Quantum computing at the frontiers of biological sciences

Computing plays a critical role in the biological sciences but faces increasing challenges of scale and complexity. Quantum computing, a computational paradigm exploiting the unique properties of quantum mechanical analogs of classical bits, seeks to address many of these challenges. We discuss the potential for quantum computing to aid in the merging of insights across different areas of biological sciences.

Prashant S. Emani, Jonathan Warrell, Alan Anticevic, Stefan Bekiranov, Michael Gandal, Michael J. McConnell, Guillermo Sapiro, Alán Aspuru-Guzik, Justin T. Baker, Matteo Bastiani, John D. Murray, Stamatios N. Sotiropoulos, Jacob Taylor, Geetha Senthil, Thomas Lehner, Mark B. Gerstein and Aram W. Harrow

Understanding complex biological phenomena has required concurrent advances in experiment, theory and computing power. The increasing need for computing infrastructure has led to expansions of current supercomputing and other massively parallel computing facilities, but also considerations of entirely new computing paradigms. Here we consider the potential of quantum computing (QC) to address complex biological questions. QC is an approach to computation in which an algorithm is defined by a series of operations on quantum states that results in a solution to a problem. Recent technological developments have carried QC capabilities from the realm of academic exploration to commercial opportunities[1,2]. While the scale is not yet competitive with classical technologies, there is substantial excitement in its eventual promise, and we hope to provide an entry point for biologists to certain aspects of the discussion surrounding QC. This effort is especially timely given recent policy efforts at a national or international level, such as the US National Quantum Initiative Act of 2018[3] (the implementation of a National Quantum Initiative for quantum information science and technology[4]), the European Quantum Technologies Flagship, and efforts in the United Kingdom and China[5].

We first present a primer on quantum computation to familiarize the reader with the basic concepts and language of QC. The remainder is focused on the study of the human brain through genetics, genomics, neuroimaging and deep behavioral phenotyping, a multidisciplinary effort that falls under the term "convergent neuroscience." We highlight these areas as they exemplify two sources of complexity: separately, each field presents a rich set of problems that often push the limits of classical computational capability; in combination, they offer a multiscale challenge leading from the molecular scale through the cellular and tissue levels to brain architecture and, eventually, to complex human behaviors and disorders. The study of the emergent properties of the brain, such as cognition and behavior, is a uniquely challenging multilevel endeavor that demands pioneering approaches in computation. Accordingly, we discuss how quantum algorithms that map onto methodological issues in the biological sciences may provide much needed improvements in computational efficiency, and we posit open questions for eventual development of new computational solutions.

Classical versus quantum circuits: state of the art

Quantum computing uses the laws of quantum mechanics to perform computations. Quantum mechanics is the physical theory that governs all matter but is particularly relevant at the molecular scale and below. It states that particles have wave-like properties and waves have particle-like properties. If a quantum computer could be built, then this wave-like behavior could be harnessed for computational benefit: in a conventional (classical) computer using randomness, different random choices can lead to different outcomes, and the total probability of an outcome is the sum of the probabilities of each computational path leading to that outcome; by contrast, a quantum computer can have complex amplitudes along computational paths, just as a wave can have different amplitudes in different modes. Measuring will 'collapse' the state and yield a specific outcome with probability equal to the squared absolute value of the amplitude. Thus, quantum computers promise a new form of computing that would be qualitatively different from any previous (classical) form of computation by allowing interference between computational paths, analogous to the interference between waves[6]. While quantum computers are technically more difficult to build, and the best current general-purpose quantum computers have only 50–100 qubits, they can solve some problems in a time that grows more slowly as a function of the input size than classical computation. The term "qubit" refers to a quantum two-level system, such as a photon that can travel down one of two optical fibers. Qubits can be thought of as a generalization of classical bits (cbits): cbits can be in states 0 or 1, while the state of a single qubit is described by complex numbers α_0 and α_1 satisfying $|\alpha_0|^2 + |\alpha_1|^2 = 1$. The power of quantum computers comes from scaling. A system of n cbits can be in one of 2^n possible states at any time, while the state of n qubits is described by a complex unit vector of dimension 2^n (Fig. 1a,b). These vectors (also called wavevectors or wavefunctions) can be transformed by multiplying them by unitary matrices, and in many cases this can be done efficiently. For example, the wavevector can be Fourier transformed using $\mathcal{O}(n^2)$ elementary quantum gates. However, not all transformations can be done efficiently. The laws of quantum measurement also limit the amount of information that can be extracted from a quantum state. A full measurement of the state yields outcome x with probability $|\alpha_x|^2$, destroying the state in the process. Thus, even though describing the quantum state of n qubits requires an amount of information that scales exponentially with n,

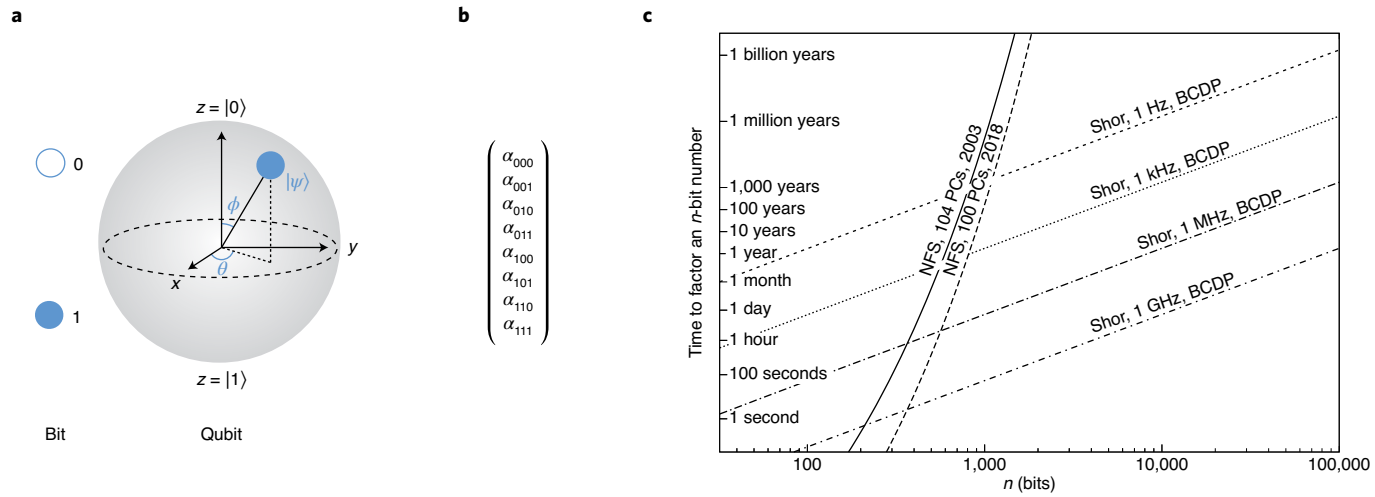

Fig. 1 | Concepts in quantum computing. a, Conceptual illustration of bit versus qubit. The state of a qubit can be represented by a point on the unit sphere with the north and south poles corresponding to the states 0 and 1 of a classical bit. **b**, The state space of 3 qubits is a 2^3-dimensional complex vector. **c**, Classical (number field sieve (NFS) algorithm) and quantum (Beckman–Chari–Devabhaktuni–Preskill (BCDP) implementation of Shor's algorithm) run times for factoring integers. Shor's algorithm for quantum computers yields an exponential speedup over the best known classical algorithm (panel **c** adapted with permission from ref. [61], R. Van Meter et al.).

measurement can only extract n bits of information. Finding a way to benefit from the exponential state space of quantum computers despite this limitation and others is the central challenge of quantum algorithm design[7].

