MODELING GALACTIC EXTINCTION WITH DUST AND "REAL" POLYCYCLIC AROMATIC HYDROCARBONS

Giacomo Mulas[1], Alberto Zonca[2], Silvia Casu[1], and Cesare Cecchi-Pestellini[1]
[1] INAF–Osservatorio Astronomico di Cagliari, Strada n.54, Loc. Poggio dei Pini, I-09012 Capoterra, Italy;
gmulas@oa-cagliari.inaf.it, silvia@oa-cagliari.inaf.it, ccp@oa-cagliari.inaf.it
[2] Dipartimento di Fisica, Università di Cagliari, Strada Prov.le Monserrato–Sestu Km 0.700, I-09042 Monserrato (CA), Italy; azonca@oa-cagliari.inaf.it

ABSTRACT

We investigate the remarkable apparent variety of galactic extinction curves by modeling extinction profiles with core-mantle grains and a collection of single polycyclic aromatic hydrocarbons. Our aim is to translate a synthetic description of dust into physically well-grounded building blocks through the analysis of a statistically relevant sample of different extinction curves. All different flavors of observed extinction curves, ranging from the average galactic extinction curve to virtually "bumpless" profiles, can be described by the present model. We prove that a mixture of a relatively small number (54 species in 4 charge states each) of polycyclic aromatic hydrocarbons can reproduce the features of the extinction curve in the ultraviolet, dismissing an old objection to the contribution of polycyclic aromatic hydrocarbons to the interstellar extinction curve. Despite the large number of free parameters (at most the 54×4 column densities of each species in each ionization state included in the molecular ensemble plus the 9 parameters defining the physical properties of classical particles), we can strongly constrain some physically relevant properties such as the total number of C atoms in all species and the mean charge of the mixture. Such properties are found to be largely independent of the adopted dust model whose variation provides effects that are orthogonal to those brought about by the molecular component. Finally, the fitting procedure, together with some physical sense, suggests (but does not require) the presence of an additional component of chemically different very small carbonaceous grains.

Key words: dust, extinction – ISM: abundances – ultraviolet: ISM

Online-only material: color figures, figure set, machine-readable tables

1. INTRODUCTION

For most of the 20th century, the presence of interstellar dust mixed with gas has been a widely accepted concept in astronomy. However, only over the last half century has it become increasingly accepted that dust has many very important roles to play in astronomy, and that dust is a crucial component of the Milky Way and other galaxies. For astrophysicists to understand and describe these roles accurately, it is necessary to know not only the composition of dust grains but also their detailed structure (e.g., Draine 2003).

The most frequently used observational data are those of the wavelength-dependent extinction curves along the lines of sight to individual stars. There is a basic similarity in the shapes of these curves: a general rise in normalized extinction from low values in the infrared to high values in the far-ultraviolet, a near linear portion in the optical region, and a pronounced and broad "bump" near 217.5 nm. While much attention has been given to the average InterStellar Extinction Curve (ISEC), it is evident that there is considerable variation from one line of sight to another (e.g., Fitzpatrick & Massa 2007, 2009; Gordon et al. 2009). Both the general behavior of the average curve and the details of specific curves along particular lines of sight contain useful information and generally indicate that grains of a wide range of sizes, roughly nanometers to microns, are required. Other important indicators of the nature of interstellar dust include the pattern of elemental depletions, the distribution of the linear polarization of starlight, the properties of scattered light, absorption and emission features, broadband emission in the visible and infrared regions, and polarization of starlight (Williams 2000). Although there is considerable uncertainty

about the physical nature of dust particles, in particular their shape and morphology (e.g., Wright 1987; Mathis & Whiffen 1989; Iatí et al. 2001), it is clear from the spectroscopic evidence, and consistent with the depletion data, that important components of interstellar dust are chemically in the form of silicates, predominantly amorphous but with a crystalline component (see Li & Draine 2001 for a summary and Li 2009), and carbonaceous materials, including C atoms with both graphitic, sp^2, and polymeric, sp^3, valencies (e.g., Jones et al. 1990). Carbonaceous materials appear to be amorphous, though there is some evidence from collected interstellar dust particles that some crystalline graphite may also be present.

Current dust models are based on the Mathis et al. (1977) model. Those authors assumed the existence of two distinct types of interstellar dust: graphite grains and silicate grains. All grains were assumed to be spherical and to have a continuous distribution of radii weighted by a power law in size with a negative exponent, so that there are many more small grains in number compared with large ones. Such a power-law size distribution of dust particles has been widely used, and universally referred to as an "MRN distribution (Mathis–Rumpl–Nordsieck)." The calculations were performed using the analytical Mie theory and this model was successful in reproducing the average ISEC in the optical and ultraviolet very well. Many of the distinctive features of the Mathis et al. (1977) model are retained in more sophisticated models.

A model with significant differences from the Mathis et al. (1977) model (and subsequent manifestations) has been developed by researchers at Cagliari and Messina (Iatí et al. 2008; Cecchi-Pestellini et al. 2008). The components of this model include astronomical silicates, solid carbon, and polycyclic

aromatic hydrocarbons (PAHs). The carbon is assumed to be hydrogenated amorphous carbon containing both sp^2 (H-poor graphitic) and sp^3 (H-rich polymeric) bondings. The observed spatial variations in dust extinction are interpreted as a natural evolution of dust optical properties with time in response to local physical conditions (Cecchi-Pestellini & Williams 1998; Cecchi-Pestellini et al. 2010). Carbon (other than the PAHs) is assumed to be deposited on the silicate cores rather than in a separate population of carbon grains. Such a description was suggested by the work of Jones et al. (1990), who argued that there is a cycle of carbon in the interstellar medium; initially, the silicate cores are assumed to be bare, then carbon is deposited kinetically from the gas with a rate depending on gas density and temperature. The resulting solid carbon layer is assumed, at the deposition, to be hydrogen-rich and sp^3 bonded, and is converted to sp^2 bonding at a rate, the so-called photodarkening rate, proportional to the intensity of the ambient radiation field; thus, the outer layer of the mantle is sp^3 and the inner layer is sp^2.

The model assumes distinct populations of small and large grains, separated by a gap in grain sizes. This feature is shared with some other dust descriptions that explicitly consider separate populations for "big" (bare or coated) silicate grains and carbonaceous "very small grains." The assumption of missing grain sizes is also somewhat implicitly present in models in which the distribution of grain sizes, even if formally continuous, shows a deficiency of medium-sized grains (see, e.g., the size distribution of Li & Draine 2001 and Zubko et al. 2004). The Cagliari/Messina core-mantle ($[CM]^2$) model, in its simplified form in which PAHs are represented by two Lorentzian profiles, was proven to be able to fit a wide variety of extreme Milky Way extinction curves (Zonca et al. 2011) as well as some unusual extinction curves from external galaxies. This simplified version of the $[CM]^2$ model has 15 free parameters (9 for the classical dust grains; 6 for the 2 Lorentzians).

Zonca et al. (2011) successfully modeled a number of Galactic extinction curves whose shapes are extremely different from the average ISEC. The main result of this work is that the so-called peculiar curves do not need exotic properties of dust grains to be accurately described. The authors concluded that the wide dispersion in the ISEC morphology may merely reflect the grain size distributions and the chemical response of dust materials to the environmental conditions along their sightlines, within the general circulation of matter into and out of clouds and stars. Indeed, such a conclusion is implicitly embedded in the general observational result that ISECs can be phenomenologically described by means of a very limited number of empirical parameters (Cardelli et al. 1989; Valencic et al. 2004; Fitzpatrick & Massa 2007). Fitzpatrick & Massa (2007), and previous works referenced therein, showed that all ISECs can be accurately described by a six-parameter empirical representation. Cardelli et al. (1989) showed that (with few notable exceptions) these six parameters can be expressed as (empirically determined) functions of a single free one.

In this work we use the $[CM]^2$ model, both in its simplified form and in its version including "real" PAH cross-sections (see Section 2) to systematically fit the *whole* sample of galactic extinctions provided by Fitzpatrick & Massa (2007). We do not mean to assess whether this model is better than other current alternatives, which might reproduce the same data equally well; we would actually welcome similar systematic fitting efforts with other dust models as this would help to separate model-independent trends from model-dependent ones. This work aims at reconciling all the observed varieties of ISECs within

a unified model of interstellar dust to unfold the synthetic description of dust into physically well-grounded properties. This is not a trivial task since, given the complexity of the underlying physics and chemistry, *any* physical dust model requires a relatively large number of parameters, typically not fewer than ∼12 even despite drastic simplifying assumptions. To try to understand how and why the physical parameters defining the populations of interstellar dust and macromolecules may respond to a smaller number of parameters defining the environmental conditions, one must begin by determining the former ones on a large statistical sample of observations so that further analysis becomes possible. One must do so in a way that is

1. *homogeneous*: using the same model and assumptions, so that variations between lines of sight are less masked by systematic effects;
2. *realistic*: using a model simplified enough to be tractable, yet still containing enough physical information to be worth the effort; and
3. *validated*: able to match all observations within observational errors.

This work represents a first step in this direction. We associate a set of well-defined *physical* dust parameters to each line of sight in a statistically relevant sample of different ISECs, with a solid analysis of how tightly constrained (or not) they are by the fit. This makes the search for trends and correlations with other potentially related observable quantities possible.

We present the model and the input data construction in Section 2. In this section we also report a detailed discussion of the fitting technique (Section 2.3). This discussion is relevant to understanding why and how one can squeeze significant results out of an underdetermined fit and how we determine their accuracy. Uninterested (and trusting) readers can skip this section as well as the Appendix, containing some more mathematical details. An extensive collection of results is reported in Section 3, while the final paragraph contains discussion, integration, and combination of data into a synthesis essay.

2. THE MODEL

The model being used here includes two main components: a population of silicate core, carbon-coated dust grains, and a "molecular" component of free-flying PAHs. This second component can be represented in two ways: either in a simplified way as a superposition of two Lorentzians constructed to match the average spectral properties of PAHs or in a detailed way by an actual linear combination of the PAH ab initio cross-sections. The descriptions of these components are detailed below.

2.1. Dust Grains

Model dust grains are composed of four concentric spherical layers: a central void within a silicate shell and the core, embedded in a mantle composed of an inner sp^2 carbon layer and an outer sp^3 layer. The optical constants of the various materials that we have used are: Draine (1985) for the silicate core, Rouleau & Martin (1991), Be1 sample, for the sp^2 amorphous carbon, and Ashok et al. (1991) for the sp^3 polymeric carbon. We shall discuss this choice of optical constants in Section 4.

In order to calculate the optical properties of stratified spheres, we use the extension of the Mie theory to radially non-homogeneous spheres put forward by Wyatt (1962) as modified by Borghese et al. (1987) exploiting the *T*-matrix approach (Waterman 1971). In the region external to the sphere

the field is the superposition of incident and scattered fields. Both fields are expanded in a series of spherical vector multipole fields (Rose 1957). The expansion of the incident field has to be regular at its origin while the expansion of the scattered field has to satisfy the appropriate radiation conditions at infinity (Jackson 1975). The theory has been reported elsewhere (Cecchi-Pestellini et al. 2005; Iatì et al. 2008). Details for the general case are found in Borghese et al. (2007). We truncate the multipole expansions following the convergence criterion suggested by Wiscombe (1980). However, we carefully test the convergence by monitoring the changes in the quantities of interest throughout the computations.

The carbonaceous mantle thickness w, while being a variable fit parameter of the extinction model, is assumed to be the same over the whole range of grain sizes, to mimic mantle accretion in the interstellar medium (Whittet 2002). Likewise, the ratio of the inner cavity volume to the total silicate core volume (the vacuum fraction in volume) is also assumed, somewhat arbitrarily, to be a unique, single-fit parameter, independent of grain size. For dust grain radii we adopt a power-law size distribution $1/(a + w)^q$, where a is the outer silicate core radius, so that $(a+w)$ is the total radius of the dust particle. Such distribution allows for a "gap" in particle sizes, so that two populations of dust grains, "big" and "small" ones, may be present, each characterized by lower and upper size limits in the distribution. There is no qualitative difference between these two populations whose components have the same structure described by the same parameters.

Summing up, the free parameters of the dust grain component of the extinction model are: a_- the lower limit of the "small" grain component, a_+ the upper limit of the "small" grain component, b_- the lower limit of the "big" grain component, b_+ the upper limit of the "big" grain component, q the exponent of the MRN power-law size distribution, f_v the vacuum fraction of silicate core, w the total carbonaceous mantle thickness, f_{sp^2} the fraction of sp^2 material of the carbonaceous mantle width, and N_d the normalization factor of the size distribution proportional to the total amount of dust grains.

The above parameters are subject to a number of physical constraints, namely, $5\,\text{nm} \leqslant a_- \leqslant a_+ \leqslant b_- \leqslant b_+$, $0 \leqslant f_v \leqslant 1$, $0 \leqslant f_{sp^2} \leqslant 1$, $0\,\text{nm} \leqslant w \leqslant 3\,\text{nm}$, and $N_d \geqslant 0\,\mu\text{m}^{(q-3)}$. We impose an upper limit of $w \leqslant 3\,\text{nm}$ on the mantle thickness in order to avoid an unreasonable amount of carbon atoms locked in dust mantles.

The present dust model, including only spherically symmetric particles, does not account for polarization effects, which require some kind of anisotropy. However, the purpose of this paper is to consider accurately the effect of layers with different optical properties on a grain, and to consider such an effect in the description of observed ISECs. Several methods to solve the difficult problem of the calculation of the optical properties of non-spherical grains are present in the literature (e.g., Saija et al. 2001). We might also apply a correction, e.g., to arbitrarily assume that dust grains are approximately oblate spheroids (Mathis 1996). Still, the shape of dust grains is unknown except to the extent that it departs from a sphere. Thus, adding any correction for non-sphericity would imply some additional arbitrary choice, while adding little in the way of the scope of this paper, i.e., modeling extinction curves and relating their variations to physical quantities in a uniform way. We therefore choose to stick to spherical symmetry for classical dust grains in this work. The features of dust grain models including spherical/non-spherical, homogeneous/inhomogeneous particles have been recently discussed by Voshchinnikov (2012).

