Astronomy & Astrophysics manuscript no. LRD_stack_v1
September 10, 2025

Letter to the Editor

AGN-heated dust revealed in "Little Red Dots"

I. Delvecchio[1,⋆], E. Daddi[2], B. Magnelli[2], D. Elbaz[2], M. Giavalisco[3], A. Traina[1], G. Lanzuisi[1], H.B. Akins[4], S. Belli[5], C.M. Casey[6], F. Gentile[2,1], C. Gruppioni[1], F. Pozzi[5], and G. Zamorani[1]

[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 Citè, 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

ABSTRACT

Little Red Dots (LRDs) are a puzzling population of extragalactic sources whose origin is highly debated. In this Letter, we stack NIRCam, MIRI and ALMA images of a large and homogeneously-selected sample of LRDs from multiple JWST Legacy fields. We find clear evidence for hot-dust emission, seen as a rising mid-infrared continuum up to rest-frame $\lambda_{\rm rest}\sim 3\mu$m, that is best explained by a standard dusty AGN structure. This scenario holds regardless of whether the Optical/Ultraviolet (UV) continuum is stellar or AGN-dominated. Either ways, we argue that AGN-heated dust must have a geometry that leaves at least part of the Optical/UV continuum and broad-line region emission unobscured, which explains the observed blue UV slopes and prominent Balmer features. Lack of detections in the deep stack of X-ray images suggests that Compton-thick ($N_{\rm H}$>3×10^{24} cm^{-2}) gas obscuration is common, and likely confined within the dust sublimation radius ($R_{\rm sub}$~0.1 pc). While this gas-dust displacement is in line with recent models (e.g. the "Black Hole Star"), how pre-enriched hot dust can form around a nearly-pristine black hole environment remains to be explained.

Key words. galaxies: nuclei – galaxies: active – galaxies: evolution

1. Introduction

The James Webb Space Telescope (JWST) has ushered extragalactic astronomy in a transformative era, offering unprecedented access to the high-redshift universe. One of the most puzzling discoveries is a population of sources dubbed as "Little Red Dots" (LRDs, e.g. Matthee et al. 2024, Greene et al. 2024, Labbe et al. 2025). LRDs are extremely compact (often unresolved at ≲100 pc scales) and show red optical colors combined with blue UV slopes, resulting in a characteristic "V-shaped" spectral energy distribution (SED), with a sharp turnover near the Balmer break at rest 3645 Å (e.g. Setton et al. 2024, Hainline et al. 2025). Remarkably, LRDs appear to be two orders of magnitude more abundant than luminous quasars at similar redshifts (Kokorev et al. 2024, Akins et al. 2024), making them accessible even in narrow-field JWST surveys.

However, an observational consensus of the physical nature of LRDs is still lacking. The red optical colors and broad emission lines point toward obscured Active Galactic Nuclei (AGN) activity (Inayoshi et al. 2020, Volonteri et al. 2021, Ji et al. 2025). In this case, the AGN continuum would be moderately reddened by dust (A_V~2–3 mag; Casey et al. 2024; Akins et al. 2024; Wang et al. 2025). Yet, 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. 2024; Gloudemans et al. 2025). To alleviate this tension, some theoretical models postulate LRDs as supermassive black holes (SMBHs) accreting at super-Eddington luminosities and enshrouded in a turbulent, dust-free, and highly dense atmosphere: the "BH-star" scenario. This explains the strong Balmer break and the red continuum through pure gas opacity, and the broad Balmer lines via resonant electron scattering (e.g. Naidu et al. 2025; Inayoshi & Maiolino 2025; Rusakov et al. 2025; de Graaff et al. 2025). Similar gas conditions are postulated in other models, with the presence of a metal-free halo surrounding a pre-collapse "supermassive star" (Nandal & Loeb 2025), or a late-stage "Quasi star" (Begelman & Dexter 2025). All these models share the absence of prior metal enrichment and the presence of a gas-dense, dust-free region at sub-pc scales. So far, this dust paucity is supported by the flat LRD SED observed 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, conveying the idea that MIR emission in LRDs is dominated by dust-poor stellar populations and not AGN dust (Williams et al. 2024; Pérez-González et al. 2024, Leung et al. 2024). In addition, the non-detection from deep ALMA stacking poses a stringent dust mass limit of $M_{\rm dust}<10^6$ $M_\odot$ at z~6 (Casey et al. 2025; see also Chen et al. 2025; Setton et al. 2025). Nevertheless, we caution that previous NIRCam/MIRI SED studies of LRDs might have been biased towards bright 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 sub-samples (e.g. H-dropouts; Williams et al. 2024), whose SED could intrinsically differ from that of "typical" LRDs. To explore this possibility, in this Letter we stack a large and homogeneously-selected sample of LRDs with broad NIRCam/MIRI coverage. Throughout, we assume a Chabrier (2003) IMF and a flat ΛCDM cosmology with H_0 = 70 km/s/Mpc, Ω_m = 0.30, and Ω_Λ = 0.70.

