

Python-Based Quantum Chemistry Calculations with GPU Acceleration

Xiaojie Wu^{1,*}, Qiming Sun^{1,†}, Zhichen Pu¹, Tianze Zheng¹, Wenzhi Ma¹, Wen Yan¹, Xia Yu¹, Zhengxiao Wu¹, Mian Huo¹, Xiang Li¹, Weiluo Ren¹, Sheng Gong¹, Yumin Zhang¹, and Weihao Gao¹

¹ByteDance Inc.

*xiaojie.wu@bytedance.com

†osirpt.sun@gmail.com

April 16, 2024

Abstract

To meet the increasing demand of quantum chemistry calculations in data-driven chemical research, the collaboration between industrial stakeholders and the quantum chemistry community has led to the development of GPU4PySCF, a GPU-accelerated Python package. This open-source project is accessible via its public GitHub repository at <https://github.com/pyscf/gpu4pyscf>. This paper outlines the primary features, innovations, and advantages of this package. When performing Density Functional Theory (DFT) calculations on modern GPU platforms, GPU4PySCF delivers 30 times speedup over a 32-core CPU node, resulting in approximately 90% cost savings for most DFT tasks. The performance advantages and productivity improvements have been found in multiple industrial applications, such as generating potential energy surfaces, analyzing molecular properties, calculating solvation free energy, identifying chemical reactions in lithium-ion batteries, and accelerating neural-network methods. To make the package easy to extend and integrate with other Python packages, it is designed with PySCF-compatible interfaces and Pythonic implementations. This design choice enhances its coordination with the Python ecosystem.

1 Introduction

Quantum chemistry is essential in various fields such as drug discovery, materials science, chemical engineering, and environmental science. It primarily relies on *ab initio* simulations for atomic interactions, serving as a crucial source of data beyond experimental findings. These fields often require extensive *ab initio* simulations to explore chemical spaces, identify equilibrium geometries, and determine chemical properties. Density Functional Theory (DFT) stands out as the most commonly used method in quantum chemistry for generating physical observables. According to a 2018 workload analysis by the National Energy Research Scientific Computing Center (NERSC), DFT algorithms account for 30% of total computational resource usage¹. The surge in demand for quantum chemistry calculations is not only due to traditional scientific simulations but also the emergence of deep learning models. These models significantly increase the demand for data from *ab initio* calculations and use quantum features as input. Consequently, the development of DFT algorithms for GPUs has become increasingly beneficial for these industries.

The utilization of GPUs in quantum chemistry has emerged as a transformative tool, thanks to their parallel processing capabilities. GPUs are adept at managing large data volumes and conducting calculations across multiple cores simultaneously, making them exceptionally efficient for quantum chemistry simulations. The initiative to accelerate two-electron integral evaluations using GPUs can be traced back to 2008 [1, 2], a period when gaming cards exhibited limited double-precision computing capabilities. Today’s GPUs not only offer substantial floating-point performance but also boast high

¹https://portal.nersc.gov/project/m888/nersc10/workload/N10_Workload_Analysis.latest.pdf

memory bandwidth and extensive memory capacity, effectively minimizing the bottlenecks associated with quantum chemistry algorithms. Traditional CPU-based quantum chemistry software packages, such as GAMESS [3, 4], GAUSSIAN, Q-Chem [5], and Psi4 [6] now delegate computationally intensive tasks to GPUs through interfaces like BrainQC [7, 8]. Furthermore, the past decade has seen the development of GPU-based quantum chemistry packages, including Terachem [9, 10] and Quick [11, 12, 13], which notably reduce CPU-GPU data transfers.

Achieving optimal performance in quantum chemistry computations on Graphics Processing Units (GPUs) necessitates an approach that extends significantly beyond the simple porting of algorithms from Central Processing Units (CPUs) to GPU architectures. This complex process is hindered by several technical challenges that can substantially impact the acceleration of quantum chemistry simulations on GPUs. 1) *Limited memory*. GPUs are equipped with high-speed memory that, despite its rapid access capabilities, is limited in capacity and substantially smaller than the system RAM found in CPU-based systems. The quantum chemistry domain, characterized by its reliance on large matrices and complex wave functions, frequently demands more memory than GPUs can provide. This limitation often requires computational strategies such as offloading portions of the data back to the CPU or the implementation of advanced memory management techniques, both of which can negatively affect computational throughput. 2) *CPU-GPU communication*. The bandwidth between CPU memory and GPU memory is restricted and can become a critical bottleneck, especially in computational scenarios where the GPU necessitates frequent access or updates to data stored in the main system memory. This bottleneck necessitates the re-implementation of a significant number of traditional CPU algorithms directly on the GPU to minimize data transfer overhead. 3) *Complexities of GPU architecture*. The process of maximizing GPU computational efficiency involves the fine-tuning of several coding and architectural aspects, including but not limited to, optimizing memory access patterns, managing thread execution, and configuring kernels. The same optimization strategies do not necessarily apply to the new architecture. This level of optimization is highly specialized and demands a considerable investment of time, highlighting the intricate and expert-driven nature of leveraging GPU capabilities for the advancement of quantum chemistry calculations.

Optimizing a quantum chemistry workflow that balances chemical accuracy and computational efficiency requires extensive domain expertise and practical experience. Initially, a quantum chemist is tasked with selecting an appropriate basis set and exchange-correlation functional from a broad dataset, exemplified by the Basis Set Exchange, which catalogs over 600 basis sets [14], and LibXC, comprising more than 400 exchange-correlation functionals [15]. Subsequent considerations include evaluating methods for dispersion correction, solvent model selection, basis set superposition error (BSSE) correction, and establishing convergence criteria for various computational tasks. The adoption of recent composite DFT protocols, such as PBEh-3c [16], ω B97X-3c [17], and r2SCAN-3c [18], necessitates precise configurations of basis sets, exchange-correlation functionals, BSSE corrections, and dispersion corrections. Additionally, contemporary quantum chemistry workflows frequently incorporate integration with other software packages, including RDKit, PyTorch, and Qiskit, to extend their computational capabilities. Notably, the open-source communities behind PySCF [19] and Psi4 [6] have contributed user-friendly Python interfaces, offering a wide range of functionalities. PySCF, in particular, is recognized for its straightforward, Pythonic interface, facilitating seamless integration with other Python libraries and enabling the scripting of customized quantum chemistry workflows. Its adaptable design also supports acceleration within the modern GPU ecosystem, underscoring the importance of flexible, integrated approaches in the development of efficient quantum chemistry workflows.

In this work, we have developed a Python-based and GPU-optimized framework, building upon the foundations of PySCF. This initiative aims to advance quantum chemistry calculations on GPUs by focusing on several key areas: 1) Prioritizing density fitting algorithms, which are inherently more GPU-friendly, while offering limited support for direct Self-Consistent Field (SCF) methods. 2) Adapting our codebase to leverage modern GPU architectures, such as tensor cores, to achieve up to twice the speed in tensor contractions. 3) Integrating a wide array of open-source packages from the quantum chemistry community within our GPU4PySCF package, thus supporting sophisticated DFT methods, popular basis sets, solvent models, chemical property calculations, geometry optimization, and transition state search methodologies. This comprehensive approach signifies a notable advancement in the field of quantum chemistry, bridging the gap between traditional computational methods and the cutting-edge capabilities offered by modern GPU technology.

We underscore the notable advantages of GPU4PySCF as follows:

- **Speed:** GPU4PySCF provides performance that is equivalent to that of 600-1000 CPU cores for running conventional CPU-based quantum chemistry software. As a result, it offers substantial cost savings, for instance, up to 90% when utilizing the NVIDIA A100-80G cards.
- **Python Integration:** The framework allows direct access to quantum features of atoms at the Python level, ensuring seamless compatibility with other Python-based packages such as PyTorch, Jax, TensorFlow, RDkit, and more.
- **Open Source:** GPU4PySCF benefits from the support of the PySCF community, the largest open-source quantum chemistry community, fostering collaborative development and innovation.
- **Application Driven:** Specifically optimized and tailored for industrial applications, GPU4PySCF excels in tasks such as DFT calculations, offering practical advantages for commercial research and development.

Detailed performance metrics and benchmarks are provided in Section 2. These calculations have been rigorously cross-validated with Q-Chem 6.1, ensuring reliability and accuracy. Furthermore, Section 3 briefly illustrates how GPU4PySCF can be utilized to navigate the potential energy landscape, analyze quantum features, compute thermodynamic properties, estimate solvation free energies, quantify chemical reactions, and integrate with neural network models. This comprehensive approach showcases the framework’s versatility and potential to revolutionize quantum chemistry computations through GPU acceleration.

2 Performance and Benchmark

2.1 Dataset and Performance

We focus on the benchmarks with density fitting scheme, which is efficient for small molecules with less than 200 atoms. A small and diverse dataset is constructed for the benchmarking the fundamental modules in GPU4PySCF. We select small molecules (< 100 atoms) from the supplement information in [20] and two medium-sized molecules from the supplement information in [8]. The newly constructed dataset includes 8 elements (H, C, O, N, Mg, S, Cl, P), and 13 small molecules with 20-168 atoms. The details of the dataset are provided in Appendix A. The xyz files and benchmarking details can be found on [GitHub](#). Other transition metals, such as Nb, Ti, Ru and so on, are also supported in GPU4PySCF, although they are not included in this dataset. This capability will be shown in Sec. 3.4.1. Larger molecules with more than 168 atoms still can be calculated with the direct SCF scheme in GPU4PySCF. The detailed performance of direct SCF scheme will be disclosed in the future work. The availability of direct SCF scheme can be found in Sec. 4.2.

For a typical DFT protocol with def2-TZVPP, B3LYP, and (99,590) grids, we time SCF, gradient, and Hessian calculations on NVIDIA A100-80G. For the smallest molecule in the dataset, SCF, gradient, and Hessian take 1.7, 0.68, and 43 seconds respectively. For the largest molecule in the dataset, SCF, gradient, and Hessian take 309 seconds, 58 seconds, and 21 hours respectively. Hessian calculation is the most expensive module of the vibrational analysis. With GPU4PySCF, one can perform the vibrational analysis for large molecules with 168 atoms in one day. The vibrational analysis of moderately sized (100-200 atoms) molecules is prohibitively expensive with traditional quantum chemistry softwares [21]. The time can be further reduced with the recent GPUs such as H100. The time scaling curve of SCF, gradient, and Hessian calculations is shown in Figure 1. Theoretically, SCF and gradient calculations are quartic scaling with respect to the system size, while Hessian calculation is $O(N^5)$. The slope of the scaling curve decreases as the system size gets larger, due to sparsity. The bottleneck of calculating a large molecule is the CPU memory storage. A CPU-GPU hybrid strategy is employed to store GTO integrals in this work. For small molecules, all the intermediate variables are stored in GPU memory. For the large system, the variables are stored in CPU memory. The data in CPU memory is asynchronously transferred to GPU memory for consumption. The performance does not drop significantly in the scaling curve.

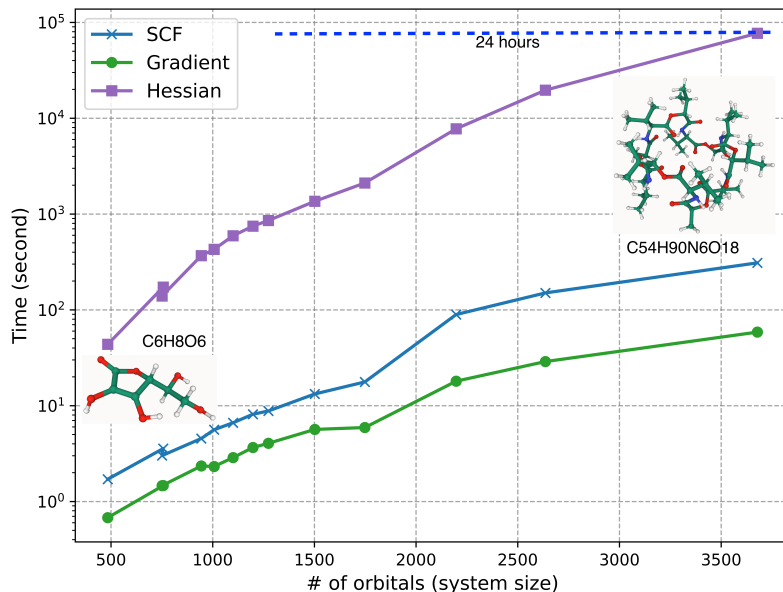


Figure 1: Scaling curve of SCF, gradient, and Hessian calculations. def2-TZVPP basis set, def2-universal-JKFIT auxiliary basis, B3LYP, (99,590) grids, on A100-80G.

2.2 Speedup of SCF, Gradient, and Hessian Calculations

Achieving complete alignment between various quantum chemistry software settings presents a significant challenge. Fundamental configurations like the threshold for two-electron integrals, exchange-correlation functionals, basis sets, and Lebedev grids are generally standardized within the quantum chemistry community. However, more intricate adjustments—such as mitigating the linear dependence of atomic orbitals, defining SCF convergence criteria, and pruning DFT grids—often exhibit subtle discrepancies. Moreover, benchmarking these results against commercial software treated as a "black box" adds an additional layer of complexity. In our analysis, we strive to minimize these discrepancies by uniformly applying stringent criteria across basic settings. Given the restricted access to commercial quantum chemistry software, our benchmarking efforts are focused exclusively on Q-Chem, a prominent commercial platform renowned for its cutting-edge algorithms designed for large systems. The discrepancies between Q-Chem 6.1 and GPU4PySCF for exchange-correlation functionals and different basis sets are shown in Appendix B and C.

Our benchmarking endeavors concentrate on three foundational modules: Self-Consistent Field (SCF), gradient, and Hessian calculations. These modules are the most expensive components in numerous quantum chemistry computations. In this section, all computations are conducted on closed-shell systems using density fitting. The functionalities of open-shell calculations will be shown in Sec. 3.4.2 and Sec. 3.4.3. Notably, due to the memory management issue within the density fitting coupled perturbed SCF iteration in Q-Chem, the reference Hessian calculations are performed using finite difference methods. But the Hessian matrices are calculated analytically in GPU4PySCF. We anticipate that resolving these issues could enhance the efficiency of Hessian calculations in Q-Chem by a factor of 2-3. Hessian calculations for molecules exceeding 84 atoms become prohibitively expensive for traditional quantum chemistry packages. We only benchmark the speedup across molecules containing 20-84 atoms.