The challenges in building quantum hardware and mitigating noise are considerable and are not addressed in this paper, since our focus is principally on algorithm development and potential biological applications. Large-scale quantum computers are likely to rely on error-correcting codes and other error mitigation strategies that will result in additional overhead; for example, needing to use many physical qubits to store one logical qubit. However, quantum algorithms can be built out of a universal set of quantum gates in a way that does not depend on the underlying hardware, just like classical algorithms.

Given the ubiquity of classical computers, the natural way to understand the strengths of quantum computers is by comparing their run-time scaling with that of the best-known classical algorithms. In some cases, these speedups are exponential: a quantum computer with a few thousand error-corrected qubits could factor numbers that could not be factored using existing classical computers and currently known algorithms in time less than the age of the universe. In other cases, provable polynomial speedups are known: for example, given the ability to compute a function $f(x)$ where x takes on N values, a quantum computer can find the minimum value of $f(x)$ in only $\mathcal{O}(\sqrt{N})$ evaluations of $f(x)$ while a classical computer would require $\mathcal{O}(N)$ steps (assuming that $f(x)$ has no other structure we can exploit)[8]. By contrast, for some problems, quantum computers are known to be no stronger than classical computers. And in many other cases, plausible heuristic algorithms have been proposed for quantum computers, whose performance is only incompletely understood.

The source of quantum speedup. There is not a simple description of what accounts for speedups, although the most plausible explanation is the difference between interference of amplitudes and addition of probabilities. For example, a qubit can have states $|0\rangle$ and $|1\rangle$, which correspond to cbit values 0 and 1 and, in the representation of Fig. 1a, are the north and south poles. Qubits can also be in superpositions (see Box 1) such as $\frac{|0\rangle+|1\rangle}{\sqrt{2}}$ and $\frac{|0\rangle-|1\rangle}{\sqrt{2}}$, which lie on the equator in the figure; these correspond to having amplitude $\frac{1}{\sqrt{2}}$ in the $|0\rangle$ state and amplitude $\frac{\pm 1}{\sqrt{2}}$ in the $|1\rangle$ state. To see that these differ from each other, and also from a random mixture of $|0\rangle$ and $|1\rangle$, consider the $\sqrt{\text{NOT}}$ gate, which maps $|0\rangle$ and $|1\rangle$ to $\frac{|0\rangle+|1\rangle}{\sqrt{2}}$ and $\frac{|0\rangle-|1\rangle}{\sqrt{2}}$, respectively. Starting with the $|0\rangle$ state, applying $\sqrt{\text{NOT}}$ once yields $\frac{|0\rangle+|1\rangle}{\sqrt{2}}$. This state could be thought of as analogous to a random mixture of 0 and 1, as we would expect if $\sqrt{\text{NOT}}$ means applying NOT with probability one-half. However, applying $\sqrt{\text{NOT}}$ twice yields $|1\rangle$, just as we would expect from a NOT gate, whereas applying the randomized version twice would yield the same uniform mixture of 0 and 1. More generally, quantum computers and randomized computers can both be thought of as taking different paths through the 2^n possible bit strings, but for randomized computers we sum the non-negative-valued probabilities of these paths to get the final output distribution, while for quantum computers we sum the complex-valued amplitudes of these paths. Adding complex numbers of roughly the same phase (for example, 1 + 1) corresponds to constructive interference while adding ones of opposite phases (for example, 1 + (−1)) corresponds to destructive interference, analogous to the way that light and other waves can exhibit interference.

While we often do not know how to take advantage of the rich possibilities offered by quantum interference, in some cases we can use them to achieve asymptotic speedups. Algorithms like Grover's unstructured search algorithm[9] are simple examples of this. Grover's algorithm takes a subroutine with a small success probability p, which would need to be repeated $\mathcal{O}(1/p)$ times on a classical computer to obtain a successful outcome, and obtains an answer on a quantum computer using only $\mathcal{O}(1/\sqrt{p})$ repetitions. This makes use of the fact that probabilities are obtained by taking the square of quantum amplitudes. The quantum Fourier transform (used in period

Box 1 | Glossary

Biological

- **Single-nucleotide polymorphism (SNP)**. Germline (inherited) variant in a genome where the identity of a single nucleotide is changed relative to a reference genome; the prevalence of a SNP in a population is dependent on the pattern of its inheritance.
- **Genetic recombination**. Exchange of segments between separate genomes or chromosomes, or different regions of the same chromosome, by the creation of single-stranded (in, for example, viruses) or double-stranded (in, for example, humans) breaks and subsequent ligation of the crossed segments.
- **Genome-wide association study (GWAS)**. Identification of variants in a population with statistically significant associations to the occurrence of a studied phenotype.
- **Quantitative trait locus (QTL)**. Variant in a genome or population with statistically significant association to the occurrence of a studied phenotype, including but not limited to endophenotypes (that is, phenotypes at the suborganismal level; for example, cell- or tissue-level gene expression).

Machine learning

- **Hidden Markov model (HMM)**. Stochastic latent-state method to model a linear sequence of observations as a probabilistic sequence of underlying state transitions and state-to-observation emissions.
- **Boltzmann machine**. Generative classical neural network model, based on an energy function containing local (unary) and pairwise terms over an underlying undirected graph. Recently, the model has been extended to replace the classical energy with a quantum Hamiltonian to form a quantum Boltzmann machine (QBM)[39].
- **Variational autoencoder (VAE)**. Generative neural network model, incorporating a latent space that is mapped to observed variables by a learned feedforward classical neural network. Latent space can be a classical (Gaussian) or quantum (QBM)[40] distribution.

