Piece-wise multipole-expansion implicit solvation for arbitrarily shaped molecular solutes

Jakob Filser,*,† Karsten Reuter,†,‡ and Harald Oberhofer†

†*Chair for Theoretical Chemistry and Catalysis Research Center, Technische Universität München, Lichtenbergstr. 4, 85747 Garching, Germany*

‡*Fritz-Haber-Institut der Max-Planck-Gesellschaft, Faradayweg 4-6, 14195 Berlin, Germany*

E-mail: jakob.filser@tum.de

Abstract

The multipole-expansion (MPE) model is an implicit solvation model used to efficiently incorporate solvent effects in quantum chemistry. Even within the recent direct approach, the multipole basis used in MPE to express the dielectric response still solves the electrostatic problem inefficiently or not at all for solutes larger than ≈ 10 non-hydrogen atoms. In existing MPE parameterizations, the resulting systematic underestimation of the electrostatic solute-solvent interaction is presently compensated for by a systematic overestimation of non-electrostatic attractive interactions. Even though the MPE model can thus reproduce experimental free energies of solvation of small molecules remarkably well, the inherent error cancellation makes it hard to assign physical meaning to the individual free energy terms in the model, raising concerns about transferability. Here, we resolve this issue by solving the electrostatic problem piece-wise in 3D regions centered around all non-hydrogen nuclei of the solute, ensuring reliable convergence of the multipole series. The resulting method thus allows for a much improved reproduction of the dielectric response of a medium to a solute. Employing a reduced non-electrostatic model with a single free parameter, in addition

1

to the density isovalue defining the solvation cavity, our method yields free energies of solvation of neutral, anionic and cationic solutes in water in good agreement with experiment.

1 Introduction

With many—if not most—chemical processes of interest in biochemistry,[1] electrochemistry,[2,3] catalysis[4] and related fields[5] taking place in solution or at liquid interfaces with other phases, it is not surprising that the modelling of liquid environments is a major topic in computational chemistry.[4,6–9] The fundamental issue is that, from a first-principles point of view, the separation of a liquid system into one or multiple 'solutes' and a 'solvent' is artificial. Electrons and nuclei of the solvent and the solute follow the same quantum mechanical (QM) and statistical laws, and the straightforward approach would be to treat both at an explicit, quantum mechanical level.[6] In practice, however, treating both at the same level of theory quickly turns out as computationally intractable already for rather trivial systems. The numbers of electrons and nuclei in this complete picture are much larger than those of the solute alone, as usually a large number of solvent molecules would have to be included in the model system to obtain converged results. In addition, the solvent typically consists of small, highly mobile molecules. While some short-range ordering may occur, there is generally no long-range order and a thermodynamic sampling would be necessary to accurately account for all solvent degrees of freedom.[4,6]

This great effort to accurately account for solvent effects is in stark contrast to the empirically known fact that typically, if one is interested in the properties of a molecular system in solution (but not the properties of the solution under the influence of the solute), the solvent plays only a minor role compared to the atomic structure of the solute itself.[6] It is therefore reasonable—and necessary—to simulate only the solute at a detailed, quantum mechanical level and treat the solvent merely as an environment for the solute, using a simpler model.[7]

There are, of course, different practical approaches to this general idea.[4,8] One highly popular group of methods, which we focus on in this work, are the so-called implicit solvation or continuum embedding models.[6,7,9] The general idea is to ignore any atomistic detail of the solvent altogether—although in some cases one or few individual solvent molecules may still need to be included explicitly.[10,11] Instead, the time-averaged effect of the solvent on the solute is modelled using the former's macroscopic properties, most notably its relative dielectric permittivity, and/or empirical models fitted to experimental reference data, typically free energies of solvation in the limit of infinite dilution.[12]

At the heart of many implicit solvent models stands the interaction of the solute charge density with a dielectric medium. The difference between the individual methods lies mostly in the way this electrostatic problem is solved in practice. The here discussed multipole-expansion (MPE) methods expand the dielectric response of the solvent in a multipole series. Earlier versions[13–18] can solve the electrostatic problem exactly only when the charge density is completely localized within the cavity,[19] which is the region around the solute assumed to be inaccessible to the solvent. This leads to the so-called outlying charge error when solute charge density and model dielectric overlap, as is often necessary to accurately describe the strong interaction with solvents such as water.[20,21]

This rather severe limitation of the MPE method was recently lifted through a novel regularization of the electrostatic potential and the so-called direct approach by Sinstein *et al.*[19,20] This generalization now allows the modelling of stronger electrostatic solute-solvent interactions, opening up the method for a more general use. Unfortunately, though, practical applications of both the original and the direct approach of the MPE model so far suffered from convergence and accuracy issues for larger solutes. The reason for this is that in these models the expansion of the electrostatic potential inside the cavity relies on a single expansion center, typically in the geometric center of the solute. Thus solutes that deviate strongly from a roughly spherical shape tend to necessitate rather high multipole expansion orders.[20] Indeed in this work, we show that for some solutes the multipole series might

fail to converge altogether. As a remedy, we here present a new approach, based on the partition of space into small subcavities, similar in philosophy to the domain decomposition conductor-like screening model (ddCOSMO).[22–24] This adaptation now allows for a true multi-center expansion of the electrostatic potential. We show that this new variant of the MPE model lifts size restrictions of the solute and allows for a much more stable convergence with expansion order.

This work is organized as follows: After briefly recapitulating the current state of the art of the MPE method in section 2, we thoroughly test the model on a popular benchmark set[12] described in section 3, in section 4. We develop the modified MPE subcavities method in section 5 and demonstrate its efficacy in section 6. There we further present a new parameterization of the model for water as a solvent. Lastly, we discuss how the success of the original method[20] in describing solution in water is based on a size-dependent systematic error cancellation. Our modified method is largely free of such phenomena, improving transferability.

2 Theoretical and technical background

2.1 Scope of the model

We consider a cluster-like (i.e. not periodic in any real-space direction) molecular system (henceforth referred to as 'solute') surrounded by a liquid solvent. One of the main characteristics of the solvent in any implicit solvation model is its dielectric permittivity. Therefore, the terms 'solvent' and 'dielectric (medium)' are used interchangeably in the present work. In contrast, the term '(in) vacuum' will refer to the absence of a solvent. The solute is treated at a density-functional theory (DFT) level of theory,[20,25–29] while the molecules of the solvent are not treated explicitly. Instead, their time-averaged dielectric response is approximated in a classical electrostatic model.[6,7,9,20] Space is divided into a finite number of regions X_i, where the index i can take arbitrary values, depending on the specific problem definition.

Each X_i is assigned to either the dielectric medium or vacuum, with the latter describing the inside of the so-called solvation cavity, i.e. the region around the solute which is assumed to be inaccessible by solvent molecules. Any dielectric effects within this solvation cavity are explicitly accounted for at the DFT level of theory. Within the region belonging to the implicit dielectric, the dielectric response is instead assumed to be local, linear, homogeneous and isotropic. It can thus be described by a scalar relative permittivity ε_i, which is constant within a particular dielectric.

$$\varepsilon_i = \varepsilon(\mathbf{r} \in X_i) = \text{const.} \quad . \tag{1}$$

There are, however, transitions between the different dielectric regions. There, the permittivity may show a complicated spatial behavior,[30] which implicit models tend to replace with simple smeared out step-functions for the sake of efficiency.[21,31,32] In this spirit, the MPE method, together with other popular approaches such as COSMO,[33] SMx,[10,11] or the original PCM model,[7] defines a sharp transition on the 2D boundaries B_{ij} between 3D regions X_i and X_j. In section 5 below we will generalize this concept to an arbitrary number of regions in space X_i. In the following, we therefore keep the discussion at the level of general regions X_i which obey eq. (1). The exact definition of their shape will be given in the next section.

Finally, we restrict ourselves to the time-independent and non-periodic case in the absence of external fields other than that of the nuclei for the time being. Furthermore, the solvent contains no ions (the solute may, however, be ionic). This means that ionic liquids as well as saline solutions[34] lie beyond the scope of the present model.

2.2 Electrostatic problem

The electrostatic problem is solved identically to previous work, using the ansatz[20]

$$\Phi(\mathbf{r}) = \varepsilon^{-1}(\mathbf{r})\Phi_{\text{H}}(\mathbf{r}) + \Phi_{\text{MPE}}(\mathbf{r}) \quad , \tag{2}$$

where the electrostatic potential Φ_{MPE} is due to the solvent response. Φ_{H} is the classical Coulomb—or, in a DFT[20,25–29] context, Hartree—potential[35] of the solute charge density ρ

$$\Phi_{\text{H}}(\mathbf{r}) = \int_{\mathbb{R}^3} d^3\mathbf{r}' \, \frac{\rho(\mathbf{r}')}{\|\mathbf{r} - \mathbf{r}'\|} \quad . \tag{3}$$

The latter obeys Poisson's equation in vacuum,

$$\nabla^2 \Phi_{\text{H}}(\mathbf{r}) = -4\pi\rho(\mathbf{r}) \quad . \tag{4}$$

The net charge density ρ includes both electrons and nuclei of the solute,

$$\rho(\mathbf{r}) = -\rho_{\text{el}}(\mathbf{r}) + \sum_{N=1}^{N_{\text{max}}} Q_N \delta(\mathbf{r} - \mathbf{R}_N) \quad . \tag{5}$$

Here, ρ_{el} is the solute electron density, δ is the delta function, and $\mathbf{R}_N$ and Q_N are respectively the nuclear coordinates and charges of the N_{max} nuclei constituting the explicitly treated solute.[20]

As mentioned above, implicit solvation models generally partition space into a dielectric medium X_{Q} and a so-called 'cavity' X_{R} in which the solute resides.[6,7,9,20] For these regions, the relative dielectric permittivity is defined by

$$\varepsilon(\mathbf{r} \in X_{\text{R}}) = 1 \tag{6a}$$

$$\varepsilon(\mathbf{r} \in X_{\text{Q}}) = \varepsilon_{\text{bulk}} \quad , \tag{6b}$$

where $\varepsilon_{\text{bulk}}$ is the experimentally determined macroscopic permittivity of the solvent. The spatial extents of these regions are defined by an isovalue ρ_{iso} of the solute electronic density

through

$$X_{\mathrm{R}} = \{\mathbf{r} : \widetilde{\rho}_{\mathrm{el}}(\mathbf{r}) > \rho_{\mathrm{iso}}\} \tag{7a}$$

$$X_{\mathrm{Q}} = \{\mathbf{r} : \widetilde{\rho}_{\mathrm{el}}(\mathbf{r}) < \rho_{\mathrm{iso}}\} \tag{7b}$$

An electronic density isosurface is per definition smooth, continuous and follows the shape of the solute. It is thus a convenient way to define solvation cavities and is commonly used in various implicit solvation models, implemented in different electronic structure packages.[20,21,34,36] The isovalue ρ_{iso} thereby takes the form of a free parameter that needs to be determined e.g. by fitting to a suitable training set, cf. section 6.2. However, other definitions, e.g. based on atom-centered spheres,[11,22,37] are also commonly used. The method described in section 5 makes no assumptions about the cavity shape other than that it roughly follows the shape of the solute and that it is sharp in the sense of eqs. (6a) and (6b). It can thus, in principle, be used together with any such cavity definition. An electronic density isosurface is merely the choice which we use to demonstrate the method in the present work.

Being an effective model, though, the exact choice of which electron density $\widetilde{\rho}_{\mathrm{el}}$ to use in the above definition is ambiguous. Different possibilities have been proposed[19,20]. In the present work we use the superposition of free atom electron densities ρ_{free}, i.e. the electron densities of the solute's individual atoms in vacuum.[25] We refer the reader to the cited publications for the other options. It has been shown before[19,20] that those choices lead to a systematic error in the free energy of solvation for anionic solutes. The superposition of ρ_{free} does not suffer from this issue and is able to describe neutral, cationic and anionic solvents all with the same set of parameters.[19] It should be noted here that in practice, when using a hybrid exchange correlation functional, the DFT software `FHI-aims`[20,25–29] which we use throughout the present work calculates ρ_{free} from its generalized gradient approximation (GGA) part, so e.g. the PBE[38] and HSE06[39,40] functionals will in practice yield the exact same cavity.

Note that the non self-consistent electron density $\widetilde{\rho}_{\mathrm{el}}$ is used solely for the cavity definition. The electrostatic potentials Φ_{H} and Φ_{MPE}, which depend on the former, are, of course, calculated self-consistently from the instantaneous electron density ρ_{el} of the solute at each SCF step.

Next, we need to choose a basis to represent the solvent response to the electrostatic potential given in eq. (2). It has been shown that in a region of constant $\varepsilon(\mathbf{r})$, Φ_{MPE} is a harmonic function, lending itself to a series expansion in regular solid harmonic functions $\mathcal{R}_m^l$ or their irregular counterparts $\mathcal{I}_m^l$.[13–18,20]

$$\mathcal{R}_m^l(\mathbf{r}) = r^l\, Y_m^l(\theta, \varphi) \tag{8a}$$

$$\mathcal{I}_m^l(\mathbf{r}) = r^{-(l+1)}\, Y_m^l(\theta, \varphi) \tag{8b}$$

$$l, m \in \mathbb{Z}, \quad l \geq 0, \quad -l \leq m \leq l$$

Here, Y_m^l are the real valued spherical harmonics (cf. ref. 25). We thus use a piece-wise definition of Φ_{MPE}, treating each region X_i separately.[20] We define

$$\Phi_i = \Phi_{\mathrm{MPE}}(\mathbf{r} \in X_i) \quad . \tag{9}$$

If X_i is bounded, i.e. if

$$\exists r \in \mathbb{R}^+ : r > \|\mathbf{r}\| \,\forall \mathbf{r} \in X_i \quad , \tag{10}$$

as is the case inside the solvation cavity X_{R}, then Φ_i can be expanded in a set $\{\mathcal{R}_m^l\}_K$ of $\mathcal{R}_m^l(\mathbf{r} - \mathbf{r}_K)$ around a center $\mathbf{r}_K$[13–18,20]

$$\Phi_i(\mathbf{r}) = \sum_{l=0}^{l_{\mathrm{max}}^{\mathrm{R}}} \sum_{m=-l}^{l} R_K^{(l,m)} \mathcal{R}_m^l(\mathbf{r} - \mathbf{r}_K) \quad , \tag{11}$$

with expansion coefficients $R_K^{(l,m)}$. If X_i is unbounded, i.e. if eq. (10) is not satisfied, as is the case in X_{Q}, then Φ_i can be expanded in a union of sets $\{\mathcal{I}_m^l\}_J$ of $\mathcal{I}_m^l(\mathbf{r} - \mathbf{r}_J)$ with multiple

centers $\mathbf{r}_J$[13–18,20]

$$\Phi_i(\mathbf{r}) = \sum_{J=1}^{J_{\max}} \sum_{l=0}^{l_{\max}^{\mathrm{Q}}} \sum_{m=-l}^{l} Q_J^{(l,m)} \mathcal{I}_m^l(\mathbf{r} - \mathbf{r}_J) \quad , \tag{12}$$

and with expansion coefficients $Q_J^{(l,m)}$. The only singularities are at $\mathbf{r} = \mathbf{r}_J$. By choosing $\mathbf{r}_J \notin X_i \forall J$, Φ_i is continuous in X_i by construction.[13–18,20]

Φ_{MPE} can then be uniquely defined by imposing continuity of the total electrostatic potential and of the electric flux density at B_{ij}.[7] Together with the regularization in eq. (2), this leads to the following boundary conditions[20]

$$\Phi_i(\mathbf{r}) - \Phi_j(\mathbf{r}) = \left(\varepsilon_j^{-1} - \varepsilon_i^{-1}\right) \Phi_{\mathrm{H}}(\mathbf{r}) \tag{13a}$$

$$\mathbf{n_r}\nabla\left(\varepsilon_i\Phi_i(\mathbf{r}) - \varepsilon_j\Phi_j(\mathbf{r})\right) = 0 \tag{13b}$$

$$\forall \mathbf{r} \in B_{ij}, \forall i \neq j \quad ,$$

where $\mathbf{n_r}$ is the surface normal of B_{ij} at $\mathbf{r}$.

We insert eq. (11) and eq. (12) into eq. (13), for bounded and unbounded $X_{i/j}$, respectively. Evaluating it at sets of points situated at the interfaces B_{ij} then essentially turns the problem of determining the dielectric response into an algebraic problem of solving an overdetermined system of linear equations (SLE) for the expansion coefficients of the potential.[16,20]

$$\mathbf{Ax} = \mathbf{b} \tag{14}$$

Here, $\mathbf{A}$ contains the basis functions $\mathcal{R}_m^l$ and $\mathcal{I}_m^l$ for eq. (13a) and their scaled derivates for eq. (13b), evaluated at said discrete set of points on B_{ij}. $\mathbf{b}$ contains the scaled Hartree potential for eq. (13a) and zero for eq. (13b). The solution vector $\mathbf{x}$ contains the expansion coefficients $R_K^{(l,m)}$ and $I_J^{(l,m)}$. The discretization algorithm for B_{ij} aims to achieve a certain degree of determination d_{det}, i.e. the ratio between rows and columns in $\mathbf{A}$. Equation (14) can finally be solved using a number of algorithms which all more or less reduce to a linear least-squares fit. The technical details of this numeric solution have been described in ref.

