ARTICLES
https://doi.org/10.1038/s42256-020-0198-x
nature machine intelligence

Quantum approximate Bayesian computation for NMR model inference

Dries Sels[1,2]✉, Hesam Dashti[3], Samia Mora[3,4], Olga Demler[3] and Eugene Demler[1]

Recent technological advances may lead to the development of small-scale quantum computers that are capable of solving problems that cannot be tackled with classical computers. A limited number of algorithms have been proposed and their relevance to real-world problems is a subject of active investigation. Analysis of many-body quantum systems is particularly challenging for classical computers due to the exponential scaling of the Hilbert space dimension with the number of particles. Hence, solving the problems relevant to chemistry and condensed-matter physics is expected to be the first successful application of quantum computers. In this Article, we propose another class of problems from the quantum realm that can be solved efficiently on quantum computers: model inference for nuclear magnetic resonance (NMR) spectroscopy, which is important for biological and medical research. Our results are based on three interconnected studies. First, we use methods from classical machine learning to analyse a dataset of NMR spectra of small molecules. We perform stochastic neighbourhood embedding and identify clusters of spectra, and demonstrate that these clusters are correlated with the covalent structure of the molecules. Second, we propose a simple and efficient method, aided by a quantum simulator, to extract the NMR spectrum of any hypothetical molecule described by a parametric Heisenberg model. Third, we propose a simple variational Bayesian inference procedure for estimating the Hamiltonian parameters of experimentally relevant NMR spectra.

One of the central challenges for quantum technologies during the past few years has been a search for useful applications of near-term quantum machines[1]. Although considerable progress has been made in increasing the number of qubits and improving their quality[2,3], in the near future we expect the number of reliable gates to be limited by noise and decoherence—the so-called noisy intermediate-scale quantum era. As such, hybrid quantum–classical methods have been proposed to make the best of the available quantum hardware and supplement it with classical computation. Most notably, there has been the development of the quantum approximate optimization algorithm (QAOA)[4] and the variational quantum eigensolver (VQE)[5]. Both algorithms use the quantum computer to prepare variational states, some of which might be inaccessible through classical computation, but use a classical computer to update the variational parameters. A number of experiments have already been performed, demonstrating the feasibility of these algorithms[6–8], yet their bearing on real-world problems remains unclear. In model-based statistical inference one is often faced with similar problems. For simple models one can find the likelihood and maximize it, but for complex models the likelihood is typically intractable[9,10]. NMR spectroscopy is a perfect example. There is a good understanding of the type of model that should be used (equation (1)) and one only needs to determine the appropriate parameters. However, computing the NMR spectrum for a specific model requires performing computations in the exponentially large Hilbert space, which makes it extremely challenging for classical computers. This feature has been one of the original motivations for proposing NMR as a platform for quantum computing[11]. Although it has been shown that no entanglement is present during NMR experiments[12,13], strong correlations make it classically intractable; that is, the operator Schmidt rank grows exponentially, which, for example, prohibits efficient representation through tensor networks[14]. Its computational power is between that of classical computation and that of deterministic quantum computation with pure states[15], which makes it an ideal candidate for hybrid quantum–classical methods. As we argue below, the required initial quantum states can be prepared by low-depth circuits and the problem is robust against decoherence. By simulating the model on a quantum computer, it runs efficiently while the remaining inference part is simply solved on a classical computer. One can think of this as an example of quantum approximate Bayesian computation, placing it in the broader scope of quantum machine learning methods[16]. In contrast to most of the proposed quantum machine learning applications, the present algorithm does not require challenging routines such as amplitude amplification[17,18] or the Harrow–Hassidim–Lloyd (HHL) algorithm[19].

NMR spectroscopy

NMR spectroscopy is a spectroscopic technique that is sensitive to local magnetic fields around atomic nuclei. Typically, samples are placed in a high magnetic field while driving (radiofrequency, RF) transitions between the nuclear magnetic states of the system. Because these transitions are affected by the intramolecular magnetic fields around the atom and the interaction between the different nuclear spins, one can infer details about the electronic and thus chemical structure of a molecule in this way. One of the main advantages of NMR is that it is non-destructive (in contrast to X-ray crystallography or mass spectrometry, for example). This makes NMR one of the most powerful analytical techniques available to biology[20], as it is suited for in vivo and in vitro studies[21]. NMR can, for example, be used for identifying and quantifying small molecules in biological samples (serum, cerebral fluid and so on)[22–24]. On the other hand, NMR experiments have limited spectral resolution. As such, there is a challenge in interpreting the data, because the extracted information

[1]Department of Physics, Harvard University, Cambridge, MA, USA. [2]Theory of Quantum and Complex Systems, Universiteit Antwerpen, Antwerpen, Belgium. [3]Division of Preventive Medicine, Brigham and Women's Hospital, Harvard Medical School, Boston, MA, USA. [4]Division of Cardiovascular Medicine, Brigham and Women's Hospital, Harvard Medical School, Boston, MA, USA. ✉e-mail: dsels@g.harvard.edu

is quite convoluted. We only directly observe the magnetic spectrum of a biological sample, whereas our goal is to learn the underlying microscopic Hamiltonian and ultimately identify and quantify the chemical compounds. Although this inference is tractable for small molecules, it quickly becomes problematic, making inference a slow and error-prone procedure[25]. The analysis can be simplified by incorporating a priori spectral information in the parametric model[26]. For that purpose, considerable attention has been devoted to determining NMR model parameters for relevant metabolites such as those found in plasma, cerebrospinal fluid and mammalian brains[27–33].

