

ChemVox: Voice-Controlled Quantum Chemistry

Umberto Raucci,^{1,2} Alessio Valentini,^{1,2} Elisa Pieri^{1,2}, Hayley Weir^{1,2}, Stefan Seritan^{1,2} and
Todd J. Martínez^{1,2,*}

¹Department of Chemistry and The PULSE Institute, Stanford University, Stanford, CA 94305

²SLAC National Accelerator Laboratory, 2575 Sand Hill Road, Menlo Park, CA 94025

Abstract:

Over the last decade, artificial intelligence has been propelled forward by advances in machine learning algorithms and computational hardware, opening up myriad new avenues for scientific research. Nevertheless, virtual assistants and voice control have yet to be widely utilized in the natural sciences. Here, we present ChemVox, an interactive Amazon Alexa skill that uses speech recognition to perform quantum chemistry calculations. This new application interfaces Alexa with cloud computing and returns the results through a capable device. ChemVox paves the way to making computational chemistry routinely accessible to the wider community.

Artificial Intelligence (AI) is so successfully embedded in our daily activities that it is almost easy to forget its presence. Virtual assistants such as Apple's Siri, Microsoft's Cortana, Amazon's Alexa and Google Assistant[1] make the interaction with cutting-edge technology intuitive and accessible for everyone, concealing complex computational workflows and replacing cumbersome input formats with voice communication. Alexa, the AI that powers Amazon's smart assistants, was sold in more than 200 million Amazon devices in 5 years,[2] a testament to the public's desire for technology that is easy to interact with.

In the last decade, the natural sciences have begun to exploit artificial intelligence to tackle specific problems. Among many examples, we point out that AI has been employed for classifying 3D particle showers[3] and nanoscale images from electron microscopes,[4] predicting the three-dimensional structure of proteins from DNA sequences,[5] weather forecasting,[6] assessing climate change,[7] detecting patterns in astronomical data,[8] and predicting the outcome of organic chemical reactions.[9] Despite this, the virtual assistant territory is, to our knowledge, largely unexplored by the natural science research community. A notable exception is Helix, a lab assistant that can retrieve chemical reactivity information and run stoichiometric calculations.[10]

The field of computational and theoretical chemistry has flourished with the emergence of new algorithms and more powerful hardware, allowing machine learning,[11] cloud computing and GPU-accelerated electronic structure codes[12,13] to produce almost real-time quantum calculations results. This remarkable speed allows for the pairing of such methods with intelligent systems to design an interactive and accessible framework.

Despite this, quantum chemistry software has a high barrier to entry, with the need for coding expertise, understanding complicated input file structures and access to high-performance computing clusters. Therefore, a low barrier to entry alternative that does not require such specialized knowledge or tools would allow the wider chemistry community, in both research and education, to access quantum chemical properties. This would make theoretical chemistry routine in labs and classrooms to support and facilitate chemical understanding, exploration and discovery.

In this article, we present ChemVox, a voice-controlled quantum chemistry application. ChemVox is a free, ready to use Alexa skill that interfaces Amazon Web Services (AWS) with the PubChem[14] chemical database and the TeraChem Cloud[12] framework for cloud based

quantum chemistry. ChemVox, which is usable on all Alexa capable devices, can answer questions related to the electronic properties of small and medium (up to ≈ 100 atoms) sized molecules in a matter of seconds, offering universal easy access to quantum chemistry simulations.

The ChemVox skill workflow is outlined in Figure 1. The back-end employs AWS Lambda, an AWS serverless computing service that allows functions to be constructed and run in a container on a multi-tenant computer cluster. The ChemVox user interaction is managed through the core Lambda event-based function with voice control. When ChemVox is prompted for a calculation on a molecule, the requested molecular geometry is retrieved as an xyz cartesian coordinates file from the PubChem database, along with the molecular charge. The calculation input file for TeraChem Cloud is assembled by ChemVox using a default set of parameters, as specified in the Supporting Information (SI). The input file and geometry file are then dispatched to TeraChem Cloud, which runs the calculation requested by the user. The results of the calculation are returned via Alexa along with a skeletal image of the molecule (for display capable devices), and an e-mail is sent to the user (if the required permissions have been granted) containing all the details of the calculation. The user can also ask clarifying questions regarding the theoretical parameters used in the calculation and the capabilities of the ChemVox skill.

