


Automated preparation of nanoscopic structures: Graph-based sequence analysis, mismatch detection, and pH-consistent protonation with uncertainty estimates

Journal Article

Author(s):

Csizi, Katja-Sophia; [Reiher, Markus](#) 

Publication date:

2024-04-30

Permanent link:

<https://doi.org/10.3929/ethz-b-000649816>

Rights / license:

[Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International](#)

Originally published in:

Journal of Computational Chemistry 45(11), <https://doi.org/10.1002/jcc.27276>

Funding acknowledgement:

182400 - Exhaustive First-Principles Exploration of Chemical Reaction Networks for Catalysis Design (SNF)

Automated preparation of nanoscopic structures: Graph-based sequence analysis, mismatch detection, and pH-consistent protonation with uncertainty estimates

Katja-Sophia Csizi | Markus Reiher

Department of Chemistry and Applied Biosciences, ETH Zurich, Zurich, Switzerland

Correspondence

Markus Reiher, Department of Chemistry and Applied Biosciences, ETH Zurich, Vladimir-Prelog-Weg 2, Zurich, Switzerland.
Email: mreiher@ethz.ch

Funding information

Schweizerischer Nationalfonds zur Förderung der Wissenschaftlichen Forschung

Abstract

Structure and function in nanoscale atomistic assemblies are tightly coupled, and every atom with its specific position and even every electron will have a decisive effect on the electronic structure, and hence, on the molecular properties. Molecular simulations of nanoscopic atomistic structures therefore require accurately resolved three-dimensional input structures. If extracted from experiment, these structures often suffer from severe uncertainties, of which the lack of information on hydrogen atoms is a prominent example. Hence, experimental structures require careful review and curation, which is a time-consuming and error-prone process. Here, we present a fast and robust protocol for the automated structure analysis and pH-consistent protonation, in short, ASAP. For biomolecules as a target, the ASAP protocol integrates sequence analysis and error assessment of a given input structure. ASAP allows for pK_a prediction from reference data through Gaussian process regression including uncertainty estimation and connects to system-focused atomistic modeling described in Brunken and Reiher (*J. Chem. Theory Comput.* 16, 2020, 1646). Although focused on biomolecules, ASAP can be extended to other nanoscopic objects, because most of its design elements rely on a general graph-based foundation guaranteeing transferability. The modular character of the underlying pipeline supports different degrees of automation, which allows for (i) efficient feedback loops for human-machine interaction with a low entrance barrier and for (ii) integration into autonomous procedures such as automated force field parametrizations. This facilitates fast switching of the pH-state through on-the-fly system-focused reparametrization during a molecular simulation at virtually no extra computational cost.

KEYWORDS

atomistic simulation, Gaussian process, machine learning, protein structure

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs](https://creativecommons.org/licenses/by-nc-nd/4.0/) License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2023 The Authors. *Journal of Computational Chemistry* published by Wiley Periodicals LLC.

1 | INTRODUCTION

The generation of a high-fidelity initial structure of a nanoscopic chemical system for atomistic modeling is an important, but nontrivial task. Especially for functional biopolymers, for which extensive experimental structural databases exist, minimal structural rearrangements within an active site or through allosteric effects can change, drive, or even impede chemical processes such as enzyme catalysis.^{1,2}

Of special usefulness and need is a protein-oriented framework that enables control and building of structural models in a fast, reproducible, robust, physically realistic, standardized, and automated manner, making it possible to process structures with minimal and unbiased human intervention. Only by such frameworks, computational studies of large structures (for instance, relying on molecular quantum mechanics (QM), molecular mechanics (MM), QM/MM hybrid simulations, or molecular dynamics (MD)) can be efficiently carried out in a comprehensive and reproducible way.³ While efforts have been made to predict structures from sequence information alone,^{4–6} most simulations rely on experimentally curated input structures taken from X-ray crystallography or NMR spectroscopy, which are collected in the protein data bank (PDB).⁷

Leveraging experimental data for simulations is appealing and provides dense molecular information such as coordinates, atom types, residue and chain information, connectivity, conformers, and more. However, these structures often suffer from significant uncertainties in the atomic coordinates and other structural errors. For instance, incompletely resolved residues, unreasonable connectivities, and atomic clashes can occur. These mismatches must be detected and corrected, which can be handled in an automated fashion for most of the above-mentioned structural issues through implementation of simple chemical rules. Chemical valences of an atom may, for example, be evaluated by deriving connectivity information from simple distance-based rules, at least for organic and main-group chemical compounds, which make up for a large part of many nanoscopic bio-macromolecules.

Even accurately resolved experimental data that are not plagued by these mismatches require careful processing. For instance, multiple conformers or ensembles of molecular models (determined by NMR experiments) of the same structure can be condensed into a single PDB file depending on the time resolution of the underlying measurements. For enzymes, the conformers that reflect chemically relevant structural intermediates of a catalytic cycle might be only a small portion of the recorded structures in the file. Moreover, proteins may be subjected to modifications (such as mutations) to facilitate the crystallization process,⁸ which might require correction for simulations of generic parent systems.

The multitude of the above-mentioned structural uncertainties and mismatches are challenging to detect but propagate into the property prediction in a simulation, because every single atom and its position will have a decisive effect on the electronic structure of the compound. Therefore, it is vital to review carefully every molecular structure prior to simulation.

Apart from automated mismatch detection, the second challenge of structure evaluation is that X-ray diffraction usually lacks resolution

of hydrogen atoms (which often can only be recovered by more advanced techniques such as neutron scattering), and the many vacant valences need to be occupied in a chemically consistent way, reflecting also chemical constraints such as pH value and solvation. Whereas the overall pH-stability of a nanoscopic system is primarily determined by the global pK_a value of the macromolecule, specific residues adopt site-specific pK_a values based on their local environment, guiding the protonation state of this site. However, the assignment of meaningful discrete protonation configurations in a protein is by no means trivial, because the many close and long-range electrostatic interactions that affect the protonation state of one specific site must be accurately captured by the model that is employed to predict the negative logarithm of the acid dissociation constant, the pK_a . Although the protonated configuration of highest probability could, in principle, be found by extensive search in the space of protonated configurations, this will rapidly become unfeasible because of the combinatorial growth of their number with the number of protonation sites and the number of protons to be added.

In this work, we present a physically meaningful and robust pipeline for automated structure analysis and pH-consistent protonation (denoted ASAP) based on quantum chemical reference data. Our pipeline combines sequence analysis of a protein, mismatch detection, addition of hydrogen atoms, solvation, and pK_a prediction from a quantum reference. It has been implemented as part of our open-source software project SWOOSE⁹ within the SCINE¹⁰ infrastructure. While there are multiple computer programs that can complete individual steps of the structure preparation pipeline, only few unify many of the ASAP features, for instance, the protein preparation wizard¹¹ and the PDBFIXER of OpenMM.¹² Compared to them, our ASAP pipeline offers additional features as described in the following.

We accomplish error detection through a graph-isomorphism ansatz that extracts faulty occurrences of structural building blocks within a large structure. We derive protonation probabilities through a machine learning approach. In particular, a regression ansatz based on a Gaussian process is trained on experimental pK_a values and corresponding quantum chemical reference data. As any physicochemical property model should be equipped with confidence intervals,^{13,14} we provide uncertainty estimates for predicted pK_a values directly delivered by the Gaussian process, such that special consideration is given to residues whose pK_a values are close to the system pH because they are more prone to be misclassified. This delivers pK_a values from the deprotonation free energy in solution, which can be obtained with a quantum chemical method.

Our workflow is seamlessly integrated into our pipeline for automated construction of quantum classical hybrid models.¹⁵ This implies that the system-focused atomistic model¹⁶ (SFAM) for the given structure can be generated, and the corresponding quantum region can be selected in an automated fashion.¹⁵ The design of ASAP also enables automated, on-the-fly (re)-parametrization of the SFAM force field for a specific protonation state configuration, which we will discuss in detail in future work. Although our approach is optimized for biomolecular systems, it can be extended to other nanoscopic

architectures due to its graph-based and first-principles foundations, both being agnostic to the specific atomic composition.

This work is organized as follows: In Section 2, we discuss how machine learning protonation probabilities can be derived from quantum chemical reference data and corresponding confidence intervals from Gaussian processes, followed by an overview of the computational methodology applied. In Section 3, we present the conceptual design and all components of the ASAP pipeline. In Section 4, we demonstrate the structure preparation protocol for chain A of the peptide hormone insulin and subsequently study our approach for a set of protein structures extracted from the PDB. Then, we illustrate that pK_a values in solution can be derived from a linear regression ansatz, for which the isolated amino acids represent an ideal test case as their pK_a values are well known from experiment. Finally, we derive protonation probabilities for selected sites in chain A of insulin.

2 | GENERAL CONSIDERATIONS AND COMPUTATIONAL METHODOLOGY

2.1 | Protonation probabilities from a quantum reference

The pK_a value is a quantitative measure of the propensity of an acidic molecule AH to dissociate into a proton (H^+) and the corresponding base (A^-) in aqueous solution, measuring the protonation probability at a given pH. It is defined as the negative base-10 logarithm of the acidity constant K_a , which is the ratio of the concentrations of proton H^+ and base A^- and the acid AH in solution:

$$pK_a = -\log_{10}(K_a) = -\log_{10} \frac{[A^-][H^+]}{[AH]}. \quad (1)$$

Note that Equation (1) is only valid at low proton concentrations. In cases of high proton concentrations, the pK_a value should be expressed as a function of activities and activity coefficients of the respective species instead of concentrations. The acidity constant relates to the Gibbs free energy ΔG_a of AH deprotonation in solution:

$$K_a = \exp\left(-\frac{\Delta G_a}{RT}\right). \quad (2)$$

ΔG_a in solution is given by:¹⁷

$$\Delta G_a = G^{aq}(A^-) + G^{aq}(H^+) - G^{aq}(AH) = \Delta G_0^{aq} + G^{aq}(H^+) \quad (3)$$

with $G^{aq}(H^+)$, $G^{aq}(A^-)$, and $G^{aq}(AH)$ being the free energy of a proton, a base A^- , and the corresponding acid (AH) in aqueous solution. The difference of the latter two is denoted ΔG_0^{aq} . Insertion of Equation (2) into Equation (1) yields¹⁷

$$pK_a = \frac{\Delta G_a}{RT \ln(10)}. \quad (4)$$

For a reliable prediction of pK_a values, the energy contributions in Equation (3) must be calculated with sufficient accuracy, because an error of $\sim 1.36 \text{ kcal mol}^{-1}$ in ΔG_a translates into one pK_a unit of error on the logarithmic scale at $T = 293.15 \text{ K}$. This accuracy is hard to achieve for standard quantum chemical methods such as density functional theory (DFT). Moreover, it is difficult to model the free energy of a proton in solution reliably (last term in Equation (3)). Most studies rely on the indirect calculation of ΔG_a from a thermodynamic cycle (see, for instance, References 18–21). Alternatively, this quantity can be measured in experiments, which, however, can result in a mismatch with the theoretical model chosen for the calculation of the other free energy contributions.