Quantum computing

- **Quantum superposition**. A fundamental principle of quantum mechanics whereby the overall state of a system (for example, an electron in an atom or qubit) is in a linear combination of orthogonal basis states (for example, the lowest energy state, next excited state and so forth). For example, if $|0\rangle$ denotes the lowest energy state of a qubit and $|1\rangle$ an excited state of a qubit, the state of the qubit, $|\psi\rangle$, can be a superposition of basis states: $|\psi\rangle = \alpha_0|0\rangle + \alpha_1|1\rangle$.
- **Quantum random-access memory (qRAM)**. In analogy with random access memory (RAM), which uses n bits to address 2^n distinct memory cells, qRAM would use n qubits to address any quantum superposition of 2^n memory cells[18].
- **Quantum annealing (QA)**. A technique for minimizing a function f using a low-temperature quantum system whose energy corresponds to f, along with an auxiliary field that is slowly turned off. The auxiliary field attempts to create superpositions between nearby qubit strings, similarly to equally weighting possible solutions, and facilitates quantum tunneling (that is, transition of a quantum state between nearby low-energy strings even through regions of higher energy) to arrive at a minimum of f relatively efficiently once turned off.
- **Hidden quantum Markov model (HQMM)**. The quantum analog of HMMs, where the sequence of quantum operations is such that information of the state transition and emission probabilities of the qubits can be retained even after partial measurement of the system (that is, measurements do not collapse the entire system)[21].

finding and Shor's factoring algorithm; Fig. 1c) is a more sophisticated example of how complex-weighted transitions can be useful, and in some cases this can give rise to exponential speedups. By contrast, some problems are known to not admit any quantum speedup: for example, finding the parity of N numbers requires time $\mathcal{O}(N)$ on either a quantum or classical computer[10]. It is a major open research problem to determine when quantum speedup does or does not exist, and it is unlikely to ever be fully resolved, just as there is still no single theorem describing which problems can be solved by efficient classical algorithms. We next discuss some examples of potential quantum speedups.

Exponential speedup. The main exponential speedups known are for code breaking (dramatic but unlikely to be relevant here) and quantum simulation of molecules or other large quantum systems. If the properties of a molecule are not well captured by simple classical approximations, then there is a good case to be made for using a quantum computer to make a better-quality approximation computationally tractable. The advantage of a quantum computer here arises from the exponentially growing dimension of quantum states. As a result, some promising cases for quantum advantage involve molecules with large numbers of active electrons, such as organometallic compounds[11].

Polynomial speedup. Typical polynomial speedups can be thought of as direct improvements of some classical algorithms. The best known of these is Grover's square-root search speedup, described above[12]. Other, more sophisticated algorithms also admit provably quadratic improvements. For example, a classical algorithm might search over a tree of possibilities in a manner that can improve over brute-force search by sometimes being able to quickly prune entire subtrees. Such searches can also be quadratically improved quantumly; that is, if the classical search process explores N nodes, then the quantum algorithm requires effort roughly equal to $\sqrt{N}$ times the effort to evaluate one node[13]. The strength of these algorithms is that they apply under very general conditions, such as needing to minimize an easily computable function. They also do not usually need more qubits than are already needed to compute the function.

Heuristic speedups. Many of the most important algorithms for classical computers either lack formal proofs of correctness or are often run outside of the regime in which these proofs of correctness apply. These include Markov chain Monte Carlo (where rigorous upper bounds on mixing time are usually not known) and gradient descent applied to non-convex problems such as deep neural networks. For quantum computers, heuristic algorithms include adiabatic optimization[14] — or, more generally, quantum annealing (QA)[15] — and the quantum approximate optimization algorithm (QAOA)[16]. The level of speedup provided by these algorithms over classical algorithms is in general unknown. It is expected that as quantum computers are

Box 2 | Computational opportunities for the future

Existing quantum algorithms — for example, function minimization — are often written in terms of abstract and highly general functions. If biological applications can help motivate specific, mathematically well-posed tasks, then it may be the case that targeted quantum algorithm development can lead to improvement. While this promise is discussed at length in the main text in the context of the study of the human brain, here we briefly introduce some of the key areas of ongoing research in quantum computing, related to and providing the context for applications in biology.

Optimization in biomolecular problems
There has been considerable interest in extending QC to biomolecular and biological problems[63]. In several cases, small examples of biological problems have been mapped to combinatorial optimization problems. A QA approach was employed in the exploration of the coarse-grained folding landscape of a six-amino-acid peptide, within a 2D lattice framework[64]. QA was also evaluated against a set of classical methods on an optimization problem involving the search for the consensus DNA sequence motif of transcription factor binding[35]. In this instance, Li et al. trained a classifier (sequence binds or does not bind) and a ranking algorithm (ranking sequences by binding affinity), finding a slight improvement of QA over classical approaches in the classification problem and similar performance for the ranking task.

Simulation of classical and quantum systems
There have been successful demonstrations of the application of quantum computation to problems in chemistry. A variational quantum eigensolver (VQE) approach was used[65] to estimate the ground state energies of small molecules as a function of their component atomic separations. Briefly, short quantum circuits define a variational ansatz of trial solutions for the ground state, and the circuit parameters are varied to minimize the energy using algorithms such as gradient descent. While the complexity of simulating quantum dynamics on quantum computers is well understood and is usually tractable, the success of VQE will depend on the quality of the ansatz and is an active area of ongoing research.

Quantum simulation of chemical reactions is known in principle to be possible on a quantum computer, and, as the practical details are fleshed out, this is expected to be an important application of quantum computers for applications both inside and outside of biology. One particular strength is in modeling dynamics, and there is evidence that energy transport and electron transport in biological molecules involves quantum effects that could potentially be more accurately modeled by a quantum simulation[66].

built, our understanding of the performance of these heuristics will improve, just as much of our understanding of the performance of classical heuristics comes from empirical evidence and not only theory. In the following sections, we refer to this class of methods as "quantum heuristics."

Interfacing with classical algorithms. There is an important caveat about quantum algorithms. Suppose for concreteness that we are minimizing a function $f(x)$. For a speedup, a quantum computer would need to interfere computational paths that compute $f(x)$ for different values of x. If information about the value of x leaked to an outside classical system, then this would prevent those paths from coherently interfering, and we would be left with $f(x)$ for a random choice of x. This would limit its ability to share the computation with a classical computer. Suppose, for example, that the evaluation of $f(x)$ were a memory- and time-intensive calculation for which quantum speedups were not known. Then using quantum computers to improve the minimization of f would need to use qubits to perform this evaluation and could not offload the computation to a classical computer. This means that the overall speedup would be less than quadratic.

Big data and quantum RAM. A related limitation of current models of quantum computers is that they cannot access large classical datasets in superposition; attempting to do so would amount to measuring the qubit register containing the address being queried, which would collapse any superposition there into a random mixture. This means that quantum computers may be able to speed up complicated calculations on small datasets (for example, finding the best Bayesian network) but have less advantage in solving problems on large datasets. One way to address this is with filtering or data-reduction techniques, which select a small but hopefully representative sample of the data and use that as input to the optimization problem[17]. Or the quantum computer could be used for 'small data' problems where the difficulty comes from the complexity of the analysis. A more speculative possibility is a quantum hardware solution known as a qRAM (quantum RAM)[18], which would give a quantum computer the ability to coherently query a large classical dataset as a superposition of qubits: a superposition of input memory addresses would yield an output consisting of a superposition of memory cell contents (see Box 2). A qRAM would enable powerful quantum algorithmic primitives[18], but there are no proposals for scalable error-corrected qRAM, and it is not clear whether it would ultimately be easier than making a large quantum computer[19].

