

Multiscale Biomolecular Simulations in the Exascale Era

David Carrasco-Busturia^{a,1}, Emiliano Ippoliti^b, Simone Meloni^c, Ursula Rothlisberger^d, Jógvan Magnus Haugaard Olsen^{a,*}

^a*DTU Chemistry, Technical University of Denmark (DTU), Kongens Lyngby, DK-2800, Denmark*

^b*Computational Biomedicine, Institute for Neuroscience and Medicine INM-9, Forschungszentrum Jülich GmbH, Jülich, DE-52428, Germany*

^c*Dipartimento di Scienze Chimiche, Farmaceutiche ed Agrarie (DOCPAS), Università degli Studi di Ferrara (Unife), Ferrara, I-44121, Italy*

^d*Laboratory of Computational Chemistry and Biochemistry, Institute of Chemical Sciences and Engineering, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, CH-1015, Switzerland*

Abstract

The complexity of biological systems and processes, spanning molecular to macroscopic scales, necessitates the use of multiscale simulations to get a comprehensive understanding. Quantum mechanics/molecular mechanics (QM/MM) molecular dynamics (MD) simulations are crucial for capturing processes beyond the reach of classical MD simulations. The advent of exascale computing offers unprecedented opportunities for scientific exploration, not least within life sciences, where simulations are essential to unravel intricate molecular mechanisms underlying biological processes. However, leveraging the immense computational power of exascale computing requires innovative algorithms and software designs. In this context, we discuss the current status and future prospects of multiscale biomolecular simulations on exascale supercomputers with a focus on QM/MM MD. We highlight our own efforts in developing a versatile and high-performance multiscale simulation framework with the aim of efficient utilization of state-of-the-art supercomputers. We showcase its application in uncovering complex biological

*Corresponding author

Email address: jmho@kemi.dtu.dk (Jógvan Magnus Haugaard Olsen)

¹Present address: Division of Theoretical Chemistry and Biology, School of Engineering Sciences in Chemistry, Biotechnology and Health, KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden

Introduction

Biological processes extend across wide scales in space and time due to the hierarchical organization of biological matter [1]. The characteristic dimensions span from the molecular scale of a few ångströms with ultrafast electronic processes on the order of atto- and femtoseconds and rapid chemical reactions that occur within pico- to microseconds, to the macroscopic scale of cells and organs that are visible to the naked eye and where processes extend to seconds and even days and years. The complexity in living organisms largely stems from this hierarchical structure, where a local process may trigger a cascade of events across multiple spatial and temporal scales. Thus, multiscale approaches integrating different resolutions and methodologies are essential for capturing the entire spectrum of biological events [2, 3].

The continuous development of multiscale methods in computational structural biology is driven by simultaneous advancements in algorithms, software implementations, and hardware technology that push the boundaries of molecular simulations in terms of accessible time scales and system sizes, capturing biological systems from the atomistic to the cellular level. In particular, quantum mechanics/molecular mechanics (QM/MM) molecular dynamics (MD) simulations have become increasingly important in the last decades for studying processes where a description of electronic degrees of freedom is paramount and thus beyond the capabilities of classical MD based on standard analytical force fields. Examples of such events encompass all types of chemical reactions, including proton-coupled electron transfers [4, 5], photochemistry [6], and the manifold of chemical transformations observed in enzymatic reactions [7, 8], especially those involving transition metals [9]. A detailed characterization of these phenomena, including reactants, transition states, intermediates, and products, as well as the involved relative free energies, reaction rates, and binding affinities in both electronic ground and excited states, also has direct implications for drug design [10]. In QM/MM models of biomolecular systems, a smaller part (the QM subsystem), such as the active site of an enzyme or a chromophore embedded in a protein or lipid bilayer, is described at the quantum mechanical level, while the remainder (the MM subsystem) is modeled using molecular mechanics. This multiscale strategy balances a detailed and accurate but computationally

costly description of the essential part of a system with a coarser but computationally expedient approach for the much larger MM part. The dynamic aspect of QM/MM MD is crucial to accurately capture the behavior and function of complex proteins [11, 12]. However, QM/MM MD simulations have a substantially higher computational cost than classical (i.e., MM) MD simulations, severely limiting the accessible time scales. Typically, for a density functional theory (DFT)-based QM/MM MD simulation with around 100 atoms in the QM subsystem, accessible time scales are limited to a few hundred picoseconds [8, 13].

The advent of exascale computing marks a pivotal moment for all simulation-based scientific fields [14, 15]. This remarkable technological achievement was accomplished by connecting thousands of computing nodes through high-speed network interconnects. Each node combines traditional general-purpose central processing units (CPUs) with powerful graphics processing units (GPUs). However, exascale supercomputer architectures also introduce new challenges since programming software applications for these heterogeneous machines requires a judicious decomposition of the computational work to exploit each component of the machines optimally [16]. Therefore, to fully exploit the computational power of present and future heterogeneous HPC architectures, a high degree of concurrent parallelism is needed, where different parts of the supercomputer work on different subdomains of the computational model, each described within a different theoretical method. Together, exascale computing and the development of novel computational methods and software can enable longer and more accurate simulations on larger and more complex systems. This opens up new opportunities for discovery and innovation in the life and health sciences, potentially revolutionizing areas such as drug design and bioengineering.

The heterogeneous CPU/GPU technology highlighted above is taken to the next level with modular supercomputer architectures that integrate a variety of hardware technologies into interconnected partitions [17, 18]. A prime example of this is the LUMI supercomputer [19], along with the upcoming exascale JUPITER supercomputer [20], both procured by the European HPC Joint Undertaking (EuroHPC JU). Currently, LUMI consists of two primary *number-crunching* partitions: a general-purpose CPU partition and a high-performance GPU-accelerated partition. Additionally, it features an interactive data analytics partition and an accelerated storage partition. Looking ahead, the integration of prospective technologies like quantum and neuromorphic computing into traditional HPC infrastructure

is expected [17, 18]. Indeed, the procurement of a quantum computing partition for LUMI is already underway [21], and the Jülich Supercomputing Center in Germany is already experimenting with this kind of integration [22]. Modular supercomputer architectures offer substantial benefits, particularly in allowing calculations to run on the most suitable hardware for the specific problem. Moreover, existing software packages that have yet to be optimized for new hardware remain useful for solving important scientific problems. Fully exploiting the capabilities of complex modular supercomputers for exceptionally challenging scientific problems, which require the full computing power of the machine, necessitates a new computational paradigm. Here, we highlight an approach recently introduced in the field of multiscale biomolecular simulations.