Methods for pK_a prediction can roughly be grouped into microscopic, macroscopic, and empirical techniques. Macroscopic techniques are rooted in grid-based continuum electrostatic schemes^{22,23} solving the Poisson–Boltzmann equation. Limitations of the Poisson–Boltzmann ansatz are the neglect of specific protein-ion interactions and that properties are usually calculated for a single structure, which will only be valid if this structure is representative of protonated and deprotonated states. Moreover, solvation effects described by a dielectric continuum in an implicit solvation model will lack the explicit description of hydrogen bonding between solute and solvent, although they can be crucial.²⁴ Microscopic methods, by contrast, derive the pK_a value from quantum chemical calculations,^{20,25–29} all-atom molecular dynamics simulations,^{30–33} or quantum classical hybrid models.^{34,35} For these methods, the different routes to pK_a prediction build on the calculation of acidity in the gas or solution phase, either with a dielectric continuum or with an explicit solvent model.³⁶ Empirical pK_a prediction is mostly based on adding an empirical correction to a baseline model- pK_a value.^{37–40} This approach is valuable because of its unmatched readiness, allowing for pK_a prediction in nanoscale structures within seconds, but it lacks a rigorous physical foundation. Various pipelines implement different routes towards pK_a prediction (see, for example, PYPKA,⁴¹ PDB2PQR,^{42,43} and DELPHI²³ for Poisson–Boltzmann-based pK_a prediction, the empirical approaches $H^+ +$ ⁴⁴ and PROPKA,⁴⁵ and JAGUAR,⁴⁶ which is DFT-based).

Klicic et al. have shown that errors in the computation of terms required in Equation (4) are systematic²⁵ (see also References 17,24,47), which can be exploited to leverage pK_a prediction through a linear regression ansatz. In such an ansatz, the pK_a value of a given site can be derived from free energies of acid and base only (condensed in ΔG_0^{aq} on the right-hand side of Equation (3)). These quantities can be calculated with standard electronic structure methods, whereas the fitting parameters absorb both, the value of $G(H^+)$ and (ideally) all other deficiencies of the QM method employed:

$$pK_a = \frac{s\Delta G_0^{aq}}{RT \ln(10)} + \frac{sG^{aq}(H^+)}{RT \ln(10)} = k\Delta G_0^{aq} + C. \quad (5)$$

The scaling factor s corresponds to an error correction for some given computational method.⁴⁸ An appropriate training set is then needed to derive the regression parameters k and C . Not only does the approach then circumvent the need to calculate all quantities in

the thermodynamic cycle, but it also delivers an estimate for the free energy of a proton for the applied method:

$$G^{\text{aq}}(H^+) = \frac{C}{k}. \quad (6)$$

We apply Gaussian process (GP) regression to the regression problem in Equation (5) to derive confidence intervals for a given prediction. The data points comprise the ΔG_0^{aq} values and the corresponding experimental pK_a values. GP regression is a probabilistic Bayesian machine learning technique that predicts continuous quantities from a given prior data set. An important advantage of GP models is that they deliver an estimate for the prediction uncertainty,^{13,14,49–53} which can be leveraged for physicochemical models provided that suitable measures for molecular similarity can be found (cf., References 54–57). For a theoretical introduction to GPs, we refer to Section S1 in the Supporting Information. Although a linear relationship is known for our case, the application of GPs instead of ordinary linear regression provides several advantages here. First of all, the homoscedasticity assumption, namely that the residual terms in ordinary linear regression have the same standard deviation, is insufficient for the present problem. If our data point requires extrapolation from the range of training data, the uncertainty of the predictive mean will be high, resulting in an overall heteroscedastic error distribution. Second, GPs offer an easy way to learn (or incorporate) the uncertainty in both our variables x (precisely, the accuracy of the employed QM method to calculate ΔG_0^{aq}) and our function values y (the accuracy of the experimentally measured pK_a value). The latter is defined by the hyperparameter σ_y^2 (for details, see Section S1 in the Supporting Information). Note that Bayesian linear regression and bootstrapped linear least-squares regression⁵⁰ are alternatives for the present regression problem.

2.2 | Computational methodology

For all DFT calculations with the ORCA program package, version 4.2.0.,^{58,59} we applied the Perdew–Becke–Ernzerhof (PBE) pure⁶⁰ and the corresponding hybrid PBE0 hybrid^{61,62} density functionals with an Ahlrich's def2-TZVP basis set.⁶³ The resolution-of-the-identity density-fitting technique was selected to accelerate all PBE calculations.^{64,65} Grimme's D3 dispersion correction with Becke–Johnson damping was turned on in all DFT calculations.^{66,67} We applied the implicit solvation model SMD⁶⁸ with water as solvent. The dielectric constant was chosen to be 80.4 and the solvent probe radius was set to 1.3 Å. Thermodynamic quantities at 298.15 K were obtained by application of the standard gas-phase model of statistical thermodynamics; that is, by approximating the translational, rotational, and vibrational degrees of freedom by the ideal gas, rigid rotor, and harmonic oscillator models, respectively. We applied our MOLASSEMBLER library^{69,70} for all graph-theoretical algorithms. As part of this work, GP regression has been made freely available within our open-source SCINE¹⁰ UTILITIES library.⁷¹ All timings were measured on a computer equipped with two Intel Xeon E5-2670 CPU with 2.60 GHz base

frequency on one single thread. Structures were visualized using PyMOL 2.5.2.⁷²

3 | AUTOMATED STRUCTURE PREPARATION PIPELINE ASAP

The ASAP pipeline generates accurate structural models from some given structural input data (with the preferred file format being the PDB format). The preparation is implemented as a sequence of three consecutive steps: In step 1, sequence analysis and error detection are carried out through a subgraph isomorphism ansatz that sorts out potentially faulty parts of the overall structure. In step 2, automated protonation is performed. In step 3, the protonated structure is solvated and the corresponding topology information for subsequent MM parametrization step is generated.

As part of SCINE SWOOSE,⁹ we provide both, a fully automated and a semi-automated variant of the ASAP pipeline, which offers high flexibility with regard to the degree of interference between operator and software. In the semi-automated supervised variant, the result of each step is provided as temporary structural output data. These data can be viewed, corrected, and fed back into the pipeline for subsequent steps. The workflow is shown in Figure 1.

The building blocks for which the nanoscopic structure is searched must be sorted such that the largest building block is searched for first (to avoid false detection of occurrences in cases where smaller building blocks are elements of larger building blocks, that is, glycine is a substructure of almost all other amino acids). We discuss the individual steps of ASAP in the following sections. For a detailed step-by-step description, we refer to the SCINE SWOOSE manual on GitHub.

3.1 | Graph-based automated structure analysis

In step 1 of the ASAP pipeline, a raw molecular structure is subjected to error detection. The principle idea of our structure analysis algorithm relies on the assumption that large parts of a nanoscopic structure often consist of recurring building blocks. Examples are proteinogenic amino acids in proteins, nucleotides in nucleic acids, triglycerides in lipid double layers, monosaccharides in oligosugar molecules.

We start by automatically converting the molecular input structure to a graph employing the MOLASSEMBLER⁷³ library. Furthermore, all building blocks of a nanoscopic structure must be encoded as molecular graphs. Then, occurrences can be efficiently detected by application of graph-isomorphism techniques, that is, by querying the molecular graph of the overall structure for occurrences of each building block, which has been successfully demonstrated for related problems.^{74,75} For the graph construction, we apply flexible edge (element type) and vertex (topography; in our case: distance-based connectivity) matching criteria, and do not evaluate three-dimensional features of the building blocks (stereochemistry, conformers, and so forth).

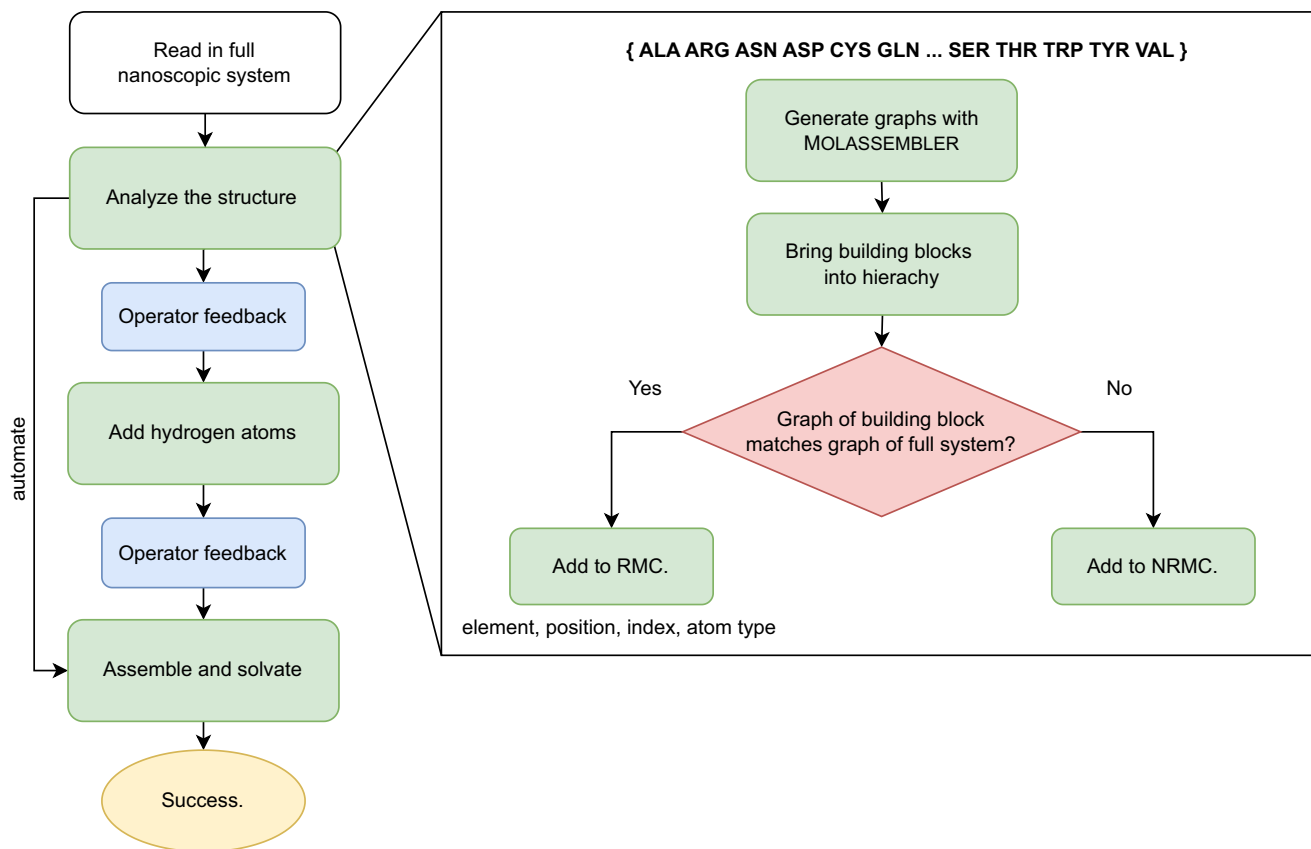


FIGURE 1 Schematic workflow of the ASAP pipeline (left) and the graph-based structure analysis step 1 (right) as implemented in SCINE Swoose (see text for further explanation). The preparation steps are highlighted in green.

