Photodynamic Therapy as a Driver for Useful Quantum Advantage in Chemistry: Perspective from the Q4Bio program

Q4Bio Team Algorithmiq

Abstract

Can quantum computing meaningfully impact biology and health in the near term? This question lies at the core of Wellcome Leap's Q4Bio program and motivates the need to define what constitutes *useful quantum advantage* for real-world applications. Here, we introduce a framework based on three simultaneous requirements: executability on current hardware, relevance to biologically meaningful systems, and validation against state-of-the-art classical methods under realistic resource assumptions. We instantiate this framework in photodynamic therapy (PDT), a clinically relevant application that exposes a clear computational bottleneck in excited-state electronic structure. Algorithmiq's approach combines scalable circuit construction driven by Majorana Propagation, efficient measurement protocols, and a hybrid post-processing method, quantum boosted density matrix renormalization group (QB-DMRG), which integrates quantum-generated data into tensor-network optimization. Within this pipeline, we demonstrate that quantum hardware can improve a leading classical electronic-structure method for a physically relevant system. More broadly, our results establish a setting in which quantum devices can, with performance guarantees, provide correlated information that enhance classical algorithms, forming a concrete pathway toward practical quantum advantage in chemistry.

1 Executive Summary

The mission of Wellcome Leap Q4Bio program was to identify health-relevant applications where emerging quantum hardware can begin to shift practical limits, and to co-develop the algorithms and workflows required to realize that impact. Despite substantial progress, it has to date remained unclear what would constitute such an impact in practice. In this perspective, we introduce a framework for assessing useful quantum advantage based on the simultaneous satisfaction of three conditions: executability on near term and fault tolerant hardware, relevance to biologically meaningful systems, and validation against state-of-the-art classical methods under fair resource assumptions. Satisfying any one of these in isolation is not sufficient, and progress along individual dimensions can give the appearance of advancement without addressing the underlying computational bottleneck.

The first condition is executability on current hardware. Within Q4Bio, this required methods that remain compatible with noisy, resource-constrained quantum devices rather than assuming fault-tolerant hardware, while simultaneously guaranteeing that the algorithms and methods scale as quantum devices become larger and more performant. To address this, Algorithmiq developed scalable approaches to circuit construction together with measurement strategies that make it possible to extract useful information from realistic hardware under finite depth and finite sampling budgets.

The second condition is relevance to biologically meaningful systems. For this, we use the simulation of drug candidates for photodynamic therapy (PDT) as a representative benchmark. PDT is not chosen because it is narrowly specific, but because it provides a particularly clear and demanding instance of a broader health-relevant computational bottleneck. Designing improved photosensitizers requires predictive access to excited-state electronic structure in chemically realistic systems, precisely where classical approaches become unreliable or prohibitively expensive and where the design loop reverts to empirical exploration. This makes PDT a natural setting in which to test whether quantum methods can address a real bottleneck rather than a simplified proxy.

The third condition is validation of any quantum advantage claims against state-of-the-art classical methods. Within Q4Bio, this meant building a framework in which quantum resources are assessed directly against the strongest available classical approaches under fair assumptions. In our case, this led to the development by Algorithmiq of a hybrid workflow centered on boosting DMRG, the leading tensor-network method for strongly correlated electronic structure [1–5]. The goal was therefore not merely to produce quantum outputs competitive with classical methods, but to determine whether quantum hardware can provide samples that can be used to improve the best classical chemistry methods in regimes relevant to the application.

Taken together, these elements define an end-to-end, hardware-in-the-loop framework for assessing useful quantum advantage in practice. Within this framework, we find first evidence that quantum hardware resources can boost state-of-the-art classical chemistry methods, providing information that improves over standard DMRG under matched classical resource constraints at system sizes accessible on today's hardware. More broadly, this work argues that useful quantum advantage is not demonstrated by isolated components, but by bringing executability, relevance, and rigorous validation together in a single coherent workflow.

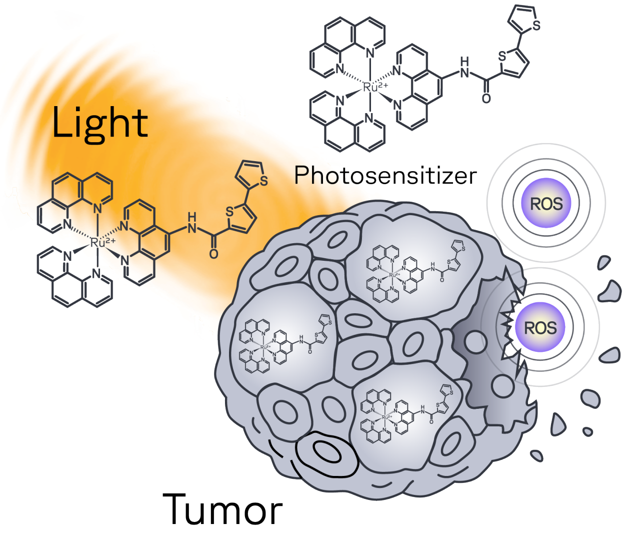

Figure 1: Photodynamic Therapy (PDT) is a targeted, minimally invasive medical treatment that uses a light-activated drug to selectively destroy diseased (cancerous) cells. Upon uptake of a photosensitizer, a drug that remains inactive until exposed to light, the drug naturally accumulates more in diseased or rapidly dividing cells. A specific wavelength of light is then directed to the affected area. When activated, the drug produces reactive oxygen species (ROS) that kill targeted cells while sparing surrounding tissue. This "light-triggered chemistry" naturally gives PDT its precision and low systemic toxicity.

2 The Bio-Problem: Photodynamic Therapy and the Photosensitizer Design Bottleneck

PDT, as sketched in Figure 1, is a clinically approved treatment in which a photosensitizer (PS) is activated by light to generate reactive chemical species at a precisely defined location and time [6]. This mode of action distinguishes PDT fundamentally from systemic therapies, such as chemotherapy, that expose the entire body to cytotoxic agents and are often associated with severe acute toxicity as well as long-term side effects that can persist for years after treatment. These effects include organ damage, secondary malignancies, immune dysfunction, and long-lasting reductions in quality of life. By confining cytotoxic chemistry to illuminated tissue, PDT offers a therapeutic strategy that can, in principle, achieve effective tumor control while avoiding many of the shortand long-term harms inherent to systemic treatment [7].

Because of this spatial and temporal control, PDT has the potential to support a broader therapeutic role than localized tumor ablation alone. It can be combined with surgery to target residual disease at resection margins, and with imaging modalities that guide light delivery and treatment planning in real time, enabling highly precise, image-guided intervention [8]. In addition, PDT has been shown to induce immunogenic cell death and to modulate the tumor microenvironment in ways that can stimulate antitumor immune responses, opening opportunities for combination with immunotherapy and for systemic benefit beyond the illuminated region [9]. In this sense, PDT offers a rare convergence of localized physical control and biologically meaningful immune engagement.

Despite its clinical promise, the broader impact of photodynamic therapy has remained limited. The main obstacle has not been the therapeutic concept itself, but the lack of photosensitizers whose behavior can be predicted and optimized reliably. Most PDT agents currently used in the clinic are based on chemical scaffolds identified several decades ago, and in more than twenty years there have been very few fundamentally new photosensitizers adopted into routine clinical practice. This stagnation reflects a persistent difficulty in translating chemical innovation into predictable therapeutic performance.

At its core, PDT is governed by a sequence of electronic processes which we summarize in Figure 2. Upon light absorption, a photosensitizer is promoted to an excited singlet state, which may undergo intersystem crossing to form long-lived triplet states. These triplet states drive the chemical reactions responsible for therapeutic action, either through energy transfer to molecular oxygen to generate singlet oxygen (Type II pathways) or through electron-transfer reactions that produce radical species (Type I pathways) [10, 11]. The efficiency, selectivity, and robustness of these processes are determined by the electronic structure of the photosensitizer, in particular by the relative energies, couplings, and character of low-lying excited states. For example, the probability of the absorption and ISC processes, respectively, is directly proportional to the magnitude of the transition dipole moment $\vec{\mu}_{el}$, and the spin-orbit coupling (SOC) matrix element $\langle S| \hat{H}_{\mathrm{SOC}} |T\rangle$ between the singlet and triplet states in question. Hence these processes effectively steer how efficiently a photosensitizer can be activated by light and how the latter triggers distinct therapeutic pathways under different tissue conditions.

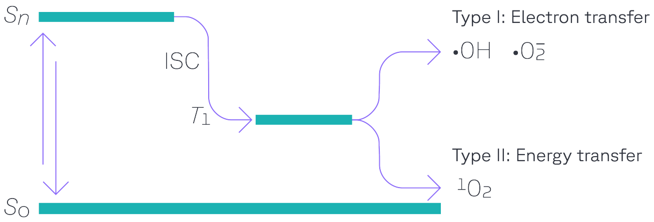

Figure 2: llustration of the key photophysical processes that determine PDT performance. The lowest singlet and triplet states S_0, S_1, and T_1 are the key players in the photophysical processes of light absorption, intersystem crossing, and the generation of either singlet oxygen or radical species.

