



Article

---

# Defining Quantum Agents: Formal Foundations, Architectures, and NISQ-Era Prototypes

---

Eldar Sultanow, Madjid Tehrani, Siddhant Dutta, William J. Buchanan and Muhammad Shahbaz Khan

Special Issue

Beyond Classical Limits: Quantum Machine Learning for Multi-Field Research


Edited by

Dr. Fahad Ahmad and Prof. Dr. Vincenzo Tamma



Article

# Defining Quantum Agents: Formal Foundations, Architectures, and NISQ-Era Prototypes

Eldar Sultanow <sup>1</sup>, Madjid Tehrani <sup>2</sup>, Siddhant Dutta <sup>3</sup>, William J. Buchanan <sup>2,\*</sup>  
and Muhammad Shahbaz Khan <sup>2,4</sup>

<sup>1</sup> Capgemini Germany, 90402 Nuremberg, Germany; eldar.sultanow@capgemini.com

<sup>2</sup> School of Computing, Engineering and the Built Environment, Edinburgh Napier University, Edinburgh EH10 5DT, UK; madjid\_tehrani@gwu.edu (M.T.); m.khan2@napier.ac.uk (M.S.K.)

<sup>3</sup> College of Computing and Data Science, Nanyang Technological University, Singapore 639798, Singapore; siddhant010@e.ntu.edu.sg

<sup>4</sup> School of Computer Science and Digital Technologies, Aston University, Birmingham B4 7ET, UK

\* Correspondence: b.buchanan@napier.ac.uk

## Abstract

Quantum computing offers potential computational advantages, yet its integration into autonomous decision-making systems remains largely unexplored. This paper addresses the need for a unified framework that systematically combines quantum computation with agent-based artificial intelligence. We examine how quantum technologies can enhance the capabilities of autonomous agents and, conversely, how agentic AI can support the advancement of quantum systems. We analyze both directions of this synergy and present conceptual and technical foundations for future quantum–agentic platforms. Our work introduces a formal definition of quantum agents and outlines architectures that integrate quantum computing with agent-based systems. As concrete proof-of-concept implementations, we develop and evaluate three quantum agent prototypes: (i) a Grover-based decision agent for quantum search-driven action selection, (ii) a variational quantum reinforcement learning agent for adaptive policy learning in a multi-armed bandit setting, and (iii) an adaptive quantum image encryption agent that autonomously selects encryption strategies based on entropy-driven feedback. These prototypes demonstrate practical realizations of quantum agency in decision-making, learning, and security contexts under NISQ-era constraints. Furthermore, we discuss application domains including quantum-enhanced optimization, hybrid quantum–classical orchestration, autonomous quantum workflow management, and secure quantum information processing. By bridging these fields, we introduce a structured theoretical and architectural framework for quantum–agentic systems, providing formal definitions, system models, and early operational prototypes that illustrate the feasibility of quantum-enhanced agency under NISQ constraints.



Academic Editors: Fahad Ahmad and Vincenzo Tamma

Received: 29 January 2026

Revised: 05 March 2026

Accepted: 10 March 2026

Published: 13 March 2026

**Copyright:** © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\) license](https://creativecommons.org/licenses/by/4.0/).

**Keywords:** quantum agents; quantum ML; agentic AI

## 1. Introduction

The convergence of quantum computing and agentic AI raises two fundamental systems questions: how can quantum resources be systematically embedded into autonomous decision-making architectures and, conversely, how can agentic control improve the scalability of quantum systems?

While prior research demonstrates isolated advances in quantum-enhanced learning, optimization, and orchestration, existing work remains fragmented across domains. A

coherent formal definition, architectural perspective, and control taxonomy for quantum–agentic systems is currently lacking.

Current agentic AI systems are fundamentally classical in their internal decision mechanisms. While they may invoke quantum subroutines as external tools, there exists no formalized framework for embedding quantum computational primitives directly into perception–decision–action loops in a principled and architecturally consistent manner. As a result, quantum resources are typically treated as isolated accelerators rather than as integrated components of autonomous systems. This creates a capability gap between (i) quantum algorithms developed in isolation and (ii) agentic architectures that require structured control, orchestration, and run-time adaptation of quantum processes.

To address this gap, we propose a formal definition of *quantum agents*, grounded in both quantum information theory and AI agency principles. We then introduce a set of conceptual architectures for quantum–agentic platforms that integrate quantum processors with agent-based reasoning frameworks.

