



Quantifying stochasticity-driven uncertainties in H II region metallicities

Raghav Arora,^{1,2,3★} Mark R. Krumholz^{1,4★} and Christoph Federrath^{1,4}

¹Research School of Astronomy and Astrophysics, The Australian National University, Canberra, ACT 2611, Australia

²BITS, Pilani, Department of Physics, Goa Campus, Goa 403726, India

³Hamburger Sternwarte, Universität Hamburg, Gojenbergsweg 112, D-21029 Hamburg, Germany

⁴ARC Centre of Excellence for Astronomy in Three Dimensions (ASTRO-3D), Canberra, ACT 2611, Australia

Accepted 2021 September 14. Received 2021 September 11; in original form 2020 September 22

ABSTRACT

With the advent of integral field units (IFUs), surveys can now measure metallicities across the discs of nearby galaxies at scales $\lesssim 100$ pc. At such small scales, many of these regions contain too few stars to fully sample all possible stellar masses and evolutionary states, leading to stochastic fluctuations in the ionizing continuum. The impact of these fluctuations on the line diagnostics used to infer galaxy metallicities is poorly understood. In this paper, we quantify this impact for six most commonly used diagnostics. We generate stochastic stellar populations for galaxy patches with star formation rates varying over a factor of 1000, compute the nebular emission that results when these stars ionize gas at a wide range of densities, metallicities, and determine how much inferred metallicities vary with fluctuations in the driving stellar spectrum. We find that metallicities derived from diagnostics that measure multiple ionization states of their target elements (e.g. electron temperature methods) are weakly affected (variation < 0.1 dex), but that larger fluctuations (~ 0.4 dex) occur for diagnostics that depend on a single ionization state. Scatter in the inferred metallicity is generally largest at low star formation rate and metallicity, and is larger for more sensitive observations than for shallower ones. The main cause of the fluctuations is stochastic variation in the ionization state in the nebula in response to the absence of Wolf–Rayet stars, which dominate the production of $\gtrsim 2$ –3 Ryd photons. Our results quantify the trade-off between line brightness and diagnostic accuracy, and can be used to optimise observing strategies for future IFU campaigns.

Key words: methods: statistical – stars: Wolf–Rayet – ISM: abundances – H II regions – galaxies: abundances.

1 INTRODUCTION

The gas-phase metallicity of a galaxy is an important parameter for understanding fundamental properties concerning the build-up of chemical elements in galaxies (e.g. Somerville & Primack 1999; Nagamine et al. 2001; De Lucia, Kauffmann & White 2004; Davé, Finlator & Oppenheimer 2012). However, it is not just the mean metallicity that carries information. The metallicity distribution and gradient within galaxies also carry important information that is sensitive to factors such as the star formation history (Chiappini, Matteucci & Romano 2001), mass accretion and mass-loss via winds (Wang et al. 2019), and galaxy environment (Kewley et al. 2010; Rupke, Kewley & Barnes 2010; Rosa et al. 2014; Torres-Flores et al. 2014; Taylor & Kobayashi 2017). Beyond simple gradients, the azimuthal variation of the metallicity also provides important information regarding various processes that drive gas mixing and inflows, including bar-driven radial mixing (Di Matteo et al. 2013), spiral-arm driven systematic streaming motions (Grand et al. 2016; Sánchez-Menguiano et al. 2016), kiloparsec-scale mixing-induced dilutions due to spiral density wave passage (Ho et al. 2017), thermal and gravitational instabilities (Yang & Krumholz 2012; Petit et al.

2015), and interstellar turbulence (de Avillez & Mac Low 2002; Krumholz & Ting 2018).

Despite the importance of quantifying the metallicity distribution in galaxies, detailed studies of the azimuthal variations in the metallicity have so far been limited (e.g. Kreckel et al. 2016, 2019, 2020; Vogt et al. 2017; Ho et al. 2018). This is mainly due to the observational challenges in measuring nebular emission lines across a substantial portion of the galactic disc with sufficient spatial resolution (\lesssim a few hundred pc). However, with the advent of the wide-field optical integral field unit (IFU) spectrographs, we have gained the ability to obtain high-resolution line emission maps of large samples of nearby galaxies at high resolution. Current IFU-based mapping campaigns include The Physics at High Angular Resolution in Nearby Galaxies¹ (PHANGS; Schinnerer et al. 2019), the Calar Alto Legacy Integral Field Area Survey (CALIFA; Sánchez et al. 2012a), Mapping Nearby Galaxies at Apache Point Observatory (Bundy et al. 2015), and the Sydney-AAO Multi-Object Integral-field spectrograph survey (Bryant et al. 2015; Scott et al. 2018). The highest resolution of these studies, the PHANGS-MUSE (Multi Unit Spectroscopic Explorer) survey, resolves < 100 pc scales in tens of nearby galaxies (Kreckel et al. 2019, 2020). These studies have turned up some interesting anomalies. For example, PHANGS-

* E-mail: raghav.sofos@gmail.com (RA); mark.krumholz@anu.edu.au (MRK)

¹<http://www.phangs.org>

MUSE reports systematic disagreement between the metallicities inferred using different diagnostics coupled with a scatter of up to 0.3 dex, depending on the diagnostics being compared. This level of uncertainty has the potential to compromise the science we can extract from resolved metallicity maps, and thus it is crucial that we understand its origin. In order to do so, we must understand the assumptions that are baked into metallicity diagnostics.

There are many metallicity diagnostics, each with its own set of limitations and caveats, and with large systematic disagreements (~ 1 dex) amongst them (Stasińska 2005; Kewley & Ellison 2008; Peimbert, Peimbert & Delgado-Inglada 2017; Kewley, Nicholls & Sutherland 2019). The most direct methods of measuring metallicity rely on either recombination lines or auroral lines (also known as direct or electron-temperature methods), and these techniques rely primarily on atomic physics, with relatively few additional assumptions or steps. Unfortunately, because the lines they exploit are weak, these methods are often unavailable. For this reason, studies of metallicity beyond the Milky Way and its nearest neighbours often rely on either theoretical or empirically calibrated diagnostics. The former are derived using models for photoionization, stellar evolution, and stellar atmospheres, while the latter rely on empirical correlations between the ratios of strong lines and metallicities established by other means.

For both theoretical and strong line methods, one possible source of uncertainty is variation in the shape of the ionizing continuum driving the H II region. This dependence is clear in the case of theoretical diagnostics, which require the ionizing spectrum as an input, and where variations in the shape of that spectrum (for example due to variations in the age of the stellar population; Byler et al. 2017) can induce corresponding variations in diagnostic line ratios. Even for empirical calibrations, however, variation in spectral shape is a possible concern. While the level of uncertainty in these diagnostics can be assessed by examining the scatter between true metallicity and diagnostic line ratio, and this in principle should include uncertainties due to variation in the spectral shape of the driving source, in practice this method of estimating the uncertainty is valid only to the extent that the sample used to derive the calibration covers the full range of spectral variation that might be encountered. This is generally not the case, since the parent samples used to derive empirical correlations are necessarily biased toward those bright enough to allow determination of the metallicity using auroral lines. As an example, several of the empirical metallicity diagnostics we discuss in this paper have been calibrated from H II region samples drawn mainly from the CALIFA survey (Sánchez et al. 2012a), for which the median luminosity in the H α line is $10^{39.0}$ erg s $^{-1}$, and fifth percentile luminosity (i.e. 95 per cent of the catalogue is brighter than this value) is $10^{37.4}$ erg s $^{-1}$ (Sánchez et al. 2012b). By contrast, PHANGS-MUSE includes H II regions with luminosities down to 10^{36} erg s $^{-1}$ (Kreckel et al. 2019), and the next generation of IFUs will go even deeper.

One potential source of variation in ionizing spectral shape is stochastic sampling of the initial mass function (IMF). While this effect is relatively small for observations of an entire galaxy, with the advent of IFU surveys, this situation has changed. To see why, a simple back of the envelope estimate suffices. Consider a patch of a galaxy forming stars at a rate per unit area comparable to that of the solar neighbourhood ($\Sigma_* \approx 2.5 \times 10^{-3}$ M $_{\odot}$ Myr $^{-1}$ pc $^{-2}$; Fuchs, Jahreiß & Flynn 2009) that is resolved into pixels (or spaxels) that are 100 pc \times 100 pc in size, comparable to those available in the highest resolution published studies. The expected star formation rate (SFR) in the spaxel is $\dot{M}_* = 25$ M $_{\odot}$ Myr $^{-1}$, and over the ≈ 4 Myr lifetime of an O star, ≈ 100 M $_{\odot}$ worth of new stars should be formed. By contrast,

for a star cluster mass function $dN/dM \propto M^{-2}$ from 10^2 to 10^5 M $_{\odot}$, as commonly observed in spiral galaxies (Krumholz, McKee & Bland-Hawthorn 2019), the median cluster mass is $10^{3.5}$ M $_{\odot}$. This is also the approximate cluster mass required to fully sample the stellar IMF (da Silva, Fumagalli & Krumholz 2012).² Thus, a typical 100 pc \times 100 pc galaxy patch similar to the solar neighbourhood has only a $100/10^{3.5} \sim 1$ in 30 chance of containing a cluster young and massive enough to fully sample all the possible ionizing stars that go into producing a fully sampled stellar spectrum. Clearly for such an observation, stochastic variation of the stellar population is a first-order effect that cannot be ignored. We do not expect to approach full sampling until the combination of spaxel size and SFR is a factor of ~ 30 larger.

Investigations of this issue thus far have been limited. da Silva, Fumagalli & Krumholz (2014) considered the effects of stochastic variation on inferences of the SFR from H α in low-luminosity systems, and showed that errors $\gtrsim 1$ dex are common for systems with true SFRs $\lesssim 10^{-4}$ M $_{\odot}$ yr $^{-1}$. Krumholz et al. (2015b) show that a similar-sized stochastic uncertainty affects photometric estimation of the masses and ages of low-mass star clusters, and Ashworth et al. (2017) extended this analysis to studies of the IMF in star clusters. In the most direct precursor to our work here, Paalvast & Brinchmann (2017) use Monte Carlo simulations to study how stochastic variation affects determinations of metallicity via nebular lines in dwarf galaxies with low SFRs. They find significant stochastic errors for galaxies with SFRs $\lesssim 10^{-3}$ M $_{\odot}$ yr $^{-1}$, an effect they attribute to undersampling of the IMF. In this work, we extend their analysis to the case of IFU observations, using a substantially larger simulation library and considering a larger range of metallicity diagnostics; we also compare the effects of undersampling of the mass distribution and the age distribution of the stellar population and show that the latter is more important than the former, a distinction not made in Paalvast & Brinchmann (2017). Both of these effects contribute when a stellar population does not sample the IMF completely (da Silva et al. 2012).

