



Universiteit
Leiden
The Netherlands

**A low [CII]/[NII] ratio in the center of a massive galaxy at $z = 3.7$:
evidence for a transition to quiescence at high redshift?
(Corrigendum)**

Schreiber, C.; Glazebrook, K.; Papovich, C.; Díaz-Santos, T.; Verma, A.; Elbaz, D.; ... ;
Wang, T.

Citation

Schreiber, C., Glazebrook, K., Papovich, C., Díaz-Santos, T., Verma, A., Elbaz, D., ... Wang, T. (2021). A low [CII]/[NII] ratio in the center of a massive galaxy at $z = 3.7$: evidence for a transition to quiescence at high redshift? (Corrigendum). *Astronomy And Astrophysics*, 650, C2. doi:10.1051/0004-6361/201936460e

Version: Accepted Manuscript

License: [Leiden University Non-exclusive license](#)

Downloaded from: <https://hdl.handle.net/1887/3264002>

Note: To cite this publication please use the final published version (if applicable).

A low [CII]/[NII] ratio in the center of a massive galaxy at $z = 3.7$: witnessing the transition to quiescence at high-redshift?

C. Schreiber¹, K. Glazebrook², C. Papovich³, T. Díaz-Santos⁴, A. Verma¹, D. Elbaz⁵, G. G. Kacprzak²,
T. Nanayakkara⁶, P. Oesch^{7,8}, M. Pannella⁹, L. Spitler^{10,11}, C. Straatman¹², K.-V. Tran^{13,14}, and T. Wang^{15,16}

¹ Astrophysics, Department of Physics, Keble Road, Oxford OX1 3RH, United Kingdom
e-mail: cschreib@orange.fr

² Centre for Astrophysics and Supercomputing, Swinburne University of Technology, Hawthorn, VIC 3122, Australia

³ George P. and Cynthia W. Mitchell Institute for Fundamental Physics and Astronomy, Department of Physics and Astronomy,
Texas A&M University, College Station, TX 77843, USA

⁴ Núcleo de Astronomía, Facultad de Ingeniería y Ciencias. Universidad Diego Portales, Ejército Libertador 441, Santiago, 8320000,
Chile.

⁵ AIM-Paris-Saclay, CEA/DSM/Irfu – CNRS – Université Paris Diderot, CEA-Saclay, pt courrier 131, 91191 Gif-sur-Yvette, France

⁶ Leiden Observatory, Leiden University, NL-2300 RA Leiden, The Netherlands

⁷ Department of Astronomy, University of Geneva, 51 Ch. des Maillettes, 1290 Versoix, Switzerland

⁸ International Associate, Cosmic Dawn Center (DAWN)

⁹ Faculty of Physics, Ludwig-Maximilians Universität, Scheinerstr. 1, 81679 Munich, Germany

¹⁰ Research Centre for Astronomy, Astrophysics & Astrophotonics, Macquarie University, Sydney, NSW 2109, Australia

¹¹ Department of Physics & Astronomy, Macquarie University, Sydney, NSW 2109, Australia

¹² Max-Planck Institut für Astronomie, Königstuhl 17, D-69117, Heidelberg, Germany

¹³ Australia Telescope National Facility, CSIRO Astronomy and Space Science, PO Box 76, Epping, NSW 1710, Australia

¹⁴ School of Physics, University of New South Wales, Sydney, NSW 2052, Australia

¹⁵ Institute of Astronomy, The University of Tokyo, Osawa, Mitaka, Tokyo 181-0015, Japan

¹⁶ National Astronomical Observatory of Japan, Mitaka, Tokyo 181-8588, Japan

Received 5 August 2019; accepted 26 November 2020

ABSTRACT

Understanding the process of quenching is one of the major open questions in galaxy evolution, and crucial insights may be obtained by studying quenched galaxies at high redshifts, at epochs when the Universe and the galaxies were younger and simpler to model. However, establishing the degree of quiescence in high redshift galaxies is a challenging task. One notable example is *Hyde*, a recently discovered galaxy at $z_{\text{spec}} = 3.709$. As compact ($r_{1/2} \sim 0.5$ kpc) and massive ($M_* \sim 10^{11} M_{\odot}$) as its quenched neighbor *Jekyll*, it is also extremely obscured yet only moderately luminous in the sub-millimeter. Panchromatic modeling suggested it could be the first galaxy found in transition to quenching at $z > 3$, however the data were also consistent with a broad range of star-formation activity, from fully quenched to moderate star formation rates (SFR) in the lower scatter of the galaxy main-sequence. Here, we describe Atacama Large Millimeter Array (ALMA) observations of the [C II] 157 μm and [N II] 205 μm far-infrared emission lines. The [C II] emission within the half-light radius is dominated by ionized gas, while the outskirts are dominated by photo-dissociation regions or neutral gas. This suggests that the ionization in the center is not primarily powered by on-going star formation, and could come instead from remnant stellar populations formed in an older burst, or from a moderate active galactic nucleus (AGN). Accounting for this information in the multi-wavelength modeling provides a tighter constraint on the star formation rate of $\text{SFR} = 50_{-18}^{+24} M_{\odot}/\text{yr}$. This rules out fully quenched solutions, and favors SFRs more than factor of two lower than expected for a main-sequence galaxy, confirming the nature of *Hyde* as a transition galaxy. These results suggest that quenching happens from inside-out, and starts before the galaxy expels or consumes all its gas reservoirs. Similar observations of a sample of massive and obscured galaxies would determine whether this is an isolated case or the norm for quenching at high-redshift.

Key words. Galaxies: evolution – galaxies: high-redshift – galaxies: star formation – sub-millimeter: galaxies

1. Introduction

Massive galaxies in the low-redshift Universe are observed to be mostly quiescent, with current star formation rates less than 1% of their past average (e.g., Pasquali et al. 2006). But how galaxies quench their star formation and turn into massive, red early-type galaxies is one of the key unresolved questions of galaxy evolution.

While a number of mechanisms have been proposed in the literature to cause galaxies to stop, reduce, or prevent star formation – including black hole feedback, strong outflows, or gas

stripping (e.g., Silk & Rees 1998; Birnboim & Dekel 2003; Croton et al. 2006; Gabor & Davé 2012; Martig et al. 2009; Förster Schreiber et al. 2014; Genzel et al. 2014; Peng et al. 2015), it is still unknown which of these mechanism is most important to explain emergence of quenched galaxies. Crucial insight can be gained by studying quiescent galaxies at higher redshift, where the available time for feedback processes to act is shorter, limiting the range of possible mechanisms. Recent work has shown that massive quiescent galaxies already existed as early as $z = 4$ (e.g. Labbé et al. 2005; Kriek et al. 2009; Gobat

et al. 2012; Merlin et al. 2018; Schreiber et al. 2018b; Belli et al. 2019), at an epoch where state-of-the-art numerical simulations predict all galaxies to be forming stars (e.g., Wellons et al. 2015; Davé et al. 2016). These galaxies must have had massive star-forming progenitors at even higher redshifts, and because the age of the Universe was then comparable to their estimated stellar ages (of order a billion years or less), they must have quenched shortly before being observed. The actor of this abrupt quiescence may therefore be easier to identify than in lower redshift objects, which are seen after several billion years of passive evolution.

An interesting case is that of the most distant known quiescent galaxy to date, at $z_{\text{spec}} = 3.715$ (Glazebrook et al. 2017). Detection of sub-millimeter emission toward this object (Simpson et al. 2017) was later found to arise from a neighboring extremely obscured galaxy at the same redshift (Schreiber et al. 2018c, hereafter S18). This was demonstrated by the detection with ALMA of the [C II] line blueshifted by 550 km/s compared to the Balmer absorption lines of the quiescent galaxy, and most convincingly by the improved spatial resolution ($0.4''$) and depth of the sub-millimeter imaging, showing that both the [C II] and dust emissions are produced by a separate, rotating galaxy, mostly unresolved ($0.1''$ radius) and located $0.5''$ away from the quiescent galaxy. The pair was dubbed “*Jekyll* and *Hyde*”, with *Jekyll* being the quiescent galaxy, and *Hyde* the obscured galaxy.

In S18, we performed an extensive analysis of the rich multi-wavelength data at hand to understand the physical conditions in the obscured galaxy *Hyde*. We showed that, despite its extreme obscuration and sub-mm detection, this galaxy appears to form stars at a relatively slow pace, with a star-formation rate $\text{SFR} < 100 M_{\odot}/\text{yr}$ and a stellar mass $M_{*} \approx 10^{11} M_{\odot}$. Given the strong obscuration ($A_V \approx 3$) and the galaxy’s large stellar mass, the infrared luminosity is in fact low enough ($L_{\text{IR}} \approx 10^{12} L_{\odot}$) that it may be entirely powered by intermediate-age stars, such that its current SFR (averaged over the last 10 Myr) could be as low as zero. Although these data could not exclude that *Hyde* is simply a normal galaxy in the lower envelope of the main sequence, the possibility of it being quenched or in transition to quenching is particularly interesting. Indeed, given that this galaxy is still obscured and is therefore likely to contain substantial gas reservoirs, this would be at odds with a number of proposed quenching mechanisms which require full removal or consumption of the gas reservoirs prior to quenching.

This surprising conclusion is nevertheless independently supported by a number of pieces of evidence, first reported in S18 and summarized here for the reader’s convenience: a) its dust temperature ($T_{\text{dust}} \approx 30$ K) is almost 10 K lower than the average for $z \sim 4$ galaxies (Schreiber et al. 2018a), which suggests a softer-than-average radiation field, b) its compact dust continuum size of about 0.5 kpc is smaller than the rest-ultraviolet size of all $z \sim 4$ star-forming galaxies in the same field, and is instead similar to that of quiescent galaxies (Straatman et al. 2015), and c) its $L_{[\text{C II}]} / L_{\text{FIR}}$ ratio is abnormally low for a galaxy of this luminosity, which we showed could be explained by a recent truncation of star formation, although other explanations for the observed [C II] deficit could not be excluded. Considered independently, none of these facts is conclusive or extremely unusual, and indeed there are other examples of galaxies which share at least one of these properties (e.g., similar compactness, dust temperature, or $L_{[\text{C II}]} / L_{\text{FIR}}$ ratio). It is, however, the sum of these facts which suggests an abnormal process is at play in this galaxy; to our knowledge, such a combination of observables is unique among distant massive galaxies. Although all these signs may point toward *Hyde* being caught in a very specific phase,

possibly in transition to quiescence, none of the data available at the time allowed us to prove this conclusively. In light of the results presented in this paper, we describe these points further in our discussion.

