

Reproducing the CO-to-H₂ conversion factor in cosmological simulations of Milky-Way-mass galaxies

Laura C. Keating ¹  ¹★ Alexander J. Richings, ² Norman Murray, ^{1,3} Claude-André Faucher-Giguère ^{4,5}  Philip F. Hopkins ⁶  Andrew Wetzel ⁷  Dušan Kereš ⁸ Samantha Benincasa ⁹  Robert Feldmann ⁹  Sarah Loebman ⁷  and Matthew E. Orr ¹ 

¹ Canadian Institute for Theoretical Astrophysics, 60 St George Street, University of Toronto, ON M5S 3H8, Canada

² Department of Physics, Institute for Computational Cosmology, Durham University, South Road, Durham DH1 3LE, UK

³ Canadian Research Chair in Theoretical Astrophysics, Queen's University Kingston, ON K7L 3N6, Canada

⁴ Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA), Northwestern University, CIERA, Evanston, IL 60201, USA

⁵ Department of Physics and Astronomy, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208, USA

⁶ TAPIR, MC 350-17, California Institute of Technology, Pasadena, CA 91125, USA

⁷ Department of Physics, University of California, Davis, CA 95616, USA

⁸ Department of Physics, Center for Astrophysics and Space Sciences, University of California at San Diego, La Jolla, CA 92093, USA

⁹ Institute for Computational Science, University of Zurich, Zurich CH-8057, Switzerland

Accepted 2020 September 3. Received 2020 August 19; in original form 2020 January 22

ABSTRACT

We present models of CO(1–0) emission from Milky-Way-mass galaxies at redshift zero in the FIRE-2 cosmological zoom-in simulations. We calculate the molecular abundances by post-processing the simulations with an equilibrium chemistry solver while accounting for the effects of local sources, and determine the emergent CO(1–0) emission using a line radiative transfer code. We find that the results depend strongly on the shielding length assumed, which, in our models, sets the attenuation of the incident UV radiation field. At the resolution of these simulations, commonly used choices for the shielding length, such as the Jeans length, result in CO abundances that are too high at a given H₂ abundance. We find that a model with a distribution of shielding lengths, which has a median shielding length of ~ 3 pc in cold gas ($T < 300$ K) for both CO and H₂, is able to reproduce both the observed CO(1–0) luminosity and inferred CO-to-H₂ conversion factor at a given star formation rate compared with observations. We suggest that this short shielding length can be thought of as a subgrid model, which controls the amount of radiation that penetrates giant molecular clouds.

Key words: methods: numerical – ISM: molecules – galaxies: evolution – galaxies: ISM.

1 INTRODUCTION

Understanding the molecular gas content of galaxies is crucial for understanding the process of star formation (e.g. McKee & Ostriker 2007; Kennicutt & Evans 2012; Krumholz 2014b). Molecular gas tends to form in high-density regions, where self-shielding and shielding by dust from the interstellar radiation field is most effective. Much of this molecular gas is in the form of giant molecular clouds (GMCs), which typically have average densities larger than 10^2 cm⁻³ and temperatures in the range 10–20 K (Ferrière 2001). The most abundant molecule is molecular hydrogen (H₂). However, it is challenging to observe the H₂ in emission, as it requires high temperatures ($T \sim 500$ K) to be excited and therefore cannot be seen in emission at the low temperatures of GMCs. Instead, carbon monoxide (CO) can be used as a convenient tracer of this molecular gas, as it has its first rotational transition at 5.5 K (Carilli & Walter 2013) – low enough to probe the cold interstellar medium (ISM).

However, it has long been known that CO is not a perfect tracer of H₂, as CO is more easily dissociated than H₂ (van Dishoeck & Black 1988; Wolfire, Hollenbach & McKee 2010). This results in regions of ‘dark gas’ in a phase that is lower in density and warmer than the CO-bright gas (see, e.g. the model of Seifried et al. 2019). Dark gas can be probed observationally by γ -rays produced from interactions between cosmic rays and hydrogen (Grenier, Casandjian & Terrier 2005; Remy et al. 2017), with maps of thermal dust emission (Planck Collaboration XIX 2011) and from emission from singly ionized carbon atoms (Pineda et al. 2013; Langer et al. 2014).

The CO(1–0) emission is usually related to the total H₂ mass by a CO-to-H₂ conversion factor X_{CO} , defined as

$$X_{\text{CO}} = \frac{N_{\text{H}_2}}{W_{10}}, \quad (1)$$

where N_{H_2} is the H₂ column density and W_{10} is the velocity-integrated CO(1–0) brightness temperature, which is related to the CO(1–0) intensity. In the disc of the Milky Way, this conversion factor is estimated to be $X_{\text{CO}} \approx 2 \times 10^{20}$ cm⁻² K⁻¹ km⁻¹ s using a range of different techniques (Bolatto, Wolfire & Leroy 2013). This conversion factor is however also thought to depend

* E-mail: lkeating@cita.utoronto.ca

† Hubble Fellow.

on environmental factors, such as the metallicity (Israel 1997; Leroy et al. 2011; Sandstrom et al. 2013). The most appropriate choice for a given galaxy, or whether a single value for an individual galaxy is valid, is a subject of much observational (Solomon et al. 1987; Downes & Solomon 1998; Accurso et al. 2017) and theoretical discussion (Wolfire, Hollenbach & Tielens 1993; Narayanan et al. 2011, 2012; Shetty et al. 2011a,b; Feldmann, Gnedin & Kravtsov 2012a, b; Gong, Ostriker & Kim 2018; Li et al. 2018; Richings & Faucher-Giguère 2018).

One way of understanding how CO emission relates to properties of the ISM and galaxy is through simulations. However, modelling the emission of CO from galaxies is extremely challenging, due to the wide range of spatial scales involved. CO has been observed in the Milky Way tens of kpc from the Galactic Centre (Heyer, Carpenter & Snell 2001), but the molecular gas resides in GMCs with sizes of a few to ~ 100 pc (Miville-Deschénes, Murray & Lee 2017). These clouds further contain clumps and cores that have sizes less than 0.1 pc (Bergin & Tafalla 2007). Some works have therefore focused on modelling individual GMCs (e.g. Glover & Mac Low 2011; Shetty et al. 2011a,b; Clark & Glover 2015; Peñaloza et al. 2018). This allows for high spatial resolution in the regions hosting the molecular gas, but this comes at the cost of having to neglect the impact of the larger scale galactic environment. Another approach is to model the emission coming from entire or parts of isolated galaxies (e.g. Duarte-Cabral et al. 2015; Glover & Smith 2016; Richings & Schaye 2016; Gong et al. 2018). This still neglects the fact that galaxy formation occurs in a cosmological context. At the largest scales, semi-analytic schemes can be used to quickly compute CO luminosities for large numbers of galaxies (Obreschkow et al. 2009; Lagos et al. 2012; Popping et al. 2019), which provide large samples of CO-emitting galaxies but cannot be used to study small-scale environmental effects. An intermediate approach is to post-process cosmological hydrodynamic simulations by making some assumptions about the internal structure of GMCs (Olsen et al. 2016; Li et al. 2018; Vallini et al. 2018). Ideally, one would calculate the molecular gas abundances on the fly. However, such simulations are extremely costly due to the increased number of species that need to be tracked and rate equations that must be integrated. These simulations have so far been limited to high-redshift studies (Katz et al. 2017; Lupi et al. 2019; Pallottini et al. 2019), isolated galaxies (Kannan et al. 2019), or dwarf galaxies (Lupi & Bovino 2019).

In this paper, we continue the attempt to bridge the gap between simulations of individual GMCs or isolated galaxies, and fully cosmological simulations. We use cosmological zoom-in simulations from the FIRE project¹ (Feedback in Realistic Environments; Hopkins et al. 2014, 2018). These simulations have high mass resolution and an explicit treatment of the multiphase ISM, and can resolve the most massive GMCs ($M \sim 10^5 M_\odot$; Benincasa et al. 2019) where the bulk of star formation occurs (Williams & McKee 1997). We post-process the redshift zero simulation outputs with an equilibrium chemistry solver and a line radiative transfer code to calculate the CO and H₂ abundances, and the CO(1–0) emission. The structure of this paper is as follows. In Section 2, we describe how we model the CO(1–0) line emission in our simulations. In Section 3, we compare our models with observations of local galaxies and describe the effects of varying some of the assumptions in our chemical modelling. In Section 4, we discuss the interpretation of our results. We present our conclusions in Section 5.

¹<https://fire.northwestern.edu/>

2 MODELLING CO(1–0) EMISSION

2.1 Cosmological zoom simulations of Milky-Way-type galaxies

In this work, we make use of the FIRE-2 simulations first presented in Hopkins et al. (2018). The FIRE-2 simulations are a set of cosmological zoom-in simulations that explicitly model a multiphase ISM. The simulations were run with GIZMO, using the meshless-finite mass hydrodynamic solver (Hopkins 2015). GIZMO uses a gravity solver and domain decomposition algorithm descended from the Tree-PM solver in GADGET-3 (last described in Springel 2005) with some modifications, including adaptive softening for gas.

The physics included in the FIRE-2 simulations is described in detail in Hopkins et al. (2018) and only a brief summary is provided here. Star formation occurs using a probabilistic criteria for gas that is self-gravitating, Jeans-unstable, molecular, and has density $n > 1000 \text{ cm}^{-3}$. During the course of the simulation, the molecular fraction of the gas particles is determined using the relation presented in Krumholz & Gnedin (2011). Each star particle is treated as a single stellar population assuming a Kroupa (2001) initial mass function. The stellar feedback of each star particle is estimated from tables computed using the stellar population models of STARBURST99 (Leitherer et al. 1999). The stellar feedback mechanisms include Type Ia and II supernovae, stellar winds from evolved stars, photoionization/heating, and radiation pressure. Photoionization and photoheating from the extragalactic UV background is also included (Faucher-Giguère et al. 2009), as well as photoelectric heating and heating by cosmic rays. Heating and cooling rates are estimated from pre-computed CLOUDY tables (Ferland et al. 2013) that include metal-line and molecular cooling, allowing the gas to cool down to 10 K. Eleven elements are explicitly tracked (H, He, C, N, O, Ne, Mg, Si, S, Ca, and Fe). Subgrid metal mixing is included using the prescription described in Hopkins (2017) and Escala et al. (2018).

