



Joint Survey Processing. I. Compact Oddballs in the COSMOS Field—Low-luminosity Quasars at $z > 6$?

Andreas L. Faisst¹ , Ranga Ram Chary¹ , Sergio Fajardo-Acosta¹ , Roberta Paladini¹ , Benjamin Rusholme¹ , Nathaniel Stickley¹ , George Helou¹ , John R. Weaver^{2,3} , Gabriel Brammer^{2,3} , Anton M. Koekemoer⁴ , and Hironao Miyatake^{5,6,7}

¹ Caltech/IPAC, MS314-6, 1200 E. California Blvd. Pasadena, CA 91125, USA; aafaisst@ipac.caltech.edu

² Cosmic Dawn Center (DAWN), Copenhagen, Denmark

³ Niels Bohr Institute, University of Copenhagen, Jagtvej 128, DK-2200 Copenhagen, Denmark

⁴ Space Telescope Science Institute, 3700 San Martin Dr., Baltimore, MD 21218, USA

⁵ Institute for Advanced Research, Nagoya University, Nagoya 464-8601, Japan

⁶ Division of Particle and Astrophysical Science, Graduate School of Science, Nagoya University, Nagoya 464-8602, Japan

⁷ Kavli Institute for the Physics and Mathematics of the Universe (Kavli IPMU, WPI), University of Tokyo, Chiba 277-8582, Japan

Received 2021 March 8; revised 2022 January 19; accepted 2022 February 28; published 2022 April 13

Abstract

The faint-end slope of the quasar luminosity function at $z \sim 6$ and its implication on the role of quasars in reionizing the intergalactic medium at early times has been an outstanding problem for some time. The identification of faint high-redshift quasars with luminosities of $< 10^{44.5} \text{ erg s}^{-1}$ is challenging. They are rare (few per square degree), and the separation of these unresolved quasars from late-type stars and compact star-forming galaxies is difficult from ground-based observations alone. In addition, source confusion becomes significant at > 25 mag, with $\sim 30\%$ of sources having their flux contaminated by foreground objects when the seeing resolution is $\sim 0''.7$. We mitigate these issues by performing a pixel-level joint processing of ground and space-based data from Subaru/Hyper-SuprimeCam (HSC) and Hubble Space Telescope (HST) Advanced Camera for Surveys (ACS). We create a deconfused catalog over the 1.64 deg^2 of the COSMOS field, after accounting for spatial varying point-spread functions and astrometric differences between the two data sets. We identify twelve low-luminosity ($M_{UV} \sim -21$ mag) $z > 6$ quasar candidates through (i) their red color measured between ACS/F814W and HSC/*i* band and (ii) their compactness in the space-based data. Nondetections of our candidates in Hubble DASH data argues against contamination from late-type stars. Our constraints on the faint end of the quasar luminosity function at $z \sim 6.4$ suggest a negligibly small contribution to reionization compared to the star-forming galaxy population. The confirmation of our candidates and the evolution of number density with redshift could provide better insights into how supermassive galaxies grew in the first billion years of cosmic time.

Unified Astronomy Thesaurus concepts: [Surveys \(1671\)](#); [Quasars \(1319\)](#); [High-redshift galaxies \(734\)](#); [Photometry \(1234\)](#); [Astronomical techniques \(1684\)](#)

1. Introduction

Quasi-stellar objects (QSOs) or quasars, are powered by the accretion of gas onto a supermassive ($> 10^9 M_\odot$) black hole. Since the detection of luminous quasars within the first Gyr ($z > 6$) of the Big Bang using wide-area surveys such as the Sloan Digital Sky Survey, PanSTARRS, CFHTLS, and VIDEO/VIKING (Becker et al. 2001; Fan et al. 2001; Bañados et al. 2016; McGreer et al. 2018), it has been challenging to explain their origin and existence. Have their central engines built up their mass through sporadic, Eddington-limited accretion of gas, or do they build up through more continuous accretion processes onto a relatively massive black hole seed (Trakhtenbrot et al. 2017; Bañados et al. 2018)? If so, are the massive black hole seeds primordial in nature, or are they end stages of early epochs of massive star formation? Tracing the evolution of the quasar luminosity function, particularly at the low-luminosity end, at the earliest cosmic times, can potentially shed light on the origin of these systems. In addition, studying the number of quasars in the early universe can help quantify

their contribution to the reionization of the universe (see, e.g., Fan et al. 2006).

While identification and spectroscopic confirmation of luminous $z \sim 6$ quasars has become relatively straightforward, measuring the faint end ($< 10^{44.5} \text{ erg s}^{-1}$, $M_{UV} < -23$ mag) of the quasar luminosity function is much more challenging (see, for example, the comprehensive work by Akiyama et al. (2018), Niida et al. (2020), and Matsuoka et al. (2018) for quasar luminosity function measurements at intermediate brightness, $M_{UV} \sim -23$ mag, at $z \sim 4, 5$, and 6 , respectively). First, the identification of faint quasars requires a significant color difference across the Lyman-break/Ly α forest. This implies that, in order to measure the break, the blue band must be at least a magnitude more sensitive. Second, faint quasars are point sources and can lack multiwavelength information due to the sensitivity differences between the different bands. They can be mistaken as either late-type stars in the local universe or compact, strong emission line galaxies at intermediate redshifts. This is especially true at seeing-limited spatial resolution of $\sim 0''.7$. Space-resolution data with $< 0''.1$ point-spread function (PSF) full width at half maximum can help avoid this misidentification, but wide-area surveys from space at such resolution are not yet available except in the COSMOS field with the Hubble Space Telescope



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

(HST; Koekemoer et al. 2007; Scoville et al. 2007). The combined analysis of ground- and space-based data is therefore crucial to be able to robustly identify candidate, low-luminosity QSOs.

The next problem arises due to source confusion. At optical wavelengths, the classical confusion limit at seeing-limited resolution is 25 AB mag (5×10^{44} erg s $^{-1}$ at $z \sim 6$ is 25 AB mag) as estimated from source counts in deep surveys (see Appendix A). From this we can estimate $\sim 30\%$ of objects at that brightness are contaminated by the presence of a neighboring brighter source. This makes it challenging to reliably measure colors or color limits for the sources. Disentangling the relative contributions of the confusing sources requires priors from deep, higher spatial resolution data and joint pixel-level processing, taking into account accurate astrometry, and accurate position/source-dependent PSFs; this is beyond the capability of current cataloging algorithms. Future deep, wide-area surveys conducted from space using the Euclid and the Nancy Grace Roman space telescopes as well as from the ground using the *Vera C. Rubin* observatory are well in the regime of deep, confusion-limited, wide-area imaging. Although they promise to push the identification and characterization of *luminous* quasars out to redshifts beyond $z = 8$, the infrastructure for precise, joint analysis of those data sets does not yet exist.

In this paper, we demonstrate the scientific benefits of joint pixel-level processing over the relatively small area of 1.64 square degrees, to assess the faint end of the quasar luminosity function at $z \sim 6$. We combine two data sets that are very similar to the future surveys. These are the Subaru/HyperSuprimeCam (HSC; Miyazaki et al. 2018) *i*-band survey (Aihara et al. 2018), and the 2004 Hubble Advanced Camera for Surveys (ACS) F814W imaging survey on the roughly 2 deg 2 COSMOS field (Koekemoer et al. 2007; Scoville et al. 2007). These data have similar seeing and PSF properties as future Rubin and Roman/Euclid imaging observations.

This paper is structured as follows. In Section 2, we present in detail the different data sets used and their preparation for extracting photometry (astrometric calibration and measurement of the PSF kernel). In Section 3, we detail the generation of a photometry catalog based on a joint modeling of the ACS and HSC images. Specifically, we use the high-resolution ACS images as priors to mitigate blending and confusion issues in the HSC data. In Section 4, we present the sample selection and discuss different possible sources of contamination of our sample. We discuss our derived number density of low-luminosity quasars and the impact on reionization in Section 5 and conclude in Section 6.

Throughout this work, we assume a Λ CDM cosmology with $H_0 = 70$ km s $^{-1}$ Mpc $^{-1}$, $\Omega_\Lambda = 0.7$, and $\Omega_m = 0.3$. All magnitudes are given in the AB system (Oke 1974), and stellar masses and star formation rates (SFRs) are normalized to a Chabrier (2003) initial mass function.

2. Data Sets and Preparation

2.1. Imaging Data

Identification of faint $z \sim 6$ QSOs requires having adequately deep imaging data for a selection similar to that of Lyman-break galaxies (e.g., Steidel et al. 1996). Furthermore, we need at least one band to have high spatial resolution so that a separation between resolved galaxies and pointlike quasars is

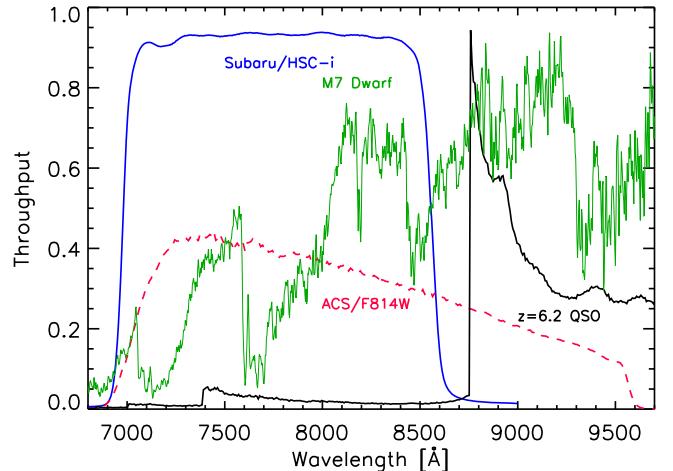


Figure 1. Throughput curves of the HSC-*i*-band (blue solid) and ACS/F814W filter (red dashed). The F814W filter extends to the red, hence the flux difference acts similar to a medium-band filter. This allows the selection of quasars through a [HSC-*i*]–[F814W] color difference. The black line shows an average quasar template made from a stack of ~ 2000 quasars in the Sloan Digital Sky Survey (Vanden Berk et al. 2001), redshifted to $z = 6.2$ with absorption by the intergalactic medium applied (Madau et al. 1999). We also show the spectrum of an M7 brown dwarf from Fajardo-Acosta et al. (2016) for comparison (green).

possible. We use the Hubble/ACS F814W images (Koekemoer et al. 2007; Scoville et al. 2007) and the 2018 data release 1 (Aihara et al. 2018) of the HSC-*i*-band ultradeep images (called *calexp*, calibrated stacked visits) from the Subaru Strategic Program survey in COSMOS. As shown in Figure 1, the bandpasses and depths of these data sets are well matched, allowing for the selection of red objects between $5.9 < z < 6.7$. The ACS data were taken between 2003 October 15 and 2005 May 21 while the HSC data were taken between 2014 March and 2015 November. Stellar proper-motions result in significant astrometric mismatches between these data sets, which need to be corrected for. Furthermore, these data have not been aligned to the Gaia astrometric reference frame, which needs to be remedied.

We start with the v2.0 ACS mosaics available within the Infrared Science Archive (IRSA)⁸, which have a PSF full width at half maximum of 0''.095 and a spatial scale of 30 mas pixel $^{-1}$ (Koekemoer et al. 2007). In contrast, the HSC data have a PSF FWHM between 0''.6 and 0''.9 and a spatial scale of 168 mas pixel $^{-1}$ (Aihara et al. 2018).

These data are advantageous for the search of high-redshift objects, including quasars. The F814W bandpass overlaps entirely with the Subaru/HSC-*i*-band filter and extends redward (see Figure 1 and Section 4.1). This allows high- z sources (at a median $z \sim 6.4$) to be selected based on a large color difference in these bandpasses due to their Lyman-break and Ly α emission. The configuration acts similarly to the selection in a medium-band filter. The sensitivity of the surveys should allow the detection of objects down to $M_{UV} = -22.1$ mag, at least 3–4 mags deeper than DESI Legacy imaging surveys, over much smaller areas (Dey et al. 2019), with the added advantage of high spatial resolution in one of the bands.

⁸ <https://irsa.ipac.caltech.edu/Missions/cosmos.html>

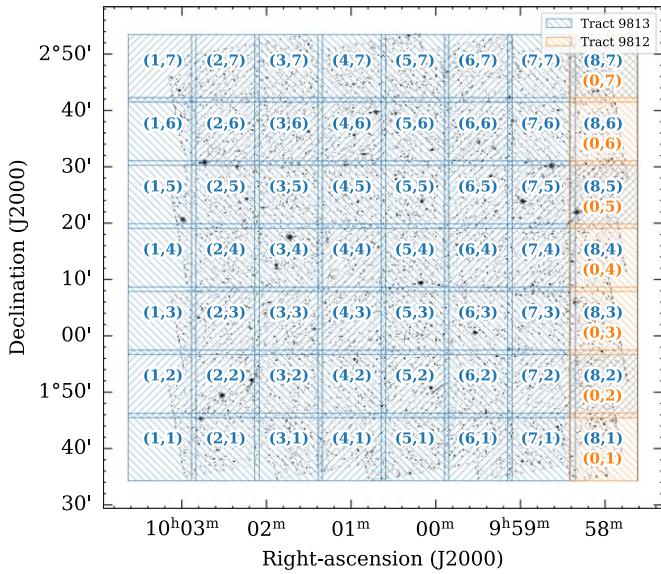


Figure 2. Spatial organization of *patches* on the COSMOS field. Shown are only the $63 \times 12'$ patches overlapping with the ACS/F814W observations (gray scale, background). COSMOS is covered by *tract* 9813 (blue) but also has a small coverage from *tract* 9812 (orange). The *patch* numbers are indicated as (X,Y) pairs. Each of the *patches* is divided into 16 *sub-patches* (size $3' \times 3'$) to facilitate multiprocessor computing.

2.2. Data Organization Terminology

For efficient handling and computing of joint catalogs between these data sets that are a total of ~ 1 TB, we organized the data in *tracts*, *patches*, and *sub-patches*. The HSC data is organized in $1^{\circ}7 \times 1^{\circ}7$ *tracts*. Each of them is split in 81 $12' \times 12'$ *patches* (see Aihara et al. 2018). The *tracts* have a $1'$ overlapping region, while the *patches* overlap within 100 px, which corresponds to $16''.8$. The *patches* themselves are stacks of different visits.

The ACS observations on the COSMOS field are covered by *tract* numbers 9812 and 9813, including 63 *patches* of the total 133 *patches* contained in the ultra-deep data. To allow multiprocessing of these data in the following, we cut the *patches* into 16 *sub-patches* of size $3' \times 3'$ and overlap of $10'$. Hence, in total there are 1008 *sub-patches* to be processed (Figure 2).

While the HSC images are already cut into *patches*, we scripted a Python wrapper to use the online IRSA cutout tool⁹ to cut and retrieve the ACS images at the correct size and location from the IRSA server. The *sub-patches* are then created using the Cutout2D Python package provided by Astropy¹⁰ (Astropy Collaboration et al. 2013, 2018).

2.3. Astrometric Calibration of Imaging Data

The alignment of images (i.e., their relative astrometric calibration) is crucial to perform any successful joint pixel-level processing. Among others, an accurate astrometric calibration of the images allows us to use the ~ 11 yr time baseline between ACS and HSC data to study the proper motions of faint stars (see Fajardo-Acosta et al. 2022). For this work, an accurate relative astrometric calibration is needed to use the priors on location and sizes from the Hubble/ACS

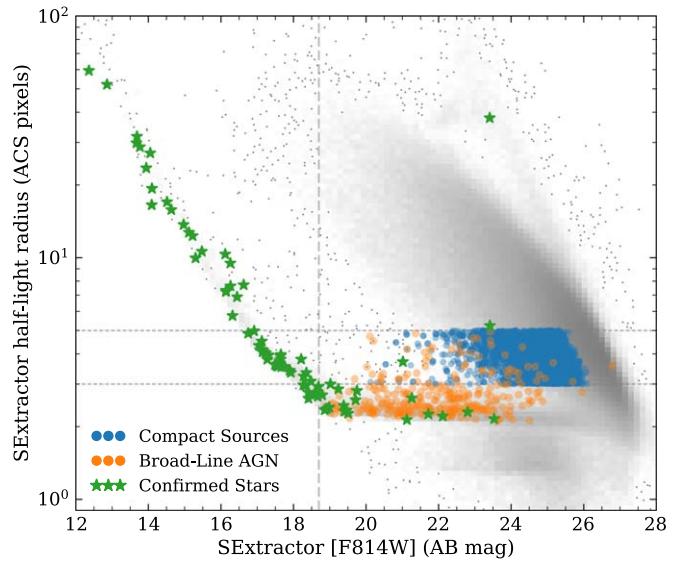


Figure 3. Selection of compact sources (blue) for astrometric calibration of imaging data. Also shown are confirmed broad-line AGN (orange, Civano et al. 2016; Marchesi et al. 2016; Lanzuisi et al. 2018) and spectroscopically confirmed stars (green). The gray symbols (“background cloud”) are all the objects in COSMOS. The “stellar locus,” where point sources lie, is horizontal from right to left, and turns up around 18.7 mag (indicated by the vertical dashed line), which is the magnitude where point sources become saturated. The horizontal dashed lines indicate sizes of 3 and 5 ACS pixels. See text for more details on the selection of compact sources.

imaging, to mitigate the effects of blending and confusion in ground-based images (see Section 3). The HSC data used here had been aligned with PanSTARRS DR1, while the astrometric reference frame for the ACS data was defined by CFHT *i*-band mosaics (Capak et al. 2007) that covered the full COSMOS field to a 5σ limiting depth of 26.2 mag. It provided ~ 300 –600 sources on each ACS tile, which could be used for relative registration. Each ACS tile was then registered to the CFHT grid, and the overall relative alignment precision was found to be ~ 5 –10 milliarcseconds (see also Koekemoer et al. 2011). However, this relative precision between the ACS and the CFHT sources does not account for possible absolute errors in the CFHT astrometry, which we explore.