⋆ ivan.delvecchio@inaf.it

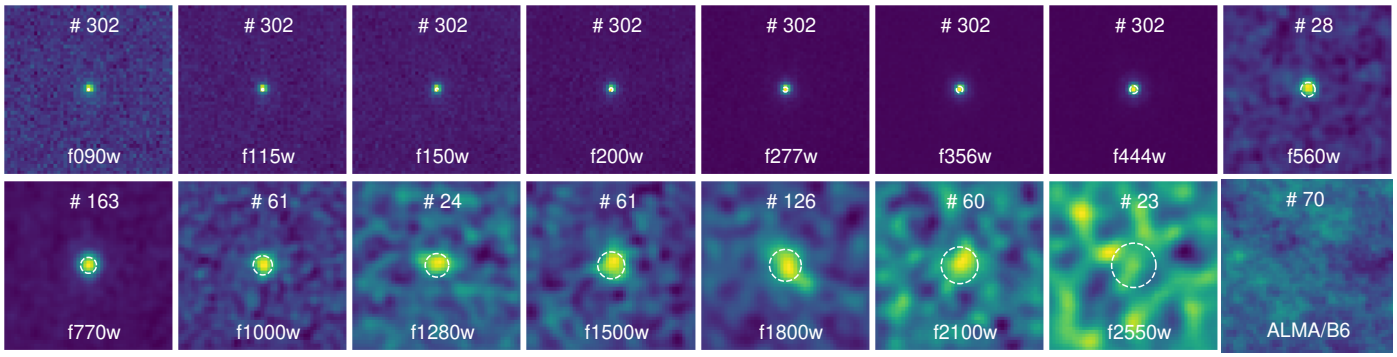

Fig. 1: Median stacked cutouts (3″×3″) of LRDs in NIRCam (7; from f090w to f444w), MIRI (8; from f560w to f2550w) and ALMA/B6 (1.1 mm) bands. All NIRCam and MIRI images have the same pixel size (0.06″). Each cutout reports the corresponding PSF's FWHM (white circle) and the number of stacked sources. To ease visualization, a smoothing with a Gaussian kernel is applied to all MIRI images, with a larger radius for larger FWHM: 2 pixels for f560w, f770w and f1000w; 3 pixels for f1280w and f1500w; 4 pixels for f1800w and f2100w and 5 pixels for f2550w. The stacked ALMA/B6 image is obtained in the *uv*-plane and imaged at 0.1″/px scale. Except for MIRI/f2550w and ALMA/B6, all other median stacks lead to a S/N>3 detection.

2. Photometric selection of LRDs

We select LRDs from the latest compilation of Kocevski et al. (2025), consisting of 341 photometrically-selected LRDs in the redshift range z∼2–11, and retrieved from several JWST/NIRCam surveys. These LRDs were homogeneously selected to display compact morphologies and "V-shaped" continuum from spectral slope fitting between blueward and redward of the rest-frame 3645 Å (i.e. close to the Balmer break). 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 replace photometric redshifts ($z_{\rm phot}$) with NIRSpec-based spectroscopic redshifts ($z_{\rm spec}$) taken from the public Dawn JWST Archive (DJA; Brammer 2023b; de Graaff et al. 2024; Heintz et al. 2024), whenever of high-quality (grade=3, see archive). This increases the spectroscopic fraction from 11% (39/341) to 29% (100/341). The mutual agreement between $z_{\rm phot}$ and $z_{\rm spec}$ estimates is generally good, with only five sources showing an offset $|z_{\rm spec} - z_{\rm phot}|/(1+z_{\rm spec})$>0.2.

Finally, we restrict ourselves to the NIRCam fields with (at least partial) MIRI coverage in two or more bands (including 18μm and/or 21μm; see Table A.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 GOODS-S, 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 (PRIMER[1]) survey in the COSMOS and UDS fields. Out of 302 LRDs, 163 (54%) are covered by MIRI 7.7μm, 126 (42%) also by MIRI-18μm, and 85 (28%) have a robust $z_{\rm spec}$. Moreover, we exploit archival ALMA/B6 (1.1 mm) data from the combined A3COSMOS+A3GOODSS surveys (Liu et al. 2019; Adscheid et al. 2024; Magnelli et al. 2024), which are available for a total of 70 LRDs in the COSMOS and GOODS-S fields. The number of sources split down by field, photometric coverage and spectroscopic redshifts is summarized in Table A.1.

[1] https://primer-jwst.github.io/

3. Stacking results

We perform median stacking over all available NIRCam bands (7) and all MIRI bands (2 in PRIMER, 7 in CEERS and 8 in SMILES), whenever available, at the optical position of all LRDs. We use a fixed pixel scale of 0.06″, and we extract fluxes from a circular aperture with fixed radius r=0.25″. ALMA *mean* stacking, instead, is performed in the *uv*-plane, in order to homogeneously combine archival data at different frequency, angular resolution and sensitivity. Further details of the stacking methodology are given in Appendix A. We refer the reader to Appendix B, to verify the robustness of our stacking results against possible biases. In Fig. 1 we show the stacked cutouts (3″×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 smooth them with a Gaussian kernel, adopting a larger radius for larger PSF FWHM (see caption of Fig. 1). The stacked ALMA/B6 image is obtained in the *uv*-plane and imaged at 0.1″/px scale. Except for MIRI/f2550w and ALMA/B6, all other median stacks lead to a detection with 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. 1). All fluxes, errors, rms of the stacked images are listed in Table A.2 along with the effective number of sources stacked in each band.

4. Discussion and Conclusions

4.1. SED-fitting with vs without AGN

In order to interpret the stacked photometry of LRDs, we perform SED-fitting using the Code Investigating GALaxy Emission (CIGALE; Boquien et al. 2019). We fix the input redshift to the median value of the LRD sample ($\langle z\rangle$∼6.2). In order to account for the broad redshift distribution of the sample (mostly at 4.5<z<8; see Fig. A.1) and the correspondingly wide rest-frame range sampled by each filter, we convolve the filter response curves at $\langle z\rangle$∼6.2 with the band-specific redshift distribution of the underlying LRD sample. Then, we run 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 is taken from Bruzual & Charlot (2003), including a delayed star formation history with random bursts. En-