9

20. As a measure for the quality of the solution to the linear algebra problem we use the adjusted coefficient of determination $\bar{R}^2$,[41–43] which essentially measures how much of the variability of the dependent variable of an equation is due to a linear relationship with the independent variable. The calculation of $\bar{R}^2$ is described in the supporting information (SI) of the present work.

With Φ_{MPE} defined in this manner, it is possible to calculate the free energy $G^{\mathrm{elstat}}(\{\mathbf{R}_N\})$ of a molecule electrostatically embedded into the solvent by replacing the Hartree potential with eq. (2) in a DFT[20,25–29] calculation. Using the total energy of the solute in vacuum $E_{\mathrm{vac}}(\{\mathbf{R}_N\})$, this then defines the electrostatic contribution to the solvation free energy of the solute, $\Delta G_{\mathrm{solv}}^{\mathrm{elstat}}(\{\mathbf{R}_N\}) = G^{\mathrm{elstat}}(\{\mathbf{R}_N\}) - E_{\mathrm{vac}}(\{\mathbf{R}_N\})$. Note that all these quantities refer to the respective electronic ground state and we have here explicitly written the dependence on the nuclear positions $\{\mathbf{R}_N\}$. For clarity of notation, this dependence will henceforth be dropped.

2.3   Non-electrostatic free energy contributions

There are, of course, further contributions to the free energy of a solute in solvent, beyond the hitherto discussed electrostatics. Specifically, solute and solvent also interact via Pauli repulsion and dispersion forces, and there are entropic and even grand-canonical contributions to consider due to the displacement of solvent molecules (cavity formation). We refer to recent reviews for a detailed overview over these, so-called non-electrostatic contributions generally considered in implicit solvation models.[7–9] Here, we only point out that these different contributions are often lumped together into effective expressions with a minimum number of free parameters. The solvation free energy is correspondingly written as arising from two contributions,

$$\Delta G_{\mathrm{solv}} = \Delta G_{\mathrm{solv}}^{\mathrm{elstat}} + \Delta G_{\mathrm{solv}}^{\mathrm{non-elstat}} \quad , \tag{15}$$

10

with $G_{\text{solv}}^{\text{non}-\text{elstat}}$ containing all non-electrostatic contributions. One frequent functional expression for the latter is a simple linear function of surface area A and volume V of the solvation cavity[21]

$$\Delta G_{\text{solv}}^{\text{non-elstat}} = \alpha A + \beta V \quad , \tag{16}$$

with α and β empirical parameters that are fitted to a set of experimental free energies of solvation. While this functional form was initially established for the self-consistent continuum solvation (SCCS) model,[21] it turns out that such a treatment of non-electrostatics can also be used with other electrostatic models such as MPE.[20]

Both A and V are a measure of molecular size and are therefore typically correlated. Considering the effective character of the functional form employed for the non-electrostatic free energy contribution, it is correspondingly possible to also use only one of its two terms, typically the surface area one, αA, at generally insignificant increases in the mean error of computed solvation free energies with respect to experimental references.[19–21] For our present purpose of assessing and removing error cancellation between the electrostatic and non-electrostatic contributions to ΔG_{solv}, having less flexibility in a non-electrostatic term with only one fit parameter is actually even beneficial. In this work, we correspondingly use as simple non-electrostatic model,

$$\Delta G_{\text{solv}}^{\text{non-elstat}} = \alpha A \quad . \tag{17}$$

Following earlier work,[20] we also treat $\Delta G_{\text{solv}}^{\text{non-elstat}}$ as a post-SCF correction instead of including it self-consistently in the Kohn-Sham[44] (KS) operator.

It is worth noting that in fitting the parameter α to experimental reference data,[12] see below, the free energy contributions due to nuclear degrees of freedom are to some degree accounted for in $\Delta G_{\text{solv}}^{\text{non-elstat}}$, at least in a statistical manner.[19] For unpolar organic solvents, the problem may also emerge that the optimized α becomes negative. With $\Delta G_{\text{solv}}^{\text{elstat}}$ negative by construction and then $\Delta G_{\text{solv}}^{\text{non-elstat}}$ also negative, the solvation free energy would

necessarily result as exothermic. The case of a solute being repelled by the solvent thus lies outside the scope of such a model. Notwithstanding, in the present work we focus on solvation in water, where α is typically positive[20,21] and this problem does not play a role.

2.4 Technical details

The computational setup used to calculate $\Delta G_{\mathrm{solv}}^{\mathrm{elstat}}$ is identical to the one of preceding work and we refer to the corresponding publication for all technical details.[20] Unless mentioned otherwise, all DFT calculations are performed with `FHI-aims`[20,25–29] (version `210415`) into which the present implicit solvent model was implemented. The *tight* default electronic basis sets and integration grids included in the `FHI-aims`[20,25–29] package are employed throughout. We use the Perdew, Burke, Ernzerhof (PBE) exchange-correlation (xc) functional[38] with collinear spin and an atom-wise scalar zeroth order regular approximation (ZORA)[25] for relativistic effects of the core electrons. Kohn-Sham levels are occupied through a Gaussian broadening scheme[45] with a rather narrow width of 0.01eV to account for the molecular nature of our solutes. Equation (14) was solved in a direct solver by QR factorization of the left-hand side matrix and subsequent singular value decomposition of the R matrix,[46,47] as described before.[20]

3   Benchmark systems and experimental reference data

All molecular geometries and experimental reference values for ΔG_{solv} used in the present work were taken from the Minnesota Solvation Database, version 2012.[12] We use all 790 small molecules included in the database to investigate the convergence of $\Delta G_{\mathrm{solv}}^{\mathrm{elstat}}$ with numeric parameters in the respective models, especially with the expansion order. We furthermore use this set to test the computational performance of our method. In order to keep these convergence studies as unbiased and to make the resulting parameters as generally applicable

as possible, we randomize the two physical model parameters ε and ρ_{iso}. For each of the 790 test solutes, two random numbers $\mathfrak{R}_{\varepsilon/\rho}$ are drawn from a uniform distribution between 0 and 1. The parameters are then calculated according to

$$\varepsilon = \frac{2}{\mathfrak{R}_{\varepsilon}} \tag{18a}$$

$$\rho_{\mathrm{iso}} = 10^{-(1+3\cdot\mathfrak{R}_{\rho})}\,\mathrm{e}\,\mathring{\mathrm{A}}^{-3} \quad , \tag{18b}$$

yielding $\varepsilon > 2$ with the distribution becoming increasingly sparse for higher values. A molecular geometry together with its ε and ρ_{iso} will henceforth be referred to as a system.

While any $\varepsilon \geq 1$ is physically possible, this distribution is chosen to roughly mimic the range of permittivities of real-life solvents, which are rarely ever much smaller than 2.[12] The isodensities are uniformly distributed on a logarithmic scale, ranging from 10^{-1} to $10^{-4}\,\mathrm{e}\,\mathring{\mathrm{A}}^{-3}$ which is roughly the range in which we expect to find optimized isodensities for different solvents based on previous experience with the original direct version of the MPE model.[19,20] Note that as the cavity tends to become smoother for lower isodensities, it is reasonable to assume that any set of converged parameters within the sampled isodensity range can also be applied to lower isodensities. In fact, we will see in the following section, specifically in table 1, that systems for which the electrostatic problem is hard to solve typically have high ρ_{iso}. To ensure comparability between calculations, the randomly chosen permittivities and isodensities differ only between solutes, but for each solute the same values are used for all convergence studies.

DFT calculations with implicit solvation may occasionally fail for various reasons. In our case, this happened mostly due to issues with the cavity surface discretization algorithm (cf. SI as well as ref. 20). Addressing these issues is beyond the scope of the present work. Due to such failures occurring only rarely, we simply ignore problematic systems for the time being. Unless mentioned otherwise, within each convergence study, performance test or parameterization, solutes for which at least one calculation failed for any reason are discarded

from further analysis.

For the parameterizations in section 6.2 we used all solutes for which ΔG_{solv} in water is available in the Minnesota Solvation Database. For some ionic solutes, a version clustered by one explicit water molecule and an unclustered version is available.[10,11] For these solutes, only the clustered versions were used. This set of training data is slightly larger than in previous work where only some subsets of the database were used.[20]

4   Convergence of the original MPE model

The original version of the direct MPE model[20] partitioned space only into the cavity and the dielectric medium, as defined in eqs. (7a) and (7b). The basis set for the dielectric response inside the cavity Φ_R were the regular solid harmonics $\{\mathcal{R}_m^l\}$ centered around the geometric center of the molecule. Within the perspective of the present generalization, we will henceforth call this model MPE-1c, with 1c indicating the single cavity X_R. Correspondingly, we will also denote electrostatic solvation free energies as $\Delta G_{\text{solv}}^{\text{elstat}}(nc)$ to indicate that they have been computed with an n cavity model. In this section, we explore the capabilities and limitations of MPE-1c, making no conceptual modifications with respect to the original publication.

4.1   Insufficiency of the solid harmonic basis

To highlight the initially mentioned multipole expansion issues of the MPE-1c model, we analyze the convergence behavior of $\Delta G_{\text{solv}}^{\text{elstat}}(l_{\text{max}}^R; 1c)$ with increasing expansion order l_{max}^R of Φ_R using the test systems described in section 3. Specifically, we test $l_{\text{max}}^R = 2, 4, 6, 8, 12, 16$ and 20. For Φ_Q, we use $l_{\text{max}}^Q = 6$.[20] As a reference, we use the $\Delta G_{\text{solv}}^{\text{elstat}}(l_{\text{max}}^R = 20; nc)$ obtained for the same systems with our new method MPE-nc at $l_{\text{max}}^R = 20$ and $l_{\text{max}}^Q = 8$. Further explaining this new method in section 5, we show in section 6.1 that these reference values themselves are converged up to on average $\lesssim 1\,\text{meV}$ and thus constitute a firm reference.

Additionally, we test the convergence of the adjusted coefficient of determination $\bar{R}^2$ of the solution to the discretized boundary conditions eqs. (13a) and (13b).
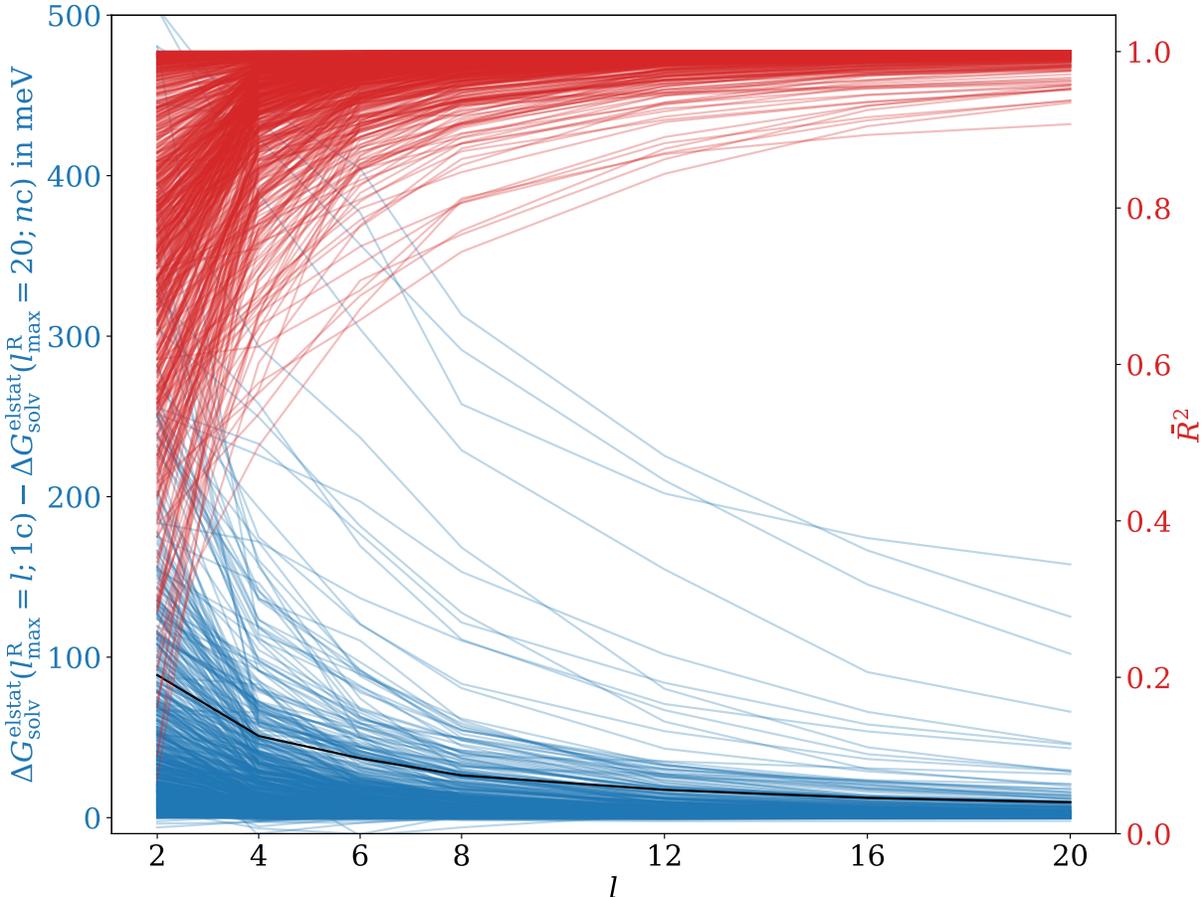
Figure 1: Convergence of the electrostatic contribution to the solvation free energy $\Delta G_{\text{solv}}^{\text{elstat}}(l_{\text{max}}^{\text{R}}; 1c)$ (blue) and adjusted coefficient of determination $\bar{R}^2$ (red) of the discretized boundary conditions eqs. (13a) and (13b) with expansion order $l_{\text{max}}^{\text{R}}$ of the MPE potential inside the cavity Φ_{R} in MPE-1c. Individual lines refer to individual systems contained in the test set as described in section 3. The root mean-square error over all systems is shown as a thick black line. As reference, we use $\Delta G_{\text{solv}}^{\text{elstat}}(l_{\text{max}}^{\text{R}} = 20; nc)$ calculated with the MPE-nc model (see section 5), the same reference as also in fig. 9.

The results are compiled in fig. 1 and show clearly that MPE-1c systematically underestimates the electrostatic interaction between solute and solvent at the originally published $l_{\text{max}}^{\text{R}} = 8$.[20] Even at the significantly higher $l_{\text{max}}^{\text{R}} = 20$, there are several test systems with errors in the order of $100\,\text{meV}$. The root mean-square error (RMSE) at that expansion order, in contrast, is only $9.6\,\text{meV}$. Below we will show that the larger errors are not outliers, but

15

rather a systematic error occurring for larger solutes. The RMSE is only relatively low, because smaller molecules are over-represented in the test set. Furthermore, for some systems,
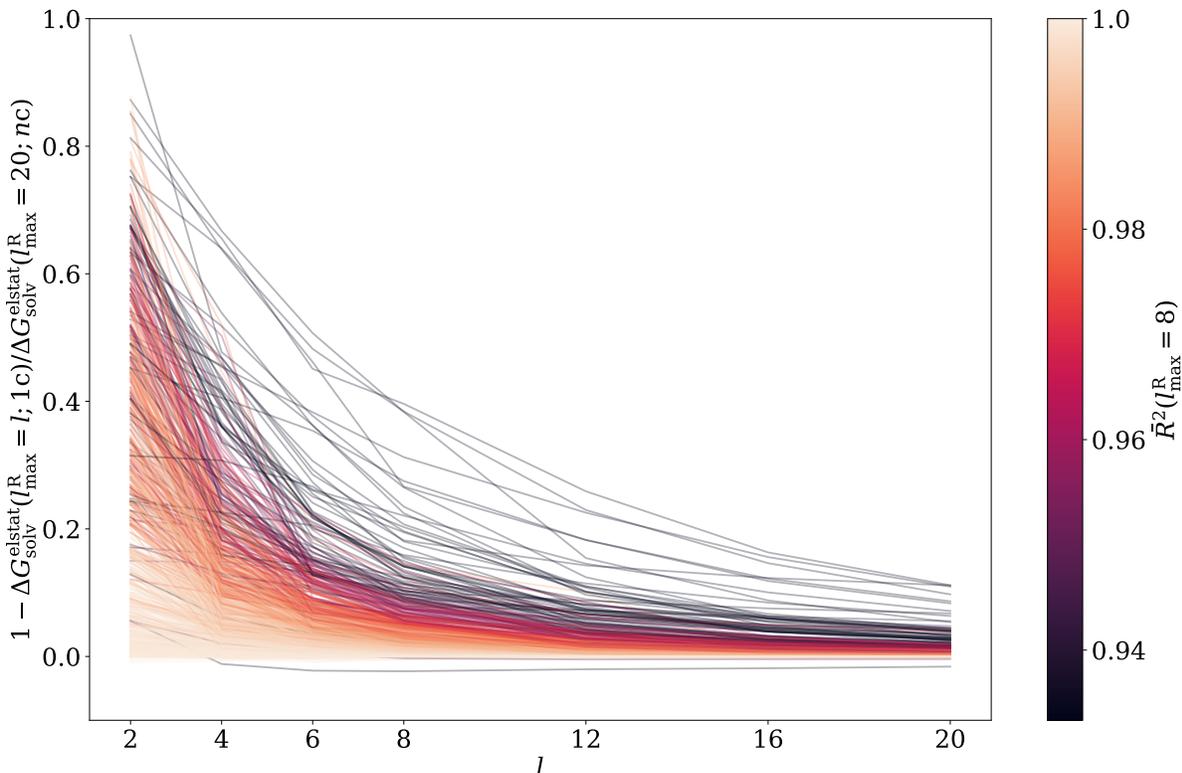
Figure 2: Relative error of the electrostatic contribution to the solvation free energy $\Delta G_{\text{solv}}^{\text{elstat}}(1\text{c})$ in MPE-1c vs. the MPE-nc reference. Individual lines reflect the individual systems contained in the test set. The lines are colored according to the adjusted coefficient of determination $\bar{R}^2$ of the discretized boundary conditions eqs. (13a) and (13b) at $l_{\text{max}}^{\text{R}} = 8$, which is the originally published value for this parameter.[20]

even at $l_{\text{max}}^{\text{R}} = 20$ the solid harmonic basis is insufficient to solve the electrostatic boundary conditions eqs. (13a) and (13b), as indicated by $\bar{R}^2 < 1$. It can straightforwardly be seen that the error in $\Delta G_{\text{solv}}^{\text{elstat}}$ is related to an insufficient solution of the electrostatic problem, by relating the relative error of $\Delta G_{\text{solv}}^{\text{elstat}}(1\text{c})$ to $\bar{R}^2$ as depicted in fig. 2.