We focus most of our analysis here on the m12i galaxy, presented as part of the Latte suite of FIRE-2 simulations in Wetzel et al. (2016). This is a simulation of a Milky-Way-mass galaxy, with a virial mass of $1.1 \times 10^{12} M_\odot$. Maps of the projected density, mass-weighted temperature, mass-weighted metallicity, and mass-weighted UV flux we measure in m12i are shown in Fig. 1. The halo was selected for its final mass from a dark matter only simulation of a 85.5-Mpc volume, with initial conditions generated using MUSIC (Hahn & Abel 2011) as part of the AGORA project (Kim et al. 2014). A Lambda cold dark matter cosmology was assumed, with $\Omega_\Lambda = 0.728$, $\Omega_m = 0.272$, $\Omega_b = 0.0455$, $h = 0.702$, $\sigma_8 = 0.807$, and $n_s = 0.961$. The gas elements have mass $7100 M_\odot$ and the force resolution reaches pc scales in gas resolution elements with densities a few times 10^3 cm^{-3} . Using the FIRE-2 physics model described above, this simulation has been shown to agree with many observed properties of our Galaxy, such as the properties of satellite galaxies (Wetzel et al. 2016), the stellar mass–halo mass relation (Hopkins et al. 2018), the mass–metallicity relation (Ma et al. 2017) and the Kennicutt–Schmidt relation (Orr et al. 2018), and properties of the stellar thin and thick disc (Sanderson et al. 2018), among others. Most relevant for this work is that the simulation also reproduces many observed properties of GMCs (Benincasa et al. 2019; Guszejnov et al. 2020; Lakhani et al., in preparation)

2.2 Chemical equilibrium modelling

As we wish to compute the H₂ and CO abundances, we post-process the simulation snapshots with the chemistry solver CHIMES (Richings, Schaye & Oppenheimer 2014a,b). CHIMES is capable of

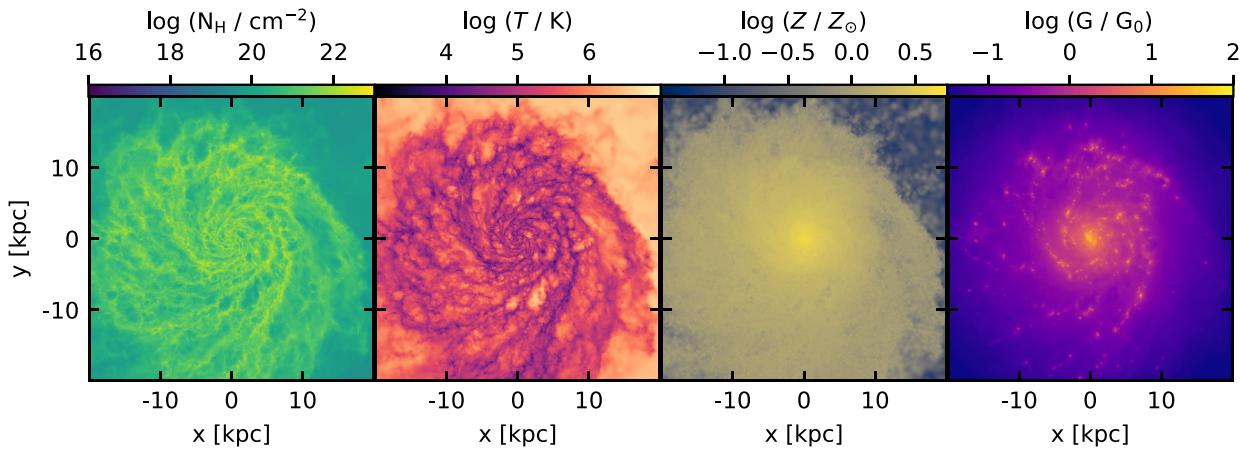


Figure 1. First row, from the left- to right-hand side: maps of the total hydrogen density, the mass-weighted gas temperature, the mass-weighted metallicity in solar units, and mass-weighted UV flux in Habing units for the m12i simulation.

computing non-equilibrium abundances, but in this work, we use it in equilibrium mode (i.e. we evolve the network until chemical equilibrium is reached). One benefit of using CHIMES is that the same library can be used to calculate equilibrium chemical abundances by post-processing simulation outputs, as well as calculating the full non-equilibrium abundances on the fly in simulations. This allows for a direct comparison between equilibrium and non-equilibrium methods (e.g. Richings & Schaye 2016). However, calculating the abundances on the fly adds a substantial computational cost and so we only present results assuming equilibrium here. CHIMES calculates the chemical abundances of 157 species. These include the ionization states of the 11 elements tracked in the simulation, as well as 20 molecular species. The CHIMES chemical network includes collisional reactions, photochemical reactions, ionization by cosmic rays, and dust grain reactions. We do not evolve the temperature of the gas particles, since the FIRE-2 simulations already account for molecular cooling. This potentially introduces an inconsistency into our analysis, as the metal-line and molecular cooling tables used in the hydrodynamic simulation assume that the gas is optically thin but in our chemical modelling the gas is allowed to self-shield. However, we confirmed that also allowing the temperature to evolve to equilibrium along with the chemical abundances did not change our results significantly, with the mass of CO differing by a factor of 2 between the two cases. As we will discuss later, this is a much smaller change than we see when changing other parameters in our modelling (in particular the shielding length), which can vary the results by orders of magnitude.

One assumption that goes into our chemical modelling is the choice of UV background, which regulates the photoionization and photodissociation rates in the gas. We compute the UV background due to local sources (which dominates over the extragalactic UV background in the disc of the galaxy) by restarting each snapshot of the underlying hydrodynamic simulation and printing out the far-UV fluxes. These are computed using the LEBRON radiative transfer approximation included in the FIRE-2 simulations (Hopkins et al. 2018, 2019). The emissivity of the sources is calculated by computing the luminosity of each star particle based on its age, metallicity, and mass using the STARBURST99 stellar population models. The far-UV fluxes are defined as the 6–13.6 eV band, so this is not the ionizing radiation field. However, it gives us an estimate of the inhomogeneous UV background we are interested in. Attenuation in the vicinity of the source by dust is accounted for in the simulation, and the radiation

is then transported under the assumption that the gas is optically thin. When the photoionization rates are calculated in the simulation, the radiation is attenuated again for each gas particle to account for local self-shielding. We print out the fluxes before this step, as this is already accounted for in CHIMES (see below).

Using these UV fluxes, we can now account for spatial variations in the amplitude of the UV background due to local sources. We do not however account for spatial fluctuations in the shape of the background, and use photoionization cross-sections calculated using the shape of the Black (1987) interstellar radiation field everywhere. This means we are not accounting for the effects of spectral hardening, and changes in the shape of the spectrum due to different stellar populations.

We assume that the cosmic ray ionization rate scales linearly with the spatially varying far-UV flux, and use a fiducial cosmic ray ionization rate $\zeta_{\text{HI}} = 1.8 \times 10^{-16} \text{ s}^{-1}$ (Indriolo & McCall 2012) at the point where the far-UV flux is equal to the average Milky Way value. We also explore the effect of raising/lowering the cosmic ray ionization rate, as well as a model where the far-UV flux and cosmic ray ionization rate are held constant (described in Section 3.5).

In CHIMES, each species i is shielded from the UV background by a shielding factor calculated based on the local column density. This shielding regulates the photoionization and photodissociation rate of that species. However, in the simulations, only a volume density for each particle is tracked. We therefore need to relate the volume density to a column density. We use the definition

$$N_i = n_i L_{\text{shield}}, \quad (2)$$

where N_i is the column density of species i , n_i is the number density of the particle, and L_{shield} is the shielding length. We assume that L_{shield} is constant for all species, but note that in reality, different species may be associated with different shielding lengths, depending on how they are distributed throughout the galaxy (e.g. the molecular gas may be clumpier than the dust). We will discuss how L_{shield} is chosen in more detail in the following section. Once this column density is known, the shielding factor for each species is calculated and used to attenuate the optically thin rates. This provides us with the rates that are appropriate for optically thick gas. Shielding by dust is accounted for each species i using the shielding factor,

$$S_{\text{dust}}^i = \exp(-\gamma_{\text{dust}}^i A_v), \quad (3)$$

where γ_{dust}^i is a constant taken from van Dishoeck, Jonkheid & van Hemert (2006) and Glover et al. (2010). A_v is the extinction, defined as $A_v = 4.0 \times 10^{-22} (N_{\text{H}}/\text{cm}^{-2}) (Z/Z_{\odot}) \text{ mag cm}^{-2}$. The species we are most interested in here are H₂ and CO. H₂ can additionally self-shield, and this is accounted for in a temperature-dependent relation outlined in detail in Richings et al. (2014b). Additional Doppler broadening with line width 7.1 km s⁻¹ due to subgrid turbulence at the scale of the resolution element is assumed (Krumholz 2012) when calculating the H₂ shielding factor. For CO, self-shielding is taken into account and there is also a contribution from shielding due to H₂. These shielding factors are taken from Visser, van Dishoeck & Black (2009).

2.3 Line radiative transfer

Once we have the CO abundances, the final step is to calculate the observed CO luminosity. We focus here on the $J = 1-0$ transition at 2.6 mm. We compute the luminosity by further post-processing the simulations using the non-LTE line radiative transfer code RADMC-3D (Dullemond et al. 2012). We perform the radiative transfer on an AMR grid, refining when more than two particles are present in a cell. We use a cubic spline kernel to interpolate the resolution elements on to the grid. We take a 40 kpc cube, centred on the halo, and perform the radiative transfer after rotating the galaxy to be face-on. We perform the radiative transfer twice for each model: once with the line emission plus the dust continuum, and again with the dust continuum alone. Subtracting these provides us with the CO(1–0) emission only. Our final maps have 2048² cells in the spatial plane, with 19.5 pc resolution, and 1 km s⁻¹ resolution along the frequency axis.

To calculate the level populations, we use the large velocity gradient method mode in RADMC-3D (Castor 1970; Goldreich & Kwan 1974; Shetty et al. 2011a). This calculates an optical depth for each cell based on the gas velocities in neighbouring cells, assuming that a photon will eventually be Doppler shifted away from line centre and hence yielding an escape probability for each photon. Gas can also be excited based on its temperature. We use Einstein and collisional rate coefficients taken from the LAMDA data base (Schöier et al. 2005) and also account for the effects of the cosmic microwave background.

As dust is not explicitly modelled as a separate species in these simulations, we assume that a fraction of the metals in gas with temperature less than 10⁵ K are in dust and use fixed dust-to-metal gas mass ratios for the silicate and graphite species. We note that we have not implemented a correction for dust depletion into our modelling, which means that we are double-counting metals that would otherwise not be in the gas phase. Estimates of dust depletion factors (Jenkins 2009; De Cia et al. 2016) suggest that carbon and oxygen can be depleted by ∼40–60 per cent, relevant for the CO abundances we are interested in here. This could potentially reduce our modelled CO(1–0) luminosities by a factor of ∼2. This is not at the level that will change any of the main conclusions of this paper (which, as described later, seeks to resolve order of magnitude discrepancies in the line emission) and is something that we will investigate more carefully in future work. We account for local turbulent broadening by adding a microturbulent line width of 7.1 km s⁻¹, consistent with what is assumed in Section 2.2. We assume an external radiation field while performing the radiative transfer. The role of this background radiation is to calculate the dust temperatures and continuum emission. As we will later subtract away the continuum emission to isolate the line emission, we do not use

an inhomogeneous UV background for this step (in contrast with our chemical modelling as described in Section 2.2). We instead use a fixed uniform background with the shape and amplitude of the Black (1987) model for the interstellar radiation field.

3 CO(1–0) EMISSION AND SHIELDING LENGTH

3.1 Shielding length approximations

As discussed above, one of the fields that must be input to CHIMES is the shielding length of each particle (L_{shield}), which controls whether the optically thick or optically thin photodissociation rates should be applied. Ideally, one would use a full ray-tracing scheme; however, this is computationally expensive so an alternative is to assume some local approximation. However, it is not obvious how this approximate shielding length should be defined. This uncertainty has been noted in previous works, exploring molecular gas formation in different regimes. In a network that includes shielding by dust, H₂ and CO self-shielding, and shielding of CO by H₂, Safranek-Shrader et al. (2017) performed an analysis of seven different approximations of this shielding length. They compared the resulting H₂ and CO abundances to a full ray-tracing simulation. They found that using the Jeans length with a temperature ceiling of 40 K provided the closest solution to the ray-tracing scheme. However, Wolcott-Green, Haiman & Bryan (2011), again comparing to full 3D ray-tracing, found that the Sobolev length was a better approximation, although this was in the case of H₂ self-shielding alone in gas of primordial composition. There is therefore no consensus on the best approximation to make, and we note that this likely depends on the properties of the simulation, in particular the resolution, as well as the physics that is included. Here we also take the approach of testing different approximations of the shielding length (as well as other parameters assumed in the chemical modelling), and compare the results directly against observations.