2.2. The Molecular Component

The presence of a molecule-sized component of carbonaceous dust was inferred as far back as when the first observations of the 3.3 μm emission band became available. Sellgren (1984) showed that this emission would require unrealistically large temperatures if arising from dust in thermal equilibrium. Soon afterward, Lèger & Puget (1984) and Allamandola et al. (1985) discovered that the whole group of correlated emission bands at 3.3, 6.2, 7.7, 8.6, and 11.3 μm matched the main infrared bands of the family of PAHs. These molecules are known to efficiently convert electronic excitation, due to the absorption of UV-visible photons, into vibrational excitation, and then IR emission, and are very photostable. Their contribution to extinction is energetically consistent with the observed emission in the above set of IR bands (hereafter known as the Aromatic Infrared Bands—AIBs). From then on, PAHs have been considered a necessary ingredient of any interstellar dust model (see, e.g., Draine & Li 2001). Nonetheless, to date, no specific PAH has been identified with certainty. AIBs are highly non-specific, being essentially due to vibrations of functional groups inside the molecules. Far-IR features involving vibrations of the whole molecule can distinguish individual species, but none were detected so far in *Herschel* data. Similarly, $\pi^\star \leftarrow \pi$ electronic transitions of neutral PAHs can be narrow and specific enough to permit an identification, but none were detected despite targeted observations (see, e.g., Clayton et al. 2003). The joint constraints that PAHs as a family must be abundant enough to produce the observed AIB intensities (containing at least $\sim$10% of the available interstellar carbon atoms), yet that no single one be abundant enough for its weak individual spectral signature to stand out, were interpreted to point to a very large chemical diversity with (at least) hundreds of different PAH species in the interstellar mixture.

2.2.1. Simplified Version

The contribution of PAHs to extinction is represented in the simplified version of the $[\text{CM}]^2$ model by two Lorentz profiles, the first at lower energies mimicking the $\pi^\star \leftarrow \pi$ resonance in PAHs, accounting for most of the bump in the extinction, the second mimicking the low-energy tail of the $\sigma^\star \leftarrow \sigma$ plasmon resonance, producing the far-ultraviolet nonlinear extinction rise. Since the electronic photo-absorption cross-sections of PAHs fall to zero more quickly than Lorentzians at low energies, Iatì et al. (2008) also introduced an exponential damping factor on the low-energy tail of the $\sigma^\star \leftarrow \sigma$ resonance. Details can be found in Zonca et al. (2011). For a characterization of the molecular component we rely on the database of spectral properties of PAHs computed by Malloci et al. (2007), to which we refer for details, including 54 molecules in the size range 10–66 carbon atoms, and charge states 0, $\pm$1, and +2.

A simple physical interpretation of the cutoff in terms of properties such as size and charge distribution of the undetermined mixture of PAHs is given in Duley (2006). Guided by quantum-chemical calculations performed by Malloci et al. (2007), we slightly change the exponential damping factor to $1/\{1+\exp[(C_1 - \lambda^{-1})/C_2]\}$, in order to remove any contribution of the electronic photo-absorption by the molecular component to the extinction longward of visual wavelengths. The free parameters of the simplified molecular component are therefore Li_1 the integral of the first Lorentzian, Lp_1 the peak wavenumber of the first Lorentzian, Lw_1 the FWHM of the first Lorentzian,

Li_2 the integral of the second Lorentzian, Lp_2 the peak wavenumber of the second Lorentzian, Lw_2 the FWHM of the second Lorentzian, C_1 the center of the exponential cutoff, and C_2 the width of the exponential cutoff.

These parameters also are constrained to be physically sensible, i.e., $Li_1 \geqslant 0$, $Li_2 \geqslant 0$, $Lw_1 \geqslant 0\,\mu^{-1}$, $Lw_2 \geqslant 0\,\mu^{-1}$, $Lp_1 \leqslant Lp_2$, $C_2 \geqslant 0\,\mu^{-1}$.

2.2.2. Detailed Version

In the detailed version of the extinction model, instead of the molecular component being represented by two Lorentzian profiles parameterized to fit the average of the ab initio photo-absorption cross-sections of a sample of 54 PAHs in 4 charge states (see above), it is represented by a straightforward linear combination of the same cross-sections. In this case, therefore, the formal free parameters become the 216 individual column densities of each PAH in each charge state considered, each of them subject to the obvious physical constraint of being non-negative. This is a more demanding constraint than it might appear at first sight: from a mathematical point of view, it reduces the available space for these parameters by a factor of 2^{216}; from a physical point of view, it means that it is impossible to "subtract away" an unobserved absorption band of a molecule with a negative amount of other molecules.

2.3. Fitting Details

The data to be fit are taken from the published data of Fitzpatrick & Massa (2007). Such data consist of tabulated extinction curves, expressed in units of normalized color excess $E_{\lambda-V}/E_{B-V}$, and the corresponding tabulated errors. Since the extinction model used here is expressed in units of normalized total extinction A_λ/A_V, we first convert all tabulated data to these units, using the $R_V = A_V/E_{B-V}$ values given in the same work.

We then use a slightly modified version of the C MPFIT[3] implementation of the Levemberg–Marquardt nonlinear least-squares algorithm to actually perform the fits, with both the simplified and the detailed versions of the extinction model. The iterative fit procedure seeks to minimize the numerical χ^2 starting from a given (hopefully sensible) first guess of the free parameters, making sure all parameter constraints are respected. The apparent drawback of this kind of procedure is that the algorithm will fall in the nearest minimum which it can reach by descending the χ^2 hypersurface in the free parameters from the starting point. If the hypersurface is ill-behaved enough, the algorithm may get stuck on a local minimum, with a much larger χ^2 than the global one, or even on a saddle point. The curvature of the χ^2 hypersurface at the minimum can be used, in principle, to yield information on the covariance matrix of the parameter values derived from the fit. This evaluation of the covariance is indeed returned as one of the results of the MPFIT function. However, this is a good approximation only if the second-order Taylor expansion of the χ^2 hypersurface around its minimum $\chi^2_{\min}$ is accurate for the entire region around the minimum in which $\chi^2 \leqslant \chi^2_{\min} + 1$.

Since this is not necessarily true, and indeed the covariance values returned by MPFIT are sometimes clearly not sensible in our case, we use another strategy, which is useful both to provide a robust evaluation of the covariance matrix of the fitted parameters and to get out of local minima and saddle points.

After performing a successful fit on one of the original data samples, we perturb it by a random Gaussian noise proportional to the errors, also taken from Fitzpatrick & Massa (2007), and repeat the fit on the perturbed data. The procedure is iterated to obtain some synthetic statistics. In some cases, the fit to the perturbed data is a better match even for the unperturbed data, meaning that the first fit ended on a local minimum or saddle point. In that case, the whole procedure is restarted using the new "best guess" as the initial point. The synthetic statistics on the fitting parameters are then used to directly estimate their covariance matrix. In some cases, some of the parameters are constrained by the fit much more tightly than others (i.e., the derivative of χ^2 with respect to one parameter was much smaller than that with respect to another) that the formal variance of the better constrained parameter results in zero, i.e., the fit converged on the same value of that parameter in all perturbed fits, within numerical accuracy. This does not mean that this parameter is really determined with infinite accuracy, but merely that its error is below the capability of the algorithm to evaluate it. The random error we introduce in perturbed fits is completely uncorrelated, i.e., adjacent pixels have the same probability of being perturbed in the same as in opposite directions. This implies that our error estimation cannot account for any possible *systematic* observational error in some observed extinction curve, whereby, e.g., the whole visible part of the ISEC was overestimated. Therefore, in such cases, we would somewhat underestimate the errors in the fit parameters.

Fits using the detailed extinction model pose an additional problem, which required a slight modification of the MPFIT implementation. As detailed in the Appendix, the Levemberg–Marquardt method proceeds by linearizing the non-linear least-squares problem around the previous guess of the parameter values. It builds and solves a system of linear equations (with some clever modifications to ensure robust and rapid convergence) to obtain the next improved guess of the fitting parameters. If the linear system is quasi-singular,[4] its solution is numerically unstable and the iterative step may produce useless results with the standard MPFIT implementation. Such behavior reflects a real problem, i.e., the fit is strongly underconstrained when the detailed extinction model is used. The implication is not that the fit is without solution, but rather that the solution is far from unique. In particular, there is a large subspace of the parameter space that fits each extinction curve being fitted. In other words, there is not a unique mixture of the PAHs in our sample that best fits a given extinction curve, but a *family* of PAH mixtures that all fit it very nearly equally well. This entails two problems. First, there is the practical, numerical problem of finding one representative of this family of solutions, overcoming the instability of the standard Levemberg–Marquardt implementation in MPFIT. This is addressed by solving the linear problem in each iteration via singular value decomposition (SVD), and discarding "singular" values at, or below, the level of numerical noise before inverting the matrix to solve the linear problem. In such a way we provide a robust method to converge on *one* solution of the family. The results are remarkably insensitive to the precise threshold one chooses to accept or discard singular values. The minimum number of singular values that must be retained without adversely affecting the quality of the fit is between 20 and 30. Keeping many more than 30 includes some that are just noise; keeping less than 20 discards too much

---

[3] http://www.physics.wisc.edu/~craigm/idl/cmpfit.html

[4] It cannot be *really* singular due to the way the Levemberg–Marquardt method builds it, but it can be *very close* to singular.

Table 1
Classical Dust Component Parameters, Obtained from the Simplified and Detailed Models for the Average ISEC and for the Lines of Sight Plotted in Figure 1

LoS	N_d $(\mu m^{(q-3)})$	f_v	w (nm)	f_{sp2}	a_- (nm)	a_+ (nm)	b_- (nm)	b_+ (nm)	q	χ^2_{red}
Average ISEC	1.013 (0.016)	0.45 (0.02)	1.3 (0.5)	0.98 (0.05)	5.10 (0.24)	11.9 (2.6)	63.3 (2.3)	491 (12)	3.469 (0.002)	$<\varepsilon_4$
	0.999711 ($<\varepsilon_1$)	0.45 ($<\varepsilon_2$)	0.97 ($<\varepsilon_2$)	0.92 ($<\varepsilon_2$)	5.00 ($<\varepsilon_2$)	27.00 ($<\varepsilon_2$)	70.40 ($<\varepsilon_2$)	493.06 ($<\varepsilon_2$)	3.4727 ($<\varepsilon_3$)	$<\varepsilon_4$
HD 12323	1.152 (0.011)	0.44 (0.01)	0.95 (0.23)	0.5 (0.3)	5.00 ($<\varepsilon_2$)	13.2 (1.3)	56.8 (2.8)	465 (6)	3.4487 (0.0013)	$<\varepsilon_4$
	1.128 (0.022)	0.48 (0.03)	0.75 (0.13)	0.67 (0.13)	5.00 ($<\varepsilon_2$)	42 (6)	133 (28)	492 (11)	3.48 (0.01)	$<\varepsilon_4$
HD 13659	0.966 (0.004)	0.47 (0.01)	0.67 (0.06)	0.87 (0.06)	5.00 ($<\varepsilon_2$)	14.45 (0.08)	55.2 (1.2)	479.2 (2.7)	3.4664 (0.0004)	0.4
	1.0246 (0.0018)	0.36 ($<\varepsilon_2$)	0.51 ($<\varepsilon_2$)	0.00 ($<\varepsilon_2$)	5.00 ($<\varepsilon_2$)	55.5 (0.4)	129.4 (1.8)	485.99 (0.24)	3.4945 (0.0005)	0.2
HD 147165	0.9978 (0.0022)	0.32 ($<\varepsilon_2$)	3.00 ($<\varepsilon_2$)	1.00 (0.01)	5.05 (0.05)	13.05 (0.23)	67.96 (0.19)	517.5 (1.3)	3.4714 (0.0004)	0.1
	0.998617 ($<\varepsilon_1$)	0.37 ($<\varepsilon_2$)	3.00 ($<\varepsilon_2$)	0.78 ($<\varepsilon_2$)	6.37 ($<\varepsilon_2$)	13.65 ($<\varepsilon_2$)	65.00 ($<\varepsilon_2$)	566.37 ($<\varepsilon_2$)	3.4717 ($<\varepsilon_3$)	0.2
HD 151346	1.024 (0.007)	0.35 (0.01)	1.55 (0.14)	1.00 ($<\varepsilon_2$)	5.38 (0.05)	7.8 (0.4)	69.4 (0.8)	515.0 (3.0)	3.4659 (0.0007)	0.1
	0.969 (0.004)	0.27 (0.01)	0.98 (0.04)	0.91 (0.01)	5.00 ($<\varepsilon_2$)	26.1 (0.7)	107.7 (1.5)	492.2 (1.0)	3.4874 (0.0015)	0.1
HD 185418	0.968582 ($<\varepsilon_1$)	0.50 ($<\varepsilon_2$)	1.00 ($<\varepsilon_2$)	0.50 ($<\varepsilon_2$)	8.00 ($<\varepsilon_2$)	20.00 ($<\varepsilon_2$)	70.00 ($<\varepsilon_2$)	459.00 ($<\varepsilon_2$)	3.3948 ($<\varepsilon_3$)	0.9
	0.933621 ($<\varepsilon_1$)	0.52 ($<\varepsilon_2$)	3.00 ($<\varepsilon_2$)	0.07 ($<\varepsilon_2$)	5.00 ($<\varepsilon_2$)	24.97 ($<\varepsilon_2$)	85.87 ($<\varepsilon_2$)	470.13 ($<\varepsilon_2$)	3.3957 ($<\varepsilon_3$)	0.1
HD 197702	0.852 (0.007)	0.39 (0.02)	3.00 ($<\varepsilon_2$)	1.00 ($<\varepsilon_2$)	5.00 ($<\varepsilon_2$)	17.1 (0.4)	55.7 (1.7)	523.8 (2.7)	3.4930 (0.0011)	0.8
	0.864020 ($<\varepsilon_1$)	0.09 ($<\varepsilon_2$)	3.00 ($<\varepsilon_2$)	1.00 ($<\varepsilon_2$)	5.10 ($<\varepsilon_2$)	28.00 ($<\varepsilon_2$)	45.41 ($<\varepsilon_2$)	567.33 ($<\varepsilon_2$)	3.5098 ($<\varepsilon_3$)	0.1
HD 210121	1.198 (0.011)	0.43 (0.03)	0.00 ($<\varepsilon_2$)	1.00 ($<\varepsilon_2$)	5.01 (0.02)	11.7 (0.7)	65.1 (0.9)	400 (10)	3.4420 (0.0026)	0.1
	1.860 (0.009)	0.46 (0.03)	0.61 (0.06)	0.21 (0.12)	5.00 ($<\varepsilon_2$)	47.2 (1.1)	252.6 (2.4)	386.8 (2.4)	3.4501 (0.0009)	$<\varepsilon_4$
VCT10	0.970 (0.029)	0.44 (0.04)	1.24 (0.27)	0.94 (0.07)	5.00 ($<\varepsilon_2$)	13.3 (1.2)	74 (6)	600 (60)	3.487 (0.007)	1.4
	0.9854 (0.0016)	0.33 (0.01)	0.45 (0.02)	0.66 (0.02)	5.00 ($<\varepsilon_2$)	36.34 (0.17)	160 (4)	561.1 (1.3)	3.5054 (0.0015)	0.4

Note. We adopt $\varepsilon_1 = 0.000001$, $\varepsilon_2 = 0.01$, $\varepsilon_3 = 0.0001$, $\varepsilon_4 = 0.1$.

