High-Resolution Radiative Transfer Modelling of M33

Thomas G. Williams,[1]★ Maarten Baes,[2] Ilse De Looze,[2,3] Monica Relaño,[4,5] Matthew W. L. Smith,[1] Sam Verstocken,[2] and Sébastien Viaene[2,6]

1 *School of Physics & Astronomy, Cardiff University, Queens Buildings, The Parade, Cardiff, CF24 3AA, UK*
2 *Sterrenkundig Observatorium, Universiteit Gent, Krijgslaan 281, B-9000 Gent, Belgium*
3 *Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK*
4 *Departamento de Física Teórica y del Cosmos, Facultad de Ciencias, Universidad de Granada, E-18071 Granada, Spain*
5 *Instituto Universitario Carlos I de Física Teórica y Computacional, Universidad de Granada, E-18071 Granada, Spain*
6 *Centre for Astrophysics Research, University of Hertfordshire, College Lane, Hatfield, AL10 9AB, UK*

ABSTRACT
In this work, we characterise the contributions from both ongoing star formation and the ambient radiation field in Local Group galaxy M33, as well as estimate the scale of the local dust-energy balance (i.e. the scale at which the dust is re-emitting starlight generated in that same region) in this galaxy through high-resolution radiative transfer (RT) modelling, with defined stellar and dust geometries. We have characterised the spectral energy distribution (SED) of M33 from UV to sub-mm wavelengths, at a spatial scale of 100 pc. We constructed input maps of the various stellar and dust geometries for use in the RT modelling. By modifying our dust mix (fewer very small carbon grains and a lower silicate-to-carbon ratio as compared to the Milky Way), we can much better fit the sub-mm dust continuum. Using this new dust composition, we find that we are able to well reproduce the observed SED of M33 using our adopted model. In terms of stellar attenuation by dust, we find a reasonably strong, broad UV bump, as well as significant systematic differences in the amount of dust attenuation when compared to standard SED modelling. We also find discrepancies in the residuals of the spiral arms versus the diffuse interstellar medium (ISM), indicating a difference in properties between these two regimes. The dust emission is dominated by heating due to the young stellar populations at all wavelengths (~80% at 10 μm to ~50% at 1 mm). We find that the local dust-energy balance is restored at spatial scales greater than around 1.5 kpc.

Key words: galaxies: individual: M33 – galaxies: ISM – galaxies: star formation – dust, extinction – radiative transfer

1 INTRODUCTION

Despite only contributing around 1% of the mass of the interstellar medium (ISM) of a galaxy, dust absorbs, scatters, and reprocesses around 30% of the starlight in star-forming galaxies (e.g. Popescu & Tuffs 2002; Viaene et al. 2016). An understanding of the processes governing the interactions of stars and dust is, therefore, essential to understanding how galaxies evolve, their dust properties, and extracting important intrinsic parameters such as the star formation rate (SFR) and initial mass function (IMF). The starlight absorbed in UV and optical is re-emitted by the dust at far-infrared (FIR) and sub-mm wavelengths. Assuming only absorption of light from younger stars, the total infrared (TIR) luminosity can therefore be used as a proxy for star-formation (see, e.g. Murphy et al. 2011). Alternatively, by understanding the (wavelength-dependent) amount of dust attenuation, wavebands that suffer from attenuation can be

corrected using some combination of dust measurements (e.g. Leroy et al. 2008; Hao et al. 2011), or by assuming some dust model (e.g. Charlot & Fall 2000).

One method of modelling the light from a galaxy is by fitting an SED across these wavelengths, often using a large library of models, and several tools are available for this purpose (e.g. da Cunha et al. 2008; Noll et al. 2009; Chevallard & Charlot 2016). However, these tools assume a local dust-energy balance (i.e. that the dust emission per unit area comes from light originating from stars in that same area), which may be unsuitable for modelling sub-kpc regions (Boquien et al. 2015; Smith & Hayward 2018). These tools also neglect the 3D geometry of a galaxy, and do not consider the propagation vectors of photons through this medium. For a complete study of the interactions of the dust and stellar components of a galaxy, 3D radiative transfer (RT) models are required, which take into account this 3D geometry and are not beholden to a per-pixel local dust-energy balance. There are a number of codes available for this purpose (see Steinacker et al. 2013 for a review of these, as well as an overview of the RT mathematics). Due to the complexity of the

★ Email: thomas.williams@astro.cf.ac.uk

RT calculations, and the fact that these calculations are both non-linear and non-local, most of these codes make use of Monte Carlo (MC) or ray-tracing techniques. Unlike traditional SED fitting, RT is computationally very expensive, and thus faces its own series of challenges, such as loss of information due to projection effects, and the difficulty of applying traditional solution algorithms to these problems.

Previous work in this area has tended to focus on "simpler", better-behaved galaxies such as edge-on (or nearly edge-on) spirals (e.g. Misiriotis et al. 2001; Bianchi 2008; Baes et al. 2010; De Looze et al. 2012b,a; De Geyter et al. 2014, 2015; Mosenkov et al. 2016, 2018). Galaxies at lower inclinations have also been modelled, including the spiral galaxy M51 (De Looze et al. 2014) and very nearby galaxy M31 (Viaene et al. 2017b), finding significant variations in dust heating by old and young stellar populations; these works also find that the relative contributions to dust heating are both wavelength- and position-dependent. A large step in increasing the complexity of these simulations was employed by De Looze et al. (2014), using observed images to describe the distribution of stars and dust. A framework for modelling face-on galaxies is currently in development by Verstocken et al. (in prep.), which will be applied to a number of the DustPedia (Davies et al. 2017) galaxies. Different approaches to RT modelling of galaxies have also been employed – in particular, taking an axisymmetric approach, Popescu et al. (2017) have produced an RT model for the Milky Way (MW).

In this work, we perform a high-resolution RT simulation of the third massive spiral galaxy in our Local Group, M33 (the Triangulum Galaxy). Being the third largest spiral on the sky, smaller only than our own MW and M31, and with a close proximity of 840kpc (Madore & Freedman 1991), it is an excellent target of choice for high-resolution observations. M33 has been mapped across many wavelengths with a variety of observatories. Due to the wealth of high-resolution data, this galaxy is therefore naturally suited for detailed RT simulations. M33 has a roughly half-solar metallicity $(12 + \log(O/H) = 8.36 \pm 0.04$, Rosolowsky & Simon 2008a), and a shallow metallicity gradient. This lower metallicity makes M33 a very different environment to M31 and the MW, more analogous to younger, higher redshift galaxies. As the RT model is 3D, the data is necessarily deprojected, and a third dimension modelled, but with a moderate inclination of 56°(Regan & Vogel 1994), the deprojection degeneracies are not as pronounced as in M31. M33 has significant star-formation across its disc (Heyer et al. 2004), with SFRs between $0.2\,M_\odot\,\mathrm{yr}^{-1}$ and $0.45\,M_\odot\,\mathrm{yr}^{-1}$, depending on the SFR tracer used (Verley et al. 2007; Williams et al. 2018). Given its relatively small size $(R_{25} = 30.8'$, ~7.4 kpc, Paturel et al. 2003), this means that M33 has a much higher star formation efficiency than other Local Group galaxies (with a gas depletion timescale of 1.6-3.2×10^8 yr; Gardan et al. 2007). Because of this active star formation, we may expect a higher contribution to the overall dust heating by younger stellar populations, but it is important not to neglect the effect of dust heating by older stellar populations.

Earlier RT studies of M33 have focussed on the nucleus (Gordon et al. 1999), and in modelling the global SED (Hermelo et al. 2016). Gordon et al. (1999) modelled only the ultraviolet to near-infrared (UV-NIR) SED of this nucleus, finding evidence of strong dust attenuation. Hermelo et al. (2016) applied the RT model of Popescu et al. (2011) which uses a series of axisymmetric models to describe the various geometries of the galaxy, and produced a global SED from UV-sub-mm wavelengths. The goal of this study was to investigate the "sub-mm excess", which appears to be present in many low-metallicity environments (e.g. Bot et al. 2010; Galametz et al. 2011; Kirkpatrick et al. 2013; Rémy-Ruyer et al. 2013). The

main conclusion of this work was that likely, the sub-mm excess could be accounted for by modifying the dust grain composition. Our study seeks to build on these previous works, studying the attenuation of M33 on a global level with a richer data-set than Gordon et al. (1999), as well as to modify the dust grain properties in our input model to better fit the data, and to use input geometries based on observables to produce a resolved study of many of these properties.

The layout of this paper is as follows: we present an overview of the dataset we use in this work (Sect. 2), before an overview of the setup of our RT model (Sect. 3). We then fit this model to the observed SED of M33 (Sect. 4), before investigating some of the global and resolved properties of M33 (Sect. 5). Finally, we summarise our main conclusions in Sect. 6.

2 DATA

The data used in this work is largely the same as in Williams et al. (2018), and we refer the reader to that work for a more detailed description. A brief description is given here. Both FUV and NUV data was obtained (Thilker et al. 2005) by the Galaxy Evolution Explorer (GALEX; Martin et al. 2005). In the optical, Sloan Digital Sky Survey (SDSS; York et al. 2000) data was mosaicked together using primary frames from the SDSS DR13 (Alam et al. 2015), using MONTAGE[1]. A 3 square degree mosaic was created for all of the u, g, r, i, and z bands, to allow us to accurately model background variations. We also make use of Hα data (Hoopes & Walterbos 2000), which was not included in the previous work of Williams et al. (2018). This map has a pixel size of around 2 arcsec, and covers a total field-of-view of 1.75 deg². This map has also been continuum-subtracted. Corrections for contamination from [N II] emission have not been carried out, although it is estimated that a maximum of 5% of the flux could result from [N II] emission in any region of the galaxy (Hoopes & Walterbos 2000).

For near- and mid-infrared, we make use of *Spitzer* and Wide-field Infrared Survey Explorer (WISE; Wright et al. 2010) data. The former of these was taken as part of the Local Volume Legacy (LVL, Dale et al. 2009) survey, with Infrared Array Camera (IRAC, Fazio et al. 2004) data at 3.6, 4.5, 5.8 and 8 μm and Multiband Imaging Photometer (MIPS, Rieke et al. 2004) data at 24 and 70 μm. The latter covers a similar wavelength range to the former, with data at 3.4, 4.6, 12, and 22 μm, and has been mosaicked together from the ALLWISE data release, which includes both the WISE cryogenic and NEOWISE (Mainzer et al. 2011) post-cryogenic phase.

Far-infrared and sub-mm data was obtained from the *Herschel* space observatory (Pilbratt et al. 2010) and the Submillimetre Common-User Bolometer Array 2 (SCUBA-2, Holland et al. 2013) on the James Clerk Maxwell Telescope (JCMT). As part of the HerM33es (Kramer et al. 2010) open time key project, M33 was mapped by both the Photodetector Array Camera and Spectrometer (PACS, Poglitsch et al. 2010) at 100 and 160 μm, and the Spectral and Photometric Imaging REceiver (SPIRE, Griffin et al. 2010) at 250, 350, and 500 μm. The SCUBA-2 data is at 450 and 850 μm, and we use the technique presented in Smith et al. (in prep.), to maintain the high resolution offered by SCUBA-2, but add back in the large-scale structure that is lost in the data reduction process. Details of this SCUBA-2 data reduction are given in Williams et al. (2019). We note that this SCUBA-2 data does not cover the entirety

---

[1] http://montage.ipac.caltech.edu

of M33, so for global flux values we use the SPIRE 500 μm and Planck 850 μm fluxes. In total, our dataset covers almost 4 orders of magnitude in wavelength, from 1516 Å to 850 μm.

For each of these images, we have performed a number of steps to make this diverse dataset homogeneous. For frames in which foreground star emission is present, we masked this using UV colours (Leroy et al. 2008). We performed a Galactic extinction correction using the prescription of Schlafly & Finkbeiner (2011), using extinction values calculated for the central position of M33, although we note that due to M33's large angular extent, this correction varies across the face of the disc, which is taken into account in our treatment of the uncertainties. However, for the GALEX bands, which are most affected by this variation in extinction correction, the difference in flux is maximally ~3%, which is negligible when combined with the other errors considered (see Williams et al. 2018, for more details on this error analysis). We then convolved all of the data to our worst working resolution, the SPIRE 350 μm beam (which has a FWHM of 25 arcsec, corresponding to 100 pc at the distance of M33). This data is then regridded to pixels of 25 arcsec, so that they can be considered statistically independent.

With this dataset homogenized, we performed pixel-by-pixel SED fitting for the ~19000 pixels within a radius of 60 arcmin × 70 arcmin, using the SED fitting tool MAGPHYS (da Cunha et al. 2008), and we refer readers to this work for details on the MAGPHYS model details. This allows us to calculate a number of intrinsic quantities of the galaxy, and provides both an attenuated and unattenuated SED for each pixel, with the attenuation following the model of Charlot & Fall (2000). This modelling technique has previously been employed by Viaene et al. (2016) for their modelling of M31, and means that we can make immediate comparison with this earlier work. We also note that Viaene et al. (2014) and Williams et al. (2018) find that MAGPHYS produces similar results to more conventional measures of, e.g., dust mass and SFR with observational data at resolutions of 130 and 100 pc, respectively. Smith & Hayward (2018) find statistically acceptable fits to many key galaxy properties, when compared to simulated data at resolutions of 200 pc to 25 kpc.