The remainder of this paper is as follows: Section 2 details the method we use to produce a synthetic library of stochastically sampled stellar populations and corresponding synthetic line emission data. We analyse the properties of this library, and in particular the uncertainties in metallicity inference resulting from stochastic sampling, in Section 3, and discuss their physical origins in Section 4. We summarize our conclusions in Section 5.

2 METHODS

Our overall goal in this paper is to characterise how much inferred metallicities, derived from a sample of six commonly used metallicity diagnostics, fluctuate in response to stochastic variations in the ionizing spectrum of the stellar population that drives the H II region. Our method for reaching this goal is to generate a library of ionizing spectra from a stochastically sampled stellar population, predict the nebular line fluxes these would produce, and compare the metallicities derived from these line fluxes to the fully sampled case and the true inputs. In this section, we describe the procedure we use to generate the required library and the metallicity diagnostics that

²The definition of the mass required to ‘fully sample’ the IMF is somewhat arbitrary, since the more rare a particularly stellar mass or type is, the larger the total cluster mass that will be required to sample it. The mass of $10^{3.5}$ M $_{\odot}$ that we have quoted is obtained using the condition that the expected number of stars with initial mass > 100 exceeds unity.

we apply to it. Section 2.1 describes how we generate our synthetic spectra. Section 2.2 describes our physical model of H II regions, and how we simulate nebular emission lines within the context of this model. We describe the metallicity diagnostics and their calibrations in Section 2.3.

2.1 SLUG simulations

We are interested in generating stochastically sampled realizations of the ionizing spectra produced by stellar populations that reside within a spatial pixel (spaxel), whose size is of order $\sim 100 \text{ pc} \times 100 \text{ pc}$ in a nearby galaxy. For this purpose, we make use of the Stochastically Lighting Up Galaxies code (SLUG; da Silva et al. 2012, 2014; Krumholz et al. 2015a), and for the remainder of this section we describe the underlying physical model we adopt, and that SLUG implements.

We assume that the average SFR in our region of interest has been constant for longer than the lifetime of a massive star, and that stars within our sample pixel form as one or more discrete ‘clusters’, each of which can be described as a simple stellar population (SSP). The masses of these clusters are drawn from the observed cluster mass function (CMF), which we parametrize as $dN/dM \propto 1/M^2$ from $M = 10^2$ to $10^5 M_\odot$ (Krumholz et al. 2019). Fumagalli, da Silva & Krumholz (2011) show that this hypothesis leads to predicted distributions of H α and FUV luminosity that are in very good agreement with observations. For practical computational purposes, we only include in our calculation clusters with ages from 0 to $t_{\text{max}} = 20 \text{ Myr}$, since SSPs older than this produce negligible ionizing luminosity. The total number of clusters in the stellar population is controlled by the SFR used in the simulation, which is determined by the star formation per unit area and the pixel size. We adopt a fiducial pixel size of $100 \times 100 \text{ pc}^2$, and consider SFRs per unit area $\dot{\Sigma}_* = (0.1 - 100) \dot{\Sigma}_{*,0}$, where $\dot{\Sigma}_{*,0} = 2.5 \times 10^{-3} M_\odot \text{ Myr}^{-1} \text{ pc}^{-2}$ is the SFR in the solar neighbourhood (Fuchs et al. 2009). Thus, we have a range of SFR from 2.5 to $2500 M_\odot \text{ Myr}^{-1}$.

We generate our sample stellar populations using SLUG’s ‘Poisson’ sampling method (for full details see Krumholz et al. 2015a). The procedure operates in three steps. We first decide how many clusters to draw. For our chosen CMF, the mean cluster mass is $\langle M \rangle = 690 M_\odot$, so the expected number of clusters with age $< t_{\text{max}}$ in our sample pixel is $\langle N \rangle = t_{\text{max}} \text{SFR} / \langle M \rangle$. We draw our number of clusters N from a Poisson distribution with expectation value $\langle N \rangle$. The second step is to assign each cluster an age; since the SFR is by assumption constant, all ages are equally likely, and we therefore draw the age of each cluster from a uniform probability distribution from 0 to t_{max} . The third and final step is to draw individual stellar masses within each cluster; our procedure for this step mirrors that we use to draw cluster masses. Specifically, we compute the mean stellar mass $\langle m \rangle$ from the IMF; for this paper, we adopt a Chabrier (2005) IMF from 0.08 to $120 M_\odot$. We then compute the expected number of stars in the cluster, $\langle n_* \rangle = M / \langle m \rangle$, where M is the cluster mass, and draw the actual number of stars n_* from a Poisson distribution with expectation value $\langle n_* \rangle$. Finally, we draw the mass of each individual star from the IMF. At the end of this procedure, we have a complete description of the mass and age of every star in the stellar population.

Our next step is to calculate the spectrum that these stars produce. We compute the physical properties of every star in our population (including the possibility that some stars will already have completed their lives) using the MIST v1 stellar evolution tracks (Choi et al. 2016), and we use SLUG’s ‘STARBURST99’ stellar atmospheres method to generate the stellar atmospheres and spectra (Leitherer et al. 1999a); this method uses WM-Basic model atmospheres for

O stars (Pauldrach, Hoffmann & Lennon 2001), CMFGen atmospheres for Wolf–Rayet stars (Hillier & Miller 1998) and Lejeune, Cuisinier & Buser (1997), mainly based on Kurucz models, for the remainder of the stellar population.

We perform 2×10^5 runs for each value of metallicity in the range $\log Z/Z_\odot = -1.5$ to 0 in steps of 0.5 dex; here, Z is the metallicity normalized to the solar metallicity [$12 + \log(\text{O}/\text{H}) = 8.69$]. These runs are sampled uniformly in $\log(\text{SFR}/\text{SFR}_0)$ from $\text{SFR}/\text{SFR}_0 = 0.1$ to 100, where $\text{SFR}_0 = (2.5 \times 10^{-3} M_\odot \text{ Myr}^{-1} \text{ pc}^{-2}) \times (100 \text{ pc})^2$, i.e. SFR_0 is the SFR of a $100 \text{ pc} \times 100 \text{ pc}$ pixel with an SFR per unit area equal to that of the solar neighbourhood. We supplement these with an additional $\approx 2.5 \times 10$ runs at each metallicity with SFRs sampled uniformly in $\log(\text{SFR}/\text{SFR}_0)$ over the range $\text{SFR}/\text{SFR}_0 = 0.1 - 1$. We require these extra samples at low SFR because a significant fraction of the runs yield line luminosities that fall below the observability cuts we impose below; the additional cases ensure that our statistical errors remain small despite this. Note that our method includes stochastic variation in both the age of the stellar population and the masses of individual stars. As one increases the input SFR, both these distributions become fully sampled.

To provide a baseline for comparison, we also carry out two additional SLUG runs where we disable part or all of the stochastic sampling. In the first of these, to which we refer below as the fully sampled case. Note that this is an analytical calculation and is different from the stochastic runs with $M \geq 10^{3.5} M_\odot$ that are approaching a fully sampled population. We generate a spectrum for a system forming stars continuously at a rate of $0.001 M_\odot \text{ yr}^{-1}$ with no stochasticity in either stellar age or mass. That is, in this case we assume that the light is produced by a stellar population containing an infinite number of stars with ages uniformly distributed from 0 to 20 Myr, and masses distributed following the Chabrier (2005) IMF; in this case the total SFR enters the problem only as a normalization factor on the overall luminosity. Run in this mode, SLUG operates identically to a conventional, non-stochastic stellar population synthesis code, and thus the resulting spectrum is derived using the same assumptions usually made in generating theoretical photoionization models.

In our second comparison case, we consider a stellar population that is continuously (i.e. non-stochastically) sampled in stellar mass, but where all the stars have the same age, i.e. where the stars form a single simple stellar population (SSP); we view this as being representative of a case where the stellar population is fully sampled in mass, but not in age. In this case, we adopt a stellar population mass of $10^4 M_\odot$ and generate spectra for this population at ages from 0 to 10 Myr in steps of 10^4 yr . Run in this mode, SLUG acts identically to a conventional stellar population synthesis code simulating the case of star formation in an instantaneous burst. We refer to this case below as the SSP case.

We generate nebular spectra and inferred metallicities for both the fully sampled and SSP cases using the same procedure as that used for our main library of SLUG simulations. We describe this procedure next.

2.2 CLOUDY simulation

The radiation from the stellar population is incident on a surrounding H II region. We use the C17 version of the CLOUDY radiative transfer code, last described by Ferland et al. (2017), to calculate the resultant nebular spectrum. In our calculation, we adopt the photospheric solar abundances ($Z = 0.0134$, $X = 0.7381$) for the nebula for the case $Z = Z_\odot$ from Asplund et al. (2009). For other metallicity values, we adjust the abundances of each element using the non-linear and

element-specific empirical scaling described in Nicholls et al. (2017). In addition to abundances, the nebular emission will depend on the geometry and structure of the H II region. To specify these choices, we run CLOUDY in its spherical, isobaric mode, for which the input parameters to the calculation are the inner radius r_0 and number density n_0 of the hydrogen nuclei at that radius; CLOUDY then computes the structure of the remainder of the photoionised region by requiring that the pressure be constant. However, while the parameter r_0 is required computationally, it is not necessarily very descriptive of the physical properties of the H II region. We therefore choose to define an alternative parametrization, based on a greatly simplified H II region model, from which we can derive r_0 .

In our simplified model, the H II region surrounding the stellar population is assumed to be isobaric and isothermal, which makes the density uniform as well, i.e. $n_0 = n_{\text{II}}$, where n_{II} is the mean density. All hydrogen present in the region is assumed to be fully ionized, helium is singly ionized and the radiation pressure is negligible. The H II region is a spherical shell, with an inner radius r_0 and an outer radius r_1 . The inner radius is set by the presence of a bubble of shocked stellar wind material at a temperature $\gtrsim 10^6$ K that produces negligible optical or UV emission, and the outer radius by the location where all the ionizing photons have been absorbed.