Potential applications for quantum computing in biology

Sequence analysis. We first consider QC algorithms implementable on near-term quantum processors. An essential initial step in genetics and genomics is the matching of sequences of nucleotides and amino acids to organism databases and, more specifically, the mapping of sequencing reads from experimental assays to reference genomes. Any approach needs to contend with both memory (holding a representation of the reference and information on the mapping) and speed concerns. Dynamic programming methods, such as the Smith–Waterman algorithm[20], enable queries of sequence strings against immense databases and could be cast as hidden Markov models (HMMs). The recent development of hidden quantum Markov models (HQMMs)[21,22] opens the possibility of simulating classical HMMs on available quantum circuits[22], as well as extending model space beyond classical HMMs[21]. In fact, the potential advantage of HQMMs stems from this extension of the model space to yield more efficient representations of sequence generators[21]. However, it is unclear how and to what extent this increased efficiency

would translate into speedups. Hybrid approaches are attractive prospects: the iteration through hyperparameter space in HMMs could be classical, with quantum optimization of the maximal trajectory through state space. Given that dynamic programming methods have mostly been supplanted by the approximate but faster *k*-mer-based BLAST algorithm[20] for database searches, a QC-based improvement in efficiency could reopen the case for their utility.

We next explore problems whose QC solutions may depend on the availability and storage in memory of superpositions of qubits (qRAM). For genomic read mapping, state-of-the-art classical algorithms include the exploitation of the Burrows–Wheeler transform to efficiently perform DNA sequence alignments[23], and seed-based approaches to map RNA reads to exon boundaries separated by large genomic distances[24]. Both methods rely on lexicographically sorted suffixes constructed from the reference genome, followed by scanning for matches of the query read. The classical complexity of sequence matching depends on whether exact ($\mathcal{O}(n+m)$; n = length of reference sequence, m = query read length) or inexact matches ($\mathcal{O}(nm)$), including gaps, are considered. Grover's-algorithm-based improvements in string-matching speeds[25] could be exploited ($\tilde{\mathcal{O}}(\sqrt{n}+\sqrt{m})$ for exact matches) to aid the scanning process. Recent work has demonstrated the potential for even further QC speed gains under the assumption of unique membership of a query string within a reference database[26]. The scaling of the problem is such that a reduction in complexity of even simpler mapping problems would be highly beneficial, although the need to generate superpositions of the entire reference string also creates potential problems: given the need to store a large reference database in superposition, the current lack of qRAM is an issue. Furthermore, speed gains from Grover's-algorithm-based methods could be reduced by the cost of evaluating the function being searched, if done classically.

Genetics. As in the previous section, a possible problem for near-term QC algorithms to tackle is the imputation of individual-specific mutations, especially single-nucleotide polymorphisms (SNPs). Given shared sets of haplotypes across subpopulations, a relatively sparse set of SNPs can be expanded by inferring additional SNPs that co-occur with the original set with high probability. This imputation usually involves an HMM-based likelihood maximization[27], which could be cast as HQMMs.

While imputation depends on inherited SNPs within populations (germline mutations), cells also contain postconception de novo variants, called "somatic variants." Every neuron in the human brain is likely to contain private somatic variants, including single-nucleotide variants and large structural variants that alter allelic diversity for dozens of genes. Identifying their functional impact is essential. Machine-learning classifiers have been trained on case–control datasets to identify psychiatric-disorder-associated variants[28]. However, given the high-dimensional parameter search space for the classification problem, classical computation frequently runs into search efficiency issues. These issues might be ameliorated using near-term implementable QC machine learning methods[29], discussed at length in subsequent subsections.

Another important category of genetic analyses is the construction of optimal trees that describe the relative proximity of genetic sequences, including ancestral recombination graphs (ARGs)[30], depicting ancestral relationships between individual genomes while accounting for genetic recombination; pathogen evolutionary trees in epidemiological studies; and tumor cell mutational lineages, as could be relevant to malignancy and medical response. Tree reconstruction algorithms optimize across the similarity constraints between genomic segments, mainly involving sampling from the space of possible genealogies with heuristics and simplifications[31]. For smaller input sequence sets, the massive tree-search space makes this an open candidate problem for speedup using available quantum heuristic optimization methods[14–16].

SNP association and heritability analyses are problematic for near-term quantum approaches, given the need to manipulate large matrices to solve systems of linear equations. In association studies, SNPs can be statistically associated with individual-level phenotypes in genome-wide association studies (GWAS) or to quantitative molecular traits (cell or tissue gene expression, methylation, epigenetic markers, cell fractions) and other quantitative traits (loci designated as QTLs). The evaluation of total SNP heritability often involves linear mixed effects models, with genetic variance estimations carried out through techniques such as the restricted maximum likelihood (REML) method[32]. With qRAM, algorithms such as quantum least squares[33,34] could offer up to exponential speedups through the ability to perform fast linear-algebraic operations, under certain assumptions of sparseness and condition number, although it is unclear to what extent any advantages would be undercut by the time cost of querying the qRAM. For lower-dimensional regression problems, there is some potential for near-term quantum heuristic optimizers to tackle these tasks[35].

Functional genomics. The chain of factors that leads from genetic variation to higher-level behaviors such as cognitive traits includes complex intermediate links, such as the molecular regulatory framework within cells, cell-to-cell interactions, heterogeneity in cellular composition and behavior in tissues, and inter-regional connectivity patterns in the brain, among many others (Fig. 2). These factors are further governed by complex developmental processes and gene–environment interactions in an individual-specific manner. Despite this complexity, recent studies have shown that genetic risk for particular traits can be partitioned across 'intermediate' phenotypes, such as gene expression or chromatin binding profiles; a direct approach to such analysis is to impute intermediate molecular phenotypes first and then link the imputed phenotypes to high-level traits[36]. However, intermediate molecular phenotypes are typically high dimensional and interdependent; for example, bulk transcriptome expression profiles can be ~22,000 dimensional. Possible models that can learn joint probability distributions over such levels of analyses include Bayesian networks, undirected models such as Boltzmann machines[37], and recent deep-learning approaches such as variational autoencoder (VAEs). Exact optimization of such models, however, is intractable: structure learning in Bayesian networks requires optimization over a search space of all directed acyclic graphs, which is super-exponential ($\mathcal{O}\left(n!2^{\frac{n!}{2!(n-2)!}}\right)$, where n is the dimensionality[38]). Inference in Boltzmann machines requires a search over $\mathcal{O}(2^n)$ states after binarization to calculate a gradient, and training VAEs requires the optimization of a non-convex objective function. Such problems may be potential candidates for quantum approaches: for smaller input sizes, near-term approaches without qRAM may be developed to perform exact searches across the space of Bayesian networks, while for moderate-sized problems, approximate quantum analogs of Boltzmann machines (QBMs) and VAEs (QVAEs) have been tested in simulation and experimentally[39,40], with the optimization being conducted through QA. We note also

