High-resolution, 3D radiative transfer modelling

IV. AGN-powered dust heating in NGC 1068

S. Viaene[1,2], A. Nersesian[3,4,1], J. Fritz[5], S. Verstocken[1], M. Baes[1], S. Bianchi[6], V. Casasola[7,6], L. Cassarà[8,3],
C. Clark[9], J. Davies[10], I. De Looze[1,11], P. De Vis[10], W. Dobbels[1], M. Galametz[12], F. Galliano[12], A. Jones[13],
S. Madden[12], A. Mosenkov[14,15], A. Trcka[1], E. M. Xilouris[3], and N. Ysard[13]

[1] Sterrenkundig Observatorium, Universiteit Gent, Krijgslaan 281, 9000 Gent, Belgium
    e-mail: sebastien.viaene@ugent.be
[2] Centre for Astrophysics Research, University of Hertfordshire, College Lane, Hatfield AL10 9AB, UK
[3] National Observatory of Athens, Institute for Astronomy, Astrophysics, Space Applications and Remote Sensing, Ioannou Metaxa
    and Vasileos Pavlou GR-15236, Athens, Greece
[4] Department of Astrophysics, Astronomy & Mechanics, Faculty of Physics, University of Athens, Panepistimiopolis, GR15784
    Zografos, Athens, Greece
[5] Instituto de Radioastronomía y Astrofísica, UNAM, Campus Morelia, A.P. 3-72, 58089 Morelia, Mexico
[6] INAF – Osservatorio Astrofisico di Arcetri, Largo E. Fermi 5, 50125 Firenze, Italy
[7] INAF – Istituto di Radioastronomia, Via P. Gobetti 101, 40129 Bologna, Italy
[8] INAF – Istituto di Astrofisica Spaziale e Fisica cosmica, Via A. Corti 12, 20133 Milano, Italy
[9] Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, Maryland 21218, USA
[10] School of Physics and Astronomy, Cardiff University, The Parade, Cardiff CF24 3AA, UK
[11] Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK
[12] Laboratoire AIM, CEA/DSM – CNRS – Université Paris Diderot, IRFU/Service d'Astrophysique, CEA Saclay, 91191
    Gif-sur-Yvette, France
[13] Institut d'Astrophysique Spatiale, CNRS, Univ. Paris-Sud, Université Paris-Saclay, Bât. 121, 91405 Orsay Cedex, France
[14] Central Astronomical Observatory of RAS, Pulkovskoye Chaussee 65/1, 196140 St. Petersburg, Russia
[15] St. Petersburg State University, Universitetskij Pr. 28, 198504 St. Petersburg, Stary Peterhof, Russia

ABSTRACT

The star formation rate and the mass of interstellar medium (ISM) have a high predictive power for the future evolution of a galaxy. Nevertheless, deriving such properties is not straightforward. Dust emission, an important diagnostic of star formation and ISM mass throughout the Universe, can be powered by sources unrelated to ongoing star formation. In the framework of the DustPedia project we set out to disentangle the radiation of the ongoing star formation from that of the older stellar populations. This is done through detailed 3D radiative transfer simulations of face-on spiral galaxies. We take special care in modelling the morphological features present for each source of radiation. In this particular study, we focus on NGC 1068, which in addition contains an active galactic nucleus (AGN). The effect of diffuse dust heating by an AGN (beyond the torus) has so far only been investigated for quasars. This additional dust heating source further contaminates the broadband fluxes that are used by classic galaxy modelling tools to derive physical properties. We aim to fit a realistic model to the observations of NGC 1068 and quantify the contribution of the several dust-heating sources. Our model is able to reproduce the global spectral energy distribution of the galaxy. It matches the resolved optical and infrared images fairly well, but deviates in the UV and the submillimetre (submm). This is partly due to beam smearing effects, but also because the input dust distribution is not sufficiently peaked in the centre. We find that AGN contamination of the broadband fluxes has a strong dependency on wavelength. It peaks in the mid-infrared, drops in the far-infrared, and then rises again at submm wavelengths. We quantify the contribution of the dust-heating sources in each 3D dust cell and find a median value of 83% for the star formation component. The AGN contribution is measurable at the percentage level in the disc, but quickly increases in the inner few hundred parsecs, peaking above 90%. This is the first time the phenomenon of an AGN heating the diffuse dust beyond its torus is quantified in a nearby star-forming galaxy. NGC 1068 only contains a weak AGN, meaning this effect could be stronger in galaxies with a more luminous AGN. This could significantly impact the derived star formation rates and ISM masses for such systems.

Key words. galaxies: individual: NGC 1068 – galaxies: ISM – dust, extinction

1. Introduction

Dust can be found in all star-forming galaxies, and even in a significant fraction of early-type galaxies (De Looze et al. 2010; Smith et al. 2012a; di Serego Alighieri et al. 2013). Dust grains play a vital role in the formation of stars by regulating the gas temperature, catalysing the creation of molecules such as H_2 and CO, and subsequently shielding them. Due to this close connection with both stars and gas, dust is increasingly used as a tool to study the evolution of galaxies through time (see Galliano et al. 2018, for a review). Dust and gas are strongly intertwined in the interstellar medium (ISM) and usually follow a metallicity-dependent gas-to-dust ratio (Lisenfeld & Ferrara 1998; Sandstrom et al. 2013; Rémy-Ruyer et al. 2014; De Vis et al. 2019; Casasola et al. 2020). As such, dust emission can not only be used to trace its own mass, but also the total mass of the ISM (Eales et al. 2012;

Scoville et al. 2014; Hughes et al. 2017). Accurate measurements of dust mass are an important extra metric for evolutionary studies of ISM in galaxies, as neutral gas is hard to detect beyond the local Universe, and molecular gas tracers rely on empirical conversion factors (Bolatto et al. 2013). The most obvious effect of dust is of course its capacity to absorb and scatter UV and optical light. In the local Universe, roughly one-third of starlight is absorbed by dust grains (Skibba et al. 2011; Viaene et al. 2016; Bianchi et al. 2018). This absorbed energy is re-emitted again at mid-infrared (MIR) to millimetre (mm) wavelengths where it can even be observed up to high redshifts with telescopes such as ALMA and *Herschel*. Dust is therefore also a popular star formation tracer across cosmic time (see e.g. da Cunha et al. 2010; Madau & Dickinson 2014; Casey et al. 2014; Scoville et al. 2015).

A big caveat of using dust as a tracer of either ISM mass or the star formation rate (SFR), is the assumed (often linear) conversion of far-infrared (FIR) to submillimetre (submm) emission into either of these quantities. At this point, it becomes important to understand how dust grains are heated to their equilibrium temperature. In practice, the most common assumption is that of a single heating source (e.g. star formation) producing a single radiation field, which is then parametrised with a single–average–temperature component. Dust masses can then be estimated from a modified blackbody function (Hildebrand 1983; Bianchi 2013) or a dust emission template (e.g. Dale et al. 2001; Draine et al. 2014; Jones et al. 2017). However, in reality, dust in galaxies is a mix of grain types and sizes, with multiple temperatures, and multiple heating sources. A more accurate determination of dust masses (and thus total ISM masses) requires more complex models for the dust emission. Such models also benefit the determination of the SFR from dust emission. In this case, it is important to only count the absorbed energy from new stars that is re-emitted by dust.

The main dust heating sources in galaxies have been qualitatively derived from correlations of dust emission with SFR or stellar mass tracers (Galametz et al. 2010; Boquien et al. 2011; Smith et al. 2012b; Bendo et al. 2012; Foyle et al. 2013; Hughes et al. 2014) or from panchromatic models for a galaxy's spectral energy distribution (SED; Groves et al. 2012; Aniano et al. 2012; Dale et al. 2012; Mentuch Cooper et al. 2012; Ciesla et al. 2014; Rémy-Ruyer et al. 2015; Boquien et al. 2016; Viaene et al. 2016). The broadest of such studies was performed by Nersesian et al. (2019) as they modelled the global SED of 814 nearby galaxies of different morphologies. Nersesian et al. (2019) found a wide variety in the individual contributions of dust heating sources, with a significant and sometimes dominant contribution of the old stellar populations. In a dedicated resolved study, Bendo et al. (2015) mapped the dust heating sources for 24 nearby face-on spiral galaxies, and found that in some galaxies the dust is predominantly heated by the ongoing star formation, while some are heated by the old stellar population. However, most galaxies again show a mixed heating pattern with energy from both young and older stars contributing to the dust emission.

In a pioneering study by De Looze et al. (2014), high-resolution 3D radiative transfer simulations were used to model the dust heating in M 51. The novelty in this study was to infer realistic, non-axisymmetric distributions for both stars and dust from observed broadband images. De Looze et al. (2014) found strong changes in the dust heating fractions for varying FIR/submm wavelengths, with an overall contribution of 37% from the old stellar populations. A similar model was created for the Andromeda galaxy (M 31) by Viaene et al. (2017), who found that the old stellar populations dominate the dust heating (91% globally), even in the main star-forming ring of

M 31. This effect was attributed to the large and old bulge of Andromeda. Most recently, M33 was also modelled in this way, finding old population heating contributions between 50% and 80% (Williams et al. 2019). In the context of the DustPedia project (Davies et al. 2017), we set up a best-practice framework for such radiative transfer modelling and applied it to M 81 (Verstocken et al. 2020) and to a set of nearby barred star-forming spirals (M 83, M 95, M 100 and NGC 1365; Nersesian et al. 2020).

All these studies assume a stellar origin for the energy absorbed by dust. In the radiative transfer models, three stellar populations are needed to match the observed panchromatic SED and images: an old (5–12 Gyr) stellar component, a non-ionising (~100 Myr) component, and a population of ongoing obscured star formation. While this is acceptable for many star-forming spiral galaxies, there is evidence for other sources of dust heating. Dust grains can gain energy by capturing hot electrons, protons, or even ions (Jones et al. 2004). In the latter two cases this can also partially destroy a grain (a process called sputtering). As such, cosmic rays produced by recent supernovae may be a secondary dust heating mechanism following recent star formation. In a similar way, hot gas can produce a flux of energetic electrons which directly heat the dust grains (Bocchio et al. 2013).