In Fig. 2, we present the distribution of shielding lengths measured in the simulation for the four different cases we test. The first is the smoothing length of the gas resolution elements, which is a commonly used approximation for L_{shield} . A caveat of this definition is that the shielding length will be a function of the resolution of the simulation. In particular, in the limit of infinite resolution, the smoothing length will become very small and our model would approach the optically thin limit. We nevertheless include it in our analysis as it is a commonly used approach. In the FIRE-2 simulations, this length scale is also identical to the gravitational force softening of the gas and so sets the spatial resolution of the particle. It can be written as

$$L_{\text{Sm}} = 16 \text{ pc} \left(\frac{m_{\text{gas}}}{1000 M_{\odot}} \right)^{1/3} \left(\frac{n_{\text{H}}}{10 \text{ cm}^{-3}} \right)^{-1/3}, \quad (4)$$

where m_{gas} is the mass resolution and n_{H} is the hydrogen number density. The second length scale we consider is the Jeans length, which is defined as

$$L_{\text{Jeans}} = \sqrt{\frac{\pi c_s^2}{G \rho}}, \quad (5)$$

where L_{Jeans} is the Jeans length, c_s is the sound speed, G is the gravitational constant, and ρ is the density of each resolution element. This corresponds to the length scale at which a system becomes self-gravitating (e.g. Schaye & Dalla Vecchia 2008). The third and fourth cases we test are based on a Sobolev-like length scale defined using

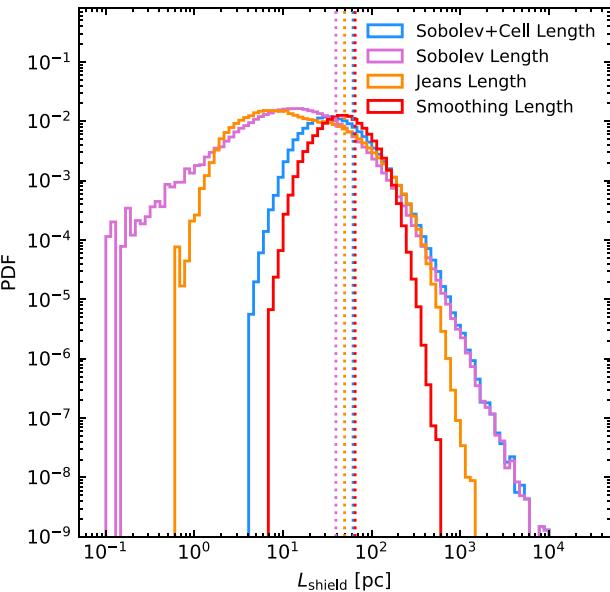


Figure 2. Histograms of the different fiducial shielding length scales we consider here: the smoothing length (red), the Jeans length (orange), the Sobolev length (blue), and the Sobolev length plus a contribution from the local cell (pink). Only values for cold gas ($T < 300$ K) in the galactic disc of m12i (which we define as gas with radius $r < 20$ kpc and vertical height $|z| < 3$ kpc) are shown. The vertical dotted lines mark the median of the distributions.

local density gradients (Sobolev 1957; Gnedin, Tassis & Kravtsov 2009). This definition of the Sobolev length is given by

$$L_{\text{Sob}} = \frac{\rho}{|\nabla \rho|}. \quad (6)$$

We also explore the effect of using the length scale that is used in GIZMO to convert densities into surface densities. These surface densities are used in the Krumholz & Gnedin (2011) fitting function to calculate H₂ fractions, as well as to attenuate the background radiation of local sources when determining the radiative feedback (Hopkins et al. 2018). This is a combination of the Sobolev length defined in equation (6), plus a contribution to the shielding length from the local cell. This is defined as

$$L_{\text{Sob+cell}} = \frac{\rho}{|\nabla \rho|} + \frac{L_{\text{sm}}}{N_{\text{ngb}}^{1/3}}, \quad (7)$$

where N_{ngb} is the number of nearest neighbours. This additional term is an approximation for the interparticle distance, which is smaller than the smoothing length of the cell. As for the local UV fluxes, we output this directly from the simulation when restarting the snapshot. We note that none of these length scales are equivalent to the velocity-gradient length scale we are using in RADMC-3D to calculate the level populations, which make a moderate difference to the results (Péñaloza et al. 2018). As demonstrated in Fig. 2, although the medians of the distributions of shielding lengths are similar, the detailed shapes can be quite different.

We use these different assumptions about the shielding lengths as input for CHIMES and calculate the resulting CO abundances for each particle, and hence the total CO(1–0) luminosity for the galaxy. Results for three of our shielding lengths are shown in Fig. 3. Looking first at the case where we use the Sobolev length plus a contribution from the local cell, we find that the results are unsurprising: The H₂ and CO clearly trace the spiral arms of the galaxy (first and second rows). The spiral arms are also home to the young stars, and

hence where the UV flux is highest (Fig. 1). This demonstrates the importance of accounting for the shielding from the UV background and local sources. In the third row of Fig. 3, we show the velocity-integrated CO(1–0) brightness temperature (W_{10}). This is calculated as

$$W = \frac{1}{2k_B} \left(\frac{c}{v} \right)^2 \int I_v dv, \quad (8)$$

where k_B is the Boltzmann constant, c is the speed of light, v is the frequency, I_v is the line intensity computed by RADMC-3D, and v is the velocity. It is clear that the emission is dominated by the high-column-density regions, and also that there is a significant fraction of H₂ that is ‘CO-dark’. In the fourth row, we show the CO-to-H₂ conversion factor (X_{CO}), defined as in equation (1). This is low throughout the disc compared to the Milky Way value (here we find an emission-weighted $X_{\text{CO}} \sim 10^{19} \text{ cm}^{-2} \text{ K}^{-1} \text{ km}^{-1} \text{ s}$, a factor of 10 smaller than expected). This discrepancy will be further quantified and discussed in more detail below.

Next, comparing the different columns, we investigate the effect of changing the definition of shielding length (columns 1–3) or increasing the cosmic ray ionization rate. We do not see a large difference between choosing the Sobolev length plus a contribution from the local cell, Jeans length or smoothing length (perhaps because the median values are similar; Fig. 2). The most striking difference arises when we increase the cosmic ray ionization rate by a factor of 10. This dramatically decreases the CO abundance and hence the CO(1–0) emission, while only having a moderate effect on the H₂ abundance. This increases the amount of ‘CO-dark’ gas and therefore raises the CO-to-H₂ conversion factor. This will be discussed further in Section 3.5.

3.2 Comparison with observations

We next wish to check how these models compare with observations of CO in nearby galaxies (Fig. 4). In the first panel, we plot the CO(1–0) luminosity as a function of star formation rate and compare with the observed xCOLD GASS sample (Saintonge et al. 2017). For all cases of our assumed shielding lengths, the CO(1–0) luminosity appears to be overestimated by almost an order of magnitude. We also find very little difference between the different definitions of shielding length, with the CO(1–0) luminosity varying by only 5 per cent depending on the definition that we use. Likewise, we see similar results for different zoom simulations of Milky-Way-mass galaxies (the m12b and m12c simulations), highlighting that this issue is not specific to an individual simulation. We note that we are plotting the instantaneous star formation rate in the simulation here, which we compute by summing the star formation rate over all the gas in the disc, and that this may not be a fair comparison with the observations (e.g. Hayward et al. 2014; Sparre et al. 2015). However, this effect is unlikely to be significant enough to explain our overestimates of the CO(1–0) luminosity, and m12i in particular has a very stable recent star formation history, with an instantaneous star formation rate that is close to the average over the last 100 Myr.

In the middle panel, we plot the H₂ gas masses we find in the simulations (filled triangles). We also show the observational estimates, but emphasize that these are derived quantities and that a conversion factor must be assumed (shown here in the right-hand panel). We find that our H₂ masses are somewhat low at a given star formation rate, but they do fall within the scatter of the observations. For comparison, we also show the H₂ masses estimated using the Krumholz & Gnedin (2011) method, which is used in the simulations to calculate the molecular fraction of gas.

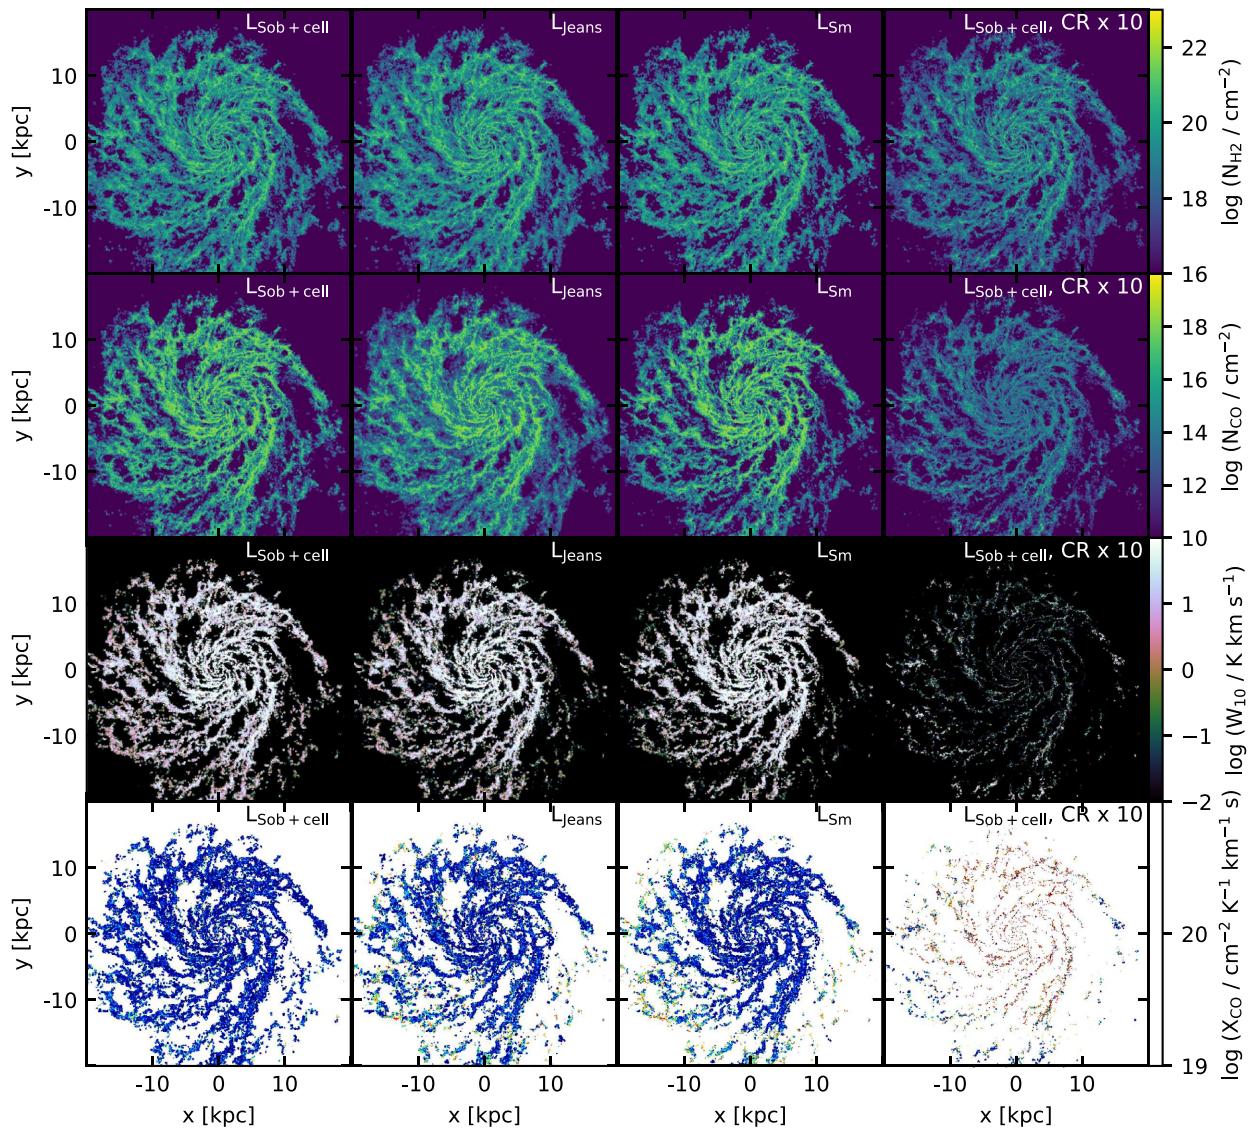


Figure 3. Maps of different quantities resulting from our chemical modelling and line radiative transfer in the m12i simulation, in which we have explored different definitions of the shielding length and normalizations of the cosmic ray ionization rate we assume. The first row shows the H₂ column density for four cases: assuming the shielding length is the Sobolev length plus a contribution from the local cell, assuming the shielding length is the Jeans length, assuming the shielding length is the smoothing length, and assuming the shielding length is the Sobolev length plus a contribution from the local cell, but also increasing the cosmic ray ionization rate by a factor of 10. The second row shows the CO column density for the four cases shown in the second row. The third row shows the velocity-integrated CO(1–0) brightness temperature for the four cases shown in the second row. The fourth row shows the CO-to-H₂ conversion factor (X_{CO}) for the four cases shown in the second row for pixels that have $W_{10} > 0.1 \text{ K km s}^{-1}$.