Multiscale methods, especially the QM/MM approach, have proven indispensable for exploring complex biological phenomena. Interestingly, QM/MM is not only a robust technique in itself but also offers a route to overcome some of the existing limitations in the development of software for atomistic simulations on modern hybrid computer architectures [14]. Beyond QM/MM, multiscale methods can also be extended to incorporate multiple layers ranging from coarse-grained (CG) to continuum models of matter where parts of a system, e.g., membrane regions that are sufficiently distant from an embedded protein of interest, are treated by techniques usually employed for meso- and macroscopic systems [23–27].

To reach the full potential offered by exascale supercomputers, it is imperative that multiscale interfaces scale efficiently, fully leveraging the extensive network of CPUs and GPUs. That is one of the main objectives of MiMiC (multiscale modeling in computational chemistry), a high-performance and versatile framework for multiscale simulations [28]. The *program-agnostic* MiMiC framework is designed to combine virtually any QM and MM (or other) program without compromising the computational efficiency and scalability of the simulation. Indeed, MiMiC, although still in its nascent stages, has demonstrated its ability to efficiently scale QM/MM MD simulations across thousands of CPU-based computing nodes [13, 29]. Here we will showcase the MiMiC framework for biomolecular simulations, illustrating its potentiality in leveraging state-of-the-art supercomputing resources and its capability to address complex biochemical and biophysical problems.

Current Status of Multiscale Biomolecular Simulation Software

Multiscale QM/MM capabilities are often implemented by either extending dedicated QM or MM programs with the functionalities of the other [30–37] or by creating ad hoc interfaces between stand-alone QM and MM programs [38–41]. Additionally, some programs offer flexible interfaces that facilitate easy coupling with other programs [42–44]. Beyond these, there are integrative frameworks that do not contain inherent QM, MM, or similar functionalities, but instead depend entirely on other programs for these capabilities [45–59].

The interface between QM and MM components can be classified as tightly or loosely coupled, reflecting the degree of interdependence between the components. Loose coupling is preferred for its flexibility and ease of maintenance, enabling straightforward integration of independent programs. This approach supports the quick incorporation of new functionalities and advancements within individual programs without necessitating alterations to other coupled components. Crucially, loose coupling permits independent optimization of each program for peak performance. However, this flexibility can come at the cost of computational efficiency due to slower inter-program communication, particularly when compared to the tight coupling approach that allows for direct in-memory data sharing.

There are three communication mechanisms for exchanging data between programs, namely the file-, library-, and network-based approaches. Most commonly utilized for general interfaces and integrative frameworks is the file-based approach, which does not require modifications to the source code of the interfaced programs, making it the simplest to implement. Here, the main driver program generates input files for the interfaced programs, executes them, and subsequently reads their output files. However, it is notably the least efficient communication mechanism, primarily due to slow disk write/read operations and overhead associated with executing and terminating the interfaced programs. Both drawbacks are avoided in the library-based approach, where the interfaced programs are converted into a library that is linked to the main program, thus enabling efficient in-memory data exchange. The disadvantages are the rather intrusive modifications needed in the source codes of the interfaced programs, and the risk of ending up with a tight coupling focused on a few specific programs. Moreover, a critical aspect of multiscale implementations is parallel scalability and efficiency. The optimal parallel algorithms for QM and MM components, or even among dif-

ferent QM methods, may differ significantly and may not seamlessly integrate within the more tightly coupled environments.

The network-based approach, on the other hand, offers a solution that potentially circumvents the limitations inherent in both file- and library-based methods. It achieves a balance between flexibility and communication efficiency by enabling data exchange over a network. This facilitates seamless communication between programs running on different processors (CPUs and GPUs), across computing nodes, or even among supercomputer partitions. Such an approach retains the benefits of loose coupling, i.e., ease of integration and maintenance, while surpassing file-based approaches in performance through the use of fast network protocols and remote direct memory access (RDMA) technologies. Crucially, it can be implemented so as not to disrupt the normal execution of interfaced programs, thus preserving their performance. However, the network-based approach requires a more sophisticated design and management strategy to reduce latency and enhance throughput. Despite these challenges, with proper implementation, the network-based approach substantially improves the efficiency and scalability of multiscale simulations by facilitating high-speed communication between diverse computational models.

Toward Multiscale Biomolecular Simulations on Exascale Architectures

For multiscale simulations to run efficiently on state-of-the-art supercomputers, first and foremost, it is imperative that the programs dedicated to a specific methodology, such as QM and MM, are able to take advantage of the strengths of both CPUs and GPUs. Moreover, these programs must be capable of using the vast parallel processing capabilities of exascale architectures. Considerable efforts are underway to push QM- and MM-based programs towards exascale [60–67] with support from EuroHPC JU through its HPC Centres of Excellence [68] and the Exascale Computing Project (ECP) led by the US Department of Energy [69].

For a multiscale simulation framework to fully exploit the capabilities of these QM and MM programs, it must seamlessly integrate with their optimal parallelization strategies without introducing inefficiencies. This entails enabling the concurrent execution of interfaced programs, eliminating the need for their repeated startup and shutdown, and reducing communication overhead to a minimum. Crucially, the calculation of subsystem interactions,

such as those between QM and MM particles, needs to be highly parallel and efficient to prevent it from becoming a computational bottleneck. Implementing effective and automatic load balancing is also critical to ensure that computational resources are utilized optimally, thereby maximizing the throughput of individual simulations.

The versatile and high-performance MiMiC framework for multiscale simulations was designed with the aim of addressing the criteria outlined above [28]. The strategy used by the MiMiC framework is illustrated in Figure 1. On one side, MiMiC connects to a simulation driver that manages the overall simulation process, including the integration of the equations of motion and the maintenance of temperature and pressure. On the other side, it interfaces multiple external programs using a client-server approach combined with a multiple-program multiple-data (MPMD) model. Each external program is tasked with calculations belonging to a given subsystem, while MiMiC calculates subsystem interactions. Importantly, the external programs run concurrently on separate computational resources using their own optimal parallelization capabilities. Communication between MiMiC and the external programs is facilitated by the MiMiC Communication Library (MCL), a dedicated lightweight library that ensures efficient network-based communication and simplifies the interfacing with external programs. An illustration of a MiMiC-based simulation workflow is shown in Figure 2. The scope of the MiMiC framework is broad, extending beyond QM/MM. This includes QM/QM, which integrates different levels of QM theory, QM/MM/CG, and even models that incorporate machine-learning (ML) techniques like ML/MM or QM/ML.