Note that we define a chemical bond based on the sum of the relevant van der Waals radii (including an additional threshold of 0.4 Å). We provide the graph information of each building block as string representations of the corresponding molecular graph.

The graph-based approach offers the advantage of a flexible redefinition of the “building block”, because any molecular structure can be converted to a molecular graph. In this work, our set of building blocks consists of the 20 proteinogenic amino acids, complemented by their selenium-containing analogs of methionine and cysteine (selenomethionine, selenocysteine), and of pyrrolysine.

We fragment the input structure into a container for building blocks and into a container for all substructures that either consist of non-motif structures (containing non-amino-acid structures or amino acids with incorrect connectivity or missing atoms). We denote the resulting containers “regular module container” (RMC) and “nonregular module container” (NRMC), respectively. If applied to a protein structure, the RMC containers will collect all amino acids that form one or many polypeptide chains and the NRMC container will collect all non-amino acid structure subunits. This search algorithm efficiently characterizes the great majority of the input structure, while only a small remainder (the NRMC) is to be carefully checked for errors by the operator. This transfers the needle-in-a-haystack search, which comprises an exhaustive manual enumeration of the nanoscopic structure, to a tailored, system-specific and local evaluation of erroneous

and nonstandard sites that are buried within a large structure. For instance, amino acids with missing atoms will be transferred to the NRMC, and can be corrected by the operator. This procedure is currently favored over automatic procedures in order to ensure a high fidelity of the structures. After finalization of step 1, a temporary output is generated, where the substructures are stored in separated files.

Although currently optimized for protein-type nanoscopic systems, our software can in principle easily be extended towards other building blocks. For instance, a DNA strand can be searched for occurrences of the nucleobases by storing string representations of adenine, cytosine, guanine, and thymine. Another example would be the treatment of membrane proteins with ASAP, where our algorithms can be extended by incorporating string representations of lipid building blocks into ASAP, which are then automatically checked for in the large structure through the subgraph-isomorphism ansatz.

Moreover, we implemented two additional features into the pipeline: First, if the input structure is given in PDB format, multiple states can be separated from one another if identified through a marker in the file (for instance, different conformers or snapshots that are condensed in a single file). With this information at hand, one can either process all input structures separately or pick one or a few of them for further preprocessing. Note that we recommend discarding most or all solvent molecules in this step (except those that might




FIGURE 2 The recursive bond search algorithm implemented in the analysis step of the structure preparation pipeline.

contribute to chemical reactivity because they are in the vicinity of the active site), and reinsert them in step 3 through a more systematic explicit cluster-solvation approach (such as the one developed by us in Reference 76).

Second, we apply a recursive bond search strategy in which we iteratively transfer all correct building blocks that are covalently bound to at least one non-regular module atom (for instance, a metal center) also to the NRMC. This can be illustrated with a simple example: Structure processing of plastocyanin (PDB-ID: 1AG6⁷⁷) separates the overall structure into 737 atoms in the regular-module polypeptide chain and one single Cu^{2+} atom stored in NRMC. After recursive bond search, the amino acids Cys₈₄, His₃₇, and His₈₇ are transferred to NRMC, as shown in Figure 2. The inclusion of the direct coordination environment around the metal site into NRMC facilitates the assignment of charges and spin states, which is required to generate the SFAM model from QM reference data after structure preparation.

3.2 | Addition of hydrogen atoms

In step 2 of ASAP, hydrogen atoms (which are usually not resolved in crystallography) are added to both substructures in RMC and NRMC separately. For this purpose, we exploit the redundancy of the recurring building-block structure motifs, which were fully characterized through the graph analysis in the previous step. In an amino acid, the residue name and the atom type of each atom together form a unique key from which the protonation pattern can be extracted. We apply the IUPAC-IUB rules for atom typing in amino acids,⁷⁸ which is the default format in PDB files. In this format, atoms are labelled according to their distance from the carboxylate carbon atom; for example, the directly bound C atom is denoted “CA”. In phenylalanine (PHE), for instance, atom “CG” is a phenylic carbon atom, whereas in lysine (LYS), it is an aliphatic carbon atom, so that the protonation pattern for this atom type is only well-defined in combination with the respective amino acid. With this information at hand, all atoms in the structure can be grouped according to how many hydrogen atoms must be added and in which geometric arrangement they should be placed (for the above example, PHE-CG is tetrahedrally coordinated, whereas

LYS-CG is trigonal planar). This is schematically shown for a peptide chain as an example in Figure 3. For the 23 amino acids, the above criteria define *tetrahedral* groups (with one to three hydrogen atoms to be added), *bent* groups (with one hydrogen atom added), *trigonal planar* groups (with one hydrogen atom added) and *pseudo-tetrahedral* groups (with one or two hydrogen atoms added and at least one polyhedral corner remains vacant). Assigning atoms into one of the above groups is trivial at first glance, but it can become complicated if the protonation state of a site depends on the chemical environment (that is, the pH value and the solvent).

The NRMC substructures are protonated by the external library OPENBABEL.⁷⁹ We emphasize that OPENBABEL fills open valences with hydrogen atoms based on a hybridization assignment according to bond lengths and angles, so the result must be inspected with care. For this purpose, this step generates temporary output holding the protonated substructures in a human-readable format.

3.3 | Assembly and solvation

The third and last step in ASAP reassembles the final atomistic structure from the protonated RMC and NRMC substructures. Once merged, the substructure boundaries are cleaned up, because superfluous hydrogen atoms might have been added as an artifact of the fragmentation. Then, charge information are generated for the full structure. For the RMC, this is realized through application of simple chemical rules that derive atomic charges from the coordination environment of an atom. Charges and corresponding spin states for the NRMC need to be provided by the operator. Also, topological information about the system's connectivity are generated, and the final structure is solvated by application of a systematic microsolvation approach developed in our group.⁷⁶ As a success criterion, we evaluate with Swoose whether the SFAM MM model¹⁶ can be generated for the resulting structure: During SFAM model parametrization, the atomistic structure is fragmented, and for each fragment, basic validity checks are carried out (e.g., it is analyzed whether the assigned charge and multiplicity form a fitting pair).




FIGURE 3 Hydrogen atom addition algorithm for atom typing and grouping applied to a peptide chain of proline, cysteine, tryptophane, aspartic acid, lysine, and threonine as an example.

While the RMC is handled automatically by ASAP and always delivers chemically correct structures (i.e., valences, charges, and multiplicities are reasonable and valid), errors in the assembled full system can only be introduced by the operator in the NRMC; for instance, by assigning incorrect charge-multiplicity pairs. For that reason, inspection and manipulation of the NRMC must be carried out carefully. Any introduced errors will later be detected during the parametrization of SFAM through aforementioned validity checks.

3.4 | Machine learning pK_a prediction from a quantum reference

The assignment of a discrete protonation state for specific sites buried in a large structure as a function of environmental conditions (pH, solvent) is a long-standing problem. As discussed above, the pK_a can be predicted from quantities calculated with electronic structure methods, that is the free energy of deprotonation in solution. The unfavorable scaling of these methods with the number of electrons, however, hampers pK_a prediction through Equation (4) for nanoscopic structures.

ASAP implements a parallelizable and physically consistent (i.e., derived from quantum mechanical calculations, typically DFT, instead of empirical rules) pK_a prediction protocol that is feasible for nanoscale structures by virtue of a divide-and-conquer strategy. This is achieved by coupling the linear regression procedure reviewed in section 2.1 to the fragmentation scheme developed in the context of the SFAM model parametrization.¹⁶ The fragmentation scheme allows for the automated generation of spherical fragments around some

central atom, which, in the present context, is a pH-sensitive atom. The radius of the sphere defining the size of the generated fragment should be chosen by two criteria: First, the QM reference calculations for the generated fragment must be computationally affordable. Second, the size of the fragment sets a cut-off radius that determines to what extent long-range electrostatic interactions, which affect the protonation state, are explicitly evaluated. Every interaction beyond this cut-off radius will be disregarded, as it is not part of the fragment structure. For each site, both a protonated (AH, acidic) and a corresponding deprotonated (A, basic) fragment are automatically constructed, as is shown for one cysteine residue in chain A of insulin in Figure 4.

The two fragments are subjected to structure optimizations and subsequent Hessian calculations. These reference calculations are carried out during the MM model parametrization step and controlled by the SCINE database module.⁸⁰ This procedure offers several advantages: First, it exploits the efficient parallelization of the MM model parametrization at negligible additional cost. For example, to generate reference data for the SFAM model of insulin chain A, which consists of 328 atoms, 328 structure optimizations, and 328 subsequent Hessian calculations need to be performed, because a fragment is generated around each atom (the resulting redundancy can be exploited for internal consistency checks). The peptide chain contains only six pH-sensitive sites, that is, 12 additional reference calculations need to be performed for the corresponding second state of the site. Also, the force constants required for MM parametrization can be extracted from these reference calculations for both species. Information about them enables a change of the pH-value during a molecular simulation with fast reparametrization of SFAM without the need to




FIGURE 4 Fragmentation protocol applied to chain A of insulin, shown at the example of the pH-sensitive site in residue CYS₄₈. The corresponding acid (HA) and base (A) fragments are generated. Atom coloring: carbon, gray; hydrogen, white; oxygen, red; nitrogen, blue; and sulfur, yellow.

generate additional reference data for a different protonation state. In transferable force fields, the parameters can easily be adjusted as the atom type changes upon addition or removal of a proton. In the context of system-focused atomistic models, however, the force field parameters have been optimized for a specific molecular configuration, so that an efficient strategy for parameter re-optimization will be needed if minimal structural rearrangements occur.

4 | RESULTS AND DISCUSSION

4.1 | Automated structure processing at the example of insulin chain A

We selected chain A of the peptide hormone insulin as an example to demonstrate the capabilities of ASAP. The overall workflow, the resulting structures, and the generated output information are compiled in Figure 5. The starting structure was a hexameric hormone that featured 10 conformers resolved in the experimental PDB file and was taken from PDB-ID 1AI0.⁸¹ The zinc active site (Zn²⁺) coordinates three histidine ligands and one water ligand in a tetrahedral configuration.

We applied step 1 of ASAP to separate the overall structure into its 10 conformers, all being candidates for further preprocessing. In all conformers, only two solvent molecules were resolved in the structure close to the active site, and they were therefore kept. After separation of the 10 candidates, we selected candidate 4 for further preprocessing. In all other candidates, at least two of the three histidine ligands were coordinated to the metal site at a distance of at least 2.4 Å, which

clearly exceeds both the tight range of 1.95–2.10 Å for this type of bonds known from experiment⁸² and our definition of a chemical bond derived from van der Waals radii (see above). Note that none of the 10 candidate structures exhibited interatomic distances that fit either one of the two criteria for all Zn-coordinating ligands.