We find that our H₂ masses are consistently smaller than the Krumholz & Gnedin (2011) estimate by about a factor of 2, although these also fall towards the lower end of the H₂ masses inferred from observations. This is consistent with what was found in Orr et al. (2018), which showed that these simulations somewhat underestimate the surface density of cold and dense gas at a given star formation rate surface density. Part of the difference between the Krumholz & Gnedin (2011) estimates and our chemical modelling could be due to differences in the assumed radiation field. When computing the Krumholz & Gnedin (2011) estimate, we assume the radiation field is a function of density and metallicity, which may differ from the inhomogeneous UV background computed in the FIRE-2 simulations that we are using in our chemical modelling (see the right-hand panel of Fig. 1). The mean amplitude of this inhomogeneous UV background is $1.04 G_0$ in m12i. This is close to

the Milky Way average, and perhaps somewhat lower than expected given that the star formation rate in the simulation is a few times higher than the Milky Way value. We find that the amplitude of this background is highly correlated with density (Fig. 3) as assumed in the Krumholz & Gnedin (2011) model.

In the right-hand panel of Fig. 4 we compare the CO-to-H₂ conversion factor assumed for the xCOLD GASS sample using the metallicity-dependent CO-to-H₂ conversion factor determined by Accurso et al. (2017). The CO-to-H₂ conversion factor plotted for our models here is an emission-weighted average, defined following Narayanan et al. (2012) as

$$\langle X_{\text{CO}} \rangle = \frac{\int \Sigma_{\text{H}_2} dA}{\int W_{10} dA}, \quad (9)$$

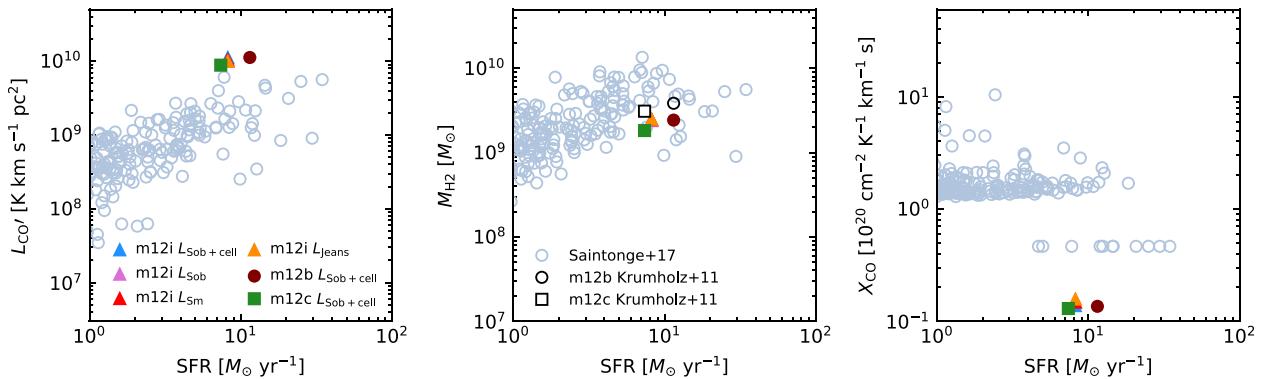


Figure 4. CO(1–0) luminosity (left-hand panel), H₂ gas mass (middle panel), and CO-to-H₂ conversion factor (right-hand panel) at a given star formation rate. Plotted in light blue are the observations from the xCOLD GASS sample (Saintonge et al. 2017) assuming the metallicity-dependent CO-to-H₂ conversion factor from Accurso et al. (2017). The blue, orange, red, and pink triangles show the results for the m12i galaxy assuming different shielding lengths. Note that these points overlap with each other almost perfectly. The brown circle and green square show results from other simulations, assuming the Sobolev shielding length plus a contribution from the local cell. In the middle panel, the black open points are obtained from calculating the H₂ fraction for each particle using the Krumholz & Gnedin (2011) fitting function (which is used to determine the molecular fraction of gas in FIRE-2).

where Σ_{H_2} is the H₂ surface density. We find that our models have a CO-to-H₂ conversion factor that is too low by almost an order of magnitude compared to what is assumed in the Accurso et al. (2017) estimates. It is even lower than the value for merging galaxies assumed in the xCOLD GASS sample ($4.7 \times 10^{19} \text{ cm}^{-2} \text{ K}^{-1} \text{ km}^{-1} \text{ s}$). While there are expected to be trends in the conversion factor based on environmental parameters (Sandstrom et al. 2013), the conversion factor we recover here is more consistent with values usually assumed for ULIRGs (Bolatto et al. 2013), suggesting that something is missing in our chemical modelling. This will be investigated in more detail in Section 3.4. Using the shielding length favoured by Safranek-Shrader et al. (2017) (the Jeans length with a temperature cap of 40 K) did reduce the CO(1–0) emission, but only by a factor of 2. Accounting for dust depletion would further reduce the CO(1–0) emission by another factor of 2. This would bring the CO(1–0) luminosity into better agreement with the observations (although still towards the high end), but would still underpredict the CO-to-H₂ conversion factor.

As mentioned above and demonstrated in the middle panel of Fig. 4, we found that our H₂ masses are reasonably consistent with the data. We further find that the radial profile of the H₂ surface density is in good agreement with the measurement of Miville-Deschénes et al. (2017) in the central 10 kpc of the galaxy, and in fact overpredicts the H₂ surface density by almost an order of magnitude at larger radii. This excess of H₂ surface density at large radii is also present when we assume the Krumholz & Gnedin (2011) relation for calculating the molecular fraction. We further find that the atomic to molecular transition in our simulations is in agreement with observations from Gillmon et al. (2006), Wolfire et al. (2008), and Rachford et al. (2009). As increasing the H₂ abundance further would break the agreement with the observations of the H₂ surface density at small radii, this suggests that our models have a CO(1–0) luminosity that is too high for their H₂ gas mass.

3.3 Modelling the CO-to-H₂ conversion factor

We next test the CO and H₂ abundances in our models against absorption line measurements at UV wavelengths of the CO and H₂ column densities. This allows us to neglect the radiative transfer effects, and isolate any effects in our chemical modelling. To measure the column densities in the simulations, we map the CO and H₂

masses of each particle on to a 2048^2 regular Cartesian grid (chosen to be at the same spatial resolution as the line emission maps we produce with RADMC-3D). We interpolate assuming a cubic spline, rotate the galaxy to be face-on, and treat each pixel of this map as a single sightline from which we take the CO and H₂ column densities. The results of this are shown in Fig. 5. In the left-hand panel, we show the CO and H₂ column densities we measure in the Jeans length, smoothing length, Sobolev length, and the length scale used in the FIRE-2 simulations ($L_{\text{Sob+cell}}$) compared to the absorption line data. We find that, consistent with Fig. 4, we seem to overproduce CO at a given H₂ abundance. In particular, the point where the CO column density begins to steeply rise appears to begin at a lower H₂ column density in our models than what is favoured by the observations. This leads to CO column densities that are too high for their equivalent H₂ column density compared with the observations. The same is true for different haloes in the FIRE-2 suite of simulations (shown for the Sobolev length plus local cell correction model, $L_{\text{Sob+cell}}$, in the middle panel), confirming that this is an issue in the chemical modelling and is not specific to the properties of the m12i simulation.

After exploring the parameter space of assumptions in our modelling (discussed in more detail below), we found that by far the most effective way to reduce the CO abundance while not strongly changing the H₂ abundance was by assuming a smaller shielding length. This is perhaps indicative of unresolved substructure in the gas distribution. We parametrized this in terms of a constant factor times the density-based Sobolev length plus the cell correction ($L_{\text{Sob+cell}}$) measured for each resolution element, the same length scale used to convert volume densities to column densities in the hydrodynamic simulation. Reducing the shielding length means that at a given volume density, the associated column density will be smaller. Since CO primarily depends on being shielded from photodissociation to form, while H₂ is more sensitive to the volume density and metallicity (e.g. Glover & Mac Low 2011), changing this shielding length changes the ratio of CO to H₂.

As shown in the right-hand panel of Fig. 5, lowering the shielding length has the desired result of reducing the abundance of CO in our models at a given H₂ column density. As the shielding length is decreased, the point at which CO can form effectively is shifted to increasingly higher H₂ column densities. Decreasing the shielding length further then leads us to undershoot the CO column densities at large H₂ column densities. It is therefore possible to scale the

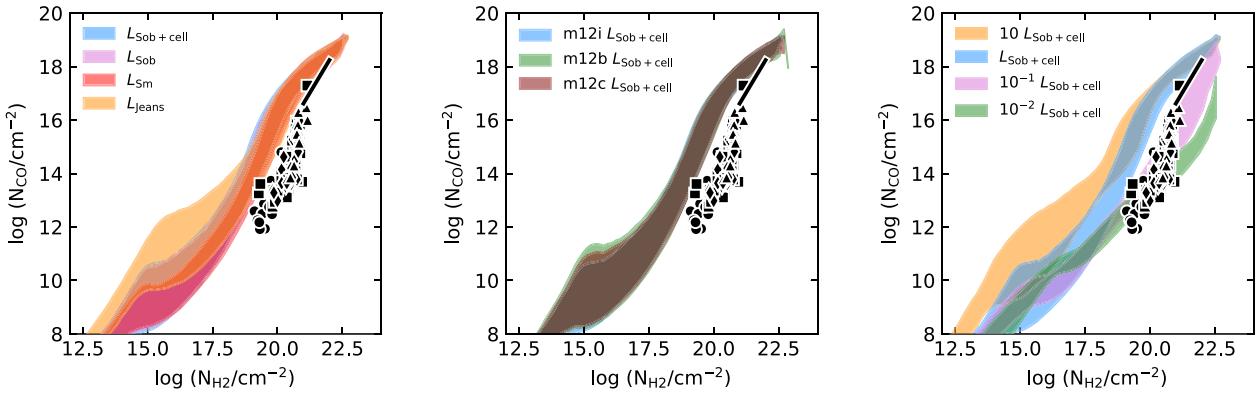


Figure 5. Relation between CO and H₂ column densities for different definitions of shielding length (left-hand panel), different haloes (middle panel), and models where we have multiplied the distribution of shielding lengths by a constant factor (right-hand panel). The shaded regions show the range spanned by the 15th and 85th percentiles of the models. Also plotted are constraints from UV absorption lines in the Milky Way including a fit to a compilation of observations from Federman, Huntress & Prasad (1990) (solid line), as well as observations from Rachford et al. (2002) (triangles), Crenny & Federman (2004) (circles), Sheffer et al. (2008) (squares), and Burgh, France & Jenkins (2010) (diamonds).

distribution of shielding lengths in our simulation such that we can find good agreement with the absorption line data. This smaller shielding length can perhaps be thought of as a subgrid model for the unresolved structure in the density field of GMCs in the simulation, as discussed in Section 4.