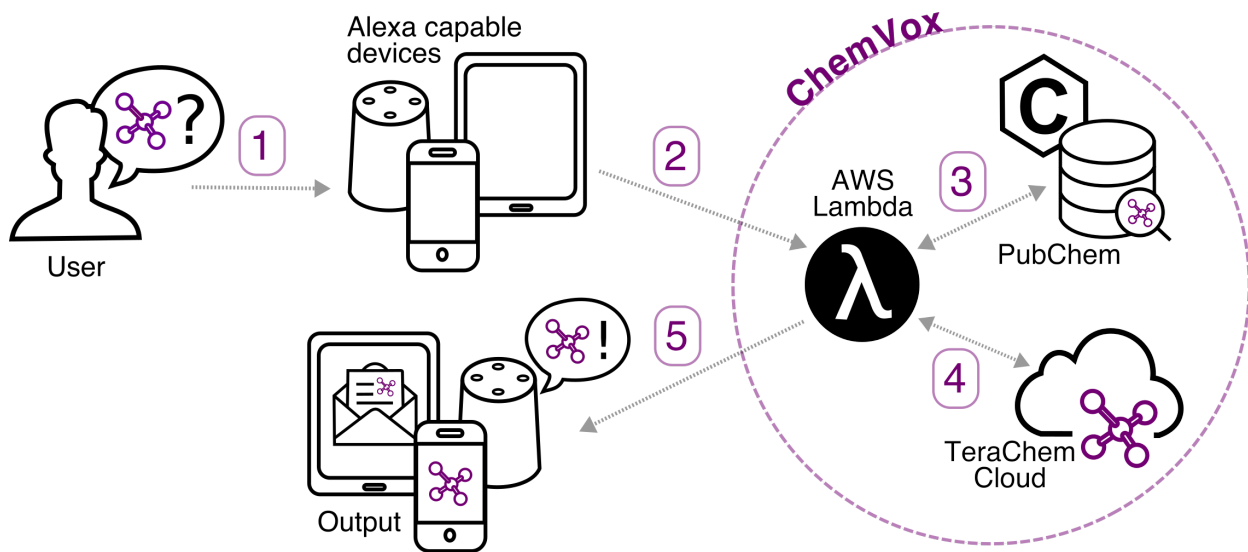


Figure 1. ChemVox workflow: 1) the user issues a request to the application through Alexa; 2) the request is passed to a Lambda function processing the information; 3) the requested molecular structure is searched and retrieved from the PubChem database; 4) the Lambda function sends the geometry and the assembled input file to TeraChem Cloud, which performs the wanted calculation; 5) the results are returned to the user through the device and via e-mail.

In its current form, ChemVox has the ability to calculate the dipole moment, the excitation energies for the two lowest lying singlet excited states and the solvatochromic shift. The calculations can be performed in gas phase or implicit solvent (spanning a broad range of dielectric constants, as discussed in the SI). Quantum chemistry calculations are performed using density functional theory for electronic ground states and time-dependent density functional theory for electronic excited states.

The ChemVox application is invoked by speaking “Open ChemVox” to Alexa. Once the application is open, the user formulates his/her request, for example “Calculate the absorption maximum of benzene in gas phase,” or “What is the dipole moment of thymine in methanol?” Follow up questions such as “What level of theory are you using?” are possible. On screen capable devices, ChemVox will display the chemical structure of a molecule if asked “Can you show me the structure of retinal?” Figure 2 shows a series of examples of such interactions. The application is terminated by saying “Goodbye.”

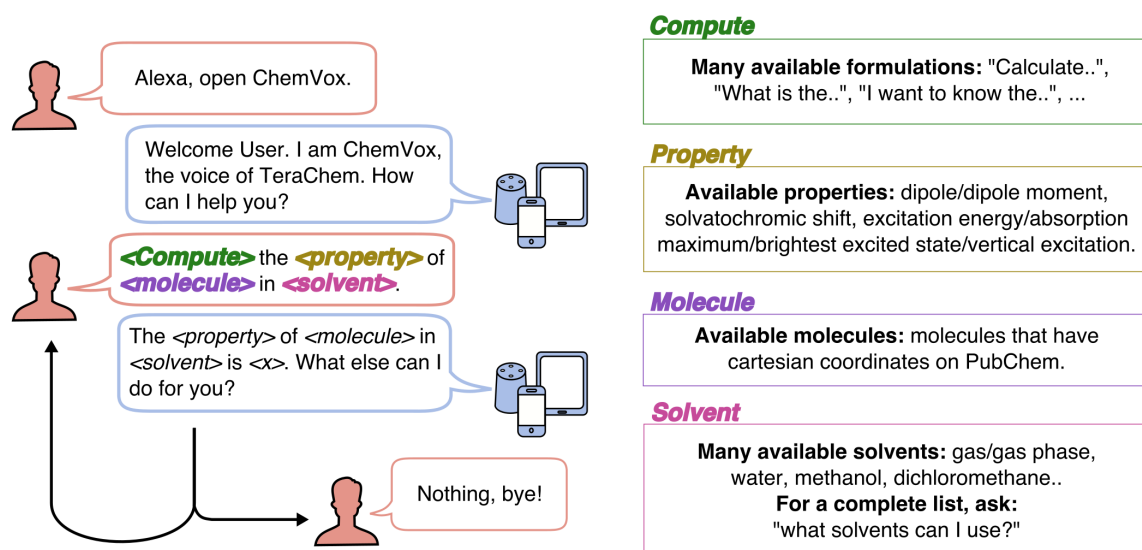


Figure 2. Capabilities of the ChemVox skill.

