

Unsupervised selection and characterisation of Little Red Dots in JWST surveys with manifold learning

Michele Ginolfi^{1,2}, Filippo Mannucci², Alessandro Marconi^{1,2}, Francesco D'Eugenio^{3,4}, Giacomo Venturi²,
Francesco Belfiore^{2,5}, Giovanni Cresci², Caterina Bracci^{1,2,5}, Stefano Carniani⁵, Alessandra Cozzi¹,
Roberto Maiolino^{3,4}, and Guido Risaliti^{1,2}

¹ Dipartimento di Fisica e Astronomia, Università di Firenze, Via G. Sansone 1, I-50019, Sesto F.no (Firenze), Italy

² INAF — Osservatorio Astrofisico di Arcetri, Largo E. Fermi 5, I-50125, Florence, Italy

³ Kavli Institute for Cosmology, University of Cambridge, Madingley Road, Cambridge, CB3 0HA, UK

⁴ Cavendish Laboratory - Astrophysics Group, University of Cambridge, 19 JJ Thomson Avenue, Cambridge, CB3 0HE, UK

⁵ European Southern Observatory, Karl-Schwarzschild-Str. 2, 85748, Garching bei München, Germany

⁶ Scuola Normale Superiore, Piazza dei Cavalieri 7, I-56126 Pisa, Italy

Received XX; accepted YY

ABSTRACT

Little Red Dots (LRDs) are compact, red sources discovered at high redshift by JWST whose physical nature and selection function remain debated. We investigate whether an unsupervised machine-learning approach applied to multi-band photometry can identify LRD-like objects, and other populations, without relying on predefined colour cuts. Using UMAP, a manifold-learning (dimensionality-reduction) method, we place $\sim 242,000$ isolated, well-measured sources from the ASTRODEEP-JWST catalogue on a two-dimensional map, where objects with similar broadband colours, morphology, and photometric redshift lie close together. We then use spectroscopically confirmed LRDs to identify where LRD-like objects lie within this map, compare the resulting areas with published colour cuts, and validate our data-driven selection against archival NIRSpec spectra from the DJA. We find that the spectroscopically selected LRDs concentrate in two well-defined regions with no colour cut imposed, tracing populations that differ mainly in redshift, a difference imprinted in their broadband colours. The main region reaches a purity of ≈ 0.78 at ≈ 0.82 completeness on the spectroscopically classified subset, competitive with, or cleaner than, literature colour cuts, and yields ~ 100 additional candidates. We also test the method as a general tool for population discovery: the manifold recovers the locations of brown dwarfs and broad-line AGN with no explicit criterion, and isolates rare pathological outliers. Overall, unsupervised manifolds, anchored by sparse high-confidence spectroscopic labels, provide an efficient, assumption-light framework for characterising populations, comparing selection methods on a common basis, and discovering rare objects in large photometric datasets.

Key words. galaxies: active – galaxies: high-redshift – galaxies: photometry – methods: data analysis – techniques: photometric – techniques: spectroscopic

1. Introduction

Among the early surprises of the James Webb Space Telescope (JWST) there has been the discovery of a large population of compact sources with red colours, now widely referred to as Little Red Dots (LRDs; Labbé et al. 2023; Matthee et al. 2024; Kokorev et al. 2024). Their spectral energy distributions (SEDs) are characteristically “V-shaped”, with a blue rest-frame ultra-violet and a steeply rising rest-frame optical continuum. A large fraction of LRDs show broad Balmer emission lines indicative of active galactic nuclei (AGN, Greene et al. 2024; Kocovski et al. 2025; Matthee et al. 2024). Despite rapid observational progress, their physical nature remains debated.

A clearer picture is now emerging, based on physically motivated scenarios. In one, many LRDs are accreting black holes embedded in dense gas, where electron scattering (among other effects) broadens the Balmer lines, and the continuum is either due to thermal emission at ~ 5000 K or nebular emission (with some dust reddening; Matthee et al. 2026; Rusakov et al. 2026; de Graaff et al. 2025; Sneppen et al. 2026); in another, the “super-Eddington unification” model of Madau & Maiolino (2026), LRDs are the dust-obscured (type-2-like) counterparts

of blue “Little Blue Dots”, the broad lines remaining visible because their obscuring column is modest (e.g., $A_V \approx 2\text{--}5$ mag instead of $20\text{--}40$ mag as in type-2 AGN). Alternative interpretations have also been proposed, including dust-reddened AGN (Greene et al. 2024; Kocovski et al. 2025) and compact dusty stellar systems (Labbé et al. 2023); however, the growing spectroscopic diversity still favours a heterogeneous, possibly composite population (Zhang et al. 2026; Pérez-González et al. 2026; Pan et al. 2026). Establishing a consistent observational definition of LRDs, followed by a clean, well-understood census and the identification of possible LRD sub-classes, is therefore an essential prerequisite for interpreting the population.

In practice, photometric LRD samples are typically assembled by combining colour–colour cuts, a compactness requirement, and a step designed to reject cool stars (brown dwarfs), whose molecular absorption mimics the LRD colours (Kokorev et al. 2024; Greene et al. 2024; Barro et al. 2024; Rinaldi et al. 2026; Hainline et al. 2026). The specific thresholds, however, vary substantially from study to study: the adopted rest-optical “redness” threshold alone moves the recovered numbers by factors of several, and different choices preferentially select different parts of the underlying popula-

tion. As a consequence, published LRD samples differ appreciably, and the inferred properties depend on the selection function nearly as much as on the data. Spectroscopic selections (e.g. [de Graaff et al. 2025](#)) avoid broadband-colour priors at the classification stage, but are limited to sources with high-quality spectroscopy and inherit the selection function of the parent spectroscopic sample. This is not a minor caveat: archival spectroscopic samples tend to over-represent LRDs, because many observing programmes deliberately targeted colour-selected candidates, so a photometric prior re-enters through the target selection.

A complementary strategy, which we pursue here, is to characterise the data *before* imposing thresholds, using unsupervised manifold learning (for a recent review of unsupervised machine-learning methods in an astronomical context, see [Fotopoulou 2024](#)). The aim of this paradigm is to map complex, high-dimensional measurements into a low-dimensional representation (a “manifold”) that captures the intrinsic diversity of the data, without human-induced priors or restrictive empirical boundaries. In astronomy, these methods are increasingly used to organise and classify large datasets without predefined labels. For example, they have been used to compress and explore galaxy spectra and SEDs ([Reis et al. 2021](#); [Portillo et al. 2020](#)) and to build data-driven taxonomies of galaxy morphology ([Hocking et al. 2018](#)); in a semi-supervised setting close to the one we adopt here, they have also been used to classify AGN in DESI spectra beyond traditional emission-line diagnostics ([Alcolea et al. 2026](#)). This is precisely because structure in the data can emerge naturally rather than being imposed. Several dimensionality-reduction algorithms exist, including *t*-SNE ([van der Maaten & Hinton 2008](#)) and a range of (variational) autoencoders. We adopt the Uniform Manifold Approximation and Projection (UMAP; [McInnes et al. 2018](#)) method for its practical advantages: it scales well to large catalogues with relatively low-dimensional feature vectors, tends to preserve local *and* a degree of global structure, is comparatively stable to its hyperparameters, and produces an embedding, the low-dimensional map of the sources (the “manifold” introduced above, realised here in two dimensions), well suited to both visualisation and downstream quantitative analysis. We regard the specific choice of algorithm as exploratory and return to it in our assessment of robustness.

Once the sources are placed on this unsupervised map, we adopt a semi-supervised strategy: we *anchor* the representation with a small number of high-confidence external labels (here the spectroscopically selected LRDs of [de Graaff et al. 2025](#), classified from spectral shape rather than broadband colours, although the parent spectroscopic targeting is not entirely colour-independent; see Section 5) and propagate that information to the unlabelled sources nearby in the representation. This has three useful consequences. First, it provides a selection that is driven by where confirmed sources actually live in the data rather than by chosen thresholds. Second, it makes the boundaries of the population, and the contaminants adjacent to it, explicit and measurable. Third, the same representation naturally supports the identification of outliers and the discovery of rare populations, simply as regions of the manifold that are sparsely populated or unlike the bulk of the population.

Such a data-driven perspective is timely: the recent literature spans selections that impose increasingly restrictive criteria (e.g. [Park et al. 2026](#); [Lin et al. 2026](#)), at the risk of cutting away any redshift evolution of the population, and selections that admit X-ray- and mid-infrared (MIR)-detected sources from very large parent samples (e.g. [Casey et al. 2026](#); [Fu et al. 2026](#)), at the risk of including interlopers; a manifold-based representation instead

shows directly whether such candidates fall within the LRD region or elsewhere, which we exploit in Section 4.

We note that, framed this way, our approach has a concrete advantage over a fully supervised search for LRD-like objects in the present data regime. Supervised classifiers require large, representative training sets, which for a rare and still ill-defined population do not yet exist; here a few dozen high-confidence labels suffice, because the unsupervised representation has already organised the data and the labels are needed only to indicate *where* in that organisation the population lies. The same property lets us recover sources, and even entire sub-populations, collectively (as coherent regions of the manifold) rather than one object at a time.

In this work we apply the approach to public JWST imaging surveys through the homogeneously reduced ASTRODEEP-JWST photometric catalogue ([Merlin et al. 2024](#)). We construct a feature space from eight-band optical-to-MIR photometry, two morphological indicators and the photometric redshift for ~242,000 well-measured, isolated sources across six extragalactic fields, embed it with UMAP, and anchor the embedding with the spectroscopically selected LRDs of [de Graaff et al. \(2025\)](#). We deliberately restrict the inputs to broadband photometry, morphology and photometric redshift as a proof of concept; the same framework extends naturally to spectra and resolved imaging, which we consider in future work (see Section 5). We carefully state the assumptions and limitations of each step to allow for critical assessment of the method and its results.

The paper is organised as follows. Section 2 describes the parent sample, the features, the pre-processing, and the external catalogues. Section 3 presents the embedding, the definition of the LRD region, and the spectroscopic measurements. Section 4 reports the structure of the manifold, the comparison with literature selections, the candidate sample, and the archival-spectroscopy results. Section 5 discusses the implications of our findings, including which of the input features actually carry the LRD signature and the potential of the manifold to reveal other astrophysical populations and outliers. Finally, Section 6 summarises our conclusions.