Finally, fig. 3 shows that these errors in the electrostatic potential are related to the size of the molecules. The fraction of solutes with $\bar{R}^2 \approx 1$ decreases with increasing number of non-hydrogen (non-H) atoms in the solute. Given the dominance of small molecules (≤ 10 non-H atoms), the relatively low RMSE in $\Delta G_{\text{solv}}^{\text{elstat}}$, apparent in fig. 1, is thus only

16
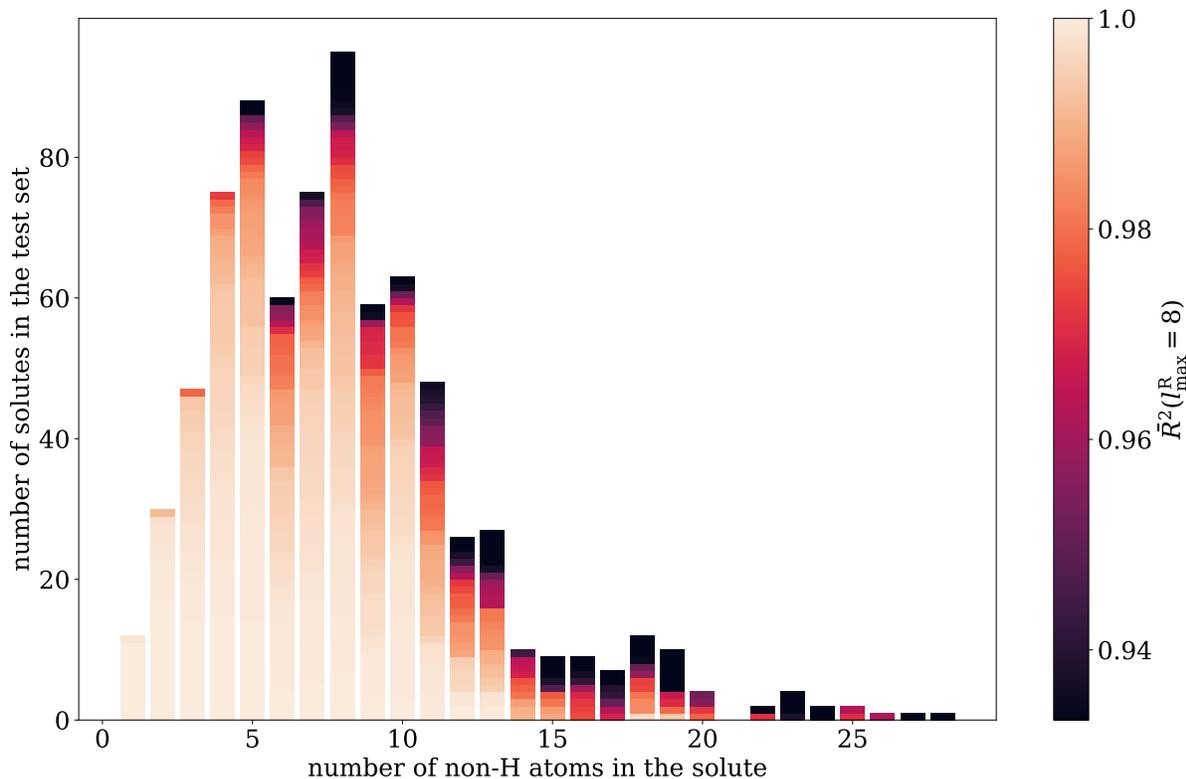
Figure 3: Histogram of the number of non-H atoms per solute in the test set, colored according to the adjusted coefficient of determination $\bar{R}^2$ of the discretized boundary conditions eqs. (13a) and (13b) at $l_{\max}^{\mathrm{R}} = 8$.

representative for such smaller molecules, with increasingly larger errors to be expected for larger molecules.

Taking all of these findings into account, we conclude that MPE-1c with the originally published parameters[20] suffers from a systematic underestimation of the electrostatic solute-solvent interaction. While negligibly small for small molecules, this error increases significantly for larger molecules.

4.2 The wrong basis functions or too few?

An intuitive explanation why MPE-1c fails for larger molecules would be that arbitrary harmonic potentials $\nabla^2 \Phi_{\mathrm{harm}}(\mathbf{r} \in X) = 0$ can be approximated to arbitrary accuracy in a finite series in $\{\mathcal{R}_m^l\}$, if X is strictly convex. This follows from the translation theorem[48,49]

for arbitrary $\mathcal{R}_m^l$ being exact, truncated at the same order l, and converging everywhere in $\mathbb{R}^3$. For larger molecules, however, we expect the cavities X_R to deviate further from convex shapes. It is then not generally possible to know beforehand if the expansion will converge fast, or at all. It is worth noting here that except for the simplest systems, cavities are never perfectly convex in practice, yet most systems do actually converge.

Theoretical considerations about whether or not arbitrary harmonic potentials can be expanded in $\{\mathcal{R}_m^l\}$ in any given cavity X_R are of limited use. On one hand, even if we know that the infinite series converges, we do not generally know if it converges to desired accuracy at a practically applicable expansion order. On the other hand, even if harmonic functions exist for which the expansion does not converge, for Φ_R in particular it may still converge.
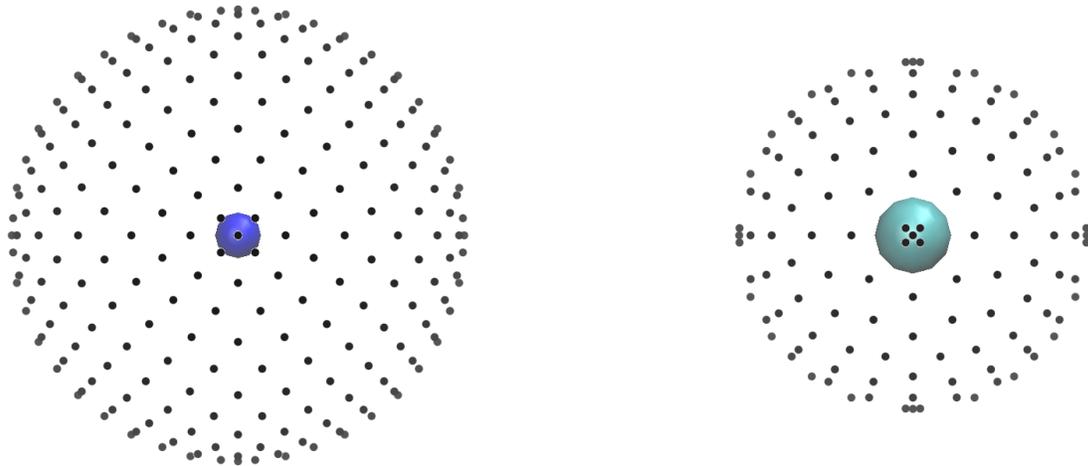
Figure 4: Na and Cl atom at $10\,\text{Å}$ separation with discretized isodensity cavity surfaces at $\rho_{\text{iso}} = 0.01\,\frac{e}{\text{Å}^3}$. Despite the two cavities being completely disjoint, one single expansion in regular solid harmonics around the center between the two is sufficient to solve the electro-static problem. Note that by 'one single expansion' we mean not only that the same basis is used in both cavities, but also the expansion coefficients are the same.

The latter aspect can be illustrated by the following example: We consider one Na and one Cl atom at $10\,\text{Å}$ separation, as shown in fig. 4. We run a DFT calculation in solution

with $\varepsilon = 78.36$ and $\rho_{\mathrm{iso}} = 0.01\,\frac{e}{\mathrm{\AA}^3}$. Based on the above considerations, we would expect MPE-1c to fail for this system. The cavities of the two atoms are completely disjoint and there is no straightforward reason to assume that Φ_{R} in both cavities could be described by the same multipole expansion with the same expansion coefficients. In practice, however, at $l_{\mathrm{max}}^{\mathrm{R}} = 8$, we get a $\bar{R}^2 = 0.9957$, which suggests that the electrostatic boundary conditions are almost exactly fulfilled. When we use separate multipole expansions for Φ_{R} in both cavities (a specific case of the general method described in the following sections), we, of course, obtain a $\bar{R}^2 = 1$, but the error in the electrostatic contribution to the solvation free energy made by the single expansion in comparison to the separate expansion is only $12\,\mathrm{meV}$. Using two different multipole expansions effectively doubles the number of basis functions for Φ_{R}, from 81 to 162 (although at each point $\mathbf{r}$ only 81 are used, they still amount to 162 degrees of freedom in eq. (14)). Going back to the single expansion, we can use $l_{\mathrm{max}}^{\mathrm{R}} = 12$ to reach 169 basis functions. With this, we then get $\bar{R}^2 = 0.9982$ and the error with respect to the two-center expansion is reduced to a mere $1.8\,\mathrm{meV}$. By simply increasing the number of basis functions to approximately the same amount, we have thus achieved almost the same improvement as by using separate multipole expansions in both cavities.

This raises the question if the error for larger solutes in MPE-1c might simply be due to the basis set being too small. Indeed, already in the original publications[19,20] it was noted that for larger solutes, higher expansion orders than the originally published $l_{\mathrm{max}}^{\mathrm{R}} = 8$ may be necessary. Specifically, some number of basis functions per non-H atom might exist at which $\Delta G_{\mathrm{solv}}^{\mathrm{elstat}}$ is reliably converged and $\bar{R}^2 \approx 1$. We test this hypothesis by picking all test systems (58 in total) with $\bar{R}^2(l_{\mathrm{max}}^{\mathrm{R}} = 20) \leq 0.99$ and running additional calculations at $l_{\mathrm{max}}^{\mathrm{R}} = 40, 60$ and 90. The data from fig. 1 together with these additional calculations, plotted against the number of basis functions per non-H atom, is shown in fig. 5.

While a remarkable portion of test systems converges somewhere around 100 basis functions per non-H atom, there are some systems which do not converge to reasonable accuracy even within the enlarged range of $l_{\mathrm{max}}^{\mathrm{R}}$. It is worth looking at the systems for which the model
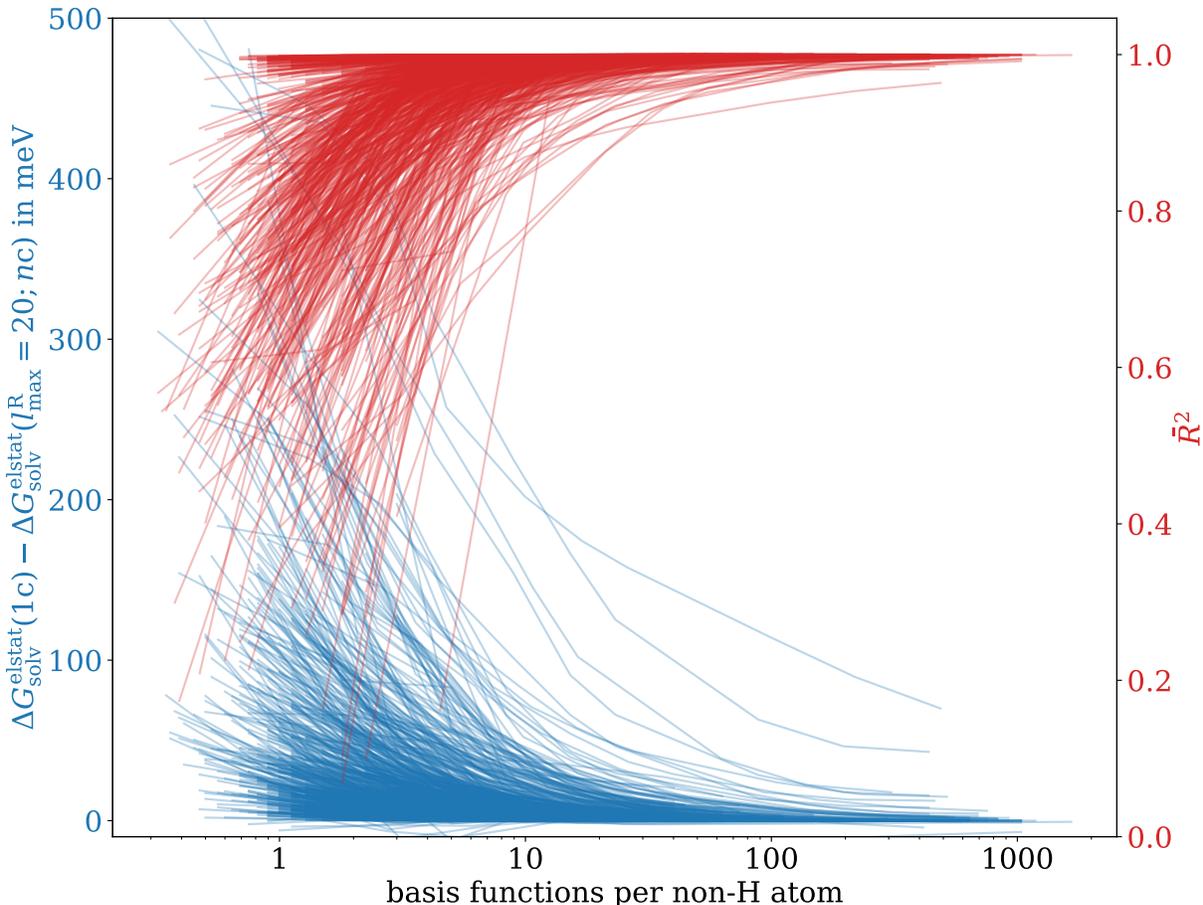
Figure 5: Same as fig. 1, but now showing the convergence with respect to the number of basis functions per non-H atom (on a logarithmic scale) for Φ_R in MPE-1c.

performs worst in some detail. We identify the systems for which at least one calculation was successfully conducted which fulfills the following two conditions: First, the number of basis functions for Φ_R per non-H atom is at least 81, and second, $\bar{R}^2 < 0.98$ and/or the error in $\Delta G_{\text{solv}}^{\text{elstat}}$ is $> 20\,\text{meV}$. We will show in section 6.1 that these conditions do not apply for any of our test systems studied with the new model described in the following sections.

The six systems identified in this way are shown in fig. 6 and their parameters and results listed in table 1. For all these systems, X_R deviates far from a convex shape, confirming our explanation from the beginning of this section at least in the sense that when the original method[20] fails, it is linked to cavities being strongly non-convex. Yet, the inverse conclusion, that non-convex cavities necessarily lead to failure, is not true.

From this data alone it is not entirely clear if the non-convergence is due to incompleteness of the basis for Φ_R alone. As described at the beginning of section 4.1, we compare MPE-1c with $l_{max}^Q = 6$ to a reference with $l_{max}^Q = 8$. The objective was to assess the performance of MPE-1c as originally published in ref. 20, where a value of 6 was proposed for use with *tight* integration grids in `FHI-aims`.[20,25–29] We note that the lower expansion order does not automatically mean a smaller basis for Φ_Q, because MPE-1c places expansion centers on all solute nuclei, whereas our reference places them only on non-H nuclei, cf. section 5.1. Nonetheless, we verify that the basis for Φ_Q does not cause the observed errors by rerunning the six systems in table 1 with $l_{max}^Q = 8$, ensuring a larger (or equally large in the case of octafluorocyclobutane) basis for Φ_Q than in the reference. While minor improvements to both $\bar{R}^2$ and G_{el} are observed, the larger basis for Φ_Q clearly does not solve the issue, confirming that it is, indeed, tied to an incomplete basis for Φ_R.