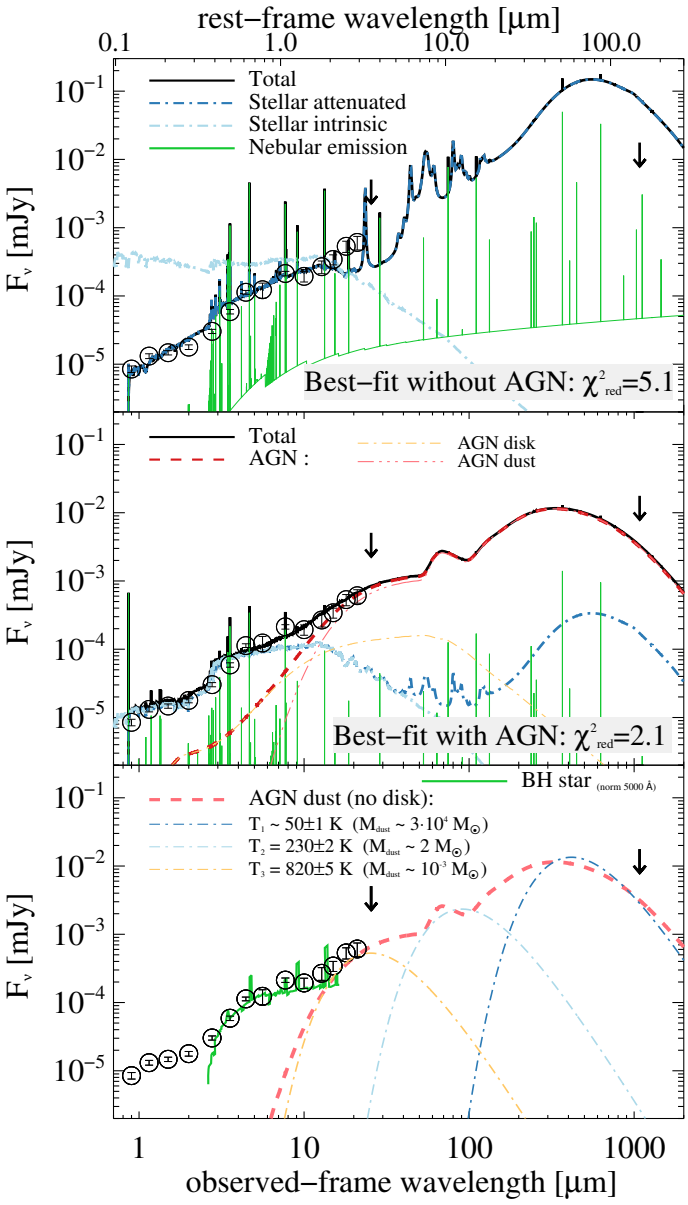

Fig. 2: Best SED fits obtained with CIGALE on the stacked photometry (open circles) at z=6.2. *Top panel:* Fit including only galaxy templates. *Middle panel:* fit including both galaxy and AGN templates. This fit is strongly favoured based on χ^2 arguments and ALMA constraints. *Bottom panel:* Decomposition of AGN dust as multi-temperature greybody models (dot-dashed lines). The normalized BH-star SED (Naidu et al. 2025) is shown for comparison (green). Details are given in Section 4.1.

ergy 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 Appendix C and D. In Fig. 2, 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) or with AGN (middle panel).

We caution the reader that our SED-fitting analysis is not primarily aimed at interpreting the rest optical or far-IR properties of LRDs. For instance, CIGALE assumes a single-zone obscuration between optical, UV and IR light, which might be oversimplistic in case of no co-spatiality (see Sect. 4.2). Nonetheless, regardless of the origin of Optical/UV emission in LRDs, we argue that AGN-heated dust is needed to reproduce the monothonic increase up to rest-frame 2–3μm, which pure dust-attenuated stellar light fails to match. This is further confirmed 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), which translates into $M_{dust}\lesssim10^6$ $M_\odot$ for T∼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 1–3μm. 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 favours the AGN solution at >95% significance.

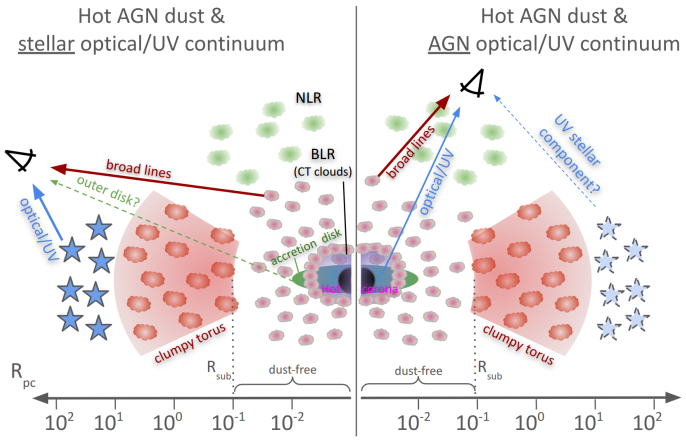

Fig. 3: Sketch of the proposed hot AGN-dust scenario, both for stellar-dominated (left) and AGN-dominated (right) Optical/UV continuum.

4.2. Are LRDs typically dust-obscured AGN?

By considering the best-fit solution with AGN, the rest-frame 6 μm luminosity of the AGN torus can be used as a proxy for the intrinsic AGN X-ray (rest [2–10] keV) luminosity ($L_{X,intr}$ (e.g. Stern 2015). From $L_{6\mu m}$=(3.8±0.6)×10^{43} erg/s, we derive $L_{X,intr}$=2×10^{43} erg/s, which nicely matches that expected from the best-fit AGN bolometric luminosity, if assuming standard X-ray bolometric corrections (e.g. Lusso et al. 2012). By stacking the deepest X-ray images of the 44 LRDs falling inside the *Chandra* Deep Field South (Luo et al. 2017), we obtain a 3σ [2–10] keV luminosity limit $L_X\lesssim10^{42}$ erg/s at z=6.2 (Comastri et al. subm). In order to match it with the $L_{X,intr}$ expected from the AGN SED, the hydrogen column density has to be $N_H\gtrsim3\times10^{24}$ cm^{-2}, implying Compton-thick gas obscuration (Ananna et al. 2024; Yue et al. 2024; Maiolino et al. 2025).