2. Data and parent sample

2.1. The ASTRODEEP-JWST catalogue

We build on the ASTRODEEP-JWST photometric catalogue ([Merlin et al. 2024](#)), a homogeneous reduction and multi-band catalogue of the principal public JWST extragalactic fields. It provides forced, PSF-matched aperture photometry on a common detection image, total-flux corrections, SExtractor morphological parameters, and photometric redshifts computed with EAZY ([Brammer et al. 2008](#)), processed identically across fields so that colours and derived quantities are directly comparable. The catalogue spans seven fields: the lensing cluster Abell 2744 (GLASS-JWST and UNCOVER; [Treu et al. 2022](#); [Bezanson et al. 2024](#)), CEERS in the EGS ([Bagley et al. 2023](#); [Finkelstein et al. 2023](#)), the two JADES fields in GOODS-North and GOODS-South ([Rieke et al. 2023](#); [Eisenstein et al. 2025](#); [D’Eugenio et al. 2025](#); [Eisenstein et al. 2026](#)), PRIMER-COSMOS and PRIMER-UDS ([Dunlop et al. 2021](#)), and NGDEEP ([Bagley et al. 2024](#)). We use the six blank/wide fields and exclude NGDEEP, because it lacks coverage in the F410M medium band that we require for a uniform feature set (Section 2.2); retaining a common set of bands across all fields is necessary so that position in the feature space reflects the SED shape rather than which filters happen to be available.

Table 1. The eleven features spanning the manifold, all built from the ASTRODEEP-JWST catalogue (Merlin et al. 2024). Seven bands enter as logarithmic flux ratios to the F356W reference band; the F356W flux itself is retained as an overall-brightness feature. Morphology comes from the two SExtractor quantities measured on the F356W+F444W detection stack, and the redshift is the ASTRODEEP EAZY photometric redshift. All fluxes are PSF-matched aperture photometry; features are robust-scaled before embedding (Section 2.3)

Feature	Band	Feature	Observable
<i>Photometry</i>		<i>Morphology</i>	
$\log(f_{814}/f_{356})$	F814W	ClassStarSE	stellarity
$\log(f_{115}/f_{356})$	F115W	$\log r_{50}$	half-light radius
$\log(f_{150}/f_{356})$	F150W	<i>Redshift</i>	
$\log(f_{200}/f_{356})$	F200W	$\log(1+z_{\text{phot}})$	photo-z
$\log(f_{277}/f_{356})$	F277W		
$\log(f_{410}/f_{356})$	F410M		
$\log(f_{444}/f_{356})$	F444W		
$\log f_{356}$	F356W (ref.)		

2.2. Features: photometry, morphology and redshift

For each source we use eight bands spanning the optical to MIR: the HST/ACS F814W band (central wavelength $\lambda_c \approx 0.81 \mu\text{m}$) and the JWST/NIRCam bands F115W, F150W, F200W, F277W, F356W and F444W (central wavelengths 1.15, 1.50, 1.99, 2.76, 3.57 and $4.40 \mu\text{m}$, respectively; broad filters with relative widths $\Delta\lambda/\lambda \sim 0.2\text{--}0.3$), together with the F410M medium band ($\lambda_c \approx 4.08 \mu\text{m}$, $\Delta\lambda/\lambda \sim 0.1$), which helps isolate strong rest-frame optical emission lines from continuum at the relevant redshifts. This set brackets the rest-frame ultraviolet-to-optical break that defines the LRD SED across $z \sim 3\text{--}8$.

Because compactness is a defining property of LRDs, we deliberately include morphology in the representation through two SExtractor quantities, both measured on the ASTRODEEP detection image, a stack of the F356W and F444W mosaics (Merlin et al. 2024), so that compactness is assessed in the rest-frame optical where the LRD morphology is defined: the stellarity index ClassStarSE, which runs from 0 (resolved) to 1 (point-like), and the half-light radius $r_{50\text{SE}}$ (the radius enclosing 50% of the flux).

The final feature is the photometric redshift. This carries an important caveat: it assumes the photometric redshifts are reliable, whereas JWST photometric redshifts are known to include a fraction of catastrophic outliers (e.g. Arrabal Haro et al. 2023). Such failures displace the affected sources on the map; indeed, one of the anomalous regions we recover (Section 5.4) consists precisely of catastrophic photo-z failures.

Photometry enters as colours: fluxes are normalised by the F356W band and log-scaled, so that the representation responds to SED *shape* rather than apparent brightness; the reference flux in the F356W band itself is retained as a separate feature so that brightness information is not discarded.

These eleven features are summarised in Table 1.

2.3. Pre-processing and the parent sample

We require a signal-to-noise ratio $S/N > 2$ in each band, with one deliberate exception: the bluest band, HST/ACS F814W, is allowed to be undetected. A source that fails the threshold in F814W is retained, with its F814W flux replaced by a 1σ estimate from its measurement error, whereas detection is enforced in the seven JWST/NIRCam bands (see Section 2.4 for why this is required and its effect on completeness). This exception keeps

the sample sensitive to Lyman-break drop-outs: at $z \gtrsim 5$ the Lyman break redshifts into or beyond F814W, leaving genuine high-redshift sources faint or undetected there. Since F814W then probes the rest-frame ultraviolet, its faintness reflects the break, not the rest-optical “redness” that defines LRDs; imposing a formal F814W detection would therefore preferentially discard exactly the high-redshift LRDs of interest here. In contrast, a detection is still required in F115W, so genuine drop-outs in that band ($z \gtrsim 8$, where the Lyman break falls redward of F115W) are excluded. This sets an effective redshift ceiling of $z \lesssim 8$: sources near this limit survive only because the break lies *within* F115W, suppressing but not extinguishing its flux, so that the permissive $S/N > 2$ threshold still admits them. The high-redshift population (median $z_{\text{phot}} \approx 8.2$) we recover and discuss in Section 5.2, sits precisely at this boundary; objects at higher redshift, for which F115W vanishes entirely, fall outside the present sample.

We note that the $S/N > 2$ threshold is itself deliberately permissive. Because our analysis is driven by the overall *shape* of the SED, encoded in a learned representation, rather than by hard cuts on individual colours, it is comparatively robust to the per-band noise that scatters sources across a fixed threshold, and can therefore operate at lower signal-to-noise than is typical for cut-based selections. We verified that lowering the requirement further mainly admits noise-dominated photometry that carries little SED information, without altering any of our conclusions, while raising it shrinks the sample without sharpening the results; $S/N > 2$ is thus a pragmatic compromise between sample size and SED reliability.

We additionally apply an isolation criterion, removing any source with a catalogued neighbour within $0.5''$. This step guards against blended or neighbour-contaminated photometry, which would otherwise scatter sources in the feature space for reasons unrelated to their intrinsic SED. The isolation cut has a clear cost, in that it removes genuine sources, in particular in crowded regions and close pairs or mergers, and so introduces incompleteness that we do not attempt to correct; in return, the retained sources have cleaner, more trustworthy SEDs, while the sample remains large enough for our purposes. One concern deserves particular attention: because many LRDs host close star-forming companions (Baggen et al. 2026), whose small separations are exactly what the isolation criterion targets, the cut could in principle remove them preferentially and bias our completeness. We tested this directly by cross-matching the Baggen et al. (2026) sample, split into objects with candidate companions and solitary ones, into our sequence of selection steps. Within our footprint, that is the sky covered by our six fields, we recover 81% of the companion-hosting LRDs and 78% of the solitary ones, a difference consistent with noise, and the residual $\sim 20\%$ loss is dominated by the signal-to-noise pre-processing and is common to both groups. The reason lies in the catalogue construction rather than in the imaging resolution: ASTRODEEP detects sources on a single common detection image and then measures forced, PSF-matched photometry (Section 2.1), so a close companion that is not split off as a separate entry at the detection stage is never counted as a neighbour by our $0.5''$ cut, even though the native NIRCam imaging resolves finer separations. The isolation criterion therefore does not bias our completeness against companion-hosting LRDs.

After these steps the features are scaled with a robust (median and interquartile) scaler to limit the influence of outliers. The resulting parent sample contains 242,327 sources.

2.4. Completeness through the selection

A central point for interpreting our results is that the parent sample is itself the product of a selection function, and we quantify its effect on the spectroscopic LRDs explicitly. Of the 116 unique spectroscopically selected LRDs of [de Graaff et al. \(2025\)](#), 89 (77%) fall within the footprint and detection limits of our parent catalogue. Of these, 68 (59% of the original 116) additionally satisfy the pre-processing requirements above and enter the final sample. The losses are not random in redshift: the sources removed by pre-processing have a higher median redshift ($z \simeq 5.6$) than those retained ($z \simeq 5.1$). Waiving the detection requirement in F814W, i.e., allowing it to be undetected (Section 2.3), already recovers the most common dropout cases, but a detection is still required in the bluer NIRCcam bands, so the very reddest and highest-redshift sources (faint or absent in F115W and F150W) continue to be preferentially removed; this residual imprint of the detection requirement on the high-redshift tail is the dominant bias of the parent selection.

We therefore report completeness at two complementary levels. The *intrinsic* completeness measures how well a given selection recovers the spectroscopic LRDs that are present in our sample (relative to the 68), and isolates the performance of the selection itself. The *end-to-end* completeness folds in the $\sim 59\%$ pre-processing retention and describes the fraction of all known spectroscopic LRDs recovered by the full procedure. The distinction matters: the pre-processing, not the manifold selection, is the dominant source of incompleteness, and it is the part most readily relaxed in future work. The natural route is to require detections only redward of the Lyman break rather than in every band, of which the F814W exemption adopted here is a first step. Extending the same logic to the NIRCcam bands, for instance by anchoring the redshift on medium-band photometry such as F335M where available (as in much of the JADES footprint), would let the method follow the receding redshift frontier and, in principle, admit F115W- and F150W-drop-out LRD candidates that the current detection requirement removes. Doing so properly is non-trivial: our features are log-flux ratios (colours), which are undefined for an undetected band, so admitting non-detections requires reconstructing the feature space to handle upper limits and censored fluxes in a principled way, together with a field-dependent choice of which bands lie redward of the break. We therefore leave it to a dedicated follow-up.

2.5. External catalogues and cross-matching

We cross-match the parent sample against several external catalogues, using a $0.4''$ positional tolerance throughout. The spectroscopically selected LRD sample of [de Graaff et al. \(2025\)](#) (116 sources; selected from broken power-law fits to NIR-Spec/PRISM continua and a compactness requirement, independent of broadband colour at the classification stage) provides the labelled *anchor* of our analysis (68 in our sample). For comparison we use two photometrically selected catalogues cross-matched into our sample, [Barro et al. \(2026a\)](#) (94 sources; colour-colour selected and spectroscopically confirmed) and [Kokorev et al. \(2024\)](#) (187; colour-colour selected), together with the published colour criteria of [Rinaldi et al. \(2026\)](#) and [Akins et al. \(2025\)](#), which we re-apply to the same parent sample (Section 3.3). We also match the brown-dwarf compilation of [Hainline et al. \(2024\)](#) and [Hainline et al. \(2026\)](#) (12 in our sample) and a sample of broad-line AGN ([Baccus & Xu 2025](#)). Finally, we record spectroscopic coverage from the DAWN JWST Archive (DJA; [Valentino et al. 2023](#); [Heintz et al. 2024](#); Sec-

tion 4.7): 14,982 sources in our sample have a public NIR-Spec spectrum, and 8,774 have a high-quality PRISM spectrum (grade= 3, which denotes the most reliable spectra and secure redshifts; [Heintz et al. 2024](#)). The latter subset is important because it constitutes the input over which the spectroscopic anchor selection was performed essentially exhaustively, which we exploit to measure purity directly in Section 4.3.