We use two methods to align the ACS and HSC images *patch* by *patch*. In the first method, we compute the absolute astrometric alignment of the images using 3937 stars from the Gaia DR2 catalog (Lanzuisi et al. 2018), matched to ACS sources within a $0''.2$ radius. Since Gaia stars are generally bright, saturation is an issue especially for the deep ground and space-based images used here. To explore the saturation limit, we run SExtractor (Bertin & Arnouts 1996) on the ACS images to generate a catalog with approximate source sizes (half-light radii). We define a “stellar locus” in a magnitude versus size plot as shown in Figure 3. Typically, stars would be unresolved and would have FWHM sizes of ~ 2 –3 ACS pixels. However, bright/saturated stars have a larger fraction of their PSF profile above the noise threshold, which is the reason for their size increasing as a function of brightness. From the upturn of the “stellar locus” on this diagram, as well as the Gaussian fits to individual stars, we estimate a saturation threshold of about 18.7 mag. The faintest Gaia stars in our ACS sample were 20.9 mag in our catalog, and so we use sources in the magnitude range 18.7–20.9 mag that have less than 3 pixels half-light radius, to estimate the astrometry. We identified the

⁹ https://irsa.ipac.caltech.edu/data/COSMOS/index_cutouts.html

¹⁰ <http://www.astropy.org>

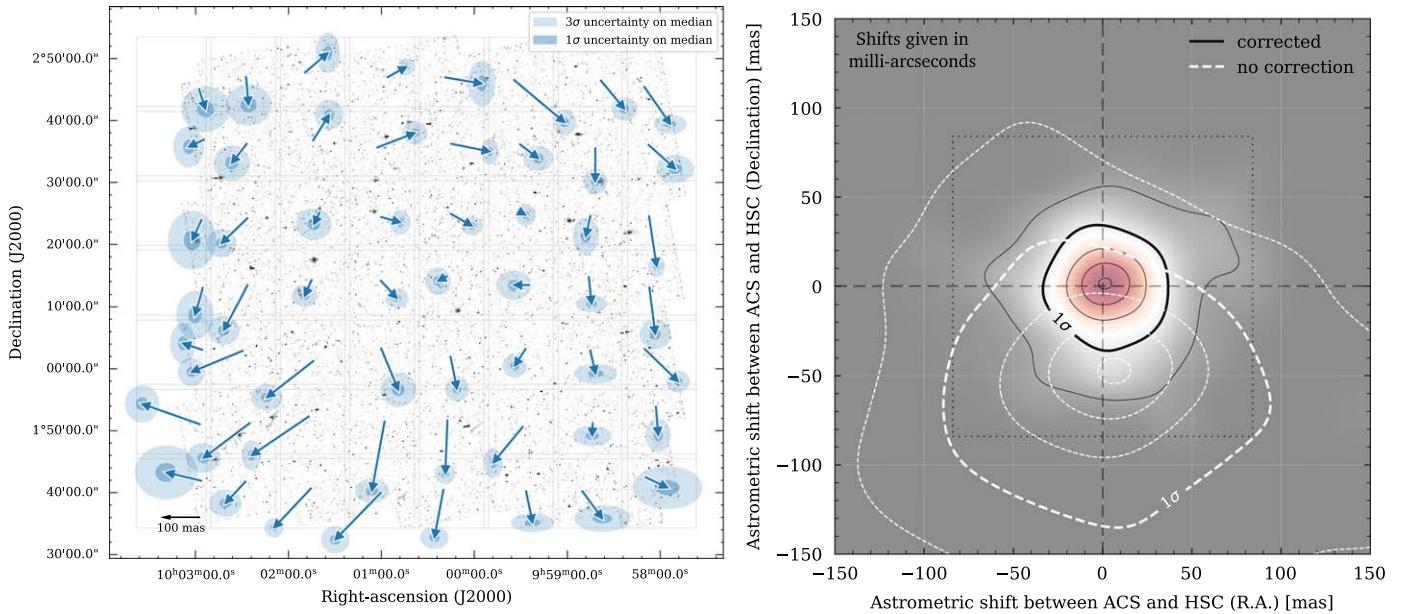


Figure 4. Astrometric calibration of ACS and HSC images. Left: astrometric offset between ACS (from Koekemoer et al. 2007, originally registered to the ground-based astrometric catalog from Capak et al. 2007) and HSC, for each of the $12' \times 12'$ patches. The ellipses show 1σ and 3σ uncertainties in the direction and length of the vectors. The 100 mas length is shown on the lower left. Right: accuracy of astrometric corrections (in milliarcseconds) tested on a sample of 150 spectroscopically confirmed broad-line AGN (Civano et al. 2016; Marchesi et al. 2016; Lanzuisi et al. 2018). The white contours show the distribution of the astrometric offsets between ACS and HSC before correction. The black contours show the offset after applying the corrections. We find no significant residual offset, and the accuracy is $\sim 30 - 40$ mas.

positions of these objects on the HSC images using two-dimensional Gaussian fits. This was not necessary for the ACS images because the centroiding accuracy is much smaller than the typical astrometric scatter that we expect. For each of the different *patches*, we then computed the true position of the Gaia stars from their DR2 catalog proper motion. While the epochs of the HSC images are within 2015 (close to the epoch of Gaia, 2015.5), the ACS images were taken between 2003 and 2005. We therefore computed the mean epoch of each ACS source from the individual exposures that covered it. For HSC, all sources within a *patch* were measured simultaneously. The mean of the offsets between the true and measured positions of the Gaia stars are then used to compute the astrometric offset between ACS and HSC of a given *patch*.

For the second method, we use compact extragalactic sources on both images to compute their relative astrometric offset (Figure 3). Ideally, one would use quasars or active galactic nucleus (AGN) for this, as they are point sources and do not have proper motions. However, the density of bright quasars is less than a few per square degree. Instead, we select our compact sources to have (i) a minor-to-major (B/A) axis ratio of more than 0.9 on the ACS images, (ii) a signal-to-noise ratio (S/N) of more than 25 in ACS, (iii) an ACS magnitude fainter than 20 mag, and (iv) sizes between 3 and 5 ACS pixels (corresponding to $0''.09 - 0''.15$). In addition, we avoid blended sources by restricting the sample to those with SExtractor flag `FLAG = 0` (on ACS and HSC images). Finally, we apply a S/N threshold of 10 for their HSC photometry measured by SExtractor. In total, we select about 5000 compact sources. Their position on the HSC and ACS images are then compared to compute the astrometric offset between these images.

We find that both methods lead to similar astrometric offset corrections (within $\sim 3 - 5$ mas) and therefore choose to use the second method in what follows. Furthermore, we note that

some of the bright Gaia stars could be in the nonlinear/saturation regime even at 20 mag. This could potentially effect the astrometric correction. However, since we find almost identical astrometric offsets using Gaia stars and faint compact sources, we think that this effect is negligible.

The left panel of Figure 4 shows the average relative astrometric offset between ACS and HSC images per $12' \times 12'$ patch. The ellipses show the 1σ and 3σ uncertainties in direction, and the length of a 100 mas shift is indicated on the lower left. Note the clear clockwise circular pattern that could be due to inhomogeneous astrometric calibration due to distortions on either the ACS images or the HSC images. Subsequent data releases both by HSC and a rereduction of the ACS data (G. Brammer et al., private communication) have fixed these offsets, but at the time when this work was initiated, these astrometric offsets were present.

The right panel of Figure 4 is an assessment of the accuracy of our astrometric calibration. Specifically, we apply the astrometric shifts to a sample of 150 spectroscopically confirmed broad-line AGN (see Figure 3; Civano et al. 2016; Marchesi et al. 2016; Lanzuisi et al. 2018). These are point sources and do not have proper motions, hence are ideal to test the astrometric calibration. The white contours show the shift in R.A. and decl. before applying the astrometric correction between HSC and ACS. The black contours show the distribution after correction. It is centered at zero (as it should be because AGN do not have motions) with a 1σ width of 30–40 mas. The latter is our precision of astrometric calibration. For more details, we refer to our companion paper on proper motions of faint stars by Fajardo-Acosta et al. (2022). We anticipate that, in space-based data with larger fields of view resulting in higher source numbers, the astrometric precision can be improved further.

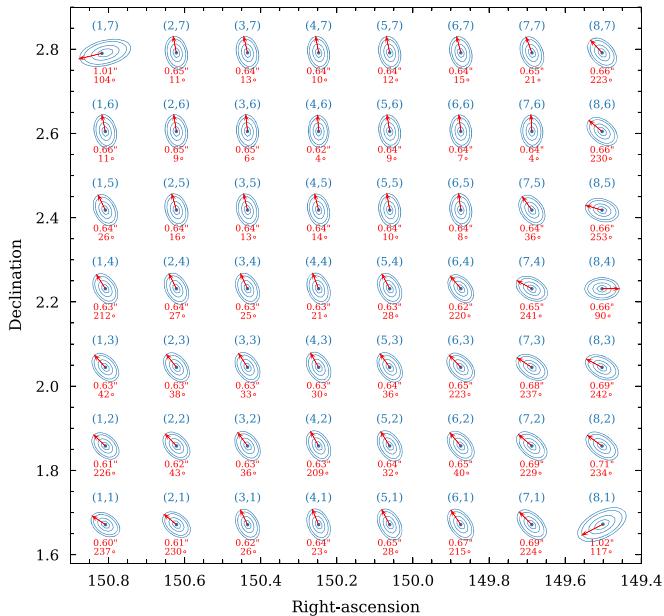


Figure 5. Variations in the HSC PSF across the COSMOS field. Generally, the variations are on a very small level; however we notice that the PSFs have a preferred north–south direction (ellipticity exaggerated by 50%). The contours show the 20%, 50%, 80%, and 90% enclosed flux of the PSF. The red arrows show the semimajor axis (with FWHM and position angle indicated). The *patch* numbers are indicated in blue.

2.4. PSF Estimation

In addition to the astrometry, the PSF of both HSC and ACS images has to be known accurately in order to produce reliable photometry. To create a spatially varying PSF, we stack unsaturated Gaia stars as well as fainter stars. We select the latter of the ACS images by extrapolating the “stellar locus,” fit by the Gaia stars on the magnitude versus size diagram (Figure 3), to lower magnitudes of 23 mag. To create the stacks, we use the code PSFex (Bertin 2011), which creates a magnitude-dependent model PSF based on a linear combination of (subpixel centered) stars in a given catalog. It also outputs the residual after subtracting a scaled PSF for each of the stars in the sample, which is useful to access the quality of the fits. We found that the residuals are smaller by a few percent for a magnitude-dependent fit compared to a simple stack. Depending on the ACS coverage, we are able to use between 30 and 100 stars per $12' \times 12'$ *patch* for the fit. The ACS PSF FWHM is less than $0''.1$, while the HSC PSF FWHM varies between $0''.60$ and $0''.75$ with a median around $0''.64$. We find that the HSC PSFs have a preferred direction, which is approximately north–south. Other variations in rotation and ellipticity between the different *patches* are relatively minor but nonetheless have to be taken into account to measure robust photometry (Figure 5).

3. Joint Cataloging with TRACTOR

Once the astrometric offset between the ground- and space-based data sets is established, and the PSFs kernels for each data set are calculated, we can undertake joint pixel-level photometry. This takes into account the location and morphological extent of the sources in the space-resolution data to alleviate the role of source confusion and accurately measure the photometry in the ground-based data. Here we use the code TRACTOR¹¹ (Lang et al. 2016a, 2016b; Weaver et al. 2022).

In brief, TRACTOR performs a parametric shape fit to a source in an image by using a maximum likelihood analysis including a weight map of the image. Thereby multiple sources can be fit simultaneously on an image, which is the preferred way to run TRACTOR to obtain robust photometric measurements (see also detailed description and testing in J.R. Weaver et al. 2022).

TRACTOR has two major advantages compared to classical photometry codes such as SExtractor or other aperture-based methods. First, it can use shape priors derived from a high-resolution image to photometer low-resolution images. To do so, TRACTOR can be forced to fix the shape and position and only vary the normalization (i.e., total flux) to minimize the residuals. In overlapping bands such as HSC – *i* and ACS F814W, the morphological k – correction is negligible; hence this is a reasonable assumption. Since TRACTOR can fit nearby sources of light simultaneously, this forced-photometry approach is valuable in fields with high confusion or blending. Second, because of the parametric fitting approach and the inclusion of the PSF (in the model step), TRACTOR provides PSF-corrected total fluxes for the fitted objects; this is especially valuable for extracting the contribution of low surface brightness regions beyond the isophotal size of objects.

On the other hand, there are regimes where the performance of TRACTOR is significantly reduced. Although TRACTOR provides a variety of models to parameterize the light distribution, including point sources (i.e., δ – functions), simple Gaussians, Sérsic models, and even multi-Gaussian representations, it would fail to measure robust fluxes of nonsmooth light distributions. Examples are galaxies with extended spiral structure or bulge and disk components (specifically when fit with a single Sérsic profile), or irregular galaxies such as lumpy galaxies at high redshifts. As we later demonstrate through simulations (Section 3.2), we found that, for our scientific goal (measurement of the photometry of compact sources), TRACTOR results in accurate photometry.

In Appendix B, we compare the performance of our method to another extensively used photometry package called TPHOT (Merlin et al. 2015, 2016), which is currently being upgraded in preparation for Euclid joint photometry. We find very good agreement between both methods.

3.1. ACS and HSC Photometer Pipeline

Our pipeline takes as input an ACS and HSC *patch* together with the corresponding PSF (Section 2.4) and astrometric shifts between the images (Section 2.3). It then runs TRACTOR on the ACS image to create a parametric model for each source in the image. The model for each source is then convolved by the (positional dependent) HSC PSF and scaled to fit the data in the HSC image (see Figure 6). We detail the different steps in the following. Multiple *sub-patches* are processed in parallel (see Appendix C).

Step 1—Preparation. We first cut a $3' \times 3'$ *sub-patch* with $10''$ overlap from the HSC *patch* together with the corresponding ACS *sub-patch*. The overlap ensures a good fitting of sources at the edges. To keep the World Coordinate System information, we are using the Python command Cutout2D from the AstroPy package. The following steps are applied to a single *sub-patch*.