(This table is available in its entirety in a machine-readable form in the online journal. A portion is shown here for guidance regarding its form and content.)

real information. Synthetic statistics then provide a way to obtain a *sample* of the family of solutions since each perturbed fit converges to a different individual solution.

The second problem is related to the physical interpretation of the results: while it is impossible to constrain the precise composition of PAHs using only the extinction curve, we can instead hope to constrain rather precisely some of the PAH properties as a group (e.g., total amount of C atoms in PAHs, charge per C atom, and so on). In a way, the minimum of the χ^2 hypersurface looks like a trough in the PAH parameters, which is shallow or flat in many directions (representing PAH properties poorly constrained by the fit), but steep in some (representing, on the contrary, properties of the PAH family strongly constrained by the fit). A hint about the number of global parameters well constrained by the fit can be obtained from the minimum number of singular values in the SVD that need to be retained (i.e., 20–30). The 9 parameters of classical dust grains are all rather well constrained by the fit, so this leaves $\lesssim 15$ parameters defining the PAH mixture (out of the 216 column densities). So, while on the one hand this clearly reveals that the information one can hope to attain by such a huge and cumbersome fit is limited, on the other hand, the fact that it can constrain some *global* properties of the PAH mixtures fitting the extinction lends credibility to these results, making them far less dependent on the specific species used for the fit. Indeed, almost the same results are obtained even if one uses half of the 54 PAHs chosen randomly or worse (see discussion in Section 3.4). In any case, the physical problem now becomes that of linking these relatively few global PAH parameters to meaningful physical/chemical properties characterizing the mixtures fitting the extinction. This is a task that simplified descriptions of PAHs are unable to address since molecular properties and electronic structures are smeared out into a homogeneous mixture.

3. RESULTS AND DISCUSSION

Figure 1 shows a small, representative sample of the extinction curves and their fits using the detailed model. They are chosen to represent all different "flavors" of observed extinction curves, ranging from what is considered the "average" galactic extinction curve to virtually "bumpless" ones, similar to the ones frequently observed in galaxies at high redshifts, or in parts of the Magellanic Clouds. Moreover, the eight lines of sight in Figure 1 present widely varying relative intensities of the different components of the Fitzpatrick & Massa (2007) parameterization, as well as different bump widths. Indeed, we chose by and large the most extreme cases we could find. The observed extinction curves are represented in Figure 1 by the dotted lines, the fitted curves obtained with our model (detailed version) are the continuous lines, and the error range reported by Fitzpatrick & Massa (2007) is the lightly shaded area in each plot. In the same plots, the dashed lines show the contribution of the classical dust component in our model; the dash-dotted lines show the contribution from the mixture of PAHs.

In all cases the fitted extinction curve is virtually indistinguishable from the observed one. The fitted extinction curves are also remarkably "smooth," with no detectable fine structure. While only eight lines of sight are shown here, this is the rule for *all* cases, with no exception (see the left panel in Figures 11.1 to 11.329 in the online version of the journal). This unambiguously proves that a mixture of a relatively small number of PAHs *can* reproduce the features of the extinction curve in the ultraviolet, putting to rest an old objection to the contribution of PAHs to the ISEC. While PAHs are not the *unique* solution to represent ultraviolet extinction features (and we do not claim they are), they are proven to be a viable one. In Table 1 for each line of sight we report the fitted values of the parameters of the classical dust component, resulting from both the simplified and detailed versions.
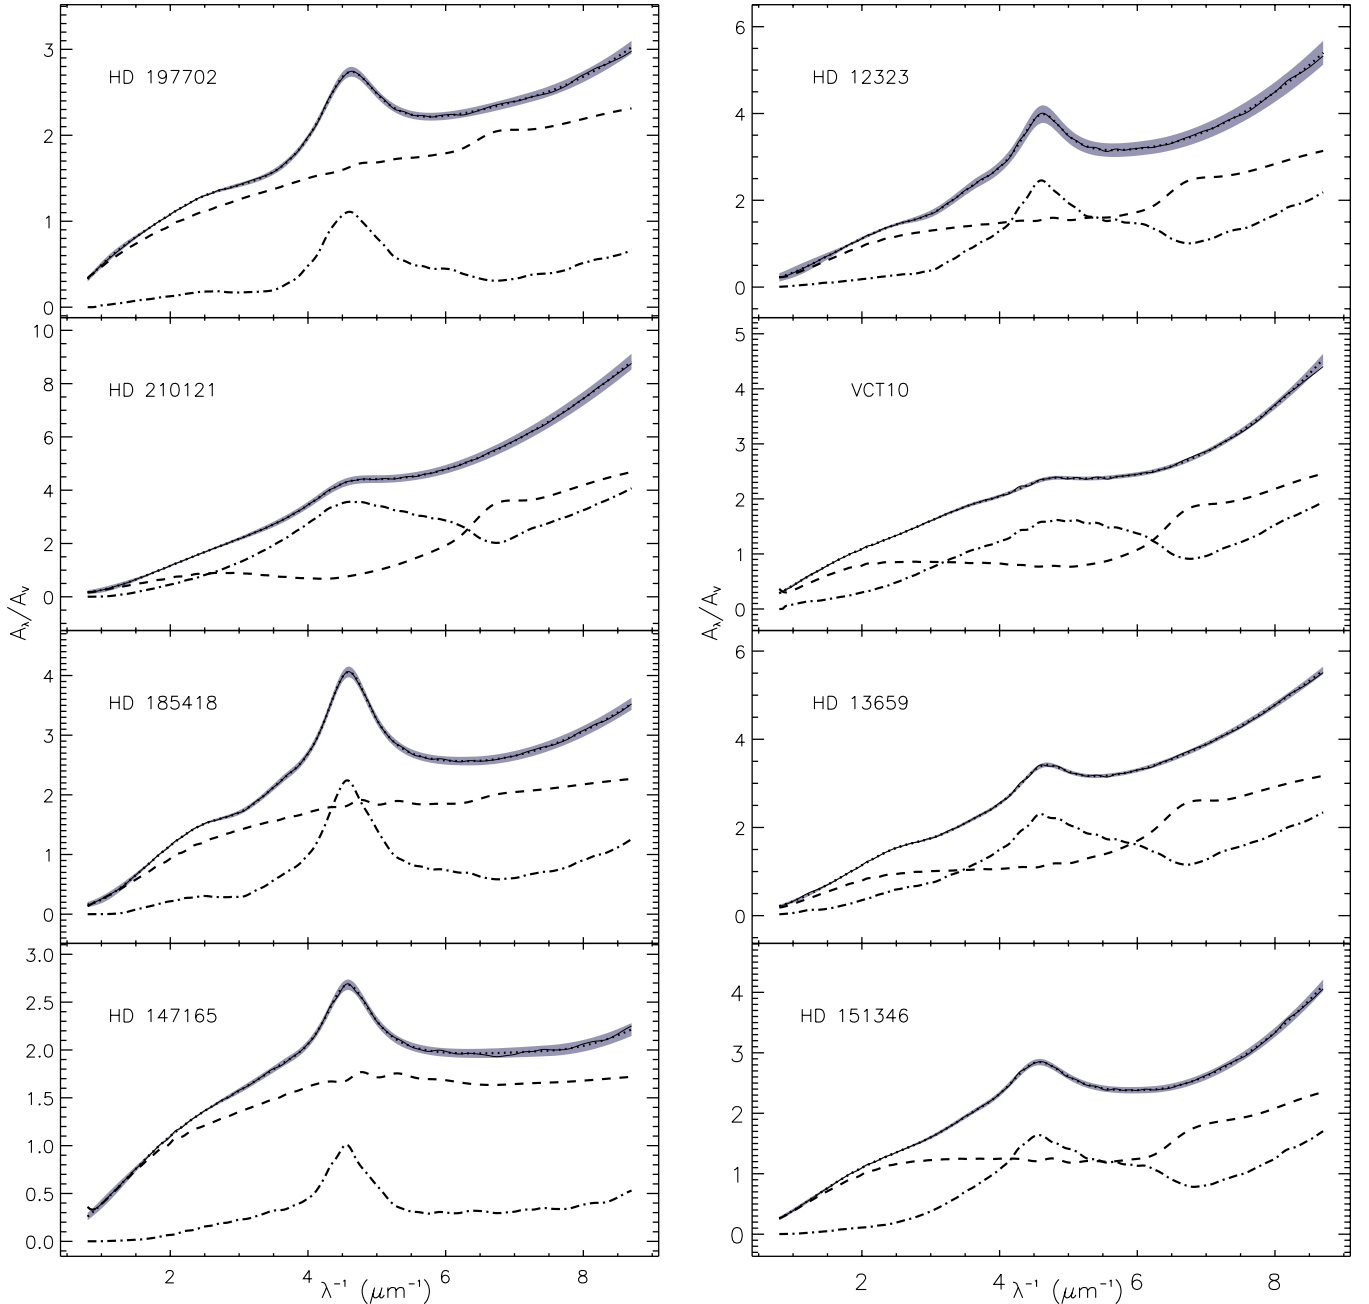
Figure 1. Fits to the normalized extinction curves for eight lines of sight chosen to be representative of the features carried by the Galactic curves. Solid lines: fit; dashed lines: classical dust contribution; dot-dashed lines: PAHs. The lightly shaded areas are the observational error ranges given in Fitzpatrick & Massa (2007). (A color version of this figure is available in the online journal.)

3.1. Classical Dust Distribution Parameters

Some of the parameters determining the distribution of classical dust grains appear to be remarkably constant across the whole sample of lines of sight. This is the case for the q exponent in the power-law size distribution, which is almost always very close to 3.45, with few exceptions. Also the lower limit a_- of the "small" grain sizes is very nearly always fixed on a limit value of 5 nm. This is the minimum size included in our model, and indeed the fit, if unconstrained, would in most cases try to move to smaller sizes. The upper limit b_+ of the "big" grain sizes appears to show relatively little variation as well. In this case, this is because the modeled extinction curve (in the interval considered) is remarkably insensitive to the population of dust grains on the bigger end of the size distribution. Arbitrarily setting this limit to extend to pebble-sized grains would not be visible in the modeled extinction curves here. The values of b_+ that we obtain should therefore be regarded as lower limits. Such an upper size limit would be better constrained by extending the fit to infrared portions of ISECs (e.g., Fitzpatrick & Massa 2009), including silicate absorption features at 10 and 18 μm.

The fraction of vacuum f_v in the classical dust grains appears to cluster rather tightly around $\sim 0.4 \pm 0.15$. We therefore obtain relatively hollow particles. Silicate porosity may be a relevant parameter when we study the optical behavior of large grains because it strongly affects the strength and the shape of silicate absorption features (e.g., Iatí et al. 2001, 2011; Voshchinnikov & Henning 2008).
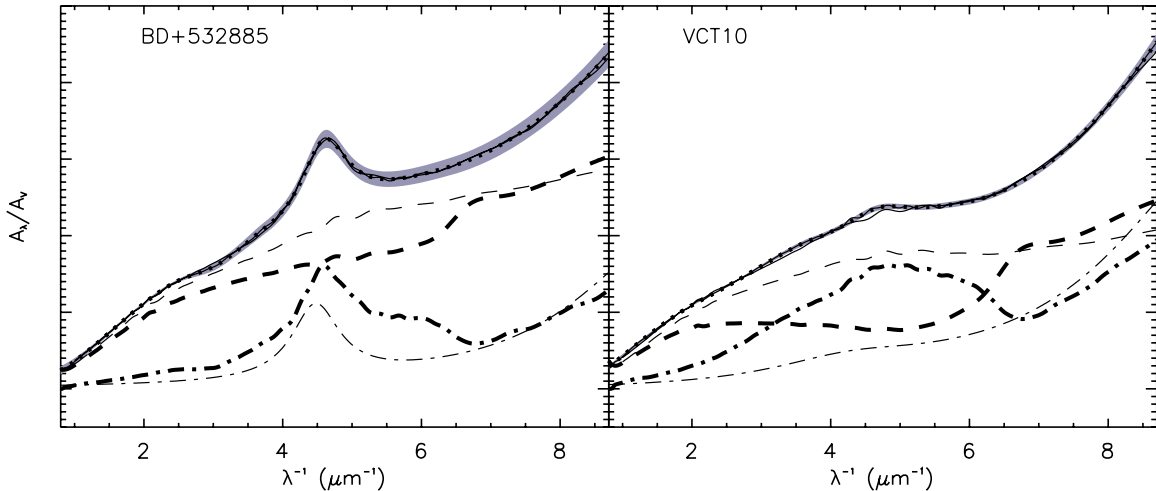
Figure 2. Comparison of synthetic ISECs obtained with simplified (thin lines) and detailed (thick lines) models for the lines of sight toward BD+532885 and VCT10. Solid lines: fit; dashed lines: classical dust contribution; dot-dashed lines: PAHs. The lightly shaded areas are the observational error ranges given in Fitzpatrick & Massa (2007).

(A color version of this figure is available in the online journal.)

Carbonaceous mantles appear to be mostly ~1 nm thick (parameter w), more tightly so in the detailed models and with more scatter in the simplified models. In a non-negligible number of lines of sight we obtain vanishing mantle thickness, but only with the simplified model, whereas this essentially never happens with the detailed model. This might pose some problems if the lines of sight with negligible w also have a non-negligible population of very small grains: if naked, they would be expected to emit in the silicate bands when transiently heated by the absorption of energetic photons and such emission is not observed. This does not happen for the size limits adopted here since even the smallest grains, with a 5 nm radius, would have relatively small temperature fluctuations. It might become a problem if one wants to allow silicate cores to reach sizes around ~1 nm (or smaller). Also, a small but non-negligible number of lines of sight show w values of 3 nm. This was adopted as an upper limit in the fit to avoid using an unphysically large amount of carbon in grain mantles. If left free, these fits would have converged to somewhat larger values of w. The other parameter defining the carbonaceous mantles is f_{sp^2}, the fraction of sp^2 carbon. This parameter is less tightly determined by the fit, but in most cases is close to 1, with only a few cases of mostly aliphatic (polymeric) carbonaceous material.