Given the SED-based estimate of the AGN bolometric luminosity, $L_{bol,AGN}$ = (2.5 ± 0.4) × 10^{44} erg s^{-1} and the best-fit stellar mass from the AGN solution ($M_\star$∼2.5×10^9 $M_\odot$; i.e. ≈2× 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_{BH}/M_\star\gtrsim$1:100, implying strongly overmassive BHs compared to local scaling relations (see also Pacucci et al. 2023; Juodžbalis et al. 2024; Maiolino et al. 2024). Moreover, our $L_{bol,AGN}$ estimate corresponds to a dust sublimation radius of $R_{sub}\approx$ 0.1 pc at $T_{sub}\simeq$ 1500 K (Barvainis 1987), meaning that any dusty AGN structure should begin at scales $\gtrsim$ 0.1 pc from the central SMBH. In the bottom panel of Fig. 2, we try to decompose the AGN-dust template with three multi-temperature

greybody curves (with T free to vary). Our rising MIRI SED up to rest-frame 2–3μm indicates hot dust at $T \simeq 820 \pm 5$ K, though it could be even hotter if optically-thick. Instead, the bulk dust mass ($M_{dust} \approx 3 \times 10^4$ $M_\odot$) comes from colder polar AGN dust at $\approx$50 K (though not constrained by our data). While a dusty structure could induce X-ray obscuration, it would also strongly suppress the Optical/UV continuum, in tension with the blue UV slope and strong Balmer features that imply relatively modest dust attenuation ($A_V \sim 2-3$; Akins et al. 2024; Casey et al. 2025).

This apparent contradiction can be resolved if X-ray absorption in LRDs occurs within R_{sub}, in a gas-dense and dust-free region. This gas-dust displacement is formally consistent with the recently-proposed models of Quasi-star (Begelman & Dexter 2025), BH–star (Naidu et al. 2025), and super-massive star (Nandal & Loeb 2025) scenarios. However, the presence of pre-enriched dust around a first-generation (or metal-poor) star is hard to justify. Moreover, the normalized BH–star SED (green line in Fig. 2, bottom panel) appears to flatten out at rest-frame $\lambda \gtrsim 1$ μm, contrary to the rising MIRI fluxes.

Our findings might fit better with an alternative hypothesis suggested by previous studies (Setton et al. 2024; Maiolino et al. 2025; D'Eugenio et al. 2025; Ji et al. 2025; Inayoshi & Maiolino 2025), in which X-ray obscuration arises from Compton-thick gas clouds within the BLR (at $R < R_{sub}$, i.e. dust-free), which covers the compact X-ray corona without suppressing the broad Balmer lines. Yet, the origin of the Optical/UV continuum remains open. Although we caution that LRDs might not consist of a single population, we speculate two simple scenarios in Fig. 3, depending on what powers the Optical/UV continuum in LRDs. In either cases, the bulk of Optical/UV continuum is not dust-attenuated, despite the presence of hot AGN dust, and the outer BLR is directly visible along the line of sight. (i) If the Optical/UV comes from stellar light (Fig. 3, left), the LRD resembles a "type-2" AGN configuration, in which the torus covers the AGN accretion disk, which would otherwise outshine the host. In case of high dust clumpiness, mildly-attenuated light from the outer accretion disk could add to the optical emission. (ii) Instead, if the Optical/UV comes from AGN light (Fig. 3, right), the LRD resembles a "type-1" unobscured AGN configuration, though the origin of the sharp Balmer break remains unclear. Fainter Optical/UV light from kpc-scale star formation cannot be ruled out (e.g. Rinaldi et al. 2025).

In summary, we find clear evidence for a rising MIRI continuum up to $\lambda_{rest} \sim 3\mu$m, which we interpret via SED fitting as hot (820±5 K) AGN dust, regardless of whether the Optical/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. We acknowledge that such homogeneous selection does not necessarily translate into a single LRD population with common *physical* properties. Our findings mainly provide en *empirical* benchmark for the presence of AGN-heated dust in the majority of LRDs, or (at least) in a significant sub-population of them. For instance, if LRDs are standard AGN, we speculate that the above selection criteria might be biased against the most dust-obscured ("type-2") LRD analogs, which could be recovered by lifting the blue-UV slope requirement.

Acknowledgements. 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. (Some of) the data producs 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 DNRF140.