Another potential dust heating source is the accretion disc around supermassive black holes. In galaxies hosting an AGN, the accretion disc certainly heats the dusty torus surrounding it and produces strong MIR emission (see e.g. Fritz et al. 2006). However, higher dust temperatures and anomalously bright nuclear FIR emission are also found for galaxies hosting an AGN (Wu et al. 2007; Bendo et al. 2012; Kirkpatrick et al. 2012, 2015; Verstappen et al. 2013; Roebuck et al. 2016). This circumstantial evidence suggests that the accretion disc may be powerful enough to heat ISM dust beyond the torus. Schneider et al. (2015) and Duras et al. (2017) investigated this phenomenon using radiative transfer simulations of quasars, and found that the AGN may contribute 30−70% to the heating of diffuse dust. This phenomenon was confirmed through analysis of observations of a sample of QSOs by Symeonidis et al. (2016), Symeonidis (2017). While this may be an extreme case, it shows that AGN can contribute to the FIR-submm SED in galaxies. This effect influences the total mass and SFR estimates of galaxies hosting an AGN, but is not taken into account by SED-fitting models.

The goal of this paper is to address the question of AGN-powered dust heating in the local Universe. We choose to take a quantitative approach to separate the contributions from star formation, old stellar populations, and the AGN. Therefore, we focus on a single galaxy, but construct a detailed model based on 3D radiative transfer simulations. This is the fourth paper in a series of radiative transfer studies on individual galaxies. Within the DustPedia framework, the method was homogenised as much as possible to allow comparisons between the models. Paper I presents a radiative transfer model and dust heating analysis in the grand-design spiral M51 (De Looze et al. 2014). Paper II (Verstocken et al. 2020) is focussed on a streamlined modelling pipeline and its application to the spiral galaxy M81. A sample of barred galaxies was subsequently studied in Paper III (Nersesian et al. 2020). There have been two spin-offs studies targeting the large Local Group galaxies M31 (Viaene et al. 2017), and M33 (Williams et al. 2019). In the future we will also test this method on simulated galaxies to gain further insight into the dust-heating mechanisms and identify opportunities to improve our method.

This paper focuses on NGC 1068 (M 77), a nearby ($D = 10.1$ Mpc) Seyfert 2 SAb galaxy (Osterbrock & Martel 1993). This AGN has no broad-line emission features, despite being a nearly face-on ($i = 28.1°$) disc galaxy. Instead, optical polarimetry reveals broad reflection lines (Antonucci & Miller 1985). These observations suggest that the accretion disc is obscured by the torus for an observer on Earth (see also Marin 2018). Very-long-baseline interferometry (VLBI) measurements of a H_2O megamaser (Greenhill et al. 1996) confirm that the torus is in fact seen edge-on. Recent ALMA observations of the core of NGC 1068 reveal a complex ISM, with the torus being inclined at only $34-66°$ at larger radii (García-Burillo et al. 2016). The orientation of the torus suggests that the accretion disc is in effect beaming its radiation straight into the galactic disc, which makes NGC 1068 a well-suited target to study AGN-powered dust heating. In addition, the galaxy is actively forming stars in its inner regions and hosts a relatively quiescent outer disc (D'Agostino et al. 2018). We can therefore also compare the relative importance of stellar heating mechanisms to each other, and to the AGN contribution.

In Sect. 2 we describe DustPedia's panchromatic dataset for NGC 1068. The model preparation based on these data is outlined in Sect. 3. We describe our fitting process in Sect. 4 and the best-fit high-resolution model in Sect. 5. The dust heating in NGC 1068 is analysed in Sect. 6. We finally discuss and summarise our findings in Sect. 7.

2. Data

This investigation relies mostly on broad-band imaging from the DustPedia[1] database (Davies et al. 2017). This comprehensive database contains the 875 nearby (<40 Mpc) galaxies that were observed by the *Herschel* Space Observatory. It includes uniformly reduced PACS (Poglitsch et al. 2010) and SPIRE (Griffin et al. 2010) maps, and broad-band images from major ground and space-based survey telescopes up to the UV. In addition, aperture-matched photometry in the UV-mm is available. For more details on the image processing and photometry we refer to Clark et al. (2018).

In particular, for NGC 1068, we make use of UV maps from GALEX (Martin et al. 2005), optical images from SDSS (York et al. 2000), MIR observations from WISE (Wright et al. 2010), and submm data from *Herschel* (Pilbratt et al. 2010). We do not rely on the standard photometry of DustPedia in this case, as our model focuses on a smaller aperture. We therefore recomputed the photometry measurements based on these new apertures (see Sect. 3). The AGN in NGC 1068 is especially bright in the MIR. Images from $3.4\,\mu$m till $24\,\mu$m are severely distorted by the PSF signature of their respective instruments. This renders any morphological study almost impossible in this wavelength domain. We note that the images are not over-exposed in the MIR, meaning the integrated flux measurements are still useful and serve as the most important constraint on the AGN model.

There are only two *Spitzer*-IRAC bands ($3.6\,\mu$m and $4.5\,\mu$m) for NGC 1068 in DustPedia. They unfortunately have strong blooming artefacts and do show signs of an over-exposed centre. We therefore prefer the WISE images, which are of comparatively much better quality in this case, and cover a wider wavelength range ($3.3\,\mu$m–$22\,\mu$m). The *Spitzer*-MIPS bands do not have artefacts, but are still heavily PSF dominated (especially at $24\,\mu$m). Moreover, the PACS bands surpass them in spatial reso-

lution. For these reasons, we do not include any *Spitzer* observations in our dataset.

The 2MASS observations (also in the DustPedia database) turned out to be too shallow to capture the entire disc of NGC 1068. We instead use the UKIDSS (Lawrence et al. 2007) H and K band observations which are deeper and of higher spatial resolution. Most importantly, the H band is the reddest band without a significant PSF signature and is vital to determining the spatial distribution of the old stellar populations. For a similar reason, we also use the narrow-band $H\alpha$ map of NGC 1068 (Knapen et al. 2004). The map is an important constraint on the spatial distribution of the ongoing star formation in the absence of useable $22\,\mu$m or $24\,\mu$m images. We discuss this in more detail in the following section.

3. Model preparation

Our modelling strategy closely follows the prescriptions outlined in Verstocken et al. (2020), which, in turn were based on the work by De Looze et al. (2014). The requirements for a panchromatic, 3D, and resolved radiative transfer fit are three stellar components and one dust component. Here, we limit ourselves to a brief description at each modelling step and highlight where we deviate from Verstocken et al. (2020).

3.1. Dust component

The dust component requires a 3D geometry plus the choice of a dust mix with associated optical properties. While similar studies in the literature mostly used analytical density profiles as input geometries, it is now possible to derive them from observations directly. This allows the use of non-axisymmetric and more clumpy distributions for stars and dust, leading to more realistic models.

The most obvious way to trace the dust morphology is through a map of the dust mass surface density. However, this requires resolved SED fitting up to the SPIRE $350\,\mu$m band and drastically limits the working resolution. However, the input dust distribution does not need to be normalised as the total dust mass will be a free parameter. We therefore construct a map of FUV attenuation, which also traces the dust distribution but importantly attains higher spatial resolution (De Looze et al. 2014). The per-pixel determination of A_{FUV} is the same for all galaxies studied in this series. It relies on the calibrations of Cortese et al. (2008). These latter authors derive the attenuation based on the total infrared-to-FUV ratio (TIR/FUV) with second-order adjustments based on the observed FUV $- r$ colour. The A_{FUV} map is shown in Fig. 1 (top left panel) labelled as the dust disc. To convert this to a 3D density distribution, we deproject the 2D image along the minor axis and stretch the emission in each deprojected pixel out into the vertical (z-axis) direction according to an exponential profile. We set the scale height of this exponential profile to 76 pc, which is half that of the old stellar scale height as suggested by De Geyter et al. (2014). For the spectral properties of dust absorption, scattering, and emission we adopt the THEMIS dust model for diffuse Milky Way dust (Ysard et al. 2015; Jones et al. 2017), in line with other DustPedia studies. A typical emission SED of this dust model is shown in Fig. 2.