Step 2—Obtain positions and preliminary shapes. We first run SExtractor on the ACS image. This provides the locations of the sources as well as the initial shape parameters (such as *A_IMAGE*, *B_IMAGE*, *FLUX_AUTO*, *FLUX_RADIUS*, and *THETA_IMAGE*) that are needed as initial guesses for

¹¹ <http://thretractor.org/>

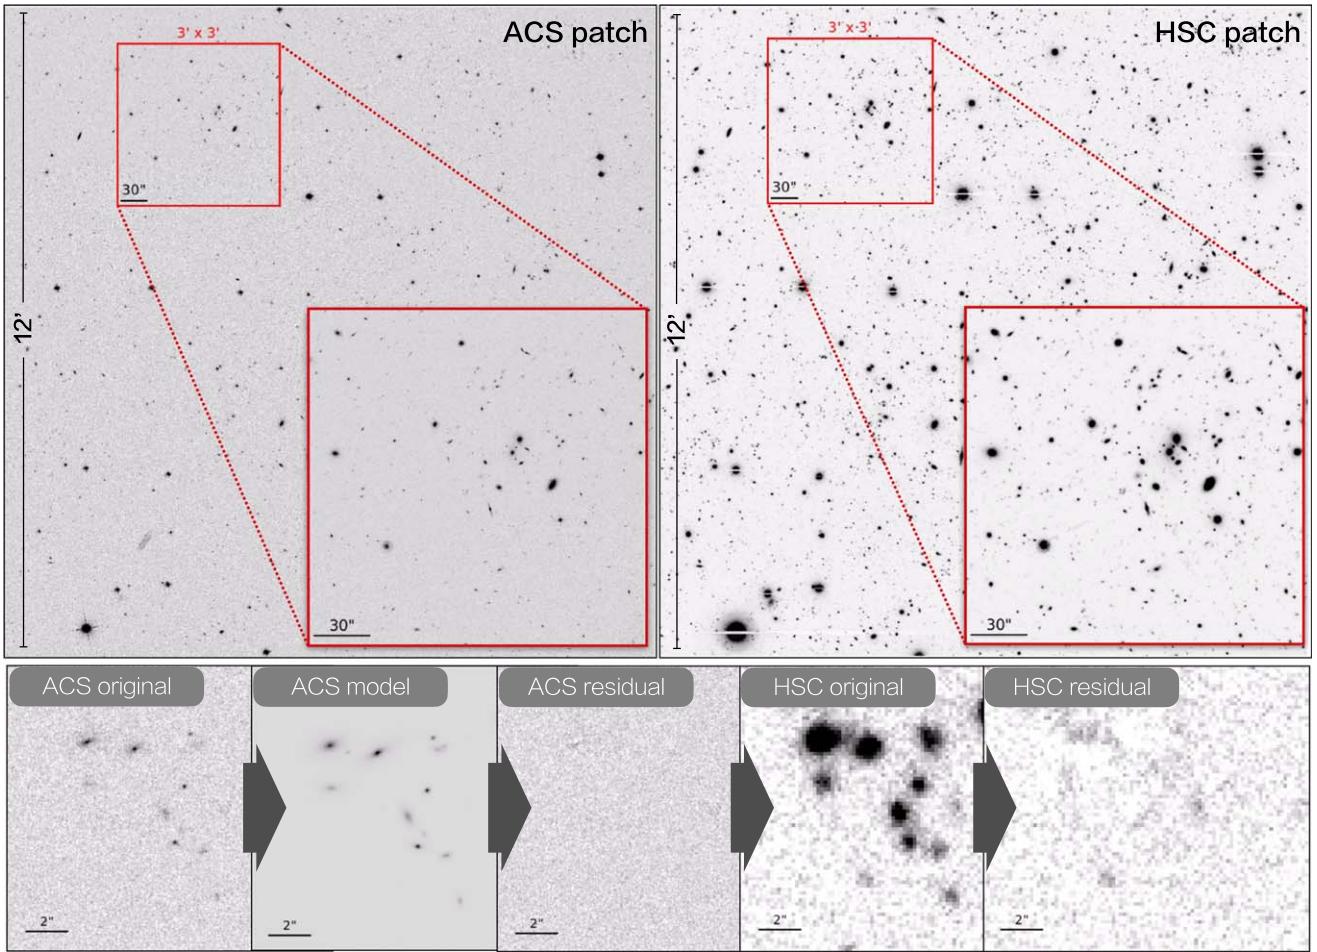


Figure 6. Illustration of our prior-based photometry approach with TRACTOR. The top large panels show the $12' \times 12'$ patches of the ACS and HSC observations. The inset shows a $3' \times 3'$ region, which is the size of a sub-patch. The bottom row of panels shows different stages of our approach. From left to right, these are (i) original ACS image, (ii) ACS model image created with TRACTOR, (iii) ACS residual image (original minus model), (iv) HSC original image, and (v) HSC residual image after convolving the ACS models with the difference in the PSFs and scaling to the HSC images of the sources. This prior-based approach has the advantage that it can efficiently deblend sources when applied to deep, low-resolution images.

TRACTOR. It also produces a segmentation map that identifies the pixels out to the isophotal radius of each source. The ACS data has very few repeats per pixel on the sky as a result of which there are residual cosmic rays in the mosaic. We perform a removal of such spurious sources in the SExtractor catalog, which are characterized by sizes smaller than the diffraction limit, with `FLUX_RADIUS < 1.7` ACS pixels. We found that this cut removes most cosmic rays and spurious sources clustered at the edges of the ACS coverage.

Step 3—Obtain ACS models and photometry. TRACTOR is run on each source extracted in step 2. Specifically, we create a cutout for each source using the SExtractor `XMIN`, `XMAX`, `YMIN`, and `YMAX` keywords. We increase the size of this cutout by 50% and fit all objects within this cutout (also the ones not covered entirely) simultaneously. It is important to distinguish between the unresolved and resolved sources for obtaining the best possible fits. Specifically, the unresolved sources (or point sources) are fit using TRACTOR’s `PointSource`¹² class. All other sources are fit using the `SérsicGalaxy`¹³ class. We assign a point-source flag to all sources with either (i) `CLASS_STAR` greater than 0.8, brighter

than 23 mag, and axis ratio¹⁴ greater than 0.8 or (ii) fainter than 23 mag, axis ratio greater than 0.8, and half-light radius smaller than the F814W PSF FWHM. The parameters measured by SExtractor in step 2 and a Sérsic index $n = 2$ are used as initial guesses. The TRACTOR `Image` object is created using the `PixelizedPSF` class, which converts the input PSF from `FITS` format to a format suitable for TRACTOR. Furthermore, we feed in the per-pixel noise measured from a 3σ clipping on the image. The background is fixed at zero level as the images are background subtracted. During fitting, we allow the position to wander within ± 1 px (corresponding to ± 30 mas) from its initial guess. All other parameters, except the position angle (given by `THETA_IMAGE`), are set free to vary.

Step 4—Forced-photometry on the HSC image. Finally, the best-fit parameters (position, flux, and shape parameters in the case of resolved sources) obtained in step 3 are used to photometer the HSC image. In a similar manner as in the previous step, cutouts of each source are created. However, as we expect significant blending and confusion compared to the ACS image, we cannot simply take the segmentation image created in step 2. Instead, we first run SExtractor on the HSC image and enlarge the resulting segmentation map by

¹² The `PointSource` class requires a position and total flux.

¹³ The `SérsicGalaxy` class requires a position, total flux, axis ratio, position angle, half-light radius, and Sérsic index n .

¹⁴ In the following, defined as the ratio of `B_IMAGE` and `A_IMAGE`.

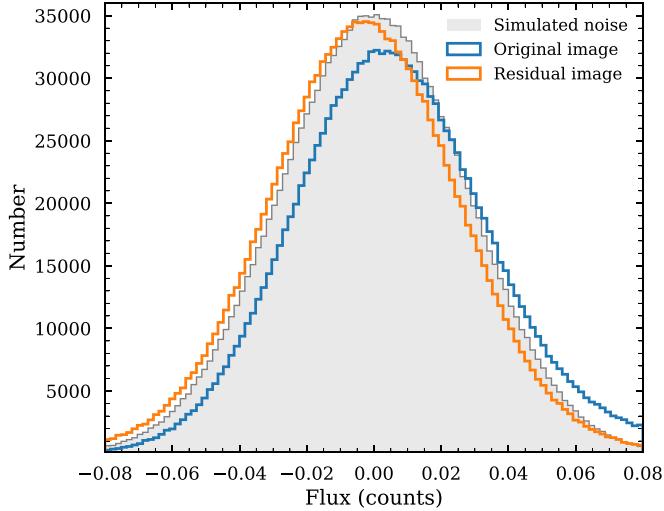


Figure 7. Comparison of the pixel flux distribution on the original HSC image (blue) and the residual HSC image after photometry with TRACTOR (orange). The gray histogram shows the expected distribution for uniform noise across the image. The peak of the residuals is shifted slightly negative, which indicates a slight *oversubtraction* of sources (i.e., overestimation of flux).

convolving it with a 1'' boxcar filter. This will combine the segmentation areas of close-by objects, creating one single large area.¹⁵ From this, we recalculate the XMIN, XMAX, YMIN, and YMAX. For each source in the ACS TRACTOR catalog (from step 3), we check in which enlarged segmentation area it falls and create a cutout accordingly. All sources within this cutout are fit simultaneously using their ACS position and shape priors. If no segmentation area was found (e.g., due to faintness or color gradient), we apply a 2'' × 2'' cutout size. Note that we add the astrometric offsets between the HSC and ACS images obtained in Section 2.3 to the ACS prior position in this step. During the fit, we allow the positions to vary within ±0.5 px (corresponding to ±84 mas) and fix all other parameters except the total flux. We note that because we are applying a median astrometric offset between the ACS and HSC images per *patch*, allowing the positions to vary a small amount is crucial to account for astrometric scatter.

Figure 7 shows the histograms of the pixel flux distribution on the original HSC image as well as the residual image after fitting all sources with TRACTOR. While the original image (blue line) shows a tail toward positive fluxes (these are the detected sources), the residual image (orange line) is in good agreement with the expected noise distribution (simulated uniform Gaussian noise across image without sources). There is a small imbalance on the residual image toward negative fluxes, which could hint toward a slight oversubtraction of fluxes by TRACTOR. Statistically, this offset is less than 0.01 mag at 25.5 mag.

Appendix C provides detailed descriptions of the setup and the process to run the pipeline efficiently at the National Energy Research Scientific Computing Center (NERSC).

3.2. Simulations

To demonstrate the performance of our pipeline in measuring point-source fluxes and to compute the sensitivity limits

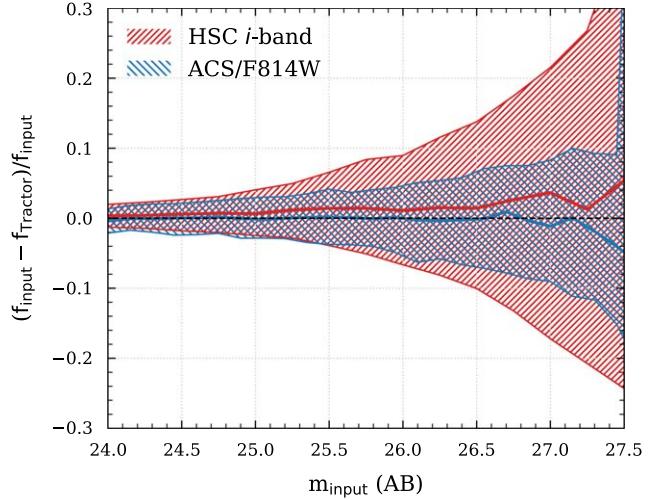


Figure 8. Recovery of point-source photometry as a function of magnitude for ACS (blue) and HSC (red). The hatched regions show the 16th/84th percentiles of the scatter and the solid lines show the running median. Undetected point sources are included with $f_{\text{TRACTOR}} = 0$. We do not find any significant biases in the point-source photometry.

and detection completeness, we created simulated images including point sources of different brightness using the software SKYMAKER¹⁶ (Bertin 2009). In total, we simulate 15,000 point sources in the magnitude range 24–27 mag and 30,000 point sources in the range 27–29 mag arranged in a grid to avoid blending (we discuss the effects of blending in Appendix A). The increased number of faint sources provides us with a more uniform distribution of sources across this magnitude range. The pixel noise properties (specifically the per-pixel rms measured in empty apertures) of the simulated images match the ones of the real ACS and HSC images. We found that correlated noise is negligible and is therefore not included. We also use the measured PSF of ACS and HSC. The generated simulated images are then run through our pipeline with the same setup as used for the real images.

Figure 8 shows the difference between the input and recovered magnitudes for F814W (blue) and HSC-*i* band (red) as a result of this test. The median (solid lines) and 68% percentiles (hatched regions) are computed in running bins of size 0.2 mag. As expected, the scatter increases toward fainter magnitudes. The 1σ width of the distribution (here computed as the width around the median that contains 68% of the points) in each magnitude bin reaches $0.21 \text{ mag} = 2.5/\ln(10)/5 \text{ mag}$ (corresponding to 5σ) at 26.9 mag for HSC-*i* band and 27.4 mag for F814W, respectively. These are our obtained 5σ sensitivities. The latter is consistent with the number quoted in Koekemoer et al. (2007). Note that the scatter includes undetected point sources, for which we set their flux value to zero. Generally, we find biases of less than 0.02 mag in the measured fluxes of compact sources down to the 5σ sensitivity limit of the HSC-*i* band.

As we photometer the HSC images using an ACS prior, the detection completeness of our catalog depends on the detection on the ACS image entirely. The latter is performed by SExtractor. It is therefore straightforward to compute the overall detection completeness for point sources. Using the same simulated image as above, we find a detection

¹⁵ Note that the resulting segmentation map is a binary map as the knowledge of pixel ownership for the extracted sources is lost in the convolution step.

¹⁶ <https://www.astromatic.net/software/skymaker>

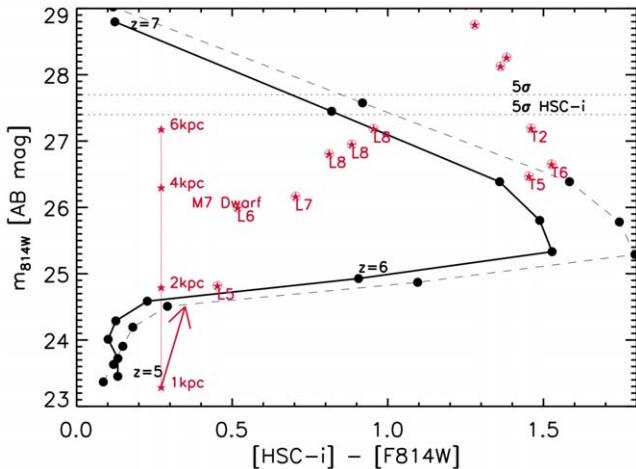


Figure 9. High-redshift objects at $z > 6$ can be selected through a color difference between photometry in the *Subaru*/Hyper-SuprimeCam *i*-band and the *Hubble*/ACS F814W filter (see also Figure 1). The lines show the expected $[\text{HSC-}i] - [\text{F814W}]$ color for an average quasar template at redshifts $z = 5 - 7$ and normalized to a B -band luminosity of $10^{45} \text{ erg s}^{-1}$ (black line and dots). The black dots are separated by $\Delta z = 0.2$. The solid black lines shows the template from Telfer et al. (2002), while the dashed line shows the template from Vanden Berk et al. (2001). For galaxies and quasars at $5.9 < z < 6.8$, the $[\text{HSC-}i] - [\text{F814W}]$ colors will be > 0.5 mag (for both templates used) mainly due to absorption by the Ly α forest blueward of 1215 \AA as well as the Ly α emission line. The red symbols denote different spectral types of brown dwarfs at their native distances from DwarfArchives.org. The red line shows the M7 dwarf from Figure 1 at distances between 1 and 6 kpc. The impact of a $A_V = 2$ mag extinction on the stellar spectra is shown by the red arrow. We discuss the contamination by stars in Section 4.5 in more detail.

completeness of 100% down to the HSC 5σ limit and a decrease to 90% at 27.5 mag in F814W.

4. Selection of Quasar Candidates

Quasars at high redshift can be identified primarily by their compactness and color. By definition, quasars should be pointlike and therefore unresolved, even in space-based observations. Furthermore, the ACS F814W filter extends redward of the *Subaru*/HSC-*i*-band filter, which allows the measurement of the flux difference across rest-frame wavelengths of 1216 \AA (Figure 1). For quasars (or any high-redshift galaxy), a red color is expected due to the drop in flux caused by the Ly α forest (absorption of ionizing photons by intervening neutral Hydrogen, see, e.g., Steidel et al. 1996; Fan et al. 2001).

In the following, we describe the selection of our high-redshift quasar candidates in more detail.

4.1. Selection by Color and Compactness

First, we perform a selection of high- z quasars from the joint catalogs, based on color and compactness.

As shown in Figure 1, the combination of the HSC-*i*-band and the ACS F814W filter allows the selection of high-redshift objects by the $[\text{HSC-}i] - [\text{F814W}]$ color.

Figure 9 shows the expected $[\text{HSC-}i] - [\text{F814W}]$ colors for the quasar template in Figure 1 as well as a template from Telfer et al. (2002) at different redshifts, normalized to a rest-frame B -band luminosity of $10^{45} \text{ erg s}^{-1}$ and with a Madau et al. (1999) intergalactic medium (IGM) absorption applied. We tested more realistic realizations of IGM absorption (e.g., Inoue et al. 2014) and found changes of less than 0.01 mag in

color, which would not change the results in the following. Note that luminosity (an unknown parameter here) does not affect the color, only the value on the y -axis.