On average, GPU4PySCF achieves a performance that is 20 times faster than identical SCF calculations performed using Q-Chem 6.1 on 32 CPU cores. This speedup increases dramatically to 50 times for Hessian calculations, although the acceleration observed in gradient calculations is relatively modest. However, since gradient calculations are generally quicker than SCF iterations, this does not significantly affect the overall efficiency of geometry optimization or transition state search tasks. An evaluation of the actual costs based on AWS pricing for comparable machinery shows that

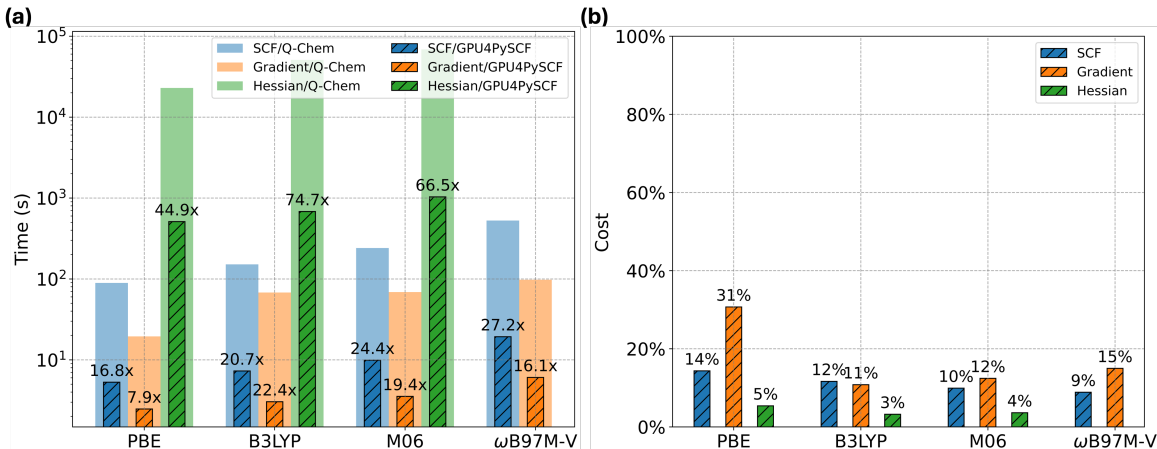


Figure 2: Speedup and actual cost saving of GPU4PySCF on A100-80G over Q-Chem v6.1 on 32-core CPU. The cost is estimated based on AWS pricing. \$40.966/hr for **p4de.24xlarge** GPU instance with 8 A100-80G GPUs, 96 vCPUs and 1152 GiB CPU memory. \$2.117/hr for **r7i.8xlarge** CPU instance with 32 vCPUs and 256 GiB memory. def2-TZVPP basis set, def2-universal-JKFIT auxiliary basis, (99,590) xc grids, (50,194) nlc grids.

GPU4PySCF can reduce expenses by about 90% across most tasks. For the gradient calculation with pure DFT, the saving decreases to 70%. This task technically can be accelerated with a more efficient algorithm [22]. We leave this implementation as future work. It is observed that larger molecules typically experience a more substantial speedup compared to smaller molecules. This is due to the high GPU occupancy of a large molecule.

2.3 Speedup of DFT with Implicit Solvent Models

Solvent models in quantum chemistry critical for simulating and understanding the behavior of molecules in solution, an essential aspect given most chemical reactions occur in some type of solvent. The Polarizable Continuum Model (PCM) models are the most common implicit solvent models. But the computational efficiency of those models are rarely discussed. Especially when the density fitting scheme is employed, the computational cost of PCM models is not negligible. Since Q-Chem 6.1 does not support density fitting SCF with PCM models, we use the standard SCF scheme as a reference. Thus, the following speedups should be interpreted as two factors: GPU speedup and density fitting speedup. We stress that the different algorithms are used in Q-Chem and GPU4PySCF. The comparison is performed to mimic the actual usage of the solvent models. The discrepancies between Q-Chem 6.1 and GPU4PySCF for solvent models are shown in Appendix D.

We benchmark DFT calculations employing two PCM models, C-PCM [23, 24] and IEF-PCM [25]. IEF-PCM model is slightly more expensive than C-PCM model since more entries are calculated for the linear system. The SMD model [26] introduces the additional computation of CDS contributions on the top of IEF-PCM model. Yet, the computational cost of these contributions is almost negligible. The efficiency of CDS contributions will not be discussed in this section, but the accuracy benchmarks will be presented in Section 3.3. Other types of implicit solvent models, such as generalized Born models, are not implemented in the current version.

We evaluate the performance of molecules with up to 42 atoms. The computational cost of larger molecules is prohibitively expensive with Q-Chem, although GPU4PySCF is able to handle large molecules. With the A100-80G GPU and density fitting, GPU4PySCF significantly accelerates the SCF calculations by 40-80 times, gradient calculations by 20-40 times, and Hessian calculations by 100-170 times, respectively. The speedup of the calculations in the gas phase is slightly lower than the results in Table 1. For instance, the speedup of SCF for Vitamin C in gas phase is around 33. The speedup difference between gas phase and liquid phase suggests the high efficiency of PCM models in GPU4PySCF.

Molecule	C-PCM			IEF-PCM		
	SCF	Gradient	Hessian	SCF	Gradient	Hessian
Vitamin C	41.1	28.4	106.4	41.6	27.3	104.9
Inosine	65.5	38.9	147.4	65.6	40.2	146.2
Bisphenol A	76.2	43.3	162.7	76.7	45.4	161.3
Mg Porphin	84.0	49.0	162.2	87.3	49.8	160.1
Penicillin V	83.2	44.1	170.5	84.0	43.4	169.4

Table 1: Speedup of GPU4PySCF on A100-80G over Q-Chem 6.1 on 32-core CPUs for DFT tasks with PCM solvent models. B3LYP, def2-TZVPP, def2-universal-JKFIT, (99,590), 302 Lebedev grids for both H atoms and Heavy atoms.

3 Applications

In this section, we showcase the capabilities of GPU4PySCF with several industrial applications. A complete list of capabilities will be shown in Section 4.1.

3.1 Exploring Potential Energy Surface

3.1.1 Torsion Scan

The development of molecular mechanics force fields significantly benefits from torsion scans conducted at precise quantum chemistry levels, serving as a crucial source of training data [27, 28, 29, 30, 31, 32, 33]. However, torsion scans are notably time-consuming, as they require at least 24 independent constrained geometry optimizations for each proper torsion angle in a molecule. These optimizations determine the potential energy of the molecule with the torsion angle fixed at specific values within the $[-180^\circ, +180^\circ]$ range. Depending on the complexity of the potential energy surface, each optimization might necessitate 10-100 sequential quantum chemistry calculations, underscoring the importance of computational efficiency for this task. Historically, such scans have been restricted to small molecule fragments, typically using small basis sets at the double-zeta level.

The advent of GPU4PySCF, coupled with its seamless integration with the Pythonic interface of the geometric optimizer geomeTRIC, revolutionizes this process, allowing for torsion scans on a scale previously deemed unfeasible. This enables routine computations of torsion scans for real-world drug molecules employing large basis sets.

An illustrative example of this capability is provided by our work on Enzalutamide (Figure 3). Utilizing GPU4PySCF, we conducted a torsional energy profile scan using ω B97M-V/def2-TZVPP at 10° step (37 angles in total). This comprehensive analysis entailed 1089 DFT evaluations, with the entire scan completed in 34,679.9 seconds, averaging approximately 32 seconds per DFT evaluation. This ability to scan large, complex molecules using high-precision quantum chemistry algorithms significantly broadens the scope of chemical spaces accessible through torsion scan methodologies. All calculations in this subsection were performed using grid level 3 in PySCF with density fitting on an A100-80G GPU.

3.1.2 Dimer Interaction Energies

Noncovalent interactions are paramount in molecular modeling, especially for the accurate prediction of thermodynamics and transport properties [34, 35]. Additionally, robust noncovalent interactions, such as hydrogen bonding and π - π stacking, are essential for ligand-protein binding [36, 37, 31, 38]. The assessment of dimer interaction energies through quantum mechanical approaches serves as a prevalent method to explore noncovalent interactions. However, attaining precise noncovalent interaction energies demands high-fidelity techniques, which invariably elevate computational costs. Thanks to the efficiency of GPU4PySCF, interaction energies for systems comprising hundreds of atoms are now attainable within practical timeframes, even when utilizing sophisticated DFT methodologies that incorporate non-local correlation effects.

We offer an illustrative example involving the oral COVID-19 antiviral Nirmatrelvir [39] (PubChem CID 155903259, Figure 4 (a)), a SARS-CoV-2 main protease inhibitor developed by Pfizer. In its binding pocket, Nirmatrelvir forms a hydrogen bond with a histidine side chain. The interaction

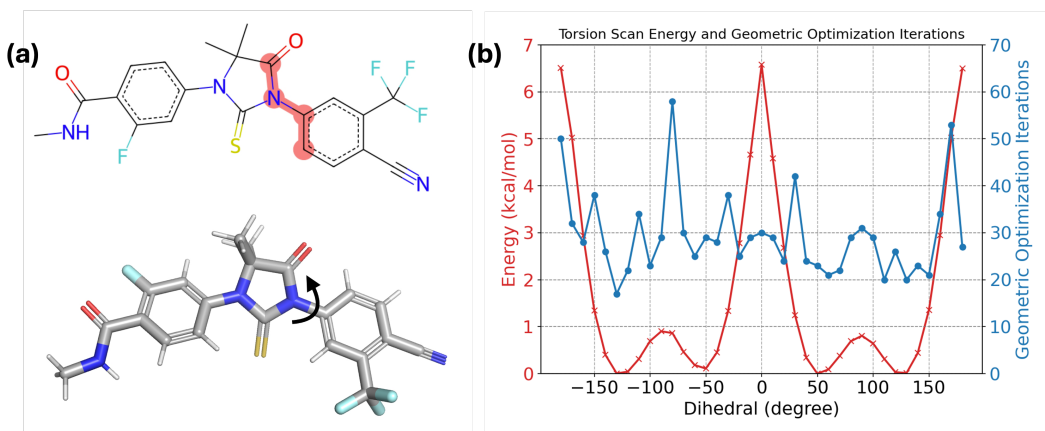


Figure 3: An example of torsion scan. (a) The molecule Enzalutamide (PubChem CID 15951529). (b) Torsion scan energy profile and number of geometry optimization iterations (number of DFT calculations) at each scanned proper torsion angle degree.

energy between Nirmatrelvir and the histidine residue (Figure 4 (b)) is calculated using various bases. For the direct evaluation of interaction energy, 42 single-point energy calculations were conducted. The total wall times for def2-SVP, def2-SVPD, and def2-TZVPPD bases were 2370, 3082, and 5823 seconds, respectively. Moreover, the counterpoise method (CP) [40] was applied to mitigate basis set superposition error (BSSE) [41, 42], involving 120 single-point energy calculations for each base, with total wall times of 4020, 5510, and 11892 seconds, respectively. All these calculations were performed using the ω B97M-V functional and grid level 3 in PySCF with density fitting on an A100-80G GPU.

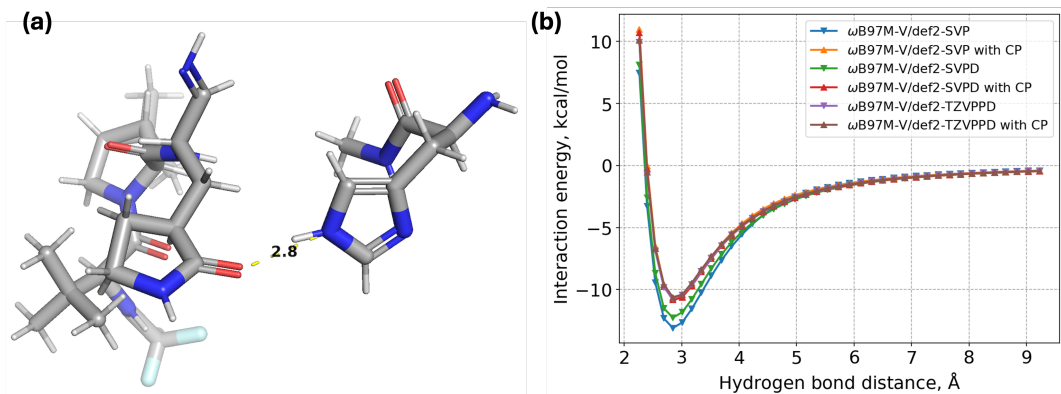


Figure 4: An example of noncovalent interaction. (a) Nirmatrelvir (left) and the histidine residue (right). The hydrogen bond distance (between donor and acceptor atoms) is 2.8 Å in the shown geometry, which is truncated from PDB 7yrz. (b) Interaction energy at varying hydrogen bond distances.

3.1.3 Lithium Ion Solvation Structure

The concept of solvation structure engineering has been utilized in recent studies, such as high-concentrated [43], localized high-concentrated [44], and fluorinated ether-based electrolytes [45, 46]. By engineering solvation structures, these studies show improvements in ion transport, solid electrolyte interphase, and electrochemical stability [47, 48, 49, 50]. Binding energy of solvation structures is an indicator of solvents' ability to dissolve salt [51]. In addition, binding energy plays a crucial role in solvation and desolvation, which is the mechanism for Li^+ plating/stripping and intercalation in Li batteries [52, 53]. As such, investigating solvation structures and their corresponding binding energies facilitate the realization of high energy density, fast charging and improved cell lifetime.

In this work, we examine the solvation structures of Li^+ ions in the electrolyte by computing the structure and binding energy of Li^+ ion clusters. Through molecular dynamics simulations of the

battery electrolyte containing 2.25M LiFSI(lithium bis(fluorosulfonyl)imide) in DMC(dimethyl carbonate):EC(ethylene carbonate) with a weight percentage ratio of 51:49, we captured typical solvation structures of Li^+ ions and computed the binding energies of these structures using GPU4PySCF, as presented in Figure 5 and Table 2.

Most of the clusters we examined have a coordination number of 4, consisting of either solvent or anion molecules, or a combination of the two. Their binding energies range from -0.065 eV to -1.078 eV, indicating significant differences in the energy stability of the clusters. We observe that clusters composed solely of Li^+ and FSI^- exhibit the most positive binding energies compared to those containing solvents. This observation aligns with chemical intuition, suggesting that the salt dissolution of LiFSI in an EC-DMC polar solvent mixture is preferred, manifesting the fundamental requirement of solvent functionality in Li-battery liquid electrolyte applications. Besides, clusters with the same composition can exhibit different energies if their structures vary, as shown in Figure 6 for the $\text{Li}^+(\text{DMC})(\text{EC})(\text{FSI}^-)_2$ cluster, suggesting adequate geometric sampling for DFT calculations is necessary for accurate representation of solvation structure stabilities, emphasizing the importance of computational efficiency provided in our code.

All calculations are performed with 6-311++G(3df,3pd) basis set and cc-pVTZ-RI auxiliary basis for the density fitting technique. The IEF-PCM with UFF radii was utilized for modeling solvent effects, with the dielectric constant being 20.0. Additionally, M06-2X functional together with Grimme's D3 version of dispersion correction with zero damping was used. Note that our examples in this section are merely illustrative and do not provide a complete study of the issue of Li^+ ion solvation structures.

