ANNUAL REVIEWS

Annual Review of Physical Chemistry

Interactive Quantum Chemistry Enabled by Machine Learning, Graphical Processing Units, and Cloud Computing

Umberto Raucci,[1,2,3] Hayley Weir,[1,2] Sukolsak Sakshuwong,[4] Stefan Seritan,[1,2,5] Colton B. Hicks,[1,2] Fabio Vannucci,[6] Francesco Rea,[6] and Todd J. Martínez[1,2]

[1]Stanford PULSE Institute, SLAC National Accelerator Laboratory, Menlo Park, California, USA

[2]Department of Chemistry, Stanford University, Stanford, California, USA; email: Todd.Martinez@stanford.edu

[3]Current affiliation: Italian Institute of Technology, Genova, Italy

[4]Department of Management Science and Engineering, Stanford University, Stanford, California, USA

[5]Current affiliation: Sandia National Laboratories, Livermore, California, USA

[6]Italian Institute of Technology, Genova, Italy

Annu. Rev. Phys. Chem. 2023. 74:313–36

First published as a Review in Advance on February 7, 2023

The *Annual Review of Physical Chemistry* is online at physchem.annualreviews.org

https://doi.org/10.1146/annurev-physchem-061020-053438

Keywords

interactive quantum chemistry, natural user interfaces, chemical structure recognition, cloud computing, graphical processing units, extended reality

Abstract

Modern quantum chemistry algorithms are increasingly able to accurately predict molecular properties that are useful for chemists in research and education. Despite this progress, performing such calculations is currently unattainable to the wider chemistry community, as they often require domain expertise, computer programming skills, and powerful computer hardware. In this review, we outline methods to eliminate these barriers using cutting-edge technologies. We discuss the ingredients needed to create accessible platforms that can compute quantum chemistry properties in real time, including graphical processing units–accelerated quantum chemistry in the cloud, artificial intelligence–driven natural molecule input methods, and extended reality visualization. We end by highlighting a series of exciting

applications that assemble these components to create uniquely interactive platforms for computing and visualizing spectra, 3D structures, molecular orbitals, and many other chemical properties.

1. INTRODUCTION

Hartree and Fock could have scarcely imagined in the 1930s that their equations would eventually be solved in a matter of seconds for molecules with hundreds of atoms, much less that these calculations could be launched through voice commands using a virtual assistant such as Amazon's Alexa. Advances in computing power, artificial intelligence (AI)/machine learning (ML), networking communication, and scientific algorithms have made this a reality. We are in the midst of an exciting transition from an era where chemical computation was restricted to experts willing to struggle with algorithm/code development, input deck formatting, and 2D visualizations to one where students can launch calculations by speaking to their smartphones and then visualize the results in three dimensions using augmented reality (AR). The rapid advances in all of the above areas promise to make interactions with chemical computation even more ubiquitous and natural in the future. It is hard to predict how this will change chemical research and chemical education, but we eagerly anticipate the day when chemical computation is as ubiquitous and frictionless as arithmetic on handheld calculators.

Increasingly powerful computational hardware, affordable smart devices, and versatile programming languages laid the ground for this technological shift. Graphical processing units (GPUs) had a remarkable effect on computational capabilities by allowing data to be efficiently processed in parallel (1, 2). This resulted in a dramatic reduction in computation time for parallelizable tasks. Numerous computational chemistry software packages have now been developed to exploit the massively parallel architecture of modern graphics cards, both for empirical force-field molecular dynamics (3–6) and for quantum chemistry and/or ab initio molecular dynamics (7–10). This has enabled empirical force-field simulations with $O(10^8)$ atoms (11) and both quantum chemical and ab initio molecular dynamics simulations with $O(10^3)$ atoms (12–16). It has also enabled quantum mechanics/molecular mechanics condensed phase simulations of photochemical dynamics (including electronic excited states) with $O(10^2)$ atom quantum mechanics regions for aggregate times of up to 1 ns (17–21). Additionally, it has opened up the possibility for molecular dynamics–driven reaction discovery on ground (22) and excited (23) electronic states.

Meanwhile, the advent of cloud computing provided dynamically scalable virtual resources as a service over the web (24). This offers a nearly barrierless platform for accessing powerful computational hardware and managing data. GPUs and cloud computing have been fundamental to the growth of AI by allowing scientists and engineers to efficiently train and deploy complex ML models without purchasing and housing expensive hardware. In the last decade, ML methods have been widely applied in chemistry to predict new synthetic routes (25, 26), discover new drugs (27, 28), and design new materials (29). In computational chemistry, ML-based potentials have emerged as a possible solution to attain the accuracy of density functional theory (DFT)-based potentials with the computational efficiency of force fields (30, 31) and there are continuing efforts to use ML to find better functionals to increase the accuracy of DFT (32–34).

ML-based technologies can also enhance the interaction between chemistry students and researchers, and computational chemistry tools. Quantum chemistry software typically requires coding expertise to install and run the software, specialized knowledge to prepare input files and parse the output, and access to high-performance computing (HPC) resources. This results in a

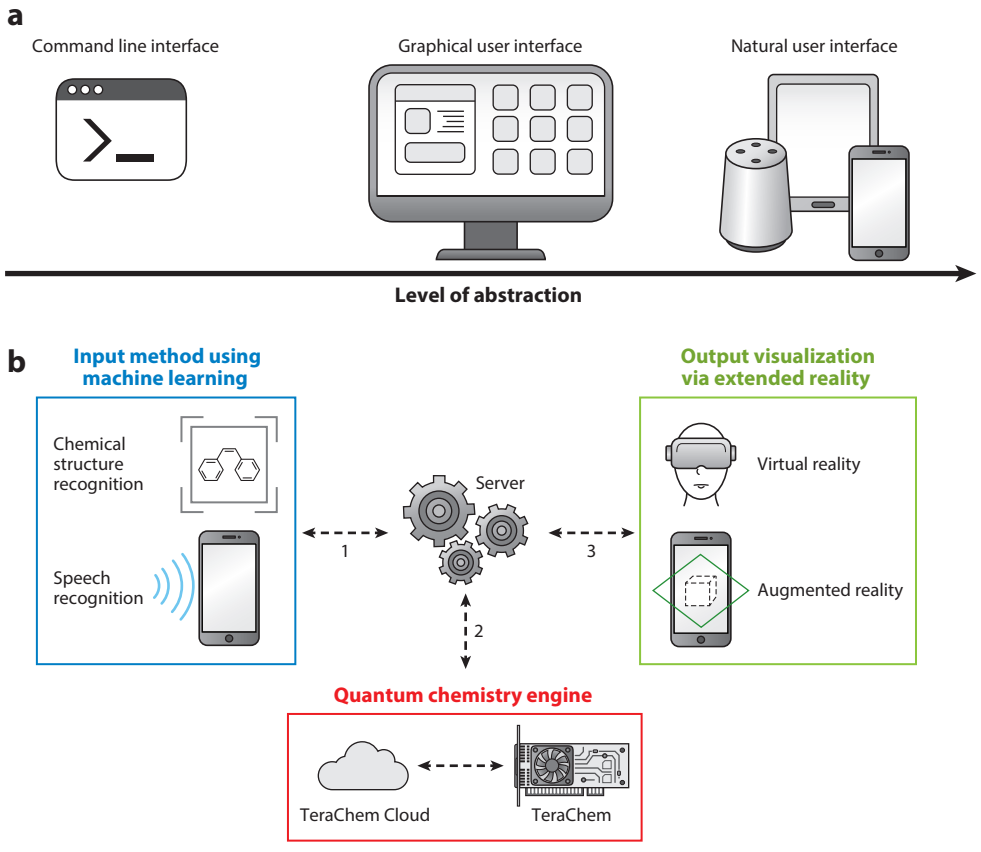

Figure 1

(*a*) Different modes of input for quantum chemistry packages. (*b*) Workflow showing how natural user interfaces, graphical processing unit–accelerated quantum chemistry in the cloud, and extended reality can work together to create accessible, user-friendly interactions with quantum chemistry. (①) A user takes a picture of a chemical structure or prompts a vocal request, which is passed to a server. The server calls external application programming interfaces (i.e., ChemPix) to predict the SMILES representation and retrieves the molecular information (i.e., atomic positions). (②) The server sends a request to perform the calculation in the cloud. (③) The results are finally communicated to the user via virtual reality or augmented reality. Abbreviation: SMILES, simplified molecular input line-entry system.

high barrier to entry for nonexpert users. Calculations are most commonly launched and analyzed using the command-line or limited graphical user interfaces (35) (**Figure 1*a***). The recent ML-driven advances in image and speech recognition offer new routes to input molecular structures and initiate calculations. Moreover, versatile programming languages (e.g., Python) allow these technologies to be readily interfaced with cloud-computing resources and accessible platforms such as mobile devices, web browsers, and virtual assistants. Applying these methods can lower or completely eliminate the barrier to interact with quantum chemistry software.

The development of such natural user interfaces (NUIs), which can listen and see, offers an exciting new avenue toward a virtual chemical laboratory. A typical workflow of a NUI is shown schematically in **Figure 1*b***. A user takes a picture of a chemical structure with a smart device or prompts a vocal request to a virtual assistant, which is then passed to a server that feeds external application programming interfaces (APIs). These APIs translate the image or the vocal input to a

machine-readable format and return it to the server. The server can then call a web database (e.g., PubChem; 36) to retrieve the molecular structure (atomic positions in Cartesian coordinates). The server can then launch a quantum chemistry calculation in the cloud. Finally, the server returns the results to the device, which relays them to the user via voice, standard display, or virtual reality (VR) or AR.

In this review, we summarize the progress to date in this field, from the enabling technology of GPU-accelerated quantum chemistry to the various interfaces and workflows layered on top of this and natural user input schemes that have been enabled. An overview of GPU-accelerated, cloud-based quantum chemistry is provided in Section 2, intuitive molecule input mechanisms driven by AI are described in Section 3, and Section 4 covers molecular visualization in VR or AR. Finally, in Section 5, we discuss how these cutting-edge tools can be combined to build accessible, user-friendly services for human interaction with computational chemistry (see 37).