In what follows we will be concerned with one-dimensional (1D) proton NMR, but generalization to other situations is straightforward. For liquid ^{1}H NMR, a Heisenberg Hamiltonian

$$H(\theta) = \sum_{i,j} J_{ij}\mathbf{S}_i \cdot \mathbf{S}_j + \sum_i h_i \mathbf{S}_i^x \tag{1}$$

yields a reasonable effective description for the nuclear spins, where θ explicitly denotes the dependence of the Hamiltonian on its parameters $\theta = \{J_{ij}, h\}$. Here J_{ij} encodes the interaction between the nuclear spins $\mathbf{S}$ and h_i is the effective local magnetic field. Note that this Hamiltonian contains two essential approximations: (1) the interactions are chosen to be $SU(2)$ invariant and (2) the local magnetic fields—called chemical shifts in the NMR literature—are unidirectional. The rationale for the latter is that most of these local magnetic fields are caused by diamagnetic screening due to electronic currents induced by the large external magnetic field. This field will tend to oppose the external field and is hence largely uniaxial. For liquid-state NMR, the rapid tumbling of the molecules averages out the dipolar coupling between the nuclei, approximately resulting in isotropic exchange interactions between nuclear spins[34]. The fact that the interactions are rotationally invariant allows us to remove the average (external) field from the Hamiltonian, that is $\mathbf{S}_{\mathrm{tot}}^x = \sum_i \mathbf{S}_i^x$ commutes with Hamiltonian (1) and will therefore only shift the NMR spectrum.

Within the linear response, the evolution of the system subject to a RF z-magnetic field is determined by the response function:

$$S(t|\theta) = \mathrm{Tr}\left[\mathrm{e}^{iH(\theta)t}\mathbf{S}_{\mathrm{tot}}^z \mathrm{e}^{-iH(\theta)t}\mathbf{S}_{\mathrm{tot}}^z \rho_0\right] \tag{2}$$

where ρ_0 denotes the initial density matrix of the system and $\mathbf{S}_{\mathrm{tot}}^z = \sum_i \mathbf{S}_i^z$. The measured spectrum is simply given by

$$A(\omega|\theta) = \mathrm{Re}\int_0^\infty \mathrm{d}t \mathrm{e}^{i\omega t - \gamma t} S(t|\theta) \tag{3}$$

where γ is the effective decoherence rate. For room-temperature ^{1}H NMR, the initial density matrix can be taken to be an infinite temperature state, that is

$$\rho_0 \approx \frac{\mathbb{1}}{\mathrm{Tr}[\mathbb{1}]} \tag{4}$$

Indeed, even a 20 T magnetic field will only lead to a bare proton resonance frequency of ~900 MHz. In contrast, room temperature is ~40 THz, so for all practical purposes we can consider it equally likely for the spin to be in the excited state or in the ground state. Chemical shifts h_i are of the order of a few parts per million, resulting in local energy shifts of a few kHz, while the coupling or interaction strength J is of the order of a few Hz. Despite these low frequencies and the high temperature of the system, one can typically still infer the parameters due to the small decoherence rate of the proton nuclear spin. Owing to the absence of a magnetic quadrupole moment, the protons do not decohere from the electric dipole fluctuations caused by the surrounding water molecules. This gives the proton nuclear spin a coherence time of the order of seconds to tens of seconds, sufficiently long to create some correlations between the various spins. The remaining part of this work is concerned with the question of how to infer the model parameters J_{ij} and h_i of our effective Hamiltonian (1) from measured spectrum (3).

Clustering

Given real NMR data, summarized by the experimentally acquired spectrum $\mathcal{A}(\omega)$, our goal, in general, is to learn a parametrized generative model that explains how the NMR data are generated. Fortunately, we have a good idea about the physics, which allows us to write down a model—equation (3)—that is close to reality, thereby ensuring a small misspecification error. The drawback, however, is that the model is analytically intractable and becomes increasingly complex to simulate with increasing number of spins. In the next section we will discuss how to alleviate this problem by using a programmable quantum simulator to simulate the problem instead. Even if we can simulate our model (3), we still have to find a reliable and robust way to estimate parameters θ. Physical molecules have far from typical parameters θ (see Supplementary Information for a mathematical description). After all, if they do not, how could we infer any structural information out of the spectrum? To extract NMR spectral features, we first perform unsupervised learning on a dataset containing 69 small organic molecules, all composed of four ^{1}H atoms, observable in NMR 1D-^{1}H experiments. Their effective Hamiltonian parameters θ have been determined previously, which provides us with a labelled dataset to test our procedure. Furthermore, by only using the spectra themselves, we can use any relevant information as an initial prior for inference on unknown molecules. The dataset was compiled using the GISSMO library[30,31,35]. To extract the structure in the dataset, we perform a t-distributed stochastic neighbourhood embedding (t-SNE)[36,37] to visualize the data in two dimensions. Figure 1b shows the 2D t-SNE embedding of the dataset based on the Hellinger distance shown in Fig. 1a (a detailed comparison of different metrics is presented in the Supplementary Information). The colour scale in Fig. 1b shows the inverse participation ratio (IPR) of each sample, $\mathrm{IPR} = \int_{-\infty}^{\infty} \mathrm{d}\omega A(\omega|\theta) / \int_{-\infty}^{\infty} \mathrm{d}\omega A^2(\omega|\theta)$, a measure for the total number of transitions that contribute to the spectrum. At least four well-defined clusters are identified, and density-based spatial clustering of applications with noise (DBSCAN)[38] was used to perform the clustering. Using the clusters as indicated in Fig. 1b, we can sort the molecules per cluster and have a look at the spectra. The sorted distance matrix is shown in Fig. 2a, which clearly shows that we have managed to find most of the structures in the system. In fact, a closer look at the spectra of each of the clusters indeed reveals they are all very similar. Figure 2b shows a representative spectrum for each of the clusters and, as expected, the IPR goes up if we go from cluster 1 to cluster 4. All spectra in cluster 1 have the property of containing two large peaks and two small peaks, where the larger peak is about three times higher than the small peak. This is indicative of molecules with a methyl group (CH_3) with its protons coupled with a methine proton (CH). One example of such structures can be seen in acetaldehyde oxime (BMRB ID[39] bmse000467) (as shown to the left in Fig. 2b). The fact that the three protons are equivalent results in the 3:1 ratio of the peaks. Molecules from cluster 2 are highly symmetric and have two pairs of two methine protons (CH) where the protons are on neighbouring carbon atoms. The symmetry in the molecule makes the spectrum highly degenerate. In contrast, cluster 3 has molecules where there are two neighbouring methylene groups (CH_2). The interacting splitting causes a spectrum as shown in Fig. 2b. Finally, cluster 4 has four inequivalent protons with different chemical shifts and interactions between them. As a result, there are a plethora of possible transitions and the spectrum has an erratic form, as shown in Fig. 2b. In that sense, cluster 4 is most like a disordered quantum spin chain.