3   THE 3D MODEL

For the radiative transfer simulations, we make use of SKIRT[2] (Baes et al. 2003; Camps & Baes 2015), a publicly available, Monte Carlo RT code. This code was originally developed to investigate the effects of dust extinction on the photometry and kinematics of galaxies, but has developed to accurately model the absorption, scattering and emission of starlight by dust. It has also been tested against the major benchmarks published that the code is applicable to (e.g Camps et al. 2015). SKIRT can accept an arbitrary number of components to model, where each of these components are defined by a 3D geometry, an intrinsic spectrum, and a normalisation of this spectrum (either at a given wavelength, or a bolometric luminosity). This code allows for panchromatic RT simulations, using a wide variety of geometry models and optional modifiers for these geometries (Baes & Camps 2015). It also provides a number of options for efficient dust grids (Saftly et al. 2014), for which we use a binary tree adaptive grid method. This means that we can effectively increase the resolution in dense regions (such as spiral arms), while minimising the computational cost of this increased resolution. The

code can also model stochastically heated dust grains (Camps et al. 2015). It is also provided with parallelisation, to allow these computationally expensive simulations to run efficiently (Verstocken et al. 2017). Finally, it allows for the input of a 2D FITS image as a geometry, which was first employed by De Looze et al. (2014) in the grand-design spiral galaxy M51, and which we use to define our various geometries in this work. SKIRT deprojects and derotates this image given an inclination and position angle, and assumes that the distribution of pixel values in this input image corresponds to the density in a linear way. It then scales this map to a total density provided when setting up the geometry, and conserves total flux during deprojection. This 2D model is then given extra dimensionality by assuming an exponential profile with a provided vertical scale height (which will vary for each input geometry).

To make our notation consistent throughout this work, but comparable to earlier studies, we refer to flux densities using the symbol S, and luminosities as L. Fractions of these quantities will be referred to with the symbol $\mathcal{F}$.

3.1   Model Components

A typical galaxy model setup for RT simulations composes of a bulge and thick disc containing old stars, with a thin star-forming disc containing dust and young stars (e.g. Xilouris et al. 1999; Popescu et al. 2000). We use this model with one alteration – M33 does not appear to have a bulge, at least in the traditional sense (Bothun 1992). This claim is somewhat controversial, but for the purposes of our work we treat M33 as bulge-less. This means that we assume all of the old stars reside within the same exponential disc, rather than a population at the centre extending much further above the plane of the galaxy. We use three stellar components in our model: the first represents the old stellar populations (stars of ages around ~8 Gyr; Sect. 3.1.1). The second stellar component consists of the young stars that are UV bright but dissociated from their birth clouds, and have ages around 100 Myr (Sect. 3.1.2). Our final stellar component are the young stars still present in their birth clouds, and producing hard, ionizing radiation (Sect. 3.1.3). We refer to the combination of these young non-ionizing and ionizing stellar populations as "young" throughout this work. We also provide a map of the dust mass surface density, which traces the dust distribution within the galaxy (Sect. 3.1.4). The details of this modelling approach are based on Verstocken et al. (in prep.).

For each component, we specify an input geometry, a particular SED type and a luminosity normalisation. Along with this, we provide an input FITS image, where we have truncated the disc to $1.2\,R_{25}$ and set to 0 any pixels that correspond to those that have signal-to-noise (S/N) < 5 in the SPIRE 350 μm map (the map that defines our working resolution). For each stellar component, we specify a metallicity of $12 + \log(\text{O/H}) = 8.36$, corresponding to the central metallicity of M33 (Rosolowsky & Simon 2008b). Whether M33 has a radial gradient in its metallicity is a topic of contention. Whilst Rosolowsky & Simon (2008b) find a slight radial gradient, Bresolin (2011) find no such significant gradient. In either case, the practical effect this would have on the form of the SED is minor. Finally, in all cases for the geometries we assume a position angle of 22.5°(de Vaucouleurs 1959) and an inclination of 56°(Regan & Vogel 1994). A summary of the major parameters of our model are given in Table 1.

---

[2] skirt.ugent.be

Table 1. Overview of parameters in the model. Although most parameters are fixed, for free parameters we indicate the parameter search range and the wavelength these luminosities are normalised at. All luminosities are given in $L_\odot$ at the normalisation wavelength. The dust mass is given in $M_\odot$. The error in each parameter is calculated by sampling the likelihood distributions, and is quoted as half the bin width if the sampled error is smaller than a single bin.

Component	Parameter	Value	Best Fit Luminosity/Mass ($L_\odot/M_\odot$)
Old Stellar Disc	2D Geometry Total Luminosity Vertical Scale Height	IRAC 3.6 μm $0.4 - 4 \times 10^8$ (3.6 μm) 200 pc	$(2.8^{+1}_{-0.5}) \times 10^8$
Non-Ionizing Stellar Disc	2D Geometry Total Luminosity Vertical Scale Height	GALEX FUV[1] $0.8 - 5 \times 10^9$ (0.15 μm) 100pc	$(1.7 \pm 0.5) \times 10^9$
Ionizing Stars	2D Geometry Total Luminosity Vertical Scale Height	Hα + 24 μm[2] $0.3 - 3.3 \times 10^7$ (0.66 μm) 50 pc	$(3.3 \pm 1.5) \times 10^7$
Dust	2D Geometry Total Dust Mass Vertical Scale Height	MAGPHYS Dust Mass Map[3] $2.5 - 7 \times 10^6$ M$_\odot$ 100 pc	$(3.6 \pm 0.6) \times 10^6$

[1]Corrected for attenuation and diffuse emission. [2]Corrected for diffuse stellar emission. [3]Obtained from pixel-by-pixel MAGPHYS fitting.

3.1.1 Old Stellar Disc

The geometry of our old stellar component is set by the IRAC 3.6 μm image, which is generally considered to be a pure tracer of stellar mass (e.g. Zhu et al. 2010). In our initial testing, we found significant contribution from the young stellar populations at this wavelength, and so using the MAGPHYS star formation history (SFH) we make a first-order correction to separate out the contribution of these younger populations from the total luminosity. We note that as we leave the luminosity of each stellar component as a free parameter, this is only needed for a first guess. There may also be a contribution at this wavelength from hot dust heated by the young stars, but leaving the stellar luminosities as free parameters in our fitting will effectively account for this. Also, in practice the contribution from young populations is likely position-dependent, but given the coarse nature of the MAGPHYS SFH, performing robust corrections of this nature is beyond the scope of this work. We normalise the luminosity of the old stellar disc at 3.6 μm.

For a panchromatic simulation, we require an emitted luminosity at each wavelength for each component. We do this by taking a template SED, and matching the observed emission to this. In the case of this old stellar population, we make use of the Bruzual & Charlot (2003) simple stellar populations (SSPs), at an age of 8 Gyr, which we assume is the average age of these older stars. Finally, to make this geometry 3D, we assume an exponential profile for the disc, characterised by a vertical scale height. Generally, the scale height of the old stellar populations is taken to be 1/8.86 the scale length (De Geyter et al. 2013). In M33 this scale length is 1.82 ± 0.02 kpc (Kam et al. 2015), giving us a scale height of ~ 200 pc.

3.1.2 Non-Ionizing Stellar Disc

The first of our young stellar populations are the stars of age ~ 100 Myr, which are UV bright but unable to ionize hydrogen. These stars are only attenuated by dust in the diffuse ISM, and so suffer much less from dust attenuation than those stars in the birth clouds. We used as our initial input geometry the GALEX FUV image, which traces unobscured star formation over the last 10-100 Myr (Meurer et al. 1999). We calculated an unattenuated flux

for each pixel by convolving the unattenuated MAGPHYS SED with the GALEX FUV filter response, which effectively corrects for the effects of dust attenuation.

Although the FUV is dominated by these young stars, there can also be a significant amount of UV flux from more diffuse, older, stellar populations, which we correct for using the prescription of Leroy et al. (2008):

$$S_{\text{FUV, young}} = S_{\text{FUV, unatten}} - \alpha_{\text{FUV}} S_{3.6}, \tag{1}$$

where $\alpha_{\text{FUV}} = 3 \times 10^{-3}$, and S_x is in Jy. Given that Leroy et al. (2008) do not correct the 3.6 μm flux for young stars when calculating this factor, we use the uncorrected 3.6 μm flux. As these young stars are expected to reside within a thinner disc than the old stars, we adopt a scale height of 100 pc, half that of the old stellar component, and we normalise the extinction-corrected luminosity at the FUV wavelength.

3.1.3 Ionizing Stars

Our final stellar component consists of very young (<10 Myr) stars that are still embedded in their birth clouds and produce hard, ionizing radiation. This radiation is difficult to trace directly, but can be inferred from Hα emission, and dust grains heated to high temperatures. To create a map of the ionizing radiation, we used a continuum-subtracted Hα map, and combined this with a map of the hot dust. The hot dust is traced via the 24μm emission, which, much like the FUV map, we corrected for a diffuse stellar component:

$$S_{24, \text{ion}} = S_{24} - \alpha_{24} S_{3.6}. \tag{2}$$

These fluxes are again in Jy, and α_{24} was determined by Leroy et al. (2008) to be 0.1. This factor was calculated from a sample of nearby galaxies, but appears to be robust throughout the sample (see the discussion in their appendix D.2.4), and thus should be applicable to M33. The input geometry for the ionizing map is then

$$S_{\text{H}\alpha, \text{ion}} = S_{\text{H}\alpha} + 0.031 S_{24, \text{ion}}, \tag{3}$$

where fluxes are in ergs/s (Calzetti et al. 2007).

The SED we used for this input geometry was the MAPPINGS III (Groves et al. 2008) nebular modelling code, and we refer readers

to that work for definitions of the various parameters of the model. Generally, we adopt the same parameters for this SED as De Looze et al. (2014), with a compactness of $\log C = 5.5$, and a surrounding ISM pressure of 1×10^{12} K m^{-3}. However, we choose a slightly lower cloud covering factor of 0.1, half that of De Looze et al. (2014), which we found in initial testing to be a slightly better fit to the data. A lower covering factor leads to slightly colder dust, and a higher fraction of UV flux escaping, which is likely the case in low-metallicity environments. We normalise this luminosity at the wavelength of Hα. We expect the ionizing component to be in a thinner disc than the older stars, and so we used a vertical scale height of 50 pc, half that of the 100 Myr stars, and similar to the scale height of the UV discs (Combes et al. 2012).

3.1.4 Dust System

We created a map of the dust mass as the input component for the dust geometry. For this, we made use of pixel-by-pixel MAGPHYS fits. The dust model is more thoroughly described in da Cunha et al. (2008), but as a brief overview, MAGPHYS models Polycyclic Aromatic Hydrocarbons (PAHs) using a fixed template based on M17, which dominate at MIR wavelengths. The hot dust is modelled as a series of modified blackbodies with temperatures of 850, 250 and 130 K. We use a MAGPHYS library with extended priors (Viaene et al. 2014), meaning that the warm dust is modelled as a modified blackbodies (MBBs) with a fixed β of 1.5, and can vary from 30 to 70 K. The cold dust has a fixed β of 2 and can vary from 10 to 30 K. We use a vertical scale height of 100 pc, the same as the young, non-ionizing stellar population, and similar to the 100 μm scale height found by (Combes et al. 2012). This is also consistent with model predictions of edge-on galaxies (e.g. Xilouris et al. 1999; Bianchi 2008; De Geyter et al. 2013).

This model implicitly assumes a per-pixel local dust-energy balance. As we discuss in Sect. 5.5, this is not an acceptable assumption at scales <1500 pc. However, as shown in Williams et al. (2018), there is a very tight relationship between the dust masses obtained from MAGPHYS, and from a single-temperature MBB, with a median offset of 0.02 dex and an RMS scatter of 0.10 dex for the same data (their Fig. 9). As an additional check, we also performed this comparison on a pixel-by-pixel dust map fitted using our modified THEMIS fitting routine (Appendix A), and find a similar relationship. Thus, this choice of dust map will have a negligible impact on the simulation.

We use the THEMIS dust model (Jones et al. 2013) to describe the dust in M33. This model consists of small and large amorphous hydrocarbons (sCM20 and lCM20), along with silicates (aSilM5, Köhler et al. 2014) to model the diffuse ISM of the MW. This model is primarily laboratory-based, and can naturally explain most of the features of the dust SED in the MW. However, in our initial testing we found that the default THEMIS parameterisation was insufficient to fit the dust SED of M33, particularly at sub-mm wavelengths, where M33 is known to have a sub-mm excess (Hermelo et al. 2016; Relaño et al. 2018). We therefore modified THEMIS from its default parameters, which is described in more detail in Appendix A. The main results of this modification are to use a dust mix with fewer very small carbon grains, which we might expect in a low-metallicity environment where the dust grains have less shielding from the interstellar radiation field (ISRF). The fit also modifies the silicate-to-carbon ratio. In the MW, this is ~10 but we find a mass ratio of ~0.3, very similar to the ratios found in the LMC and SMC by Chastenet et al. (2017). This would imply that either silicate grains are readily destroyed, or do not form in great numbers. This

is unlikely, and so more likely is that the silicate grains are not emissive enough in the current THEMIS model, as inferred from more recent laboratory studies (Anthony Jones, priv. comm.). Given more emissive silicate grains, a smaller mass of carbon grains would be required to explain the flatter sub-mm slope, and this ratio would be closer to that of the MW. It is, therefore, not necessarily a much higher mass of carbon grains that are required, but simply a higher mass of more emissive dust grains. The ratio of small-to-large grains is very similar to the MW, however (0.4 in the MW, 0.3 in our fitting). As shown in Fig. A1, the parameters of the THEMIS dust model can be adjusted to fit well in the sub-mm range. The results of this fitting confirm the hypothesis of Hermelo et al. (2016), who suggest different physical grain properties as the most plausable explanation for the observed sub-mm excess in M33. We use these recalculated abundances and size distributions in our SKIRT model, but within the RT simulation keep this dust mix constant throughout the entire galaxy.

4 MODEL FITTING

To find the best fit model, we ran a series of simulations with a variety of luminosities exploring the parameter space around our initial guesses (see Table 1). As these simulations are computationally expensive, and we can only explore the parameter space using a grid method, we fix all of the parameters apart from the various normalisations. For each of these parameters, we use 5 equally spaced values between our minima and maxima, for a total of 625 simulations.