References

Adscheid, S., Magnelli, B., Liu, D., et al. 2024, A&A, 685, A1
Akins, H. B., Casey, C. M., Lambrides, E., et al. 2024, arXiv e-prints, arXiv:2406.10341
Alberts, S., Lyu, J., Shivaei, 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, arXiv e-prints, arXiv:2412.01887
Barvainis, R. 1987, ApJ, 320, 537
Begelman, M. C. & Dexter, J. 2025, arXiv e-prints, arXiv:2507.09085
Boquien, M., Burgarella, D., Roehlly, Y., et al. 2019, A&A, 622, A103
Brammer, G. 2023a
Brammer, G. 2023b
Bruzual, G. & Charlot, S. 2003, MNRAS, 344, 1000
Carranza-Escudero, M., Conselice, C. J., Adams, N., et al. 2025, arXiv e-prints, arXiv:2506.04004
Casey, C. M., Akins, H. B., Finkelstein, S. L., et al. 2025, arXiv e-prints, arXiv:2505.18873
Casey, C. M., Akins, H. B., Kokorev, V., et al. 2024, ApJ, 975, L4
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, arXiv e-prints, arXiv:2505.22600
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, arXiv e-prints, arXiv:2503.16600
D'Eugenio, F., Maiolino, R., Perna, M., et al. 2025, arXiv e-prints, arXiv:2503.11752
Eisenstein, D. J., Johnson, B. D., Robertson, B., et al. 2023, arXiv e-prints, 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, arXiv e-prints, arXiv:2506.03256
Gloudemans, 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
Hainline, K. N., Maiolino, R., Juodžbalis, I., et al. 2025, ApJ, 979, 138
Harish, S., Kartaltepe, J. S., Liu, D., et al. 2025, arXiv e-prints, arXiv:2506.03306
Heintz, K. E., Watson, D., Brammer, G., et al. 2024, Science, 384, 890
Inayoshi, K. & Maiolino, R. 2025, ApJ, 980, L27
Inayoshi, K., Visbal, E., & Haiman, Z. 2020, ARA&A, 58, 27
Ji, X., Maiolino, R., Übler, H., et al. 2025, arXiv e-prints, arXiv:2501.13082
Juodžbalis, I., Maiolino, R., Baker, W. M., et al. 2024, Nature, 636, 594
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
Labbe, I., Greene, J. E., Bezanson, R., et al. 2025, ApJ, 978, 92
Labbe, I., Greene, J. E., Matthee, J., et al. 2024, arXiv e-prints, arXiv:2412.04557
Leung, G. C. K., Finkelstein, S. L., Pérez-González, P. G., et al. 2024, arXiv e-prints, arXiv:2411.12005
Liu, D., Schinnerer, E., Groves, B., et al. 2019, ApJ, 887, 235
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
Magnelli, B., Adscheid, S., Wang, T.-M., et al. 2024, A&A, 688, A55
Maiolino, R., Risaliti, G., Signorini, M., et al. 2025, MNRAS, 538, 1921
Maiolino, R., Scholtz, J., Curtis-Lake, E., et al. 2024, A&A, 691, A145
Matthee, J., Naidu, R. P., Brammer, G., et al. 2024, ApJ, 963, 129
Naidu, R. P., Matthee, J., Katz, H., et al. 2025, arXiv e-prints, arXiv:2503.16596
Nandal, D. & Loeb, A. 2025, arXiv e-prints, arXiv:2507.12618
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
Rieke, G. H., Alberts, S., Shivaei, I., et al. 2024, ApJ, 975, 83
Rinaldi, P., Rieke, G. H., Wu, Z., et al. 2025, arXiv e-prints, arXiv:2507.17738
Rusakov, V., Watson, D., Nikopoulos, G. P., et al. 2025, arXiv e-prints, arXiv:2503.16595
Setton, D. J., Greene, J. E., de Graaff, A., et al. 2024, arXiv e-prints, arXiv:2411.03424
Setton, D. J., Greene, J. E., Spilker, J. S., et al. 2025, arXiv e-prints, arXiv:2503.02059
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
Valentino, F., Brammer, G., Gould, K. M. L., et al. 2023, ApJ, 947, 20
Volonteri, M., Habouzit, M., & Colpi, M. 2021, Nature Reviews Physics, 3, 732
Wang, B., de Graaff, A., Davies, R. L., et al. 2025, ApJ, 984, 121
Wang, T.-M., Magnelli, B., Schinnerer, E., et al. 2024, A&A, 681, A110
Wang, T.-M., Magnelli, B., Schinnerer, E., et al. 2022, A&A, 660, A142
Williams, C. C., Alberts, S., Ji, Z., et al. 2024, ApJ, 968, 34
Xiao, M., Oesch, P. A., Bing, L., et al. 2025, arXiv e-prints, arXiv:2503.01945
Yue, M., Eilers, A.-C., Ananna, T. T., et al. 2024, ApJ, 974, L26

Appendix A: Stacked samples and fluxes

Here we provide a more detailed outline of the input images and stacking technique used in this work.

Image and noise maps are retrieved from the corresponding mosaics of each field (see below). In this step, we discard narrow or medium-band filters (whenever available) to mitigate potential flux boosting in the case of strong emission lines.

We retrieve 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 retrieve NIRCam and MIRI images from the JADES and SMILES surveys publicly available on the MAST[2] archive. While SMILES maps are already at 0.06″/px, JADES maps (with a native pixel scale of 0.03″) are re-scaled to 0.06″/px for consistency with the other fields.

For CEERS and PRIMER-UDS, which are unavailable on MAST, we leverage the DJA (Brammer 2023a; Valentino et al. 2023), where the full reduced mosaics are publicly available in each NIRCam and MIRI band[3], all with a pixel scale of 0.04″ and in units of [10 nJy]. We do a resampling of all mosaics to 0.06″/px for consistency with SMILES and PRIMER-COSMOS maps, while preserving astrometry and total flux over a given aperture. For PRIMER-UDS, we also merge the East and West mosaics (listed separately in the DJA) into a full mosaic.

We only consider 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 create a $N\times N$ pixel image (with N=51, 0.06″/px) from the corresponding map, each centred on the NIRCam position of the LRD. We then stack all $N_{\rm stack}$ sources and retrieve the median flux at each pixel.

The uncertainty on the stacked flux density is quantified through a bootstrapping technique: being $N_{\rm stack}$ the number of input targets in a given band, in each random realisation we reshuffle the input sample, preserving the same $N_{\rm stack}$ by allowing source duplication. We repeat this 100 times, and take 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 is computed by stacking the inverse-variance weighted noise map of each target. Our median stacks yield a detection in each band up to 21 μm with signal-to-noise[4] S/N>3 (see Table A.2).

