WARPFIELD population synthesis: the physics of (extra-)Galactic star formation and feedback-driven cloud structure and emission from sub-to-kpc scales

Eric W. Pellegrini,[1]★ Stefan Reissl,[1] Daniel Rahner,[1] Ralf S. Klessen,[1,2] Simon C. O. Glover,[1]
Rüdiger Pakmor ●,[3] Rodrigo Herrera-Camus[4] and Robert J. J. Grand ●[3]

[1]*Zentrum für Astronomie, Institut für Theoretische Astrophysik, Universität Heidelberg, Albert-Ueberle-Str. 2, D-69120 Heidelberg, Germany*
[2]*Interdisziplinäres Zentrum für Wissenschaftliches Rechnen, Universität Heidelberg, Im Neuenheimer Feld 205, D-69120 Heidelberg, Germany*
[3]*Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, D-85748 Garching, Germany*
[4]*Departamento de Astronomía, Universidad de Concepción, Casilla 160-C, Concepción, Chile*

ABSTRACT

We present a novel method to model galactic-scale star formation and emission of star clusters and a multiphase interstellar medium (ISM). We combine global parameters, including star formation rate and metallicity, with the 1D cloud evolution code WARPFIELD to model the sources of feedback within a star-forming galaxy. Within individual star-forming regions, we include stellar evolution, stellar winds, radiation pressure, and supernovae, all coupled to the dynamical evolution of the 1D parental cloud in a highly non-linear fashion. Heating of the diffuse galactic gas and dust is calculated self-consistently with the age-, mass-, and density-dependent escape fractions of photons from these fully resolved local star-forming regions. We construct the interstellar radiation field, and we employ the multifrequency radiative transfer code POLARIS to produce synthetic emission maps for a one-to-one comparison with observations. We apply this to a cosmological simulation of a Milky-Way-like galaxy built-up in a high-resolution MHD simulation of cosmic structure formation. From this, we produce the multiscale/phase distribution of ISM density and temperature and present a synthesized all-sky H α map. We use a multipole expansion to show that the resulting maps reproduce all observed statistical emission characteristics. Next, we predict [S III] 9530 Å, a key emission line that will be observed in several large forthcoming surveys. It suffers less extinction than other lines and provides information about star formation in very dense environments that are otherwise observationally inaccessible optically. Finally, we explore the effects of differential extinction, and discuss the consequences for the interpretation of H α emission at different viewing angles by an extragalactic observer.

Key words: radiative transfer – Galaxy: evolution – Galaxy: formation.

1 INTRODUCTION

Recent years have seen dramatic improvements in our ability to model the formation and evolution of realistic spiral galaxies within large cosmological simulations (e.g. Grand et al. 2017; Hopkins et al. 2018). At the same time, the advent of ALMA and of large integral field unit (IFU) spectrographs such as MUSE (Bacon et al. 2010) or SITELLE (Grandmont et al. 2012) has, for the first time, made it possible to map both gas and star formation on small (∼100 pc) scales within a large sample of local galaxies. (e.g. Kreckel et al. 2018; Rousseau-Nepton et al. 2018a). The role that superstar clusters play in determining the emission properties of galaxies has also become increasingly clear, thanks to advances in spatially resolved observations that lead to their identification as important sources of emission (Turner et al. 2017; Oey et al. 2017b). It is therefore important to have a single overarching framework to explore star formation that could be used to understand the evolution of clusters

with a wide range of masses and how they contribution to the observed emission from a galaxy.

An obvious next step is to compare the predictions of simulations with observations of real galaxies, but this remains a highly challenging problem. Although high-precision cosmological simulations can now resolve individual star-forming regions on scales of 10–100s of parsecs, this is not the same as being able to do the same with individual young stellar clusters. In addition, the small-scale physics of the interstellar medium (ISM) and stellar feedback is often treated in a highly idealized fashion in these simulations, and so the information necessary for making realistic synthetic emission maps based on the simulations is frequently not directly available.[1]

---

[1] See e.g. the Illustris or Illustris-TNG simulations, introduced in Vogelsberger et al. (2014) and Springel et al. (2018), respectively, which use the effective equation of state from Springel & Hernquist (2003) to model the pressure of the dense interstellar medium and hence do not directly follow the dense gas temperature.

★ E-mail: eric.pellegrini@uni-heidelberg.de

On the observational side, even at the resolution achievable with ALMA and MUSE, individual star-forming regions are not fully resolved (pc scale), unless one focuses solely on very nearby galaxies such as the Magellanic Clouds or M31. Observations of more distant galaxies convolve together light from multiple stellar clusters of different masses, ages, and potentially also metallicities within a single aperture. Observations of emission lines within these apertures therefore probe gas with a range of different physical conditions, exposed to a variety of radiation fields, greatly complicating efforts to compare the emission-line strengths with the predictions of simple single-component photoionization and photodissociation region (PDR) models. Depending on the wavelength considered, these measurements are often also contaminated by diffuse emission powered by the field star population or by ionizing photons that manage to escape from the immediate vicinity of star-forming clouds (e.g. Medling et al. 2018; Tomičić et al. 2019).

In addition, some quantities of great observational interest, such as the distribution of Faraday rotation measures (see e.g. Oppermann et al. 2012), are sensitive to the small-scale distribution of young massive clusters via their impact on the galactic free electron distribution, but also depend directly on large-scale features such as the structure and strength of the magnetic field. Making realistic predictions for these quantities therefore requires us to take a holistic view of a model galaxy, rather than one focused on individual star-forming regions.

Since cosmological simulations do not currently have sufficient resolution to directly predict the locations or masses of young massive clusters, if we want to use these simulations to make observational predictions of e.g. star formation diagnostics or Faraday rotation measures, it is necessary to adopt a population synthesis approach in which we add a population of clusters to the galactic gas distribution provided by the simulation and then use the resulting combined model of stars and gas to generate our synthetic observables. However, an important consideration when constructing such a population synthesis model is the relationship between the gas and the young stars. Given an overall star formation rate (SFR) for a model galaxy, or a coarse-resolution map of the SFR surface density, it is straightforward to generate a population of young clusters by sampling from an appropriate cluster mass function and depositing the clusters randomly within the model galaxy. However, such an approach results in a cluster distribution that takes no account of the gas distribution. For old stellar clusters, this may be reasonable, but for the young clusters that contribute most of the stellar feedback, it is a poor approximation, since we know that these clusters must have formed within gravitationally unstable clouds of gas.

In this paper, we present a new method for building up a population of clusters within a simulated galaxy that accounts for the link between the gas distribution and the cluster locations, and that allows us to model not just the direct emission from the clusters but also the diffuse emission from the ISM. Our method is based upon the WARPFIELD-EMP code (Pellegrini et al. 2020). This combines the WARPFIELD stellar feedback model, described in Rahner et al. (2017, 2018a), with the CLOUDY PDR code[2] (Ferland et al. 2017) and the POLARIS radiative transfer (RT) code[3] (Reissl, Wolf & Brauer 2016; Reissl et al. 2019b). WARPFIELD models the impact of a stellar cluster on its surrounding cloud, accounting for a wide range of different feedback processes (radiation in ionizing and non-ionizing wavebands, stellar winds, supernovae) and solving for

the dynamical evolution of the gas in spherical symmetry. The results of the model are then post-processed using CLOUDY (to generate emissivities) and POLARIS (to account for line RT, dust absorption, and synthetic observations), yielding predictions for the emission from the cloud/cluster system in the continuum and a large number of lines (Pellegrini et al. 2020). The method is fast and computationally efficient, and thus allows us to put together a large data base of cloud/cluster models that cover the entire parameter space relevant for normal spiral galaxies. Hence, for any combination of cloud masses, gas densities, and most importantly star formation efficiencies, we can describe the state and lifetime of individual clouds.

Here, we connect our WARPFIELD-EMP models to a Milky-Way-like galaxy produced within a cosmological simulation taken from the Auriga project (Grand et al. 2017). However, we note that, in principle, our new population synthesis method is compatible with any type of mock or simulated galaxy. In the proof of concept presented here, we restrict our attention to a subset of the observational tracers that can be studied using the model.

We generate synthetic maps of Hα, Hβ, and [S III] 9530-Å line emission as well as the Faraday rotation measure (published in a companion paper), considering both what would be seen by an observer within the galaxy and also what would be seen by an external observer. We focus on [S III] rather than the more commonly used [O III] 5007-Å line because although both have similar ionization potentials and hence trace similar regions of massive star formation, the longer wavelength of the [S III] line means that it is less affected by dust extinction, allowing it to probe more distant or more embedded H II regions than [O III]. Nevertheless, our method can also be used to make maps of the [O III] line, as well as many other different observational tracers, ranging from polarized dust emission to atomic and molecular line emission.

Using Fig. 1 as a guide, our paper is structured in the following manner. In Section 2, we describe the cosmological simulation of a Milky-Way-type galaxy (diffuse gas) that we populate with WARPFIELD star cluster models, shown as clusters surrounded by expanding shells. In Section 3, we present the main features of WARPFIELD-POP. In particular, in Section 3.2, we describe our population synthesis model, which depends on an input SFR and cluster mass function, and in Section 3.3, we describe how we use the emergent radiation (red arrows emerging from the model) computed by the WARPFIELD-EMP models to photoionize the disc. Line emission from the H II region complexes and diffuse gas, as well as its transfer through the galaxy is treated in Section 4, followed an in-depth analysis of the synthetic observations in Section 5. We discuss some of our main results in Section 6 and close with a brief summary in Section 7.

2 COSMOLOGICAL SIMULATION

The model galaxy that we post-process using the WARPFIELD population synthesis method is taken from a simulation carried out as part of the Auriga project. It comes from a set of 30 cosmological magnetohydrodynamical (MHD) zoom-in simulations of isolated Milky-Way-like galaxies. The simulations assume Lambda cold dark matter, with cosmological parameters taken from the Planck Collaboration (2014). They begin at a redshift of $z = 127$ and are evolved all the way to $z = 0$. The initial conditions were generated by selecting regions in the dark matter only version of the EAGLE simulation (McAlpine et al. 2016) that form dark matter haloes with $M_{vir} \sim 10^{12} M_\odot$, with the restriction that the haloes should not be too close to other haloes of comparable mass. The selection criteria (described in more detail in Grand et al. 2017) yielded 174 candidate

[2]http://www.nublado.org/.
[3]http://www1.astrophysik.uni-kiel.de/~polaris.
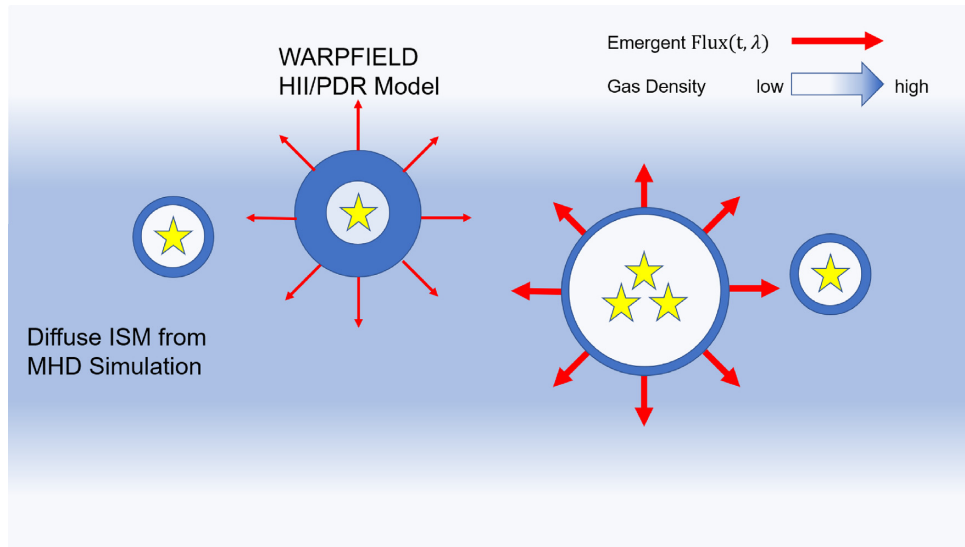
Figure 1. A schematic of the model presented. A distribution of cluster ages is derived from a star formation history and a cluster mass function. The associated H II region and PDR evolution is determined at each evolutionary point with WARPFIELD-EMP, including the line emission. Finally, the emergent radiation from each region is then injected back into the galaxy as a source of heating and ionization, added to a diffuse galactic radiation field.

regions, of which 30 were selected for further study using a standard zoom-in approach. The zoom-in simulations include both gas and dark matter and account for a wide variety of physical processes, including primordial and metal line cooling (with a correction for self-shielding), the influence of the extragalactic UV background, star formation, stellar evolution, and metal return (Vogelsberger et al. 2013). Moreover, the simulations employ an effective model for Galactic winds (Marinacci, Pakmor & Springel 2014; Grand et al. 2017), and a prescription for the formation and growth of active galactic nuclei (AGN) and their feedback (Grand et al. 2017), although we do not account for the possible presence of an AGN in the galaxy in our post-processing treatment. All simulations include magnetic fields that are seeded with small amplitudes at $z = 127$ and are self-consistently evolved until $z = 0$ (Pakmor & Springel 2013; Pakmor, Marinacci & Springel 2014; Pakmor et al. 2017).

The simulations employ the moving mesh code AREPO that solves the equations of MHD coupled with self-gravity on an unstructured Voronoi grid that evolves with time in a quasi-Lagrangian fashion (Springel 2010). The Auriga simulation suite focuses on two sets of simulations at different resolution: 30 galaxies at standard resolution (level 4, $M_{baryon} = 2 \times 10^4 - 5 \times 10^4 \, M_\odot$) and six haloes at high resolution (level 3 and level 4, $M_{baryon} = 3 \times 10^3 - 6 \times 10^3 \, M_\odot$). Explicit refinement and de-refinement is used to keep cells in the high-resolution region within a factor of 2 of the target mass resolution. The high-resolution region is made sufficiently large that there is no contamination within 1 Mpc of the main halo at $z = 0$, i.e. there are no low-resolution elements closer than 1 Mpc.

The highest gas density at $z = 0$ reached within the high-resolution (level 3) simulations is $n \sim 10 \, cm^{-3}$, corresponding to a cell size $\sim 25 \, pc$. Resolving structure in the gas distribution requires a few cells per dimension, and so the effective spatial resolution of the simulation is at best around 100 pc (comparable to the gravitational softening length for the gas) at the highest densities. Because of this limited resolution (which is, nevertheless, very good by the standards of cosmological simulations), the Auriga simulations cannot follow the formation of individual star clusters and cannot model the effects of stellar feedback with the same fidelity as our WARPFIELD models. Instead, the impact of stellar feedback is accounted for in a

subgrid fashion following the prescription introduced by Springel & Hernquist (2003). Gas above a density threshold of $n = 0.13 \, cm^{-3}$ is artificially pressurized and is taken to represent some unresolved mix of cold/warm gas with medium to high densities and hot, supernova-heated gas with a low density.

Since we are performing a 3D photoioniziation raytrace (see Section 4 below), the assumed geometric configuration of the different densities within a cell will affect the result. For example, if higher density gas is placed closer to a source, the low density gas may be shielded from ionizing radiation due to the higher recombination rate of the higher density material. The hydrostatic equation of state (EOS) employed in CLOUDY and used in WARPFIELD is a physically motivated and self-consistent method for describing density changes within the portions of the galaxy represented by the WARPFIELD models. However, outside of these regions, in the portion of the calculation where we use the densities provided by the cosmological simulation, there is no good way to constrain the geometric distribution of different densities on subgrid scales in the gas. While we could assume a distribution, this would introduce multiple poorly constrained free parameters into the problem (since we would need to specify not only the range of different densities but also their spatial distribution). For this reason, outside of the regions represented by the WARPFIELD models, we assume that the true density is the same as the density in the cosmological simulation, since we lack a simple physically motivated alternative.

For the purposes of this paper, we post-process only one of the 30 galaxies simulated in the Auriga project. Our selected galaxy, Au-6, is modelled in one of the high resolution (level 3) runs and has a halo mass of $10^{12} \, M_\odot$ and a stellar mass of $6 \times 10^{10} \, M_\odot$. This galaxy is similar to the Milky Way in many respects, including the properties of the stellar disc (Grand et al. 2017, 2018), the gas disc (Marinacci et al. 2017), the stellar halo (Monachesi et al. 2016), the magnetic field structure (Pakmor et al. 2017), and the population of satellite galaxies (Simpson et al. 2018). This makes the Auriga galaxy Au-6 an excellent test-case for understanding the physical processes that govern the formation and evolution of the Milky Way. It is also a good starting point for our population synthesis model. We take the gas density and other galaxy properties of Au-6 as input for deriving

the cluster mass distribution as well as for synthesizing electron densities, electron temperatures, and emissivities, and, finally, the calculation of synthetic line emission, as we discuss in the sections below.