Given a new spectrum of an unknown molecule, we can find out whether the molecule belongs to any of the identified molecular

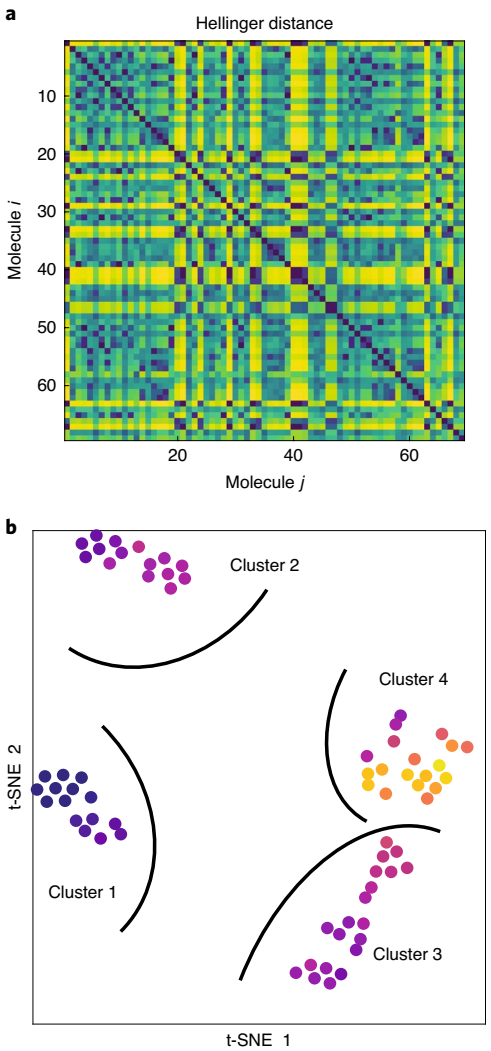

Fig. 1 | Clustering analysis to identify whether naturally occurring molecules have an atypical NMR spectrum. a, The distance between the various NMR spectra, where the Bhattacharyya coefficient is used to measure similarity. To obtain a meaningful comparison, spectra are shifted and scaled such that they are all centred around the same frequency and have the same bandwidth. **b**, To extract clusters we perform a t-SNE with perplexity of 10. This is chosen because it has minimal Kullback–Leibler (KL) loss (the KL loss was 0.145).

sub-structures; that is, by computing the mean Hellinger distance to each of the identifed clusters one can robustly classify the spectra. In the Supplementary Information we present results in which we randomly choose 39 samples and consider those as clustered, while we use the other 30 samples to test the procedure. Samples that belong to clusters 1 and 2 are always correctly classified. One sample from cluster 4 was misclassified for cluster 3. Because we know the spin matrix θ_i for each of the molecules in the dataset, we have a rough estimate of what the Hamiltonian parameters are and where the protons are located with respect to each other. However, there is still a lot of fine structure within clusters, in particular in clusters 3 and 4, as can be seen in Fig. 2a. In what remains, we are concerned with finding an algorithm to further improve the Hamiltonian parameter estimation.

Quantum computation

Although our model is microscopically motivated, thereby capturing the spectra very well and allowing for a physical interpretation of the model parameters, it has the drawback that, unlike simple models such as Lorentzian mixture models[40,41], there is no analytic form for the spectrum in terms of the model parameters. Moreover, even simulating the model becomes increasingly complex when the number of spins increases. Before we solve the inference problem, let us present an efficient method to extract the simulated NMR spectrum on a quantum simulator–computer. The basic task is to extract spectrum (3) by measuring (2). Recall that we work at infinite temperature, so by inserting an eigenbasis of the total z-magnetization $\mathbf{S}^z_{\mathrm{tot}} = \sum_j m_j |z_j\rangle\langle z_j|$, we find

$$S(t|\theta) = \sum_{i,j} m_i m_j P_{\mathrm{t}}(i|j,\theta) P_0(j) \tag{5}$$

with the transition probability $P_{\mathrm{t}}(i|j,\theta) = \left|\langle z_i|U_\theta(t)|z_j\rangle\right|^2$, initial distribution $P_0(j) = 2^{-N}$ and m_j is the total z-magnetization in the eigenstate $|z_j\rangle$. Consequently, we can extract the spectrum by initializing our system in a product state of z-polarized states, after which we quench the system to evolve under $U_\theta(t)$, generated by Hamiltonian $H(\theta)$, and then finally perform a projective measurement in the z-basis again at time t. By repeating the procedure by uniformly sampling the initial eigenstates and estimating the product of the initial and final magnetization $m_i m_j$, one obtains an estimate of $S(t|\theta)$ (Fig. 3). Note that, at this stage, the problem is entirely classical and all quantum physics is hidden in the transition probablity $P_{\mathrm{t}}(i|j,\theta)$. It is the intractability of this transition probability that forms the basis of recent quantum supremacy experiments[42].