Under the assumptions made above, and further assuming photoionization equilibrium, the ionizing photon luminosity at a distance r from the stellar population located at the centre of the H II region is given by (Draine 2011)

$$Q(r) = Q(\text{H}^0) \left[1 - \left(\frac{r}{r_s} \right)^2 + \left(\frac{r_0}{r_s} \right)^3 \right], \quad (1)$$

where $Q(\text{H}^0)$ is the hydrogen-ionizing luminosity of the source, defined as

$$Q(\text{H}^0) = \int_0^{\lambda_{\text{H}^0}} \frac{f_\lambda}{hc/\lambda} d\lambda, \quad (2)$$

$\lambda_{\text{H}^0} \approx 912 \text{ \AA}$ is the cut-off wavelength corresponding to the 13.6 eV ionization potential of H^0 , f_λ is the ionizing spectrum of the stellar population that is incident on the H II region, and r_s is the Strömgren radius, given by

$$r_s = \left(\frac{3Q(\text{H}^0)}{4\pi\alpha_B f_e n_{\text{II}}^2} \right)^{1/3}. \quad (3)$$

Here, $\alpha_B = 2.59 \times 10^{-13} \text{ cm}^3 \text{ s}^{-1}$ is the case B recombination coefficient, $f_e = 1.1$ is the abundance of free electrons per H nucleus, and n_{II} is the number density of hydrogen present in the region. The values of α_B and f_e we have chosen are appropriate for temperatures of 10^4 K and for a region where He is singly ionized. Since r_s depends only on n_{II} and $Q(\text{H}^0)$, equation (1) allows us to solve for the inner radius r_0 in terms of the outer radius r_1 , defined implicitly by the condition $Q(r_1) = 0$. However, rather than using the outer radius as our alternative parameter, it is more physically meaningful to use the wind parameter defined by Yeh & Matzner (2012) as

$$\Omega = \frac{r_0^3}{r_1^3 - r_0^3}. \quad (4)$$

This quantity is the ratio of the volume occupied by the wind gas to the volume occupied by the ionised gas. Values of Ω greater than unity signify that winds are dynamically important. For given values of Ω , n_{II} , and $Q(\text{H}^0)$, we can solve for r_0 by simultaneously solving the system consisting of equation (4) and equation (1) [with $r = r_1$ and $Q(r_1) = 0$]. We therefore use $n_{\text{II}} = n_0$ and Ω as our parameters, in place of n_0 and r_0 .

Yeh & Matzner (2012) show that the choice $\Omega \approx 0.3$ produces good fits to the infrared forbidden lines observed in nearly all nearby H II regions, so we adopt this value. However, there are a range of densities present in observed H II regions, and thus we also treat n_{II} as a stochastic variable. We draw n_{II} from a distribution $P(n_{\text{II}}) \propto 1/n_{\text{II}}$ (i.e. uniform probability in $\log n_{\text{II}}$) from $n_{\text{II}} = 10\text{--}10^4 \text{ cm}^{-3}$.

2.3 Diagnostics

The final step in our method is to calculate the gas-phase oxygen abundances [$\log Z \equiv 12 + \log(\text{O}/\text{H})$] of H II regions. For this work, we select six commonly used optical line diagnostics that span a range from highly accurate but reliant on faint lines, to methods that use only bright lines, but are less accurate. Following the classification scheme used in Kewley et al. (2019), we use one auroral line method (the direct T_e method), two theoretical optical diagnostics (Kobulnicky and Kewley R_{23} and D16), and three empirical diagnostics (O3N2, PP04–N2, and S-calibration). Recently, there has been some interest in using rest-frame ultraviolet emission lines (Byler et al. 2020) for high-redshift galaxies ($z \geq 4$), as the optical lines are expected to be redshifted out of observational windows for ground-based optical and near-infrared facilities. However, we omit these diagnostics because they are still experimental and because the focus of our work is local galaxies. We also omit recombination-line methods, since the lines are very weak and are only detected in nearby Galactic regions. For all of the diagnostics we discuss, our goal is not to provide a comprehensive review, simply to collect the relevant formulae and the uncertainties we use into one place for reader's convenience. We refer readers to the original papers referenced below for further detail.

2.3.1 Direct T_e method

The direct electron temperature (T_e) method, which was suggested in Garnett (1992), has been the traditional approach to calculating metallicities. However, this method relies on weak auroral lines, which are rarely detected in galaxies with $12 + \log(\text{O}/\text{H}) > 8.7$ (Peimbert 1967; Rubin 1969; Stasińska 2005); indeed, in their study of eight nearby galaxies, Kreckel et al. (2019) report an order of magnitude difference between the number of H II regions that have observable auroral lines and the number with detectable optical strong lines. None the less, because this method is the most reliable, and because this method does not rely on making any implicit or explicit assumptions about the shape of the ionizing spectrum driving the H II region, we start with it to provide a baseline.

The T_e method relies on the calculation of the electron temperature T_e and density of the H II region. We first estimate the electron temperature as (their equation 4; Nicholls, Kewley & Sutherland 2020)

$$\log t = \frac{7.2939x + 3.5363}{-0.0074x^3 - 0.1221x^2 + 1.6298x + 1}, \quad (5)$$

where $t = T_e/10^4 \text{ K}$ and

$$x = \log_{10} \left(\frac{[\text{O III}]\lambda 4363}{[\text{O III}]\lambda 5007 + [\text{O III}]\lambda 4959} \right). \quad (6)$$

We then use the calculated T_e to estimate the abundances of the first two ionization stages of oxygen (their equation 3; Izotov et al. 2006)

$$12 + \log(\text{O}^+/\text{H}^+) = \log \left(\frac{[\text{O II}]\lambda 3727}{\text{H}\beta} \right) + 5.961 + \frac{1.676}{t} - 0.40 \log t - 0.034t + \log(1 + 1.35y), \quad (7)$$

$$12 + \log(\text{O}^{2+}/\text{H}^+) = \log\left(\frac{[\text{O II}]\lambda 4959, 5007}{\text{H}\beta}\right) + 6.200 + \frac{1.251}{t} - 0.55 \log t - 0.014t, \quad (8)$$

where $y = 10^{-4}n_e t^{-0.5}$ and n_e is the electron density of the H II region. Since n_e for our models is by construction $\ll 10^4 \text{ cm}^{-3}$, we take $y = 0$. The total oxygen abundance is then calculated as

$$\left(\frac{\text{O}}{\text{H}}\right) = \left(\frac{\text{O}^+}{\text{H}^+}\right) + \left(\frac{\text{O}^{2+}}{\text{H}^+}\right), \quad (9)$$

neglecting the contribution from the O^0 and the O^{3+} ions, which are negligible for H II regions produced by massive stars (Berg et al. 2016).

The uncertainties in this method arise due to the intrinsic inaccuracies in the atomic data, correction for unseen stages of ionization, and the collisional cross-section data (Kewley et al. 2019). While the total level of uncertainty is not known, it is likely small.

2.3.2 Kobulnicky and Kewley R_{23}

The Kobulnicky and Kewley R_{23} diagnostic (KKR₂₃ hereafter), similar to direct T_e , relies on the auroral lines for estimating the metallicity. However, unlike T_e , this diagnostic is calibrated theoretically on a combination of stellar population synthesis models and photoionization models. Kewley & Dopita (2002) provide an iterative scheme for the R_{23} diagnostic (Pagel et al. 1979), which accounts for its dependence on the ionization parameter, and was further improved by Kobulnicky & Kewley (2004). We use the Kobulnicky & Kewley method for our experiments. In this method, we first calculate the ionization parameter as (their equation 13; Kobulnicky & Kewley 2004)

$$\log q = \frac{32.81 - 1.153y^2 + a[12 + \log(\text{O}/\text{H})]}{4.603 - 0.3119y - 0.163y^2 + b[12 + \log(\text{O}/\text{H})]}, \quad (10)$$

where

$$a = 0.1444y^2 - 0.025y - 3.396,$$

$$b = 0.02037y^2 + 0.0271y - 0.48,$$

and $y = \log \text{O}_{32}$ is defined as

$$\log \text{O}_{32} = \log\left(\frac{[\text{O III}]\lambda 4959 + [\text{O III}]\lambda 5007}{[\text{O II}]\lambda 3727}\right). \quad (11)$$

As we can see, equation (10) needs an initial guess for the value of metallicity in the first iteration. For this purpose, we use the input metallicity as the initial guess, because this does not affect the final result, but speeds up convergence. From the ionization parameter, we estimate the metallicity depending upon the branch into which the initial value falls. For the lower branch, $12 + \log(\text{O}/\text{H}) < 8.4$, we have (their equations 16 and 17; Kobulnicky & Kewley 2004)

$$12 + \log(\text{O}/\text{H})_{\text{lower}} = 9.40 + 4.65x - 3.17x^2 - \log q(0.272 + 0.547x - 0.513x^2), \quad (12)$$

while for the upper branch, $\log(\text{O}/\text{H}) + 12 > 8.4$, we have

$$12 + \log(\text{O}/\text{H})_{\text{upper}} = 9.72 - 0.777x - 0.951x^2 - 0.072x^3 - 0.811x^4 - \log q(0.0737 - 0.0713x - 0.141x^2 + 0.0373x^3 - 0.058x^4), \quad (13)$$

where $x = \log R_{23}$ is given by

$$\log R_{23} = \log\left(\frac{[\text{O II}]\lambda 3727 + [\text{O III}]\lambda 4959, 5007}{\text{H}\beta}\right). \quad (14)$$

This yields a new estimate of O/H, which we can in turn use in equation (10) to generate a new estimate of $\log q$, which then yields a new O/H estimate. We therefore iterate between computing q from equation (10) and O/H from equation (12) or equation (13) until the values of q and O/H converge, which usually takes about 2–3 iterations.

This diagnostic has been calibrated on some physical assumptions of the H II regions – e.g. the dust depletion factor, simplified geometry, uniform composition, and isobaric conditions. Since we do not know how much a real H II region will deviate from these assumptions, it is difficult to take this into account in the calibration.