In this paper, we exploit new spectroscopic data from the Atacama Large Millimeter Array (ALMA) to address this question, combining observations of the [C II] and [N II] ($205.178 \mu\text{m}$) emission lines. [N II], owing to its high ionization potential (14.5 eV), is exclusively found in ionized gas regions, where it can be found alongside [C II] with an almost constant line ratio $[\text{C II}]_{\text{ion}} / [\text{N II}]_{\text{ion}} \sim 3$ (see Oberst et al. 2006). In contrast, photo-dissociation regions (PDRs) surrounding stellar birthclouds are a neutral medium that emits [C II] but no [N II]. The observed [C II]/[N II] ratio can therefore be used as a probe of the fraction of the [C II]-emitting gas which is associated with ionized gas. While star-forming regions will typically be comprised of both neutral and ionized gas, the former is found to dominate the [C II] emission in star-forming galaxies (Pavesi et al. 2016; Díaz-Santos et al. 2017). In contrast, nearby early-type galaxies, which are forming stars at much lower rates, were shown to have a significantly larger fraction of their [C II] associated with ionized gas (Lapham et al. 2017). This implies that the ionized component of [C II] is mostly not associated with star-formation, hence that the [C II]/[N II] ratio can be used as an independent tracer of star formation activity, which we apply here to the galaxy *Hyde*.

The structure of this paper is as follows. In section 2 we describe the new ALMA observations and detections, in section 3 we describe the method used to analyze them, in section 4 we describe our results and the ionization state of the gas, and use these observations to refine our estimate of the current SFR in *Hyde*. In section 5 we discuss how other observables support these results, and we finally conclude in section 6.

In the following, we assumed a Λ CDM cosmology with $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_M = 0.3$, $\Omega_{\Lambda} = 0.7$ and a Chabrier (2003) initial mass function (IMF), to derive both star-formation rates and stellar masses.

2. Data

The data we use in this paper consist of ALMA observations of *Hyde* obtained in two different programs. The first data set was obtained in the Director’s Discretionary Time (DDT) program 2015.A.00026.S (PI: Schreiber) to measure the [C II] emission; the galaxy was observed in band 8 (TDM correlator, covering 401.05–416.68 GHz with four 1.875 GHz spectral windows at 31.25 MHz resolution) for 1.2 hours (on-source), with a synthesized beam size of $0.52 \times 0.42''$ (natural weighting). These observations were already presented in S18, and were reprocessed here for homogeneity. The second data set was obtained in the regular call program 2018.1.00216.S (PI: Schreiber) to measure the [N II] emission; the galaxy was observed in band 7 (TDM correlator, covering 295.34–311.20 GHz with four 1.875 GHz spectral windows at 31.25 MHz resolution) for 0.8 hours (on-source), with a synthesized beam size of $0.29 \times 0.25''$ (natural weighting).

Both data sets were reduced with the same procedure, using the ALMA pipeline to produce dirty images with natural weighting (to maximize S/N) and a spectral averaging of three elements (so that the spectral response function is effectively one channel, see Section A.6.1 in the ALMA Proposer’s Guide). The pixel size was left to its default value of $0.05''$ and $0.085''$ in band 7 and band 8, respectively (corresponding to 0.36 and 0.61 kpc), which generously samples the core of the dirty beam; finer pixel sizes would not benefit the quality of the analysis.




Fig. 1. [C II] (green) and [N II] (red) spectra extracted at the peak pixel of the line emission, after continuum subtraction. The [N II] spectrum was rescaled upward by a factor five for easier comparison to the [C II] spectrum. The integration window used to create the line maps is shown with vertical dotted lines.

Since both data sets were observed with the TDM correlators, the final cubes have a spectral resolution of 35 and 45 km/s, respectively, which is sufficient to resolve the broad line profiles expected in massive compact galaxies.

Compact continuum emission was clearly detected in both data sets (S_{peak}/N of 46 and 41, respectively), which had to be removed prior to analyzing the line emission. To optimize the S/N , this was done in the image domain by fitting the best continuum spectral model obtained in S18 to the spectral data at each pixel of the image, excluding spectral elements with expected line emission. The resulting spectra, extracted at the peak pixel, are shown on Fig. 1. Line maps were finally created by spectrally-averaging the continuum-subtracted cubes over 800 km/s around the expected line frequencies at $z = 3.7087$ (see Fig. 1), the redshift obtained from the [C II] emission in S18. The lines were detected with a S_{peak}/N of 17 and 5, respectively, and the maps are shown in Fig. 2. Despite these moderate S_{peak}/N , the detection of [N II] is still highly significant: since we did not fit for the position of the line emission in neither the spatial nor the frequency domains, the null hypothesis of a non-detection can be safely rejected.

3. Spatial profile modeling

3.1. Description of the modeling

Using a custom-build software, we then modeled each line map independently in the image domain. We did not attempt to model the maps directly in the visibility domain for two reasons; first, fitting complex profiles other than Gaussians and point sources in the visibility domain is particularly challenging and time-consuming, and second, it has been shown that image-based and visibility-based methods actually provide similar results for ALMA data (Hodge et al. 2016). Nevertheless, we double check in section 3.2 the result of our modeling against the observed visibilities.

Our model consists of a large grid of two-component exponential profiles. The first “central” component was given a small scale length of 0.1 to 0.5 kpc to represent an unresolved central region (S18 showed the continuum emission extends over a ~ 0.7 kpc half-light radius), while the scale length of the sec-




Fig. 2. Dirty beam (top), continuum (center) and line maps (bottom) of *Hyde* in band 8 (left) and 7 (right) produced by the ALMA pipeline, without any cleaning applied (“dirty” images). Line maps are continuum-subtracted and were produced by summing the flux in a 800 km/s velocity window centered on the [C II] mean velocity. The half-intensity area of the corresponding dirty beams are shown in the bottom-right corner of each panel. Contours shown are 3σ (dotted line) and 10σ (solid line).

ond “extended” component was allowed to vary between 0.15 and 3 kpc, with the constraint that its scale length must be at least 0.05 kpc larger than that of the central component¹. We let the flux of each component vary independently and freely, with the sole constraint that the combined light profile must remain positive at all positions. We also varied the axis ratio and the position angle, however to save on computation time we forced the values of both components to match. Finally, the position of both components was fixed to the centroid of the continuum emission. Since the size of the galaxy is small compared to the image pixel size, we generated models with a $\times 9$ oversampling factor, and a further $\times 3$ oversampling for the central 3×3 pixels to properly sample the exponential core. This corresponds to an over-sampled pixel size of 0.01 (band 6) and 0.02 kpc (band 8)

¹ We tried extending this grid to a maximum of 1.5 kpc for the “central” component and 5 kpc for the “extended” component; this had no significant impact on our flux measurements, suggesting the original grid was large enough.

at the core. We then convolved the model with the ALMA dirty beam, and selected the best-fit model using maximum likelihood estimation, accounting for spatially-correlated noise in the computation of the likelihood (see Appendix C).

To estimate the uncertainty on the model, we used Monte Carlo simulations: the whole procedure was repeated 200 times on mock images that were created by perturbing the observed images with realistic noise (same amplitude and same covariance as in the real data, see Appendix A), leading to 200 other “acceptable” models from which we estimated confidence intervals using the 16th and 84th percentiles (see Appendix E). Because the distributions of the observables derived from this modeling (e.g., the ionized gas fraction) can be strongly non-Gaussian, we carried out all our calculations on each of the 200 Monte Carlo realizations, and derived confidence intervals for all derived quantities in the same fashion throughout the paper. With 200 Monte Carlo simulations, 1σ asymmetric error bars have an accuracy of 10%.

3.2. Validation of the image-domain modeling

Since we performed all our analysis in the image-domain, we checked that the same trends we detect in this analysis are present in the raw ALMA visibilities. To do so, we used CASA to perform continuum subtraction in the visibility domain using `uvcontsub` (first order polynomial), excluding channels containing the line in the fit. From the resulting measurement set, we then extracted and averaged the frequency channels covered by the line using `splitt`. We note that this data set is not strictly equivalent to the image-domain data we used in our analysis, where the continuum subtraction and frequency averaging were performed in the image domain (with slightly different weights, computed from the image RMS in each channel). The images created from these alternative measurement sets have 10-20% larger noise RSM; the S/N will therefore be lower than in our main analysis, but this will be sufficient for our purpose.

We finally extracted the visibilities using `ms.getdata()` and manually performed the remaining averaging (polarization) and the binning by (u, v) distance. The real and imaginary parts of the visibilities were averaged separately in each bin, using the weights in the measurement set, and combined to form the binned amplitude. The uncertainties on the binned amplitudes were estimated from the weighted standard deviation of the data in the bin, and scaled down by the square root of the number of points.

In parallel, we also produced mock visibilities for our best fit and Monte Carlo models and binned them in a similar way. The corresponding visibilities were obtained from the $\times 9$ over-sampled model images created by our fitting procedure. These visibilities were computed and injected into the real measurement set using the CASA task `ft`, which generates mock measurement sets with the exact same (u, v) coverage and weights as the real data. To correct for any discrepancy the data reduction between these visibilities and the images from which we derived the models, we applied a global rescaling factor to our best-fit model amplitudes to best match the observed visibilities.

The binned amplitudes are shown in Fig. 3. These figures show the [C II] emission is clearly resolved. The [N II] emission is more noisy, but appears more compact and consistent with being unresolved. The best-fit models obtained from the image-domain analysis provide good matches to the visibility data, but seem mutually inconsistent between the two lines. In fact, trying to fit the [C II] data with the [N II] model, we obtained a worse χ^2 (22.4, vs 20.9 for the best fit [C II] model), and similarly for

the [N II] data and the [C II] model (15.0, vs 8.0 for the best fit [N II] model). This implies that the two emission lines have different spatial distributions, a point we analyze in greater detail in section 4.

3.3. Aperture fluxes

We first applied our model to the new band 7 continuum image, which has a higher S/N and sharper angular resolution than the band 8 data used in S18, and updated the dust continuum half-light radius to 0.51 ± 0.07 kpc (major axis), which is consistent with the value obtained from the band 8 data. This value was used to separate the emission into a “center” and “outskirts” components, as we now describe.

Coming back to the other images (line and continuum maps), we summed the flux of every model (both for the best fit and the 200 Monte Carlo simulations) in two circular annuli: from 0 to 0.5 kpc for the “central” annulus, and from 0.5 to 3.5 kpc for the “outskirts” annulus. We note that the central annulus actually contains more than half ($\sim 75\%$) of the continuum emission since the model axis ratio is lower than one.