2. THE QUANTUM CHEMISTRY ENGINE

2.1. GPU-Accelerated Quantum Chemistry

Solution of the electronic structure problem, i.e., quantum chemistry, is a key enabling technology in computational chemistry. The resulting electronic energies and forces allow for optimization of molecular geometries and ab initio molecular dynamics to directly model time-resolved experiments (38) or to explore potential energy surfaces, including the possibility for bond rearrangements (22, 39). Even for large-scale problems involving many atoms where ab initio molecular dynamics may be computationally infeasible, the quantum chemistry toolkit provides the data needed to fit empirical force fields (40). Unfortunately, the expense of quantum chemistry often presents a major computational bottleneck. Traditional ab initio approaches are often hampered by the evaluation of the many two-electron repulsion integrals needed to represent the fundamental Coulomb interactions. As the number of these integrals scales with the fourth power of molecular size, a naive treatment will lead to at best $O(N^4)$ scaling of computational effort (where N is the number of electronic basis functions, which is generally linear in the number of atoms). However, many of these integrals are numerically negligible (for example, those corresponding to the interaction of electrons that are far apart in the molecule). This led to the introduction of screening approaches that identify the negligible integrals beforehand and avoid computing them entirely (41–43). Screening reduces the scaling of computational effort attributable to integral handling for the simplest self-consistent field (SCF) approaches (Hartree–Fock theory or DFT) to $O(N^3)$, which is still problematic. The other aspect of SCF approaches that consumes a large share of computational resources is linear algebra associated with matrices of the size of the molecular basis set (i.e., $N \times N$), also leading to $O(N^3)$ scaling.

Traditional central processing units (CPUs) have been used to solve the electronic structure problem since the advent of quantum chemistry. As these CPUs evolved they have focused on providing performance for general-purpose tasks needed by most users. Unfortunately, quantum chemistry has not been among the most popular of these general-purpose tasks, so it is not surprising that CPUs have not always been optimal for the requirements of quantum chemistry. Furthermore, the advances in CPU processing power due to increasing clock speeds (i.e., Moore's law) have slowed considerably since the turn of the century. Fortunately, the twenty-first century witnessed the rise of special-purpose GPUs aimed at the videogame market. These GPUs were designed for the demands of videogames—fast data movement to get millions of pixels on the screen at frame rates of 60 Hz, fast matrix computations to rotate and stretch objects in a 3D scene, efficient data parallelism to apply transformations to millions of pixels at a time, and efficient collision detection and ray-tracing to achieve realistic lighting effects.

As videogames are primarily physics-based simulations, it should not be surprising that the GPU architecture maps well onto the requirements of quantum chemistry. The data parallelism and efficient linear algebra can be put to use immediately to (*a*) compute many integrals in parallel, (*b*) contract these into the needed matrix operators, and (*c*) diagonalize matrices to find eigenvalues and eigenvectors.

We do not go into the technical details of this mapping of quantum chemistry algorithms to GPU architectures (2, 7, 8, 10, 44–49) but instead only highlight a few key points. Our efforts to utilize GPUs started with videogame consoles such as the Sony PlayStation2 (PS2). These efforts were challenging because videogame consoles use proprietary architectures and closed software ecosystems, with very little generally available documentation. Nevertheless, we were able to execute quantum chemistry programs [a locally modified version of GAMESS (General Atomic and Molecular Electronic Structure System); 50] on multiple PS2 consoles in parallel (51). Nvidia's introduction of CUDA in 2008 simplified GPU programming considerably, and this was the catalyst for increased progress in the area, including applications of GPUs to both quantum chemistry and empirical force field–based molecular dynamics (52).

To utilize GPUs effectively, one should follow a few key principles. Equations for important quantities should be grouped by type, so that one can leverage data parallelism most effectively. In the context of two-electron integrals, this means that one needs to group them by the angular momenta of the basis functions involved in the integral (44–46). As an example, the code required for (ss|ss) integrals is different from that for (ss|sp) integrals and therefore one should have different kernels (i.e., GPU subroutines) for each of these. This is easily accomplished by sorting the basis functions by angular momentum, which leads to the desired grouping of the integrals. One then needs to specify how GPU threads should walk through the set of two-electron integrals, generating the integrals and contracting them with density matrix elements. The amount of fast memory (registers) on the GPU is limited, and this can make high-angular-momentum integrals difficult to handle. To cope with the limited memory, one is often forced to split the computation of high-angular-momentum integral classes into subgroupings, for example, $(d_{xx}d_{xx}|dd)$ and $(d_{yy}d_{yy}|dd)$, instead of handling all the (dd|dd) integrals with a single kernel (47). This implies redundant computation (as integrals of the same angular momentum class for the same atoms will involve some reused intermediates), but the calculation is then possible (whereas it would not be if the number of available registers was insufficient). Often, the bottleneck in GPU computations is not the arithmetic operations but rather the data movement. Thus, it can often be the case that redundant work has little effect on performance.

Another key aspect of efficient GPU algorithms is the avoidance of any logic. Ideally, integral screening is arranged in a way that completely avoids any consideration of the negligible integrals. This turns out to be rather straightforward: One sorts the pairs of basis function indices representing either of the two electrons in the integral so that the corresponding charge densities, i.e., $\chi_i(r)\chi_j(r)$, are ordered by decreasing self-repulsion. Charge densities with very small self-repulsion (usually arising because the atom centers are very distant) are then discarded completely (pruned). When a charge density is pruned, the corresponding one-electron integrals are also discarded (53).

Finally, the most affordable GPUs are typically capable of only limited double-precision arithmetic. This motivates the use of mixed and dynamic precision schemes (44, 54, 55), where the smallest integrals are treated in single precision and only the largest integrals are stored and manipulated in double precision. Even when full double-precision arithmetic is available (e.g., in the GPUs designed for HPC), it can still be advantageous to adopt dynamic and mixed precision schemes, since single-precision numbers consume half of the memory (and thus can be moved at twice the speed). Similar conclusions have also recently been reached for CPUs (56, 57).

Application of the above principles to quantum chemistry algorithms on GPUs has led to speedups of one to three orders of magnitude (44–47). At least one order of magnitude of this speedup comes from the improved performance of GPUs (which are 10–15 times faster for both floating-point operations and memory bandwidth). The remainder of the acceleration comes from improvements in the algorithms (which might also be beneficial on CPUs, although this has not been investigated in detail).

2.2. Quantum Chemistry in the Cloud

Cloud computing provides a mechanism to rent computational resources on demand without incurring the costs of purchasing and maintaining HPC resources. Commercial cloud platforms often provide the most affordable access to specialty hardware such as GPUs, field-programmable gate arrays, and tensor processing units. Combined with containerization technologies (58, 59), entire scientific workflows can be deployed reproducibly at scale across a plethora of cutting-edge hardware. This has led to considerable interest from the scientific community, including chemistry (60–66).

In the context of quantum chemistry, there have been several efforts to use cloud computing. In one mode of operation, cloud resources are rented and used for batch computing similar to traditional HPC clusters managed by job schedulers such as SLURM (Simple Linux Utility for Resource Management). The first computational chemistry calculations on public cloud resources (60, 61) were carried out in this way. Alternatively, one can take full advantage of the flexibility allowed by the on-demand nature of cloud computing by deploying entire workflows that act more like traditional web services. In this model, the computing service (*a*) provides an interface for users to submit jobs; (*b*) farms out those jobs on cloud resources, scaling the number of workers up and down as needed; and (*c*) returns results to the user. This latter strategy has been adopted by projects such as TeraChem Cloud (63) and OpenEye's Orion platform (64).

We illustrate how cloud computing can be used to provide on-demand HPC calculations by reviewing the TeraChem Cloud (TCC) framework, which is shown schematically in **Figure 2**. TCC provides GPU-accelerated electronic structure calculations as a service and is organized into

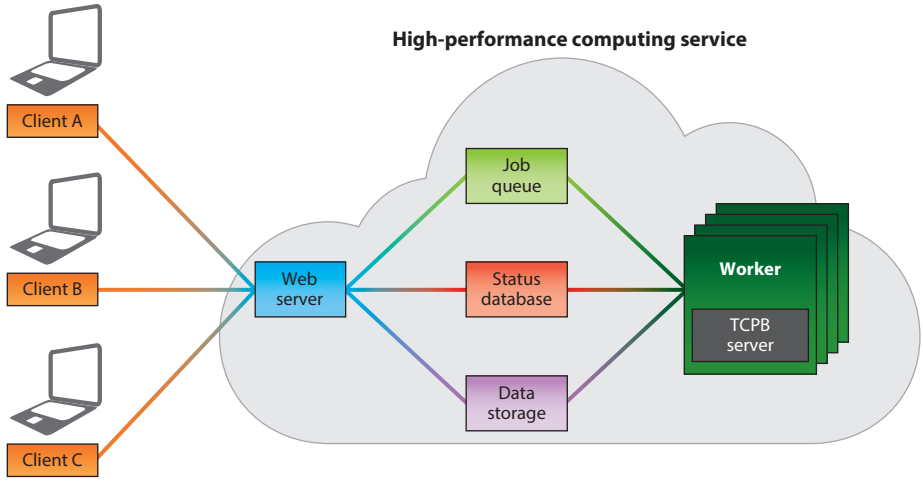

Figure 2

Overview of the TeraChem Cloud framework. Figure adapted with permission from Reference 63; copyright 2020 American Chemical Society. Abbreviation: TCPB, TeraChem Protocol Buffer.

three main layers. The first layer is a client-facing web server that handles job submission, status, and retrieval. The second layer is built from components that manage the state of the calculations: a job queue, status database, and long-term data storage for completed jobs. Finally, the third layer is composed of backend workers that compute the requested electronic structure calculation. In our case, this is done with an instance of a TeraChem Protocol Buffer server (7) running on up to 16 GPUs on a single node, but the TCC framework can in principle use any quantum chemistry package with a Python interface.

The encapsulated structure of the TCC framework is designed to be scalable. In principle, additional web servers could be launched if there was heavy client traffic, or more backend workers could be added if the queue was full and cloud resources were available or advantageously priced. Combining the scalable nature of cloud computing with the GPU-accelerated single-node performance of the TeraChem package results in a platform that can provide real-time electronic structure calculations for many users on systems containing hundreds of atoms. This capability, when combined with intuitive ways to input, display, and interact with chemical data, serves as the foundation for a host of applications that greatly increase the accessibility of quantum chemistry.

3. NATURAL MOLECULE INPUT METHODS USING MACHINE LEARNING

Currently, the standard molecule input method for quantum chemistry packages consists of XYZ or PDB files, which contain Cartesian coordinates of the 3D structure. These files are not always straightforward to generate or interpret without visualization software. In this section, we discuss two alternative molecule input mechanisms designed to be user-friendly and intuitive: sketches of hand-drawn chemical structures and voice commands. Both techniques use ML to automatically extract information from natural human input and to translate this information to the appropriate quantum chemistry software input format.

3.1. Chemical Structure Recognition

Diagrams of skeletal structure are the universal language of organic chemistry. They describe the complete 2D structure of a molecule (including stereochemistry) in an exceptionally simple way by representing carbon and hydrogen atoms implicitly. Any reasonably sized molecule or chemical reaction can be quickly sketched out in this way with a pen and paper or a mouse, allowing humans to efficiently communicate the arrangement of atoms and bonds. Because of this, skeletal formulas serve as perhaps the most natural molecule input method for quantum chemistry software.