2.3.3 S-calibration (Scal)

This is a relatively new empirical diagnostic, calibrated by Pilyugin & Grebel (2016) from a sample of 313 H II regions with T_e -based metallicities, with the goal of still correcting for the ionization parameter as in the T_e and KKR₂₃ methods, but using the [S II] $\lambda 6717, 6731$ doublet in place of the weak [O II] $\lambda 3727$ auroral line for this purpose – hence, the name S calibration (Scal) – since the availability of the latter line is often the limiting factor in metallicity measurements using T_e or KKR₂₃. Scal has been used by Kreckel et al. (2019) in place of the T_e method because of the unavailability of the [O II] $\lambda 3727$ line in MUSE observations.

For this diagnostic, we first use N II lines to decide whether a particular H II region falls in the upper or the lower branch. For the upper branch, $\log N_2 > -0.6$, we have (their equation 6; Pilyugin & Grebel 2016)

$$12 + \log(\text{O}/\text{H})_{\text{upper}} = 8.424 + 0.030 \log(R_3/S_2) + 0.751 \log(N_2) + \log(S_2) \times (-0.349 + 0.182 \log(R_3/S_2) + 0.508 \log(N_2)), \quad (15)$$

and for the lower branch, $\log N_2 < -0.6$, we have (Pilyugin & Grebel 2016); their

$$12 + \log(\text{O}/\text{H})_{\text{lower}} = 8.072 + 0.789 \log(R_3/S_2) + 0.726 \log(N_2) + \log(S_2) \times (1.069 - 0.170 \log(R_3/S_2) + 0.022 \log(N_2)), \quad (16)$$

where N_2 , S_2 , and R_3 are defined as

$$N_2 = [\text{N II}]\lambda 6548, 6584/\text{H}\beta,$$

$$S_2 = [\text{S II}]\lambda 6717, 6731/\text{H}\beta,$$

$$R_3 = [\text{O III}]\lambda 4959, 5007/\text{H}\beta.$$

Pilyugin & Grebel find systematic errors of ≤ 0.1 dex in the Scal metallicities when compared with metallicities derived using the T_e method, and their calibration sample covers a relatively large metallicity range, from $\log(\text{O}/\text{H}) \approx 7-8.8$.

2.3.4 O3N2

The O3N2 diagnostic, proposed in Pettini & Pagel (2004), uses the line ratios and [N II] $\lambda 6584/\text{H}\alpha$. It has been calibrated on the metallicities derived by the direct T_e calculations and photoionization models. The diagnostic was originally formulated for high-redshift galaxies. However, it has also been used for H II regions in nearby galaxies (Kreckel et al. 2019). It is somewhat less preferred than the other diagnostics since it has been shown to have a high dependence upon the ionization parameter (Kewley et al. 2019), but we include it because it is commonly used in cases where the auroral lines are too faint to be of use. We use the empirical calibration given in

Marino et al. (2013), which the authors derive from a sample of 309 H II regions with metallicities obtained using the T_e method. The diagnostic is given by Marino et al. (2013, their equation 2)

$$12 + \log(\text{O}/\text{H}) = 8.533 - 0.214 \times \log \left(\frac{[\text{O III}] \lambda 5007 / \text{H} \beta}{[\text{N II}] \lambda 6584 / \text{H} \alpha} \right). \quad (17)$$

One caveat of this method is the limited sample available to the authors for the calibration, for $\text{O3N2} = -1.1$ to 1.7 , which gives a metallicity range of $8.169 \leq 12 + \log(\text{O}/\text{H}) \leq 8.768$. The authors report a range of 0.18 dex for 68 per cent of the observations.

2.3.5 PP04-N2

Pettini & Pagel (2004) also propose an optical-line method that just uses the line ratio $[\text{N II}] \lambda 6584 / \text{H} \alpha$. The availability of the $[\text{N II}] \lambda 6584$ line in moderate-redshift samples (Byler et al. 2020), especially nearby galaxies, makes it a preferred choice over many other diagnostics. We use the improved calibration of the diagnostic described in Marino et al. (2013), where they have used T_e -based metallicities derived from 452 H II regions for calibration of the coefficients. The metallicity estimate for this diagnostic is given by (their equation 4; Marino et al. 2013)

$$12 + \log(\text{O}/\text{H}) = 0.462 \log \left(\frac{[\text{N II}]}{\text{H} \alpha} \right) + 8.743. \quad (18)$$

Marino et al. find that 68 per cent of the observations lie within 0.16 dex of the T_e measurements. This calibration too has been derived from a limited sample of H II regions and thus is only valid for the metallicity range $8.00 \leq 12 + \log(\text{O}/\text{H}) \leq 8.65$. For the rest of the metallicity range, we use the original calibration of Pettini & Pagel (2004, their equation 1)

$$12 + \log(\text{O}/\text{H}) = 0.57 \log \left(\frac{[\text{N II}]}{\text{H} \alpha} \right) + 8.90, \quad (19)$$

For this metallicity range, 68 per cent of the measurements have been reported to lie within an interval of 0.18 dex.

2.3.6 D16

Recently, Dopita et al. (2016) proposed a new theoretical diagnostic that uses lines that are not only easy to detect, but are also close to each other so that reddening corrections can also be neglected. The calibration is derived from a grid of photoionization models run with the MAPPINGS code (Sutherland et al. 2018), which are computed using STARBURST99 (Leitherer et al. 1999b) stellar spectra (identical to what we use here) to describe the shape of the driving radiation field. The lines used in this diagnostic are $\text{H} \alpha$, $[\text{N II}] \lambda 6584$, and the $[\text{S II}] \lambda 6717, 31$ doublet. Defining

$$y = \log \left(\frac{[\text{N II}]}{[\text{S II}]} \right) + 0.264 \log \left(\frac{[\text{N II}]}{\text{H} \alpha} \right), \quad (20)$$

we get the metallicity estimate for this diagnostic, using (their equation 3; Sutherland et al. 2018)

$$12 + \log(\text{O}/\text{H}) = 8.77 + y + 0.45(y + 0.3)^5. \quad (21)$$

Like other theoretical diagnostics (e.g. KKR_{23}), this diagnostic suffers from the limitations of the physical assumptions of the model. Dopita et al. (2016) find that the variation in the dust depletion factor is the major source of systematic uncertainties that can be as high as ± 0.12 dex for non-solar metallicities.

2.3.7 Observability limits

A final consideration that we must apply to all our metallicity diagnostics is observability: especially for low SFRs, the total ionizing luminosity of our stellar population may be low enough that some or all of the lines used in various diagnostics will be unobservably faint. Since we are interested in understanding observations, we must include this effect. Of course the observability limits will depend on details such as the integration time, the size of the aperture, etc., and thus are not the same from one survey to another. Except where otherwise stated, for all the analysis presented here we adopt limits that roughly match the sensitivity of PHANGS-MUSE. Specifically, we only calculate an inferred metallicity for H II regions for which the luminosity of the $\text{H} \alpha$ line is $> 10^{36} \text{ erg s}^{-1}$, which is similar to the minimum $\text{H} \alpha$ luminosity reported by Kreckel et al. (2019).³ We also require that the other lines used in diagnostics be bright enough relative to this so that they too will be observable. To this end, we adopt a minimum line luminosity of $10^{-1.5}$ times the $\text{H} \beta$ flux for oxygen lines, comparable to the limiting line ratios found in Pilyugin, Grebel & Mattsson (2012) and Kreckel et al. (2019), and $10^{-2.5}$ times the $\text{H} \alpha$ line for nitrogen lines, comparable to the limiting line ratio found in Marino et al. (2013). If one of our sample H II regions has a line luminosity below our observability cut, we do not compute any of the metallicity diagnostics that rely on that line.

There are some important implications of our choice of sensitivity limit. One is that, at low SFRs, the sample of H II regions for which we derive an inferred metallicity is not the same for all diagnostics. For example, in our lowest SFR bin (see below), roughly five times as many H II regions have detectable $[\text{N II}] \lambda 6584$ lines (and thus can have their metallicities inferred using the D16 diagnostic) as have detectable $[\text{O III}] \lambda 4364$ lines (as required for KKR_{23} or T_e). Compared to the cases with detectable $[\text{N II}] \lambda 6584$, those that also have detectable $[\text{O III}] \lambda 4364$ are present in cases where the stochastic draw happened to produce particularly massive stars, or stars at stages of stellar evolution where they produce the hard photons needed to produce O III emission. Secondly, we also observe the well-known effect (Peimbert 1967; Rubin 1969; Stasińska 2005) that the number of H II regions that have an observable $[\text{O III}] \lambda 4364$ line falls tremendously at high metallicities. This effect is also seen for the $[\text{N II}] \lambda 6584$ line at low metallicities. We emphasise that these are real physical effects that can potentially account for some of the observed systematic differences between metallicities inferred with different diagnostics. We therefore make no attempt to correct for or remove these effects in our analysis.

In order to verify that our sampled H II regions have reasonable line luminosities, and that they would not be mistaken for active galactic nucleus (AGN) in an observed sample, in Fig. 1 we plot the BPT diagram for our simulation library. This diagram includes only those H II regions that pass our observability screen, as described above. We see that the locus occupied by the H II regions in this diagram is below the Kewley et al. (2001) starburst line with a few exceptions, and generally populates a region consistent with observed H II regions.

³It is worth noting that a more accurate treatment of observability limits would be to phrase these limits in terms of line equivalent widths (EWs), rather than absolute luminosities, since in practice what limits the detectability of a line is its brightness relative to the background stellar continuum. However, our modelling pipeline does not enable us to calculate the EWs and so we use the absolute line luminosity as a rough proxy.




Figure 1. The BPT diagrams of our models after applying the observational cuts of $H\beta > 10^{36} \text{ erg s}^{-1}$, $[O\text{ III}]\lambda 5007/H\beta > 10^{-1.5}$ and $[N\text{ II}]\lambda 6584/H\alpha > 10^{-2.5}$; to make the plot readable, we show only ≈ 5 per cent of our realizations. The four different colours show four different values of the input metallicity, as indicated in the legend. The dotted line represents the extreme starburst line from Kewley et al. (2001). Only a small fraction of H II regions cross over to the AGN region above the dotted line.

3 RESULTS

Our library of H II regions models varies with three input parameters:

- (i) SFR (SFR), which varies continuously from 0.1 to $100 \times \text{SFR}_0$, where $\text{SFR}_0 = 25 \text{ M}_\odot \text{ Myr}^{-1}$ is our fiducial SFR for a $100 \times 100 \text{ pc}^2$ spaxel forming stars at the same rate as the solar neighbourhood.
- (ii) Metallicity, which takes on four discrete values, $\log Z/Z_\odot = -1.5, -1.0, -0.5$, and 0.
- (iii) Ionised gas number density, which varies continuously from 10 to 10^4 cm^{-3} .