This project serves as a first step towards voice-controlled computational chemistry. There is certainly plenty of scope for improvements. The back-end can be extended to allow for larger molecules (up to hundreds of atoms) and enable the use of more computationally demanding theoretical methods. Additionally, the voice recognition on the front-end will

continuously become more robust as new AI algorithms come online. ChemVox is currently able to recognize ~5000 organic compounds, taken from a database containing names of the most common organic molecules. Improvement of speech recognition APIs will allow more cumbersome and complex IUPAC names to be recognized. Furthermore, the ability to interface mobile applications with virtual assistants presents the opportunity to provide alternative or additional input information and visualize the results of the requested quantum calculations (e.g. molecular orbitals, spectra). We plan to include the results of TeraChem calculations initiated by ChemVox in a web searchable database. This will reduce repetitive calculations (as ChemVox can simply retrieve the previously computed result) and ultimately provide an open access database for the scientific community.

In conclusion, we present a new tool that allows users to perform voice-activated quantum chemistry calculations in real-time. ChemVox can be used to aid scientific research and as a teaching tool to help students and the wider community interact with theoretical chemistry in an easy and accessible way. We hope that ChemVox will serve as a building block in a larger platform that combines the most advanced and easily accessible technologies (e.g. virtual assistants, virtual reality and mobile devices) to make quantum chemistry a readily accessible tool for the community.

Code Availability

Source code including video tutorial showing how to build the code into an Alexa skill is available at https://bitbucket.org/mtzcloud/chemvox_tutorial/src/master.

Acknowledgments

This work was supported by the Office of Naval Research (N00014-18-1-2659 and N00014-18-1-2624).

References

1. Hoy, M. B. Alexa, Siri, Cortana, and More: An Introduction to Voice Assistants. *Med. Ref. Serv. Quart.* **37**, 81-88, (2018).
2. Fox Rubin, B. *Amazon sees Alexa devices more than double in just one year*, <<https://www.cnet.com/news/amazon-sees-alexa-devices-more-than-double-in-just-one-year/>> (2020).
3. Paganini, M., de Oliveira, L. & Nachman, B. CaloGAN: Simulating 3D high energy particle showers in multilayer electromagnetic calorimeters with generative adversarial networks. *Phys. Rev. D* **97**, 014021, (2018).
4. Modarres, M. H., Aversa, R., Cozzini, S., Ciancio, R., Leto, A. & Brandino, G. P. Neural Network for Nanoscience Scanning Electron Microscope Image Recognition. *Sci. Rep.* **7**, 13282, (2017).
5. Wei, G.-W. Protein structure prediction beyond AlphaFold. *Nature Mach. Int.* **1**, 336-337, (2019).
6. Vandal, T., Kodra, E. & Ganguly, A. R. Intercomparison of machine learning methods for statistical downscaling: the case of daily and extreme precipitation. *Theo. Appl. Climat.* **137**, 557-570, (2019).
7. Crane-Droesch, A. Machine learning methods for crop yield prediction and climate change impact assessment in agriculture. *Environ. Res. Lett.* **13**, 114003, (2018).
8. Fluke, C. J. & Jacobs, C. Surveying the reach and maturity of machine learning and artificial intelligence in astronomy. *WIREs Data Min. Know. Disc.* **10**, e1349, (2020).
9. Maryasin, B., Marquetand, P. & Maulide, N. Machine Learning for Organic Synthesis: Are Robots Replacing Chemists? *Ang. Chem. Int. Ed.* **57**, 6978-6980, (2018).
10. Halford, B. Meet your new lab assistant. *Chem. Eng. News* **95**, 26-27 (2017). <<https://cen.acs.org/articles/95/i19/Meet-your-new-lab-assistant.html>>.
11. Schütt, K. T., Gastegger, M., Tkatchenko, A., Müller, K. R. & Maurer, R. J. Unifying machine learning and quantum chemistry with a deep neural network for molecular wavefunctions. *Nature Comm.* **10**, 5024, (2019).
12. Seritan, S., Thompson, K. & Martínez, T. J. TeraChem Cloud: A High-Performance Computing Service for Scalable Distributed GPU-Accelerated Electronic Structure Calculations. *J. Chem. Inf. Model.* **60**, 2126-2137, (2020).
13. Thackston, R. & Fortenberry, R. C. The performance of low-cost commercial cloud computing as an alternative in computational chemistry. *J. Comp. Chem.* **36**, 926-933, (2015).
14. Kim, S., Chen, J., Cheng, T., Gindulyte, A., He, J., He, S., Li, Q., Shoemaker, B. A., Thiessen, P. A., Yu, B., Zaslavsky, L., Zhang, J. & Bolton, E. E. PubChem 2019 update: improved access to chemical data. *Nucl. Acids Res.* **47**, D1102-D1109, (2018).