3 POPULATION SYNTHESIS MODELLING

Our objective in this paper is to present a method for forward modelling the stellar population and emission of a region, be it an entire galaxy or a kpc-scale subregion within a galaxy, that is described by an SFR, a metallicity, and a characteristic environmental density and instantaneous star formation efficiency. Within such a region, we expect to find many different clouds and star clusters[4] with a range of different masses and ages, and so a key part of this method is the generation of an appropriate sample of clouds and clusters.

This task is made more difficult by the potential complexity of the evolution of the individual star-forming regions. As we have explored in earlier papers, the interplay of cooling, gravity, and the time-varying energy and momentum input from an evolving stellar population yields a range of different dynamical outcomes (see e.g. Rahner et al. 2017, 2019) that are not self-similar between objects of different mass, density, or metallicity. While all star-forming regions undergo an initial feedback-driven expansion, the combination of hot gas cooling and escape of radiation alone are enough to cause the expansion of some regions to stall. Clouds in this regime will often recollapse under their own self-gravity, forming a second generation of stars. As a result, the star formation efficiency of a given cloud can depend on whether it is destroyed by its initial burst of star formation, or whether feedback is initially unable to destroy it, resulting in star formation continuing over a more extended period. Because of this, it is difficult to predict a priori the contribution of a given cloud to the global SFR, as this depends on the cloud's dynamical history and on whether it undergoes one or multiple bursts of star formation.

In this section, we outline how we address this problem and generate a sample of clouds and clusters that are consistent with a specified global SFR, while still accounting for the fact that some clouds may form multiple populations of stars. Armed with this sample, we then explore how we can use it to make predictions about the observable properties of the region as a whole.

3.1 Subgrid models of star-forming clouds

To calculate the evolution of the natal cloud around each of our model clusters we use WARPFIELD. In this introductory study, we assume for simplicity that all clouds share the same average natal cloud density $n = 100 \, \text{cm}^{-3}$ and star formation efficiency $\epsilon = 1$ per cent, but we note that the model can easily be extended to consider complex distributions of both of these properties. Since we know the cluster mass M_{cl}, the cloud mass then follows simply as $M_{\text{cloud}} = \epsilon^{-1} M_{\text{cl}}$. The remaining WARPFIELD input parameter is the metallicity. This could, in principle, be adopted from the cosmological simulation, but in this paper, we assume, again for simplicity, that it has the solar value. The time evolution of the cloud, as modelled by WARPFIELD, provides us

with the internal pressure and the radiation field, which uniquely determine the properties of the hydrostatic H II region as it expands into the natal cloud structure and beyond. These properties are then fed into CLOUDY, along with the spectral energy distribution of the star cluster, allowing us to solve simultaneously for the local volume emissivity at each point in the cloud. These local emissivities can then be used, along with the dust distribution, as inputs for POLARIS, allowing us to obtain the full dust attenuated/reprocessed spectrum emerging from the star-forming region, as described in more detail in our companion paper on WARPFIELD-EMP (Pellegrini et al. 2020). Note that although CLOUDY itself can be used to account for the effects of dust attenuation on the emission lines produced by a given star-forming region, we do not make use of this option here, in order to avoid double counting the effects of the dust.

Our model makes the assumption that the duration of any individual burst of star formation within a cloud is short compared to the evolutionary time-scale of the cloud, so that we can treat it as instantaneous.[5] We justify this assumption as a consequence of the effectiveness of feedback in cluster-forming regions, as explored in Rahner et al. (2017, 2019). The young massive clusters that we are primarily concerned with here quickly clear out gas from their immediate vicinity, driving expanding shells into the surrounding cloud and the larger-scale ISM. During the expansion phase, the cloud is subjected to intense radiation, and we find that molecular gas is destroyed rapidly, temperatures rise, and the cloud becomes partially to fully ionized. In these extreme conditions, star formation is unlikely to occur over extended periods of time. See also Li et al. (2019) and Chevance et al. (2020) for additional numerical and observational support for this point.

For more massive clusters, it can sometimes happen that feedback from the initial burst of star formation is unable to completely disrupt the cloud. In that case, as the cluster ages and its feedback becomes less effective, the expansion of the feedback-driven shell stalls. Following this, the gas undergoes renewed collapse due to its own self-gravity, a phenomenon we refer to as 'recollapse' (Rahner et al. 2017; see also Rahner et al. 2018b; Rugel et al. 2019 for examples of some clusters where there is good evidence that recollapse is occurring or has occurred). In mass bins in which recollapse occurs, equilibrium between cluster formation and cluster death takes longer to establish, but for the current experiment, tailored to Milky Way conditions, we find the cluster population and the ionizing output reach equilibrium after around 20 Myr, which is the time we choose to present here.

### 3.2 Generating the cluster mass and age distribution

In order to generate a sample of clusters in a region of interest, we need to know two things: the initial mass function of the star clusters and the rate at which gas is converted into stars within that region. We assume that the initial cluster mass function is a power law with exponent β, where

$$\log_{10} \left(\frac{\text{d}N_{\text{cl}}}{\text{d}M_{\text{cl}}} \right) \propto -\beta \log_{10}(M_{\text{cl}}), \quad (1)$$

and which extends between a minimum cluster mass $M_{\text{cl, min}} = 10^2 \, \text{M}_\odot$ and a maximum cluster mass $M_{\text{cl, max}} = 10^7 \, \text{M}_\odot$, consistent with observations in nearby galaxies (Zhang & Fall 1999; Lada &

---

[4]We use the words 'star cluster' or 'cluster' for brevity, but in actual fact, our model is agnostic on the issue of whether the stars are located in a gravitationally bound cluster or a gravitationally unbound association, so long as they remain reasonably localized in space for the ∼20-Myr period during which they contribute to the emission of the galaxy in the tracers of interest in this study.

[5]Strictly speaking, we only require that the massive stars that dominate the feedback form rapidly, and so our assumption is not inconsistent with scenarios in which low-mass stars form over a more extended period.
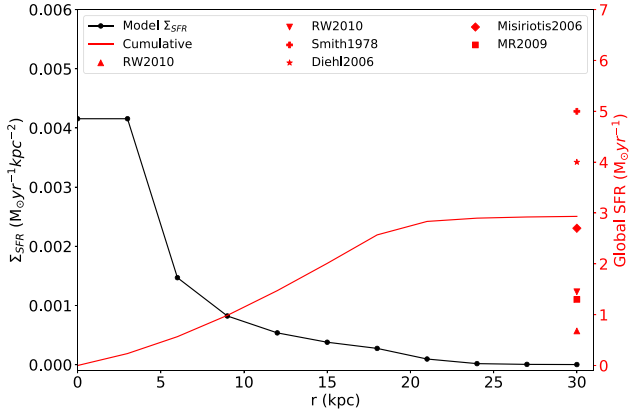
Figure 2. The radial SFR surface density (black) of our model galaxy, computed using equation (2) for 10 evenly spaced radial bins. Also shown is the cumulative SFR (red) of the entire galaxy as a function of radius. Various estimates of the global Milky Way SFR are indicated as symbols at the outermost radii for comparison. These are taken from Robitaille & Whitney (2010, hereafter RW2010), who give both upper and lower limits, Smith, Biermann & Mezger (1978), Diehl et al. (2006), Misiriotis et al. (2006), and Murray (2009, hereafter MR2009).

Lada 2003; Portegies Zwart, McMillan & Gieles 2010; Gouliermis 2018; Krumholz, McKee & Bland-Hawthorn 2019) Observationally determined values of β typically lie in the range $\beta = 2.0 \pm 0.2$ and so in our fiducial model we set $\beta = 2.0$. However, our method is readily generalizable to the case where $\beta \neq 2.0$.

The other important parameter is the SFR. One option is to take this directly from the adopted cosmological simulation, but it can also be specified directly by the user. For example, in the models presented in this paper, we make the simplifying assumption that the SFR is constant over times that are short compared to the molecular gas depletion time ($\tau_{\rm depl} \equiv M_{\rm mol}/{\rm SFR}$, where $M_{\rm mol}$ is the molecular gas mass). We also assume that the SFR is directly proportional to the local gas mass. We measure this gas mass by dividing up the galaxy in the Au-6 simulation into a set of linearly spaced radial bins originating at the galactic centre of mass and summing up the gas mass within each bin. We limit the vertical extent of each bin to ± 5 kpc from the disc mid-plane.

If the gas mass in the ith bin is $M_{{\rm gas},i}$, then the corresponding SFR is given by

$${\rm SFR}_i = M_{\rm gas,i} \times f_{\rm dense}/\tau_{\rm depl}, \qquad (2)$$

where $f_{\rm dense}$ is fraction of dense, cold gas, and $\tau_{\rm depl}$ is the assumed depletion time due to star formation. We take values of $f_{\rm dense} = 25$ per cent (Klessen & Glover 2016) and $\tau_{\rm depl} = 2.0$ Gyr (Bigiel et al. 2008), reasonable values for a Milky Way like system. In the Au-6 simulation, roughly 40 per cent of the gas mass is above the Springel & Hernquist (2003) density threshold and hence is potentially dense and cold. However, the total mass of the Au-6 galaxy also is approximately twice that of the MW, so by using a smaller value for $f_{\rm dense}$, we obtain a total SFR in better agreement with the Galactic value. This enables us to make a more meaningful comparison between the results we obtain for this galaxy and observations of the real Milky Way (see Section 5).

In the model presented in this paper, we calculate the SFR in 10 annular bins, resulting in the radial profile depicted in Fig. 2. Also shown are various observational estimates for the total SFR in the Milky Way. These span values from ~ 1 to $5\,{\rm M}_\odot\,{\rm yr}^{-1}$, and the number of $2.9\,{\rm M}_\odot\,{\rm yr}^{-1}$ that we obtain with our simple SFR prescription lies

well within this range. Indeed, it agrees well with the value that Diehl et al. (2006) infer based on their Hα observations of the Milky Way.

Given the initial cluster mass function and the SFR, we proceed as follows. We consider only clusters formed within the last 50 Myr[6] and split up this period into 50 uniform time-steps, each with length $\Delta t = 1.0$ Myr. We calculate the gas mass converted into stars in each time-step. For time-step n, spanning the period $t_n \rightarrow t_n + \Delta t$, we have

$$\Delta M = \int_{t_n}^{t_n + \Delta t} {\rm SFR}(t)\,{\rm d}t, \qquad (3)$$

where ${\rm SFR}(t)$ is the value of the SFR at time t. In the simple case of a constant SFR, this reduces to $\Delta M_n = {\rm SFR} \times \Delta t$.

We next assign this mass to a set of discrete mass bins of 0.25 dex, evenly spaced in logarithm between $M_{\rm cl, min}$ and $M_{\rm cl, max}$. We use discrete mass bins matched to our existing WARPFIELD-EMP data base instead of randomly sampling the mass function to limit the computational cost of the model. Each bin receives a fraction of ΔM corresponding to the slope of the initial cluster mass function that is located in that bin mass. In the simple case of $\beta = 2.0$ considered in this paper, each logarithmic bin has an equal fraction of the mass in the initial cluster mass function and hence receives an equal fraction of ΔM on each time-step.

At the end of each time-step, we check for each bin M_i whether the accumulated mass is sufficient to form a cluster of mass M_i. If so, then we add a new cluster of that mass to our population at that time-step and decrease the mass in the bin by M_i. The newly created cluster is assigned an age drawn randomly from the time interval [$t \rightarrow t + \Delta t$]. If the mass in the cluster is sufficient to form more than one new cluster, then we simply repeat this procedure as many times as necessary. On the other hand, if the mass in the bin is insufficient to form even one new cluster, then we retain all of it for the next time-step.

Up to this point in calculating an SFR, we have assumed that each cloud undergoes only a single burst of star formation. As shown in Rahner et al. (2018b) and Rugel et al. (2019), this is not necessarily the case. Once a stellar population becomes older than ~ 3 Myr, its output of energy in the form of stellar winds and radiation steadily declines as its most massive stars begin to die. Consequently, the cluster becomes a less effective source of stellar feedback as it ages. This is of little importance if the cloud has already been destroyed, but for some combinations of parameters, the cloud may not yet have been completely dissolved at the time that the cluster feedback becomes ineffective. In this case, recollapse of the cloud may occur, leading to a new burst of star formation. To determine whether this occurs, we run a WARPFIELD model for each new cluster, using the appropriate cluster and cloud masses. This allows us to identify the cases in which recollapse occurs and also provides us with the time at which it occurs and the mass of new stars formed in each recollapsing cloud.

We chose a fiducial reference point of 20 Myr to compute the SFR, having found that at this point, the cluster population has reached a roughly steady state. The contribution to the total SFR coming from recollapsed clusters is

$${\rm SFR}_{\rm rec,20} = \frac{M_{\rm rec,20}}{20\,{\rm Myr}}. \qquad (4)$$
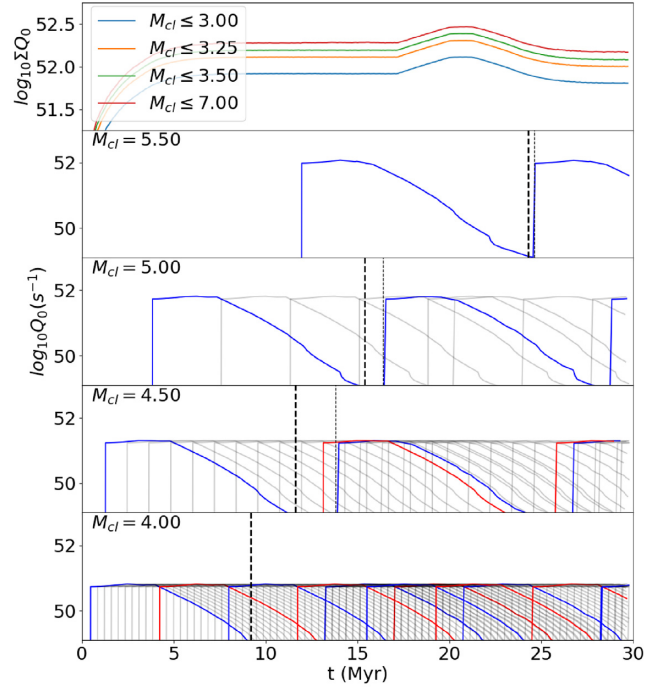
---

[6]Older clusters make a negligible contribution to the emission of the galaxy in the tracers of interest in this study.

This needs to be added to the SFR of newly formed clusters. Since this is initially chosen to be the same as our desired SFR, whenever there is recollapse, the total SFR is initially larger than our desired value. To account for this effect, we adjust the input SFR downwards and repeat the whole procedure, resulting in a new value of $SFR_{rec, 20}$. We continue like this using an iterative shooting method until we match the desired SFR calculated from the gas distribution. We note that this is a highly non-linear process: Changing the SFR changes the integer number of clusters at a given mass that have formed. Since the recollapse time-scale and the question of whether or not a given cloud recollapses both depend on cloud mass, instantaneous star formation efficiency, and cloud density, changing the number of new clusters formed during a given time-step inevitably has a knock-on effect on $SFR_{rec, 20}$. With each iteration, we therefore use the first derivative of the change in $SFR_{rec, 20}$ to anticipate the input SFR, which will produce the desired total SFR. In most cases, we find that no more than four iterations are necessary to reproduce a desired SFR to 1 per cent accuracy. For clouds with low densities, as studied here, the contribution made by recollapsing clouds to the total SFR is below 40 per cent.

It is important to note that recollapse does not represent a periodic evolution in the cluster/cloud system. In the illustrative example presented in this paper, we keep the star formation efficiency fixed and so each new burst of star formation in a given cloud consists of a similar (but slightly smaller) number of stars to the last. However, importantly, these stars add to the cluster that is already present. Although the existing stars are obviously older than the newly formed ones, and so provide less feedback per unit mass, this does not mean that their contribution to the feedback is negligible. In particular, if the interval between bursts of star formation is less than a B star lifetime, stars from the older population will continue to explode as SNe even after the formation of the younger population. In addition, following recollapse, the feedback from the cluster also has less gas to accelerate, since more has been consumed by star formation, and so subsequent star-forming events typically produce more effective cloud dispersal. A full exploration of the effects of recollapse on the shape and normalization of the resulting cluster mass function will be the subject of a follow-up study.