3.2. Stellar components

Each of the stellar components has a 3D spatial distribution associated with it, and a template shape for its SED. The templates
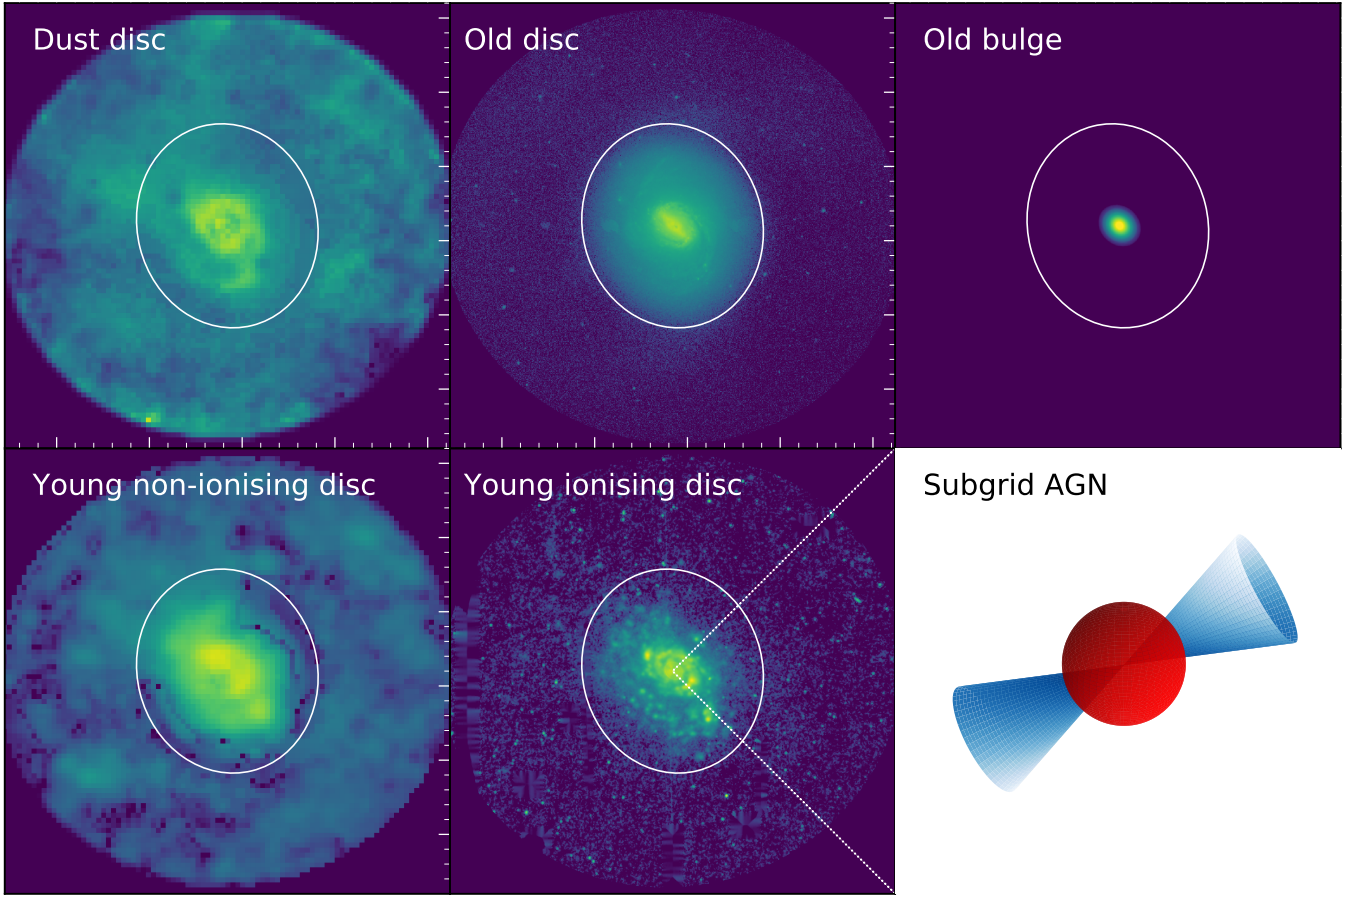
---

[1] http://dustpedia.astro.noa.gr

Fig. 1. Two-dimensional representations (on the plane of the sky) used as input morphologies for the 3D geometries for the model of NGC 1068. The old bulge and AGN model are based on analytical prescriptions while the other morphologies are derived from observed imagery. The model χ^2 is only evaluated within the white ellipse to ensure a sufficient signal-to-noise ratio. The ellipse has a semi-major axis of 83 arcsec or 4 kpc. The AGN cone orientation in the *bottom right panel* is arbitrary and is for display purposes only.

are simple stellar populations taken from the Bruzual & Charlot (2003) library assuming a Chabrier (2003) initial mass function (IMF; see the blue lines in Fig. 2). We split the stellar components into three age groups. The old stellar population contains stars that are much older than 100 Myr and therefore no longer contribute to the classic SFR tracers (Kennicutt & Evans 2012). Typically, one can trace this component best with either WISE 3.4 μm or IRAC 3.6 μm observations. However, in the case of NGC 1068, these images are dominated by the PSF shape and do not provide clear morphological information. We therefore rely on the UKIDSS *H* band to trace the old stellar component. The additional advantage of using this shorter waveband is that the AGN contamination in the centre is minimal. We perform a bulge/disc decomposition on the *H* band image to further split the geometry into a Sérsic bulge and a disc. The bulge model (the only analytical density profile) is already three-dimensional by construction. The disc is deprojected in the same way as the dust map, but now given an exponential scale height of 151 pc, which is 1/8.26 of the scale length as suggested by De Geyter et al. (2014). The second stellar component is a young non-ionising stellar population with average age of 100 Myr and is based on the attenuation-corrected GALEX-FUV map. This map was deprojected and given an exponential scale height equal to the dust scale height.

The third stellar component models the ongoing star formation in the form of young (10 Myr) ionising stars still in their birth clouds. The MAPPINGS III (Groves et al. 2008) model

of stars and dust in a birth cloud with this average age was used as SED template for this component (see Fig. 2). To trace the ongoing star formation, we use the narrow-band continuum-subtracted Hα emission. This image still has an unusually bright nucleus. D'Agostino et al. (2018) showed that the AGN is responsible for 25% of the Hα emission in the galaxy, and for 42% of the [NII] line emission. Both emission lines fall within the narrowband filter image used here. To mitigate the contamination, we mask out the nuclear Hα peak in the inner 350 pc or 7 arcsec (see also Casasola et al. 2015, for a similar approach). The astropy routine interpolate_replace_nans[2] was used to infer the Hα emission from the surrounding disc. We note that this procedure preserves the exponential flux profile of the Hα disc. The total flux in the image is reduced by 41% after this correction, which is at the high end of the range estimated by D'Agostino et al. (2018). The resulting image is converted to a 3D geometry by deprojecting it and assuming an exponential scale height of 38 pc, which is half that of the dust scale height and represents the embedded nature of the star-forming regions.

3.3. The AGN component

An important novelty is the addition of an AGN in our model. Radiative transfer simulations of accretion discs surrounded by a dusty torus have been executed in the past (see e.g.

---

[2] http://docs.astropy.org/en/stable/convolution/

Fig. 2. Overview of the different SED templates adopted in the model. Templates are scaled for improved visualisation.

Nenkova et al. 2008; Stalevski et al. 2012). However, given the enormous difference in scale between the torus and the galactic disc, it is not necessary to solve the transfer of radiation in the same simulation. As such, we construct a subgrid model for the AGN in NGC 1068. The two AGN emission components (the accretion disc and the dusty torus) are both modelled as point-like sources located at the centre of the galaxy.

The accretion disc spectrum is modelled as a composition of power laws with indexes varying as a function of wavelength. We used the Schartmann et al. (2005) prescription, which combines observed and theoretical evidence and is commonly used as a heating source in radiative transfer models for dusty tori in AGNs (Fig. 2). Its emission peaks at UV–optical wavelengths and its normalisation is tied to the total torus luminosity. A part of the radiation of the accretion disc is blocked locally by dust within the torus. We implemented this as an anisotropic (cone-shaped) emission profile with an opening angle of 20° around the polar direction (see Fig. 1, bottom right panel). This simple geometrical implementation mimics the same internal absorption fraction as our preferred choice of torus model (see below): only 6% of the accretion disc light escapes the torus. The orientation of the accretion disc is such that it beams directly into the galactic disc, along the minor axis (inclination $i = 0°$, azimuth $\phi = 90°$), with the obscured part pointing to the observer, as expected for a Type-2 Seyfert galaxy with a hidden Type-1 core like NGC 1068 (Marin 2018).

We assume the torus emission to be isotropic, which is a fair approximation given our limited spatial resolution (44 pc) and the fact that its MIR radiation is far more optically thin than the obscured accretion disc (see also Fig. 1, bottom right panel). To choose an emission model from the dusty torus, we performed a fit to the photometric points used by Lopez-Rodriguez et al. (2018), using the upgraded model grid of Fritz et al. (2006) pre-

sented in Feltre et al. (2012). The former is a collection of spatially resolved data that are dominated by the torus emission and include crucial features such as the near-infrared (NIR) emission and the 9.7 μm silicate feature, observed in slight absorption for this AGN (Fig. 2). The model that provides a good fit to these points has a dust full-opening angle of 140°, and an outer-to-inner radius of 30. The equatorial optical depth is 6 measured at $\lambda = 9.7\,\mu$m, and the dust density $\rho(r, \theta)$ varies within the torus according to:

$$\rho(r, \theta) = \alpha r^\beta e^{-\gamma |\cos(\theta)|}, \qquad (1)$$

where α is directly related to the value of the equatorial dust optical depth, and the two parameters β and γ regulate the radial and height density profile (see Eq. (3) in Fritz et al. (2006) for further details, where the authors assume the value of 0 and 4, respectively, for the best fit model). The AGN model is normalised by the 4.6 μm luminosity since this wavelength is particularly sensitive to the presence of very hot dust. It is not the goal of this paper to further constrain the torus or accretion disc properties. We simply require a realistic subgrid implementation for the AGN. As such, the total luminosity is the only free parameter of the AGN component (with a fixed ratio between accretion disc and torus luminosity).

4. Radiative transfer SED fitting

The different components described in the previous section define the framework for our model of NGC 1068. Radiative transfer simulations are computationally demanding, and so the free parameters in this model need to be limited. We therefore rely on well-informed estimates for the spatial and spectral distribution of each component as outlined above. For the dust density distribution the only free parameter is the total dust mass, $M_{\rm dust}$. For the other components, the free parameters are the normalisation luminosities in certain wavebands: The spatial and spectral distribution of the young non-ionising (yni) and young ionising (yi) stellar populations are normalised by their FUV luminosity: $L_{\rm FUV}^{\rm yni}$ and $L_{\rm FUV}^{\rm yi}$, respectively. The old stellar components are normalised together (assuming a fixed bulge-to-disc ratio) by their total luminosity in the H band: $L_H^{\rm old}$. Finally, the subgrid AGN SED is normalised by the total emission of accretion disc and torus at 4.6 μm: $L_{4.6}^{\rm AGN}$. Here again the ratio between disc and torus luminosity is fixed to guarantee the internal energy balance. In total, the model therefore has five free parameters.

We make use of SKIRT (Baes et al. 2011; Camps & Baes 2015) to perform the radiative transfer calculations in the 3D simulation space. The stellar and AGN components serve as sources of radiation. The dust component acts as a sink of UV–optical–NIR radiation, which is then reprocessed and emitted at longer wavelengths. Light is also anisotropically scattered by the dust grains. SKIRT can perform these simulations in a highly efficient way thanks to multiple optimisation techniques for the transfer of radiation towards the observer (Baes et al. 2011), effective quantisation of the dust density distribution using a binary tree dust grid (Saftly et al. 2014), and support for various types of parallel computing (Verstocken et al. 2017).