We ran our simulations using a wavelength grid of 90 points, spaced for effective convolution with various filters and weighted depending on the importance of photons in that particular energy regime (Verstocken et al., in prep). We also use a small number of photons (10^6), to reduce the computational time for each model. Our dust grid is a binary tree dust grid (see Saftly et al. 2014 for more details on this grid method), and cells are no longer split when their mass fraction is less than 10^{-5}. This means that cells are not equal in size, with smaller cells in areas of higher density. In total, each of these "low-resolution" simulations takes around 3 CPU hours, and contain around 150,000 dust cells. To determine the best fit model, we defined six wavelength regimes – UV (GALEX), optical (SDSS), NIR (3.4 − 4.6 μm), MIR (5.8 − 24 μm), FIR (70 − 250 μm), and submm (250 − 850 μm). We calculated the reduced chi-squared, χ_ν^2, for each of these wavelength regimes, as there are an uneven number of points in each wavelength range. Our best fit is then the minimum of the sum of each of the χ_ν^2 values, including an extra 10% error in each of the points to account for uncertainties in the modelling, which is often used in other studies (e.g. Noll et al. 2009). Table 1 gives our best fit values for each free parameter. The likelihood of each model is given by $\mathcal{L} \propto \exp^{-\chi_\nu^2/2}$, and we sample from this distribution, quoting our errors as the 16[th] and 84[th] percentiles. When this error is smaller than the width of the bin, we instead quote the error as half the bin width. We find that given the low luminosity of the ionizing stars, we cannot well constrain this parameter, and it has a flat χ_ν^2 distribution across our parameter range, but our other parameters are reasonably well constrained. The normalised χ_ν^2 distributions of these free parameters are shown in Fig. 1.

Having found a best-fit model, we then re-simulated this model using a higher resolution wavelength grid with 553 points, and 2×10^7 photons to produce images with more reliable filter convolution and higher S/N. The cell maximum mass fraction is decreased by
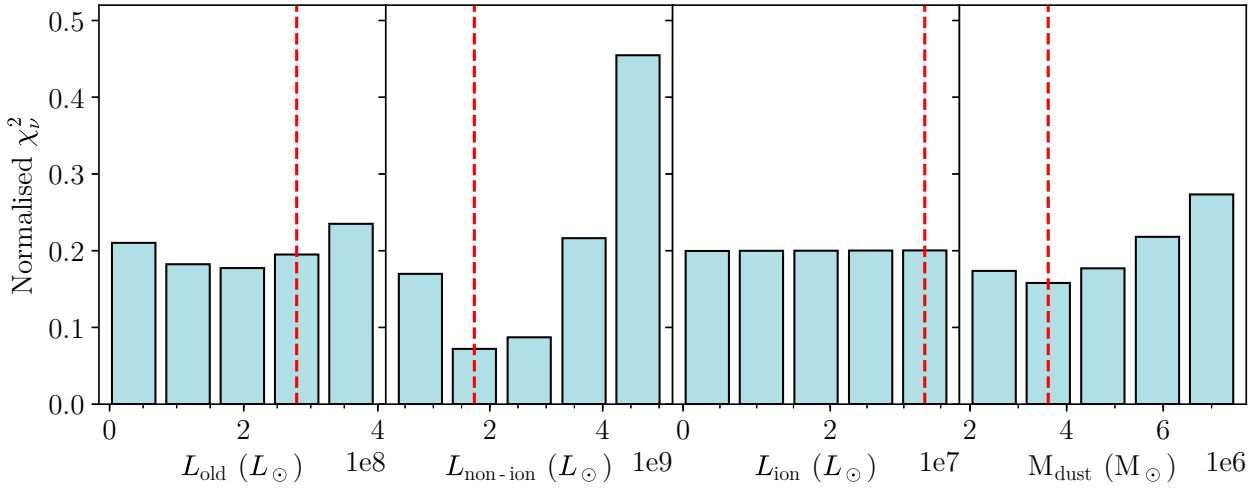
Figure 1. Normalised χ_ν^2 distributions of the four free parameters in our model. From left to right, these are the old stellar luminosity, the non-ionizing young stars, the ionizing stars and the dust mass. The distribution is normalised such that the total sum of the bars is equal to one, so that the four distributions are at approximately the same scale. The dashed red line shows the best fit parameter used in the high-resolution simulation.
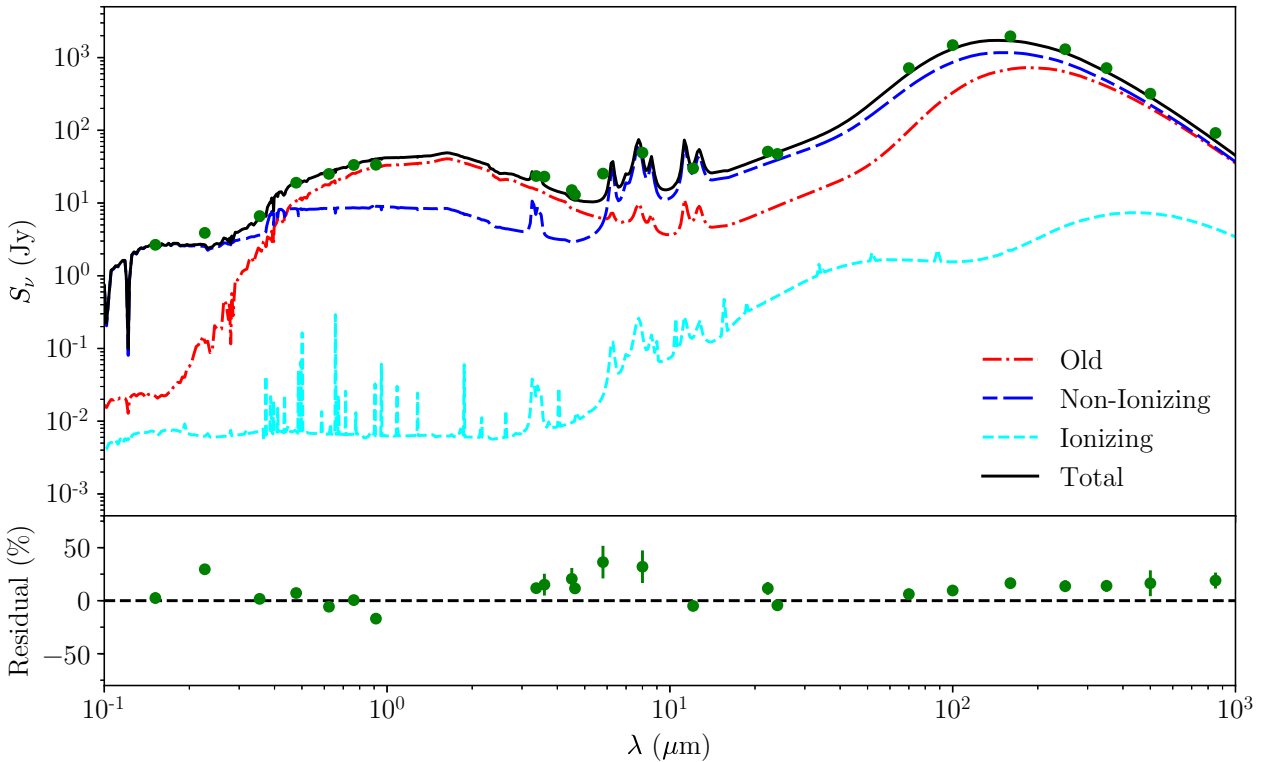

Figure 2. *Top:* High-resolution, best-fit RT simulation for M33. The red dot-dash line indicates the contribution from old stellar populations, the short-dash long-dash blue line young non-ionizing stars, light blue dashed line the ionizing population and the solid black line the total. Since the simulation considers the dust heating from the three stellar components simultaneously, this black line is not simply the sum of the three component lines. In this sense, the component decomposition should be taken as indicative only. *Bottom:* residuals for this fit.

a factor of 10, to 10^{-5}, which leads to around 1 million dust cells. These cell sizes vary from 37 pc^3 to 1500 pc^3, with an average size of 51 pc^3. With a maximum optical depth in a cell of 0.47, and an average V-band optical depth of 0.01, we can be confident that this grid is well sampled, even in regions of high optical depth. We also break the simulation up into its various stellar components to see

the relative contribution from each across our wavelength range. We use an instrument referred to as a FULLINSTRUMENT in SKIRT, which separates the total recorded flux, at every pixel and at every wavelength, in contributions due to direct stellar flux, scattered

Table 2. SFRs for M33, calculated using a variety of SFR tracers. In each case, we give the model SFR ($\mathrm{SFR_{model}}$), the SFR as calculated from the data ($\mathrm{SFR_{obs}}$) and references for the SFR prescription used. SFRs are in $M_\odot\,\mathrm{yr^{-1}}$.

Band(s)	$\mathrm{SFR_{model}}$	$\mathrm{SFR_{obs}}$	Reference
$24\,\mu\mathrm{m}$	0.11	0.10 ± 0.01	a
$70\,\mu\mathrm{m}$	0.14	0.15 ± 0.02	b
FUV+$24\,\mu\mathrm{m}$	$0.23^{+0.04}_{-0.02}$	$0.25^{+0.10}_{-0.07}$	c
TIR	0.23	0.17 ± 0.06^1	d, e

(a) Rieke et al. (2009), (b) Calzetti et al. (2010), (c) Leroy et al. (2008), (d) Hao et al. (2011), (e) Murphy et al. (2011). [1]Including an error of 30% to estimate uncertainty in IMF, dust attenuation; single temperature modified blackbody.

stellar flux, direct dust flux, and scattered dust flux, i.e.

$$S_\lambda^{\mathrm{tot}} = S_\lambda^{\star,\mathrm{dir}} + S_\lambda^{\star,\mathrm{sca}} + S_\lambda^{\mathrm{dust,dir}} + S_\lambda^{\mathrm{dust,sca}}. \tag{4}$$

This instrument also calculates S_λ^{tra}, i.e. the flux that would be obtained if the galaxy were completely dust-free. These simulations take around 600 CPU hours each.

5 RESULTS AND DISCUSSION

5.1 Global SED

The high-resolution, best-fit SED can be seen in Fig. 2. We also repeated the simulation for each stellar component individually, to illustrate the contribution from each of these components. As discussed in Sect. 5.4, this decomposition will not take into account the fact that the dust is simultaneously heated by each component. Thus, the sum of these lines will not be equal to the overall SED. We calculated the residuals by convolving the overall SED with the respective filter response for each waveband to produce a model flux, and then

$$\mathrm{residual} = \frac{\mathrm{observation} - \mathrm{model}}{\mathrm{observation}}. \tag{5}$$

In this sense, a negative residual means the model overestimates the observed data, and vice-versa. In the UV, the emission is dominated by light from the young stellar populations. In the optical and NIR, the emission is dominated by the old stellar populations. The dust emission is, in general, dominated by heating from the young stars, but is formed of a complex interplay of heating from the stellar components – a warmer component from heating due to the young stellar populations and a colder component from heating due to the older stars. The dust heating from the ionizing stars forms two distinct bumps, one from warmer dust heated from within the molecular clouds, and a cooler component from the emission of the more diffuse dust in the ISM surrounding these birth clouds.

We find a median absolute deviation (MAD) across all wavebands of 12%. We find that the NUV point is underestimated in the model (with a residual value of 30%). The RT model underestimating the NUV point is common across similar studies (see, e.g. De Looze et al. 2014; Mosenkov et al. 2016; Viaene et al. 2017b), and is likely caused by a UV attenuation bump that is too strong (see Sect. 5.3). We also find that many of the MIR points are underestimated. The MIR points are dominated by aromatic features, and so producing an adequate fit in this wavelength range is strongly dependent on the properties of the small carbon grains. Increasing

the weighting to these points can produce a better fit at these wavelength ranges, but a much poorer fit to the UV/optical points. Given the complex nature of this wavelength range, and our particular interest in the local dust-energy balance of M33, we prefer the current fit. Finally, the longer wavelengths are underestimated, potentially indicating a dust mass that is too low, or an incorrect dust emissivity. However, an increase in dust mass leads to increased dust attenuation and a poorer fit to the short wavelength points. Considering the uncertainty on the power-law slope for the small carbon grains (4.26 ± 0.13), and the fact that this has a large effect on the dust emissivity, the emissivity could well be underestimated. Given the fact the UV/optical and FIR/sub-mm points are given equal weight, this is the preferable fit. Due to recent work on the sub-mm excess (Hermelo et al. 2016; Relaño et al. 2018), we explore the 450 and $850\mu\mathrm{m}$ wavelengths in more detail in Sect. 5.2.1.

We calculated the SFR that this model produces using a variety of single- and multi-band SFR tracers ($24\,\mu\mathrm{m}$, $70\,\mu\mathrm{m}$, a combination of FUV+$24\,\mu\mathrm{m}$, and TIR luminosity), and compared these to the values calculated from the data. The results of this can be seen in Table 2. We see a good correspondence between the modelled and observed SFRs in M33. The TIR SFRs marginally agree within error, but we note a difference in the way these are calculated – for SKIRT, we integrate the emission from 3-1100 micron to calculate a TIR flux. In the case of the observed data, we fit a single-temperature modified blackbody (MBB) to the cold dust continuum emission, which will have negligible contributions at shorter wavelengths. As the longer wavelength regime is more affected by dust heating from older stellar populations, with a higher fraction of dust heating at shorter wavelengths from the young stellar populations (see Sect. 5.4), our model TIR luminosity is likely more representative of the TIR luminosity. We also highlight the importance of including the unattenuated starlight here – compared to the monochromatic $24\,\mu\mathrm{m}$ and $70\,\mu\mathrm{m}$ calculated SFR, the FUV+$24\,\mu\mathrm{m}$ SFR is nearly a factor of 2 higher. This is also true for our observed SFRs. The calculated SFR is consistently lower (by a factor of 2-3) than those calculated by Verley et al. (2009). The reason for this is twofold – firstly, they use SFR prescriptions similar to that of Kennicutt (1998). Our updated SFR measurements are generally around a factor of two lower (see Kennicutt & Evans 2012, their Table 1). Secondly, in truncating and masking the disc, we remove a significant amount of flux in the outer disc. The values given in Table 2, therefore, should be treated as a consistency check between pixels considered in the simulation and observations, and not as a measure of the true SFR of M33.