3.4 Effect of decreasing the shielding length

In this section, we explore the effects of changing the shielding length on the CO and H₂ masses, the CO(1–0) luminosity and the CO-to-H₂ conversion factor. As above, we change the distribution of shielding lengths by rescaling the distribution of $L_{\text{Sob}} + \text{cell}$ (the same shielding lengths used in the simulation) by a constant value, i.e. each resolution element has its shielding length reduced by the same factor. This was the simplest way to implement this, but, in reality, the factor we multiply by may depend on the local properties of the gas. The results of changing this shielding length are shown in Fig. 6. In the left-hand panel of Fig. 6, we show how the total CO and H₂ masses in the simulation are changed by making the shielding length smaller. Above a certain threshold, there is no effect as all the gas that has a high enough density is fully self-shielded. As the shielding length is decreased further, both the CO and H₂ masses begin to decline. However, the CO is more sensitive to this effect and the CO mass declines faster than the H₂ mass. This allows us to reduce the CO abundance in our model without changing the H₂ abundance by much, exactly as required.

A similar effect is seen in the CO(1–0) luminosity (middle panel). For the unscaled Sobolev lengths with local cell correction, we recover our result from above that we overpredict the CO(1–0) emission compared to the observations of Saintonge et al. (2017). As the shielding length is decreased, the CO(1–0) luminosity falls as less CO is now being formed. For a small enough shielding length, we underproduce the CO(1–0) emission compared to the observations. The behaviour of the CO-to-H₂ conversion factor is slightly more complicated. For large shielding factors, we underpredict the CO-to-H₂ conversion factor favoured for the Milky Way by Bolatto et al. (2013). As the shielding length is decreased, the CO-to-H₂ conversion factor rises, as we are producing less CO in the models while the H₂ abundance stays relatively constant. Eventually, the relation between shielding length and CO-to-H₂ conversion factor turns over, as for small enough shielding lengths, the production of H₂

is also affected. However, this occurs in the regime where the CO(1–0) emission is underpredicted, and is therefore not important here.

To build some intuition for what is happening when we decrease the shielding length, we plot the temperature and density covered by the H₂ and CO in for three different cases of the rescaled shielding length distribution (Fig. 7). This shows that the gas phase traced by CO changes significantly with the assumed shielding length. For the model that uses the unscaled Sobolev shielding length plus local cell correction, we find that CO exists in gas with densities as low as $n_{\text{H}} = 1 \text{ cm}^{-3}$ and temperatures as high as $T = 10^4 \text{ K}$. As the shielding length decreases, the CO in particular is restricted to increasingly higher densities and lower temperatures. This is also true for the H₂, but the effect is not as strong.

By rescaling the shielding length by an appropriate factor, it is therefore possible to ‘tune’ our model to match both the observed CO(1–0) luminosity and estimated CO-to-H₂ conversion factor at fixed star formation rate. One such model is shown in Fig. 6 and is represented by the star. This model rescales the distribution of Sobolev lengths with the local cell correction by a factor of 0.05, which results in a median shielding length of 3 pc in the cold ($T < 300 \text{ K}$) gas in the disc. It is reassuring that it is possible to find a rescaled shielding length distribution that simultaneously reproduces the expected CO(1–0) luminosity and estimated CO-to-H₂ conversion factor. Maps of the resulting H₂ and CO column densities, CO(1–0) brightness temperature, and CO-to-H₂ conversion factor for the model with 0.05 $L_{\text{Sob}} + \text{cell}$ are shown in Fig. 8. In contrast to the panels of Fig. 3 that use the same cosmic ray ionization rate, we now find a lower CO column density and much of the diffuse CO outside the spiral arms has been removed. Likewise the CO(1–0) emission is now restricted to the spiral arms of the galaxy. We also find a higher CO-to-H₂ conversion factor throughout the disc, with the regions that are brightest in CO(1–0) emission showing a conversion factor that is in line with observational estimates (e.g. Bolatto et al. 2013). Using this same value for the reduced shielding length also produces a similar level of agreement between the data and other simulations (Fig. 9).

3.5 Effects of changing other parameters

There are of course other assumptions that we make when calculating the chemical abundances, and we investigate this further in Fig. 10.

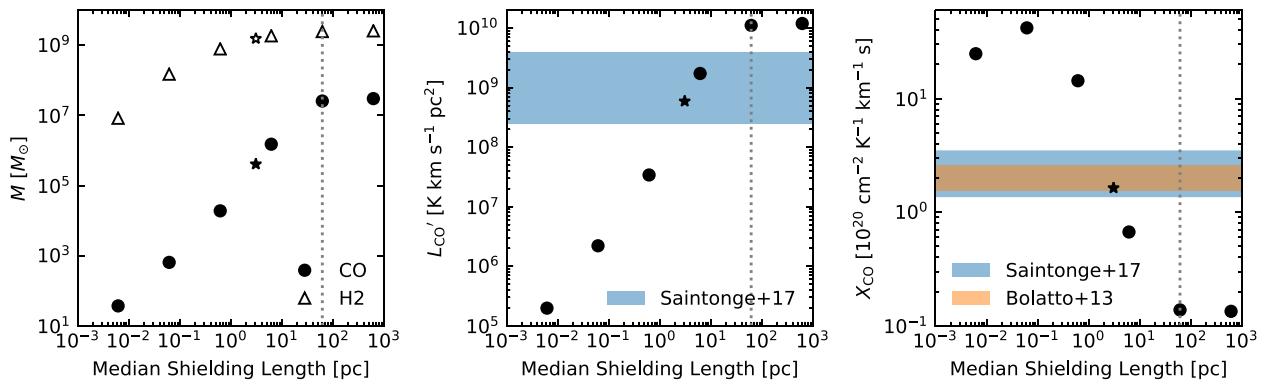


Figure 6. The effects of changing the shielding length on the CO and H₂ masses (left-hand panel), the CO(1–0) luminosity (middle panel), and the emission-weighted CO-to-H₂ conversion factor (right-hand panel). The shielding lengths here are calculated using the Sobolev length plus a contribution from the local cell, rescaled by a constant factor. The grey dotted vertical line shows median shielding length of the original distribution. The shielding length plotted here is the median in shielding length for cold ($T < 300$ K) gas in the disc. The star represents a model that falls within the observed scatter in CO(1–0) luminosity and CO-to-H₂ conversion factor for the star formation rate in the simulation. This model is equivalent to 5 per cent of the distribution of Sobolev lengths plus a contribution from the local cells, which gives a median shielding length of 3 pc in the cold gas. We compare with the observations from the xCOLD GASS survey (Saintonge et al. 2017) and Bolatto et al. (2013).

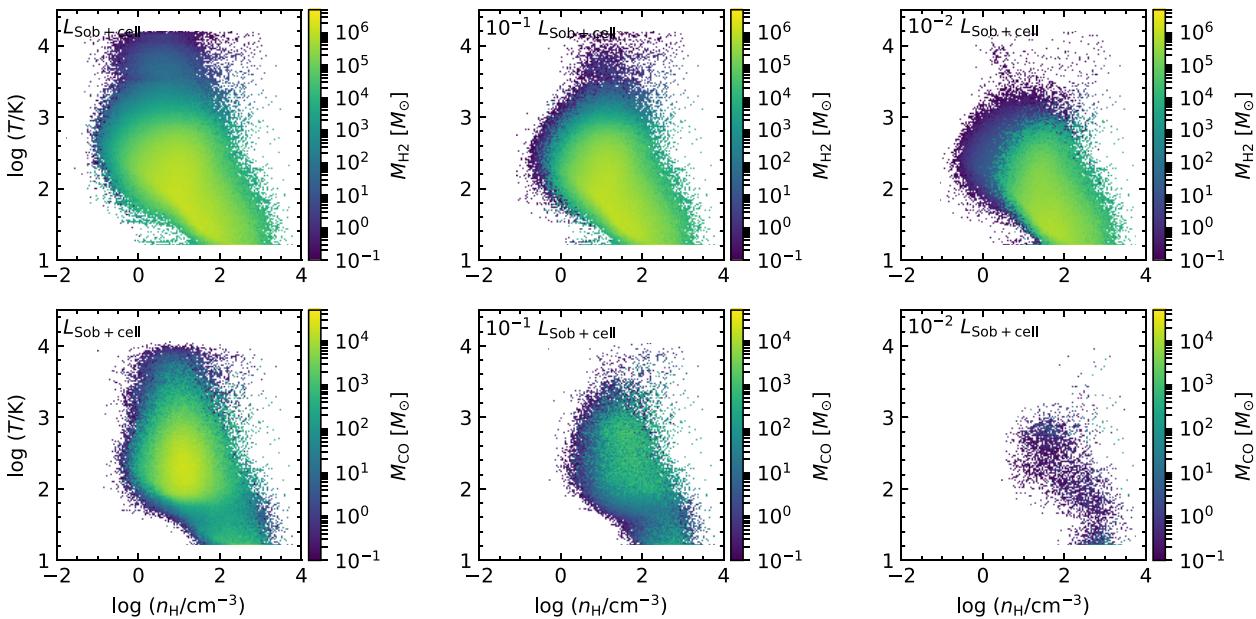


Figure 7. We show the temperature and hydrogen number density of all gas particles that contain at least $0.1 M_{\odot}$ of H₂ (top panel) and CO (bottom panel). When computing the temperature, we assume a mean molecular weight $\mu = 2$ appropriate for molecular gas. Each column shows the distribution of shielding lengths rescaled by a different factor. The colour of the points reflects the mass of that species in a given bin. From the left- to right-hand panels: assuming that the shielding length is the Sobolev length plus a correction from the local cell ($L_{\text{Sob} + \text{cell}}$), 0.1 times $L_{\text{Sob} + \text{cell}}$ and 0.01 times $L_{\text{Sob} + \text{cell}}$. These length scales correspond to median shielding lengths 61.6, 6.16, and 0.616 pc in cold ($T < 300$ K) gas. As the shielding length is reduced, we see that the mass of CO strongly declines. The H₂ mass also declines, but not as steeply.

As in Fig. 6, we show how the CO and H₂ masses, the CO(1–0) luminosity, and the CO-to-H₂ conversion factor change as the shielding length is varied. We present models that make different assumptions compared to our fiducial model (black line). First, we explore the effects of changing the cosmic ray ionization rate, raising/lowering it by a factor of 10 (blue and orange lines). Some observations suggest that the cosmic ray ionization rate could be a factor of 10 lower (Williams et al. 1998). We found that using a lower cosmic ray ionization rate increased the CO and H₂ abundances by a small amount, and did not have a large effect on the CO(1–0) luminosity or CO-to-H₂ conversion factor.