Our goal is to find the optimal combination of free parameters that reproduce the observations. Despite the efficiency of SKIRT, this is not a trivial task. We therefore run a suite of low-resolution simulations and use the global broadband SED of NGC 1068 as an observational constraint since it is efficient to generate the synthetic line-of-sight model fluxes. The radiation transfer is evaluated in a set of 3D dust cells of variable volume,

Table 1. Overview of the free parameters of our model setup together with their initial guess values.

Parameter	Initial guess	Batch 1 (50th pct)	Batch 2 (50th pct)	Best fit
$\log(M_{\rm dust}/M_\odot)$	7.23	$7.43^{+0.15}_{-0.19}$	$7.40^{+0.11}_{-0.12}$	7.31
$\log(L_H^{\rm old}/W\,\mu{\rm m}^{-1})$	36.65	$36.64^{+0.17}_{-0.25}$	$36.54^{+0.15}_{-0.19}$	36.44
$\log(L_{\rm FUV}^{\rm yni}/W\,\mu{\rm m}^{-1})$	37.56	$37.85^{+0.22}_{-0.32}$	$37.65^{+0.21}_{-0.28}$	37.58
$\log(L_{\rm FUV}^{\rm yi}/W\,\mu{\rm m}^{-1})$	37.46	$37.76^{+0.22}_{-0.33}$	$37.59^{+0.23}_{-0.31}$	37.50
$\log(L_{4.6}^{\rm AGN}/W\,\mu{\rm m}^{-1})$	35.47	$35.63^{+0.30}_{-0.32}$	$35.41^{+0.16}_{-0.24}$	35.41

Notes. After each fitting batch, the 50th percentile corresponds to the median value for that fit. The uncertainties on the 50th percentiles reflect the width of the corresponding bin. The last column lists the parameter values for the model with the lowest overall χ^2.


Fig. 3. Probability density distributions for the free parameters in our model configuration. The spacing between bars represents the spacing of the parameter grid. The width of the bars is set so they do not overlap due to the unequal spacing. The blue distribution corresponds to the first batch of models. The parameter grid was refined based on this batch. After the second fit iteration, the red distributions are found. The median values of the second batch are indicated with vertical dashed lines. The solid vertical lines indicate the values for the model with the lowest overall χ^2.

but roughly equal in mass. With a dust grid of 1.3×10^6 cells, we reach an effective resolution of 44 pc in each dimension. In each simulation, the model is evaluated in 134 wavelength bins by 1×10^6 photon packages per wavelength. Doing so, each set of parameters can be evaluated in under one hour (on average) on a 16 CPU node.

With each point in the 5D parameter space taking 16 CPU hours, it becomes impractical to run a classic non-linear fitting algorithm. We therefore adopt the same two-step fitting approach as for the other DustPedia-modelled galaxies (Verstocken et al. 2020; Nersesian et al. 2020). In a first run, we explore a wide parameter space, sparsely sampled around a well-motivated initial guess (see below). For each parameter set, the observed and model broadband SED is compared by summing the χ^2 metric. The fluxes are measured in 2D (sky projection) images inside an elliptical aperture with semi-major axis of 4 kpc (see Fig. 1). Each wavelength regime (UV, optical, NIR, MIR, FIR and submm) is given equal weight to the total squared sum (see Verstocken et al. 2020, for a detailed description). The parameter values for each fit can then be weighted by $\exp(-\chi^2/2)$ to sample the probability density distribution.

As an initial guess, we use the global properties derived by Nersesian et al. (2019) which are based on SED modelling with CIGALE (Boquien et al. 2019). Nersesian et al. (2019) find

$$
\begin{cases}
M_{\rm dust}^{\rm CIGALE} = (1.71 \pm 0.23) \times 10^7 \, M_\odot \\[2mm]
{\rm SFR}^{\rm CIGALE} = 13.4 \pm 3.6 \, M_\odot \, {\rm yr}^{-1} \qquad\qquad (2) \\[2mm]
A_{\rm FUV}^{\rm CIGALE} = 3.64 \pm 0.31.
\end{cases}
$$

This dust mass can directly be used as an initial guess. From the SFR, we compute an initial guess for the UV luminosity of

the ionising stellar populations. The $A_{\rm FUV}$ can be used to de-redden the observed FUV luminosity, which we assume here for simplicity to be the sum of the ionising and non-ionising stellar components. Furthermore, we use the observed H band luminosity as an initial guess for the total (bulge+disc) old stellar populations as this band is relatively free from other contributions. For the AGN torus, our initial guess normalisation was set at $2.50 \times 10^{35}\,W\,\mu{\rm m}^{-1}$ and tied to this an accretion disc luminosity $0.15 \times 10^{35}\,W\,\mu{\rm m}^{-1}$. These luminosities provide the best fit of our model to the estimated fluxes of the central 20 pc from Lopez-Rodriguez et al. (2018). The initial parameter set for the model is listed in Table 1.

For galaxies without a strong AGN, radiative transfer fits with three parameters ($M_{\rm dust}$, $L_{\rm FUV}^{\rm yni}$ and $L_{\rm FUV}^{\rm yi}$) are sufficient. The challenge of adding an AGN is that the MIR SED becomes even more convoluted and difficult to disentangle into separate contributions. To deal with this added complexity, we first introduce $L_{4.6}^{\rm AGN}$ as a free parameter to scale the AGN template SED. Secondly, we also allow the old stellar population template SED to scale up and down. This is set by $L_H^{\rm old}$ since the H band is a good tracer of this stellar population and the contamination of the AGN is expected to be minor. An adequate coverage of this parameter space therefore requires many more sample points. We build a sparse but wide 5D Cartesian parameter grid around our initial guesses. The grid spans more than one order of magnitude in $M_{\rm dust}$, $L_{4.6}^{\rm AGN}$, and $L_H^{\rm old}$, and more than two orders of magnitude in $L_{\rm FUV}^{\rm yni}$ and $L_{\rm FUV}^{\rm yi}$. Each dimension is uniformly sampled in log space with five points during this first run, which already amounts to 3125 simulations.

The resulting probability distributions of the first batch are shown in Fig. 3 by the blue histograms. These have a clear peak for all free parameters, which importantly does not occur at the edge of the sampled space. This means that we are already able

Fig. 4. Spectral energy distributions from the best-fit model for NGC 1068. *Left*: global SED (black line) fits to the observed (orange) broadband fluxes. The grey shading indicates the 16th to 84th percentile area around the fit. Triangles indicate lower limits. The corresponding residuals are shown in the *bottom panel*. The solid green line shows the total emission for a model with only the AGN + dust component. The dotted green line then corresponds to the light reprocessed by the diffuse dust only in this simulation. *Right*: same best-fit SED and data points. The solid blue and red lines correspond to simulations where only one component (young or old, respectively) is present alongside the dust. A CIGALE model is also shown (dashed black line). This model can be decomposed into young and old populations (blue and red dashed lines, respectively). The intrinsic young population (blue dotted line) of the CIGALE fit is also shown.

to put coarse constraints on all parameters. The median values (50th percentile of the PDFs) are summarised in Table 1.

Based on the distributions of the first batch, we can narrow the parameter space significantly and adopt a finer sampling. We ensure that over 90% of the probability is captured in the parameter space for the second batch. We again use five sample points per parameter, but now uniformly distributed in linear space. The probability distributions for this second batch of 3125 simulations are also shown in Fig. 3 (red histograms). Each distribution shows a clear peak, with the exception of $L_{\mathrm{FUV}}^{\mathrm{yi}}$. It seems we are reaching the limit of precision by which we can constrain this parameter. The PDF for $L_{4.6}^{\mathrm{AGN}}$ is also rather broad for the second batch. The parameter values for the best-fitting model (lowest χ^2) lie close to the median parameters (50th percentile values) and are listed together in Table 1. These values also correspond well with our initial guess. We explore the 2D probability distributions in Appendix A and find no strong degeneracies or suspicious features in these figures. We conclude that at this point the free parameters of our model are reasonably constrained. Adding another iteration would again require significant computational resources and would not significantly reduce the uncertainties. In the following section, we analyse the best-fitting model in further detail.

5. A 3D high-resolution model for NGC 1068

5.1. Comparison of SEDs

The best-fit model obtained in the previous section forms the basis for our dust heating analysis. We first produce a higher quality version of the best-fit model by shooting five times as many photons per wavelength (5×10^6), and almost twice the number of wavelength points (252). This increases the spectral sampling and reduces the inherent Monte Carlo noise in the dust cells and in the resulting images. To separate the contributions of each source of radiation, each simulation is run again with the same dust distribution, but only one source component each time: the old (bulge and disc), young (non-ionising and ionising discs), and AGN (accretion disc and torus) components, respectively. This process takes another 1280 CPU hours. The resulting

SED, along with its uncertainty levels, is shown in Fig. 4 (left panel, black line).

Globally, our model follows the observed SED fairly well. The main stellar and dust peak match the datapoints closely and within the uncertainties. However, there are points of discrepancy. First, the FUV point is overestimated by the model, while the NUV point is underestimated. This issue has been noted before in galaxies with a significant contribution of the young ionising templates (De Looze et al. 2014; Verstocken et al. 2020; Nersesian et al. 2020). This signature in the template for ionising stellar populations is amplified by a broad bump in the THEMIS attenuation curve. Nersesian et al. (2020) compared the attenuation curve (normalised to A_V) of our model to their galaxies and found that NGC 1068 exhibits the flattest attenuation curve of the sample. It also lies closest to the extinction curve corresponding to the THEMIS diffuse Milky Way dust model, which is also used here in our radiative transfer simulations.