The promise of PDT lies in the possibility of engineering photosensitizers whose electronic properties are tuned to their biological context. In an ideal scenario, photosensitizers could be designed to respond selectively to local oxygen concentration, redox environment, or tissue-specific conditions, enabling therapies that adapt to the physiological landscape rather than acting uniformly. Achieving such control requires predictive access to excited-state energies, spin-state structure, transition strengths, and redox properties across entire families of related compounds, not just for isolated molecules.

In practice, this level of predictive control has proven elusive. For clinically relevant photosensitizers—particularly transition-metal complexes and extended conjugated systems—the electronic structure is strongly correlated and often near-degenerate. In this regime, small inaccuracies in computed excited-state energies or state ordering can lead to qualitatively incorrect predictions of triplet yield, intersystem crossing efficiency, or dominant reaction pathways. As a result, many standard electronic-structure methods become unreliable, while more accurate approaches quickly become computationally prohibitive as system size grows.

Consequently, PDT drug discovery remains largely empirical [12]. Optimization proceeds through slow, trial-and-error modification of molecular scaffolds, guided by limited experimental feedback rather than by predictive, transferable design principles. This empirical approach is fundamentally misaligned with the complexity of the underlying physics: the therapeutic outcome is set by subtle balances between excited-state energetics, spin conversion pathways, and redox processes that cannot be tuned reliably without quantitative electronic-structure insight. The result is a development pipeline that advances incrementally, despite the existence of a therapeutic concept capable, in principle, of far greater precision and selectivity.

This stands in stark contrast to the level of molecular optimization achieved by biological systems, where function emerges from finely tuned electronic and structural interactions shaped by evolution and continuously refined by environmental feedback [13]. Bridging this gap—moving from empirical exploration to predictive, mechanism-driven design—requires electronic-structure methods that can accurately resolve strongly correlated excited states while remaining scalable to chemically realistic systems. Such methods must not only reproduce individual molecular properties, but enable systematic optimization across families of compounds, reflecting the way biological systems adapt function at the molecular level.

Within the Q4Bio program, we therefore treated PDT not as a downstream application of quantum chemistry, but as a defining challenge that exposes the fundamental limits of current computational approaches. Our objective was not to showcase isolated quantum algorithms or chemically inspired toy models, but to construct an end-to-end computational pipeline capable of delivering the excited-state and spin-resolved information required for PDT-relevant decision-making. In this sense, PDT serves simultaneously as a concrete biomedical target and as a stringent benchmark for what it would mean to optimize drugs at a level approaching the adaptability, specificity, and efficiency of biological systems themselves.

3 Why Classical Techniques Have Not Been Adequate

The classical toolbox for excited-state electronic structure is well-stocked [14], but the requirements imposed by PDT push excited-state approaches into regimes where their limitations become decisive. PDT demands reliable access to multiple low-lying excited states, accurate singlet–triplet energetics, and state-to-state properties for molecules containing transition metals, whose electronic structure is often strongly correlated and sensitive to small perturbations. It is precisely this combination of accuracy, complexity, and scale that pose a challenge for classical excited-state electronic structure approaches.

At the lower end of the accuracy–cost spectrum, mean-field approaches such as density functional theory and time-dependent density functional theory offer favorable computational scaling but are known to be unreliable for many PDT-relevant use-case applications [14]. For photosensitizers containing transition metals or extended conjugated ligands, small errors in excited-state energies or state ordering can lead to qualitatively incorrect predictions of relevant chemistry. As a result, these approaches cannot be relied upon as predictive design tools in the regimes that matter most for PDT.

By contrast, more systematic (multiconfigurational) wavefunction-based methods offer improved accuracy by explicitly capturing electronic correlation, but at the cost of rapidly increasing computational expense. As the active orbital space required to qualitatively describe the correlated excited states grows, the scaling of these methods becomes prohibitive. This limitation is especially acute for PDT, where the relevant observables are not restricted to a single ground state but span entire manifolds of excited singlet and triplet states.

Tensor network methods, most notably DMRG, have significantly extended the reach of classical computation for strongly correlated systems [15, 16] and hence serve as indispensable benchmarks for our work. However, their favorable scaling relies on controlling entanglement through a so-called finite bond dimension

χ. As active orbital spaces grow larger and excited-state manifolds become more entangled, the bond dimension required for an accurate representation increases rapidly, leading to steep computational costs.

A concrete illustration of these limitations is provided by the photosensitizer TLD1433 (Figure 3), which we have used as a benchmark system throughout the Q4Bio project due to its clinical relevance [12, 17–19]. Its extended conjugation and metal–ligand structure give rise to a manifold of low-lying singlet and triplet states with significant multi-reference character. State-of-the-art DMRG calculations show that the bond dimension required to reach chemical accuracy grows rapidly with active space size (Figure 4), exhibiting a steep increase within a relevant active orbital space range until finite-size effects kick in. This reflects the increasing entanglement in the system and implies substantially growing computational cost for molecules of this kind. More importantly, this behavior is not expected to represent a worst-case scenario: as suggested by the shaded region in Figure 4, photosensitizers with increased conjugation or stronger correlation effects are likely to exhibit a delayed onset of finite-size effects and a prolonged unfavorable scaling regime, pushing classical methods further into computationally intractable territory.

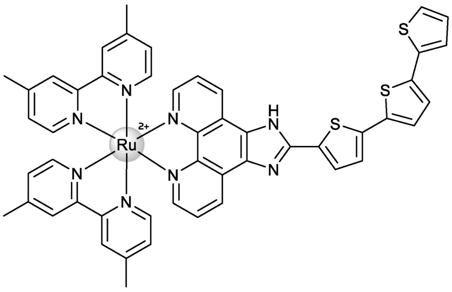

Figure 3: Molecular structure of TLD1433, a Ru(II)-based photosensitizer (Ruvidar™) developed for photodynamic therapy (PDT), currently under clinical evaluation in a Phase II study for Bacillus Calmette–Guérin (BCG)-unresponsive non-muscle invasive bladder cancer (NMIBC).

The consequence of these limitations is that classical electronic-structure methods do not support the iterative design loop required for PDT drug discovery. In principle, one would like to propose candidate photosensitizer structures, predict their excited-state and redox properties

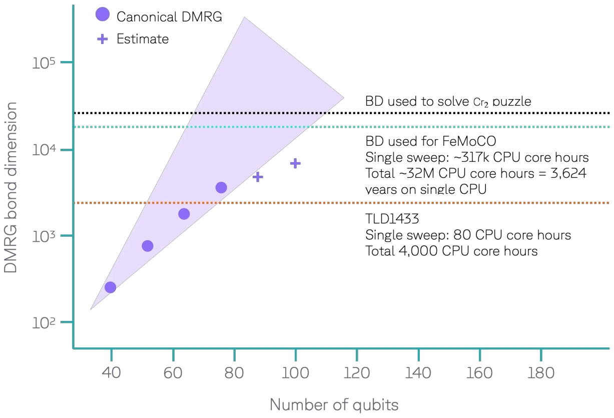

Figure 4: Bond dimension required for DMRG simulations of the S_0 electronic ground state of TLD1433 as a function of a given active orbital space comprising a corresponding number of qubits. The shaded region indicates a projected scaling regime for related photosensitizers with extended conjugation or stronger correlation effects. In such systems, the onset of finite-size effects is expected to occur at larger active space sizes, leading to a prolonged unfavorable growth of the required bond dimension.

with sufficient accuracy to guide chemical modification, and iterate efficiently across families of compounds. In practice, when the available computational tools are either unreliable in the relevant electronic regime or too expensive to deploy at scale, this rational design loop breaks down. PDT development therefore reverts to empirical exploration, with computational modeling playing at best a qualitative or retrospective role rather than serving as a predictive engine.

4 How Quantum Computing Can Change PDT Discovery

PDT-relevant photosensitizers often sit in a regime of electronic structure that is governed by strong correlation and near-degeneracies between low-lying states. Quantum computers are well suited to this setting because they can represent many-body electronic states faithfully, without relying, for example, on approximations that tend to break down for classical mean-field approaches when multiple correlated states compete within a narrow energy window. For PDT, this matters because the design quantities of interest—excited-state energies, singlet–triplet splittings, and the structure of the low-lying manifolds that determine triplet yield and reactivity—are exactly the quantities that become most sensitive to (strong) correlation effects [20, 21].

The usefulness of quantum computing for PDT also depends on the development timeline of the hardware. Fault-tolerant quantum computers are expected to enable high-precision chemistry calculations, but the scale of error correction required for realistic excited-state problems remains beyond current devices. The Wellcome Leap Q4Bio program was explicitly designed around this reality: it was focused on identifying, developing, and demonstrating biology and health applications that can benefit from the quantum computers expected to emerge in the three to five years following the beginning of the program in mid 2023, and on co-developing health applications together with new quantum computing algorithms that enable early optimization with emerging hardware systems. In this near-term setting, the central question is not whether fault tolerance will eventually solve chemistry, but what can be achieved in the coming years with noisy devices, limited circuit depth, and finite measurement budgets.

A practical PDT pipeline requires improvements across several parts of the stack. Circuit preparation must generate correlated reference states while remaining shallow enough to execute. Measurement must be efficient enough to estimate chemical Hamiltonians with controlled statistical error. Noise must be handled in a way that preserves quantitative accuracy, rather than only improving qualitative trends. These requirements are coupled: increasing circuit expressivity is not helpful if measurement cost dominates, and a sophisticated estimator cannot recover accuracy if the quantum signal is overwhelmed by device errors. A workable approach therefore depends on co-design across state preparation, measurement, and post-processing.

