, K. N., Maiolino, R., Juodžbalis, I., et al. 2025, *ApJ*, **979**, 138
- Harish, S., Kartaltepe, J. S., Liu, D., et al. 2025, *ApJ*, **992**, 45
- Heintz, K. E., Brammer, G. B., Watson, D., et al. 2025, *A&A*, **693**, A60
- Hviding, R. E., de Graaff, A., Miller, T. B., et al. 2025, *A&A*, **702**, A57
- Inayoshi, K., & Maiolino, R. 2025, *ApJ*, **980**, L27
- Inayoshi, K., Visbal, E., & Haiman, Z. 2020, *ARA&A*, **58**, 27
- Inayoshi, K., Murase, K., & Kashiyama, K. 2025, *ApJ*, submitted [arXiv:2509.19422]
- Ji, X., Maiolino, R., Übler, H., et al. 2025, *MNRAS*, submitted [arXiv:2501.13082]
- Juodžbalis, I., Maiolino, R., Baker, W. M., et al. 2024a, *Nature*, **636**, 594
- Juodžbalis, I., Ji, X., Maiolino, R., et al. 2024b, *MNRAS*, **535**, 853
- Juodžbalis, I., Marconcini, C., D’Eugenio, F., et al. 2025, *Nature*, submitted [arXiv:2508.21748]
- King, A. 2025, *MNRAS*, **536**, L1
- Kocevski, D. D., Finkelstein, S. L., Barro, G., et al. 2025, *ApJ*, **986**, 126
- Kokorev, V., Caputi, K. I., Greene, J. E., et al. 2024, *ApJ*, **968**, 38
- Labbé, I., van Dokkum, P., Nelson, E., et al. 2023, *Nature*, **616**, 266
- Labbe, I., Greene, J. E., Matthee, J., et al. 2024, *ApJ*, submitted [arXiv:2412.04557]
- Labbe, I., Greene, J. E., Bezanson, R., et al. 2025, *ApJ*, **978**, 92
- Lambrides, E., Garofali, K., Larson, R., et al. 2024, *ApJ*, submitted [arXiv:2409.13047]
- Lambrides, E., Larson, R., Hutchison, T., et al. 2025, *ApJ*, submitted [arXiv:2509.09607]
- Leung, G. C. K., Finkelstein, S. L., Pérez-González, P. G., et al. 2025, *ApJ*, **992**, 26
- Lin, X., Fan, X., Wang, F., et al. 2025, *ApJ*, submitted [arXiv:2504.08039]
- Liu, D., Schinnerer, E., Groves, B., et al. 2019, *ApJ*, **887**, 235
- Liu, H., Jiang, Y. F., Quataert, E., Greene, J. E., & Ma, Y. 2025, *ApJ*, submitted [arXiv:2507.07190]
- Loiacono, F., Gilli, R., Mignoli, M., et al. 2025, *A&A*, **703**, A36
- Luo, B., Brandt, W. N., Xue, Y. Q., et al. 2017, *ApJS*, **228**, 2
- Lusso, E., Comastri, A., Simmons, B. D., et al. 2012, *MNRAS*, **425**, 623
- Madau, P. 2025, *ApJ*, submitted [arXiv:2501.09854]
- Madau, P., & Haardt, F. 2024, *ApJ*, **976**, L24
- Magnelli, B., Adscheid, S., Wang, T.-M., et al. 2024, *A&A*, **688**, A55
- Maiolino, R., Scholtz, J., Curtis-Lake, E., et al. 2024, *A&A*, **691**, A145
- Maiolino, R., Rinaldi, G., Signorini, M., et al. 2025, *MNRAS*, **538**, 1921