Also driven by the dust model are the strong aromatic peaks in the MIR SED in Fig. 4. The adopted emission model was calibrated on diffuse dust in the Milky Way. The strength of MIR emission features is often reduced in AGNs compared to star-forming galaxies (see e.g. Sales et al. 2010) although this may be attributed to dilution by the bright MIR continuum (Alonso-Herrero et al. 2014). Mid-infrared spectra of NGC 1068 do reveal line emission even close to the nucleus (Mason et al. 2006; Howell et al. 2007) due to ongoing star formation there. Environmental changes in dust composition and grain mixture are capable of significantly changing the dust emission profile (Köhler et al. 2015) but this is currently not possible in our model. Destruction of the smallest amorphous hydrocarbon (a-C) grains would for example reduce the $3\,\mu m - 13\,\mu m$ emission band intensities. At the same time, a lack of a-C grains would also reduce the UV bump and provide a better match to the GALEX data points. This underlines the need to allow for dust evolution in the next generation of radiative transfer codes. A framework for such simulations is currently being developed (Camps, in prep.). We did not wish to skew our model by fitting to wavebands that contain MIR emission peaks and therefore do not include the WISE $3.4\,\mu m$ and $12\,\mu m$ flux in the χ^2 weighting.

As already mentioned in Sect. 2, good-quality maps of the main disc of NGC 1068 are lacking in the MIR. Particularly at $12\,\mu m$ and $22\,\mu m$, the PSF signature (a blend of the disc and AGN emission) is larger that the field of view we consider. We therefore treat these two points as lower limits. The most discrepant point is the WISE $22\,\mu m$ flux, which is underestimated by the model. We looked into models with a stronger AGN contribution, but these lift the MIR continuum at the same time. Changing the ratio between the stellar components also did not improve the fit without introducing discrepancies at other wavelengths. The observed WISE $22\,\mu m$ flux in our aperture is higher than the MIPS $24\,\mu m$ flux (not shown and also suffering from strong PSF signatures). Part of this discrepancy could therefore be attributed to the observational uncertainty, but still our model seems unable to capture the exact SED shape in the MIR. This may partly be due to the subgrid implementation of the torus. A non-isotropic torus model or a more clumpy one could boost the emission at $22\,\mu m$ in the AGN SED (see Feltre et al. 2012). However, as indicated earlier, it is not the goal of this paper to further constrain the torus properties themselves, but instead we focus on how radiation escaping the AGN affects the host galaxy.

Alternatively, changing one of the stellar emission templates could bridge the discrepancy. The MIR emission in particular can be boosted by tweaking the young ionising component template. The covering factor used in the Groves et al. (2008) template library can be increased to boost the emission from warm dust surrounding star-forming regions, while at the same time keeping the escaping UV radiation relatively low. Unfortunately, introducing yet another free parameter in the optimisation procedure would be computationally difficult. In addition, one of the main goals of this paper series is to provide a suite of radiative transfer models for different galaxies with a consistent set of fixed parameters (Verstocken et al. 2020). Only the template luminosities and dust mass are considered free parameters, which allows us to investigate heating mechanisms across the sample (Nersesian et al. 2020).

We also run separate simulations containing only one emission component: the young (non-ionising+ionising) stellar populations, the old (bulge+disc) stellar populations, and the AGN. All of these simulations still contain the dust component and are simply re-runs of the best-fit simulation with fewer emission sources. The diffuse dust is now heated solely by the components that are present in the simulation. We note that the sum of these SEDs will not add up to the global SED (with all components) as dust emission is a non-linear effect.

In the left panel of Fig. 4 we plot the SED of the simulation with the AGN (accretion disc+torus) + dust only. The AGN only contributes significantly to the MIR, but does not dominate the global SED in any regime. However, an important observation here is the broad tail of the AGN SED in the submm. This is not direct emission from the AGN, but light reprocessed by the diffuse dust in the galaxy (see green dotted line in Fig. 4). This already indicates that the AGN radiation is influencing the surrounding ISM. Nevertheless, the effect is rather small as the dust emission lies typically one order of magnitude below the emission generated by the old or young components. We compute a bolometric luminosity $L_{\mathrm{AGN}}^{\mathrm{bol}} = 0.4 \times 10^{37}\,W$, compared to $L_{\mathrm{dust}}^{\mathrm{bol}} = 2.7 \times 10^{37}\,W$. For the total observed SED, the bolometric luminosity is $L_{\mathrm{tot}}^{\mathrm{bol}} = 4.3 \times 10^{37}\,W$. Following Bianchi et al. (2018), we can compute $f_{\mathrm{abs}} = L_{\mathrm{dust}}^{\mathrm{bol}}/(L_{\mathrm{tot}}^{\mathrm{bol}} - L_{\mathrm{AGN}}^{\mathrm{bol}}) = 0.69$ the fraction of light that is reprocessed into dust emission. Our estimate corresponds well with the 0.68 found by these latter authors from a CIGALE model of the global SED.

The SEDs of the old+dust and young+dust simulations are plotted in the right-hand panel of Fig. 4. The dominance of the young+dust simulation (blue line) is immediately clear. Its SED lies well above the SED of the other simulations in the UV, MIR, and FIR regimes. This is already an indication that the FIR flux is powered by ongoing star formation. The second-most-luminous SED comes from the old+dust simulation, although it only really dominates in the NIR. Still it seems to contribute some energy to the submm emission.

We can compare our model and subcomponents with 1D SED models like the ones fitted by CIGALE. This method is well suited for fast modelling of statistical samples, but as a trade-off it lacks 3D information and radiative transfer effects. We adopt the modelling method of Nersesian et al. (2019) and apply this to the fluxes of NGC 1068 in the ellipse region under consideration (see Fig. 1). The resulting SEDs of the total and subcomponents of the CIGALE model are shown as dashed lines in Fig. 4. The global SED fits the observed data very well. The number of free parameters (and their priors) in CIGALE is sufficient to alleviate the discrepancies present in our radiative transfer model. Especially interesting is the emergence of a second peak in the MIR to match the $22\,\mu m$ point. This does seem to be at the cost of a lower continuum at shorter wavelengths, underestimating the $4.6\,\mu m$ point.

A breakup of the CIGALE model into subcomponents provides more insight into the global fit. Figure 4 shows that the observed SED model is dominated by the old component. This can be compared to the old stellar component of the radiative transfer model, bearing in mind that the latter SED also contains dust emission. The CIGALE subcomponents shown here do not include dust emission. Similarly, we can compare the young components and find a vast discrepancy between the models. The CIGALE young component has a negligible contribution to the observed SED. To inspect this mode closely we also plot the unattenuated SED of the young component (blue dotted line). The unattenuated SED is much more luminous pointing towards extreme levels of attenuation. This is likely related to the emission peak at $22\,\mu m$. As CIGALE tries to fit this data point, it needs to boost the continuum and PAH luminosity, which requires strong attenuation of UV and optical light. It is important to note that this model does not contain an AGN. Additional freedom of the AGN component could help to fit the MIR without skewing the attenuation for the young population.

From this comparison we conclude that 1D SED fitting tools can provide good fits to the global fluxes of galaxies hosting an AGN. However, caution is needed in interpreting these models. In large samples it is a priori not clear which galaxies require a low-luminosity AGN component (like NGC 1068) in the model. Looking at the subcomponents provides useful insight into this case. Standard CIGALE fits with an extremely attenuated young stellar component and a double IR bump can identify AGNs in a sample of galaxies. This initial result motivates more comparisons between 3D RT SED fits and 1D tools to highlight deficiencies in both methods and act upon them.

5.2. Image comparison

We generate synthetic images to further assess the quality of our model. We convolve the 3D datacube produced by SKIRT using the spectral response curves of five representative broadband filters across the SED. We do not convolve the images with a corresponding PSF as it is difficult to determine the model kernel. The PSFs of the images that were used to extract the input density distributions, as well as the resolution of the dust grid, result

in a complex effective PSF of the model image. As such we simply compare the spectrally convolved model images with their observed counterparts in Fig. 5 and indicate the instrument PSF as a yellow circle. We also produce residual maps in dex units as $\log_{10}(\text{observation}) - \log_{10}(\text{model})$, and plot the kernel density estimates (KDEs) for the pixel distributions in these residual maps.

The model images reproduce the observed counterparts within 0.5 dex. Visually, the correspondence is quite good in the UV, optical, and NIR maps. However, the residuals reveal significant asymmetries corresponding to the spiral arms. This is due in part to the deprojection of 2D input images to create a 3D density distribution, where light is smeared out in the vertical direction. On top of that, we do not resolve the dust extinction at sufficient resolution, which introduces residual features at smaller scales. The effect is worst in the SDSS g band image, where the spatial resolution is the highest. In this band the flux is also systematically underestimated by the model as highlighted by the KDE plot. In the NIR, there are two strong positive residual features on opposite sides of the galaxy. These correspond to star-forming regions which are over-luminous in the model (represented by the ionising component). However, globally, the NIR regime shows the best correspondence with the observations.

The model FIR images do not show the bright central area, but the PACS map does peak sharply in the nucleus. This can mainly be attributed to the difference in PSF, but the model also underestimates the flux in the centre (while matching the total flux in the FIR regime). We speculate that this may be due to an insufficient central dust density in the input map. The current dust geometry map is based on pixel-by-pixel determinations of A_{FUV}. This relies on the calibrations of Cortese et al. (2008), which derive the stellar age bin from the FUV−r colour. It is possible that the AGN emission in NGC 1068 can locally skew these age estimates, resulting in a less reliable A_{FUV} estimate in the centre of the galaxy. We have tried to fit a model where the dust geometry followed the PACS 70 μm map more closely. This dust map is more centrally peaked than the current one based on A_{FUV} and is presented in Appendix B. However, this alternative map causes severe attenuation in the UV and optical bands, too much dust emission in the FIR, and a less favourable fit to the observed global fluxes. We therefore did not analyse this model further and stick to the model based on the A_{FUV} dust map. Knowing the qualities and caveats to the best-fitting model, we analyse the dust heating sources and energy balance in NGC 1068 in the following section.