We therefore apply an initial $[\text{HSC-}i] - [\text{F814W}]$ color cut of > 0.5 mag, which selects objects between redshifts $5.9 \leq z \leq 6.8$. Note that this selection also removes stars of spectral types earlier than L5 (even dust obscured, see Figure 9), which are spatially unresolved and therefore would pass the compactness criteria. The contamination by cooler stars will be discussed in Section 4.5. For the color selection, we include the uncertainties and limits of the magnitude measurement as well as a 3σ clearance. Specifically, we require

$$m_{\text{HSC}} - m_{\text{ACS}} > 0.5 + 3 \times \sqrt{\sigma_{\text{HSC}}^2 + \sigma_{\text{ACS}}^2}, \quad (1)$$

where σ_{HSC} and σ_{ACS} are the 1σ uncertainties on the HSC-*i* and F814W magnitudes, respectively. For HSC-detected sources, we take the 1σ flux errors output by TRACTOR. For sources detected at less than a S/N of 5, we assume the 1σ limit measured from our simulations (Section 3.2) for a S/N = 5 source. Note that galaxies or quasars at $z \geq 6$ should generally not be detected in the HSC-*i* band (they are so-called *i*-band dropouts). To be most inclusive in our initial selection, we include sources detected in HSC-*i*, which will later be removed in our final visual inspection (Section 4.2).

LBGs, as well as strong emission line galaxies at lower redshifts (such as [OIII] at $0.72 < z < 0.88$), will have similar colors and therefore could be included in this selection. We estimate that the latter would result in $[\text{HSC-}i] - [\text{F814W}]$ colors of ~ 0.2 or less, hence would be excluded by our color cut. Furthermore, LBGs and low- z galaxies would be more spatially extended than quasars given the high resolution of the HST observations (see Section 4.4). We therefore require an additional cut in ACS size.

Specifically, we select sources that are compact and unresolved in the F814W images. In the following, we use the SExtractor-derived CLASS_STAR, FLUX_RADIUS (half-light radius), and axis ratio to perform this selection. From our point-source simulations described in Section 3.2, we find that SExtractor’s star classification neural network can reliably tag point sources with CLASS_STAR > 0.9 down to a magnitude of 25.5 mag. For fainter magnitudes, CLASS_STAR > 0.5 includes point sources, but with a significant amount of contamination from extended sources. We therefore add an additional FLUX_RADIUS cut of 3 ACS pixels, which corresponds to the average PSF FWHM. In addition, we require an axis ratio (B_IMAGE/A_IMAGE) of > 0.8 , which, according to our point-source simulations, is expected for point sources brighter than 27 mag in F814W.

After applying these two selection criteria, we remove sources at the edges of the ACS coverage, which likely are spurious due to the reduced number of frames causing lower sensitivities and less-reliable cosmic ray/artifact removal. We end up with a sample of 555 sources, which are shown in Figure 10, along with the other objects in our TRACTOR catalog.

4.2. Visual Inspection

Next, we visually inspect the ACS and HSC images and the corresponding TRACTOR residuals of the 555 candidates. The main purpose is to remove obvious stars from our sample and

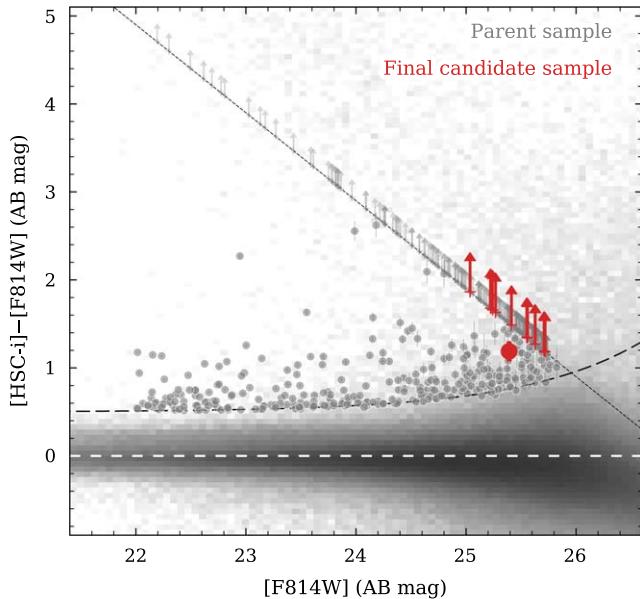


Figure 10. Initial $[HSC-i] - [F814W]$ color selection. The gray background cloud shows all extracted sources in our TRACTOR catalog, while the gray symbols show the 555 galaxies from the initial selection by color and compactness. The points denote $>5\sigma$ detections in $HSC-i$, while the arrows show upper limits. The final sample of 12 candidates is indicated in red. The long-dashed line shows Equation (1), assuming the theoretical relation between magnitude uncertainty and 1σ limit. The dotted line shows the 5σ point-source limit in $HSC-i$.

to ensure a nondetection in $HSC-i$ and bluer bands (see below). The latter is expected for quasars and galaxies at $z > 6$.

We notice that some stars with proper motion are included in our candidate selection due to an underestimated flux measurement in HSC, hence resulting in a red color. This is because we let the prior position centroid only vary by ± 0.5 HSC pixels (corresponding to about 80 mas) to improve deblending; hence the stars with proper motions of more than ~ 8 mas yr^{-1} (80 mas over the 10 yr baseline of ACS and HSC) are expected to be fit poorly by TRACTOR. Figure 11 visualizes this in the extreme case of a 23 mag star with a proper motion of roughly ~ 15 mas yr^{-1} .

To ensure a nondetection in $HSC-i$ and bluer bands, we created stacks of the ancillary data in various broad-band and narrow-band filters, which we downloaded directly from IRSA. Specifically, we create two stacks; a “blue” stack including the Subaru B , V , r , g , and $NB816$ (narrow-band) images (Taniguchi et al. 2007, 2015), and a “red” stack consisting of the Subaru z -band and the UltraVISTA Y -, J -, H -, and K -band images (McCracken et al. 2012).¹⁷ If the sources are truly at high redshifts, the blue stack should result in a nondetection. We therefore exclude candidates that have a visual detection in the blue stack as well as the $HSC-i$ filter. On the other hand, a detection in the red stack is possible and could serve as confirmation that the source is real, despite the significantly lower sensitivity in those bands. For example, a $z = 6.0$ quasar with ~ 25.5 mag in $F814W$ would be expected to have a magnitude of ~ 25 in the UltraVISTA filters and could be marginally detected in the stack. For this calculation, we used the same quasar templates as in Figure 9, which we normalized

¹⁷ The 3σ point-source sensitivities in magnitudes (taken from Laigle et al. 2016 and Capak et al. 2007) are 27.6, 26.9, 27.0, 27.0, 25.8, 26.4, 25.8, 25.4, 25.0, 25.2 for the B , V , r , g , $NB816$, z , Y , J , H , and K bands, respectively.

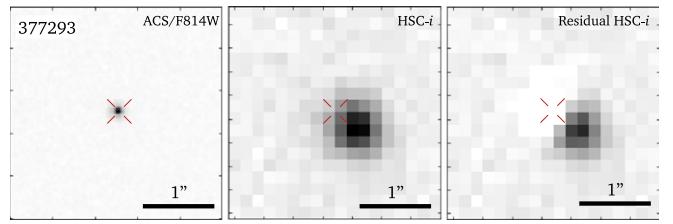


Figure 11. Example of a 23 mag star with significant proper motion (~ 15 mas yr^{-1} measured between 2004 and 2015). This star was included in our selection due to the underestimation of the $HSC-i$ flux by TRACTOR. The latter is introduced by semi-fixed prior positions.

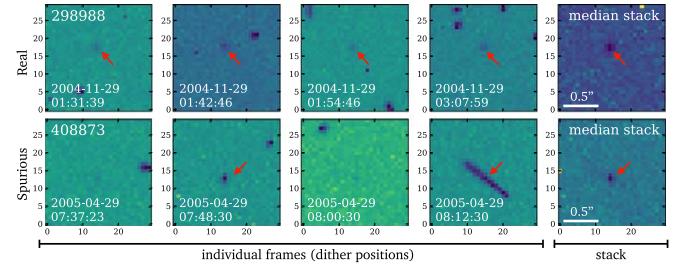


Figure 12. Example of a real source (top row) and a spurious source (cosmic ray, bottom row) on images taken in the ACS/F814W filter. The left four panels show the individual frames (i.e., dither positions, date, and time of observation are indicated). The right-most panel shows a simple median stack of the individual frames. The red arrow marks detections. We remove sources from our candidate sample if they are detected in fewer than three frames.

accordingly and then convolved with the UltraVISTA filters. We therefore do not impose any selection criteria on the red stack.

Finally, we also remove some of the obvious spurious sources (such as objects on the edges of the field or diffraction spikes of stars) and candidates with unreliable photometry where the contribution from a bright foreground source could not be accurately removed.

After this selection step, we end up with 33 candidates that obey the color cut (Equation (1)), are compact, and are undetected in the “blue” stack.

4.3. Removal of Spurious Sources (Cosmic Rays)

Our best candidates from above are only detected in ACS/ $F814W$ and therefore may include spurious artifacts in the HST images. Specifically, due to the preference of surveying a large area, each ACS/F814W pointing on COSMOS only used a 4-point dither pattern, making cosmic ray removal more challenging. To alleviate this, we require that candidates are detected in at least 3 frames (out of 4). Note that, due to the superior depth of the ACS images compared to HSC, our candidates are detected at $S/N \gtrsim 30$ in $F814W$, hence should be detected in individual frames as well at a $S/N \gtrsim 15$. Figure 12 shows an example of a real detection and a spurious detection. The latter is likely caused by cosmic ray hits on frames two and four, while the former is at consistent flux levels in all four dither positions. After the visual inspection and rejection of spurious detections, we end up with 12 final candidates in our sample.

As an additional check, we compared the sizes of our candidates to the sizes of simulated point sources. For this, we injected 30,000 point sources at various magnitudes convolved with the $F814W$ PSF into real ACS $F814W$ images and extracted them using the cataloging procedure. Figure 13

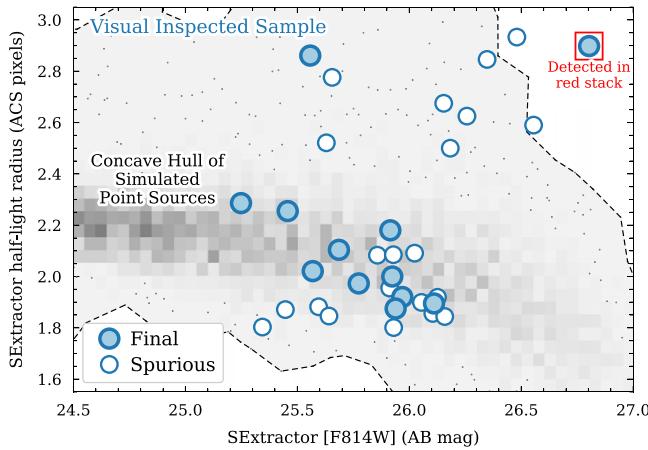


Figure 13. Visually inspected sample of 33 candidates on the `MAG_AUTO` vs. `FLUX_RADIUS` diagram. The final 12 candidates are shown in filled blue circles. Candidate 772319, detected in the red stack and with $z_{\text{phot}} = 5.92$, is highlighted with a red box. Visually identified spurious detections (Section 4.3) are shown as empty blue circles. The envelope of all simulated point sources (see text) is shown in gray. It includes our final sample of candidates thus suggesting that they are real.

shows the results of these simulations on the `FLUX_RADIUS` (half-light radius) versus `MAG_AUTO` diagram. Also shown are our 33 visually inspected candidates. At bright magnitudes, the measurements converge to the half-light radius of the PSF (~ 2.5 pixels). At fainter magnitudes, the scatter in half-light radius increases and sizes are generally underestimated due to surface brightness effects, leading to values of 2 pixels or less. Our 12 final candidates (solid blue) are consistent with the sizes of the simulated point sources within their scatter, a further indication that they are real.

Figure 14 shows the $3'' \times 3''$ cutouts in F814W and HSC-*i* including the corresponding TRACTOR residuals of the final sample of 12 candidates. We also show the blue and red stack for each source (see Section 4.2). None are detected in the blue stack by construction, and candidate 772319 is the only one detected in the red stack. This candidate (maybe slightly more extended compared to simulated point sources, see Figure 13) is also detected in the latest *COSMOS2020* catalog (Weaver et al. 2022) and has a reported photometric redshift of $z_{\text{phot}} = 5.92$. Note that the other candidates are not in the this catalog, which is based only on ground-based imaging.

Some properties of the 12 candidates are listed in Table 1.

4.4. Size Comparison to Galaxies at Similar Redshift

Quasars, outshining their host galaxies, are expected to be unresolved point sources even in observations with space-based observatories such as the HST. In Figure 15, we compare the PSF-corrected half-light radii of our candidates (red) to $z \sim 6$ galaxies from the literature (gray), which would have similar [HSC-*i*]-[F814W] colors. We report the median sizes with scatter of narrow-band selected Ly α emitters (LAEs) at $z \sim 5.5$ (Paulino-Afonso et al. 2018) and *i*-band dropouts ($z \sim 6$ LBGs Mosleh et al. 2012). Generally, both measurements (especially the one of the LAEs) suggest sizes considerably larger than for our candidates. However, it is worth noting that, while the former study measures the sizes on F814W images, the latter uses WFC3/IR F160W. Furthermore, the images (mostly taken from deep Hubble fields) may differ in depth and reduction from the COSMOS F814W observations. Therefore we also

show several samples in the COSMOS field, which allows a direct comparison to our candidates. Specifically, we show the F814W size distributions of photometric galaxies between $6 < z_{\text{phot}} < 7$ (Weaver et al. 2022), spectroscopically confirmed $z \sim 5.7$ LAEs (Shibuya et al. 2018), and spectroscopic galaxies between $5 < z_{\text{spec}} < 6$ (these include also Ly α undetected galaxies, Hasinger et al. 2018). The latter are at slightly lower redshift as we would expect for our candidates. The expected size evolution between these redshift ranges is however less than 10% (see, for example, Mosleh et al. 2012). Note that all these samples are matched in apparent magnitude to our candidates. This comparison confirms the picture that our candidates, showing sizes of ~ 0.2 kpc (assuming $z = 6$), are among the most compact objects at these redshifts. A population of low-luminosity quasars or exceptionally compact star-forming galaxies (see Section 5.3) can explain this. Note that the dust-reddened galaxies at lower redshifts (e.g., $z \sim 2$) could show red colors as well. However, such galaxies would be ruled out as they would be resolved in the Hubble images. Furthermore, even extremely dusty low- z galaxies would not be able to produce such a red color in two essentially overlapping bands. For example, to reach a color difference of 0.5 mag at $z \sim 2$, an obscuration $A_V > 17$ mag would be necessary.

4.5. Quantifying the Contamination by Stars

Late-type stars are likely the major source of contaminants as they can have red [HSC-*i*]-[F814W] colors and are also unresolved point sources in the HST images. As discussed in the previous sections and shown in Figure 9, the applied [HSC-*i*]-[F814W] color cut of 0.5 mag removes stars of spectral type warmer than L5 even for a *V*-band extinction of 2 mag—we note that this is smaller than the typical Galactic extinction in the COSMOS field. However, the same figure shows that cooler stars can enter our selection.

Removing stars by their proper motion is not possible in our case as our candidates are only detected in one single band, and no other deep space-based observations pre or post 2004 (the epoch of the ACS images) are available. At bright magnitudes, Bayesian approaches using the variation of light concentration relative to that of the PSF, have been successful at identifying stars, but at faint magnitudes, due to surface brightness limitations, these approaches are challenging to apply (Scranton et al. 2002). With the current data in hand, we can therefore only investigate the contamination of stars by quantifying their number densities as a function of spectral type and magnitude.

Figure 16 shows the expected cumulative number density of stars with spectral types cooler than L5 over the survey area in the direction of the COSMOS field as a function of F814W magnitude. The numbers are based on the volume number densities of the brown dwarfs in Kirkpatrick et al. (2021) and the absolute magnitudes in the *I* band from the PanSTARRS “Three Pi Survey” (Best et al. 2017, 2021). We converted the PanSTARRS *I*-band magnitude to ACS/F814W magnitudes for consistency by computing the color of a set of real dwarf stars in the corresponding filters (Burgasser et al. 2003, 2004, 2006; Kirkpatrick et al. 2008; Burgasser et al. 2010; Kirkpatrick et al. 2010, 2011).¹⁸ For the brown dwarfs of spectral types L5 through T5 and apparent magnitudes 18.5 through 28, we computed their surface number densities in a cone toward the COSMOS field subtending 1.64 square degree

¹⁸ See also https://roman.ipac.caltech.edu/sims/Brown_Dwarf_Spectra.html.