- Matthee, J., Naidu, R. P., Brammer, G., et al. 2024, *ApJ*, **963**, 129
- Mazzolari, G., Scholtz, J., Maiolino, R., et al. 2025, *A&A*, **700**, A12
- Naidu, R. P., Matthee, J., Katz, H., et al. 2025, *Nature*, submitted [arXiv:2503.16596]
- Nandal, D., & Loeb, A. 2025, *ApJ*, submitted [arXiv:2507.12618]
- Pacucci, F., & Narayan, R. 2024, *ApJ*, **976**, 96
- Pacucci, F., Nguyen, B., Carniani, S., Maiolino, R., & Fan, X. 2023, *ApJ*, **957**, L3
- Pérez-González, P. G., Barro, G., Rieke, G. H., et al. 2024, *ApJ*, **968**, 4
- Reines, A. E., & Volonteri, M. 2015, *ApJ*, **813**, 82
- Rieke, G. H., Alberts, S., Shivaee, I., et al. 2024, *ApJ*, **975**, 83
- Rinaldi, P., Rieke, G. H., Wu, Z., et al. 2025, *ApJ*, submitted [arXiv:2507.17738]
- Ronayne, K., Papovich, C., Kirkpatrick, A., et al. 2025, *ApJ*, submitted [arXiv:2508.20177]
- Rusakov, V., Watson, D., Nikopoulos, G. P., et al. 2025, *Nature*, submitted [arXiv:2503.16595]
- Sacchi, A., & Bogdán, Á. 2025, *ApJ*, **989**, L30
- Santarelli, A. D., Campbell, C. B., Farag, E., et al. 2025, *AJ*, submitted [arXiv:2510.11772]
- Setton, D. J., Greene, J. E., de Graaff, A., et al. 2024, *ApJ*, submitted [arXiv:2411.03424]
- Setton, D. J., Greene, J. E., Spilker, J. S., et al. 2025, *ApJ*, **991**, L10
- Stalevski, M., Fritz, J., Baes, M., Nakos, T., & Popović, L. Č. 2012, *MNRAS*, **420**, 2756
- Stalevski, M., Ricci, C., Ueda, Y., et al. 2016, *MNRAS*, **458**, 2288
- Stern, D. 2015, *ApJ*, **807**, 129
- Taylor, A. J., Kokorev, V., Kocevski, D. D., et al. 2025, *ApJ*, **989**, L7
- Torralba, A., Matthee, J., Pezzulli, G., et al. 2025, *A&A*, submitted [arXiv:2510.00103]
- Tripodi, R., Bradač, M., D’Eugenio, F., et al. 2025, *ApJ*, submitted [arXiv:2507.20684]
- Valentino, F., Brammer, G., Gould, K. M. L., et al. 2023, *ApJ*, **947**, 20
- Valentino, F., Heintz, K. E., Brammer, G., et al. 2025, *A&A*, **699**, A358
- Volonteri, M., Habouzit, M., & Colpi, M. 2021, *Nat. Rev. Phys.*, **3**, 732
- Wang, T.-M., Magnelli, B., Schinnerer, E., et al. 2022, *A&A*, **660**, A142
- Wang, T.-M., Magnelli, B., Schinnerer, E., et al. 2024, *A&A*, **681**, A110
- Wang, B., de Graaff, A., Davies, R. L., et al. 2025, *ApJ*, **984**, 121
- Williams, C. C., Alberts, S., Ji, Z., et al. 2024, *ApJ*, **968**, 34
- Xiao, M., Oesch, P. A., Bing, L., et al. 2025, *A&A*, **700**, A231
- Yue, M., Eilers, A.-C., Ananna, T. T., et al. 2024, *ApJ*, **974**, L26