Even in the fault-tolerant regime, the efficiency of quantum phase estimation (QPE) is fundamentally limited by the quality of the initial state. QPE requires input states with non-negligible overlap with the target eigenstate, and if this overlap becomes exponentially small with system size—a phenomenon often referred to as the *orthogonality catastrophe*—the number of repetitions required for success grows exponentially, eliminating any practical advantage. In this sense, state preparation is not a peripheral step, but a central bottleneck that determines whether QPE can be used efficiently for chemistry. Our results indicate that chemically relevant states can be prepared with circuits whose size grows only polynomially while maintaining significant overlap with the target state. This provides evidence that the orthogonality catastrophe can be avoided in this setting, establishing a viable path toward efficient fault-tolerant quantum simulation.

It is also important to be clear-sighted about what quantum computing does and does not have the potential to change. Even when quantum hardware reaches regimes where it can accurately represent electronic states that are beyond exact classical limits, chemistry calculations will not become instantaneous. For example, while near-term simulations are limited by device and sampling noise, fault-tolerant simulations will be limited by algorithmic depth and the cost of logical operations. Hence, quantum computing expands the set of tractable problems, but accurate electronic-structure prediction remains computationally demanding.

For PDT discovery, this implies that quantum chemistry will be a high-value component of a broader workflow rather than a high-throughput screening tool on its own. The most practical route to impact is to integrate quantum simulation with AI-driven discovery pipelines that can explore chemical space efficiently, identify promising regions, and allocate expensive quantum calculations to the candidates and properties where they provide the most information [22]. In this setting, quantum computing provides access to the correct correlated physics, and AI provides the scale required to turn that physics into an iterative design process.

5 Q4Bio: Working Toward Practical Quantum Advantage Under Constraints of Current Hardware

Wellcome Leap's Q4Bio program set a clear expectation: the goal was to identify biology and health applications that could plausibly benefit from the quantum computers expected to emerge in the three to five years following the beginning of the project, and to co-develop the algorithms required to make such early optimizations possible. This framing shaped our work from the start. We selected PDT because it is clinically relevant, because it exposes a concrete and well-known computational bottleneck in excited-state electronic structure, and because it naturally connects to a discovery workflow in which predictive simulation could directly accelerate photosensitizer design.

This choice immediately defined the kind of team and working style required. From the outset, our objective was to run large-scale experiments on real quantum hardware, and our prior experience with hardware made it clear that algorithms cannot be developed in a vacuum. The near-term quantum devices available today are noisy, heterogeneous, and far from standardized. As a result, methods that look promising in simulation can behave unexpectedly once compilation constraints, calibration drift, correlated errors, and device-specific idiosyncrasies are taken into account. For this reason, our development process was built around a continuous feedback loop between algorithm design, full-stack simulation, and execution on available hardware. We repeatedly tested the complete pipeline end to end, identified the true bottlenecks in practice rather than in theory, and iterated accordingly.

This workflow also depended on being in constant contact with specialists working on the real biomedical and chemical bottlenecks that motivate the problem. Collaboration with Cleveland Clinic was particularly important in this regard. Their perspective helped ensure that our choices of target systems, observables, and performance criteria remained anchored in the requirements of clinically meaningful PDT design, rather than drifting toward abstract benchmarks. Just as importantly, this close connection enabled us to engage directly with experts who have dedicated their careers to discovering and developing novel photosensitizers, and to shape our technical roadmap around the concrete obstacles they face in day-to-day discovery work. This continuous feedback from domain specialists was essential for keeping the program aligned with the real design bottleneck and for steering our efforts toward methods that can genuinely shift what is possible in practice.

A second guiding principle was to avoid the common compromise of reducing chemical complexity in order to fit within near-term hardware constraints. In quantum chemistry, it is tempting to switch to toy Hamiltonians, minimal basis sets, or problem instances that are easy to run but do not capture the features that make the target chemistry difficult. We did not adopt this approach. Throughout Q4Bio, we focused on chemically non-trivial systems and realistic active spaces that preserve the key physics relevant to PDT, even when this made experiments harder. This decision ensured that each methodological advance translated into progress on the real problem, rather than on simplified proxies.

A third principle was to hold ourselves to a standard of fair comparison. In a fast-moving field, it is easy to report progress against either weak baselines or classical methods that are not competitive for the problem at hand. From the beginning, we benchmarked against the strongest classical techniques available for the relevant physics, including high-quality tensor-network reference calculations, and we evaluated quantum workflows under resource assumptions that reflect what can realistically be executed on hardware. This high bar shaped both the technical direction of the work and the interpretation of results: success was defined not by outperforming a convenient baseline, but by making measurable progress against the best methods that currently exist.

These three requirements, hardware executability, chemical relevance, and fair comparison, are not independent. Neglecting any one of them can produce results that appear encouraging but do not translate into progress on the real PDT bottleneck.

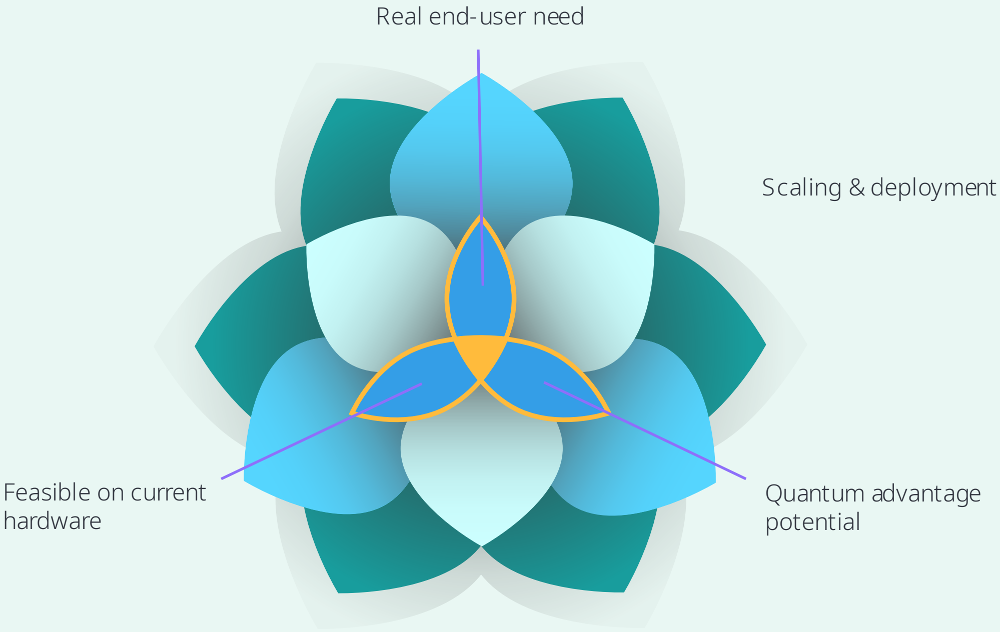

The Lotus of Useful Quantum Advantage

Although unconventional in this context, we use the metaphor of a lotus flower to describe a meaningful attempt toward useful quantum advantage. The inner petals represent three conditions that must be satisfied simultaneously: executability on current hardware, relevance to real-world biologically meaningful chemistry, and validation against state-of-the-art classical methods under fair resource assumptions. Each of these can be achieved in isolation, and significant effort has gone into advancing them individually, but none is sufficient on its own, and progress along any single dimension does not guarantee meaningful overall progress.

The outer petals represent deployment, integration, and scaling. These are essential for translating advances into practice and can be systematically engineered, often leading to systems that appear complete and compelling. However, it is possible to build such systems while avoiding the central challenge. These efforts may look polished and convincing, yet remain at the level of structure and presentation, without advancing the field toward useful quantum advantage.

The central challenge lies in the center of the lotus, where the three conditions come together. Demonstrating this convergence is the difficult step, as it requires satisfying all constraints simultaneously within a single coherent approach.

In a real lotus, the central structure contains the seeds that enable new flowers to grow. In the same spirit, we view this core as generative, capable of supporting multiple applications once established.

Finally, the need for continuous end-to-end testing made access to reliable hardware resources a practical requirement, not an afterthought. For our purposes, IBM's quantum hardware ecosystem was the most natural choice. In addition to offering some of the best available systems in terms of error rates and sampling throughput, IBM provides access to a large number of devices through the cloud. This breadth was critical for running rigorous test loops of the full pipeline, for validating that improvements were robust across hardware instances, and for exposing failure modes that only appear when moving from idealized simulation to real-device execution.

6 State of the Art at the Start of Q4Bio

At the beginning of Q4Bio, it was clear that PDT-relevant chemistry would be more demanding than the ground-state problems that dominated many early quantum chemistry demonstrations on quantum computers. The key observables in PDT depend on low-lying excited-state manifolds, singlet–triplet structure, and spin-dependent pathways, and these requirements are inherently more complex than computing a single ground-state energy. At the same time, it was not obvious from the outset where the true computational boundary lay. In particular, while classical methods could be pushed remarkably far for many ground-state problems, it remained an open and practically important question whether transition-metal complexes relevant for PDT would become genuinely difficult even for the most competitive classical techniques, including tensor-network approaches such as DMRG. Establishing this difficulty was therefore part of what needed to be verified rather than assumed.