The MiMiC framework has been used to implement a DFT-based electrostatic embedding QM/MM method, employing the CPMD program [70] as both the MD driver and QM engine, and GROMACS [61] as the MM engine [28, 29]. In electrostatic embedding, the (external) electric field from the point charges in the MM subsystem is included in the Hamiltonian of the QM




Figure 1: Illustration of the strategy used by the MiMiC framework

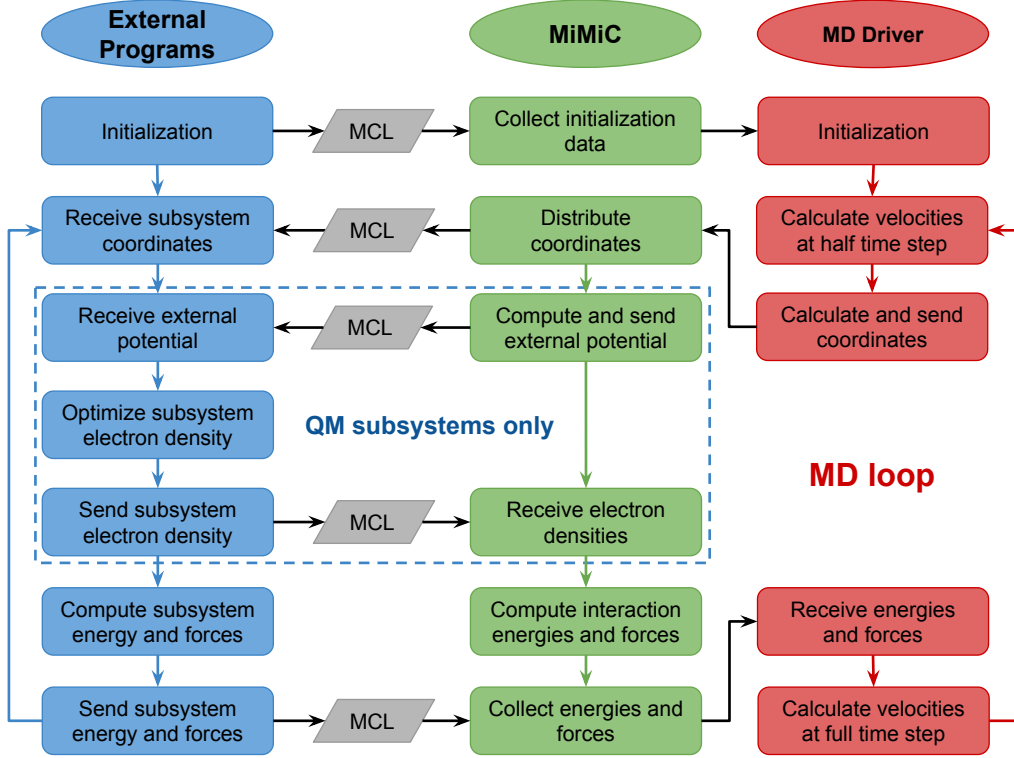


Figure 2: Illustration of MiMiC-based simulation workflow

subsystem, thus directly polarizing its electronic density. The calculation of electrostatic QM/MM interactions is based on a dense grid representation of the electronic density, which is computationally expensive, especially for large systems, because the number of integrals over the electron density scales linearly with the number of atoms in the MM subsystem. The MiMiC framework has implemented an efficient approach that substantially speeds up the calculation essentially without compromising the accuracy of the forces, thus reducing the computational cost by about 80 % for small systems (e.g., small solvated proteins) and up to 99 % for larger systems (e.g., membrane-embedded proteins) [28]. Furthermore, MiMiC implements a hybrid shared- and distributed-memory (OpenMP/MPI) parallelization strategy to ensure that these calculations remain efficient and do not hinder performance, even under highly parallel conditions [29]. A powerful example of this is shown in Figure 3, where the extreme scalability of the CPMD program is harnessed to




Figure 3: **Parallel scalability and efficiency of MiMiC-based QM/MM MD simulations of IDH1.** The IDH1 system has a total of 130,828 atoms with 142 QM atoms [13]. The simulations were run on the CPU partition of JUWELS [71]. The speedup is given in terms of the time per MD step normalized against a reference run using seven nodes. Adapted from Raghavan et al. [13], licensed under CC BY 4.0.

achieve strong scalability beyond 80,000 CPU cores with a parallel efficiency of 70 % for a single MiMiC-based QM/MM MD simulation of the human isocitrate dehydrogenase-1 (IDH1) [13]. Table 1 shows examples of the computational performance and cost of MiMiC-based QM/MM MD simulations using CPMD and GROMACS. It is clear that the simulation throughput depends first and foremost on the chosen exchange-correlation functional and the size of the QM subsystem. This reflects the fact that the QM calculation is by far the most computationally demanding part of a QM/MM MD simulation. Hybrid exchange-correlation functionals are particularly expensive in a plane-wave basis set such as the one employed by the CPMD program. Still, due to the excellent parallelism in CPMD, the simulation throughput can be pushed to 4.8 ps/day for the small QM region (46 atoms) and 0.7 ps/day for the larger QM subsystem (142 atoms). Using instead a non-hybrid functional pushes the performance to 21 and 5.4 ps/day for the small and large QM subsystems, respectively.

Biological Applications

The MiMiC framework has enabled a number of recent computational studies that demonstrate its utility and efficiency across a broad spectrum of systems with biological relevance. In this section, we showcase select examples that highlight the impactful contributions of MiMiC-based QM/MM MD simulations in advancing our understanding of complex biological processes.

Among the pioneering uses of MiMiC-based QM/MM MD simulations

Table 1: **Computational performance and cost of MiMiC-based QM/MM MD simulations.** All simulations were run with a 0.5 fs timestep on the CPU partition of JUWELS [71] at the scaling limit (parallel efficiency $\geq 70\%$). The systems are p38 α mitogen-activated protein kinase (169,550 atoms) and human isocitrate dehydrogenase-1 (130,828 atoms). We refer to the original work for full computational details [13].

System	p38 α		IDH1	
QM atoms	46		142	
XC functional	BLYP	B3LYP	BLYP	B3LYP
Throughput (ps/day)	21	4.8	5.4	0.7
Cost (node-hours/ps)	9	1280	480	60480

were studies of CLC proteins, a large family of anion channels and transporters. Chiariello et al. studied the molecular mechanism of fluoride inhibition of the anion/proton exchanger CLC-ec1 from *E. coli* (Fig. 4A) [72]. The use of QM/MM for this study was mandatory as ion translocation involves proton transfer processes. On the basis of QM/MM MD and well-tempered metadynamics (wtMTD) simulations at the B3LYP and BLYP levels of theory, Chiariello et al. were able to report proton affinities of the fluoride ion and the gating glutamate residue E148, thus providing valuable insights into transport inhibition. In a second study, the mechanisms of proton transfer and release by the fluoride/proton antiporter CLC^F-eca were investigated (Fig. 4B) [73]. Employing the same simulation techniques, it could be shown that a triad is formed between fluoride, glutamate E318, and the gating glutamate E118, eventually releasing protons and fluoride as hydrogen fluoride.