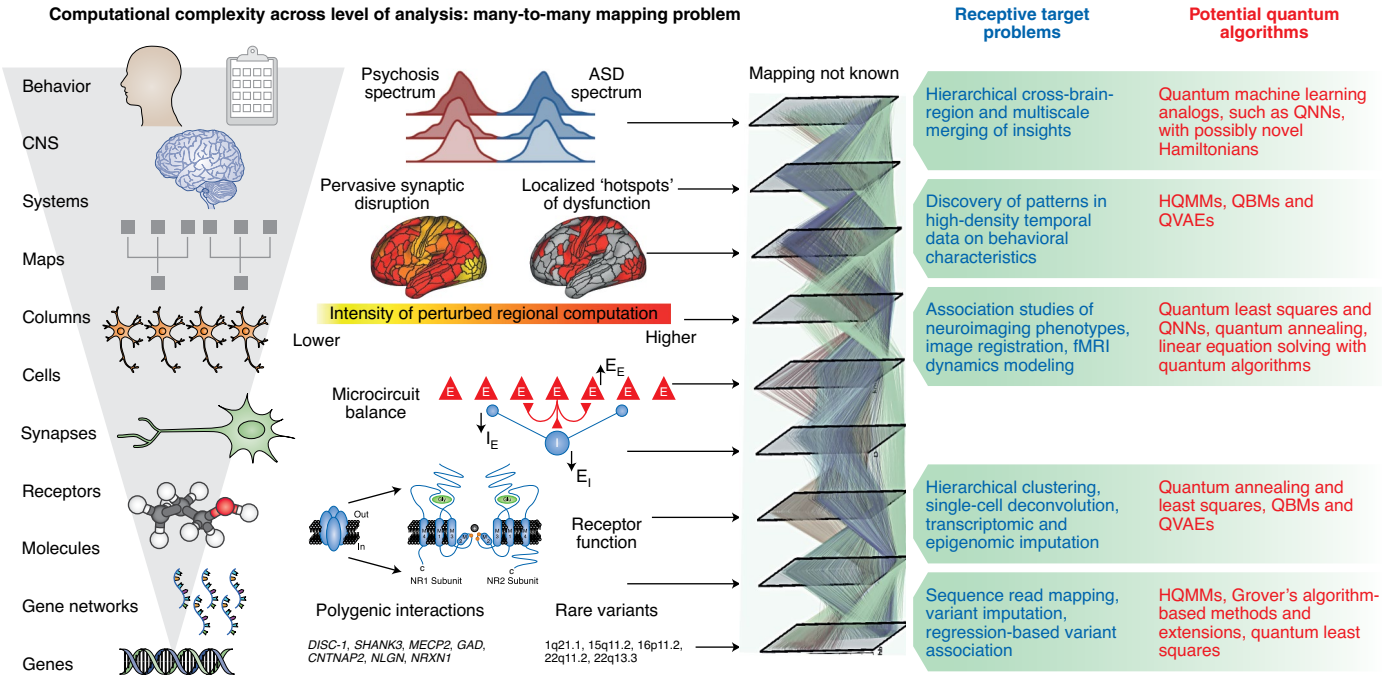

Fig. 2 | Complexity of linking levels of analyses from genetics to human behavior. The challenge consists, in part, of the need to interrogate the enormous search space for determining the mapping across levels, which constitutes a many-to-many probabilistic problem. Computational innovation will be a key effort to help close these gaps. Portion of figure adapted with permission from ref. [62], Elsevier. Also shown are some of the ways in which QC can aid in the interrogation of these levels.

that, for all these models, prior knowledge of molecular interactions may be used during training to suggest causal network interpretations.

In contrast to direct imputation of molecular phenotypes, intermediate phenotypes may be derived at the level of sets of genes (such as functional pathways) and cell-type proportions. For instance, weighted gene correlation network analysis (WGCNA) performs a version of hierarchical clustering to derive coexpression modules, which are enriched in gene pathways[41], and non-negative matrix factorization (NMF) based on marker gene profiles can be used to decompose bulk transcriptome data into components corresponding to cell-type fractions[37]. Exact optimization of these models is again intractable; exact hierarchical clustering would require a search over a large space of trees, and NMF is a non-convex optimization problem. The former may be a candidate for an exact quantum solution for small-scale problems while both may benefit from quantum heuristic approaches (a QA approach to NMF is found in ref. [42], and quantum speedups for approximate clustering are described in ref. [17]). While clustering ~1,000 to ~20,000 features is common in genomics, there are a number of applications where a relatively small number of features, ~100, are clustered across samples (for example, protein-array data). Clustering associated with global minimization of objective functions is of great interest in these small-feature-number cases. More generally, comparison of clusters (and solutions to other genomic algorithms) derived from exact and approximate greedy minimization would inform the nature of the errors associated with applying greedy algorithms to large numbers of features and samples, as well as suggest possible approaches to improving the greedy algorithms in the short term. Application of these methods at full genomic scale, however, would require further technical developments in qRAM or quantum processor size.

Mapping neurobehavioral variation via neuroimaging and deep phenotyping. The overarching goal of convergent neuroscience is to link cellular-level mechanisms to system-level observations and ultimately behavior. Multimodal neuroimaging provides rich high-dimensional data that can map neural and behavioral mechanisms in humans. While many quantitative optimizations remain, one of the core challenges is accurate identification and alignment across people of brain anatomy to reference atlases. For instance, one widespread approach implemented in FreeSurfer software[43] employs a sequence of registration steps involving the minimization of an energy function over the spatial transformation field. Here, potential quantum heuristic approaches could be brought to bear for images of moderate resolution if the corresponding energy function (Hamiltonian) can be mapped to an Ising-type model. A related challenge involves training statistical models to rapidly and accurately quantify neurobehavioral variation. For instance, the presence of active psychotic symptoms in previously unseen individuals diagnosed with schizophrenia and bipolar illness can be predicted using dynamic functional connectome features derived from fMRI[44]. Quantum analogs (such as HQMMs[21,22]; see "Sequence analysis" and "Genetics") may help train such predictive models more efficiently.

Computational neuroscience has used circuit models to inform and constrain experimental observations. Dynamical neural models operate at the local circuit or global level and use parameterizations based on known constraints (for example, biophysical parameters) or learned de novo. Local and global neural dynamics are typically highly nonlinear, producing difficult optimization problems in the case

of parametric model fitting[45] and requiring a rich model class for de novo learning methods. Fluctuations at equilibrium exhibit complex interdependencies. Furthermore, the hierarchical relationships between genetics, anatomy, function and the equilibrium connectivity neural state are, in general, highly nonlinear and only partially captured by available computational models. Current classical models relate such simulations to equilibrium distribution features (or to resting state characteristics) — for instance, Ising models and second-order mean-field regional models of resting-state fMRI observations[46,47]. These differential-equation-based analyses of global brain dynamics represent regional firing rates using a mean-field approximation[46]. Such models can be fitted to functional neuroimaging data by linearizing the initial stochastic nonlinear system of differential equations around a fixed point using the method of moments[46] and using methods such as approximate Bayesian computation to fit parameters[45]. In the QC domain, quantum algorithms have been developed that have the potential to offer exponential speedups in the solving of linear differential equations[48,49]. Furthermore, models such as the QBM[39] and QVAE[40], as discussed in the previous subsection, may be naturally applied to model complex distributions such as those found in neurodynamics datasets.

General-purpose quantum solvers for nonlinear systems of differential equations have also been proposed[50], although so far these seem unlikely to offer speedups over classical methods. Efficient general-purpose solvers would eliminate the need for linear approximations and allow more accurate fitting of neural dynamical models, particularly out of steady state (for example, transitions between resting-state and task-based fMRI). This application may help motivate finding fast quantum algorithms for nonlinear differential equations.