Against this background, quantum computing offered a fundamentally different representation of correlated electronic structure by encoding the many-body wavefunction directly on qubits. In principle, this created a route to treat strongly correlated regimes that challenge classical approximations. However, at the start of Q4Bio, quantum chemistry calculations on real hardware remained limited to small chemical instances and tightly constrained workflows. The main barrier was not a lack of proposed algorithms, but the inability to execute them at meaningful chemical scale while maintaining physical relevance and quantitative reliability. This gap was especially consequential for PDT, because the need to resolve multiple low-lying states amplified every hardware limitation compared to ground-state-only tasks.

Three bottlenecks dominated quantum chemistry calculations on hardware at that stage. The first was state preparation. Circuits expressive enough to capture multi-reference character and strongly correlated manifolds typically became too deep to run reliably on noisy devices, particularly when targeting not a single state but a structured set of low-lying singlets and triplets. Hardware-efficient ansätze reduced depth, but often at the cost of chemical structure and controllability, leading to optimization landscapes that were difficult to navigate and results that were not robust across related molecular systems.

The second bottleneck was measurement. Even when a correlated state could be prepared, extracting chemically meaningful quantities required estimating expectation values of Hamiltonians with number of terms scaling fast as the system sizes increased. Shot requirements grew quickly as the number of measured operators increased and as statistical error targets became tighter. For PDT this measurement burden compounded, because one needed not only a ground-state energy but multiple excited-state energies, gaps, and additional observables that determine photophysical pathways. In practice, measurement cost became a dominant limiter of scale and accuracy long before the qubit count itself became the primary obstacle.

The third bottleneck was not merely noise, but the failure of conventional error-mitigation strategies to scale to chemistry-grade accuracy. A broad class of mitigation methods attempts to reconstruct noiseless expectation values from noisy measurements by learning or extrapolating a mapping from calibration circuits that remain classically simulable. For quantum chemistry, this is problematic for two distinct reasons. First, the relation between noisy and noiseless observables is generally unknown and system dependent, so even the calibration stage introduces systematic errors that are already too large on the scale of millihartree accuracy. Second, and more fundamentally, these methods suffer from an intrinsic scaling breakdown: as circuit size and depth increase, noise exponentially suppresses the physical signal, causing noisy expectation values to concentrate into an increasingly narrow numerical range, while the finite-shot uncertainty does not shrink accordingly. The relative size of the error bars therefore grows exponentially with system size, making the inferred mapping increasingly unstable. This problem has been shown to affect a broad class of calibration-based mitigation methods, including zero-noise extrapolation and related approaches, and renders them fundamentally incompatible with the accuracy requirements of realistic electronic-structure simulations at scale [23].

Taken together, these constraints meant that at the outset of Q4Bio, quantum chemistry on hardware was not limited by a single missing ingredient, but by the absence of an end-to-end workflow that balanced expressivity, executability, measurement efficiency, and reliable inference from noisy data. Building such a workflow was therefore a prerequisite for making PDT a credible driver problem for near-term quantum advantage.

7 Our PDT Pipeline Toward Practical Quantum Advantage

We now describe the simulation pipeline developed by Algorithmiq and the scientific advances that made it workable at chemically meaningful scale. The central message is that PDT is not unlocked by a single algorithmic breakthrough. The bottleneck is end-to-end. One must prepare accurate state-specific quantum circuits for strongly correlated systems, extract energies and state-resolved properties under severe measurement constraints, and convert noisy hardware outcomes into chemically meaningful predictions through post-processing that is both statistically controlled and computationally efficient. Each of these steps is difficult in isolation; together they define a systems problem.

7.1 From PDT Requirements to Quantum Circuits: Active Spaces and State-Specific Hamiltonians

Mapping PDT chemistry to a quantum computer begins with a deceptively simple question: what is the smallest electronic-structure problem that still contains the physics that makes a photosensitizer work? A real PDT candidate is a large molecule in a complex environment, but the photophysics and photochemistry that determine performance are governed by a comparatively small set of frontier orbitals and near-degenerate electronic configurations. The role of the mapping step is therefore not to "simplify the molecule," but to isolate the correlated electronic subspace that controls absorption, triplet formation, and downstream reactivity, and then express that subspace in a form that can be represented on qubits.

In practice this is done through an *active-space* construction. The active space defines which orbitals are treated explicitly as a correlated many-electron problem, while the remaining electrons provide an effective mean-field background. For PDT photosensitizers, the choice of active space is unusually demanding because the relevant states are not only excited states but often strongly mixed in character. A photosensitizer typically absorbs light through a singlet excitation that may have metal-centered, ligand-centered, or charge-transfer contributions, and then forms long-lived triplet states that drive Type II or Type I chemistry. Capturing this pathway requires an active space that can represent both the optically accessible singlet manifold and the triplet manifold, together with the configuration mixing that connects them. For transition-metal complexes this means that chemically meaningful active spaces must simultaneously include metal-centered orbitals and the ligand π/π^* system responsible for absorption and charge transfer. This requirement pushes the problem beyond the small "toy" regimes where a single-reference picture remains adequate, and it is precisely this chemically forced complexity that makes PDT a natural driver problem for quantum simulation.

To address the active-space challenge in a systematic and reproducible way, Algorithmiq developed the AEGISS workflow (Atomic orbital and Entropy-based Guided Inference for Space Selection), which combines orbital entanglement measures with chemical intuition to guide active-space construction that is both physically meaningful and hardware-conscious [24]. Rather than relying purely on energy thresholds or heuristic orbital lists, AEGISS uses a combination of orbital entropies [25], chemical locality, and diagnostic metrics to curate a set of orbitals that are likely to contribute to the correlated physics of interest. This semi-automated selection process produces active spaces that remain tractable (in qubit count) while preserving the strong correlation effects that drive PDT behavior.

A second challenge is how to define and compute the excited states themselves. In classical quantum chemistry, many methods approximate excited states as perturbations of a ground-state [14, 21], but such approaches can fail when orbital relaxation and charge-transfer character are strong. In Q4Bio we adopted a state-specific Δ-based approach [26], in which an initial state-specific reference is obtained for each target (e.g. S_0, S_1, T_1) by promoting electrons into a non-Aufbau configuration and optimizing the orbitals and determinants accordingly, in the spirit of ΔSCF. The quantum variational algorithm is then guided not toward the global minimum of a single Hamiltonian but toward the minimum of an effective Hamiltonian corresponding to the target state. This state-specific construction enforces state identity and suppresses variational collapse, which is particularly important in dense spectra such as those found in transition-metal systems.

Once an active space has been chosen and state-specific references have been defined, the mapping to a quantum computer becomes concrete. The electronic Hamiltonian restricted to the active orbitals is written in second quantization and transformed into a qubit Hamiltonian via a fermion-to-qubit mapping. At that point, the quantum device is asked to do a well-defined task: prepare approximate eigenstates of the active-space Hamiltonian and estimate expectation values needed to recover energies and properties. The key is that the excited-state problem has been reformulated into a sequence of state-specific quantum tasks, rather than a single ground-state calculation with an uncertain route to the excited manifold.

This mapping step is scientifically essential for PDT. The photophysical outcome can change qualitatively under small changes in state ordering: whether the molecule efficiently populates a reactive triplet, whether charge-transfer pathways dominate, and whether redox-active channels become accessible all depend on relative energetics at a resolution that is unforgiving to uncontrolled approximations. By combining chemically grounded active spaces with a Δ-based state-specific construction such as ΔADAPT-VQE, we obtain a mapping that preserves state identity and remains aligned with the real photosensitizer design bottleneck: predicting and controlling the low-lying singlet–triplet landscape that determines PDT performance.

7.2 Circuit Generation: Scalable State Preparation Beyond Error Mitigation

For PDT relevant active spaces, the central challenge in circuit generation is not defining an ansatz in principle, but constructing one that is both expressive and executable under realistic constraints. Early approaches to quantum chemistry on near term devices focused on variational energy minimization combined with increasingly sophisticated error mitigation techniques. However, it has become clear that this path does not scale. Error mitigation introduces rapidly growing sampling overheads and instability as system size increases, making chemistry grade accuracy unattainable in practice [23, 27].

In this work, we take a fundamentally different approach. Rather than attempting to extract accurate energies directly from noisy hardware, we focus on scalable state preparation. The goal is to construct quantum circuits that prepare states with high overlap with the true correlated solution, while remaining shallow enough to be executed reliably. This shift is motivated both by near-term considerations and by the requirements of fault-tolerant algorithms. In particular, quantum phase estimation relies critically on the availability of initial states with non negligible overlap with the target eigenstate [28, 29]. If this overlap decays exponentially with system size, the success probability becomes exponentially suppressed, leading to the orthogonality catastrophe [30].