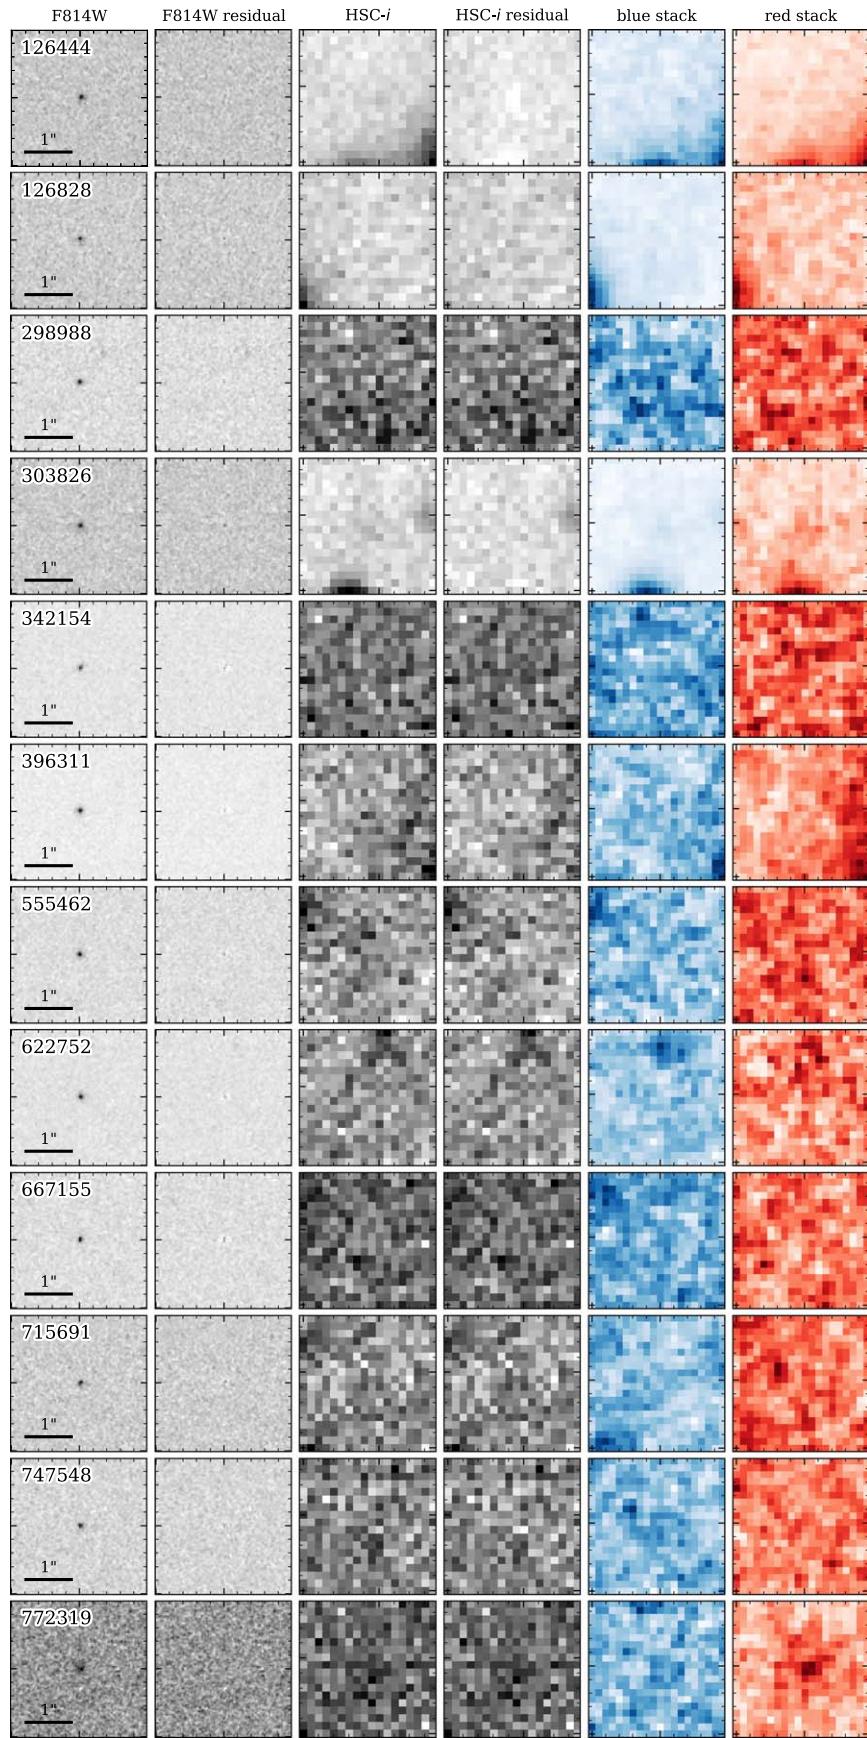


Figure 14. Cutouts of the final 12 candidates (rows). The columns show (from left to right) F814W, F814W residual, HSC-*i*, HSC-*i* residual, blue stack, and red stack. The cutouts have a size of $3'' \times 3''$, the black bar denotes $1''$. Only candidate 772319 is detected in the red stack and has a photometric redshift of $z_{\text{phot}} = 5.92$ according to the COSMOS2020 catalog. None of the candidates are detected in the blue stacks.

Table 1
Summary of Candidates

ID	R.A. (J2000)	Decl. (J2000)	[F814W] (AB mag)	[HSC- <i>i</i>]-[F814W] (AB mag) ^a	[HSC- <i>i</i>]-[F814W] (AB mag) ^b	R_e ($''$) ^c	z_{phot}	Detection in Red Stack
126444	150.6622	1.926	25.39 ± 0.02	26.58 ± 0.11	1.19 ± 0.11	0.04	...	no
126828	150.589	1.809	25.63 ± 0.03	>26.90	>1.27	0.04	...	no
298988	150.4077	2.7466	25.71 ± 0.02	>26.90	>1.19	0.03	...	no
303826	150.3563	2.7593	25.27 ± 0.03	26.92 ± 0.17	>1.63	0.07	...	no
342154	150.1967	2.0339	25.41 ± 0.02	>26.90	>1.49	0.03	...	no
396311	150.2752	2.5275	25.04 ± 0.01	27.55 ± 0.26	>1.86	0.05	...	no
555462	149.8889	2.238	25.24 ± 0.02	27.59 ± 0.25	>1.66	0.03	...	no
622752	149.7843	1.7169	25.22 ± 0.01	>26.90	>1.68	0.05	...	no
667155	149.6293	2.2474	25.63 ± 0.02	>26.90	>1.27	0.03	...	no
715691	149.5934	1.6758	25.56 ± 0.02	>26.90	>1.34	0.03	...	no
747548	149.5442	2.2755	25.55 ± 0.02	27.73 ± 0.36	>1.35	0.03	...	no
772319	149.4493	2.5234	25.72 ± 0.01	>26.90	>1.18	0.07	5.92 ^d	yes

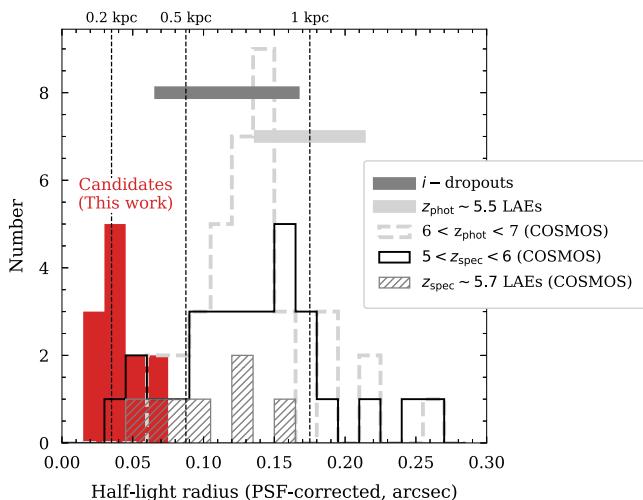
Notes.^a TRACTOR nondetections have been replaced with the 5σ limit (26.9 mag).^b Limits are reported for [HSC-*i*] measurements at $<5\sigma$ and TRACTOR nondetections.^c PSF-corrected half-light radii.^d Has match in the COSMOS2020 catalog (Weaver et al. 2022).

Figure 15. Comparison of PSF-corrected sizes of our 12 candidates (red histogram) with galaxies at similar redshifts selected with various methods (see text for details). We show median sizes and scatter of *i*-band dropouts (dark-gray horizontal line; Mosleh et al. 2012) and narrow-band selected $z \sim 5.5$ LAEs (light-gray horizontal line; Paulino-Afonso et al. 2018), galaxies selected by photometric redshifts at $6 < z_{\text{phot}} < 7$ (light-gray dashed histogram; Weaver et al. 2022), spectroscopically selected galaxies at $5 < z_{\text{spec}} < 6$ (black solid histogram; Hasinger et al. 2018), and spectroscopically confirmed LAEs at $z_{\text{spec}} \sim 5.7$ (light-gray hatched histogram; Shibuya et al. 2018). The latter three samples are magnitude matched and in the COSMOS field, hence allowing a comparison using a data set consistent with that used for our candidates. This comparison shows that our candidates are among the most compact objects at these redshifts.

and with a vertex at Earth. We previously scaled the volume number densities as a function of distance to Earth by using galactic scale height-dependent stellar density profiles of the thin and thick disk components of the galaxy (based on Reid & Majewski 1993; Buser 2000; and Binney et al. 1997, their Sections 3.6 and 10.4.3). The thick green line shows the cumulative F814W magnitude distribution for our final 12 candidates.

From this figure, we can see that the expected total number of L6-T5 dwarf stars over the surveyed area in our magnitude range is between 5 and 10 (thick black line). The largest

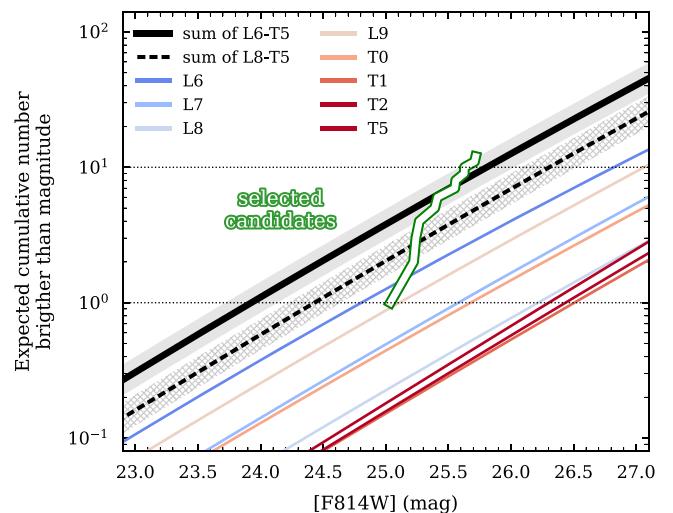


Figure 16. Expected cumulative number of brown dwarfs of different spectral types as a function of ACS/F814W magnitude. The total cumulative number of L6-T5 (L8-T6) dwarf stars is shown as a black (dashed black) line together with the uncertainty in number counts from Kirkpatrick et al. (2021). The horizontal dotted lines show numbers of 1 and 10. The green thick line shows the cumulative magnitude distribution of the 12 final candidates.

contribution comes from the warmest stars of L6 spectral type (contributing 1–3 stars). However, all our candidates have [HSC-*i*]-[F814W] colors of more than 1 magnitude (see Figure 10), which is too red for L6 and L7 dwarf stars (Figure 9). Taking this into account lowers the contamination significantly due to the lower number density of cooler stars at these magnitudes — we expect on the order of 5 ± 2 L8 to T5 stars in our sample (thick black dashed line) for the faintest magnitudes.

In a next step, we can investigate if such stars are compatible with the nondetections of our candidates in other ancillary data redward of the *i* band. Figure 17 shows a quasar spectral energy distribution (SED; Vanden Berk et al. 2001) at $z \sim 6$ as well as observed L8 and T5 star templates normalized to a magnitude of 25.5 AB in the ACS/F814W filter. The vertical bands show the available data and their 5σ sensitivity limits. Shown are the

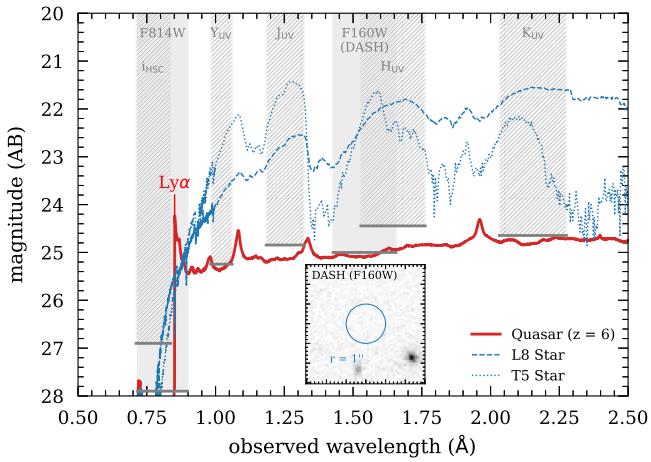


Figure 17. If our candidates would be dwarf stars (such as L8 or T5 spectral types, blue dashed and dotted lines, respectively), they should be detected in infrared data such as from the DASH HST program (5σ point-source sensitivity limits indicated by vertical bars). However, for the 7 candidates with DASH observations, we do not detect any either individually or in the stack (inset, blue circle has a radius of $1''$). The red line shows a quasar template redshifted to $z = 6$. All SEDs are normalized to 25.5 AB in F814W.

UltraVISTA Y , H , J , and K_s bands as well as the WFC3/F160W (H) band from the DASH HST program (Momcheva et al. 2017; Mowla et al. 2019). The latter is available for 7 of our candidates. The main stellar contaminants would be easily detected in the infrared bands; however, stacking the DASH F160W data for the 7 candidates quasars does not show to a significant detection (see inset). From this, we conclude that stellar contamination is very unlikely for 7 out of 12 candidates — for the remaining 5 candidates we cannot, yet, draw final conclusions.

Summarizing, the red [HSC-i]-[F814W] color cut and the absence of a detection in the DASH data suggests that stellar contamination of our sample is very unlikely — at least for 7 out of our 12 candidates. However, infrared spectroscopy, for example with James Webb Space Telescope (JWST), will be required for final confirmation.

5. Discussion

5.1. Number Density of Low-luminosity $z > 6$ Quasar Candidates

Assuming that all or a fraction of our candidates are true low-luminosity quasars at $z > 6$, we can compute their number density and compare it to estimates of the QSO luminosity function derived at higher luminosities.

We derive the number density using a V_{\max} method (Schmidt 1968). Because we do not know the redshifts of our candidates, we carried out a Monte-Carlo sampling of 5000 galaxies assuming a flat redshift prior between $6 < z < 6.7$ and a Gaussian F814W magnitude distribution (with σ equal to the standard deviation of our candidate sample). Faint galaxies with redshifts close to $z \sim 6$ in our Monte Carlo sample see a smaller volume, hence have higher $1/V_{\max}$ (i.e., number densities) associated with them (which causes an asymmetry in the distribution of number densities).

The resulting distribution is shown as red points together with the median and uncertainty in Figure 18. We find a number density of $1.4^{+0.8}_{-0.6} \times 10^{-6} \text{ Mpc}^{-3} \text{ mag}^{-1}$. The error budget includes the uncertainty from the unknown redshift distribution (as derived by our Monte-Carlo simulation), shot

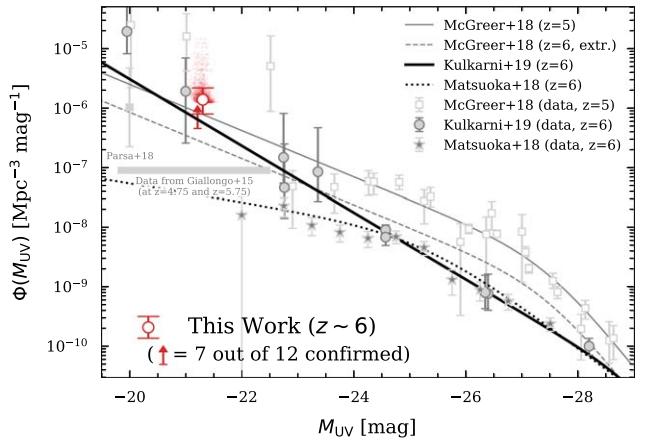


Figure 18. Constraints on the faint end of the $z \sim 6$ quasar luminosity function by our candidates (red symbols). The distribution from our V_{\max} sampling is shown as small red dots. The error bars include the uncertainty from redshift (computed by a Monte-Carlo simulation, see text) as well as shot-noise and cosmic variance ($\sim 30\%$). The red arrow shows the lower limit if 7 out of 12 candidates are stars (see Section 4.5). The lines show double power-law fits to $z = 5$ quasars (gray solid line, data as empty gray symbols) from McGreer et al. (2018), their extrapolation to $z = 6$ (gray dashed line) as well as at $z \sim 6$ from Kulkarni et al. (2019; black thick line, data as filled gray symbols) and Matsuoka et al. (2018; spectroscopically confirmed, dotted thick line, data as filled stars). Note that these fits are derived from data brighter than $M_{UV} = -22.8$. The data fainter than that (from Giallongo et al. 2015 at $z \sim 5$ (empty squares) and $z \sim 6$ (filled circles), indicated by the horizontal bar) are not used in the fits by McGreer et al. and Kulkarni et al. For $z = 5$, Parsa et al. (2018) suggest lower values (filled light-gray square).

noise, as well as the uncertainty introduced by cosmic variance. The latter is estimated by the frame work discussed in Trenti & Stiavelli (2008) using their online cosmic variance calculator.¹⁹ We assumed a redshift of $z = 6.5 \pm 1.0$, a square-area of 1.64 deg^2 , a halo filling factor of 1, and a Press-Schechter bias. We obtain an uncertainty of 30% ($\pm 5\%$ for different assumptions). While we have DASH HST observations for 7 candidates, we cannot, yet, rule out stellar contamination for the remaining 5. Therefore, we also show a lower limit ($0.5 \times 10^{-6} \text{ Mpc}^{-3} \text{ mag}^{-1}$) in the case of only 7 confirmed candidates with a red arrow.