For ALMA, we perform a stacking analysis in the *uv*-domain using all publicly available Band-6 observations covering our LRDs in the COSMOS and GOODS-S fields. 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.

Fig. A.1 displays the redshift distribution of our final LRD sample (302 sources, grey dashed), split by field (coloured open histograms). Filled histograms highlight the spectroscopic subsets. Corresponding numbers are listed in Table A.1.

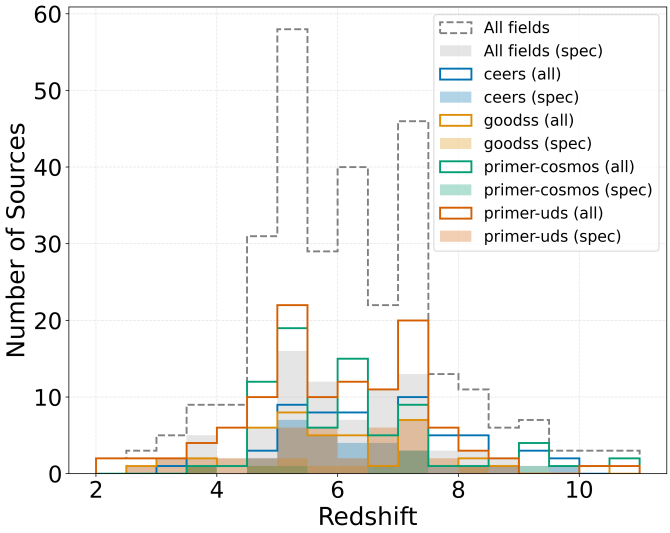

Fig. A.1: Redshift distribution of the final LRD sample (black dashed histogram) used in this work. The full sample is split by field (CEERS, GOODSS, PRIMER-COSMOS, PRIMER-UDS), as highlighted by the coloured open histograms, while filled histograms mark the corresponding subset with spectroscopic redshifts.

[2] SMILES: https://archive.stsci.edu/hlsp/smiles; JADES: https://archive.stsci.edu/hlsp/jades

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

[4] 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 A.1: This table lists the survey fields considered in this study. For each survey, the sky area, the available filters, and the corresponding number of LRDs (also with a spectroscopic redshift) are reported. The range of $N_{\rm stack}$ reflects the varying area across different filters of the same instrument. The global sample of LRDs reaches up to 302 for NIRCam, 163 for MIRI, and 70 for ALMA/B6 coverage.

Survey/Field	Instrument	Filters	Area (arcmin^2)	$N_{\rm stack}$ range (spec)
CEERS/EGS	NIRCam	[f090w, f115w, f150w, f200w, f277w, f356w, f444w]	88.6	63 (30)
	MIRI	[f560w, f770w, f1000w, f1280w, f1500w, f1800w, f2100w]	55.0	2–39 (2–19)
JADES/GOODSS	NIRCam	[f090w, f115w, f150w, f200w, f277w, f356w, f444w]	62.1	44 (7)
SMILES/GOODSS	MIRI	[f560w, f770w, f1000w, f1280w, f1500w, f1800w, f2100w, f2550w]	34.0	22–23 (2)
A3-GOODSS	ALMA	Band-6 (1080μm)	86.8	34 (4)
PRIMER/COSMOS	NIRCam	[f090w, f115w, f150w, f200w, f277w, f356w, f444w]	138.9	79–80 (7)
	MIRI	[f770w, f1800w]	73.0	41–42 (3)
A3-COSMOS	ALMA	Band-6 (1080μm)	127.9	36 (2)
PRIMER/UDS	NIRCam	[f090w, f115w, f150w, f200w, f277w, f356w, f444w]	243.0	116 (41)
	MIRI	[f770w, f1800w]	130.0	61–62 (23)

Notes. Field acronyms: EGS = Extended Groth Strip; GOODS-S = Great Observatories Origins Deep Survey–South; UDS = UKIDSS Ultra Deep Survey; COSMOS = Cosmic Evolution Survey.

Table A.2: Median stacked fluxes (open grey circles in Fig. B.1). For each wavelength, the table reports the number of stacked sources ($N_{\rm stack}$), the rms of the stacked image, the median stacked flux and its 1σ uncertainty from bootstrapping, respectively. In case of non-detection, the 3σ upper limit is given.

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

Appendix B: 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. B.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. B.1 from top to bottom.

B.1. Mean vs median stacking

In the first panel of Fig. B.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 (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 NIRCam 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). A similar effect is seen at MIRI 7.7μm and 18μm, which include about 100 more targets than at other MIRI bands (see Table A.1), thanks to a complete coverage across all the fields studied here. Even a handful of bright detections at 7.7μm and 18μm (some have S/N>100 in the PRIMER fields) are able to boost the mean MIRI flux upwards. Moreover, we opt for median stacking also because it strongly mitigates contamination from bright neighbors and catastrophic outliers, and thus reduces the confusion noise for the faint sources. This comparison, however, demonstrates that the mean stacked LRD SED reveals a similar trend as the median stacked SED, showing a rising MIR slope that is indicative of dusty AGN activity.

B.2. Field-to-field variations

An important aspect that our analysis relies on is the combination of multiple NIRCam 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 second panel in Fig. B.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 NIRCam/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 in different fields and with heterogeneous spectral coverage share a common median SED.