Somewhat more successful was the model where we increased the cosmic ray ionization rate (see also Fig. 3). There is little evidence that the average value for the Milky Way should be higher; however, it may scale with the star formation rate of the galaxy (Lacki, Thompson & Quataert 2010). The star formation rate in the simulation is almost a factor of 4 higher than estimates for the Milky Way (e.g. Chomiuk & Povich 2011), which provides some motivation that the cosmic ray ionization rate could be higher. Apart from changing the shielding length, we find that increasing the cosmic ray ionization rate was the most effective way to produce less CO in the models. This is because cosmic rays are able to penetrate

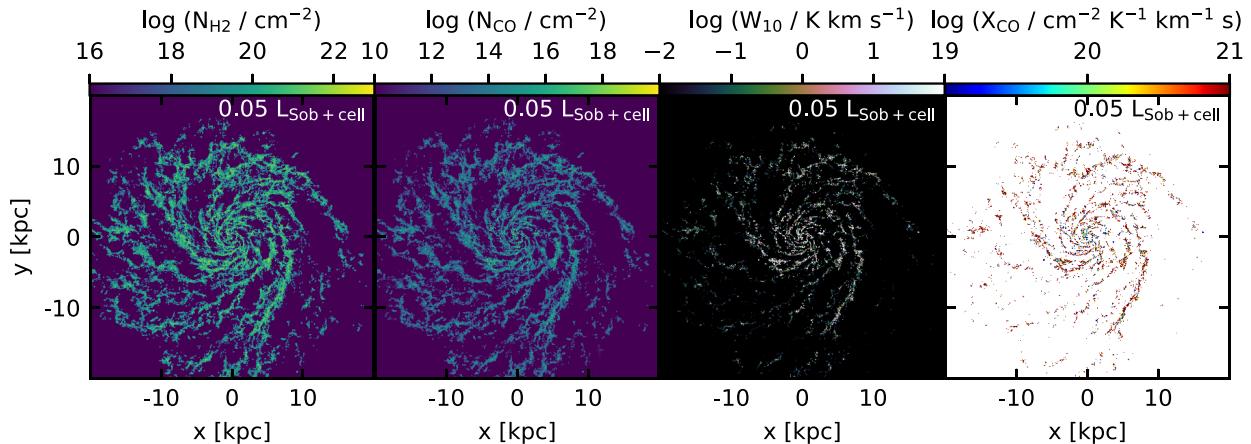


Figure 8. From the left- to right-hand side: For a shielding length of $0.05 L_{\text{Sob}} + \text{cell}$, a model that has integrated CO(1–0) luminosity and emission-weighted CO-to-H₂ conversion factor in line with observations, given its star formation rate, we show maps of the H₂ column density, the CO column density, the velocity-integrated CO(1–0) brightness temperature, and the CO-to-H₂ conversion factor for pixels that have $W_{10} > 0.1 \text{ K km s}^{-1}$ in the m12i simulation.

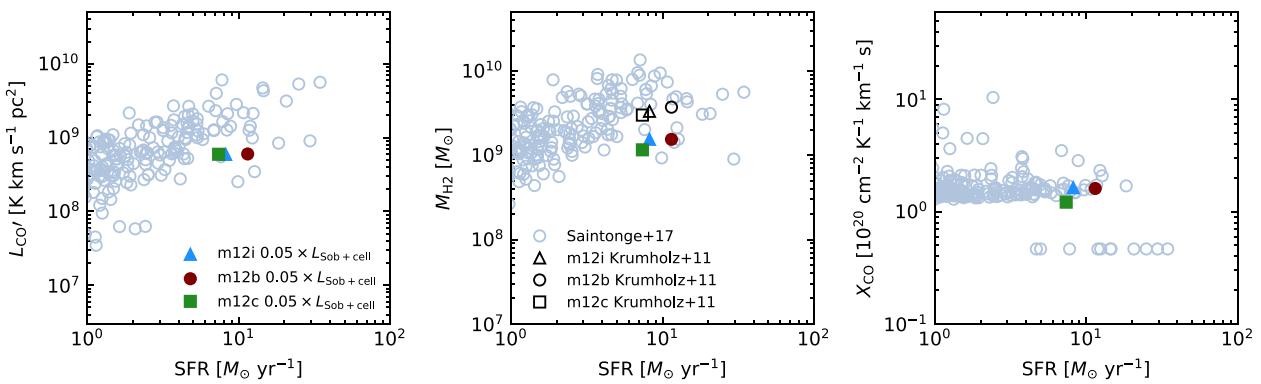


Figure 9. CO(1–0) luminosity (left-hand panel), H₂ gas mass (middle panel), and CO-to-H₂ conversion factor (right-hand panel) at a given star formation rate. Plotted in light blue are the observations from the xCOLD GASS sample (Saintonge et al. 2017), assuming the metallicity-dependent CO-to-H₂ conversion factor from Accurso et al. (2017). The coloured points are results from the simulations all assuming that the distribution of Sobolev shielding lengths plus local cell correction has been rescaled by a constant factor of 0.05, corresponding to a median length of 3 pc for cold ($T < 300$ K) gas. In the middle panel, the black open points are obtained from calculating the H₂ fraction for each particle using the Krumholz & Gnedin (2011) fitting function.

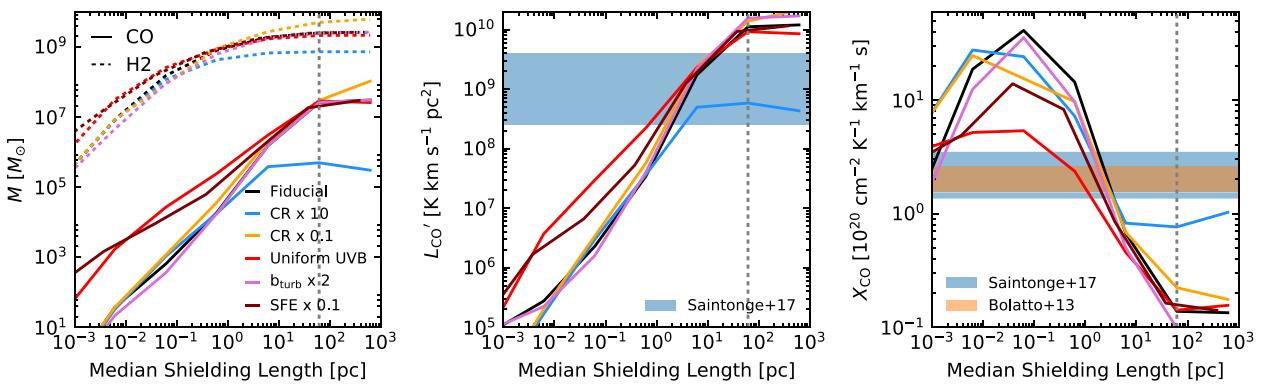


Figure 10. The effects of the changing the shielding length on the CO and H₂ masses (left-hand panel), the CO(1–0) luminosity (middle panel), and the emission-weighted CO-to-H₂ conversion factor (right-hand panel). The grey dotted vertical line shows the median of the original distribution of Sobolev lengths plus local cell correction. The black lines are the same as the black points in Fig. 6. The other coloured lines show the effects of changing different assumptions in our modelling: increasing the cosmic ray ionization rate by a factor of 10 (blue line), lowering the cosmic ray ionization rate by a factor of 10 (orange line), using a homogeneous UV background set to the interstellar radiation field from Black (1987) (red line), increasing the subgrid turbulence in our model by a factor of 2 (pink line) and restarting the simulation with a lower local star formation efficiency to allow more dense gas to build up (maroon line). We compare with the observations from the xCOLD GASS survey (Saintonge et al. 2017) and Bolatto et al. (2013).

the shielded regions and destroy the CO via a reaction with He II (Bisbas, Papadopoulos & Viti 2015; Narayanan & Krumholz 2017). Increasing the cosmic ray ionization rate has however a smaller effect on the H₂, although the H₂ abundance is still decreased. We found that, when using the unscaled Sobolev length plus local cell correction as our shielding length, increasing the cosmic ray ionization rate resulted in a CO(1–0) luminosity that was compatible with the observations. However, we do not favour this solution as the emission-weighted CO-to-H₂ conversion factor we measure was still too small by a factor of a few.

Next, we explore the effects of using a homogeneous UV background set to the interstellar radiation field from Black (1987). This is obviously not a realistic case, compared with our fiducial model for the UV background that accounts for the effects of local sources, and in reality we expect at least some of the GMCs to be sitting near young stars and therefore in regions of enhanced UV flux (see, e.g. Fig. 1). In practice, however, we find little difference between our fiducial model (which includes a inhomogeneous UV background) and this model with a uniform UV background. This is because the choice of UV background makes little difference at high shielding lengths, as the CO and H₂ are so effectively shielded that their abundances are almost independent of the assumed UV flux. At lower shielding lengths, we see that the CO abundances are increased when the uniform background is used compared with the inhomogeneous case due to the overall lower amplitude of the UV background. However, these differences only arise when the CO(1–0) luminosities are already underproduced compared to the observations.

We also explore the effect of changing the subgrid turbulence, which accounts for unresolved turbulent motion in the simulation. As described in Section 2.2, our fiducial subgrid turbulence has a line width of 7.1 km s^{−1}, corresponding to a velocity dispersion of 5 km s^{−1}. This value is also used in Krumholz (2012) to model the H I to H₂ transition and its implications for star formation, and is close to the observed velocity dispersion in nearby GMCs. We test the effect of increasing this Doppler parameter by a factor of 2 (pink line, Fig. 10), but find it only has a minimal effect.

The final parameter we vary is the local star formation efficiency in the underlying cosmological simulation. This parameter sets the fraction of molecular gas that turns into stars per free-fall time. The default FIRE-2 physics assumes that this efficiency is 100 per cent in dense star-forming gas, as the feedback automatically regulates the star formation without having to force a lower efficiency by hand. Hopkins et al. (2018) show that assuming that the local star formation efficiency in dense gas is 100 times lower or higher has a negligible effect on the star formation rate history of the galaxy. Lowering the star formation efficiency in dense star-forming gas however does have an effect on the amount of dense gas present in the simulation: If it is lowered, then more dense gas is accumulated to achieve the same star formation rate. As we are particularly interested in the dense, molecular phase, it is worth investigating this parameter. We restarted a snapshot using a star formation efficiency in dense star-forming gas $\epsilon = 0.1$ (compared to the fiducial $\epsilon = 1$) and reran the simulation for 900 Myr down to redshift zero to allow the dense gas to build up. Looking at the density PDF, we did see a increased amount of dense gas in the distribution compared to the fiducial run. However, changing this efficiency parameter only affected a small fraction of resolution elements in the simulation. This is reflected in our CO and H₂ models that are very similar to the fiducial case at high shielding lengths, although an increase in CO and H₂ abundances is seen at lower shielding lengths. We therefore find that our results are not sensitive to the local star formation efficiency we assume,

as changing this parameter only impacts a small fraction of the gas resolution elements.

In summary, although changing the default parameters in our modelling can affect the CO and H₂ abundances we predict, these effects are generally small at fixed shielding length (Fig. 10). Apart from the shielding length, we found that the most significant parameter was the cosmic ray ionization rate. However, at the fiducial Sobolev plus local cell shielding length, increasing the cosmic ray ionization rate did not produce a CO-to-H₂ conversion factor that agrees with the expected Milky Way value. We therefore conclude that changing the shielding length is the most effective mechanism we found for producing CO and H₂ abundances that are in agreement with the observations, although we note that there could also be alternative methods that may produce the same results.

4 DISCUSSION

By tuning the shielding length to reduce the CO(1–0) luminosity in the simulation, we have introduced a model that matches the observations quite well. This model appears to recover the expected global molecular gas properties of the galaxy, such as its integrated luminosity and emission-weighted CO-to-H₂ conversion factor. Reducing the shielding length may not be the only way to recover these properties, and there may be a combination of different parameters that we could vary to achieve the same result. However, varying the shielding length was the most effective solution that we found. There are two connected reasons that this is an effective strategy. In the models that use common approximations for the shielding length, there is an overabundance of CO for a given amount of H₂ (see, e.g. Fig. 5). This results in an emission-weighted CO-to-H₂ conversion factor that is too low, as most of the emission is coming from low density regions, and the CO-to-H₂ conversion factor is low at these densities due to the overproduction of CO. As we reduce the shielding length, we lower the abundance of CO at lower densities. This raises the CO-to-H₂ conversion factor in that region, but the densities responsible for the bulk of the CO emission also change. This changes the weighting in our quoted CO-to-H₂ conversion factors, and restricts the gas it is probing to denser regions where the CO-to-H₂ conversion factor was initially higher.