Our key advance is the combination of the ADAPT-VQE algorithm with Algorithmiq's Majorana Propagation, dubbed ADAPT-VMPE [31]. Majorana Propagation [32] enables scalable optimization of Fermionic quantum circuits using polynomial classical resources. It evaluates expectation values of observables for circuits composed of Fermionic operators without simulating the full quantum state, instead relying on low order expansions and a truncation strategy with provable exponential error suppression. This allows the expensive inner loop of adaptive circuit construction to be carried out entirely classically, eliminating the measurement bottleneck that limits standard ADAPT-VQE workflows [33].

The resulting state preparation procedure produces circuits whose size grows polynomially with system size. As shown in Figure 5, the number of Fermionic operators required to achieve fixed overlap with the ground state exhibits approximately quadratic scaling across the range of active spaces considered. Increasing the target overlap from 70% to 85% leads only to a modest increase in circuit size. This demonstrates that high overlap states can be prepared without exponential growth in circuit resources, directly addressing the initial state bottleneck that limits both near term algorithms and quantum phase estimation.

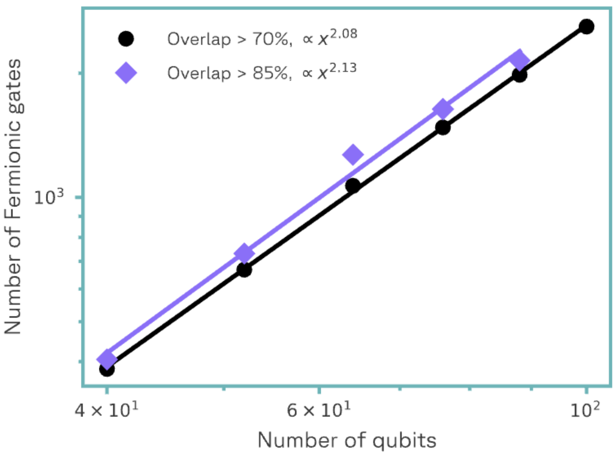

Figure 5: Scaling of ADAPT-VMPE circuit size for TLD1433 with increasing active space sizes. The number of Fermionic operators required to reach fixed overlap with the target state grows polynomially. This favorable scaling implies polynomial circuit depth after transpilation, since active rotations can be absorbed into the Hamiltonian and do not contribute to hardware depth.

This favorable scaling has two key consequences. First, it shows that preparing chemically relevant states does not inherently require exponentially large circuits, providing evidence that the orthogonality catastrophe can be avoided in practice for these systems. Second, it enables the construction of circuits that remain shallow enough to be executed on current hardware, even at system sizes approaching 100 qubits. In contrast to classical tensor network methods such as DMRG, where the required bond dimension grows rapidly with entanglement, the quantum circuits produced by ADAPT-VMPE maintain controlled depth while capturing the relevant correlations.

It is important to emphasize that the goal of this stage is not to produce chemically accurate energies on quantum hardware. Instead, the objective is to prepare high quality quantum states that can be sampled efficiently. These samples contain information about correlated electronic structure that is difficult to obtain classically, even if individual expectation values are noisy. By combining these samples with classical post-processing techniques, such as tensor network or configuration interaction methods [34, 35], we obtain a hybrid workflow in which quantum hardware provides non classical input that enhances classical computation.

In this sense, the role of the quantum computer is not to replace classical quantum chemistry, but to act as a generator of correlated samples that augment it. Majorana Propagation and ADAPT-VMPE make this possible by removing the circuit construction bottleneck, enabling the generation of large expressive circuits with high overlap at polynomial cost. The remaining challenge is therefore not state preparation, but the extraction and integration of useful information from quantum measurements.

7.3 Efficient Measurement Schemes for Chemistry Grade Estimation

Building on the previous section, where we showed that high overlap quantum states can be prepared efficiently, the role of measurement in our pipeline is not to extract chemically accurate energies directly from hardware. Instead, the objective is to extract the most relevant information from these states in a form that can be used effectively in downstream classical post processing.

For quantum chemistry on quantum computers, measurement is often the dominant computational bottleneck [36]. The challenge is structural: electronic structure Hamiltonians contain a large number of terms, and chemically meaningful predictions require extremely high precision in expectation values. As a result, the relevant cost is not a single circuit execution, but the number of samples required to suppress statistical uncertainty to chemically meaningful levels.

A wide range of techniques has been developed to address this problem, including Hamiltonian structured measurements and randomized measurement protocols such as classical shadows [37–39]. While these approaches can reduce the number of measurement settings or enable the estimation of many observables simultaneously, they do not fundamentally resolve the scaling problem. In particular, achieving chemistry grade precision for realistic systems typically requires a number of samples that grows rapidly with system size, making direct energy estimation on hardware impractical.

Algorithmiq's measurement framework is designed around a different objective. Rather than minimizing the variance of energy estimates in isolation, we aim to maximize the usefulness of the measurement data for downstream algorithms. This leads to a measurement strategy based on informationally complete data acquisition [40] combined with variance optimization targeted at physically relevant observables [41–43].

The first ingredient is locally biased informationally complete, readout-error-aware measurements [44]. Instead of sampling measurement outcomes uniformly, we bias the measurement distribution toward operators that contribute most strongly to the target observables. This concentrates sampling effort where it has the largest impact on the extracted information content.

The second ingredient is dual optimization. For a given informationally complete measurement, one can construct estimators for observables through a choice of dual operators. While many such choices exist, their statistical efficiency can differ dramatically. By optimizing the dual operators with respect to the Hamiltonian and measurement statistics, we obtain unbiased estimators whose variance is reduced by many orders of magnitude compared to standard reconstruction methods [41–43].

The impact of this approach is most clearly seen in the scaling of the required sampling cost. As shown in Figure 6, standard classical shadow estimators require an astronomically large number of measurement shots to reach chemical precision for the TLD1433 Hamiltonian, far beyond what can be achieved on current hardware. In contrast, locally optimized dual estimators reduce the sampling cost by multiple orders of magnitude. Increasing the degree of nonlocality in the dual reconstruction leads to a systematic reduction in variance, bringing the required number of shots into a regime compatible with realistic hardware execution times. Extrapolations indicate reductions of up to 18 orders of magnitude in shot requirements compared to canonical shadow estimators [44].

Despite these improvements, we emphasize that reaching chemistry grade precision in energy estimates on hardware remains impractical in the near term. Crucially, however, this is not the goal of our pipeline. The purpose of the measurement stage is to extract structured information from the quantum state that can be used as input to classical quantum chemistry methods.

In this hybrid setting, the quantum processor acts as a generator of correlated samples rather than a standalone solver. Efficient measurement protocols ensure that the most relevant information is extracted from these samples with minimal overhead, making them suitable for integration into downstream methods such as tensor network optimization. The value of the quantum device therefore lies not in producing final observables directly, but in providing high quality data that enhances classical computation.

7.4 QB-DMRG: Post-Processing as Variational Map Learning on Quantum Data

Building on the previous stages of the pipeline, where quantum hardware is used to prepare high-overlap states and extract structured measurement data, the role of post-processing is to convert this information into lower-energy variational states within a controlled tensor-network ansatz. In this final stage, we do not attempt to reconstruct a noiseless quantum state or compute chemically accurate energies directly on hardware. Instead, quantum data is used to enhance a classical variational optimization. Algorithmiq's Quantum Boosted DMRG (QB-DMRG) implements this idea by learning a variational map acting on the measured state: given

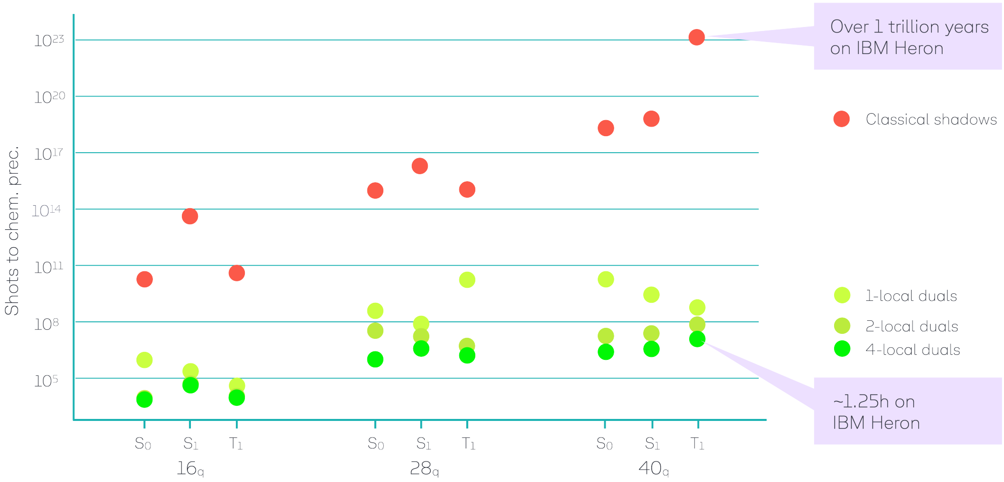

Figure 6: Measurement cost scaling for the TLD1433 Hamiltonian. Standard classical shadow estimators require an impractically large number of samples to reach chemical precision. Locally optimized dual estimators reduce the sampling cost by many orders of magnitude. Increasing the degree of nonlocality in the dual reconstruction leads to systematically improved variance reduction and enables practical measurement costs for realistic systems.

an estimate obtained from informationally complete measurements, it searches for a transformation that lowers the energy with respect to the target Hamiltonian while remaining within a fixed bond-dimension tensor-network ansatz.