Appendix A: Stacking tests

We test our stacking results against multiple potential biases related to sample selection, stacking method, field-to-field variations and photometric redshift uncertainties. The main outcome of these tests is shown in Fig. A.1. In general, these tests further corroborate the robustness of our stacking results and the need for an obscured AGN component to reproduce the MIR excess. Each of the following tests is discussed in a dedicated subsection, numbered as the panels of Fig. A.1 from top to bottom.

A.1. Mean vs median stacking

In the first panel of Fig. A.1 we compare our median stacked fluxes (grey open circles) against mean stacked fluxes (red starred symbols). These latter values are computed from the inverse-variance weighted flux density map, accounting for the different sensitivity (i.e. exposure time) of each target. Mean stacking leads to higher fluxes with respect to median stacking. This is likely caused by the fact that the resulting (rms-weighted) linear mean flux can be boosted upwards by relatively bright individual detections (we note that upper limits at f2550w and ALMA/B6 are unchanged as they are calculated from the rms map). This is the case in NIRCcam bands, where all LRDs are individually detected by definition over multiple bands, both blueward and redward the Balmer break (in order to assess the V-shape). An even stronger boost is seen at MIRI 7.7 μ m and 18 μ m, which count about 100 more targets than at other MIRI bands (see Table 1 and Appendix A.2). In these two MIRI bands, we investigate the origin of these fluctuations, finding that the f770w and f1800w fluxes are over-boosted relative to other bands because of two strong outliers at $z < 4.5$ in the PRIMER-UDS field. These are: RUBIES-BLAGN1 (or PRIMER-UDS 38223, at $z_{\text{spec}}=3.1$; Wang et al. 2025), which is a bright X-ray detected BLAGN; and PRIMER-UDS 9235 ($z_{\text{phot}}=4.45$), previously flagged in Leung et al. (2025) to be the strongest f1800w emitter of the PRIMER-UDS field. We check that after removing only those two PRIMER-UDS outliers, the ratio between mean and median stacked fluxes (over all fields) re-aligns to the average factor of $\times 2$ -3 seen in all other bands. We do not remove these outliers in the rest of our work, in order to not fine-tune the results. However, this is again a caveat on using mean stacks without pre-removing bright individual detections. Moreover, we opt for median stacking also because it strongly mitigates contamination from bright neighbors, and thus reduces the confusion noise for the faint sources. This comparison, however, demonstrates that mean and median stacked SEDs display broadly-similar shapes.

A.2. Contribution of MIRI detections to the stacking

We test the robustness of our stacking results against the contribution of individual detections. We repeat both mean (open stars) and median (open circles) stacking after removing individual MIRI detections (with $S/N > 3$), in each band. In Figure A.1 (2nd panel from top) we display the new stacking results. The individual MIRI detections are shown as small blue circles, and the corresponding detection rate is reported at each MIRI band. We notice that mean and median stacking of MIRI non-detections lead to well-consistent results, also in f560w and f770w where the detection rate is above 70%. In addition, for MIRI non-detections we still observe a significant rise at rest-frame 1–3 μ m, while individual MIRI detections stand out at much higher fluxes. This check strongly suggests that the differ-

