

Compressive Sensing in Quantum Chemistry: A Little Computation Goes a Long Way

Gregory J. O. Beran

Department of Chemistry, University of California, Riverside, California 92521, United States

Computational quantum chemistry has advanced dramatically in recent decades, and it is now routinely used in concert with experimental research. This progress stems from massive gains in computer hardware, advances in physical models (effective density functionals, for example), and creative numerical algorithms that maximize the efficiency with which those models can be solved. The exploitation of inherent sparsity, or the pattern of zeros or near-zero terms in a matrix, is a common theme in many of these numerical algorithms. Simply put: why waste time computing quantities that do not contribute significantly?

The challenge lies in the fact that the sparsity of some particular matrix in quantum chemistry depends on the representation of the matrix (the basis), and it is often difficult to identify an efficient representation *a priori*. Physical insight sometimes suggests an efficient representation, as in so-called local-correlation methods¹ which exploit the relatively short-range nature of electron–electron

Simply put: why waste time computing quantities that do not contribute significantly?

correlations that produces zeros in the integrals involving orbitals on opposite ends of a large molecule.

In this issue, Sanders, Andrade, and Aspuru-Guzik have published an interesting new approach for the prediction of molecular vibrations in electronic structure theory using compressive sensing.² Compressive sensing rose to prominence in signal processing over the past decade, and it has quickly spread to applications such as medical imaging, audio processing, and even DNA microarrays. Compressive sensing allows one to construct sparse matrices inexpensively, even if one does not know the specific sparsity pattern beforehand, by randomly sampling elements in a non-sparse representation.

Gregory Beran reflects on the role of sparse sampling techniques introduced by Aspuru-Guzik et al. to make chemistry calculations faster.

In essence, compressive sensing asks: what is the minimum number of measurements that need to be made in order to obtain a high-quality approximation for a given signal? The

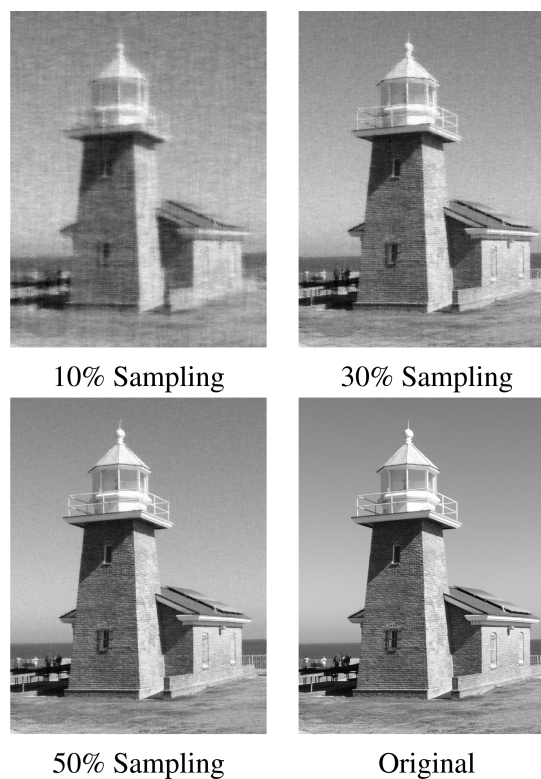


Figure 1. Compressive sensing reproduces the original photograph well even when only a fraction of the pixels are sampled.

reconstruction of a high-resolution graphical image from a low-resolution one provides a nice demonstration of the method. As shown in Figure 1, compressive sensing (using code kindly

Published: March 23, 2015

provided by Jacob Sanders) provides a very good approximation to the original picture even when only a fraction of the pixels are sampled randomly.

Compressive sensing is not magic, of course. The ability to construct an accurate approximation to the desired matrix with minimal sampling requires the existence of underlying structure in the data that produces sparsity in the optimal representation. Fortunately, real-world images contain structures and patterns that enable compressive sensing to reconstruct the true image effectively. Similarly, the matrices

[This work] is most interesting not for the specific application of compressive sensing to the calculation of molecular vibrations, but rather, for how it might alter our perspective on how quantum chemistry algorithms should work.

that arise in quantum chemistry algorithms and other scientific applications also exhibit structure, such as the stronger couplings between adjacent atoms in molecular vibrations, that lends them to compressive sensing.