This workflow is illustrated schematically in Figure 7. A low-depth quantum circuit prepares a correlated state that captures relevant structure of the problem but is not sufficiently accurate on its own. Informationally complete measurements produce a quasistate, i.e., an unbiased finite-shot estimator of the quantum state, which encodes this correlation structure. QB-DMRG uses this quasistate as input to a tensor-network optimization, learning a transformation that maps it to a lower-energy state's quasistate within the chosen variational class. The advantage of this approach lies in the structure of the quantum reference state: even at low depth, quantum circuits can generate entanglement corresponding to an effective bond dimension that grows exponentially with circuit depth. While such states cannot be directly represented within a fixed bond-dimension ansatz, their quasistates contain correlations that would require a significantly larger bond dimension to reproduce classically. QB-DMRG exploits this by compressing and refining these correlations into a tensor-network representation of fixed bond dimension, effectively accessing states that would otherwise require larger classical resources.

A key structural property of QB-DMRG is that it cannot underperform classical DMRG at matched bond dimension, since its variational class contains all states accessible to DMRG at the same bond dimension. It therefore recovers the classical result in the worst case, while allowing strictly lower energies when the quantum reference state provides additional useful correlations. The computational cost follows from this hybrid structure: contraction scales linearly in the number of measurement shots, while tensor-network optimization scales polynomially in system size and bond dimension, so that the overall cost remains polynomial for polynomial shot counts. The comparison with classical DMRG is thus governed by the bond dimension required to reach a given energy. Classical DMRG must increase bond dimension to capture additional correlation, leading to rapidly growing cost, whereas QB-DMRG can achieve comparable or lower energies at fixed bond dimension by incorporating quantum-generated correlations, corresponding to an effective increase in representational power without increasing classical resources.

An important consequence is that QB-DMRG avoids the scaling limitations of conventional error mitigation. Rather than reconstructing noiseless expectation values, which incurs exponential overhead, it uses quantum measurement data directly within a variational optimization. Its performance therefore does not de-

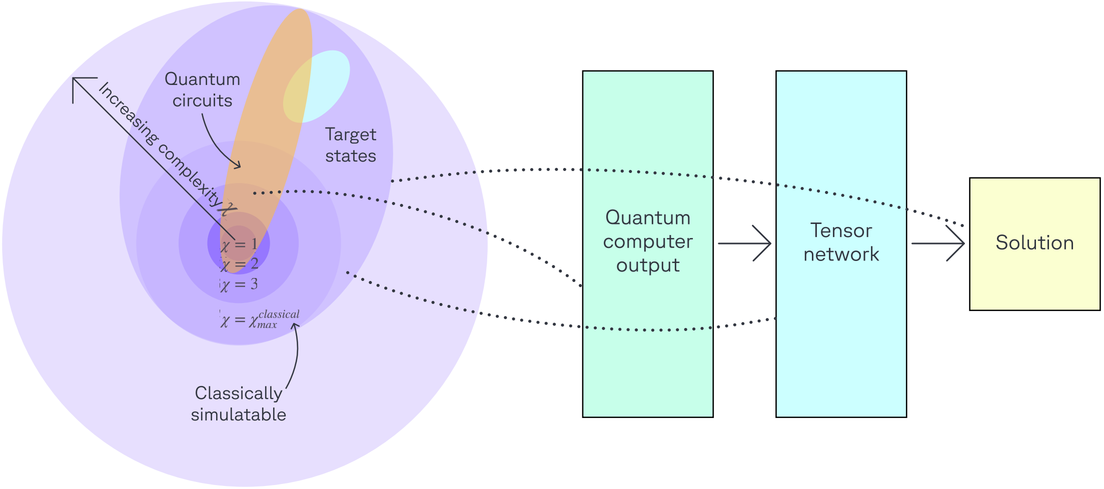

Figure 7: Conceptual illustration of the QB-DMRG workflow. A low-depth quantum circuit prepares a correlated state with large effective bond dimension. Informationally complete measurements produce a quasistate, an unbiased finite-shot estimator of the quantum state, capturing this correlation structure. QB-DMRG learns a tensor-network map that transforms this quasistate into a refined state within a fixed bond-dimension ansatz, effectively compressing quantum-generated correlations into a classically tractable representation.

pend on achieving chemistry-grade accuracy at the level of raw measurements, but on the prepared state having non-negligible overlap with the target state and capturing relevant entanglement structure. Under these conditions, the optimization can incorporate quantum information to produce lower-energy states within the fixed bond-dimension ansatz. Together with the previous stages, this yields a consistent scaling picture: circuit preparation remains polynomial while maintaining high overlap, measurement cost is controlled through variance optimization, and post-processing remains polynomial while enabling improvements over classical tensor-network methods at fixed bond dimension, providing a concrete route toward quantum advantage in strongly correlated quantum chemistry.

7.5 Towards Quantum Enhancement of Tensor-Network Methods

A central objective of Wellcome Leap's Q4Bio program is to determine whether quantum computing can meaningfully impact biology and health in the near term by improving real computational workflows. In this context, the relevant question is not whether quantum devices can outperform classical methods on synthetic benchmarks, but whether they can provide measurable benefit when combined with state-of-the-art techniques used in practice.

In this work, we address this question for strongly correlated electronic structure calculations, where tensor-network methods such as DMRG define the classical state of the art. Algorithmiq's hybrid approach, comprising state preparation, informationally complete measurements, and QB-DMRG, is designed to augment these methods by incorporating information obtained from quantum hardware directly into the tensor-network optimization.

To test this idea, we executed circuits prepared using the ADAPT-VMPE procedure on IBM's `ibm_cleveland` quantum processor (156 qubits, heavy-hex topology). Measurement data were collected using informationally complete, locally biased POVMs, producing quasistates—unbiased finite-shot estimators of the prepared quantum states. These data serve as input to the QB-DMRG optimization, which refines the tensor-network state at fixed bond dimension.

The central observation is that, for a 52-qubit active-space instance of the TLD1433 Hamiltonian (see Fig. 8), QB-DMRG produces energies lower than those obtained by classical DMRG at the same bond dimension. Since both methods operate under similar classical resource constraints, determined by the bond dimension, this establishes that quantum-generated data can directly improve a leading classical electronic-structure method.

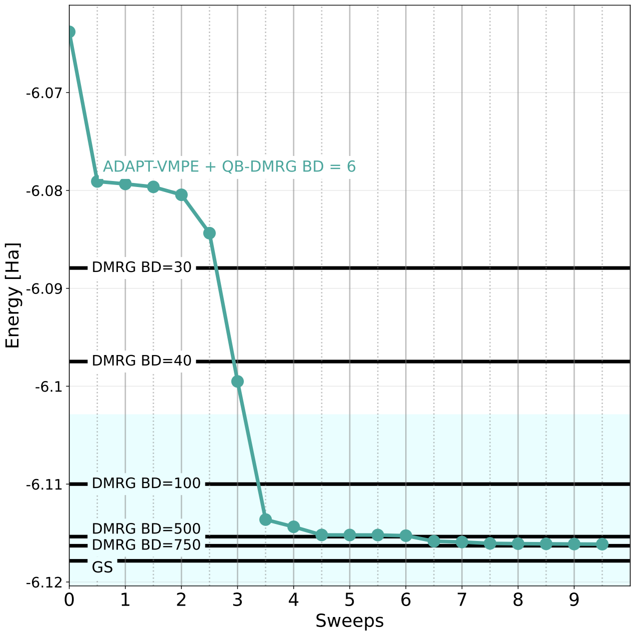

Figure 8: Algorithmiq's hybrid pipeline improves over classical DMRG at fixed bond dimension. Energy obtained during sequential QB-DMRG optimization (sweeps) for a 52-qubit active-space instance of the TLD1433 Hamiltonian. The reference quantum state was prepared using an ADAPT-VMPE circuit and sampled on quantum hardware, producing quasistates (unbiased finite-shot estimators of the state) used as input to the tensor-network optimization. The horizontal dashed lines indicate classical DMRG energies at increasing bond dimensions. Starting from the same tensor-network resources, QB-DMRG incorporates quantum-generated correlation information and lowers the energy below the corresponding classical DMRG values at fixed bond dimension.

The mechanism underlying this improvement is clear. Classical DMRG encodes correlations through bond dimension, and restricting this dimension limits the entanglement that can be represented. In contrast, the quantum circuit prepares a correlated many-body state whose effective bond dimension grows with circuit depth. Measurement samples therefore contain correlation information that is inaccessible to classical tensor networks at the same bond dimension. QB-DMRG incorporates this information into the variational optimization, effectively extending the representational power of the tensor-network ansatz without increasing its classical cost.

At the same time, these results should be interpreted in the appropriate regime. Classical DMRG remains more powerful when unconstrained resources are available, and increasing bond dimension will eventually recover the same correlations. The significance of the present result lies instead in the regime where bond dimension becomes the limiting resource. In this regime, quantum-generated data can provide access to correlations that are otherwise inaccessible at comparable classical cost.