Alternatively, the SFR can also be calculated directly from the SED templates inputted into SKIRT for the young stellar populations, as these are scaled from a known normalisation factor. For the non-ionizing stars (i.e. the SFR over 100 Myr), this gives a value of $0.15\ M_\odot\,\mathrm{yr^{-1}}$, similar to the single-band SFR prescriptions in Table 2. For the ionizing stars (the SFR over the last 10 Myr), this produces an SFR of $0.58\ M_\odot\,\mathrm{yr^{-1}}$, but given that the ionizing stellar luminosity is not well constrained, this is unlikely to be a good measure of the true SFR.

From this model, we can calculate the fraction of the stellar radiation absorbed by dust, $\mathcal{F}_{\mathrm{abs}}$, which is given by

$$\mathcal{F}_{\mathrm{abs}} = \frac{L_{\mathrm{dust}}}{L_{\mathrm{dust}} + L_{\mathrm{stars}}}. \tag{6}$$

For the DustPedia galaxies, Bianchi et al. (2018) find this value to be 19% on average (25% if only considering late-type galaxies). For M33, we find this value to be 21%, in agreement with the findings of this earlier study.
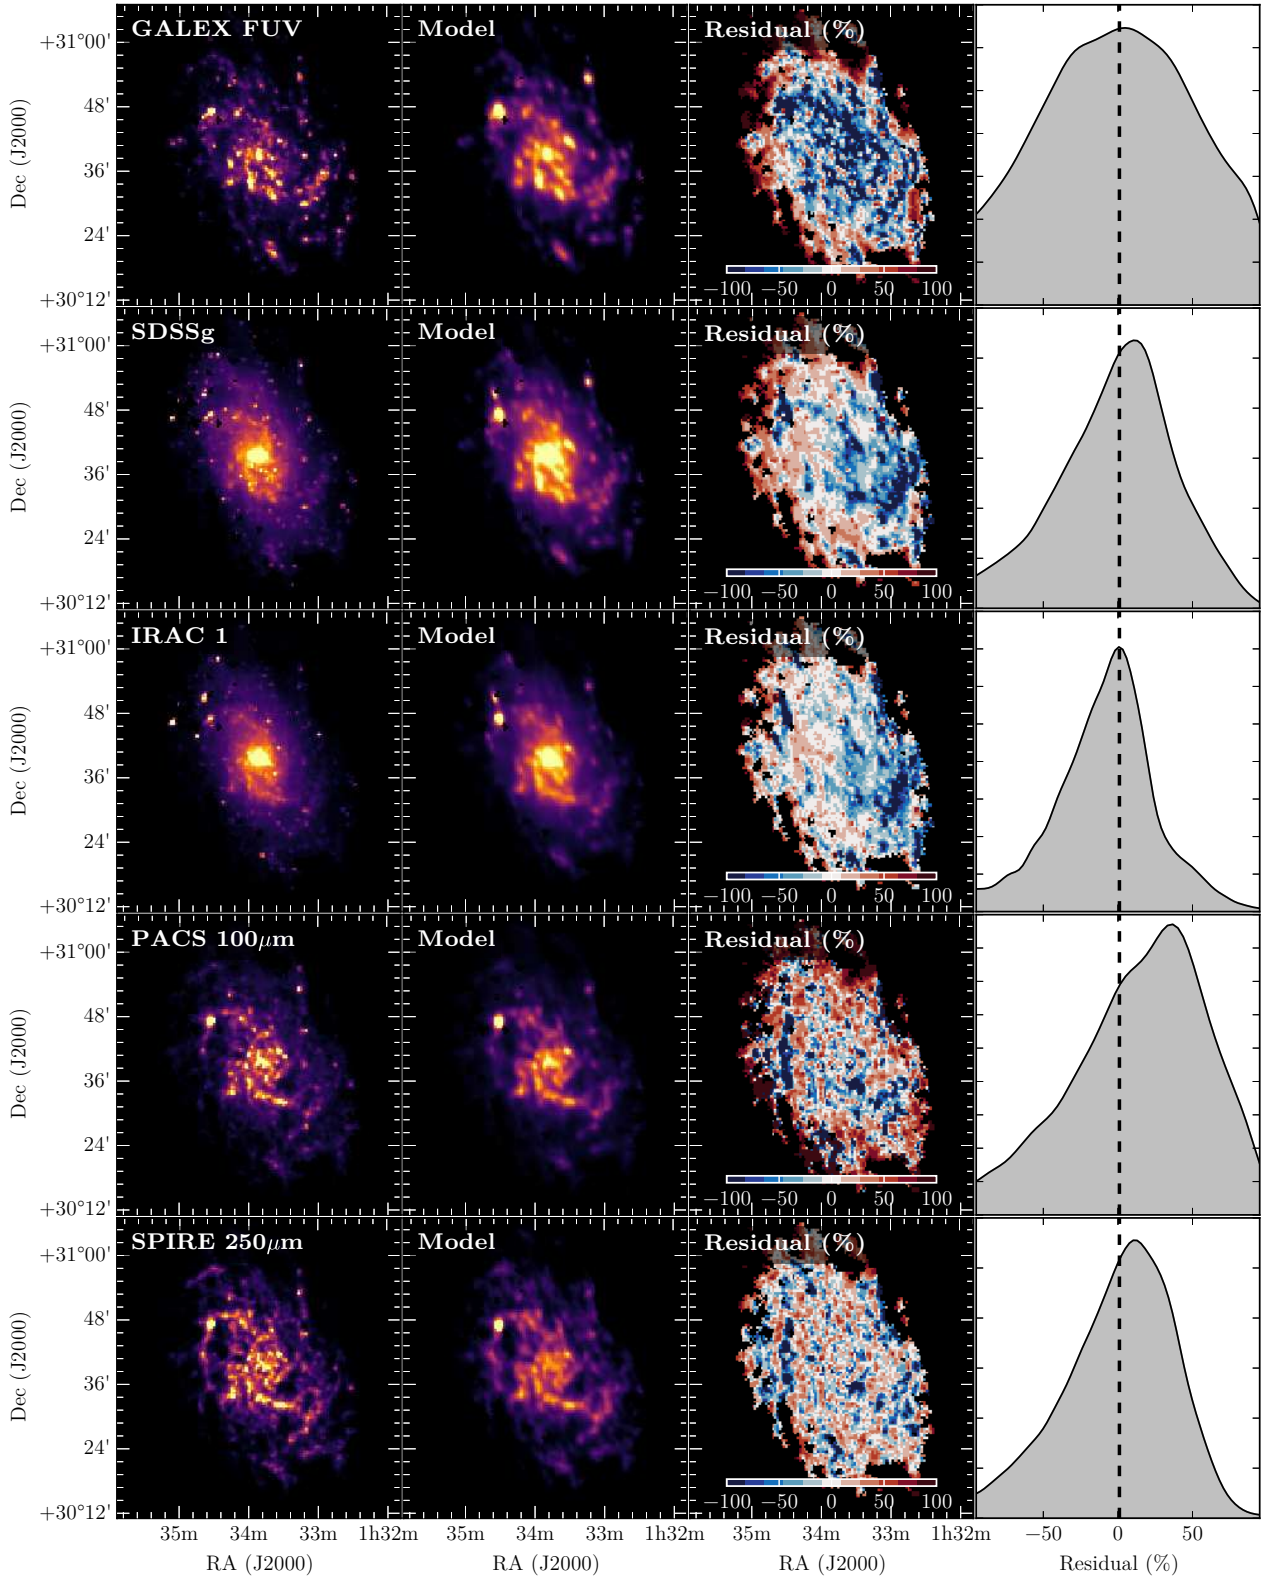
Figure 3. Residual plots of *first row:* GALEX FUV, *second row:* SDSS g-band, *third row:* IRAC 3.6μm, *fourth row:* PACS 100μm, *fifth row:* SPIRE 250μm wavebands. In each case, we show *first column:* the observed image, *second column:* the model image from our RT simulation, *third column:* the residuals, and *fourth column:* a KDE plot of these residuals.
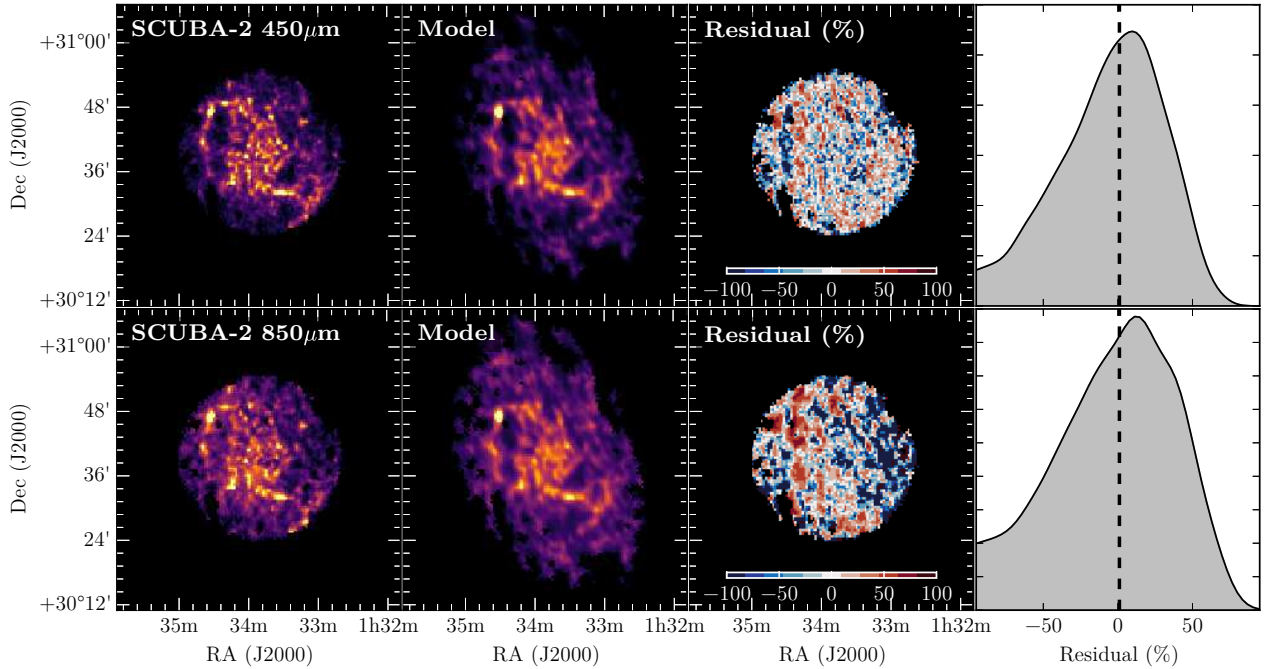
Figure 4. Residual plots of *top:* SCUBA-2 450µm and, *bottom:* SCUBA-2 850µm wavebands. We show *first column:* the observed image, *second column:* the model image from our RT simulation, *third column:* the residuals, and *fourth column:* a KDE plot of these residuals.

We find a dust mass of $(3.6 \pm 0.6) \times 10^6 \, \mathrm{M_\odot}$. This is very comparable to Williams et al. (2018), with a dust mass calculated from MAGPHYS of $4 \times 10^6 \, \mathrm{M_\odot}$, for a similarly good fit to this wavelength regime (see Fig. A1 and Fig. 2 of Williams et al. 2018). However, Hermelo et al. (2016) find a significantly higher dust mass, of around $1.7 \times 10^7 \, \mathrm{M_\odot}$ (albeit with a large uncertainty). However, as discussed in their work, this leads to a much lower gas-to-dust ratio (GDR) than expected in this low-metallicity environment, and so this dust mass estimate is likely too high, potentially due to grain properties or the dust attenuation assumed in their models. Assuming that the GDR scales with metallicity as Z^{-1} (e.g. Draine et al. 2007; Leroy et al. 2011), we would expect a value of the GDR between 200-450 (Hermelo et al. 2016). Using a total gas mass of $1.7 \times 10^9 \, \mathrm{M_\odot}$ (Gratier et al. 2010; Druard et al. 2014), we get a GDR of 470, significantly higher than Hermelo et al. (2016), but in agreement with the radial profiles of Relaño et al. (2018, their Fig. 10).

5.2 A Resolved Comparison of M33

SKIRT also produces a data cube which provides a 2D view of the galaxy at each wavelength in the wavelength grid. This means that we can compare the model on a spatial scale at any given wavelength. To this end, we produced residual images at a number of wavelengths. We note that due to images going through rotation and projection within the SKIRT routine, comparing these images directly may be an unfair comparison. This is discussed in more detail in Appendix B, but the effect of this on a moderately inclined galaxy like M33 is minor, and so we opted to compare directly to the original images. Given the resolution of our input geometries, we first spectrally convolve these datacubes with the relevant filter response, before spatially convolving with the point-spread function (PSF) of that waveband and regridding to pixels of 25 arcsec (using

MONTAGE) to make these images comparable. We also mask any pixels not considered in our simulation.