In contrast to the latter, we are only interested in estimating a simple statistic, namely the average $m_i m_j$. Note that this quantity is bounded by $N^2/4$, so, according to Hoeffding's inequality, one needs to sample at most $O(N^4/\epsilon^2)$ times to get a precision of ϵ on $S(t|\theta)$. At present, the structure of equation (5) allows one to bound the variance of $m_i m_j$ by $3(N/4)^2$, such that $O(N^2/\epsilon^2)$ would suffice. As shown in detail in the Supplementary Information, one can in general not improve on this scaling with N unless one uses additional structure of the transition probability P_{t}. At short times, one benefits from importance sampling, for example. Although we have no control over the transition probability P_{t}, we can control the initial probability out of which we sample states, as long as those states are easy to prepare. Because equation (5) is diagonal in the z-basis, it is sufficient to consider sampling product states in the z-basis; that is, one can equivalently write the response function as

$$S(t|\theta) = \sum_{i,j} \left(\frac{m_i m_j P_0(j)}{Q_0(j)}\right) P_{\mathrm{t}}(i|j,\theta) Q_0(j) \tag{6}$$

where Q_0 is the distribution from which we sample. By minimizing the variance of estimand $r = m_i m_j P_0/Q_0$, one obtains an optimal sampling distribution. The true optimal depends on time through P_{t} and, given that this is unknown to us, we must settle for a good, albeit suboptimal, distribution Q_0. Various approximations might be considered, but the distribution

$$Q_0(j) = \frac{4}{N}\frac{m_j^2}{2^N} \tag{7}$$

is particularly interesting because it gives zero variance for r at $t = 0$ and at any other time the variance is smaller than $(N/4)^2$. Consequently, we can estimate $S(t|\theta)$ with precision ϵ by taking at most $O(N^2/\epsilon^2)$ samples. Given the finite decoherence rate γ and the fact that the energy bandwidth of the many-body spectrum scales linearly with N, one needs to measure $S(t|\theta)$ at worst in time steps of the order of $1/N$ up to a time that scales as $1/\gamma$. One thus has to repeat the entire circuit at worst $O(N^3/\epsilon^2)$ times. Furthermore, if the time evolution is implemented as an analogue simulation, this takes a time of $O(1/\gamma)$. The gate complexity is at worst a factor of N^2 more because one, at worst, has to

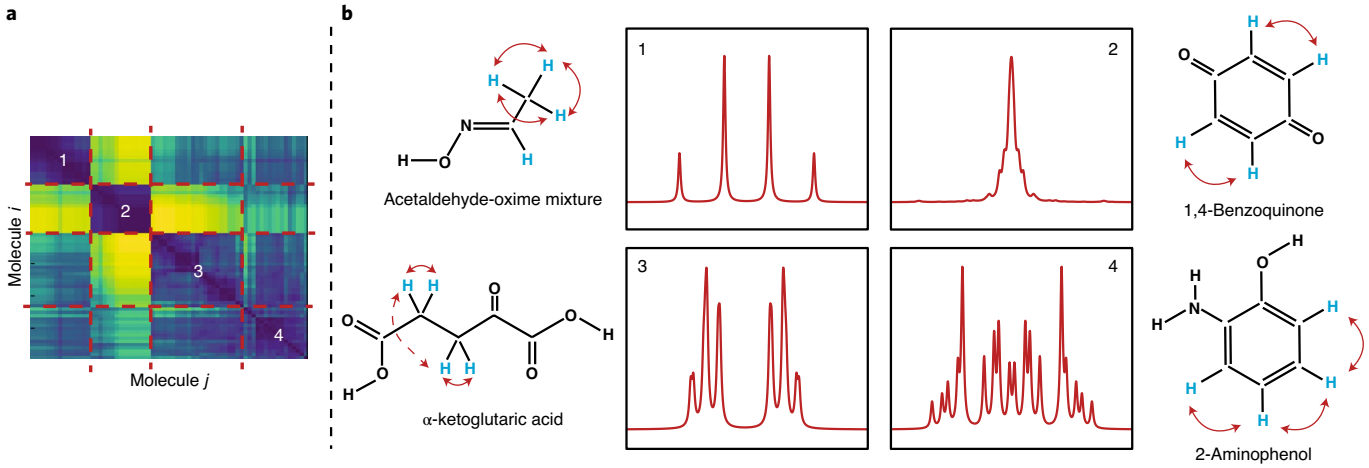

Fig. 2 | NMR spectra. a, By clustering the molecules according to the Hellinger distance t-SNE clusters, we can reorganize the distance matrix as shown. For each of the clusters, we look at the different spectra, which indeed show great similarity. **b**, A representative spectrum for each of clusters 1–4, where the spectra are labelled according to the t-SNE clusters shown in Fig. 1b. In addition, we show an example small molecule from this cluster next to the associated spectrum. The atoms and interactions responsible for the shown portions of the spectra are indicated by blue and red arrows, respectively.