For each model case, our output consists of the metallicity that we would infer from each of our six sample line diagnostics at 20 Myr from the time our model began forming stars. In the following sections, we discuss how the level of stochastic fluctuation in inferred metallicity varies in response to changes in the SFR, true metallicity, and survey sensitivity. For the purposes of this analysis we limit our sample to H II regions within which the ionized gas number density $n_{\text{H}} = 10\text{--}100 \text{ cm}^{-3}$, the range most commonly observed for optically selected H II regions (Tremblin et al. 2014). We show the effect of having an ionized gas number density far from the commonly observed values in Appendix A.

3.1 Variation with SFR and metallicity

We begin by examining the effects of stochasticity at different metallicities and SFRs (SFRs). To this end, we select from our library all models with a particular discrete metallicity and density $n_{\text{H}} = 10\text{--}100 \text{ cm}^{-3}$, and bin the models in SFR/SFR_0 ; our bins are 0.5 dex wide, and run from $\text{SFR}/\text{SFR}_0 = 0.1$ to 100. For each bin, we characterise the level of stochastic fluctuation in the inferred metallicity by calculating half the 16th to 84th percentile range (equivalent to the 1σ dispersion for a Gaussian distribution) in metallicity, $12 + \log(\text{O}/\text{H})$, derived from each diagnostic, including in this calculation only those H II regions in our sample library that are bright enough to pass our observability cuts. We show the result in Fig. 2. Each panel of Fig. 2 is for a particular diagnostic, and within each panel, the position along the vertical axis shown by dotted lines, marks the true input metallicity of our models and position along the

horizontal axis shows SFR/SFR_0 . The centre of each point lies on the median of the metallicity values and the size of the point and colour indicates the spread in the metallicity values, as indicated in the colourbar, given by a diagnostic at a particular input metallicity value and bin of SFR/SFR_0 .

First, we see that the metallicity, as derived from different diagnostics, have their own systematic offsets. These have been studied before in detail (Stasińska 2005; Kewley & Ellison 2008; Peimbert et al. 2017; Kewley et al. 2019). Unless stated otherwise, we will be focusing on the uncertainties of these diagnostics. We see that, as expected, there are almost no fluctuations in the metallicity derived from the T_e method, independent of the SFR or metallicity. This is as expected: the T_e method is based on fundamental atomic physics, and relies on essentially no assumptions about the shape of the ionizing spectrum. Thus, this diagnostic offers a useful control for our experiment.

For all other diagnostics, we see a generic trend that, at almost all true metallicities, the dispersion in the inferred metallicity increases as we lower the SFR, though the amount of increase varies by diagnostic and by true metallicity – for example, KKR_{23} shows only a weak dependence on SFR, whereas D16 and Scal show a very strong dependence, particularly at low metallicity; the others are intermediate.

We can interpret the results as follows. A larger value of SFR/SFR_0 simply entails more stars in our fiducial patch and a better sampling of the IMF and stellar ages. A value of $\text{SFR}/\text{SFR}_0 \sim 100$, on average, will give us $\approx 10^4 \text{ M}_\odot$ worth of stars in 4 Myr, more than the $\sim 10^{3.5} \text{ M}_\odot$ that we take to be the mass required for a full sampling of the IMF (da Silva et al. 2012). Thus, for our highest SFR bin, the 16th to 84th percentile range is in almost all cases < 0.1 dex. [The primary exception to this statement is KKR_{23} at a true metallicity of $12 + \log(\text{O}/\text{H}) = 8.19$, which has a large dispersion because this metallicity lies close to the break between the upper and lower branches of the diagnostic.] Stochasticity has little effect in this limit, and the uncertainty budget is instead dominated by other sources.

By contrast, for $\text{SFR}/\text{SFR}_0 < 10$, the expected mass is $< 10^3 \text{ M}_\odot$, and we are far from the fully sampled regime. Depending upon the diagnostic and the true metallicity value, this can induce large or small fluctuations in the inferred metallicity. In the worst cases, e.g. Scal at low SFR and metallicity, the dispersion can be ~ 0.3 dex, substantially greater than the $\approx 0.1\text{--}0.15$ dex nominal uncertainty. We emphasize that these large dispersions, even for empirically calibrated diagnostics like Scal, are not surprising in light of our discussion in Section 1: the calibration sample for Scal comes from CALIFA, and therefore consists of H II regions $\gtrsim 30$ times more luminous than the sensitivity limit that we are now exploring. Stochasticity is expected to have a larger effect in this regime.

3.2 Variation with survey sensitivity

Our fiducial choice of sensitivity limit is typical of modern IFU surveys such as PHANGS-MUSE (Section 2.3.7). However, it is also of interest to ask whether we expect to find similar results for a deeper survey, as might be possible with a next-generation IFU on a 20–30 m class telescope, e.g. GMTIFS on the GMT or HARMONI on the E-ELT. To this end, we repeat our analysis in the previous section, but using an absolute cut-off on the $H\alpha$ line luminosity of $10^{35} \text{ erg s}^{-1}$, i.e. a factor of 10 deeper than current limits; we do not change the observability cut-offs for other lines, which are expressed as brightness relative to the $H\beta$ line, rather than absolute luminosity.

We show the results of this analysis in Fig. 3. We can see from the figure that the trend of stochastic variation with metallicity and




Figure 4. The joint distribution of metallicities and $Q(X)$ values of stochastic runs ($\log Z_{\text{stoch}}$) with true metallicity $\log Z/Z_{\odot} = -1.0$ and density in the range $n_{\text{II}} = 10\text{--}100\text{ cm}^{-3}$. Each panel is for a different diagnostic, as indicated, and colours are proportional to the logarithm of the number of realizations in a particular bin in the xy -plane. The value on the vertical axis of each panel is the difference between the metallicity predicted by that diagnostic for a particular stochastic realization, Z_{stoch} , and the true metallicity. The horizontal axis is the ratio of luminosity in photons capable of ionizing species X , $Q(X)$, to luminosity of H^0 -ionizing photons, $Q(\text{H}^0)$, and we plot the value in a particular stochastic run relative to the value $[Q(X)/Q(\text{H}^0)]_{\text{fs}}$ for the case of a fully sampled stellar population. The species X for which we plot is O^+ for the T_e , Scal, KKR₂₃, and O3N2 diagnostics, and N^0 for PP04-N2 and D16.

SFR remains largely the same as that shown in Fig. 2, but that the absolute spread in inferred metallicity increases for some diagnostics. In general, T_e , KKR₂₃, and O3N2 are unaffected by the observability limits. By contrast, D16 and PP04-N2 show a significant increase (>0.1 dex) in dispersion at solar metallicity, such that the dispersion at higher metallicity becomes comparable to that at low metallicities. For Scal, on the other hand, the already high dispersion at lower metallicities and SFRs increases even more, reaching 0.4 dex for the lowest SFR bin at $12 + \log(\text{O}/\text{H}) = 7.89$. This is again much larger than the stated $\approx 0.1\text{--}0.15$ dex uncertainty of the diagnostic.

4 DISCUSSION

So far, we have explored how metallicities inferred from optical line diagnostics vary in response to changes in the degree of stochasticity (SFR/SFR_0), the true metallicity of the underlying nebula (Z_{in}), and the sensitivity of the observation. Depending on these parameters and on the choice of diagnostic, the stochasticity-induced spread can be as small as ~ 0.01 dex, or as large as ~ 0.4 dex. We now seek to understand the physical mechanisms that ultimately drive these fluctuations. We also try to separate out the two complementary effects that make up the stochasticity: inadequate sampling of the mass distribution and the age distribution of the stellar population.

4.1 Mechanisms for stochastic variation

The primary driver of stochastic variation in the inferred metallicity is variation in the ionizing spectral shape, and in particular variations

in the spectral shape that induce changes in the ionization states of the metal atoms upon which metallicity diagnostics rely. To demonstrate this, we examine the quantity $Q(X)$, as defined for $X = \text{H}^0$ in equation (2): the photon luminosity of the stellar population, counting only photons that are capable of ionizing X to X^+ . For the choice of species X , we consider the highest energy ionization transition relevant to the lines used by a given diagnostic – thus for T_e , KKR₂₃, Scal, and O3N2 we compute $Q(\text{O}^+)$, since these diagnostics all use one or both components of the $[\text{O III}] \lambda 4959, 5007$ doublet, while for PP04-N2 and D16 we compute $Q(\text{N}^0)$, since these use $[\text{N II}] \lambda 6584$. We show the joint distribution of $Q(X)$ and inferred metallicity from our stochastic runs, $\log Z_{\text{stoch}}$, in Fig. 4. In this figure, we show one panel for each diagnostic, and in these panels, the vertical axis shows $\log Z_{\text{stoch}} - \log Z_{\text{in}}$, the difference between the inferred and true metallicities, and the horizontal axis shows $\log [Q(X)/Q(\text{H}^0)]_{\text{stoch}} - \log [Q(X)/Q(\text{H}^0)]_{\text{fs}}$, where ‘fs’ refers to the fully sampled case; thus a value of 0 on the horizontal axis corresponds to a spectrum where the ratio of X -ionizing photons to H^0 ionizing photons is the same as in the fully sampled case, and values larger (smaller) than zero indicate an excess (deficit) of such photons. Here, we show the subset of stochastic runs with $(\log Z_{\text{in}}/Z_{\odot} = -1.0)$ and with densities in the range $n_{\text{II}} = 10\text{--}100\text{ cm}^{-3}$. Note that this figure includes only runs where the emission would be observable, as described by our observability criteria in Section 2.3.7.