The obtained “central” and “outskirts” fluxes (for both lines and continua) are the main measurements we use in the rest of this analysis. The other parameters of the model, such as the respective size and flux of both exponential components, are effectively marginalized over and considered unimportant. Aperture fluxes are typically simpler to constrain than other shape-related model parameters, because they tend to be less model-dependent. For example, the flux inside our “central” aperture depends only slightly on the adopted size for the central component; if the aperture is larger or comparable in size, it then contains all the flux from that component, no matter how that flux is distributed internally. Furthermore, the total flux in that aperture has a natural upper bound from the data. To illustrate this, we can consider the peak pixel of the image, which for [C II] entirely contains our aperture. Because of the convolution with the dirty beam, this pixel contains flux from both inside and outside of our aperture. Yet, even if the relative contribution of one versus the other is uncertain, neither can exceed the observed peak flux.

In this light, it may seem overly complex to introduce a rich two-component model to, in the end, simply measure two aperture fluxes. We stress that this model complexity is in fact crucial to assess the reliability of our measurement. Indeed, the constraining power of our data (owing both to S/N and angular resolution) is not sufficient to allow us to determine the exact shape of the intensity profile at all radii. Therefore, by exploring as wide a range of models as possible – including numerous models which are barely distinguishable on the ALMA images – we make sure our uncertainties on the aperture fluxes encompass all credible scenarios allowed by the data.

To test our method and the accuracy of our uncertainties, we applied our full measurement procedure to simulated images with injected sources, as described in Appendix D. We found that our method can recover the total, central, and outskirts fluxes in all our simulated images with no detectable systematic bias. The estimated uncertainties from Monte Carlo simulations were found to correctly capture the noise in our measurements, with at most a 10% under-estimation.




Fig. 3. Binned amplitude of the ALMA visibilities as a function of the UV distance ($u^2 + v^2$). The [C II] visibilities are shown on the left (green), and the [N II] visibilities are shown on the right (red). Observed visibilities are displayed as filled circles with error bars; for clarity, the visibilities on long baselines, which have large uncertainties, are not displayed. The best fit models obtained in this paper (from the image-domain analysis) are shown as solid lines (green for [C II], red for [N II]). For comparison purposes, on each panel, the model of the other line is also displayed, but re-normalized to fit the observed visibilities. The 1σ confidence interval of the model (determined from our Monte Carlo simulations) is shown as a shaded region in the background.

4. Results

4.1. Spatial distribution of the lines

In Fig. 4 (left) we show the surface brightness profile of *Hyde* inferred from our two-component model for the dust continuum and the two emission lines, [C II] and [N II]. Although the S/N of the [N II] data is low, the angular resolution is sufficient to show that most of the emission is centrally-concentrated, with a profile that is similar to that of the dust continuum. The [C II] emission, however, appears more extended (as is commonly found in the literature, see, e.g., Gullberg et al. 2018; Rybak et al. 2019; Tadaki et al. 2019). The intrinsic half-light radii we obtained for the continuum, [C II], and [N II] are (respectively) $r_{1/2} = 0.51 \pm 0.07$, 1.2 ± 0.3 , and $0.3^{+0.3}_{-0.1}$ kpc (along the major axis). The lower error bar on the [N II] size is limited by the minimum source size in our model grid, 0.1 kpc; the data are consistent with the [N II] emission being point-like.

Martí-Vidal et al. (2012) quantify what we should be able to measure in our data given the angular resolution and S/N : based on their Eq. (7), our data should be able to distinguish, with more than 95% reliability (2σ), between a point source and an extended profile of size 0.30, 0.96, and 0.94 kpc, respectively. The sizes of the continuum and [C II] are both larger than these values and are indeed “measured” ($r/\sigma_r = 7$ and 4, respectively), and the size we report for [N II] indeed has a 2σ upper limit of 0.9 kpc. Our measurements are thus consistent with these theoretical expectations.

To further confirm this difference in half-light radii with a simpler independent method, we used the simulated noise maps described in the previous section and injected sources of sizes matching the above values, and convolved with the dirty beam. For each source, we computed its observed radial profile on the noisy image, and located the radius at which the emission falls below half of the peak. Comparing this simulation to the values observed on the real images, we found that the [C II] emission

is significantly (3.6σ) larger than the continuum, and that [N II] is significantly (2.4σ) smaller than [C II], while the difference between [N II] and the continuum is only marginal (1.2σ).

This confirms that the difference in size between the [C II] and [N II] emission is not simply due to noise, and implies the presence of a strong gradient in the [C II]/[N II] ratio, with inner regions having lower [C II]/[N II] and therefore a higher proportion of ionized gas.

To quantify the gradient in the line ratio, we now turn to the measured fluxes. We list the measured values and their uncertainties in Table 1, and provide in Appendix E the full observed distributions in the Monte Carlo simulations, which we use to determine and propagate uncertainties throughout this paper. We found no significant flux in [N II] beyond our fiducial 0.5 kpc radius, and conversely, we found less than half of the [C II] flux is located inside this radius. This leads to a line ratio in the central region of $[C II]_{\text{center}}/[N II]_{\text{center}} = 4.8^{+2.4}_{-1.7}$, and a lower limit on the line ratio in the outskirts of $[C II]_{\text{outskirts}}/[N II]_{\text{outskirts}} > 15$ (1σ limit).

4.2. Ionized gas distribution

Using the fluxes estimated in the previous section, we estimated the fraction of [C II] emission associated with ionized gas (or “ionized gas fraction” for short), $f_{[C II],\text{ion}}$, following a method similar to Díaz-Santos et al. (2017) and Lapham et al. (2017):

$$f_{[C II],\text{ion}} = \frac{[C II]_{\text{ion}}/[N II]_{\text{ion}}}{[C II]/[N II]}. \quad (1)$$

Here we adopt a fixed $[C II]_{\text{ion}}/[N II]_{\text{ion}} = 2.80 \pm 0.18$, which is the mean value found by Lapham et al. (2017). In principle this ratio has a weak dependence on the electron density, which can be measured using the $[N II]_{122}/[N II]_{205}$ ratio, but the $[N II]_{122}$ line is not observable for *Hyde* due to poor atmospheric transmission at that redshifted wavelength. If *Hyde* turns out to have an unusual

Table 1. Measured line fluxes from the image analysis, corresponding ionized gas fractions and infrared luminosities.

Flux origin	$I_{[\text{C II}]}$ Jy.km/s	$I_{[\text{N II}]}$ Jy.km/s	$[\text{C II}]/[\text{N II}]$ $_e$	$f_{[\text{C II}],\text{ion}}$ %	$f_{\text{IR,ISM}}$ %	$\log_{10}(L_{\text{IR}}/L_{\odot})$ $_f$	$\log_{10}(L_{\text{IR,BC}}/L_{\odot})$	$\log_{10}(L_{[\text{C II}]} / L_{\text{FIR}})$
Total ^a	$1.73^{+0.13}_{-0.10}$	$0.227^{+0.053}_{-0.046}$	$9.8^{+2.6}_{-1.8}$	28^{+8}_{-6}	60^{+8}_{-10}	$12.03^{+0.14}_{-0.14}$	$11.52^{+0.17}_{-0.24}$ ^g	$-2.89^{+0.23}_{-0.26}$
Central ^b	$0.69^{+0.25}_{-0.19}$	$0.187^{+0.049}_{-0.039}$	$4.8^{+2.4}_{-1.7}$	58^{+32}_{-20}	84^{+12}_{-15}	$11.93^{+0.15}_{-0.16}$	$11.12^{+0.30}_{-0.69}$	$-3.17^{+0.26}_{-0.28}$
Outskirts ^c	$1.04^{+0.25}_{-0.27}$	$0.022^{+0.072}_{-0.022}$	$60^{+\text{inf}}_{-46}$	5^{+14}_{-5}	13^{+33}_{-13}	$11.39^{+0.13}_{-0.19}$	$11.26^{+0.21}_{-0.21}$	$-2.46^{+0.33}_{-0.29}$
Peak pixel ^d	$1.24^{+0.07}_{-0.07}$	$0.190^{+0.038}_{-0.038}$	$8.6^{+2.4}_{-1.8}$	32^{+11}_{-9}	64^{+12}_{-11}	$11.95^{+0.14}_{-0.14}$	$11.37^{+0.17}_{-0.17}$	$-2.96^{+0.22}_{-0.23}$

^a Summed flux inside a 3.5 kpc radius. ^b Summed flux inside a 0.5 kpc radius. ^c Summed flux between 0.5 and 3.5 kpc radius. ^d Flux of the peak pixels on the ALMA image (warning: because of the poorer resolution, this method measures flux on a larger scale for [C II] than for [N II]). ^e Ratio of the [C II] to [N II] flux, with fluxes expressed in W/m². ^f Total infrared luminosity (8–1000 μm), computed from the total L_{IR} estimated in Schreiber et al. (2018c) scaled by the fraction of [N II] continuum flux in each region. ^g Computed as the sum of the central and outskirts luminosities.




Fig. 4. Left: modeled intrinsic surface brightness profiles of *Hyde* in the dust continuum (black, horizontal stripes), [C II] (green, +45 degrees stripes), and [N II] (red, -45 degrees stripes). The dust continuum half-light radius is indicated with a vertical dotted black line. The PSF HWHM of the [C II] (resp. [N II]) line map is indicated with vertical dot dashed green (resp. red) line. The hashed regions show the 1σ confidence intervals obtained from the 200 mock noise realizations, and are centered on the best-fit model. **Right:** Ionized gas fraction ($f_{[\text{C II}],\text{ion}}$) determined from the [C II]/[N II] line ratio in different regions of *Hyde* (green circles; center, outskirts, and total), compared to other values reported in the literature. For all galaxies, the ionized gas fraction was computed assuming a fixed $[\text{C II}]_{\text{ion}}/[\text{N II}]_{\text{ion}} = 2.80 \pm 0.18$ (see text). We show literature values for local ETGs (red circles) and KINGFISH galaxies (purple circles) from Lapham et al. (2017), local infrared luminous galaxies in GOALS (black circles) from Zhao et al. (2016) and Díaz-Santos et al. (2017) and HERCULES from Rosenberg et al. (2015) and Kamenetzky et al. (2016), and high-redshift SMGs from Pavesi et al. (2016). We also show the upper limit of Zhang et al. (2018) obtained by stacking high-redshift lensed SMGs (downward-pointing orange triangle).

electron density, based on the modeling of Lapham et al. (2017) its $[\text{C II}]_{\text{ion}}/[\text{N II}]_{\text{ion}}$ could only be higher, up to $[\text{C II}]_{\text{ion}}/[\text{N II}]_{\text{ion}} = 4$, which would only increase the value of $f_{[\text{C II}],\text{ion}}$ we infer from the data.