We find that across the 23 wavebands that form the high-resolution dataset for our simulation, we have a MAD of 33%, higher than the deviation seen in our global fluxes. Plots of the residuals at five wavelengths (FUV, SDSSg band, 3.6 µm, 100 µm, and 250 µm) are shown in Fig. 3. In general, our residuals are centred around 0 and most of the values lie within ±50% of the observed values. We see strong structure in many of our residuals, with the model often overestimating in the spiral arms and underestimating in the more diffuse ISM. In the regimes where we are observing mainly starlight, this is likely due to our temporal and spatial resolution. Whilst we assume 3 discrete, average ages, the actual star-formation history is much more complex, with stars of similar ages clumping together (e.g. Lewis et al. 2015), and these variations are on scales smaller than we are able to model in our simulation. At the wavelengths when we are dominated by dust emission, these spatial variations might indicate a variation in dust grain properties. Previous work has shown that there can be significant variation in the dust properties across a galaxy (e.g. Smith et al. 2012 in M31, Tabatabaei et al. 2014 in M33). Relaño et al. (2018) also suggest regional variations in dust properties to better explain the sub-mm excess. This work, however, has assumed an average dust grain property and mix throughout the whole of M33. We do not believe the features present in the residuals are an artefact of the use of MAGPHYS, as these features are also present in the study of De Looze et al. (2014), where the geometries are defined in a completely independent way to our analysis. There is also noise inherent both in the observations and the simulation, which makes a resolved comparison difficult. However, despite the simplicity of the model, the simulations well resemble the observations.
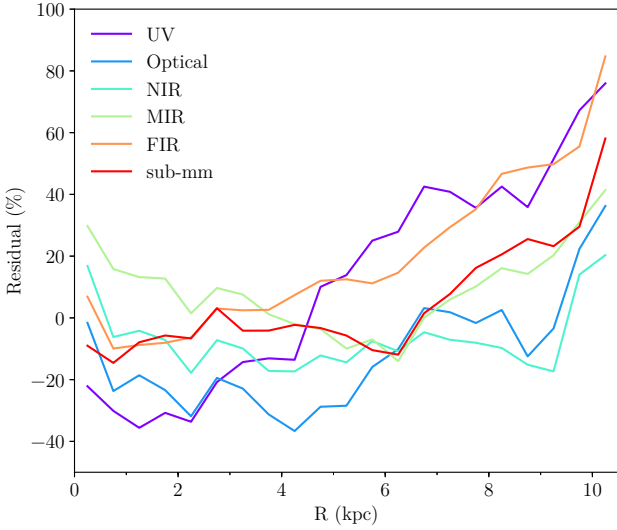
Figure 5. Median residual with galactocentric radius. Each differently coloured line indicates a wavelength regime as defined in Sect. 4.

5.2.1 The sub-mm excess

Given the sub-mm excess present in M33 (e.g. Hermelo et al. 2016; Relaño et al. 2018, i.e. that the observed fluxes are higher than the model), we have also produced residuals for the model at $450\,\mu m$, and $850\,\mu m$ wavelengths. These are compared to our SCUBA-2 images, and can be seen in Fig. 4. Unlike Hermelo et al. (2016), we find no significant sub-mm excess in our model (any higher than the excess we have at all long-wavelength points), and we also find no clear radial dependence on our residuals (consistent with those seen in any other wavelength regime), unlike that of Relaño et al. (2018). However, we note that in our earlier THEMIS fitting we modify the dust grain properties *specifically* to fit the sub-mm excess by removing many of the small carbon grains, and so the fact that we do not see this excess is not surprising.

5.2.2 Radial variation of the residuals

M33 has been shown to have a warped disc, both in the optical (e.g. Sandage & Humphreys 1980), and in its H_I 21cm line (e.g. Corbelli & Schneider 1997). If the inclination and position angle assumed in the deprojection of our model inputs has some radial variation, we would expect that to be seen as some radial dependence in the residuals. To quantify this, we calculate the median residual for each wavelength regime as defined in Sect. 4 for bins of $0.5\,kpc$ width in deprojected galactocentric radius (assuming a position angle of $22.5°$, and inclination of $56°$). The results of this are shown in Fig. 5. We see that in general, the model tends to underestimate at low galactocentric radius, and increasingly overestimate with increasing galactocentric radius. Given that the trends seen are broadly similar between each wavelength regime, we therefore conclude that M33 is similarly warped across all wavelengths considered in this study.

5.3 Dust Attenuation

5.3.1 Global Attenuation

By using the FullInstrument instrument class type in skirt, a view of the system with no dust attenuation is produced as a result
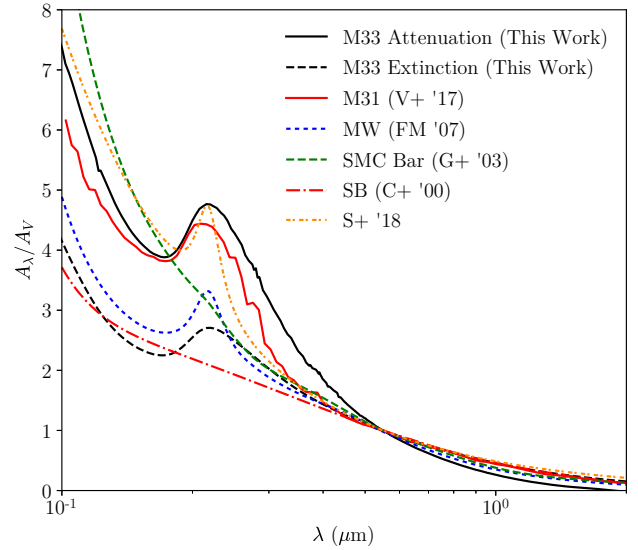

Figure 6. Dust attenuation for our RT model (solid black line) and underlying extinction curve (dashed black line), compared to extinction curves for the M31 (solid red line, Viaene et al. 2017b), the MW (short-dashed blue line; Fitzpatrick & Massa 2007), the SMC bar region (long-dashed green line; Gordon et al. 2003), the Calzetti et al. (2000) law for starburst galaxies (dot-dash red line), and the derived attenuation curve of Salim et al. (2018) for the stellar mass of M33 (short dash-dot orange line). All of these curves are normalised at V-band.

of the simulation. From this, we can directly calculate the dust attenuation at each wavelength, as

$$A_\lambda = 2.5\log_{10}\left(\frac{S_\lambda^{\mathrm{unatten}}}{S_\lambda^{\mathrm{atten}}}\right), \tag{7}$$

where $S_\lambda^{\mathrm{unatten}}$ and $S_\lambda^{\mathrm{atten}}$ are the total unattenuated and attenuated fluxes from the simulation at a given wavelength, respectively. One important caveat for the attenuation is that the MAPPINGS III SED does not truly provide a transparent view of the system with no dust, as the dust is built-in to this SED type. We make an estimate of the effect this will have on the flux by comparing a "transparent" MAPPINGS SED (i.e., a covering factor of 0) with the SED produced by our adopted covering factor of 0.1. This makes an average difference of 10% to the fluxes, which given the much lower luminosity of the ionizing stellar populations will make a negligible difference to our results. Along with calculating an attenuation curve, we also extract the intrinsic dust extinction curve, to see the effects of geometry and scattering on this attenuation curve. The results of this are shown in Fig. 6, and are compared to several literature extinction curves.

Visually, the underlying extinction curve lies somewhere between the MW (Fitzpatrick & Massa 2007) and starburst galaxies (Calzetti et al. 2000). We find a reasonably strong UV bump, which is broader than that of the MW. This is due to the adopted dust grain properties in the THEMIS model, which are discussed in more detail in Jones et al. (2013). This figure also highlights the effect of the treatment of geometry and dust scattering in shaping the difference between attenuation and extinction curves. These properties have previously been shown to have an important role in shaping attenuation curves (e.g. Granato et al. 2000; Witt & Gordon 2000; Baes & Dejonghe 2001; Panuzzo et al. 2007; Viaene et al. 2017a). Our M33 dust attenuation curve appears much more similar to the SMC bar region of Gordon et al. (2003), albeit with a strong UV
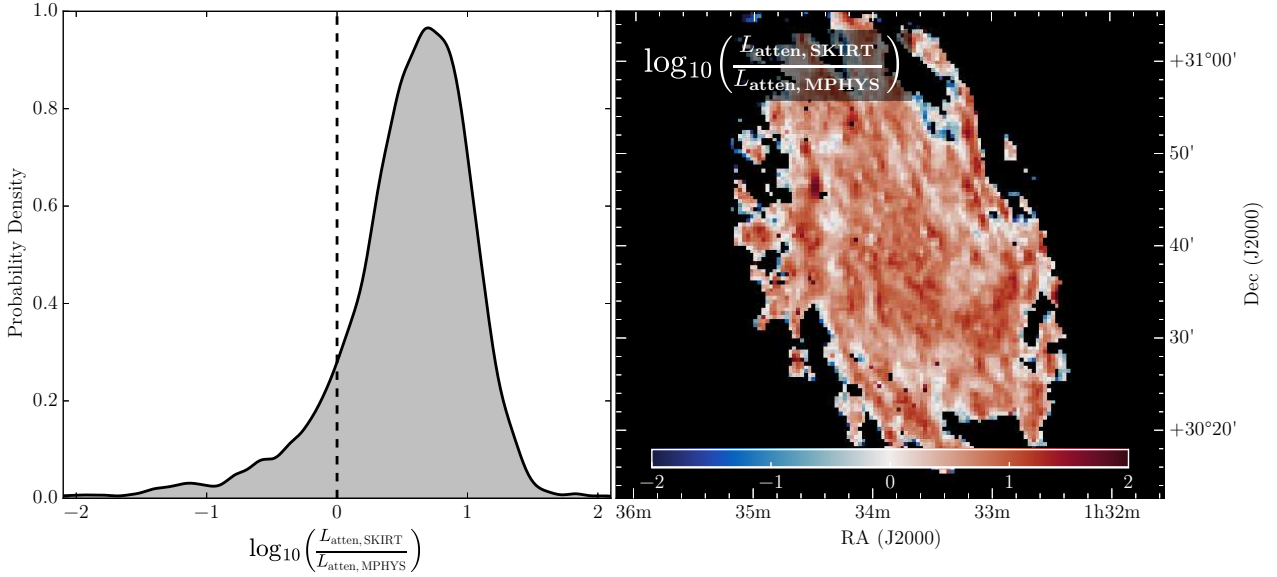
Figure 7. Comparison of SKIRT and MAGPHYS attenuated luminosities. *Left:* KDE plot of the ratio of the SKIRT attenuated luminosity to that of MAGPHYS. The black dashed line indicates where these two values are equal. *Right:* this same ratio, but instead plotted positionally.

bump. This is somewhat counter-intuitive, as our fitting technique alters the size distribution of the very small carbon grains, which are responsible for this bump. However, the larger of these small grains ($1\,\mathrm{nm} < r < 20\,\mathrm{nm}$) also contribute to this bump (Jones et al. 2013, their Table 1), so removing the smallest grains will not necessarily eliminate this UV bump. The attenuation curve is also very similar to the attenuation curve calculated by Viaene et al. (2016) for M31, derived in a similar way. This is somewhat surprising, given the very different intrinsic properties and geometry between M31 and M33. Given the stellar mass of M33 ($3\text{-}6\times10^9$ M$_\odot$, Corbelli 2003), the average dust attenuation curve for this stellar mass from the work of Salim et al. (2018) is very similar to that obtained in our simulations, although our UV bump is wider. An analysis of the nuclear region of M33 (Gordon et al. 1999) finds evidence of strong attenuation, along with a strong 2175Å bump. Our RT simulation shows that this may be the case across the whole of M33. However, we note that as the NUV flux is underestimated in our simulation, the strength of this NUV bump may be overestimated.

5.3.2 *Comparison to SED Modelling*

We can also compare, positionally, the amount of dust attenuation in our RT simulation to more traditional SED fitting. For our comparison, we take the pixel-by-pixel MAGPHYS fitting from Williams et al. (2018). MAGPHYS uses the dust attenuation model of Charlot & Fall (2000), and we refer readers to that work for details of the model. Essentially, this model assumes two populations of stars – one in their birth clouds, and others that have drifted away from these birth clouds. The light from both of these populations is subject to attenuation from dust in the ISM, and the stars in their birth clouds are additionally attenuated by the dusty clouds they reside within. This attenuation has a power-law type dependence on the wavelength, and the V-band optical depth is one of the parameters MAGPHYS fits, as well as the fraction of attenuation by dust in the ISM compared to birth clouds. The Charlot & Fall (2000) model, or variations of it, are typically used in panchromatic SED fitting tools. As MAGPHYS employs a dust-energy balance, the modelled dust lu-

minosity is by definition the attenuated luminosity. For SKIRT, in terms of its FULLINSTRUMENT output, the attenuated luminosity is

$$L^{\mathrm{atten}} = 4\pi D^2 \int \left(S_\lambda^{\mathrm{tra}} - S_\lambda^{\star,\mathrm{dir}} - S_\lambda^{\star,\mathrm{sca}} \right) d\lambda \qquad (8)$$

Given that this does not include the flux directly from the dust, this is not simply the transparent flux minus the total flux in the simulation. We calculate the ratio of the SKIRT to MAGPHYS attenuated luminosities, and show this in Fig. 7. We find a median offset of 0.56 dex for the SKIRT luminosity compared to the MAGPHYS attenuated luminosity, and a clear positional dependence in this offset, with much higher values for SKIRT in the spiral arms, and regions of more intense star formation, as compared to MAGPHYS. The reason for this may be two-fold – firstly, the pixel-by-pixel MAGPHYS fitting uses pixels of $100\,\mathrm{pc}^2$, where the local dust-energy balance may not hold (i.e. the amount of dust luminosity and attenuated luminosity may not be the same). With simulations of a galaxy, Smith & Hayward (2018) find acceptable fits to the V-band attenuation on scales of 0.2-25 kpc in ~99% of pixels modelled with MAGPHYS, so this is unlikely to explain the large discrepancy between these two attenuated luminosities. Secondly, the geometry can play an important role in affecting dust attenuation – given the positional dependence on the discrepancies between MAGPHYS and SKIRT, this is more likely to be the case.