The parameters a_+ and b_- give a "hole" in the size distribution of classical dust grains. a_+ shows a systematic difference between the simplified ($a_+ \sim 13$ nm) and detailed ($a_+ \sim 25$ nm) models. In both cases there is relatively large scatter around these averages, going from ~5 nm (no small grains since 5 nm $\leqslant a_- \leqslant a_+$) to twice the average. The values of b_- cluster around ~70 nm, again with considerable scatter, ranging from one-half to two to three times the average. This range is larger for the detailed model, which requires larger gaps in the size distribution to account for some peculiar cases (see Sections 3.2 and 3.3).

3.2. Comparison between Simplified and Detailed Model Results

The simplified model (Iatì et al. 2008) was derived from, and calibrated on, the same PAH photo-absorption cross-sections that are explicitly included in the detailed model. The two

models are therefore expected to produce very similar results, both in the PAH abundances and even more so in the classical part, which is exactly the same in both models. Still, some differences remain and should be taken as a lower bound to the systematic uncertainty in our model-derived quantities since they stem from different model choices. Much of this difference results from the very different "flexibility" of the two representations of PAH extinction. In the simplified model, the two Lorentzians representing the $\pi^\star \leftarrow \pi$ and the $\sigma^\star \leftarrow \sigma$ broad resonances have complete freedom to vary their relative intensities, while being restricted to be strictly Lorentzian in shape. Conversely, in the detailed model the intensity ratio between the $\pi^\star \leftarrow \pi$ and the $\sigma^\star \leftarrow \sigma$ resonances can have limited variability, mainly due to charge effects, whereas the $\pi^\star \leftarrow \pi$ resonance can change its shape more freely depending on the specific species included in the mixture, since the exact positions of $\pi^\star \leftarrow \pi$ transitions are rather variable in individual species.

When fitting the same extinction curve, especially if it is a relatively non-standard one, the extinction component due to classical dust has to compensate for these differences, resulting in the observed difference in classical dust distributions between the two models (see Figure 2).

Figure 3 compares the fitted values for homologous parameters obtained with the two models. These parameter values can be straightforwardly translated into abundances for the Si and C locked into classical dust grains. These are listed in Table 2. The same table lists the abundance of C locked in PAHs resulting from the simplified and detailed models and the total abundance of C. On average, the simplified model seems to require a bit more C in PAHs. Again, Figures 4 and 5 compare Si and C abundances resulting from the simple and detailed models for the whole set of lines of sight.

With our statistically large sample we find that abundances of elements locked in the classical dust component obtained with the simplified model are within a factor of ~2 of the ones derived using the detailed description. In particular, the simplified model yields Si abundance in dust almost always very close to 50 ppM, whereas this is more variable in the detailed model, ranging mostly between ~20 and 70 ppM. The abundance of C locked in classical dust mantles shows larger variations in the statistical
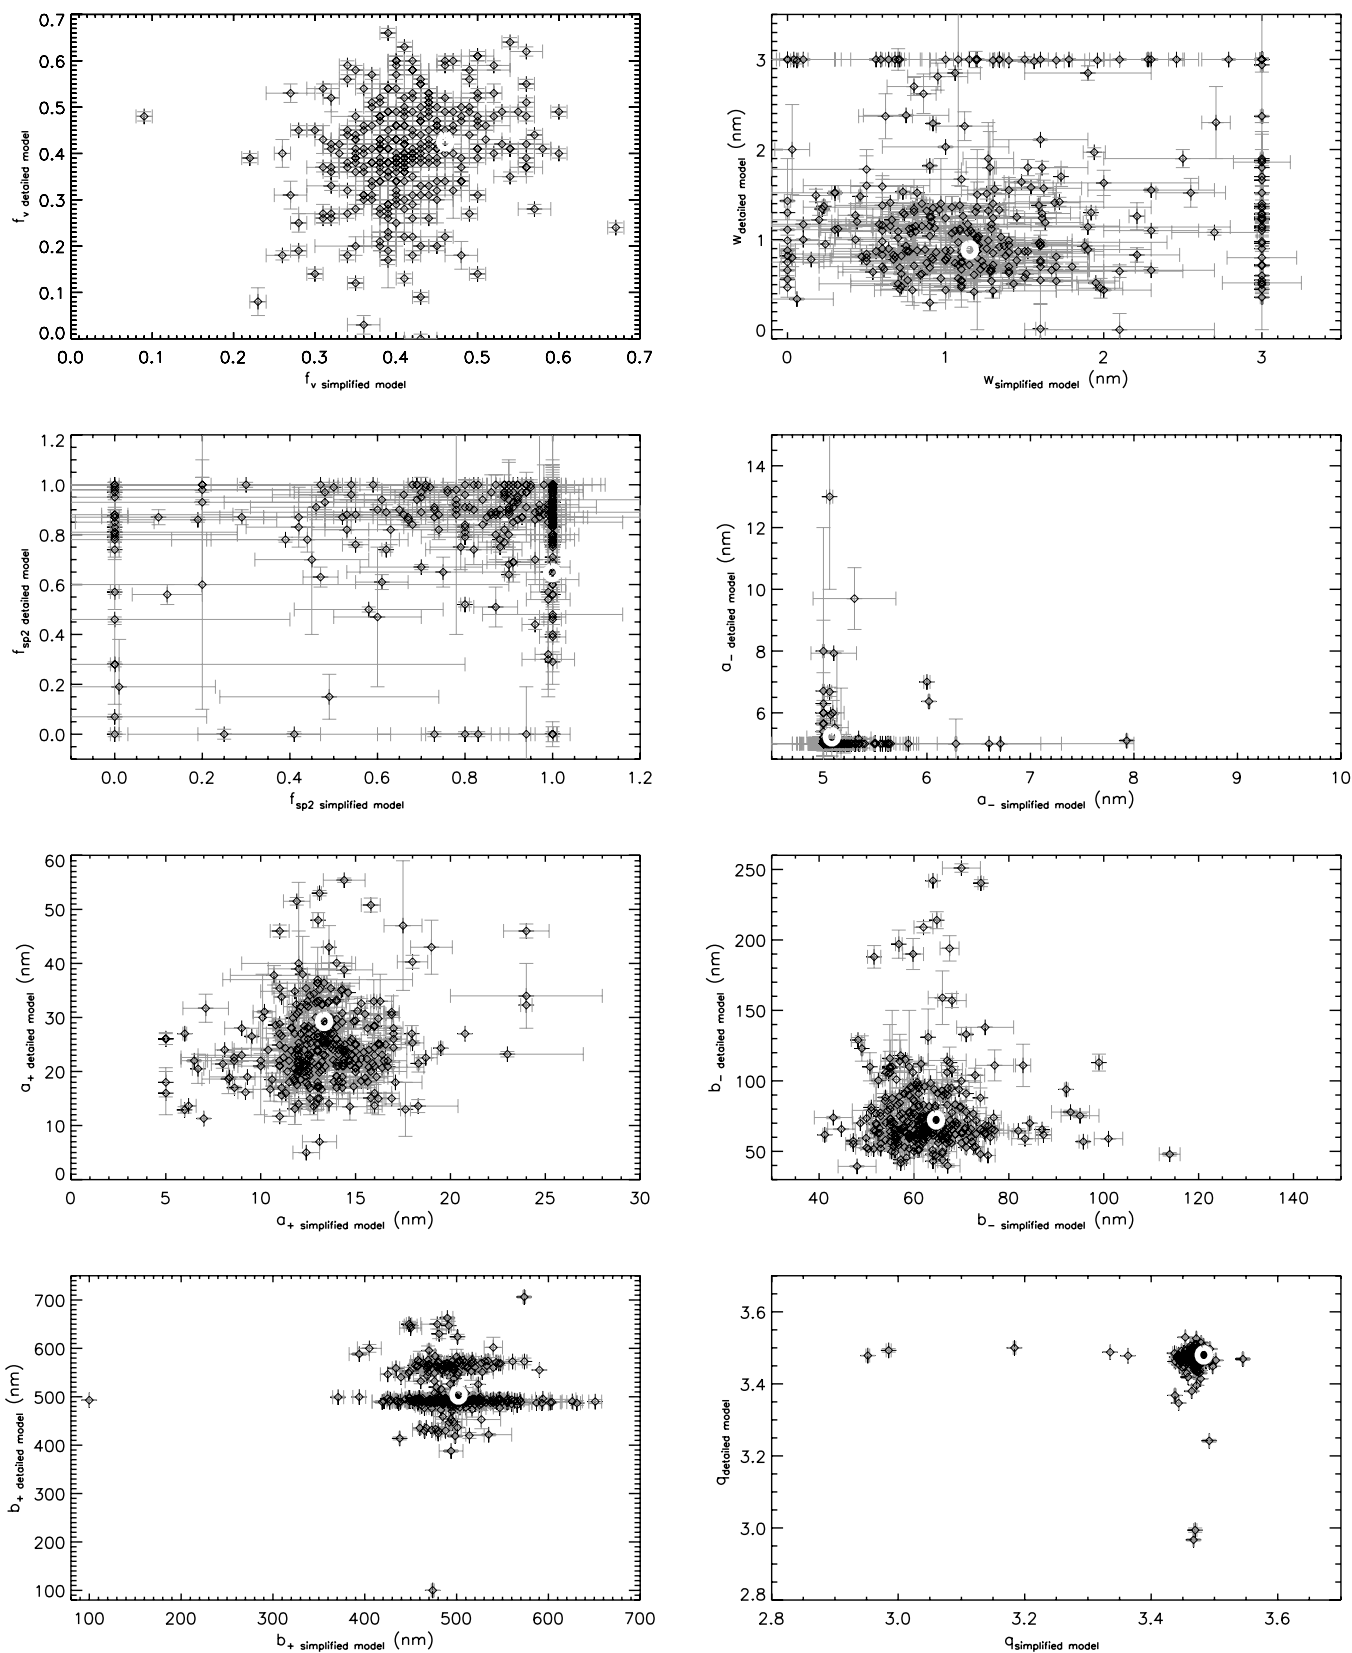
Figure 3. Comparison between homologous parameters obtained with simplified and detailed models. Values for the average ISEC are overplotted with a circle.

sample, but still the simplified versus detailed values are within a factor of ∼2 of each other. In cases in which such accuracy (a factor of two) is acceptable, it is therefore safe to use the simplified model to determine such parameters.

The abundance of C in PAHs is calibrated in the simplified model as proportional to the coefficient of the first Lorentzian (i.e., the bump). Here they are also determined from the population of PAHs in the detailed model fits, and Figure 6 shows

Table 2
Si and C Abundances Resulting from the Simple and Detailed Models for the Average ISEC and for the Lines of Sight Plotted in Figure 1

LoS	Si/H (ppM)			C/H (ppM)			
	Small Grains	Large Grains	Total	Small Grains	Large Grains	PAHs	Total
Average ISEC	1.9 (0.7)	40.6 (1.2)	42.5 (1.7)	9.8 (3.9)	8.8 (3.2)	126.9 (9.3)	145.5 (12.2)
	6.0 ($<\varepsilon$)	39.0 ($<\varepsilon$)	45.0 ($<\varepsilon$)	14.5 ($<\varepsilon$)	5.7 ($<\varepsilon$)	64.3 (2.3)	84.5 (2.3)
HD 12323	2.6 (0.3)	44.2 (2.1)	46.8 (2.2)	7.3 (3.3)	5.6 (2.3)	141.0 (6.9)	153.9 (7.3)
	8.6 (1.1)	27.2 (3.2)	35.8 (2.2)	10.3 (1.9)	1.9 (0.5)	281.1 (46.1)	293.2 (46.0)
HD 13659	2.6 ($<\varepsilon$)	33.6 (0.7)	36.2 (0.7)	5.8 (0.7)	3.8 (0.6)	85.1 (1.1)	94.7 (2.3)
	11.4 ($<\varepsilon$)	26.0 (0.3)	37.5 (0.2)	4.1 ($<\varepsilon$)	0.6 ($<\varepsilon$)	70.5 (0.5)	75.2 (0.5)
HD 147165	1.9 (0.1)	59.3 (0.1)	61.2 (0.1)	27.3 (0.4)	27.7 (0.2)	82.8 (1.4)	137.8 (1.6)
	1.9 (0.1)	59.3 ($<\varepsilon$)	61.1 (0.1)	21.7 (1.1)	23.3 ($<\varepsilon$)	40.2 (0.6)	85.3 (1.3)
HD 151346	1.0 (0.2)	62.6 (0.7)	63.6 (0.6)	9.4 (1.0)	15.0 (1.4)	119.0 (3.4)	143.4 (3.8)
	9.4 (0.3)	52.3 (0.4)	61.7 (0.3)	23.3 (1.0)	6.2 (0.3)	71.7 (1.8)	101.3 (2.5)
HD 185418	3.9 (0.2)	42.1 ($<\varepsilon$)	46.0 (0.2)	8.2 (0.7)	4.7 ($<\varepsilon$)	154.2 ($<\varepsilon$)	167.1 (0.7)
	3.1 ($<\varepsilon$)	37.6 ($<\varepsilon$)	40.8 ($<\varepsilon$)	11.8 ($<\varepsilon$)	7.1 ($<\varepsilon$)	78.6 (0.9)	97.5 (0.9)
HD 197702	3.1 (0.2)	61.5 (1.7)	64.6 (1.9)	36.8 (1.7)	31.7 (1.5)	158.8 (1.8)	227.3 (4.9)
	8.7 ($<\varepsilon$)	91.8 ($<\varepsilon$)	100.5 ($<\varepsilon$)	72.1 (0.2)	50.4 ($<\varepsilon$)	46.4 (0.6)	168.8 (0.6)
HD 210121	2.3 (0.3)	27.7 (0.5)	30.0 (0.8)	0.0 ($<\varepsilon$)	0.0 ($<\varepsilon$)	87.0 (5.7)	87.0 (5.7)
	7.7 (0.3)	8.2 (0.2)	15.9 (0.5)	4.6 (0.5)	0.2 ($<\varepsilon$)	257.3 (6.4)	262.1 (6.2)
VCT10	3.1 (0.3)	58.3 (3.0)	61.4 (3.2)	14.7 (1.6)	9.5 (1.1)	265.9 (81.1)	290.1 (81.3)
	14.0 (0.1)	44.5 (0.5)	58.5 (0.4)	10.9 (0.4)	1.6 (0.1)	67.2 (9.9)	79.6 (9.9)

Note. All the abundances are expressed in ppM. We adopt $\varepsilon = 0.1$.