In Fig. 4 (right), we display the ionized gas fraction of *Hyde* in the “central” and “outskirts” annuli (which we recall correspond to inside and outside of the dust half-light radius, respectively). We find a clear difference between the two regions. The outskirts of the galaxy has a low $f_{[\text{C II}],\text{ion}}$ of 0–19%, which is typical of infrared-luminous local galaxies and distant SMGs (Pavesi et al. 2016; Díaz-Santos et al. 2017). Likewise, the $L_{[\text{C II}]} / L_{\text{FIR}}$ ratio in the outskirts, $\log_{10}(L_{[\text{C II}]} / L_{\text{FIR}}) = -2.46^{+0.33}_{-0.29}$, is similar to that observed in other distant galaxies of this luminosity (e.g., Capak et al. 2015). This suggests that the majority of the [C II] emission in the outskirts is associated with cold star-forming gas.

On the other hand, the center of the galaxy has a high² $f_{[\text{C II}],\text{ion}}$ of 38–90%, which has never been observed in a distant SMG and is rarely observed in local infrared-luminous galaxies. Finding such high $f_{[\text{C II}],\text{ion}}$ is instead not uncommon in more normal or quiescent local galaxies, such as main-sequence galaxies and ETGs (Lapham et al. 2017). It should be noted, however, that $f_{[\text{C II}],\text{ion}}$ or [C II]/[N II] values from the literature are typically only quoted for the galaxies as a whole, and few studies prior to this one have attempted to separate the emission from the center and outskirts (see, e.g., Parkin et al. 2013 where a $f_{[\text{C II}],\text{ion}}$ gradient was found in M51). It is possible that higher $f_{[\text{C II}],\text{ion}}$ values could also be found in the center of some SMGs and other IR-

² As a sanity check, we also computed the $f_{[\text{C II}],\text{ion}}$ using the peak fluxes only; these peak fluxes are not model-dependent, but they will contain flux from both inside and outside of our central aperture; we still found a high $f_{[\text{C II}],\text{ion}}$ of 23–43%, see Table 1.

luminous galaxies if they were observed with a sufficient angular resolution. This would be an interesting avenue for future observations, in particular as it could allow finding analogs to *Hyde* in the local Universe, which would provide invaluable insights on the physical processes at play.

The $\log_{10}(L_{[\text{C II}]} / L_{\text{FIR}}) = -3.17^{+0.26}_{-0.28}$ in the center of *Hyde* is also a factor five lower than in the outskirts, which is expected if the gas is predominantly ionized (Díaz-Santos et al. 2017). This supports the hypothesis that the majority of the [C II] emitting gas in the galaxy center is ionized and not star-forming, and therefore that an energy source other than on-going star formation is significantly contributing to the energy budget in the center of the galaxy. If *Hyde* is indeed on the path to quenching, this would strongly suggest an inside-out quenching channel.

We cannot determine with certainty the nature of this central ionizing source with the available data, but we can propose two likely candidates. First, the source could be a compact population of intermediate-age stars born in a recent past. Indeed, *Hyde* is believed to contain about $4 \times 10^{10} M_{\odot}$ of stars within its half light radius (S18), and a Bruzual & Charlot (2003) single stellar population of this mass should generate an [N II]-ionizing flux with total energy $\geq 10^8 L_{\odot}$, even at advanced ages of several hundred million years, which is larger than the observed central [N II] luminosity, $(6.5^{+1.7}_{-1.4}) \times 10^7 L_{\odot}$. However, hydrogen and carbon and will also consume part of this ionizing flux; determining whether this is indeed a viable hypothesis would require dedicated modeling, a measure of the gas metallicity, and a better understanding of the geometry of the system, which are all presently lacking.

Second, the ionizing source could be an AGN. Given that none of the available data indicate the presence of an AGN in this galaxy ($L_X < 10^{44}$ erg/s; Civano et al. 2016, $L_{1.4\text{GHz}} < 10^{24}$ W/Hz; Smolčić et al. 2017), this may seem less likely, yet a weak or obscured AGN cannot be ruled out. The presence of an AGN would make the interpretation of the observed line ratios more difficult, and could invalidate some of the assumptions and finer calculations presented in the following section, which only account for stellar ionization flux. Yet, the implication on the galaxy's SFR estimate would follow a similar path: if an AGN exists in the center of this galaxy, with a luminosity large enough to affect the continuum and line emission, it would necessarily contribute to the observed L_{IR} . This, in turn, would imply that SFR estimates using L_{IR} would be biased high. This would include our initial estimate from S18, which already tentatively placed the galaxy below the galaxy main sequence. If an AGN is indeed present, and given the large amount of dust and the compactness of the galaxy, the AGN radiation could be trapped (e.g., Costa et al. 2018) and couple efficiently with the gas to suppress star formation.

4.3. Revised star formation rate

We now use the characterization of the ionizing emission presented in the previous section to provide a refined estimate of the galaxy's SFR. We do so by estimating the infrared luminosity produced in birth clouds by young stars, and feed this luminosity into our SED modeling as a prior.

The total infrared luminosity, L_{IR} , is a very useful observable in SED modeling, as it offers an independent constraint on the obscured luminosity of a galaxy. This is particularly crucial to constrain the SFR. However, like [C II], the L_{IR} can originate from different regions in the galaxy. Here we will consider two main contributors: birthclouds (BC), which are heated exclusively by stars younger than 10 Myr (Charlot & Fall 2000),




Fig. 5. Likelihood of the photometric data for *Hyde* as a function of $\text{sSFR} = \text{SFR}/M_{\star}$. The likelihood is shown without any constrain on the dust luminosity (gray), using the total infrared luminosity L_{IR} (orange), or using the birthcloud infrared luminosity $L_{\text{IR,BC}}$ (red). The horizontal red lines show the 68% (thick) and 90% (thin) confidence intervals. The vertical lines indicate the locus of the galaxy main-sequence and its scatter from Schreiber et al. (2017).




Fig. 6. Updated location of *Hyde* with respect to the galaxy main-sequence. The dashed (resp., dotted) line indicate the locus of the galaxy main-sequence (resp., its scatter) from Schreiber et al. (2017). The dark (resp., light) red region shows the final 68% (resp., 90%) confidence region for *Hyde* obtained in this paper. The red cross indicates the location of our best-fit solution.

and the inter-stellar medium (ISM), which is heated by the older stars (e.g., da Cunha et al. 2008). We note the corresponding infrared luminosities $L_{\text{IR,BC}}$ and $L_{\text{IR,ISM}}$, respectively. $L_{\text{IR,BC}}$ is the quantity we aim to estimate, as by definition it is most directly tracing the recent SFR.

Next, we need to estimate the fraction of the infrared luminosity produced in the ISM, $f_{\text{IR,ISM}} = L_{\text{IR,ISM}}/L_{\text{IR}}$. Intuitively, we

can expect this fraction to be closely related to $f_{[\text{C II}],\text{ion}}$, which we computed earlier, since PDRs are co-located with birthclouds, and since ionized gas is part of the ISM. In what follows we will therefore assume $L_{[\text{C II}],\text{ion}} = L_{[\text{C II}],\text{ISM}}$ and $L_{[\text{C II}],\text{PDR}} = L_{[\text{C II}],\text{BC}}$.

With these assumptions, we can relate $f_{\text{IR,ISM}} = L_{\text{IR,ISM}}/L_{\text{IR,ISM}}$ to $f_{[\text{C II}],\text{ion}}$, using the following relation:

$$\begin{aligned} \frac{1}{f_{\text{IR,ISM}}} - 1 &= \frac{L_{\text{IR,BC}}}{L_{\text{IR,ISM}}} = \frac{L_{\text{IR,BC}}}{L_{[\text{C II}],\text{BC}}} \frac{L_{[\text{C II}],\text{ISM}}}{L_{\text{IR,ISM}}} \frac{L_{[\text{C II}],\text{BC}}}{L_{[\text{C II}],\text{ISM}}} \\ &= \frac{([\text{C II}]/\text{IR})_{\text{ISM}}}{([\text{C II}]/\text{IR})_{\text{BC}}} \left(\frac{1}{f_{[\text{C II}],\text{ion}}} - 1 \right) \\ &= \alpha \left(\frac{1}{f_{[\text{C II}],\text{ion}}} - 1 \right) \end{aligned} \quad (2)$$

We can determine α empirically using a set of reference galaxies for which we can estimate both $f_{[\text{C II}],\text{ion}}$ and $f_{\text{IR,ISM}}$. Here we used the local infrared-luminous galaxies from GOALS (Díaz-Santos et al. 2017), which have direct measurements of [C II] and [N II] from *Herschel* and an average of $f_{[\text{C II}],\text{ion}} = 12.6 \pm 0.7\%$, similar to high- z SMGs. Based on the short gas depletion times typically observed for starbursting galaxies (e.g., Béthermin et al. 2015), we assumed that the GOALS galaxies were mostly formed in a brief burst lasting 50 to 200 Myr³. We then used Bruzual & Charlot (2003) stellar populations and the uniform Calzetti et al. (2000) dust screen to compute the expected fraction of their bolometric luminosity produced by young stars for such star formation histories, $f_{\text{bol,BC}} = 65 \pm 5\%$. Young stars were defined as stars younger than 10 Myr, as above. Given the strong attenuation in these galaxies (Howell et al. 2010), the bolometric luminosity is practically equal to the infrared luminosity, and this latter fraction can thus be identified as $1 - f_{\text{IR,ISM}}$ for GOALS galaxies. Feeding these estimates back to Eq. 2 gives $\log_{10}(\alpha) = -0.57 \pm 0.10$ (in other words, the [C II]/IR ratio is a factor of four times lower in the ambient ISM than in birthclouds).

With the knowledge of $f_{\text{IR,ISM}}$, which we can compute independently in the center ($f_{\text{IR,ISM}} = 69\text{--}96\%$) and outskirts ($f_{\text{IR,ISM}} = 0\text{--}46\%$) of *Hyde*, we can estimate the summed infrared luminosity produced in birth clouds for both regions: $\log_{10}(L_{\text{IR,BC}}/L_{\odot}) = 11.52 \pm 0.20$. We then ran FAST++⁴ v1.3 to model the multi-wavelength photometry with the same setup as in S18, but now using $L_{\text{IR,BC}}$ as a prior on the obscured luminosity of stars younger than 10 Myr. Briefly, in S18 we used a flexible model for the star-formation history, formulated as an exponentially rising SFR followed by an exponential decline, with a variable time of transition between the two phases and variable exponential timescales.

The outcome is illustrated in Figs. 5 and 6. The stellar mass was essentially unaffected, but the updated model produced an $\text{SFR} = 50_{-18}^{+24} M_{\odot}/\text{yr}$ (averaged over the last 10 Myr), which is non-zero. This leads to $\text{sSFR} = 0.71_{-0.34}^{+0.58} \text{Gyr}^{-1}$. With a main-sequence locus at $\text{sSFR} = 2.8 \text{Gyr}^{-1}$ (Schreiber et al. 2017), this places the galaxy a factor 4.0 below the main sequence, with a lower limit of >2.3 at 68% confidence (and just >1.3 at 90% confidence).