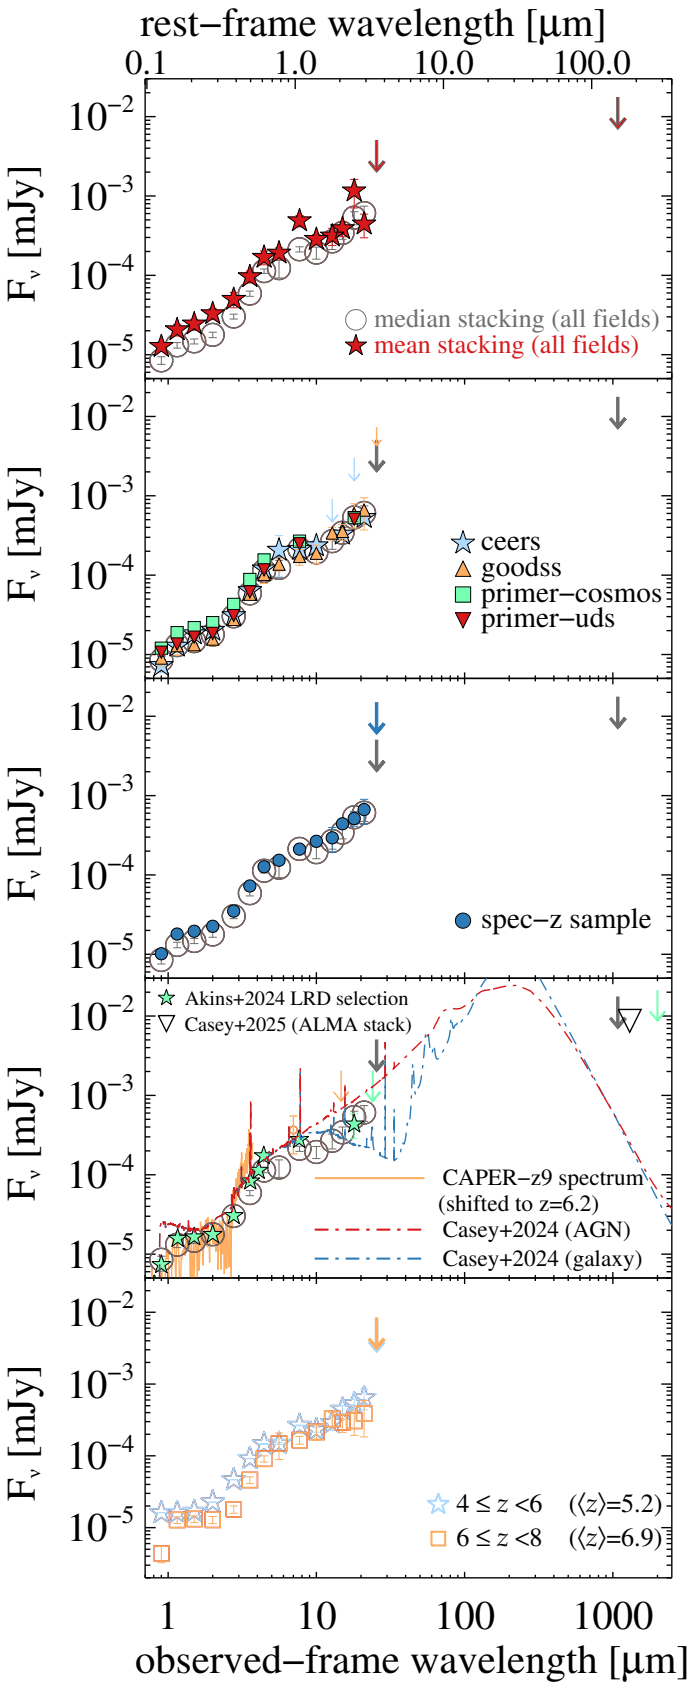

Fig. B.1: Summary plot comparing our median stacked fluxes (grey open circles) against multiple effects. From top to bottom: mean vs median stacking (Appendix B.1); field-to-field variations (Appendix B.2); effect of photometric redshifts (Appendix B.3); effect of different LRD selections (Appendix B.4);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

B.3. 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 (85/302 objects) there is good agreement between photometric and spectroscopic redshifts (Sect. 2). However, we test this directly by stacking NIRCam and MIRI images solely for the spectroscopic subsample. The third panel in Fig. B.1 displays the corresponding median stacked SED (blue filled circles). Numbers are given in Table A.1 for each field: when combining all fields there are between 21 and 47 objects in MIRI (except for f560w, f1280w and f2550w with ≤5 objects due to a smaller areal coverage) and 85 objects in NIRCam. Despite the S/N being smaller than that of the full sample, the two median stacked SEDs agree remarkably well with each other.

B.4. Effect of LRD selection

In the literature, LRDs are inevitably selected with distinct criteria, depending on the availability of multi-band NIRCam 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. The fourth panel of Fig. B.1 shows for comparison: the median stacked SED from Akins et al. (2024; green symbols); the corresponding galaxy and AGN best-fit LRD SED from Casey et al. (2024; dot-dashed lines); the ALMA upper limit from Casey et al. (2025; downward triangle) and the observed CAPER-z9 NIRSpec spectrum (at z=9.288; Taylor et al. 2025) shifted to z=6.2 (orange symbols).

Akins et al. (2024) selected 434 LRDs in the COSMOS-Web area (0.5 deg^2) that simultaneously meet a compactness (f444w) cut, a very red colour cut (f444w/f277w) and a bad brown drarf fit, without explicitly requiring a "V-shaped" SED due to limitation in NIRCam imaging. Their median SED agrees remarkably well with ours, though their statistics in MIRI is limited to 7.7μm (over 1/3 of COSMOS-Web) and 18μm (from the inner PRIMER-COSMOS footprint, ≈4% of the full sample), this latter leading to tentative detection (0.43±0.15 mJy). We show that their 18μm flux is fully consistent with the rising MIR slope seen longward 10μm. Their upper limits are here re-scaled to 3σ (*Herschel* and SCUBA limits are too shallow to enter the y-scale). Despite the different LRD selection and survey field (≲25% of our LRDs are in Akins et al. 2024), this agreement further suggests that the compactness and the conservative colour selection by Akins et al. (2024) broadly compensate the need for a V-shaped SED, leading to self-consistent LRD samples.