3. Methods

3.1. The embedding

Each source is represented by a 11-dimensional feature vector: seven F356W-normalised log-colours, the log reference flux, the stellarity, the log half-light radius, and $\log(1 + z_{\text{phot}})$ (Section 2.2; Table 1). After robust-scaling the features (Section 2.3), we reduce the representation to two dimensions with the manifold-learning algorithm UMAP ([McInnes et al. 2018](#)), using a Manhattan metric, 15 nearest neighbours and a minimum separation of zero. We explored the principal hyper-parameters (the number of neighbours, the minimum distance, the distance metric) and several random seeds over a grid; the qualitative structure relevant here (the concentration of the spectroscopic anchor, the presence of two associated regions, and the separation of brown dwarfs) is recovered in all cases, and the quantitative results reported below do not change appreciably. We use the two-dimensional embedding for discovery and visualisation: a low-dimensional view is what makes the population structure inspectable at a glance and supports rapid, interactive exploration of the manifold for discovery, which we develop into a public interactive tool in Section 5.4. We verify the key conclusions directly in the original feature space (Section 4.2), where distances are not distorted by the projection. We adopt two dimensions for legibility and ease of presentation in print; we verified that a three-dimensional embedding recovers the same locus structure and leaves our conclusions unchanged, so the choice of two dimensions is not what drives our results.

The resulting manifold is not featureless: it carries clear physical structure (Figure 1). Photometric redshift, reference-band flux and stellarity all vary smoothly across the embedding, and the photometric and spectroscopic LRDs occupy specific, well-defined regions rather than being scattered at random.

3.2. Defining the LRD region

We locate the LRD region directly from the positions of the spectroscopically confirmed LRDs (the labelled anchor). We first group the anchor sources themselves by applying a density-based clustering (DBSCAN) to their standardised embedding coordinates. This robustly returns two compact groups (hereafter the main and secondary regions, containing 56 and 11 anchor sources; for brevity we also call them the main and secondary *islands*, and use the terms islands, regions and loci interchangeably in what follows), together with a single isolated source. With only one object we cannot define an anchor region around it as we do for the two loci, so for the present analysis we treat it as an outlier and set it aside; we return to it in Section 5.2, where its neighbourhood suggests a possible high-redshift extension of the main locus. The two groups are clearly separated on the manifold (Figure 2a).

We then describe each group with a two-dimensional Gaussian in the UMAP plane and define an enclosing region as a Mahalanobis ellipse, i.e., an iso-probability contour of that Gaussian (a distance measured in units of the local spread along each axis).




Fig. 1. The two-dimensional embedding colour-coded by (a) photometric redshift, (b) reference-band (F356W) flux, (c) stellarity, and (d) spectroscopic coverage with the locations of the spectroscopic LRDs. Smooth gradients in (a)–(c) show that the embedding encodes physical structure; panel (d) shows where the spectroscopic follow-up and the labelled anchor sit on the manifold.

For the main region the size of the ellipse is set by the fraction f of anchor sources it encloses: scaling the Mahalanobis radius to the f -th percentile of the member distances yields a region containing a fraction f of the anchors. The completeness with respect to the anchor is therefore a *design parameter* that we set and vary (we consider f between 0.5 and 1.0, and adopt $f = 1.0$, which encloses all 56 main-region anchors, as the fiducial value); this is meaningful precisely because the anchors are tightly clustered. The secondary region contains too few sources for a percentile to be meaningful, and is elongated, so a single Mahalanobis radius would be set by the most extreme source in any direction and would over-extend the long axis. We therefore enclose it with the minimum-volume (Löwner–John) ellipse, which jointly optimises the centre, shape and orientation for the smallest area containing the members. One of the eleven secondary anchors is a clear outlier (much brighter than the rest and offset from their narrow locus; Section 4.3), which we exclude when

defining the fiducial secondary region, quantifying below the effect of re-including it. The zoomed panels of Figure 2 show that both regions are populated by compact, high-redshift sources, as expected for LRDs.

3.3. A comparison with literature selections

Our aim here is to compare the data-driven region against the established photometric LRD selections in a controlled, like-for-like way. Published LRD samples are drawn from different catalogues, fields and depths, so their reported numbers cannot be compared directly with ours, or with one another. We therefore take the colour criteria that define each method and re-apply them to our *own* parent sample, so that every selection acts on exactly the same objects and is judged against the same spectroscopic anchor. This is what allows us to place each method, to-




Fig. 2. (a) The two anchor regions on the embedding. The spectroscopic anchors are shown on top (de Graaff et al. 2025 main and secondary loci); the photometrically selected samples of Barro et al. (2026a) (orange squares) and Kokorev et al. (2024) (purple triangles) are overlaid beneath them, illustrating that the latter are more dispersed across the manifold, and the black cross marks the lone anchor discarded by the clustering. To keep the map uncluttered, the enclosing ellipses are drawn only in the two insets, one per locus (main, lower right; secondary, upper right): the main-region ellipse and the fiducial (ten-anchor) secondary-region ellipse are solid, while the full secondary region obtained by re-including the single bright outlier (crossed; Section 4.3) is dotted. (b–e) Zooms on each region, with the field colour-coded by photometric redshift and by stellarity (each panel with its own colourbar), showing that both loci are populated by compact, high-redshift sources. In panel (b) the main-region ellipse is drawn at several values of f , the fraction of main-island anchors it encloses (the completeness knob swept in Section 4.3; see Section 3.2). In the secondary zooms (d,e) the bright outlier anchor is marked with a black cross and the dotted curve is the full region.

gether with our region, as a point in the common completeness–purity plane of Section 4.3.

Each of these selections combines colour–colour cuts with a compactness requirement and a brown-dwarf rejection. To isolate the axis that actually differs between studies (the colour cuts), we keep the other two ingredients fixed and identical for all methods. For compactness, since our catalogue does not provide the small-aperture flux ratios used in some of these works, we adopt a single empirically calibrated proxy (a stellarity $\text{ClassStarSE} > 0.8$, which retains 99% of the spectroscopic anchor). For brown-dwarf rejection we use $\text{F115W} - \text{F200W} > -0.5$, following Greene et al. (2024), which cleanly separates the known brown dwarfs in our sample (Section 4.5). Any per-study magnitude or signal-to-noise floors (for instance the $\text{F444W} < 27$ limit of Barro et al. 2026b) are likewise superseded by the common detection requirement of our parent sample (Section 2.3), so that every method acts on identical objects. With these common ingredients we implement the following colour selections:

- Barro et al. (2026b): $\text{F200W} - \text{F444W} > 1$ and $\text{F200W} - \text{F444W} > (\text{F115W} - \text{F200W}) + 0.25$, with $\text{F115W} - \text{F200W} > -0.5$;
- Kokorev et al. (2024): a two-branch cut, either $\text{F115W} - \text{F150W} < 0.8$, $\text{F200W} - \text{F277W} > 0.7$ and $\text{F200W} - \text{F356W} > 1.0$, or $\text{F150W} - \text{F200W} < 0.8$, $\text{F277W} - \text{F356W} > 0.6$ and $\text{F277W} - \text{F444W} > 0.7$;
- Rinaldi et al. (2026): $\text{F150W} - \text{F200W} < 1.0$ and $\text{F277W} - \text{F444W} > 0.5$;
- Akins et al. (2025): $\text{F277W} - \text{F444W} > 1.5$.

3.4. Spectroscopic measurements from archival NIRSpec data

We retrieve public NIRSpec spectra from the DAWN JWST Archive¹ (DJA), which provides uniformly reduced one-dimensional extractions and an emission-line catalogue (Heintz et al. 2024). Our manifold selection is purely photometric, so we use archival spectroscopy both to test it independently and to exploit it for discovery. Concretely, the spectra serve two purposes: to confirm that sources falling in our regions show the spectroscopic signatures expected of LRDs and AGN, rather than merely their broadband colours, and to search for confirmed LRDs or AGN among in-region sources that are *not* part of the labelled anchor; that anchor was defined only over PRISM-classified spectra, leaving the sources with grating-only coverage as an untapped reservoir. We apply two diagnostics, each calibrated on the anchor as a consistency check, so that the same measurement applied to the anchor recovers it; both are deployed and their outcomes reported in Section 4.7.

First, for sources with a high-quality PRISM spectrum we measure the continuum shape: we shift to the rest frame, convert to f_{λ} , and fit power-law slopes in a rest-UV (1400–2600 Å) and a rest-optical (4100–6700 Å) window. We require the characteristic V-shape ($\beta_{\text{opt}} > 0$, $\beta_{\text{UV}} < -0.2$, $\beta_{\text{opt}} - \beta_{\text{UV}} > 0.5$; de Graaff et al. 2025) to hold in at least 95% of Monte-Carlo realisations of the flux uncertainties, so that low signal-to-noise spectra are flagged as inconclusive rather than spuriously selected. Second, for sources with a medium-resolution grating spectrum covering H α we fit a linear continuum plus narrow

¹ https://s3.amazonaws.com/msaexp-nirspec/extractions/nirspec_public_v4.4.html

$H\alpha+[N\text{ II}]$ and an optional broad $H\alpha$ component, and record a broad line when it is statistically preferred (with the broad-line fit lowering the Bayesian information criterion by $\Delta\text{BIC} > 10$), significant in integrated flux ($S/N > 3$), and broader than 1000 km s^{-1} but narrower than 7000 km s^{-1} (the upper bound preventing an over-broad component from absorbing the continuum and low-level noise into a spurious line). We restrict the broad-line measurement to the signal-to-noise regime in which it reliably recovers the anchor (Section 4.7).