(This table is available in its entirety in a machine-readable form in the online journal. A portion is shown here for guidance regarding its form and content.)
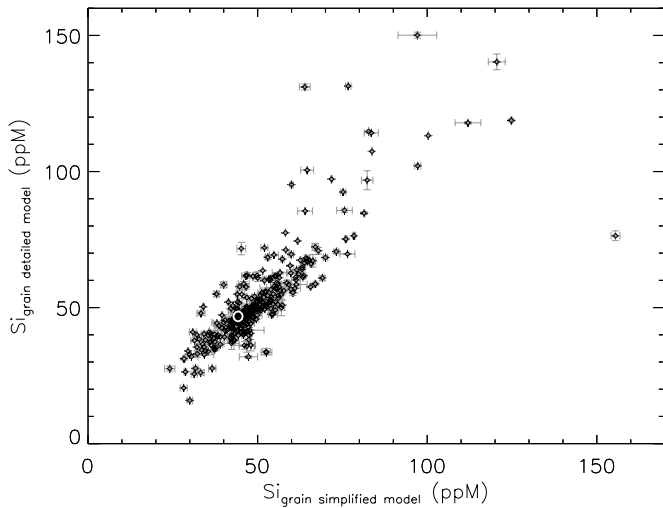

Figure 4. Comparison between Si abundances locked into classical dust grain obtained from the simplified and the detailed models. Values for the average ISEC are overplotted with a circle.

Figure 5. Comparison between C abundances locked into classical dust grains obtained from the simplified and detailed models. Values for the average ISEC are overplotted with a circle.

the comparison, with the best-fitting straight line superimposed. That calibration relied solely on the integrated area of the first Lorentzian, disregarding the second Lorentzian and the cutoff at low energies (assumed because PAHs mostly do not absorb at wavelengths longer than the visible). It may be closer to physics to correlate the abundance of C in PAHs to their *integrated* absorption in the whole modeled range. We therefore computed this integral between 1 and 8.7 μm^{-1} (the observational range), and plotted it against the abundance of C in PAHs obtained from the detailed model in Figure 7, along with the best-fitting straight line. Indeed, points in Figure 7 are more closely packed toward
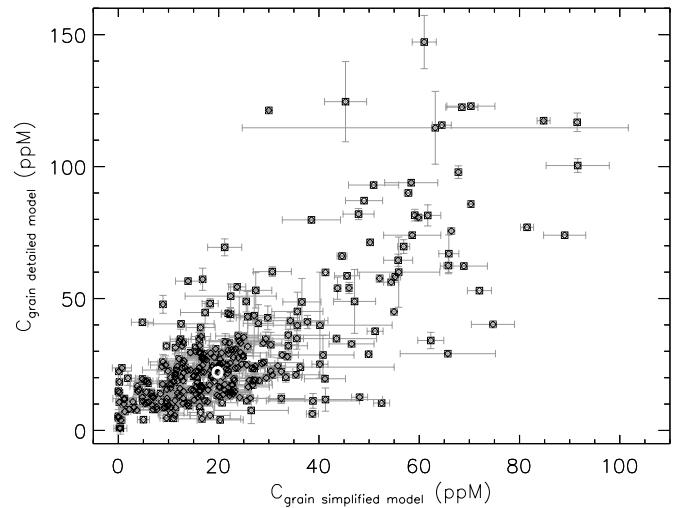
the regression line with respect to Figure 6. The regression line has a coefficient of $\sim$5.6 ppM μm mag^{-1}.

Abundances of C in PAHs obtained from the simplified model using only the fitted coefficient and the integral of both Lorentzians including cutoff are, respectively, within a factor of $\sim$4 and $\sim$2 of those obtained with the detailed model, across the statistical sample. Again, if such additional systematic uncertainty is acceptable, one can safely choose the simplified model over the detailed one for this purpose. In this case, however, we recommend using the calibration on the integral of the simplified PAH absorption (including both Lorentzian
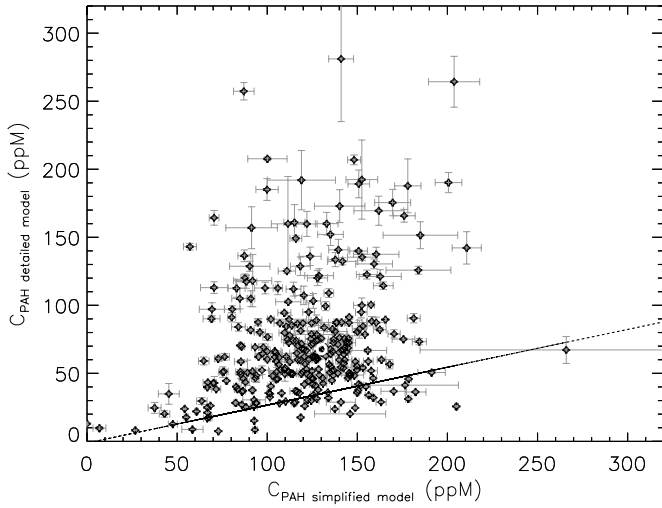
Figure 6. Comparison between C abundances locked into PAHs obtained from the simplified and detailed models. In the simplified version carbon abundances are proportional to $Lp1$. Values for the average ISEC are overplotted with a circle. Solid line: best-fitting straight line.
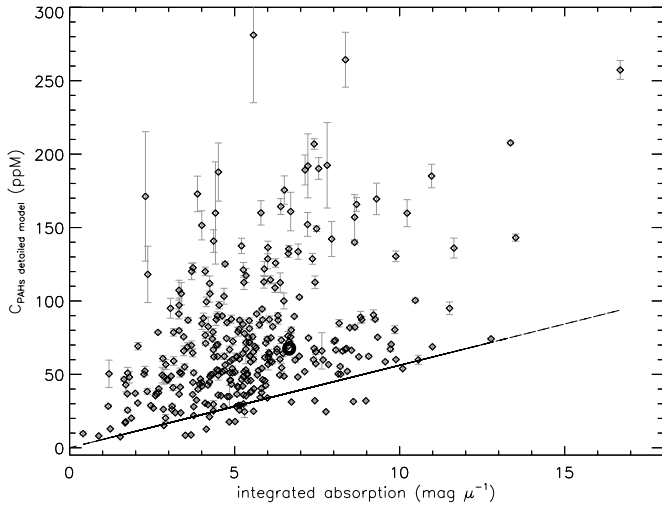

Figure 7. Comparison between C abundances locked into PAHs obtained from the detailed models and the integral of the "PAH" component of the simplified modes throughout the extinction curve, namely between 1 and $8.7\,\mu m^{-1}$. Values for the average ISEC are overplotted with a circle. Solid line: best-fitting straight line.

and cutoff functions) instead of just the coefficient of the first Lorentzian, since this is both more accurate and unbiased, whereas using only the first Lorentzian will systematically overestimate C abundance in PAHs.

3.3. The Charge of PAH Mixture

Cecchi-Pestellini et al. (2008) first introduced the detailed model, and, applying it to a few lines of sight, found that whereas the detailed PAH composition is poorly constrained by the fit, some global properties of the mixture are relatively strongly constrained. This is most remarkably true for the total abundance of C atoms in PAHs, and for the average charge per carbon atom. Here we go further and try to understand whether, in each line of sight, PAHs appear to concentrate only in the charge state closest to the average charge per C atom, or whether, instead, they spread over all charge states. To study this, for each line of sight we computed the rms dispersion of the distribution of
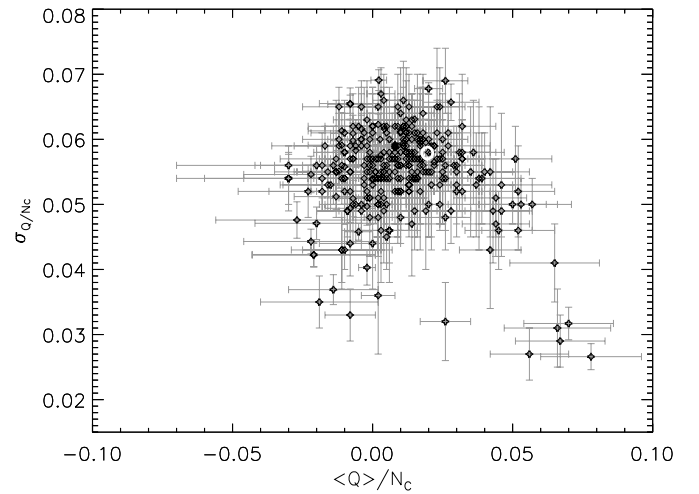

Figure 8. Average charge per C atom vs. its rms dispersion. Values for the average ISEC are overplotted with a circle.

Table 3
Average Charges Per C Atoms and rms Dispersions Obtained with the Detailed Model for the Average ISEC and for the Lines of Sight Plotted in Figure 1

LoS	$\langle Q \rangle / N_C$	σ_{Q/N_C}
Average ISEC	0.017 (0.012)	0.057 (0.010)
HD 12323	0.010 (0.013)	0.056 (0.010)
HD 13659	−0.008 (0.008)	0.0654 (0.0014)
HD 147165	0.028 (0.007)	0.049 (0.005)
HD 151346	0.020 (0.006)	0.060 (0.007)
HD 185418	0.032 (0.008)	0.057 (0.006)
HD 197702	0.032 (0.008)	0.057 (0.003)
HD 210121	−0.008 (0.011)	0.0655 (0.0024)
VCT10	0.0042 (0.0019)	0.0616 (0.0014)

(This table is available in its entirety in a machine-readable form in the online journal. A portion is shown here for guidance regarding its form and content.)

the charge per carbon atom. The obtained values are reported for the full set of lines of sight in Table 3. Figure 8 plots the average charge per C atom versus its rms dispersion. The rms dispersion is mostly clustered around $\sim 0.6 e^-/N_C$, whereas the average charge ranges from ~ -0.5 to $\sim 0.9 e^-/N_C$. Since the PAHs in our mixture range in size from 10 to 66 C atoms, this corresponds to any fitted mixture always including a little more than two charge states, with the average charge changes between lines of sight being significant with respect to its intrinsic dispersion.

Cecchi-Pestellini et al. (2008) noted that on their small set of lines of sight the average charge per C atom appeared to correlate with the ratio c_4/c_3 of the c_4 and c_3 parameters of the Fitzpatrick & Massa (2007) parameterization. In this parameterization, the c_3 parameter is the bump intensity, while c_4 is the intensity of the nonlinear far-ultraviolet rise. The ratio c_4/c_3 is therefore highest for lines of sight with a bumpless extinction and large nonlinear far-ultraviolet rise (e.g., VCT10), whereas it is low for lines of sight with the reverse behavior (e.g., HD 185418). Indeed, it is a known systematic spectroscopic trend that PAHs, with increasing positive ionization, have a larger gap between the $\pi^\star \leftarrow \pi$ and the $\sigma^\star \leftarrow \sigma$ resonances (Cecchi-Pestellini et al. 2008). This trend has been found to become weaker for increasingly larger PAHs (Steglich et al. 2011) because, predictably, one more or less electron makes less difference for a very big molecule with hundreds of electrons; however, this is
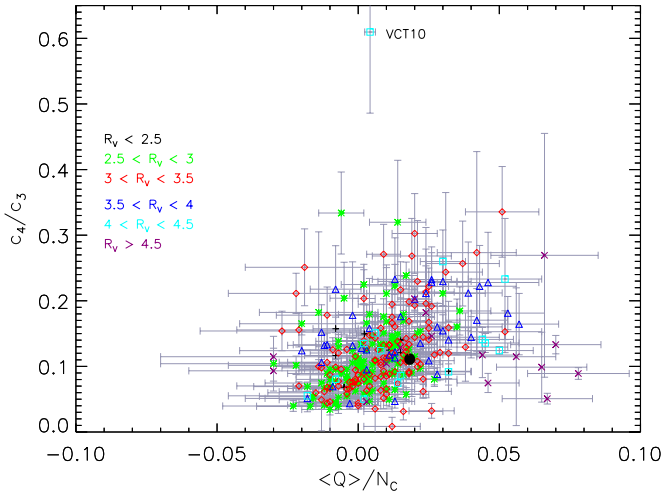
Figure 9. Relation between the average charge per C atom and c_4/c_3 for all lines of sight. The different symbols and colors refer to different intervals of R_V. Values for the average ISEC are overplotted with a black circle.

(A color version of this figure is available in the online journal.)

Figure 10. Occupation volume in the parameter space of the average PAH charge per C atom, the "missing grain" interval in the dust size distribution, and the observational ratio c_4/c_3. Values for the average ISEC are overplotted with an X.

balanced if one normalizes charge to the number of C atoms in the molecule, as we do here.

Figure 9 shows the relation between the average charge per C atom and c_4/c_3 for all lines of sight. Lines of sight in different intervals of R_V are shown with different symbols and colors. Even simple visual inspection of Figure 9 shows that there is *some* positive correlation, and that this changes somewhat with R_V. It is also apparent that there is quite a real scatter, which cannot be dismissed as just due to measurement and fitting errors. This tells us that PAH charge *does* play a role in determining the ratio between bump and far-ultraviolet nonlinear rise intensities, but it also tells us that it does not *solely* determine it. Indeed, it is instructive to examine the lines of sight with very weak bumps but steep far-ultraviolet extinction in Figure 1, and look in detail at the fitted PAH and classical dust contributions. One might naively expect to find a vanishing contribution from PAHs in these cases, which indeed are fitted, in the simplified model, by a vanishing first Lorentzian and a simultaneously large second one. Exploiting the full PAH representation, it is impossible to achieve such extreme ratios between the $\sigma^\star \leftarrow \sigma$ and $\pi^\star \leftarrow \pi$ resonances, which have only limited variability due to charge effects. Thus, even in bumpless extinction curves the PAH component is well present, but the $\pi^\star \leftarrow \pi$ PAH resonance is masked by the weak underlying classical extinction. Along bumpless lines of sight, in going from small to high energies, the extinction due to dust flattens under the bump region, and resumes its linear rise at higher energies, when extinction due to small particles kicks in. This happens because there is a wavenumber range in which extinction due to big grains saturates to its optical limit, and small particles do not contribute much yet.