Our estimates are consistent with the extrapolation of the McGreer et al. (2018) luminosity function at $z = 5$ (gray solid line) and the Kulkarni et al. (2019) luminosity function at $z = 6$ (black thick line, data points shown as gray filled symbols) to fainter magnitudes. On the other hand, the $z = 5$ McGreer et al. (2018) luminosity function extrapolated to $z = 6$ (gray dashed line) underestimates our number density counts by almost a factor of 10. The study from Matsuoka et al. (2018) at $z \sim 6$ shows an even lower number of low-luminosity quasars, which could be due to their spectroscopic selection function. The latter two studies are also not consistent with our lower limit (assuming 7 confirmed low-luminosity quasars). Note that these fits shown in Figure 18 are heavily constrained by the bright end of the luminosity functions ($M_{UV} \leq -22.8$ mag) and thus need to be extrapolated to the luminosities of our candidates. At fainter magnitudes ($M_{UV} > -22.8$ mag, indicated by the gray horizontal bar), Giallongo et al. (2015) provide data at $z \sim 5$ (empty squares) and $z \sim 6$ (filled circles) from the putative X-ray detections in the GOODS-S field with photometric redshifts. The measurements of that study at $z \sim 5$ are not consistent with McGreer et al., and in fact Parsa

¹⁹ <https://www.ph.unimelb.edu.au/~mtrenti/cvc/CosmicVariance.html>

et al. (2018), who reanalyzed the data from Giallongo et al. (2015), suggest that their number densities could be up to a factor of ~ 3 lower, more in agreement with a turnover at faint magnitudes of the $z = 5$ luminosity function. On the other hand, their number densities at $z \sim 6$ are consistent with Kulkarni et al. and our measurement, supporting the steep incline in number density.

5.2. Estimated Contribution to Reionization

Having placed constraints on the faint end of the QSO luminosity function, we can now calculate how much quasars contribute to the total budget of ionizing photons that led to the reionization of our universe. The answer to this question is still under debate, although it is generally thought that star-forming galaxies are the dominant contributor (e.g., Madau & Haardt 2015; Robertson et al. 2015). The most luminous quasars likely created large bubbles of ionized hydrogen around the most massive dark matter halos (e.g., Mesinger & Furlanetto 2007; Furlanetto & Mesinger 2009). This likely increased the efficiency with which ionizing photons escape from nearby star-forming galaxies resulting in the further growth of these bubbles to encompass the entire IGM. The role of low-luminosity quasars, which reside in lower-mass halos, is, however, not clear. Our measurements add an additional data point to the low-luminosity end of the quasar luminosity function, and we can make reasonable assumptions to compute the contribution of such quasars to reionization.

Rather than fitting the data in Figure 18, we adopt the Kulkarni et al. (2019) luminosity function parameterization but with a slightly different faint-end slope $\alpha = -2.35$ (instead of $\alpha = -2.40$,²⁰ as it matches our limiting case in which only 7 out of 12 candidates are confirmed low-luminosity quasars (as we cannot rule out the contribution of dwarf stars for the other 5, see Section 4.5). We find that faint-end slopes of $\alpha = 2.45 \pm 0.10$ are generally consistent with our measurements.

We integrate the luminosity function described by the (modified) Kulkarni et al. (2019) parameters (see above) and find a rest-frame 1500 Å luminosity density for these compact sources down to $M_{UV} < -20.5$ mag at $z \sim 6.4$ to be $\sim 1.1 \times 10^5 L_\odot \text{ Mpc}^{-3}$. We stress that these estimates are lower limits in the case that 7 out of our 12 candidates are actual low-luminosity quasars. Compared to other studies at $z \sim 6$, the luminosity densities obtained from our new constraints are similar to those obtained from the Kulkarni et al. (2019) and Giallongo et al. (2015) luminosity function but more than an order of magnitude higher compared to Matsuoka et al. (2018; see also Section 5.1). As reviewed in Chary et al. (2016), the star-forming galaxy population produces a rest-frame UV luminosity density of $\sim 10^8 L_\odot \text{ Mpc}^{-3}$, which is almost three orders of magnitude larger than the luminosity density estimated in this work. If these galaxies indeed show strong nebular line emission as inferred from their multi-wavelength SEDs, it would imply an ionizing photon production rate that is a factor of 10 higher than those among star-forming galaxies in the local universe. Even with a modest escape fraction of 10%, such galaxies can easily reionize the IGM by $z \sim 6$. If the Lyman-continuum production rate in the

QSO hosts is similar, as seems to be the case based on their inferred nebular emission line properties (B. Lee & R. Chary, private communication), and the Lyman-continuum escape of quasars is not luminosity dependent (as suggested by Iwata et al. 2022), it implies that the faint end of the QSO luminosity function has a negligible impact on reionization. However, this has the indirect impact of an increased escape fraction of ionizing photons from the star-forming galaxies that are within the Strömgren sphere of a given QSO (Furlanetto et al. 2006; Mesinger & Furlanetto 2007; Furlanetto & Mesinger 2009). We therefore conclude that even low-luminosity quasars such as the ones studied here, which are more numerous but several orders of magnitude fainter than the luminous quasars that have been well characterized in earlier studies, play a relatively small role in reionization. However, low-luminosity quasars can indirectly increase the escape fraction of ionizing photons from star-forming galaxies that are within the Strömgren sphere of a given QSO. This may enhance the role of nearby star-forming galaxies depending on their clustering strength.

5.3. Alternative Interpretation: Unusually Compact Star-forming Galaxies?

Current estimates of the UV luminosity function (e.g., Bouwens et al. 2015) indicate ~ 500 star-forming $z \sim 6$ galaxies in the same area and at the same magnitude as our candidates. As shown in Figure 15, most of these would be resolved in Hubble imaging. However, our candidates, as they are unresolved, could be ultra compact star-forming galaxies. In that case, it would be interesting to estimate the star formation surface density to put them in context with other galaxies at $z \sim 6$.

For the estimation of the SFR, the contamination of Ly α to the flux in the F814W filter has to be taken into account. From the F814W magnitudes, we estimate a UV-based SFR for our candidate sample of $70^{+18}_{-10} M_\odot \text{ yr}^{-1}$ using the relation between UV luminosity and SFR in Kennicutt (1998a). This estimate includes the likely contamination by Ly α in the F814W filter, which depends on the EW of the line. We estimate the fraction of flux from Ly α by creating an intrinsic spectrum including Ly α of different EWs, then applying IGM absorption, and finally convolving it with the F814W filter transmission. Assuming an intrinsic EW of 500 Å for Ly α (which is rather a conservative upper limit), we expect Ly α to contribute to about half of the measured UV flux. A more common EW for star-forming galaxies at these redshifts of < 50 Å (see Schenker et al. 2014) results only in $< 10\%$ contamination.

Together with a measurement of a physical size (radius) of our candidates (~ 0.2 kpc, see Figure 15), we estimate a SFR surface density of $\Sigma_{\text{SFR}} = 554^{+139}_{-79} M_\odot \text{ yr}^{-1} \text{ kpc}^{-2}$. Including the contamination by a Ly α line with EW = 500 Å, this number can be up to a factor 2 smaller. However, as mentioned above, this is a very conservative estimate. Also, note that this estimate describes the unobscured star formation. A moderate dust attenuation of $E(B-V) = 0.1$ would result in a total SFR density that is a factor ~ 3 higher. The most likely lower limit for the SFR surface density is therefore likely around $500 M_\odot \text{ yr}^{-1} \text{ kpc}^{-2}$.

We can compare this to the value for typical main-sequence galaxies at the same redshift. From Figure 15, typical high-z galaxies have a size of around 0.7 kpc. The average main-

²⁰ Our limit and scarcity of measurements at $z \sim 6$ make the parameterization of the luminosity function rather uncertain. The adjustment of the faint-end slope is therefore done by eye (to match our lower limit) without the use of more sophisticated fitting methods.

sequence SFR (total dust corrected) of a typical $10^9 M_{\odot}$ galaxy²¹ at $z \sim 6$ is at $10\text{--}100 M_{\odot} \text{yr}^{-1}$ (e.g., Speagle et al. 2014; Faisst et al. 2020; Schaefer et al. 2020). This results in a SFR surface density of around $5\text{--}70 M_{\odot} \text{yr}^{-1} \text{kpc}^{-2}$, which is well consistent with the gas mass estimated with Atacama Large Millimeter/submillimeter Array (ALMA) at these redshifts (Dessauges-Zavadsky et al. 2020) assuming the Kennicutt-Schmidt relation (Kennicutt 1998b; Kennicutt & De Los Reyes 2021).

This estimate indicates that our candidates, if they are compact star-forming galaxies, show $7\text{--}100 \times$ higher SFR surface densities than typical galaxies at the same redshifts, hence would themselves be interesting targets to follow up spectroscopically.

6. Conclusions

In this work, we demonstrate the power of a joint pixel-by-pixel analysis of different data sets with different depth and PSF properties. The approach of prior-based photometry mitigates issues of blending and provides robust limits for nondetections. This work is an important stepping stone toward identifying challenges in the analysis steps, and paves the way for applying this method to future data sets such as from the Rubin, Roman, and Euclid missions. A pixel-by-pixel joint analysis of these data sets will result in precision photometric properties and enhance the science output of the individual missions. The Joint Survey Processing (JSP) initiative at Caltech/IPAC is working toward building the techniques and infrastructure to carry out such an analysis in the future.

In this work, we have applied JSP to a unique combination of a space- and ground-based set of filters in the COSMOS field, resembling in resolution and depth the future imaging data from Rubin and Roman/Euclid. We use this technique to place new constraints on the number density of low-luminosity ($M_{UV} \sim -21$ mag) quasars.

We perform a pixel-by-pixel analysis of the space-based ACS/F814W filter and ground-based HSC-*i*-band filter using TRACTOR. Specifically,

1. we use the fact that the ACS/F814W filter extends redward of the HSC-*i*-band filter, which allows us to select galaxies at $z \gtrsim 6$ using the [HSC-*i*]-[F814W] colors;
2. we leverage the high spatial resolution of the space-based F814W filter to select point sources (as expected for quasars).

With this technique, we are able to identify 12 robust candidate sources. Their [HSC-*i*]-[F814W] colors of >1 mag suggest them to be between $5.9 < z < 6.7$. Their unresolved nature in Hubble imaging and average PSF-corrected half-light radius of $0''.03$ (~ 0.2 kpc at $z = 6$) confirms the compactness expected for quasars. Number density estimates of L and T dwarf stars identify them as the main likely contaminants (up to 50%). However, deep DASH HST observations argue against stellar contamination for 7 of the 12 candidates. Taking the potential stellar contamination into account, we estimate a number density of $1.4^{+0.8}_{-0.6} \times 10^{-6} \text{ Mpc}^{-3} \text{ mag}^{-1}$ at $M_{UV} \sim -21.2$. If only 7 out of 12 candidates are true low-luminosity quasars, we find a 1σ limiting number density of

$0.5 \times 10^{-6} \text{ Mpc}^{-3} \text{ mag}^{-1}$. Both numbers may be underestimates. Given the unknown redshift of our candidates, we estimate a potential underestimation of the number densities of a factor 3–4. Furthermore, the quasar host could dominate the light at these low quasar luminosities, resulting in an extended object (Bowler et al. 2021). Such configurations would be missed as we limit our selection to point sources.

The low number densities of these sources (even including the aforementioned caveats) imply a contribution to reionization that is about three orders of magnitude below those of the star-forming galaxy population. However, low-luminosity quasars may enhance the escape of ionizing photons from neighboring star-forming galaxies and would therefore be interesting environments to study the growth of ionizing bubbles through next-generation spectroscopy.

Alternatively, our candidates could be compact star-forming galaxies at $z \sim 6$. If true, their SFR surface densities can be estimated to be about $500 M_{\odot} \text{yr}^{-1} \text{kpc}^{-2}$, which is $7\text{--}100 \times$ higher than a typical main-sequence galaxy at these redshifts at $10^9 M_{\odot}$. This would suggest an interesting but rare population of compact star-forming galaxies, which future spectroscopy with JWST will help characterize.

We thank E. Merlin for providing an updated TPHOT version and technical support, and the anonymous referee for a constructive feedback on our manuscript. This research is partially funded by the Joint Survey Processing effort at IPAC/Caltech through NASA grant NNN12AA01C. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

The Hyper Suprime-Cam (HSC) collaboration includes the astronomical communities of Japan and Taiwan, and Princeton University. The HSC instrumentation and software were developed by the National Astronomical Observatory of Japan (NAOJ), the Kavli Institute for the Physics and Mathematics of the Universe (Kavli IPMU), the University of Tokyo, the High Energy Accelerator Research Organization (KEK), the Academia Sinica Institute for Astronomy and Astrophysics in Taiwan (ASIAA), and Princeton University. Funding was contributed by the Faint Images of the Radio Sky at Twenty cm (FIRST) program from Japanese Cabinet Office, the Ministry of Education, Culture, Sports, Science and Technology (MEXT), the Japan Society for the Promotion of Science (JSPS), Japan Science and Technology Agency (JST), the Toray Science Foundation, NAOJ, Kavli IPMU, KEK, ASIAA, and Princeton University.

This paper makes use of software developed for the Large Synoptic Survey Telescope. We thank the LSST Project for making their code available as free software at <http://dm.lsst.org>.

Based in part on data collected at the Subaru Telescope and retrieved from the HSC data archive system, which is operated by Subaru Telescope and Astronomy Data Center at National Astronomical Observatory of Japan.

Software: ASTROPY 3 (Astropy Collaboration et al. 2013, 2018), TRACTOR (Lang et al. 2016a, 2016b; Weaver et al. 2022), PHOTUTILS (Bradley et al. 2019), SEXTRACTOR (Bertin & Arnouts 1996), PSFEX (Bertin 2011), TPHOT (Merlin et al. 2015, 2016), SKYMAKER (Bertin 2009).

²¹ Since we do not detect most candidates in near-IR bands such as UltraVISTA, we can set a conservative upper limit on stellar mass.