The graph-based analysis in step 1 resulted in a separation into 2380 RMC atoms and 96 NRMC atoms, as depicted in the first gray box in Figure 5. The NRMC container was filled with the metal-center active site and the corresponding ligands. In the second step, hydrogen atoms were added to the RMC substructure by applying the clustering strategy based on atom typing and amino acid categorization as described above. All pH-sensitive sites were protonated such that the site is uncharged, by default. After the addition of hydrogen atoms to the NRMC substructure, we assigned a charge of +2 to both zinc atoms (central gray box in Figure 5). Ultimately, the final processing step 3 assembled both substructures and corrected the boundaries (because superfluous hydrogen atoms were added as a result of cutting bonds at the substructure interface), and determined the topology by a distance-based connectivity evaluation. All corresponding information was stored in separate files. This yielded the final structure depicted in the right gray box in Figure 5.

4.2 | ASAP preparation benchmark and performance evaluation

To investigate the overall efficiency and stability, we applied ASAP to three different data sets: The RS126 data set,⁸³ the CB396 data set,⁸⁴

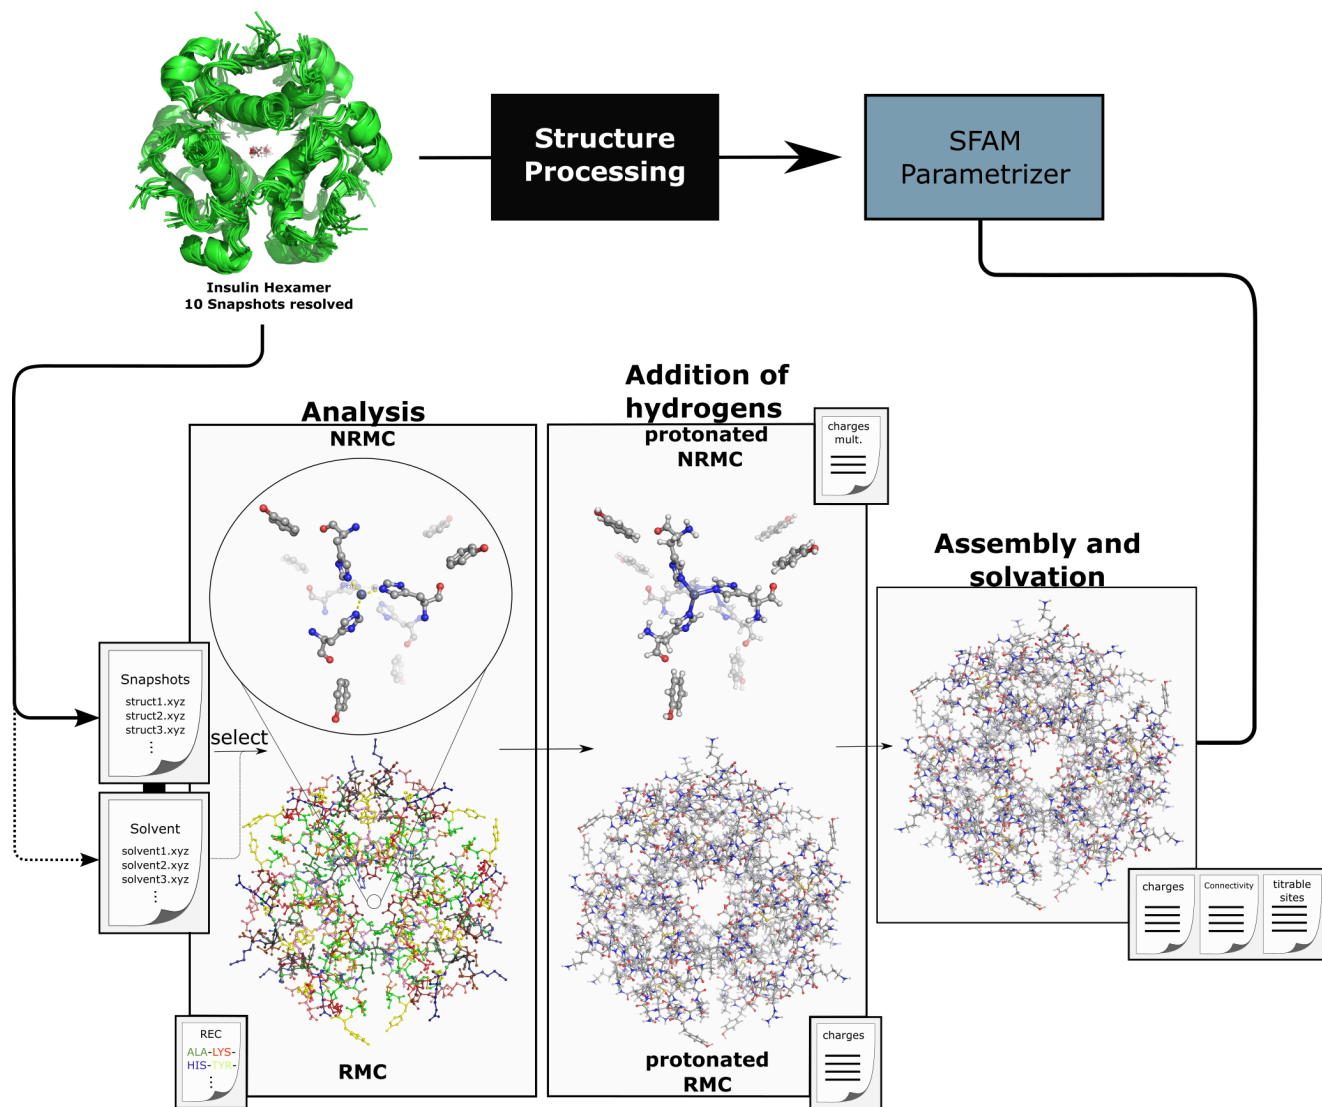


FIGURE 5 Schematic representation of the ASAP pipeline applied to chain A of insulin.

and the EVA6⁸⁵ data set. The corresponding PDB identification codes are compiled in Table S1 of the Supporting Information. We combined these three data sets to one large test set, which consists of, in total, 589 structures. After ASAP processing, we found that only 210 structures required no further preprocessing, which corresponds to 35% of the data set. That is, they consisted exclusively of amino-acid motifs that have a chemically correct composition and connectivity. Overall, 153 structures have at least two conformations resolved. For 88 % of the test-set compounds (518 structures), we found that the NRM C substructure accounted for less than 5 % in terms of atom count of the full system. Hence, in most cases, only a very small portion of the full structure requires manual inspection and curation.

For completeness, we compared the results obtained with ASAP for selected PDB-IDs to the PDBFIXER software provided by OpenMM. PDBFIXER also realizes different degrees of automation and precise control of every part of the process. The main difference to ASAP is

that the pH-specific protonation states are derived empirically, and that it cannot be directly coupled to the parametrization of a system-focused atomistic model.

Regarding performance, our focus in this discussion is primarily on assessing the time required for the graph isomorphism step. This is driven by two key factors. First, it is the most time-consuming step in ASAP. Second, a majority of the structures in the test set is unsuitable for fully automated processing because they contain multiple non-standard atoms or residues, as illustrated in Figure 6. Therefore, an analysis of the times of the fully automated preparation step would not be representative of the philosophy of ASAP. The results are summarized in Figures 6 and 7.

We emphasize that structure preparation is computationally negligible compared to subsequent SFAM parametrization, which can consume days to weeks for large systems. For the largest error-free structure (i.e., a structure that can be processed in full automation), all three steps were completed within few minutes.




FIGURE 6 (i) Atom-count ratios between REC and NREC substructures for the test sets.




FIGURE 7 Distribution of wall times for the preparation of the CB396, RS126, and EVA6 datasets with PDBFIXER (purple) and ASAP (blue). The ASAP data refer to the graph-based analysis step only.

4.3 | Quality assessment of structures prepared with ASAP

To assess the structural integrity and validity of structures generated by ASAP, we subjected a subset of the dataset to short MD simulations (0.3 ns pressure equilibration and 1 ns NPT production run), and evaluated selected features from these molecular dynamics trajectories. We furthermore compared the standard deviation of the total energy as a criterion for overall structural stability to results from structures processed with PDBFIXER or OpenMM.¹² Currently, our system-focused atomistic model is not fully compatible with available MD software. This is due to the fact that we generate system-focused atom-types and force field parameters for every individual structure, together with customized functional forms (see Reference 16), which can not directly be interfaced to a standard MD engine. Therefore, we only selected those structures in the dataset that solely consist of

RMC substructures (which are, in total, 243 structures of our dataset). In this case, the ASAP-structures (provided as XYZ coordinates) can be converted into PDB files that are compatible with transferable force fields. For all MD calculations, we employed the OPENMM program package and the ff14SB force field.⁸⁶ The pressure was set to 1 bar, the temperature to 298.15 K, the time step to 2.0 fs, the friction to 1.0 ps^{-1} , and every 100th step being recorded. All respective structures have been processed by ASAP in a fully automated manner, that is, no operator interference was required, but the basic ASAP suggestions were applied.

As illustrated in Figure 8, all inspected ASAP-structures remain stable over a simulation time of 1 ns, maintaining structural integrity. Furthermore, density and temperature quickly converge, which is indicative of effective calibration and high-quality structures. The standard deviation of the total energy during the simulation is on par or similar to the well-established PDBFIXER.

4.4 | Automated pK_a prediction

4.4.1 | Construction of a suitable training set

To train the GP model, we assemble six small-molecule data sets for each of the six different pH-sensitive functional groups that can be found in amino acids (*vide infra*). These groups comprise thiol groups (–SH), alcohol groups (–OH), phenol groups (Ph–OH), carboxylate groups (–COOH), imidazole groups (Im–H), and primary amines (–NH₂). Once ΔG_0^{aq} is calculated for all training-set molecules, the GP must be trained only once for each combination of functional-group class and electronic structure method. We selected candidate molecules for the six training sets according to two criteria: (i) Accurate experimentally measured pK_a values should be available for (ii) preferably small and rigid molecular structure. The second criterion mitigates the computational cost for training data evaluation for different conformers. The training-set molecules and corresponding pK_a values for phenol and alcohol groups were taken from Reference 28, for thiol groups from Reference 87, for imidazole groups from References 88 and 48, for amines from References 48 and 89, and for carboxylate groups from Reference 90; see Figure 9. Their Cartesian coordinates are given in Tables S2–S7 in the Supporting Information.

We evaluate the relationship between calculated ΔG_0^{aq} and experimental reference pK_a through comparison of ordinary linear regression and GP regression. As shown in Figure 10, the GP efficiently learns the linear relationship from only six data points each. Compared to linear regression, which assumes a homoscedastically distributed error, we find that the uncertainty of the predicted pK_a from the GP is higher in regions where our training-set data are sparse (and grows even larger outside the boundaries of the training set, which is omitted in Figure 10). GP models work well for the given regression problem, because they can be trained on small training data sets. However, the training data presented in this work may be extended to cover an even broader pH region.