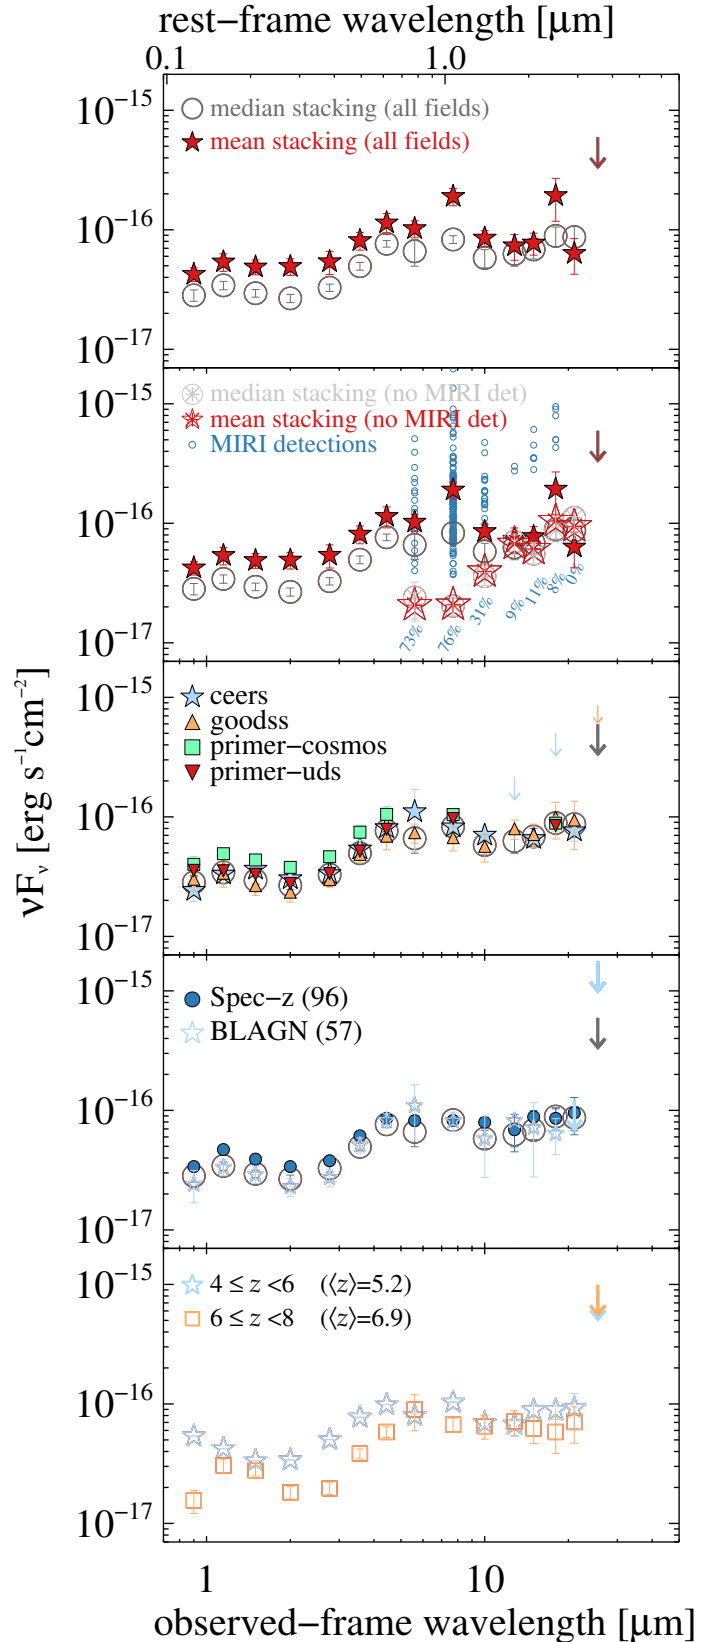


Fig. A.1. Summary plot comparing our median stacked fluxes (grey open circles) against multiple effects. From top to bottom: mean vs median stacking (Appendix A.1); contribution of MIRI detections (Appendix A.2); field-to-field variations (Appendix A.3); effect of photometric redshifts (Appendix A.4); effect of a broad redshift range (Appendix A.5). Details are given in the corresponding sub-sections.

ence between median and mean stacks boils down to the number of individual detections, and that the rising near-IR slope is driven by the bulk of non-detections, which is sensitive to more common and intrinsically fainter sources.

A.3. Field-to-field variations

An important aspect that our analysis relies on is the combination of multiple NIRC*am* and MIRI survey fields. While this is necessary to improve the statistics and the stacked S/N, it might potentially introduce systematics related to different JWST coverage and depth, which could affect the properties of photometrically-selected LRDs. We thus verify that our stacking results are robust against field-to-field variations, by repeating the same analysis separately for each field. The third panel in Fig. A.1 displays the median stacked SED obtained in each field (represented by different filled symbols) overlaid to our full stacking results (grey open circles). For visual purposes, here we relax the detection threshold to $S/N > 2$ (hence upper limit fluxes are set to twice the noise level if $S/N < 2$). Despite their heterogeneous spectral coverage across MIRI bands, all single-field median SEDs are consistent with the combined median SED. Moreover, we verify that stacking the same exact 22 objects having simultaneously NIRC*am* and MIRI coverage in all bands (from the GOODS-S field), the resulting SED is fully consistent with that obtained via stacking the full LRD sample. These checks strengthen that similarly-selected LRDs across different fields share a common median SED, despite their and heterogeneous NIRC*am*/MIRI spectral coverage.