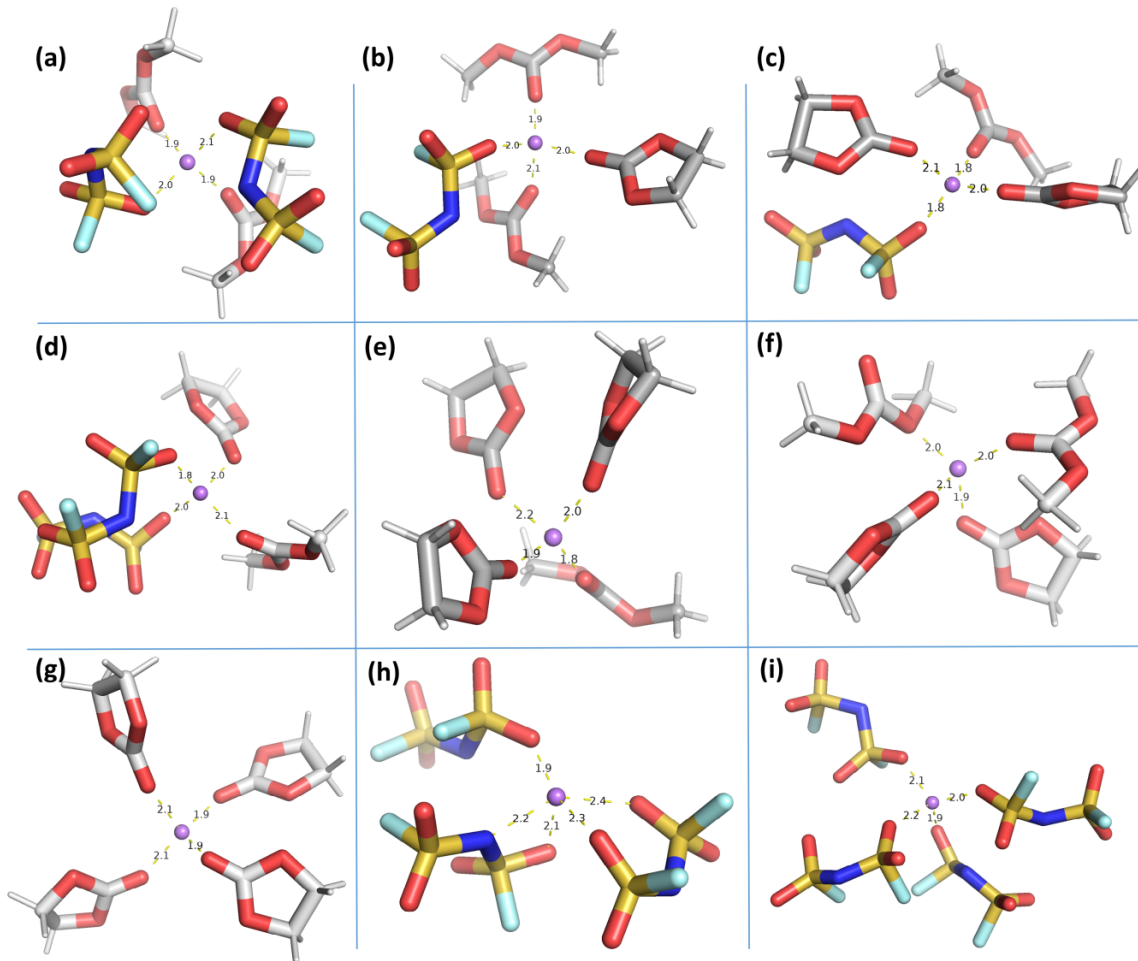


Figure 5: Solvation structures of Li^+ ion. The clusters depicted in images (a) to (i) correspond respectively to the clusters listed in Table 2. The distances between Li^+ ion and coordinated atoms are plotted in Å.

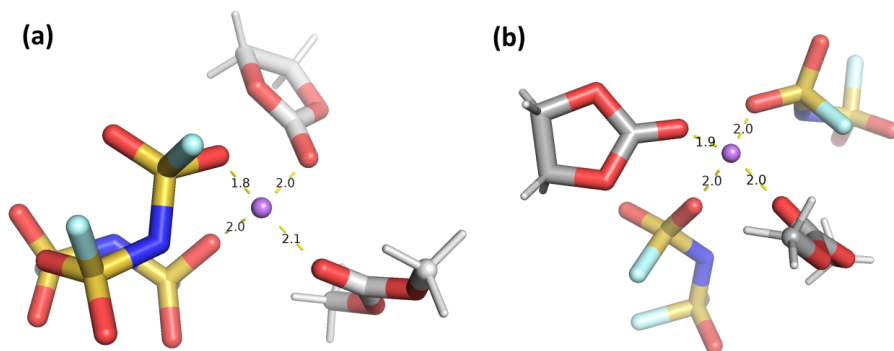


Figure 6: Two different structures of $\text{Li}^+(\text{DMC})(\text{EC})(\text{FSI}^-)_2$. The energies of clusters in (a) and (b) are -1.078 eV and -0.941 eV, respectively. The distances between Li^+ ion and coordinated atoms are plotted in Å.

Cluster	Binding Energy
$\text{Li}^+(\text{DMC})_2(\text{FSI}^-)_2$	-1.077
$\text{Li}^+(\text{DMC})_2(\text{EC})(\text{FSI}^-)$	-1.020
$\text{Li}^+(\text{DMC})(\text{EC})_2(\text{FSI}^-)$	-1.018
$\text{Li}^+(\text{DMC})(\text{EC})(\text{FSI}^-)_2$	-1.004
$\text{Li}^+(\text{DMC})(\text{EC})_3$	-0.893
$\text{Li}^+(\text{DMC})_2(\text{EC})_2$	-0.857
$\text{Li}^+(\text{EC})_4$	-0.735
$\text{Li}^+(\text{FSI}^-)_3$	-0.656
$\text{Li}^+(\text{FSI}^-)_4$	-0.065

Table 2: Binding energies (eV) of various clusters.

3.2 Analyzing Molecular Properties

3.2.1 CHELPG Charge Calculation

CHELPG (charges from electrostatic potentials using a grid-based method) is a computational method designed to calculate atomic charges within a molecule based on the electrostatic potential generated by its electrons [54, 55]. Introduced by Breneman and Wiberg in 1990, this method aims to derive atomic charges that are reflective of the molecule’s electrostatic potential on a grid surrounding the molecule [54]. These charges can be then utilized in simulations to model intermolecular forces and other electrostatic properties. The essence and the most time-consuming part of CHELPG lies in calculating the electrostatic potential at grids around the molecule. The method then optimizes the atomic charges by fitting these calculated potentials to the actual electrostatic potential observed, using a least-squares method. This optimization process ensures that the derived atomic charges closely represent the true electrostatic potential around the molecule.

Followed by the work of Herbert and co-workers [56], CHELPG is implemented in our package, and the computational efficiency has been enhanced with the aid of GPUs. Our implementation uses the Cartesian grid with the default spacing 0.3 Å and the radial extent of the CHELPG grid 2.8 Å as recommended in Ref. [54]. Besides, the smoothing function used in Ref. [56] is also implemented, where the Bondi radius is used [57] (the radius of hydrogen is changed to 1.1 Å). Consistent results can be obtained when employing identical settings with Q-Chem. As shown in Figure 7, in GPU4PySCF, the calculation speed of CHELPG has been greatly improved approximately 10 times versus CPU-based quantum chemistry software [5]. This improvement in computational efficiency is due to the efficient implementation of molecular integrations (calculating the electrostatic potential at grids) on GPU. All the calculations in this subsection are performed using B3LYP functional and def2-SVPD basis with density fitting used.

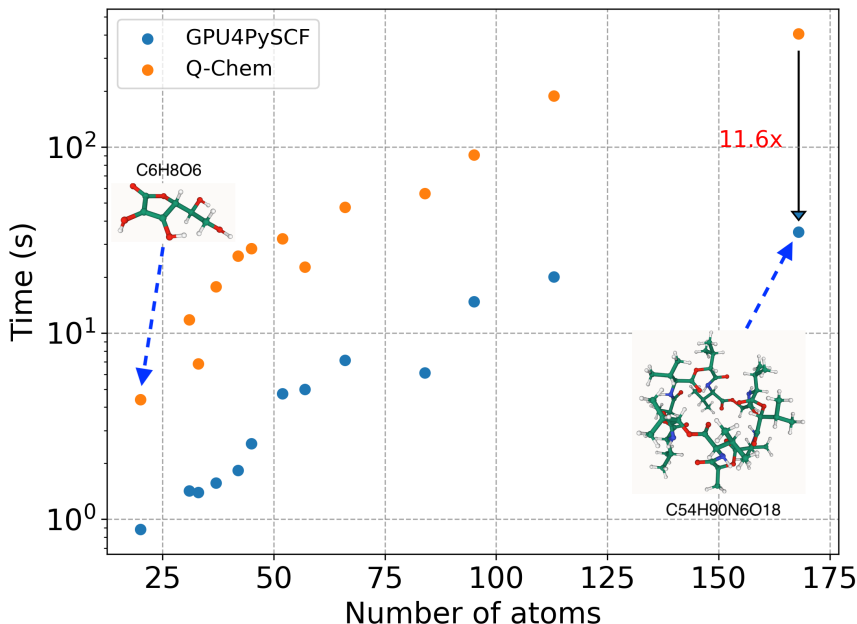


Figure 7: Comparison of the time required for CHELPG calculation between GPU4PySCF and Q-Chem. The computations by Q-Chem used 64 cores and 128G memory, while GPU4PySCF used 15 cores, 245G memory, and one A100-SXM-80GB GPU. The tested systems are benchmark molecules listed in Table 5.

3.2.2 NMR Shielding Constants

Nuclear Magnetic Resonance (NMR) Shielding Constants represent a pivotal spectroscopic parameter in NMR spectroscopy, utilized for determining the electronic structure and geometric configuration of molecules [58]. During NMR experiments, the absorption signals through nuclear spin transitions under the external magnetic field are detected. Actually, the magnetic field experienced by the nucleus is not the external magnetic field per se, but rather the residual field post electron shielding. The shielding constant elucidates the behavior of the wave function in the vicinity of the nucleus, serving as a coefficient between the electron shielding effect and the external magnetic field. Its ability to provide detailed information about the molecular framework and atomic-level interactions makes it indispensable for chemists. For example, NMR is particularly valuable in organic chemistry for identifying the composition and structure of small organic molecules, aiding in the elucidation of complex organic compounds and reaction mechanisms [59, 60]. Employing DFT to calculate NMR shielding constants offers the significant advantage of a level of precision that closely aligns with experimental results with efficiency, yet without the extensive resource demands of experiments [61]. This balance makes calculations of NMR shielding constants an invaluable tool for guiding experimental design and interpreting complex NMR spectra, bridging the gap between theoretical insights and practical applications in stereochemistry. In this work, the NMR shielding constants are implemented using the gauge-including atomic orbitals [62] (GIAO) to address the uncertainty of the gauge origin, introduced by the introduction of the uniform external magnetic field. The specific working equations may be referred to in the discussion section on general NMR calculations as presented in Ref. [63].

Nuclear Magnetic Resonance (NMR) Shielding Constants represent a pivotal spectroscopic parameter in NMR spectroscopy, utilized for determining the electronic structure and geometric configuration of molecules [58]. During NMR experiments, the absorption signals through nuclear spin transitions under the external magnetic field are detected. Actually, the magnetic field experienced by the nucleus is not the external magnetic field per se, but rather the residual field post electron shielding. The shielding constant elucidates the behavior of the wave function in the vicinity of the nucleus, serving as a coefficient between the electron shielding effect and the external magnetic field. Its ability to

provide detailed information about the molecular framework and atomic-level interactions makes it indispensable for chemists. For example, NMR is particularly valuable in organic chemistry for identifying the composition and structure of small organic molecules, aiding in the elucidation of complex organic compounds and reaction mechanisms [59, 60]. Employing DFT to calculate NMR shielding constants offers the significant advantage of a level of precision that closely aligns with experimental results with efficiency, yet without the extensive resource demands of experiments [61]. This balance makes calculations of NMR shielding constants an invaluable tool for guiding experimental design and interpreting complex NMR spectra, bridging the gap between theoretical insights and practical applications in stereochemistry. In this work, the NMR shielding constants are implemented using the gauge-including atomic orbitals [62] (GIAO) to address the uncertainty of the gauge origin, introduced by the introduction of the uniform external magnetic field. The specific working equations may be referred to in the discussion section on general NMR calculations as presented in Ref. [63].

An illustrative example of NMR shielding constants of hydrogens on toluene is displayed in Figure 8. When the toluene is oriented perpendicular to the external magnetic field (and for non-perpendicular orientations, it can be decomposed into horizontal and vertical magnetic fields), its delocalized π electrons will generate a ring current, which in turn induces a magnetic field. The direction of the induced magnetic field is opposite to that of the external magnetic field on the top and bottom sides of the phenyl ring. However, on the sides of the phenyl ring (where the hydrogen atoms are positioned on the sides of the ring), the directions are the same, i.e., the induced magnetic field enhances the effect of the external magnetic field, deshielding the hydrogen nuclei and causing the chemical shift to move to lower field values; whereas the hydrogen atoms on the methyl group are in a stronger shielding environment, leading to the chemical shift to move to higher field values. Due to the geometrical configuration, the three hydrogen atoms on the methyl group exhibit differing chemical shifts, although theoretically, they should be identical.

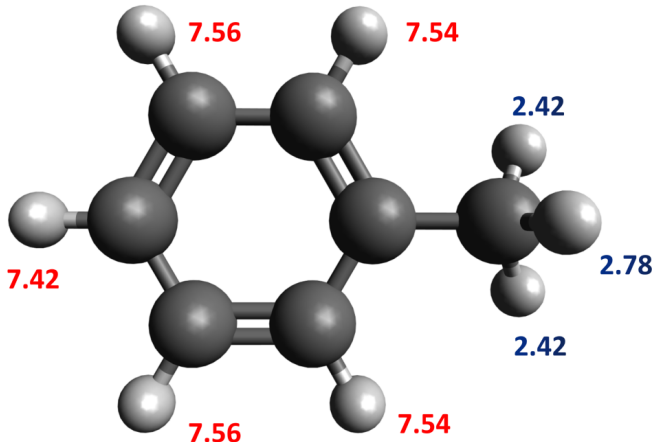


Figure 8: Chemical shifts of hydrogens on toluene (in units of p.p.m.). Hydrogens on the methyl group are indicated in blue text, while hydrogens on the benzene ring are in red. This calculation employs the B3LYP functional, the def2-TZVPP basis set, and the def2-universal-JKFIT auxiliary basis set, and the structure of toluene is taken from PubChem [64].

3.3 Solvation Free Energy

Implicit solvent models are computational tools designed to efficiently simulate the effect of a solvent on molecular systems, emphasizing the solvent’s overall impact rather than detailing every interaction between solvent and solute molecules. In the realm of quantum chemistry, electrostatic interactions between solvent and solute molecules are captured using the Polarizable Continuum Model (PCM), which offers a nuanced understanding of molecular properties in solution. In GPU4PySCF, users can access four variants of PCM models: C-PCM, COSMO, IEF-PCM, and SS(V)PE. For accurate prediction of solvation energies, accounting for non-electrostatic interactions such as cavitation, Pauli repulsion, dispersion, and hydrogen bonding is essential. These contributions are encompassed within

more advanced models like SMD [26], CMIRS [65], and COSMO-RS [66] models, with the SMD model being particularly prevalent in quantum chemistry calculations and available from GPU4PySCF.