The computational challenge in human neuroscience is particularly acute in the case of 'deep' behavioral phenotyping (for example, digital real-time measures), which can generate massive amounts of continuously measured dynamical behavioral variables with varied granularity. In this situation, there is clear potential for 'very deep' optimization and the opportunity for massive state-space exploration. Relevant use-case scenarios include in-the-moment clinical decisions that may require rapid computation. This becomes challenging for longitudinal real-time digital phenotyping, which may require rapid and precise data reduction. For instance, rich individualized phenotypic characterization using high-resolution video and audio datasets have yet to be leveraged since they are identifiable in raw form and present operational challenges to data reduction and protection of participant privacy.

Collectively, the complexity of human neurobehavioral data tests the boundaries of learning algorithms, which have to deal with the high dimensionality of data needed to robustly link nonlinear dynamics of brain states (for example, fMRI) and the influence of time-related variables relevant to behavioral mapping. Recent deep learning approaches using interpretable recurrent networks have provided a powerful means of learning such brain-state/behavior associations de novo by jointly modeling fMRI and behavioral data[51]. Quantum analogs of neural network frameworks (such as QNNs[52], QBMs[39] and QVAEs[40]) have the potential to discover novel structure in these datasets. Models such as HQMMs provide alternative dynamical models with intrinsically quantum representations[22], which have been shown to have comparable or possibly improved performance relative to classical methods on small-scale problems through classical simulations. Further, there is evidence that HQMMs allow complex dynamics to be modeled in a reduced state space[21] compared to classical models. The application of such methods to behavioral data, though, is a long-term goal, since reliable qRAM appears necessary to handle large dataset sizes.

Integration across disciplines. Stitching together insights across fields and levels of analyses to yield a complete picture of brain function is an ongoing challenge. Quantum machine learning may help elucidate the interdependencies between levels through its ability to learn and simulate nonlinear, potentially classically intractable models. One promising avenue involves mechanism-agnostic machine learning methods like deep neural networks, where biological insights are gained by interpreting the model a posteriori. Such an interpretable framework would involve connections between modules such as gene regulatory networks, on the one hand, and structural or functional neuroimaging parameters (for example, cortical thickness, white matter integrity and dynamic functional connectivity) on the other. The exact nature of these connections could be altered in competing hypotheses. One could imagine a hierarchical network with molecular phenotypes at the base, emergent neuroimaging-based parameters at a higher layer, and behavioral phenotypes as prediction targets. An alternative framework would treat the molecular and neural-systems-level components as parallel factors in determining behavior, with the latter having been influenced at a developmental stage and not directly emerging from the molecular phenotypes per se but rather operating in dependent lockstep. Thus, different architectures of relationships between levels of analysis may be constructed. The National Institute of Mental Health in the United States has recently supported efforts at building such multiscale, convergent neuroscience approaches (https://grants.nih.gov/grants/guide/pa-files/par-17-176.html). Such an analysis could be aided by QNNs[52] and quantum variational classifiers[53] designed for use on non-qRAM, gate-based quantum computers. Quantum variational classifiers are able to successfully classify states designed to be hard to simulate classically[53]. This hints at the greater generality of such circuits than their classical counterparts. Here the challenge lies in scaling up the available number of qubits.

Epilogue

While the field of QC is undergoing notable development and progress in both hardware and software, knowledge gaps and challenges remain. To surpass classical computers, quantum computer architectures will need to improve numbers of and connectivity between qubits, reduce error rates both for operations and storage, and expand algorithmic development into all areas where classical computing faces inherent bottlenecks. These challenges are all significant and are partially conflicting; indeed, the central experimental QC challenge is to create quantum systems that are both highly decoupled from unwanted environmental degrees of freedom yet subject to fast and precise control and measurement. While there has been steady experimental progress over the past two decades, it is not easy to predict the rate of future improvements in QC. A recent consensus study on the progress and prospects of QC from the National Academies of Sciences, Engineering and Medicine estimates that to effectively break current internet security protocols (that is, find a private key in a 1,024-bit RSA encrypted message) using Shor's algorithm requires building a quantum computer that is five orders of magnitude larger and has error rates that are two orders of magnitude lower than existing machines[54]. More than 100 academic and government laboratories around the world are working to address these challenges with a variety of hardware solutions[54]. These include ion-trap quantum computers with 20–100 qubits that are likely to become available by the early 2020s[54].

Leveraging the power of lithographic technology, superconducting quantum computers hold great promise, and 5-, 16- and 20-qubit machines are available to users via the web. Other promising approaches include developing quantum computers based on photonic, neutral-atom and semiconductor qubits[54].

As mentioned above, many algorithmic quantum speedups depend on qRAM, but there is no practical implementation of this technology. In fact, this reliance on qRAM, in part, stems from attempts to arrive at algorithms that are essentially quantum versions of classical algorithms. An alternative approach is to design intrinsically quantum algorithms that take advantage of quantum features such as interference. We think that this alternative approach offers the added benefit that small-scale versions of problems are readily implementable on existing hardware. Indeed, recent advances in quantum machine learning algorithms exploit the exponentially large quantum state space to estimate kernel functions[53,55], as well as the natural ability of quantum computers to execute kernel-based classification[56,57]. We believe that generalizations of these algorithms for genomics applications hold great promise and will allow assessment of the current capabilities of publicly available quantum computers[29]. Given the potential of quantum computers to efficiently explore a vast state space, we think that the natural applications to neuroscience problems are largely associated with optimization and machine learning, as detailed above. We feel that another potentially fruitful path is to identify computational problems that can be naturally cast into a quantum framework. For example, the minimum free energy among all possible protein folds is an important problem with an exponentially large search space and thus a compelling target. Another natural set of problems are those associated with quantum biology — the study of chemical processes including formation of excited electron states within molecules (for example, proteins) in living cells, along with their functional effects[58]. These processes are inherently quantum mechanical and may involve an exponentially vast set of excitation states, which can only be efficiently modeled by applying transformations to an exponentially large state space afforded by a quantum computer. However, we are not sure whether such processes can be relevant to higher-levels of brain function (and consciousness[59]); the algorithms used by the brain at David Marr's algorithmic or representational level may ultimately be classical[60], although the advent of quantum machine learning means that increasingly this need not be the case for artificial agents. While a cautious albeit optimistic estimation associated with steady progress of quantum hardware development (for example, applying Moore's law) puts the availability of sufficiently powerful, universal quantum computers years in the future, sudden, orders-of-magnitude breakthroughs in resolution, noise reduction and so forth are not unprecedented in experimental physics. We strongly believe that such unforeseen breakthroughs would unleash the power of quantum computing to address pressing computational challenges in biology. ❐

Prashant S. Emani[1,2,27], Jonathan Warrell[1,2,27], Alan Anticevic[3], Stefan Bekiranov[4], Michael Gandal[5], Michael J. McConnell[6], Guillermo Sapiro[7], Alán Aspuru-Guzik[8,9,10,11], Justin T. Baker[12,13], Matteo Bastiani[14,15], John D. Murray[16,17], Stamatios N. Sotiropoulos[14,18], Jacob Taylor[19,20], Geetha Senthil[21], Thomas Lehner[21,22,26] ✉, Mark B. Gerstein[1,2,23,24] ✉ and Aram W. Harrow[25] ✉