We see that most of our runs are positioned close to zero in the horizontal direction, indicating, as expected, that the ratio $Q(X)/Q(\text{H}^0)$ in our stochastic runs scatters about the value expected for the fully sampled case. However, it is also clear that it is




Figure 5. The evolution of the metallicity one would infer by applying a particular diagnostic to a fully sampled SSP cluster of mass $10^4 M_\odot$; the quantity on the vertical axis is the difference between the metallicity $Z_{\text{SSP}}(t)$ inferred for a SSP population of age t and that inferred for a population of age zero. The first and the second vertical lines represent the time at which the O^+ - and H^0 -ionizing luminosities $Q(\text{O}^+)$ and $Q(\text{H}^0)$ fall to 10 per cent of their initial values, respectively.

the runs that are farther away from the zero value on the x -axis that bring about the observed spread in inferred metallicity. That is, stochastic runs where the inferred metallicity is far from the correct value are preferentially those where the ionizing spectral shape, as parametrized by $Q(X)/Q(\text{H}^0)$, differs most strongly from that expected in the fully sampled case. The case of O3N2 is particularly clear, and, thanks to its simplicity, provides a useful lens for understanding the origin of the effect. In this diagnostic, the inferred metallicity simply scales with the strength of the $[\text{O III}] \lambda 5007$ line as $\log(\text{O}/\text{H}) \propto -0.32 \log([\text{O III}] \lambda 5007 / \text{H}\beta)$.⁴ Thus, a 1 dex decrease in the luminosity of the $[\text{O III}] \lambda 5007$ line translates to an $\approx 1/3$ dex increase in inferred metallicity. Examining Fig. 4, we see that the error in the inferred metallicity as a function of $Q(\text{O}^+)/Q(\text{H}^0)$ shows precisely this behaviour: a 1 dex decrease in $Q(\text{O}^+)/Q(\text{H}^0)$ leads to an $\approx 1/3$ dex rise in $\log Z_{\text{stoch}}$. The physical explanation is simply that the $[\text{O III}] \lambda 5007$ line comes from O^{2+} ions, which are produced by O^+ -ionizing photons; thus the ratio $[\text{O III}] \lambda 5007 / \text{H}\beta$ is close to linearly proportional to the ratio $Q(\text{O}^+)/Q(\text{H}^0)$.

With this mechanism understood, we can now also understand the behaviour of other diagnostics. The T_e and KKR_{23} diagnostics are, for the case of input metallicity ($\log Z_{\text{in}}/Z_\odot = -1.0$), much less vulnerable to stochasticity than O3N2 and Scal because they attempt to correct for the unknown ionization state of oxygen in the nebula – by directly measuring lines produced by both O^+ and O^{2+} ions (e.g. both $[\text{O II}] \lambda 3727$ and $[\text{O II}] \lambda 5007$). Because they include an

explicit correction for the ionization state, these diagnostics are less easily fooled when the oxygen ionization balance shifts unexpectedly in response to stochastic variations in the ionizing spectrum. The T_e and KKR_{23} methods continue to work reasonably well even at metallicities close to the solar metallicity. On the other hand, the Scal diagnostic attempts to correct for the unknown ionization states of oxygen using lines produced by N^+ and S^+ as proxies for the hard-to-observe auroral lines of O^+ . This works at solar metallicities but encounters problems at low metallicity because the N^+ and S^+ lines on which it relies cease to provide a good approximation to the unobserved O^+ . Thus Scal is relatively immune to stochasticity at high metallicity, but not at low metallicity.

By contrast, O3N2, PP04-N2, and D16 do not include an implicit or explicit ionization correction at all. For this reason, they can make substantial errors when the ionizing spectrum contains a deficit or excess of O^+ -ionizing photons, unexpectedly shifting the ionization balance in the nebula. This error is not strongly dependent on the underlying metallicity of the nebula.

4.2 Mass sampling versus age sampling

We now explore the relative importance of the two complementary factors that contribute to the spread in inferred metallicities: mass and age sampling. For this purpose, we compare the spread in inferred metallicities for a stellar population that is fully sampled in mass but not in age (the SSP case) with the fully stochastic runs that are not fully sampled in either of the two. In order to explore how imperfect sampling in age affects metallicity diagnostics, we plot

⁴This diagnostic also depends on the ratio $[\text{N II}] \lambda 6584 / \text{H}\alpha$, but this varies by much less than the $[\text{O III}] \lambda 5007 / \text{H}\beta$ ratio, so we can focus on the latter.




Figure 6. Top: inferred metallicity for an SSP as a function of age, using our six sample diagnostics, for a cluster of true metallicity $\log Z_{\text{in}}/Z_{\odot} = -1.0$; this panel is identical to the lower left panel of Fig. 5. Bottom: fraction of the stellar population by mass in various WR spectral classes; here, WNL and WNE are late and early type WN stars, Total WR indicates the sum of all WR sub-types, and $0.1 \times (M > 10 M_{\odot})$ indicates the total mass of all stars above $10 M_{\odot}$, reduced by a factor of 10 so that it is visible on the same scale.

the metallicity derived by applying each diagnostic to an SSP run as a function of age in Fig. 5. In the figure, each panel shows SSP runs at true metallicity values $\log Z_{\text{in}}/Z_{\odot} = -1.5, -1.0, -0.5, 0$, as indicated in the legend. The horizontal axis shows the age and the vertical axis shows $\log Z_{\text{SSP}}(t) - \log Z_{\text{SSP}}(t=0)$, i.e. the difference between the metallicity inferred from the population at age t and one of age 0. Lines of a given diagnostic end at an age when it is no longer possible to measure the metallicity for an SSP of total mass $10^4 M_{\odot}$ using that diagnostic, based on our adopted observability limits⁵ (see Section 2.3.7). The dotted vertical lines show the ages at which the $Q(\text{H}^0)$ and $Q(\text{O}^+)$ fall to 10 per cent of their initial values ($t=0$).

In Fig. 5, we can immediately see that the inferred metallicities for all diagnostics stay nearly constant at ages $\lesssim 3$ –5 Myr, depending upon the metallicity, before drifting away from their initial values at older ages. The initial drift away from zero and the subsequent rise in the uncertainty at later times can be associated with the Wolf-Rayet and massive stars. We can see this explicitly in Fig. 6, where we plot the fraction of the population (by mass) in WR stellar classes as a function of SSP age for example case $\log Z_{\text{in}}/Z_{\odot} = -1.0$. We construct this plot using the WR assignments provided in the MIST evolutionary tracks that we use throughout this work (Choi et al. 2016). As is clear from the plot, the onset of the small fluctuations in $\log Z_{\text{SSP}}(t)$ coincides with the initial appearance of WR stars. However, the rapid growth in fluctuations coincides with the disappearance of the Wolf-Rayets and massive stars from the population. It is straightforward to understand why: the first

appearance of Wolf-Rayet stars significantly alters the ionizing spectral shape by increasing the hard ionizing photon flux. A few Myr later these stars begin to die off, and there is a rapid drop in $Q(\text{O}^+)$. The resulting fluctuations in inferred metallicity are in opposite directions, depending on whether a particular diagnostic increases or decreases as the nebular ionization state changes. Clearly the appearance of WR stars, more importantly, their disappearance as a function of age is capable of inducing significant shifts in the metallicity derived by at least some diagnostics, even for stellar populations that fully sample the IMF.

In order to quantify the importance of partial IMF sampling versus age sampling, we next investigate how much of the stochastic variations result from inadequate age sampling alone. For this purpose, we generate a stochastic-in-time library of models by randomly selecting SSP models at different ages, from age 0 to the age at which $Q(\text{H}^0)$ falls to 10 per cent of its initial value (denoted by the second vertical line in Fig. 5), which we take to be the lifetime of the ionizing stars. This library represents what we would observe if all stars were formed in SSP clusters, and each spaxel in an IFU observation contained exactly one such cluster, but each of these clusters perfectly sampled the IMF. For each of these libraries, we have a distribution of inferred metallicities for each diagnostic (as usual, discarding cases where the lines for a particular diagnostic would be too faint to see), and we can compute the 16th to 84th percentile range of the resulting inferred metallicities, which we denote ΔZ_{SSP} . We can then compare this range to the 16th to 84th percentile range produced by our fully stochastic runs (or some subset thereof), ΔZ_{stoch} , which are stochastic in both age and stellar mass.

We compare the uncertainties observed in the SSP and fully stochastic runs in Fig. 7. In the figure, each panel represents a metallicity diagnostic. On the vertical axis we plot $\Delta Z_{\text{stoch}}/\Delta Z_{\text{SSP}}$, while on the horizontal axis we plot the true metallicity (Z_{in}) of our models. Note that the stochastic runs for which we compute ΔZ_{stoch} are those with densities in the range $n_{\text{H}} = 10$ – 100 cm^{-3} and that we have divided the stochastic runs into three bins of SFR/SFR_0 , running from 0.1–1, 1–10, and 10–100.

The first thing that we notice is that the blue lines, indicating fully stochastic runs with $\text{SFR}/\text{SFR}_0 = 10$ – 100 , generally show a spread ΔZ_{stoch} that is quite close to ΔZ_{SSP} , i.e. $\Delta Z_{\text{stoch}} \approx \Delta Z_{\text{SSP}}$ to within ≈ 50 per cent for most diagnostics at most metallicities (one exception being KKR₂₃). This means that the stochasticity present in our high-SFR runs is mostly a result of imperfect time sampling. A simple physical interpretation is that, at a relatively high SFR, we have enough stars to sample the range of possible stellar masses reasonably well. However, because stars are clustered in time (present in SSP clusters), not all stages of stellar evolution are present in any given patch – indeed, this is consistent with our earlier estimate that, even at $\text{SFR}/\text{SFR}_0 \approx 30$, we expect a given galaxy patch to contain only one cluster with age < 4 Myr that is massive enough to fully sample the IMF. Consequently, stochasticity in time dominates in this range of SFR.

Now, looking at different diagnostics separately, we see for the diagnostics that do not use the oxygen lines at all – D16 and PPO4-N2 – the uncertainties in stochastic runs and SSP runs are comparable. Thus, in this case, inadequate time-sampling is the major contributor to the uncertainty. In contrast, for the diagnostics that use the $[\text{O II}]$ and/or $[\text{O III}]$ lines directly – T_e , Scal, KKR₂₃, and O3N2 – we see that mass-sampling plays a major role at lower SFRs, because ΔZ_{stoch} is generally larger than ΔZ_{SSP} for $\text{SFR}/\text{SFR}_0 < 10$. We tentatively interpret this as indicating the importance of WR stars, which are the only stars capable of producing appreciable numbers of photons capable of ionizing O^+ to O^{2+} . At any given age, only stars born

⁵For the SSP, we use only the relative line cut-offs because the absolute line luminosities can be scaled by increasing the cluster mass.