A.4. Effect of photometric redshifts

A relevant caveat in our analysis is the photometric selection of LRDs. Specifically, photometric redshift uncertainties can be notably high, potentially biasing the LRD selection. We iterate that for the spectroscopic subsample (96/302 objects) there is good agreement between photometric and spectroscopic redshifts (Sect. 2). However, we test this directly by stacking NIRC*am* and MIRI images solely for the spectroscopic subsample. The fourth panel in Fig. A.1 displays the corresponding median stacked SED (blue filled circles). Numbers are given in Table 1 for each field: when combining all fields there are between 36 and 53 objects in MIRI (except for f560w, f1280w and f2550w with ≤ 5 objects due to a smaller areal coverage) and 96 objects in NIRC*am*. Despite the S/N being smaller than that of the full sample, the two median stacked SEDs agree remarkably well with each other.

We also explore whether a cleaner selection of “bona fide” AGN based on broad emission lines yields similar trends. From the NIRS*pec* Merged Table of the DJA, we isolate 57/96 (~59%) objects with broad H_α or H_β emission lines, as inferred from the best-fit continuum fluxes, line intensity and equivalent widths. This fraction agrees with that reported in the literature (e.g. Greene et al. 2024; Hviding et al. 2025). The corresponding median stacked SED of BL AGN (empty starred symbols), displays no significant difference with respect to the full spec- z (96) subsample, albeit a lower S/N in all bands due to smaller statistics.

A.5. Effect of a broad redshift range

As shown in Fig. 1, our targets span a very broad redshift range, from $z \sim 2$ to $z \sim 11$, with over 90% of the sample being

within the range ($4 < z < 8$). Because we perform stacking in the observed frame, the contribution to each stacked point comes from very different parts of the rest-frame spectrum. As mentioned in Sect. 4.1, we already account for this by convolving the filter response with the underlying redshift distribution. However, this approach assumes that each galaxy contributes evenly to the stack, regardless of its redshift, whereas the flux dimming can vary significantly across the full range. We thus test the impact of the redshift range to the stacked fluxes by splitting our LRDs into two similarly-populated z -bins at $4 \leq z < 6$ (126 LRDs) and $6 \leq z < 8$ (121 LRDs). We show this comparison in the bottom panel of Fig. A.1. The lower- z bin shows slightly brighter fluxes and steeper MIRI slope compared to the higher- z bin, as expected given the smaller distance and the longer rest-frame wavelength probed (e.g. $3.5 \mu\text{m}$ instead of $2.5 \mu\text{m}$ at MIRI-21 μm), respectively. We also note that fluxes are in energy units, hence the apparent flattening of the $6 \leq z < 8$ NIR SED translates into a monothonic rise in flux density (F_ν) units, as seen in the full sample. Overall, this check suggests that LRDs selected at different redshifts have broadly-consistent stacked SEDs, with no evidence for a galaxy-like drop at rest $1-2 \mu\text{m}$. We further note that emission lines are unlikely to systematically boost the median stacks, as this would be the case if a large fraction of sources were clustered in a narrow redshift range (such that a rest-frame emission line would systematically fall into a given filter). Given the broad similarity of the two SEDs in two different z -bins, we believe this is unlikely the case.

Appendix B: Input SED-fitting parameters

Table B.1. Main input parameters for SED fitting with CIGALE (Section 4.1).