In Fig. 3, we show the time evolution of the ionizing output from our cluster population. To make it visually understandable, in the bottom four panels, we show the time evolution of a single mass bin. Each cluster is shown with a new line. Statistically, the combination of an SFR and a cluster mass function will inevitably lead to an average formation rate of a cluster at a given mass. For our purposes, the cluster 'dies' when its ionizing luminosity drops to undetectable levels.[7] In the present case, we assume that this occurs when the ionizing photon flux Q_0 drops below $Q_0 = 10^{49.0} \, \text{s}^{-1}$, the equivalent of a single O6.5 star, and roughly corresponding to the ionizing photon flux of the Orion Nebula. This time is marked with a heavy vertical dashed line in each panel. It occurs earlier for lower mass, less luminous clusters. Note that we use this threshold only in order to avoid spending large amounts of time computing the Hα and [S III] emission produced by clusters that we know, a priori, will not be significant sources of emission in either line. It plays no special

Figure 3. The time evolution of the cluster population in the central annulus of the galaxy. From the bottom to top are clusters in the $\log_{10} M_{cl} = 4$, 4.5, 5, and 5.0 bins. For visual clarity, every 10th cluster alternates between blue and red to make it easier to distinguish a recollapse event from a new cluster. The heavy dashed line marks the time when the ionizing luminosity of the first cluster formed in each bin decreases to equal $Q_0 = 10^{49.0} \, \text{s}^{-1}$, equivalent to a single O star. Clusters with ionizing luminosities below this value do not contribute significantly to the total emission produced in the lines of interest in this study. This is also the time when the ionizing photon flux of a system with constant SFR reaches a steady state. The vertical dotted line indicates the occurrence of a recollapse event forming a new stellar population in an existing cluster.

role in the underlying WARPFIELD models. In addition, if we were interested in observational tracers of star formation that are less directly dependent on the presence of the most massive stars, e.g. [C II] fine structure emission, then a different choice of threshold would be appropriate.

3.3 Photoionization calculation

We now have the age and mass distribution of the clusters as a function of galactic radius, but we still need to determine their spatial distribution. To do this for a given cluster, we begin by picking a random location in the disc, with a probability distribution set by the SFR/gas mass profile (see Fig. 2). We then check whether the total gas mass within a distance of 50 pc fulfills the condition,

$$M_{r < 50 \, \text{pc}} \geq M_{\text{cloud}}. \tag{5}$$

If this condition is fulfilled, we calculate the centre of mass of the gas within 50 pc of our randomly selected point and place our cluster there. Otherwise, we repeat this procedure with a new randomly selected point. This algorithm ensures that our star clusters are distributed in positions close to, but not necessarily on top of, density peaks of the gas distribution. This is consistent with observations where young clusters are seen in the vicinity of dense molecular gas, but are no longer deeply embedded in their parental clouds owing to efficient stellar feedback.

---

[7]Note that this is a separate issue from whether or not the cluster survives for an extended period of time as a gravitationally bound structure following gas expulsion. Note also that the emission from old clusters is considered to be accounted for by our diffuse interstellar radiation field. In order to avoid overcounting their contribution, once they become faint in ionizing radiation we stop tracking them as individual sources.

Once we have placed each of our clusters in the density distribution, we next compute the density profile of the gas. For each cluster, we randomly select a set of piecewise orthogonal basis vectors. This yields six cardinal directions around each cluster along which we trace $n_g(r)$, stopping once we reach the boundary of the grid.[8] Our WARPFIELD models provide us with the flux of ionizing photons escaping from each cluster at any given moment. We use this as input to a set of photoionization calculations in which we obtain the ionization state, and hence the thermal electron number density n_{th} and electron temperature T_e, as a function of distance from the cluster. Finally, we assume that in other radial directions from the cluster, we can obtain n_{th} and T_e as a function of distance by interpolating between the results of our six calculations. We do this by projecting the unit vector in the radial direction of interest on to our set of basis vectors. This yields a set of three coefficients: the dot products of the radial unit vector with the three basis vectors. We use the magnitude of each coefficient as the weight for that component in the interpolation, while the sign tells us whether we should take the solution in the positive or the negative direction along that basis vector.

The photoionization calculations are carried out using the patched version of CLOUDY v17 (Ferland et al. 2017), and so as a byproduct, we also obtain the local unattenuated emissivities of any emission lines of interest within this region. As an example, in this paper, we present results for Hα and the [S III] 9530-Å line. We stress that the method itself is not limited to these lines, but instead can be used to model any line emission process that CLOUDY or POLARIS can internally handle. Each CLOUDY calculation assumes Milky-Way-like values for the atomic hydrogen cosmic ray ionization rate ($\zeta_H = 2 \times 10^{-16}$ s^{-1}; Indriolo et al. 2007), and the strength and spectrum of the diffuse interstellar radiation field from old stars. Each calculation also adopts solar values for the metal abundances. The actual sphere of influence of the individual clusters typically does not extend beyond ≈ 75 pc in the plane of the disc, outside of which cosmic ray ionization and interstellar radiation from older stars dominate.

At this point, we have a prediction for what the temperature, ionization state, and local volume emissivity of the gas would be as a function of distance from each cluster if that cluster were the only source of radiation (besides the diffuse interstellar background radiation field). We now want to translate this information back on to the Voronoi grid from the Au-6 simulation, which means that we need to decide how to combine these individual radial predictions. To do this, we loop over the Voronoi grid cells. For each cell, we first check whether it overlaps a region directly represented by one of our WARPFIELD models. If it does, then we populate it with the appropriate profile of emission taken from that model. If it does not, then we identify which cluster contributes the largest emissivity to this point in space (see Appendix A for details) and take the necessary values from the radial solution[9] that we have computed for that cluster. In other words, we make the approximation that at any given point in space, the local temperature and ionization are determined solely by the contribution from a single cluster. In practice, we find that this is generally a very good approximation in our models, since only a few

cells see roughly equal ionizing photon fluxes from different clusters. We have experimented with using a more accurate iterative approach to combine the effects of the different clusters on their surroundings, but find that although it is vastly more expensive to calculate, it does not offer significant improvement for our results.

3.4 Galaxy-wide spatial distribution of electrons, temperatures, and emissivities

In the top panels of Fig. 4, we present the PDFs of electron number density n_{th} and electron temperature T_e that we obtain by applying our post-processing scheme to the Au-6 galaxy at redshift $z = 0$. We see immediately that there is a clear bimodality in the electron density distribution. The high electron density branch corresponds to gas that is almost completely ionized. This component is produced by ionizing photons escaping from the clusters and hence is found in close vicinity to them. The low electron density branch, on the other hand, arises primarily from gas that is located at large distances from young clusters and that is only slightly ionized. The ionization of this component is brought about primarily by cosmic rays. A similar bifurcation is recognizable for the distribution of T_e. However, compared to the density PDF the high-temperature branch is less pronounced.

Fig. 5 shows the PDFs for the synthesized emissivities $j_{H\alpha}$ for the Hα line and j_S for [S III]. For the total data set (top panels), there is a range of emissivity values at constant gas density resulting from a range of ionization fractions for Hα. [S III] is more complex due to its higher ionization potential, and the fact that it is not a recombination line, but rather is collisionally excited. In contrast, diffuse interstellar gas between clusters (Fig. 5, bottom panels) shows a well-defined correlation between emissivity and gas density. This is expected because this gas component is dominated by the diffuse ISRF, and thus there is a one-to-one correlation between density and ionization state.

In Fig. 6, we present cuts in the $z = 0$ and $y = 0$ planes through the gas density distribution in our chosen snapshot from the Auriga Au-6 simulation. The corresponding distributions of electron density and electron temperature are consistent with photoionization by stellar sources. The large-scale spiral morphology of the galaxy is apparent in all three plots, but close examination of Fig. 6 reveals a number of additional spots with high electron density, corresponding to the high-ionization regions in the immediate vicinity of each star cluster. This ionization structure strongly influences the local emissivities of the different species. In Fig. 7, we show the local emissivity of Hα (left) and [S III] (right), respectively. While the spiral arms are visible in diffuse Hα emission because they are partially ionized, the gas is too dense to allow for an appreciable abundance of S^{++}, and so they are not visible in [S III] emission. Within the inner 20 kpc of the galaxy, the [S III] emission is completely dominated by the bright spots of emission associated with the individual H II regions.

In Fig. 8, we illustrate how the gas density, electron density, and electron temperature vary as a function of radius and as a function of the height above the mid-plane. The radial plots show azimuthally averaged values computed in the mid-plane ($z = 0$ pc), while the vertical plots show azimuthally averaged values computed at a radial distance $r = 8$ kpc (i.e. a position roughly corresponding to the location of the solar neighborhood in the Milky Way).

The gas density in the mid-plane in Au-6 is roughly comparable with that in the Milky Way at $r \sim 12.5$ kpc, but is clearly higher at both larger and smaller r. This simply reflects the fact that Au-6 is somewhat more gas-rich than the Milky Way. However, the size of the discrepancy is not large: We find at most a factor of a 40 per cent

---

[8]A more accurate approach would be to carry out the photoionization calculation along a much larger set of rays sampling the space around each cluster. Unfortunately, the computational cost of the calculation and the fact that we need to repeat it for a large number of clusters renders this approach computationally unfeasible at the present time.

[9]We use the radius measured in the plane of the disc with an emissivity solution dependent also on z-height above the disc.
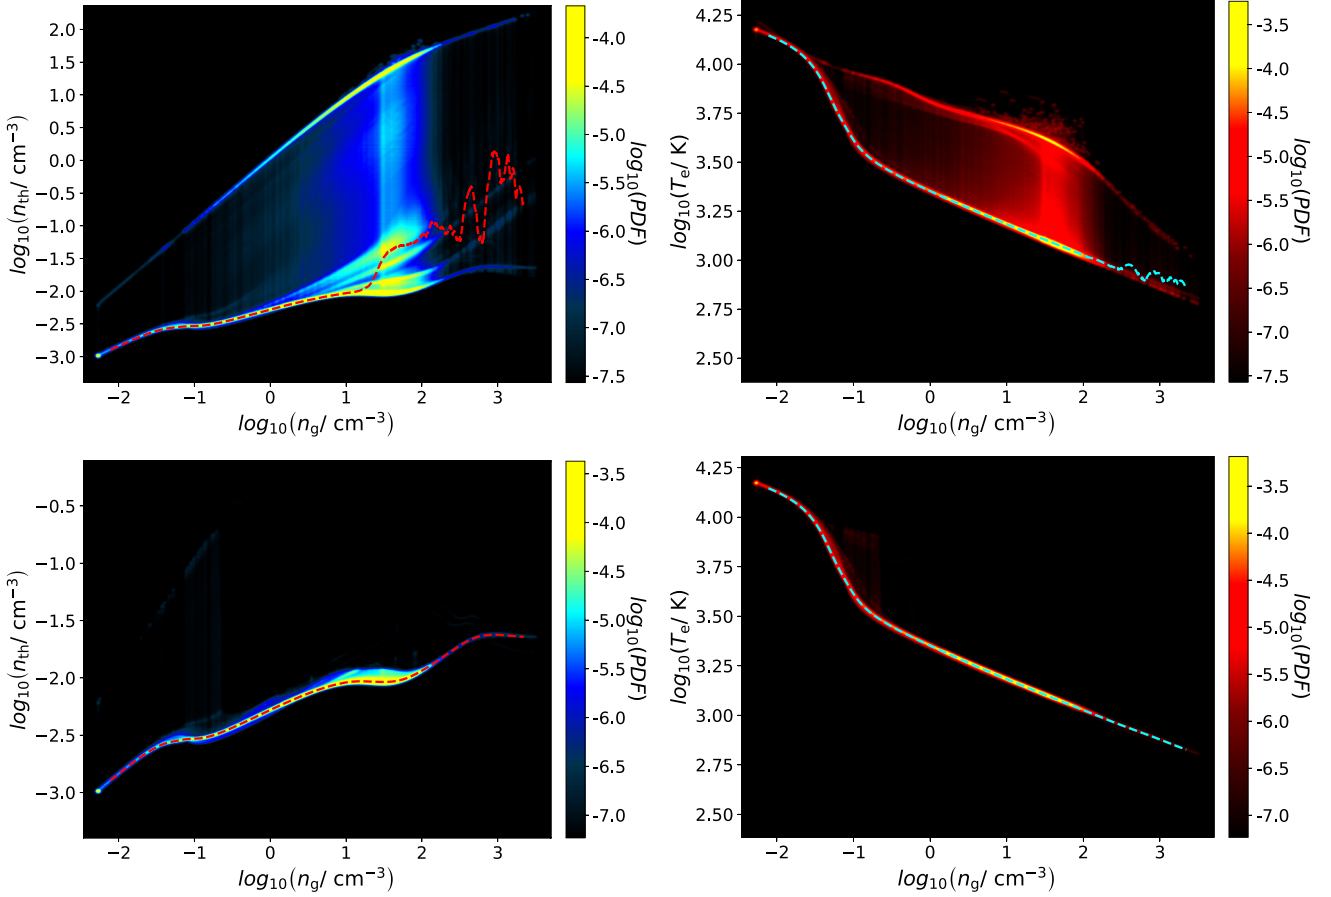
Figure 4. Probability distribution of thermal electron number densities n_{th} (left-hand panels) and electron temperatures T_e (right-hand panels) as a function of the gas number density n_g. Dashed lines indicate the median values. The top panels show the full parameter set, whereas the bottom panels show the results for all regions with a distance $d \geq 75$ pc away from any cluster. We refer to the latter as the ISM condition.

difference between the mid-plane gas density in the model and the value inferred for the Milky Way by Wolfire et al. (2003), meaning the effects of porosity and an increased scale height spread out the extra mass.

For the thermal electron distribution, we find very good agreement with the results of Yao et al. (2017) at radii $7.5 < r < 15$ kpc, but we do not reproduce the increase in the electron density that they find at smaller radii. However, we note that at small radii, any comparison between Au-6 and the real Milky Way will be strongly affected by the lack of a bar in the simulated galaxy, and so it is perhaps not surprising that we do not get good agreement in this regime.

The range of thermal electron temperatures matches well with the parametrization for our own Milky Way presented in Sun et al. (2008) up to $r = 20$ kpc away from the Galactic Centre, but we note an underestimation of the temperature along z. However, the overall magnitude and trends demonstrate the predictive capability of our population synthesis model.

4 PRODUCING SYNTHETIC OBSERVATIONS

We create synthetic maps of line emission in order to further assess the quality of our population synthesis model. For this, we make use of the RT code POLARIS (Reissl et al. 2016), which is capable of dust polarization calculations (Reissl et al. 2017, 2018; Seifried et al. 2019) as well as RT with atomic and molecular lines including

Zeeman splitting (Brauer et al. 2017a,b; Reissl et al. 2018). POLARIS solves the RT problem on the native Voronoi grid of the post-processed Au-6 data set considering both plane external detectors as well as observations inside the grid on a spherical detector, pixellated using Healpix (Górski et al. 2005).