Accounting for the effects of dust depletion will likely reduce the CO(1–0) emission and allow us to adopt a somewhat larger shielding length. However, we estimate that this will change things by a factor of 2 and we emphasize that matching the CO-to-H₂ conversion factor requires us to reduce the CO(1–0) emission by a factor of ~ 20 . We do note however that the effects of dust depletion are dependent on both the metallicity, density, and molecular fraction of the gas (see, e.g. Chiang et al. 2018) and we intend to investigate this more carefully in future work.

Decreasing the shielding length is effectively a ‘subgrid’ model that we can tune to compensate for the finite resolution of the simulation. It is therefore interesting to think about the substructure that is not resolved in these simulations, and how it may relate to this model. CO emission comes from GMCs, which are known to have extensive substructure that may effect the optical depth in that region. In fact, there is some observational evidence that optical depths in star-forming regions may be lower than expected. The free-free emission in our galaxy is dominated by the contribution from the extended low density region (ELD; Mezger 1978; Guesten & Mezger 1982). Using observations from *WMAP*, Murray & Rahman (2010) showed that the ELD was associated with ionizing photons leaking out of massive star clusters rather than with photons coming from H II regions as previously thought (Lockman 1976; Anantharamaiah

1985a,b). If ionizing photons are able to escape from star-forming regions, then they must also be able to enter them, equivalent to lowering the optical depth in that region. By using a smaller shielding length scale in our model, we reduce the self-shielding experienced by each particle and effectively allow more of the radiation field to leak into each resolution element, lowering its optical depth.

An alternative interpretation is that the small shielding length we require corresponds to the ‘coherence length’ of the gas. This length scale corresponds to the length over which the velocity of the gas changes by about the width of the shielding line. Gneden & Kravtsov (2011) and Feldmann et al. (2012a) have used a coherence length of 1 pc in their modelling of molecular gas, which allows them to successfully reproduce a range of observations, such as the molecular and atomic gas fraction as a function of gas column density and the dependence of the CO-to-H₂ conversion factor on metallicity. This value of 1 pc is close to the median of the distribution of rescaled shielding lengths in our preferred model.

It is useful to place our work in context by comparing it to other studies of modelling CO in the literature. There are of course other published subgrid models for computing the CO emission from cosmological simulations (e.g. Krumholz 2014a; Olsen et al. 2016; Narayanan & Krumholz 2017; Vallini et al. 2018). In some respects, these models are more sophisticated than the simple model we have constructed here, and include subgrid models for the internal structure of gas resolution elements in the simulation based on properties such as the surface density of a cell or particle. The benefit of these subgrid models is that higher densities can be attained in the imposed substructure, which allows for the prediction of, e.g. higher order CO lines that would not be possible with the model we present here. The benefit of our modelling, however, is that it places a target length scale on the resolution that should be achieved in a cosmological simulation to self-consistently model CO(1–0) emission. Our work is therefore complementary to these other subgrid models, and we do not necessarily suggest it as the favoured approach.

We find that our preferred value of the median shielding length is also approaching the spatial scale of simulations where CO abundances can be computed without the aid of explicit subgrid modelling. Simulations of isolated galaxies, or segments of galaxies, run with chemical networks computed on the fly such as Smith et al. (2014), Richings & Schaye (2016), and Gong et al. (2018), or run with chemical post-processing such as Fujimoto et al. (2019), and have successfully reproduced the expected value of the CO-to-H₂ conversion factor or have matched the observed CO-based Kennicutt–Schmidt relation (which would require a correct X_{CO}). Even the lowest resolution of these simulations has a fixed spatial scale of 8 pc. This is almost a factor of 8 higher in resolution than the median smoothing length in the cold ($T < 300$ K) gas in the simulations we study here (Fig. 2), although the resolution in the densest regions in our simulations can be much higher. The fact that our preferred median shielding length of ~ 3 pc falls in the range of spatial resolution scales in the above works is perhaps indicative that our modelling is placing a limit on the resolution that must be attained to self-consistently model CO in cosmological simulations. However, we note that Joshi et al. (2019) showed that scales of 0.04 pc are required for converged CO formation, which is still far below our preferred median shielding length scale.

Since our claim is that we need this small shielding length to compensate for the finite resolution of the simulations, we could test this with a higher resolution simulation that has a median spatial resolution of the order of 1 pc in the disc. In this case, our work predicts that the simulation should reproduce the observed CO(1–

0) emission and CO-to-H₂ conversion factor using the smoothing length as the shielding length. This is a very ambitious resolution requirement, but the idea could perhaps be tested by restarting snapshots of existing simulations and splitting the particles. In this case, the simulation only has to be run for a dynamical time to allow the gas to settle down, which is computationally more feasible. Such a run has already been performed for m12i, with the spatial resolution increased by a factor of 2 (mass resolution higher by a factor of 8). In our preliminary investigations, we have found that the H₂ mass was increased by 1 per cent, but the CO mass decreased by 53 per cent compared to the original simulation. In both cases, here we assumed the shielding length was equal to the smoothing length. This decrease is in line with what we would expect from our reduced shielding lengths, suggesting that it is indeed the resolution of the simulation that requires us to use a subgrid model. Another test would be to see how much ionizing radiation escapes from young star clusters in the hydrodynamic simulations. If little ionizing radiation escapes, in contrast to what is observed, this would suggest that the radiation is being overattenuated due to the large shielding lengths assumed in the simulation.

Another avenue of investigation is how well the simulations reproduce other emission lines. One possibility would be to look at higher order CO lines. However, these lines have increasingly high critical densities. The critical density of the $J = 1-0$ transition is already at $\sim 10^3$ cm^{−3}, which is close to the highest density gas present in our simulation (although this critical density decreases with increasing optical depth). The $J = 2-1$ critical density is almost an order of magnitude higher so our predictions for this line are unlikely to be accurate. Another possibility is to model line emission from other species, such as the [C II] 158 μm line. Our preliminary results from modelling the [C II] emission are that, as with the CO(1–0) emission, we overpredict the [C II] emission in a model that assumes the shielding length is equal to the Sobolev length plus the correction for the local cell. However, unlike the CO, we do not find better agreement with observations when we rescale the shielding length. This suggests that there may be still something missing in our modelling, and perhaps motivates a more explicit implementation of the ISM substructure (e.g. Olsen et al. 2016; Li et al. 2018; Vallini et al. 2018). We note that the [C II] 158- μm emission and CO(1–0) emission are probing very different phases, with the CO coming from GMCs while the [C II] can trace both ionized and neutral gas. We would therefore not necessarily expect the same approach to reproduce both lines, and leave a more careful study to future work, preferably in simulations with higher mass resolution.

5 SUMMARY AND CONCLUSIONS

In this paper we have presented models of CO(1–0) emission in cosmological zoom simulations of Milky-Way-mass galaxies which reproduce many of the observed properties of our Galaxy. We have used a chemical equilibrium solver to predict the H₂ and CO abundances, and a line radiative transfer code to find the emergent line intensity. We compare these models against observations, specifically testing them against CO(1–0) emission from galaxies in the xCOLD GASS survey (Saintonge et al. 2017) and UV absorption line constraints on CO and H₂ column densities along different lines of sight through the Galaxy.

We explore the effect of changing many of the assumptions that go into our modelling, and find that our results depend most strongly on the shielding length we assume. This shielding length relates the volume densities in the simulation to the column densities required to calculate an optical depth for the gas. We investigate four definitions

for the shielding length derived from the simulation: the Jeans length, the smoothing length, a density-gradient based Sobolev length, and the Sobolev length plus a correction for the local cell (which is used in the FIRE-2 simulations to estimate column densities). These definitions all produce a CO(1–0) luminosity that is too high given the star formation rate in the simulation. By reducing the shielding length, until the median of the distribution is 3 pc for the cold gas, we find that we can reproduce both the global CO(1–0) luminosity and CO-to-H₂ conversion factor expected for the simulated galaxies based on their star formation rates. However, we find that the [C II] 158 μm emission is still overpredicted, suggesting that our modelling can be improved further in the future. Further progress will be therefore be made using higher resolution simulations, simultaneously modelling multiple emission lines and accounting for the effects of dust depletion. Comparing with spatial variations of these quantities within a galaxy may also provide new insights, e.g. comparing with the variations of the CO-to-H₂ conversion factor for individual GMCs as in Narayanan & Hopkins (2013).

Our results suggest that it is important to account for the substructure in GMCs when making predictions for the molecular component of galaxies. This substructure occurs at a spatial scale as yet unresolved in the simulations we post-process here. The exquisite quality of existing and upcoming observations demand realistic theoretical models to aid the interpretation of the detailed ISM physics, and to relate the observable line emission to fundamental quantities such as the total molecular gas mass and star formation rate. However, it is interesting to explore simple models such as the one presented here, as they may provide some insight into the relevant physical processes.

ACKNOWLEDGEMENTS

We thank Desika Narayanan and Xiangcheng Ma for their comments on this draft, and Yuxuan Yuan for his assistance in post processing the particle split snapshot. LCK acknowledges the support of a CITA fellowship and a Beatrice and Vincent Tremaine Fellowship. AJR was supported by a COFUND/Durham Junior Research Fellowship under EU grant 609412, and by the Science and Technology Facilities Council [ST/P000541/1]. We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC). This research was undertaken, in part, because of funding from the Canada Research Chairs program. CAFG was supported by NSF through grants AST-1517491, AST-1715216, and CAREER award AST-1652522; by NASA through grant 17-ATP17-0067; and by a Cottrell Scholar Award from the Research Corporation for Science Advancement. Support for PFH was provided by an Alfred P. Sloan Research Fellowship, NSF Collaborative Research Grant 1715847 and CAREER grant 1455342, and NASA grants NNX15AT06G, JPL 1589742, and 17-ATP17-0214. AW received support from NASA, through ATP grant 80NSSC18K1097 and *HST* grants GO-14734 and AR-15057 from STScI, the Heising-Simons Foundation, and a Hellman Fellowship. DK was supported by NSF grant AST-1715101 and the Cottrell Scholar Award from the Research Corporation for Science Advancement. RF acknowledges financial support from the Swiss National Science Foundation (grant no. 157591). Support for SRL was provided by NASA through Hubble Fellowship grant HST-JF2-51395.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS5-26555. Computations were performed on the Niagara supercomputer at the SciNet HPC Consortium. SciNet is funded by the Canada Foundation for Innovation, the Government of Ontario, Ontario

Research Fund – Research Excellence, and the University of Toronto (Loken et al. 2010; Ponce et al. 2019)

DATA AVAILABILITY

The data underlying this paper will be shared on reasonable request to the corresponding author.