Module	Parameter	Symbol	Values
Star formation history [SFR $\propto t \exp(-t/\tau)$]	Stellar e -folding time	τ_{star} (10^6 yr)	150, 200, 300, 500, 700, 900*
	Stellar age	t_{star} (10^6 yr)	200, 300, 500, 700, 900*
	Age of the burst	t_{burst} (10^6 yr)	10, 30, 50*, 100
	SFR ratio after/before the burst	r_{SFR}	0.1, 0.5, 1*, 2, 5, 10, 30, 50
Single stellar population [(Bruzual & Charlot 2003)]	Initial mass function	–	Chabrier (2003)
	Metallicity	Z	0.0004, 0.008, 0.02*
Nebular emission	Ionization parameter	$\log U$	-2.0*, -1.0
	Gas metallicity	Z_{gas}	0.0004*, 0.008, 0.019
	Electron density	n_e	100, 1000*
Dust attenuation [mod. (Charlot & Fall 2000)]	V-band attenuation (ISM)	$A_{V,\text{ISM}}$	0.01*, 0.1, 0.3, 0.5, 0.7, 1.0, 1.5, 2, 2.5, 3, 3.5
	Fractional ISM attenuation	$A_{V,\text{ISM}} / (A_{V,\text{BC}} + A_{V,\text{ISM}})$	0.44*
	Power-law slope (ISM)	–	-0.7*
	Power-law slope (BC)	–	-1.3*
Galactic dust emission [(Dale et al. 2014)]	Mass fraction of PAH	q_{pah}	0.47, 2.5, 3.9*
	Minimum radiation field	U_{min}	1.0, 10, 20, 40*
	Slope in $dM_{\text{dust}} \propto U^{-\alpha} dU$	α	2.0*
	Illuminated fraction [Umin–Umax]	γ	0.02*
AGN disk+torus+polar dust: SKIRTOR [Stalevski et al. (2012; 2016)]	Edge-on optical depth at $9.7\mu\text{m}$	$\tau_{9.7}$	3*, 7
	Slope of radial dust density profile	–	1.0*
	Slope of polar dust density profile	–	1.0*
	Torus opening angle (from equator)	Θ	40*
	Ratio outer/inner torus radius	–	20*
	Viewing angle (from vertical axis)	i	50*, 70
	AGN fraction	f_{AGN}	0.0, 0.1, 0.3, 0.5, 0.7, 0.9, 0.99*
	Extinction law for polar dust	–	SMC
	$E(B - V)$ in the polar direction	–	0.00, 0.05*, 0.10
	Temperature of polar dust	–	100 K*
Emissivity of polar dust	–	1.6*	

Notes. (*) Asterisks mark the best-fit parameters obtained from the AGN run.

Appendix C: Output SED-fitting parameters

Table C.1. List of relevant best-fit output parameters obtained with CIGALE (Section 4.1).

Parameter	Value and $\pm 1\sigma$ error
Stellar mass (fit without AGN)	$(5.0 \pm 1.0) \times 10^9 M_{\odot}$
Stellar mass (fit with AGN)	$(2.5 \pm 0.3) \times 10^9 M_{\odot}$
Total dust luminosity (fit without AGN)	$(5.4 \pm 1.2) \times 10^{44}$ erg/s
Total dust luminosity (fit with AGN)	$(1.4 \pm 0.2) \times 10^{44}$ erg/s
SFR [100 Myr] (fit without AGN)	$(8.2 \pm 3.6) M_{\odot}/\text{yr}$
SFR [100 Myr] (fit with AGN)	$(0.6 \pm 0.1) M_{\odot}/\text{yr}$
SFR [10 Myr] (fit without AGN)	$(15.9 \pm 5.2) M_{\odot}/\text{yr}$
SFR [10 Myr] (fit with AGN)	$(0.5 \pm 0.2) M_{\odot}/\text{yr}$
AGN bolometric luminosity	$(2.5 \pm 0.4) \times 10^{44}$ erg/s
AGN $6\mu\text{m}$ luminosity	$(3.8 \pm 0.6) \times 10^{43}$ erg/s

Notes. Uncertainties are reported at $\pm 1\sigma$ and retrieved from the likelihood distribution function recorded in the fitting.