Although diagrams of chemical structures are easily decoded by humans, they are not machine readable and so cannot be directly input into computational chemistry software. Converting images of chemical structures to computer-readable representations, termed optical chemical structure recognition (OCSR), has been a goal for three decades (67–69). Until recently, OCSR systems were rule based and consisted of a series of handcrafted algorithms for identifying lines, edges, and characters to compile molecular graphs (70–76). In the past few years, however, the rise of deep learning led to a shift toward data-driven approaches, with neural networks learning end-to-end mappings from images of chemical structures to machine-readable representations (77–81). ChemGrapher is an interesting case that takes ideas from classical and modern approaches to create a modular pipeline of neural networks that break down the recognition task into three subtasks: semantic segmentation, segment classification, and graph building (77).

Most OCSR tools are developed with the goal of mining the literature for chemical structures and so focus on recognizing ChemDraw-type images. Since the topic of this article is creating barrierless ways for scientists to perform quantum chemistry calculations, we are concerned with

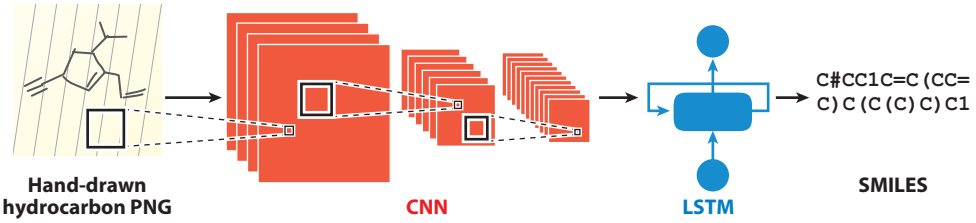

Figure 3

Image-to-SMILES neural network for hand-drawn hydrocarbon recognition consisting of a CNN encoder and a LSTM decoder. Figure adapted with permission from Reference 82. Abbreviations: CNN, convolutional neural network; LSTM, long short-term memory; PNG, portable network graphics; SMILES, simplified molecular input line-entry system.

designing the most natural input method possible. With this in mind, we focus on recognizing photographs of hand-drawn chemical structures, which can be sketched by any chemist and require no specialized software or cumbersome procedures. This increases the difficulty of the recognition task considerably, as large variability of molecule drawing styles, backgrounds, and picture quality is introduced.

Using deep learning methods, we recently developed ChemPix for offline hand-drawn hydrocarbon recognition (82). We chose to represent molecules as simplified molecular input line-entry system (SMILES) strings (83) so that we could directly apply natural language processing (NLP) networks. A neural image captioning network consisting of a convolutional neural network encoder and a long short-term memory (LSTM) decoder with attention and beam search was repurposed to translate photographs of hand-drawn chemical structures to SMILES representations (**Figure 3**). Transformer models (84) have now replaced LSTMs as the network of choice for NLP tasks and should likely be explored for future expansions of ChemPix. Our image-to-SMILES network was trained with a synthetic data set of augmented and degraded images of RDKit hydrocarbon structures and fine-tuned with the addition of a small subset of real-world-photographed hand-drawn structures. By forming a committee of the trained neural networks, we achieved over 85% accuracy for the top three predicted structures using only 400 real-world training samples. The output SMILES labels can be converted to an XYZ file using OpenBabel (85).

Shortly after ChemPix was released, Mathpix (86), which originally developed software for converting images of mathematical equations to LaTeX code, released similar software that can recognize molecules (including heteroatoms) with reasonable accuracy. A template script to perform TCC calculations starting from a picture of a hand-drawn chemical structure is available (87). In Section 5.2, we discuss how these chemical structure recognition tools have been combined with other components described in this article to create MolAR (88), an accessible, user-friendly, and interactive quantum chemistry phone application.

3.2. Speech Recognition

Since speech is one of the most natural methods of human communication, voice-enabled technologies hold great promise as intuitive user interfaces (89, 90). In the last decade, huge strides have been made in the development of efficient voice user interfaces (VUIs) thanks to progress in NLP (91). NLP uses AI methods to interpret speech and text (92) for virtual assistants, chatbots, language translators, and spam filters.

NLP algorithms have exciting potential for chemistry applications; for example, voice-activated, hands-free software could be used to read and write electronic notebooks while

performing experiments. In computational chemistry, vocal prompts could be used to input the commands or as support for programming tasks; Hocky, White, and colleagues (93, 94) recently showed that Codex (95) can generate code for chemical applications from natural language prompts. Further applications of NLP in chemistry are certainly promising. However, chemical language is even richer in ambiguities, synonyms, and context-dependent interpretation than most spoken languages are. This can lead to challenges in chemical adaptations of NLP algorithms designed for the spoken or written word.

In the following section, we describe how to design a basic VUI to launch molecular calculations with TCC and perform single-point energy calculations at a standard level of theory (e.g., PBE0/6-31G*). As with chemical structure recognition described above, this VUI must recognize the name of a molecule and convert it to a machine-readable format (e.g., SMILES).

Extended discussions of speech recognition can be found in the literature (96). For our purposes here, the speech recognition algorithm returns a text transcription of a recorded audio signal. If the vocal prompt contains just the molecular name, it can be directly converted to the corresponding SMILES string using public online databases (e.g., PubChem). The molecular structure can then be retrieved and passed along with the method for the calculation to be performed in the cloud. The result of the calculation can be vocally communicated back to the user using text-to-speech APIs (a template script to perform TCC calculations via voice control is available; 87). Vocal input can also facilitate user interaction in VR. We recently combined vocal prompts with VR in InteraChem (97) using Unity's built-in voice recognition module (see Section 4.1).

The common challenges of voice recognition APIs include variation in the pronunciation of words, accents, and background noise. In chemistry applications, recognizing chemical jargon adds another level of difficulty. Common speech recognition algorithms (e.g., Google Voice API; 98) can recognize only commonly employed chemical names such as water or ammonia, although their ability to understand complex IUPAC names can be improved by designing a custom vocabulary. A more detailed discussion of how to build a robust VUI is provided in Section 5.1. Below, we describe the use of virtual assistants in computational chemistry, including ChemVox (66), an Alexa skill for voice-controlled quantum chemistry.

4. OUTPUT VISUALIZATION VIA EXTENDED REALITY

4.1. Virtual Reality

Molecules are inherently 3D objects; therefore, it should come as no surprise that computational chemists have long been searching for a visualization solution better than standard monitors and displays. Even as early as the 1990s, immersive VR environments such as the CAVE (cave automatic virtual environment) (99) system were being explored as an alternate way to visualize 3D graphics. CAVE-like systems utilize a series of projectors to display stereoscopic images on multiple walls of the room. The user wears a pair of 3D stereoscopic glasses, providing the illusion that a 3D object is floating in the air before them. Although full CAVE setups are expensive and difficult to install, a single stereoscopic display is more affordable, portable, and still provides some 3D visualization capabilities. Despite the promise of a more intuitive visual experience for users, these displays have seen only limited adoption among molecular visualizers.

However, the ability to visualize molecules in 3D has risen dramatically in the past decade due to the mass availability of VR headsets. These head-mounted displays contain two high-resolution screens that serve the same purpose as the spectroscopic glasses in the CAVE setup: By presenting two slightly different perspectives, the user interprets the two separate images as a single 3D scene. The VR headsets are typically paired with a set of handheld controllers that allow the user to manipulate the 3D scene. Due to the popularity of VR headsets in gaming, the hardware is

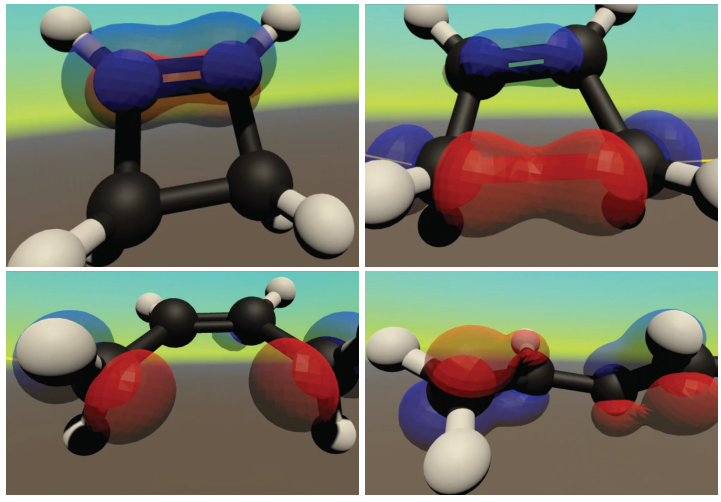

Figure 4

Several snapshots from an ab initio interactive molecular dynamics simulation of the ring opening of cyclobutene in the InteraChem visualizer, with the highest occupied molecular orbital visualized. Figure reproduced with permission from Reference 97; copyright 2021 American Chemical Society.

relatively affordable. Additionally, the existence and ease-of-use of game engines (e.g., Unity and Unreal Engine) simplify the development of new applications for these peripherals.

As a result, a host of molecular visualizers for VR headsets have been developed over the past few years. The well-known visualizer VMD (visual molecular dynamics) (100) has VR capabilities through the VRPN (Virtual Reality Peripheral Network) library (101). VRChem (102) is a molecular builder for organic molecules. Nanome (103) and Molecular Rift (104) are VR visualizers targeting drug design applications. A group at the University of California, San Francisco, has developed a trio of applications for biomolecular visualization, collaboration, and education: ChimeraX VR, AltPDB, and Molecular Zoo (105). Narupa (106) is a multiuser visualizer designed for interactive molecular dynamics simulations, and InteraChem (97, 107) is designed for ab initio interactive molecular dynamics.

Many of the VR visualizers share similar features: Small molecules are shown in ball-and-stick representation, while larger biomolecules such as proteins may be rendered with cartoon structures to represent α-helices and β-sheets. The handheld controllers can be used to manipulate the scene, pull on individual molecules or atoms, or interact with various menus or displays. **Figure 4** shows several snapshots from an ab initio interactive molecular dynamics simulation of the ring opening of cyclobutene in the InteraChem visualizer. This simulation was performed in real time by connecting to the same TeraChem Protocol Buffer server that powers the TCC framework described in Section 2.2. Taken together, accessible computation and visualization can enable new kinds of tools for research and education.

4.2. Augmented Reality

AR is an exciting extended reality technology that superimposes virtual objects onto real-world scenes (108, 109). The popularity of AR is due in large part to its interactive nature: Users can directly manipulate their point of view of the virtual 3D object by rotating, translating, and zooming.