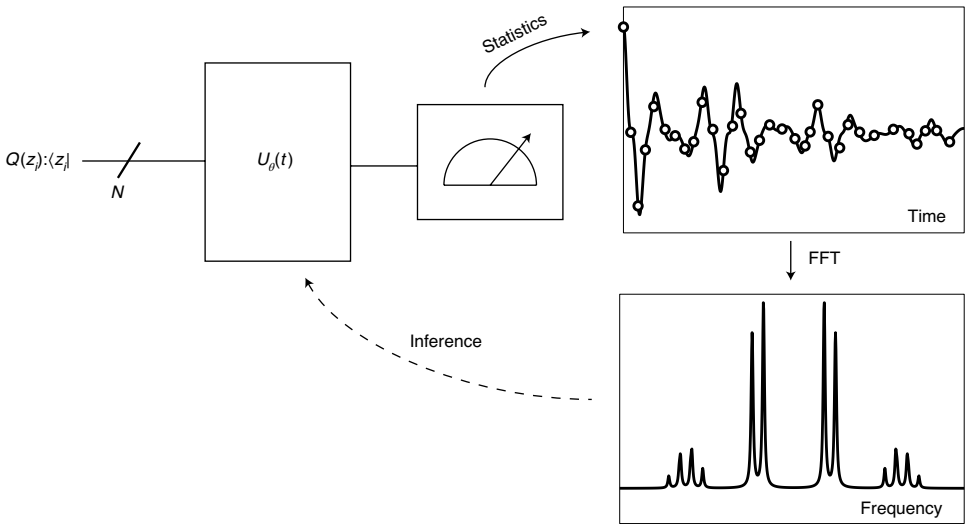

Fig. 3 | Method overview. We take a product state with a given total magnetization m_i, according to $Q_0(j)$. The latter can be chosen to minimize the variance of the estimand. After this initial preparation, we evolve the state under the Hamiltonian $H(\theta)$ and measure the project back onto the z-basis at time t. By applying a fast Fourier transform to the estimate $S(t|\theta)$, one obtains the spectrum, which can be used to infer the parameters of the Hamiltonian.

implement a Heisenberg interaction between all possible qubits, yielding $O(N^2/\gamma)$. Note that these are worst-case scalings; for an extensive spectrum one actually expects linear scaling of the gate complexity with N. Typical transitions happen between states that only differ by an energy of $O(1)$, such that the typical sampling complexity is only quadratic with N.

Variational Bayesian inference

Now that we have a procedure for efficiently obtaining the spectra of hypothetical molecules, how do we solve the inference problem? The standard approach would be to do maximum likelihood estimation of the parameters given the experimental spectrum or minimize one of the aforementioned cost functions. This cannot be done analytically and the problem can clearly be highly non-convex. We thus require a method to numerically minimize the error; gradient descent seems an obvious choice but is not well suited for this task. First, there is the obvious problem that additional measurements will need to performed to estimate the gradients. Those estimates are not easy to obtain because they require the measurement of three-point correlators in time. Moreover, using a quantum simulator, one only obtains a statistical estimate of the cost function and its gradient, because we only perform a finite number of measurements. To move down the optimization landscape we thus need to resolve the signal from the noise, meaning gradients have to be sufficiently large to be resolved. However, we find extremely small gradients for this problem. Taking, for example, the Hellinger distance, D_H, used to construct Fig. 1, we find the gradient satisfies

$$\left|\partial_\theta D_H^2\right| = \left|\int \frac{d\omega}{2\pi}\sqrt{\frac{\mathcal{A}(\omega)}{A(\omega|\theta)}}\partial_\theta A(\omega|\theta)\right| \leq \sqrt{I_{\theta\theta}} \tag{8}$$

where $I_{\theta\theta}$ is the diagonal component of the Fisher information. The bound simply follows from the Cauchy–Schwarz inequality. As