5.3. Light fractions

The global SED of our model further allows us to estimate the light contamination of sources unrelated to star formation in NGC 1068. We therefore compute, in each band, the emission ratio between the AGN-only and the total simulation (still including the same dust component), F_λ^{AGN}, and similarly for the old stellar population: F_λ^{Old}. It is useful to quantify these contributions since several broadband colours or luminosities are used to directly derive physical properties of galaxies. An adequate correction for star-forming galaxies containing an AGN is necessary before applying such recipes. Likewise, a correction for light coming from older stellar populations can improve estimates of star-formation and attenuation curves (see e.g. Roebuck et al. 2019).

Figure 6 (top panel) shows F_λ^{AGN} for all bands in our dataset, and for several concentric elliptical apertures. As expected, there is a general increase in the AGN light ratio towards the innermost

regions of the galaxy. F_λ^{AGN} averages to only 0.12 inside a 4 kpc major axis radius and across the wavelength range. In the innermost aperture (0.5 kpc), this average increases to 0.32. Along the wavelength axis, all apertures exhibit the same behaviour. There is virtually no AGN contamination in the UV bands and only a minor contamination in the optical bands ($F_\lambda^{\text{AGN}} = 0.01 - 0.05$). There is a strong increase in the NIR and the peak F_λ^{AGN} is reached in the WISE 4.6 μm band and varies from 0.50 within 4 kpc to 0.83 within the inner 0.5 kpc. The peak actually extends well into the MIR, but drops off quickly towards the PACS 70 μm band. In the FIR, an interesting positive slope points to an increasing contribution with increasing wavelength. The intrinsic AGN (disc+torus) emission model does not have this behaviour, as is evident from the right panel in Fig. 4. This is actually an indirect effect as AGN light is reprocessed by the diffuse dust. It is important to realise that even SPIRE bands can therefore be contaminated by the AGN.

The radial and spectral variation in F_λ^{Old} is displayed in the bottom panel of Fig. 6. There is less variation between the different apertures than for the AGN light fraction. Still a similar pattern emerges: the highest fractions are found in the centre and can be attributed to the bulge population of older stars. Interestingly, there are two peaks along the spectral axis. The first peak around 1 μm is expected and corresponds to the peak of direct emission of the old component. The second peak lies around 350 μm where F_λ^{Old} still reaches 0.5. This can be attributed to reprocessed light from the old stellar populations by the diffuse dust. We note that this number does not suggest a dust heating fraction of 50% because the conversion of dust emission SED to the input energy is non-linear. We study the input energies further in the following section but the light fractions already qualitatively confirm that both old stellar populations and the AGN affect the total dust emission in this galaxy.

6. Dust-heating sources

The main advantage of 3D radiative transfer simulations is the inherent information on the energy balance in every dust cell. Our model, based on NGC 1068, contains the absorbed energy in each cell and thus summarises how the dust in that cell is heated before re-emitting this energy at longer wavelengths. In line with our previous work (Verstocken et al. 2020; Nersesian et al. 2020), we make use of the total (wavelength-integrated) absorbed energy per stellar component and per volume element (dust cell). We derive three quantities: f_{young}, the combined heating fraction of the young (non-ionising + ionising) component, f_{old} the heating fraction of the old stellar component, and f_{AGN}, the AGN heating fraction. These reflect dust heating through ongoing star formation, the general radiation field, and through AGN emission, respectively, and are defined as

$$f_{\text{young}} = \frac{L_{\text{yni}}^{\text{abs}} + L_{\text{yi}}^{\text{abs}}}{L_{\text{tot}}^{\text{abs}}},$$

$$f_{\text{old}} = \frac{L_{\text{old}}^{\text{abs}}}{L_{\text{tot}}^{\text{abs}}}, \qquad (3)$$

$$f_{\text{AGN}} = \frac{L_{\text{AGN}}^{\text{abs}}}{L_{\text{tot}}^{\text{abs}}},$$

where the L_k^{abs} refers to the absorbed luminosity for component k. These quantities are computed per dust cell. The total absorbed luminosity is the sum of all components per dust cell: $L_{\text{tot}}^{\text{abs}} = L_{\text{yni}}^{\text{abs}} + L_{\text{yi}}^{\text{abs}} + L_{\text{AGN}}^{\text{abs}} + L_{\text{old}}^{\text{abs}}$.

Fig. 5. Spatial comparison of our best-fit model to observed broad-band images (*left column*). The *second column* contains the model images. We note that the model images are not convolved by any PSF and only hold the intrinsic model PSF (a composite of the input images). Residual images are shown in the *third column*. Corresponding residual distributions are represented by KDE plots in the *right column*. The area under the KDE curves is normalized to 1.

Fig. 6. Fraction of emission originating from the AGN component to the total emission per band, F^{AGN} (*top panel*) and from the old component F^{Old} (*bottom panel*). Different lines correspond to different concentric elliptical apertures. The light fractions are derived from the integrated light within the full aperture.

Figure 7 summarises the dust heating fraction of each component. The left panels show a cut through the mid-plane of dust cells in a face-on view of the galaxy model. In this view, the direction of the AGN beam is aligned with the horizontal axis of the figure. The colour coding is set by the histograms on the right. The histograms contain all dust cells (including the ones above and below the mid-plane). Globally, a median $f_{\text{young}} = 0.83$ was found, which underlines the dominance of star formation as a heating source in this galaxy model. In contrast, the old stellar populations have a median $f_{\text{young}} = 0.16$ and the AGN effect is minor with a median $f_{\text{AGN}} = 0.006$.

More detail is visible on a local level in Fig. 7. The young component peaks in a ring around the centre, and in two zones above and below it. These structures can also be distinguished in optical images of NGC 1068 and are associated with enhanced star formation (see the ionising disc map in the bottom middle panel of Fig. 1 or the NUV map in Fig. 5). The base level of f_{young} remains high across the disc, but drops below 0.5 in some patches at larger radii. This is where the ongoing star formation fades and the radiation field created by the old stellar population takes over in heating the dust. In many ways, the map of f_{old} appears complementary to the f_{young} map. This indicates that stars (young or old) are still very much the dominant heating sources in this galaxy.

Quite remarkable is the central dip in f_{young}, which cannot be fully explained by the peak in f_{old} due to the bulge. This central

dip spatially correlates with the peak in f_{AGN} in the bottom left panel of Fig. 7. The central ∼0.5 kpc is the main zone of influence of the AGN. Here, the AGN-powered dust heating even dominates that of the surrounding star-forming disc. The f_{AGN} map also shows a clear directional preference for the AGN heating (roughly along the horizontal axis of Fig. 7). The main direction is guided by our implementation of the AGN accretion disc, beaming straight into the disc with a conical opening angle of 20° (see Sect. 3). The map also shows a slight left-right asymmetry (highlighted by the white contour), where the AGN heating can be noted further out towards the right side of the map. The star-forming disc fades off more quickly in that direction, leaving room for the AGN to increase its contribution to the dust heating.

To further quantify the AGN-powered dust heating, we plot the radial profiles of f_{AGN} and f_{young} in Fig. 8. The same trend as in Fig. 7 arises. The ongoing star formation dominates the dust heating in most of the dust cells in the disc with fractions between 60 and 95% and with appreciable scatter at every radius. The radial distribution for f_{AGN} exhibits a steep decline from the centre towards the outskirts. We can approximate the radial decline of the AGN heating fraction quite well with a power law:

$$f_{\text{AGN}} = 0.0117 \left(\frac{R}{\text{kpc}} \right)^{-1.19}. \tag{4}$$

In the above fit we excluded data points inside the inner 50 pc to avoid limited sampling and our spatial resolution limit. We also excluded the points beyond 3000 pc to avoid edge effects. The power law thus only holds within these radial limits. Further models of galaxies hosting an AGN are required to verify whether or not this power-law behaviour is universal. Our analysis suggests that AGN-powered dust heating is appreciable (>5%) inside the inner few hundred parsecs, and remains at the percentage level throughout the galaxy.

7. Discussion and summary

We created a radiative transfer model based on observed imagery of NGC 1068. This galaxy contains an extended old disc and bulge, and a smaller star-forming disc with spiral arms. NGC 1068 also bears an AGN and its accretion disc appears to be edge-on while the galaxy is roughly face-on. As such the accretion disc beams directly into the star-forming disc. This particular configuration makes the galaxy an ideal target to investigate a scenario where the AGN heats diffuse dust beyond its torus. While there is circumstantial evidence for such AGN-powered dust heating in the Local Universe (Wu et al. 2007; Bendo et al. 2012; Kirkpatrick et al. 2012, 2015; Verstappen et al. 2013; Roebuck et al. 2016), the phenomenon has thus far only been simulated for quasars (Schneider et al. 2015; Duras et al. 2017). For such objects, observational evidence suggests that the AGN heating fraction is at least 5−10% (Symeonidis et al. 2016; Symeonidis 2017). However, QSOs are extreme environments with AGNs that are orders of magnitude more luminous that in the Local Universe. To our knowledge, we present the first full 3D radiative transfer fit to a local star-forming galaxy with an AGN.

The complex geometry of structure of NGC 1068 makes it difficult to disentangle the main components of our model. We used linear combinations of observed broad- and narrow-band images to obtain the purest possible 2D maps of old, non-ionising, and ionising stellar populations. The main challenge

Fig. 7. Dust heating fractions in a slice through the mid-plane of the face-on view of our model (*left panels*). The direction of the AGN beam is aligned with the horizontal axis of the figure. The colour code is set on the histograms on the *right*. The histograms contain all dust cells (including the ones above and below the midplane). *Top row* corresponds to f_{young}, middle row to f_{old}, and *bottom row* to f_{AGN}. The white contour corresponds to the 0.5% level in f_{AGN} and highlights an asymmetry in the heating of dust in the disc.