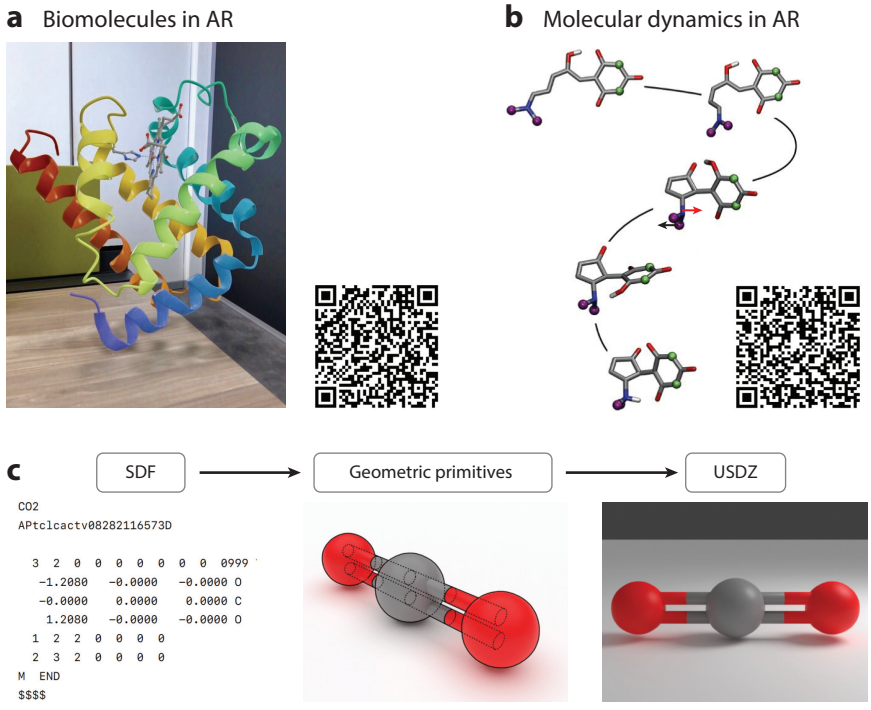

Figure 5

Visualization of (*a*) biomolecule and (*b*) molecular dynamics trajectory in AR. These can be viewed in AR using an iOS mobile device (e.g., iPhone or iPad) by scanning the respective QR codes. (*c*) The workflow to create an AR molecular structure, from SDF file to geometric primitives to USDZ file. Panels *b* and *c* adapted with permission from References 39 and 88, respectively. Abbreviations: AR, augmented reality; QR, quick response; SDF, spatial data file; USDZ, universal scene description ZIP.

Thanks to the ubiquity of smartphones and more powerful mobile chips and sensors, AR has now become much more accessible (110): One can simply point a phone's camera at a real-world scene and see a virtual object appear in the scene. Major mobile operating systems (Android and iOS) now have software development kits (SDKs), allowing developers to easily create AR applications.

In chemistry, AR allows users to visualize molecules in 3D and manipulate them in a natural, intuitive way (111–113). For example, the visualization of large biomolecules in AR is a unique experience where users can see the 3D molecular structure appear in the real world in front of them and interactively inspect its structural features, active site, and solvent-accessible channels (**Figure 5*a***). Visualizing animations in AR is a unique way to follow chemical reactions. Understanding a molecular dynamics simulation from 2D plots is extremely difficult, whereas viewing an animation of the trajectory in AR offers an immediately intuitive experience. Embedding quick response (QR) codes into the figures of scientific publications that link to AR visualizations can therefore enhance accessibility (39) (**Figure 5*b***).

Our workflow to generate a 3D molecule in AR uses ARKit (114), an iOS SDK tool that implements AR (**Figure 5*c***). The starting point is a spatial data file (SDF) containing the XYZ Cartesian coordinates of each atom and the type of bond between pairs of atoms. This file is converted to a list of geometric primitives, where each atom is represented by a sphere (radius

proportional to the van der Waals radius) and each bond is represented by cylinders (the bond order is denoted by the number of cylinders used). The geometric primitives are then used to create a universal scene description ZIP (USDZ) file. USDZ is a file format created by Pixar that can be displayed in AR on iOS devices without the need to install any third-party software (115). To create smoother 3D models, we subdivide the surfaces of spheres and cylinders into small triangles and rectangles and write them as meshes.

Recently, many AR applications for chemistry education have been developed, and a comprehensive overview has been provided (112, 113). Below, we discuss MolAR (88), which combines chemical structure recognition, cloud quantum chemistry, and AR to create a seamless experience for visualizing and interacting with molecular structures.

5. ASSEMBLING THE PARTS: INTEGRATION INTO COMMON PLATFORMS

In this section we consider how the technologies discussed so far can be combined to dramatically increase the accessibility and usability of computational chemistry tools. GPU-accelerated quantum chemistry software makes it possible to carry out real-time computations in seconds, the cloud-computing framework performs these calculations as a versatile service, and NUIs offer user-friendly ways to interact with these technologies on easily accessible platforms (e.g., smartphones and virtual assistants). The primary design focus of the following tools is to provide a natural and interactive human interface for easy access to computational chemistry tools.

5.1. Virtual Assistants

Virtual assistants such as Amazon's Alexa, Apple's Siri, and Google's Assistant simplify the interaction with technology by allowing users to communicate with devices via natural spoken language. They are commonly used in smart homes but also have great potential to facilitate routine operations in the scientific setting. For example, Helix (116) and Vitro (117) are lab assistants that can walk users through predefined experimental procedures with step-by-step instructions.

We recently released ChemVox (66) (**Figure 6**), an Alexa skill to perform real-time quantum chemistry calculations from voice commands (see 118). ChemVox can answer questions such as "What is the dipole moment of caffeine?" and return results in a matter of seconds. This is possible thanks to GPU-accelerated quantum chemistry (discussed in Section 2.1), which carries out requested computations. Since Alexa's response time is restricted to only 5 seconds, the calculations need to be very fast.

Currently, ChemVox can calculate the dipole moment, the excitation energy, and the solvatochromic shift for small- and medium-sized molecules (up to ~100 atoms). The ChemVox backend employs an AWS (Amazon web service) Lambda function that is interfaced with PubChem and TCC. The ChemVox workflow proceeds as follows: (*a*) The user issues a request via an Alexa skill; (*b*) the request is passed to the Lambda function, which processes the information; (*c*) the requested molecular structure is retrieved from the PubChem database; (*d*) the input specifications are sent to TCC, which performs the calculation; and (*e*) the results along with a skeletal image of the molecule (in screen powered devices) are returned via Alexa and the full calculation details are emailed to the user.

To overcome the limits of VUIs discussed in Section 3.2 and improve the recognition of molecule names, we built a custom vocabulary containing ~5,000 unique chemical words, such as dichloromethane and cyclohexane. Furthermore, to improve the ability of ChemVox to recognize user's intentions, we defined hundreds of sample utterances for each intent. This allows the skill to understand the same question with different phrasings.

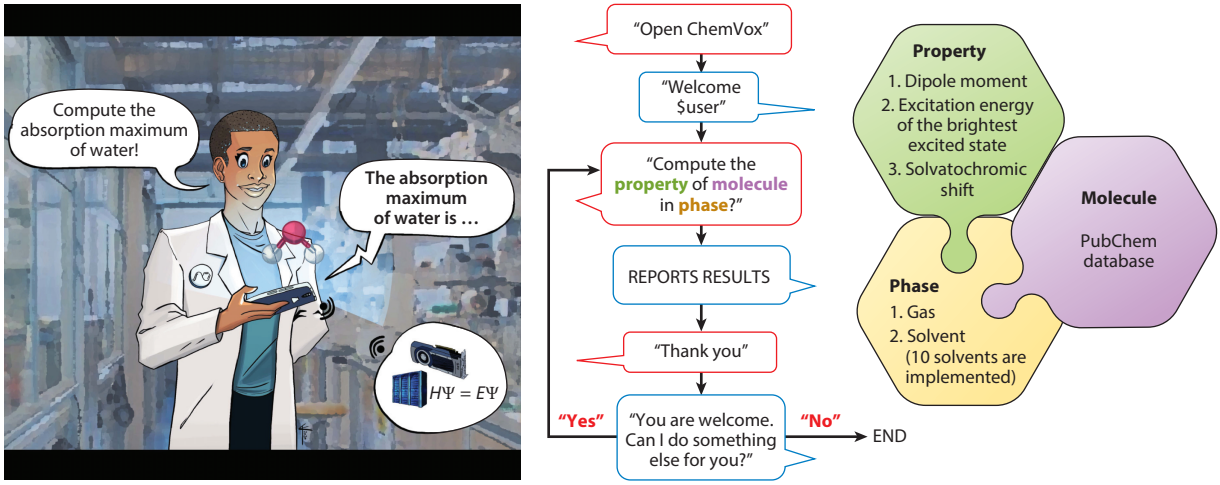

Figure 6

(*Left*) Diagram of user interacting with ChemVox, and (*right*) conversation between ChemVox and a user performing quantum chemistry calculations with vocal commands.

Integrating virtual assistants with mobile applications, such as MolAR (described in the following section), is an exciting future avenue. This would allow molecules to be visualized in AR directly from speech commands.

5.2. Mobile Applications

Considering the universal use of mobile smart devices (e.g., smartphones and tablets) among students and researchers, mobile applications are an ideal platform for delivering computational chemistry tools to the community (119, 120).

Recently, several applications for interactive learning and molecular visualization have been developed (121, 122). For example, MILAGE LEARN+ integrates several teaching strategies, like gaming and autonomous learning, to make learning organic chemistry concepts more interactive and engaging (123). Molecular viewer applications such as Atomdroid (124) or Molecules (125) provide platforms for 3D molecular visualization, enhancing student understanding of molecular structure.

Relatively few applications can compute and visualize quantum mechanical properties. MO-Cubed (126) allows users to build molecular structures and compute their chemical properties via semiempirical calculations, while WebMO (127) offers visualization of molecular orbitals/vibrational modes and can submit jobs to a WebMO server and view the results.

As discussed in Section 4.2, AR is the frontier of molecular visualization on mobile smart devices. The Android and iOS mobile operating systems now support AR natively, leading to several applications for viewing molecules in AR (128–130). MoleculARweb (131) is a web-based application that displays the chemical structure of a built molecule in AR, whereas users of BiochemAR (132) can visualize and manipulate potassium channels.