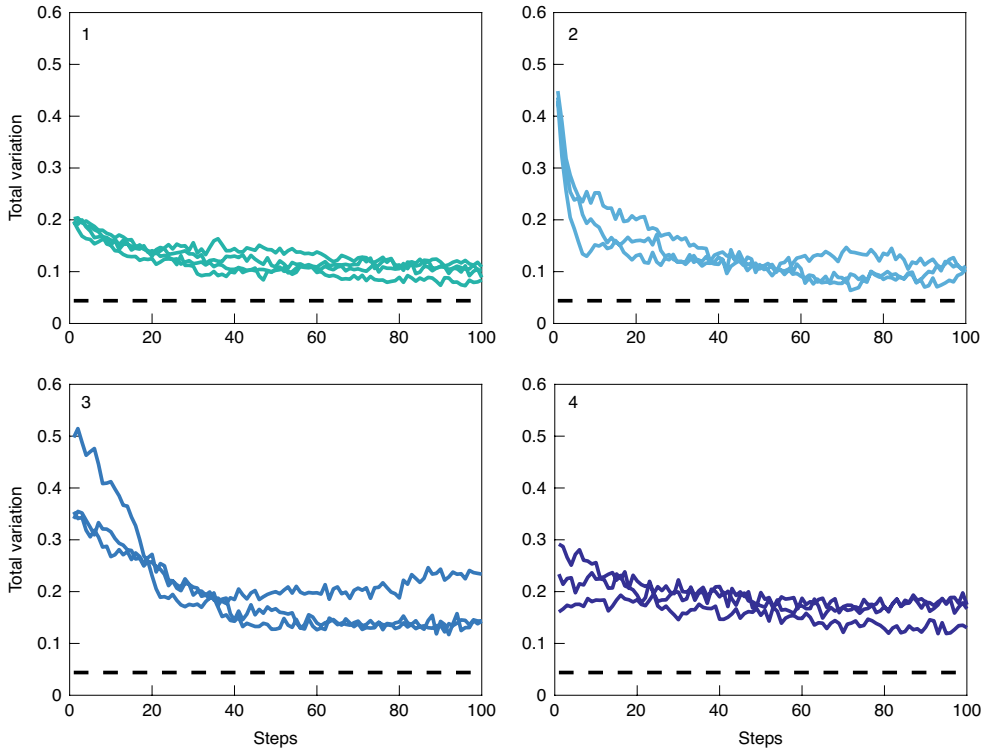

Fig. 4 | Inference. For each of the clusters, labelled according to Fig. 2, we investigate the convergence of the parameter inference in our variational Bayesian inference scheme by looking at the total variation distance between the spectra. Dashed lines indicate the shot noise limit, set by the finite number of acquired quantum measurements.

shown in the Supplementary Information, the Fisher information, even for optimal values, is very small, typically of the order 10^{-4}–10^{-6} for our four spin molecules. We are thus in a situation where there is a very shallow rough optimization landscape. The problem is of similar origin as the vanishing gradient problem in quantum neural networks[43]. A gradient free method seems advisable. Here, we adopt a Bayesian type of approach to update our estimated parameters. Alternative approaches, such as the DIRECT method adopted in ref. [44], are expected to work as well, but more research on the structure of the optimization landscape is required to understand the hardness of this inference problem. Recall that Bayes theorem, in the current notation, reads as

$$P(\theta|\omega) = \frac{A(\omega|\theta)P(\theta)}{A(\omega)} \tag{9}$$

where $P(\theta|\omega)$ is the conditional probability to have parameters θ given that we see spectral weight at frequency ω, $A(\omega|\theta)$ is the NMR spectrum for fixed parameters θ, $P(\theta)$ is the probability to have parameters θ and $A(\omega)$ is the marginal NMR spectrum averaged over all θ. If we acquire some data, say a new spectrum $\mathcal{A}(\omega)$, and we have some prior belief about the distribution $P(\theta)$, we can use it to update our belief about the distribution of the parameters; that is

$$P_{i+1}(\theta) = \int \frac{\mathrm{d}\omega}{2\pi}\mathcal{A}(\omega)\frac{A(\omega|\theta)}{A_i(\omega)}P_i(\theta) \tag{10}$$

with $A_i(\omega) = \int \mathrm{d}\theta A(\omega|\theta)_{P_i}(\theta)$. Note that the above rule indeed conserves positivity and normalization. Moreover, it simply reweights the prior distribution with some weight

$$w_i(\theta) = \int \frac{\mathrm{d}\omega}{2\pi}\mathcal{A}(\omega)\frac{A(\omega|\theta)}{A_i(\omega)} \tag{11}$$

that is directly related to the log-likelihood, because the Jensen inequality gives

$$\log\left(\mathrm{w_i}(\theta)\right) \geq \int \frac{\mathrm{d}\omega}{2\pi}\mathcal{A}(\omega)\log\frac{\mathrm{A}(\omega|\theta)}{\mathrm{A_i}(\omega)} = \mathcal{L}(\theta) + \mathrm{c} \tag{12}$$

where $\mathcal{L}(\theta)$ is the log-likehood and c is a constant independent of θ. Consequently, iterating equation (10) is expected to converge to a distribution of parameters that is highly peaked around the maximum likelihood estimate. Although it avoids the use of any gradients, it requires us to sample from the current parameter distribution $_{P_i}(\theta)$. This by itself could become intractable and so we make an additional approximation. To be able to sample from the parameter distribution, we approximate it by a normal distribution at every step. That is, given that we have obtained some Monte Carlo samples out of $_{P_i}(\theta)$, we can estimate all the weights $w_i(\theta)$ by simply simulating the model and obtaining $A(\omega|\theta_i)$ for all the samples. Next, we approximate $_{P_{i+1}}(\theta)$ with a normal distribution that is as close as possible to it; that is, it has minimal KL distance. The latter is simply the distribution with the same sample mean and covariance as $_{P_{i+1}}(\theta)$. We use an atomic prior, $P_0(\theta) = \sum_i \frac{1}{N_s}\delta(\theta - \theta_i)$, consisting of all the samples that belong to the same cluster to which the spectrum is identified to belong. The result of this procedure for some randomly chosen test molecules is shown in Fig. 4. We observe steady, albeit noisy, convergence of the molecular spectra. Convergence is limited by three factors: (1) shot noise from the quantum measurements, (2) sampling noise from the Monte Carlo procedure and (3) the Gaussian variational approximation. Although both noise sources can be made smaller by using more computational resources, more advanced methods ultimately seem to be needed.

Summary and outlook

We have presented a method to improve model inference for NMR with a relatively modest amount of quantum resources. Similar to

generic generative models such as Boltzmann machines, for which a more efficient quantum version has been constructed[45,46], we have constructed an application-specific model from which a quantum machine can sample more efficiently than a classical computer. Model parameters are determined through a variational Bayesian approach with an informative prior, constructed by applying t-SNE to a dataset of small molecules. As a consequence of the noisy nature of the generative model, as well as the absence of significant gradients, both the initial bias as well as the derivative-free nature of Bayesian inference are crucial to tackling the problem. This situation, however, is generic to any hybrid quantum–classical setting that is sufficiently complicated. A similar approach might thus be useful to improve convergence of QAOA or VQE. For example, heuristic optimization strategies for QAOA have been developed in ref. [47]. Both the classical and quantum parts of our approach can be extended further. On the quantum side, one can envision developing more efficient approaches for computing the spectra—trading computational time for extra quantum resources. On the classical side, improvements on the inference algorithm might be possible by combining or extending the variational method with Hamiltonian Monte Carlo techniques[48].