Excluding polarization effects such as non-spherical dust grains or line Zeeman splitting, the RT equation for a velocity channel and along a certain path-length $d\ell$, with emissivity j_ν and opacity κ simply reads

$$\frac{dI_\nu}{d\ell} = j_\nu - \kappa_\nu I_\nu. \tag{6}$$

In this paper, we produce synthetic maps of H α, H β, and [S III] line emission. All of these lines are optically thin and so, in this case, line attenuation is dominated by dust extinction and $\kappa = \kappa_{dust}$. Here, we apply the canonical ISM dust grain mixture (Mathis, Rumpl & Nordsieck 1977; Draine & Li 2001) with 37.5 per cent graphite and 62.5 per cent silicate grains following a power-law size distribution $n(a) \propto a^{-3.5}$ distributed over a grain size range of $a \in [5\,\text{nm}; 250\,\text{nm}]$. We calculate the dust density in each cell with

$$\rho_{dust} = \delta_{DGR}\rho_{gas}\left(1.0 - 0.9 \times \frac{\rho_{H^+}}{\rho_H}\right), \tag{7}$$

where ρ_{H^+} is the mass density of H$^+$ and ρ_H is the total mass density of hydrogen in all forms (H$^+$, H, or H$_2$).
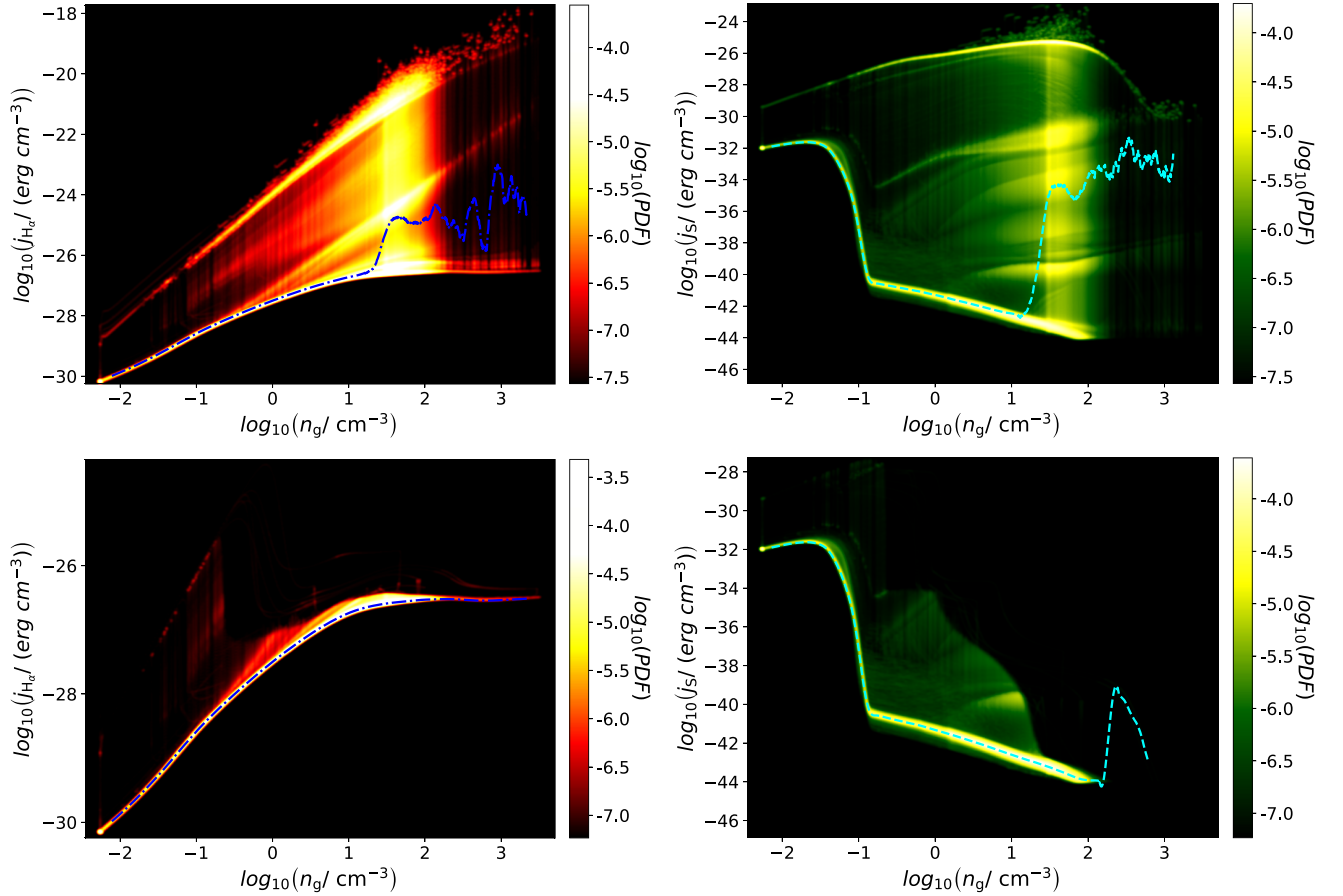
Figure 5. The same as Fig. 4 for the emissivities of Hα (left-hand panels) and [S III] (right-hand panels). The large scatter at constant density reflects the partially ionized nature of the diffuse gas between H II regions. Of note, much of the scatter in the [S III] emissivity comes from partially ionized gas distributed above and below the disc. Ionizing radiation from massive clusters can travel much larger distances in this direction than in the disc mid-plane, and so one finds gas here with a much wider range of densities and ionization parameters than in the mid-plane.

Here, we account for the fact that the amount of dust is lower in close proximity to ionizing sources (Pellegrini et al. 2009; Pellegrini, Baldwin & Ferland 2011). One modification we need to make to the simulation concerns the dust-to-gas ratio δ_{DGR}. We apply $\delta_{DGR} = 0.003$ to reproduce the magnitude of line emission and structure observed in the Milky Way (see Section 5.1), which is close to the lower end of the usual ratios of 0.01–0.003 (see e.g. Whittet 1992; Boulanger, Cox & Jones 2000). However, the simulated Au-6 galaxy is 2.4 times more massive than the Milky Way. Thus, in order to get the correct dust attenuation, the reduction in dust abundance is needed to match the Milky Way opacity per unit length. Note that if we were not interested in comparing with Milky Way observations, or were using a model galaxy that was a closer match to the Milky Way in terms of gas mass, this modification of the dust-to-gas ratio would not be necessary. Further details regarding our treatment of the dust are given in Reissl et al. (2016a,b).

To generate these maps with POLARIS, we use emissivities calculated as described in Section 3.3. These are the frequency-integrated values, and so to get the emissivity at any particular frequency, we need to multiply them by an appropriate line profile function. We assume that Doppler broadening determines the line profile within each cell and also account for the bulk velocity of the cell using the velocity field from the Au-6 simulation. For any given line, we

therefore have

$$j_\nu = j_X \frac{c}{\sqrt{\pi} a_{tot} \nu_{X,0}} \exp\left(-\frac{c^2(\nu_X(v) - \nu)^2}{a_{tot}^2 \nu_{X,0}^2}\right), \qquad (8)$$

where j_X is the frequency-integrated emissivity of the line, with the X representing either Hα, Hβ, or [S III] emission, $\nu_{X,0}$ is the line-centre frequency, $\nu_X(v)$ is the line-centre frequency shifted by the velocity v of the cell relative to the observer, and a_{tot} is the line broadening parameter.

For a_{tot}, we consider a mix of thermal and micro-turbulent broadening and assume that the two contributions are in equipartition:

$$a_{tot}^2 = a_{turb}^2 + a_{th}^2 = 2a_{th}^2. \qquad (9)$$

We therefore have

$$a_{tot}^2 = \frac{4k_B T_{gas}}{m_X}, \qquad (10)$$

where T_{gas} is the gas temperature (which we assume to be equal to T_e) and m_X is the mass of the atom. Note that the assumption of equipartition between turbulent and thermal motions would be a bad approximation if we were treating molecular emission (e.g. CO) from cold gas, but it is a much better approximation for the ionized gas tracers we are interested in here, since their emission
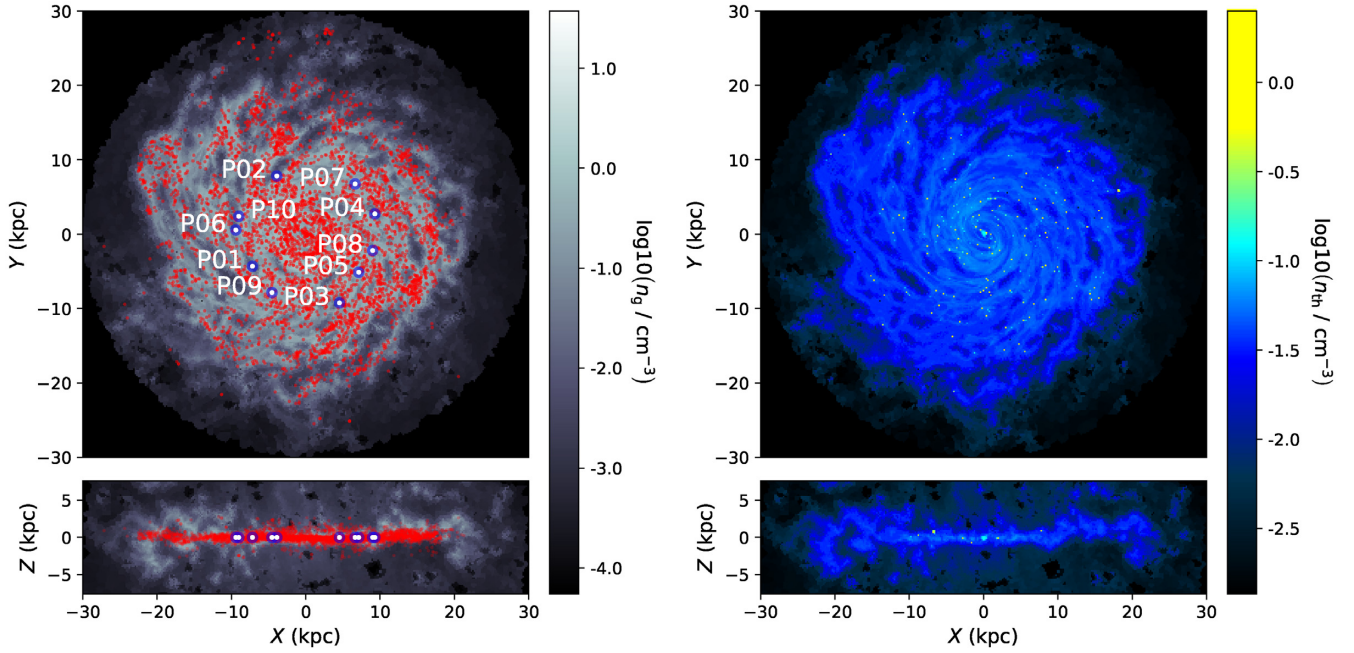
Figure 6. Left-hand panel: mid-plane gas distribution n_g of the Auriga Au-6 galaxy (Grand et al. 2017) at $z = 0$ pc in the x–y plane (top) and $y = 0$ pc in the x–z plane (bottom). Red dots represent all the sampled cluster positions projected on the plane while the blue dots with labels indicate the distinct observer positions (see Section 5). Right-hand panel: the same as the left-hand panel, but for the thermal electron number densities n_{th} as derived in Sections 3.2 and 3.3.
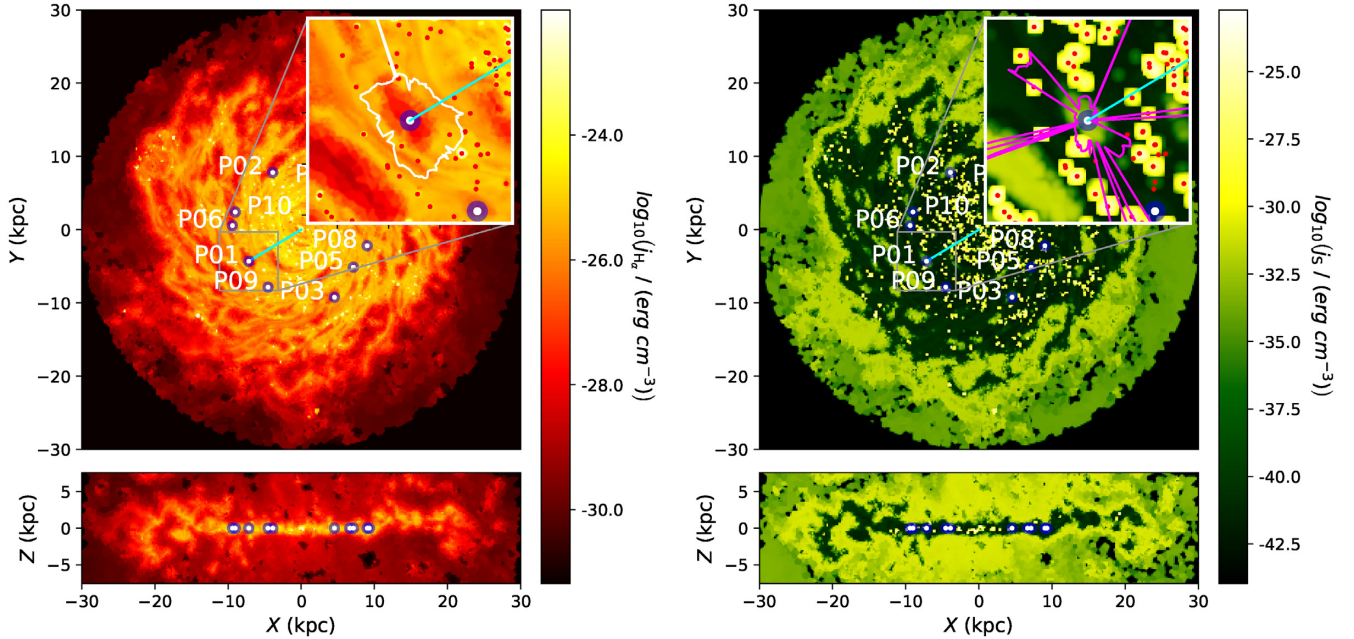

Figure 7. Left-hand panel: the same as Fig. 6 for the H α emissivity. The cyan lines represent the direction towards the Galactic Centre. Zoom-in panel: red dots indicate the sampled cluster positions in the very mid-plane ($|z| < 80$ pc). The marked area around the observer position P01 is the profile of the H α emissivity-weighted distance $\langle d_{H\alpha} \rangle$ (see equation 11) along the galactic longitude (see Section 5.1.1 and compare Fig. 13). Right-hand panel: the same as the left-hand panel for the [S III] emissivity. For the pink marked region in the zoom-in panel, see also Fig. 14.

comes primarily from much hotter regions in which we expect the turbulence to be transonic or only mildly supersonic.

One simplification that we are making here is the assumption that the emissivity within each grid cell is constant. In practice, this is a reasonable assumption. Regions with higher emissivity gradients typically also have higher densities and hence are well

resolved in the simulation. On the other hand, in the large cells above and below the plane, the spatial resolution is poor but the emissivity gradient is small, so we still make little error by assuming a constant emissivity. We note, however, that the WARPFIELD models of the individual star-forming regions still retain their full subgrid resolution.
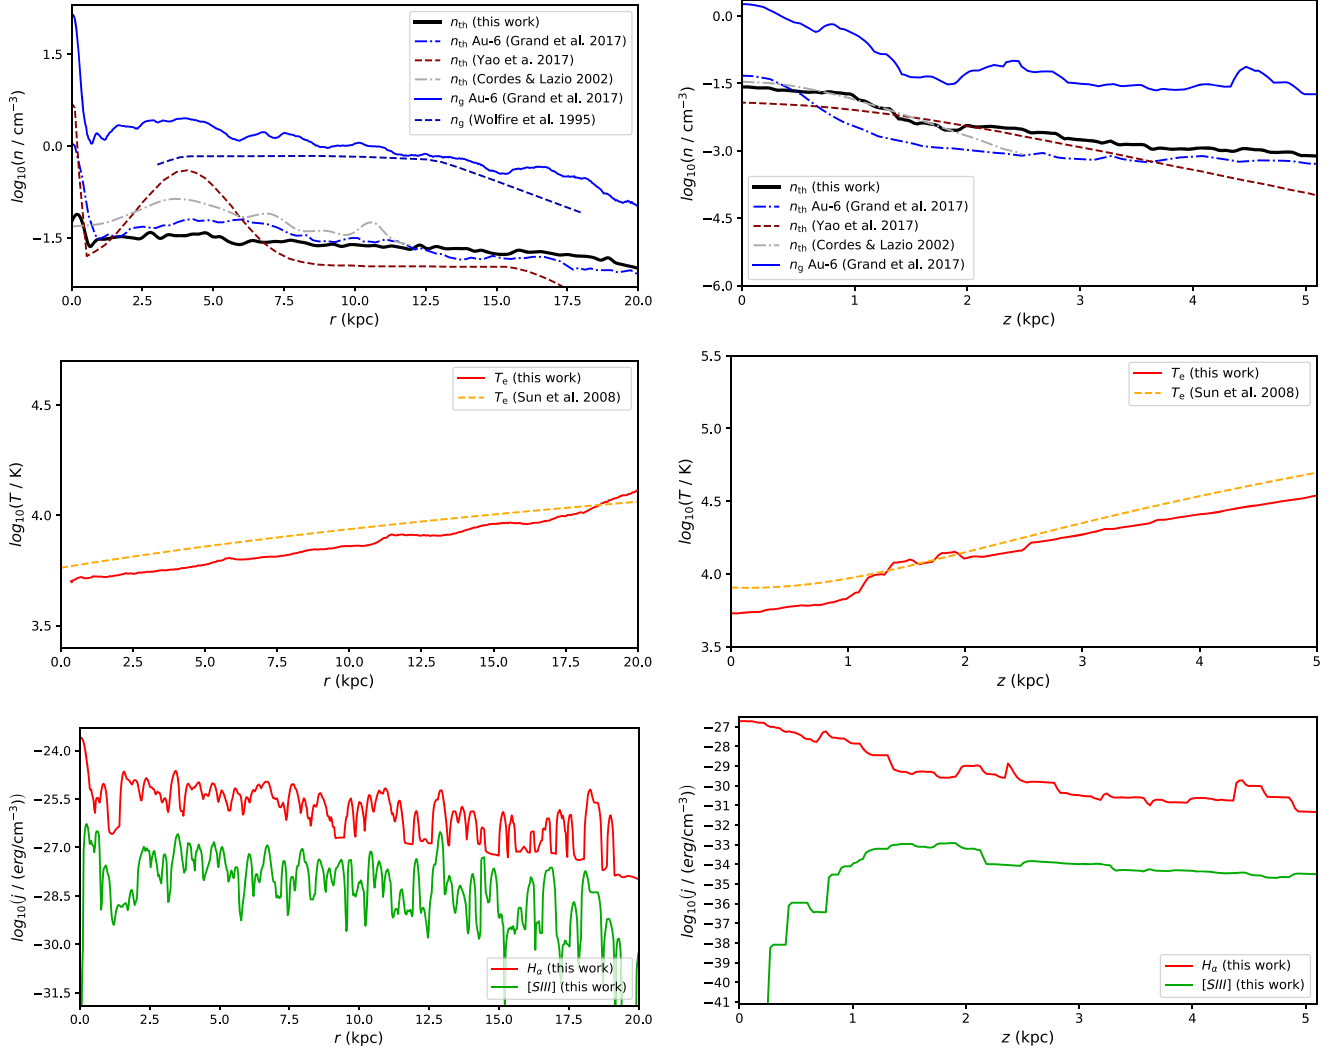
Figure 8. Top left-hand panel: averaged distributions of gas density n_g (blue solid line) and thermal electron density n_{th} (black solid line), corresponding to Fig. 6, along the Galactic radius r at $z = 0$ pc. For comparison, we also show the radial profile presented in Wolfire et al. (1995) for the gas density in the Milky Way (blue dashed line) and the models of Yao, Manchester & Wang (2017) (dash–dotted grey line) and Cordes & Lazio (2002) (dashed red line) for the radial distribution of n_{th} in the Milky Way. Top right-hand panel: as in the left-hand panel, but showing the averaged vertical distribution at a radius of $r = 8$ kpc. Middle panels: the electron temperature T_e (red) and the Galactic temperature parametrization (dashed yellow) presented in Sun et al. (2008). Bottom panels: the Hα emissivity (red) and [S III] emissivity (green), respectively.