We recently developed MolAR (88), an iOS open-source application that combines chemical structure recognition, AR, and quantum chemistry cloud computing to transform hand-drawn chemical structures into 3D molecules in AR: The AR structure appears to float above the page as if it has come to life (**Figure 7**) (for a demonstration, see 133). MolAR is connected to TCC, allowing users to compute and visualize chemical properties (e.g., frontier molecular orbitals

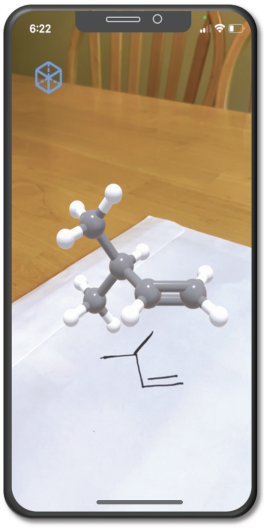

Hand-drawn chemical structure recognition

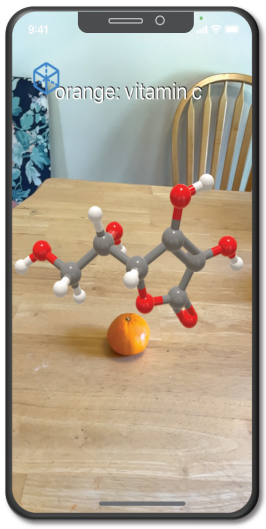

Molecule hunt

Molecule and biomolecule visualization

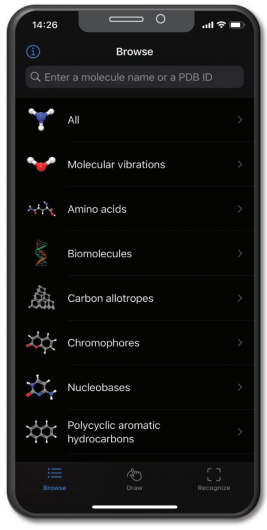

Molecule gallery

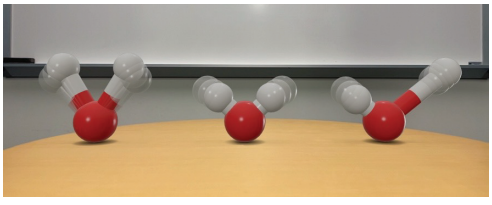

Molecular vibrations

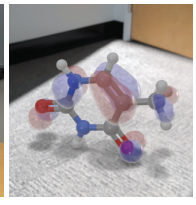

Molecular orbitals

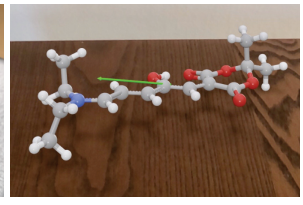

Dipole moment

Figure 7

Features of the MolAR app: hand-drawn chemical structure recognition, recognition of molecules in objects (molecular hunt), visualization of molecules and biomolecules from public databases (i.e., Protein Data Bank), a molecular gallery with selected template molecules for users not familiar with chemistry, visualization of vibrational normal modes, and calculation of electronic properties such as dipole moment and molecular orbitals. Figure adapted with permission from Reference 88; copyright 2022 AIP Publishing.

and dipole moments) in real time. AR animations of vibrational normal modes can also be viewed.

Unlike other mobile applications, MolAR does not require the use of markers, printouts, or preparation of 3D model files prior to use. Indeed, the application supports several molecular input formats: common/IUPAC names, SMILES strings, Protein Data Bank (PDB) IDs (for proteins), or skeletal structures. Furthermore, in the spirit of Pokémon Go, where users collect virtual creatures in real-world scenes, a hunt feature allows one to discover representative molecules inside common household objects such as food and drink; for example, when users take a picture of coffee, the 3D model of caffeine appears above it. Work is in progress to incorporate dynamic motion so that reaction mechanisms can be viewed in AR.

5.3. Web User Interfaces

The internet has become a universal tool for communication and a democratic source of information. The continuous growth of users has been accelerated by web user interfaces (web UIs). A

web UI is a software application based on a client–server architecture; it runs inside a web browser (client) and often communicates with a remote web server (server).

Web UIs offer several desirable features for interfacing to computational chemistry software. They do not need to be installed locally and can run across multiple platforms, allowing users to collaborate using different operating systems (e.g., Linux, macOS, Microsoft Windows). Web UIs have low requirements for the user's computer (just an updated web browser) since the expensive computational operations are performed on a remote server. Furthermore, they are accessible from any device with web browsing capabilities, such as desktop computers, smartphones, and tablets. The classroom is a particularly compelling setting for web UIs displayed on the newest generation of projectors, which can turn any flat surface into an interactive touchscreen (for an example, see 134).

In the last few years, many web UIs have been developed for computational chemistry. Rose and coworkers (135) developed Mol* for 3D molecular visualization and analysis of large biomolecular structures; MoleculARweb (131) and MolAR web UI (136) display chemical structures in AR, whereas ProteinVR (137) and Autodesk Molecular Viewer (138) allow molecular visualization in VR. Other web UIs aid in the construction of input structures; for example, MolView (139) is an intuitive application to build chemical structures, whereas CHARMM-GUI (140) simplifies the construction of complex biological simulation systems (e.g., membrane structures).

Furthermore, some web UIs set up, launch, and analyze calculations directly in a web browser. Durrant and coworkers (141) developed Webina, a JavaScript/WebAssembly library to run docking calculations in a web browser with AutoDock Vina. Polik & Schmidt (142) developed WebMO, a web-based interface for computational chemistry programs, and Entos released Envision (143) to calculate and compare properties of molecules.

As we have seen with ChemVox and MolAR so far, TCC can be easily interfaced to different front-end applications. In the following, we expand the platform choice further and describe how to set up a web UI to build input and perform quantum chemistry calculations in the cloud (144). This web UI is written in browser-supported languages (e.g., JavaScript, HTML, and CSS), and its workflow is as follows: (*a*) The user selects input specifications (e.g., method and basis set) for the quantum mechanical calculations in the browser using a 3D editor [we use J(S)mol web applet; 145] and intuitive menus; (*b*) a request to perform the calculation is sent to TCC from the web browser; (*c*) TCC accepts the request and input data, performs the calculations, and sends back a JSON object with the results; and (*d*) the results are sent back to the client and displayed to the user in the browser. Based on this workflow, we are currently developing TeraChem Web Services (144), a web UI that uses chemical structure recognition as the input method and AR as the visualization tool (for a demonstration, see 146).

5.4. Twitter Bots

Nowadays, Twitter is an important communication medium for scientists to share information (e.g., journal publications, conferences, and job openings) (147). Posting multimedia content (e.g., videos or pictures of chemistry experiments) can reach a diverse audience of nonscientists interested in advancing their scientific curiosity. Interfacing chemistry software with Twitter is a unique way for nonexperts to access the tools, and it can be done by developing Twitter bots, which are automated Twitter accounts controlled through the Twitter API (148).

Recently, Yamamoto and coworkers (149) presented Twitter bots for cheminformatics applications such as retrosynthetic analysis, 2D chemical structure editing, and 3D molecule viewing. These bots reply to a tweet request that includes the SMILES string of a molecule, returning, for example, the retrosynthetic analysis for that compound.

We developed an interactive application to perform quantum chemical calculations on Twitter. By tweeting @TerachemBot compute [property][molecule][functional], the Twitter bot picks up the request and replies to the tweet with the results of the prompted calculation. TerachemBot can perform single-point energy calculations with different density functionals (the basis set is fixed for simplicity to 3-21G) and return the self-consistent energy or dipole moment. Its workflow is simple: (*a*) retrieve mentions via Twitter API, (*b*) parse the input information and launch the calculation on TCC, and (*c*) send back the calculation results via the Twitter API. Currently the bot is offline; however, a demonstration video can be found at Reference 150 and the code is freely available (87).

5.5. Humanoid Robotics

The flexibility of the discussed APIs lends itself to innovative human interactions with cutting-edge technologies, including those described above, and robotic platforms. Robotic platforms allow users to explore complex scientific concepts with support from intelligent systems that are specifically trained to interact with humans. A new generation of robots have been shown to successfully participate in complex social scenarios where they interact closely with multiple individuals for prolonged periods of time. The complexity and variability of these scenarios require the robot to learn advanced social skills autonomously and incrementally. The robots frequently learn how to interact with others in the world in a natural manner directly from their experience as an embodied agent.

Acquiring skills is crucial in human daily and professional life; however, learning new scientific topics, including quantum chemistry, is challenging for many. An expert tutor in the skill is therefore highly valuable for explaining challenging concepts. The physical presence of an embodied robot can pick up on implicit signals via communication channels that can help create an enhanced, personalized learning process. Indeed, the interaction between the robotic expert and the learner could be seen as a continuous flow of physical and emotional signals, which leads the tutor to build a complex understanding of the human partner and act accordingly. Several studies have demonstrated the potential of social robots to positively contribute to users' learning experience for skill acquisition (151, 152). Furthermore, the presence of physical robots may have advantages in sensing and using affective data by inducing higher degrees of emotional expressiveness (153). These results suggest not only that embodied social robots may be a more effective medium for developing intelligent virtual interfaces but also that integrating affect awareness in the tutoring model can lead to important benefits.

Several works have attempted to provide solutions for social robots that naturally interact with users during free interactions. Examples are robots developed to entertain large audiences in museums or shopping malls: A socially intelligent robot has been designed to interact with the general public in open spaces by integrating methods from audiovisual scene processing, social signal processing, and conversational AI (154). These studies proved that even nonexperts can successfully engage in short-term interactions with a robot. The next challenge for the community is to support long-term personalization.

We aim to exploit interactive tasks between the humanoid iCub robot (155, 156) (**Figure 8*a***) and a human partner to effectively tutor quantum chemistry. The humanoid robot has 53° of freedom distributed in a body structure that resembles a human's. The numerous degrees of freedom and the accurate motor control enable whole-body behaviors and actions that are rich and expressive. This facilitates natural and instantaneous communication between the robot and the user. As with many other human–robot interaction platforms, iCub's software architecture allows the robot to perceive the environment and act appropriately. Sensorial perception allows the robot

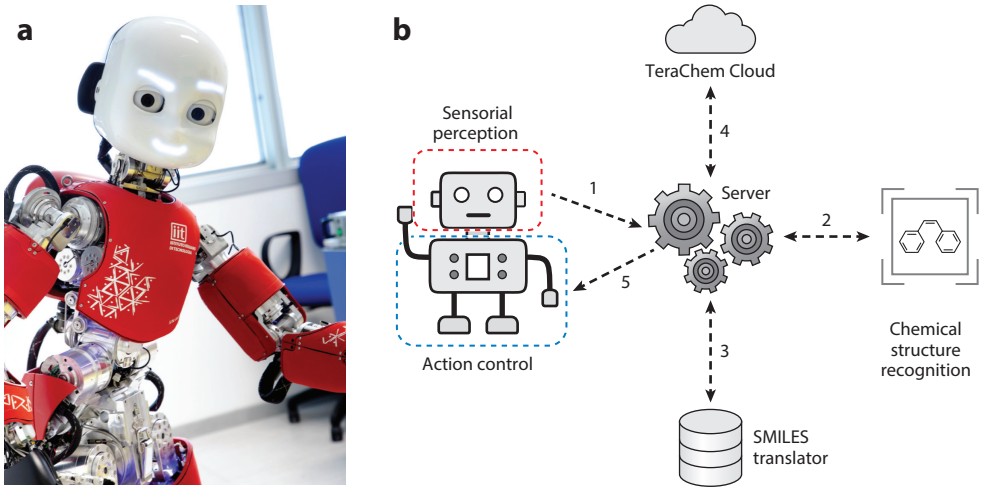

Figure 8

(*a*) Photograph of an iCub humanoid robot. (*b*) Outline of the iCub quantum chemistry workflow: (①) iCub takes a picture of the hand-drawn molecule and sends the image to the server, which then (②) feeds it to Mathpix to predict the SMILES representation. (③) The SMILES is then sent to PubChem to retrieve the XYZ Cartesian coordinates of the molecule, which (④) are passed to the TeraChem Cloud to perform the calculation. (⑤) The results are finally communicated by voice. Panel *b* adapted with permission from Reference 88; copyright 2022 AIP Publishing. Abbreviation: SMILES, simplified molecular input line-entry system.

to understand discussion points and respond with the whole body, head, and voice. Typical body language and communicative movements (e.g., nodding to show understanding and engagement) are replicated in the action control system of the robot, which presents the robot as an expressive learning partner. This enhances the learning experience of the human by leveraging the robot's communicative familiarity.