Fig. 8. Scatter density plots showing the radial profiles of the dust heating fraction. The colour bar corresponds to the density in points of f_{young}. The red points correspond to f_{AGN}. The power-law fit of Eq. (4) is shown as a black line.

was to construct a dust mass map at a workable resolution. We deviated from the standard recipes provided by Verstocken et al. (2020) as we could not use bands that are dominated by the instrumental PSF signature (notably the MIR). Our model can reproduce the global SED of NGC 1068 and many of the resolved features (see Figs. 4 and 5). However, it is less successful in retrieving the FIR wavebands due to a slight shortage of dust mass towards the centre. Adding more dust in these areas would increase the heating fraction of the AGN and would produce a less favourable fit to the global SED. With these caveats in mind, our model should be seen as a representation of a star-forming galaxy based on NGC 1068. The subsequent dust heating results are therefore still realistic.

Our best-fit radiative transfer model has a dust mass and an SFR (derived from young stellar luminosities), which are in line with a classic CIGALE SED fit (Table 1 and Nersesian et al. 2019). However, major differences are found between the relative importance of the attenuated stellar components. In the CIGALE model, the young stellar population is extremely attenuated in order to boost the MIR emission and compensate for the lack of an AGN in the model. When modelling larger samples of galaxy SEDs it is therefore worthwhile to check the subcomponents of 1D SED models. Signatures similar to the one found here could point to a low-luminosity AGN which requires more tailored SED models. This is an example of how 3D RT models can inform 1D SED models of large samples even though it is not efficient to perform a RT fit for each individual galaxy.

The AGN in our model has a bolometric luminosity of $L_{\mathrm{bol}}^{\mathrm{AGN}} = 3.97 \times 10^{43}\,\mathrm{erg\,s^{-1}}$, which is somewhat lower than previous AGN models for NGC 1068 (e.g. Lopez-Rodriguez et al. 2018). This could in part be due to our assumption of an isotropic point source for the subgrid implementation of the AGN torus. The AGN in NGC 1068 is in any case quite weak for Seyfert 2 galaxies (see e.g. Lusso et al. 2012). It is therefore possible that the reported AGN dust heating fractions are on the low end of the potential spectrum. Still, even for this relatively weak AGN, we find percentage-level dust heating out to 4 kpc from the centre. The AGN heating fraction declines radially following a powerlaw with index -1.19. In the inner 500 pc, the heating

fraction rises quickly above 10% and peaks above 90% in the inner resolution element of our simulation (40 pc). This is also reflected in the contribution of AGN light to each broadband flux in NGC 1068 (Fig. 6). Even in the SPIRE bands there is contamination in the emission due to reprocessed AGN energy.

Globally, we find that most of the dust in the NGC 1068 model is heated by ongoing star formation. The median of this fraction is 83% when all dust cells are considered. This corresponds to a reduction of the L_{dust}-derived SFR from $11.7\,M_{\odot}\,\mathrm{yr^{-1}}$ in the full model to $8.4\,M_{\odot}\,\mathrm{yr^{-1}}$ in the model where only young (non-ionising and ionising) stellar populations heat the dust. These estimates were obtained using the Kennicutt & Evans (2012) conversion from L_{dust} to SFR, which adopts the IMF from Kroupa & Weidner (2003). Relatively speaking, this is a fairly small correction as NGC 1068 is the galaxy with the highest f_{young} in our project sample (containing M 81 and four barred spirals M 83, M 95, M 100 and NGC 1365). This is clearly shown already in Fig. 8 of Nersesian et al. (2020), where we plotted f_{young} as a function of sSFR (the SFR divided by stellar mass). The dust cells of our model lie in line with the general increasing trend that other star-forming galaxies exhibit. Here, SFR and stellar mass were computed per dust cell, already taking into account the fact that not all dust emission arises from the young stellar populations.

Most of the additional energy to heat the dust actually comes from the old stellar populations ∼16% in our model, with only a negligible part from the AGN (on global scales). This is also in line with the global energy balance study of the DustPedia sample (Bianchi et al. 2018), where AGN do not stand out from the rest of this sample of local galaxies. However, Fig. 7 clearly shows that there are significant local differences. In the extreme case, considering the dust cells in the inner 100 pc, the median heating fractions average out at 34% and 28% for the young and AGN heating sources, respectively. The remaining energy can be attributed to the small bulge of old stars. FIR-based SFR indicators or dust mass estimates which include this part of the galaxy are therefore susceptible to significant bias due to both the AGN and the radiation field generated by the old stellar populations.

This study is the third in a series after Verstocken et al. (2020) and Nersesian et al. (2020) where we investigated the effect of dust heating in face-on large galaxies in the DustPedia database. The novelty of this study is that we added an AGN in the model as a source of radiation. We have pushed our models to the limit by requiring high-resolution images from FUV to submm wavelengths, and by fitting 3D radiative transfer simulations to these data. We are thus for the first time able to quantify the dust heating fraction for separate heating sources in a sample of galaxies. Our main conclusion is that star formation alone is inadequate to explain the observed FIR emission. In all cases, a significant contribution from the old stellar population is needed.

In the particular case of NGC 1068, a low-luminosity AGN is also required to match the observations. In systems hosting a stronger AGN or in quasars, the AGN-powered heating of the diffuse dust will be even higher. Hence, simply removing a (fitted) AGN torus component from the FIR flux will still lead to an overestimation of the SFR of the host galaxy.

These findings, in particular the necessity for multiple heating sources, should be considered when building the next generation of galaxy SED modelling tools.

References

Alonso-Herrero, A., Ramos Almeida, C., Esquej, P., et al. 2014, MNRAS, 443, 2766
Aniano, G., Draine, B. T., Calzetti, D., et al. 2012, ApJ, 756, 138
Antonucci, R. R. J., & Miller, J. S. 1985, ApJ, 297, 621
Astropy Collaboration (Robitaille, T. P., et al.) 2013, A&A, 558, A33
Astropy Collaboration (Price-Whelan, A. M., et al.) 2018, AJ, 156, 123
Baes, M., Verstappen, J., De Looze, I., et al. 2011, ApJS, 196, 22
Bendo, G. J., Boselli, A., Dariush, A., et al. 2012, MNRAS, 419, 1833
Bendo, G. J., Baes, M., Bianchi, S., et al. 2015, MNRAS, 448, 135
Bianchi, S. 2013, A&A, 552, A89
Bianchi, S., De Vis, P., Viaene, S., et al. 2018, A&A, 620, A112
Bocchio, M., Jones, A. P., Verstraete, L., et al. 2013, A&A, 556, A6
Bolatto, A. D., Wolfire, M., & Leroy, A. K. 2013, ARA&A, 51, 207
Boquien, M., Calzetti, D., Combes, F., et al. 2011, AJ, 142, 111
Boquien, M., Kennicutt, R., Calzetti, D., et al. 2016, A&A, 591, A6
Boquien, M., Burgarella, D., Roehlly, Y., et al. 2019, A&A, 622, A103
Bruzual, G., & Charlot, S. 2003, MNRAS, 344, 1000
Camps, P., & Baes, M. 2015, Astron. Comput., 9, 20
Casasola, V., Hunt, L., Combes, F., & García-Burillo, S. 2015, A&A, 577, A135
Casasola, V., Bianchi, S., De Vis, P., et al. 2020, A&A, 633, A100
Casey, C. M., Narayanan, D., & Cooray, A. 2014, Phys. Rep., 541, 45
Chabrier, G. 2003, PASP, 115, 763
Ciesla, L., Boquien, M., Boselli, A., et al. 2014, A&A, 565, A128
Clark, C. J. R., Verstocken, S., Bianchi, S., et al. 2018, A&A, 609, A37
Cortese, L., Boselli, A., Franzetti, P., et al. 2008, MNRAS, 386, 1157
da Cunha, E., Eminian, C., Charlot, S., & Blaizot, J. 2010, MNRAS, 403, 1894
D'Agostino, J. J., Poetrodjojo, H., Ho, I. T., et al. 2018, MNRAS, 479, 4907
Dale, D. A., Helou, G., Contursi, A., Silbermann, N. A., & Kolhatkar, S. 2001, ApJ, 549, 215
Dale, D. A., Aniano, G., Engelbracht, C. W., et al. 2012, ApJ, 745, 95
Davies, J. I., Baes, M., Bianchi, S., et al. 2017, PASP, 129, 044102
De Geyter, G., Baes, M., Camps, P., et al. 2014, MNRAS, 441, 869
De Looze, I., Baes, M., Zibetti, S., et al. 2010, A&A, 518, L54
De Looze, I., Fritz, J., Baes, M., et al. 2014, A&A, 571, A69
De Vis, P., Jones, A., Viaene, S., et al. 2019, A&A, 623, A5
di Serego Alighieri, S., Bianchi, S., Pappalardo, C., et al. 2013, A&A, 552, A8
Draine, B. T., Aniano, G., Krause, O., et al. 2014, ApJ, 780, 172
Duras, F., Bongiorno, A., Piconcelli, E., et al. 2017, A&A, 604, A67
Eales, S., Smith, M. W. L., Auld, R., et al. 2012, ApJ, 761, 168
Feltre, A., Hatziminaoglou, E., Fritz, J., & Franceschini, A. 2012, MNRAS, 426, 120
Foreman-Mackey, D., Vousden, W., Price-Whelan, A., et al. 2016, corner.py: corner.py v2.0.0
Foyle, K., Natale, G., Wilson, C. D., et al. 2013, MNRAS, 432, 2182
Fritz, J., Franceschini, A., & Hatziminaoglou, E. 2006, MNRAS, 366, 767
Galametz, M., Madden, S. C., Galliano, F., et al. 2010, A&A, 518, L55
Galametz, M., Kennicutt, R. C., Calzetti, D., et al. 2013, MNRAS, 431, 1956
Galliano, F., Galametz, M., & Jones, A. P. 2018, ARA&A, 56, 673
García-Burillo, S., Combes, F., Ramos Almeida, C., et al. 2016, ApJ, 823, L12
Greenhill, L. J., Gwinn, C. R., Antonucci, R., & Barvainis, R. 1996, ApJ, 472, L21
Griffin, M. J., Abergel, A., Abreu, A., et al. 2010, A&A, 518, L3