This is consistent with observations sometimes showing quite non-negligible PAH *emission* from regions with bumpless extinction curves (see, e.g., Paradis et al. 2011; Sandstrom et al. 2012). We conclude that there is an interplay between PAH charge state and classical dust size distributions in determining the observed c_4/c_3 ratios. This also explains the differences in Figure 9 between subsets of lines of sight in different R_V intervals, since larger R_V values roughly translate into larger average grain sizes. From an observational point of view,
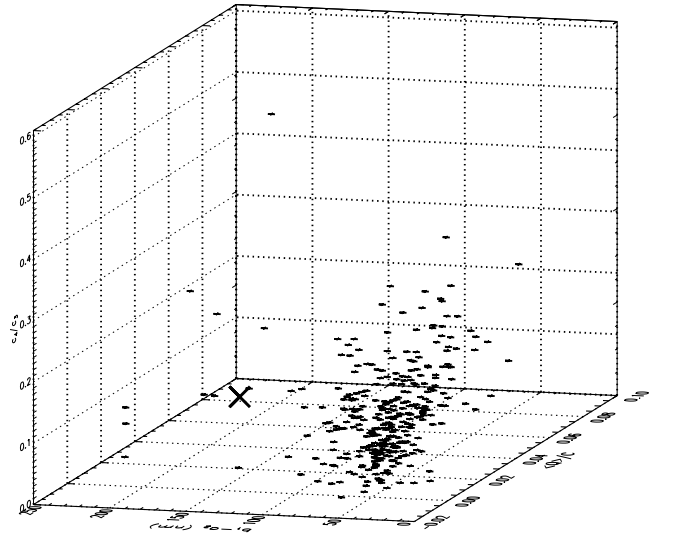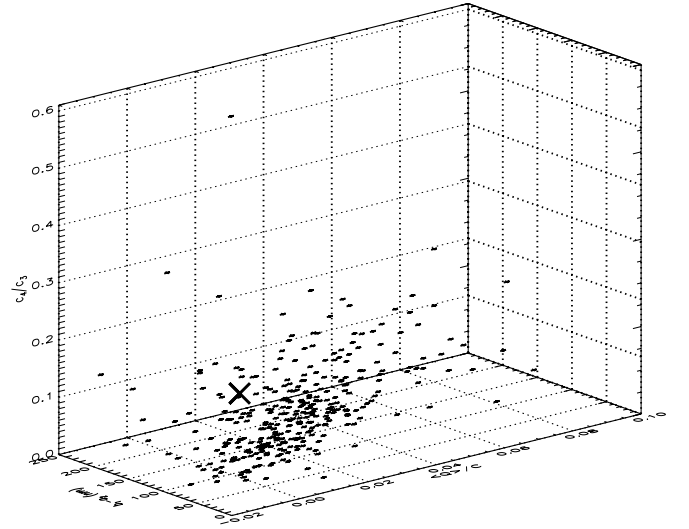
bumpless extinction curves typically correspond to regions with harsher irradiation conditions. From a physical standpoint, this is expected to affect PAH charge state and dust size distributions. A large ultraviolet radiation density may destroy smaller dust particles, thus increasing R_V and the size gap between small and large dust grains. At the same time, an increase in the ultraviolet radiation density tends to move average PAH charge toward more positive values. However, in this case such an effect can be balanced and/or masked by other environmental conditions, e.g., electron density and selective destruction of smaller PAHs (thereby decreasing the absolute value of average charge per C atom). It is therefore not surprising that it is not possible to find simple, perfect correlations relating individual Fitzpatrick & Massa (2007) empirical parameters and the more physical (albeit model-dependent) results from our fits.

As an example, Figure 10 shows a three-dimensional plot with average PAH charge per C atom, gap of the dust size distribution, and c_4/c_3, seen from two different viewing angles. The points clearly do not occupy the whole space, but all three quantities are related in an interdependent way. The source
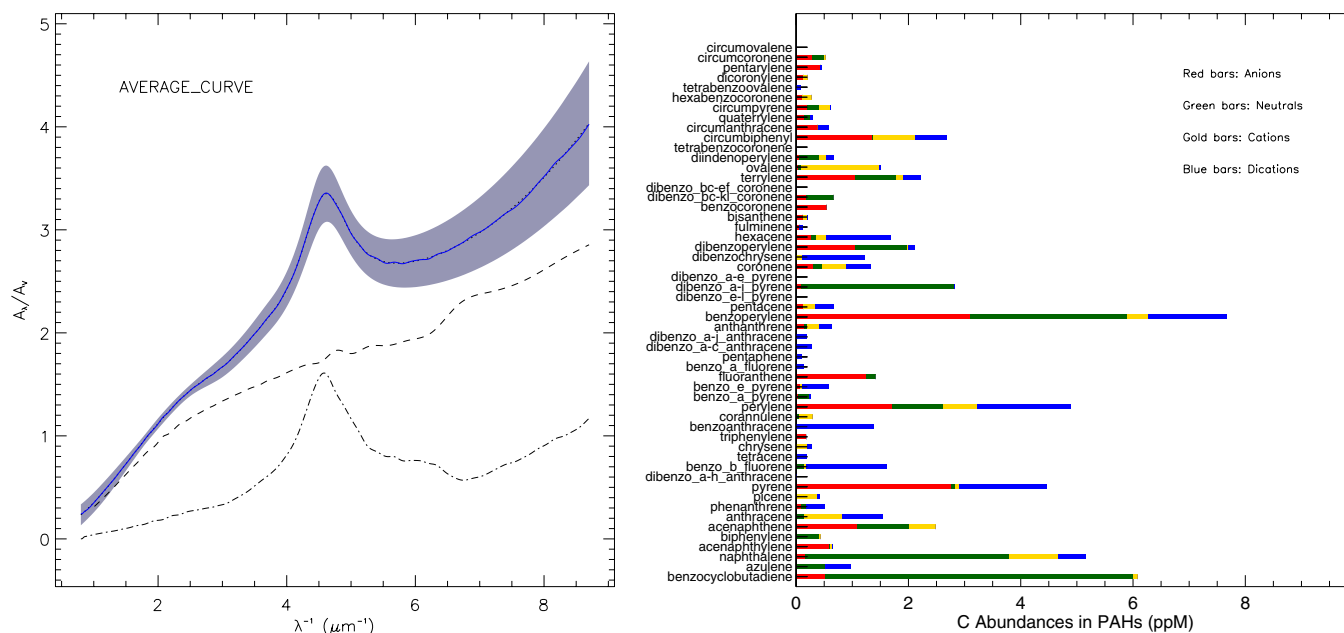
Figure 11. Left panel: fit to the normalized extinction curve. Solid lines: fit; dashed lines: classical dust contribution; dot-dashed lines: PAHs. The lightly shaded area is the 1σ observational error range given in Fitzpatrick & Massa 2007). Right panel: histogram of carbon distribution among synthetic PAH populations, averaged over their statistics. .eps files of the images in the figureset are available from the authors upon request.

(The color version and complete figure set (329 images) are available in the online journal.)

code for the interactive (rotatable and zoomable) version of this example figure can be used with any javascript-enabled browser. It depends on the open source CanvasXpress toolkit, which can be obtained separately.[5] An updated usable version of the interactive figure is also made available and will be maintained by the authors.[6]

A thorough analysis of the complex relations relating observational, empirical, and "physical" (fitted) parameters requires applying appropriate statistical multivariate analysis techniques. This is definitely out of the scope of the present paper, and will be presented in future works that are in preparation.

3.4. Detailed PAH Composition: Unconstrained

In Figures 11.1 to 11.329 in the online journal, the right panel shows a histogram with the average distribution of C among the PAHs we include in our model, in all four charge states, for each line of sight, along with the individual fits in the left panel.

These kinds of histograms can be useful to explore trends among different PAH subgroups for specific lines of sight. However, one should be well aware that *individual* PAH abundances are very poorly (if at all) determined by the extinction curve alone. This is clear if one inspects the covariance matrices of the C abundance in individual PAHs for an individual line of sight (all available on request from the authors). To make this visually apparent, Figure 12 shows one typical individual extinction curve (dotted line), along with its best fit (blue line), and the fit obtained by removing from the PAH mixture the species determined to be most abundant in the best fit, including ~98% of the C in PAHs in the best fit. The left panel of Figure 13 shows the histogram corresponding to the best fit; the right panel shows the same histogram for the "degraded" fit. The degraded fit is therefore obtained by deliberately including
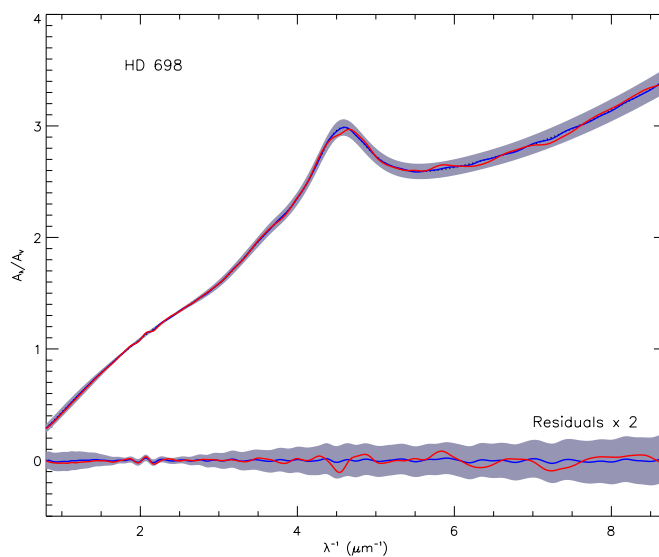

Figure 12. Best fit, its degradation, and residuals for the line of sight toward HD 698. Dotted lines: observational data; red lines: best fit: blue lines: degraded fits. Lightly shaded area: observational error range taken from Fitzpatrick & Massa (2007).

(A color version of this figure is available in the online journal.)

only those PAHs which were essentially absent in the best fit. In Figure 12, with the same color coding used for the fitted curves, we show the residuals (multiplied by 2 to make them more visible) of the two fits with respect to the observed extinction curve. It is clear that the red fit is still perfectly acceptable, despite our best efforts to deliberately sabotage it. However, the total amount of C in PAHs and the average charge per C atom in PAHs is very nearly unchanged in the two fits, showing again that these quantities *are* well constrained by the extinction curve alone within our detailed model.
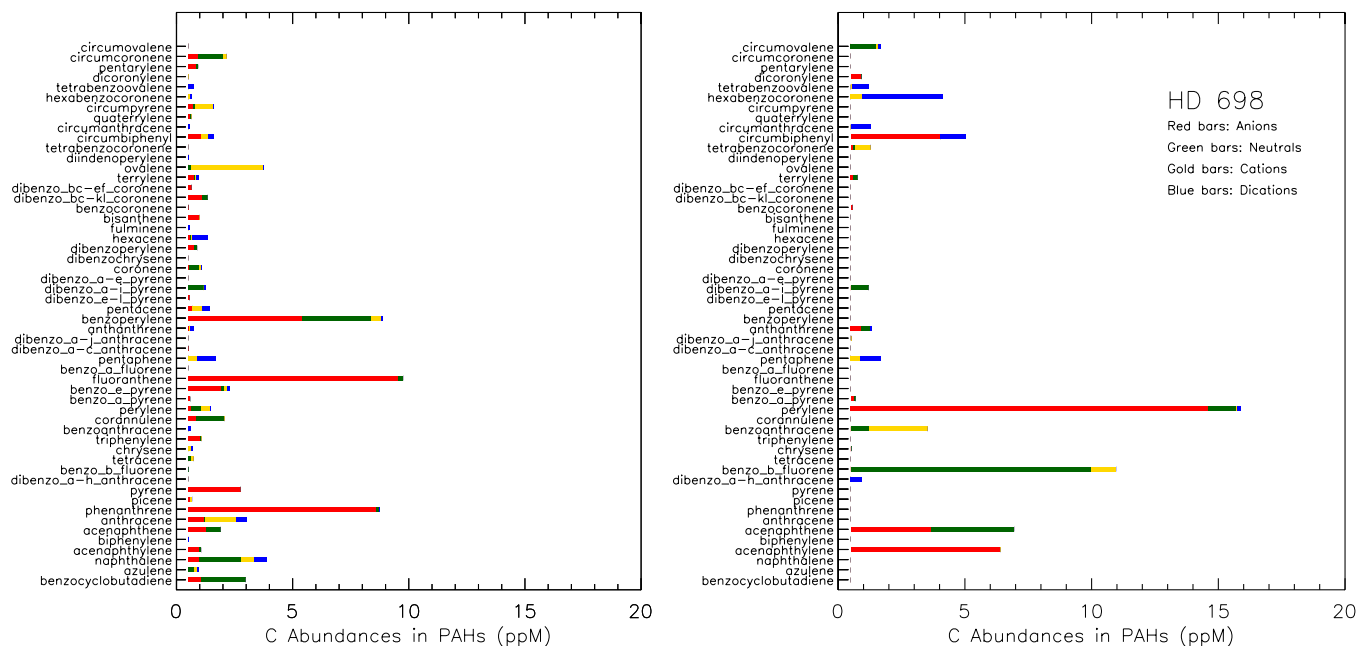

Figure 13. Distribution of carbon among synthetic PAH populations for HD 698. Left panel: best fit; right panel: degraded fit (see the text).

(A color version of this figure is available in the online journal.)

4. CONCLUSIONS

In this work we systematically applied the [CM]^2 model to the 328 individual Galactic extinction curves observed by Fitzpatrick & Massa (2007). Most previous attempts to apply interstellar dust models to observations have been limited to the average Galactic ISEC or to a few individual directions. We demonstrate that our physical extinction model can successfully fit a very large sample of extinction curves without exceptions. It also highlights the need for future modeling of interstellar matter to consider the number of free parameters as an important issue that should not be neglected.