4. Results

4.1. Spectroscopic LRDs concentrate on the manifold

Before turning to the LRDs, we recall that the manifold is physically structured rather than arbitrary: photometric redshift, reference-band flux and stellarity all vary smoothly across it (Figure 1a–c), so the position of a source in the embedding already carries physically meaningful information. In this context, the first and most fundamental result is that the spectroscopically selected LRDs are strongly localised on a manifold that was constructed without any knowledge of their labels (Figure 1d). We quantify the localisation in two complementary, threshold-free ways. Considering, for each anchor source, its 50 nearest neighbours in UMAP space among all sources, the fraction that are themselves anchors exceeds the value expected if the anchors were randomly distributed by a factor of ~ 570 ; a label-permutation test rejects the null hypothesis of random placement at very high significance. Equivalently, a kernel-density core enclosing 90% of the anchors contains fewer than 10^3 of the $\sim 2.4 \times 10^5$ sources, i.e., an over-density of ~ 230 . That such a compact, over-dense region emerges from an unsupervised representation indicates that LRDs occupy a genuinely distinct part of the data space, and motivates using their location to define a selection.

The concentration is a property of the population, not of any single field: the manifold combines six independent JWST fields, and leaving one or two of them out leaves the spectroscopic LRDs clustered in the same region. The absolute completeness and purity shift slightly, because the set of anchors matched into the sample changes, but the overall result (that LRDs occupy a compact, over-dense locus) is unchanged. As noted in Section 3.2, the anchors form two separated groups rather than one (Figure 2); we treat the larger as the main island and return to the smaller in Section 4.6, where we show that the two differ mainly in redshift while remaining distinct in their broadband colours (Section 5.3).

4.2. The structure is intrinsic to the feature space

Because distances in a UMAP projection are not strictly metric, we verify that the concentration is a property of the data and not of the two-dimensional embedding. Repeating the nearest-neighbour over-density measurement directly in the original ~ 11 -dimensional feature space (the input features themselves: flux ratios, reference flux, morphology and photometric redshift; see Table 1) yields a comparable enhancement (~ 520), confirming that the anchors are close to one another in the full representation. As a complementary supervised validation, we trained a small neural-network classifier directly in the original 11-dimensional feature space, using five-fold cross-validation. This test is not used to define the LRD region, but only to ask whether the feature-space location of the anchors generalises to objects that were held out during training. The out-of-fold

predictions recover $\sim 90\%$ of held-out anchor sources at a fixed threshold, and the high-probability sources overlap substantially with the two-dimensional region. This confirms that the localisation is present in the original feature space rather than being introduced by the UMAP projection. The details of this test and the diagnostic probability distributions are given in Appendix A.

The same feature space also tells the two anchor loci apart from each other: they sit somewhat closer to one another than to the field, yet a classifier separates their members at a cross-validated area under the curve (AUC) of 0.996, dropping only to 0.991 when redshift is removed. The two loci are therefore genuinely distinct in broadband colours (largely the imprint of their different redshifts, which the colours encode) and only marginally closer without the explicit redshift feature; we return to what distinguishes them in Section 5.3.

4.3. Completeness and purity relative to the literature

We now compare the data-driven selection with the literature colour cuts in the completeness–purity plane (Figure 3, Table 2), all evaluated on the common parent sample and judged against the same spectroscopic anchor. Completeness is the fraction of anchor LRDs recovered. Purity is the fraction of confirmed LRDs among the selected sources with a high-quality (grade-3) PRISM spectrum, that is, the subset already classified in the literature from PRISM as LRD or not. A source counts as a confirmed LRD if de Graaff et al. (2025) (or Barro et al. 2026a) spectroscopically classified it as one, that is, from the rest-frame continuum V-shape (a blue ultraviolet slope and a red optical slope, the $\beta_{\text{UV}}/\beta_{\text{opt}}$ criterion of Section 3.4) rather than from broadband colours. Because de Graaff et al. (2025) applied this classification essentially exhaustively to that subset, a selected, PRISM-classified source that is not a confirmed LRD is a genuine interloper rather than merely unobserved; purity here is therefore a measured quantity, not a lower limit.

For the data-driven selection we trace a curve by varying the enclosed anchor fraction; each literature method is a single point, shown both as its published, positionally matched catalogue (“catalogue”) and as its colour cuts re-applied to our common sample (“criteria”). We omit the spectroscopic Barro et al. (2026a) catalogue point, which is $\approx 100\%$ pure by construction because its members are confirmed LRDs; the same would hold for any spectroscopically confirmed sample and so carries no information about the photometric selection.

At its fiducial size the main-region selection recovers 82% of the in-sample anchors at a purity of ≈ 0.78 measured over the PRISM-classified subset (the selection contains 282 sources in total). We stress that this is not the purity of the whole photometric sample, which we cannot measure directly since most members lack spectroscopy; it is the confirmed-LRD fraction among the selected sources that carry a high-quality PRISM classification. The comparison between the different photometric selections shows the expected trade-off between purity and completeness. The broadest cut of Rinaldi et al. (2026) recovers ~ 0.76 of the anchors at ~ 0.35 purity from ~ 2000 sources, while the reddest cut of Akins et al. (2025) is purer (~ 0.66) but recovers only ~ 0.40 . The most effective colour selections are the two-branch cut of Kokorev et al. (2024) and the V-shape cut of Barro et al. (2026b), which reach high completeness (~ 0.93 and ~ 0.90) at moderate purity (~ 0.66 and ~ 0.59). To capture both axes at once, the dotted red diagonals in Figure 3 mark lines of constant *quality* $Q \equiv (\text{completeness} + \text{purity})/2$, a selection lying on a higher diagonal being better in this combined sense. On this measure the data-driven main locus reaches the

highest value, $Q \approx 0.80$, marginally above the best colour selection (Kokorev et al. 2024, $Q \approx 0.79$) and well above the others ($Q \approx 0.53$ – 0.75). More importantly, at a given completeness the data-driven region is the purest of all the selections, and it is a tunable curve rather than a single operating point, obtained without hand-designing any colour boundary.

However, we emphasise two caveats: the region is defined using the anchor and then evaluated against it, which favours the data-driven selection, and our purity is measured only over the PRISM-classified subset. At the same time, the colour cuts are not a fundamentally different exercise but a lower-information version of the same one, using a few filters and redshift bins rather than the full SED, morphology and redshift jointly; the gap in Figure 3 largely reflects that difference in information content.

Adding the secondary region raises completeness towards unity, but its photometric definition hinges on a single discrepant anchor. In fact, ten of the eleven secondary anchors trace a narrow, filamentary locus; the eleventh, on the other hand, is a clear outlier, and a well-studied source in its own right (GN-28074; Juodžbalis et al. 2024; Loiacono et al. 2025; Brazzini et al. 2026). It lies off this narrow axis and is about an order of magnitude brighter than the others (a factor of ≈ 14 in the F356W normalisation flux and ≈ 11 in observed F444W, i.e. ≈ 2.6 – 2.8 mag), and at a somewhat lower redshift ($z \approx 2.3$ versus a median of ≈ 3.4). We therefore adopt the ten-anchor filament as the fiducial secondary region, enclosed by its minimum-volume ellipse: a thin locus (axis ratio ~ 17) of 110 sources at $\sim 41\%$ purity. Re-including the bright outlier (chosen objectively as the anchor whose exclusion most reduces the ellipse area) stretches the ellipse off-axis and inflates it to 698 sources at only $\sim 7\%$ purity. Both versions of the region are shown in Figure 2; in the completeness–purity plane (Figure 3) we plot only the fiducial, ten-anchor filament as the “both loci” track. We adopt the filament as fiducial because the eleventh anchor is atypical of the population (Section 4.6); its photometric boundary nonetheless remains only weakly constrained by the handful of available anchors, and a larger spectroscopic sample will be needed to delineate it firmly.

The quoted purity is an average over the selection, but it depends on the observed properties of the sources, so sub-samples can be cleaner still. We find essentially no dependence on redshift: within the main locus the purity stays at ≈ 0.78 from $z \approx 4$ to $z \approx 5$ and, if anything, declines slightly beyond $z \approx 6$. It does, however, increase with luminosity. The F444W band (the reddest broad filter, which samples the rest-frame optical at these redshifts, i.e. the red side of the LRD V-shape, and in which LRDs are brightest and best measured) is the natural luminosity proxy: restricting the main locus to $F444W < 26$ raises the purity to ≈ 0.85 – 0.90 at a modest cost in completeness (orange curve in Figure 3), because the fainter members carry noisier photometry and admit more contaminants. Compactness provides a further, independent purity lever, which we return to in Section 5.3.

4.4. A sample of new candidates

At the fiducial size the main region contains 282 sources. Of these, 164 are already identified as LRDs in at least one published catalogue, and a further 11 have a high-quality PRISM spectrum but were not selected by the spectroscopic anchor; that is, they were examined and not classified as LRDs, and we treat them as interlopers and remove them. They fall in the main region because their broadband colours, compactness and photometric redshifts resemble those of LRDs, even though their




Fig. 3. Completeness versus purity on the common parent sample, judged against the spectroscopic anchor. The data-driven selections are curves (the enclosing ellipse grown over a range of enclosed anchor fractions); the literature methods are points, as published catalogues (filled) and as their colour cuts re-applied to our sample (open). The lime track adds the fiducial (ten-anchor) secondary region to the main locus, and the orange track restricts the main locus to bright sources ($F444W < 26$), which are purer (Section 4.3). Dotted red diagonals are lines of constant quality $Q = (\text{completeness} + \text{purity})/2$; selections towards the upper right are better in this combined sense.

Table 2. Literature selections on the common sample, judged against the spectroscopic anchor (common compactness proxy and brown-dwarf rejection; purity over the PRISM-classified subset).

Selection	Compl.	Purity	N
Barro et al. (2026b) criteria	0.90	0.59	792
Rinaldi et al. (2026) criteria	0.76	0.35	2027
Kokorev et al. (2024) criteria	0.93	0.66	604
Akins et al. (2025) criteria	0.40	0.66	139
Data-driven (main, $f=1.0$)	0.82	0.78	282

PRISM spectra lack the defining features (the V-shaped continuum and/or broad Balmer lines); thus, they are likely photometric analogues rather than genuine LRDs, and it is exactly these spectroscopically identified contaminants that set the measured purity.

The remaining 107 sources constitute a new candidate sample: they sit within the LRD region of the manifold but are absent from existing photometric LRD catalogues. We rank them by locus centrality, defined objectively as the Mahalanobis distance to the main-region centre, and provide the full list as a machine-readable table (see Appendix B). The most locus-central candidates are shown individually, as RGB cutouts and observer-frame SEDs, in Appendix B.

Two checks establish that this sample is not an artefact. First, it is not a by-product of the cross-match tolerance: the 107 candidates lie a median of $48''$ (minimum $3.0''$) from the nearest catalogued LRD, so none would be re-classified as “known” even if the matching radius were widened to $2''$. Second, although we do not impose a stringent signal-to-noise cut, the candidates are nonetheless well detected: all 107 have F444W signal-to-noise above 8 (median ≈ 56), because the all-bands requirement of Sec-

tion 2.3 already sets an effective floor. The candidate sample is also stable: resampling the anchor sources (bootstrap) leaves the candidate set largely unchanged, with a mean Jaccard overlap of 0.94 (the size of the intersection divided by the size of the union of the two candidate sets) and 103 of 107 candidates recovered in at least 80% of resamples.