Groves, B., Dopita, M. A., Sutherland, R. S., et al. 2008, ApJS, 176, 438
Groves, B., Krause, O., Sandstrom, K., et al. 2012, MNRAS, 426, 892
Hildebrand, R. H. 1983, QJRAS, 24, 267
Howell, J. H., Mazzarella, J. M., Chan, B. H. P., et al. 2007, AJ, 134, 2086
Hughes, T. M., Baes, M., Fritz, J., et al. 2014, A&A, 565, A4
Hughes, T. M., Ibar, E., Villanueva, V., et al. 2017, MNRAS, 468, L103
Jones, A. P. 2004, in Astrophysics of Dust, eds. A. N. Witt, G. C. Clayton, & B. T. Draine , ASP Conf. Ser., 309, 347
Jones, A. P., Köhler, M., Ysard, N., Bocchio, M., & Verstraete, L. 2017, A&A, 602, A46
Kennicutt, R. C., & Evans, N. J. 2012, ARA&A, 50, 531
Kirkpatrick, A., Pope, A., Alexander, D. M., et al. 2012, ApJ, 759, 139
Kirkpatrick, A., Pope, A., Sajina, A., et al. 2015, ApJ, 814, 9
Knapen, J. H., Stedman, S., Bramich, D. M., Folkes, S. L., & Bradley, T. R. 2004, A&A, 426, 1135
Köhler, M., Ysard, N., & Jones, A. P. 2015, A&A, 579, A15
Kroupa, P., & Weidner, C. 2003, ApJ, 598, 1076
Lawrence, A., Warren, S. J., Almaini, O., et al. 2007, MNRAS, 379, 1599
Lisenfeld, U., & Ferrara, A. 1998, ApJ, 496, 145
Lopez-Rodriguez, E., Fuller, L., Alonso-Herrero, A., et al. 2018, ApJ, 859, 99
Lusso, E., Comastri, A., Simmons, B. D., et al. 2012, MNRAS, 425, 623
Madau, P., & Dickinson, M. 2014, ARA&A, 52, 415
Marin, F. 2018, MNRAS, 479, 3142
Martin, D. C., Fanson, J., Schiminovich, D., et al. 2005, ApJ, 619, L1
Mason, R. E., Geballe, T. R., Packham, C., et al. 2006, ApJ, 640, 612
Mentuch Cooper, E., Wilson, C. D., Foyle, K., et al. 2012, ApJ, 755, 165
Nenkova, M., Sirocky, M. M., Ivezić, Ž., & Elitzur, M. 2008, ApJ, 685, 147
Nersesian, A., Xilouris, E. M., Bianchi, S., et al. 2019, A&A, 624, A80
Nersesian, A., Verstocken, S., Viaene, S., et al. 2020, A&A, 637, A25
Osterbrock, D. E., & Martel, A. 1993, ApJ, 414, 552
Pilbratt, G. L., Riedinger, J. R., Passvogel, T., et al. 2010, A&A, 518, L1
Poglitsch, A., Waelkens, C., Geis, N., et al. 2010, A&A, 518, L2
Rémy-Ruyer, A., Madden, S. C., Galliano, F., et al. 2014, A&A, 563, A31
Rémy-Ruyer, A., Madden, S. C., Galliano, F., et al. 2015, A&A, 582, A121
Roebuck, E., Sajina, A., Hayward, C. C., et al. 2016, ApJ, 833, 60
Roebuck, E., Sajina, A., Hayward, C. C., et al. 2019, ApJ, 881, 18
Saftly, W., Baes, M., & Camps, P. 2014, A&A, 561, A77
Sales, D. A., Pastoriza, M. G., & Riffel, R. 2010, ApJ, 725, 605
Sandstrom, K. M., Leroy, A. K., Walter, F., et al. 2013, ApJ, 777, 5
Schartmann, M., Meisenheimer, K., Camenzind, M., Wolf, S., & Henning, T. 2005, A&A, 437, 861
Schneider, R., Bianchi, S., Valiante, R., Risaliti, G., & Salvadori, S. 2015, A&A, 579, A60
Scoville, N., Aussel, H., Sheth, K., et al. 2014, ApJ, 783, 84
Scoville, N., Sheth, K., Aussel, H., et al. 2015, ArXiv e-prints [arXiv:1505.02159]
Skibba, R. A., Engelbracht, C. W., Dale, D., et al. 2011, ApJ, 738, 89
Smith, M. W. L., Eales, S. A., Gomez, H. L., et al. 2012a, ApJ, 756, 40
Smith, M. W. L., Gomez, H. L., Eales, S. A., et al. 2012b, ApJ, 748, 123
Stalevski, M., Fritz, J., Baes, M., Nakos, T., & Popović, L. Č. 2012, MNRAS, 420, 2756
Symeonidis, M. 2017, MNRAS, 465, 1401
Symeonidis, M., Giblin, B. M., Page, M. J., et al. 2016, MNRAS, 459, 257
Verstappen, J., Fritz, J., Baes, M., et al. 2013, A&A, 556, A54
Verstocken, S., Van De Putte, D., Camps, P., & Baes, M. 2017, Astron. Comput., 20, 16
Verstocken, S., Nersesian, A., Baes, M., et al. 2020, A&A, 637, A24
Viaene, S., Baes, M., Bendo, G., et al. 2016, A&A, 586, A13
Viaene, S., Baes, M., Tamm, A., et al. 2017, A&A, 599, A64
Williams, T. G., Baes, M., De Looze, I., et al. 2019, MNRAS, 487, 2753
Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, AJ, 140, 1868
Wu, H., Zhu, Y.-N., Cao, C., & Qin, B. 2007, ApJ, 668, 87
York, D. G., Adelman, J., Anderson, J. E., Jr, et al. 2000, AJ, 120, 1579
Ysard, N., Köhler, M., Jones, A., et al. 2015, A&A, 577, A110

Appendix A: Two-dimensional probability distributions

**Fig. A.1.** Probability distributions of our model parameters (M_{dust}, L_H^{old}, L_{FUV}^{yni}, L_{FUV}^{yi} and $L_{4.6}^{AGN}$) for the second optimization iteration (red histograms in Fig. 3). Darker shades of blue indicate higher probability. The histograms show probability associated with the parameter on the abscissa, with the total probability in the bins normalized to 1.

Figure 3 shows the probability distributions of the individual parameters. To explore the probability space further, we can look at 2D probability distributions. Figure A.1 shows this in the form of a corner plot[3], where the 1D histograms correspond again to the ones in Fig. 3. Because our parameter grid only samples five values per parameters, the 2D probability distributions are only sparsely sampled, making it ineffective to compute probability contours. It is therefore difficult to draw strong conclusions from this plot, however some features can be noted.

The luminosity of the old stellar component (L_H^{old}) is well constrained with no significant link to the other parameters. We also do not see any correlation between the AGN luminosity and the other properties, although it is evident that the probabilities for $L_{4.6}^{AGN}$ are spread out more evenly across the parameter space. This is different for M_{dust}, L_{FUV}^{yni} and L_{FUV}^{yi}, which do show vague correlations among each other. This points to a mild degeneracy between the UV luminosities of the young stars and the

---

[3] Produced using the python `corner` package (Foreman-Mackey et al. 2016).

dust component. This is not surprising because they influence the observed UV flux in opposite ways. In principle, this degeneracy could be broken by perfect sampling of the infrared SED. As we only sparsely sample this regime, some residual correlation is possible between these properties. However, we can still sufficiently constrain each parameter individually.

Appendix B: Alternative dust map

Our standard model for NGC 1068 includes a dust component that is based on a 2D dust mass surface density map. This map was derived from the per-pixel A_{FUV} estimate following the recipe of Cortese et al. (2008). This recipe is based on TIR/FUV and FUV/r colours. It is possible that the AGN affects these measurements, resulting in a less reliable A_{FUV} estimate in the centre of the galaxy. For this reason we ran the fitting procedure with an alternative dust surface density map (see Fig. B.1).

As the standard model lacks dust emission in the centre of the galaxy, we constructed a dust map that is strongly peaked in the centre. To do so we normalise the PACS 70 μm flux density map, which is a good tracer of the total infrared emission (Galametz et al. 2013). At this wavelength, we are more sensitive to warmer dust, but we importantly attain sufficient spatial resolution ($\sim$300 pc) in our model. If we were to derive a dust mass map based on the SPIRE bands or based on an modified black body fit to the *Herschel* data, the effective resolution would quickly drop to 1 kpc. The PACS 70 μm based dust map is significantly different from the A_{FUV} based dust map, which at least allows a qualitative assessment between both models.

In Fig. B.2 we compare the observed and model images of the alternative model. This figure can be compared to Fig. 5 for the standard model. We see a roughly opposite trend compared to the standard model. Instead of a central dip in the dust emission, the emission now peaks in the center, which was the goal of this new model. However, the central peak is too strong and overestimates the FIR observations. Consequently, the attenuation in the UV and optical bands is severe and causes a strong underestimation of the flux in the central regions. At the same time the attenuation in the spiral arms is not strong enough causing a large negative residuals there. The KDE distributions shown in the right column of Fig. B.2 as a result show broad and often double-peaked residuals. We conclude that the alternative model is highly unbalanced and not suitable for further analysis.


Fig. B.1. Two-dimensional representations (on the plane of the sky) of the dust mass surface density used as input morphologies for the 3D dust distribution. The *top panel* is the same A_{FUV} based map as in Fig. 1 but on a logarithmic colour scale. The *bottom panel* is the alternative dust mass map based on the PACS 70 μm image, which peaks strongly in the centre.

Fig. B.2. Same as Fig. 5 but for the alternative dust mass map. Spatial comparison of the model to observed broad-band images (*left column*). The *second column* contains the model images. We note that the model images are not convolved by any PSF and only hold the intrinsic model PSF (a composite of the input images). Residual images are shown in the *third column*. Corresponding residual distributions are represented by KDE plots in the *right column*. The area under the KDE curves is normalised to 1.