The implementation of the SMD model varies across different quantum chemistry software platforms. And the algorithm has been updated with minor modifications over time. We benchmark our implementation against Q-Chem and Gaussian 03, and we note several unique aspects in our approach:

- The cavity surface of PCM models is smoothed using SWIG methods [67], enhancing the stability of geometry optimization and molecular dynamics simulations.
- Atomic SASA and molecular SASA are calculated numerically using Lebedev quadrature, diverging from the analytical formulation found in [68].
- Modified Coulomb radii are employed with SMD18 [69], with newly fitted parameters that improve the accuracy of halogen bonding interactions.

We employ the Minnesota Solvation Database – version 2012 [70] for the benchmarks. The database has been updated to version 2020 since the publication of SMD. The corresponding solvent descriptor database has been updated twice. We should take the changes into consideration for the following results. For neutral solutes, the error in solvation free energies consistently remains below 1 kcal/mol across various protocols (Table 3), with the errors decreasing further when employing larger basis sets or advancing up Jacob’s Ladder. These errors are comparable to those from corresponding protocols in Q-Chem and Gaussian 03. For achieving the best accuracy, a protocol employing a larger basis set and a higher-level Rung exchange-correlation functional is recommended. For ions, the overall error in solvation free energies is around 4 kcal/mol, with cases involving acetonitrile and dimethyl sulfoxide contributing significantly to this discrepancy. Notably, SMD calculations using Hartree-Fock generally outperform those utilizing DFT methods such as B3LYP and M06-2X, a conclusion supported by other studies [71, 72, 26].

Method	Neutrals		Ions			
	Aqueous	NAQ	Acetonitrile	DMSO	Methanol	Water
PySCF/M06-2X/TZ	0.69	0.68	5.9	4.2	2.4	5.6
PySCF/M06-2X/6-31G*	0.71	0.65	5.9	4.6	2.1	4.8
PySCF/B3LYP/TZ	0.79	0.69	5.7	3.8	3.1	6.3
PySCF/B3LYP/6-31G*	0.89	0.70	5.8	4.3	2.5	5.4
PySCF/HF/6-31G*	0.84	0.74	6.2	5.4	3.1	3.8
Q-Chem/HF/6-31G*	0.90	0.70		5.4/4.1		2.9/3.9
G03/M05-2X/cc-pVTZ	0.68	0.67	5.9	4.4	2.3	4.6
G03/B3LYP/6-31G*	0.80	0.67	5.7	4.3	2.9	4.9
G03/HF/6-31G*	0.90	0.73	5.9	5.2	2.7	3.4

Table 3: Mean unsigned error (MUE) in solvation free energies (kcal/mol) with different protocols. The detailed protocols for G03 are described in the original SMD paper [26], Table 13. The data for Q-Chem is taken from [71]. In all calculations, the bulk electrostatic contribution is calculated with IEF-PCM. Density fitting is applied for all PySCF calculations. 2346 solvation free energies for neutrals and 332 solvation free energies for ions are used for G03. 2369 solvation free energies and 361 solvation free energies are used for PySCF results. The uncertainty of the reference data are ± 0.2 kcal/mol for neutral solutes and ± 3 kcal/mol for ions. ‘TZ’ is short for def2-TZVPP. ‘NAQ’ is short for nonaqueous.

3.4 Chemical Reactions

3.4.1 Transition State Search

The activation energy, the energy of the transition state relative to the reactants, is a critical parameter for determining the chemical reaction rate. A higher activation energy means the fewer molecules have enough energy to react at a given temperature, leading to a slower reaction. Knowing the activation energy helps in designing and optimizing chemical processes. Transition state searches in quantum chemistry are complex and require careful planning and execution. Computational tools

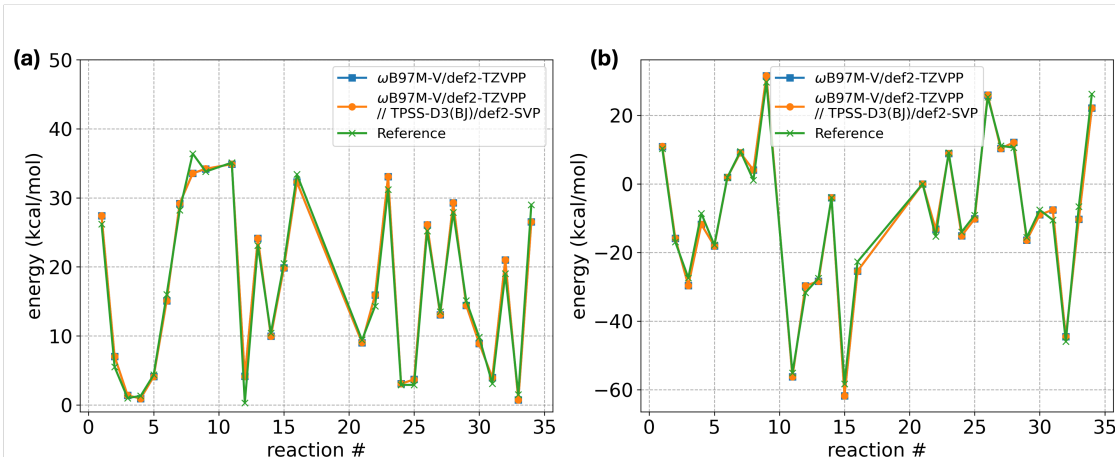


Figure 9: Activation energy (left) and reaction energy (right). ω B97M-V/def2-TZVPP: direct single-point energy calculations with the geometries in Ref. [76] using ω B97M-V functionals and def2-TZVPP basis set. ω B97M-V/def2-TZVPP//TPSS-D3(BJ)/def2-SVP: re-optimize geometries using the same protocols as [76] TPSS-D3(BJ)/def2-SVP and calculate the single-point energy with using ω B97M-V functionals and def2-TZVPP basis set. All the calculations employ (200,1202) grids for transition metals, (99,590) grids for other elements, and density fitting.

play an important role for finding the optimized geometries of the reactants and products, optimizing the transition state, and verifying the transition state. Those computations need a lot of computational resources for calculating SCF, gradient, and Hessian. An efficient quantum chemistry tool can accelerate the finding of the reaction pathway. Successfully finding transition states requires not just computational resources but also a deep understanding of the chemistry involved and experience with the computational methods and software. There are fully or partially automated tools [73, 74, 75] for generating reaction profiles. This work does not provide the automated approaches. But one can build GPU4PySCF in the existing automated workflows.

In this work, we try to reproduce the result in MOBH35 dataset [76]. For simplicity, the reactions involving bimolecular reactants or products are removed [73]. This subset of MOBH35 includes various transition metals Sc, Ti, Mn, Fe in the first row, Nb, Mo, Ru, Rh, Pd in the second row, and Ta, W, Re, Os, Ir, Pt in the third row. The transition metals are generally challenging for computation because they not only requires high-angular momentum GTO integrals, but also need ECP in the large basis set. Ref. [73] tried to modify the geometries of some reactions based on the chemical intuition. However, we are not able to verify the transition states in the modified geometries. We still use the original geometries in [76] as the initial guess for re-optimizing transition state. Both geometry optimizations and transition state optimizations are performed with geomeTRIC [77].

We benchmark the results with ω B97M-V, which is one of the best performer in MOBH35 database [76]. In the dataset, the reference energies are calculated with CCSD(T)/CBS_{w1} + Δ (T)/TZVPP. The barrier heights and reaction energies are presented in Fig. 9. The average absolute deviation (MAD) of the reaction energy we calculated in the subset of MOBH35 is 1.75 kcal/mol, which is similar to MAD (1.7 kcal/mol) of the entire database reported in [76]. The re-optimized transition states are almost identical to the geometries reported in [76]. The difference of the reaction energy between the original geometry and re-optimized geometries is 0.01 kcal/mol in MAD. That verifies the robustness of transition state optimization scheme.

3.4.2 Electrochemical Stability Windows of Electrolytes

To guarantee thermodynamic stability of a lithium-ion battery, it’s critical that the electrochemical potential of the electrodes falls within the electrolyte’s electrochemical stability window, which is the range between the oxidation and reduction potential of the electrolyte [78, 79, 80]. If an anode’s electrochemical potential exceeds the reduction potential of the electrolyte, it can lead to the reduction of the electrolyte; similarly, if a cathode’s electrochemical potential is below the oxidation potential of the electrolyte, it could oxidize the electrolyte, despite that the formation of a passivation layer

can prevent continuous electron transfer and mitigate these effects. A broader electrochemical window signifies a larger operational voltage range, which directly translates to higher energy density and potentially higher power output of the device [81, 82]. Thus, in the study of electrolyte in lithium-ion batteries, electrolyte’s electrochemical window directly influences the selection of solvent molecules in electrolyte for high-performance energy storage solutions.

A straightforward approach to determining the electrochemical stability window of electrolytes involves calculating the gap between HOMO and LUMO [81, 82], but this method has its shortcomings [83]. Thus, our analysis adopts a more accurate thermodynamic approach to define the electrochemical stability of the electrolyte, which relies on calculating the redox potentials rather than molecular orbital gaps [80, 84, 79, 85, 86, 83]. The determination of oxidation and reduction potentials can be rigorously approached through thermodynamic cycles [84, 83]. Besides, for the evaluation of reduction potentials in particular, the introduction of Li^+ ions plays a pivotal role for considering the polarization effects of Li^+ on solvent molecules, which significantly influence the electrochemical behavior of the system, thereby affecting the solvent’s reduction potential [86, 80]. Thus, in this work, we consider the following oxidation and reduction reactions for some specific solvent S



The oxidation and reaction potential of those reactions can be calculated as

$$E_{\text{oxidation}} = (G_{\text{gas}}(\text{oxidized}) - G_{\text{gas}}(\text{initial}) + \Delta G_{\text{solv}}(\text{oxidized}) - \Delta G_{\text{solv}}(\text{initial}))/F - 1.44, \quad (3)$$

$$E_{\text{reduction}} = - (G_{\text{gas}}(\text{reduced}) - G_{\text{gas}}(\text{initial}) + \Delta G_{\text{solv}}(\text{reduced}) - \Delta G_{\text{solv}}(\text{initial}))/F - 1.44, \quad (4)$$

where G_{gas} is the free energy in gas of different species, such as initial, oxidized and reduced species. The solvation energy in Eq. (3) and Eq. (4) of some specific solvent S reads

$$\Delta G_{\text{solv}}(\text{S}) = E_{\text{solv}}(\text{S}) - E_{\text{gas}}(\text{S}), \quad (5)$$

approximated by the energy difference between solvation and gas phase. Subtraction of 1.44 V converts from the absolute electrochemical scale to the Li/Li^+ potential scale, and the value 1.44 V is referred to Ref. [80] for comparisons.

In our study, we calculated the redox potentials of some molecules reported in Ref. [80], and compared our results with those previously reported in the same reference. It should be noted that our calculations were performed using a thermodynamic cycle, with computational parameters that slightly differ from those used in Ref. [80]. As shown in Figure 10, despite differences in the computational approach and parameters from those described in the original literature, the results obtained are remarkably similar. Consequently, we posit that our package possesses the capability to calculate redox potentials, thereby enabling the prediction of electrochemical windows. Considering the computational efficiency, oxidation and reduction potential calculations involves the optimization of geometric structures in vacuum and solvent, as well as the calculation of Hessian in vacuum. The computational acceleration using GPU4PySCF with respect to Q-Chem for these processes has been discussed in previous sections. Consequently, the acceleration ratio for the calculation of oxidation-reduction potential is determined by the lowest acceleration ratio among these individual parts.

3.4.3 Fukui Function, Condensed Fukui Function and ESP

In chemical reaction research, determining the most reactive site on a compound is crucial. Typically, organic chemists make judgments based on knowledge and experience, but quantum chemistry calculations can also be helpful. The widely-used Fukui function [87, 88] predicts reactive sites and comprises three types: f^+ , f^- , and f^0 , for nucleophilic, electrophilic, and radical attacks, respectively. In most applications, the Fukui function is calculated by the finite difference formula [89]. It requires calculations for neutral (N electrons), cationic (N-1 electrons), and anionic (N+1 electrons) states and analyzing the electron density difference. Regions with larger Fukui function values imply greater reactivity at this site.

The condensed Fukui function uses atomic charges to evaluate the reactivity of individual atoms in a molecule more intuitively [90]. By calculating atomic charges for N, N-1, and N+1 electron states, differences between states can be obtained for each atom, enabling reactivity comparisons.

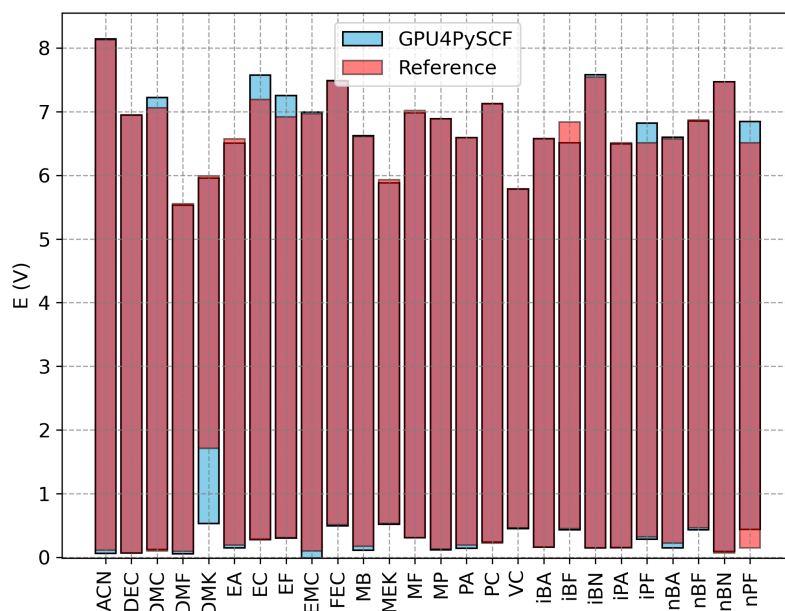


Figure 10: Electrochemical stability windows calculated by GPU4PySCF (blue bars) and from Ref. [80] (red bars). It should be noted that the largest discrepancy lies in the reduction potential of the DMAc molecule. Using the same structure and calculation method as Ref. [80], we get 0.55 V, which differs from the reported results. Computational details are the same as in section 3.1.3.

For example, consider an electrophilic aromatic substitution using Anisole to examine the orientational effect of the directing group. After optimizing the initial structure for the neutral state, single point calculations are performed for N, N-1, and N+1 electron states at the B3LYP/def2-TZVPP level. The Fukui function is obtained from electron density differences in different states and saved to *.cub files, which can be visualized using the open-source tool VMD [91]. Finally, the condensed Fukui function is calculated using CHELPG atomic charges [54].