The results depicted in the preceding sections provide the first global characterization of the galactic extinction in terms of the physical parameters of a well-defined dust model, including a mixture of individual PAHs, described by their actual quantum properties. We found that solids and molecules are tightly coupled and both of them concur in the determination of interstellar extinction. Such a conclusion is much less obvious than model assumptions suggest. For instance, while in most of the cases dust grains provide the "perfect" baseline (in the spirit of the Fitzpatrick & Massa 2007 parameterization) for the PAH ensemble resonances, in a few "peculiar" lines of sight their extinction properties appear to hide the 217.5 nm feature.

The first and most evident result of this work is that the [CM]^2 model can adequately reproduce *all* the extinction curves in the Fitzpatrick & Massa (2007) sample. Such a conclusion is not terribly exciting per se, and is to some extent expected. Still, it is interesting that a physically simple model (despite the apparently large number of 225 free parameters) can match all observed variations. The physical simplicity of the model is promising for investigating how a small number of physical quantities, defining the environment, can drive the model to respond as a whole. This is the missing link needed to reconcile the unavoidable complexity of any physical model with the small number (less than six) of empirical parameters necessary to account for observations (Cardelli et al. 1989; Fitzpatrick & Massa 2007; Krełowski & Strobel 2012). This will also make
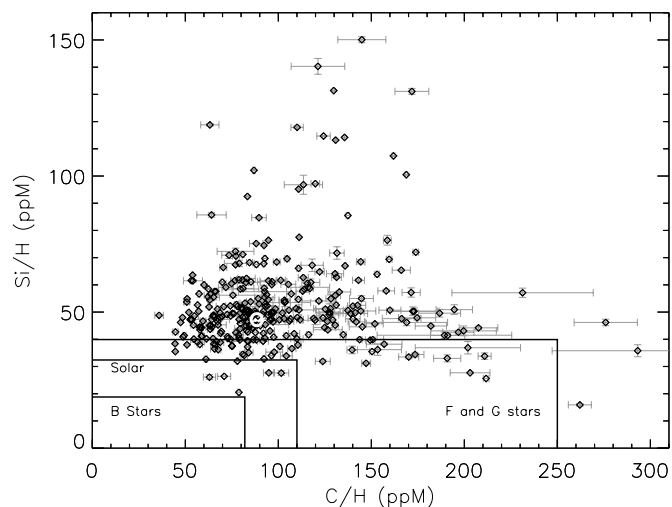

Figure 14. Total Si and C abundances in ppM. The circle represents the average ISEC. The boxes refer to observational constraints: solar taken in Asplund et al. (2009), Mathis (2000), and Be, F, and G stars from Zubko et al. (2004).

it possible to investigate the conditions that characterize the observational exceptions, i.e., the lines of sight that fail to be fitted by empirical parameterizations (e.g., HD 210121).

The present modeling effort already produced a number of direct and indirect results that provide a starting point for future investigations. Our results can be summarized as follows.

1. The total number of C atoms locked in both solids and molecules is largely consistent with the available interstellar budget. Conversely, this model requires somewhat too much Si for a consistent fraction of the 328 observed lines of sight (see Figure 14). The average Galactic ISEC falls close to elemental constraints. However, our use of old optical constants for dust materials (i.e., Draine 1985; Rouleau & Martin 1991; Ashok et al. 1991) probably contributes to exhausting the available atoms. Indeed, Zubko et al.

13

(2004), using a revision of the Rouleau & Martin (1991) optical constants, were able to propose an interstellar dust model whose results are consistent with interstellar element abundances. Moreover, we emphasize that our fitting algorithm aims at the very best match with the observed ISEC, with no regard for abundance constraints, obtaining an unrealistically perfect match. It would be easy to substantially reduce the required abundance of C and Si at the price of a slightly worse fit, but still largely compatible with observations within measurement errors. As an example we consider the case of BD+56518; Zonca et al. (2011) derived a fit much less constrained than the one derived here (but still smoothly within 1σ); such a description led to the exploitation of 5 ppM of Si, 6 ppM of carbon in classical grains, and 48 ppM of carbon in PAHs. On the other hand, from the data in Table 1, we find corresponding abundances of 34 ppM (Si), 13 ppM (C), and 127 ppM (C), respectively. In future works, we may introduce in the fit an increasing χ^2 penalty term for abundances increasingly exceeding observational constraints.

In any case, the use of the sole average Galactic ISEC does not provide a reliable validation of dust models; in particular, chemical and physical characteristics of specific lines of sight, such as, e.g., metal abundances,[7] should be included in the endorsement process.

2. In the ultraviolet, classical dust grains contribute an almost gray baseline to the extinction for the totality of lines of sight explored in this work. Again, such a result is not totally unexpected, and it reflects the a priori choice of the bump carriers with respect to the Fitzpatrick & Massa (2007) parameterization of ISECs. Despite the apparently large number of free parameters, the model is far from being infinitely flexible because many such parameters control the extinction properties of the mixture within a stiff class of morphological profiles. The visible part of any ISEC is mainly produced by the larger sized dust, while the ultraviolet rise is due to the contribution of small grains and PAHs floating on the top of the flat, saturated extinction of the large ones. In other words, dust grains are "forced" to behave in this way in response to the shape of the convolved $\pi^\star \leftarrow \pi$ and $\sigma^\star \leftarrow \sigma$ resonances of the PAH ensemble. This implies that the description of classical dust grains is far from unique, and is poorly constrained by the data, as long as the "correct" size distribution is employed. This is not a new concept, but here it is validated against a statistically relevant sample of lines of sight.

3. In spite of the detailed PAH chemical composition being poorly constrained by extinction data, some global properties of the molecular ensemble *are* well determined, e.g., the column density of carbon atoms locked in PAHs and the mean charge of the molecular mixture. By contrast, individual species are interchangeable and expendable, and no individual property is highlighted by the fitting procedure. We conclude that many loose mixtures of a relatively small number of PAHs may reproduce the features of the extinction curve in the ultraviolet without significant losses in the accuracy of the description. Unfortunately, this conclusion raises the more general question of the stability of

the 217.5 nm feature: why should a random assembling of different molecules provide the same extinction profile, modulo the differences in the profile width, along different lines of sight? We know that not all the linear combinations of PAHs are consistent with the interstellar features (e.g., Malloci et al. 2007). A similar problem is posed by models considering graphite grains responsible for the 217.5 nm feature. Extinction calculations indicate that graphite particles must exist in a relatively narrow size range and be almost spherical in shape (Draine 1988). As a consequence, the straight assumption of a set of PAHs as bump carriers appears to be at best somewhat artificial. If a set of PAHs (such as we find) were indeed to produce the observed extinction, there must be a (heretofore unknown) chemical selection mechanism at work.

4. We found that, in the framework of our model, bumpless extinction curves are not necessarily lines of sight devoid of PAHs, but rather that the $\pi^\star \leftarrow \pi$ PAH absorption can be very effectively masked by a large gap between the maximum size limit of small dust grains and the minimum size limit of large ones. Thus, the present model was proven to be sufficient, without additional components, to fit all observed extinction curves. On the other hand, the previous discussion makes such a conclusion suspicious unless bumpless extinction curves could be defined *only* with respect to the properties of classical grains. A possible observational test would be to look for some peculiar ISEC with a "dip" in the bump regions, which would be observable if the "hole" in the dust grain size distribution and the mixture of PAHs do not perfectly cancel each other out in a smooth extinction curve.

When extinction curves are very steep (e.g., HD 62542) the weakness of the bump may be apparent, being an artifact of the normalization (Fitzpatrick 2004). However, this suggestion may fail to explain the weakening of the bump in the flattest (highest R_V) extinction curves such as HD 29647. Whatever the case may be, suppression of the $\pi-$plasmon resonance in small carbon particles can be obtained by increasing the density of sp^3 bonds (e.g., Jones et al. 1990; Duley 2006). This may suggest super-hydrogenation (Hecht 1986) or the presence of a fraction of carbon in nanodiamonds (e.g., Rai & Rastogi 2010).

5. The numerical tendency to go below the lowest limit of 5 nm in particle size in a significant number of fitted lines of sight (and some physical sense) suggests the presence of an additional component of very small grains. Such particles might be carbon mantle fragments, possibly produced during some destructive events, or an independent population of very small carbonaceous grains (Dèsert et al. 1986). Such very small particles are unlikely to include silicates, since they would be stochastically heated by single photon absorption events and would produce significant emission in silicate bands, which is not observed; still, Li & Draine (2001) argued that a small amount of very small silicatic grains might be present. In a dynamical scenario, this "mesoscopic" carbonaceous component, heated in a hydrogen plasma (e.g., shocks), may rearrange some of its carbon skeletons from sp^2 to sp^3 bondings (e.g., Landstrass & Ravi 1989; Furton & Witt 1993). Such processed particles would be partially aliphatic and therefore would have optical properties qualitatively different from a superposition of PAH cross-sections; they might provide an efficient extinction in the far-ultraviolet without the bump at 217.5 nm (as

---

7 Gudennavar et al. (2012) reported the existence of a tight relation between the total hydrogen column density $N(H)$ and the $B-V$ color excess, $N(H)/E_{B-V} = 6.12 \times 10^{21}$ cm^{-2} mag^{-1}, along a great number of diffuse and translucent lines of sight.

implicitly suggested by the results of a simplified version model). In different physical conditions, some of the very small carbonaceous grains might instead evaporate, producing PAHs, as suggested by Cesarsky et al. (2000) and supported by observational evidence in some environments (Rapacioli et al. 2005; Berné et al. 2007).

6. We highlighted that the ratio between the observed intensities of the nonlinear far-ultraviolet rise and of the ultraviolet bump results from the combined effect of PAH charge on PAH absorption features and from changes in the classical dust size distributions in a complex interplay. This behavior, together with the variations of the spectral properties of PAHs in different charge states (see Figure 1 in Cecchi-Pestellini et al. 2008), naturally leads to the observed weak anti-correlation between nonlinear far-ultraviolet rise and ultraviolet bump intensities.

In conclusion, we show in this work the potentialities of our approach, and make the resulting data available to the community, but we did not fully exploit them here: multi-dimensional correlations simultaneously involving several parameters and observed quantities will require appropriate multi-variate statistical analysis to be unraveled. Since the ultraviolet extinction curve alone underconstrains some parameters, additional physics, additional observational constraints, or both, must be employed to shrink the parameter space and the number of independent free parameters. However, this requires a pre-selection of lines of sight according to severe observational constraints and cannot be done in a systematic way on the whole sample of Fitzpatrick & Massa (2007).

The most obvious piece of complementary information would be to constrain the PAH mixture to simultaneously fit the extinction curve and the observed AIBs in emission for the same line of sight, for those few cases in which both are available. It has been proven that the mid-infrared AIB spectra can be fitted by mixtures of PAHs (Boersma et al. 2013). Here we proved that PAH mixtures can fit any extinction curve, but it remains to be seen whether the two observational features can be *simultaneously* fitted by *the same* PAH mixture for the same line of sight. This would also provide the most direct observational test for the need of a different population of particles contributing to the far-UV nonlinear rise without contributing to the bump (see point 4 in this discussion) or lack thereof.

We might also include an ionization equilibrium condition for PAHs, instead of treating all ionization states as independent parameters as we did here. This would obviously have a huge effect, adding only a handful of free parameters (at most electron density, kinetic temperature, and ultraviolet radiation density) to remove ~150. However, this would make physical sense *only* for lines of sight whose extinction is produced by relatively homogeneous interstellar material, which is not the common case. Observations of common tracers of interstellar material, e.g., the neutral sodium lines, almost always show a complicated structure of many intervening clouds with different radial velocities and velocity dispersions, hinting at different physical conditions among them. In such cases, trying to enforce a unique ionization equilibrium for all PAHs along the line of sight would just be a meaningless exercise, which we hereby avoided. Again, we do plan to select a small subset of lines of sight which are observationally known to be dominated by a single, relatively homogeneous interstellar cloud (or a small number of well-defined ones) and redo the fit for these using the detailed model

with the added constraint of ionization equilibrium of PAHs (or a different one for each observed component).

Finally, *Spitzer Space Telescope* 24 μm images of extragalactic survey fields have revealed extremely dust-obscured galaxies (e.g., Dey et al. 2008). If, certainly, the presence of dust can significantly change the chemistry and dynamics of early galaxies, we can also hope to learn more about the nature of dust and about its long cosmic history. Thus, we shall not limit our research to the Milky Way. We plan to extend our analysis to external galaxies, starting with our closest neighbors, the Small and Large Magellanic Clouds.

APPENDIX

THE INVERSE PROBLEM

The problem of retrieving interstellar dust properties from an extinction curve falls within the general class of inverse problems that do not have a direct solution (such as to describe a dragon from its tracks; Bohren & Huffman 1983). In this Appendix, we briefly outline the formalism to handle nonlinear inverse problems. Interested readers are referred to the exhaustive treatments of Twomey (1977) and Rodgers (2000) for details.