It is interesting to extend our technique to other types of experiment. NMR is hardly the only problem where one performs inference on spectroscopic data. For example, one can imagine combining resonant inelastic X-ray scattering (RIXS) data from strongly correlated electron systems[49] with Fermi–Hubbard simulators based on ultracold atoms[50,51]. Currently, RIXS data are analysed by performing numerical studies of small clusters on classical computers (see ref. [52] for a review). A dynamical mean field theory-based hybrid algorithm was recently proposed in ref. [53]. With cold atoms in optical lattices one may be able to create larger systems and study their non-equilibrium dynamics corresponding to RIXS spectroscopy.

Data availability

The data and code to numerically generate the NMR data sets used in this manuscript can be found at https://github.com/dsels/QuantumNMR.

Received: 17 September 2019; Accepted: 2 June 2020;
Published online: 6 July 2020

References

1. Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79–98 (2018).
2. Bernien, H. et al. Probing many-body dynamics on a 51-atom quantum simulator. *Nature* **551**, 579–584 (2017).
3. Friis, N. et al. Observation of entangled states of a fully controlled 20-qubit system. *Phys. Rev. X* **8**, 021012 (2018).
4. Farhi, E., Goldstone, J. & Gutmann, S. A quantum approximate optimization algorithm. Preprint at https://arxiv.org/abs/1411.4028 (2014).
5. Peruzzo, A. et al. A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.* **5**, 4213 (2014).
6. Kokail, C. et al. Self-verifying variational quantum simulation of lattice models. *Nature* **569**, 355–360 (2019).
7. Kandala, A. et al. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* **549**, 242–246 (2017).
8. Colless, J. I. et al. Computation of molecular spectra on a quantum processor with an error-resilient algorithm. *Phys. Rev. X* **8**, 011021 (2018).
9. Diggle, P. J. & Gratton, R. J. Monte Carlo methods of inference for implicit statistical models. *J. R. Stat. Soc. B* **46**, 193–227 (1984).
10. Beaumont, M. A., Zhang, W. & Balding, D. J. Approximate Bayesian computation in population genetics. *Genetics* **162**, 2025–2035 (2002).
11. Gershenfeld, N. A. & Chuang, I. L. Bulk spin-resonance quantum computation. *Science* **275**, 350–356 (1997).
12. Braunstein, S. L. et al. Separability of very noisy mixed states and implications for NMR quantum computing. *Phys. Rev. Lett.* **83**, 1054–1057 (1999).
13. Menicucci, N. C. & Caves, C. M. Local realistic model for the dynamics of bulk-ensemble NMR information processing. *Phys. Rev. Lett.* **88**, 167901 (2002).
14. Datta, A. & Vidal, G. Role of entanglement and correlations in mixed-state quantum computation. *Phys. Rev. A* **75**, 042310 (2007).
15. Knill, E. & Laflamme, R. Power of one bit of quantum information. *Phys. Rev. Lett.* **81**, 5672–5675 (1998).
16. Biamonte, J. et al. Quantum machine learning. *Nature* **549**, 195–202 (2017).
17. Brassard, G. & Hoyer, P. An exact quantum polynomial-time algorithm for Simon's problem. In *Proceedings of the Fifth Israeli Symposium on Theory of Computing and Systems* 12–23 (IEEE, 1997).<q>
18. Grover, L. K. Quantum computers can search rapidly by using almost any transformation. *Phys. Rev. Lett.* **80**, 4329–4332 (1998).
19. Harrow, A. W., Hassidim, A. & Lloyd, S. Quantum algorithm for linear systems of equations. *Phys. Rev. Lett.* **103**, 150502 (2009).
20. Bothwell, J. H. F. & Griffin, J. L. An introduction to biological nuclear magnetic resonance spectroscopy. *Biol. Rev.* **86**, 493–510 (2011).
21. Hwang, J.-H. & Choi, C. S. Use of in vivo magnetic resonance spectroscopy for studying metabolic diseases. *Exp. Mol. Med.* **47**, e139 (2015).
22. Beckonert, O. et al. Metabolic profiling, metabolomic and metabonomic procedures for NMR spectroscopy of urine, plasma, serum and tissue extracts. *Nat. Protoc.* **2**, 2692–2703 (2007).
23. Larive, C. K., Barding, G. A. & Dinges, M. M. NMR spectroscopy for metabolomics and metabolic profiling. *Anal. Chem.* **87**, 133–146 (2015).
24. Napolitano, J. et al. Proton fingerprints portray molecular structures: enhanced description of the ^{1}H NMR spectra of small molecules. *J. Org. Chem.* **78**, 9963–9968 (2013).
25. Ravanbakhsh, S. et al. Accurate, fully-automated NMR spectral profiling for metabolomics. *PLoS ONE* **10**, e0124219 (2015).
26. De Graaf, A. A. & Boveé, W. M. M. J. Improved quantification of in vivo ^{1}H NMR spectra by optimization of signal acquisition and processing and by incorporation of prior knowledge into the spectral fitting. *Magn. Reson. Med.* **15**, 305–319 (1990).
27. Wevers, R. A., Engelke, U. & Heerschap, A. High-resolution ^{1}H-NMR spectroscopy of blood plasma for metabolic studies. *Clin. Chem.* **40**, 1245–1250 (1994).
28. Wevers, R. A. et al. Standardized method for high-resolution ^{1}H-NMR of cerebrospinal fluid. *Clin. Chem.* **41**, 744–751 (1995).
29. Govindaraju, V., Young, K. & Maudsley, A. A. Proton NMR chemical shifts and coupling constants for brain metabolites. *NMR Biomed.* **13**, 129–153 (2000).
30. Dashti, H. et al. Applications of parametrized NMR spin systems of small molecules. *Anal. Chem.* **90**, 10646–10649 (2018).
31. Dashti, H. et al. Spin system modeling of nuclear magnetic resonance spectra for applications in metabolomics and small molecule screening. *Anal. Chem.* **89**, 12201–12208 (2017).
32. Pickard, C. J. & Mauri, F. All-electron magnetic response with pseudopotentials: NMR chemical shifts. *Phys. Rev. B* **63**, 245101 (2001).
33. Paruzzo, F. M. et al. Chemical shifts in molecular solids by machine learning. *Nat. Commun.* **9**, 4501 (2018).
34. Levitt, M. H. *Spin Dynamics*: *Basics of Nuclear Magnetic Resonance* (Wiley, 2008).
35. Dashti, H. Guided Ideographic Spin System Model Optimization (GISSMO)<q> (2019); http://gissmo.nmrfam.wisc.edu/
36. van der Maaten, L. & Hinton, G. Visualizing data using t-SNE. *J. Mach. Learn. Res.* **9**, 2579–2605 (2008).
37. van der Maaten, L. https://lvdmaaten.github.io/tsne/ (2019).
38. Ester, M., Kriegel, H.-P., Sander, J. & Xu, X. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge*, *Discovery and Data Mining* 226–231 (AAAI Press, 1996).
39. Ulrich, E. L. et al. Biomagresbank. *Nucleic Acids Res.* **36**, D402–D408 (2008).
40. Sokolenko, S. et al. Robust 1D NMR lineshape fitting using real and imaginary data in the frequency domain. *J. Magn. Reson.* **298**, 91–100 (2019).
41. Xu, K., Marrelec, G., Bernard, S. & Grimal, Q. Lorentzian-model-based Bayesian analysis for automated estimation of attenuated resonance spectrum. *IEEE Trans. Signal Process.* **67**, 4–16 (2019).
42. Arute, F. et al. Quantum supremacy using a programmable superconducting processor. *Nature* **574**, 505–510 (2019).
43. McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R. & Neven, H. Barren plateaus in quantum neural network training landscapes. *Nat. Commun.* **9**, 4812 (2018).
44. Kokail, C. et al. Self-verifying variational quantum simulation of lattice models. *Nature* **569**, 355–360 (2019).
45. Kieferová, M. & Wiebe, N. Tomography and generative training with quantum Boltzmann machines. *Phys. Rev. A* **96**, 062327 (2017).
46. Amin, M. H., Andriyash, E., Rolfe, J., Kulchytskyy, B. & Melko, R. Quantum Boltzmann machine. *Phys. Rev. X* **8**, 021050 (2018).
47. Zhou, L., Wang, S.-T., Choi, S., Pichler, H. & Lukin, M. D. Quantum approximate optimization algorithm: performance, mechanism and implementation on near-term devices. Preprint at https://arxiv.org/abs/1812.01041 (2018).
48. Radford, N. in *Handbook of Markov Chain Monte Carlo* (eds Brooks, S. et al.) 116–162 (CRC, 2011).
49. Murakami, Y. & Ishihara, S. (eds) *Resonant X-ray Scattering in Correlated Systems* (Springer, 2017).