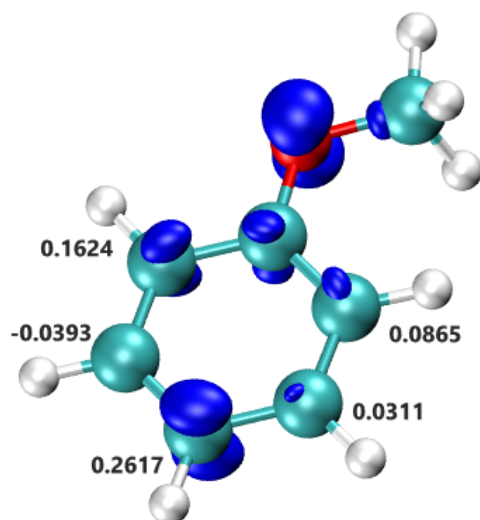


Figure 11: The f^- Fukui function isosurface (isolevel 0.008) of Anisole is shown in blue, and the condensed f^- values are marked near the corresponding carbon atoms.

As the study concentrates on electrophilic substitution sites, only the f^- function is relevant. As shown in Figure 11, the f^- isosurface (isolevel 0.008) occurs mainly on the *ortho*- and *para*-position carbon atoms, and less in the *meta*-position. The condensed Fukui function values also suggest that *ortho* and *para* carbon atoms with higher f^- values are more susceptible to electrophilic substitution reactions. This result is consistent with organic chemistry knowledge that the $-\text{OCH}_3$ group is an *ortho*, *para* director.

The electrostatic potential (ESP) is a commonly used real-space function that provides important information about the interaction sites of compounds. It helps understand the anisotropic distribution of charges on the molecular surface, thereby assisting in inferring possible reaction sites. The wave function obtained from the single-point calculation of compounds using GPU4PySCF can be further exported as the ESP’s *.cub file, which can also be visualized using VMD. The color map of the electrostatic potential for Anisole on the electron density isosurface ($\rho=0.001$) is displayed in Figure 12.

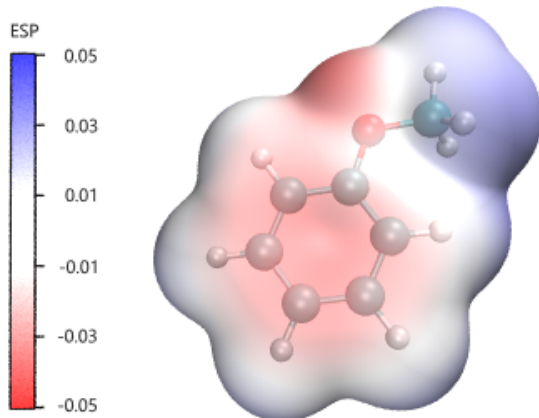


Figure 12: ESP color map for Anisole on the electron density isosurface ($\rho=0.001$).

3.5 Incorporating with Neural Networks

Recently, the application of machine learning has been extended to the quantum chemistry community [92, 93], demonstrating substantial potential to address quantum mechanic problems. The traditional quantum chemistry techniques are still important ingredients of the recipe. Currently, most researchers employ the PySCF routines for accessing the quantum chemistry calculations. This work further offers the same functionalities with GPU acceleration, which is crucial as most deep learning models are running on GPU.

We take Fermionic neural networks [94, 95, 92] as an example. The main philosophy of this research realm, namely neural network-based Quantum Monte Carlo, is to learn highly accurate ground state wavefunctions using a variational Monte Carlo approach. For this approach, atomic orbital (AO) evaluation is an important component that encapsulates traditional chemistry knowledge. Specifically, utilizing the Hartree-Fock AO to ‘pretrain’ neural networks is a common strategy in this research area, as introduced in [92]. This process guides the networks to an acceptable initial state, after which they evolve based on variational principles. However, AO evaluation, invoked at every ‘pretraining’ step, is typically executed on CPUs, which quickly becomes the computational bottleneck of the whole process. Another factor contributing to the inefficiency of this implementation is substantial data transfer between CPU memory and GPU memory, especially for large systems. With GPU4PySCF developed, AO evaluation can be efficiently performed on GPUs and achieves a hundredfold speed improvement as shown in Fig. 13.

Furthermore, GPU4PySCF can be effortlessly integrated with existing neural network frameworks within quantum chemistry and manifest significant acceleration. For instance, by integrating Lap-Net [96] with GPU4PySCF, we achieved approximately three times faster pretraining on Benzene molecule. Beyond pretraining, GPU4PySCF also enables researchers to design and implement neural

network wavefunction ansatzes that incorporate traditional chemistry knowledge, like AO, to better characterize the desired quantum state.

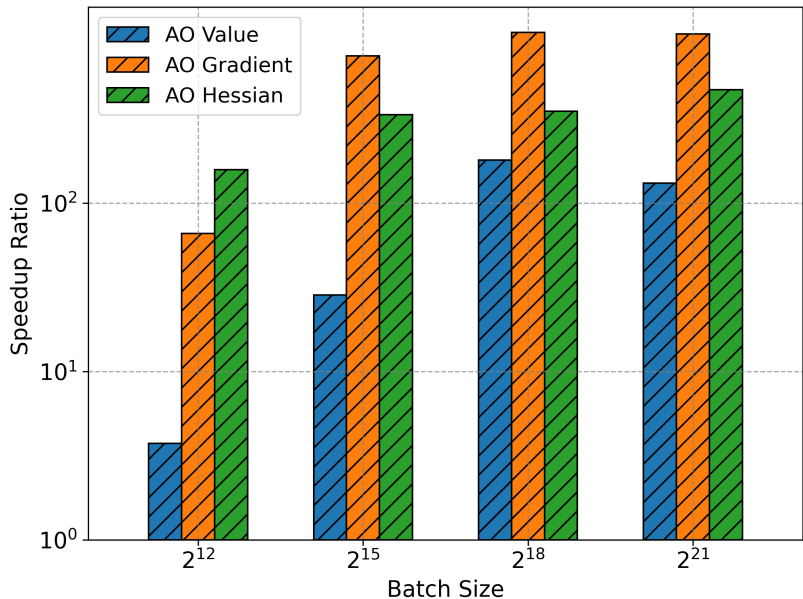


Figure 13: Speedup ratio of A100-80G card over 32-core CPUs in atomic orbital evaluations. Benzene with cc-pVDZ basis is used for benchmark. AO values, gradients and Hessian matrices are chosen for comparison.

4 Community and Ecosystem

PySCF has been a constantly growing open-source community. It has been integrated into various open-source packages as the interface to quantum chemistry. For example,

- **Quantum Computing.** For instance, Google’s quantumlib, IBM’s Qiskit, NVIDIA’s CUDA Quantum, Tencent’s TenCirChem.
- **Deep Learning.** For instance, DeepMind’s FermiNet [92] and DM21, DeepModeling’s DeePKS-kit [97], ByteDance’s DeepSolid and LapNet.
- **Quantum Chemistry Workflows.** For instance, Microsoft’s AiiDA.

Without changing too many code, those works can be accelerated with GPU4PySCF.

4.1 Compatibility with PySCF

GPU4PySCF objects are designed to be compatible with the corresponding PySCF objects. A PySCF object can be converted into a GPU4PySCF object with calling `to_gpu` function when the functionality is implemented on GPU. The `kernel` execution will be accelerated with GPU. On the other hand, a GPU4PySCF object can be converted into a PySCF object with `to_cpu`. Then PySCF features if not accelerated still can be applied. Here is an example

```
1 import pyscf
2 from pyscf import lib
3 from pyscf.dft import rks
4
5 atom = '''
6 0      0.0000000000      -0.0000000000      0.1174000000
```

```

7 H      -0.7570000000    -0.0000000000    -0.4696000000
8 H       0.7570000000     0.0000000000    -0.4696000000
9 '''
10
11 mol = pyscf.M(atom=atom, basis='sto3g')
12 mol.verbose = 4
13 mf = rks.RKS(mol, xc='b3lyp').density_fit()
14
15 # Move PySCF object to GPU
16 mf_GPU = mf.to_gpu()
17
18 # Compute Energy
19 e_dft = mf_GPU.kernel()
20
21 # Compute Gradient
22 g = mf_GPU.nuc_grad_method()
23 g_dft = g.kernel()
24
25 # Compute Hessian
26 h = mf_GPU.Hessian()
27 h.auxbasis_response = 2
28 h_dft = h.kernel()

```

Listing 1: Example of PySCF with GPU support

GPU4PySCF has been meticulously engineered to maintain a high degree of compatibility with PySCF, ensuring that any settings related to accuracy in PySCF are seamlessly applicable in GPU4PySCF. This compatibility is facilitated by the use of matching APIs, which we will explore in Section 4.1. Furthermore, any custom configurations set in PySCF are automatically applied to GPU4PySCF. We have conducted benchmarks on SCF energy, gradient, and Hessian calculations for the water molecule using well-established exchange-correlation functionals across Rungs 2-4 of Jacob’s Ladder [98]. For all examined exchange-correlation functionals, the discrepancies in energy relative to PySCF calculations are less than 10^{-11} Ha, falling below the SCF convergence threshold of 10^{-10} . Moreover, the 2-norm differences for gradients and Hessians are below 10^{-7} Ha/Bohr and 10^{-6} Ha/Bohr², respectively, demonstrating exceptional accuracy and consistency with PySCF.

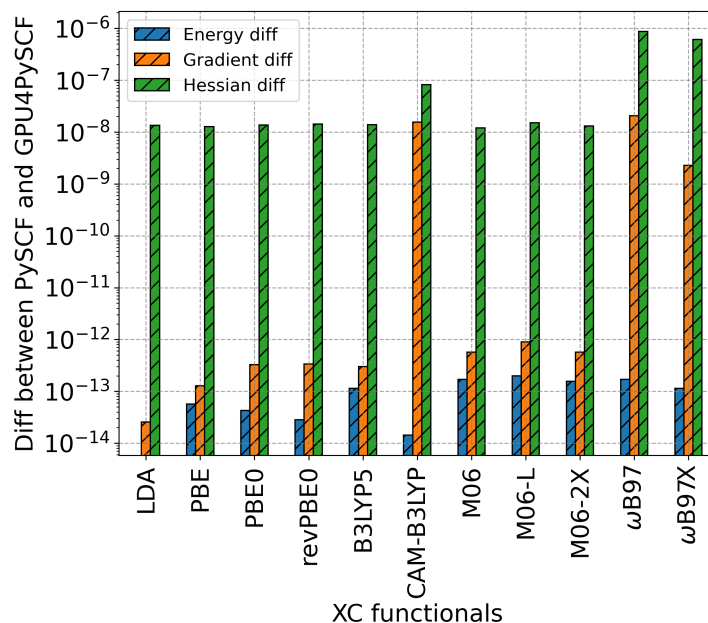


Figure 14: 2-norm of the differences between PySCF and calculations with different DFT functionals for H₂O molecule, grids level = 5, def2-QZVPP atomic basis, def2-universal-JKFIT auxiliary basis, convergence tolerance = 10^{-12} . Different units (Ha for energy, Ha/Bohr for gradient, and Ha/Bohr² for Hessian) are used in y-axis.

4.2 Accelerated Functionalities

Practical DFT calculations not only involve the contributions of foundational DFT but also advanced features such as dispersion correction, solvent effect, ECP, and etc. Those contributions are essential to accurately model the electronic structure. This study adheres to a principle of utilizing existing open-source implementations for these advanced features wherever possible. In scenarios where these implementations are either non-existent or inadequate, we undertake the redevelopment of these features, emphasizing the enhancement of computational efficiency through GPU acceleration, particularly for processes that exhibit slower performance than SCF iterations.

Dispersion corrections are usually much cheaper than DFT self-consistent iterations. And these contributions are calculated with dftd3/simple-dftd3 and dftd4/dftd4 packages without GPU acceleration. Although DFT-D3 has been implemented in Torch and Jax, these implementations will introduce the dependencies of heavy-weight machine learning framework. For robustness, we build simple-dftd3 and dftd4 in GPU4PySCF, and disable OpenMP. Nonlocal corrections are carefully optimized at CUDA level due to its expensive computational cost.

PCM/SMD solvent models require solving Poisson-Boltzmann equation in each SCF iteration. The computational complexity of solving Poisson-Boltzmann equation is cubic-scaling with respect to the number of grids on molecular surface. Although it is cheaper than the quartic-scaling hybrid DFTs, it would be the bottleneck without GPU acceleration. To address this, we have re-engineered the PCM/SMD solvent models for GPU execution using Python, recognizing the potential need for further optimizations to minimize the memory requirements.

ECP contributions are calculated only once in SCF. We reuse PySCF routines on CPU to calculate the ECP contributions in SCF, gradient, and Hessian. Other contributions of DFT are still accelerated with GPU for a molecular system with ECP. The computational cost is acceptable if the system only has a few heavy atoms. We leave this as future work to further accelerate ECP calculations.

The functionalities are summarized in the Table 4.

Method	SCF	Gradient	Hessian
direct SCF	O	GPU	CPU
density fitting	O	O	O
LDA	O	O	O
GGA	O	O	O
mGGA	O	O	O
hybrid	O	O	O
unrestricted	O	O	O
PCM solvent	GPU	GPU	FD
SMD solvent	GPU	GPU	FD
dispersion correction	CPU*	CPU*	FD
nonlocal correlation	O	O	NA
ECP	CPU	CPU	CPU
MP2	GPU	CPU	CPU
CCSD	GPU	CPU	NA

Table 4: A summary of GPU4PySCF functionalities and level of optimizations. ‘O’: carefully optimized for GPU. ‘CPU’: only cpu implementation. ‘GPU’: drop-in replacement or naïve implementation. ‘FD’: use finite-difference gradient to approximate the exact Hessian matrix. ‘NA’: not available. ‘CPU*’: DFTD3 [99]/DFTD4 [100] on CPU.

4.3 Coordinating with other packages

The development of PySCF/GPU4PySCF benefits from the contributions from the open-source community. GPU4PySCF, acting as an extension of PySCF, can be seamlessly incorporated with the open-source packages compatible with PySCF. We highlight the roles of the following fundamental packages working with PySCF/GPU4PySCF.

- BSE [14]: basis set exchange for newly developed basis set.

- GeomeTRIC [77]: geometric optimization and transition state search.
- LibXC [15]: DFT exchange correlation functionals.
- DFTD3 [99]/DFTD4 [100]: dispersion corrections.
- Libcint [101]: fundamental GTO integrals.
- CuPy [102]: A drop-in replacement of Numpy & SciPy for GPU.

5 Limitations and Future development

Quantum chemistry is a field characterized by its complex algorithms and significant computational demands. Accelerated computing technologies, such as GPUs, offer promising ways to tackle these challenges, yet their adoption in quantum chemistry has been slower than in other areas. Although the package delivers significant speedup for various tasks, we admit our implementations are not the optimal. We recognize the current limitations and suggest focused areas for future development to enhance computational efficiency and capabilities.