Finally, we note that POLARIS itself also allows one to compute atomic and molecular level populations and line emissivities (see e.g. Brauer et al. 2017a), although we do not make use of this capability in our present study.

5 ANALYSIS OF SYNTHETIC OBSERVATIONS

Using the method outlined in the previous sections, we compute synthetic maps of Hα and [S III] emission as seen by observers located at different locations within the simulation volume and we examine the properties of these 'all-sky' maps. Furthermore, we discuss what would be seen by distant observers looking along a line of sight perpendicular to the disc or at any other arbitrary angle.
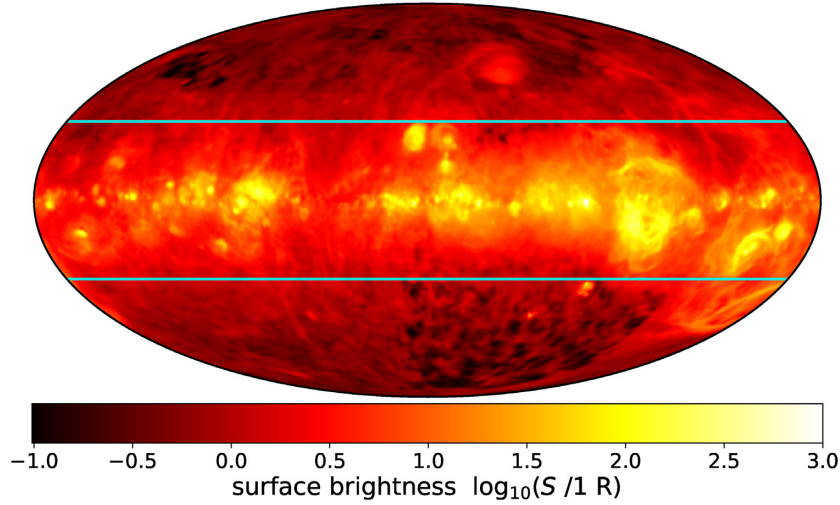
5.1 Galactic all-sky emission maps

For the all-sky maps, we select 10 distinct positions within the Auriga Au-6 simulation in environments with parameters close to our own

solar neighborhood. We identify regions with density comparable to the local bubble our Sun is located in (see e.g. Fuchs et al. 2009; Liu et al. 2017; Alves et al. 2018). These lie within the Galactic plane and at a distance of about $8 \leq r \leq 10$ kpc from the centre. The exact positions are indicated by blue circles in Fig. 6. Our selection ensures that the resulting all-sky maps can be meaningfully compared to real all-sky observations on Earth, and that they are not dominated by signals from nearby dense clouds or young massive clusters that are not present in the real data.

There are three resolutions to consider when comparing synthetic emission maps to real observations. First is the resolution of the observed Hα map, which is a combination of data from multiple surveys, with a range in angular resolution from a few arcmin to $1°$. Next is the resolution of the Healpix pixelation scheme used in POLARIS. The all-sky maps are calculated with rays distributed according the Healpix pixelation scheme with $N_{side} = 256$ subdivisions per side of the 12 base pixels, resulting in a total of $N_{pix} =$

Figure 9. Observed all-sky Hα map as presented in Finkbeiner (2003). Blue horizontal lines indicate the disc region within $|b| < 30°$ that we consider when analysing the structure of the emission maps.
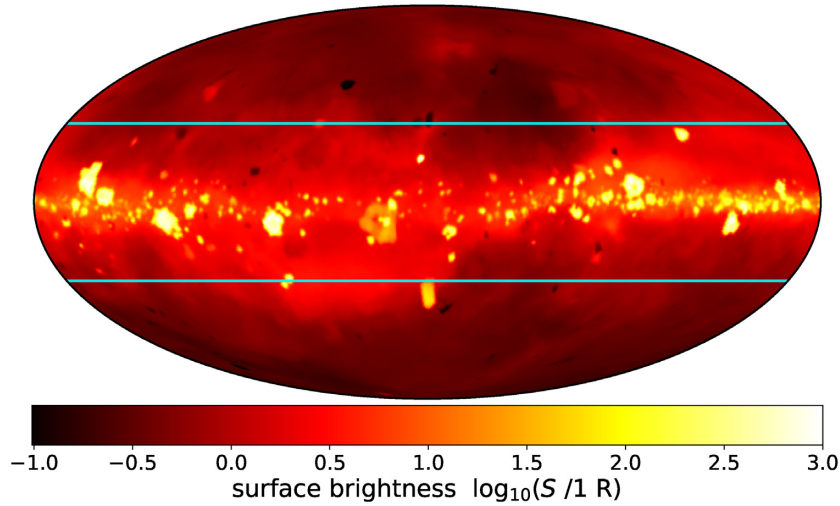
Figure 10. As Fig. 9, but showing the synthetic all-sky Hα emission map as seen by an observer located at position P01 in the model disc. The map is integrated over a velocity range of $\pm 200\,\mathrm{km\,s^{-1}}$.

$12 \times 256^2 = 786\,432$ pixels over the entire sky. This corresponds to an angular resolution of about 13.7 arcmin, or a physical scale ~ 40 pc at a distance of 10 kpc. The last resolution is the resolution of the Voronoi grid from the Auriga simulation. This varies as a function of position and is also seen in projection at different distances, so that the same physical size of grid cell corresponds to very different angular sizes depending upon whether it is close to or far from the observer. In practice, our Healpix resolution is sufficient to ensure that even the smallest cells in the Au-6 Voronoi mesh are sampled with one or more rays.

Fig. 9 shows a reprocessed all-sky Hα map based on data from the VTSS and SHASSA surveys and centred towards the centre of Milky Way, as presented by Finkbeiner (2003). The map has a resolution of about 3.4 arcmin ($N_{\mathrm{side}} = 1024$) and shows characteristic multiscale patches of glowing hot ionized hydrogen gas surrounding star-forming regions. Fig. 10 presents the corresponding synthetic Hα map of the post-processed Au-6 galaxy for the observer position P01. The overall structure of Hα patches as well as the magnitude of the emission match rather well. However, the angular

nature of some of the bright and dark patches in the synthetic maps still reflects the underlying Voronoi grid geometry rather than any actual physical effect. This demonstrates that the grid resolution of the underlying numerical simulation is one of the limiting factors in this kind of synthetic image calculation. We note that in the anticentre directions, a large fraction of the disc is observed at lower resolution than our predictions, and this could account for the higher contrast seen in some of our star-forming regions.

Of course, our post-processing method is not limited to lines such as Hα that have already been widely observed in the Milky Way. As an example, in Fig. 11, we show our prediction for what an all-sky map of [S III] emission would look like. The large-scale structure is similar to the Hα maps. However, the [S III] emission is less affected by the diffuse line emission and dust extinction, allowing it to probe deeper into the Galactic disc (see also Section 5.1.1). We emphasize that this result comes without any fine-tuning of our post-processing pipeline. Instead, it is a direct result of Galactic population synthesis modelling from first principles.
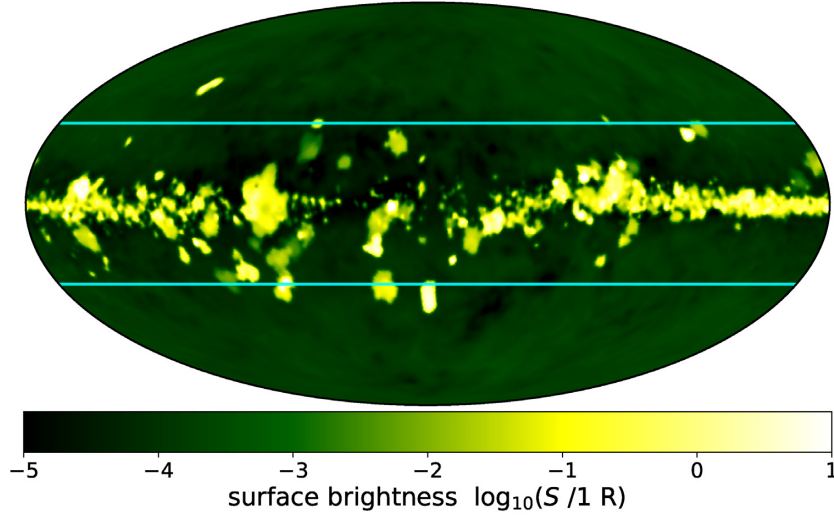
Figure 11. Same as Fig. 9, but showing the synthetic all-sky [S III] emission map as seen by an observer located at position P01 in the model disc.
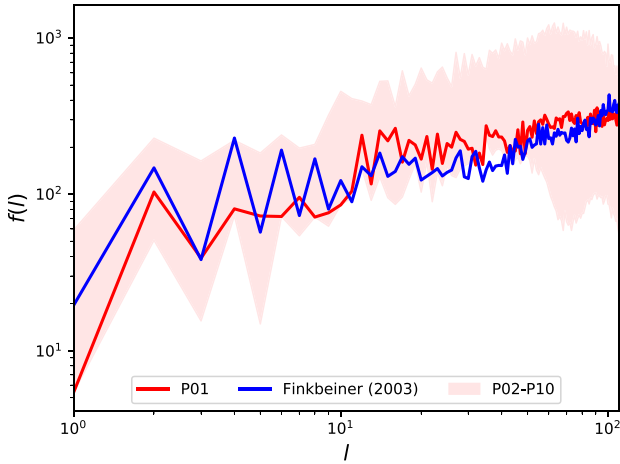
Figure 12. Multipole spectrum as a function of multipole moment l of the synthetic Hα map for the different observer positions P01–P10. For comparison, we also show the spectrum derived from the map presented in Finkbeiner (2003). The analysis is done for the designated Galactic disc with a latitude of $|b| < 30°$.

A quantitative analysis of the structure of the synthetic Hα all-sky emission maps is provided in Fig. 12. Here, we apply a multipole analysis as outlined in Appendix B. The multipole moment is plotted down to a resolution of $1°$. Due to the variable angular resolution in the observed map, it is difficult to meaningfully compare our synthetic map with the observations on scales smaller than this. Probing this issue farther will be possible with forthcoming Hα surveys at half-arcmin resolution.

For the analysis, we disregard regions above and below the plane with a galactic latitude of $|b| > 30°$. The inner region is defined to be the Galactic disc region. We have three motivations for focusing on this region:

(i) The Milky Way observations farther from the disc begin to become dominated by a combination of noise and survey artifacts related to tiling.

(ii) The Milky Way observations outside the disc are contaminated by extragalactic sources such as the Large and Small Magellanic Clouds that are obviously absent from our synthetic maps.
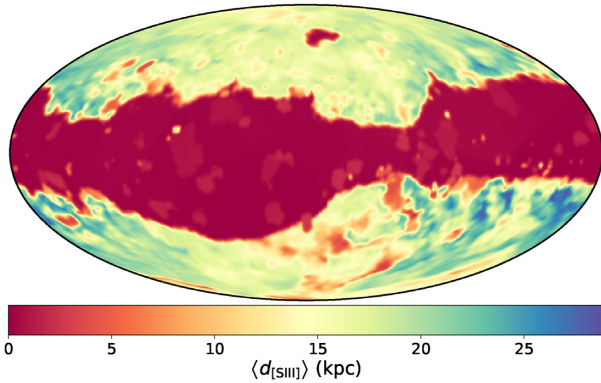
(iii) In the Au-6 simulation, lines of sight toward the poles are dominated by low density gas, which is represented by a relatively small number of physically large Voronoi cells. We therefore see an increase in the number of grid artefacts as we look towards the poles in the synthetic all-sky maps.

The resulting multipole spectra in Fig. 12 reveal a characteristic zigzag pattern for the large-scale structure. A similar pattern can be observed in observations and synthetic images of Milky Way synchrotron emission (Haslam et al. 1981; Reissl et al. 2019a), with the magnitude slowly declining towards small scales. In contrast to Galactic synchrotron emission, the overall trend for Hα emission is the opposite with a minimum at a multipole moment of $l = 0$ and an increase towards smaller scales. This trend and the zigzag pattern is common to all considered observer positions, although some positions have more small-scale power compared to the Milky Way. We attribute this difference in structure to the local conditions surrounding each observer position. The number of clusters in the proximity of each observer ranges from one to several dozens. This influences the multipole fitting significantly. Another contributory factor is the dust extinction. The depth to which one can see in the mid-plane is sensitive to the local dust distribution, and so from some of our example observer positions, one can see clusters at much greater distances than from other observer positions. As more distant clusters have small angular scales, this translates into a considerable variation in the amount of small-scale power seen from each position. It explains the excess in magnitude from $l \approx 20$ onward compared to the spectrum of the Milky Way that we find for many (but not all) observer positions.

This result also demonstrates the necessity of accounting for dust extinction when making this kind of synthetic image. Ignoring the dust extinction in the RT calculations would lead to an increase in the small-scale structure of the Hα maps. Without extinction, the Hα emission penetrates the entire disc, meaning that all of the clusters present within the disc would contribute to the image.
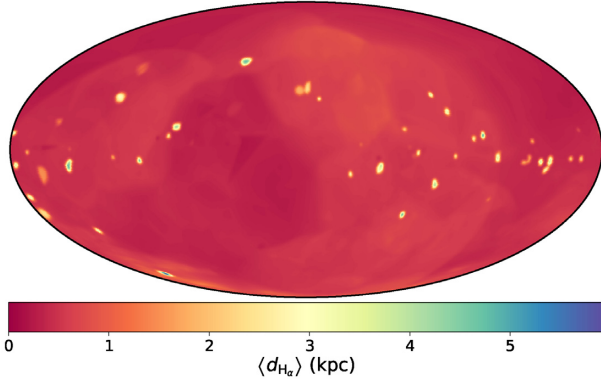
Figure 13. All-sky map of the Hα emissivity-weighted $\langle d_{H\alpha} \rangle$ as seen from position P01.

Figure 14. The same as Fig. 13 for $\langle d_{[S\,III]} \rangle$.

Figure 15. All-sky map of the ratio $\langle d_{[S\,III]} \rangle$ to $\langle d_{H\alpha} \rangle$.
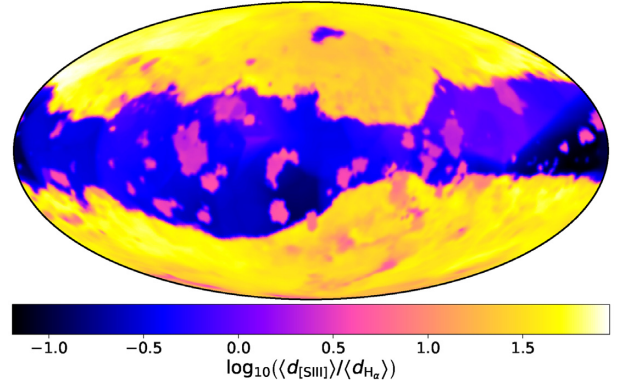
5.1.1 Source of the Galactic emission

In Fig. 11, we show a synthetic all-sky map for Galactic [S III] emission. The large-scale structure is different than in the Hα maps, as it is less extended due to the lower relative emissivity of diffuse gas. Individual star-forming regions have high contrast because there is less diffuse emission and lower dust extinction, allowing the [S III] line to probe deeper into the Galactic disc. As a measurement of the origin of the Hα and [S III] emission within the grid, we define the emissivity-weighted distance along the line of sight as

$$\langle d \rangle = \frac{\int_{\ell'}^0 \mathrm{d}(\ell) \times j_\nu \exp\left(-\tau_\nu(\ell)\right)\,\mathrm{d}\ell}{\int_{\ell'}^0 j_\nu \exp\left(-\tau_\nu(\ell)\right)\,\mathrm{d}\ell}, \tag{11}$$

with a distance d(ℓ) between the observer and the particular position ℓ and an optical depth of $\tau_\nu(\ell)$ from the observer position up to ℓ.