Appendix A

Demonstration of Deblending with TRACTOR

Blending is a significant concern in confusion-limited imaging. The left panel of Figure 19 shows the fraction of blended sources as a function of optical i -band brightness in deep, space-based, seeing-limited CANDELS fields (GOODS-N and GOODS-S, Grogin et al. 2011; Koekemoer et al. 2011). Blending was estimated by assuming that the seeing was $0''.8$ FWHM, and if two sources are closer than $1''.6$, then the photometry of each is biased by more than 5% due to overlapping isophotal areas (assuming Gaussians). The CANDELS catalogs are complete to 26.5 mag as shown, so to measure the blending fraction at fainter magnitudes (dotted-dashed steps), the source counts were extrapolated using a power-law fit to the values at brighter magnitudes. These faint sources were distributed randomly across the sky and the same criterion used to identify the blended fraction. As can be seen in the left panel of Figure 19, almost 50% of sources at Rubin/LSST deep survey depths (i -band 5σ magnitude limit over 10 yr of 26.8 mag, Ivezić et al. 2019) will be confused and prior-based source fitting has to be undertaken to obtain more robust photometry. In reality the sources are extended; hence the above estimate is a lower limit. The general benefit of prior-based photometry over simple catalog matching or aperture photometry is suggested in Weaver et al. (2022; see also Nyland et al. 2017). Specifically, their study shows an increase in the precision of $\sim 20\%$ and slight decrease in the outlier fraction of the measurement of photometric redshifts (see their Figure 15). However, we note that the TRACTOR analysis in that study makes use of ground-based priors, which themselves are prone to blending and confusion. The benefit of prior-based photometry is therefore suppressed. A better demonstration is provided in Lee et al. (2012), comparing the HST/F850LP – Spitzer/3.6 μm colors derived from SExtractor (aperture-based) and TFIT (the precursor of TPHOT using space-based priors) of simulated galaxies. Their figure shows a significant

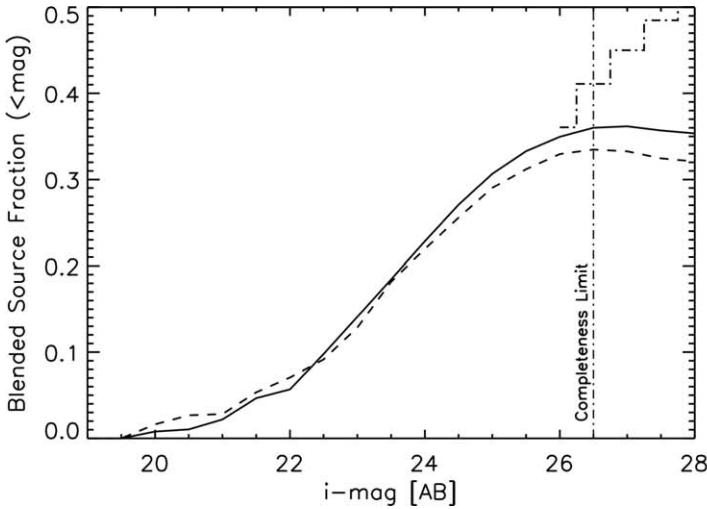
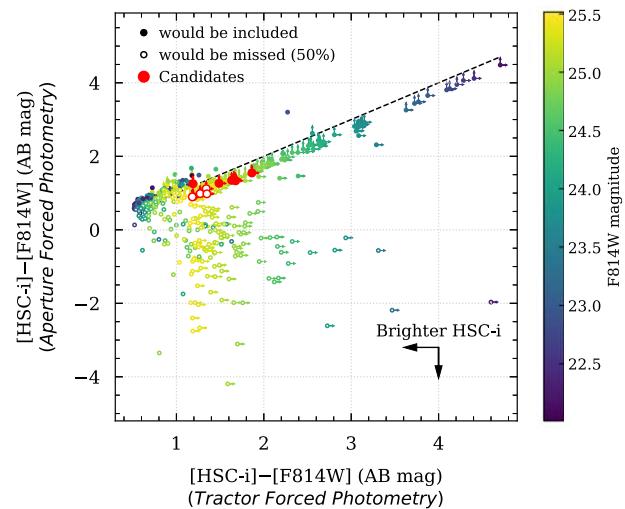


Figure 19. Left: fraction of galaxies blended with others as a function of optical i -band brightness measured in deep, seeing-limited extragalactic surveys (GOODS-N, solid line, and GOODS-S, dashed line). These catalogs are complete to 26.5 mag (vertical line). Numbers beyond that (dotted-dashed steps) are found by extrapolation using a power-law fit to the values at brighter magnitudes (see text). At our magnitude limit, we expect $>30\%$ of the sources to be blended. Right: comparison of $[HSC-i] - [F814W]$ colors of our 555 initial candidates (see Section 4) measured from TRACTOR and aperture photometry (see text). The color of each symbol indicates its $F814W$ magnitude. While the colors of bright objects are consistent between the two methods, the colors (and limits) computed via aperture photometry for fainter objects are significantly bluer in some cases. This is due to flux of neighboring sources bleeding into the aperture around the main target as well as blending of objects within the PSF FWHM. About half of the initial candidates (and 4 of our final 12 candidates) would have been missed. This shows the need for prior-based photometry methods (such as TRACTOR or TPHOT) that allow for the fitting and subtracting of neighboring sources to provide more robust flux measurements and limits.

improvement in the measurement of the colors measured with space-based priors.

In the case of our 5σ HSC- i -band sensitivity limit, we would expect a $>5\%$ photometry bias in more than 30% of the sources (see left panel of Figure 19). The use of TRACTOR (or a similar software such as TPHOT, see Appendix B) allows the mitigation of these effects to provide more robust photometry by using high-resolution prior images and fitting the neighboring sources simultaneously. Performing a prior-based measurement of photometry is especially important in our case, as we are seeking sources that are fainter (or even undetected) on the low-resolution images.

To demonstrate the impact of blending statistically on our results, we compare the $[HSC-i] - [F148W]$ colors obtained by our method with simple aperture photometry. The latter would be the choice to search for nondetections — by measuring the flux on the low-resolution image in an aperture around a centroid from a high-resolution image. In detail, we measure the fluxes in $0''.8$ diameter apertures on the HSC image at the position of the source detected on the ACS image (thereby we apply the necessary astrometric shifts). We apply an aperture correction, which is calculated by computing the percent of flux of a PSF enclosed within the same diameter. The local background as well as error are calculated from σ -clipped pixels in an annulus around the source. The right panel of Figure 19 shows the results of this comparison for the 555 initial candidate galaxies that are already color selected (see also Figure 10). Generally, sources show aperture-based $[HSC-i] - [F814W]$ colors lower than the TRACTOR-based colors. This is because of flux from neighboring sources *bleeding* into the aperture and causing an overestimation of aperture flux (hence bluer color). For isolated and bright sources, both methods return a similar result (within 0.2 mag). If aperture-based colors were used, we would miss 50% of potential candidates (open symbols). In addition, four out of our twelve final candidates (red) would have been missed.



Appendix B Comparison to TPHOT

We used TRACTOR to measure photometry as it allows proper deblending by using high-resolution positional and shape priors. The use of TRACTOR with simple parametric models (as we do here) has the disadvantage that the models often do not represent the true shapes of galaxies well. This is expected to happen for bright extended galaxies with significant additional substructure or lumpy galaxies, for example at high redshifts. In principle, more complicated models can be used at a higher computational cost. TPHOT (Merlin et al. 2015, 2016) offers an alternative tool to perform forced-photometry measurements. Since it uses nonparametric structural models, we expect it to perform better for bright and extended galaxies.

In brief, TPHOT uses the high-resolution image directly by creating cutouts of the galaxies, which are then convolved by a kernel and scaled in flux to fit the galaxies on the low-resolution image. The kernel K (on high-resolution pixel scale) has to be constructed using the low-resolution (lr) and high-resolution (hr) PSF such that $\text{PSF}_{\text{lr}} = K \otimes \text{PSF}_{\text{hr}}$. This allows TPHOT to use the exact shape of galaxies as priors in a nonparametric way.

We ran TPHOT on each of the $3' \times 3'$ sub-patches. We first identified sources on the high-resolution ACS/F814W image using SExtractor. To account for the different PSF sizes, TPHOT creates a dilated segmentation map based on the segmentation map created by SExtractor. We chose the following parameters: $\text{minarea_dilate} = 2.0$, $\text{maxarea_dilate} = 10000.0$, $\text{minfactor_dilate} = 4.0$, $\text{maxfactor_dilate} = 12.0$, $\text{dilation_factor} = 2.0$, $\text{dilation_threshold} = 0.0$. Because TPHOT requires the pixel scale of the low-resolution image ($0''.168/\text{px}$) to be an integer multiple of the pixel scale of the high-resolution image, we resized the ACS images from $0''.03/\text{px}$ to $0''.028/\text{px}$. To create the convolution kernel, we took the following steps. First, we resized the HSC PSF to the new pixel scale of the

high-resolution image using the function `resize_psf` from the PhotUtils Python package.²² We made sure that the resulting PSF is properly normalized. Second, we used the function `create_matching_kernel` (same package) to create the convolution kernel. For the window function, we found that a top-hat filter with an inner diameter of 0.4 works best. The convolution kernel is verified by convolving the ACS PSF and comparing its profile to the HSC PSF.

Figure 20 compares the residuals from TPHOT (middle) and TRACTOR (right) for two different $3''$ sub-patches. The original HSC-*i*-band image is shown on the left. While fainter point sources are fit equally well by both codes, substantial differences in the residuals can be seen for brighter extended galaxies. Specifically, the TPHOT residuals are more symmetric while the TRACTOR residuals show a *butterfly* pattern. This is because the former uses the actual shape of the ACS image while the latter is not able to capture the same amount of details in a parametric fit to the galaxies (for example a pronounced bulge + disk or spiral pattern).

Figure 21 shows this more quantitatively. For each measured source, we computed the per-pixel residual in the same aperture (defined by the segmentation map) on the residual maps produced by the different codes. The figure shows that the residuals left by TRACTOR are larger at brighter magnitude where more extended galaxies contribute. However, at fainter magnitude (dominated by compact sources), both codes perform equally well.

Summarizing, both codes have advantages and disadvantages. While TPHOT shows superior performance by reducing residuals for bright extended galaxies, both codes seem to perform equally well for compact fainter sources (>23 mag). However, we found that TRACTOR is more versatile. Importantly, TRACTOR has no constraints on pixel scale or size of the images, and no convolution kernel has to be created (which needs some tinkering). Furthermore, image artifacts such as diffraction spikes in the vicinity of bright stars affect the functionality of the codes and require dedicated corrections.

²² <https://photutils.readthedocs.io/en/stable/>

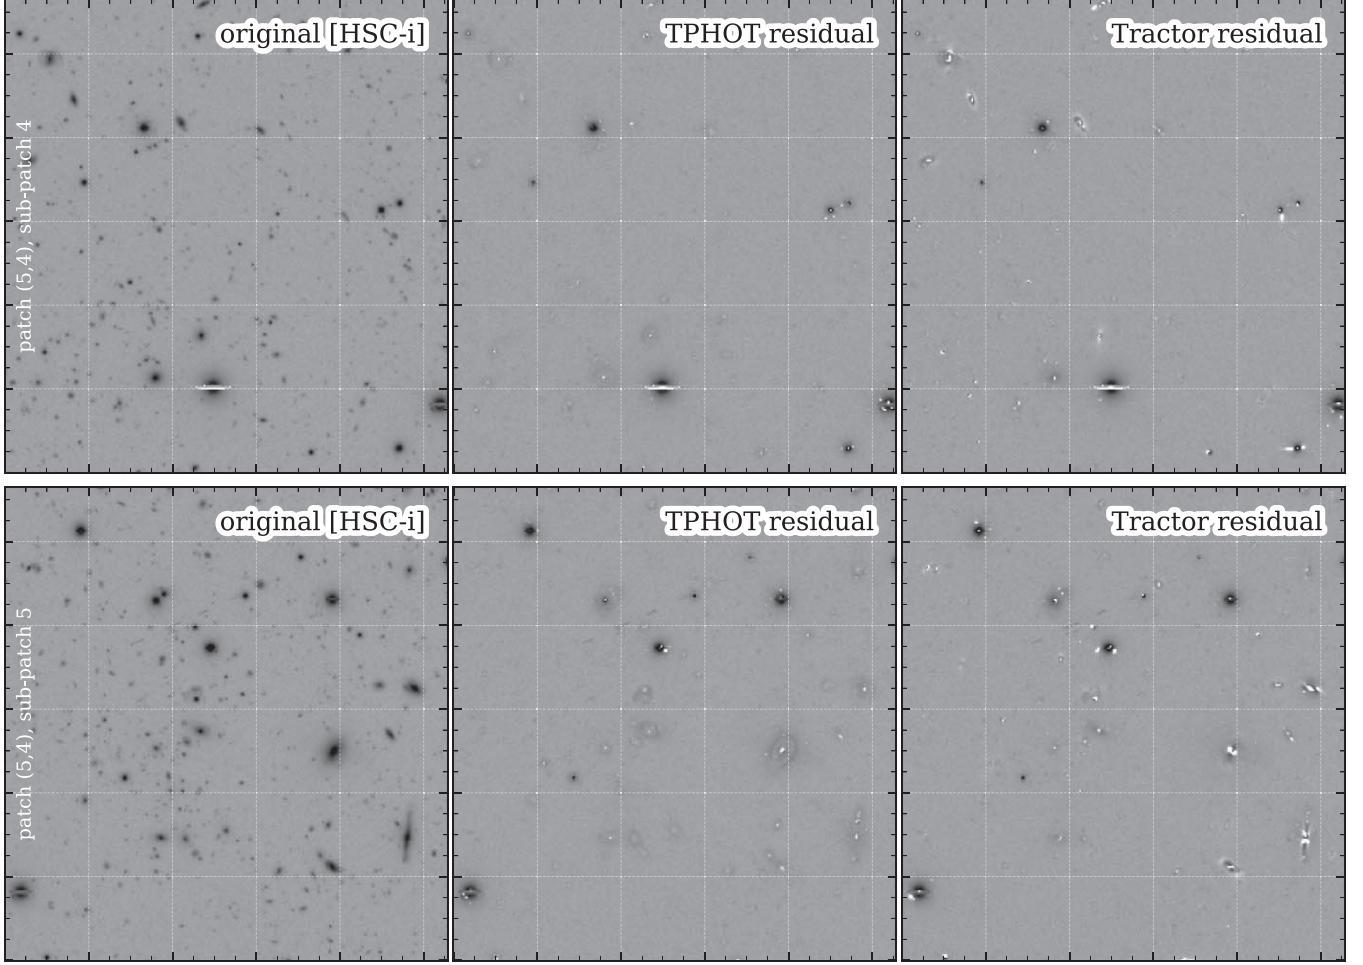


Figure 20. Comparison of TRACTOR (right) and TPHOT (middle) residuals of two different sub-patches (top and bottom row). The original HSC-*i*-band image is shown on the left. The scaling is the same in all images.

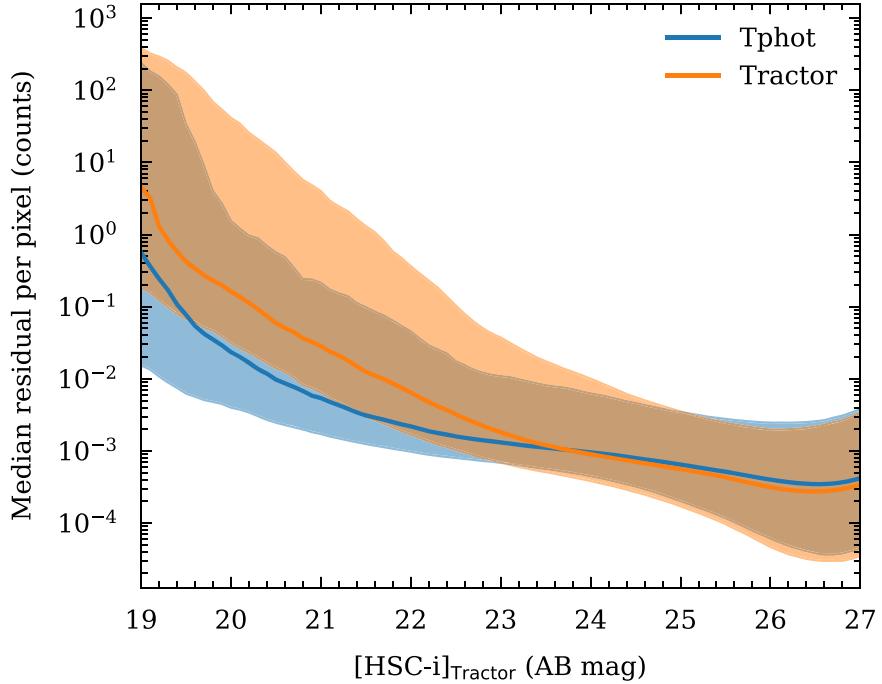


Figure 21. Comparison of TRACTOR (orange) and TPHOT (blue) per-pixel residuals as a function of [HSC-*i*] magnitude. The residuals of both photometry fitting methods are similar at faint magnitudes (preferentially point sources), while TPHOT is doing better for bright and extended galaxies due to its nonparametric nature. The shaded area shows the 1σ range around the median (solid lines).

Appendix C

Details on the Parallelization at NERSC

An efficient use of supercomputing facilities will be crucial when undertaking the joint pixel-level analysis of 1000 s of square degrees of observations. A single *sub-patch*, $3' \times 3'$ in size, with about 1500 sources, requires about 20–50 minutes to run on a high-end 2019 laptop or desktop computer. Scaling this up to something as small as the COSMOS field would require about two weeks to one month of computing time on a single desktop machine for a single pair of wavelengths. Joint pixel analysis is however intrinsically parallelizable. We made use of the supercomputer *Cori* at the NERSC²³ to execute our pipeline.