Table 1: Systems with $\bar{R}^2 < 0.98$ and/or electrostatic solvation free energy error $\Delta\Delta G_{solv}^{elstat} > 20\,\mathrm{meV}$ in MPE-1c at number $n_{bf,rel} \geq 81$ of basis functions for Φ_R per non-H atom. N is the number of non-H atoms. Solutes are identified by their entry number in the Minnesota Solvation Database.[12] For each system, of all calculations fulfilling these conditions, only the one with the lowest l_{max}^R is shown.

solute	N	$\frac{\rho_{iso}}{1\,\mathrm{meÅ}^{-3}}$	ε	$n_{bf,rel}$	l_{max}^Q	$\bar{R}^2$	$\frac{\Delta\Delta G_{solv}^{elstat}}{1\,\mathrm{meV}}$
0445pho	16	59.8	6.51	105.1	6	0.981	23.9
					8	0.982	22.3
0925dec	17	82.4	3.19	98.9	6	0.939	114.4
					8	0.943	110.3
test1010	19	94.7	4.12	88.5	6	0.983	28.4
					8	0.984	26.4
test1020	27	30.1	5.18	137.8	6	0.986	22.3
					8	0.988	20.3
test1031	19	64.4	346.30	88.5	6	0.975	63.0
					8	0.977	58.0
test2023	12	49.5	6.61	140.1	6	0.975	5.2
					8	0.979	3.7

Overall, six systems with unsatisfactory solutions to the electrostatic problem may not sound like a lot. All of these system do, however, lie in the sparsely sampled upper end of our

test set regarding molecular size (cf. fig. 3). Furthermore, their ρ_{iso} all lie in the upper end of the sampled range, where we generally expect more complicated cavity shapes. Looking only at systems with 16 or more non-H atoms and $\rho_{\text{iso}} \geq 30\,\text{meÅ}^{-3}$, we find that of only 12 such systems in our test set, 5 failed in the above described sense. Thus, even after weighting the basis set size by the number of non-H atoms, larger molecules in combination with high isodensities still show a much worse convergence (or possibly none at all) with expansion order than the smaller molecules or systems with lower isodensities.

One may argue that these errors are due to small voids in the solutes. Such voids can be assigned to the dielectric medium X_{Q} following only the simplistic cavity definition in eqs. (7a) and (7b). If they are smaller than a solvent molecule, the resulting model is obviously unphysical. Indeed, such voids are observed in some of the problematic cases described above, probably most evidently in decamethyltetrasiloxane (entry 0925dec in the Minnesota Solvation Database,[12] top right in fig. 6). There exist solutions to this issue, e.g. using so-called solvent-aware interfaces.[32] Implementing such a method is, however, beyond the scope of the present work.

For the time being, we state that although the convergence with expansion order may improve with a more physically motivated cavity definition, the issues presented in this section still persist. First, we observe the issue also in some systems without obviously un-physical cavity shapes, particularly chlorimuron-ethyl (test1020 in the Minnesota Solvation Database[12]) and octafluorocyclobutane (test2023 in the Minnesota Solvation Database[12]) (center right and bottom right in fig. 6). Second, our model is agnostic of the solvent's molecular size and ideally the solution ansatz should work independent of cavity choice. On a similar notion, when going to larger solutes than our test set, e.g. biomolecules or catalytic surfaces, the issue of non-convex cavities will inevitably emerge again.

We conclude that for small molecules, MPE-1c exhibits remarkable robustness once $l_{\text{max}}^{\text{R}}$ is chosen based on the number of basis functions per non-H atom. For a significant portion of the larger solutes, however, the basis expansion converges slowly or not at all. In the

22
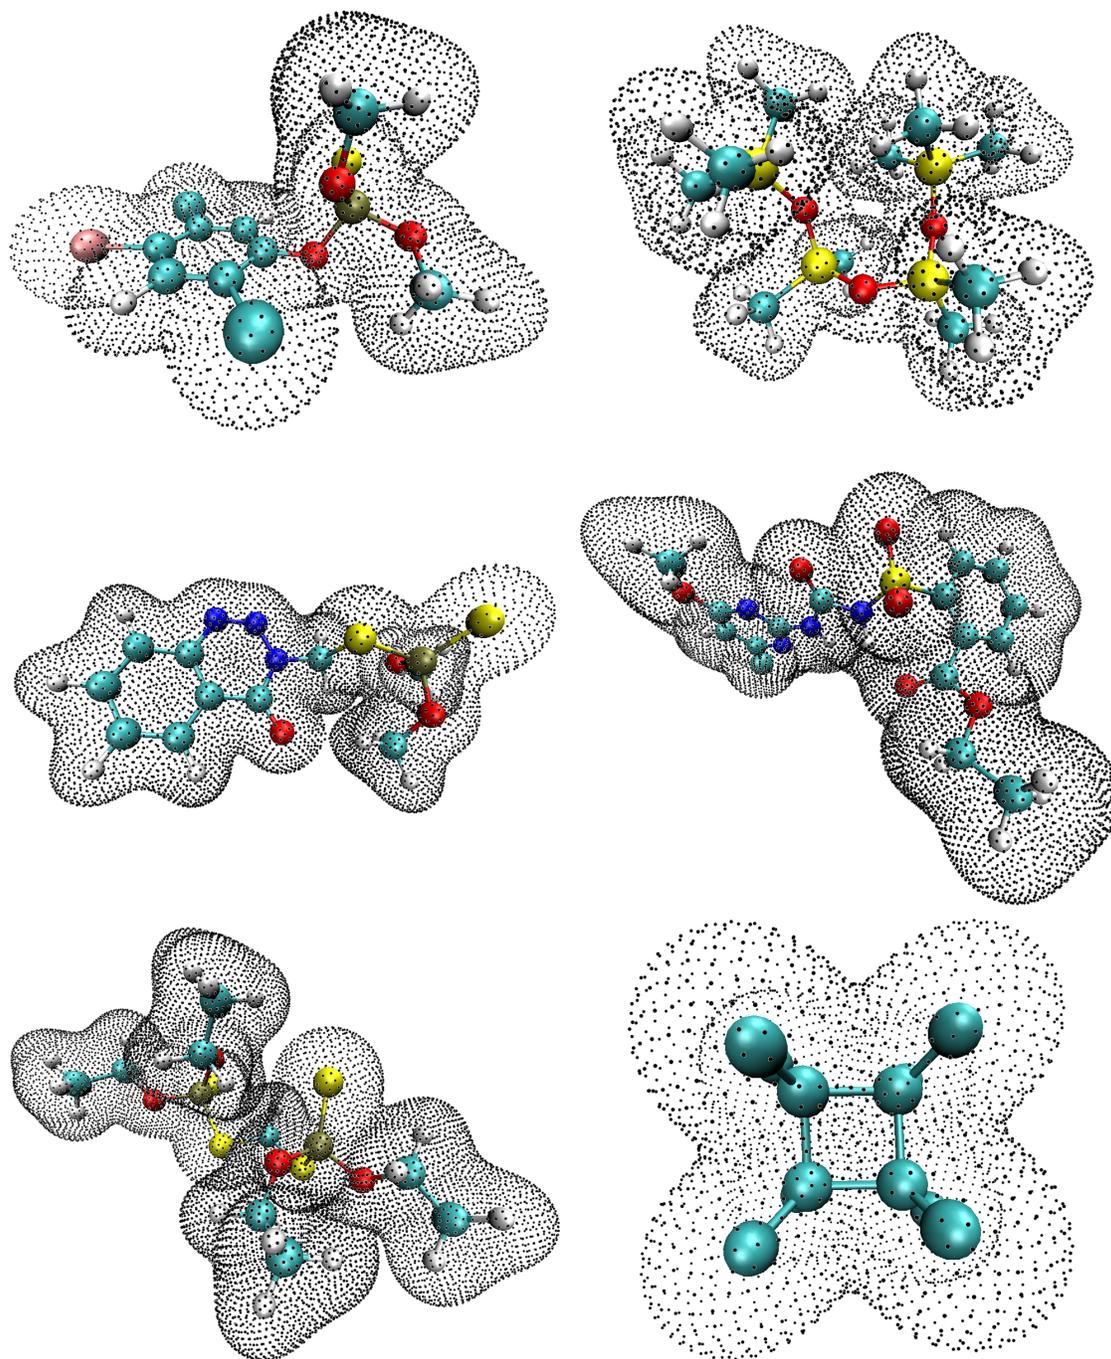
Figure 6: Systems described in table 1, in the same order (row first). Cavity surface sampling density is arbitrary and does not represent the sampling density in the calculations shown in the table.

following sections we therefore develop a modified MPE model which solves this issue, aiming at reliable convergence with expansion order.

5 Subcavity method

As a remedy for the issues discussed above, we modify the MPE method,[20] putting particular emphasis on the solution of the electrostatic problem in arbitrarily shaped cavities. The main idea of our method, which we call MPE-nc, is to formally not solve the problem for one single cavity X_R, but for multiple small 'subcavities' $X_{R,K}$. A minimal example for the user input required to use our method in `FHI-aims`[20,25–29] is provided in the SI.

5.1 Separation into subcavities

We separate X_R into subcavities $X_{R,K}$ around multiple centers $\mathbf{r}_K$.

$$X_{R,K} = \{\mathbf{r} : \mathbf{r} \in X_R, \|\mathbf{r} - \mathbf{r}_K\| < \|\mathbf{r} - \mathbf{r}_{K'}\| \, \forall K' \neq K\} \quad . \tag{19}$$

In principle, fast convergence of a series expansion in $\{\mathcal{R}_m^l\}_K$ in these subcavities is to be expected for any choice of $\mathbf{r}_K$, as long as the cavity can approximately be described as a superposition of spheres around these centers. The straightforward choice is to use the positions of all non-H nuclei of the solute. We use this set of $\mathbf{r}_K$ throughout the rest of this work. An example of such a partition is shown in fig. 7. The basic idea of decomposing an implicit solvent cavity into atom-centered domains and solving the individual electrostatic problems in a mutually consistent way has previously been applied to COSMO.[22–24] Here, we discuss the application of this idea to the MPE model.

Φ_{MPE} is defined piece-wise in each of these subcavities

$$\Phi_{MPE}(\mathbf{r} \in X_{R,K}) = \Phi_{R,K}(\mathbf{r}) \tag{20}$$

The reaction field $\Phi_{R,K}$ in the respective subcavities can be expanded according to eq. (11),[20]

$$\Phi_{R,K}(\mathbf{r}) = \sum_{l=0}^{l_{max}^R} \sum_{m=-l}^{l} R_K^{(l,m)} \mathcal{R}_m^l(\mathbf{r} - \mathbf{r}_K) \quad . \tag{21}$$
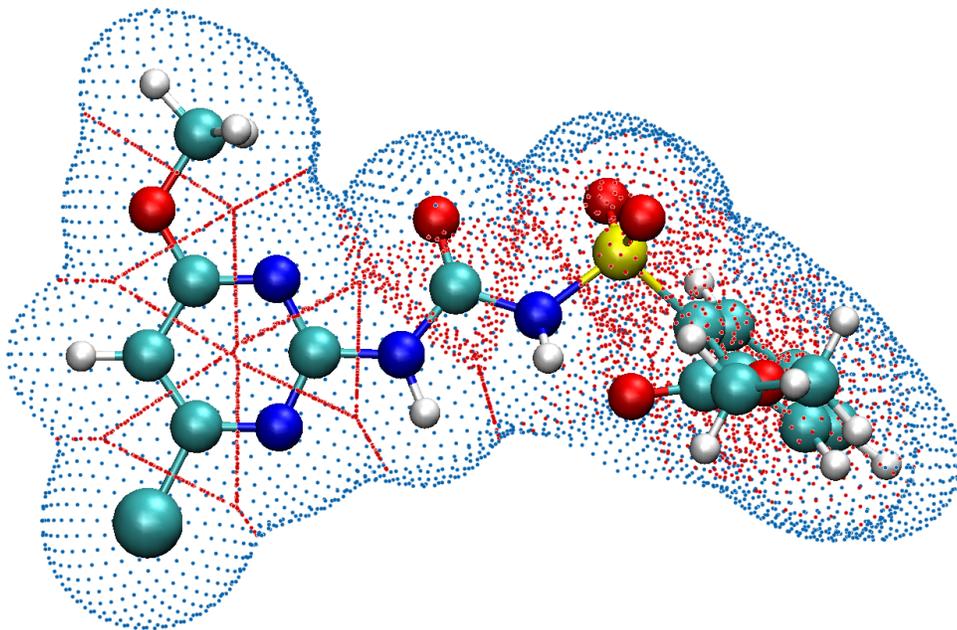
Figure 7: Orthographic view of chlorimuron-ethyl (entry test1020 in Minnesota Solvation Database,[12] cf. fig. 6, center right) with discretized cavity surface (blue) and subcavity-subcavity boundaries (red). Interfaces cut along the plane of the heteroaromatic ring.

Similarly, the external potential Φ_Q can be expanded according to eq. (12),[20]

$$\Phi_Q(\mathbf{r}) = \sum_{J=1}^{J_{\max}} \sum_{l=0}^{l_{\max}^Q} \sum_{m=-l}^{l} Q_J^{(l,m)} \mathcal{I}_m^l(\mathbf{r} - \mathbf{r}_J) \quad . \tag{22}$$

While the choice of expansion centers is, once again, arbitrary within certain limits, we use the same set of expansion centers $\{\mathbf{r}_J\} = \{\mathbf{r}_K\}$ as for the reaction field in the subcavities, i.e. the non-H nuclear positions of the solute. This is different from the original version[20] of the method, which placed expansion centers for Φ_Q also on hydrogen cores. We show in section 6.1 that Φ_Q still converges using this smaller basis set, saving computational resources. Similarly, it is in principle possible to use different expansion orders for different centers. Since all our centers are, however, qualitatively similar, there is no need to do this. Therefore, we use only two different expansion orders - one for all $\Phi_{R,K}$ and one for Φ_Q.

In complete analogy to MPE-1c, the boundaries B_{ij} are discretized into a finite set of points, aiming at some target degree of determination $d_{\det}$ of the overall linear system

eq. (14). Additionally, the algorithm tries to sample all boundaries with approximately the same density of points. The exact discretization procedure for the subcavity-subcavity boundaries $B_{\mathrm{R}K,\mathrm{R}K'}$, and a modification to the discretization of the subcavity surfaces $B_{\mathrm{Q},\mathrm{R}K}$ with respect to the original method[20] are described in the SI. There, we also describe how the uniform coordinate scaling used to improve matrix conditioning[19] is adjusted to MPE-nc.

The matrix $\mathbf{A}$ in the central linear system eq. (14) becomes sparse with this piecewise multipole expansion, because the columns corresponding to the basis functions for one region X_i have non-zero elements only in the rows corresponding to points on boundaries B_{ij} which are in contact with X_i.[19] The exact structure of $\mathbf{A}$, as well as an implementation of the iterative LSQR[50] solver that exploits this sparsity, are elaborated in the SI. Alternatively, one can use one of the direct solution algorithms reported in ref. 20. First, QR factorization of $\mathbf{A}$ is performed. For well-conditioned systems, eq. (14) can afterwards be solved straightforwardly. We call this procedure the 'QR' solver henceforth. If the resulting matrix $\mathbf{R}$ is rank deficient, however, it can be further factorized in a singular value decomposition (SVD), and a regularized solution can be obtained by applying a cutoff in the inversion of the singular values. We call the latter option the 'QR+SVD' solver henceforth. Note that the QR and QR+SVD solvers do not exploit sparsity. The implications of this issue are discussed in section 6.3.

5.2 Charge conservation

We take the opportunity of generalizing MPE implicit solvation to address another, minor issue which is not directly related to those described in sections 4.1 and 4.2.

From Gauss's law it follows that

$$\oint_{\mathbf{A}\subset X_{\mathrm{Q}}} \nabla\Phi_{\mathrm{Q}}\, d\mathbf{A} = 0 \tag{23}$$

for arbitrary closed surfaces $\mathbf{A}$ lying entirely in X_{Q}. In principle, this should be fulfilled

automatically, if eq. (13b) is fulfilled exactly on all boundaries B_{ij}. Our method is, however, approximate by nature and boundary conditions are fulfilled only in a least-squares sense. Usually, this does not lead to any severe or systematic errors. In the case of eq. (23), however, the error formally amounts to a charge, altering the effective charge $\oint_{\mathbf{A} \subset X_{\mathrm{Q}}} \mathbf{D} \, d\mathbf{A}$ of the system in solution. We note that this is unrelated to the so-called outlying charge error concerning explicit electron density, which was already discussed and resolved in refs. 19,20. The error discussed here concerns solely the implicit potential. The existence of a spurious charge alone may, in some cases, be enough to introduce a systematic error in the electrostatic energy. It is therefore worth enforcing eq. (23) not in a least-squares fashion, but exactly (up to numeric precision). The technical details on the constrained solution algorithm are explained in the SI.

6   Results and discussion

With the details of the method now in place we turn to the performance of MPE-nc and provide numeric parameters for practical applications.