50. Hofstetter, W., Cirac, J. I., Zoller, P., Demler, E. & Lukin, M. D. High-temperature superfluidity of fermionic atoms in optical lattices. *Phys. Rev. Lett.* **89**, 220407 (2002).
51. Bloch, I., Dalibard, J. & Zwerger, W. Many-body physics with ultracold gases. *Rev. Mod. Phys.* **80**, 885–964 (2008).
52. Ament, L. J. P., van Veenendaal, M., Devereaux, T. P., Hill, J. P. & van den Brink, J. Resonant inelastic X-ray scattering studies of elementary excitations. *Rev. Mod. Phys.* **83**, 705–767 (2011).
53. Kreula, J. M. et al. Few-qubit quantum-classical simulation of strongly correlated lattice fermions. *EPJ Quantum Technol.* **3**, 11 (2016).

Acknowledgements

D.S. acknowledges support from the FWO as post-doctoral fellow of the Research Foundation—Flanders and from a 2019 grant from the Harvard Quantum Initiative Seed Funding programme. S.M. is supported by a research grant from the National Heart, Lung, and Blood Institute (K24 HL136852). O.D. and H.D. are supported by a research award from the National Heart, Lung, and Blood Institute, (5K01HL135342) and (T32 HL007575) respectively. E.D. acknowledges support from the Harvard–MIT CUA, ARO grant number W911NF-20-1-0163, the National Science Foundation through grant number OAC-1934714, AFOSR Quantum Simulation MURI.The authors acknowledge useful discussions with P. Mehta and M. Lukin.

Author contributions

D.S. and E.D. conceived the presented idea in consultation with H.D., S.M. and O.D. D.S. developed the theoretical formalism, performed the analytic calculations and performed the numerical simulations. H.D. compiled the NMR data used in the manuscript. All authors provided critical feedback and helped shape the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at https://doi.org/10.1038/s42256-020-0198-x.

Correspondence and requests for materials should be addressed to D.S.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.