In the same fourth panel of Fig. B.1, we compare our stacks against the median LRD SED by Casey et al. (2024), that is inferred from a joint LRD sample (675 LRDs) from Akins et al. (2024) and Kokorev et al. (2024). This latter consists of a compactness (in f444w) criterion and various colour cuts that mimic a V-shape requirement (see their Sect.3.1 in Kokorev et al. 2024 for details). In their work, Casey et al. (2024) performed stacking in rest-frame wavelength bins, following 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 (red) or a pure galaxy template (blue). We emphasize that our median stacks are fully consistent with their median AGN SED within their uncertainties (here not reported). If any, the AGN SED from Casey et al. (2024) stands slightly higher than our MIRI points, that partly reflects the MIRI upper limits in their study (~20 objects per bin at $\lambda_{\rm rest}$>0.8μm). Our stacks and their SED are also mutually consistent with the ALMA upper limit derived in Casey et al. (2025) by stacking 60 LRDs (among the Akins et al. 2024 sample) at λ=1.3 mm.

In addition, we compare our stacks against the NIRSpec spectrum of CAPER-z9 at z=9.288 (Taylor et al. 2025), which is complemented by MIRI data at 10μm (detection) and 21μm (upper limit). This continuous spectrum enables a direct comparison with our SED shape over a broad wavelength range. After moving the CAPER-z9 spectrum and photometry to z=6.2 (orange symbols), we again find a reassuring consistency with our data.

Another meaningful comparison with MIRI studies of LRDs is against the sample of Williams et al. (2024), who stacked LRDs in all MIRI bands within the GOODS-SMILES footprint. They found a flat IR-SED up to 18μm (and upper limits longward), thus excluding a dominant contribution from AGN-heated dust. Besides the small numbers (only 9 objects) and the slightly higher redshift compared to our sample ($\langle z\rangle$~7.7 instead of $\langle z\rangle$~6.2), their LRDs were pre-selected from a pool of H-band dropouts (i.e. HST-dark galaxies), hence likely biased towards the most obscured sources. We repeat the stacking for the same 9 sources, finding 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 9 LRDs are not representative of our photometrically-selected LRDs (from Kocevski et al. 2025), albeit they are certainly an interesting sub-sample to follow-up.

B.5. Effect of a broad redshift range

As shown in Fig. A.1, our targets span a very broad redshift range, from z~2 to z~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≤z<6 (126 LRDs) and 6≤z<8 (121 LRDs). We show this comparison in the bottom panel of Fig. B.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 μm instead of 2.5 μm at MIRI-21 μm), respectively. 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 μm. 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 C: Input SED-fitting parameters

In the following Table, we outline the main input parameters used for SED-fitting with CIGALE (Section 4.1).

Table C.1: Model parameters for SED fitting with CIGALE.

Module	Parameter	Symbol	Values
Star formation history [SFR $\propto$ t exp($-$t/τ)]	Stellar e-folding time	$\tau_{\rm star}$ (10^6 yr)	150, 200, 300, 500, 700, 900*
	Stellar age	$t_{\rm star}$ (10^6 yr)	200, 300, 500, 700, 900*
	Age of the burst	$t_{\rm burst}$ (10^6 yr)	10, 30, 50*, 100
	SFR ratio after/before the burst	$r_{\rm SFR}$	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	$Z_{\rm gas}$	0.0004*, 0.008, 0.019
	Electron density	n_e	100,1000*
Dust attenuation [mod. (Charlot & Fall 2000)]	V-band attenuation (ISM)	$A_{\rm V,ISM}$	0.01*, 0.1, 0.3, 0.5, 0.7, 1.0, 1.5, 2, 2.5, 3, 3.5
	Fractional ISM attenuation	$A_{\rm V,ISM}$ / ($A_{\rm V,BC}$+$A_{\rm V,ISM}$)	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	$q_{\rm pah}$	0.47, 2.5, 3.9*
	Minimum radiation field	$U_{\rm min}$	1.0, 10, 20, 40*
	Slope in $dM_{\rm dust} \propto U^{-\alpha} dU$	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 9.7μm	$\tau_{9.7}$	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	$f_{\rm AGN}$	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 from the AGN run.

Appendix D: Output SED-fitting parameters

.

Table D.1: We report some relevant best-fit output parameters obtained with CIGALE (Section 4.1). Uncertainties are reported at $\pm1\sigma$ and retrieved from the likelihood distribution function recorded in the fitting.

Parameter	Value and $\pm1\sigma$ error
Stellar mass (fit without AGN)	$(5.0\pm1.0)\times10^9$ $M_\odot$
Stellar mass (fit with AGN)	$(2.5\pm0.3)\times10^9$ $M_\odot$
Total dust luminosity (fit without AGN)	$(5.4\pm1.2)\times10^{44}$ erg/s
Total dust luminosity (fit with AGN)	$(1.4\pm0.2)\times10^{44}$ erg/s
SFR [100 Myr] (fit without AGN)	(8.2±3.6) $M_\odot$/yr
SFR [100 Myr] (fit with AGN)	(0.6±0.1) $M_\odot$/yr
SFR [10 Myr] (fit without AGN)	(15.9±5.2) $M_\odot$/yr
SFR [10 Myr] (fit with AGN)	(0.5±0.2) $M_\odot$/yr
AGN bolometric luminosity	$(2.5\pm0.4)\times10^{44}$ erg/s
AGN 6 μm luminosity	$(3.8\pm0.6)\times10^{43}$ erg/s