The previous estimate of the SFR from S18, using instead the total L_{IR} as a constraint, only allowed us to obtain an upper

³ As shown in da Cunha et al. (2010), starbursting galaxies can also contain a component of older stars, which would also contribute to the IR emission. If true, this would increase our estimate of α from the GOALS galaxies, which would in turn decrease the final value of $L_{\text{IR,BC}}$ and SFR for *Hyde*. Our assumption is therefore conservative.

⁴ <https://github.com/cschreib/fastpp>

limit on the SFR (see Fig. 5). The new data, however, exclude the possibility that *Hyde* has fully quenched ($\text{SFR} \ll \text{SFR}_{\text{MS}}$). While the 90% upper limit on the SFR is actually unchanged compared to our earlier estimates, we are now able to constrain both sides of the SFR probability distribution, and we can therefore give more credit to the maximum probability solution and the standard 68% confidence interval. In this light, we can claim that the SFR is now constrained to intermediate values, a factor of 2–10 lower than the main-sequence level. Accounting for the observational uncertainty and the log-normal main-sequence scatter, the probability of observing $\text{sSFR} < 0.71 \text{Gyr}^{-1}$ when randomly drawing from the main-sequence distribution is 4%.

A factor limiting the strength of this conclusion is that the location of the galaxy main sequence is not yet perfectly determined at $z \sim 4$. Here we based our comparison on the estimate from Schreiber et al. (2017) since it is also based on ALMA-derived SFRs and with stellar masses derived in a similar way; this mitigates the impact of systematic biases on both M_{\ast} and SFR estimations, and makes the relative SFR difference between *Hyde* and the main sequence more robust. In fact, the most robust comparison would be obtained by applying the exact same method to determine the SFR for *Hyde* and main-sequence galaxies. Unfortunately, this cannot be achieved until high-quality [C II] and [N II] data are available for a representative sample of other normal galaxies. Yet, even setting aside all our SFR modeling, it is still true that *Hyde* has a high overall [N II]/[C II] ratio compared to high- z SMGs, and that it is “special” in a number of ways (high attenuation, compact size, etc.). Based on this, we argue that most galaxies are not like *Hyde*, and therefore that on average the SFR of typical main-sequence galaxies should be correct.

Nevertheless, we can quantify how the above result depends on the estimated locus and scatter of the main-sequence. If the main-sequence scatter is increased to 0.4 dex (resp. 0.5 dex), we find that the probability P of observing *Hyde*’s sSFR when drawing from the main-sequence distribution increases to 9% (resp. 13%). This does not appear to be supported by observations however, as other references in the literature typically report a lower scatter. For example, if the scatter is decreased to 0.25 dex (resp. 0.15 dex; Speagle et al. 2014; Pearson et al. 2018), P drops to 2.5% (resp. 0.5%). Similarly, if the main-sequence mean sSFR is decreased by 0.1 dex (Pearson et al. 2018), P increases to 8%. However, most references in the literature using FIR, sub-mm, or radio-based SFR estimates actually report a similar or higher mean sSFR at $z \sim 4$ and $\log_{10}(M_{\ast}) \sim 10.8$; with a difference of +0.1 dex (Speagle et al. 2014), +0.02 dex (Tomczak et al. 2016), +0.03 dex (Bourne et al. 2017), +0.07 dex (Leslie et al. 2020). Considering an increase of the main-sequence sSFR by 0.1 dex would decrease P to 2.3%.

5. Our result in context

5.1. The transition to quiescence?

The modeling in the previous section allowed us to obtain a revised SFR estimate for *Hyde*, which confirms and refines the earlier estimate from S18. This SFR, together with the estimated large stellar mass, would place the galaxy significantly outside of the standard galaxy main-sequence and its scatter (e.g., Schreiber et al. 2017). This supports the hypothesis that the galaxy does not belong to the main-sequence, and is instead in transition to quiescence.

As in S18, we are still not able to determine the future of this galaxy with certainty; although its SFR at the time of obser-

vation does appear to be low, we cannot exclude that this only corresponds to a temporary pause in its activity. At best, we can bring forward two arguments which disfavor (but not disprove) this possibility. Firstly, the galaxy is already among the most massive individual known at high redshift, and sits beyond the knee of the stellar mass function. This implies that, if its SFR does increase in the future, it cannot do so for very long. Secondly, if high-redshift galaxies commonly experience large (but temporary) fluctuations in their SFR, this would be detectable as an increase in the scatter of the galaxy-main sequence. This does not appear to be the case at least at $z \sim 4$ (e.g., Schreiber et al. 2017).

5.2. Other evidence

As pointed out in our introduction, *Hyde* compiles an unusual combination of observables, when compared to other massive galaxies (or SMGs) at a similar epoch. We expand on these aspects in the following paragraphs, and discuss what additional observations could help confirm or contradict our results.

One particularly unusual combination is that of a lower-than-average dust temperature and a compact geometry. *Hyde* is indeed unusually compact; in recent sub-millimeter surveys, only 4 to 6% of SMGs turn out to have sizes as small as *Hyde* (e.g., Ikarashi et al. 2017; Gullberg et al. 2019). For massive star-forming galaxies in general, compact sizes are typically observed in starburst galaxies located above the galaxy main-sequence (e.g., Elbaz et al. 2011). These galaxies also tend to have a higher dust temperature (e.g., Elbaz et al. 2011; Magnelli et al. 2014; Béthermin et al. 2015), which can be seen as a natural consequence of their high luminosity and compact geometry. Based on the very compact size of *Hyde*, we could have therefore expected to see an enhanced dust temperature, but the opposite is observed.

Indeed, when comparing to the average dust temperature of massive $z \sim 4$ galaxies from Schreiber et al. (2018a), *Hyde* displays a temperature about 10 K lower than average. Although dust temperatures are notoriously difficult to measure (e.g., Casey 2012), in this case we can compare *Hyde*'s temperature to a reference which has been measured with the same method (SED fitting and SED templates) and the same wavelength coverage (from *Herschel* to ALMA band 7). This eliminates most of the systematics, and allows us to robustly quantify the relative difference.

In an attempt to understand these possibly conflicting observations, we can try to draw a comparison to other known galaxies. Here, we have selected the few known massive galaxies at high-redshift with a similar dust temperature of ~ 30 K, and with a known spectroscopic redshift; the latter being required to measure the dust temperature accurately. The two most famous examples include GN20 (Daddi et al. 2009; Hodge et al. 2013; Tan et al. 2014) and HDF850.1 (Walter et al. 2012). These two galaxies have a half-light radius of 6–7 kpc, which is an order of magnitude larger than *Hyde*. This may on its own explain their lower-than-average dust temperature, although recent results suggest this could also be caused by optically-thick dust, at least for GN20 (Cortzen et al. 2020). To our knowledge, the only other reported instance of a low dust temperature combined with a compact geometry can be found in the interacting pair SGP 38326 (Oteo et al. 2016); unfortunately this system only has poor wavelength coverage, which renders the temperature uncertain ($T_{\text{dust}} = 33\text{--}55$ K). Although comparably massive to *Hyde*, these galaxies are also an order of magnitude brighter, hence would fit in the traditional picture of a starburst galaxy.

To date, the case of *Hyde* therefore appears to be unique. Its low dust temperature could correspond to a softer radiation field (or, equivalently, to a low star formation efficiency), which would match the results in this paper. However, as for GN20, it could also be caused by optically-thick dust. Disentangling the two possibilities would require an alternative measurement of the temperature (e.g., using the CO or [C I] line ratios; Cortzen et al. 2020).

6. Conclusions

We have obtained new ALMA observations toward the galaxy *Hyde* to observe its [N II]₂₀₅ emission. We showed that the line emission is more concentrated in [N II] than in [C II], which implies a gradient in the ionized gas fraction. We found the center of the galaxy to be predominantly ionized, which suggests that young stars in their birth-clouds are not the dominant ionization source in the galaxy center, leaving the room for other sources such as intermediate-age-stars, or an AGN. In contrast, the outskirts are dominated by neutral gas which suggests the [C II] emission there is mostly associated with star-forming regions. Using these new insights, we obtained an updated estimate of the galaxy's star formation rate, placing it securely below the galaxy main sequence, hence possibly in transition to quiescence.

These results point toward an ongoing inside-out quenching mechanism, and show that this process may start before the galaxy has fully expelled or consumed its gas reservoirs, as demonstrated by the strong obscuration of this galaxy. This confirms the importance of studying the properties of this galaxy in greater details, to better understand the process of quenching at high redshifts.

Further observations of this system would enable a better understanding of the state of the gas, and of the mechanism responsible for quenching, in particular whether the galaxy hosts an AGN or not. New observations with ALMA are already scheduled, including high-resolution imaging of the dust continuum emission to establish the morphology and geometry of the galaxy, as well as [C I] for measuring the cold gas mass to study the interplay between the ionized and neutral gas reservoirs. Otherwise, the *James Webb Space Telescope* (*JWST*) could be used to look at the rest-frame optical lines, in particular H α , to better constrain the SFR of the galaxy and obtain an estimate of the gas-phase metallicity. Unfortunately, even *JWST* may not have a sharp enough resolution to resolve the optical emission line profiles, which would then require larger telescopes, such as the European Extremely Large Telescope (E-ELT). Finally, deep low-frequency radio observations, for example with the LOw Frequency ARray (LOFAR), would reveal even moderate AGN activity and thus determine the role of AGNs in the high-redshift quenching process.

To go beyond the case study of this single object and determine the rate of occurrence of these features, similar observations of a sample of massive, obscured galaxies would be required. As illustrated here, finding transitioning galaxies is not an easy task, especially if most of them are strongly obscured and would not be included in the *H*-band selected catalogs produced from *Hubble* imaging. Samples of *Spitzer*-IRAC-selected galaxies may be more adequate (see Caputi et al. 2015; Wang et al. 2016; Wang et al. 2019), and in the near future *JWST* will hopefully open this search to fainter, more numerous galaxies.