In a recent study, the ligand iperoxo (routinely used in neuroimaging) targeting the human muscarinic acetylcholine receptor 2 was investigated [75]. The work focuses on the calculation of the drug unbinding rate constant k_{off} , a very difficult parameter to correctly estimate with force field-based approaches. In fact, while methods based on modern force fields are nowadays able to predict accurate binding free energies and affinities, this is often not the case for rate constants such as k_{off} , which also require a correct estimation of the free energy of transition states. The study shows how sensitive this estimation is to values of the partial charges of the ligand and that while results obtained with a QM/MM MD simulation are in good agreement with experimental findings, standard force field-based procedures lead to qualitatively wrong results, most likely due to the lack of explicit electronic polarization and charge transfer, which are included in QM/MM.




Figure 4: **Illustrations of biological systems that have been studied using the MiMiC framework.** A: ClC-ec1 anion/proton antiporter embedded in a solvated lipid bilayer with a total of 150,925 atoms (19 QM atoms) [72]. Reprinted with permission from Ref. [72]. Copyright 2020 American Chemical Society. B: CLC^F-eca fluoride/proton antiporter embedded in a solvated lipid bilayer with a total of around 174,000 atoms (36 QM atoms) [73]. Reprinted with permission from Ref. [73]. Copyright 2021 American Chemical Society. C: AMPAR cation channel embedded in a solvated lipid bilayer where the transmembrane domain was included in the simulations (22 QM atoms) [74]. Figure by Schackert et al. [74], licensed under CC BY 4.0.

Another important application of QM/MM MD simulations is in scenarios where parts of a system cannot be adequately described using simplified an-

alytical force fields. Notable examples are metal ions, in particular transition metals or divalent alkaline-earth ions. The absence of accurate parameters for the latter hinders a thorough understanding of many biological processes related to, e.g., ion channels, transporters, and pumps. Recently, Schackert et al. studied the mechanism of calcium permeation in α -amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid receptors (AMPA receptors), which are key to rapid synaptic transmission in the central nervous system (Fig. 4C) [74]. In that study, the calcium-binding sites within the channel, initially identified through classical MD, were further confirmed using QM/MM MD. This step was crucial for validating a newly developed classical force field specifically designed for calcium ions. Interestingly, they observed charge transfer between the calcium ion and the water molecules in the first solvation shell, which underlines the importance of a QM/MM description for such systems.

Outlook

The exascale era is poised to radically change life sciences. With the ability to conduct longer and more accurate simulations on larger and more complex molecular systems than ever before, we will be able to tackle scientific questions that are currently beyond our reach, deepening our understanding of biological processes. This opens up new opportunities for discovery and innovation, potentially revolutionizing fields such as drug design and bioengineering. However, fully realizing the potential of exascale computing is contingent upon overcoming substantial challenges in software design, algorithm optimization, and the efficient utilization of modular and heterogeneous HPC architectures.

The versatility of the MiMiC framework makes it ideal for pushing multiscale biomolecular simulations towards the exascale. It is capable of exploiting the results of the large-scale initiatives by EuroHPC JU and ECP for individual domains such as QM and MM. Indeed, work is already underway to couple a diverse set of QM programs to MiMiC, namely, Quantum ESPRESSO, CP2K, and DFT-FE, all of which are being developed to exploit state-of-the-art HPC technology [67]. Additionally, it is likely that MD simulations will benefit greatly from future ML-based force fields [76, 77], which MiMiC is fully able to integrate into advanced simulation workflows [78]. For instance, hybrid ML/MM models facilitate long simulations of biological systems using ML-based force fields, achieving near QM/MM precision at a substantially reduced computational cost [79], while ML-enhanced free-

energy methods significantly accelerate QM/MM calculations of ligand binding affinities for drug discovery [80]. These and other novel methods are expected to aid in accurately predicting key biophysical properties like drug-protein binding free energies and complete free-energy profiles. We anticipate that the integration of MiMiC-based multiscale simulations with ML and enhanced sampling techniques will further catalyze a *quantum leap* in the fundamental atomistic understanding of biological processes.

Acknowledgments

JMHO thanks VILLUM FONDEN (Grant No. VIL29478) for financial support. UR gratefully acknowledges funding from the Swiss National Science Foundation via the NCCR MUST and the individual grants 200020-185092 and 200020-219440, as well as computing time from the Swiss National Computing Centre CSCS. EI thanks the Helmholtz European Partnering program (“Innovative high-performance computing approaches for molecular neuromedicine”) for funding.

References

- [1] L. V. Schaffer, T. Ideker, Mapping the multiscale structure of biological systems, *Cell Syst.* 12 (2021) 622–635. doi:10.1016/j.cels.2021.05.012.
- [2] G. Palermo, A. M. J. J. Bonvin, M. Dal Peraro, R. E. Amaro, V. Tozzini, Multiscale Modeling From Macromolecules to Cell: Opportunities and Challenges of Biomolecular Simulations, *Front. Mol. Biosci.* 7 (2020). doi:10.3389/fmolb.2020.00194.
- [3] F. Trovato, J. Trylska, P. J. Bond, P. G. Wolynes, Combining Simulations, Theory, and Experiments into Multiscale Models of Biological Events, *Front. Mol. Biosci.* 8 (2021). doi:10.3389/fmolb.2021.797754.
- [4] S. Hammes-Schiffer, Proton-Coupled Electron Transfer: Moving Together and Charging Forward, *J. Am. Chem. Soc.* 137 (2015) 8860–8871. doi:10.1021/jacs.5b04087.
- [5] M. Nottoli, L. Cupellini, F. Lipparini, G. Granucci, B. Mennucci, Multiscale Models for Light-Driven Processes, *Annu. Rev. Phys. Chem.* 72 (2021) 489–513. doi:10.1146/annurev-physchem-090419-104031.
- [6] J. M. Toldo, M. T. do Casal, E. Ventura, S. A. do Monte, M. Barbatti, Surface hopping modeling of charge and energy transfer in active environments, *Phys. Chem. Chem. Phys.* 25 (2023) 8293–8316. doi:10.1039/d3cp00247k.
- [7] S. F. Sousa, A. J. M. Ribeiro, R. P. P. Neves, N. F. Brás, N. M. F. S. A. Cerqueira, P. A. Fernandes, M. J. Ramos, Application of quantum mechanics/molecular mechanics methods in the study of enzymatic reaction mechanisms, *Wiley Interdiscip. Rev. Comput. Mol. Sci.* 7 (2017) e1281. doi:10.1002/wcms.1281.
- [8] V. Vennelakanti, A. Nazemi, R. Mehmood, A. H. Steeves, H. J. Kulik, Harder, better, faster, stronger: Large-scale QM and QM/MM for predictive modeling in enzymes and proteins, *Curr. Opin. Struct. Biol.* 72 (2022) 9–17. doi:10.1016/j.sbi.2021.07.004.
- [9] C. E. Tzeliou, M. A. Mermigki, D. Tzeli, Review on the QM/MM Methodologies and Their Application to Metalloproteins, *Molecules* 27 (2022) 2660. doi:10.3390/molecules27092660.