6.1   Convergence with numeric parameters

We test the convergence of $\Delta G_{\mathrm{solv}}^{\mathrm{elstat}}(n\mathrm{c})$ and $\bar{R}^2$ with expansion orders $l_{\mathrm{max}}^{\mathrm{R/Q}}$ and target degree of determination d_{det}. For each of the three parameters, values are sampled on respective grids $l_{\mathrm{max}}^{\mathrm{R}} = 2, 4, 6, 8, 12, 16, 20$, $l_{\mathrm{max}}^{\mathrm{Q}} = 0, 2, 4, 6, 8, 12, 16$ and $d_{\mathrm{det}} = 3, 4, 5, 6, 8, 10$. We vary each of the three convergence parameters separately and set the respective other two to their '*really tight*' defaults, as discussed below. These are $l_{\mathrm{max}}^{\mathrm{R}} = 12$, $l_{\mathrm{max}}^{\mathrm{Q}} = 8$ and $d_{\mathrm{det}} = 5$. These values were determined from preparatory runs, which were then confirmed in the here reported results. We therefore do not give an extra account of these initial runs, but rather focus on the full results summarized in figs. 8 to 10. In order to illustrate the convergence with these values, we use $\Delta G_{\mathrm{solv}}^{\mathrm{elstat}}(n\mathrm{c})$ at $l_{\mathrm{max}}^{\mathrm{R}} = 20$ for the $l_{\mathrm{max}}^{\mathrm{R}}$ convergence study, at $l_{\mathrm{max}}^{\mathrm{Q}} = 16$

for the $l_{\max}^Q$ convergence study and at $d_{\det} = 10$ for the $d_{\det}$ convergence study as a converged reference, respectively.
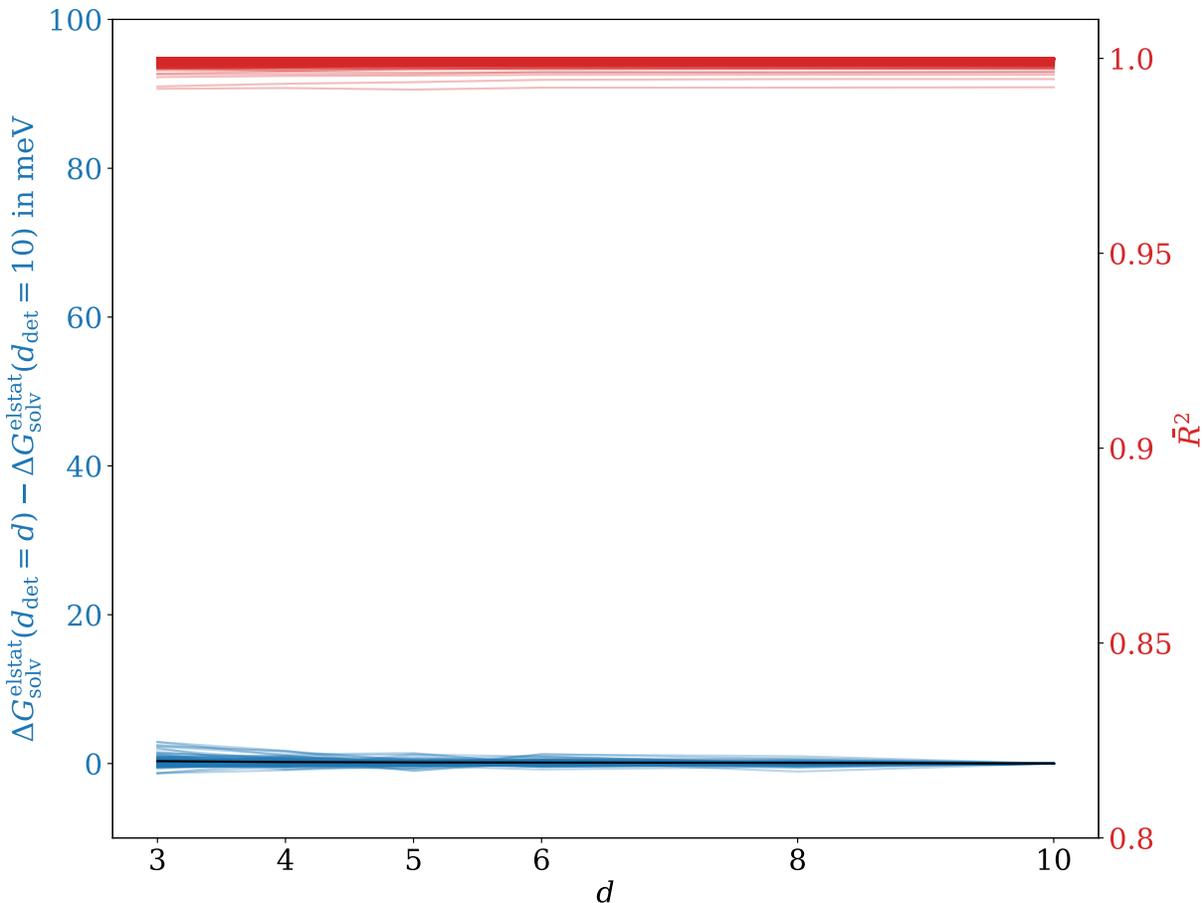
Figure 8: Convergence of the electrostatic contribution to the solvation free energy $G_{\text{solv}}^{\text{elstat}}$ (blue) and adjusted coefficient of determination $\bar{R}^2$ (red) of the discretized boundary conditions eqs. (13a) and (13b) with target degree of determination $d_{\det}$ of the SLE.

As depicted in fig. 8, the target degree of determination $d_{\det}$ appears to have no noticeable influence. The RMSE with respect to the reference value is consistently smaller than $1\,\text{meV}$ and shows little dependence on $d_{\det}$. Furthermore, no significant outliers are observed. Nonetheless, to avoid parts of the SLE becoming underdetermined, $d_{\det}$ should not be chosen too small.[19] Indeed, the ratio of non-zero rows to columns for individual potentials Φ_i may be smaller than $2d_{\det}$, as regions X_i have different numbers of basis functions and discretized points in adjacent interfaces B_{ij}. Choosing a safe, large enough $d_{\det}$ does not

increase the computational cost too dramatically, as it enters the matrix size only linearly. All in all, d_{det} can be considered a fairly uncritical parameter and can be chosen higher or lower where needed. For the rest of the present work, we use the previously established [20] value of $d_{\text{det}} = 5$, which yielded an RMSE of $0.15\,\text{meV}$ with respect to the reference value in our convergence study.
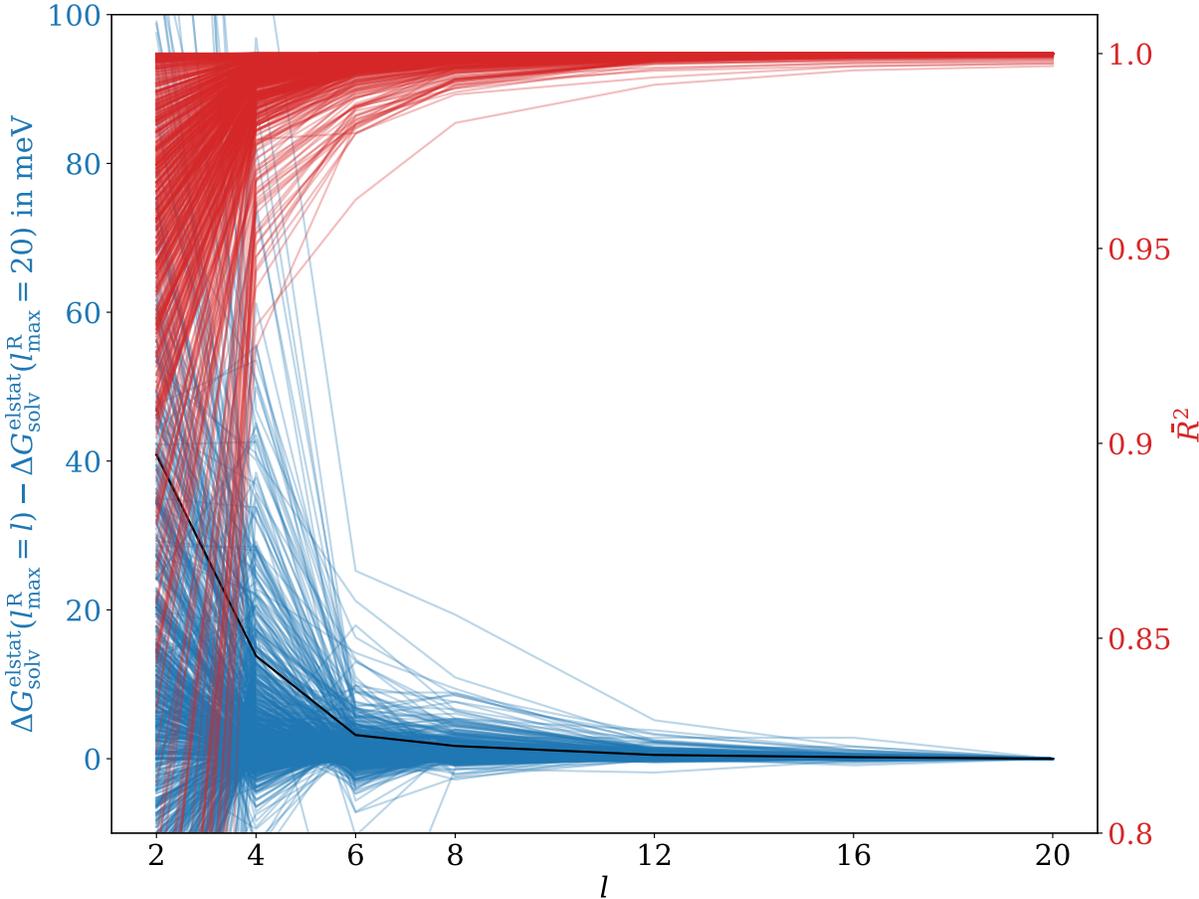
Figure 9: Same as fig. 8, but showing convergence with expansion order $l^{\text{R}}_{\text{max}}$ of the MPE potential inside the subcavities.

A more significant convergence behavior of $\Delta G^{\text{elstat}}_{\text{solv}}$ is instead obtained with increasing expansion orders $l^{\text{R/Q}}_{\text{max}}$ as depicted in figs. 9 and 10. In both cases, no noticeable changes in $\Delta G^{\text{elstat}}_{\text{solv}}$ occur at the upper end of the range of sampled values, indicating that the solution to the electrostatic problem is then, indeed, converged. This is further supported by $\bar{R}^2$ of the discretized boundary conditions eqs. (13a) and (13b) converging to 1. MPE-nc shows
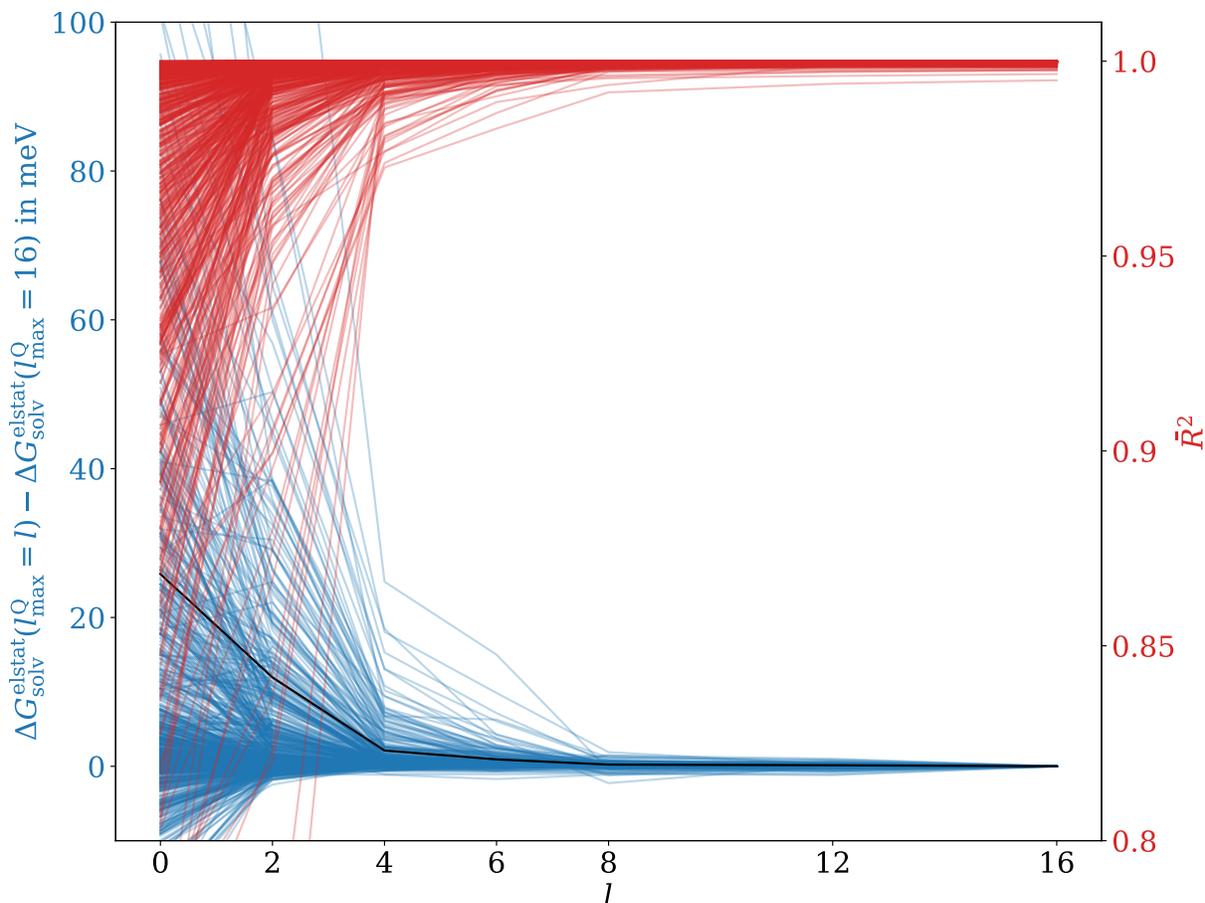
Figure 10: Same as fig. 8, but showing convergence with expansion order l_{max}^{Q} of the MPE potential outside the cavity.

convergence for all of the test cases, even those that failed to converge with MPE-1c. To illustrate the advantage of MPE-nc over MPE-1c we plot in fig. 11 the convergence with overall basis set size of Φ_{R}. Contrary to MPE-1c, the new subcavity method starts to converge for all geometries at around 50 basis functions per non-H atom.

Similarly to the single cavity case, solvation free energies at too low expansion orders are typically higher than the reference values, indicating a systematic underestimation of the electrostatic solute-solvent interaction at these expansion orders. The error is, however, not strictly variational as overestimation can be observed occasionally.

Assuming that the errors introduced by the individual parameters are uncorrelated, we
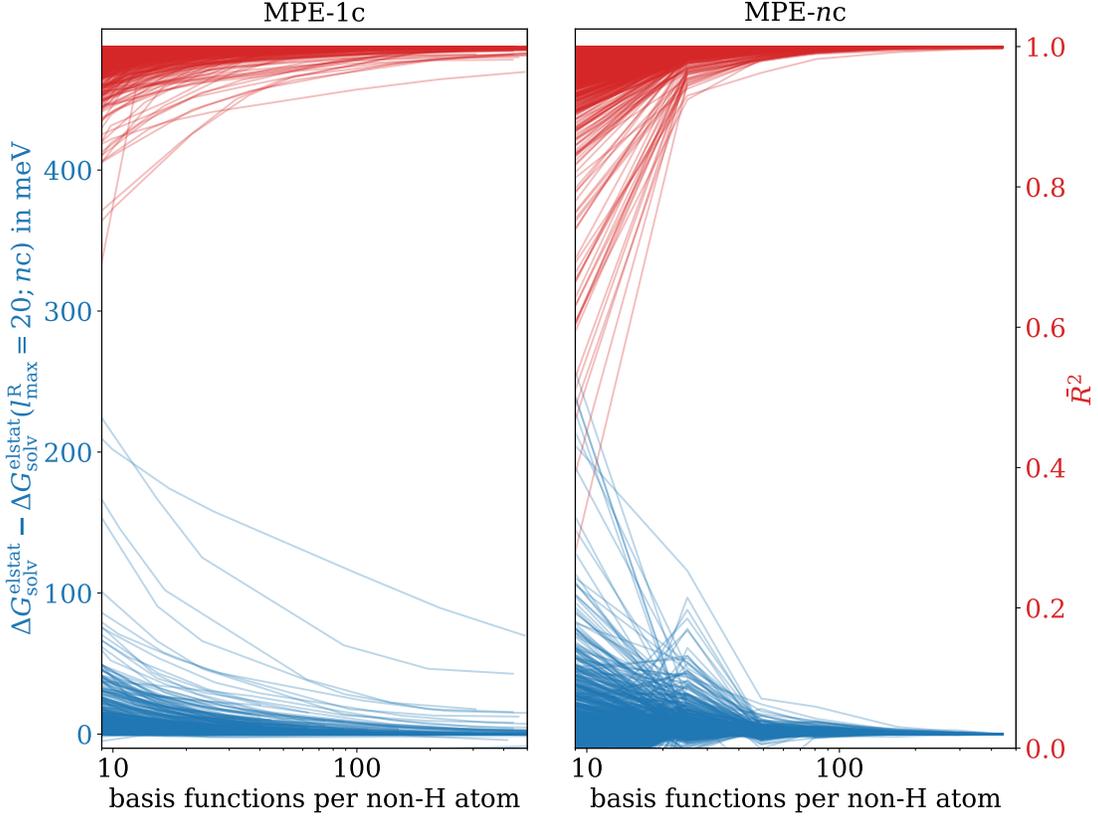
Figure 11: Convergence of the electrostatic contribution to the solvation free energy $\Delta G_{\mathrm{solv}}^{\mathrm{elstat}}$ (blue) and adjusted coefficient of determination $\bar{R}^2$ (red) of the discretized boundary conditions eqs. (13a) and (13b) with number of basis functions per non-H atom (on a logarithmic scale) for Φ_{R} in the original (left) and our modified method (right).

can estimate the total root mean-square error as

$$
\begin{aligned}
\mathrm{RMSE}_{\mathrm{est}}(l_{\mathrm{max}}^{\mathrm{R}}, l_{\mathrm{max}}^{\mathrm{Q}}, d_{\mathrm{det}}) = \\
\sqrt{\mathrm{RMSE}(l_{\mathrm{max}}^{\mathrm{R}})^2 + \mathrm{RMSE}(l_{\mathrm{max}}^{\mathrm{Q}})^2 + \mathrm{RMSE}(d_{\mathrm{det}})^2}
\end{aligned} \quad , \tag{24}
$$

where $\mathrm{RMSE}(x)$ is the root mean-square error of $\Delta G_{\mathrm{solv}}^{\mathrm{elstat}}$ at the respective value of x with respect to $\Delta G_{\mathrm{solv}}^{\mathrm{elstat}}$ at the largest tested value of x. Given that the error with respect to d_{det} stays mostly constant we omit if from further consideration. With respect to the other two error components, though, we can use $\mathrm{RMSE}_{\mathrm{est}}$ to estimate the expected errors for given combinations of $l_{\mathrm{max}}^{\mathrm{R}}, l_{\mathrm{max}}^{\mathrm{Q}}$. Given that the computational cost of the method rises steeply with increased MPE expansion order it useful to define default values for different levels of

expected accuracy. In analogy to the electronic basis sets of the `FHI-aims`[20,25–29] code itself we therefore define three convergence levels

- *'really tight'*: $l_{max}^R = 12$, $l_{max}^Q = 8$

  solver: QR+SVD

  $RMSE_{est} = 0.58\,meV$

  $RMSE(l_{max}^R) = 0.51\,meV$, $RMSE(l_{max}^Q) = 0.23\,meV$

- *'tight'*: $l_{max}^R = 8$, $l_{max}^Q = 6$

  solver: any

  $RMSE_{est} = 2.0\,meV$

  $RMSE(l_{max}^R) = 1.7\,meV$, $RMSE(l_{max}^Q) = 0.93\,meV$

- *'light'*: $l_{max}^R = 6$, $l_{max}^Q = 4$

  solver: any

  $RMSE_{est} = 3.8\,meV$

  $RMSE(l_{max}^R) = 3.2\,meV$, $RMSE(l_{max}^Q) = 2.1\,meV$

As further explained in the SI, for *really tight* defaults, the error introduced by non-regularized solvers for ill-conditioned matrices can reach the order of some $10^{-5}\,eV$. Aiming for some $10^{-4}\,eV$ accuracy, the regularized QR+SVD solver is then the safer choice. For the smaller basis sets, the error introduced by the solver is smaller and the error from the basis set is larger, so any solver can be used in that case.