The workflow of the iCub quantum chemistry tutor is outlined in **Figure 8*b***: (*a*) iCub takes a picture of the hand-drawn molecule and sends the image to the server, which then (*b*) feeds it to Mathpix to predict the SMILES representation. (*c*) The SMILES is then sent to PubChem to retrieve the XYZ Cartesian coordinates of the molecule, which (*d*) are passed to TCC to perform the calculation. (*e*) The results are finally communicated by voice. For a video demonstration, see Reference 157.

6. CONCLUSIONS

The democratization of quantum chemistry among students and nonexpert users has been hindered by barriers to usability and accessibility. Cutting-edge technologies in GPU-accelerated quantum chemistry, cloud computing, and AI are gradually enabling the development of interactive interfaces to remove these barriers and make computational chemistry easily accessible for the wider community. Real-time interactions, intuitive and engaging input/output technologies, and use of widespread platforms are key factors for success.

Certainly, there is room for improvement. Image and speech recognition can be made more robust to allow a wider variety of molecular drawings and accent styles to be recognized. Furthermore, they can be extended to recognize stereocenters, charged or radical molecules, and metal complexes. The computational capabilities of these NUIs can also be extended to more complex problems such as simulation of molecular reactivity. As an example, one could envision an

application that proceeded from pictures of hand-drawn reactants and products to computing the minimum energy pathways between them to providing a visualization of the reaction path and energy profile in AR. A serious challenge for the future is ensuring that the results are reliable for a broad range of problems and molecules. This requires the development of robust methods and workflows for automating computational chemistry tasks and assessing the expected uncertainty in the results.

Another serious challenge is the maintenance of applications such as ChemVox or MolAR. For example, the continual software updates to the Alexa API have made it challenging to ensure that ChemVox is always operating as expected. Indeed, this is a major challenge in the development of any complex software involving many interacting (and changing) APIs. It remains to be seen whether the academic environment can provide the long-term attention needed to keep these complex services running.

Although these technologies are still at an early stage, we strongly believe they can affect the molecular vision of students and academics, serving as powerful and convenient educational and research tools that encourage scientific curiosity and facilitate learning. Quoting John Dewey, "If we teach today as we taught yesterday, we rob our children of tomorrow." Integrating these interactive technologies in the classroom could deliver a unique learning experience that forms an enriched foundation for students' future scientific endeavors.

NOTE ADDED IN PROOF

MolAR is now available for both iPhone (Apple App Store) and Android (Google Play Store).

DISCLOSURE STATEMENT

T.J.M. is a co-founder of PetaChem, LLC. The other authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

LITERATURE CITED

1. Owens JD, Houston M, Luebke D, Green S, Stone JE, Phillips JC. 2008. GPU computing. *Proc. IEEE* 96:879–99
2. Ufimtsev IS, Martínez TJ. 2008. Graphical processing units for quantum chemistry. *Comput. Sci. Eng.* 10:26–34
3. Salomon-Ferrer R, Case DA, Walker RC. 2013. An overview of the Amber biomolecular simulation package. *WIREs Comput. Mol. Sci.* 3:198–210
4. Arthur EJ, Brooks CL 3rd. 2016. Efficient implementation of constant pH molecular dynamics on modern graphics processors. *J. Comput. Chem.* 37:2171–80
5. Kutzner C, Páll S, Fechner M, Esztermann A, de Groot BL, Grubmüller H. 2015. Best bang for your buck: GPU nodes for GROMACS biomolecular simulations. *J. Comput. Chem.* 36:1990–2008
6. Eastman P, Swails J, Chodera JD, McGibbon RT, Zhao Y, et al. 2017. OpenMM 7: rapid development of high performance algorithms for molecular dynamics. *PLOS Comput. Biol.* 13:e1005659
7. Seritan S, Bannwarth C, Fales BS, Hohenstein EG, Isborn CM, et al. 2021. TeraChem: a graphical processing unit-accelerated electronic structure package for large-scale ab initio molecular dynamics. *WIREs Comput. Mol. Sci.* 11:e1494
8. Seritan S, Bannwarth C, Fales BS, Hohenstein EG, Kokkila-Schumacher SIL, et al. 2020. TeraChem: accelerating electronic structure and ab initio molecular dynamics with graphical processing units. *J. Chem. Phys.* 152:224110
9. Barca GMJ, Alkan M, Galvez-Vallejo JL, Poole DL, Rendell AP, Gordon MS. 2021. Faster self-consistent field (SCF) calculations on GPU clusters. *J. Chem. Theory Comput.* 17:7486–503

10. Kussmann J, Laqua H, Ochsenfeld C. 2021. Highly efficient resolution-of-the-identity density functional theory calculations on central and graphics processing units. *J. Chem. Theory Comput.* 17:1512–21
11. Casalino L, Dommer AC, Gaieb Z, Barros EP, Sztain T, et al. 2021. AI-driven multiscale simulations illuminate mechanisms of SARS-CoV-2 spike dynamics. *Int. J. High Perform. Comput. Appl.* 35:432–51
12. Sisto A, Stross C, van der Kamp MW, O'Connor M, McIntosh-Smith S, et al. 2017. Atomistic non-adiabatic dynamics of the LH2 complex with a GPU-accelerated ab initio exciton model. *Phys. Chem. Chem. Phys.* 19:14924–36
13. Kulik HJ, Zhang J, Klinman JP, Martínez TJ. 2016. How large should the QM region be in QM/MM calculations? The case of catechol O-methyltransferase. *J. Phys. Chem. B* 120:11381–94
14. Liu F, Luehr N, Kulik HJ, Martínez TJ. 2015. Quantum chemistry for solvated molecules on graphical processing units using polarizable continuum models. *J. Chem. Theory Comput.* 11:3131–44
15. Liu F, Sanchez DM, Kulik HJ, Martínez TJ. 2019. Exploiting graphical processing units to enable quantum chemistry calculation of large solvated molecules with conductor-like polarizable continuum models. *Int. J. Quantum Chem.* 119:e25760
16. Isborn CM, Luehr N, Ufimtsev IS, Martínez TJ. 2011. Excited-state electronic structure with configuration interaction singles and Tamm–Dancoff time-dependent density functional theory on graphical processing units. *J. Chem. Theory Comput.* 7:1814–23
17. Jones CM, List NH, Martínez TJ. 2021. Resolving the ultrafast dynamics of the anionic green fluorescent protein chromophore in water. *Chem. Sci.* 12:11347–63
18. Jones CM, List NH, Martínez TJ. 2022. Steric and electronic origins of fluorescence in GFP and GFP-like proteins. *J. Am. Chem. Soc.* 144:12732–46
19. Yu JK, Liang R, Liu F, Martínez TJ. 2019. First-principles characterization of the elusive I fluorescent state and the structural evolution of retinal protonated Schiff base in bacteriorhodopsin. *J. Am. Chem. Soc.* 141:18193–203
20. Liang R, Liu F, Martínez TJ. 2019. Nonadiabatic photodynamics of retinal protonated Schiff base in channelrhodopsin 2. *J. Phys. Chem. Lett.* 10:2862–68
21. Liang R, Yu JK, Meisner J, Liu F, Martinez TJ. 2021. Electrostatic control of photoisomerization in channelrhodopsin 2. *J. Am. Chem. Soc.* 143:5425–37
22. Wang L-P, Titov A, McGibbon R, Liu F, Pande VS, Martínez TJ. 2014. Discovering chemistry with an ab initio nanoreactor. *Nat. Chem.* 6:1044–48
23. Pieri E, Lahana D, Chang AM, Aldaz CR, Thompson KC, Martínez TJ. 2021. The non-adiabatic nanoreactor: towards the automated discovery of photochemistry. *Chem. Sci.* 12:7294–307
24. Armbrust M, Fox A, Griffith R, Joseph AD, Katz R, et al. 2010. A view of cloud computing. *Commun. ACM* 53:50–58
25. Strieth-Kalthoff F, Sandfort F, Segler MHS, Glorius F. 2020. Machine learning the ropes: principles, applications and directions in synthetic chemistry. *Chem. Soc. Rev.* 49:6154–68
26. Coley CW, Green WH, Jensen KF. 2018. Machine learning in computer-aided synthesis planning. *Acc. Chem. Res.* 51:1281–89
27. Vamathevan J, Clark D, Czodrowski P, Dunham I, Ferran E, et al. 2019. Applications of machine learning in drug discovery and development. *Nat. Rev. Drug Discov.* 18:463–77
28. Ekins S, Puhl AC, Zorn KM, Lane TR, Russo DP, et al. 2019. Exploiting machine learning for end-to-end drug discovery and development. *Nat. Mater.* 18:435–41
29. Butler KT, Davies DW, Cartwright H, Isayev O, Walsh A. 2018. Machine learning for molecular and materials science. *Nature* 559:547–55
30. Behler J, Parrinello M. 2007. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.* 98:146401
31. Noé F, Tkatchenko A, Müller K-R, Clementi C. 2020. Machine learning for molecular simulation. *Annu. Rev. Phys. Chem.* 71:361–90
32. Dick S, Fernandez-Serra M. 2020. Machine learning accurate exchange and correlation functionals of the electronic density. *Nat. Commun.* 11:3509
33. Kirkpatrick J, McMorrow B, Turban DHP, Gaunt AL, Spencer JS, et al. 2020. Pushing the frontiers of density functionals by solving the fractional electron problem. *Science* 374:1385–89