Extending this behavior to larger system sizes remains an important open challenge. In particular, demonstrating consistent improvements in regimes where the bond dimension becomes a dominant bottleneck for classical methods will be essential for establishing practical impact. While the current results clearly establish the underlying mechanism, further work is required to make the QB-DMRG optimization robust at larger scales and to fully exploit the information contained in quantum-generated datasets. In addition to the 52-qubit benchmark discussed above, we carried out QB-DMRG simulations for larger instances corresponding to 76- and 100-qubit representations of the TLD system. These simulations confirm that the full hybrid pipeline can be executed for substantially larger system sizes. At the same time, in the current implementation we have not yet succeeded in converging the QB- DMRG optimization for the larger instances to energies below those obtained from classical DMRG at the same bond dimension. Achieving stable convergence in this regime requires further refinement of the circuit depth selection criterion in order to determine the depth at which two competing forces, entanglement creation and noise accumulation, find an ideal balance.

Taken together, these results provide a concrete step toward the Q4Bio objective: demonstrating that quantum hardware can contribute to scientifically relevant workflows by enhancing state-of-the-art classical methods. Rather than demonstrating advantage in isolation, this work establishes a pathway toward useful quantum advantage through tightly integrated hybrid algorithms.

8 Road to Practical Quantum Advantage for PDT

Our view is that practical quantum advantage in chemistry will not emerge as a single dramatic crossover, but rather as a sequence of capability thresholds, each expanding the class of problems that can be addressed in practice. In this context, there is a clear need for a community-driven effort, analogous to existing quantum advantage trackers [45], to systematically benchmark progress against the strongest classical methods under realistic resource assumptions.

For PDT, a first critical milestone is reaching chemically relevant active spaces at the 100-qubit scale. This requires continued scaling of the full hybrid pipeline, including circuit generation, measurement, and post-processing, together with improvements in robustness and efficiency. The results presented here indicate that this regime is already within reach, with encouraging scaling behavior observed across all components of the workflow. Beyond system size alone, PDT design depends on a broader set of quantities than ground-state energies, including transition dipoles, intersystem crossing pathways, and excited-state redox energetics. Extending the present framework to multi-state QB-DMRG and map-based property extraction therefore provides a pathway to compute these properties at scale, moving from energy estimation toward predictive modeling of photophysical processes.

The ultimate impact of quantum computing in PDT will be realized when it becomes part of a closed discovery loop. We have already demonstrated an AI-guided screening framework for Type I PDT agents (see Ref. 46), combining high-fidelity quantum-chemical data with active learning over large chemical spaces, including million-scale libraries of transition metal complexes. Notably, even at the level of classical quantum chemistry, this pipeline already provides useful input to the PDT drug discovery process by identifying promising candidates and structure–property relationships. In this setting, quantum methods act as a source of higher-fidelity data in regimes where classical methods become unreliable or prohibitively expensive. As quantum chemistry calculations become more accurate through hybrid quantum-classical approaches, they can be directly integrated into this discovery pipeline, further improving model accuracy and enabling the exploration of chemical space beyond the reach of current methods. The value of practical quantum advantage is therefore not only in improving individual computations, but in enabling the design of new molecules by systematically expanding what can be predicted reliably.

Beyond Q4Bio, we expect this hybrid stack to generalize across a broader class of photochemical and materials discovery problems, including OLEDs [47], photocatalysis [48, 49], and solar-energy conversion [50–52]. In the longer term, as fault-tolerant quantum hardware becomes available, the same framework is expected to play a complementary role. Near-term methods provide scalable state preparation and validation tools, which can serve as initialization and verification layers for fault-tolerant algorithms such as quantum phase estimation. In this way, the present approach not only addresses near-term applications, but also forms a bridge toward the fault-tolerant regime.

References

[1] S. Wouters and D. van Neck. The density matrix renormalization group for ab initio quantum chemistry. *Eur. Phys. J. D*, 68:272, 2014.

[2] Sz. Szalay, M. Pfeffer, V. Murg, G. Barcza, F. Verstraete, R. Schneider, and Ö. Legeza. Tensor product methods and entanglement optimization for *Ab Initio* quantum chemistry. *Int. J. Quantum Chem.*, 115:1342–1391, 2015.

[3] S. Knecht, E. D. Hedegaard, S. Keller, A. Kovyrshin, Y. Ma, A. Muolo, C. J. Stein, and M. Reiher. New approaches for ab initio calculations of molecules with strong electron correlation. *Chimia*, 70:244–251, 2016.

[4] Roberto Olivares-Amaya, Weifeng Hu, Naoki Nakatani, Sandeep Sharma, Jun Yang, and Garnet Kin-Lic Chan. The ab-initio density matrix renormalization group in practice. *J. Chem. Phys.*, 142:034102, 2015.

[5] A. Baiardi and M. Reiher. The density matrix renormalization group in chemistry and molecular physics: Recent developments and new challenges. *J. Chem. Phys.*, 152:040903, 2020.

[6] Dennis EJGJ Dolmans, Dai Fukumura, and Rakesh K Jain. Photodynamic therapy for cancer. *Nat. Rev. Cancer*, 3(5):380–387, 2003.

[7] Patrizia Agostinis, Kristian Berg, Keith A Cengel, Thomas H Foster, Albert W Girotti, Sandra O Gollnick, Stephen M Hahn, Michael R Hamblin, Asta Juzeniene, David Kessel, et al. Photodynamic therapy of cancer: An update. *CA Cancer J. Clin.*, 61(4):250–281, 2011.

[8] Girgis Obaid, Jonathan P. Celli, Mans Broekgaarden, Anne-Laure Bulin, Petteri Uusimaa, Brian Pogue, Tayyaba Hasan, and Huang-Chiao Huang. Engineering photodynamics for treatment, priming and imaging. *Nat. Rev. Bioeng.*, 2(9):752–769, 2024.

[9] Marta Warszynska, Pawel Repetowski, and Janusz M. Dabrowski. Photodynamic therapy combined with immunotherapy: Recent advances and future research directions. *Coord. Chem. Rev.*, 495:215350, 2023.

[10] Claude Schweitzer and Reinhard Schmidt. Physical mechanisms of generation and deactivation of singlet oxygen. *Chem. Rev.*, 103(5):1685–1758, 2003.

[11] Maurício S. Baptista, Jean Cadet, Paolo Di Mascio, Ashwini A. Ghogare, Alexander Greer, Michael R. Hamblin, Carolina Lorente, Silvia Cristina Nunez, Martha Simões Ribeiro, Andrés H. Thomas, et al. Type I and type II photosensitized oxidation reactions: Guidelines and mechanistic pathways. *Photochem. Photobiol.*, 93(4):912–919, 2017.

[12] Susan Monro, Katsuya L. Colon, Huimin Yin, John Roque III, Prathyusha Konda, Shashi Gujar, Randolph P. Thummel, Lothar Lilge, Colin G. Cameron, and Sherri A. McFarland. Transition metal complexes and photodynamic therapy from a tumor-centered approach: Challenges, opportunities, and highlights from the development of TLD1433. *Chem. Rev.*, 119(2):797–828, 2018.

[13] Alberto Baiardi, Matthias Christandl, and Markus Reiher. Quantum computing for molecular biology. *ChemBioChem*, 24:e202300120, 2023.

[14] Andreas Dreuw. Why computational photochemistry is challenging and will probably remain so: A quantum chemist's perspective. *Advanced Science*, page e21012, 2026.

[15] Henrik R. Larsson, Huanchen Zhai, C. J. Umrigar, and Garnet Kin-Lic Chan. The chromium dimer: Closing a chapter of quantum chemistry. *J. Am. Chem. Soc.*, 144:15932–15937, 2022.

[16] Huanchen Zhai, Chenghan Li, Xing Zhang, Zhendong Li, Seunghoon Lee, and Garnet Kin-Lic Chan. Classical solution of the femo-cofactor model to chemical accuracy and its implications, 2026. arXiv: 2601.04621.

[17] Lothar Lilge, Mark Roufaiel, Savo Lazic, Pavel Kaspler, Manjunatha A. Munegowda, Mark Nitz, Jay Bassan, and Arkady Mandel. Evaluation of a ruthenium coordination complex as photosensitizer for pdt of bladder cancer: Cellular response, tissue selectivity, and in vivo response. *Transl. Biophotonics*, 2(1-2):e201900032, 2020.

[18] Lothar Lilge, Jenny Wu, Yiwen Xu, Angelica Manalac, Daniel Molenhuis, Fynn Schwiegelshohn, Leonid Vesselov, Wayne Embree, Michael Nesbit, Vaughn Betz, et al. Minimal required pdt light dosimetry for nonmuscle invasive bladder cancer. *J. Biomed. Opt.*, 25(6):068001–068001, 2020.

[19] Girish S. Kulkarni, Lothar Lilge, Michael Nesbitt, Roger J. Dumoulin-White, Arkady Mandel, and Michael A. S. Jewett. A phase 1b clinical study of intravesical photodynamic therapy in patients with bacillus

calmette–guérin–unresponsive non–muscle-invasive bladder cancer. *Eur. Urol. Open Sci.*, 41:105–111, 2022.

[20] Michael Jäger, Leon Freitag, and Leticia González. Using computational chemistry to design Ru photosensitizers with directional charge transfer. *Coord. Chem. Rev.*, 304-305:146–165, 2015.