REFERENCES

Accurso G. et al., 2017, *MNRAS*, 470, 4750
 Anantharamaiah K. R., 1985a, *J. Astrophys. Astron.*, 6, 177
 Anantharamaiah K. R., 1985b, *J. Astrophys. Astron.*, 6, 203
 Benincasa S. M. et al., 2019, *MNRAS*, 497, 3993
 Bergin E. A., Tafalla M., 2007, *ARA&A*, 45, 339
 Bisbas T. G., Papadopoulos P. P., Viti S., 2015, *ApJ*, 803, 37
 Black J. H., 1987, in Hollenbach D. J., Thronson H. A., eds, *Interstellar Processes Heating and Cooling of the Interstellar Gas, Interstellar Processes: Astrophysics and Space Science Library (A Series of Books on the Recent Developments of Space Science and of General Geophysics and Astrophysics Published in Connection with the Journal Space Science Reviews*, Vol. 134, Springer, Dordrecht
 Bolatto A. D., Wolfire M., Leroy A. K., 2013, *ARA&A*, 51, 207
 Burgh E. B., France K., Jenkins E. B., 2010, *ApJ*, 708, 334
 Carilli C. L., Walter F., 2013, *ARA&A*, 51, 105
 Castor J. I., 1970, *MNRAS*, 149, 111
 Chiang I. D., Sandstrom K. M., Chastenet J., Johnson L. C., Leroy A. K., Utomo D., 2018, *ApJ*, 865, 117
 Chomiuk L., Povich M. S., 2011, *AJ*, 142, 197
 Clark P. C., Glover S. C. O., 2015, *MNRAS*, 452, 2057
 Crenn T., Federman S. R., 2004, *ApJ*, 605, 278
 De Cia A., Ledoux C., Mattsson L., Petitjean P., Srianand R., Gavignaud I., Jenkins E. B., 2016, *A&A*, 596, A97
 Downes D., Solomon P. M., 1998, *ApJ*, 507, 615
 Duarte-Cabral A., Acreman D. M., Dobbs C. L., Mottram J. C., Gibson S. J., Brunt C. M., Douglas K. A., 2015, *MNRAS*, 447, 2144
 Dullemond C. P., Juhasz A., Pohl A., Sereshti F., Shetty R., Peters T., Commercon B., Flock M., 2012, *Astrophysics Source Code Library*, record ascl:1202.015
 Escala I. et al., 2018, *MNRAS*, 474, 2194
 Faucher-Giguère C.-A., Lidz A., Zaldarriaga M., Hernquist L., 2009, *ApJ*, 703, 1416
 Federman S. R., Huntress W. T. J., Prasad S. S., 1990, *ApJ*, 354, 504
 Feldmann R., Gnedin N. Y., Kravtsov A. V., 2012a, *ApJ*, 747, 124
 Feldmann R., Gnedin N. Y., Kravtsov A. V., 2012b, *ApJ*, 758, 127
 Ferland G. J. et al., 2013, *Rev. Mex. Astron. Astrofis.*, 49, 137
 Ferrière K. M., 2001, *Rev. Mod. Phys.*, 73, 1031
 Fujimoto Y., Chevance M., Haydon D. T., Krumholz M. R., Kruijssen J. M. D., 2019, *MNRAS*, 487, 1717
 Gillmon K., Shull J. M., Tumlinson J., Danforth C., 2006, *ApJ*, 636, 891
 Glover S. C. O., Mac Low M. M., 2011, *MNRAS*, 412, 337
 Glover S. C. O., Smith R. J., 2016, *MNRAS*, 462, 3011
 Glover S. C. O., Federrath C., Mac Low M. M., Klessen R. S., 2010, *MNRAS*, 404, 2
 Gnedin N. Y., Kravtsov A. V., 2011, *ApJ*, 728, 88
 Gnedin N. Y., Tassis K., Kravtsov A. V., 2009, *ApJ*, 697, 55
 Goldreich P., Kwan J., 1974, *ApJ*, 189, 441
 Gong M., Ostriker E. C., Kim C.-G., 2018, *ApJ*, 858, 16
 Grenier I. A., Casandjian J.-M., Terrier R., 2005, *Science*, 307, 1292
 Guesten R., Mezger P. G., 1982, *Vistas Astron.*, 26, 159
 Guszejnov D., Grudić M. Y., Offner S. S. R., Boylan-Kolchin M., Faucher-Giguère C.-A., Wetzel A., Benincasa S. M., Loebman S., 2020, *MNRAS*, 492, 488
 Hahn O., Abel T., 2011, *MNRAS*, 415, 2101
 Hayward C. C. et al., 2014, *MNRAS*, 445, 1598
 Heyer M. H., Carpenter J. M., Snell R. L., 2001, *ApJ*, 551, 852
 Hopkins P. F., 2015, *MNRAS*, 450, 53

Hopkins P. F., 2017, *MNRAS*, 466, 3387

Hopkins P. F., Kereš D., Oñorbe J., Faucher-Giguère C.-A., Quataert E., Murray N., Bullock J. S., 2014, *MNRAS*, 445, 581

Hopkins P. F. et al., 2018, *MNRAS*, 480, 800

Hopkins P. F., Grudić M. Y., Wetzel A., Kereš D., Faucher-Giguère C.-A., Ma X., Murray N., Butcher N., 2019, *MNRAS*, 491, 3702

Indriolo N., McCall B. J., 2012, *ApJ*, 745, 91

Israel F. P., 1997, *A&A*, 328, 471

Jenkins E. B., 2009, *ApJ*, 700, 1299

Joshi P. R., Walch S., Seifried D., Glover S. C. O., Clarke S. D., Weis M., 2019, *MNRAS*, 484, 1735

Kannan R., Marinacci F., Vogelsberger M., Sales L. V., Torrey P., Springel V., Hernquist L., 2019, preprint(arXiv:1910.14041)

Katz H., Kimm T., Sijacki D., Haehnelt M. G., 2017, *MNRAS*, 468, 4831

Kennicutt R. C., Evans N. J., 2012, *ARA&A*, 50, 531

Kim J.-h. et al., 2014, *ApJS*, 210, 14

Kroupa P., 2001, *MNRAS*, 322, 231

Krumholz M. R., 2012, *ApJ*, 759, 9

Krumholz M. R., 2014a, *MNRAS*, 437, 1662

Krumholz M. R., 2014b, *Phys. Rep.*, 539, 49

Krumholz M. R., Gnedin N. Y., 2011, *ApJ*, 729, 36

Lacki B. C., Thompson T. A., Quataert E., 2010, *ApJ*, 717, 1

Lagos C. d. P., Bayet E., Baugh C. M., Lacey C. G., Bell T. A., Fanidakis N., Geach J. E., 2012, *MNRAS*, 426, 2142

Langer W. D., Velusamy T., Pineda J. L., Willacy K., Goldsmith P. F., 2014, *A&A*, 561, A122

Leitherer C. et al., 1999, *ApJS*, 123, 3

Leroy A. K. et al., 2011, *ApJ*, 737, 12

Li Q., Narayanan D., Davè R., Krumholz M. R., 2018, *ApJ*, 869, 73

Lockman F. J., 1976, *ApJ*, 209, 429

Loken C. et al., 2010, *J. Phys.: Conf. Ser.*, 256, 012026

Lupi A., Bovino S., 2019, *MNRAS*, 492, 2818

Lupi A., Volonteri M., Decarli R., Bovino S., Silk J., Bergeron J., 2019, *MNRAS*, 488, 4004

McKee C. F., Ostriker E. C., 2007, *ARA&A*, 45, 565

Ma X., Hopkins P. F., Wetzel A. R., Kirby E. N., Anglés-Alcázar D., Faucher-Giguère C.-A., Kereš D., Quataert E., 2017, *MNRAS*, 467, 2430

Mezger P. O., 1978, *A&A*, 70, 565

Miville-Deschénes M.-A., Murray N., Lee E. J., 2017, *ApJ*, 834, 57

Murray N., Rahman M., 2010, *ApJ*, 709, 424

Narayanan D., Hopkins P. F., 2013, *MNRAS*, 433, 1223

Narayanan D., Krumholz M. R., 2017, *MNRAS*, 467, 50

Narayanan D., Krumholz M., Ostriker E. C., Hernquist L., 2011, *MNRAS*, 418, 664

Narayanan D., Krumholz M. R., Ostriker E. C., Hernquist L., 2012, *MNRAS*, 421, 3127

Obreschkow D., Heywood I., Klöckner H. R., Rawlings S., 2009, *ApJ*, 702, 1321

Olsen K. P., Greve T. R., Brinch C., Sommer-Larsen J., Rasmussen J., Toft S., Zirm A., 2016, *MNRAS*, 457, 3306

Orr M. E. et al., 2018, *MNRAS*, 478, 3653

Pallottini A. et al., 2019, *MNRAS*, 487, 1689

Peñaloza C. H., Clark P. C., Glover S. C. O., Klessen R. S., 2018, *MNRAS*, 475, 1508

Pineda J. L., Langer W. D., Velusamy T., Goldsmith P. F., 2013, *A&A*, 554, A103

Planck Collaboration et al., 2011, *A&A*, 536, A19

Ponce M. et al., 2019, in Furlani T., ed. *Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines (Learning)*, PEARC '19, ACM, New York, NY, p. 34

Popping G., Narayanan D., Somerville R. S., Faisst A. L., Krumholz M. R., 2019, *MNRAS*, 482, 4906

Rachford B. L. et al., 2002, *ApJ*, 577, 221

Rachford B. L. et al., 2009, *ApJS*, 180, 125

Remy Q., Grenier I. A., Marshall D. J., Casandjian J. M., 2017, *A&A*, 601, A78

Richings A. J., Faucher-Giguère C.-A., 2018, *MNRAS*, 474, 3673

Richings A. J., Schaye J., 2016, *MNRAS*, 458, 270

Richings A. J., Schaye J., Oppenheimer B. D., 2014a, *MNRAS*, 440, 3349

Richings A. J., Schaye J., Oppenheimer B. D., 2014b, *MNRAS*, 442, 2780

Safranek-Shrader C., Krumholz M. R., Kim C.-G., Ostriker E. C., Klein R. I., Li S., McKee C. F., Stone J. M., 2017, *MNRAS*, 465, 885

Saintonge A. et al., 2017, *ApJS*, 233, 22

Sanderson R. E. et al., 2018, *ApJ*, 246, 24

Sandstrom K. M. et al., 2013, *ApJ*, 777, 5

Schaye J., Dalla Vecchia C., 2008, *MNRAS*, 383, 1210

Schöier F. L., van der Tak F. F. S., van Dishoeck E. F., Black J. H., 2005, *A&A*, 432, 369

Seifried D., Haid S., Walch S., Borchert E. M., Bisbas T. G., 2019, *MNRAS*, 492, 1465

Sheffer Y., Rogers M., Federman S. R., Abel N. P., Gredel R., Lambert D. L., Shaw G., 2008, *ApJ*, 687, 1075

Shetty R., Glover S. C., Dullemond C. P., Klessen R. S., 2011a, *MNRAS*, 412, 1686

Shetty R., Glover S. C., Dullemond C. P., Ostriker E. C., Harris A. I., Klessen R. S., 2011b, *MNRAS*, 415, 3253

Smith R. J., Glover S. C. O., Clark P. C., Klessen R. S., Springel V., 2014, *MNRAS*, 441, 1628

Sobolev V. V., 1957, *Soviet Ast.*, 1, 678

Solomon P. M., Rivolo A. R., Barrett J., Yahil A., 1987, *ApJ*, 319, 730

Sparre M. et al., 2015, *MNRAS*, 447, 3548

Springel V., 2005, *MNRAS*, 364, 1105

Vallini L., Pallottini A., Ferrara A., Gallerani S., Sobacchi E., Behrens C., 2018, *MNRAS*, 473, 271

van Dishoeck E. F., Black J. H., 1988, *ApJ*, 334, 771

van Dishoeck E. F., Jonkheid B., van Hemert M. C., 2006, *Faraday Discuss.*, 133, 231

Visser R., van Dishoeck E. F., Black J. H., 2009, *A&A*, 503, 323

Wetzel A. R., Hopkins P. F., Kim J.-h., Faucher-Giguère C.-A., Kereš D., Quataert E., 2016, *ApJ*, 827, L23

Williams J. P., McKee C. F., 1997, *ApJ*, 476, 166

Williams J. P., Bergin E. A., Caselli P., Myers P. C., Plume R., 1998, *ApJ*, 503, 689

Wolcott-Green J., Haiman Z., Bryan G. L., 2011, *MNRAS*, 418, 838

Wolfire M. G., Hollenbach D., Tielens A. G. G. M., 1993, *ApJ*, 402, 195

Wolfire M. G., Tielens A. G. G. M., Hollenbach D., Kaufman M. J., 2008, *ApJ*, 680, 384

Wolfire M. G., Hollenbach D., McKee C. F., 2010, *ApJ*, 716, 1191