The pipeline, written in Python, was containerized using DOCKER²⁴ and then deployed on *Cori* using SHIFTER,²⁵ which is a custom software containerization solution, developed by NERSC. SHIFTER transforms standard DOCKER images into a custom format, which can be used to automatically launch containers on the nodes of the supercomputer.

To deploy the pipeline on multiple nodes, we made use of a SLURM²⁶ batch script, as shown below.

```
#!/bin/bash
#SBATCH --job-name=JSP
#SBATCH --account=xxx
#SBATCH --time=48:00:00
#SBATCH --constraint=haswell
#SBATCH --qos=regular
#SBATCH --mail-type=BEGIN,END,FAIL,TIME_LIMIT,
    TIME_LIMIT_80
#SBATCH --image=docker:nrstickley/jsp_apps:2019-12-20
#SBATCH --nodes 1
#SBATCH --exclusive
#SBATCH --array=0-27
#SBATCH --output=%x_%A_%a.out
shifter /bin/bash run_jobs.sh ${SLURM_ARRAY_TASK_ID}
${SLURM_CPUS_ON_NODE}
```

This allocates an array of 28 computing nodes and launches the script `run_jobs.sh` in a SHIFTER container on each. Each node is labeled with an integer ID number from 0 to 27, which is stored in the `SLURM_ARRAY_TASK_ID` variable. The number of logical CPU cores per node (64, in the case of Cori) is passed to the script using the variable, `SLURM_CPUS_ON_NODE`. The total time allocation per node was set to the maximum of 48 hr.

The most compute-intensive component of the pipeline, TRACTOR, is multithreaded but scales poorly with core count. Running multiple instances of TRACTOR on each node in parallel improves efficiency, but the memory consumption becomes the limiting factor. The maximum memory needed by a single task was slightly less than 16 GB. Since each Cori compute node contained 128 GB of system memory, we were able to run 8 simultaneous tasks per node. This was done using the script `run_jobs.sh`, which used GNU PARALLEL (Tange 2015) to efficiently distribute multiple simultaneous instances of our code across the CPU cores of each individual compute node.

²³ <https://www.nersc.gov/>

²⁴ <https://www.docker.com/>

²⁵ <https://www.nersc.gov/research-and-development/user-defined-images/>

²⁶ <https://slurm.schedmd.com/>

```
#!/bin/bash
# run_jobs.sh runs multiple jobs on a single node.
NODE_ID=$1 # node ID, which is the SLURM_ARRAY_TASK_ID
SLURM_CPUS_ON_NODE=$2 # the number of logical CPUs on
the node.
# define number of simultaneous jobs
njobs=8
cpus_per_job=$((SLURM_CPUS_ON_NODE/n_jobs))
echo Task ID is: ${SLURM_ARRAY_TASK_ID}
ls /path/${NODE_ID}/*.json | parallel -j ${njobs} ./task_runner.sh ${cpus_per_job}{%} {}
```

The `run_jobs.sh` script operates on a list of JSON-formatted job definition files, each of which contains all of the input parameters that our pipeline (Section 3.1) needs to run a specific *sub-patch*. There are 1008 job files, which we have distributed into 28 directories (named 0 to 27 corresponding to the node ID, `SLURM_ARRAY_TASK_ID`). GNU PARALLEL dynamically schedules the tasks into 8 execution “slots” so that as soon as one task completes, another one is launched in the same slot. This continues until the full list of tasks assigned to the node has been executed.

The `run_jobs.sh` script, launches an instance of `task_runner.sh` for each job assigned to the node. The script `task_runner.sh` handles the details of running an individual task on a set of CPU cores.

```
#!/bin/bash
# task_runner.sh distributes the work over 8 physical CPU
cores on a single NUMA node.
cores_per_task=$1 # the number of logical cores per task
slot=$2 # The execution slot; an integer from 1 to 8
filename=$3 # JSON job file for one 3' x 3' sub-patch
# This function determines which CPU threads correspond
to the specified execution slot:
cpu_threads() {
minus=$((slot - 1))
first_cpu_thread=$((cores_per_task * minus))
final_cpu_thread=$((first_cpu_thread + cores_per_task
- 1))
threads=$(seq $first_cpu_thread $final_cpu_thread | awk
-vORS=, 'print $1')
echo $threads:::-1
}
taskset -c $(cpu_threads) python runTractor.py \
{filename}
```

The utility `taskset` specifies which specific logical CPU cores will be used to execute the `runTractor.py` script. This forces all of the threads of the task to remain on the same set CPU cores for the duration of their lifetime, rather than migrating to other cores over time, which improves the cache utilization and reduces the memory access latency. For the first execution slot, these would be logical cores 0–7. For the second slot, 8–15, and so on.

In summary, by distributing the jobs over 28 compute nodes, with each node running 8 jobs simultaneously, we were able to complete all 1008 tasks in slightly less than 24 hr. We plan to implement additional optimizations in the future. In particular, it is possible to run more tasks simultaneously on a single compute node by more intelligently scheduling the tasks to avoid running out of memory.

ORCID iDs

Andreas L. Faisst <https://orcid.org/0000-0002-9382-9832>
Ranga Ram Chary <https://orcid.org/0000-0001-7583-0621>

- Sergio Fajardo-Acosta [ID](https://orcid.org/0000-0001-9309-0102) <https://orcid.org/0000-0001-9309-0102>
- Roberta Paladini [ID](https://orcid.org/0000-0002-5158-243X) <https://orcid.org/0000-0002-5158-243X>
- Benjamin Rusholme [ID](https://orcid.org/0000-0001-7648-4142) <https://orcid.org/0000-0001-7648-4142>
- Nathaniel Stickley [ID](https://orcid.org/0000-0003-0987-5738) <https://orcid.org/0000-0003-0987-5738>
- George Helou [ID](https://orcid.org/0000-0003-3367-3415) <https://orcid.org/0000-0003-3367-3415>
- John R. Weaver [ID](https://orcid.org/0000-0003-1614-196X) <https://orcid.org/0000-0003-1614-196X>
- Gabriel Brammer [ID](https://orcid.org/0000-0003-2680-005X) <https://orcid.org/0000-0003-2680-005X>
- Anton M. Koekemoer [ID](https://orcid.org/0000-0002-6610-2048) <https://orcid.org/0000-0002-6610-2048>
- Hironao Miyatake [ID](https://orcid.org/0000-0001-7964-9766) <https://orcid.org/0000-0001-7964-9766>

References

- Aihara, H., Armstrong, R., Bickerton, S., et al. 2018, *PASJ*, **70**, S8
- Akiyama, M., He, W., Ikeda, H., et al. 2018, *PASJ*, **70**, S34
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
- Bañados, E., Venemans, B. P., Decarli, R., et al. 2016, *ApJS*, **227**, 11
- Bañados, E., Venemans, B. P., Mazzucchelli, C., et al. 2018, *Natur*, **553**, 473
- Becker, R. H., Fan, X., White, R. L., et al. 2001, *AJ*, **122**, 2850
- Bertin, E. 2009, *MmSAI*, **80**, 422
- Bertin, E. 2011, in ASP Conf. Ser. 442, Astronomical Data Analysis Software and Systems XX, ed. I. N. Evans et al. (San Francisco, CA: ASP), 435
- Bertin, E., & Arnouts, S. 1996, *A&AS*, **117**, 393
- Best, W. M. J., Liu, M. C., Magnier, E. A., & Dupuy, T. J. 2021, *AJ*, **161**, 42
- Best, W. M. J., Magnier, E. A., Liu, M. C., Aller, K. M., & Zhang, Z. 2017, *AAS*, **229**, 240.01
- Binney, J., Gerhard, O., & Spergel, D. 1997, *MNRAS*, **288**, 365
- Bouwens, R. J., Illingworth, G. D., Oesch, P. A., et al. 2015, *ApJ*, **803**, 34
- Bowler, R. A. A., Adams, N. J., Jarvis, M. J., & Häußler, B. 2021, *MNRAS*, **502**, 662
- Bradley, L., Sipőcz, B., Robitaille, T., et al. 2019, astropy/photutils: v0.7.2, v0.7.2, Zenodo, doi:[10.5281/zenodo.3568287](https://doi.org/10.5281/zenodo.3568287)
- Burgasser, A. J., Burrows, A., & Kirkpatrick, J. D. 2006, *ApJ*, **639**, 1095
- Burgasser, A. J., Cruz, K. L., Cushing, M., et al. 2010, *ApJ*, **710**, 1142
- Burgasser, A. J., Kirkpatrick, J. D., Liebert, J., & Burrows, A. 2003, *ApJ*, **594**, 510
- Burgasser, A. J., McElwain, M. W., Kirkpatrick, J. D., et al. 2004, *AJ*, **127**, 2856
- Buser, R. 2000, *Sci*, **287**, 69
- Capak, P., Aussel, H., Ajiki, M., et al. 2007, *ApJS*, **172**, 99
- Chabrier, G. 2003, *PASP*, **115**, 763
- Chary, R., Petitjean, P., Robertson, B., Trenti, M., & Vangioni, E. 2016, *SSRv*, **202**, 181
- Civano, F., Marchesi, S., Comastri, A., et al. 2016, *ApJ*, **819**, 62
- Dessauges-Zavadsky, M., Ginolfi, M., Pozzi, F., et al. 2020, *A&A*, **643**, A5
- Dey, A., Schlegel, D. J., Lang, D., et al. 2019, *AJ*, **157**, 168
- Faisst, A. L., Schaerer, D., Lemaux, B. C., et al. 2020, *ApJS*, **247**, 61
- Fajardo-Acosta, S. B., Kirkpatrick, J. D., Schneider, A. C., et al. 2016, *ApJ*, **832**, 62
- Fajardo-Acosta, S., et al. 2022, arXiv:[2203.05658](https://arxiv.org/abs/2203.05658)
- Fan, X., Carilli, C. L., & Keating, B. 2006, *ARA&A*, **44**, 415
- Fan, X., Narayanan, V. K., Lupton, R. H., et al. 2001, *AJ*, **122**, 2833
- Furlanetto, S. R., McQuinn, M., & Hernquist, L. 2006, *MNRAS*, **365**, 115
- Furlanetto, S. R., & Mesinger, A. 2009, *MNRAS*, **394**, 1667
- Giallongo, E., Grazian, A., Fiore, F., et al. 2015, *A&A*, **578**, A83
- Grogin, N. A., Kocevski, D. D., Faber, S. M., et al. 2011, *ApJS*, **197**, 35
- Hasinger, G., Capak, P., Salvato, M., et al. 2018, *ApJ*, **858**, 77
- Inoue, A. K., Shimizu, I., Iwata, I., & Tanaka, M. 2014, *MNRAS*, **442**, 1805
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, **873**, 111
- Iwata, I., Sawicki, M., Inoue, A. K., et al. 2022, *MNRAS*, **509**, 1820
- Kennicutt, R. C. J. 1998a, *ARA&A*, **36**, 189
- Kennicutt, R. C. J. 1998b, *ApJ*, **498**, 541
- Kennicutt, R. C. J., & De Los Reyes, M. A. C. 2021, *ApJ*, **908**, 61
- Kirkpatrick, J. D., Cruz, K. L., Barman, T. S., et al. 2008, *ApJ*, **689**, 1295
- Kirkpatrick, J. D.,Looper, D. L., Burgasser, A. J., et al. 2010, *ApJS*, **190**, 100
- Kirkpatrick, J. D., Cushing, M. C., Gelino, C. R., et al. 2011, *ApJS*, **197**, 19
- Kirkpatrick, J. D., Gelino, C. R., Faherty, J. K., et al. 2021, *ApJS*, **253**, 7
- Koekemoer, A. M., Aussel, H., Calzetti, D., et al. 2007, *ApJS*, **172**, 196
- Koekemoer, A. M., Faber, S. M., Ferguson, H. C., et al. 2011, *ApJS*, **197**, 36
- Kulkarni, G., Worseck, G., & Hennawi, J. F. 2019, *MNRAS*, **488**, 1035
- Laigle, C., McCracken, H. J., Ilbert, O., et al. 2016, *ApJS*, **224**, 24
- Lang, D., Hogg, D. W., & Mykytyn, D. 2016a, The Tractor: Probabilistic Astronomical Source Detection and Measurement, Astrophysics Source Code Library, ascl:[1604.008](https://ascl.net/1604.008)
- Lang, D., Hogg, D. W., & Schlegel, D. J. 2016b, *AJ*, **151**, 36
- Lanzuisi, G., Civano, F., Marchesi, S., et al. 2018, *MNRAS*, **480**, 2578
- Lee, K.-S., Alberts, S., Atlee, D., et al. 2012, *ApJL*, **758**, L31
- Madau, P., & Haardt, F. 2015, *ApJL*, **813**, L8
- Madau, P., Haardt, F., & Rees, M. J. 1999, *ApJ*, **514**, 648
- Marchesi, S., Civano, F., Elvis, M., et al. 2016, *ApJ*, **817**, 34
- Matsuoka, Y., Strauss, M. A., Kashikawa, N., et al. 2018, *ApJ*, **869**, 150
- McCracken, H. J., Milvang-Jensen, B., Dunlop, J., et al. 2012, *A&A*, **544**, A156
- McGreer, I. D., Fan, X., Jiang, L., & Cai, Z. 2018, *AJ*, **155**, 131
- Merlin, E., Fontana, A., Ferguson, H. C., et al. 2015, *A&A*, **582**, A15
- Merlin, E., Bourne, N., Castellano, M., et al. 2016, *A&A*, **595**, A97
- Mesinger, A., & Furlanetto, S. 2007, *ApJ*, **669**, 663
- Miyazaki, S., Komiyama, Y., Kawanomoto, S., et al. 2018, *PASJ*, **70**, S1
- Momcheva, I. G., van Dokkum, P. G., van der Wel, A., et al. 2017, *PASP*, **129**, 015004
- Mosleh, M., Williams, R. J., Franx, M., et al. 2012, *ApJL*, **756**, L12
- Mowla, L. A., van Dokkum, P., Brammer, G. B., et al. 2019, *ApJ*, **880**, 57
- Niida, M., Nagao, T., Ikeda, H., et al. 2020, *ApJ*, **904**, 89
- Nyland, K., Lacy, M., Sajina, A., et al. 2017, *ApJS*, **230**, 9
- Oke, J. B. 1974, *ApJS*, **27**, 21
- Parsa, S., Dunlop, J. S., & McLure, R. J. 2018, *MNRAS*, **474**, 2904
- Paulino-Afonso, A., Sobral, D., Ribeiro, B., et al. 2018, *MNRAS*, **476**, 5479
- Reid, N., & Majewski, S. R. 1993, *ApJ*, **409**, 635
- Robertson, B. E., Ellis, R. S., Furlanetto, S. R., & Dunlop, J. S. 2015, *ApJL*, **802**, L19
- Schaerer, D., Ginolfi, M., Béthermin, M., et al. 2020, *A&A*, **643**, A3
- Schenker, M. A., Ellis, R. S., Konidaris, N. P., & Stark, D. P. 2014, *ApJ*, **795**, 20
- Schmidt, M. 1968, *ApJ*, **151**, 393
- Scoville, N., Aussel, H., Brusa, M., et al. 2007, *ApJS*, **172**, 1
- Scranton, R., Johnston, D., Dodelson, S., et al. 2002, *ApJ*, **579**, 48
- Shibuya, T., Ouchi, M., Harikane, Y., et al. 2018, *PASJ*, **70**, S15
- Speagle, J. S., Steinhardt, C. L., Capak, P. L., & Silverman, J. D. 2014, *ApJS*, **214**, 15
- Steidel, C. C., Giavalisco, M., Pettini, M., Dickinson, M., & Adelberger, K. L. 1996, *ApJL*, **462**, L17
- Tange, O. 2015, Gnu Parallel 20150322 ('Hellwig') Zenodo, doi:[10.5281/zenodo.16303](https://doi.org/10.5281/zenodo.16303)
- Taniguchi, Y., Scoville, N., Murayama, T., et al. 2007, *ApJS*, **172**, 9
- Taniguchi, Y., Kajisawa, M., Kobayashi, M. A. R., et al. 2015, *PASJ*, **67**, 104
- Telfer, R. C., Zheng, W., Kriss, G. A., & Davidsen, A. F. 2002, *ApJ*, **565**, 773
- Trakhtenbrot, B., Volonteri, M., & Natarajan, P. 2017, *ApJL*, **836**, L1
- Trenti, M., & Stiavelli, M. 2008, *ApJ*, **676**, 767
- Vanden Berk, D. E., Richards, G. T., Bauer, A., et al. 2001, *AJ*, **122**, 549
- Weaver, J. R., Kauffmann, O. B., Ilbert, O., et al. 2022, *ApJS*, **258**, 11