In Figs 13 and 14, we show an all-sky map of the emissivity-weighted distances for the Hα and [S III] emission, which helps to highlight the lines of sight for which these tracers probe very different distances. In Fig. 15, we show the ratio of emissivity-weighted distances.

The Hα emission seen by an observer located (as the Solar system is) within a low density bubble is dominated by diffuse ionized gas (DIG) with relatively low extinction, and typical distances of about 1 kpc. Along lines of sight toward bright, relatively nearby star-

forming regions, their emission outshines the local diffuse emission, resulting in a higher emission-weighted distance of about 2 kpc.[10]

The morphology of the [S III] effective distance map is markedly different. In the Galactic plane, many sightlines are once again dominated by relatively nearby diffuse emission, while some are dominated by brighter, more distant star-forming regions. However, the characteristic distances of both components are larger, and so there is more structure in the diffuse map, and the compact H II regions are typically farther away. The origin of this is a matter of statistics and extinction. Statistically, the probability of finding more massive clusters, which are required for bright [S III], is lower than that of finding lower mass clusters. This translates to a larger typical distance between massive regions, which leads to a larger ratio of distances in Figs 15. We also see that there is a dramatic increase in the emission-weighted distance of the [S III] emission as we look out of the plane. This comes about because there is so little local emission in these directions that we start to become dominated by the faint signal from the hot gas in the halo, although we see from Fig. 11 that, in practice, this signal is likely too weak to be detectable.

While purely theoretical now, planned missions to map the majority of the star-forming disc in all optical emission lines (including [S III]) are under construction (e.g. SDSS-V; see Kollmeier et al. 2017). We find [S III] will make it possible to trace obscured high-mass star formation up to five times farther in the disc than Hα, partly due to extinction, and partly due to less confusion with diffuse gas, which is much fainter in [S III] than Hα, due to reduced diffuse [S III] emission (see Fig. 7, right-hand panel). Catalogues of star-forming regions, or selection functions based on [S III] will be less sensitive to galactic structure, and more sensitive to star formation and population characteristics.

5.1.2 Hα/[S III] ionization parameter mapping

Spatially resolved emission ratio maps employing high-to-low ionization potential tracers (Pellegrini et al. 2012), as well as dust tracers sensitive to photon flux (Oey et al. 2017a; Binder & Povich 2018), can be used to map local variations in the ionization parameter $U = n_{\mathrm{ion}}/n_{\mathrm{H}}$, i.e. the ratio of the ionizing photon number density to the atomic hydrogen number density. Such ratio maps depend on the intensity and spectral shape of the ionizing radiation, and ISM

---

[10]Note that as the diffuse gas still contributes along these sightlines, this is an underestimate of the actual distance to these star-forming regions.
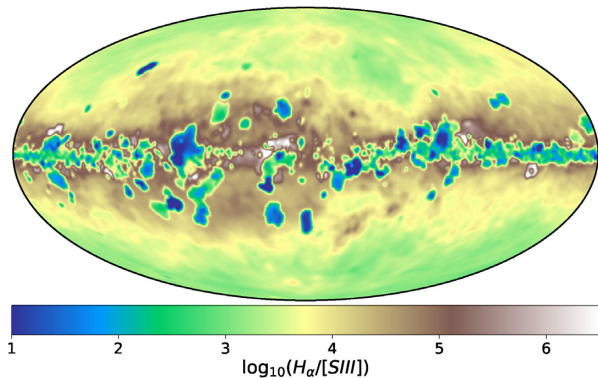
Figure 16. All-sky map of the Hα/[S III] ratio as seen from position P01.

structure in and around star-forming regions. These dependencies make such maps useful for separating distinct star-forming regions from each other (Pellegrini et al. 2012), as well as separating these regions from larger scale galactic structure. On small scales, local variations are dominated by local gas ionization structure that depends on the optical depth to Lyman continuum radiation of individual star-forming regions. This makes it possible to quantify the relative number of optically thick (radiation bounded) to optically thin (density bounded) H II regions, as well as to measure the covering fraction of blister type H II regions. H$^+$ ionization fronts have a width equal to the mean-free path of ionizing photons. Due to the high cross-section for ionizing radiation, this is typically short ($\ll$pc; see e.g. Osterbrock & Ferland 2006). As ionizing photons are depleted due to radiative transfer, the local ionization degree of the gas drops. The relative abundance of more highly ionized ions decreases first, owing to their faster recombination rates, resulting in large changes in emission-line ratios. Emission-line ratio maps can therefore be used to highlight the locations where the ionization fraction of the gas is rapidly decreasing, allowing them to reveal the existence of an ionization front in dense gas (Pellegrini et al. 2012).

Emission-line ratio maps are often interpreted in terms of large-scale ISM evolution. However, this interpretation assumes that the emission traces gas in the ISM, which has dynamically responded to feedback (e.g. by being driven into large shells). However, without the aid of simulations, there is no direct way to know if the termination of the flow of radiation in a given direction occurs in the natal gas cloud in which the massive stars have formed, or if this cloud has already been fully ionized, leading to the ionizing radiation being absorbed farther away by surrounding material not directly associated with the star formation and unaffected by mechanical feedback (winds, supernovae, etc.) from the young stars.

Fig. 16 shows the all sky Hα/[S III] ratio. Despite the spherical symmetry of the input WARPFIELD models, the ratio maps reveal low-ionization gas surrounding more highly ionized regions, with irregular morphologies. The significance of this is twofold. First, it means that even when the evolution of individual star-forming regions is idealized as spherically symmetric (which will tend to underestimate the escape fraction of ionizing photons compared to what one would find in a more irregularly structured cloud), the emergent radiation is strong enough to ionize the surrounding ISM, thus creating a DIG component that contributes to the observed total flux. The surrounding gas is dense enough to absorb the radiation over a short distance, making the DIG appear as a continuation of the now ionized natal cloud. This introduces a difficulty in testing physics by comparing models of isolated individual clouds

and their evolution to observations, even when the observations are able to spatially resolve individual star-forming regions. These comparisons require making an observational determination of what gas is involved in star formation, and distinguishing between gas that has been accelerated and swept into shells by feedback versus gas, which is only illuminated by the escaping radiation. For these reasons, a proper comparison of observations to models must not only include a self-consistent population of star-forming regions, but also the extinction and emission from the environment in which they form.

5.2 Applications to non-Milky Way galaxies

Viewed from an external position, our model galaxy should be equivalent to a typical L_* spiral galaxy. Fig. 17 and Figs C1–C6 show the galaxy as seen from different orientations, ranging from face-on (0°) to edge-on (90°). These maps were generated with a projected grid of 1024 pixels, covering a total area of 60 × 60 kpc and hence are sensitive to structures on scales down to ~60 pc. For comparison, recent IFU observations of NGC 628 provide spectroscopic information on many H II region diagnostic lines at a spatial resolution of ~40 pc (Kreckel et al. 2018; Rousseau-Nepton et al. 2018b), and the ongoing PHANGS-MUSE[11] and forthcoming SIGNALS[12] surveys will do the same for a much larger sample of nearby spiral galaxies at similar resolutions. We expect our models to play a key role in helping to guide the interpretation of these forthcoming surveys.

5.2.1 Extinction

When interpreting extragalactic observations it is necessary to understand simultaneously the role of dust, diffuse emission and crowding. Hα observations play an important role in understanding star formation (Kennicutt 1998), and the luminosity function of star-forming regions helps to constrain the initial mass function of clusters, and SFRs. In the absence of dust attenuation, Hα emission is roughly proportional to the SFR. In practice, Hα is attenuated by dust extinction, which must be corrected for in order to recover an unbiased measure of the SFR. This correction is typically carried out by observing Hβ in addition to Hα. As the Hα and Hβ lines have different frequencies, they suffer from different amounts of dust extinction. Therefore, if the intrinsic ratio of Hα to Hβ is known, the observed ratio can be used to infer the amount of dust absorption.

A full parameter space exploration of the effect of both dust and crowding is beyond the scope of this paper, and so we can only give a broad overview here. In order to quantify the effects of differential extinction for our model galaxy, we calculate the radial profile of Hα and Hβ emission measured by an external observer for inclination angles ranging from 0° to 90° at 15° intervals. We also compute the 'true' Hα and Hβ fluxes, neglecting all dust absorption. We note that we do not neglect the effects of dust absorption when calculating the ionization state of the gas. Instead, we merely set the dust opacity to zero when ray-tracing the resulting Hα and Hβ emission.

As an observer would also do, we compute the de-reddened Hα flux from the Balmer decrement. We assume that the extinction can be modelled as a foreground screen, and that an $R_V = 3.1$ reddening
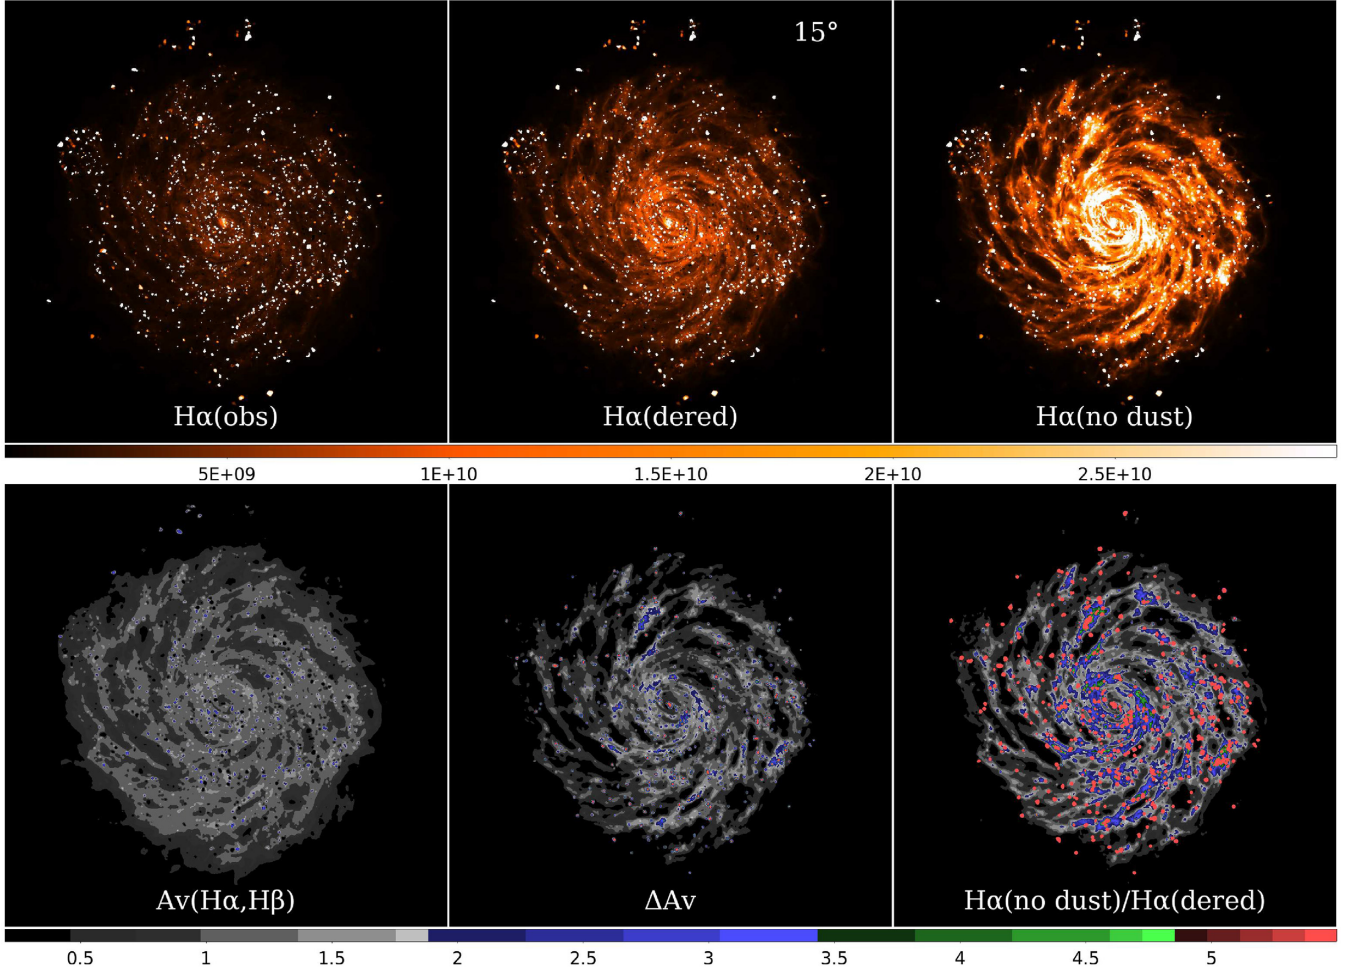
Figure 17. Log of the Hα emission maps, with a range of 10^7–5×10^{10} Jy, viewed at an inclination of $i = 15°$ (top panels) from the left- to right-hand side: observed, de-reddened using the Balmer decrement, and with no dust absorption. In the bottom panels, we show the derived extinction in A_V magnitudes (left-hand panels), ΔA_V (middle panels), and the ratio of the dust-free Hα emission to that recovered with standard Balmer decrement de-reddeing.

law applies. We also assume that the intrinsic Hα/Hβ ratio is 2.86.[13] The extinction is calculated (and de-reddening applied) at the native resolution of our ray-traced images, namely 60 pc. The different Hα emission and extinction maps are shown in Fig. 17. We define the deviation ΔA_V as the difference between the value derived from the ratio of dust-free and observed Hα fluxes and A_V derived from the Hα/Hβ ratio,

$$\Delta A_V = 2.5/f_{H\alpha} \times \log_{10}\left(F(H\alpha)_{\mathrm{no\ dust}}/F(H\alpha)_{\mathrm{obs}}\right) - A_V, \quad (12)$$

where $F(H\alpha)_{\mathrm{no\ dust}}$ and $F(H\alpha)_{\mathrm{obs}}$ are the Hα flux in the dust-free map and the full calculation, respectively, and $f_{H\alpha} = 0.818$ is the ratio of the extinction at the wavelength of Hα to the extinction in the V band, assuming a standard $R_V = 3.1$ reddening law.

We see from the figure that the standard de-reddening correction does a relatively good job in regions where the gas density is relatively low (e.g. at large galactocentric radius). However, it tends to systematically underestimate the true amount of obscuration in regions of high gas density, with this effect becoming particularly pronounced as one nears the centre of the galaxy.

[13]In reality, the ratio is a weak function of the temperature and density of the ionized gas, but we neglect this complication here.

In Fig. 18, we show how the ratio of the intrinsic to the de-reddened flux varies as a function of deprojected galactocentric radius (i.e. the radius as measured after correcting for inclination) for a range of different galactic inclinations. The values depicted are those obtained after averaging over an annulus of thickness d$r = 3$ kpc, which has the effect of smoothing out small-scale variations, e.g. associated with spiral arms. When de-projecting the 'observed' maps, we assume that all objects are in an infinitely thin plane. We also extend the outer radii to 45 kpc so that H II regions at significant heights above the galactic mid-plane do not fall outside of the deprojected image.

Several different physical effects influence the form of these profiles. Our WARPFIELD models have internal extinctions that directly affect the amount of emergent Hα emission and the intrinsic surface brightness of the H II region. As galactic inclination increases, the amount of extinction along the line of sight to these clouds decreases their flux. Simultaneously, along the same line-of-sight diffuse gas emission begins to compete until the pixel is dominated by DIG emission, and individual H II regions become hidden.

So long as the emergent light of the H II regions is brighter than the DIG, we can use the standard de-reddening technique to correct for the effects of extinction and recover a good estimate of the intrinsic Hα emission. When the inclination angle is low, the surface
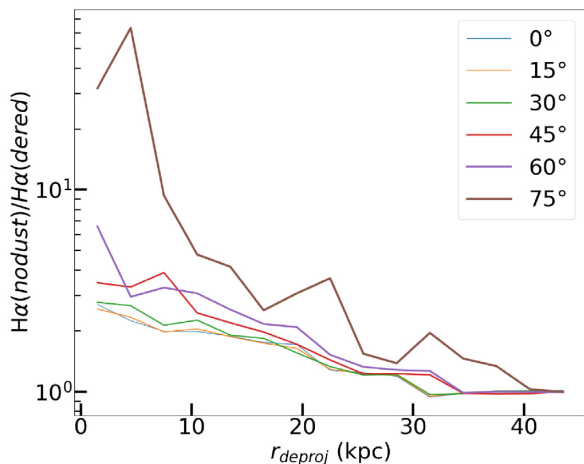
Figure 18. The ratio of de-reddened, deprojected intrinsic no-dust $H\alpha$ flux measured in annular apertures, to that recovered using the $H\alpha/H\beta$ ratio, for inclination angles ranging from $0°$ to $75°$. Note that we ratio the fluxes after integrating over the aperture, as opposed to calculating the average ratio in the aperture.

brightness of the H II regions remain high in general and the Balmer correction is effective. Nevertheless, even in this case, some clusters are so embedded that they are effectively hidden by the DIG. This is a rare occurrence in the outer reaches of the galaxy, but becomes more common as we move towards the centre, resulting in a steady increase in the ratio of intrinsic to de-reddened flux with decreasing galactocentric radius.