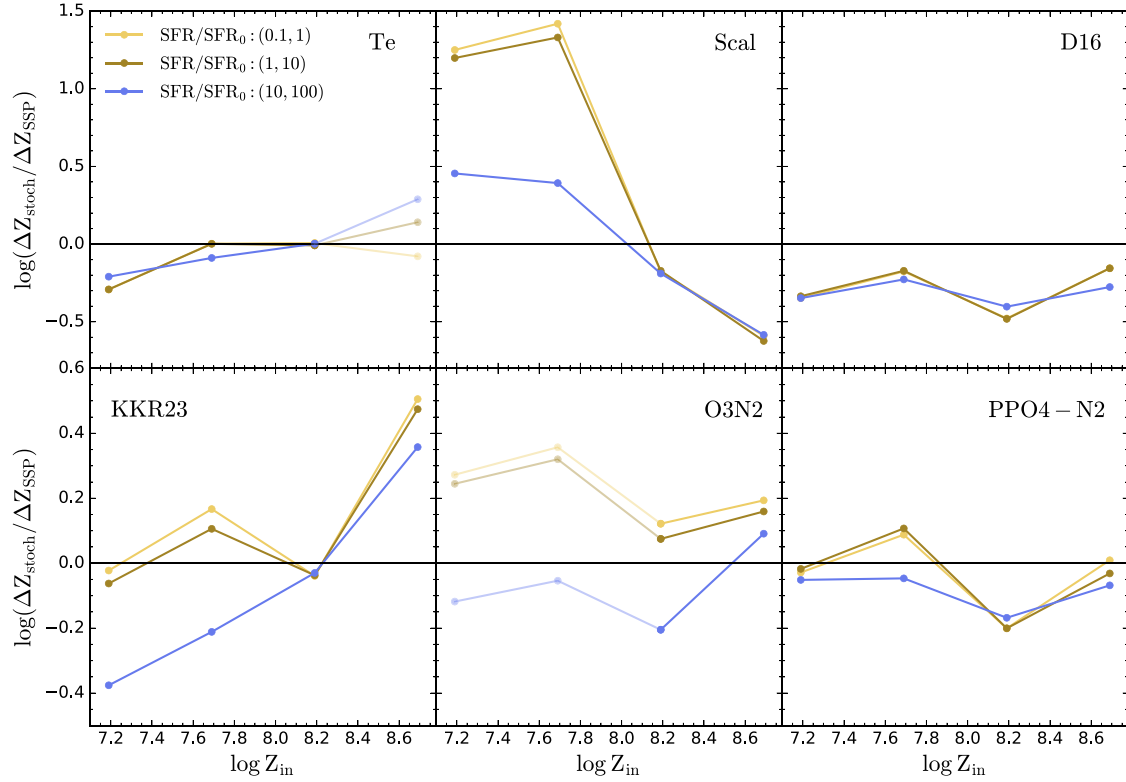


Figure 7. Each panel, corresponding to the diagnostic indicated, shows the ratio of the 16th to 84th percentile range for metallicities inferred from fully stochastic runs, ΔZ_{stoch} , divided by the corresponding range in metallicities inferred for runs where the stellar population consists of an SSP cluster that is fully sampled in mass, ΔZ_{SSP} . We plot this quantity as a function of true metallicity Z_{in} . Colours indicate stochastic runs in different bins of SFR/SFR_0 from 0.1–1, 1–10, and 10–100, as indicated in the legend. The faded points represent the regions where the diagnostic is not reliable - either due to faint lines or because of it was not calibrated in that metallicity range.

within a quite narrow range of mass will be in their WR phase. Consequently, the stochasticity in mass as well as in age is significant for diagnostics that rely on the presence of O^+ -ionizing photons.

5 SUMMARY

In this paper, we quantify how the inferred metallicities of H II regions are affected by stochastic variations in the stellar populations that drive them. More precisely, we explore to what extent metallicities inferred using a range of diagnostics fluctuate in response to stochastic fluctuations in the shape of the ionizing continuum produced when the stars driving the H II region are not fully sampled in mass (IMF) and in stellar age. Consideration of stochastic effects is important for the interpretation of modern IFU observations, because these now routinely sample ~ 100 pc-sized or even smaller patches in nearby galaxies. In such a small region, there are unlikely to be enough massive stars to represent a fully sampled stellar population, but the impact of this imperfect sampling on the metallicities that we derive has not previously been subject to extensive systematic exploration.

We use the SLUG code to generate stochastic realizations of stellar populations in patches of galaxies. We then use the photoionization code CLOUDY to calculate the line emission from nebulae driven by these stochastic stellar populations, and feed the resulting line emission into six of the most commonly used metallicity diagnostics – T_e , KKR₂₃, Scal, O3N2, PPO4–N2, and D16 – to derive inferred metallicities. We repeat this process for ~ 1 million stochastic

realizations where we vary the SFR, metallicity, and ionized gas density of our H II regions.

We find that the primary channel through which stochasticity affects metallicity inference is via the variations it causes in the highest energy parts of the ionizing spectrum, which are primarily produced by Wolf–Rayets stars, a phase of stellar evolution that is both short in duration and narrow in mass range, and thus is subject to considerable stochastic fluctuation. We have shown this results in a deficit in the high-energy part of the ionizing spectrum that directly drives a spread in metallicities that we infer by applying line diagnostics to nebular emission. We find that the size of this spread depends on the choice of diagnostic, the true metallicity (which affects both the stellar spectra and the ionization and temperature balance of the nebula), the SFR (which dictates how many massive stars are present, and thus the size of fluctuations in the spectral shape), and the survey sensitivity limit (which determines at what point low-mass, poorly sampled H II regions become so dim that we cease to observe them at all). In the most stochastic cases we consider, ~ 100 pc-sized galaxy patches with SFRs from 10 to 100 per cent of the value found in the solar neighbourhood, we observe spreads in the inferred metallicity from some of the diagnostics that can be as high as ~ 0.3 dex for sensitivity limits comparable to those of the deepest existing IFU surveys, rising to ~ 0.4 dex for hypothetical future surveys going a factor of 10 deeper. We see that at this point, the uncertainties due to the stochastic fluctuations, at times, dominate the reported uncertainties. The spread drops to ~ 0.1 dex as we increase

the SFR up to 100 times that in the solar neighbourhood and sample the age and mass distributions more adequately.

However, we also find that some diagnostics correct for stochastic effects explicitly and are thus insensitive to stochasticity, regardless of its size. The key distinction between diagnostics that are stochasticity-resistant and those that are not is the extent to which the diagnostic samples more than one ionization state of the ions being used. Diagnostics such as the electron temperature method (T_e) or KKR₂₃, which make use of lines from both O^+ and O^{2+} , for example, show stochastic variations in the derived metallicity $\lesssim 0.2$ dex, regardless of the SFR of the regions being sampled. By contrast, diagnostics such as O3N2 that rely on a single ionization state of O have stochastic variations that depend considerably on the SFR.

We further show that the primary driver of variations at the highest SFRs is stochastic sampling in age rather than in mass – that is, even though a SFR ~ 30 times that found in the solar neighbourhood is enough to guarantee a wide sampling of stellar masses, stochastic variations of ~ 0.1 – 0.2 dex in inferred metallicities still persist, because such a region is likely to contain only a single star cluster massive enough to fully sample the IMF, and thus will not contain all possible stellar ages.

Our findings present a challenge for future surveys. The best protection against stochasticity is to use diagnostics that sample multiple ionization states. Of the diagnostics we explore, T_e performs best, with KKR₂₃ coming close. However, an important caveat to note is that the auroral lines used by these diagnostics are quite faint, and thus are hard to observe in regions where the SFR is low enough for stochasticity to be a concern in the first place. $Scal$, which avoids the faintest auroral lines and uses brighter proxies to estimate the ionization state, shows strong resistance to stochastic variation at high metallicity, and thus is a good choice for solar or near-solar targets, but its S_2 -based estimate of the ionization state becomes subject to very large stochastic variations at metallicities $\lesssim 10$ per cent of solar. The strong-line diagnostics O3N2, PP04-N2, and D16 use lines that are likely to be detectable even in regions of low SFR, but are subject to quite large stochastic uncertainties, and thus the choice of relying on these methods versus more stochasticity-resistant ones that require weaker lines creates a trade-off between the number of regions that one can observe and the robustness of the derived metallicities. With the advent of IFUs and the high-resolution mapping of nearby galaxies that they enable, however, we can no longer ignore the trade-offs imposed by stochasticity, and we must take into account stochastic uncertainties when we interpret observations of metallicity trends. This study can be used as a reference for both of these purposes.

ACKNOWLEDGEMENTS

MRK acknowledges support from an Australian Research Council (ARC) Future Fellowship (award FT180100375). CF acknowledges support from the ARC Discovery Projects scheme (grant DP170100603) and the Future Fellowship scheme (grant FT180100495). This research was undertaken with the assistance of resources and services from the National Computational Infrastructure (projects jh2 and ek9), which is supported by the Australian Government.

DATA AVAILABILITY

The outputs of the SLUG and CLOUDY simulations upon which this paper is based are available for download from <https://cloudstor.aarnet.edu.au/plus/s/x2yjsHdKkRjzgWD>. The data can be read and

processed using the SLUGPY library distributed as part of the SLUG package, available from <https://bitbucket.org/krumholz/slug2/src/master/>.