5.3.3 *Face-On Optical Depth*

The optical depth of a galaxy is an important parameter to measure, as it quantifies the amount of stellar light that can be viewed directly without being obscured by dust. The question of whether galaxies are optically thin when viewed face-on is an outstanding problem, with some studies claiming the disc is nearly transparent (e.g. Xilouris et al. 1999), and some claiming that galaxies tend to be optically thick (e.g. Trewhella et al. 1997). Several works have attempted to answer this question through RT modelling of edge-on galaxies (De Geyter et al. 2014; Mosenkov et al. 2018, e.g.), but
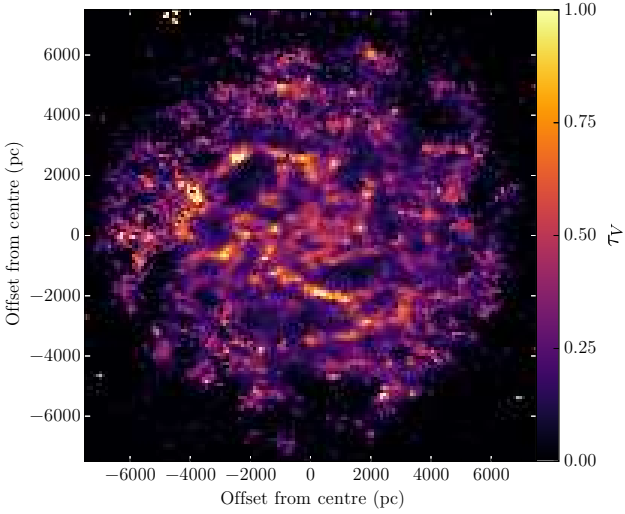
Figure 8. Face-on optical depth in V-band, derived from the RT simulation.

given degeneracies between the dust scale-length and the face-on optical depth, a reliable estimate of τ_V has been difficult to ascertain.

The optical depth is simply given by the dust column density integrated along the path length of a photon, and multiplied by the extinction coefficient, i.e.:

$$\tau_V = \kappa_V \int_0^\infty \rho_d(s)\, ds, \tag{9}$$

where κ_V for our dust mixture is calculated in the simulation to be $4625\,\mathrm{m^2\,kg^{-1}}$. Taking a deprojected column density map, this can then be trivially converted into a map of the optical depth, and we show this in Fig. 8. This map shows that the optical depth is highest in the spiral arms, and peaks in areas of active star formation. This peak can reach values >1, and thus these regions are optically thick. However, across the spiral arms the average optical depth is ∼0.3, in the interarm regions are ∼0.1, and the average V-band optical depth across the whole galaxy is ∼0.2. There is a gentle radial decline with galactocentric radius, from ∼0.5 in the centre to ∼0.2 at a radius of 5 kpc. This is well in agreement with Verley et al. (2009), who find a decrease in A_V with increasing galactocentric radius. We therefore conclude that at scales of 100pc, M33 is generally optically thin across its disc.

5.4 Dust Heating Mechanisms

From the skirt model, we can calculate the fraction of dust heating that comes from the young versus the old stars ($\mathcal{F}_{\mathrm{young}}$ and $\mathcal{F}_{\mathrm{old}}$, respectively), as simply the ratio of the total dust luminosity from the young stars divided by the total dust luminosity of the simulation. We find $\mathcal{F}_{\mathrm{young}}$ to be 72%, similar to the 63% found by De Looze et al. (2014) for M51, but significantly higher than the 9% found by Viaene et al. (2016) for M31. Globally, the dust heating of M33 is driven mainly by the young stellar populations.

We next turn to the contributions to the dust heating by the various stellar populations on a resolved level. Whilst $\mathcal{F}_{\mathrm{young}}$ + $\mathcal{F}_{\mathrm{old}}$ = 1 on a global SED scale, due to the intertwined nature of the radiation fields and the non-locality in wavelength of the dust heating, separating these quantities can only be approximated. We
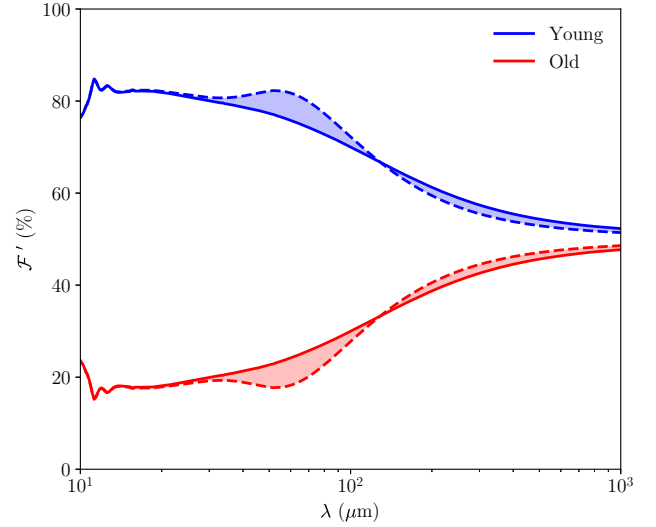

Figure 9. Dust heating fraction with wavelength from the young (blue) and old (red) stellar populations. In each case, the solid line corresponds to Eqs. 10 and 11, and the dashed line to Eqs. 12 and 13.

used the approximation of De Looze et al. (2014), where

$$\mathcal{F}'_{\lambda,\mathrm{young}} = \frac{1}{2}\frac{S_{\lambda,\mathrm{young}} + \left(S_{\lambda,\mathrm{total}} - S_{\lambda,\mathrm{old}}\right)}{S_{\lambda,\mathrm{total}}}, \tag{10}$$

and

$$\mathcal{F}'_{\lambda,\mathrm{old}} = \frac{1}{2}\frac{S_{\lambda,\mathrm{old}} + \left(S_{\lambda,\mathrm{total}} - S_{\lambda,\mathrm{young}}\right)}{S_{\lambda,\mathrm{total}}}. \tag{11}$$

Due to the non-local nature of the dust heating, $S_{\lambda,\mathrm{total}}$ is not simply $S_{\lambda,\mathrm{old}} + S_{\lambda,\mathrm{young}}$. We also included the naïve case where we do not include this non-locality, i.e.

$$\mathcal{F}_{\lambda,\mathrm{young}} = \frac{S_{\lambda,\mathrm{young}}}{S_{\lambda,\mathrm{young}} + S_{\lambda,\mathrm{old}}}, \tag{12}$$

and

$$\mathcal{F}_{\lambda,\mathrm{old}} = \frac{S_{\lambda,\mathrm{old}}}{S_{\lambda,\mathrm{young}} + S_{\lambda,\mathrm{old}}}. \tag{13}$$

The results of this are shown in Fig. 9. At all wavelengths, the dust heating is driven mainly by the young stellar populations, with a decreasing contribution towards longer wavelengths. It appears that contributions to the dust heating from the old stellar populations peak at colder dust temperatures, as they are heating the colder, more diffuse dust in the ISM (e.g. Bianchi 2008; Natale et al. 2015; Bendo et al. 2015).

It is also possible to investigate the fractional contribution to the dust heating from the young stellar populations on a resolved basis. Using Eq. 10, we calculate $F'_{\lambda,\mathrm{young}}$ across a number of wavebands, and the results of this can be seen in Fig. 10. It can be seen that at shorter wavelengths, there is a higher contribution to the dust heating from the young stellar populations in the spiral arms of M33, but this discrepancy decreases with increasing wavelength, to an almost uniform distribution at 850μm. This is similar to that observed by Viaene et al. (2017b), where the rings of M31 are clearly visible at shorter wavelengths.
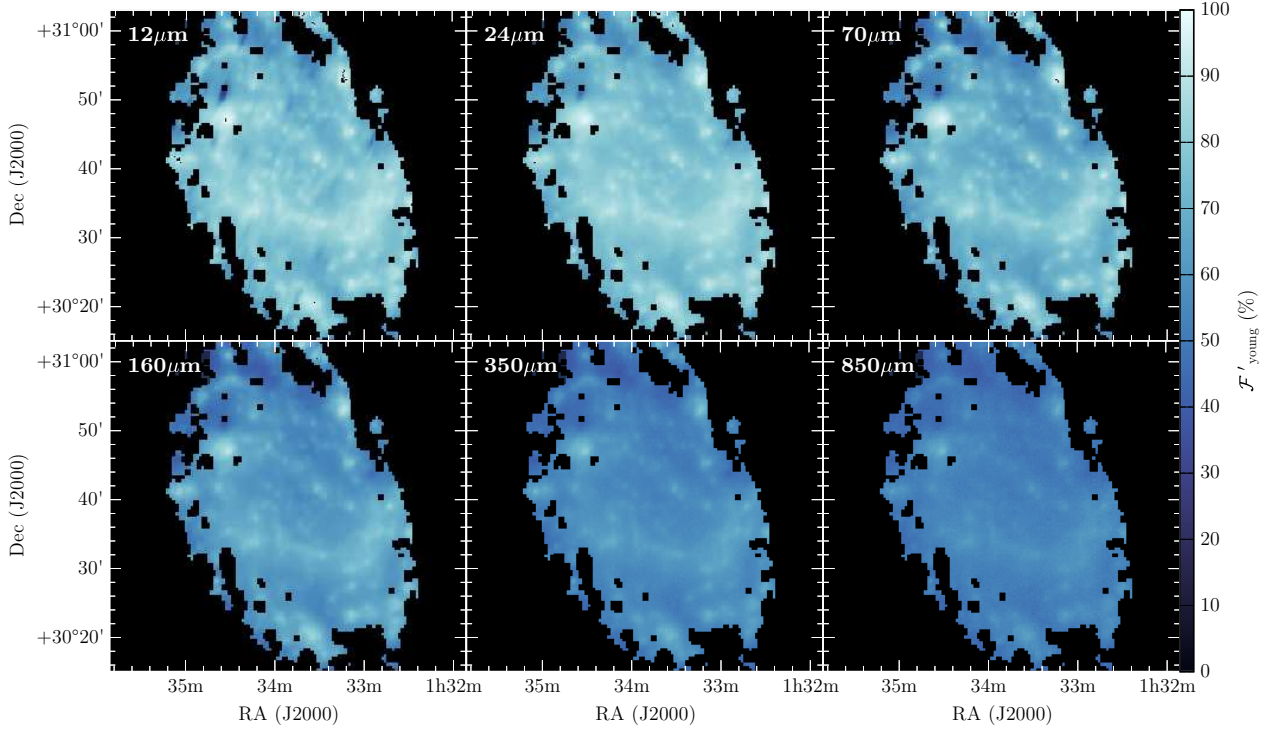
Figure 10. Dust heating fraction by young stellar populations ($F'_{\lambda,\,\text{young}}$) for a selection of wavebands. From left to right, *top row:* WISE 12 μm, MIPS 24 μm, and MIPS 70 μm. *Bottom row:* PACS 160 μm, SPIRE 250 μm, and SCUBA-2 850 μm.

5.5 Local Dust-Energy Balance

The scales at which the local dust-energy balance holds is vital for diagnosing the suitability of resolved measurements. In suitably small regions where more dust is heated by starlight originating from stars in neighbouring pixels than in the pixel being considered, traditional SED fitting tools may not recover a reliable value. We investigated the spatial scale at which the local dust-energy balance in our RT simulation becomes an acceptable assumption. This also gives an estimate of the average distance a photon travels within a galaxy. To this end, we define a "stellar luminosity excess,"

$$\delta^{\star} = \frac{L^{\text{atten}} - L^{\text{dust}}}{L^{\text{atten}}}, \qquad (14)$$

where L^{atten} is the total stellar luminosity attenuated (Eq. 8), and L^{dust} is the luminosity emitted by the dust. In terms of the SKIRT FULLINSTRUMENT output, this is

$$L^{\text{dust}} = 4\pi D^2 \int \left(S_\lambda^{\text{dust,dir}} - S_\lambda^{\text{dust,sca}} \right) d\lambda \qquad (15)$$

the integral of the direct flux from the dust. A value of 0 for $\delta^{\star}$ means the local dust-energy balance holds in that particular pixel, and increasingly positive (negative) values indicate more (less) flux attenuated than emitted by the dust in that pixel. Globally, the dust-energy balance should hold and therefore the mean of this distribution should be 0. We calculate this parameter for every 3D pixel in our data cube, and calculate the spread in these pixels, $\sigma_{\delta^{\star}}$, as the 84[th] percentile minus the 16[th] percentile. At the scale where the local dust-energy balance is a suitable assumption, $\sigma_{\delta^{\star}}$ should ideally be equal to 0. However, due to deviations between the model and observations, along with noise in the RT simulation, this is unlikely to be the case, so the point at which increasing the spatial scale causes no significant decrease in $\sigma_{\delta^{\star}}$ is the point at which we

assume the local dust-energy balance takes hold. To calculate these parameters for a variety of spatial scales, we regrid the simulation output to a number of scales, rather than re-run the simulation many times. The results of this procedure for a variety of spatial scales is shown in Fig. 11. From this, we can see that the local dust-energy balance is a suitable assumption at scales greater than ~1500 pc. This is in agreement with simulations, which show that the local dust-energy balance holds true at scales greater than around 1 kpc (Smith & Hayward 2018), as well as observational comparisons of SFR tracers (Boquien et al. 2015).

6 CONCLUSIONS

In this work, we have presented a high-resolution (100 pc) RT simulation of nearby spiral galaxy M33. Our simulation is self-consistent, includes the absorption and scattering effects of dust, and is performed in 3D. Our inputs for this simulation are based on images produced from a multi-wavelength dataset, in order to describe the young and older stellar populations in the galaxy. We also include a dust system, with the geometry informed by pixel-by-pixel SED fitting, and the dust properties from a global fit.

We find that we can well reproduce the SED of M33 to within a median deviation of 12%. This SED is dominated across almost all wavelengths by the young stellar populations (direct emission at shorter wavelengths, and dust heating at longer wavelengths), and in the optical is characterised by strong dust attenuation, with a very strong and broad UV bump. We find that we can well reproduce the SFRs given by our observed data, as well as the total dust mass. However, we find discrepancies at a resolved level, with many wavebands showing strong features in the residuals. We argue that these are due to limitations in our simple model of this galaxy – the
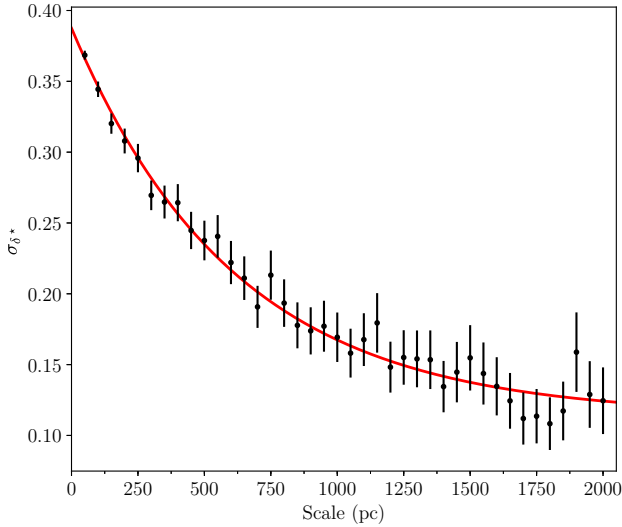
Figure 11. $\sigma_{\delta\star}$ for a variety of spatial scales. Errors are plotted assuming normal errors. The values flatten at a scale of around 1500pc, indicating this is the scale at which the local dust-energy balance holds true. The red line shows a fitted exponential, intended to guide the eye.

stellar and dust properties are not homogeneous across the disc of M33, and appear to be strikingly different in the spiral arms versus the diffuse ISM, and not due to our choices of input geometries. We are able to fit the sub-mm excess detected in previous works with a modified THEMIS dust model, showing that the excess is related to a difference in the grain properties of the dust, as suggested by Hermelo et al. (2016).