[1]Program in Computational Biology and Bioinformatics, Yale University, New Haven, CT, USA. [2]Department of Molecular Biophysics and Biochemistry, Yale University, New Haven, CT, USA. [3]Yale School of Medicine, Department of Psychiatry, New Haven, CT, USA. [4]Department of Biochemistry and Molecular Genetics, University of Virginia School of Medicine, Charlottesville, VA, USA. [5]Department of Psychiatry, Semel Institute, David Geffen School of Medicine, University of California–Los Angeles, Los Angeles, CA, USA. [6]Lieber Institute for Brain Development, Baltimore, MD, USA. [7]Department of Electrical and Computer Engineering, Duke University, Durham, NC, USA. [8]Department of Chemistry, University of Toronto, Toronto, Ontario, Canada. [9]Canadian Institute for Advanced Research (CIFAR) Artificial Intelligence Research Chair, Toronto, Ontario, Canada. [10]Vector Institute, Toronto, Ontario, Canada. [11]Department of Computer Science, University of Toronto, Toronto, Ontario, Canada. [12]Institute for Technology in Psychiatry, McLean Hospital, Belmont, MA, USA. [13]Department of Psychiatry, Harvard Medical School, Boston, MA, USA. [14]Sir Peter Mansfield Imaging Centre, School of Medicine, University of Nottingham, Nottingham, UK. [15]Wellcome Centre for Integrative Neuroimaging, FMRIB Centre, Nuffield Department of Clinical Neurosciences, University of Oxford, John Radcliffe Hospital, Oxford, UK. [16]Department of Psychiatry, Yale University School of Medicine, New Haven, CT, USA. [17]Department of Physics, Yale University, New Haven, CT, USA. [18]NIHR Nottingham Biomedical Research Centre, Queen's Medical Centre, Nottingham University Hospitals NHS Trust, Nottingham, UK. [19]Joint Center for Quantum Information and Computer Science, University of Maryland, College Park, MD, USA. [20]National Institute of Standards and Technology, Gaithersburg, MD, USA. [21]New York Genome Center, New York, NY, USA. [22]National Institute of Mental Health, Bethesda, MD, USA. [23]Department of Computer Science, Yale University, New Haven, CT, USA. [24]Department of Statistics and Data Science, Yale University, New Haven, CT, USA. [25]Center for Theoretical Physics, Department of Physics, Massachusetts Institute of Technology, Cambridge, MA, USA. [26]Present address: Neuropsychiatric Disease Genomics, New York Genome Center, New York, NY, USA. [27]These authors contributed equally: Prashant S. Emani, Jonathan Warrell.
✉*e-mail: tlehner@gmail.com; mark@gersteinlab.org; aram@mit.edu*

Published online: 4 January 2021
https://doi.org/10.1038/s41592-020-01004-3

References

1. Mohseni, M. et al. *Nature* **543**, 171–174 (2017).
2. Arute, F. et al. *Nature* **574**, 505–510 (2019).
3. U.S. House of Representatives. 115th Congress. H.R. 6227, U.S. National Quantum Initiative Act (*Government Printing Office*, 2018).
4. Monroe, C., Raymer, M. G. & Taylor, J. *Science* **364**, 440–442 (2019).
5. Thew, R., Jennewein, T. & Sasaki, M. *Quantum Sci. Technol.* **5**, 010201 (2019).
6. Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information* (Cambridge Univ. Press, 2010).
7. Alexeev, Y. et al. Quantum computer systems for scientific discovery. Preprint at https://arxiv.org/abs/1912.07577 (2019).
8. Durr, C. & Hoyer, P. A quantum algorithm for finding the minimum. Preprint at https://arxiv.org/abs/quant-ph/9607014 (1996).
9. Brassard, G., Høyer, P., Mosca, M. & Tapp, A. in *Contemporary Mathematics: Quantum Computation and Information* (eds. Lomonaco, S. J. & Brandt, H. E.) 53–74 (American Mathematical Society, 2002).
10. Farhi, E., Goldstone, J., Gutmann, S. & Sipser, M. *Phys. Rev. Lett.* **81**, 5442–5444 (1998).
11. Li, Z., Li, J., Dattani, N. S., Umrigar, C. J. & Chan, G. K. L. J. *Chem. Phys.* **150**, 024302 (2019).
12. Grover, L. K. in *Proc. Twenty-Eighth Annual ACM Symposium on Theory of Computing* 212–219 (Association for Computing Machinery, 1996).
13. Ambainis, A. & Kokainis, M. in *Proc. 49th Annual ACM SIGACT Symposium on Theory of Computing – STOC 2017* 989–1002 (2017).
14. Farhi, E., Goldstone, J., Gutmann, S. & Sipser, M. Quantum computation by adiabatic evolution. Preprint at https://arxiv.org/abs/quant-ph/0001106 (2000).
15. Kadowaki, T. & Nishimori, H. *Phys. Rev. E Stat. Phys. Plasmas Fluids Relat. Interdiscip. Topics* **58**, 5355–5363 (1998).
16. Farhi, E., Goldstone, J. & Gutmann, S. A quantum approximate optimization algorithm. Preprint at https://arxiv.org/abs/1411.4028 (2014).
17. Harrow, A. W. Small quantum computers and large classical data sets. Preprint at https://arxiv.org/abs/2004.00026 (2020).
18. Giovannetti, V., Lloyd, S. & Maccone, L. *Phys. Rev. Lett.* **100**, 160501 (2008).
19. Arunachalam, S. & de Wolf, R. *ACM SIGACT News* **48**, 41–67 (2017).
20. Li, H. & Homer, N. *Brief. Bioinform.* **11**, 473–483 (2010).
21. Monras, A., Beige, A. & Wiesner, K. Hidden quantum Markov models and non-adaptive read-out of many-body states. Preprint at https://arxiv.org/abs/1002.2337 (2010).
22. Srinivasan, S., Downey, C. & Boots, B. in *Advances in Neural Information Processing Systems 31* (eds. Bengio, S. et al.) 10338–10347 (Curran Associates, 2018).
23. Li, H. & Durbin, R. *Bioinformatics* **26**, 589–595 (2010).
24. Dobin, A. et al. *Bioinformatics* **29**, 15–21 (2013).
25. Ramesh, H. & Vinay, V. *J. Discrete Algorithms* **1**, 103–110 (2003).
26. Montanaro, A. *Algorithmica* **77**, 16–39 (2017).
27. Howie, B. N., Donnelly, P. & Marchini, J. *PLoS Genet.* **5**, e1000529 (2009).
28. McConnell, M. J. et al. *Science* **356**, eaal1641 (2017).