Acknowledgements. The authors want to thank the two anonymous referees for their comments that clearly improved the consistency, thoroughness, and overall quality of this paper. CS would like to thank Leah Morabito and

Ian Heywood for sharing their expertise on interferometric imaging. Most of the numerical analysis conducted in this work has been performed using `vi-f`, a free and open source C++ library for fast and robust numerical astrophysics (<http://cschreib.github.io/vi-f/>). This paper makes use of the following ALMA data: ADS/JAO.ALMA#2015.A.00026.S and ADS/JAO.ALMA#2018.1.00216.S. ALMA is a partnership of ESO (representing its member states), NSF (USA) and NINS (Japan), together with NRC (Canada) and NSC and ASIAA (Taiwan) and KASI (Republic of Korea), in cooperation with the Republic of Chile. The Joint ALMA Observatory is operated by ESO, AUI/NRAO and NAOJ. GJK acknowledges the support of the Australian Research Council through the Discovery Project DP170103470.

References

- Belli, S., Newman, A. B., & Ellis, R. S. 2019, *ApJ*, 874, 17
 Béthermin, M., Daddi, E., Magdis, G., et al. 2015, *A&A*, 573, A113
 Birnboim, Y. & Dekel, A. 2003, *MNRAS*, 345, 349
 Bourne, N., Dunlop, J. S., Merlin, E., et al. 2017, *MNRAS*, 467, 1360
 Bruzual, G. & Charlot, S. 2003, *MNRAS*, 344, 1000
 Calzetti, D., Armus, L., Bohlin, R. C., et al. 2000, *ApJ*, 533, 682
 Capak, P. L., Carilli, C., Jones, G., et al. 2015, *Nature*, 522, 455
 Caputi, K. I., Ilbert, O., Laigle, C., et al. 2015, *ApJ*, 810, 73
 Casey, C. M. 2012, *MNRAS*, 425, 3094
 Chabrier, G. 2003, *PASP*, 115, 763
 Charlot, S. & Fall, S. M. 2000, *ApJ*, 539, 718
 Civano, F., Marchesi, S., Comastri, A., et al. 2016, *ApJ*, 819, 62
 Cortzen, I., Magdis, G. E., Valentino, F., et al. 2020, *A&A*, 634, L14, 2002.02974
 Costa, T., Rosdahl, J., Sijacki, D., & Haehnelt, M. G. 2018, *MNRAS*, 479, 2079
 Croton, D. J., Springel, V., White, S. D. M., et al. 2006, *MNRAS*, 365, 11
 da Cunha, E., Charlot, S., & Elbaz, D. 2008, *MNRAS*, 388, 1595
 da Cunha, E., Charmandaris, V., Díaz-Santos, T., et al. 2010, *A&A*, 523, A78
 Daddi, E., Dannerbauer, H., Stern, D., et al. 2009, *ApJ*, 694, 1517
 Davé, R., Thompson, R., & Hopkins, P. F. 2016, *MNRAS*, 462, 3265
 Díaz-Santos, T., Armus, L., Charmandaris, V., et al. 2017, *ApJ*, 846, 32
 Elbaz, D., Dickinson, M., Hwang, H. S., et al. 2011, *A&A*, 533, 119
 Förster Schreiber, N. M., Genzel, R., Newman, S. F., et al. 2014, *ApJ*, 787, 38
 Gabor, J. M. & Davé, R. 2012, *MNRAS*, 427, 1816
 Genzel, R., Förster Schreiber, N. M., Rosario, D., et al. 2014, *ApJ*, 796, 7
 Glazebrook, K., Schreiber, C., Labbé, I., et al. 2017, *Nature*, 544, 71
 Gobat, R., Strazzullo, V., Daddi, E., et al. 2012, *ApJ*, 759, L44
 Gullberg, B., Smail, I., Swinbank, A. M., et al. 2019, *MNRAS*, 490, 4956
 Gullberg, B., Swinbank, A. M., Smail, I., et al. 2018, *ApJ*, 859, 12
 Hodge, J. A., Carilli, C. L., Walter, F., Daddi, E., & Riechers, D. 2013, *ApJ*, 776, 22
 Hodge, J. A., Swinbank, A. M., Simpson, J. M., et al. 2016, *ApJ*, 833, 103
 Howell, J. H., Armus, L., Mazzarella, J. M., et al. 2010, *ApJ*, 715, 572
 Ikarashi, S., Caputi, K. I., Ohta, K., et al. 2017, *ApJ*, 849, L36
 Kamenetzky, J., Rangwala, N., Glenn, J., Maloney, P. R., & Conley, A. 2016, *ApJ*, 829, 93
 Kriek, M., van Dokkum, P. G., Labbé, I., et al. 2009, *ApJ*, 700, 221
 Labbé, I., Huang, J., Franx, M., et al. 2005, *ApJ*, 624, L81
 Lapham, R. C., Young, L. M., & Crocker, A. 2017, *ApJ*, 840, 51
 Leslie, S. K., Schinnerer, E., Liu, D., et al. 2020, *ApJ*, 899, 58, 2006.13937
 Magnelli, B., Lutz, D., Saintonge, A., et al. 2014, *A&A*, 561, 86
 Martig, M., Bournaud, F., Teyssier, R., & Dekel, A. 2009, *ApJ*, 707, 250
 Martí-Vidal, I., Pérez-Torres, M. A., & Lovanov, A. P. 2012, *A&A*, 541, A135
 Merlin, E., Fontana, A., Castellano, M., et al. 2018, *MNRAS*, 473, 2098
 Oberst, T. E., Parshley, S. C., Stacey, G. J., et al. 2006, *ApJ*, 652, L125
 Oteo, I., Ivison, R. J., Dunne, L., et al. 2016, *ApJ*, 827, 34
 Parkin, T. J., Wilson, C. D., Schirm, M. R. P., et al. 2013, *ApJ*, 776, 65
 Pasquali, A., Ferreras, I., Panagia, N., et al. 2006, *ApJ*, 636, 115
 Pavesi, R., Riechers, D. A., Capak, P. L., et al. 2016, *ApJ*, 832, 151
 Pearson, W. J., Wang, L., Hurley, P. D., et al. 2018, *A&A*, 615, A146, 1804.03482
 Peng, Y., Maiolino, R., & Cochrane, R. 2015, *Nature*, 521, 192
 Rosenberg, M. J. F., van der Werf, P. P., Aalto, S., et al. 2015, *ApJ*, 801, 72
 Rybak, M., Calistro Rivera, G., Hodge, J. A., et al. 2019, *ApJ*, 876, 112
 Schreiber, C., Elbaz, D., Pannella, M., et al. 2018a, *A&A*, 609, A30
 Schreiber, C., Glazebrook, K., Nanayakkara, T., et al. 2018b, *A&A*, 618, A85
 Schreiber, C., Labbé, I., Glazebrook, K., et al. 2018c, *A&A*, 611, A22
 Schreiber, C., Pannella, M., Leiton, R., et al. 2017, *A&A*, 599, A134
 Silk, J. & Rees, M. J. 1998, *A&A*, 331, L1
 Simpson, J. M., Smail, I., Wang, W., et al. 2017, *ApJ*, 844, L10
 Smolčić, V., Delvecchio, I., Zamorani, G., et al. 2017, *A&A*, 602, A2
 Speagle, J. S., Steinhardt, C. L., Capak, P. L., & Silverman, J. D. 2014, *ApJS*, 214, 15
 Straatman, C. M. S., Labbé, I., Spitler, L. R., et al. 2015, *ApJ*, 808, L29
 Tadaki, K.-i., Iono, D., Hatsukade, B., et al. 2019, *ApJ*, 876, 1
 Tan, Q., Daddi, E., Magdis, G., et al. 2014, *A&A*, 569, A98
 Tomczak, A. R., Quadri, R. F., Tran, K. H., et al. 2016, *ApJ*, 817, 118
 Walter, F., Decarli, R., Carilli, C., et al. 2012, *Nature*, 486, 233
 Wang, T., Elbaz, D., Schreiber, C., et al. 2016, *ApJ*, 816, 84
 Wang, T., Schreiber, C., Elbaz, D., et al. 2019, *Nature*, 572, 211, 1908.02372
 Wellons, S., Torrey, P., Ma, C., et al. 2015, *MNRAS*, 449, 361
 Zhang, Z., Ivison, R. J., George, R. D., et al. 2018, *MNRAS*, 481, 59
 Zhao, Y., Lu, N., Xu, C. K., et al. 2016, *ApJ*, 819, 69

Appendix A: Simulated noise maps

To estimate parameter uncertainties, the method we used in this paper is to repeat our measurement on mock images, created by duplicating the true image and adding different realizations of noise to it. In this experiment, the observed noisy image and our best-fit model become the “truth”, and we can determine how far away our fitting procedure is from this “truth” on each mock image. The strength of this method is that it requires no assumption on the model, and no analytical calculations; the parameter probability distribution can be extracted straight away by gathering the fits to all mock images. Furthermore, it automatically takes care of parameters that are strongly correlated with one another, as well as correlated noise. The difficulty is that it strongly depends on the quality of the noise that is injected on top of the “truth” image. For uncertainties to be accurate, the noise needs to have the exact same amplitude and covariance matrix as in the real data.

While in non-interferometric data sets the noise can usually be assumed uncorrelated, this is not the case with our ALMA data; the noise observed in the [C II] line map, and its auto-correlation function (ACF), are shown in Fig. A.1 (left). One can observe that the ACF has a structure very similar to the dirty beam (see Fig. A.1, center), and we demonstrate in Appendix B that this is indeed expected when the data is imaged with natural weighting.

This makes it easy to reproduce noise with a similar covariance. Indeed, if a uniform uncorrelated random noise is convolved with a two-dimensional kernel K , its ACF will be $C = K \otimes K$, or in the Fourier domain, $\hat{C} = \hat{K}^2$ (where “hat” symbolizes the Fourier transform). If we set $C = P$, where P is the dirty beam (see Appendix B), we have:

$$\hat{K} = \sqrt{\hat{P}}. \quad (\text{A.1})$$

Since the Fourier transform introduces aliasing artifacts on the edges of K , it is desirable to use a dirty beam image P that is significantly larger than the final dimensions of the noise map. This can be achieved by padding. To stabilize this further, and since we are mostly interested in preserving the “core” of the covariance, we smoothed out the transition to the edges of the image by multiplying the dirty beam image with a broad Gaussian of FWHM $4''$.

We used this empirical convolution kernel K to generate new correlated noise realizations, shown in Fig. A.1 (right). The resulting noise maps are visually similar to the real noise map, and the core of the covariance is well reproduced. We note that, since the dirty beam side lobes in our data have a relatively low amplitude, very similar results could have been achieved by simply using a Gaussian kernel of FWHM equal to half that of the dirty beam.

Finally, to set the noise RMS, we simply measured the RMS on the real image and on the simulated noise images, and rescaled the simulated images to match the real observed RMS. Because our target is at the phase center, we did not correct for the primary beam attenuation in the real (or simulated) images, hence the RMS was constant across the entire image.