- [10] M. W. van der Kamp, J. Begum, QM / MM for Structure-Based Drug Design: Techniques and Applications, John Wiley & Sons, Ltd, 2024, pp. 119–156. doi:10.1002/9783527840748.ch6.
- [11] E. Brunk, U. Rothlisberger, Mixed Quantum Mechanical/Molecular Mechanical Molecular Dynamics Simulations of Biological Systems in Ground and Electronically Excited States, *Chem. Rev.* 115 (2015) 6217–6263. doi:10.1021/cr500628b.
- [12] C. M. Clemente, L. Capece, M. A. Martí, Best Practices on QM/MM Simulations of Biological Systems, *J. Chem. Inf. Model.* 63 (2023) 2609–2627. doi:10.1021/acs.jcim.2c01522.
- [13] B. Raghavan, M. Paulikat, K. Ahmad, L. Callea, A. Rizzi, E. Ippoliti, D. Mandelli, L. Bonati, M. De Vivo, P. Carloni, Drug Design in the Exascale Era: A Perspective from Massively Parallel QM/MM Simulations, *J. Chem. Inf. Model.* 63 (2023) 3647–3658. doi:10.1021/acs.jcim.3c00557.
- [14] A. Rimmel, What exascale computing could mean for chemistry, *C&EN Global Enterprise* 100 (2022). URL: <https://cen.acs.org/physical-chemistry/computational-chemistry/exascale-comput>
- [15] C. Q. Choi, The Beating Heart of the World’s First Exascale Supercomputer, *IEEE Spectrum* (2022). URL: <https://spectrum.ieee.org/frontier-exascale-supercomputer>.
- [16] R. Di Felice, M. L. Mayes, R. M. Richard, D. B. Williams-Young, G. K.-L. Chan, W. A. de Jong, N. Govind, M. Head-Gordon, M. R. Hermes, K. Kowalski, X. Li, H. Lischka, K. T. Mueller, E. Mutlu, A. M. N. Niklasson, M. R. Pederson, B. Peng, R. Shepard, E. F. Valeev, M. van Schilfgaarde, B. Vlasisavljevich, T. L. Windus, S. S. Xantheas, X. Zhang, P. M. Zimmerman, A Perspective on Sustainable Computational Chemistry Software Development and Integration, *J. Chem. Theory Comput.* 19 (2023) 7056–7076. doi:10.1021/acs.jctc.3c00419.
- [17] P. Carpenter, U.-H. Utz, S. Narasimhamurthy, E. Suarez, Heterogeneous High Performance Computing, <https://doi.org/10.5281/zenodo.6090425>, 2022. doi:10.5281/zenodo.6090425.

- [18] E. Suarez, N. Eicker, T. Moschny, S. Pickartz, C. Clauss, V. Plugaru, A. Hertel, K. Michielsen, T. Lippert, Modular Supercomputing Architecture, <https://doi.org/10.5281/zenodo.6508394>, 2022. doi:10.5281/zenodo.6508394.
- [19] LUMI, <https://www.lumi-supercomputer.eu/>, 2024. Date accessed: 2024-02-11.
- [20] JUPITER - Exascale for Europe, <https://www.fz-juelich.de/en/ias/jsc/jupiter>, 2024. Date accessed: 2024-02-11.
- [21] Czechia will host the European LUMI-Q quantum computer, <https://www.it4i.cz/en/about/infoservice/press-releases/czechia-will-host-the-european-lumi-q-quantum-computer>, 2022. Date accessed: 2024-02-11.
- [22] JUNIQ - Jülich UNified Infrastructure for Quantum computing, <https://www.fz-juelich.de/en/ias/jsc/systems/quantum-computing/juniq-facility/juniq>, 2024. Date accessed: 2024-02-20.
- [23] J. Jin, A. J. Pak, A. E. P. Durumeric, T. D. Loose, G. A. Voth, Bottom-up Coarse-Graining: Principles and Perspectives, *J. Chem. Theory Comput.* 18 (2022) 5759–5791. doi:10.1021/acs.jctc.2c00643.
- [24] Y. Guo, M. R. K. Mofrad, A. B. Tepole, On modeling the multiscale mechanobiology of soft tissues: Challenges and progress, *Biophys. Rev.* 3 (2022). doi:10.1063/5.0085025.
- [25] L. V. Bock, S. Gabrielli, M. H. Kolář, H. Grubmüller, Simulation of Complex Biomolecular Systems: The Ribosome Challenge, *Annu. Rev. Biophys.* 52 (2023) 361–390. doi:10.1146/annurev-biophys-111622-091147.
- [26] L. Borges-Araújo, I. Patmanidis, A. P. Singh, L. H. S. Santos, A. K. Sieradzian, S. Vanni, C. Czaplewski, S. Pantano, W. Shinoda, L. Monticelli, A. Liwo, S. J. Marrink, P. C. T. Souza, Pragmatic Coarse-Graining of Proteins: Models and Applications, *J. Chem. Theory Comput.* 19 (2023) 7112–7135. URL: <http://dx.doi.org/10.1021/acs.jctc.3c00733>. doi:10.1021/acs.jctc.3c00733.