[21] Hans Lischka, Dana Nachtigallová, Adélia J. A. Aquino, Péter G. Szalay, Felix Plasser, Francisco B. C. Machado, and Mario Barbatti. Multireference approaches for excited states of molecules. *Chem. Rev.*, 118:7293–7361, 2018.

[22] David E. Graff, Eugene I. Shakhnovich, and Connor W. Coley. Accelerating high-throughput virtual screening through molecular pool-based active learning. *Chem. Sci.*, 12(22):7866–7881, 2021.

[23] Sergey N. Filippov, Sabrina Maniscalco, and Guillermo García-Pérez. Scalability of quantum error mitigation techniques: from utility to advantage, 2024. arXiv: 2403.13542.

[24] F. Tarocco, P. A. B. Haase, F. Pavošević, V. Krishna, L. Guidoni, S. Knecht, and M. Stella. AEGISS – atomic orbital and entropy-based guided inference for space selection – a novel semi-automated active space selection workflow for quantum chemistry and quantum computing applications, 2025. arXiv: 2508.10671.

[25] K. Boguslawski and P. Tecmer. Orbital entanglement in quantum chemistry. *Int. J. Quantum Chem.*, 115:1289–1295, 2015.

[26] A. Nykänen, L. Thiessen, E.-M. Borrelli, V. Krishna, S. Knecht, and F. Pavosevic. Toward accurate calculation of excitation energies on quantum computers with ΔADAPT-VQE: A case study of BODIPY derivatives. *J. Phys. Chem. Lett.*, 15:7111–7117, 2024.

[27] Zoltán Zimborás, Bálint Koczor, Zoë Holmes, Elsi-Mari Borrelli, András Gilyén, Hsin-Yuan Huang, Zhenyu Cai, Antonio Acín, Leandro Aolita, Leonardo Banchi, Fernando G. S. L. Brandão, Daniel Cavalcanti, Toby Cubitt, Sergey N. Filippov, Guillermo García-Pérez, John Goold, Orsolya Kálmán, Elica Kyoseva, Matteo A. C. Rossi, Boris Sokolov, Ivano Tavernelli, and Sabrina Maniscalco. Myths around quantum computation before full fault tolerance: What no-go theorems rule out and what they don't, 2025. arXiv: 2501.05694.

[28] Seunghoon Lee, Joonho Lee, Huanchen Zhai, Yu Tong, Alexander M. Dalzell, Ashutosh Kumar, Phillip Helms, Johnnie Gray, Zhi-Hao Cui, Wenyuan Liu, Michael Kastoryano, Ryan Babbush, John Preskill, David R. Reichman, Earl T. Campbell, Edward F. Valeev, Lin Lin, and Garnet Kin-Lic Chan. Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry. *Nature Communications*, 14:1952, 2023.

[29] Dominic W. Berry, Yu Tong, Tanuj Khattar, Alec White, Tae In Kim, Guang Hao Low, Sergio Boixo, Zhiyan Ding, Lin Lin, Seunghoon Lee, Garnet Kin-Lic Chan, Ryan Babbush, and Nicholas C. Rubin. Rapid initial-state preparation for the quantum simulation of strongly correlated molecules. *PRX Quantum*, 6:020327, 2025.

[30] P. W. Anderson. Infrared catastrophe in fermi gases with localscattering potentials. *Phys. Rev. Lett.*, 8:1049, 1967.

[31] Rahul Chakraborty, Aaron Miller, Anton Nykänen, Özlem Salehi, Fabio Tarocco, Fabijan Pavosevic, Pi Haase, Martina Stella, and Adam Glos. Scalable quantum circuit generation for iterative ground state approximation using Majorana Propagation. 2026. arXiv: 2603.23444.

[32] Aaron Miller, Joachim Favre, Zoë Holmes, Özlem Salehi, Rahul Chakraborty, Anton Nykänen, Zoltán Zimborás, Adam Glos, and Guillermo García-Pérez. Simulation of fermionic circuits using Majorana propagation. *arXiv preprint arXiv:2503.18939*, 2025.

[33] Harper R. Grimsley, Sophia E. Economou, Edwin Barnes, and Nicholas J. Mayhall.

An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nature Communications*, 10:3007, 2019.

[34] Javier Robledo-Moreno, Mario Motta, Holger Haas, Ali Javadi-Abhari, Petar Jurcevic, William Kirby, Simon Martiel, Kunal Sharma, Sandeep Sharma, Tomonori Shirakawa, Iskandar Sitdikov, Rong-Yang Sun, Kevin J. Sung, Maika Takita, Minh C. Tran, Seiji Yunoki, and Antonio Mezzacapo. Chemistry beyond the scale of exact diagonalization on a quantum-centric supercomputer. *Science Advances*, 11(25):eadu9991, 2025.

[35] Samuele Piccinelli, Alberto Baiardi, Stefano Barison, Max Rossmannek, Almudena Carrera Vazquez, Francesco Tacchino, Stefano Mensa, Edoardo Altamura, Ali Alavi, Mario Motta, Javier Robledo-Moreno, William Kirby, Kunal Sharma, Antonio Mezzacapo, and Ivano Tavernelli. Quantum chemistry with provable convergence via randomized sample-based Krylov quantum diagonalization. *arXiv*, 2025.

[36] Smik Patel, Praveen Jayakumar, Tzu-Ching Yen, and Artur F. Izmaylov. Quantum measurement for quantum chemistry on a quantum computer. *Chem. Rev.*, 125(16):7490–7524, 2025.

[37] Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nat. Phys.*, 16(10):1050–1057, June 2020.

[38] Scott Aaronson. Shadow tomography of quantum states. *SIAM Journal on Computing*, 49(5):STOC18–368–STOC18–394, 2020.

[39] Andreas Elben, Steven T. Flammia, Hsin-Yuan Huang, Richard Kueng, John Preskill, Benoît Vermersch, and Peter Zoller. The randomized measurement toolbox. *Nature Reviews Physics*, 5(1):9–24, December 2022.

[40] Guillermo García-Pérez, Matteo A.C. Rossi, Boris Sokolov, Francesco Tacchino, Panagiotis Kl. Barkoutsos, Guglielmo Mazzola, Ivano Tavernelli, and Sabrina Maniscalco. Learning to measure: Adaptive informationally complete generalized measurements for quantum algorithms. *PRX Quantum*, 2:040342, 2021.

[41] Joonas Malmi, Keijo Korhonen, Daniel Cavalcanti, and Guillermo García-Pérez. Enhanced observable estimation through classical optimization of informationally overcomplete measurement data: Beyond classical shadows. *Physical Review A*, 109:062412, 2024.

[42] Stefano Mangini and Daniel Cavalcanti. Low variance estimations of many observables with tensor networks and informationally-complete measurements. *Quantum*, 9:1812, July 2025.

[43] Keijo Korhonen, Stefano Mangini, Joonas Malmi, Hetta Vappula, and Daniel Cavalcanti. Improving shadow estimation with locally-optimal dual frames, 2025. arXiv: 2511.02555.

[44] Keijo Korhonen, Hetta Vappula, Adam Glos, Marco Cattaneo, Zoltán Zimborás, Elsi-Mari Borrelli, Matteo A. C. Rossi, Guillermo García-Pérez, and Daniel Cavalcanti. Practical techniques for high-precision measurements on near-term quantum hardware and applications in molecular energy estimation. *npj Quantum Information*, 11(1):110, Jul 2025.

[45] `https://quantum-advantage-tracker.github.io`, accessed 25.03.2026.

[46] Alessio Fallani, Pi A. B. Haase, Julianne F. F. Eckert, Luukas Nikkanen, Sherri A. McFarland, Martina Stella, and Fabijan Pavošević. Data-efficient active learning discovery of transition metal photosensitizers for type i photodynamic therapy. *arXiv:2603.19912*, 2026.

[47] Tian-Yi Li, Shu-Jia Zheng, Peter I. Djurovich, and Mark E. Thompson. Two-coordinate thermally activated delayed fluorescence coinage metal complexes: Molecular design, photophysical characters,

and device application. *Chem. Rev.*, 124(7):4332–4392, 2024.

[48] Christopher K Prier, Danica A Rankic, and David WC MacMillan. Visible light photoredox catalysis with transition metal complexes: Applications in organic synthesis. *Chem. Rev.*, 113(7):5322–5363, 2013.

[49] Susanne Reischauer and Bartholomäus Pieber. Emerging concepts in photocatalytic organic synthesis. *iScience*, 24(3):102209, 2021.

[50] Phong D. Tran, Lydia H. Wong, James Barber, and Joachim S. C. Loo. Recent advances in hybrid photocatalysts for solar fuel production. *Energy Environ. Sci.*, 5(3):5902–5918, 2012.

[51] Ulrich Ulmer, Thomas Dingle, Paul N. Duchesne, Robert H. Morris, Alexandra Tavasoli, Thomas Wood, and Geoffrey A. Ozin. Fundamentals and applications of photocatalytic CO_2 methanation. *Nat. Commun.*, 10:3169, 2019.

[52] Shunta Nishioka, Frank E. Osterloh, Xinchen Wang, Thomas E Mallouk, and Kazuhiko Maeda. Photocatalytic water splitting. *Nat. Rev. Methods Primers*, 3:42, 2023.