34. Snyder JC, Rupp M, Hansen K, Müller K-R, Burke K. 2012. Finding density functionals with machine learning. *Phys. Rev. Lett.* 108:253002
35. Lehtola S, Kartunnen AJ. 2022. Free and open software for computational chemistry education. *WIREs Comput. Mol. Sci.* 12:e1610
36. Kim S, Chen J, Cheng T, Gindulyte A, He J, et al. 2019. PubChem 2019 update: improved access to chemical data. *Nucl. Acids Res.* 47:D1102–9
37. Raucci U. 2022. *Cutting-edge technologies in computational chemistry*. Video. Mar. 17. **https://www.youtube.com/watch?v=66nzmh3o8rc**
38. Wolf TJA, Sanchez DM, Yang J, Parrish RM, Nunes JPF, et al. 2019. The photochemical ring-opening of 1,3-cyclohexadiene imaged by ultrafast electron diffraction. *Nat. Chem.* 11:504–9
39. Sanchez DM, Raucci U, Martínez TJ. 2021. In silico discovery of multistep chemistry initiated by a conical intersection: the challenging case of donor–acceptor Stenhouse adducts. *J. Am. Chem. Soc.* 143:20015–21
40. Qiu Y, Smith DGA, Stern CD, Feng M, Jang H, Wang L-P. 2020. Driving torsion scans with wavefront propagation. *J. Chem. Phys.* 152:244116
41. Whitten JL. 1973. Coulombic potential energy integrals and approximations. *J. Chem. Phys.* 58:4496
42. Almlof J, Faegri K, Korsell K. 1982. Principles for a direct SCF approach to LCAO-MO ab-initio calculations. *J. Comput. Chem.* 3:385–99
43. Maurer SA, Lambrecht DS, Flaig D, Ochsenfeld C. 2012. Distance-dependent Schwartz-based integral estimates for two-electron integrals: reliable tightness versus rigorous upper bounds. *J. Chem. Phys.* 136:144107
44. Ufimtsev IS, Martínez TJ. 2008. Quantum chemistry on graphical processing units. 1. Strategies for two-electron integral evaluation. *J. Chem. Theory Comput.* 4:222–31
45. Ufimtsev IS, Martinez TJ. 2009. Quantum chemistry on graphical processing units. 2. Direct self-consistent-field implementation. *J. Chem. Theory Comput.* 5:1004–15
46. Ufimtsev IS, Martinez TJ. 2009. Quantum chemistry on graphical processing units. 3. Analytical energy gradients, geometry optimization, and first principles molecular dynamics. *J. Chem. Theory Comput.* 5:2619–28
47. Titov AV, Ufimtsev IS, Luehr N, Martinez TJ. 2013. Generating efficient quantum chemistry codes for novel architectures. *J. Chem. Theory Comput.* 9:213–21
48. Gordon MS, Barca G, Leang SS, Poole D, Rendell AP, et al. 2020. Novel computer architectures and quantum chemistry. *J. Phys. Chem. A* 124:4557–82
49. Götz AW, Wölfle T, Walker RC. 2010. Quantum chemistry on graphics processing units. *Annu. Rep. Comput. Chem.* 6:21–35
50. Barca GMJ, Bertoni C, Carrington L, Datta D, De Silva N, et al. 2020. Recent developments in the General Atomic and Molecular Electronic Structure System. *J. Chem. Phys.* 152:154102
51. Levine BG, Martinez TJ. 2003. Hijacking the PlayStation2 for computational chemistry. *Abstr. Pap. Am. Chem. Soc.* 226:U426
52. Stone JE, Hardy DJ, Ufimtsev IS, Schulten K. 2010. GPU-accelerated molecular modeling coming of age. *J. Mol. Graph. Model.* 29:116–25
53. Wilhite DL, Euwema RN. 1974. Charge-conserving integral approximations for ab initio quantum chemistry. *J. Chem. Phys.* 61:375
54. Yasuda K. 2008. Two-electron integral evaluation on the graphics processor unit. *J. Comput. Chem.* 29:334–42
55. Luehr N, Ufimtsev IS, Martínez TJ. 2011. Dynamic precision for electron repulsion integral evaluation on graphical processing units (GPUs). *J. Chem. Theory Comput.* 7:949–54
56. Vysotskiy VP, Cederbaum LS. 2011. Accurate quantum chemistry in single precision arithmetic: correlation energy. *J. Chem. Theory Comput.* 7:320–26
57. Pokhilko P, Epifanovsky E, Krylov AI. 2018. Double precision is not needed for many-body calculations: emergent conventional wisdom. *J. Chem. Theory Comput.* 14:4088–96
58. Docker Inc. 2022. Home page. *Docker*. **https://www.docker.com**
59. Burns B, Grant B, Oppenheimer D, Brewer E, Wilkes J. 2016. Borg, Omega, and Kubernetes: lessons learned from three container-management systems over a decade. *Queue* 14:70–93

60. Thackston R, Fortenberry RC. 2015. The performance of low-cost commercial cloud computing as an alternative in computational chemistry. *J. Comput. Chem.* 36:926–33
61. Fortenberry RC, Thackston R. 2015. Optimal cloud use of quartic force fields: the first purely commercial cloud computing based study for rovibrational analysis of $SiCH^-$. *Int. J. Quantum Chem.* 115:1650–57
62. Moghadam BT, Alvarsson J, Holm M, Eklund M, Carlsson L, Spjuth O. 2015. Scaling predictive modeling in drug development with cloud computing. *J. Chem. Inf. Model.* 55:19–25
63. Seritan S, Thompson K, Martínez TJ. 2020. TeraChem cloud: a high-performance computing service for scalable distributed GPU-accelerated electronic structure calculations. *J. Chem. Inf. Model.* 60:2126–37
64. Open Eye Scientific. 2022. Orion. *OpenEye, Cadence Molecular Sciences*. **https://www.eyesopen.com/orion**
65. Smith DGA, Altarawy D, Burns LA, Welborn M, Naden LN, et al. 2021. The MolSSI QCArchive project: an open-source platform to compute, organize, and share quantum chemistry data. *WIREs Comput. Mol. Sci.* 11:e1491
66. Raucci U, Valentini A, Pieri E, Weir H, Seritan S, Martínez TJ. 2021. Voice-controlled quantum chemistry. *Nat. Comput. Sci.* 1:42–45
67. Contreras ML, Allendes C, Alvarez LT, Rozas R. 1990. Computational perception and recognition of digitized molecular structures. *J. Chem. Inf. Comput. Sci.* 30:302–7
68. McDaniel JR, Balmuth JR. 1992. Kekule: OCR-optical chemical (structure) recognition. *J. Chem. Inf. Comput. Sci.* 32:373–78
69. Rajan K, Brinkhaus HO, Zielesny A, Steinbeck C. 2020. A review of optical chemical structure recognition tools. *J. Cheminf.* 12:60
70. Beard EJ, Cole JM. 2020. ChemSchematicResolver: a toolkit to decode 2D chemical diagrams with labels and R-groups into annotated chemical named entities. *J. Chem. Inf. Model.* 60:2059–72
71. Filippov IV, Nicklaus MC. 2009. Optical structure recognition software to recover chemical information: OSRA, an open source solution. *J. Chem. Inf. Model.* 49:740–43
72. Gkoutos GV, Rzepa H, Clark RM, Adjei O, Johal H. 2003. Chemical machine vision: automated extraction of chemical metadata from raster images. *J. Chem. Inf. Comput. Sci.* 43:1342–55
73. Ibison P, Jacquot M, Kam F, Neville AG, Simpson RW, et al. 1993. Chemical literature data extraction: the CLiDE Project. *J. Chem. Inf. Comput. Sci.* 33:338–44
74. Ouyang TY. 2007. *Recognition of hand drawn chemical diagrams*. MS Thesis, MIT, Cambridge, MA
75. Park J, Rosania GR, Shedden KA, Nguyen M, Lyu N, Saitou K. 2009. Automated extraction of chemical structure information from digital raster images. *Chem. Cent. J.* 3:4
76. Valko AT, Johnson AP. 2009. CLiDE Pro: the latest generation of CLiDE, a tool for optical chemical structure recognition. *J. Chem. Inf. Model.* 49:780–87
77. Oldenhof M, Arany A, Moreau Y, Simm J. 2020. ChemGrapher: optical graph recognition of chemical compounds by deep learning. *J. Chem. Inf. Model.* 60:4506–17
78. Staker J, Marshall K, Abel R, McQuaw CM. 2019. Molecular structure extraction from documents using deep learning. *J. Chem. Inf. Model.* 59:1017–29
79. Rajan K, Zielesny A, Steinbeck C. 2020. DECIMER: towards deep learning for chemical image recognition. *J. Cheminf.* 12:65
80. Pic2mol. 2022. The missing image-to-structure feature for ChemDraw. *Pic2mol*. **https://pic2mol.com/**
81. Clevert D-A, Le T, Winter R, Montanari F. 2021. Img2Mol – accurate SMILES recognition from molecular graphical depictions. *Chem. Sci.* 12:14174–81
82. Weir H, Thompson K, Woodward A, Choi B, Braun A, Martínez TJ. 2021. ChemPix: automated recognition of hand-drawn hydrocarbon structures using deep learning. *Chem. Sci.* 12:10622–33
83. Weininger D. 1988. SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *J. Chem. Inf. Comput. Sci.* 28:31–36
84. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, et al. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems 30 (NIPS 2017)*, pp. 5998–6008. San Diego, CA: NeurIPS
85. OpenBabel. 2022. Open Babel: the open source chemistry toolbox. *OpenBabel*. **https://openbabel.org/wiki/Main_Page**