The diagnostic breakdown of why the colour cuts miss these sources is informative (Figure 4). The candidates are compact (only 3% fail the compactness proxy) and largely consistent with the more inclusive colour selections: only 6% fall outside the Barro et al. (2026b) V-shaped region, and 28% fail the $F277W-F444W > 0.5$ threshold of Rinaldi et al. (2026). What excludes them is the strictest redness cut: 81% fail the $F277W-F444W > 1.5$ threshold of Akins et al. (2025). In other words, the candidates are compact and at LRD-like redshifts (see Section 4.6), but with somewhat bluer optical colours and slightly fainter $F444W$ magnitudes than the reddest catalogued LRDs. This is consistent with the view that the strictest colour cuts sample only the reddest extreme of the LRD population, and shows concretely which sources are left out.

That many candidates satisfy the more inclusive colour cuts yet are absent from the published catalogues shows that their colours are not exceptional: the same colour selection yields a different sample when applied to a different reduction, depth and compactness measure (see Section 5). For Barro et al. (2026b), which spans all our fields, footprint plays no role and depth only a minor one (only a minority of candidates fall below its depth, $F444W > 27$); most candidates are colour-consistent, compact and bright, and are missed instead through differences in the compactness measure (Barro et al. 2026b use an aperture flux ratio, $F444W[0.5'']/F444W[0.2''] < 1.5$, that our stellarity proxy does not reproduce) and in the photometry, which we take from the independent ASTRODEEP reduction (Merlin et al. 2024).

Consistent with this, the candidates have the broadband SED shape expected for LRDs. Figure 5 stacks the rest-frame SEDs of the region members, separating those with a spectroscopic redshift from the purely photometric candidates. In both loci the photometric candidates reproduce the characteristic LRD V-shape traced by the spectroscopically confirmed members (a blue rest-ultraviolet, a minimum near the rest-frame optical break, and a red rest-optical rise), confirming that they share the SED of the confirmed population rather than being contaminants. Such agreement is expected by construction, since both are selected as neighbours on the same manifold, but it is a useful consistency check.

4.5. Brown dwarfs separate without an explicit cut

No brown-dwarf rejection was applied when constructing the manifold, yet the known brown dwarfs occupy a region well away from the LRD loci (Figure 6): none fall within the fiducial main region, and their median Mahalanobis distance to the main-region centre is ≈ 32 , compared with ≈ 1.2 for the LRDs. The separation is driven by the blue $F115W-F200W$ colours of the cool stars, which the manifold encodes naturally; the data-driven approach therefore reproduces, without being told to, the standard colour rejection. The zoomed view of the brown-dwarf region (Figure 6b) shows that a small number of photometrically selected LRD candidates from the literature lie nearby; these are plausible stellar contaminants of purely photometric selections, and their location adjacent to, but outside, our LRD region illustrates how the representation can flag such ambiguous cases.

4.6. Two regions, two populations

The two anchor islands are not redundant, but their main apparent difference is redshift. Drawing on the publicly available spectroscopic measurements from the de Graaff et al. (2025) catalogue, the clearest differences are in redshift (Figure 7a): the main-island anchors have a median spectroscopic redshift of 5.1 and the secondary-island anchors 3.4, and the main island is the more luminous in the rest-ultraviolet (brighter M_{UV}). The remaining fitted quantities (ultraviolet slope, Balmer-break strength, characteristic continuum temperature, Balmer decrement, and the $H\alpha$ and $[O III] \lambda 5007$ equivalent widths) show no clear separation between the two (Appendix C); the distinction is thus driven by redshift and luminosity more than by emission-line properties. This is consistent with the colours-only analysis of Section 5.3: removing redshift and morphology brings the two loci only marginally closer and they remain separable in colour, but because broadband colours encode redshift, much of that colour difference reflects the redshift offset itself, and in the rest frame both loci show the same characteristic LRD V-shape (Figure 5). The photometric candidates inherit the same split, with median photometric redshifts of 5.4 and 3.5 in the main and secondary islands respectively (Figure 7b), indicating that the candidates are drawn from the same populations as the anchors rather than being a separate contaminant set. Both loci show the characteristic LRD V-shape in the rest frame (Figure 5), the secondary at systematically lower redshift; within each locus the photometric candidates and the spectroscopically confirmed members trace the same shape, though the secondary spectroscopic stack is noisy given its 11 members.

4.7. Archival spectroscopy: a consistency check and new AGN

Finally, we use archival NIRSpec spectra to both verify our picture and look for confirmed LRDs and AGN among in-region sources beyond the anchor. The relevant sources are those that lie in one of the two regions (main and secondary) and have a spectrum but no high-quality PRISM classification (the anchor having classified the PRISM-observed sources already); these are predominantly sources with only medium-resolution grating data.

The continuum diagnostic (i.e. the V-shaped continuum fit to the low-resolution PRISM spectra; Section 3.4) confirms our framework but yields no new continuum-classified LRDs: applied to the anchor it recovers the great majority of sources, but among the in-region sources without a PRISM classification it returns no confident detections, simply because almost none of them have the PRISM continuum data the test requires. This is the expected outcome: over sources with adequate continuum spectra, the anchor selection is already essentially complete.

The broad- $H\alpha$ diagnostic (i.e. a fit for a broad $H\alpha$ component in the medium-resolution grating spectra; Section 3.4) is more productive. As a calibration we apply it to the anchor sources that have grating coverage of $H\alpha$: the recovery of a broad line is set by the $H\alpha$ signal-to-noise, reaching $\approx 93\%$ (13 of 14) above a signal-to-noise of ~ 15 , comparable to the broad-line fractions reported for spectroscopic LRD samples; the lower recovery at low signal-to-noise reflects the data, not the method. We then apply the same fit to the in-region sources that have grating coverage of $H\alpha$ but no PRISM classification. As the lower-redshift, $H\alpha$ -accessible sources concentrate in the secondary region, we search within this whole region, including the area contributed by the bright outlier anchor (Section 4.3). Of the 55 such grating




Fig. 4. The new candidates (blue) compared with the spectroscopic LRDs (red) in the two colour–colour planes used by the literature, with the corresponding selection boundaries overplotted, and in photometric redshift and F444W magnitude. The candidates are compact and at LRD-like redshifts but are bluer in the rest-optical, below the strict redness thresholds.

sources, 22 have $H\alpha$ in the reliable regime ($H\alpha$ S/N ≥ 15), and of these 11 show a statistically preferred broad component. After removing marginal and floor-limited fits and any source within $1''$ of an anchor, nine are robust broad-line detections: four are broad-line AGN with no prior catalogue identification (broad $H\alpha$ widths of $1.6\text{--}2.4 \times 10^3 \text{ km s}^{-1}$), and the remaining five are spectroscopic confirmations of photometrically selected candidates. All four of the new AGN lie in the secondary, lower-redshift region ($z \approx 2.3\text{--}3.5$), reinforcing its identification as a distinct population; two of the four also fall within the fiducial ten-anchor filament, the other two in the full region only. Each missed the anchor selection because it lacks a high-quality PRISM spectrum, and missed the photometric catalogues because it falls outside their colour or footprint selection. Their cutouts, SEDs and $H\alpha$ profiles are shown in Figure 8.

5. Discussion

5.1. Little Red Dots: selection and populations

The recurring difficulty in assembling LRD samples is that the selection function shapes the inferred population. The colour cuts in current use differ mainly in the adopted rest-optical redness threshold, and our diagnostic breakdown (Section 4.4) shows directly what this implies: the sources that the strict cuts omit are not failures of compactness or signal-to-noise, but com-

pact, high-redshift objects whose optical colours are simply not red enough to pass. The data-driven region recovers many of these because it is defined by where spectroscopically confirmed LRDs actually lie in the data, not by a colour boundary. We are deliberately cautious about how far to push this comparison. Our completeness is defined relative to the spectroscopic anchor, which is itself a bright, V-shape-selected sample, and our purity is measured over the subset with high-quality PRISM spectra. Since the anchor is used both to define the region and to judge it, these in-sample numbers are optimistic; a fairer test is the cross-validation in Section 4.2, where some anchors are held out during training and the method is scored on how well it recovers them.

Even with these caveats, the outcome is encouraging: using photometry, morphology and photometric redshift alone, the data-driven selection is already competitive with established colour cuts, and because the dominant source of incompleteness is the pre-processing (Section 2.4) rather than the manifold selection itself, there is clear room to push it further, for instance by relaxing the detection requirement at the bluest wavelengths or by adding spectroscopic and imaging features.

The choice of anchor also matters. We use the [de Graaff et al. \(2025\)](#) sample because its classification is spectroscopic and independent of broadband colours, so it does not import a broadband-colour prior at the classification stage into a method whose purpose is to avoid one. One clarification is in order.




Fig. 5. Rest-frame stacked SED of the two loci (left: main; right: secondary, the fiducial ten-anchor filament). In each panel the photometric candidates (coloured line: median with 16–84th percentile band) are compared with the spectroscopically confirmed members (black points, with error bars giving the standard error of the binned median). Each source is normalised to its median band flux so the stack reflects SED *shape*, and fluxes are stacked as f_ν by binning in rest-frame wavelength. Both loci show the characteristic LRD V-shape, and in each the photometric candidates track the spectroscopic members; the secondary spectroscopic stack is noisy because it contains only 11 objects.




Fig. 6. (a) Brown dwarfs (green) relative to the LRD loci. (b) Zoom on the brown-dwarf region, with nearby photometric LRD candidates highlighted as possible contaminants. (c) The same in colour space; brown dwarfs are blue in F115W–F200W.

The [de Graaff et al. \(2025\)](#) selection is not line-based: the low-resolution PRISM cannot resolve broad H α , so their classification rests on the continuum V-shape (the UV and optical slopes; Section 3.4). It is therefore a spectroscopic, higher-resolution measurement of the SED shape rather than an emission-line classification, that is, a refined, continuum-based counterpart of a colour selection rather than one orthogonal to it. When we call it independent of broadband colours we mean independent of the *photometric* colours we embed, not of SED shape in general.