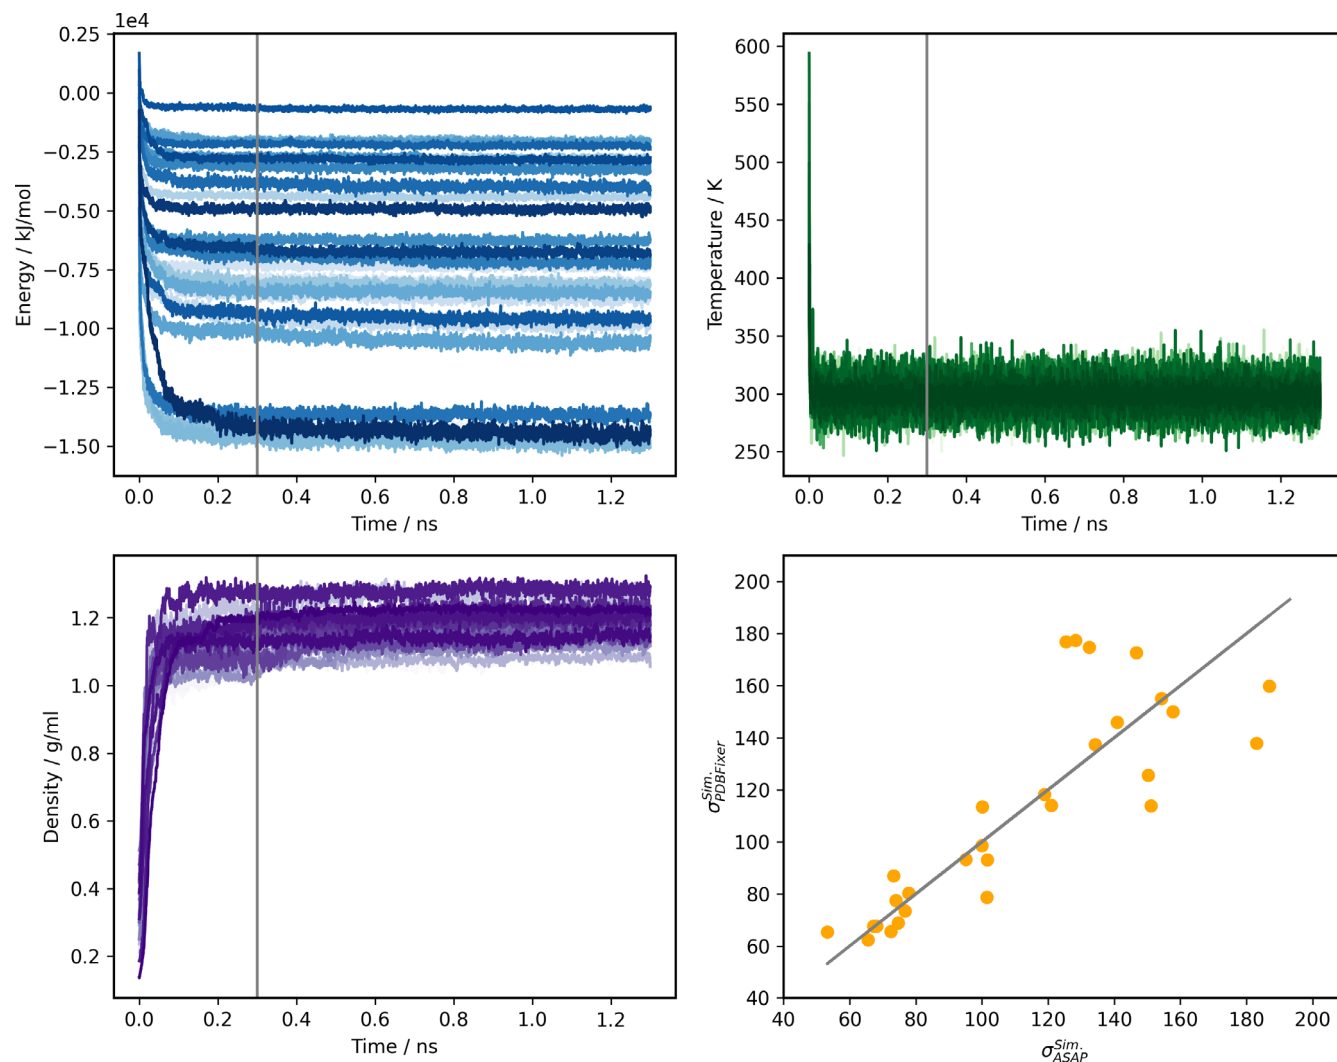


FIGURE 8 Total energy (top left), temperature (top right), and density (bottom left) as a function of simulation time for selected structures from the test set. The gray line indicates the end of NPT equilibration and the start of the actual simulation. On the bottom right, the standard deviation of total energy during the simulation ($0.3 \text{ ns} \leq t \leq 1.3 \text{ ns}$) of ASAP-structures compared to PDbFixER-structures is depicted. The respective structures and trajectories generated by ASAP and the PDbFixER can be found in the ASAPdata dataset on Zenodo.

We investigated whether semiempirical methods are a viable alternative to conventional DFT, as they allow for very fast pK_a prediction even for large structures at moderate computational cost and have therefore been widely applied to nanoscopic structures.^{91–93} The overall findings indicate that the reliability of semiempirical methods in predicting protonation probabilities should be carefully assessed. For a detailed discussion, see Section S4 Figures S2 and S3 in the Supporting Information.

4.4.2 | pK_a Prediction for isolated amino acids in aqueous solution

We applied our trained GP model to report pK_a values of the isolated amino acids glutamic acid (GLU), aspartic acid (ASP), lysine (LYS),

tyrosine (TYR), histidine (HIS) and cysteine (CYS) because there are reference pK_a values available for these amino acids at standard conditions in aqueous solution. We calculated pK_a values in two ways: 1) free energy difference of acid and base in solution were evaluated according to Equation (5), and 2) electronic energy differences, which do not contain entropy and enthalpy contributions of the nuclear scaffold in the condensed phase, were considered as an efficient approximation. The predicted pK_a values are compiled in Table 1.

We found that the predicted pK_a values for the six amino acids are overall consistent with the experimental values, with an MAE of 0.22 pK_a units for PBE-D3 and 0.36 pK_a units for PBE0-D3. The PBE0-D3 functional slightly overestimated the pK_a values compared to the PBE-D3 functional with a mean signed error of 0.19 pK_a units compared to 0.01 pK_a units. Note that we found basis sets of at least triple-zeta quality to be required for sufficiently reliable predictions. In