At a resolution of 100 pc, the galaxy is mostly transparent in the V-band, except in areas of high star-formation. This means that we should reliably be able to calculate the stellar properties in galaxies at least to these scales. We also find that the dust is heated almost solely by the young stellar populations, and so the TIR luminosity should be a reasonable tracer for star formation in this galaxy.

Finally, we estimate that the local dust-energy balance does not hold below scales of around 1500pc. This means that tools that employ this balance (e.g. MAGPHYS, CIGALE) are likely to be unsuitable at these high resolutions.

Despite the simple nature of this RT model, we find that we can broadly reproduce the characteristics of M33. Even given its simplicity, this type of RT modelling allows insights into the sub-kpc properties of galaxies that traditional SED fitting does not, and allows us to probe the complex interplay of starlight and dust in galaxies self-consistently at these small spatial scales.

REFERENCES

Alam S., et al., 2015, ApJS, 219, 12
Astropy Collaboration et al., 2013, A&A, 558, A33
Astropy Collaboration et al., 2018, AJ, 156, 123
Baes M., Camps P., 2015, Astronomy and Computing, 12, 33
Baes M., Dejonghe H., 2001, MNRAS, 326, 722
Baes M., et al., 2003, MNRAS, 343, 1081
Baes M., et al., 2010, A&A, 518, L39
Bendo G. J., et al., 2015, MNRAS, 448, 135
Bianchi S., 2008, A&A, 490, 461
Bianchi S., et al., 2018, A&A, 620, A112
Boquien M., et al., 2015, A&A, 578, A8
Bot C., Ysard N., Paradis D., Bernard J. P., Lagache G., Israel F. P., Wall W. F., 2010, A&A, 523, A20
Bothun G. D., 1992, AJ, 103, 104
Bresolin F., 2011, ApJ, 730, 129
Bruzual G., Charlot S., 2003, MNRAS, 344, 1000
Calzetti D., Armus L., Bohlin R. C., Kinney A. L., Koornneef J., Storchi-Bergmann T., 2000, ApJ, 533, 682
Calzetti D., et al., 2007, ApJ, 666, 870
Calzetti D., et al., 2010, ApJ, 714, 1256
Camps P., Baes M., 2015, Astronomy and Computing, 9, 20
Camps P., et al., 2015, A&A, 580, A87
Charlot S., Fall S. M., 2000, ApJ, 539, 718
Chastenet J., Bot C., Gordon K. D., Bocchio M., Roman-Duval J., Jones A. P., Ysard N., 2017, A&A, 601, A55
Chevallard J., Charlot S., 2016, MNRAS, 462, 1415
Combes F., et al., 2012, A&A, 539, A67
Compiègne M., et al., 2011, A&A, 525, A103
Corbelli E., 2003, MNRAS, 342, 199
Corbelli E., Schneider S. E., 1997, ApJ, 479, 244
Dale D. A., et al., 2009, ApJ, 703, 517
Davies J. I., et al., 2017, Publications of the Astronomical Society of the Pacific, 129, 044102
De Geyter G., Baes M., Fritz J., Camps P., 2013, A&A, 550, A74
De Geyter G., Baes M., Camps P., Fritz J., De Looze I., Hughes T. M., Viaene S., Gentile G., 2014, MNRAS, 441, 869
De Geyter G., et al., 2015, MNRAS, 451, 1728
De Looze I., Baes M., Fritz J., Verstappen J., 2012a, MNRAS, 419, 895
De Looze I., et al., 2012b, MNRAS, 427, 2797
De Looze I., et al., 2014, A&A, 571, A69
Draine B. T., et al., 2007, ApJ, 663, 866
Druard C., et al., 2014, A&A, 567, A118
Fazio G. G., et al., 2004, ApJS, 154, 10
Fitzpatrick E. L., Massa D., 2007, ApJ, 663, 320
Galametz M., Madden S. C., Galliano F., Hony S., Bendo G. J., Sauvage M., 2011, A&A, 532, A56

Gardan E., Braine J., Schuster K. F., Brouillet N., Sievers A., 2007, A&A, 473, 91

Gordon K. D., Hanson M. M., Clayton G. C., Rieke G. H., Misselt K. A., 1999, ApJ, 519, 165

Gordon K. D., Clayton G. C., Misselt K. A., Landolt A. U., Wolff M. J., 2003, ApJ, 594, 279

Granato G. L., Lacey C. G., Silva L., Bressan A., Baugh C. M., Cole S., Frenk C. S., 2000, ApJ, 542, 710

Gratier P., et al., 2010, A&A, 522, A3

Griffin M. J., et al., 2010, A&A, 518, L3

Groves B., Dopita M. A., Sutherland R. S., Kewley L. J., Fischera J., Leitherer C., Brandl B., van Breugel W., 2008, ApJS, 176, 438

Hao C.-N., Kennicutt R. C., Johnson B. D., Calzetti D., Dale D. A., Moustakas J., 2011, ApJ, 741, 124

Hermelo I., et al., 2016, A&A, 590, A56

Heyer M. H., Corbelli E., Schneider S. E., Young J. S., 2004, ApJ, 602, 723

Holland W. S., et al., 2013, MNRAS, 430, 2513

Hoopes C. G., Walterbos R. A. M., 2000, ApJ, 541, 597

Hunter J. D., 2007, Computing In Science & Engineering, 9, 90

Jones A. P., Fanciullo L., Köhler M., Verstraete L., Guillet V., Bocchio M., Ysard N., 2013, A&A, 558, A62

Kam Z. S., Carignan C., Chemin L., Amram P., Epinat B., 2015, MNRAS, 449, 4048

Kennicutt Jr. R. C., 1998, ARA&A, 36, 189

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

Kirkpatrick A., et al., 2013, ApJ, 778, 51

Köhler M., Jones A., Ysard N., 2014, A&A, 565, L9

Kramer C., et al., 2010, A&A, 518, L67

Leroy A. K., Walter F., Brinks E., Bigiel F., de Blok W. J. G., Madore B., Thornley M. D., 2008, AJ, 136, 2782

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

Lewis A. R., et al., 2015, ApJ, 805, 183

Madore B. F., Freedman W. L., 1991, PASP, 103, 933

Mainzer A., et al., 2011, ApJ, 731, 53

Martin D. C., et al., 2005, ApJ, 619, L1

Meurer G. R., Heckman T. M., Calzetti D., 1999, ApJ, 521, 64

Misiriotis A., Popescu C. C., Tuffs R., Kylafis N. D., 2001, A&A, 372, 775

Mosenkov A. V., et al., 2016, A&A, 592, A71

Mosenkov A. V., et al., 2018, A&A, 616, A120

Murphy E. J., et al., 2011, ApJ, 737, 67

Natale G., Popescu C. C., Tuffs R. J., Debattista V. P., Fischera J., Grootes M. W., 2015, MNRAS, 449, 243

Noll S., Burgarella D., Giovannoli E., Buat V., Marcillac D., Muñoz-Mateos J. C., 2009, A&A, 507, 1793

Panuzzo P., Granato G. L., Buat V., Inoue A. K., Silva L., Iglesias-Páramo J., Bressan A., 2007, MNRAS, 375, 640

Paturel G., Theureau G., Bottinelli L., Gouguenheim L., Coudreau-Durand N., Hallet N., Petit C., 2003, A&A, 412, 57

Pilbratt G. L., et al., 2010, A&A, 518, L1

Poglitsch A., et al., 2010, A&A, 518, L2

Popescu C. C., Tuffs R. J., 2002, MNRAS, 335, L41

Popescu C. C., Misiriotis A., Kylafis N. D., Tuffs R. J., Fischera J., 2000, A&A, 362, 138

Popescu C. C., Tuffs R. J., Dopita M. A., Fischera J., Kylafis N. D., Madore B. F., 2011, A&A, 527, A109

Popescu C. C., Yang R., Tuffs R. J., Natale G., Rushton M., Aharonian F., 2017, MNRAS, 470, 2539

Regan M. W., Vogel S. N., 1994, ApJ, 434, 536

Relaño M., et al., 2018, A&A, 613, A43

Rémy-Ruyer A., et al., 2013, A&A, 557, A95

Rieke G. H., et al., 2004, ApJS, 154, 25

Rieke G. H., Alonso-Herrero A., Weiner B. J., Pérez-González P. G., Blaylock M., Donley J. L., Marcillac D., 2009, ApJ, 692, 556

Robitaille T., Bressert E., 2012, APLpy: Astronomical Plotting Library in Python, Astrophysics Source Code Library (ascl:1208.017)

Rosolowsky E., Simon J. D., 2008a, ApJ, 675, 1213

Rosolowsky E., Simon J. D., 2008b, ApJ, 675, 1213

Saftly W., Baes M., Camps P., 2014, A&A, 561, A77

Salim S., Boquien M., Lee J. C., 2018, ApJ, 859, 11

Sandage A., Humphreys R. M., 1980, ApJ, 236, L1

Schlafly E. F., Finkbeiner D. P., 2011, ApJ, 737, 103

Smith D. J. B., Hayward C. C., 2018, MNRAS, 476, 1705

Smith M. W. L., et al., 2012, ApJ, 756, 40

Steinacker J., Baes M., Gordon K. D., 2013, ARA&A, 51, 63

Tabatabaei F. S., et al., 2014, A&A, 561, A95

Thilker D. A., et al., 2005, ApJ, 619, L67

Trewhella M., Davies J. I., Disney M. J., Jones H. G. W., 1997, MNRAS, 288, 397

Verley S., Hunt L. K., Corbelli E., Giovanardi C., 2007, A&A, 476, 1161

Verley S., Corbelli E., Giovanardi C., Hunt L. K., 2009, A&A, 493, 453

Verstocken S., Van De Putte D., Camps P., Baes M., 2017, Astronomy and Computing, 20, 16

Viaene S., et al., 2014, A&A, 567, A71

Viaene S., et al., 2016, A&A, 586, A13

Viaene S., Sarzi M., Baes M., Fritz J., Puerari I., 2017a, MNRAS, 472, 1286

Viaene S., et al., 2017b, A&A, 599, A64

Williams T. G., Gear W. K., Smith M. W. L., 2018, MNRAS, 479, 297

Williams T. G., Gear W. K., Smith M. W. L., 2019, MNRAS, 483, 5135

Witt A. N., Gordon K. D., 2000, ApJ, 528, 799

Wright E. L., et al., 2010, AJ, 140, 1868

Xilouris E. M., Byun Y. I., Kylafis N. D., Paleologou E. V., Papamastorakis J., 1999, A&A, 344, 868

York D. G., et al., 2000, AJ, 120, 1579

Zhu Y.-N., Wu H., Li H.-N., Cao C., 2010, Research in Astronomy and Astrophysics, 10, 329

da Cunha E., Charlot S., Elbaz D., 2008, MNRAS, 388, 1595

de Vaucouleurs G., 1959, ApJ, 130, 728

van der Walt S., Colbert S. C., Varoquaux G., 2011, CoRR, abs/1102.1523

APPENDIX A: MODIFYING THE THEMIS DUST MODEL

In order to better fit the THEMIS model to our data, we fit a dust emission model to the points longward of $3.4\,\mu m$ using a similar method to that of Chastenet et al. (2017). The components of this fit are the small carbon grains (sCM20), large carbon grains (lCM20), and silicates, consisting of olivines and pyroxines, tied together as aSilM5. Due to a non-negligible contribution from stars at NIR wavelengths, we also include a blackbody at 5000 K, to approximate this contribution.

Our initial fit kept the mass ratios of these various components fixed at the values calculated for the diffuse dust of the MW, and so there are only three free parameters in the model – the strength of the ISRF, a scaling factor for the stellar contribution and the overall dust mass. We generate a grid of ISRFs from $10^{-1} \le U \le 10^{3.5}$ (with 1 being the value for the local neighbourhood), equally spaced in steps of 0.01 in log space. The SEDs for this are generated using DustEM (Compiègne et al. 2011). We then fitted these three free parameters using an MCMC framework using emcee[3]. We use 500 walkers each taking 500 steps using the first half of these steps as "burn-in", and our initial guesses for the ISRF is that of the MW, the stellar scaling factor the $3.6\,\mu m$ point, and the dust mass by the $250\,\mu m$ point. We account for correlated uncertainties between bands, and use the filter responses for each waveband to calculate the flux as seen by that particular instrument. The fit and residuals for this can be seen in Fig. A1, and we find that the default THEMIS parameters consistently underestimate the bulk of the cold dust points.