However, as the inclination of the galaxy increases, a point is reached where emission from the foreground DIG dominates the emission in both the $H\alpha$ and $H\beta$ lines, with the latter being more affected. Once this occurs, the measured Balmer decrement simply traces conditions in a low optical depth layer of the DIG, and hence no longer allows us to accurately correct the flux from the H II regions. As a result, the de-reddened $H\alpha$ flux can in this case dramatically underestimate the intrinsic flux, by a factor of 10 or more (see e.g. the behaviour of the $i = 75°$ galaxy at small deprojected radii in Fig. 18). As one would expect, this is a much bigger problem at small radii, where the diffuse emission is bright and the foreground extinction is considerable, than at large radii, where the diffuse emission is fainter, and gas density is lower.

6 RESULTS

We have analysed our Milky Way analogue in three ways: first, in terms of local conditions, such as the gas temperature and ionization state, where we see significant variations resulting from the variation in the fraction of the local radiation field coming from young hot stars versus the older, evolved stellar population, and, secondly, as observed in two ionized gas tracers from locations representative of the position of the Solar system in the real Milky Way. Finally, as it would appear to an observer seeing the Milky Way from the outside. In this section, we summarize and discuss a few of the key results of this analysis.

6.1 Synthetic all-sky Milky Way emission maps

To demonstrate the power of our new approach, we have presented synthetic $H\alpha$ and [S III] emission maps produced using a simulated

Milky-Way-like galaxy. We have compared the synthetic $H\alpha$ maps to observations of $H\alpha$ in the real Milky Way. The same comparison cannot be done for [S III], since no large-scale maps of this line yet exist for the Milky Way, and so in this case, our results are predictions for what will be observed by future large surveys such as the Local Volume Mapper project.[14]

We find relatively good agreement between our synthetic $H\alpha$ maps and observations of the Milky Way. Since the observed $H\alpha$ flux depends both on the internal structure and extinction of the individual H II regions and also the larger-scale distribution of gas density and ionization within the galaxy, the fact that we find good agreement with the observations suggests that our model is doing a reasonable job of capturing both of these features of the real galaxy.

To produce these results, we adopt a mean cloud density[15] of $n = 100 \, \text{cm}^{-3}$ and a star formation efficiency on cloud scales of $\epsilon \approx 1$ per cent. With these parameters, we find that a significant number of the clouds undergo re-collapse and form multiple stellar populations with age spreads of a few Myr. We have not explored in detail the sensitivity of our results to variation of these parameters, but we can nevertheless make some qualitative statements. For example, if we had assumed a much higher average cloud density, such as $n = 1000 \, \text{cm}^{-3}$, then many more clouds would have undergone re-collapse. Moreover, cloud expansion would typically have stalled at a much earlier stage, resulting in far higher internal extinctions and little resulting $H\alpha$ emission. In this case, our average H II regions would have been significantly fainter, meaning that we would no longer match the observations on both large and small scales. In this way, we see that input parameters such as the mean cloud density that are difficult to directly constrain from observations can be indirectly constrained by finding the range of values for which our synthetic maps match the real ones.

Our maps of [S III] emission show that it typically penetrates to much greater distances in the galactic disc than $H\alpha$ and is also less confused by foreground emission. Both of these features can be easily understood in terms of the basic physics of the [S III] line. Ionizing sulphur to S^{++} requires significantly higher energy photons than ionizing hydrogen to H^+, and so [S III] emission primarily traces regions where the flux of these energetic photons is large, i.e. regions close to massive clusters, with little being produced in ionized gas lying far away from the clusters. Once emitted, the [S III] photons propagate further than $H\alpha$ photons simply because the difference in their wavelengths makes them much less susceptible to attenuation by dust.

The fact that [S III] penetrates the disc of a galaxy much more than $H\alpha$, and that it is less confused by foreground emission means that at low angular or spatial resolution, simple single object photoionization models will almost certainly fail to reproduce its relative intensity compared to other emission lines, since the observations will be probing emission produced by the superposition of many distant sources with small angular size. This is much less of an issue for $H\alpha$ because the sources probed by that line are typically much closer, and hence have larger angular sizes and suffer less from confusion.

[14]https://www.sdss.org/future/lvm/.
[15]Note that this value is the mean density of the entire cloud, including any CO-dark molecular gas or atomic gas associated with it. The mean density of any portions of the cloud traced by bright CO emission could plausibly be somewhat higher.

A variable, but significant escape fraction of ionizing radiation from individual star-forming regions can also create locally bright DIG with morphologies indistinguishable from traditional H II regions. The emission of the illuminated DIG can compete with that from our inserted WARPFIELD models, leading to potentially ambiguous interpretations of the expansion of star-forming regions, absent kinematic data. To account for this, it is important for models of the coupling of feedback with individual molecular clouds to also account for the escape of ionizing radiation and its impact on the larger-scale surrounding environment, as we aim to do in the method presented here.

6.2 Extragalactic systems

We have produced Hα and [S III] emission maps to illustrate the connection between small-scale physics and kpc-scale observations. Projection of the model galaxy from the view point of an exterior observer makes it possible to compute the emission from DIG and from individual H II regions simultaneously in a self-consistent fashion. We find the intrinsic-to-recovered Hα emission to systematically vary with galactic radii, depending on local extinction, DIG emissivity and, critically, on the fraction of objects in the deeply embedded phase, which, in turn, depends on cloud evolution and the local SFR. Winds and radiation have very different impact on the evolution of clouds with different masses. For the cloud population in an entire galaxy, this becomes a function of the cloud mass distribution and the global SFR, as outlined in Section 3.2. For some tracers, this may result in the emission from part or all of the galaxy becoming highly stochastic in the case where the emission of the tracer is dominated by the youngest and most massive clusters. This will have significant implications for the interpretation of observed emission line ratios and the derivation of physical parameters such as the overall SFR and the gas metallicity (see e.g. Kewley et al. 2001; Dopita et al. 2016; Richardson et al. 2019). The characterization of this uncertainty is highly challenging and requires the use of a time-dependent population synthesis model as introduced in this study. We plan to explore this issue in more detail in a future paper.

6.3 Impact of deeply embedded star clusters

Our model predicts that a Milky-Way-type galaxy should contain a significant population of faint and deeply embedded star-forming regions ($A_V \geq 5$) at any given time. These potentially spend up to tens of Myr-forming stars in a cycle of collapse and expansion. This result depends not only on the detailed modelling of all relevant stellar feedback processes, as implemented in WARPFIELD, but also on how nature samples density, mass, and star formation efficiency within the galaxy. Because this population of young star clusters is very difficult to measure, it makes the interpretation of the observed line fluxes more difficult and introduces uncertainty to many inferred physical parameters such as SFR or metallicity. When observing external galaxies, we could demonstrate that we are able to recover the total intrinsic dust-free Hα flux to better than 50 per cent error by using the Hα/Hβ ratio for viewing angles of less than $45°$. This is in line with expected values for de-reddened galaxies (Calzetti 2001). However, for more edge-on galaxies, the uncertainty can be as large as a factor of 10. We note that direct application of such results to observations should be done with caution. It would be desirable to derive a correction factor that only depends on the relative viewing angle to the galactic disc. However, this will be difficult as the result may also depend on natal cloud density, star formation efficiency,

and metallicity, as well as galactic morphology. This requires further investigation, and any correction formulated will likely be statistical in nature.

7 SUMMARY

Existing methods for modelling emission lines from individual star-forming regions (e.g. Kewley et al. 2001) or entire galaxies (e.g. Ceverino, Klessen & Glover 2019) most often rely on relatively unconstrained parameters to describe the regions. These simplify the distribution of complex structure, variations of the ionization parameters, densities, and object ages that make up real galaxies. However, the underlying emission properties of individual clusters are set by the relative thickness of the shell surrounding the H II region, which, in turn, is determined by the balance of winds and radiation in a complex non-linear manner (Pellegrini et al. 2020). Consequently, accurate models of the emission from entire galaxies cannot simply assume values or precomputed distributions for all of these parameters. Instead, these must be derived from accurate models of the full feedback-induced dynamical evolution of each H II region. Without such a physical basis, there is an infinite number of possibly degenerate permutations.

In this paper, we have demonstrated that it is both possible and necessary to combine physically self-consistent models of individual star-forming regions with the results of cosmological simulations of star formation and to use the result to make predictions of the corresponding line emission on scales ranging from tens of parsecs to the size of the entire galaxy. Our small-scale approach include the effects of winds, radiation, and supernova feedback, as well as the influence of gravity. The evolution and emission spectrum of each source is deterministic, depending on cloud parameters, not our detailed treatment of feedback. Our use of clouds with finite masses and physical scales (as opposed to dimensionless models of SF regions) allows us to calculate the emergent radiation into the galaxy. By combining these small-scale models with the result of a cosmological simulation in a post-processing step, we are able to handle a range of length scales and physical processes, which is beyond the reach of any large-scale simulation to date, with the important caveat that a 1D geometry is assumed for the small-scale approach.

Models that parametrize H II region emission in terms of the ionization parameter U, metal abundance Z, and number density n_H (e.g. Kewley et al. 2001) have the drawback that they may include results that are non-physical, in the sense that they would not be reached during the dynamical evolution of any real star-forming region. Because we explicitly follow the dynamical evolution of the H II regions under the influence of all of the relevant feedback mechanisms, we avoid this problem. One consequence of this is that the ionization parameter is a prediction of our model rather than a tunable parameter. Instead, the key tunable parameters in our model are the mean natal cloud density, cloud mass, metallicity, and star formation efficiency (or, alternatively, the cluster mass). Since our local cloudy/cluster evolution calculations (see e.g. Rahner et al. 2017, 2019) are relatively fast to run, it is reasonable to contemplate varying these parameters and finding which values best-fitting observations of real galaxies, but this is a topic that lies beyond the scope of this introductory paper.

We have demonstrated that our WARPFIELD-POP approach is able to reproduce with reasonable fidelity a number of features of the real Milky Way, such as the inferred radial and vertical distributions of the electron number density and electron temperature, or the measured angular power spectrum of Hα emission. We have also shown how our model can be used to make predictions for observational tracers

that have not yet been observed over large areas, such as [S III]. We find that the origin of the observable flux in the tracers that we examine in this introductory study (Hα, [S III]) arises from a complex distribution of cluster masses and ages in different galactic environments. Although Hα is a widely used tracer of star formation in obscured environments, our work suggests that [S III] is even more promising: The higher contrast of star-forming regions seen in [S III] compared to Hα will allow them to be traced to distances up to a factor of 3 larger in the Milky Way (see e.g. Fig. 14), allowing us to study star-forming regions in a much wider volume of the disc, possibly even reaching as far as the Galactic Centre. This prediction will soon be put to the test by the panchromatic survey SDSS-V LVM (Kollmeier et al. 2017), which will spectroscopically survey a range of different optical lines (including Hα and [S III]) over almost 3000 deg^2 of the Galactic disc.

Finally, we have also demonstrated how our models can be used to aid in the interpretation of observations of other galaxies. As an example, we show how they can be used to quantify the accuracy with which one can reconstruct the emitted Hα flux from the observed flux using the usual Balmer decrement technique. We argue that the standard technique works reasonably well for galaxies with low inclinations, although it always estimates the true flux, particular in the centres of galaxies. However, we also show that it fails badly for highly inclined galaxies, underestimating the true flux at small deprojected radii by up to an order of magnitude.

ACKNOWLEDGEMENTS

DATA AVAILABILITY

The data underlying this paper, including precalculated predictions of individual WARPFIELD-EMP clouds used here as well as the entire predictions for the synthesized galaxy, will be shared on reasonable request to the corresponding author.

The codes WARPFIELD and WARPFIELD-EMP underlying this paper will be shared on reasonable request to the corresponding author, or via request at GitHub at https://github.com/EricPell/WARPFIELD and https://github.com/EricPell/WARPFIELD-EMP, respectively.

REFERENCES

Alves M. I. R., Boulanger F., Ferrière K., Montier L., 2018, A&A, 611, L5
Bacon R. et al., 2010, in McLean I. S., Ramsay S. K., Takami H., eds, Proc. SPIE Conf. Ser. Vol. 7735, Ground-based and Airborne Instrumentation for Astronomy III. SPIE, Bellingham, p. 773508
Bigiel F., Leroy A., Walter F., Brinks E., de Blok W. J. G., Madore B., Thornley M. D., 2008, AJ, 136, 2846
Binder B. A., Povich M. S., 2018, ApJ, 864, 136
Boulanger F., Cox P., Jones A. P., 2000, in Casoli F., Lequeux J., David F., eds, Infrared Space Astronomy, Today and Tomorrow, Vol. 70. Les Houches Summer School, p. 251
Brauer R., Wolf S., Reissl S., Ober F., 2017a, A&A, 601, A90
Brauer R., Wolf S., Flock M., 2017b, A&A, 607, A104
Calzetti D., 2001, PASP, 113, 1449
Ceverino D., Klessen R. S., Glover S. C. O., 2019, MNRAS, 484, 1366
Chevance M. et al., 2020, MNRAS, 493, 2872
Cordes J. M., Lazio T. J. W., 2002, preprint (arXiv:astro-ph/0207156)
Diehl R. et al., 2006, A&A, 449, 1025
Dopita M. A., Kewley L. J., Sutherland R. S., Nicholls D. C., 2016, Ap&SS, 361, 61
Draine B. T., Li A., 2001, ApJ, 551, 807
Ferland G. J. et al., 2017, Rev. Mex. Astron. Astrofis., 53, 385
Finkbeiner D. P., 2003, ApJS, 146, 407
Fuchs B., Breitschwerdt D., de Avillez M. A., Dettbarn C., 2009, Space Sci. Rev., 143, 437
Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, ApJ, 622, 759
Gouliermis D. A., 2018, PASP, 130, 072001
Grand R. J. J. et al., 2017, MNRAS, 467, 179
Grand R. J. J. et al., 2018, MNRAS, 474, 3629
Grandmont F., Drissen L., Mandar J., Thibault S., Baril M., 2012, in McLean I. S., Ramsay S. K., Takami H., eds, Proc. SPIE Conf. Ser. Vol. 8446, Ground-based and Airborne Instrumentation for Astronomy IV. SPIE, Bellingham, p. 84460U
Haslam C. G. T., Klein U., Salter C. J., Stoffel H., Wilson W. E., Cleary M. N., Cooke D. J., Thomasson P., 1981, A&A, 100, 209
Hopkins P. F. et al., 2018, MNRAS, 480, 800
Indriolo N., Geballe T. R., Oka T., McCall B. J., 2007, ApJ, 671, 1736
Kennicutt R. C. Jr, 1998, in Gilmore G., Howell D., eds, ASP Conf. Ser. Vol. 142, The Stellar Initial Mass Function. Astron. Soc. Pac., San Francisco, p. 1
Kewley L. J., Dopita M. A., Sutherland R. S., Heisler C. A., Trevena J., 2001, ApJ, 556, 121
Klessen R. S., Glover S. C. O., 2016, Star Formation in Galaxy Evolution: Connecting Numerical Models to Reality. Springer-Verlag, Berlin
Kollmeier J. A. et al., 2017, preprint (arXiv:1711.03234)
Kreckel K. et al., 2018, ApJ, 863, L21
Krumholz M. R., McKee C. F., Bland-Hawthorn J., 2019, ARA&A, 57, 227
Lada C. J., Lada E. A., 2003, ARA&A, 41, 57
Li H., Vogelsberger M., Marinacci F., Gnedin O. Y., 2019, MNRAS, 487, 364
Liu W. et al., 2017, ApJ, 834, 33
McAlpine S. et al., 2016, Astron. Comput., 15, 72
Marinacci F., Pakmor R., Springel V., 2014, MNRAS, 437, 1750
Marinacci F., Grand R. J. J., Pakmor R., Springel V., Gómez F. A., Frenk C. S., White S. D. M., 2017, MNRAS, 466, 3859
Mathis J. S., Rumpl W., Nordsieck K. H., 1977, ApJ, 217, 425
Medling A. M. et al., 2018, MNRAS, 475, 5194
Misiriotis A., Xilouris E. M., Papamastorakis J., Boumis P., Goudis C. D., 2006, A&A, 459, 113
Monachesi A., Gómez F. A., Grand R. J. J., Kauffmann G., Marinacci F., Pakmor R., Springel V., Frenk C. S., 2016, MNRAS, 459, L46
Murray N., 2009, ApJ, 691, 946 (MR2009)
Oey M. S. et al., 2017a, ApJ, 844, 63
Oey M. S., Herrera C. N., Silich S., Reiter M., James B. L., Jaskot A. E., Micheva G., 2017b, ApJ, 849, L1
Oppermann N. et al., 2012, A&A, 542, A93