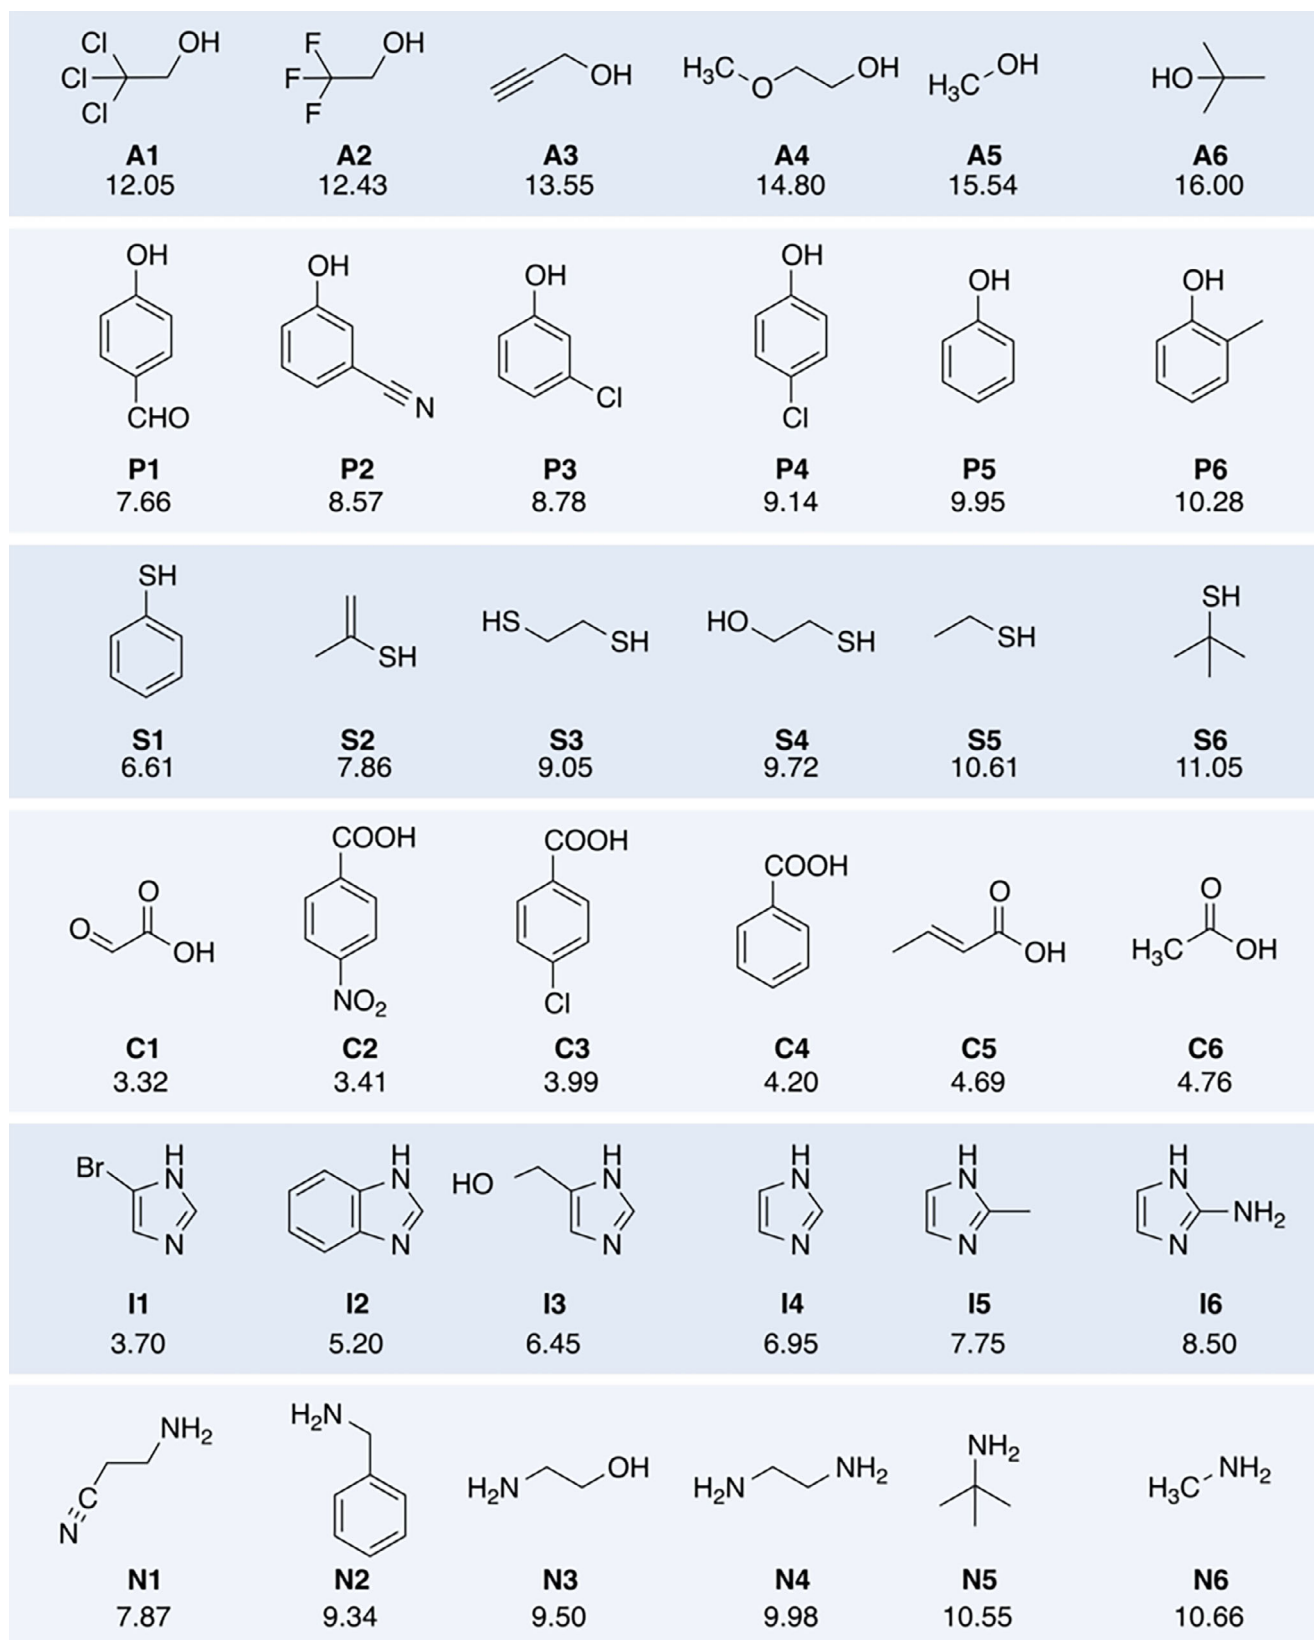


FIGURE 9 Training-set molecules with known experimental pK_a values (given below the label for the molecule) that were employed in the GP. The compounds are labeled according to their functional group: Alcohol groups: **A1-A6**, phenol groups: **P1-P6**, thiol groups: **S1-S6**, carboxylic acid groups: **C1-C6**, imidazole groups: **I1-I6**, amines: **N1-N6**.

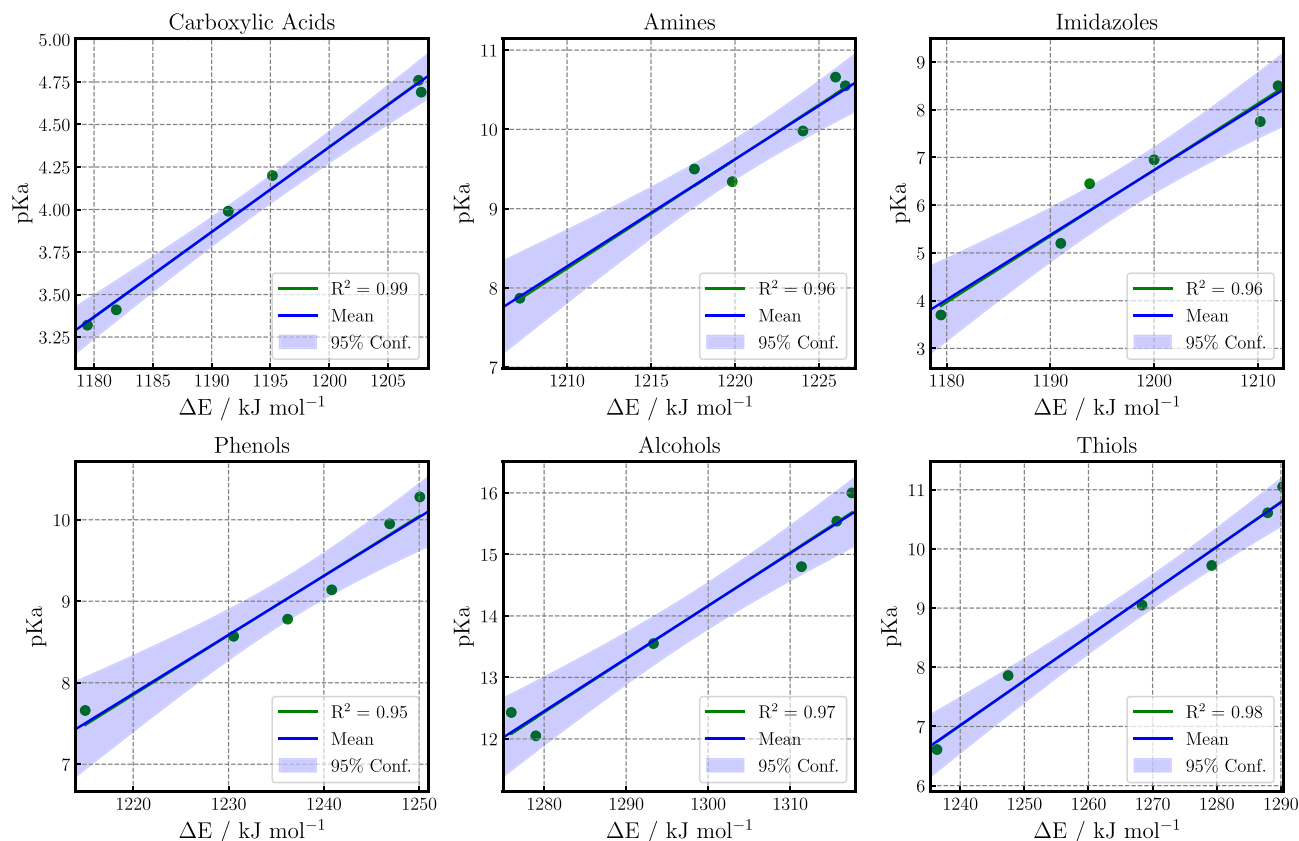


FIGURE 10 GP regression (blue) and corresponding linear regression (green) of training data from Figure 9.

TABLE 1 Predicted pK_a values of side chains of different amino acids (AA) obtained with GP regression based on PBE-D3 and PBE0-D3 results compared to model pK_a ^{94,95} data at standard conditions.

AA	Ref. pK_a	Predicted pK_a value			
		PBE-D3		PBE0-D3	
		From ΔE_0^{aq}	From ΔG_0^{aq}	From ΔE_0^{aq}	From ΔG_0^{aq}
GLU	4.4	4.20 ± 0.26	4.74 ± 0.41	4.67 ± 0.38	4.86 ± 0.45
ASP	4.0	4.37 ± 0.29	4.42 ± 0.31	4.39 ± 0.29	4.53 ± 0.34
HIS	6.3	6.52 ± 0.25	6.54 ± 0.25	6.04 ± 0.26	6.18 ± 0.26
CYS	8.5	8.38 ± 0.29	8.47 ± 0.28	8.28 ± 0.30	8.49 ± 0.28
TYR	9.6	9.73 ± 0.32	9.74 ± 0.32	9.58 ± 0.29	9.55 ± 0.29
LYS	10.4	10.07 ± 0.28	10.43 ± 0.33	9.42 ± 0.26	10.32 ± 0.32
MAE		0.22	0.26	0.36	0.20

Note: Mean absolute errors (MAE) were calculated as the sum of the absolute errors between the reference pK_a values of the six amino acids, and ASAP's predictive mean, weighted equally. For each density functional, the first column contains pK_a values learned from electronic energy differences of acid and base (ΔE_0^{aq}), while the second column derives this property from Gibbs free energy differences (ΔG_0^{aq}). For an overview of pK_a values of amino acids in proteins, we refer the interested reader to Reference 95 and references therein.

addition, we found no significant improvement in predicted pK_a values when free energy differences were evaluated compared to electronic energy differences (for a demonstration of quantitative agreement in predicted pK_a values see Table 1, and for the verification of a linear relation between ΔG and the pK_a see Section S3, Figure S1, in the Supporting Information). Hence, in the following, we chose for pK_a prediction in nanoscopic structures electronic energy differences only.

4.4.3 | pK_a prediction in nanoscopic structures

We demonstrate pK_a prediction for nanoscale structures at the example of chain A of the peptide hormone insulin. The initial structure was taken from PDB-ID 1A10.⁸¹ We calculated the pK_a value of six pH-sensitive sites in insulin chain A, namely GLU₂₇, CYS₄₈, TYR₁₀₁, GLU₁₂₇, TYR₁₄₇ and CYS₁₅₃. The subscripts denote the index of the

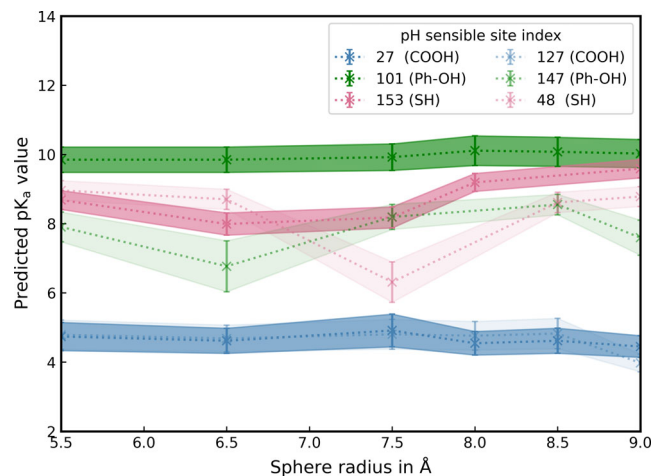


FIGURE 11 Calculated pK_a values for pH-sensitive sites in insulin chain A and the corresponding prediction uncertainty. Note that the lines between the data points should only guide the eye. Carboxylate groups are labelled in blue, phenol groups are labelled in green, thiol groups are labelled in pink.

pH-sensitive atoms that were either protonated or deprotonated. As structure optimizations are very costly for structures with more than 300 atoms, we generated different-sized fragments around the pH-sensitive atoms and optimized the structures with PBE-D3/def2-TZVP. The size of the fragment determined the degree to which long-range electrostatic interactions were explicitly considered when evaluating the pK_a value. We report the sensitivity of predicted pK_a values with respect to system size in Figure 11.

We found that especially carboxylate groups (atom indices 27 and 127 in the full structure) were very stable over a broad range of differently sized spheres, highlighting that all relevant interactions were accurately modelled even within the smallest sphere of 5.5 Å. Moreover, we found that one of the phenolic pH-sensitive atoms (atom index 101) was insensitive toward fragment size, and one (atom index 147) exhibited slightly increased sensitivity. For thiol groups, we found the largest variation with regard to system size and no convergence of the pK_a value with increasing fragment size. This observation suggests that the variations in predicted acidity constants are not a consequence of lacking electrostatic interactions, but rather a deficiency of the method or the structure. This issue may be mitigated by application of explicit solvent models, more accurate basis sets, and the generation of multiple conformers.

In the current implementation of ASAP, we do neither distinguish between the two protonation states of the two oxygen atoms in the carboxylic group of ASP and GLU, nor between those of the two nitrogen atoms in the imidazole group of HIS. This is because we assess protonation probabilities only after structure optimization, during which the R-C bond of R-COOH and R-Imidazole are not constrained, but in principle allowed to rotate, rendering the respective sites indistinguishable in principle. We stress that this rotation can very well be sterically constrained by the protein environment and would therefore require 2^n reference calculations for n potential

TABLE 2 pK_a values in chain A of insulin predicted with different available models.