In contrast to typical uses of compressive sensing for postprocessing data signals (e.g., for image reconstruction or the work by the same authors on extracting properties from molecular dynamics trajectories³ or processing data from multidimensional spectroscopies⁴), the work presented here stands out because compressive sensing is employed here to accelerate computation where the evaluation of individual matrix elements forms the computational bottleneck. In this context, the compressive sensing ideas here fit well with other recent quantum chemistry research on techniques such as orbital specific virtuals,⁵ tensor hypercontraction,⁶ multi-resolution methods,⁷ and density matrix renormalization group theory,⁸ all of which exploit sparse representations and/or low-rank approximations to model complex wave functions compactly.

The current work also provides a nice example of bootstrapping more accurate results (quantum mechanical vibrational modes) from less accurate ones (classical molecular mechanics ones). The simpler model provides a representation in which the matrix should be sparse, and that knowledge in turn reduces the amount of sampling needed to obtain an accurate approximation to the final result.

In the longer term, the article by Sanders and co-workers is most interesting not for the specific application of compressive sensing to the calculation of molecular vibrations, but rather, for how it might alter our perspective on how quantum chemistry algorithms should work—by bootstrapping our way up from simple approximations via random sampling of the matrix elements that arise in more accurate models. One can conceive of many potential applications for this technique beyond the one demonstrated here, ranging from acceleration of the repeated Fock matrix constructions required in density functional theory to cutting down on the extraordinary numbers of determinants involved in a configuration interaction calculation. This and other recent advances in sparse tensor representations in electronic structure prove there remains considerable scope for improvements in the algorithms which form the basis of modern quantum chemistry, and all chemists will benefit from the new efficiencies and physical insights yielded by these models.

Author Information

E-mail: gregory.beran@ucr.edu.

REFERENCES

- (1) Saebø, S.; Pulay, P. Local Treatment of Electron Correlation. *Annu. Rev. Phys. Chem.* **1993**, *44*, 213–236 DOI: 10.1146/annurev.pc.44.100193.001241.
- (2) Sanders, J. N.; Andrade, X.; Aspuru-Guzik, A. Compressed Sensing for the Fast Computation of Matrices: Application to Molecular Vibrations. *ACS Cent. Sci.* **2015**, DOI: 10.1021/oc5000404.
- (3) Andrade, X.; Sanders, J. N.; Aspuru-Guzik, A. Application of compressed sensing to the simulation of atomic systems. *Proc. Natl. Acad. Sci. U.S.A.* **2012**, *109*, 13928–13933 DOI: 10.1073/pnas.1209890109.
- (4) Sanders, J. N.; Saikin, S. K.; Mostame, S.; Andrade, X.; Widom, J. R.; Marcus, A. H.; Aspuru-Guzik, A. Compressed Sensing for Multidimensional Spectroscopy Experiments. *J. Phys. Chem. Lett.* **2012**, *3*, 2697–2702 DOI: 10.1021/jz300988p.
- (5) Yang, J.; Kurashige, Y.; Manby, F. R.; Chan, G. K. L. Tensor factorizations of local second-order Møller–Plesset theory. *J. Chem. Phys.* **2011**, *134*, 044123 DOI: 10.1063/1.3528935.
- (6) Hohenstein, E. G.; Parrish, R. M.; Sherrill, C. D.; Martínez, T. J. Tensor hypercontraction. III. Least-squares tensor hypercontraction for the determination of correlated wavefunctions. *J. Chem. Phys.* **2012**, *137*, 221101 DOI: 10.1063/1.4768241.
- (7) Bischoff, F. A.; Valeev, E. F. Low-order tensor approximations for electronic wave functions: Hartree–Fock method with guaranteed precision. *J. Chem. Phys.* **2011**, *134*, 104104 DOI: 10.1063/1.3560091.
- (8) Chan, G. K.-L.; Sharma, S. The Density Matrix Renormalization Group in Quantum Chemistry. *Annu. Rev. Phys. Chem.* **2011**, *62*, 465–481 DOI: 10.1146/annurev-physchem-032210-103338.