Osterbrock D. E., Ferland G. J., 2006, *Astrophysics of Gaseous Nebulae and Active Galactic Nuclei*. University Science Books, Sausalito, CA

Pakmor R., Springel V., 2013, MNRAS, 432, 176

Pakmor R., Marinacci F., Springel V., 2014, ApJ, 783, L20

Pakmor R. et al., 2017, MNRAS, 469, 3185

Pellegrini E. W., Baldwin J. A., Ferland G. J., Shaw G., Heathcote S., 2009, ApJ, 693, 285

Pellegrini E. W., Baldwin J. A., Ferland G. J., 2011, ApJ, 738, 34

Pellegrini E. W., Oey M. S., Winkler P. F., Points S. D., Smith R. C., Jaskot A. E., Zastrow J., 2012, ApJ, 755, 40

Pellegrini E. W., Rahner D., Reissl S., Glover S. C. O., Klessen R. S., Rousseau-Nepton L., Herrera-Camus R., 2020, MNRAS, 496, 339

Planck Collaboration XVI, 2014, A&A, 571, A16

Portegies Zwart S. F., McMillan S. L. W., Gieles M., 2010, ARA&A, 48, 431

Rahner D., Pellegrini E. W., Glover S. C. O., Klessen R. S., 2017, MNRAS, 470, 4453

Rahner D., Pellegrini E. W., Glover S. C. O., Klessen R. S., 2018a, Astrophysics Source Code Library, record ascl:1807.002

Rahner D., Pellegrini E. W., Glover S. C. O., Klessen R. S., 2018b, MNRAS, 473, L11

Rahner D., Pellegrini E. W., Glover S. C. O., Klessen R. S., 2019, MNRAS, 483, 2547

Reissl S., Wolf S., Brauer R., 2016, A&A, 593, A87

Reissl S., Seifried D., Wolf S., Banerjee R., Klessen R. S., 2017, A&A, 603, A71

Reissl S., Stutz A. M., Brauer R., Pellegrini E. W., Schleicher D. R. G., Klessen R. S., 2018, MNRAS, 481, 2507

Reissl S., Brauer R., Klessen R. S., Pellegrini E. W. P., 2019a, ApJ, submitted

Reissl S., Brauer R., Klessen R. S., Pellegrini E. W., 2019b, ApJ, 885, 15

Richardson C. T., Polimera M. S., Kannappan S. J., Moffett A. J., Bittner A. S., 2019, MNRAS, 486, 3541

Robitaille T. P., Whitney B. A., 2010, Highlights Astron., 15, 799 (RW2010)

Rousseau-Nepton L., Robert C., Martin R. P., Drissen L., Martin T., 2018a, MNRAS, 477, 4152

Rousseau-Nepton L., Robert C., Martin R. P., Drissen L., Martin T., 2018b, MNRAS, 477, 4152

Rugel M. R. et al., 2019, A&A, 622, A48

Seifried D., Walch S., Reissl S., Ibáñez-Mejía J. C., 2019, MNRAS, 482, 2697

Simpson C. M., Grand R. J. J., Gómez F. A., Marinacci F., Pakmor R., Springel V., Campbell D. J. R., Frenk C. S., 2018, MNRAS, 478, 548

Smith L. F., Biermann P., Mezger P. G., 1978, A&A, 66, 65

Springel V., 2010, MNRAS, 401, 791

Springel V., Hernquist L., 2003, MNRAS, 339, 289

Springel V. et al., 2018, MNRAS, 475, 676

Sun X. H., Reich W., Waelkens A., Enßlin T. A., 2008, A&A, 477, 573

Tomičić N. et al., 2019, ApJ, 873, 3

Turner J. L., Consiglio S. M., Beck S. C., Goss W. M., Ho P. T. P., Meier D. S., Silich S., Zhao J.-H., 2017, ApJ, 846, 73

Vogelsberger M., Genel S., Sijacki D., Torrey P., Springel V., Hernquist L., 2013, MNRAS, 436, 3031

Vogelsberger M. et al., 2014, MNRAS, 444, 1518

Whittet D. C. B., 1992, J. Br. Astron. Assoc., 102, 175

Wolfire M. G., Hollenbach D., McKee C. F., Tielens A. G. G. M., Bakes E. L. O., 1995, ApJ, 443, 152

Wolfire M. G., McKee C. F., Hollenbach D., Tielens A. G. G. M., 2003, ApJ, 587, 278

Yao J. M., Manchester R. N., Wang N., 2017, ApJ, 835, 29

Zhang Q., Fall S. M., 1999, ApJ, 527, L81

APPENDIX A: H II REGION OVERLAP AND IONIZING FLUX

One of the assumptions made in our current implementation of WARPFIELD-POP is that each point in the galaxy can be treated as if it were ionized by a single bright cluster plus the diffuse radiation field of the galaxy. As explained in Section 3.3, for each point in
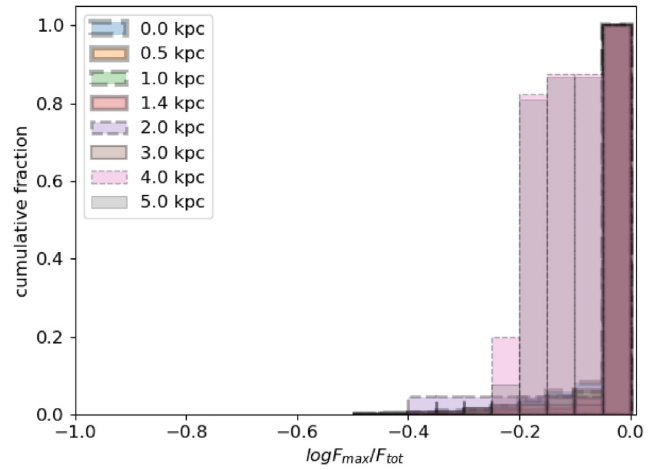
Figure A1. Cumulative histogram showing the fraction of sample points for which the ratio of the ionizing flux from the brightest cluster, F_{max}, to the total ionizing flux, F_{tot}, is equal to or less than the specified value. We consider several sets of sample points located at the listed distances above the mid-plane, and restrict our analysis to points lying within at least one ionization front, i.e. this can be interpreted as a measurement of the extent to which the regions ionized by different clusters overlap. The ionizing luminosity per cluster used here is the emergent flux, after attenuation by the WARPFIELD-EMP model.

space that is not directly represented by a WARPFIELD model, we find which cluster would produce the most emissivity at that point (i.e. which would supply the most ionizing flux), if we were to treat the effects of the clusters individually. Making this approximation allows us to avoid the very expensive iterative calculation that would be necessary in order to model the impact of all of the clusters simultaneously, but obviously introduces some inaccuracy into our solution. In this section, we attempt to quantify this inaccuracy. Our analysis here is based on the locations of the H^+/H^0 ionization fronts around the different clusters. This should be adequate for most tracers of interest, but possibly not for chemical elements with ionization potentials below than of H, such as C or Si.

To assess whether our assumption that the ionizing flux is typically dominated by a single cluster is a good one, we proceed as follows. We start by randomly selecting a set of 10^4 sample points in the mid-plane of the galaxy. We determine how many of these sample points are located within the ionization front of at least one cluster. For this subset, we then calculate for each point the ionizing flux from the cluster that makes the largest contribution at that point (F_{max}) and the total ionizing flux at the same point (F_{tot}). We next repeat the same procedure for similar distributions of sample points located at increasing distances above the mid-plane. In Fig. A1, we summarize the results of this experiment by showing a cumulative histogram of the fraction of sample points that have F_{max}/F_{tot} below the specified values. Because our sample points are randomly chosen, this can be interpreted as the fraction of the ionized disc that has F_{max}/F_{tot} below the specified value.

We see from Fig. A1 that in the galactic mid-plane, most of the volume that falls within an ionization front has $F_{max}/F_{tot} \sim 1$, i.e. in the vast majority of cases, our assumption that only a single cluster dominates is a good one. There is a small tail of regions with lower F_{max}/F_{tot}, but a single cluster contributes more than half the total ionizing flux for all but a small fraction (~ 5 per cent) of points. We also see that our assumption remains a good one up to a height

of ~4 kpc, but that it breaks down above this, with our method underestimating the radiation field by a factor of 2 for half the cells illuminated by star-forming regions at this scale height. However, we note that in any case, we do not expect our method to do a good job of modelling the ionization of gas far from the mid-plane, as in this regime, collisional ionization likely plays a significant role.

APPENDIX B: MULTIPOLE FITTING

We quantify the structure of our all-sky emission maps by computing their angular power spectrum. Here, we briefly outline how we go about this. The pattern projected on the sky can be written as a series of spherical harmonics:

$$S(\vartheta, \varphi) \simeq \sum_{l=0}^{N} \sum_{m=-l}^{l} a_{l,m} Y_{l,m}(\vartheta, \varphi). \tag{B1}$$

Here, S stands for any signal we presented in this work so far, $Y_{l,m}(\vartheta, \varphi)$ is the spherical harmonic, and $a_{l,m}$ is the fit coefficient. Mathematically, equation (B1) is exact only when we allow the sum over l to go to infinity, but in any actual fitting procedure, the computation has to stop at a distinct value N. All-sky signals can then be quantified in terms of the fit coefficients $a_{l,m}$ for each multipole l. The resulting spectrum is usually quantified by the single parameter function ,

$$f(l) = \frac{l(l+1)C_l}{2\pi}, \tag{B2}$$

where $C_l = \mathrm{Var}\left(\left|a_{l,m}\right|\right)$ is the variance of the magnitude of the complex fit parameter $a_{l,m}$ over all possible values of m. We perform this kind of analysis with the implementation provided by the PYTHON package HEALPY.[16]

APPENDIX C: EXTRAGALACTIC EMISSION MAPS

For completeness, we include the extragalactic projections at different inclinations angles used to calculate the profiles in Fig. 18 as Figs C1–C6. Apart from the differing inclinations, the details of these figures are the same as for Fig. 17, including both the physical and the intensity scale.
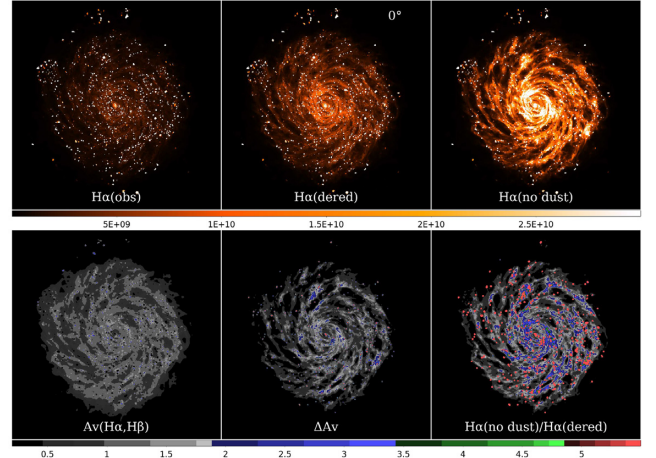
[16]http://healpix.jpl.nasa.gov.

Figure C1. Same as Fig. 17 but at 0°.

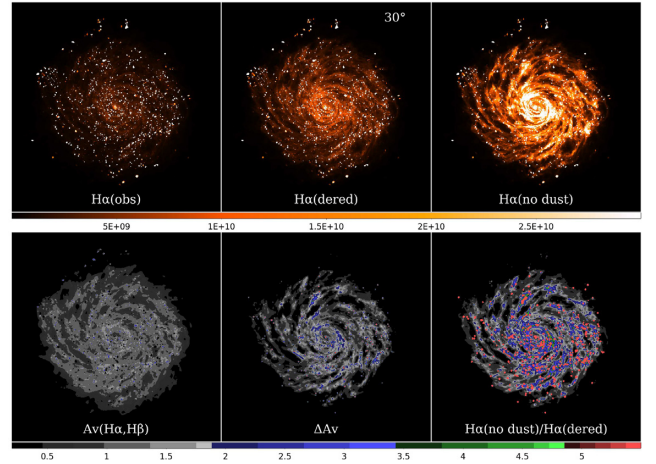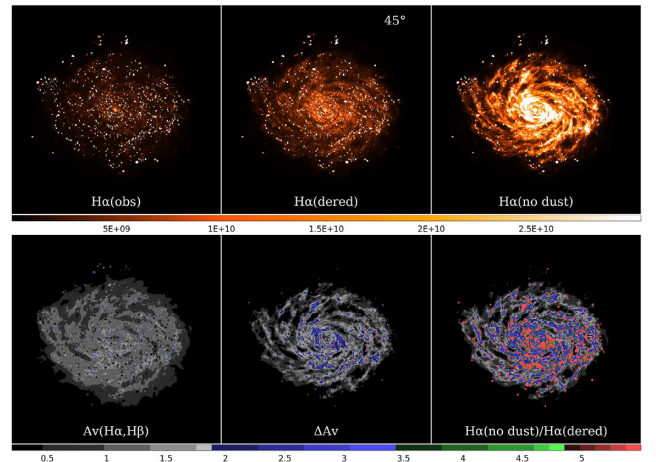**Figure C2.** Same as Fig. 17 but at 30°.

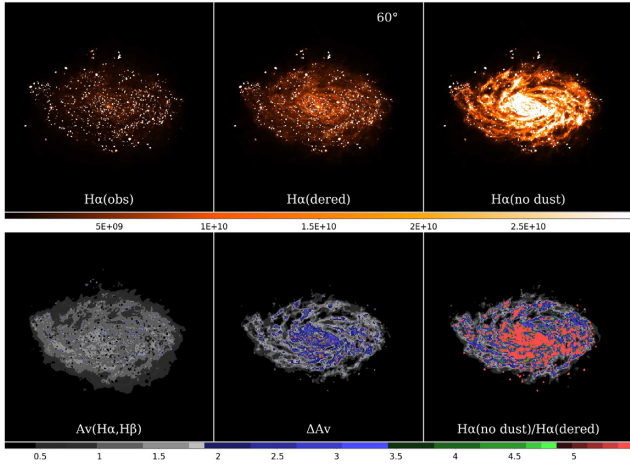**Figure C3.** Same as Fig. 17 but at 45°.

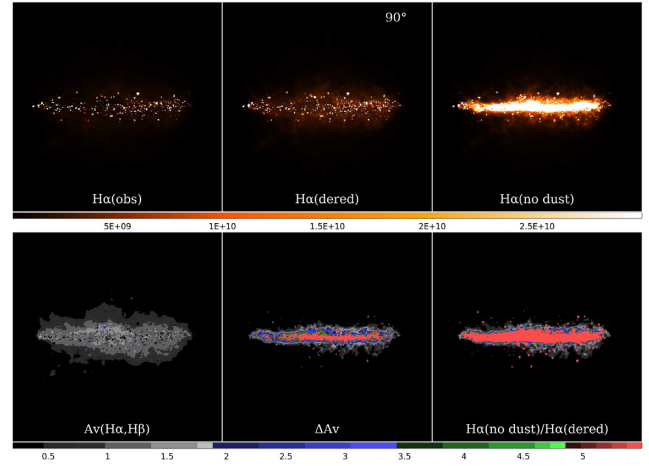**Figure C4.** Same as Fig. 17 but at 60°.
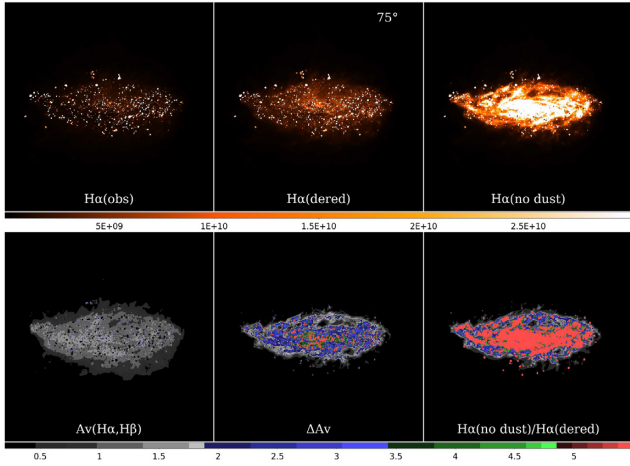

Figure C6. Same as Fig. 17 but at 90°.

Figure C5. Same as Fig. 17 but at 75°.

This paper has been typeset from a TEX/LATEX file prepared by the author.