Wavefunction-based Methods. Beyond traditional mean-field methods, sophisticated wavefunction-based techniques like MP2/RI-MP2 and CCSD(T)/DF-CCSD(T) have seen successful optimization for GPU processing. Our work in GPU4PySCF includes preliminary implementations for MP2 and CCSD methods, demonstrating the feasibility and benefits of using GPUs for these computationally intensive tasks. References such as [103, 104, 105] for MP2/RI-MP2 and [106, 107] for CCSD(T)/DF-CCSD(T) highlight the ongoing efforts and achievements in this area.

Direct SCF & Other Integral Schemes. The Direct Self-Consistent Field (SCF) method has been particularly explored for GPU acceleration due to its minimal memory requirements and scalability with system size, especially when screening techniques are applied. The efficiency of Direct SCF is largely dependent on the integral computation strategy. We utilize Rys quadrature due to its simplicity and versatility in calculating Gaussian Type Orbital (GTO) integrals on GPUs. In the current implementations, this integral scheme is applied up to g functions. While it may not be the fastest, it simplifies the integral program for various operators. The selection of the optimal algorithm for integral computations is crucial for maximizing performance.

Periodic Boundary Conditions (PBC). The PBC module in PySCF extends its capabilities to simulate extended systems using Gaussian basis sets. The main challenges in PBC calculations can be attributed to the infinite number of repeated images caused by periodicity. This leads to high computational costs and numerical instability. Although sophisticated integral screening and evaluation schemes have been implemented in the CPU code, their techniques may not be directly transferable to GPU-based programs. New techniques are under development to refine integral screening and evaluation for GPU-optimized PBC simulations, as mentioned in [108].

Multi-GPUs. The advent of data-center GPUs equipped with multi-GPU configurations significantly enhances data transfer speeds via technologies like NVIDIA’s NVLink or AMD’s xGMI, offering a substantial improvement over traditional CPU-GPU data transfer rates. This technique not only accelerates the computation of quantum chemistry problems but also enables the efficient handling of tasks that require large amounts of GPU memory. Multi-GPU implementations perhaps can achieve greater parallel efficiency than 100% for large molecules.

Mixed Precision. The mixed-precision computing approach capitalizes on the differential computational capabilities of GPUs, balancing single and double precision to optimize performance without compromising accuracy. In the context of Direct SCF algorithms, error estimation via Schwarz inequality allows for the strategic use of precision levels. This technique reduces computational overhead by performing most calculations in single precision in later SCF iterations. Such strategies, as investigated in [109, 110], showcase the ongoing innovation in leveraging GPU capabilities to improve quantum chemistry computations, focusing on maximizing performance and computational efficiency.

Acknowledgements

We thank Dr. John Herbert from The Ohio State University for providing the implementation details of CHELPG charge in Q-Chem 6.1.

A Dataset for Benchmarks

The information of the dataset used in Section 2 is as follow:

Molecule name	Number of atoms	Elements
Vitamin C	20	H,C,O
Inosine	31	H,C,O,N
Bisphenol A	33	H,C,O
Mg Porphin	37	H,C,N,Mg
Penicillin V	42	H,C,O,N,S
Ochratoxin A	45	H,C,O,N,Cl
Cetirizine	52	H,C,O,N,Cl
Tamoxifen	57	H,C,O,N
Raffinose	66	H,C,O
Sphingomyelin	84	H,C,O,N,P
Azadirachtin	95	H,C,O
Taxol	113	H,C,O
Valinomycin	168	H,C,O,N

Table 5: The information of constructed dataset.

B Accuracy for Different XC Functionals

The following calculations are using def2-TZVPP, def2-universal-JKFIT, and (99,590) grids. The energy, gradient, and Hessian difference are measured in 2-norm. The units are Hartree, Hartree/Bohr, and Hartree/Bohr² respectively.

Molecule names	Energy	Gradient	Hessian
Vitamin C	1.30E-06	3.35E-05	1.26E-03
Inosine	5.55E-08	3.44E-05	1.19E-03
Bisphenol A	4.99E-06	6.39E-05	1.51E-03
Mg Porphin	9.23E-06	4.10E-05	1.34E-03
Penicillin V	2.48E-06	1.12E-04	2.18E-03
Ochratoxin A	9.07E-06	9.93E-05	2.33E-03
Cetirizine	1.12E-05	1.30E-04	4.77E-03
Tamoxifen	7.96E-06	8.98E-05	1.77E-03
Raffinose	8.97E-06	2.15E-04	2.96E-03
Sphingomyelin	1.69E-05	1.27E-04	4.10E-03

Table 6: Difference between GPU4PySCF and Q-Chem 6.1, using PBE XC functional.

Molecule names	Energy	Gradient	Hessian
Vitamin C	5.55E-07	2.74E-05	1.04E-03
Inosine	8.17E-07	2.76E-05	9.62E-04
Bisphenol A	4.61E-06	5.21E-05	1.17E-03
Mg Porphin	8.28E-06	3.82E-05	1.20E-03
Penicillin V	1.86E-06	8.76E-05	1.69E-03
Ochratoxin A	7.67E-06	6.80E-05	1.77E-03
Cetirizine	8.40E-06	9.09E-05	3.56E-03
Tamoxifen	4.77E-06	7.09E-05	1.29E-03
Raffinose	6.43E-06	1.62E-04	2.22E-03
Sphingomyelin	1.52E-05	9.50E-05	3.28E-03

Table 7: Difference between GPU4PySCF and Q-Chem 6.1, using B3LYP XC functional.

Molecule names	Energy	Gradient
Vitamin C	1.91E-06	2.45E-05
Inosine	3.32E-06	3.96E-05
Bisphenol A	3.07E-06	6.54E-05
Mg Porphin	5.40E-06	4.38E-05
Penicillin V	3.35E-06	9.26E-05
Ochratoxin A	5.29E-06	7.22E-05
Cetirizine	1.54E-05	7.74E-05
Tamoxifen	3.78E-06	7.54E-05
Raffinose	7.88E-07	1.28E-04
Sphingomyelin	1.76E-05	9.69E-05

Table 8: Difference between GPU4PySCF and Q-Chem 6.1, using ω B97M-V XC functional, Hessian is not supported yet.

C Accuracy for Different Basis Sets

The following calculations are using B3LYP, def2-universal-JKFIT, (99,590) grids. The energy, gradient, and Hessian difference are measured in 2-norm. The units for energy, gradient, and Hessian are Hartree, Hartree/Bohr and Hartree/Bohr² respectively.

Molecule names	Energy	Gradient	Hessian
Vitamin C	1.14E-07	2.65E-05	9.88E-04
Inosine	1.06E-05	2.71E-05	1.03E-03
Bisphenol A	1.28E-06	4.68E-05	1.06E-03
Mg Porphin	6.26E-06	3.97E-05	1.01E-03
Penicillin V	5.18E-05	8.46E-05	1.65E-03
Ochratoxin A	4.72E-05	5.39E-05	1.61E-03
Cetirizine	2.46E-05	9.08E-05	3.18E-03
Tamoxifen	3.86E-05	6.43E-05	1.28E-03
Raffinose	7.80E-05	1.41E-04	2.18E-03
Sphingomyelin	8.46E-05	8.99E-05	3.14E-03

Table 9: 6-31G.

Molecule names	Energy	Gradient	Hessian
Vitamin C	6.31E-07	2.52E-05	9.24E-04
Inosine	1.25E-05	2.61E-05	9.82E-04
Bisphenol A	9.31E-07	4.68E-05	9.89E-04
Mg Porphin	7.19E-06	3.15E-05	9.59E-04
Penicillin V	3.71E-05	8.31E-05	1.52E-03
Ochratoxin A	3.06E-05	5.06E-05	1.50E-03
Cetirizine	2.26E-05	8.67E-05	2.93E-03
Tamoxifen	1.73E-05	6.38E-05	1.27E-03
Raffinose	5.60E-05	1.32E-04	2.13E-03
Sphingomyelin	4.12E-05	8.31E-05	3.13E-03

Table 10: Difference between GPU4PySCF and Q-Chem 6.1, using def2-SVP basis.

Molecule names	Energy	Gradient	Hessian
Vitamin C	5.55E-07	2.74E-05	1.04E-03
Inosine	8.17E-07	2.76E-05	9.62E-04
Bisphenol A	4.61E-06	5.21E-05	1.17E-03
Mg Porphin	8.28E-06	3.82E-05	1.20E-03
Penicillin V	1.86E-06	8.76E-05	1.69E-03
Ochratoxin A	7.67E-06	6.80E-05	1.77E-03
Cetirizine	8.40E-06	9.09E-05	3.56E-03
Tamoxifen	4.77E-06	7.09E-05	1.29E-03
Raffinose	6.43E-06	1.62E-04	2.22E-03
Sphingomyelin	1.52E-05	9.50E-05	3.28E-03

Table 11: Difference between GPU4PySCF and Q-Chem 6.1, using def2-TZVPP basis set.

Molecule names	Energy	Gradient	Hessian
Vitamin C	5.23E-07	2.79E-05	1.04E-03
Inosine	9.11E-07	2.75E-05	9.63E-04
Bisphenol A	4.78E-06	5.23E-05	1.17E-03
Mg Porphin	8.28E-06	3.83E-05	1.21E-03
Penicillin V	4.23E-06	8.74E-05	1.70E-03
Ochratoxin A	3.08E-06	6.81E-05	1.77E-03
Cetirizine	6.85E-06	9.10E-05	3.56E-03
Tamoxifen	3.36E-06	7.12E-05	1.29E-03
Raffinose	2.57E-06	1.61E-04	2.22E-03
Sphingomyelin	6.98E-06	9.57E-05	3.29E-03

Table 12: Difference between GPU4PySCF and Q-Chem 6.1, using def2-TZVPD basis set.

D Accuracy for Solvent Models

All of the following calculations are using B3LYP, def2-TZVPP, def2-universal-JKFIT, (99,590) grids. The energy, gradient, and Hessian difference are measured in 2-norm. The corresponding units are Hartree, Hartree/Bohr and Hartree/Bohr². Since Q-Chem 6.1 do not support PCM models for the density fitting scheme, we use the regular SCF scheme as the reference. The following tables reflects the density fitting errors.

Molecule names	Energy	Gradient	Hessian
Vitamin C	1.38E-04	1.20E-04	1.14E-03
Inosine	2.08E-04	1.39E-04	1.18E-03
Bisphenol A	1.99E-04	1.78E-04	1.36E-03
Mg Porphin	1.40E-04	1.35E-04	1.33E-03

Table 13: Difference between GPU4PySCF and Q-Chem 6.1, using C-PCM solvent model.

Molecule names	Energy	Gradient	Hessian
Vitamin C	1.38E-04	1.20E-04	1.14E-03
Inosine	2.08E-04	1.39E-04	1.18E-03
Bisphenol A	1.99E-04	1.78E-04	1.36E-03
Mg Porphin	1.40E-04	1.35E-04	1.33E-03

Table 14: Difference between GPU4PySCF and Q-Chem 6.1, using IEF-PCM solvent model

References

- [1] Ivan S. Ufimtsev and Todd J. Martínez. Quantum chemistry on graphical processing units. 1. strategies for two-electron integral evaluation. *Journal of Chemical Theory and Computation*, 4(2):222–231, 2008. PMID: 26620654.
- [2] Koji Yasuda. Two-electron integral evaluation on the graphics processor unit. *Journal of Computational Chemistry*, 29(3):334–342, 2008.
- [3] Giuseppe M. J. Barca, Colleen Bertoni, Laura Carrington, Dipayan Datta, Nuwan De Silva, J. Emiliano Deustua, Dmitri G. Fedorov, Jeffrey R. Gour, Anastasia O. Gunina, Emilie Guidez, Taylor Harville, Stephan Irle, Joe Ivanic, Karol Kowalski, Sarom S. Leang, Hui Li, Wei Li, Jesse J. Lutz, Ilias Magoulas, Joani Mato, Vladimir Mironov, Hiroya Nakata, Buu Q. Pham, Piotr Piecuch, David Poole, Spencer R. Pruitt, Alistair P. Rendell, Luke B. Roskop, Klaus Ruedenberg, Tosaporn Sattasathuchana, Michael W. Schmidt, Jun Shen, Lyudmila Slipchenko, Masha Sosonkina, Vaibhav Sundriyal, Ananta Tiwari, Jorge L. Galvez Vallejo, Bryce Westheimer, Marta Wloch, Peng Xu, Federico Zahariev, and Mark S. Gordon. Recent developments in the

general atomic and molecular electronic structure system. *The Journal of Chemical Physics*, 152(15):154102, April 2020.