Appendix B: Auto-correlation function of image-domain noise

For a generic interferometer, the observed dirty image is the Fourier transform of the weighted and sampled sky visibility:

$$I_D(x, y) \equiv \int A(u, v) W(u, v) V(u, v) \exp[-2\pi i (u x + v y)] du dv, \quad (\text{B.1})$$

where A is the visibility sampling of the interferometer (a sum of delta functions), W is the imaging weight, and V is the complex visibility. Without loss of generality and for simplicity of notation, we drop the time, frequency, and polarization dependence in all quantities.

The dirty beam is the response to a point source, which has constant visibilities in the Fourier domain:

$$P(x, y) \propto \int A(u, v) W(u, v) \exp[-2\pi i (u x + v y)] du dv. \quad (\text{B.2})$$

If we define the inverse Fourier transform of the dirty image

$$V_D(u, v) \equiv \int I_D(x, y) \exp[2\pi i (u x + v y)] du dv \quad (\text{B.3})$$

$$= A(u, v) W(u, v) V(u, v), \quad (\text{B.4})$$

then by definition, the image auto-correlation function (ACF) is

$$C(x, y) \equiv \int V_D(u, v) \bar{V}_D(u, v) \exp[-2\pi i (u x + v y)] du dv \quad (\text{B.5})$$

$$= \int A(u, v) W(u, v)^2 V(u, v) \bar{V}(u, v) \exp[-2\pi i (u x + v y)] du dv, \quad (\text{B.6})$$

where we used $A^2 = A$ since it is a sum of delta functions.

To study the ACF of the noise, we now assume that the sky contains no source and the visibilities are therefore only made of noise of amplitude $\sigma(u, v)$, which we assume is uncorrelated. Thus, $\langle V \bar{V} \rangle = \sigma^2$. If we now compute the expectation value of the ACF, we get

$$\langle C(x, y) \rangle = \int A(u, v) W(u, v)^2 \sigma(u, v)^2 \exp[-2\pi i (u x + v y)] du dv. \quad (\text{B.7})$$

In the case of natural weighting, the weights are chosen as $W \propto 1/\sigma^2$. Injecting this into the above equation, we finally get

$$\langle C(x, y) \rangle \propto \int A(u, v) W(u, v) \exp[-2\pi i (u x + v y)] du dv \quad (\text{B.8})$$

$$\propto P(x, y). \quad (\text{B.9})$$

Therefore, for images produced with natural weighting, the expectation value of the noise ACF is the dirty beam.

In practice the match is not perfect (see Fig. A.1), and this can be explained by at least three possible causes: a) we can only measure the ACF on a finite-sized image, so our estimate of the image ACF is noisy; b) the σ that enters in the definition of W is only an estimate of the true visibility uncertainty; and c) the noise in the visibilities can itself be correlated. These potential issues seem to have only a moderate impact however; as demonstrated in Fig. A.1 (middle, top), in our data the core of the ACF is extremely well reproduced by the dirty beam.




Fig. A.1. Observed and simulated noise maps (top) and their respective auto-correlation functions (bottom). The real [C II] map is displayed in the top left corner, with the source masked in the center. For comparison, we also display the [C II] dirty beam in the middle of the bottom panel. All the images on each row are displayed with the same color bar. The auto-correlation functions were rescaled prior to display to a peak value of unity. For easier comparison, we also show a circular average of the covariance in shown in the central column of the first row, with the observed (black) and simulated (green dashed) auto-correlation functions, and the dirty-beam profile (blue).

Appendix C: Likelihood for correlated Gaussian noise

The general formulation of the likelihood for Gaussian correlated noise is:

$$\mathcal{L} \propto \exp\left(-\frac{\chi^2}{2}\right), \quad \text{with} \quad \chi^2 = (\mathbf{d} - \mathbf{m})^t \Sigma^{-1} (\mathbf{d} - \mathbf{m}), \quad (\text{C.1})$$

where \mathbf{d} is a vector of observed data, \mathbf{m} is a vector of model data, and Σ is the data covariance matrix (symmetric and positive definite). For our data (see Appendix B), we have:

$$\Sigma_{ij} = \sigma^2 P(x_i - x_j, y_i - y_j), \quad (\text{C.2})$$

where σ is the image RMS, (x_i, y_i) are the image coordinates of pixel i , and $P(\delta x, \delta y)$ is the value of the dirty beam at an offset position $(\delta x, \delta y)$ from the peak.

In practice, inverting Σ is numerically unstable. The best way to evaluate it is to perform a Singular Value Decomposition (SVD) such that $\Sigma = U \lambda U^t$, where U is a unitary matrix ($U^{-1} = U^t$) and λ is a diagonal matrix whose entries are the singular values λ_i sorted by decreasing value. Then we have

$\Sigma^{-1} = U \lambda^{-1} U^t$. If we define $L^{-1} = (1/\sqrt{\lambda}) U^t$, then Eq. C.1 can be rewritten in a simpler form:

$$\chi^2 = (\tilde{\mathbf{d}} - \tilde{\mathbf{m}})^t (\tilde{\mathbf{d}} - \tilde{\mathbf{m}}) \quad (\text{C.3})$$

where $\tilde{\mathbf{d}} = L^{-1} \mathbf{d}$ and $\tilde{\mathbf{m}} = L^{-1} \mathbf{m}$. These can be seen as the “de-correlated” observation and model vectors, respectively.

The SVD is not sufficient to make the computation of Σ^{-1} or L^{-1} stable. In all cases, the matrix λ contains small entries that cannot be inverted safely. A workaround is to truncate the matrix $(1/\sqrt{\lambda})$ by setting to zero the inverse of those singular values that are smaller than some chosen threshold. The resulting matrix inverse is then a “pseudo-inverse” (i.e., an approximation of the true inverse), but it is usually better behaved.

The method we adopted to choose the threshold is the following. We defined the normalized singular values $\bar{\lambda}_i = \lambda_i/\lambda_0 \leq 1$, and we created a logarithmic grid of 100 threshold values v ranging from $\min(\bar{\lambda})$ to 0.1, such that all $\bar{\lambda}_i < v$ were removed from the inverse. For each value of v , we evaluated the corresponding L^{-1} , and used it to “de-correlate” one of our noise realization from the simulations (this noisy image did not contain any source). We computed the covariance matrix of the resulting image, normalized it to unit diagonal, and defined the metric k as

the sum the absolute value of entries with $|x_i - x_j| \leq 1$ and $|y_i - y_j| \leq 1$. We then picked the value of ν which minimized k , or in other words, the value which produced an image with the lowest noise covariance. Experiments showed that k is large when ν is too small, as the matrix inverse is unstable and the resulting images are degraded. On the other hand, k is also large when ν is too large, as this leads to a poorer approximation of the inverse which leaves more correlated noise.

We note that, while Eq. C.1 is formally the correct expression to compute the likelihood for our data, we obtained very similar results when we used the simpler, standard expression for uncorrelated noise, $\chi^2 = \sum_i [(d_i - m_i)/\sigma]^2$. This is because we used Monte Carlo simulations to determine the confidence intervals rather than relying on the shape of the likelihood.

Appendix D: Input/output simulations

To test the accuracy of our profile-fitting method, we performed an input/output analysis, where we placed sources of known light profiles in simulated noise maps devoid of sources, and tried to recover their profile with our method. The input sources were modeled as two-component exponential disks, with a small “core” component and an “extended” component, as is assumed in our model. The size of the “core” component was chosen randomly in the range $r_{\text{core}} = 0.1\text{--}0.5$ kpc, and the size of the extended component was chosen randomly in the range $r_{\text{extended}} = r_{\text{core}} + 0.2\text{--}2.2$ kpc. The relative flux of these components was chosen uniformly between 0% (all flux in extended component) to 100% (all flux in core component). The axis ratio was chosen randomly between 0.4 to 1.0, and the position angle could take any value. Finally, the peak flux of each source was matched to the observed peak flux in our real images. As for our real data, we then computed the true fluxes inside and outside of a 0.5 kpc aperture to obtain the “central” and “outskirts” fluxes. We note here that, since the [N II] image has a sharp angular resolution, the step of fixing the peak flux in the simulations has a strong impact on the true flux in the central aperture, as can be seen in Fig. D.1 where the range of F_{center} is limited. This is less true for [C II], where the coarser PSF means that the central pixel can contain more flux from the extended component.

The mock sources were injected directly in the image domain, by convolving the source’s known light profile with the image dirty beam at $\times 9$ oversampling. To test the accuracy of this source injection method, we also used CASA to create simulated visibilities for each mock source using `simobserve` (without thermal noise), and produced corresponding dirty images using `clean` with the same (oversampled) cell size. The resulting mock images looked identical to those produced by convolution with the dirty beam. To quantify this, for each mock source we computed the second moment (R^2) of the images produced by the two methods (visibilities imaging, and image convolution). We found a maximum relative error on R^2 of only 3×10^{-5} , i.e., close to numerical noise, with no dependence on the size of the mock sources. This implies that injecting directly in the image domain is an accurate source injection method, which is also much faster than creating mock visibilities and imaging them. We caution that this property only holds because we work with dirty images; such images can be described as a Fourier transform of an incomplete (u, v) plane, but cleaned images cannot.

We repeated this procedure with 200 different mock sources, and for each source we replicated the noise level, noise covariance matrix, peak flux, and dirty beam of the [C II], [N II], and continuum maps separately, to test the impact of the varying S/N and angular resolutions encountered in this work. Our fitting pro-

cedure was then applied to each of these mock sources, and the recovered profiles were compared to the truth. Since these simulations are accurate mocks of the real images, they allow us to reproduce exactly the measurements we perform in this paper, so we can study any bias arising from our method.

The outcome of this analysis, illustrated in Fig. D.1 for [C II] and [N II], shows that our modeling method recovers the total, central, and outskirts fluxes with no detectable bias. The measurements are noisy, particularly for the central flux, however the observed scatter is correctly captured by the error bars, which were estimated exactly the same way as for the real data. To quantify this, for each observed quantity F shown in Fig. D.1, we computed the normalized residual $\delta_F = (F_{\text{obs}} - F_{\text{true}})/\sigma_{\text{obs}}$, where σ_{obs} is our estimated uncertainty. We then computed the fraction of simulated points with $|\delta_F| < 1$. If our uncertainties are perfectly accurate, and if the noise on the quantity F is roughly Gaussian, this fraction should be equal to 68%. Since we only use 200 points to compute these fractions, the expected statistical error on this number is about 3%. We find fractions for [C II] (resp. [N II]) of 67% (resp. 66%) for the total flux, 62% (resp. 61%) for the flux in the inner 0.5 kpc, and 61% (resp. 62%) for the flux in the outskirts. These values are exactly in the expected range for the total flux, but the fractions for the central and outskirts fluxes are slightly lower than 68%. For Gaussian noise, this would imply that the uncertainties are under-estimated by at most 10%; taking this small correction into account would not affect our conclusions.