- [27] E. Fadda, O. Makshakova, S. Perez, Understanding glycobiology through multiscale molecular dynamics simulations: From basic principles to case studies, Elsevier, 2024, p. 379–396. doi:10.1016/b978-0-12-819655-7.00006-0.
- [28] J. M. H. Olsen, V. Bolnykh, S. Meloni, E. Ippoliti, M. P. Bircher, P. Carloni, U. Rothlisberger, MiMiC: a novel framework for multiscale modeling in computational chemistry, *J. Chem. Theory Comput.* 15 (2019) 3810–3823. doi:10.1021/acs.jctc.9b00093.
- [29] V. Bolnykh, J. M. H. Olsen, S. Meloni, M. P. Bircher, E. Ippoliti, P. Carloni, U. Rothlisberger, Extreme Scalability of DFT-Based QM/MM MD Simulations Using MiMiC, *J. Chem. Theory Comput.* 15 (2019) 5601–5613. doi:10.1021/acs.jctc.9b00424.
- [30] A. Laio, J. VandeVondele, U. Rothlisberger, A Hamiltonian electrostatic coupling scheme for hybrid Car–Parrinello molecular dynamics simulations, *J. Chem. Phys.* 116 (2002) 6941–6947. doi:10.1063/1.1462041.
- [31] T. Vreven, K. Morokuma, O. Farkas, H. B. Schlegel, M. J. Frisch, Geometry optimization with QM/MM, ONIOM, and other combined methods. I. Microiterations and constraints, *J. Comput. Chem.* 24 (2003) 760–769. doi:10.1002/jcc.10156.
- [32] T. Laino, F. Mohamed, A. Laio, M. Parrinello, An Efficient Real Space Multigrid QM/MM Electrostatic Coupling, *J. Chem. Theory Comput.* 1 (2005) 1176–1184. doi:10.1021/ct050123f.
- [33] B. R. Brooks, C. L. Brooks III, A. D. Mackerell Jr., L. Nilsson, R. J. Petrella, B. Roux, Y. Won, G. Archontis, C. Bartels, S. Boresch, A. Caffisch, L. Caves, Q. Cui, A. R. Dinner, M. Feig, S. Fischer, J. Gao, M. Hodoscek, W. Im, K. Kuczera, T. Lazaridis, J. Ma, V. Ovchinnikov, E. Paci, R. W. Pastor, C. B. Post, J. Z. Pu, M. Schaefer, B. Tidor, R. M. Venable, H. L. Woodcock, X. Wu, W. Yang, D. M. York, M. Karplus, CHARMM: The biomolecular simulation program, *J. Comput. Chem.* 30 (2009) 1545–1614. doi:10.1002/jcc.21287.
- [34] R. Salomon-Ferrer, D. A. Case, R. C. Walker, An overview of the amber biomolecular simulation package, *Wiley Interdiscip. Rev. Comput. Mol. Sci.* 3 (2013) 198–210. doi:10.1002/wcms.1121.

- [35] A. de la Lande, A. Alvarez-Ibarra, K. Hasnaoui, F. Cailliez, X. Wu, T. Mineva, J. Cuny, P. Calaminici, L. López-Sosa, G. Geudtner, I. Navizet, C. Garcia Iriepa, D. R. Salahub, A. M. Köster, Molecular Simulations with in-deMon2k QM/MM, a Tutorial-Review, *Molecules* 24 (2019). doi:10.3390/molecules24091653.
- [36] F. Neese, F. Wennmohs, U. Becker, C. Riplinger, The ORCA quantum chemistry program package, *J. Chem. Phys.* 152 (2020) 224108. doi:10.1063/5.0004608.
- [37] S. Seritan, C. Bannwarth, B. S. Fales, E. G. Hohenstein, C. M. Isborn, S. I. L. Kokkila-Schumacher, X. Li, F. Liu, N. Luehr, J. W. Snyder Jr., C. Song, A. V. Titov, I. S. Ufimtsev, L.-P. Wang, T. J. Martínez, TeraChem: A graphical processing unit-accelerated electronic structure package for large-scale ab initio molecular dynamics, *Wiley Interdiscip. Rev. Comput. Mol. Sci.* 11 (2021) e1494. doi:10.1002/wcms.1494.
- [38] C. M. Isborn, A. W. Götz, M. A. Clark, R. C. Walker, T. J. Martínez, Electronic Absorption Spectra from MM and ab Initio QM/MM Molecular Dynamics: Environmental Effects on the Absorption Spectrum of Photoactive Yellow Protein, *J. Chem. Theory Comput.* 8 (2012) 5092–5106. doi:10.1021/ct3006826.
- [39] J. Dziedzic, Y. Mao, Y. Shao, J. Ponder, T. Head-Gordon, M. Head-Gordon, C.-K. Skylaris, TINKTEP: A fully self-consistent, mutually polarizable QM/MM approach based on the AMOEBA force field, *J. Chem. Phys.* 145 (2016) 124106. doi:10.1063/1.4962909.
- [40] D. Loco, L. Lagardère, S. Caprasecca, F. Lipparini, B. Mennucci, J.-P. Piquemal, Hybrid QM/MM Molecular Dynamics with AMOEBA Polarizable Embedding, *J. Chem. Theory Comput.* 13 (2017) 4025–4033. doi:10.1021/acs.jctc.7b00572.
- [41] V. W. D. Cruzeiro, M. Manathunga, K. M. J. Merz, A. W. Götz, Open-Source Multi-GPU-Accelerated QM/MM Simulations with AMBER and QUICK, *J. Chem. Inf. Model.* 61 (2021) 2109–2115. doi:10.1021/acs.jcim.1c00169.
- [42] A. W. Götz, M. A. Clark, R. C. Walker, An extensible interface for QM/MM molecular dynamics simulations with AMBER, *J. Comput. Chem.* 35 (2014) 95–108. doi:10.1002/jcc.23444.

- [43] M. C. R. Melo, R. C. Bernardi, T. Rudack, M. Scheurer, C. Riplinger, J. C. Phillips, J. D. C. Maia, G. B. Rocha, J. V. Ribeiro, J. E. Stone, F. Neese, K. Schulten, Z. Luthey-Schulten, NAMD goes quantum: an integrative suite for hybrid simulations, *Nat. Methods* 15 (2018) 351–354. doi:10.1038/nmeth.4638.
- [44] V. W. D. Cruzeiro, Y. Wang, E. Pieri, E. G. Hohenstein, T. J. Martínez, TeraChem protocol buffers (TCPB): Accelerating QM and QM/MM simulations with a client–server model, *J. Chem. Phys.* 158 (2023) 044801. doi:10.1063/5.0130886.
- [45] H. L. Woodcock, B. T. Miller, M. Hodoscek, A. Okur, J. D. Larkin, J. W. Ponder, B. R. Brooks, MSCALE: A General Utility for Multiscale Modeling, *J. Chem. Theory Comput.* 7 (2011) 1208–1219. doi:10.1021/ct100738h.
- [46] C. Ma, L. Martin-Samos, S. Fabris, A. Laio, S. Piccinin, QM-MMW: A wrapper for QM/MM simulations with Quantum ESPRESSO and LAMMPS, *Comp. Phys. Commun.* 195 (2015) 191–198. doi:10.1016/j.cpc.2015.04.024.
- [47] J. Torras, E. Deumens, S. B. Trickey, Software Integration in Multi-scale Simulations: the PUPIL System, *J. Comput-Aided Mat. Des.* 13 (2006) 201–212. doi:10.1007/s10820-006-9011-3.
- [48] J. Torras, B. P. Roberts, G. M. Seabra, S. B. Trickey, Chapter One - PUPIL: A Software Integration System for Multi-Scale QM/MM-MD Simulations and Its Application to Biomolecular Systems, in: T. Karabenchewa-Christova (Ed.), *Combined Quantum Mechanical and Molecular Mechanical Modelling of Biomolecular Interactions*, volume 100 of *Adv. Protein Chem. Struct. Biol.*, 2015, pp. 1–31. doi:10.1016/bs.apcsb.2015.06.002.
- [49] J. Řezáč, Cuby: An integrative framework for computational chemistry, *J. Comput. Chem.* 37 (2016) 1230–1237. doi:10.1002/jcc.24312.
- [50] A. H. Larsen, J. J. Mortensen, J. Blomqvist, I. E. Castelli, R. Christensen, M. Dułak, J. Friis, M. N. Groves, B. Hammer, C. Hargus, E. D. Hermes, P. C. Jennings, P. B. Jensen, J. Kermode, J. R. Kitchin,