- [4] Federico Zahariev, Peng Xu, Bryce M. Westheimer, Simon Webb, Jorge Galvez Vallejo, Ananta Tiwari, Vaibhav Sundriyal, Masha Sosonkina, Jun Shen, George Schoendorff, Megan Schlinsog, Tosaporn Sattasathuchana, Klaus Ruedenberg, Luke B. Roskop, Alistair P. Rendell, David Poole, Piotr Piecuch, Buu Q. Pham, Vladimir Mironov, Joani Mato, Sam Leonard, Sarom S. Leang, Joe Ivanic, Jackson Hayes, Taylor Harville, Karthik Gururangan, Emilie Guidez, Igor S. Gerasimov, Christian Friedl, Katherine N. Ferreras, George Elliott, Dipayan Datta, Daniel Del Angel Cruz, Laura Carrington, Colleen Bertoni, Giuseppe M. J. Barca, Melisa Alkan, and Mark S. Gordon. The general atomic and molecular electronic structure system (gamess): Novel methods on novel architectures. *Journal of Chemical Theory and Computation*, 19(20):7031–7055, 2023. PMID: 37793073.
- [5] Evgeny Epifanovsky, Andrew TB Gilbert, Xintian Feng, Joonho Lee, Yuezhi Mao, Narbe Mardirossian, Pavel Pokhilko, Alec F White, Marc P Coons, Adrian L Dempwolff, et al. Software for the frontiers of quantum chemistry: An overview of developments in the q-chem 5 package. *The Journal of chemical physics*, 155(8), 2021.
- [6] Daniel G. A. Smith, Lori A. Burns, Andrew C. Simmonett, Robert M. Parrish, Matthew C. Schieber, Raimondas Galvelis, Peter Kraus, Holger Kruse, Roberto Di Remigio, Asem Alenaizan, Andrew M. James, Susi Lehtola, Jonathon P. Misiewicz, Maximilian Scheurer, Robert A. Shaw, Jeffrey B. Schriber, Yi Xie, Zachary L. Glick, Dominic A. Sirianni, Joseph Senan O’Brien, Jonathan M. Waldrop, Ashutosh Kumar, Edward G. Hohenstein, Benjamin P. Pritchard, Bernard R. Brooks, III Schaefer, Henry F., Alexander Yu. Sokolov, Konrad Patkowski, III De-Prince, A. Eugene, Uğur Bozkaya, Rollin A. King, Francesco A. Evangelista, Justin M. Turney, T. Daniel Crawford, and C. David Sherrill. PSI4 1.4: Open-source software for high-throughput quantum chemistry. *The Journal of Chemical Physics*, 152(18):184108, 05 2020.
- [7] Ádám Rák and György Cserey. The brush algorithm for two-electron integrals on gpu. *Chemical Physics Letters*, 622:92–98, 2015.
- [8] Gábor János Tornai, István Ladjánszki, Ádám Rák, Gergely Kis, and György Cserey. Calculation of quantum chemical two-electron integrals by applying compiler technology on gpu. *Journal of Chemical Theory and Computation*, 15(10):5319–5331, 2019. PMID: 31503475.
- [9] Stefan Seritan, Christoph Bannwarth, B. Scott Fales, Edward G. Hohenstein, Sara I. L. Kokkila-Schumacher, Nathan Luehr, Jr. Snyder, James W., Chenchen Song, Alexey V. Titov, Ivan S. Ufimtsev, and Todd J. Martínez. TeraChem: Accelerating electronic structure and ab initio molecular dynamics with graphical processing units. *The Journal of Chemical Physics*, 152(22):224110, 06 2020.
- [10] Stefan Seritan, Keiran Thompson, and Todd J. Martínez. Terachem cloud: A high-performance computing service for scalable distributed gpu-accelerated electronic structure calculations. *Journal of Chemical Information and Modeling*, 60(4):2126–2137, 2020. PMID: 32267693.
- [11] Madushanka Manathunga, Hasan Metin Aktulga, Andreas W. Götz, and Kenneth M. Jr. Merz. Quantum mechanics/molecular mechanics simulations on nvidia and amd graphics processing units. *Journal of Chemical Information and Modeling*, 63(3):711–717, 2023. PMID: 36720086.
- [12] Yipu Miao and Kenneth M. Jr. Merz. Acceleration of high angular momentum electron repulsion integrals and integral derivatives on graphics processing units. *Journal of Chemical Theory and Computation*, 11(4):1449–1462, 2015. PMID: 26574356.
- [13] Madushanka Manathunga, Chi Jin, Vinícius Wilian D. Cruzeiro, Yipu Miao, Dawei Mu, Kamesh Arumugam, Kristopher Keipert, Hasan Metin Aktulga, Kenneth M. Jr. Merz, and Andreas W. Götz. Harnessing the power of multi-gpu acceleration into the quantum interaction computational kernel program. *Journal of Chemical Theory and Computation*, 17(7):3955–3966, 2021. PMID: 34062061.

- [14] Benjamin P. Pritchard, Doaa Altarawy, Brett Didier, Tara D. Gibson, and Theresa L. Windus. New basis set exchange: An open, up-to-date resource for the molecular sciences community. *Journal of Chemical Information and Modeling*, 59(11):4814–4820, 2019. PMID: 31600445.
- [15] Susi Lehtola, Conrad Steigemann, Micael JT Oliveira, and Miguel AL Marques. Recent developments in libxc—a comprehensive library of functionals for density functional theory. *SoftwareX*, 7:1–5, 2018.
- [16] Stefan Grimme, Jan Gerit Brandenburg, Christoph Bannwarth, and Andreas Hansen. Consistent structures and interactions by density functional theory with small atomic orbital basis sets. *The Journal of Chemical Physics*, 143(5):054107, 08 2015.
- [17] Marcel Müller, Andreas Hansen, and Stefan Grimme. ω B97X-3c: A composite range-separated hybrid DFT method with a molecule-optimized polarized valence double- ζ basis set. *The Journal of Chemical Physics*, 158(1):014103, 01 2023.
- [18] Stefan Grimme, Andreas Hansen, Sebastian Ehlert, and Jan-Michael Mewes. r2SCAN-3c: A “Swiss army knife” composite electronic-structure method. *The Journal of Chemical Physics*, 154(6):064103, 02 2021.
- [19] Qiming Sun, Xing Zhang, Samraghi Banerjee, Peng Bao, Marc Barbry, Nick S. Blunt, Nikolay A. Bogdanov, George H. Booth, Jia Chen, Zhi-Hao Cui, Janus J. Eriksen, Yang Gao, Sheng Guo, Jan Hermann, Matthew R. Hermes, Kevin Koh, Peter Koval, Susi Lehtola, Zhendong Li, Junzi Liu, Narbe Mardirossian, James D. McClain, Mario Motta, Bastien Mussard, Hung Q. Pham, Artem Pulkin, Wirawan Purwanto, Paul J. Robinson, Enrico Ronca, Elvira R. Sayfutyarova, Maximilian Scheurer, Henry F. Schurkus, James E. T. Smith, Chong Sun, Shi-Ning Sun, Shiv Upadhyay, Lucas K. Wagner, Xiao Wang, Alec White, James Daniel Whitfield, Mark J. Williamson, Sebastian Wouters, Jun Yang, Jason M. Yu, Tianyu Zhu, Timothy C. Berkelbach, Sandeep Sharma, Alexander Yu. Sokolov, and Garnet Kin-Lic Chan. Recent developments in the PySCF program package. *The Journal of Chemical Physics*, 153(2):024109, 07 2020.
- [20] Eric D. Hermes, Khachik Sargsyan, Habib N. Najm, and Judit Zádor. Geometry optimization speedup through a geodesic approach to internal coordinates. *The Journal of Chemical Physics*, 155(9):094105, 09 2021.
- [21] Philipp Pracht, David F. Grant, and Stefan Grimme. Comprehensive assessment of gfn tight-binding and composite density functional theory methods for calculating gas-phase infrared spectra. *Journal of Chemical Theory and Computation*, 16(11):7044–7060, 2020. PMID: 33054183.
- [22] Christopher A. White and Martin Head-Gordon. A J matrix engine for density functional theory calculations. *The Journal of Chemical Physics*, 104(7):2620–2629, 02 1996.
- [23] Maurizio Cossi, Nadia Rega, Giovanni Scalmani, and Vincenzo Barone. Energies, structures, and electronic properties of molecules in solution with the c-pcm solvation model. *Journal of Computational Chemistry*, 24(6):669–681, 2003.
- [24] Vincenzo Barone and Maurizio Cossi. Quantum calculation of molecular energies and energy gradients in solution by a conductor solvent model. *The Journal of Physical Chemistry A*, 102(11):1995–2001, 1998.
- [25] E. Cancés, B. Mennucci, and J. Tomasi. A new integral equation formalism for the polarizable continuum model: Theoretical background and applications to isotropic and anisotropic dielectrics. *The Journal of Chemical Physics*, 107(8):3032–3041, 08 1997.
- [26] Aleksandr V. Marenich, Christopher J. Cramer, and Donald G. Truhlar. Universal solvation model based on solute electron density and on a continuum model of the solvent defined by the bulk dielectric constant and atomic surface tensions. *The Journal of Physical Chemistry B*, 113(18):6378–6396, 2009. PMID: 19366259.

- [27] Simon Boothroyd, Pavan Kumar Behara, Owen C. Madin, David F. Hahn, Hyesu Jang, Vytautas Gapsys, Jeffrey R. Wagner, Joshua T. Horton, David L. Dotson, Matthew W. Thompson, Jessica Maat, Trevor Gokey, Lee-Ping Wang, Daniel J. Cole, Michael K. Gilson, John D. Chodera, Christopher I. Bayly, Michael R. Shirts, and David L. Mobley. Development and benchmarking of open force field 2.0.0: The sage small molecule force field. *Journal of Chemical Theory and Computation*, May 2023.
- [28] Joshua T. Horton, Simon Boothroyd, Jeffrey Wagner, Joshua A. Mitchell, Trevor Gokey, David L. Dotson, Pavan Kumar Behara, Venkata Krishnan Ramaswamy, Mark Mackey, John D. Chodera, Jamshed Anwar, David L. Mobley, and Daniel J. Cole. Open force field bespokefit: Automating bespoke torsion parametrization at scale. *Journal of Chemical Information and Modeling*, 62(22):5622–5633, November 2022.
- [29] Yudong Qiu, Daniel G. A. Smith, Chaya D. Stern, Mudong Feng, Hyesu Jang, and Lee-Ping Wang. Driving torsion scans with wavefront propagation. *The Journal of Chemical Physics*, 152(24):244116, June 2020.
- [30] Yudong Qiu, Daniel G. A. Smith, Simon Boothroyd, Hyesu Jang, David F. Hahn, Jeffrey Wagner, Caitlin C. Bannan, Trevor Gokey, Victoria T. Lim, Chaya D. Stern, Andrea Rizzi, Bryon Tjanaka, Gary Tresadern, Xavier Lucas, Michael R. Shirts, Michael K. Gilson, John D. Chodera, Christopher I. Bayly, David L. Mobley, and Lee-Ping Wang. Development and benchmarking of open force field v1.0.0—the parsley small-molecule force field. *Journal of Chemical Theory and Computation*, 17(10):6262–6280, October 2021.
- [31] Katarina Roos, Chuanjie Wu, Wolfgang Damm, Mark Reboul, James M. Stevenson, Chao Lu, Markus K. Dahlgren, Sayan Mondal, Wei Chen, Lingle Wang, Robert Abel, Richard A. Friesner, and Edward D. Harder. Opls3e: Extending force field coverage for drug-like small molecules. *Journal of Chemical Theory and Computation*, 15(3):1863–1874, March 2019.
- [32] Junmei Wang, Romain M. Wolf, James W. Caldwell, Peter A. Kollman, and David A. Case. Development and testing of a general amber force field. *Journal of Computational Chemistry*, 25(9):1157–1174, 2004.
- [33] Bai Xue, Qingyi Yang, Qiaochu Zhang, Xiao Wan, Dong Fang, Xiaolu Lin, Guangxu Sun, Gianpaolo Gobbo, Fenglei Cao, Alan M. Mathiowetz, Benjamin J. Burke, Robert A. Kumpf, Brajesh K. Rai, Geoffrey P. F. Wood, Frank C. IV Pickard, Junmei Wang, Peiyu Zhang, Jian Ma, Yide Alan Jiang, Shuhao Wen, Xinjun Hou, Junjie Zou, and Mingjun Yang. Development and comprehensive benchmark of a high-quality amber-consistent small molecule force field with broad chemical space coverage for molecular modeling and free energy calculation. *Journal of Chemical Theory and Computation*, December 2023.
- [34] William L Jorgensen and Julian Tirado-Rives. The opl3 [optimized potentials for liquid simulations] potential functions for proteins, energy minimizations for crystals of cyclic peptides and crambin. *Journal of the American Chemical Society*, 110(6):1657–1666, 1988.
- [35] Huai Sun. Compass: an ab initio force-field optimized for condensed-phase applications overview with details on alkane and benzene compounds. *The Journal of Physical Chemistry B*, 102(38):7338–7364, 1998.
- [36] Michael H Abraham, Adam Ibrahim, Andreas M Zissimos, Yuan H Zhao, John Comer, and Derek P Reynolds. Application of hydrogen bonding calculations in property based drug design. *Drug Discovery Today*, 7(20):1056–1063, 2002.
- [37] Andrea N Bootsma, Analise C Doney, and Steven E Wheeler. Predicting the strength of stacking interactions between heterocycles and aromatic amino acid side chains. *Journal of the American Chemical Society*, 141(28):11027–11035, 2019.
- [38] Chao Lu, Chuanjie Wu, Delaram Ghoreishi, Wei Chen, Lingle Wang, Wolfgang Damm, Gregory A Ross, Markus K Dahlgren, Ellery Russell, Christopher D Von Bargen, et al. Opls4: Improving force field accuracy on challenging regimes of chemical space. *Journal of chemical theory and computation*, 17(7):4291–4300, 2021.

- [39] Bethany Halford. The path to paxlovid, 2022.
- [40] Frans B Van Duijneveldt, Jeanne GCM van Duijneveldt-van de Rijdt, and Joop H van Lenthe. State of the art in counterpoise theory. *Chemical Reviews*, 94(7):1873–1885, 1994.
- [41] Roman M Balabin. Enthalpy difference between conformations of normal alkanes: Intramolecular basis set superposition error (bsse) in the case of n-butane and n-hexane. *The Journal of chemical physics*, 129(16), 2008.
- [42] Pavel Hobza and Klaus Müller-Dethlefs. *Non-covalent interactions: theory and experiment*, volume 2. Royal Society of Chemistry, 2010.
- [43] Jianhui Wang, Yuki Yamada, Keitaro Sodeyama, Ching Hua Chiang, Yoshitaka Tateyama, and Atsuo Yamada. Superconcentrated electrolytes for a high-voltage lithium-ion battery. *Nat. Commun.*, 7(1):12032, 2016.
- [44] Shuru Chen, Jianming Zheng, Donghai Mei, Kee Sung Han, Mark H Engelhard, Wengao Zhao, Wu Xu, Jun Liu, and Ji-Guang Zhang. High-voltage lithium-metal batteries enabled by localized high-concentration electrolytes. *Adv Mater*, 30(21):1706102, 2018.
- [45] Zhiao Yu, Hansen Wang, Xian Kong, William Huang, Yuchi Tsao, David G Mackanic, Kecheng Wang, Xinchang Wang, Wenxiao Huang, Snehashis Choudhury, et al. Molecular design for electrolyte solvents enabling energy-dense and long-cycling lithium metal batteries. *Nat. Energy*, 5(7):526–533, 2020.
- [46] Chibueze V Amanchukwu, Zhiao Yu, Xian Kong, Jian Qin, Yi Cui, and Zhenan Bao. A new class of ionically conducting fluorinated ether electrolytes with high electrochemical stability. *J. Am. Chem. Soc.*, 142(16):7393–7403, 2020.
- [47] Oleg Borodin, Liumin Suo, Mallory Gobet, Xiaoming Ren, Fei Wang, Antonio Faraone, Jing Peng, Marco Olguin, Marshall Schroeder, Michael S Ding, et al. Liquid structure with nano-heterogeneity promotes cationic transport in concentrated electrolytes. *ACS nano*, 11(10):10462–10471, 2017.
- [48] Kang Xu, Yiufai Lam, Sheng S Zhang, T Richard Jow, and Timothy B Curtis. Solvation sheath of li+ in nonaqueous electrolytes and its implication of graphite/electrolyte interface chemistry. *J. Phys. Chem. C*, 111(20):7411–7421, 2007.
- [49] Oleg Borodin, Julian Self, Kristin A Persson, Chunsheng Wang, and Kang Xu. Uncharted waters: super-concentrated electrolytes. *Joule*, 4(1):69–100, 2020.
- [50] Yuki Yamada, Jianhui Wang, Seongjae Ko, Eriko Watanabe, and Atsuo Yamada. Advances and issues in developing salt-concentrated battery electrolytes. *Nature Energy*, 4(4):269–280, 2019.
- [51] Xiang Chen, Xue-Qiang Zhang, Hao-Ran Li, and Qiang Zhang. Cation- solvent, cation- anion, and solvent- solvent interactions with electrolyte solvation in lithium batteries. *Batter. Supercaps*, 2(2):128–131, 2019.
- [52] Li Zhang and Yuhui Chen. Electrolyte solvation structure as a stabilization mechanism for electrodes. *Energy Mater*, 1(1):100004, 2021.
- [53] Heejoon Moon, Ryoichi Tatara, Toshihiko Mandai, Kazuhide Ueno, Kazuki Yoshida, Naoki Tachikawa, Tomohiro Yasuda, Kaoru Dokko, and Masayoshi Watanabe. Mechanism of li ion desolvation at the interface of graphite electrode and glyme–li salt solvate ionic liquids. *J. Phys. Chem. C*, 118(35):20246–20256, 2014.
- [54] Curt M Breneman and Kenneth B Wiberg. Determining atom-centered monopoles from molecular electrostatic potentials. the need for high sampling density in formamide conformational analysis. *Journal of Computational Chemistry*, 11(3):361–373, 1990.
- [55] Thomas A Manz. Seven confluence principles: a case study of standardized statistical analysis for 26 methods that assign net atomic charges in molecules. *RSC advances*, 10(72):44121–44148, 2020.