Appendix E: Distribution in Monte Carlo simulation

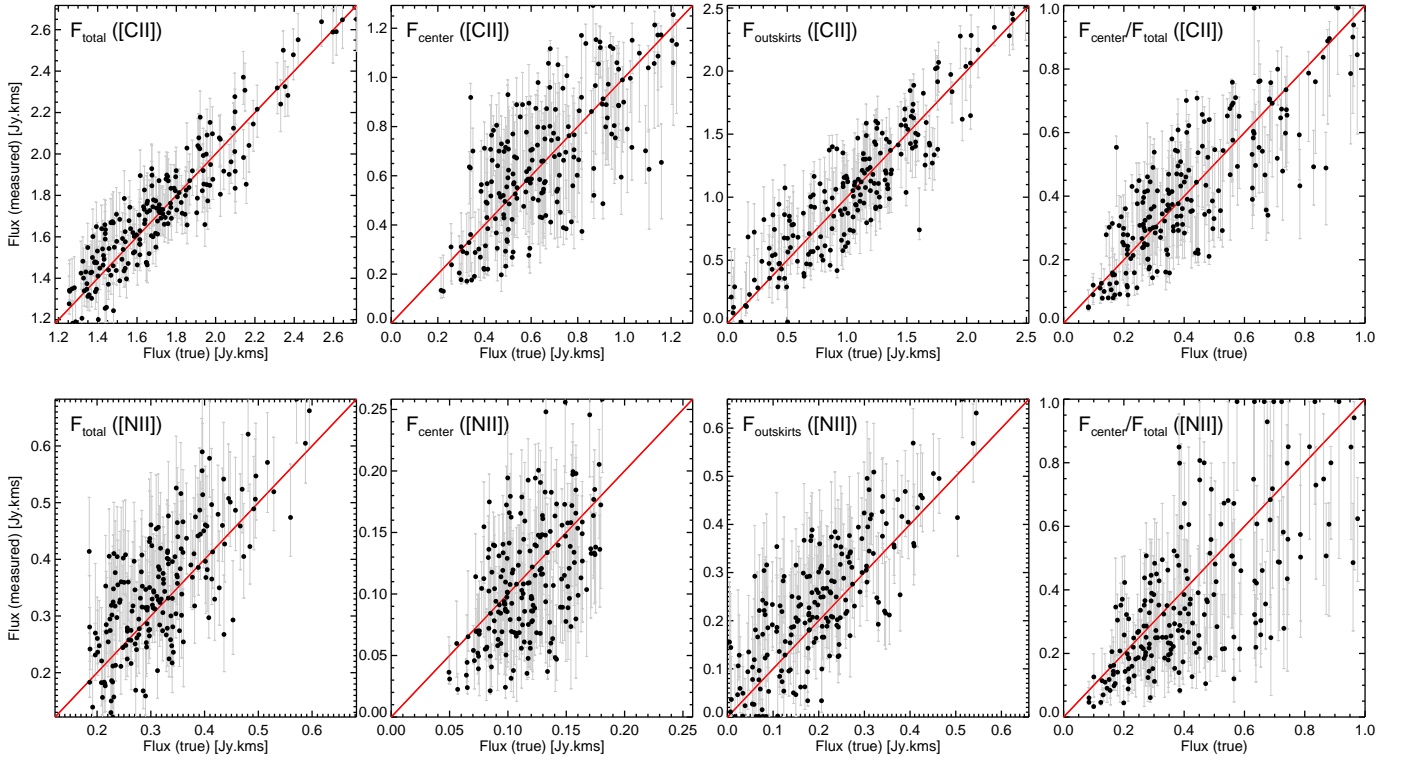


Fig. D.1. Outcome of the input/output analysis of our fitting procedure, modeling sources of known profiles in mock images. The simulations matching the [C II] and [N II] maps are shown at the top and bottom, respectively. From left to right, we show how our method recovers the total flux, the flux inside the center, the flux in the outskirts, and the ratio of center-to-total (with “center” and “outskirts” as defined in the main text). The red line is the line of perfect agreement.

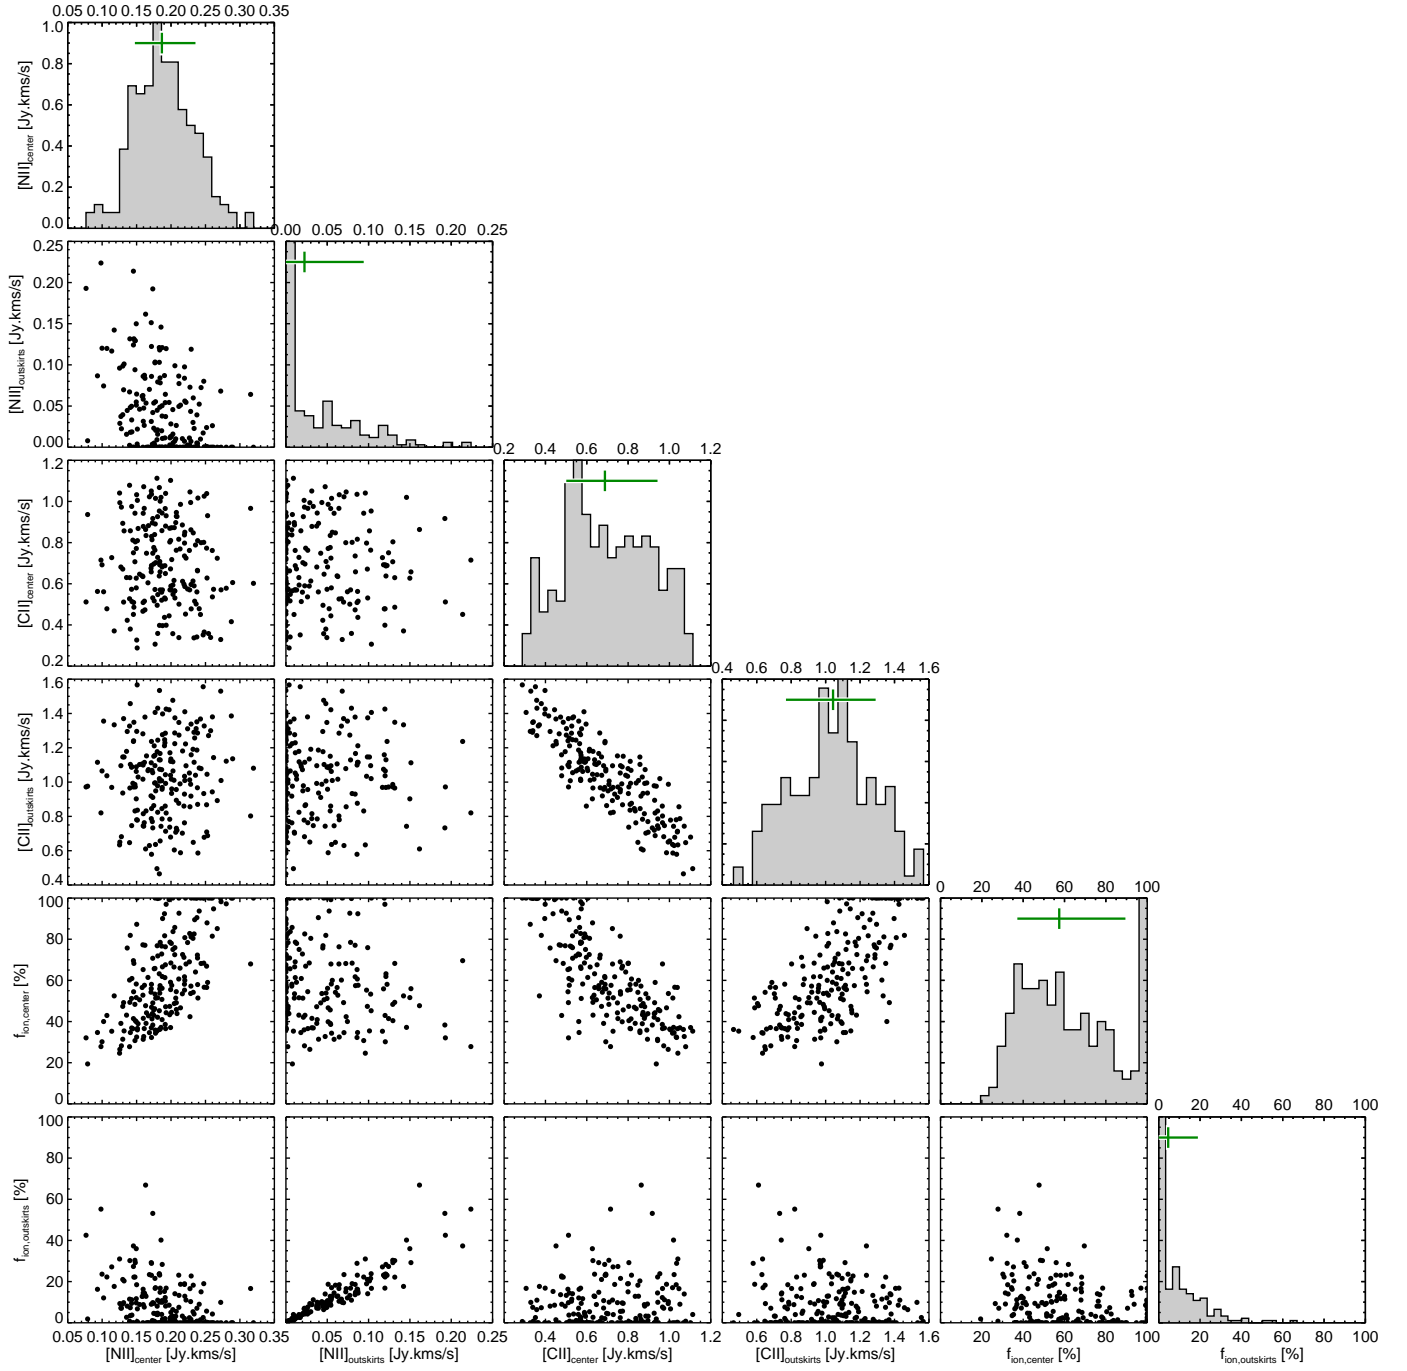


Fig. E.1. Joint distribution of the fluxes, flux ratios, and ionized gas fraction in the Monte Carlo simulations executed for the real images. As is customary, on the diagonal we show the histogram of each quantity; the y-axis for these plots is scaled automatically to include the maximum counts in a bin. The green horizontal bars correspond the 16th and 84th percentiles, which we report as uncertainties, and the horizontal green bar is the best-fit value.