29. Kathuria, K., Ratan, A., McConnell, M. & Bekiranov, S. *Quantum Mach. Intell.* **2**, 7 (2020).
30. Griffiths, R.C. & Marjoram, P. in *Progress in Population Genetics and Human Evolution* 257–270 (Springer, 1997).
31. Li, H. & Durbin, R. *Nature* **475**, 493–496 (2011).
32. Yang, J., Lee, S. H., Goddard, M. E. & Visscher, P. M. *Am. J. Hum. Genet.* **88**, 76–82 (2011).
33. Harrow, A. W., Hassidim, A. & Lloyd, S. *Phys. Rev. Lett.* **103**, 150502 (2009).
34. Wiebe, N., Braun, D. & Lloyd, S. *Phys. Rev. Lett.* **109**, 050505 (2012).
35. Li, R. Y., Di Felice, R., Rohs, R. & Lidar, D. A. *NPJ Quantum Inf.* **4**, 14 (2018).
36. Gamazon, E. R. et al. *Nat. Genet.* **47**, 1091–1098 (2015).
37. Wang, D. et al. *Science* **362**, eaat8464 (2018).
38. Robinson, R. W. in *Combinatorial Mathematics V. Lecture Notes in Mathematics, Vol. 622* (ed. Little, C. H. C.) 28–43 (Springer, 1977).
39. Amin, M. H., Andriyash, E., Rolfe, J., Kulchytskyy, B. & Melko, R. *Phys. Rev. X* **8**, 021050 (2018).
40. Khoshaman, A., Vinci, W., Denis, B., Andriyash, E. & Amin, M. H. *Quantum Sci. Technol.* **4**, 014001 (2019).
41. Zhang, B. & Horvath, S. *Stat. Appl. Genet. Mol. Biol.* **4**, 17 (2005).
42. O'Malley, D., Vesselinov, V. V., Alexandrov, B. S. & Alexandrov, L. B. *PLoS One* **13**, e0206653 (2018).
43. Fischl, B. *Neuroimage* **62**, 774–781 (2012).
44. Reinen, J. M. et al. *Nat. Commun.* **9**, 1157 (2018).
45. Demirtaş, M. et al. *Neuron* **101**, 1181–1194.e13 (2019).
46. Deco, G. et al. *J. Neurosci.* **34**, 7886–7898 (2014).
47. Deco, G., Senden, M. & Jirsa, V. *Front. Comput. Neurosci.* https://doi.org/10.3389/fncom.2012.00068 (2012).
48. Childs, A. M. & Liu, J.-P. *Commun. Math. Phys.* **375**, 1427–1457 (2020).
49. Berry, D. W., Childs, A. M., Ostrander, A. & Wang, G. *Commun. Math. Phys.* **356**, 1057–1081 (2017).
50. Leyton, S. K. & Osborne, T. J. A quantum algorithm to solve nonlinear differential equations. Preprint at https://arxiv.org/abs/0812.4423v1 (2008).
51. Dezfouli, A., Morris, R., Ramos, F., Dayan, P. & Balleine, B. W. in *Adv. Neural Information Processing Systems 31* (eds. Bengio, S. et al.) 4228–4237 (Curran Associates, 2018).
52. Farhi, E. & Neven, H. Classification with quantum neural networks on near term processors. Preprint at https://arxiv.org/abs/1802.06002 (2018).
53. Havlíček, V. et al. *Nature* **567**, 209–212 (2019).
54. National Academies of Sciences, Engineering and Medicine. *Quantum Computing: Progress and Prospects*. https://doi.org/10.17226/25196 (2019).
55. Schuld, M. & Killoran, N. *Phys. Rev. Lett.* **122**, 040504 (2019).
56. Schuld, M., Fingerhuth, M. & Petruccione, F. *Europhys. Lett.* **119**, 60002 (2017).
57. Schuld, M., Bocharov, A., Svore, K. M. & Wiebe, N. *Phys. Rev. A (Coll. Park)* **101**, 032308 (2020).
58. Lambert, N. et al. *Nat. Phys.* **9**, 10–18 (2013).
59. Tegmark, M. *Phys. Rev. E Stat. Phys. Plasmas Fluids Relat. Interdiscip. Topics* **61**, 4194–4206 (2000).
60. Marr, D. C. & Poggio, T. *Neurosci. Res. Program Bull.* **15**, 470–488 (1977).
61. Van Meter, R., Itoh, K. M. & Ladd, T. D. in *MS+S 2006 – Controllable Quantum States: Mesoscopic Superconductivity and Spintronics, Proceedings of the International Symposium* 183–188 (World Scientific Publishing, 2008).
62. Foss-Feig, J. H. et al. *Biol. Psychiatry* **81**, 848–861 (2017).
63. Outeiral, C. et al. *WIREs Comput. Mol. Sci.* e1481 (2020).
64. Perdomo-Ortiz, A., Dickson, N., Drew-Brook, M., Rose, G. & Aspuru-Guzik, A. *Sci. Rep.* **2**, 571 (2012).
65. Kandala, A. et al. *Nature* **549**, 242–246 (2017).
66. Dorner, R., Goold, J. & Vedral, V. *Interface Focus* **2**, 522–528 (2012).

Acknowledgements

This work is a product of discussions initiated during a NIMH-convened virtual workshop, addressing computational challenges in genomics and neuroscience via massively parallel computing and QC (https://www.nimh.nih.gov/news/events/2018/virtual-workshop-solving-computational-challenges-in-genomics-and-neuroscience-via-parallel-and-quantum-computing.shtml). We would also like to acknowledge the help and support of Lora Bingaman of the NIMH in overseeing the administration of this collaboration. M.B.G. acknowledges the support of NIH grant MH116492-03S1.

Author contributions

All authors contributed to discussions on the design of the manuscript. G. Senthil and T.L. led the NIMH workshop and subsequent discussions. P.S.E., J.W., A.A., S.B., M.G., M.J.M., J.T.B., M.B.G. and A.W.H wrote the manuscript. P.S.E., J.W., A.A., S.B., M.G., M.J.M., G. Sapiro, A.A.-G., J.D.M., J.T.B., M.B., S.N.S., G. Senthil, T.L., M.B.G. and A.W.H edited the manuscript. A.W.H. and S.B. contributed to the "Classical versus quantum circuits: state of the art" section. P.S.E., J.W., S.B., M.G., M.J.M., M.B.G. contributed content to the "Sequence analysis," "Genetics," "Functional genomics" and "Integration across disciplines" subsections, and A.A., G. Sapiro, J.T.B., J.W. and J.M. to the "Mapping neurobehavioral variation via neuroimaging and deep phenotyping" subsection. S.B. contributed to the Epilogue section.

Competing interests

A.A. serves as a member of the Scientific Advisory Board of, consults for, has received grants from, and holds equity in BlackThorn Therapeutics. G. Sapiro consulted for Apple, Volvo, Restore3D and SIS, and has received speaking fees from Johnson & Johnson. J.T.B. has received consulting fees from BlackThorn Therapeutics, Niraxx Therapeutics, Verily Life Sciences, AbleTo Inc. and Pear Therapeutics, and has received consulting fees and equity from Mindstrong Inc. J.D.M. consults for, has received grants from, and holds equity in BlackThorn Therapeutics. A.W.H. has recently joined the Scientific Advisory Board of Zapata Computing, from which he expects income and equity.