For neutral solutes, the solvation free energies ΔG_{solv} typically lie in the order of some tens to some hundreds of meV.[21] Tolerating even higher errors than for the suggested *light* settings would thus introduce systematic errors in the order of magnitude of the quantity we want to determine and is therefore not advisable. Furthermore, the number and maximum value of outliers in figs. 9 and 10 increase dramatically beyond this point. This was also taken into consideration when choosing $l_{max}^Q = 6$ for the *tight* settings. Although the RMSE

for $l^{\mathrm{Q}}_{\max} = 4$ is closer to that for $l^{\mathrm{R}}_{\max} = 8$, at $l^{\mathrm{Q}}_{\max} = 6$ the number of systems with errors $> 10\,\mathrm{meV}$ is significantly smaller, hinting at higher reliability.

6.2   Model parameterization

Exactly analogously to the original MPE-1c case, the physical model parameters ρ_{iso} and α of MPE-nc need to be fitted to experimental reference data. Given that in the original MPE-1c method errors in the electrostatic contribution to the solvation free energy were likely compensated by non-electrostatic contributions, we cannot re-use the parameters computed for that model. Instead, we re-fit ρ_{iso} and α following the procedure outlined in earlier work.[20] We apply the process mostly unchanged, simplifying a few technical details and sampling ρ_{iso} on a logarithmic grid. DFT calculations for all test molecules are correspondingly performed both in vacuum and in implicit solvent at isodensities $\rho_{\mathrm{iso}} = 10^x \frac{e}{\text{Å}^3}$ with x sampled on a grid from -3 to -1 in steps of 0.25. The error function

$$P(x,\alpha) = \sum_{M}^{\text{molecules}} \left(\Delta G^{\text{exp}}_{\text{solv},M} - \left(\Delta G^{\text{elstat}}_{\text{solv},M}(x) + \alpha A_M(x) \right) \right)^2 \tag{25}$$

is minimized, where $\Delta G^{\text{elstat}}_{\text{solv},M}(x)$ and $A_M(x)$ are cubic splines between the values of x for which actual DFT calculations with $\rho_{\mathrm{iso}}(x)$ were performed.[51]

All parameterization DFT calculations employed the *really tight* settings of the MPE-nc solvation model. All other settings, such as the electronic basis sets and integration grids were set to the *tight* defaults of `FHI-aims`.[20,25–29] First, parameterization with the PBE[38] functional was performed. The optimized isodensity value was found to lie in the upper part of the sampling grid. For subsequent parameterizations we thus only sampled the range from -2 to -1. In addition to PBE, parameters were optimized using the revPBE,[52] RPBE,[53] HSE06,[39,40] BLYP,[54,55] B3LYP[56] and SCAN[57] functionals. As a consistency check, we also conducted a parameterization with PBE using MPE-1c with the original parameters $l^{\mathrm{R}}_{\max} = 8$ and $l^{\mathrm{R}}_{\max} = 6$.[20] The parameterization results are shown in table 2.

Table 2: Optimized parameters and RMSE of ΔG_{solv} with respect to experimental references for water, using different solvent models and xc functionals. For the MPE-nc solvent model, the *really tight* expansion orders from the previous section were used throughout. For MPE-1c, the originally published values[20] of $l_{\max}^{\mathrm{R}} = 8$ and $l_{\max}^{\mathrm{R}} = 6$ were used. In the rightmost column, RMSEs calculated with the respective functional, but with α and ρ_{iso} optimized for HSE06 (cf. highlighted values) are shown.

xc	solvent	$\dfrac{\rho_{\mathrm{iso}}}{1\,\mathrm{me\mathring{A}^{-3}}}$	$\dfrac{\alpha}{1\,\mathrm{meV\mathring{A}^{-2}}}$	$\dfrac{\mathrm{RMSE}}{1\,\mathrm{meV}}$	$\dfrac{\mathrm{RMSE}}{1\,\mathrm{meV}}$ (p. HSE)
HSE06	MPE-nc	33.68	1.805	120	120
PBE	MPE-nc	33.75	1.609	129	132
revPBE	MPE-nc	33.35	1.557	129	134
RPBE	MPE-nc	33.11	1.527	133	138
B3LYP	MPE-nc	33.02	1.710	127	127
BLYP	MPE-nc	32.74	1.513	137	141
SCAN	MPE-nc	35.48	1.835	121	123
PBE	MPE-1c	33.36	1.281	125	

The PBE family of xc functionals yields overall better results than the BLYP family, both for the generalized gradient approach (GGA) functionals (PBE, revPBE and RPBE vs. BLYP) and for the hybrid functionals (HSE06 vs. B3LYP). Within the PBE family of GGAs, the original PBE functional and revPBE yield a smaller RMSE than RPBE. The hybrid functionals generally yield an improvement of $\approx 10\,\mathrm{meV}$ compared to their GGA counterparts. The SCAN meta-GGA functional yields remarkably good agreement with experiment, almost matching the HSE06 hybrid functional and even outperforming B3LYP.

In fig. 12, we show calculated ΔG_{solv} of the solutes in the training set plotted against the experimental reference values for some selected functionals. It becomes clear that anionic solutes receive the largest improvement going from GGA to hybrid or meta-GGA functionals. A systematic underestimation (in terms of absolute values) of ΔG_{solv}, observed using GGA, is significantly improved by the hybrid and meta-GGA functionals, although it still persists to a lesser degree.

We also performed the parameterization using the alternative non-electrostatic model eq. (16) that employs also the cavity volume. No significant improvement of our model's agreement with experiment compared to the simpler model eq. (17) was obtained. For BLYP,
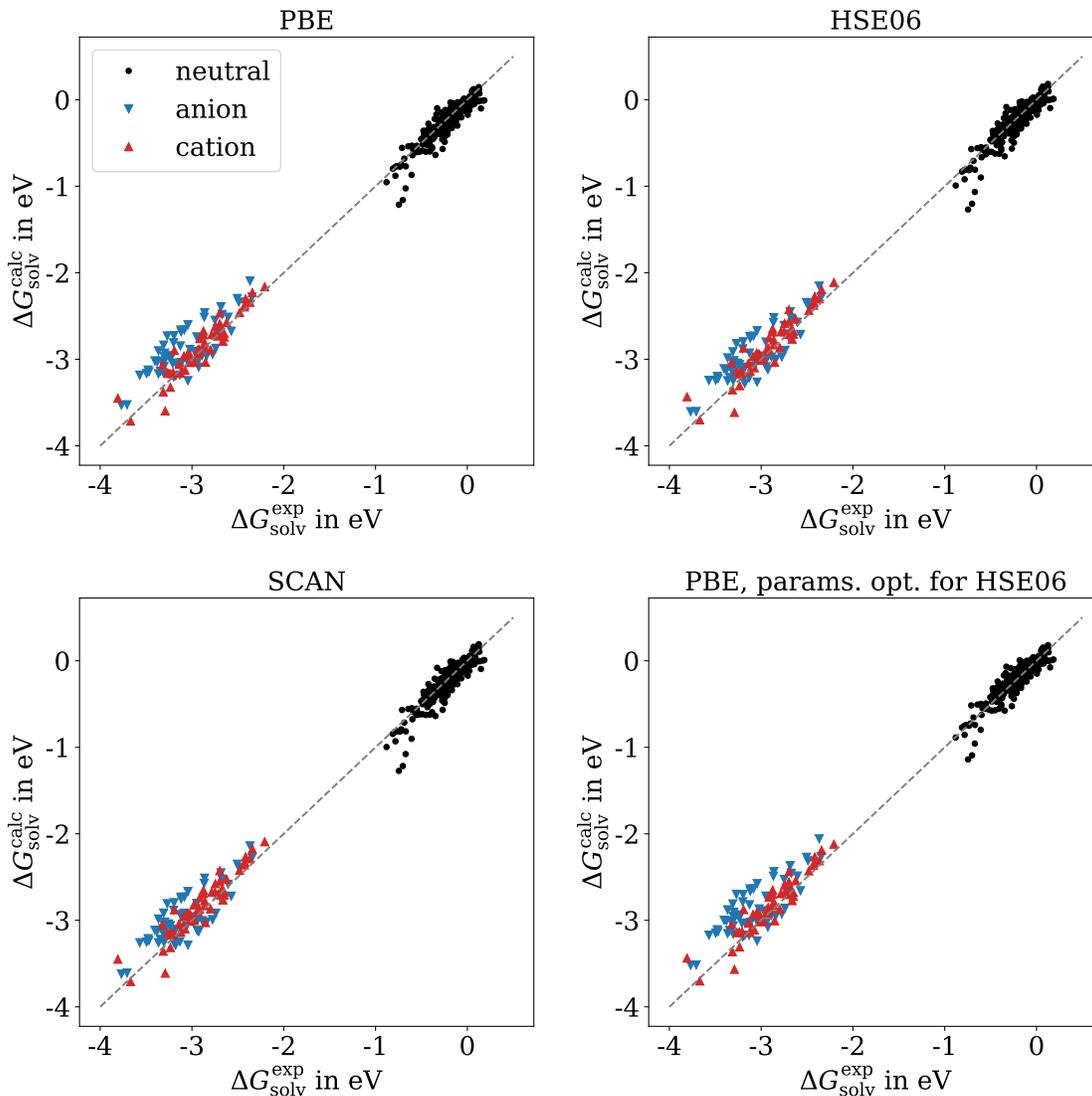
Figure 12: Calculated solvation free energies of solvation $\Delta G_{\text{solv}}^{\text{calc}}$ of our training set as calculated by MPE-*nc* plotted against experimental reference $\Delta G_{\text{solv}}^{\text{exp}}$, using 3 different xc functionals and with the optimized parameters from table 2. Grey dashed lines correspond to perfect agreement with experimental reference. The constituent terms of $\Delta G_{\text{solv}}^{\text{calc}}$, namely $\Delta G_{\text{solv}}^{\text{elstat}}$ and αA were calculated from cubic splines between the sampled values, see text. In the bottom right plot, the PBE functional was used together with parameters α and ρ_{iso} optimized for HSE06.

it would for instance reduce the RMSE of ΔG_{solv} by $2\,\text{meV}$, and for all other xc functionals, the improvement is even smaller. We do thus not recommend usage of this alternative model and refrain from reporting the respective parameters to avoid confusion.

The parameters found for MPE-1c are in very good agreement with those found in earlier

work[19] using a slightly smaller training set. Remarkably, the RMSE of ΔG_{solv} is approximately the same as with MPE-nc with the same xc functional, slightly smaller in fact. This reflects the aforediscussed combination of error cancellation and the smaller electrostatic error for the dominant smaller molecules in the test set. We have shown in section 4.1 that the error in $\Delta G_{\text{solv}}^{\text{elstat}}$ is correlated with molecular size. Our descriptor for the non-electrostatic free energy contributions, A, is also correlated with molecular size by construction. A size-dependent systematic overestimation (on a signed scale) of $\Delta G_{\text{solv}}^{\text{elstat}}$ can thus be compensated for by a size-dependent systematic underestimation of $\Delta G_{\text{solv}}^{\text{non-elstat}}$. This can be seen in the optimized parameter α being much smaller for MPE-1c than for MPE-nc, leading to such an underestimation.

Interestingly, for MPE-nc smaller RMSEs seem to be correlated with higher α, as shown in table 2. One possible interpretation is that some error cancellation of the above described kind still occurs with MPE-nc, presumably due to underestimation of the solute's polarization in the less exact functionals. Further research is needed to investigate this hypothesis.

We point out that the major part of the error cancellation in MPE-1c can be attributed to the choice of $l_{\max}^{\text{R}} = 8$ in our, as well as the original[19,20] parameterization of MPE-1c. We have seen in section 4.1 that this expansion order can be insufficient already for rather small solutes. In contrast, the number of solutes in the database[12] for which MPE-1c would not converge even at high expansion orders is small and makes only a minor contribution to the error cancellation in the original parameterization. However, this only concerns the parameterization process itself. In practical applications with large solutes, the problem of non-convergence in MPE-1c will emerge even when going to higher expansion orders.

This leads to the issue of transferability in MPE-1c. Consider a large solute. If we want to use parameters optimized at $l_{\max}^{\text{R}} = 8$, we do not know which $l_{\max}^{\text{R}}$ to use for the large solute. On one hand, the systematic error cancellation at $l_{\max}^{\text{R}} = 8$ results from the parameterization process only for relatively small molecules, which make up the major part of the training data. The error in $\Delta G_{\text{solv}}^{\text{non-elstat}}$ depends linearly on the cavity surface area A.

There is no reason to assume that the same holds strictly for $\Delta G_{\text{solv}}^{\text{elstat}}$. Therefore, the error cancellation is not guaranteed to work for large molecules at $l_{\text{max}}^{\text{R}} = 8$. On the other hand, let us assume that $\Delta G_{\text{solv}}^{\text{elstat}}$ does converge with $l_{\text{max}}^{\text{R}}$—we have shown in section 4.2 that this can in some cases work even for large solutes. Simply picking a converged expansion order does not solve the problem then either. If $\Delta G_{\text{solv}}^{\text{elstat}}$ is converged, then $\Delta G_{\text{solv}}^{\text{non-elstat}}$ compensates for a non-existent error. One would have to guess at which expansion order the errors match in absolute value. The error cancellation is thus not transferable.

MPE-nc does not suffer from these transferability problems. Some error cancellation may still occur, as described above. However, this is due to inaccuracies in the xc functional, for which the assumption of linear correlation with system size is much more reasonable. As a safe choice, one may use the parameters optimized for HSE06, of which α is fairly repulsive, and apply them also when using other functionals. The resulting $\Delta G_{\text{solv}}^{\text{elstat}}$ will then be an upper estimate—on a signed scale—but not a strict upper boundary. In fact, this does not significantly impact agreement with experiment, as shown in the rightmost column of table 2. These parameters can thus be considered general parameters, as similarly suggested in earlier work.[20] An exception is SCAN, for which the optimized α is actually larger than for HSE06. Here, the parameters optimized for SCAN itself should be at least equally as safe.

Despite these improvements, MPE-nc still calculates ΔG_{solv} only to $\approx 100\,\text{meV}$ accuracy. Possible sources of error are the simplistic model eq. (17) for non-electrostatic solute-solvent interactions, which accounts for dispersive and repulsive interactions only in a statistical manner, the assumption of solute geometries which are rigid and also identical in vacuum and solution, and finally the lack of hydrogen bonds in our model. Generally lower errors for non-aqueous solvents in a similar model[58] suggest that hydrogen bonds actually account for a considerable portion of the error, although the data sets used for these solvents in the cited study contained fewer solutes, which may also lead to lower errors.

6.3 Computational performance

We test the computational performance of our method directly in the context in which it is implemented, as a feature of `FHI-aims`[20,25–29] . To have a state-of-the-art reference, we recompiled our method in a more recent version of `FHI-aims`[20,25–29] than the one used for the calculations shown in previous sections, namely version `211010`. All tests reported in the following were conducted on an `AMD Ryzen Threadripper 3970X 32-Core Processor`. The `GNU Fortran (Gentoo 10.3.0 p1) 10.3.0` compiler with flags `-O3 -march=native -fallow-argument-mismatch -ffree-line-length-none` was used for compilation. Performance relevant libraries used are `AMD optimized ScaLAPACK 3.0`,[46] `AMD BLIS 3.0`,[59] `AMD optimized libFLAME 3.0`,[60] `ELPA 2020.05.001`[61,62] and `Open MPI v4.0.5`.[63] We find for molecules in the size range considered here, that no considerable computational speedup is gained when increasing the number of cores beyond 16, as shown below. Thus, the scaling of computational cost with molecular size is tested using only 8 cores, to avoid artifacts due to inefficient parallelization. 'Cores', in this section, refers to physical CPU cores.