Amino acid (AA)	PYPKA ⁴¹	PROPKA ⁴⁵	SCINE SWOOSSE	Isolated AA
GLU ₂₇	4.09	4.56	4.74 ± 0.41	4.4
CYS ₄₈	-	-	8.96 ± 0.28	8.5
TYR ₁₀₁	10.16	10.05	9.85 ± 0.37	9.6
GLU ₁₂₇	4.58	4.66	4.79 ± 0.42	4.4
TYR ₁₄₇	10.63	10.22	7.90 ± 0.43	9.6
CYS ₁₅₃	-	-	8.68 ± 0.27	8.5

Note: The results from our model refer to fragments with an initial radius of 5.5 Å. Note that neither with PYPKA nor PROPKA delivered pK_a values for cysteine groups.

TABLE 3 pK_a values of selected residues in PDB 1A2P compared to experimental pK_a values from References 97, 98 taken from the PKAD database.⁹⁹

Residue	Experiment	ASAP	PROPKA	PYPKA
ASP ₈	3.1 ± 0.1 ⁹⁷	3.0 ± 0.6	2.86	2.9
HIS ₁₈	7.8 ± 0.02 ⁹⁷	7.2 ± 0.3	5.37	6.79
GLU ₇₃	2.1 ± 0.1 ⁹⁸	2.0 ± 1.7	5.18	2.73
HIS ₁₀₂	6.3 ± 0.05 ⁹⁸	6.2 ± 0.3	5.87	6.11

protonation sites. The same holds true for pH-sensitive sites of different amino acids that are spatially close, and can couple (as, for instance, GLU35 and ASP52 in HEWL³³). These coupling of protonation states is currently not supported by ASAP.

Furthermore, we compared the results obtained with ASAP to those obtained with the software packages PYPKA and PROPKA for pK_a prediction. The results are summarized in Table 2.

We find that the ansatz developed by us predicts pK_a values in the same order of magnitude as the empirical approaches PYPKA and PROPKA for all functional groups except for TYR₁₄₇. Here, our model estimates a pK_a value of 7.90 ± 0.43, whereas the other models predict pK_a values close to the reference value of the isolated amino acid tyrosine. The decrease of 3 pK_a units predicted by our model is due to a stabilization of the deprotonated state through a hydrogen bond with a neighboring amino-group, which increases the acidity of this site.

To further validate the quality of pK_a prediction, we assess the predictive performance of ASAP for selected active sites in the Barnase structure (PDB-ID: 1A2P)⁹⁶, for which experimentally well-resolved pK_a values are available in the literature. The results are summarized in Table 3.

For the selected pH-sensitive sites in 1A2P, we find excellent agreement between pK_a values from experiment and predicted with ASAP. For the two histidine residues, ASAP data correctly reflect that the pK_a of HIS₁₈ is more than one magnitude larger than the one of HIS₁₀₂, whereas PROPKA predicts both values to be below 6.0 pK_a units, and PYPKA predicts both values to be greater than 6.0, respectively. Furthermore, the pK_a of ASP₈ in Barnase is

approximately one order of magnitude decreased compared to the isolated amino acid, a property that is accurately predicted by ASAP. However, the predicted uncertainty for GLU_{73} is comparatively high, owing to the fact that the considerably low pK_a value of this site exceeds the range covered by our suggested training set, which can, however, be further extended.

5 | CONCLUSIONS

Molecular simulation of chemical reactions in nanoscopic environments requires accurate, atomistically resolved molecular structures, because molecular function is tightly coupled to the electronic structure to which every individual atom and electron contribute decisively. However, experimentally determined structures are often subject to significant uncertainties. Hence, structures obtained by experiment must be carefully analyzed and processed before being considered as a starting point for computational modeling. However, this cannot be done manually for routine and high-throughput applications. Our ASAP workflow for structure analysis and protonation achieves automation to a high degree and has been optimized for biomolecular environments. It generates an accurate structural model from experimental structure information in a fast and efficient manner. Moreover, it combines sequence analysis, error detection, machine learning protonation and pK_a prediction as well as solvation in a single framework. The ASAP workflow seamlessly integrates with our automated SFAM molecular mechanics and hybrid model construction approach, where the corresponding system-focused force field and a quantum region can be generated for the ASAP-prepared structure, which offers a physically consistent protocol for QM/MM modeling.

ACKNOWLEDGMENTS

The authors gratefully acknowledge financial support through SNF project no. 200021_182400. Moreover, K.-S.C. was supported by ETH Zurich and SNF grant no. 200021_172950-1 awarded to PD Dr. Thomas Hofstetter.

DATA AVAILABILITY STATEMENT

ASAP has been implemented in our open-source software SCINE SWOOSE.⁹ SCINE UTILITIES and SCINE MOLASSEMBLER are available in References 71 and 70. The SCINE framework is also available from GitHub (<https://github.com/qcscine>). Additionally, all Cartesian coordinates of training and test set molecules, and data generated by ASAP can be found in the ASAPDATA repository on Zenodo.

REFERENCES

- [1] S. J. Wodak, E. Paci, N. V. Dokholyan, I. N. Berezovsky, A. Horovitz, J. Li, V. J. Hilser, I. Bahar, J. Karanicolas, G. Stock, P. Hamm, R. H. Stote, J. Eberhardt, Y. Chebaro, A. Dejaegere, M. Cecchini, J.-P. Changeux, P. G. Bolhuis, J. Vreede, P. Faccioli, S. Orioli, R. Rivasio, L. Yan, C. Brito, M. Wyart, P. Gkeka, I. Rivalta, G. Palermo, J. A. McCammon, J. Panecka-Hofman, R. C. Wade, A. Di Pizio, M. Y. Niv, R. Nussinov, C.-J. Tsai, H. Jang, D. Padhorny, D. Kozakov, T. McLeish, *Structure* **2019**, *27*, 566.
- [2] Y. Zhou, S. Musalgaonkar, A. W. Johnson, D. W. Taylor, *Nat. Commun.* **2019**, *10*, 1.
- [3] M. Garcia-Viloca, T. D. Poulsen, D. G. Truhlar, J. Gao, *Protein Sci.* **2004**, *13*, 2341.
- [4] A. W. Senior, R. Evans, J. Jumper, J. Kirkpatrick, L. Sifre, T. Green, C. Qin, A. Židek, A. W. Nelson, A. Bridgland, H. Penedones, S. Petersen, K. Simonyan, S. Crossan, P. Kohli, D. T. Jones, D. Silver, K. Kavukcuoglu, D. Hassabis, *Nature* **2020**, *577*, 706.
- [5] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Židek, A. Potapenko, A. Bridgland, C. Meyer, S. A. Kohli, A. J. Ballard, A. Cowie, B. Romera-Paredes, S. Nikolov, R. Jain, J. Adler, T. Back, S. Petersen, D. Reiman, E. Clancy, M. Zielinski, M. Steinegger, M. Pacholska, T. Berghammer, S. Bodenstein, D. Silver, O. Vinyals, A. W. Senior, K. Kavukcuoglu, P. Kohli, D. Hassabis, *Nature* **2021**, *596*, 583.
- [6] M. Baek, F. DiMaio, I. Anishchenko, J. Dauparas, S. Ovchinnikov, G. R. Lee, J. Wang, Q. Cong, L. N. Kinch, R. D. Schaeffer, C. Millán, H. Park, C. Adams, C. R. Glassman, A. DeGiovanni, J. H. Pereira, A. V. Rodrigues, A. A. van Dijk, A. C. Ebrecht, D. J. Opperman, T. Sagmeister, C. Buhlheller, T. Pavkov-Keller, M. K. Rathinaswamy, U. Dalwadi, C. K. Yip, J. E. Burke, K. C. Garcia, N. V. Grishin, P. D. Adams, R. J. Read, D. Baker, *Science* **2021**, *373*, 871.
- [7] H. M. Berman, *Nucleic Acids Res.* **2000**, *28*, 235.
- [8] K. L. Longenecker, S. M. Garrard, P. J. Sheffield, Z. S. Derewenda, *Acta Crystallogr. Sect. D Biol. Crystallogr.* **2001**, *57*, 679.
- [9] C. Brunken, K.-S. Csizi, M. Reiher, qcscine/swoose: Release 1.0.0. **2021**.
- [10] The Reiher Research Group, SCINE – Software for Chemical Interaction Networks. <http://scine.ethz.ch/>.
- [11] G. Madhavi Sastry, M. Adzhigirey, T. Day, R. Annabhimoju, W. Sherman, *J. Comput. Aided Mol. Des.* **2013**, *27*, 221.
- [12] P. Eastman, J. Swails, J. D. Chodera, R. T. McGibbon, Y. Zhao, K. A. Beauchamp, L.-P. Wang, A. C. Simmonett, M. P. Harrigan, C. D. Stern, R. P. Wiewiora, B. R. Brooks, V. S. Pande, *PLoS Comput. Biol.* **2017**, *13*, 1.
- [13] J. Proppe, T. Husch, G. N. Simm, M. Reiher, *Faraday Discuss.* **2016**, *195*, 497.
- [14] G. N. Simm, J. Proppe, M. Reiher, *Chimia* **2017**, *71*, 202.
- [15] C. Brunken, M. Reiher, *J. Chem. Theory Comput.* **2021**, *17*, 3797.
- [16] C. Brunken, M. Reiher, *J. Chem. Theory Comput.* **2020**, *16*, 1646.
- [17] K. S. Alongi, G. C. Shields, *Annu. Rev. Comput. Chem.* **2010**, *6*, 113.
- [18] J. Ho, M. L. Coote, *Theor. Chem. Accounts* **2009**, *125*, 3.
- [19] R. Casanovas, J. Ortega-Castro, J. Frau, J. Donoso, F. Muñoz, *Theoretical pKa calculations with continuum model solvents, alternative protocols to thermodynamic cycles.* **2014**, *114*, 1350.
- [20] J. Ho, *Phys. Chem. Chem. Phys.* **2015**, *17*, 2859.
- [21] P. G. Seybold, G. C. Shields, *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **2015**, *5*, 290.
- [22] D. Bashford, M. Karplus, *Biochemistry* **1990**, *29*, 10219.
- [23] L. Wang, L. Li, E. Alexov, *Proteins Struct. Funct. Bioinforma.* **2015**, *83*, 2186.
- [24] J. Ho, *Aust. J. Chem.* **2014**, *67*, 1441.
- [25] J. J. Klicić, R. A. Friesner, S. Y. Liu, W. C. Guida, *J. Phys. Chem. A* **2002**, *106*, 1327.
- [26] A. Klamt, F. Eckert, M. Diedenhofen, M. E. Beck, *J. Phys. Chem. A* **2003**, *107*, 9380.
- [27] A. G. Riojas, A. K. Wilson, *J. Chem. Theory Comput.* **2014**, *10*, 1500.
- [28] B. Thapa, H. B. Schlegel, *J. Phys. Chem. A* **2017**, *121*, 4698.
- [29] P. Lian, R. C. Johnston, J. M. Parks, J. C. Smith, *J. Phys. Chem. A* **2018**, *122*, 4366.
- [30] S. Donnini, F. Tegeler, G. Groenhof, H. Grubmüller, *J. Chem. Theory Comput.* **2011**, *7*, 1962.