- E. L. Kolsbjerg, J. Kubal, K. Kaasbjerg, S. Lysgaard, J. B. Maronsson, T. Maxson, T. Olsen, L. Pastewka, A. Peterson, C. Rostgaard, J. Schiøtz, O. Schütt, M. Strange, K. S. Thygesen, T. Vegge, L. Vilhelmsen, M. Walter, Z. Zeng, K. W. Jacobsen, The atomic simulation environment—a Python library for working with atoms, *J. Phys.: Condens. Matter* 29 (2017) 273002. doi:10.1088/1361-648X/aa680e.
- [51] O. Weingart, A. Nenov, P. Altoè, I. Rivalta, J. Segarra-Martí, I. Dokukina, M. Garavelli, COBRAMM 2.0 — A software interface for tailoring molecular electronic structure calculations and running nanoscale (QM/MM) simulations, *J. Mol. Model.* 24 (2018). doi:10.1007/s00894-018-3769-6.
- [52] S. Metz, J. Kästner, A. A. Sokol, T. W. Keal, P. Sherwood, ChemShell—a modular software package for QM/MM simulations, *Wiley Interdiscip. Rev. Comput. Mol. Sci.* 4 (2014) 101–110. doi:10.1002/wcms.1163.
- [53] Y. Lu, M. R. Farrow, P. Fayon, A. J. Logsdail, A. A. Sokol, C. R. A. Catlow, P. Sherwood, T. W. Keal, Open-Source, Python-Based Redevelopment of the ChemShell Multiscale QM/MM Environment, *J. Chem. Theory Comput.* 15 (2019) 1317–1328. doi:10.1021/acs.jctc.8b01036.
- [54] E. G. Kratz, A. R. Walker, L. Lagardère, F. Lipparini, J.-P. Piquemal, G. Andrés Cisneros, LICHEM: A QM/MM program for simulations with multipolar and polarizable force fields, *J. Comput. Chem.* 37 (2016) 1019–1029. doi:10.1002/jcc.24295.
- [55] H. Gökcan, E. A. Vázquez-Montelongo, G. A. Cisneros, LICHEM 1.1: Recent Improvements and New Capabilities, *J. Chem. Theory Comput.* 15 (2019) 3056–3065. doi:10.1021/acs.jctc.9b00028.
- [56] B. Zhang, D. Altarawy, T. Barnes, J. M. Turney, H. F. I. Schaefer, Janus: An Extensible Open-Source Software Package for Adaptive QM/MM Methods, *J. Chem. Theory Comput.* 15 (2019) 4362–4373. doi:10.1021/acs.jctc.9b00182.
- [57] T. A. Barnes, E. Marin-Rimoldi, S. Ellis, T. D. Crawford, The MolSSI Driver Interface Project: A framework for standardized, on-the-fly interoperability between computational molecular sciences codes, *Comp. Phys. Commun.* 261 (2021) 107688. doi:10.1016/j.cpc.2020.107688.

- [58] S. Martí, QMCube (QM3): An all-purpose suite for multiscale QM/MM calculations, *J. Comput. Chem.* 42 (2021) 447–457. doi:10.1002/jcc.26465.
- [59] H. Lin, Y. Zhang, S. Pezeshki, A. W. Duster, B. Wang, X.-P. Wu, S.-W. Zheng, L. Gagliardi, D. G. Truhlar, QMMM 2023: A program for combined quantum mechanical and molecular mechanical modeling and simulations, *Comp. Phys. Commun.* 295 (2024) 108987. doi:10.1016/j.cpc.2023.108987.
- [60] J. C. Phillips, D. J. Hardy, J. D. C. Maia, J. E. Stone, J. V. Ribeiro, R. C. Bernardi, R. Buch, G. Fiorin, J. Hénin, W. Jiang, R. McGreevy, M. C. R. Melo, B. K. Radak, R. D. Skeel, A. Singharoy, Y. Wang, B. Roux, A. Aksimentiev, Z. Luthey-Schulten, L. V. Kalé, K. Schulten, C. Chipot, E. Tajkhorshid, Scalable molecular dynamics on CPU and GPU architectures with NAMD, *J. Chem. Phys.* 153 (2020). doi:10.1063/5.0014475.
- [61] S. Páll, A. Zhmurov, P. Bauer, M. Abraham, M. Lundborg, A. Gray, B. Hess, E. Lindahl, Heterogeneous parallelization and acceleration of molecular dynamics simulations in GROMACS, *J. Chem. Phys.* 153 (2020). doi:10.1063/5.0018516.
- [62] K. Kowalski, R. Bair, N. P. Bauman, J. S. Boschen, E. J. Bylaska, J. Daily, W. A. de Jong, T. Dunning, N. Govind, R. J. Harrison, M. Keçeli, K. Keipert, S. Krishnamoorthy, S. Kumar, E. Mutlu, B. Palmer, A. Panyala, B. Peng, R. M. Richard, T. P. Straatsma, P. Sushko, E. F. Valeev, M. Valiev, H. J. J. van Dam, J. M. Waldrop, D. B. Williams-Young, C. Yang, M. Zalewski, T. L. Windus, From NWChem to NWChemEx: Evolving with the Computational Chemistry Landscape, *Chem. Rev.* 121 (2021) 4962–4998. doi:10.1021/acs.chemrev.0c00998.
- [63] S. Das, P. Motamarri, V. Subramanian, D. M. Rogers, V. Gavini, DFT-FE 1.0: A massively parallel hybrid CPU-GPU density functional theory code using finite-element discretization, *Comp. Phys. Commun.* 280 (2022) 108473. doi:10.1016/j.cpc.2022.108473.
- [64] R. Schade, T. Kenter, H. Elgabarty, M. Lass, T. D. Kühne, C. Plessl, Breaking the exascale barrier for the electronic structure problem in