In section 6.2, we used the *really tight* expansion orders from section 6.1 to obtain as exact parameters as possible. For most practical applications, however, the smaller *tight* expansion orders should be sufficient. These also allow for usage of the computationally more efficient QR or LSQR solvers, in contrast to the regularized QR+SVD solver necessary to deal with ill-conditioning at very high expansion orders. We test the scaling of computational time and memory requirement with solute size for the QR and LSQR solvers at *tight* expansion orders, using the same test systems as in our convergence studies, cf. section 3.

The scaling of computational time is compared to another commonly used implementation of an implicit solvation model, namely the SCCS model implemented in the Environ package,[21] version 2.0. Already MPE-1c found agreement with experiment comparable to SCCS, making the latter an appropriate reference for such a performance test. For both models, mean absolute errors (MAE) around $50\,\mathrm{meV}$ for neutral (SCCS)[21] or neutral and cationic (MPE-1c)[20] solutes in water were reported, depending to some degree on

the model parameters and specifics of the underlying DFT calculations. We used the electronic structure program `PWSCF v.7.0rc1` of the QUANTUM ESPRESSO package[64–66] as a host for Environ. The `GNU Fortran (Gentoo 10.3.0 p1) 10.3.0` compiler with flags `-O3 -march=native -fallow-argument-mismatch` was used for compilation. Performance relevant libraries used are `AMD optimized ScaLAPACK 3.0`, `AMD BLIS 3.0`, `AMD optimized libFLAME 3.0` and `Open MPI v4.0.5`. The `H.pbe-rrkjus_psl.1.0.0.UPF` (H) and `x.pbe-n-rrkjus_psl.1.0.0.UPF` (all other elements) ultrasoft pseudopotentials (USPP) from `http://www.quantum-espresso.org` were used. Energy cutoffs for charge density and wavefunctions were chosen as suggested in the pseudopotential files. The parabolic point-counter-charge (PCC) correction scheme[67,68] implemented in Environ was used to remove interactions with periodic images in all directions, both in vacuum and implicit solvent calculations. The PCC correction, in its current implementation, requires cubic simulation cells. Cell sizes of $a = 1.1 \, d_{\max} + 5 \, \text{Å}$ were used, with $d_{\max} = \max(d_x, d_y, d_z)$, where $d_{x/y/z}$ is the extent of the solute in $x/y/z$ direction. These extents were calculated taking into account the van-der-Waals radii of the atoms, as obtained from the Atomic Simulation Environment ASE.[69] To keep the simulation cells small, prior to determining the cell size the solute molecules were rotated such that the axis corresponding to the lowest principal moment of inertia was oriented along the space diagonal. Environ's 'full' cavity definition was used. The same randomized ε as for the MPE calculations were used. The randomized ρ_{iso} were converted to the density thresholds $\rho_{\min}, \rho_{\max}$ of SCCS using the transformation formula from ref. 58, with the generic smoothness parameter $\delta_n = 2.0$ reported therein. Pressure, surface tension and electrolyte concentration were all set to 0.

Our method is implemented for parallel execution. The scaling with physical cores is tested on the largest of the test systems (entry number test1012 in the Minnesota Solvation Database[12]), containing 28 non-H atoms.

The main results of these tests are reported in figs. 13 and 14. While the LSQR solver takes consistently more time for one SCF step than the QR solver, it is never the bottleneck
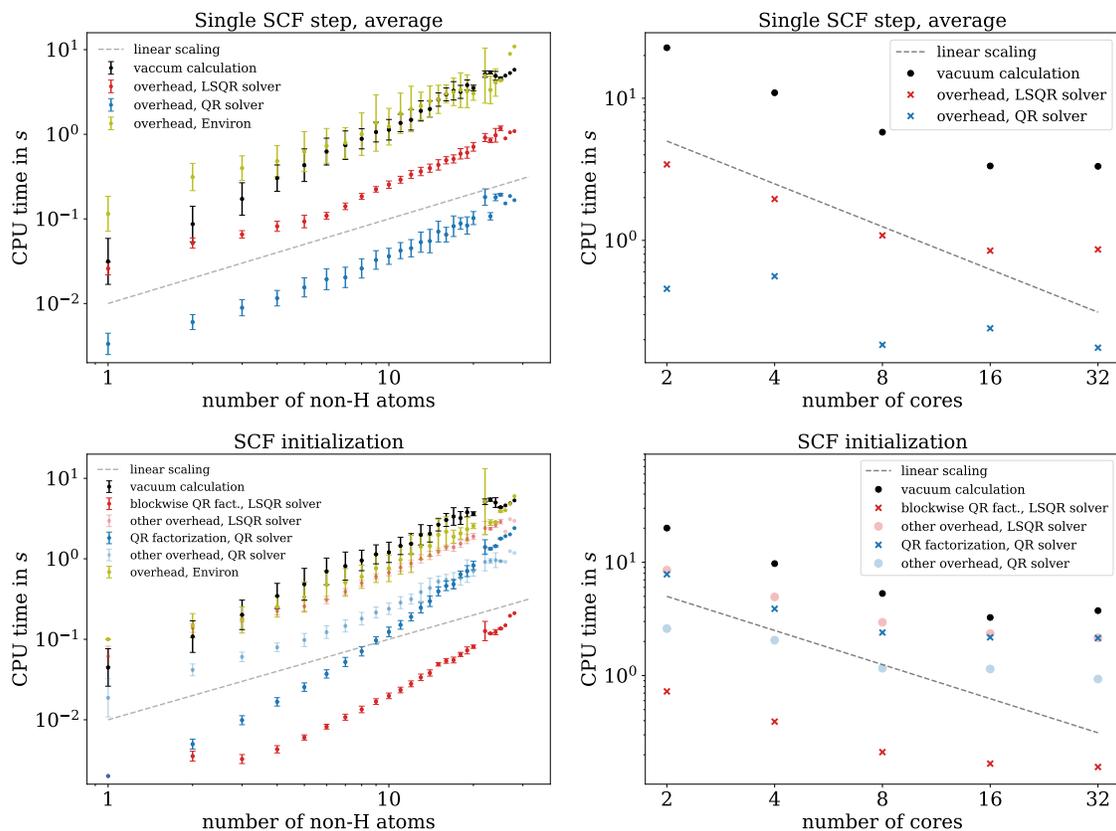
Figure 13: Scaling of CPU time with solute size on 8 cores (left) and with number of cores for largest solute in test set (right), all on a double logarithmic scale. Hypothetical linear scaling shown at arbitrary offset for reference. Negative outliers excluded from analysis. Overhead calculated as difference between values in implicit solvent and vacuum calculation. Left: Mean over solutes of same size shown as dot, root mean square deviation shown as error bar, both calculated in logarithmic space. Overhead of Environ[21] in a QUANTUM ESPRESSO[64–66] calculation shown for reference. Timing for QUANTUM ESPRESSO without implicit solvation not shown. Top: Time for 1 SCF step, averaged over all SCF steps except the last, which is typically shorter. Bottom: Time for SCF initialization. 'Other overhead' calculated as overhead minus time for (blockwise) QR factorization.

of the calculation in the sampled size range, except for very small solutes. Both solver types seem to scale almost linearly with system size, making the overhead of our method in one SCF step generally unproblematic. In the SCF initialization, both solver types cause a significant overhead. Typically, the initialization is not the bottleneck of a calculation in `FHI-aims`[20,25–29] , so the overhead here is negligible as long as it does not exceed the initialization time of a vacuum calculation by orders of magnitude. In the sampled size range, this is never the case. However, it is clear that the scaling for the QR factorization
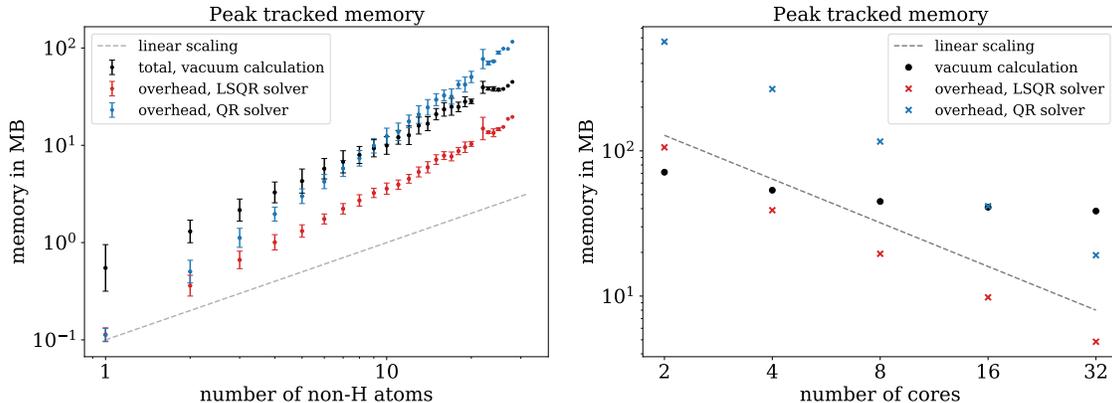
Figure 14: Scaling of highest tracked memory across all processes with solute size on 8 cores (left) and with number of cores for largest solute in test set (right), all on a double logarithmic scale. `FHI-aims`[20,25-29] does not track small arrays and scalars; real memory usage will be slightly larger.

in the QR solver is unfavorable, and this step will eventually become a bottleneck for larger solutes. The scaling of the blockwise QR factorization in the LSQR solver is not linear either, but it is generally much faster than for the QR solver, making LSQR more suitable for very large systems. Overhead other than the factorizations seems to scale no worse than the encompassing DFT initialization in both cases, with a higher offset for LSQR. Similar observations can be made in the memory requirements. Again, the scaling of the QR solver with system size is unfavorable, approximately quadratic in this case. For the LSQR solver, it is, again, much flatter, albeit not perfectly linear.

Compared to Environ,[21] MPE-*nc* with both solvers is consistently faster in the SCF steps, within the sampled size range. In the SCF initialization, the overhead is comparable in the sampled size range. This comparison is, to some degree, biased by the different design choices in the SCCS and MPE-*nc* models, as well as the characteristics of the electronic structure software used as host. Most importantly, QUANTUM ESPRESSO[64-66] is a plane wave code, formally requiring periodic boundary conditions (PBC) which are then compensated for e.g. by the PCC[67,68] scheme for non-periodic systems. The necessity of PBC can increase computational costs if significant portions of the simulation cell are empty. `FHI-aims`[20,25-29] on the other hand uses atom-centered basis functions, allowing it to conduct the computation

entirely in real space. This difference also affects the respective implicit solvation methods. Despite these limitations, this comparison illustrates the computational efficiency of MPE-nc, not only in the given context of `FHI-aims` [20,25–29] .

In terms of parallelization over cores, both solvers scale fairly well. Within one SCF step, the LSQR solver scales approximately equally as well as the encompassing DFT calculation. No considerable speedup is gained for the test system when going beyond 16 cores. This is the case both for the MPE method and the vacuum calculation. For larger solutes, parallelization over a larger number of cores will likely become more efficient. For the QR solver, no clear scaling is observed here, but the overhead is already negligibly small. In the SCF initialization, all parts of our method scale well within a certain range. Again, both for the encompassing DFT calculation and the MPE method, no speedup is gained when going beyond 16 cores, although MPE appears to reach a plateau already at slightly fewer cores. Nonetheless, all parts of both solvers consistently stay below the time for the DFT initialization. In terms of memory, the distribution is trivial and scales linearly with the number of cores for both solvers.

Finally, we briefly address computational scaling from a more theoretical point of view. The size of the central matrix $\mathbf{A}$ in eq. (14) scales $O(N^2)$ with system size N. Each non-H atom adds a constant number $(l_{\max}^{R}+1)^2 + (l_{\max}^{Q}+1)^2$ of columns to $\mathbf{A}$. The (target) number of rows depends linearly on the number of columns via $d_{\det}$. The computational time of the QR factorization will thus become the bottleneck at some point. The LSQR solver largely circumvents this issue by exploiting sparsity. In the $(l_{\max}^{R}+1)^2$ columns corresponding to the basis set of one subcavity, only approximately $2d_{\det}(l_{\max}^{R}+1)^2$ rows are filled. Each non-H atom thus only adds an approximately constant number of non-zero matrix elements in these blocks. Nonetheless, the matrix block corresponding to the basis set for Φ_Q does not get sparser in the present method. Each non-H atom adds $(l_{\max}^{Q}+1)^2$ columns, which have filled rows for all cavity surface points, the number of which also grows approximately linearly with system size. Thus, the LSQR solver will approach at least $O(N^2)$ scaling in the limit

of very large solutes, as well. However, it does so at a significantly lower offset than the QR solver. The issue of choosing the basis set for Φ_Q in a way that will sparsify $\mathbf{A}$ is beyond the scope of the present paper and left to future work.

7 Conclusions and outlook

We have shown that MPE-1c with the standard $l_{\max}^R = 8$ leads to a systematic underestimation of the electrostatic solute-solvent interaction. Already in the original publication [20] it was noted that potentially much higher expansion orders would be necessary for large solutes. We have shown, however, that the issue can sometimes occur already in relatively small molecules. Furthermore, we have seen that for larger solutes, especially using high ρ_{iso}, increasing the expansion order is not always a viable solution, with the multipole series converging slowly or not at all.

As a remedy for this issue, we have modified the MPE method, termed the MPE-nc method, by separating the solvent cavity into small subcavities, centered around the solute's non-H nuclei. The dielectric response of the solvent is expanded in an individual multipole series in each of the subcavities. In practice, this ensures a fast convergence of the multipole series. We have parameterized an implicit solvation model for water using our modified electrostatic MPE-nc approach and a simplistic model for non-electrostatic free energy contributions. Comparing our results to an equivalent parameterization using MPE-1c reveals a systematic error cancellation between electrostatic and non-electrostatic model terms in MPE-1c. While yielding surprisingly accurate results, this error cancellation can not be expected to hold for larger, more complex solutes, limiting the original method's transferability. These problems are not present in the MPE-nc approach which allows for a much more reliable reproduction of the dielectric response of a polarizable medium.

The practical applicability of our method is currently limited by the lack of atomic forces. While these lie beyond the scope of the present, self-contained work, their derivation and

implementation are an obvious next step which we plan to address in future work.

Acknowledgement

A Supporting Information

Supporting information (SI) is available for this article. There, we describe the technical details of the central linear algebra problem of our method, as well as the boundary discretization algorithm. Lastly, we provide a minimal example for the user input required for our method.

This information is available free of charge via the Internet at http://pubs.acs.org

References

(1) Levy, Y.; Onuchic, J. N. WATER MEDIATION IN PROTEIN FOLDING AND MOLECULAR RECOGNITION. *Annu. Rev. Biophys. Biomol. Struct.* **2006**, *35*, 389–415.

(2) Anantharaj, S.; Ede, S. R.; Sakthikumar, K.; Karthick, K.; Mishra, S.; Kundu, S. Recent Trends and Perspectives in Electrochemical Water Splitting with an Emphasis on Sulfide, Selenide, and Phosphide Catalysts of Fe, Co, and Ni: A Review. *ACS Catal.* **2016**, *6*, 8069–8097.

(3) Fang, G.; Zhou, J.; Pan, A.; Liang, S. Recent Advances in Aqueous Zinc-Ion Batteries. *ACS Energy Lett.* **2018**, *3*, 2480–2501.

(4) Varghese, J. J.; Mushrif, S. H. Origins of complex solvent effects on chemical reactivity and computational tools to investigate them: a review. *React. Chem. Eng.* **2019**, *4*, 165–206.

(5) Kalepu, S.; Nekkanti, V. Insoluble drug delivery strategies: review of recent advances and business prospects. *Acta Pharm. Sin. B* **2015**, *5*, 442–453.

(6) Andreussi, O.; Fisicaro, G. Continuum embeddings in condensed-matter simulations. *Int. J. Quantum Chem.* **2019**, *119*, e25725.

(7) Tomasi, J.; Mennucci, B.; Cammi, R. Quantum Mechanical Continuum Solvation Models. *Chem. Rev.* **2005**, *105*, 2999–3094.

(8) Schwarz, K.; Sundararaman, R. The electrochemical interface in first-principles calculations. *Surf. Sci. Rep.* **2020**, *75*, 100492.

(9) Ringe, S.; Hörmann, N. G.; Oberhofer, H.; Reuter, K. Implicit Solvation Methods for Catalysis at Electrified Interfaces. *arXiv preprint arXiv:2108.02461* **2021**,

(10) Kelly, C. P.; Cramer, C. J.; Truhlar, D. G. SM6: A Density Functional Theory Continuum Solvation Model for Calculating Aqueous Solvation Free Energies of Neutrals, Ions, and Solute-Water Clusters. *J. Chem. Theory Comput.* **2005**, *1*, 1133–1152.

(11) Marenich, A. V.; Cramer, C. J.; Truhlar, D. G. Universal Solvation Model Based on Solute Electron Density and on a Continuum Model of the Solvent Defined by the Bulk Dielectric Constant and Atomic Surface Tensions. *J. Phys. Chem. B* **2009**, *113*, 6378–6396.

(12) Marenich, A. V.; Kelly, C. P.; Thompson, J. D.; Hawkins, G. D.; Chambers, C. C.;

Giesen, D. J.; Winget, P.; Cramer, C. J.; Truhlar, D. G. *Minnesota Solvation Database*, version 2012. University of Minnesota: Minneapolis, 2012.