The [de Graaff et al. \(2025\)](#) selection is also distinct from a spectroscopically *confirmed* but photometrically *selected* sample such as that of [Barro et al. \(2026a\)](#): although its members are bona fide LRDs, only about half of them fall in the main region, against more than four-fifths of the [de Graaff et al. \(2025\)](#) sample. This is apparent in Figure 2a, where the [Barro et al. \(2026a\)](#) and [Kokorev et al. \(2024\)](#) photometric samples, overlaid on the manifold, are visibly more dispersed than the tightly clustered spectroscopic anchors. The difference is informative (it reflects the broader, fuzzier reach of a colour-based selection), but it also

means that using such a sample to anchor the region would re-import colour-selection effects and blur the locus. We therefore keep it as a comparison and confirmation set rather than as an anchor. We stress that the residual targeting bias of spectroscopic samples is a caveat on the anchor’s completeness, not on the method itself: the LRD locus is a structure present in the data, which the anchor *locates* rather than creates (Sections 4.1, 4.2). A colour-biased anchor would change which known LRDs populate it (an effect we quantify as the completeness relative to the [de Graaff et al. 2025](#) sample; Section 2.4), but not the existence or position of the locus.

The [de Graaff et al. \(2025\)](#) sample remains the least broadband-colour-biased set available, and hence the natural anchor.

Finally, the two islands are not merely an artefact of the method but a real feature of the data (Section 5.3): the manifold separates the spectroscopic LRDs into two regions differing mainly in redshift and rest-ultraviolet luminosity (Section 4.6), sharing the same characteristic V-shape in the rest frame (see




Fig. 7. The two loci are two populations that differ mainly in redshift and rest-ultraviolet luminosity. (a) Spectroscopic anchors, from the public de Graaff et al. (2025) fits: the secondary island (blue) lies at lower spectroscopic redshift and is fainter in M_{UV} than the main island (red). (b) Photometric candidates: their photometric redshifts follow the same two-population split, indicating they are drawn from the same populations as the anchors. The remaining de Graaff et al. (2025) fitted quantities, which show no separation between the loci, are given in Appendix C.

Figure 5). Notably, the only new spectroscopic confirmations we obtain (four broad-line active nuclei) all lie in the lower-redshift, secondary region, in part because $H\alpha$ is accessible there, marking it as a promising target for dedicated follow-up. Whether the two regions are physically distinct populations or the same population seen across a range of redshifts remains open; their full physical characterisation (the relative contributions of active nuclei and host light, dust content and emission-line properties) is beyond the present scope and is deferred to future work.

5.2. A possible high-redshift extension of the main locus

We now return, deliberately speculatively, to the single de Graaff et al. (2025) anchor that our clustering set aside as an outlier (Section 3.2). It sits well away from both loci, yet its neighbourhood is not empty: it is occupied by a coherent population with a distinctive signature, a large flux gap between the two bluest bands (a median F814W–F115W of 2.0 mag, against 0.3 in the main locus), the classic imprint of a Lyman break falling between them. Because a single object does not permit the anchor-based region-drawing used for the two loci, we adopt for this preliminary look a deliberately simple proxy, and we stress it as a caveat: a *circular* region centred on the object, with a radius equal to the mean of the semi-minor and semi-major axes of the main-locus $f=1$ ellipse. This region contains 546 sources whose photometric redshifts pile up at very high values (median $z_{\text{phot}} \simeq 8.2$, with 98% at $z > 7$), far above the main locus members (median $\simeq 5.6$); the anchor object itself lies at $z \simeq 8.4$ (Figure 9). Their rest-frame stacked SED reproduces the LRD V-shape and matches that of the anchor object (Figure 9b). Taken together, this suggests that the region is a genuine *higher-redshift continuation of the main LRD locus*, resolved by the manifold as a separate concentration simply because the Lyman break has moved redward through the filter set. We emphasise the obvious limitation: the interpretation rests on a single spectroscopic redshift and on photometric redshifts that are themselves uncertain at $z \sim 8$; confirming it, and establishing whether these are bona-fide $z \sim 8$ LRDs, will require dedicated spectroscopy. Independently of the anchor, the region is interesting in its own right: it coincides with the highest-redshift part of the manifold in the photometric-redshift map (top-left panel of Figure 1). A natural

follow-up is spectroscopic confirmation of the most promising members, prioritised by proximity to the de Graaff et al. (2025) object and by compactness; in parallel, in a companion study we will mine the archival spectroscopy already available for these sources for LRD features, noting that at $z \simeq 8$ $H\alpha$ is redshifted out of the NIRSpc range, so the search must rely on other diagnostics (rest-ultraviolet lines, the Balmer break, or [O III] and $H\beta$ where covered). We flag this region as a concrete, testable prediction of the method.

5.3. Which features carry the LRD signature?

Every photometric LRD selection in the literature treats compactness as an essential ingredient, alongside the colour cuts. The representation lets us ask a sharper question: which of the input features actually carry the information that identifies an LRD? We address it with a feature-ablation experiment, measuring the clustering of the spectroscopic anchor for different subsets of the features (grouped as in Table 1), both in the feature space directly (a k -nearest-neighbour over-density, independent of any embedding) and through a cross-validated classifier (Figure 10). The result is unambiguous: the seven broadband *colours* alone reproduce essentially the full localisation of the LRDs (an over-density of ~ 420 against the random baseline, and a classifier AUC of 0.999), whereas morphology or photometric redshift on their own are weak discriminators (AUC 0.74 and 0.62). Adding morphology, reference flux or photometric redshift on top of the colours changes the clustering only marginally. In other words, the photometric fingerprint of the LRDs is carried by the SED shape, and the additional features are largely redundant for *identification*.

We confirm this directly by rebuilding the embedding from the colour ratios alone, switching off morphology, reference flux and photometric redshift (Figure 11). The spectroscopic LRDs again form a single, compact, over-dense concentration (Figure 11a), with a k -nearest-neighbour over-density comparable to the fiducial run. Three further points follow.

First, the photometric redshift is effectively re-learned from the colours: the colours-only manifold shows a clear redshift gradient (Figure 11b), a simple regressor predicts z_{phot} from the two coordinates with a coefficient of determination $R^2 \simeq 0.3$ (the 2D

New broad-line AGN in the loci (no prior catalogue)




Fig. 8. The four new broad-line AGN with no prior catalogue identification (Section 4.7). The top two panels are the two objects that also fall within the fiducial ten-anchor secondary region (the outlier-excluded filament; see Figure 2); the lower two lie in the full secondary region only, as labelled in the leftmost panels. For each target we report a 1.5'' RGB cutout (left), the observer-frame SED with uncertainties (centre), and the rest-frame grating spectrum around H α (right).

coordinates alone capture about a third of the redshift variance, modest but non-trivial for a two-dimensional embedding built without redshift), and, most tellingly, supplying redshift as an explicit feature barely alters the LRD clustering, consistent with colours encoding redshift.

Second, the colours-only map clarifies the relationship between the two fiducial populations. They are no longer laid out as two separate islands but merge into a single connected concen-

tration (Figure 11c); this merging, however, is largely a property of the projection. In the underlying feature space the two loci stay distinct even without redshift: a classifier separates their members from the broadband colours alone at a cross-validated AUC of ≈ 0.99 , against ≈ 0.996 with all features, that is, only marginally closer. The two populations are thus genuinely different in their observed colours, not merely a photometric-redshift artefact. Because broadband colours nonetheless encode redshift




Fig. 9. The discarded de Graaff et al. (2025) outlier and its neighbourhood (speculative; a single spectroscopic redshift). (a) UMAP zoom: the field colour-coded by photometric redshift, the outlier (star), the assumed circular region (mean semi-axis of the main-locus $f=1$ ellipse), and the main locus (dashed). (b) Binned rest-frame stacked SED of the region members (normalised to the median band flux) with the outlier’s own SED overlaid. (c) Photometric-redshift distribution of the region against the main locus; the region peaks near $z \approx 8$.

(the colours-only map recovers z_{phot} as a smooth gradient, Figure 11b), much of this colour difference reflects the redshift offset between the two, and in the rest frame both show the same characteristic LRD V-shape (Section 4.6); the lower-redshift locus is the one hosting the broad-line nuclei of Section 4.7.

Third, and importantly, morphology is not useless: its role is *purity* rather than identification. To compare like with like, we define the colours-only region exactly as in the main analysis (Section 3.2): we cluster the anchors, enclose each group with an ellipse containing all of its members ($f = 1.0$), and measure purity over the same spectroscopically classified subset. At this fixed, identical region-drawing choice, the colours-only region is markedly less pure than the fiducial one (a purity of ≈ 0.4 against ≈ 0.8), and the additional interlopers are predominantly resolved sources: $\sim 70\%$ of them fall below our compactness threshold (Figure 11d), and applying that threshold *a posteriori* restores the purity to ≈ 0.7 at essentially unchanged completeness. Compactness, the staple of photometric LRD selection, therefore acts in this framework as a filter against extended red interlopers, not as part of the core identification. One caveat here is worth investigating in its own right: not all of these resolved interlopers need to be genuine contaminants. An object whose light is a compact red nucleus embedded in a more extended host can have its integrated morphology (and hence ClassStarSE, which is measured on the F356W+F444W detection stack and therefore in the rest-frame optical; Section 2.2) biased towards “resolved” when the host’s rest-optical emission is extended. Its SED can nonetheless stay nucleus-dominated, especially if the photometry is extracted within an aperture that excludes much of the host, so the source still lands among the LRDs on the manifold. Establishing how many of the colours-only interlopers are of this kind, by inspecting their cutouts individually or with dedicated morphological modelling, is left to future work. We emphasise the natural caveats: the anchor is itself selected to be compact, so the colours that identify it are the colours of compact red sources; the parent sample is moreover restricted to isolated, well-measured objects. Within these limits, the experiment shows that, given confirmed LRDs, their broadband colours alone carry the photometric signature, and points towards an even simpler, colour-only data-driven selection in which morphology enters only as a purity control.

5.4. The method as a general discovery accelerator

Beyond the specific case of LRDs, the manifold itself is a general tool for population discovery. Two modes are particularly useful. First, any external sample can be placed on the existing manifold by position, so that one can ask immediately where a new class of objects falls, which known sources it resembles, and whether it forms a coherent region. Second, the same representation makes rare objects, sparsely populated regions, outliers and “bridges” between populations apparent, without any prior definition of what to look for (see e.g., Reis et al. 2021). As a proof of concept we have made the embedding interactive through a web-based visualiser². It displays the manifold and, for any selected source, fetches the imaging cutout and the NIR-Spec spectrum from the DAWN JWST Archive (Valentino et al. 2023; Heintz et al. 2024) and shows the SED in both the observed and rest frames together with colour-colour diagrams; selecting a region returns the stacked SED and the distributions of the enclosed sources. Such interactivity turns the representation into an exploratory instrument for identifying and vetting populations directly.