REFERENCES

- Ashworth G. et al., 2017, *MNRAS*, 469, 2464
 Asplund M., Grevesse N., Sauval A. J., Scott P., 2009, *ARA&A*, 47, 481
 Berg D. A., Skillman E. D., Henry R. B. C., Erb D. K., Carigi L., 2016, *ApJ*, 827, 126
 Bryant J. J. et al., 2015, *MNRAS*, 447, 2857
 Bundy K. et al., 2015, *ApJ*, 798, 7
 Byler N., Dalcanton J. J., Conroy C., Johnson B. D., 2017, *ApJ*, 840, 44
 Byler N., Kewley L. J., Rigby J. R., Acharyya A., Berg D. A., Bayliss M., Sharon K., 2020, *ApJ*, 893, 1
 Chabrier G., 2005, in Corbelli E., Palle F., eds, *Astrophysics and Space Science Library Volume 327*, INAF Osservatorio Astrofisico di Arcetri. The Initial Mass Function 50 Years Later. Springer, Dordrecht, p. 41
 Chiappini C., Matteucci F., Romano D., 2001, *ApJ*, 554, 1044
 Choi J., Dotter A., Conroy C., Cantiello M., Paxton B., Johnson B. D., 2016, *ApJ*, 823, 102
 da Silva R. L., Fumagalli M., Krumholz M., 2012, *ApJ*, 745, 145
 da Silva R. L., Fumagalli M., Krumholz M. R., 2014, *MNRAS*, 444, 3275
 Davé R., Finlator K., Oppenheimer B. D., 2012, *MNRAS*, 421, 98
 de Avillez M. A., Mac Low M.-M., 2002, *ApJ*, 581, 1047
 De Lucia G., Kauffmann G., White S. D. M., 2004, *MNRAS*, 349, 1101
 Di Matteo P., Haywood M., Combes F., Semelin B., Snaith O. N., 2013, *A&A*, 553, A102
 Dopita M. A., Kewley L. J., Sutherland R. S., Nicholls D. C., 2016, *Ap&SS*, 361, 61
 Draine B. T., 2011, *Physics of the Interstellar and Intergalactic Medium*. Princeton Univ. Press, Princeton
 Ferland G. J. et al., 2017, *Rev. Mex. Astron. Astrofis.*, 53, 385
 Fuchs B., Jahreiß H., Flynn C., 2009, *AJ*, 137, 266
 Fumagalli M., da Silva R. L., Krumholz M. R., 2011, *ApJ*, 741, L26
 Garnett D. R., 1992, *AJ*, 103, 1330
 Grand R. J. J. et al., 2016, *MNRAS*, 460, L94
 Hillier D. J., Miller D. L., 1998, *ApJ*, 496, 407
 Ho I. T. et al., 2017, *ApJ*, 846, 39
 Ho I. T. et al., 2018, *A&A*, 618, A64
 Izotov Y. I., Stasińska G., Meynet G., Guseva N. G., Thuan T. X., 2006, *A&A*, 448, 955
 Kewley L. J., Dopita M. A., 2002, *ApJS*, 142, 35
 Kewley L. J., Ellison S. L., 2008, *ApJ*, 681, 1183
 Kewley L. J., Dopita M. A., Sutherland R. S., Heisler C. A., Trevena J., 2001, *ApJ*, 556, 121
 Kewley L. J., Rupke D., Zahid H. J., Geller M. J., Barton E. J., 2010, *ApJ*, 721, L48
 Kewley L. J., Nicholls D. C., Sutherland R. S., 2019, *ARA&A*, 57, 511
 Kobulnicky H. A., Kewley L. J., 2004, *ApJ*, 617, 240
 Kreckel K. et al., 2019, *ApJ*, 887, 80
 Kreckel K. et al., 2020, *MNRAS*, 499, 193
 Kreckel K., Blanc G. A., Schinnerer E., Groves B., Adamo A., Hughes A., Meidt S., 2016, *ApJ*, 827, 103
 Krumholz M. R., Fumagalli M., da Silva R. L., Rendahl T., Parra J., 2015a, *MNRAS*, 452, 1447
 Krumholz M. R. et al., 2015b, *ApJ*, 812, 147
 Krumholz M. R., Ting Y.-S., 2018, *MNRAS*, 475, 2236
 Krumholz M. R., McKee C. F., Bland-Hawthorn J., 2019, *ARA&A*, 57, 227
 Leitherer C. et al., 1999a, *ApJS*, 123, 3
 Leitherer C. et al., 1999b, *ApJS*, 123, 3
 Lejeune T., Cuisinier F., Buser R., 1997, *A&AS*, 125, 229
 Marino R. A. et al., 2013, *A&A*, 559, A114
 Nagamine K., Fukugita M., Cen R., Ostriker J. P., 2001, *ApJ*, 558, 497
 Nicholls D. C., Sutherland R. S., Dopita M. A., Kewley L. J., Groves B. A., 2017, *MNRAS*, 466, 4403
 Nicholls D. C., Kewley L. J., Sutherland R. S., 2020, *PASP*, 132, 033001

- Paalvast M., Brinchmann J., 2017, *MNRAS*, 470, 1612
 Pagel B. E. J., Edmunds M. G., Blackwell D. E., Chun M. S., Smith G., 1979, *MNRAS*, 189, 95
 Pauldrach A. W. A., Hoffmann T. L., Lennon M., 2001, *A&A*, 375, 161
 Peimbert M., 1967, *ApJ*, 150, 825
 Peimbert M., Peimbert A., Delgado-Inglada G., 2017, *PASP*, 129, 082001
 Petit A. C., Krumholz M. R., Goldbaum N. J., Forbes J. C., 2015, *MNRAS*, 449, 2588
 Pettini M., Pagel B. E. J., 2004, *MNRAS*, 348, L59
 Pilyugin L. S., Grebel E. K., 2016, *MNRAS*, 457, 3678
 Pilyugin L. S., Grebel E. K., Mattsson L., 2012, *MNRAS*, 424, 2316
 Rosa D. A., Dors O. L., Krabbe A. C., Hägele G. F., Cardaci M. V., Pastoriza M. G., Rodrigues I., Winge C., 2014, *MNRAS*, 444, 2005
 Rubin R. H., 1969, *ApJ*, 157, 1461
 Rupke D. S. N., Kewley L. J., Barnes J. E., 2010, *ApJ*, 710, L156
 Sánchez-Menguiano L. et al., 2016, *ApJ*, 830, L40
 Sánchez S. F. et al., 2012a, *A&A*, 538, A8
 Sánchez S. F. et al., 2012b, *A&A*, 546, A2
 Schinnerer E. et al., 2019, *The Messenger*, 177, 36
 Scott N. et al., 2018, *MNRAS*, 481, 2299
 Somerville R. S., Primack J. R., 1999, *MNRAS*, 310, 1087
 Stasińska G., 2005, *A&A*, 434, 507
 Sutherland R., Dopita M., Binette L., Groves B., 2018, Astrophysics Source Code Library, record ascl:1807.005
 Taylor P., Kobayashi C., 2017, *MNRAS*, 471, 3856
 Torres-Flores S., Scarano S., Mendes de Oliveira C., de Mello D. F., Amram P., Plana H., 2014, *MNRAS*, 438, 1894
 Tremblin P. et al., 2014, *A&A*, 568, A4
 Vogt F. P. A., Pérez E., Dopita M. A., Verdes-Montenegro L., Borthakur S., 2017, *A&A*, 601, A61
 Wang X. et al., 2019, *ApJ*, 882, 94
 Yang C.-C., Krumholz M., 2012, *ApJ*, 758, 48

Yeh S. C. C., Matzner C. D., 2012, *ApJ*, 757, 108

APPENDIX: EFFECTS OF GAS DENSITY

Here, we explore the dependence of inferred metallicities on ionized gas number density. To put this discussion in context, we note that most observed H II regions have densities in the range $\sim 10\text{--}100\text{ cm}^{-3}$ (Tremblin et al. 2014). Thus, we would typically expect an IFU survey of a galactic disc to contain only a handful of spaxels for which $n_{\text{H}} \geq 1000\text{ cm}^{-3}$. To examine how stochasticity is likely to affect these spaxels, we repeat the analysis presented in Section 3.1 for the full density range ($10\text{--}10^4\text{ cm}^{-3}$) present in our simulation library. This almost certainly exaggerates the effects of density, since our simulation library is uniformly sampled in $\log n_{\text{H}}$, making high-density regions far more common than is likely to be the case in reality.

We show the results in Fig. A1; other than the density range used to generate it, this figure is identical to Fig. 2. Comparing the two figures, we immediately notice an increase in the uncertainty range of the D16 diagnostic across the whole metallicity range. Scal is the next-most affected, though by much less than D16, and all the other diagnostics show negligible changes ($\lesssim 0.1$ dex) when compared to our analysis of the models with a narrower density range. Thus, we see that even when we include the H II regions with ionized gas densities far from the commonly observed values, we do not see major changes in our results and conclusions. The reason the D16 diagnostic is affected so strongly is that [S II] $\lambda 6717$, 31 doublet upon which it relies has a relatively low critical density, and thus can be collisionally quenched in the higher density H II regions we are now sampling. Scal, which also uses this doublet, is affected as well, but less severely.

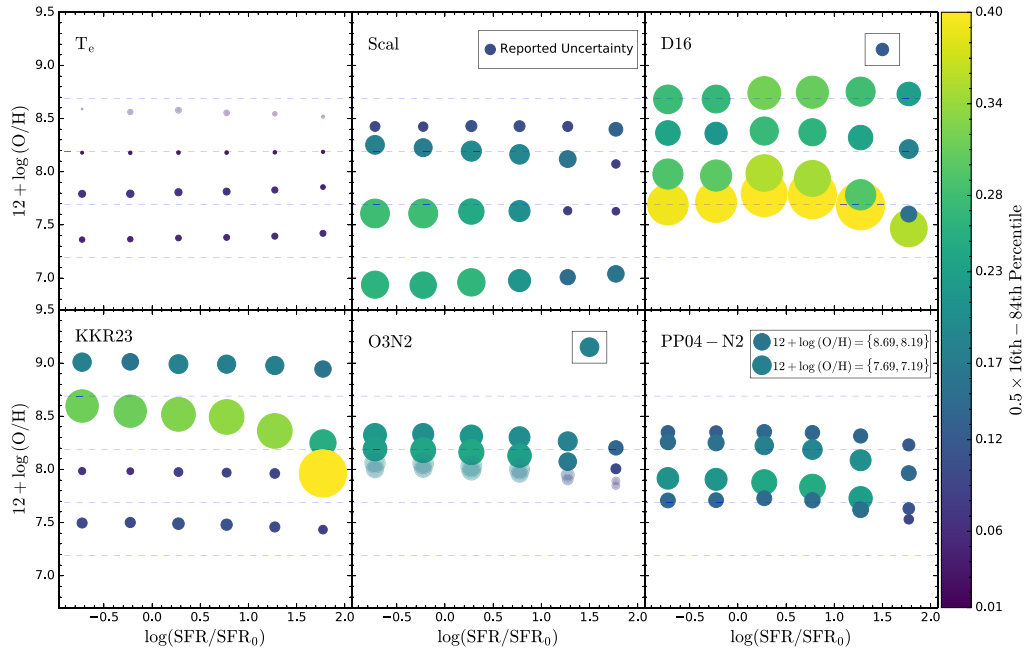


Figure A1. The same as Fig. 2, but for a range of H II region densities $n_{\text{H}} = 10\text{--}10^4\text{ cm}^{-3}$, uniformly sampled in $\log n_{\text{H}}$.