- [56] John M Herbert, Leif D Jacobson, Ka Un Lao, and Mary A Rohrdanz. Rapid computation of intermolecular interactions in molecular and ionic clusters: Self-consistent polarization plus symmetry-adapted perturbation theory. *Physical Chemistry Chemical Physics*, 14(21):7679–7699, 2012.
- [57] A van Bondi. van der waals volumes and radii. *The Journal of physical chemistry*, 68(3):441–451, 1964.
- [58] Harald Günther. *NMR spectroscopy: basic principles, concepts and applications in chemistry*. John Wiley & Sons, 2013.
- [59] Shipeng Luo, Xiao Zhang, Yu Zheng, Klaus Harms, Lili Zhang, and Eric Meggers. Enantioselective alkynylation of aromatic aldehydes catalyzed by a sterically highly demanding chiral-at-rhodium lewis acid. *The Journal of Organic Chemistry*, 82(17):8995–9005, 2017.
- [60] Joanna Siuda, Waldemar Perdoch, Bartłomiej Mazela, and Magdalena Zborowska. Catalyzed reaction of cellulose and lignin with methyltrimethoxysilane—ft-ir, ¹³c nmr and ²⁹si nmr studies. *Materials*, 12(12):2006, 2019.
- [61] Andrew M. Teale, Ola B. Lutnæs, Trygve Helgaker, David J. Tozer, and Jürgen Gauss. Benchmarking density-functional theory calculations of NMR shielding constants and spin-rotation constants using accurate coupled-cluster calculations. *The Journal of Chemical Physics*, 138(2):024111, 01 2013.
- [62] Fritz London. Théorie quantique des courants interatomiques dans les combinaisons aromatiques. *J. phys. radium*, 8(10):397–409, 1937.
- [63] Minghong Yuan, Yong Zhang, Zhi Qu, Yunlong Xiao, and Wenjian Liu. Sublinear scaling quantum chemical methods for magnetic shieldings in large molecules. *The Journal of Chemical Physics*, 150(15), 2019.
- [64] Sunghwan Kim, Jie Chen, Tiejun Cheng, Asta Gindulyte, Jia He, Siqian He, Qingliang Li, Benjamin A Shoemaker, Paul A Thiessen, Bo Yu, Leonid Zaslavsky, Jian Zhang, and Evan E Bolton. PubChem 2023 update. *Nucleic Acids Research*, 51(D1):D1373–D1380, 10 2022.
- [65] Anna Pomogaeva and Daniel M. Chipman. New implicit solvation models for dispersion and exchange energies. *The Journal of Physical Chemistry A*, 117(28):5812–5820, 2013. PMID: 23799302.
- [66] Andreas Klamt. The cosmo and cosmo-rs solvation models. *WIREs Computational Molecular Science*, 8(1):e1338, 2018.
- [67] Adrian W. Lange and John M. Herbert. A smooth, nonsingular, and faithful discretization scheme for polarizable continuum models: The switching/Gaussian approach. *The Journal of Chemical Physics*, 133(24):244111, 12 2010.
- [68] Daniel A. Liotard, Gregory D. Hawkins, Gillian C. Lynch, Christopher J. Cramer, and Donald G. Truhlar. Improved methods for semiempirical solvation models. *Journal of Computational Chemistry*, 16(4):422–440, 1995.
- [69] Elric Engelage, Nils Schulz, Flemming Heinen, Stefan M. Huber, Donald G. Truhlar, and Christopher J. Cramer. Refined smd parameters for bromine and iodine accurately model halogen-bonding interactions in solution. *Chemistry – A European Journal*, 24(60):15983–15987, 2018.
- [70] Aleksandr V Marenich, Casey P Kelly, Jason D Thompson, Gregory D Hawkins, Candee C Chambers, David J Giesen, Paul Winget, Christopher J Cramer, and Donald G Truhlar. Minnesota solvation database (mnsol) version 2012, 2020.
- [71] John M. Herbert. Dielectric continuum methods for quantum chemistry. *WIREs Computational Molecular Science*, 11(4):e1519, 2021.

- [72] Haiyang Zhang, Yang Jiang, Hai Yan, Ziheng Cui, and Chunhua Yin. Comparative assessment of computational methods for free energy calculations of ionic hydration. *Journal of Chemical Information and Modeling*, 57(11):2763–2775, 2017. PMID: 29039666.
- [73] Sebastian Dohm, Markus Bursch, Andreas Hansen, and Stefan Grimme. Semiautomated transition state localization for organometallic complexes with semiempirical quantum chemical methods. *Journal of Chemical Theory and Computation*, 16(3):2002–2012, 2020. PMID: 32074450.
- [74] Tom A. Young, Joseph J. Silcock, Alistair J. Sterling, and Fernanda Duarte. autode: Automated calculation of reaction energy profiles— application to organic and organometallic reactions. *Angewandte Chemie International Edition*, 60(8):4266–4274, 2021.
- [75] Leif D. Jacobson, Art D. Bochevarov, Mark A. Watson, Thomas F. Hughes, David Rinaldo, Stephan Ehrlich, Thomas B. Steinbrecher, S. Vaitheeswaran, Dean M. Philipp, Mathew D. Halls, and Richard A. Friesner. Automated transition state search and its application to diverse types of organic reactions. *Journal of Chemical Theory and Computation*, 13(11):5780–5797, 2017. PMID: 28957627.
- [76] Mark A. Iron and Trevor Janes. Evaluating transition metal barrier heights with the latest density functional theory exchange–correlation functionals: The mobh35 benchmark database. *The Journal of Physical Chemistry A*, 123(17):3761–3781, 2019. PMID: 30973722.
- [77] Lee-Ping Wang and Chencheng Song. Geometry optimization made simple with translation and rotation coordinates. *The Journal of Chemical Physics*, 144(21):214108, 06 2016.
- [78] Cleber FN Marchiori, Rodrigo P Carvalho, Mahsa Ebadi, Daniel Brandell, and C Moyses Araujo. Understanding the electrochemical stability window of polymer electrolytes in solid-state batteries from atomic-scale modeling: the role of li-ion salts. *Chemistry of Materials*, 32(17):7237–7246, 2020.
- [79] Samuel A Delp, Oleg Borodin, Marco Olguin, Claire G Eisner, Joshua L Allen, and T Richard Jow. Importance of reduction and oxidation stability of high voltage electrolytes and additives. *Electrochimica Acta*, 209:498–510, 2016.
- [80] David S Hall, Ahmed Eldesoky, ER Logan, Erin Marie Tonita, Xiaowei Ma, and JR Dahn. Exploring classes of co-solvents for fast-charging lithium-ion cells. *Journal of The Electrochemical Society*, 165(10):A2365, 2018.
- [81] John B Goodenough and Kyu-Sung Park. The li-ion rechargeable battery: a perspective. *Journal of the American Chemical Society*, 135(4):1167–1176, 2013.
- [82] John B Goodenough. Rechargeable batteries: challenges old and new. *Journal of Solid State Electrochemistry*, 16:2019–2029, 2012.
- [83] Pekka Peljo and Hubert H Girault. Electrochemical potential window of battery electrolytes: the homo–lumo misconception. *Energy & Environmental Science*, 11(9):2306–2309, 2018.
- [84] Oleg Borodin, Wishvender Behl, and T Richard Jow. Oxidative stability and initial decomposition reactions of carbonate, sulfone, and alkyl phosphate-based electrolytes. *The Journal of Physical Chemistry C*, 117(17):8661–8682, 2013.
- [85] Tingzheng Hou, Guang Yang, Nav Nidhi Rajput, Julian Self, Sang-Won Park, Jagjit Nanda, and Kristin A Persson. The influence of fec on the solvation structure and reduction reaction of lipf6/ec electrolytes and its implication for solid electrolyte interphase formation. *Nano Energy*, 64:103881, 2019.
- [86] Oleg Borodin, Marco Olguin, Carrie E Spear, Kenneth W Leiter, and Jaroslaw Knap. Towards high throughput screening of electrochemical stability of battery electrolytes. *Nanotechnology*, 26(35):354003, 2015.
- [87] Tian Lu and Feiwu Chen. Multiwfn: A multifunctional wavefunction analyzer. *Journal of computational chemistry*, 33(5):580–592, 2012.

- [88] Robert G Parr and Weitao Yang. Density functional approach to the frontier-electron theory of chemical reactivity. *Journal of the American Chemical Society*, 106(14):4049–4050, 1984.
- [89] Roberto Flores-Moreno, Junia Melin, JV Ortiz, and Gabriel Merino. Efficient evaluation of analytic fukui functions. *The Journal of chemical physics*, 129(22), 2008.
- [90] Weitao Yang and Wilfried J Mortier. The use of global and local molecular parameters for the analysis of the gas-phase basicity of amines. *Journal of the American Chemical Society*, 108(19):5708–5711, 1986.
- [91] William Humphrey, Andrew Dalke, and Klaus Schulten. Vmd: visual molecular dynamics. *Journal of molecular graphics*, 14(1):33–38, 1996.
- [92] D. Pfau, J.S. Spencer, A.G. de G. Matthews, and W.M.C. Foulkes. Ab-initio solution of the many-electron schrödinger equation with deep neural networks. *Phys. Rev. Research*, 2:033429, 2020.
- [93] James Kirkpatrick, Brendan McMorro, David H. P. Turban, Alexander L. Gaunt, James S. Spencer, Alexander G. D. G. Matthews, Annette Obika, Louis Thiry, Meire Fortunato, David Pfau, Lara Román Castellanos, Stig Petersen, Alexander W. R. Nelson, Pushmeet Kohli, Paula Mori-Sánchez, Demis Hassabis, and Aron J. Cohen. Pushing the frontiers of density functionals by solving the fractional electron problem. *Science*, 374(6573):1385–1389, 2021.
- [94] Jiequn Han, Linfeng Zhang, and E Weinan. Solving many-electron schrödinger equation using deep neural networks. *Journal of Computational Physics*, 399:108929, 2019.
- [95] Jan Hermann, Zeno Schätzle, and Frank Noé. Deep-neural-network solution of the electronic schrödinger equation. *Nature Chemistry*, 12(10):891–897, 2020.
- [96] Ruichen Li, Haotian Ye, Du Jiang, Xuelan Wen, Chuwei Wang, Zhe Li, Xiang Li, Di He, Ji Chen, Weiluo Ren, and Liwei Wang. A computational framework for neural network-based variational Monte Carlo with Forward Laplacian. *Nature Machine Intelligence*, 6(2):209–219, February 2024.
- [97] Yixiao Chen, Linfeng Zhang, Han Wang, and Weinan E. Deepks: A comprehensive data-driven approach toward chemically accurate density functional theory. *Journal of Chemical Theory and Computation*, 17(1):170–181, 2021.
- [98] John P. Perdew and Karla Schmidt. Jacob’s ladder of density functional approximations for the exchange-correlation energy. *AIP Conference Proceedings*, 577(1):1–20, 07 2001.
- [99] Stefan Grimme, Stephan Ehrlich, and Lars Goerigk. Effect of the damping function in dispersion corrected density functional theory. *Journal of Computational Chemistry*, 32(7):1456–1465, 2011.
- [100] Eike Caldeweyher, Sebastian Ehlert, Andreas Hansen, Hagen Neugebauer, Sebastian Spicher, Christoph Bannwarth, and Stefan Grimme. A generally applicable atomic-charge dependent London dispersion correction. *The Journal of Chemical Physics*, 150(15):154122, 04 2019.
- [101] Qiming Sun. Libcint: An efficient general integral library for gaussian basis functions. *Journal of Computational Chemistry*, 36:1664–1671, 2015.
- [102] Ryosuke Okuta, Yuya Unno, Daisuke Nishino, Shohei Hido, and Crissman Loomis. Cupy: A numpy-compatible library for nvidia gpu calculations. In *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)*, 2017.
- [103] Ryan Stocks, Elise Palethorpe, and Giuseppe M. J. Barca. High-performance multi-gpu analytic ri-mp2 energy gradients. *Journal of Chemical Theory and Computation*, 0(0):null, 0. PMID: 38456899.
- [104] Dmytro Bykov and Thomas Kjaergaard. The gpu-enabled divide-expand-consolidate ri-mp2 method (dec-ri-mp2). *Journal of Computational Chemistry*, 38(4):228–237, 2017.

- [105] JaeHyuk Kwack, Colleen Bertoni, Buu Pham, and Jeff Larkin. Performance of the ri-mp2 fortran kernel of gamess on gpus via directive-based offloading with math libraries. In *Accelerator Programming Using Directives: 6th International Workshop, WACCPD 2019, Denver, CO, USA, November 18, 2019, Revised Selected Papers 6*, pages 91–113. Springer, 2020.
- [106] Dipayan Datta and Mark S. Gordon. Accelerating coupled-cluster calculations with gpus: An implementation of the density-fitted ccSD(t) approach for heterogeneous computing architectures using openmp directives. *Journal of Chemical Theory and Computation*, 19(21):7640–7657, 2023. PMID: 37878756.
- [107] Jinsung Kim, Ajay Panyala, Bo Peng, Karol Kowalski, P. Sadayappan, and Sriram Krishnamoorthy. Scalable heterogeneous execution of a coupled-cluster model with perturbative triples. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–15, 2020.
- [108] Qiming Sun. Various integral estimations and screening schemes for extended systems in pyscf, 2023.
- [109] Andrey Asadchev and Mark S. Gordon. Mixed-precision evaluation of two-electron integrals by rys quadrature. *Computer Physics Communications*, 183(8):1563–1567, 2012.
- [110] Nathan Luehr, Ivan S. Ufimtsev, and Todd J. Martínez. Dynamic precision for electron repulsion integral evaluation on graphical processing units (gpus). *Journal of Chemical Theory and Computation*, 7(4):949–954, 2011. PMID: 26606344.