The retrieval procedure is based on the fitting of the observational data with a suitable forward model. We deal with a number $\mathcal{L}$ of unknown quantities (e.g., void fraction, mantle thickness, etc.) represented by a state vector $a \in \mathbb{R}^{\mathcal{L}}$ and a number $\mathcal{M}$ of data described by the vector $N_{\mathrm{obs}} \in \mathbb{R}^{\mathcal{M}}$, containing the extinction observed at different wave numbers. A forward model $F(n,a)$ describes the observations through a set of $\mathcal{N}$ fixed model parameters (the vector $n \in \mathbb{R}^{\mathcal{N}}$) and $\mathcal{L}$ adjustable parameters, providing a link between the observation vector and the state vector:

$$N_{\mathrm{obs}} = F(n, a) + \epsilon. \qquad (A1)$$

In the following we will omit n since by definition it is a set of fixed model parameters. If one considers a small neighborhood of some reference state vector a_0, the retrieval problem can be linearized, and Equation (A1) can be expressed as

$$R^i = \sum_l \left(\frac{\partial F^i}{\partial a^l} \right)_{a_0} \left(a^l - a_0^l \right) + \epsilon^i \sim \sum_l K_l^i \left(a^l - a_0^l \right) \quad (A2)$$

or more concisely, in matrix form

$$R = \mathbf{K}(a - a_0) + \epsilon \sim \mathbf{K}(a - a_0),$$

where $R = N_{\mathrm{obs}} - F(a_0)$ is the vector of the residuals with dimension $\mathcal{M}$. Its elements are equal to the difference between the observations and the corresponding simulation calculated using the reference state vector a_0. $\mathbf{K}$ is the Jacobian matrix of F with respect to the $\mathcal{L}$ adjustable parameters, having $\mathcal{M}$

rows and $\mathcal{L}$ columns. The goal of the retrieval operation is the determination of the difference vector $a - a_0$ in order to improve the previous estimate a_0. The problem relies on the construction of a solution matrix $\mathbf{G}$ (having $\mathcal{L}$ rows and $\mathcal{M}$ columns) which, applied to the vector $R = N_{\mathrm{obs}} - F(a_0)$, provides $a - a_0$. The retrieval procedure consists of a search for the parameters that produce the best simulation of the observations, by means of the minimization of the χ^2 merit function:

$$
\chi^2 = \sum_{i,j} (V^{-1})_{ij} \left(N_{\mathrm{obs}}^i - F^i(a) \right) \left(N_{\mathrm{obs}}^j - F^j(\boldsymbol{a}) \right)
$$

$$
\simeq \sum_{i,j} (V^{-1})_{ij} \left[R^i - \sum_l K_l^i (a^l - a_0^l) \right]
$$

$$
\times \left[R^j - \sum_l K_l^j (a^l - a_0^l) \right] \tag{A3}
$$

again, more concisely:

$$
\chi^2 = [N_{\mathrm{obs}} - F(a)]^T \mathbf{V}^{-1} [N_{\mathrm{obs}} - F(a)]
$$
$$
\simeq [R - \mathbf{K}(a - a_0)]^T \mathbf{V}^{-1} [R - \mathbf{K}(a - a_0)],
$$

where the superscript T indicates the transpose and the weight $\mathbf{V}$ is a square matrix of dimension $\mathcal{M}$, the variance-covariance matrix (VCM) associated with the vector N_{obs} and hence R (Rodgers 2000). When no further correction is to be applied to the state vector, $a - a_0$ becomes very nearly a null vector, and Equation (A3) reads as

$$
\chi^2 = R^T \mathbf{V}^{-1} R. \tag{A4}
$$

In general, the observations do not depend linearly on the unknown parameters a. As a consequence Equation (A4) is not a quadratic function of the unknowns. However, for a sufficiently close to a reference vector a_0 we may assume that the χ^2 function is well approximated by its Taylor expansion in a_0, truncated at second order, namely

$$
\chi^2(a) = \chi^2(a_0) + \sum_l \left(\frac{\partial \chi^2}{\partial a^l} \right)_{a_0} (a^l - a_0^l)
$$
$$
+ \frac{1}{2} \sum_{l,m} \left(\frac{\partial^2 \chi^2}{\partial a^l \partial a^m} \right)_{a_0} (a^l - a_0^l)(a^m - a_0^m)
$$
$$
= \chi^2(a_0) + (\nabla \chi^2)_{a_0}^T (a - a_0)
$$
$$
+ \frac{1}{2} (a - a_0)^T (\nabla^2 \chi^2)_{a_0} (a - a_0), \tag{A5}
$$

where ∇ and ∇^2 indicate the gradient and the Hessian matrix, respectively. Writing the gradient and the Hessian matrices of $F(n,a)$ explicitly, we obtain

$$
\chi^2(a) = \chi^2(a_0) - 2 \sum_{i,j,l} (V^{-1})_{ij} \left(\frac{\partial F^i}{\partial a^l} \right)_{a_0} \left(N_{\mathrm{obs}}^j - F^j(a_0) \right)(a^l - a_0^l)
$$
$$
+ \sum_{i,j,l,m} (V^{-1})_{ij} \left[\left(\frac{\partial F^i}{\partial a^l} \right)_{a_0} \left(\frac{\partial F^j}{\partial a^m} \right)_{a_0} \right.
$$
$$
\left. - \left(\frac{\partial^2 F^i}{\partial a^l \partial a^m} \right)_{a_0} \left(N_{\mathrm{obs}}^j - F^j(a_0) \right) \right] (a^l - a_0^l)(a^m - a_0^m)
$$

$$
= \chi^2(a_0) - 2 \sum_{i,j,l} (V^{-1})_{ij} K_l^i R^j (a^l - a_0^l)
$$
$$
+ \sum_{i,j,l,m} (V^{-1})_{ij} \left[K_l^i K_m^j - H_{lm}^i R^j \right] (a^l - a_0^l)(a^m - a_0^m)
$$
$$
= \chi^2(a_0) - 2 (\mathbf{K}^T \mathbf{V}^{-1} R)^T (a - a_0)
$$
$$
+ (a - a_0)^T \mathbf{K}^T \mathbf{V}^{-1} \mathbf{K} (a - a_0)
$$
$$
- \sum_{i,j,l,m} (V^{-1})_{ij} H_{lm}^i R^j (a^l - a_0^l)(a^m - a_0^m). \tag{A6}
$$

Neglecting the term with the second derivatives of $F(a)$, i.e., H_{lm}^i, in Equation (A6), the value of $a - a_0$ that minimizes the χ^2 function is the Gauss–Newton solution (e.g., Gill et al. 1981):

$$
a = a_0 + \mathbf{A}^{-1} \mathbf{K}^T \mathbf{V}^{-1} R, \tag{A7}
$$

where $\mathbf{A} = \mathbf{K}^T \mathbf{V}^{-1} \mathbf{K}$. The solution matrix of the inverse problem is thus

$$
\mathbf{G} = \mathbf{A}^{-1} \mathbf{K}^T \mathbf{V}^{-1}. \tag{A8}
$$

The improved estimate, Equation (A7), can be used as a new assumed state vector for a further iteration of the retrieval process. The iteration is successful, i.e., $\chi^2(a) < \chi^2(a_0)$, only if nonlinearities are not too large (hence justifying the neglect of the second derivative term above).

The Levemberg–Marquardt method introduces a modification to the iterative procedure outlined above, by substituting the matrix $\mathbf{A}$ with a matrix $\hat{\mathbf{A}}$ defined according to the following rule:

$$
\hat{\mathbf{A}}_{ij} = \mathbf{K}^T \mathbf{V}^{-1} \mathbf{K} \qquad \text{for } i \neq j
$$
$$
\hat{\mathbf{A}}_{ij} = \mathbf{K}^T \mathbf{V}^{-1} \mathbf{K} (1 + \lambda) \quad \text{for } i = j, \tag{A9}
$$

where λ is a damping parameter that increases the value of the diagonal elements. The modification enables a faster convergence and is recommended in the case of strongly nonlinear problems. The factor λ is initialized to a user-defined value, and, during the retrieval iterations, it is increased or decreased depending on whether the χ^2 function increases or decreases. For very small values of λ, $\hat{\mathbf{A}} \simeq \mathbf{A}$, hence going back to the simple linearized method outlined above; for large λ values, conversely, successive iterations go in the direction of the steepest descent of χ^2, with ever smaller steps for ever larger λ. While $\hat{\mathbf{A}}$ is guaranteed to be nonsingular, it may be very nearly singular if the residual changes very little in some direction of a space. In that case, $\mathbf{G}$ becomes numerically unstable, and tends to produce unrealistic (extremely large) shifts in those directions of a, thereby producing unphysical solutions. To overcome this problem, we use the SVD to invert $\hat{\mathbf{A}}$, and set to zero all "singular" values that are too low, which are exactly those directions in which the residual is almost constant. This stabilizes the procedure.

In the case of diagonal VCM, $(\mathbf{V}^{-1})_{ij} = \sigma_i^{-1} \delta_{ij}$, the χ^2 prescription becomes

$$
\chi^2(a) = \sum_i \left(\frac{N_{\mathrm{obs}}^i - F^i(a)}{\sigma_i} \right)^2, \tag{A10}
$$

where σ_i is the variance of the ith data point. In this case, $\mathbf{V_a} = \alpha^{-1}$ with $\alpha_{ij} = 1/2 \, \partial^2 \chi^2 / \partial a_i \partial a_j$ (e.g., Press et al. 1992).

REFERENCES

Allamandola, L. J., Tielens, A. G. G. M., & Barker, J. R. 1985, ApJ, 290, 25
Ashok, J., Varaprasad, P. L. H., & Birch, J. R. 1991, in Handbook of Optical Constants of Solids II, ed. E. D. Palik (New York: Academic), 957
Asplund, M., Grevesse, N., Sauval, A. J., & Scott, P. 2009, ARA&A, 47, 481
Berné, O., Joblin, C., Deville, Y., et al. 2007, A&A, 469, 575
Boersma, C., Allamandola, L. J., Bauschlicher, C., et al. 2013, BAAS, 221, 440.09
Bohren, C. F., & Huffman, D. R. 1983, Absorption and Scattering of Light by Small Particles (New York: Wiley)
Borghese, F., Denti, P., & Saija, R. 2007, Scattering by Model Nonspherical Particles (2nd ed.; Heidelberg: Springer)
Borghese, F., Denti, P., Saija, R., Toscano, G., & Sindoni, O. I. 1987, JOSAA, 4, 1984
Cardelli, J. A., Clayton, G. C., & Mathis, J. S. 1989, ApJ, 345, 245
Cecchi-Pestellini, C., Cacciola, A., Iatì, M. A., et al. 2010, MNRAS, 408, 535
Cecchi-Pestellini, C., Malloci, G., Mulas, G., Joblin, C., & Williams, D. A. 2008, A&A, 486, L25
Cecchi-Pestellini, C., Saija, R., Iatì, M. A., et al. 2005, ApJ, 624, 223
Cecchi-Pestellini, C., & Williams, D. A. 1998, MNRAS, 296, 414
Cesarsky, D., Lequeux, J., Ryter, C., & Gérin, M. 2000, A&A, 354, L87
Clayton, G. C., Gordon, K. D., Salama, F., et al. 2003, ApJ, 592, 947
Dèsert, F. X., Boulanger, F., Léger, A., Puget, J. L., & Sellgren, K. 1986, A&A, 159, 328
Dey, A., Soifer, B. T., Desai, V., et al. 2008, ApJ, 677, 943
Draine, B. T. 1985, ApJS, 57, 587
Draine, B. T. 1988, ApJ, 333, 848
Draine, B. T. 2003, ARA&A, 41, 241
Draine, B. T., & Li, A. 2001, ApJ, 551, 807
Duley, W. W. 2006, ApJL, 639, L59
Fitzpatrick, E. L. 2004, in ASP Conf. Ser. 309, Astrophysics of Dust, ed. A. N. Witt, G. C. Clayton, & B. T. Draine (San Francisco, CA: ASP), 33
Fitzpatrick, E. L., & Massa, D. 2007, ApJ, 663, 320
Fitzpatrick, E. L., & Massa, D. 2009, ApJ, 699, 1209
Furton, D. G., & Witt, A. N. 1993, ApJL, 415, L51
Gill, P. E., Murray, W., & Wright, M. H. 1981, Practical Optimization (San Diego, CA: Academic)
Gordon, K. D., Cartledge, S., & Clayton, G. C. 2009, ApJ, 705, 1320
Gudennavar, S. B., Bubbly, S. G., Preethi, K., & Murthy, J. 2012, ApJS, 199, 8
Hecht, J. H. 1986, ApJ, 305, 817
Iatí, M. A., Cecchi-Pestellini, C., Cacciola, A., et al. 2011, JQSRT, 112, 1898
Iatí, M. A., Cecchi-Pestellini, C., Williams, D. A., et al. 2001, MNRAS, 322, 749
Iatì, M. A., Saija, R., Borghese, F., et al. 2008, MNRAS, 384, 591
Jackson, J. D. 1975, Classical Electrodynamics (New York: Wiley)
Jones, A. P., Duley, W. W., & Williams, D. A. 1990, QJRAS, 31, 567
Krełowski, J., & Strobel, A. 2012, AN, 333, 60
Landstrass, M. I., & Ravi, K. V. 1989, ApPhL, 55, 975
Lèger, A., & Puget, J.-L. 1984, A&A, 137, L5
Li, A. 2009, Natur, 459, 173
Li, A., & Draine, B. T. 2001, ApJL, 550, L213
Malloci, G., Joblin, C., & Mulas, G. 2007, CP, 332, 353
Mathis, J. S. 1996, ApJ, 472, 643
Mathis, J. S. 2000, JGR, 105, 10269
Mathis, J. S., Rumpl, W., & Nordsieck, K. H. 1977, ApJ, 217, 425
Mathis, J. S., & Whiffen, G. 1989, ApJ, 341, 808
Paradis, D., Paladini, R., Noriega-Crespo, A., et al. 2011, ApJ, 735, 6
Press, W. H., Teukolsky, S. A., Wetterling, W. T., & Flannery, B. P. 1992, Numerical Recipes (Cambridge: Cambridge Univ. Press)
Rai, R. K., & Rastogi, S. 2010, MNRAS, 401, 2722
Rapacioli, M., Joblin, C., & Boissel, P. 2005, A&A, 429, 193
Rodgers, C. D. 2000, Inverse Methods for Atmospheric Sounding (Singapore: World Scientific)
Rose, E. M. 1957, Elementary Theory of Angular Momentum (New York: Wiley)
Rouleau, F., & Martin, P. G. 1991, ApJ, 377, 526
Saija, R., Iatí, M. A., Borghese, F., et al. 2001, ApJ, 539, 993
Sandstrom, K. M., Bolatto, A. D., Bot, C., et al. 2012, ApJ, 744, 20
Sellgren, K. 1984, ApJ, 277, 623
Steglich, M., Bouwman, J., Huisken, F., & Henning, T. 2011, ApJ, 742, 2
Twomey, S. 1977, Introduction to the Mathematics of the Inversion in Remote Sensing and Indirect Measurements (New York: Elsevier)
Valencic, L. A., Clayton, G. C., & Gordon, K. D. 2004, ApJ, 616, 912
Voshchinnikov, N. V. 2012, JQSRT, 113, 2334
Voshchinnikov, N. V., & Henning, Th. 2008, A&A, 483, L9
Waterman, P. C. 1971, PhRvD, 3, 825
Whittet, D. C. 2002, Dust in the Galactic Environment (Bristol: IOP Publishing)
Williams, D. A. 2000, A&G, 41, 030000
Wiscombe, W. J. 1980, ApOpt, 19, 1505
Wright, E. L. 1987, ApJ, 320, 818
Wyatt, P. J. 1962, PhRvB, 127, 1771
Zonca, A., Cecchi-Pestellini, C., Mulas, G., & Malloci, G. 2011, MNRAS, 410, 1932
Zubko, V., Dwek, E., & Arendt, G. 2004, ApJS, 152, 211