Known populations placed on the manifold fall where their physical properties would predict, not at random. The brown dwarfs are one example already shown (Section 4.5): with no rejection step they occupy a region well away from the LRD loci. Broad-line active galactic nuclei (BLAGN) are another. Projecting the JWST/NIRSpec BLAGN census of Baccus & Xu (2025) onto the map (Figure 12), 149 of the 252 sources fall within our parent sample, and they are far from uniformly distributed: about 40% (58) land in the two LRD loci and most of the rest concentrate around them and in a few extreme regions of the plane (such as the high-flux corner at the right-hand end). Colour-coding by redshift reveals the same trend seen for the LRDs themselves: the higher-redshift BLAGN gather around the main locus and the lower-redshift ones around the secondary locus. The overlap with the LRDs is substantial but partial: 26 of the 68 spectroscopic LRDs are also broad-line AGN, while more than half of the BLAGN lie outside the loci, consistent with the LRDs being one particularly compact corner of the wider accreting-black-hole population. This is exactly the question the representation answers at a glance: where a new sample sits, what it resembles, and how coherent it is.

² <https://micginolfi.github.io/compressedUniverse/>

Feature ablation: the colours carry the LRD signature




Fig. 10. Feature ablation. For each subset of the input features, the feature-space k -nearest-neighbour over-density of the spectroscopic LRDs (left) and the cross-validated classifier separability (right). The broadband colours alone carry the LRD signature; morphology and photometric redshift add little and are weak on their own.

A concrete demonstration is that the same representation that localises the LRDs also isolates other coherent groups of sources: regions of the manifold occupied by objects that share a distinctive signature. Exploring the embedding interactively with this viewer, we find several such regions; here we highlight three that illustrate the outlier- and population-finding capability, and that double as a useful data-quality diagnostic. In each case the grouping arises because the sources share a common distortion of their measured SED, which the unsupervised representation captures without being told what to look for.

The first region (Figure 13, top row) is a compact clump of bright point sources (stellarity ≈ 1 , photometric redshift ≈ 0). Their cutouts show saturated cores and strong diffraction spikes: these are bright, saturated stars, whose photometry is corrupted in a characteristic way that places them together on the manifold. The second region (Figure 13, middle row) is also dominated by point sources whose imaging is plainly stellar, yet their catalogue photometric redshifts pile up at $z \approx 6$ (median 5.9, with $\sim 80\%$ assigned $z > 4$). These are catastrophic photometric-redshift failures: Galactic stars fitted by high-redshift galaxy templates. Because the (erroneous) redshift is one of the input features, it displaces them into the high-redshift part of the map, but their stellar SEDs make them a coherent, anomalous island that stands out from the genuine high-redshift sources around it: the representation effectively flags the failure. The third region (Figure 13, bottom row) is a clump of faint sources whose cutouts reveal that they lie along the diffraction spikes of a nearby bright star; the spike contaminates their photometry and imprints a shared, non-monotonic SED distortion that again groups them together.

None of these groupings was sought in advance, and none required any threshold: they emerge because the embedding organises sources by the shape of their measured SED, so any population with a common spectral signature (astrophysical or instrumental) forms its own region. This has immediate practical value as an unsupervised quality-control step, automatically surfacing saturated sources, imaging artefacts and ancillary-catalogue failures that would otherwise contaminate a sample. More importantly, it is the same mechanism that enables genuine discovery: a previously unknown class of objects, or an outlier unlike the bulk of the population, will likewise occupy a distinct or

sparsely populated region. Combined with the interactive viewer, which allows arbitrary external samples to be projected onto the manifold by position and any region to be inspected through its stacked SED, cutouts and distributions, this makes the representation a practical instrument for both vetting and discovery. A systematic census of the additional astrophysical populations the manifold reveals is beyond the scope of this paper and is deferred to future work.

6. Conclusions

We have explored an unsupervised, data-driven route to identifying and characterising Little Red Dots in JWST surveys, embedding $\sim 242,000$ well-measured, isolated ASTRODEEP-JWST sources (Merlin et al. 2024) in a low-dimensional manifold of broadband colours, morphology and photometric redshift (Section 3.1), and anchoring it with the spectroscopically selected LRDs of de Graaff et al. (2025). Our main findings and their limitations are as follows.

- **LRDs occupy a distinct region of the manifold** (Section 4.1). Without any colour cut, the spectroscopic LRDs concentrate into a compact region that is over-dense by more than two orders of magnitude, split into two associated sub-regions (Section 3.2).
- **The structure is intrinsic, not a projection artefact** (Section 4.2). The same localisation is present in the original feature space, and a cross-validated classifier recovers $\sim 90\%$ of held-out spectroscopic LRDs, confirming that the location generalises.
- **Colours carry the LRD signature** (Section 5.3). A feature-ablation experiment shows that the seven broadband colours alone reproduce essentially the full localisation of the LRDs (classifier AUC 0.999), while morphology and photometric redshift are individually weak discriminators (AUC 0.74 and 0.62) and largely redundant for identification once the colours are included; morphology nonetheless improves purity.
- **The selection is competitive and interpretable** (Section 4.3). Defining the region with Mahalanobis ellipses, whose size sets the completeness with respect to the anchor,




Fig. 11. The colours-only embedding (morphology, reference flux and photometric redshift switched off). (a) The spectroscopic LRDs still form a single concentration and the brown dwarfs remain separate. (b) Photometric redshift is recovered as a smooth gradient. (c) In this colours-only projection the two fiducial populations merge into one concentration, although they remain separable in the underlying feature space (Section 5.3). (d) The interlopers admitted into the colours-only region are predominantly non-compact, i.e. exactly the sources a compactness criterion removes.

the main region reaches ~ 0.82 completeness at a directly measured purity of ~ 0.78 over the PRISM-classified subset (not of the whole photometric sample), placing it at the high-purity end of the completeness–purity plane, at completeness comparable to the best literature colour cuts (Table 2). For comparison, on the same sample the best hand-tuned colour selections reach high completeness but at lower purity (Kokorev et al. 2024: 0.93 completeness, 0.66 purity; Barro et al. 2026b: 0.90 completeness, 0.59 purity); the data-driven region achieves the highest purity at comparable completeness and the highest combined quality Q , without any hand-designed colour boundary.

- **A new candidate sample** (Section 4.4). The region yields 107 compact, high-redshift candidates absent from existing catalogues, robust to anchor resampling and to the cross-match tolerance. The strict literature redness cuts miss them because their rest-optical colours are not red enough; those that satisfy the more inclusive cuts are absent from the published catalogues instead through differences in the compact-

ness measure and in the photometry, not in colour (see Section 4.4).

- **Contaminants separate naturally** (Section 4.5). Brown dwarfs fall well outside the LRD region with no explicit rejection step, and the representation flags ambiguous, adjacent photometric candidates.
- **Two populations** (Section 4.6). The two regions correspond to two populations differing mainly in redshift, which shifts the V-shape break through the filter set and with it their broadband colours, a distinction inherited by the photometric candidates.
- **A predicted higher-redshift extension** (Section 5.2). The single anchor discarded as an outlier sits in a coherent neighbourhood of 546 sources whose photometric redshifts pile up at $z_{\text{phot}} \simeq 8$ and whose stacked rest-frame SED reproduces the LRD V-shape, suggesting a higher-redshift continuation of the main locus resolved separately only because the Lyman break has moved redward through the filter set. This is a concrete, spectroscopically testable prediction, and rests for now on a single spectroscopic redshift.




Fig. 12. Broad-line AGN from the Baccus & Xu (2025) JWST/NIRSpec census located on the manifold (the 149 of 252 that fall in our parent sample), colour-coded by spectroscopic redshift, with the main and secondary LRD loci ellipses for reference (the dotted curve is the full secondary region). The view is zoomed on the lower part of the plane, where the BLAGN and the loci lie. The BLAGN are not randomly scattered: they concentrate inside and around the LRD loci and in a few extreme regions, and follow the same redshift trend as the loci (high redshift near the main locus, low redshift near the secondary).

- **New broad-line AGN** (Section 4.7). A broad- $H\alpha$ analysis of archival NIRSpec spectra (Section 3.4), calibrated on the anchor ($\sim 93\%$ recovery at adequate signal-to-noise), confirms four broad-line active nuclei with no prior catalogue identification, all in the lower-redshift secondary region, plus five confirmations of photometric candidates.
- **A tool for discovery** (Section 5.4). Beyond LRDs, the same representation isolates other coherent populations and outliers as distinct regions of the manifold, including a group traceable to catastrophic photometric-redshift failures (Section 5.4), illustrating the method’s broader use as a discovery instrument.

These results come with clear caveats. Our completeness is defined relative to a bright, spectroscopically selected anchor, and the purity is measured only over the subset of sources with high-quality (grade-3) PRISM spectra (Section 4.3). The dominant incompleteness is set by the pre-processing, and in particular by the requirement of a detection down to F115W, the bluest NIRCам band (only HST/ACS F814W is exempt), which imposes an effective redshift ceiling of $z \lesssim 8$ and preferentially removes the highest-redshift, reddest sources, rather than by the manifold selection itself (Section 2.4); and the broad-line search becomes reliable only above a moderate $H\alpha$ signal-to-noise (Section 3.4). None of these limitations is fundamental to the approach, and each points to a concrete way of strengthening it.

Looking ahead, the framework extends naturally beyond photometry. The same representation can ingest spectra and resolved imaging. For instance, unsupervised methods have already been applied directly to JWST spectra (e.g. Saxena 2025), and a closely related semi-supervised approach, anchoring a learned representation with sparse labels, as we do here, has been used to classify AGN in DESI spectra (Alcolea et al. 2026). The manifold can then propagate sparse, high-confidence labels

through the learned space, and flag outliers and rare populations as they arise.

Acknowledgements. We are grateful to the organisers and participants of the hiking-workshop “Galaxy and SMBH formation, growth, and co-evolution” (Nepal, November 2025) for the fruitful discussions that inspired this work. FDE and RM acknowledge support by the Science and Technology Facilities Council (STFC), by the ERC through Advanced Grant 695671 “QUENCH”, and by the UKRI Frontier Research grant RISEandFALL. RM also acknowledges funding from a research professorship from the Royal Society. GV acknowledges financial support by the Italian National Institute for Astrophysics (INAF) under the IAF - Astrophysics Fellowships in Italy grant CUP C59J21034720001 - “AD MAJORA”. We thank Merlin et al. (2024) and the ASTRODEEP consortium for producing and releasing the ASTRODEEP-JWST photometric catalogue on which this analysis is based. Some of the data products presented herein were retrieved from the Dawn JWST Archive (DJA). DJA is an initiative of the Cosmic Dawn Center (DAWN), which is funded by the Danish National Research Foundation under grant DNR140. This work is based 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. This research made use of NUMPY (Harris et al. 2020), SCIPY (Virtanen et al. 2020), SCIKIT-LEARN (Pedregosa et al. 2011), MATPLOTLIB (Hunter 2007), ASTROPY (Astropy Collaboration et al. 2022), PANDAS (Wes McKinney 2010) and UMAP (McInnes et al. 2018).