- ab-initio molecular dynamics, *Int. J. High Perform. Comput. Appl.* 37 (2023) 530–538. doi:10.1177/10943420231177631.
- [65] I. Carnimeo, F. Affinito, S. Baroni, O. Baseggio, L. Bellentani, R. Bertossa, P. D. Delugas, F. F. Ruffino, S. Orlandini, F. Spiga, P. Giannozzi, Quantum ESPRESSO: One Further Step toward the Exascale, *J. Chem. Theory Comput.* 19 (2023) 6992–7006. doi:10.1021/acs.jctc.3c00249.
- [66] F. Zahariev, P. Xu, B. M. Westheimer, S. Webb, J. Galvez Vallejo, A. Tiwari, V. Sundriyal, M. Sosonkina, J. Shen, G. Schoendorff, M. Schlinsog, T. Sattasathuchana, K. Ruedenberg, L. B. Roskop, A. P. Rendell, D. Poole, P. Piecuch, B. Q. Pham, V. Mironov, J. Mato, S. Leonard, S. S. Leang, J. Ivanic, J. Hayes, T. Harville, K. Gururangan, E. Guidez, I. S. Gerasimov, C. Friedl, K. N. Ferreras, G. Elliott, D. Datta, D. D. A. Cruz, L. Carrington, C. Bertoni, G. M. J. Barca, M. Alkan, M. S. Gordon, The General Atomic and Molecular Electronic Structure System (GAMESS): Novel Methods on Novel Architectures, *J. Chem. Theory Comput.* 19 (2023) 7031–7055. doi:10.1021/acs.jctc.3c00379.
- [67] V. Gavini, S. Baroni, V. Blum, D. R. Bowler, A. Buccheri, J. R. Chelikowsky, S. Das, W. Dawson, P. Delugas, M. Dogan, C. Draxl, G. Galli, L. Genovese, P. Giannozzi, M. Giantomassi, X. Gonze, M. Govoni, F. Gygi, A. Gulans, J. M. Herbert, S. Kokott, T. D. Kühne, K.-H. Liou, T. Miyazaki, P. Motamarri, A. Nakata, J. E. Pask, C. Plessl, L. E. Ratcliff, R. M. Richard, M. Rossi, R. Schade, M. Scheffler, O. Schütt, P. Suryanarayana, M. Torrent, L. Truflandier, T. L. Windus, Q. Xu, V. W.-Z. Yu, D. Perez, Roadmap on electronic structure codes in the exascale era, *Model. Simul. Mat. Sci. Eng.* 31 (2023) 063301. doi:10.1088/1361-651x/acdf06.
- [68] Kick-off of 10 Centres of Excellence in HPC to support the transition towards exascale, https://eurohpc-ju.europa.eu/kick-10-centres-excellence-hpc-support-transition-towards-exascale-2023-01-26_en, 2023. Date accessed: 2024-02-11.
- [69] Exascale Computing Project, <https://www.exascaleproject.org/>, 2024. Date accessed: 2024-02-11.

- [70] CPMD, copyright 1990-2023 by IBM Corp. and copyright 1994-2001 by Max Planck Institute, Stuttgart., <http://www.cpmd.org/>, 2023. Date accessed: 2024-02-11.
- [71] JUWELS: Jülich Wizard for European Leadership Science, <https://www.fz-juelich.de/en/ias/jsc/systems/supercomputers/juwels>, 2022. Date accessed: 2024-02-11.
- [72] M. G. Chiariello, V. Bolnykh, E. Ippoliti, S. Meloni, J. M. H. Olsen, T. Beck, U. Rothlisberger, C. Fahlke, P. Carloni, Molecular Basis of CLC Antiporter Inhibition by Fluoride, *J. Am. Chem. Soc.* 142 (2020) 7254–7258. doi:10.1021/jacs.9b13588.
- [73] M. G. Chiariello, M. Alfonso-Prieto, E. Ippoliti, C. Fahlke, P. Carloni, Mechanisms Underlying Proton Release in CLC-type F⁻/H⁺ Antiporters, *J. Phys. Chem. Lett.* 12 (2021) 4415–4420. doi:10.1021/acs.jpcllett.1c00361.
- [74] F. K. Schackert, J. Biedermann, S. Abdolvand, S. Minniberger, C. Song, A. J. R. Plested, P. Carloni, H. Sun, Mechanism of Calcium Permeation in a Glutamate Receptor Ion Channel, *J. Chem. Inf. Model.* 63 (2023) 1293–1300. doi:10.1021/acs.jcim.2c01494.
- [75] R. Capelli, W. Lyu, V. Bolnykh, S. Meloni, J. M. H. Olsen, U. Rothlisberger, M. Parrinello, P. Carloni, Accuracy of Molecular Simulation-Based Predictions of k_{off} Values: A Metadynamics Study, *J. Phys. Chem. Lett.* 11 (2020) 6373–6381. doi:10.1021/acs.jpcllett.0c00999.
- [76] O. T. Unke, S. Chmiela, H. E. Sauceda, M. Gastegger, I. Poltavsky, K. T. Schütt, A. Tkatchenko, K.-R. Müller, Machine Learning Force Fields, *Chem. Rev.* 121 (2021) 10142–10186. doi:10.1021/acs.chemrev.0c01111.
- [77] E. Kocer, T. W. Ko, J. Behler, Neural Network Potentials: A Concise Overview of Methods, *Annu. Rev. Phys. Chem.* 73 (2022) 163–186. doi:10.1146/annurev-physchem-082720-034254.
- [78] F. Mouvet, J. Villard, V. Bolnykh, U. Rothlisberger, Recent Advances in First-Principles Based Molecular Dynamics, *Acc. Chem. Res.* 55 (2022) 221–230. doi:10.1021/acs.accounts.1c00503.

- [79] R. Galvelis, A. Varela-Rial, S. Doerr, R. Fino, P. Eastman, T. E. Markland, J. D. Chodera, G. De Fabritiis, NNP/MM: Accelerating Molecular Dynamics Simulations with Machine Learning Potentials and Molecular Mechanics, *J. Chem. Inf. Model.* 63 (2023) 5701–5708. doi:10.1021/acs.jcim.3c00773.
- [80] A. Rizzi, P. Carloni, M. Parrinello, Free energies at QM accuracy from force fields via multimap targeted estimation, *Proc. Natl. Acad. Sci.* 120 (2023) e2304308120. doi:10.1073/pnas.2304308120.