Next, we performed a fit where we allowed the abundances of the amorphous hydrocarbons and silicates to vary (although we lock the abundances of the two silicate populations together). This
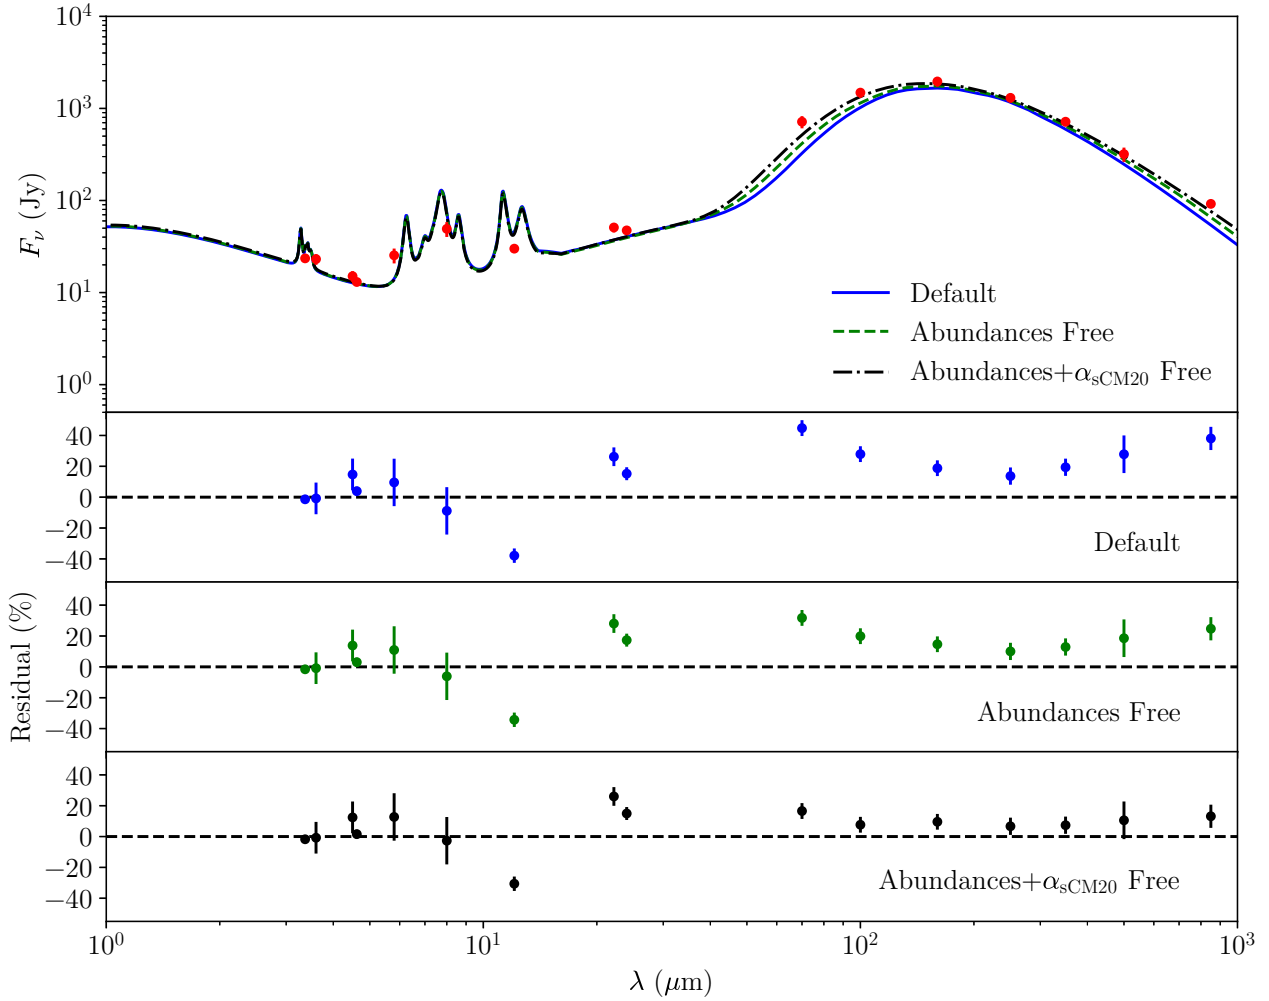
---

[3] http://dfm.io/emcee/current/

Figure A1. *Top:* Various model THEMIS fits to the dust SED of M33. The solid blue line indicates the default THEMIS parameters, the dashed green line where we allow mass abundances to vary, and the black dot-dash line where we additionally vary the size distribution of small carbon grains. The residuals for each of these fits are given in the subsequent panels.

increases our number of free parameters to 5, where compared to the total dust mass we now have the individual masses of the small and large carbon grains, and the silicates. This fit is also shown in Fig. A1, and we find that while it performs slightly better than the default parameters, the fit is still poor across the FIR/sub-mm range.

Finally, we allowed variation in the small grain size distribution. The size distribution of small amorphous hydrocarbons is given by a power-law, partly defined by $dn/da \propto a^{-\alpha_{\rm sCM20}}$, and we allow this value of $\alpha_{\rm sCM20}$ to vary. For this, we calculated a grid of $2.6 \leq \alpha_{\rm sCM20} \leq 5.4$ (where 5 is the THEMIS default) in steps of 0.01, for each value of the ISRF strength defined earlier (leading to a total grid of some 100,000 combination of parameters). The inclusion of fitting $\alpha_{\rm sCM20}$ brings our total number of free parameters to 6, and the best fit is shown, again, in Fig. A1. We also show the corner plot of this fit in Fig. A2. We find a median $\alpha_{\rm sCM20}$ of around 4.3, somewhat lower than the THEMIS default of 5. In terms of the SED, this leads to a flatter slope at longer wavelengths. Physically, this corresponds to fewer very small carbon grains, as we might expect in a lower-metallicity environment such as M33. Much like the work of Chastenet et al. (2017) on the LMC and SMC, we find a much lower value for the silicate/carbon ratio of ~0.3, compared

to the MW value of ~10. However, the ratio of small-to-large grains is very similar to the MW value of 0.4, with a value of 0.3.

Finally, we note that the fit does not perform so well in the 24-70 μm range. This can be improved by adding a second, warmer dust component (i.e. a higher ISRF strength). This produces a better fit in these wavelength ranges, but does not change the dust component masses, or $\alpha_{\rm sCM20}$ significantly. As we are only performing this fit to calculate the dust grain properties, and leave SKIRT to model the ISRF, we only show the single-temperature component fit here.

APPENDIX B: ROTATING AND PROJECTING THE DATA

To add a 3D scale to the provided images, SKIRT deprojects and de-rotates the input image, given an inclination and position angle. This means that the image becomes "smeared" as it is transformed into the plane of the galaxy, and then back into the observer frame. To test the effect that this has on our images (particularly for the purposes of residuals), we de-rotated and deprojected M33 (using a PA of 22.5° and inclination of 56°), before rotating and projecting it back into its original frame. The result of this can be seen in Fig. B1. The effects of this routine are minor and will not affect our residual
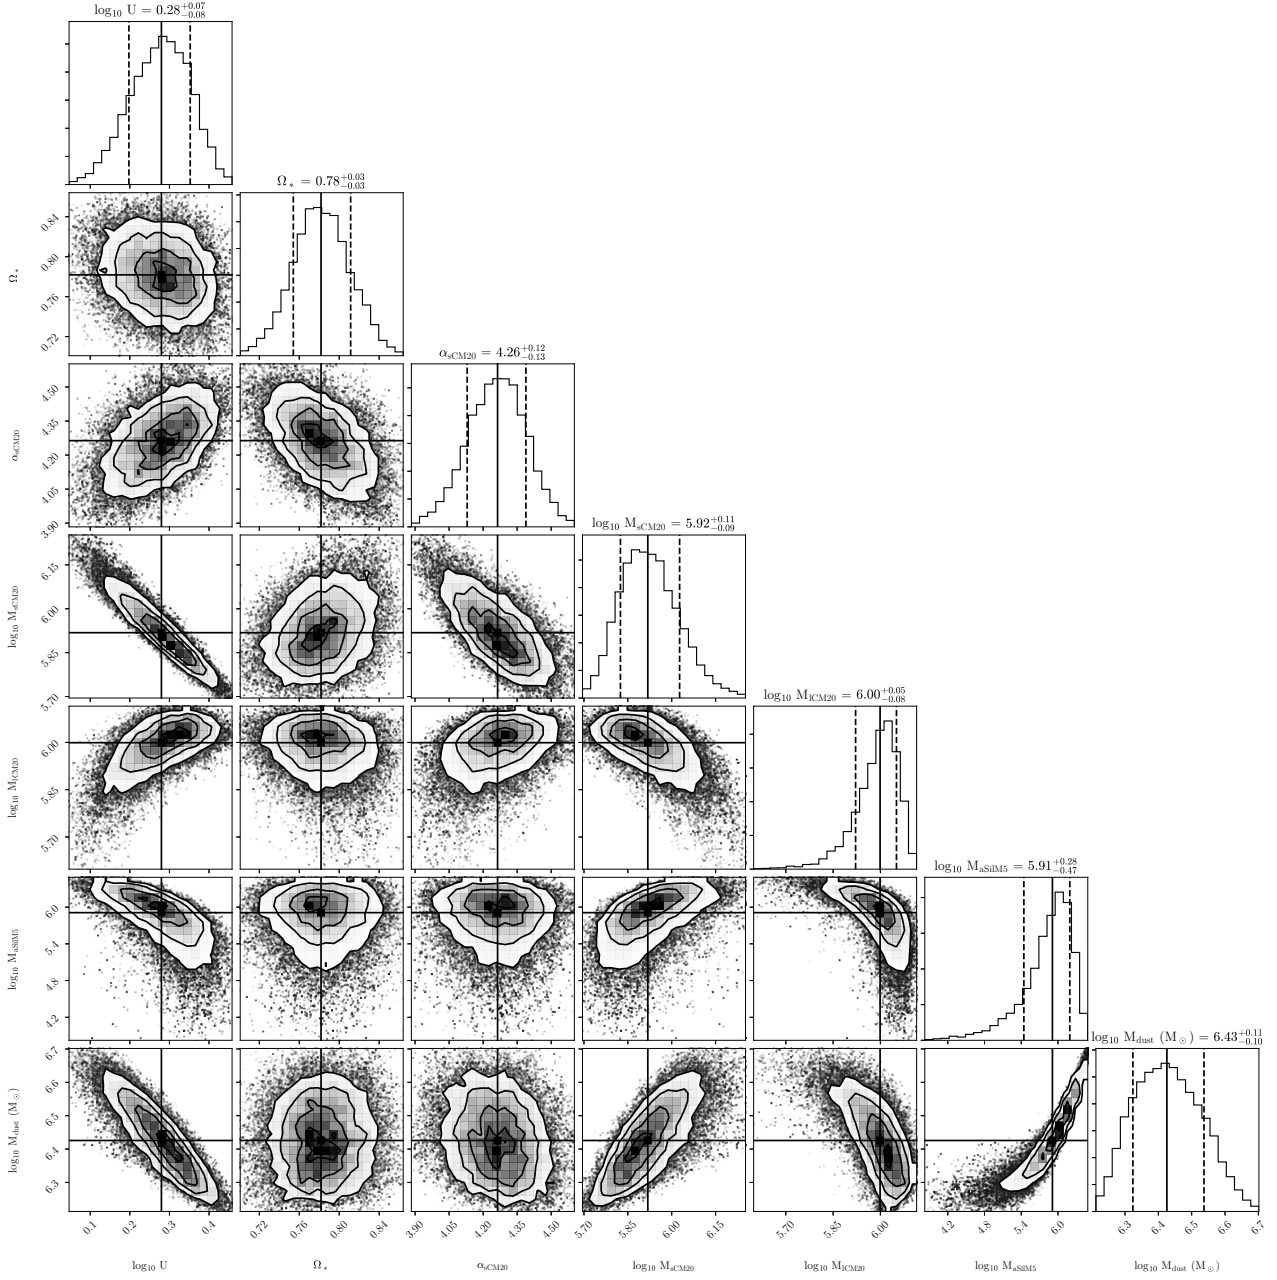
Figure A2. Corner plot for the THEMIS dust fit with free dust grain abundances and small grain size distribution. From left to right, the panels show the ISRF strength U, the stellar scaling factor (with respect to the $3.6\,\mu$m flux), the small grain power-law slope, the mass for the small carbon grains, large carbon grains and silicates, and finally the total dust mass. The solid black line in each histogram shows the median value, with the dashed lines showing the 16th and 84th percentiles.

plots in any significant way, so we opt to use the original images as-is.

This paper has been typeset from a TeX/LaTeX file prepared by the author.
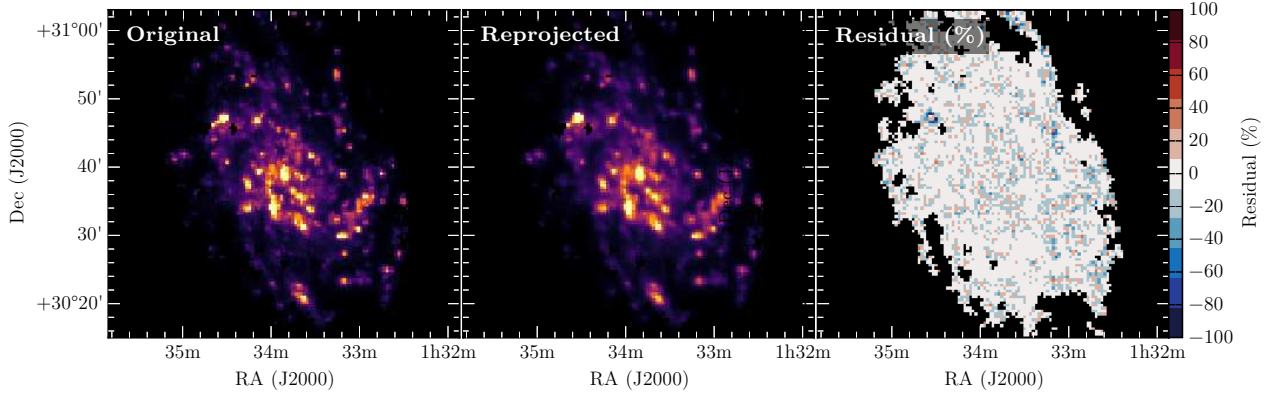
Figure B1. The effect of deprojection and derotation on an input image. From left to right, we show the original GALEX FUV image, the image after derotation, deprojection, reprojection, and rotation, and the corresponding residuals of these two maps.