- [31] S. G. Itoh, A. Damjanović, B. R. Brooks, *Proteins Struct. Funct. Bioinforma.* **2011**, *79*, 3420.
- [32] J. A. Wallace, J. K. Shen, *J. Chem. Theory Comput.* **2011**, *7*, 2617.
- [33] Y. Huang, W. Chen, J. A. Wallace, J. Shen, *J. Chem. Theory Comput.* **2016**, *12*, 5411.
- [34] H. Li, A. W. Hains, J. E. Everts, A. D. Robertson, J. H. Jensen, *J. Phys. Chem. B* **2002**, *106*, 3486.
- [35] G. Li, Q. Cui, *J. Phys. Chem. B* **2003**, *107*, 14521.
- [36] J. Ho, M. L. Coote, *WIREs Comput. Mol. Sci.* **2011**, *1*, 649.
- [37] R. Godoy-Ruiz, R. Perez-Jimenez, M. M. Garcia-Mira, I. M. Plaza del Pino, J. M. Sanchez-Ruiz, *Biophys. Chem.* **2005**, *115*, 263.
- [38] H. Li, A. D. Robertson, J. H. Jensen, *Proteins Struct. Funct. Genet.* **2005**, *61*, 704.
- [39] E. Krieger, J. E. Nielsen, C. A. Spronk, G. Vriend, *J. Mol. Graph. Model.* **2006**, *25*, 481.
- [40] D. C. Bas, D. M. Rogers, J. H. Jensen, *Proteins Struct. Funct. Genet.* **2008**, *73*, 765.
- [41] P. B. Reis, D. Vila-Vicosa, W. Rocchia, M. Machuqueiro, *J. Chem. Inf. Model.* **2020**, *60*, 4442.
- [42] T. J. Dolinsky, J. E. Nielsen, J. A. McCammon, N. A. Baker, *Nucleic Acids Res.* **2004**, *32*, W665.
- [43] T. J. Dolinsky, P. Czodrowski, H. Li, J. E. Nielsen, J. H. Jensen, G. Klebe, N. A. Baker, *Nucleic Acids Res.* **2007**, *35*, W522.
- [44] R. Anandakrishnan, B. Aguilar, A. V. Onufriev, *Nucleic Acids Res.* **2012**, *40*, 537.
- [45] M. H. Olsson, C. R. Søndergaard, M. Rostkowski, J. H. Jensen, *J. Chem. Theory Comput.* **2011**, *7*, 525.
- [46] A. D. Bochevarov, M. A. Watson, J. R. Greenwood, D. M. Philipp, *J. Chem. Theory Comput.* **2016**, *12*, 6001.
- [47] E. L. Ratkova, D. S. Palmer, M. V. Fedorov, *Chem. Rev.* **2015**, *115*, 6312.
- [48] T. Matsui, T. Baba, K. Kamiya, Y. Shigetani, *Phys. Chem. Chem. Phys.* **2012**, *14*, 4181.
- [49] P. Pernot, B. Civalleri, D. Presti, A. Savin, *J. Phys. Chem. A* **2015**, *119*, 5288.
- [50] J. Proppe, M. Reiher, *J. Chem. Theory Comput.* **2017**, *13*, 3297.
- [51] G. N. Simm, M. Reiher, *J. Chem. Theory Comput.* **2016**, *12*, 2762.
- [52] J. Proppe, M. Reiher, *J. Chem. Theory Comput.* **2019**, *15*, 357.
- [53] M. Reiher, *Isr. J. Chem.* **2022**, *62*, 1.
- [54] G. N. Simm, M. Reiher, *J. Chem. Theory Comput.* **2018**, *14*, 5238.
- [55] J. Proppe, S. Gugler, M. Reiher, *J. Chem. Theory Comput.* **2019**, *15*, 6046.
- [56] V. L. Deringer, A. P. Bartók, N. Bernstein, D. M. Wilkins, M. Ceriotti, G. Csányi, *Chem. Rev.* **2021**, *121*, 10073.
- [57] S. Gugler, M. Reiher, *J. Chem. Theory Comput.* **2022**, *18*, 6670.
- [58] F. Neese, *Wiley Int. Rev. Comput. Mol. Sci.* **2012**, *2*, 73.
- [59] F. Neese, F. Wennmohs, U. Becker, C. Riplinger, *J. Chem. Phys.* **2020**, *152*, 224108.
- [60] J. P. Perdew, K. Burke, M. Ernzerhof, *Phys. Rev. Lett.* **1996**, *77*, 3865.
- [61] M. Ernzerhof, G. E. Scuseria, *J. Chem. Phys.* **1999**, *110*, 5029.
- [62] C. Adamo, V. Barone, *J. Chem. Phys.* **1999**, *110*, 6158.
- [63] A. Schäfer, C. Huber, R. Ahlrichs, *J. Chem. Phys.* **1994**, *100*, 5829.
- [64] K. Eichkorn, O. Treutler, H. Öhm, M. Häser, R. Ahlrichs, *Chem. Phys. Lett.* **1995**, *240*, 283.
- [65] K. Eichkorn, F. Weigend, O. Treutler, R. Ahlrichs, *Theor. Chem. Accounts* **1997**, *97*, 119.
- [66] S. Grimme, J. Antony, S. Ehrlich, H. Krieg, *J. Chem. Phys.* **2010**, *132*, 154104.
- [67] S. Grimme, S. Ehrlich, L. Goerigk, *J. Comput. Chem.* **2011**, *32*, 1456.
- [68] A. V. Marenich, C. J. Cramer, D. G. Truhlar, *J. Phys. Chem. B* **2009**, *113*, 6378.
- [69] J. G. Sobez, M. Reiher, *J. Chem. Inf. Model.* **2020**, *60*, 3884.
- [70] M. Bensberg, S. A. Grimmel, J.-G. Sobez, M. Steiner, J. P. Unsleber, M. Reiher, qcscine/Molassembler: Release 2.0.0. **2023**. <https://zenodo.org/record/7928074>
- [71] A. Baiardi, M. Bensberg, F. Bosia, C. Brunken, K.-S. Csizi, R. Feldmann, N. Glaser, S. A. Grimmel, S. Gugler, M. P. Haag, M. A. Heuer, C. H. Müller, S. Polonius, G. N. Simm, J.-G. Sobez, M. Steiner, P. L. Türtcher, J. P. Unsleber, A. C. Vaucher, T. Weymuth, M. Reiher, qcscine/Utilities: Release 8.0.0. **2023**. <https://zenodo.org/record/7928050>
- [72] The PyMOL Molecular Graphics System, Version 2.0 Schrödinger, L.–The PyMOL Molecular Graphics System, Version 2.5.2. **2015**. <https://pymol.org>
- [73] M. Bensberg, S. A. Grimmel, J.-G. Sobez, M. Steiner, J. P. Unsleber, M. Reiher, qcscine/molassembler: Release 1.2.1. **2021**. <https://zenodo.org/record/10159517>
- [74] P. J. Artymiuk, A. R. Poirrette, H. M. Grindley, D. W. Rice, P. Willett, *J. Mol. Biol.* **1994**, *243*, 327.
- [75] R. V. Spriggs, P. J. Artymiuk, P. Willett, *J. Chem. Inf. Comput. Sci.* **2003**, *43*, 412.
- [76] G. N. Simm, P. L. Türtcher, M. Reiher, *J. Comput. Chem.* **2020**, *41*, 1144.
- [77] Y. Xue, M. Ökvist, Ö. Hansson, S. Young, *Protein Sci.* **1998**, *7*, 2099.
- [78] IUPAC-IUB Commission on Biochemical Nomenclature, *J. Biol. Chem.* **1970**, *245*, 6489.
- [79] N. M. O'Boyle, M. Banck, C. A. James, C. Morley, T. Vandermeersch, G. R. Hutchison, *J. Chem.* **2011**, *3*, 33.
- [80] M. Bensberg, S. A. Grimmel, J.-G. Sobez, M. Steiner, J. P. Unsleber, M. Reiher, qcscine/database: Release 1.0.0. **2022**. <https://zenodo.org/records/6695496>
- [81] X. Chang, A. M. M. Jørgensen, P. Bardrum, J. J. Led, *Biochemistry* **1997**, *36*, 9409.
- [82] S. Karlin, Z.-Y. Zhu, *Proc. Natl. Acad. Sci.* **1997**, *94*, 14231.
- [83] B. Rost, C. Sander, *J. Mol. Biol.* **1993**, *232*, 584.
- [84] J. A. Cuff, G. J. Barton, *Proteins Struct. Funct. Genet.* **1999**, *34*, 508.
- [85] B. Rost, V. A. Eylich, *Proteins Struct. Funct. Genet.* **2001**, *45*, 192.
- [86] J. A. Maier, C. Martinez, K. Kasavajhala, L. Wickstrom, K. E. Hauser, C. Simmerling, *J. Chem. Theory Comput.* **2015**, *11*, 3696.
- [87] B. Thapa, H. B. Schlegel, *J. Phys. Chem. A* **2016**, *120*, 5726.
- [88] M. Schmidt am Busch, E. W. Knapp, *ChemPhysChem* **2004**, *5*, 1513.
- [89] I. Juranić, *Croat. Chem. Acta* **2014**, *87*, 343.
- [90] T. M. Beale, M. G. Chudzinski, M. G. Sarwar, M. S. Taylor, *Chem. Soc. Rev.* **2013**, *42*, 1667.
- [91] J. C. Kromann, F. Larsen, H. Moustafa, J. H. Jensen, *PeerJ* **2016**, *2016*, e2335.
- [92] J. H. Jensen, C. J. Swain, L. Olsen, *J. Phys. Chem. A* **2017**, *121*, 699.
- [93] P. Hunt, L. Hosseini-Gerami, T. Chrien, J. Plante, D. J. Ponting, M. Segall, *J. Chem. Inf. Model.* **2020**, *60*, 2989.
- [94] C. A. Fitch, D. A. Karp, K. K. Lee, W. E. Stites, E. E. Lattman, E. B. Garcia-Moreno, *Biophys. J.* **2002**, *82*, 3289.
- [95] R. L. Thurkill, G. R. Grimsley, J. M. Scholtz, C. N. Pace, *Protein Sci.* **2006**, *15*, 1214.
- [96] C. Martin, V. Richard, M. Salem, R. Hartley, Y. Mauguen, *Acta Crystallogr. Sect. D* **1999**, *55*, 386.
- [97] M. Oliveberg, V. L. Arcus, A. R. Fersht, *Biochemistry* **1995**, *34*, 9424.
- [98] R. Loewenthal, J. Sancho, A. R. Fersht, *Biochemistry* **1991**, *30*, 6775.
- [99] S. Pahari, L. Sun, E. Alexov, PKAD: a database of experimentally measured pKa values of ionizable groups in proteins. **2019**.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: K.-S. Csizi, M. Reiher, *J. Comput. Chem.* **2024**, *45*(11), 761. <https://doi.org/10.1002/jcc.27276>