(13) Kirkwood, J. G. Theory of Solutions of Molecules Containing Widely Separated Charges with Special Application to Zwitterions. *J. Chem. Phys.* **1934**, *2*, 351–361.

(14) Onsager, L. Electric Moments of Molecules in Liquids. *J. Am. Chem. Soc.* **1936**, *58*, 1486–1493.

(15) Rivail, J.-L.; Rinaldi, D. A quantum chemical approach to dielectric solvent effects in molecular liquids. *Chem. Phys.* **1976**, *18*, 233 – 242.

(16) Rinaldi, D.; Ruiz-Lopez, M. F.; Rivail, J. Ab initio SCF calculations on electrostatically solvated molecules using a deformable three axes ellipsoidal cavity. *J. Chem. Phys.* **1983**, *78*, 834–838.

(17) Dillet, V.; Rinaldi, D.; Ángyán, J. G.; Rivail, J.-L. Reaction field factors for a multipole distribution in a cavity surrounded by a continuum. *Chem. Phys. Lett.* **1993**, *202*, 18 – 22.

(18) Rinaldi, D.; Bouchy, A.; Rivail, J.-L.; Dillet, V. A self-consistent reaction field model of solvation using distributed multipoles. I. Energy and energy derivatives. *J. Chem. Phys.* **2004**, *120*, 2343–2350.

(19) Sinstein, M. *Multipole Expansion Implicit Solvation Model in Full Potential DFT*. Dissertation, Technische Universität München, München, 2018.

(20) Sinstein, M.; Scheurer, C.; Matera, S.; Blum, V.; Reuter, K.; Oberhofer, H. Efficient Implicit Solvation Method for Full Potential DFT. *J. Chem. Theory Comput.* **2017**, *13*, 5582–5603.

(21) Andreussi, O.; Dabo, I.; Marzari, N. Revised self-consistent continuum solvation in electronic-structure calculations. *J. Chem. Phys.* **2012**, *136*, 064102.

(22) Lipparini, F.; Stamm, B.; Cancès, E.; Maday, Y.; Mennucci, B. Fast Domain Decomposition Algorithm for Continuum Solvation Models: Energy and First Derivatives. *J. Chem. Theory Comput.* **2013**, *9*, 3637–3648.

(23) Cancès, E.; Maday, Y.; Stamm, B. Domain decomposition for implicit solvation models. *J. Chem. Phys.* **2013**, *139*, 054111.

(24) Nottoli, M.; Nifosì, R.; Mennucci, B.; Lipparini, F. Energy, Structures, and Response Properties with a Fully Coupled QM/AMOEBA/ddCOSMO Implementation. *J. Chem. Theory Comput.* **2021**, *17*, 5661–5672.

(25) Blum, V.; Gehrke, R.; Hanke, F.; Havu, P.; Havu, V.; Ren, X.; Reuter, K.; Scheffler, M. Ab initio molecular simulations with numeric atom-centered orbitals. *Comput. Phys. Commun.* **2009**, *180*, 2175 – 2196.

(26) Havu, V.; Blum, V.; Havu, P.; Scheffler, M. Efficient O(N) integration for all-electron electronic structure calculation using numeric basis functions. *J. Comput. Phys.* **2009**, *228*, 8367 – 8379.

(27) zhe Yu, V. W.; Corsetti, F.; García, A.; Huhn, W. P.; Jacquelin, M.; Jia, W.; Lange, B.; Lin, L.; Lu, J.; Mi, W.; Seifitokaldani, A.; Vázquez-Mayagoitia, Á.; Yang, C.; Yang, H.; Blum, V. ELSI: A unified software interface for Kohn–Sham electronic structure solvers. *Comput. Phys. Commun.* **2018**, *222*, 267 – 285.

(28) Ren, X.; Rinke, P.; Blum, V.; Wieferink, J.; Tkatchenko, A.; Sanfilippo, A.; Reuter, K.; Scheffler, M. Resolution-of-identity approach to Hartree–Fock, hybrid density functionals, RPA, MP2 and GW with numeric atom-centered orbital basis functions. *New J. Phys.* **2012**, *14*, 053020.

(29) Ihrig, A. C.; Wieferink, J.; Zhang, I. Y.; Ropo, M.; Ren, X.; Rinke, P.; Scheffler, M.; Blum, V. Accurate localized resolution of identity approach for linear-scaling hybrid

density functionals and for many-body perturbation theory. *New J. Phys.* **2015**, *17*, 093020.

(30) Bonthuis, D. J.; Netz, R. R. Beyond the continuum: how molecular solvent structure affects electrostatics and hydrodynamics at solid-electrolyte interfaces. *J. Phys. Chem. B* **2013**, *117*, 11397–11413.

(31) Fisicaro, G.; Genovese, L.; Andreussi, O.; Mandal, S.; Nair, N. N.; Marzari, N.; Goedecker, S. Soft-Sphere Continuum Solvation in Electronic-Structure Calculations. *J. Chem. Theory Comput.* **2017**, *13*, 3829–3845.

(32) Andreussi, O.; Hörmann, N. G.; Nattino, F.; Fisicaro, G.; Goedecker, S.; Marzari, N. Solvent-Aware Interfaces in Continuum Solvation. *J. Chem. Theory Comput.* **2019**, *15*, 1996–2009.

(33) Klamt, A.; Schüürmann, G. COSMO: a new approach to dielectric screening in solvents with explicit expressions for the screening energy and its gradient. *J. Chem. Soc. Perkin Trans. 2* **1993**, 799–805.

(34) Ringe, S.; Oberhofer, H.; Hille, C.; Matera, S.; Reuter, K. Function-Space-Based Solution Scheme for the Size-Modified Poisson–Boltzmann Equation in Full-Potential DFT. *J. Chem. Theory Comput.* **2016**, *12*, 4052–4066.

(35) Jackson, J. D. In *Klassische Elektrodynamik*, 5th ed.; Witte, C., Diestelhorst, M., Müller, K., Eds.; De Gruyter: Berlin, Boston, 2013.

(36) Mathew, K.; Sundararaman, R.; Letchworth-Weaver, K.; Arias, T. A.; Hennig, R. G. Implicit solvation model for density-functional study of nanocrystal surfaces and reaction pathways. *J. Chem. Phys.* **2014**, *140*, 084106.

(37) Fisicaro, G.; Genovese, L.; Andreussi, O.; Mandal, S.; Nair, N. N.; Marzari, N.;

Goedecker, S. Soft-Sphere Continuum Solvation in Electronic-Structure Calculations. *J. Chem. Theory Comput.* **2017**, *13*, 3829–3845.

(38) Perdew, J. P.; Burke, K.; Ernzerhof, M. Generalized Gradient Approximation Made Simple. *Phys. Rev. Lett.* **1996**, *77*, 3865–3868.

(39) Heyd, J.; Scuseria, G. E.; Ernzerhof, M. Hybrid functionals based on a screened Coulomb potential. *J. Chem. Phys.* **2003**, *118*, 8207–8215.

(40) Krukau, A. V.; Vydrov, O. A.; Izmaylov, A. F.; Scuseria, G. E. Influence of the exchange screening parameter on the performance of screened hybrid functionals. *J. Chem. Phys.* **2006**, *125*, 224106.

(41) Ezekiel, M.; Fox, K. A. *Methods of Correlation and Regression Analysis: Linear and Curvilinear*, 3rd ed.; John Wiley & Sons, Inc.: New York, 1959.

(42) Huberty, C. J.; Mourad, S. A. Estimation in Multiple Correlation/Prediction. *Educ. Psychol. Meas.* **1980**, *40*, 101–112.

(43) Wherry, R. J. A New Formula for Predicting the Shrinkage of the Coefficient of Multiple Correlation. *Ann. Math. Statist.* **1931**, *2*, 440–457.

(44) Kohn, W.; Sham, L. J. Self-Consistent Equations Including Exchange and Correlation Effects. *Phys. Rev.* **1965**, *140*, A1133–A1138.

(45) Fu, C. L.; Ho, K. M. First-principles calculation of the equilibrium ground-state properties of transition metals: Applications to Nb and Mo. *Phys. Rev. B* **1983**, *28*, 5480–5486.

(46) Blackford, L. S.; Choi, J.; Cleary, A.; D'Azevedo, E.; Demmel, J.; Dhillon, I.; Dongarra, J.; Hammarling, S.; Henry, G.; Petitet, A.; Stanley, K.; Walker, D.; Whaley, R. C. *ScaLAPACK Users' Guide*. Society for Industrial and Applied Mathematics (SIAM): Philadelphia, PA, USA, 1997.

(47) Anderson, E.; Bai, Z.; Bischof, C.; Blackford, S.; Demmel, J.; Dongarra, J.; Du Croz, J.; Greenbaum, A.; Hammarling, S.; McKenney, A.; Sorensen, D. *LAPACK Users' Guide.* Society for Industrial and Applied Mathematics (SIAM): Philadelphia, PA, USA, 1999.

(48) Tough, R. J. A.; Stone, A. J. Properties of the regular and irregular solid harmonics. *J. Phys. A: Math. Gen.* **1977**, *10*, 1261.

(49) Caola, M. J. Solid harmonics and their addition theorems. *J. Phys. A: Math. Gen.* **1978**, *11*, L23.

(50) Paige, C. C.; Saunders, M. A. LSQR: An Algorithm for Sparse Linear Equations and Sparse Least Squares. *ACM Trans. Math. Softw.* **1982**, *8*, 43–71.

(51) Virtanen, P.; Gommers, R.; Oliphant, T. E.; Haberland, M.; Reddy, T.; Cournapeau, D.; Burovski, E.; Peterson, P.; Weckesser, W.; Bright, J.; van der Walt, S. J.; Brett, M.; Wilson, J.; Millman, K. J.; Mayorov, N.; Nelson, A. R. J.; Jones, E.; Kern, R.; Larson, E.; Carey, C. J.; Polat, İ.; Feng, Y.; Moore, E. W.; VanderPlas, J.; Laxalde, D.; Perktold, J.; Cimrman, R.; Henriksen, I.; Quintero, E. A.; Harris, C. R.; Archibald, A. M.; Ribeiro, A. H.; Pedregosa, F.; van Mulbregt, P.; SciPy 1.0 Contributors, SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat. Methods* **2020**, *17*, 261–272.

(52) Zhang, Y.; Yang, W. Comment on "Generalized Gradient Approximation Made Simple". *Phys. Rev. Lett.* **1998**, *80*, 890–890.

(53) Hammer, B.; Hansen, L. B.; Nørskov, J. K. Improved adsorption energetics within density-functional theory using revised Perdew-Burke-Ernzerhof functionals. *Phys. Rev. B* **1999**, *59*, 7413–7421.

(54) Becke, A. D. Density-functional exchange-energy approximation with correct asymptotic behavior. *Phys. Rev. A* **1988**, *38*, 3098–3100.

(55) Lee, C.; Yang, W.; Parr, R. G. Development of the Colle-Salvetti correlation-energy formula into a functional of the electron density. *Phys. Rev. B* **1988**, *37*, 785–789.

(56) Becke, A. D. A new mixing of Hartree–Fock and local density-functional theories. *J. Chem. Phys.* **1993**, *98*, 1372–1377.

(57) Sun, J.; Ruzsinszky, A.; Perdew, J. P. Strongly Constrained and Appropriately Normed Semilocal Density Functional. *Phys. Rev. Lett.* **2015**, *115*, 036402.

(58) Hille, C.; Ringe, S.; Deimel, M.; Kunkel, C.; Acree, W. E.; Reuter, K.; Oberhofer, H. Generalized molecular solvation in non-aqueous solutions by a single parameter implicit solvation scheme. *J. Chem. Phys.* **2019**, *150*, 041710.

(59) Van Zee, F. G.; van de Geijn, R. A. BLIS: A Framework for Rapidly Instantiating BLAS Functionality. *ACM Trans. Math. Softw.* **2015**, *41*, 14:1–14:33.

(60) Gunnels, J. A.; Gustavson, F. G.; Henry, G. M.; van de Geijn, R. A. FLAME: Formal Linear Algebra Methods Environment. *ACM Trans. Math. Softw.* **2001**, *27*, 422–455.

(61) Marek, A.; Blum, V.; Johanni, R.; Havu, V.; Lang, B.; Auckenthaler, T.; Heinecke, A.; Bungartz, H.-J.; Lederer, H. The ELPA library: scalable parallel eigenvalue solutions for electronic structure theory and computational science. *J. Phys.: Condens. Matter* **2014**, *26*, 213201.

(62) Auckenthaler, T.; Blum, V.; Bungartz, H.-J.; Huckle, T.; Johanni, R.; Krämer, L.; Lang, B.; Lederer, H.; Willems, P. Parallel solution of partial symmetric eigenvalue problems from electronic structure calculations. *Parallel Comput.* **2011**, *37*, 783–794, 6th International Workshop on Parallel Matrix Algorithms and Applications (PMAA'10).

(63) Gabriel, E.; Fagg, G. E.; Bosilca, G.; Angskun, T.; Dongarra, J. J.; Squyres, J. M.; Sahay, V.; Kambadur, P.; Barrett, B.; Lumsdaine, A.; Castain, R. H.; Daniel, D. J.;

Graham, R. L.; Woodall, T. S. Open MPI: Goals, Concept, and Design of a Next Generation MPI Implementation. Proceedings, 11th European PVM/MPI Users' Group Meeting. Budapest, Hungary, 2004; pp 97–104.

(64) Giannozzi, P.; Baroni, S.; Bonini, N.; Calandra, M.; Car, R.; Cavazzoni, C.; Ceresoli, D.; Chiarotti, G. L.; Cococcioni, M.; Dabo, I.; Corso, A. D.; de Gironcoli, S.; Fabris, S.; Fratesi, G.; Gebauer, R.; Gerstmann, U.; Gougoussis, C.; Kokalj, A.; Lazzeri, M.; Martin-Samos, L.; Marzari, N.; Mauri, F.; Mazzarello, R.; Paolini, S.; Pasquarello, A.; Paulatto, L.; Sbraccia, C.; Scandolo, S.; Sclauzero, G.; Seitsonen, A. P.; Smogunov, A.; Umari, P.; Wentzcovitch, R. M. QUANTUM ESPRESSO: a modular and open-source software project for quantum simulations of materials. *J. Phys.: Condens. Matter* **2009**, *21*, 395502.

(65) Giannozzi, P.; Andreussi, O.; Brumme, T.; Bunau, O.; Nardelli, M. B.; Calandra, M.; Car, R.; Cavazzoni, C.; Ceresoli, D.; Cococcioni, M.; Colonna, N.; Carnimeo, I.; Corso, A. D.; de Gironcoli, S.; Delugas, P.; DiStasio, R. A.; Ferretti, A.; Floris, A.; Fratesi, G.; Fugallo, G.; Gebauer, R.; Gerstmann, U.; Giustino, F.; Gorni, T.; Jia, J.; Kawamura, M.; Ko, H.-Y.; Kokalj, A.; Küçükbenli, E.; Lazzeri, M.; Marsili, M.; Marzari, N.; Mauri, F.; Nguyen, N. L.; Nguyen, H.-V.; de-la Roza, A. O.; Paulatto, L.; Poncé, S.; Rocca, D.; Sabatini, R.; Santra, B.; Schlipf, M.; Seitsonen, A. P.; Smogunov, A.; Timrov, I.; Thonhauser, T.; Umari, P.; Vast, N.; Wu, X.; Baroni, S. Advanced capabilities for materials modelling with Quantum ESPRESSO. *J. Phys.: Condens. Matter* **2017**, *29*, 465901.

(66) Giannozzi, P.; Baseggio, O.; Bonfà, P.; Brunato, D.; Car, R.; Carnimeo, I.; Cavazzoni, C.; de Gironcoli, S.; Delugas, P.; Ferrari Ruffino, F.; Ferretti, A.; Marzari, N.; Timrov, I.; Urru, A.; Baroni, S. Quantum ESPRESSO toward the exascale. *J. Chem. Phys.* **2020**, *152*, 154105.

(67) Dabo, I.; Kozinsky, B.; Singh-Miller, N. E.; Marzari, N. Electrostatics in periodic boundary conditions and real-space corrections. *Phys. Rev. B* **2008**, *77*, 115139.

(68) Dabo, I.; Kozinsky, B.; Singh-Miller, N. E.; Marzari, N. Erratum: Electrostatics in periodic boundary conditions and real-space corrections [Phys. Rev. B 77, 115139 (2008)]. *Phys. Rev. B* **2011**, *84*, 159910.

(69) Larsen, A. H.; Mortensen, J. J.; Blomqvist, J.; Castelli, I. E.; Christensen, R.; Dułak, M.; Friis, J.; Groves, M. N.; Hammer, B.; Hargus, C.; Hermes, E. D.; Jennings, P. C.; Jensen, P. B.; Kermode, J.; Kitchin, J. R.; Kolsbjerg, E. L.; Kubal, J.; Kaasbjerg, K.; Lysgaard, S.; Maronsson, J. B.; Maxson, T.; Olsen, T.; Pastewka, L.; Peterson, A.; Rostgaard, C.; Schiøtz, J.; Schütt, O.; Strange, M.; Thygesen, K. S.; Vegge, T.; Vilhelmsen, L.; Walter, M.; Zeng, Z.; Jacobsen, K. W. The atomic simulation environment—a Python library for working with atoms. *J. Phys.: Condens. Matter* **2017**, *29*, 273002.