References

- Akins, H. B., Casey, C. M., Lambrides, E., et al. 2025, *ApJ*, 991, 37
- Alcolea, J. A., Siudek, M., Eriksen, M., et al. 2026, *arXiv e-prints*, arXiv:2607.07329
- Arrabal Haro, P., Dickinson, M., Finkelstein, S. L., et al. 2023, *Nature*, 622, 707
- Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, *ApJ*, 935, 167
- Baccus, C. & Xu, X. 2025, *arXiv e-prints*, arXiv:2512.03281
- Baggen, J. F. W., Scoggins, M. T., van Dokkum, P., et al. 2026, *ApJ*, 1002, L4
- Bagley, M. B., Finkelstein, S. L., Koekemoer, A. M., et al. 2023, *ApJ*, 946, L12
- Bagley, M. B., Pirzkal, N., Finkelstein, S. L., et al. 2024, *ApJ*, 965, L6
- Barro, G., Pérez-González, P. G., Kocevski, D., et al. 2026a, *ApJ*, 1003, 96
- Barro, G., Pérez-González, P. G., Kocevski, D. D., et al. 2026b, *ApJ*, 997, 48
- Barro, G., Pérez-González, P. G., Kocevski, D. D., et al. 2024, *ApJ*, 963, 128
- Bezanson, R., Labbe, I., Whitaker, K. E., et al. 2024, *ApJ*, 974, 92
- Brammer, G. B., van Dokkum, P. G., & Coppi, P. 2008, *ApJ*, 686, 1503
- Brazzini, M., D’Eugenio, F., Maiolino, R., et al. 2026, *arXiv e-prints*, arXiv:2601.22214
- Casey, Q. O., Hickox, R. C., Cleri, N. J., et al. 2026, *arXiv e-prints*, arXiv:2606.26098
- de Graaff, A., Hviding, R. E., Naidu, R. P., et al. 2025, *arXiv e-prints*, arXiv:2511.21820
- D’Eugenio, F., Cameron, A. J., Scholtz, J., et al. 2025, *ApJS*, 277, 4
- Dunlop, J. S., Abraham, R. G., Ashby, M. L. N., et al. 2021, PRIMER: Public Release IMaging for Extragalactic Research, JWST Proposal. Cycle 1, ID. #1837
- Eisenstein, D. J., Johnson, B. D., Robertson, B., et al. 2025, *ApJS*, 281, 50
- Eisenstein, D. J., Willott, C., Alberts, S., et al. 2026, *ApJS*, 283, 6
- Finkelstein, S. L., Bagley, M. B., Ferguson, H. C., et al. 2023, *ApJ*, 946, L13
- Fotopoulou, S. 2024, *Astronomy and Computing*, 48, 100851
- Fu, S., Zhang, Z., Jiang, D., et al. 2026, *Nature Astronomy* [arXiv:2512.02096]
- Greene, J. E., Labbe, I., Goulding, A. D., et al. 2024, *ApJ*, 964, 39
- Hainline, K. N., Helton, J. M., Johnson, B. D., et al. 2024, *ApJ*, 964, 66
- Hainline, K. N., Helton, J. M., Miles, B. E., et al. 2026, *ApJ*, 1004, 223
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, 585, 357
- Heintz, K. E., Watson, D., Brammer, G., et al. 2024, *Science*, 384, 890
- Hocking, A., Geach, J. E., Sun, Y., & Davey, N. 2018, *MNRAS*, 473, 1108
- Hunter, J. D. 2007, *Computing in Science and Engineering*, 9, 90
- Juodžbalis, I., Ji, X., Maiolino, R., et al. 2024, *MNRAS*, 535, 853
- 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., van Dokkum, P., Nelson, E., et al. 2023, *Nature*, 616, 266
- Lin, X., Fan, X., Cai, Z., et al. 2026, *arXiv e-prints*, arXiv:2605.21574
- Loiacono, F., Gilli, R., Mignoli, M., et al. 2025, *A&A*, 703, A36
- Madau, P. & Maiolino, R. 2026, *arXiv e-prints*, arXiv:2605.05074
- Matthee, J., Naidu, R. P., Brammer, G., et al. 2024, *ApJ*, 963, 129
- Matthee, J., Torralba, A., Pezzulli, G., et al. 2026, *arXiv e-prints*, arXiv:2603.17667
- McInnes, L., Healy, J., & Melville, J. 2018, *arXiv e-prints*, arXiv:1802.03426




Fig. 13. Three peculiar regions that the manifold isolates without any prior definition, one per row: *top*, saturated bright stars; *middle*, catastrophic photometric-redshift failures (Galactic stars fitted with high-redshift galaxy templates, so their catalogue z_{phot} piles up near $z \simeq 6$); *bottom*, sources contaminated by the diffraction spikes of a nearby bright star. Columns: (a) the region’s location on the embedding; (b) example SEDs of indicative members, or, for the middle row, the region’s photometric-redshift distribution; (c) 3'' RGB cutouts of indicative members. Each region is a coherent group sharing a common, distinctive distortion of the measured SED, so it forms its own concentration on the manifold and is surfaced as an unsupervised data-quality diagnostic.

Merlin, E., Santini, P., Paris, D., et al. 2024, A&A, 691, A240
Pan, Z., Zhuang, M.-Y., Shen, Y., et al. 2026, arXiv e-prints, arXiv:2606.09721
Park, K., Torralba, A., Matthee, J., et al. 2026, arXiv e-prints, arXiv:2605.14233
Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, Journal of Machine Learning Research, 12, 2825
Pérez-González, P. G., Barro, G., Carniani, S., et al. 2026, arXiv e-prints, arXiv:2602.20247
Portillo, S. K. N., Parejko, J. K., Vergara, J. R., & Connolly, A. J. 2020, AJ, 160, 45
Reis, I., Rotman, M., Poznanski, D., Prochaska, J. X., & Wolf, L. 2021, Astronomy and Computing, 34, 100437
Rieke, M. J., Robertson, B., Tacchella, S., et al. 2023, ApJS, 269, 16
Rinaldi, P., Hainline, K., D’Eugenio, F., et al. 2026, arXiv e-prints, arXiv:2604.07138
Rusakov, V., Watson, D., Nikopoulos, G. P., et al. 2026, Nature, 649, 574
Saxena, A. 2025, arXiv e-prints, arXiv:2511.05439
Sneppen, A., Matthews, J. H., Watson, D., et al. 2026, arXiv e-prints, arXiv:2604.09399
Treu, T., Roberts-Borsani, G., Bradac, M., et al. 2022, ApJ, 935, 110

Valentino, F., Brammer, G., Gould, K. M. L., et al. 2023, ApJ, 947, 20
van der Maaten, L. & Hinton, G. 2008, Journal of Machine Learning Research, 9, 2579
Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nature Medicine, 17, 261
Wes McKinney. 2010, in Proceedings of the 9th Python in Science Conference, ed. Stéfan van der Walt & Jarrod Millman, 56 – 61
Zhang, Z., Jiang, L., Liu, W., Ho, L. C., & Inayoshi, K. 2026, ApJ, 998, 170

Appendix A: Feature-space classifier validation

As an additional check that the LRD localisation is not a projection artefact, we trained a supervised classifier directly on the original feature vectors. Each source was represented by the same 11-dimensional vector used to construct the embedding: seven F356W-normalised log-colours, the log F356W reference flux, the stellarity, the log half-light radius, and $\log(1+z_{\text{phot}})$. The features were robust-scaled before training. The classifier was a simple multilayer perceptron with two hidden layers and dropout regularisation, implemented in Keras/TensorFlow. We evaluated it with five-fold cross-validation: in each fold, the model was trained on four fifths of the spectroscopic anchors and the field, and predictions were recorded only for the held-out objects. The probabilities shown in Figure A.1 are therefore out-of-fold probabilities, so an anchor source does not contribute to the training set used to predict its own probability. This makes the test a validation of generalisation in feature space, rather than a restatement of the in-sample clustering.

The resulting probability distributions separate the spectroscopic LRDs from the field and show substantial overlap between the high-probability sources and the two-dimensional LRD region. We use this classifier only as a diagnostic: the candidate selection adopted in the paper is still defined geometrically on the unsupervised manifold, anchored by the spectroscopic labels.

Appendix B: Follow-up candidates

Figure B.1 shows the most locus-central candidates from the main region (Section 4.4) that have no spectrum and no prior catalogue identification, ranked by Mahalanobis distance to the region centre. For each source we show a $1.5'' \times 1.5''$ RGB cutout (F150W/F277W/F444W) and the observer-frame SED with flux uncertainties (downward triangles denote 3σ upper limits), so that the compactness and the red SED can be assessed directly. The complete ranked list, with coordinates, photometry and region membership, is provided as a machine-readable table ³, to be released publicly upon acceptance.

Appendix C: Spectroscopic properties of the two loci

Section 4.6 shows that the two anchor loci correspond to two populations differing mainly in redshift and rest-ultraviolet luminosity (Figure 7). For completeness, Figure C.1 shows the remaining quantities from the public de Graaff et al. (2025) spectral fits, matched to the anchors of each locus: the rest-ultraviolet slope β_{UV} , the Balmer-break strength, the continuum temperature T_{eff} , the Balmer decrement, and the $\text{H}\alpha$ and $[\text{O III}] \lambda 5007$ equivalent widths. None of these separates the two loci, confirming that the distinction is driven by redshift and luminosity rather than by emission-line or continuum properties.




Fig. A.1. Feature-space validation of the LRD localisation. (a) Out-of-fold probabilities from a neural-network classifier trained on the original 11-dimensional features, showing the separation between spectroscopic LRDs (red) and the field. (b) The same probabilities for sources inside and outside the two-dimensional LRD region; the high-probability sources substantially overlap the manifold-defined region.

³ <https://github.com/micginolfi/compressedUniverse>

Follow-up candidates: 1.5" RGB cutout and SED with errors




Fig. B.1. Follow-up candidates with no spectrum and no prior catalogue identification: for each, a 1.5'' RGB cutout (left) and the observer-frame SED with uncertainties (right).

de Graaff+25 fitted quantities with no separation between the two loci




Fig. C.1. Fitted quantities from de Graaff et al. (2025) that do *not* separate the two loci: rest-ultraviolet slope β_{UV} , Balmer-break strength, continuum temperature T_{eff} , Balmer decrement, and the $H\alpha$ and [O III] $\lambda 5007$ equivalent widths. Neither the main (red) nor the secondary (blue) island is distinguished by these, confirming that the two populations differ mainly in redshift and rest-UV luminosity (Figure 7).