86. Mathpix Inc. 2022. Convert printed and handwritten chemical diagrams to SMILES. *Mathpix Inc*. **https://mathpix.com/blog/handwritten-chem-diagrams**
87. Martinez Group. 2022. Interactive QC repository. *Martinez Group*. **https://github.com/mtzgroup/ARPC_InteractiveQC**
88. Sakshuwong S, Weir H, Raucci U, Martínez TJ. 2022. Bringing chemical structures to life with augmented reality, machine learning, and quantum chemistry. *J. Chem. Phys.* 156:204801
89. Söldner R, Rheinländer S, Meyer T, Olszowy M, Austerjost J. 2022. Human–device interaction in the life science laboratory. *Adv. Biochem. Eng. Biotechnol.* 182:83–113
90. Rakotomalala F, Randriatsarafara HN, Hajalalaina AR, Ravonimanantsoa NMV. 2021. Voice user interface: literature review, challenges and future directions. *Syst. Theory Control Comput. J.* 1:65–89
91. Hirschberg J, Manning CD. 2015. Advances in natural language processing. *Science* 349:261–66
92. Nadkarni PM, Ohno-Machado L, Chapman WW. 2011. Natural language processing: an introduction. *J. Am. Med. Inform. Assoc.* 18:544–51
93. Hocky GM, White AD. 2022. Natural language processing models that automate programming will transform chemistry research and teaching. *Digit. Discov.* 1:79–83
94. White AD, Hocky GM, Gandhi HA, Ansari M, Cox S, et al. 2022. Do large language models know chemistry? *ChemRxiv*. **https://doi.org/10.26434/chemrxiv-2022-3md3n**
95. Chen M, Tworek J, Jun H, Yuan Q, Pinto HPdO, et al. 2021. Evaluating large language models trained on code. arXiv:2107.03374 [cs.LG]
96. Benzeghiba M, De Mori R, Deroo O, Dupont S, Erbes T, et al. 2007. Automatic speech recognition and speech variability: a review. *Speech Commun*. 49:763–86
97. Seritan S, Wang Y, Ford JE, Valentini A, Gold T, Martínez TJ. 2021. InteraChem: virtual reality visualizer for reactive interactive molecular dynamics. *J. Chem. Educ.* 98:3486–92
98. Google Inc. 2022. GoogleCloud, cloud speech-to-text. *Google Inc*. **https://cloud.google.com/speech-to-text**
99. Cruz-Neira C, Sandin DJ, DeFanti TA, Kenyon RV, Hart JC. 1992. The CAVE: audio visual experience automatic virtual environment. *Commun. ACM* 35:64
100. Humphrey W, Dalke A, Schulten K. 1996. VMD: visual molecular dynamics. *J. Mol. Graph.* 14:33–38
101. VRPN. 2022. Virtual reality peripheral network – official github repository. *VRPN*. **https://github.com/vrpn/vrpn**
102. Pietikäinen O, Hämäläinen P, Lehtinen J, Karttunen AJ. 2021. VRChem: a virtual reality molecular builder. *Appl. Sci.* 11:10767
103. Nanome Inc. 2022. Home page. *Nanome*. **https://nanome.ai**
104. Norrby M, Grebner C, Eriksson J, Boström J. 2015. Molecular Rift: virtual reality for drug designers. *J. Chem. Inf. Model.* 55:2475–84
105. Goddard TD, Brilliant AA, Skillman TL, Vergenz S, Tyrwhitt-Drake J, et al. 2018. Molecular visualization on the Holodeck. *J. Mol. Biol.* 430:3982–96
106. O'Connor MB, Bennie SJ, Deeks HM, Jamieson-Binnie A, Jones AJ, et al. 2019. Interactive molecular dynamics in virtual reality from quantum chemistry to drug binding: an open-source multi-person framework. *J. Chem. Phys.* 150:220901
107. Wang Y, Seritan S, Lahana D, Ford JE, Valentini A, et al. 2022. InteraChem: exploring excited states in virtual reality with ab initio interactive molecular dynamics. *J. Chem. Theory Comput.* 18:3308–17
108. Azuma RT. 1997. A survey of augmented reality. *Presence Virtual Augment. Real.* 6:355–85
109. Billinghurst M, Clark A, Lee G. 2015. A survey of augmented reality. *Found. Trends Hum.-Comput. Interact.* 8:73–272
110. Rodríguez FC, Dal Peraro M, Abriata LA. 2021. Democratizing interactive, immersive experiences for science education with WebXR. *Nat. Comput. Sci.* 1:631–32
111. Singhal S, Bagga S, Goyal P, Saxena V. 2012. Augmented chemistry: Interactive education system. *Int. J. Comput. Appl.* 49:1–5
112. Jiménez ZA. 2019. Teaching and learning chemistry via augmented and immersive virtual reality. *ACS Symp. Ser.* 1318:31–52
113. Fombona-Pascual A, Fombona J, Vicente R. 2022. Augmented reality, a review of a way to represent and manipulate 3D chemical structures. *J. Chem. Inf. Model.* 62:1863–72

114. Apple Computer. 2022. ARKit. *Apple*. **https://developer.apple.com/documentation/arkit**
115. Pixar Animation Studios. 2021. Introduction to USD. *Pixar Animation Studios*. **https://graphics.pixar.com/usd/docs/Introduction-to-USD.html**
116. HelixAIInc. 2022. HelixAI – voice powered digital laboratory assistants for scientific laboratories. *HelixAIInc*. **https://www.askhelix.io/**
117. Cambre J, Liu Y, Taylor RE, Kulkarni C. 2019. Vitro: designing a voice assistant for the scientific lab workplace. In *Proceedings of the 2019 Designing Interactive Systems Conference (DIS 19)*, pp. 1531–42. New York: ACM
118. Raucci U. 2021. *ChemVox: voice-controlled quantum chemistry*. Video. Jan. 14. **https://youtu.be/Mxyw9381K-g**
119. Ranga JS. 2018. Multipurpose use of Explain Everything iPad app for teaching chemistry courses. *J. Chem. Educ.* 95:895–98
120. Williams AJ, Pence HE. 2011. Smart phones, a powerful tool in the chemistry classroom. *J. Chem. Educ.* 88:683–86
121. Libman D, Huang L. 2013. Chemistry on the go: review of chemistry apps on smartphones. *J. Chem. Educ.* 90:320–25
122. Williams AJ, Ekins S, Clark AM, Jack JJ, Apodaca RL. 2011. Mobile apps for chemistry in the world of drug discovery. *Drug Discov. Today* 16:928–39
123. Fonseca CSC, Zacarias M, Figueiredo M. 2021. MILAGE LEARN+: a mobile learning app to aid the students in the study of organic chemistry. *J. Chem. Educ.* 98:1017–23
124. Feldt J, Mata RA, Dieterich JM. 2012. Atomdroid: a computational chemistry tool for mobile platforms. *J. Chem. Inf. Model.* 52:1072–78
125. Sunset Lake Software. 2022. Molecules. *Sunset Lake Software*. **http://www.sunsetlakesoftware.com/molecules/index.html**
126. Google Play. 2022. Mobile Molecular Modeling -Mo3. *Google LLC*. **https://play.google.com/store/apps/details?id=club.amase.mocubed**
127. WebMO LLC. 2022. WebMO app features. *WebMO LLC*. **https://www.webmo.net/features/app.html**
128. Aw JK, Boellaard KC, Tan TK, Yap J, Loh YP, et al. 2020. Interacting with three-dimensional molecular structures using an augmented reality mobile app. *J. Chem. Educ.* 97:3877–81
129. Yang S, Mei B, Yue X. 2018. Mobile augmented reality assisted chemical education: insights from elements 4D. *J. Chem. Educ.* 95:1060–62
130. Jones OAH, Spichkova M, Spencer MJS. 2018. Chirality-2: development of a multilevel mobile gaming app to support the teaching of introductory undergraduate-level organic chemistry. *J. Chem. Educ.* 95:1216–20
131. Rodríguez FC, Frattini G, Krapp LF, Martinez-Hung H, Moreno DM, et al. 2021. MoleculARweb: a web site for chemistry and structural biology education through interactive augmented reality out of the box in commodity devices. *J. Chem. Educ.* 98:2243–55
132. Sung R-J, Wilson AT, Lo SM, Crowl LM, Nardi J, et al. 2020. BiochemAR: an augmented reality educational tool for teaching macromolecular structure and function. *J. Chem. Educ.* 97:147–53
133. Sakshuwong S. 2021. *MolAR: bringing chemical structures to life with augmented reality and machine learning*. Video. Aug. 18. **https://youtu.be/bLqkzz1vZL4**
134. Raucci U. 2021. *Touchscreen projector*. Video. Jan. 4. **https://youtu.be/YvrVDrN2bIg**
135. Sehnal D, Bittrich S, Deshpande M, Svobodová R, Berka K, et al. 2021. Mol* Viewer: modern web app for 3D visualization and analysis of large biomolecular structures. *Nucl. Acids Res.* 49:W431–37
136. Sakshuwong S, Raucci U. 2022. MolAR. *Martinez Group*. **https://mtzgroup.github.io/molar/**
137. Cassidy KC, Šefčík J, Raghav Y, Chang A, Durrant JD. 2020. ProteinVR: web-based molecular visualization in virtual reality. *PLOS Comput. Biol.* 16:e1007747
138. Balo AR, Wang M, Ernst OP. 2017. Accessible virtual reality of biomolecular structural models using the Autodesk Molecule Viewer. *Nat. Methods* 14:1122–23
139. MolView. 2015. Home page. **https://molview.org**

140. Lee J, Patel DS, Ståhle J, Park S-J, Kern NR, et al. 2019. CHARMM-GUI membrane builder for complex biological membrane simulations with glycolipids and lipoglycans. *J. Chem. Theory Comput.* 15:775–86
141. Kochnev Y, Hellemann E, Cassidy KC, Durrant JD. 2020. Webina: an open-source library and web app that runs AutoDock Vina entirely in the web browser. *Bioinformatics* 36:4513–15
142. Polik WF, Schmidt JR. 2022. WebMO: web-based computational chemistry calculations in education and research. *WIREs Comput. Mol. Sci.* 12:e1554
143. Entos Inc. 2022. Envision: interactive chemistry. *Entos Inc*. **https://www.entos.ai/envision**
144. TeraChem Web Services. 2021. Quantum chemistry in one click. *Martinez Group*. **https://tws.dev.mtzlab.com/**
145. Jmol. 2022. Jmol: an open-source Java viewer for chemical structures in 3D. *Jmol*. **http://www.jmol.org/**
146. Raucci U. 2022. *TeraChem web services*. Video. Aug. 17. **https://youtu.be/DxMPZTvh5Gw**
147. Vincent-Ruz P. 2022. The secret silos of #ChemTwitter. *C&EN*. **https://cen.acs.org/sections/the-secret-silos-of-chemtwitter.html#intro**
148. Twitter Inc. 2022. Twitter API. *Twitter Inc*. **https://developer.twitter.com/en/docs/twitter-api**
149. Yoshikawa N, Kubo R, Yamamoto KZ. 2021. Twitter integration of chemistry software tools. *J. Cheminform.* 13:46
150. Raucci U. 2022. *Quantum chemistry calculation on Twitter!!* Video. Jan. 16. **https://youtu.be/yrpVwqvSk-w**
151. Fasola J, Matarić MJ. 2010. *Robot motivator: increasing user enjoyment and performance on a physical/cognitive task*. Paper presented at the 2010 IEEE 9th International Conference on Development and Learning, Ann Arbor, MI, Aug. 18–21
152. Vasalya A, Ganesh G, Kheddar A. 2018. More than just co-workers: presence of humanoid robot co-worker influences human performance. *PLOS ONE* 13:e0206698
153. Spaulding S, Gordon G, Breazeal C. 2016. Affect-aware student models for robot tutors. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems (AAMAS 2016)*, ed. J Thangarajah, K Tuyls, C Jonker, S Marsella, pp. 864–72. New York: ACM
154. Foster ME, Craenen B, Deshmukh A, Lemon O, Bastianelli E, et al. 2019. MuMMER: socially intelligent human-robot interaction in public spaces. arXiv:1909.06749 [cs.RO]
155. Metta G, Sandini G, Vernon D, Natale L, Nori F. 2008. *The iCub humanoid robot: an open platform for research in embodied cognition*. Paper presented at the 8th Workshop on Performance Metrics for Intelligent Systems, Gaithersburg, MA, Aug. 19–21
156. Metta G, Natale L, Nori F, Sandini G, Vernon D, et al. 2010. The iCub humanoid robot: an open-systems platform for research in cognitive development. *Neural Net.* 23:1125–34
157. Raucci U. 2022. *iCub and quantum chemistry*. Video. Mar. 16. **https://youtu.be/0ETxoz2PDeM**