## 2. Core Contributions and Novelty

This work goes beyond a conceptual survey of quantum AI and agentic systems by introducing original theoretical, architectural, and experimental contributions. Specifically, the novelty of this paper consists of:

- A formal agent-theoretic definition of *quantum agents* based on a unified tuple formalism  $(\mathcal{Q}, \mathcal{C}, \mathcal{M}, \mathcal{P}, \mathcal{A})$ , enabling systematic classification and comparison of quantum–agentic systems.
- A structured maturity model for quantum agents aligned with quantum hardware evolution, integrating post-quantum security, hybrid quantum learning, quantum-native cognition, and fully quantum–autonomous agency.
- A bidirectional agency framework distinguishing *quantum-enhanced agency* and *agent-enabled quantum systems*, providing a unified conceptual model for control flow and decision authority.
- A modular reference architecture for quantum–agentic platforms, connecting agentic AI stacks with quantum computation, sensing, memory, and orchestration layers.
- Three implemented quantum-agent prototypes that operationalize the theory in the NISQ regime, demonstrating search-based decision-making, learning-based policy formation, and adaptive quantum encryption as concrete instantiations of quantum agency.

These contributions position the paper as a structured framework for designing and classifying emerging quantum–agentic systems.

## 3. Background and Related Work

The convergence of quantum computing and artificial intelligence (AI) has catalyzed the emergence of *quantum agents*—autonomous systems that leverage quantum resources to enhance decision-making, learning, and interaction capabilities. This section reviews seminal and contemporary works that have shaped the quantum agents’ conceptual and practical development.

### 3.1. Conceptual Models of Quantum Agents

Early conceptualizations of quantum agents laid the groundwork for integrating quantum computation into agent-based systems. Klusch’s Quantum Computational Agents (QCA) model [1] extended traditional intelligent agents by incorporating quantum processing capabilities, enabling functionalities such as quantum addressing and the execution of

quantum instructions within a master–slave architecture. This pioneering work highlighted the potential of hybrid quantum–classical systems in agent design.

### 3.2. Quantum Reinforcement Learning and Speedup

Empirical studies have demonstrated the advantages of quantum resources in reinforcement learning (RL). Saggio et al. [2] implemented a learning protocol using a programmable integrated nanophotonic processor, where the agent’s interaction with the environment was mediated through quantum channels. The results indicated a systematic quantum advantage, with the agent achieving faster learning compared to classical counterparts. This experiment underscores the potential of quantum communication to enhance learning efficiency in agentic systems.

### 3.3. Memory Efficiency and Energetic Advantages

Quantum processing can confer significant benefits in memory efficiency and energy consumption for adaptive agents. Elliott et al. [3] introduced a framework where quantum agents could encode and compress historical information more efficiently than classical agents, particularly in non-Markovian environments (environments where the future state depends on past states, breaking the memoryless Markov assumption.) requiring long-term memory retention. Additionally, Thompson et al. [4] explored the thermodynamic implications of quantum processing in agents executing complex, adaptive strategies, demonstrating that quantum agents could achieve lower energy dissipation during decision-making processes compared to classical counterparts.

### 3.4. Advancements in Quantum Multi-Agent Reinforcement Learning

Recent research has focused on scaling quantum reinforcement learning to multi-agent systems. Yun et al. [5] proposed a Quantum Multi-Agent Reinforcement Learning (QMARL) framework utilizing variational quantum circuits, enabling centralized training with decentralized execution. This approach addresses challenges in scalability and coordination among multiple quantum agents, demonstrating improved performance over classical multi-agent reinforcement learning methods.

### 3.5. Comprehensive Surveys and Future Directions

Comprehensive surveys have synthesized the current state and future prospects of quantum artificial intelligence. For instance, the survey by [6] provides an overview of achievements in quantum AI, highlighting the intersection of quantum computing and AI, and pointing to open questions for future research. These works emphasize the importance of developing standardized benchmarks, scalable architectures, and robust evaluation metrics for quantum agents.

### 3.6. Agentic Quantum Computing at Kipu Quantum

A recent industrial perspective on the convergence of quantum computing and agent-based AI is provided by Kipu Quantum, spearheaded by Solano [7]. Their approach, termed *Agentic Quantum Computing* (AQC), emphasizes the bidirectional enhancement of AI and quantum computing: quantum computing augments reasoning and agency in AI, while AI—especially generative models and agent frameworks—accelerates the application readiness of quantum algorithms.

Kipu Quantum’s architecture combines generative AI models (such as transformers and large language models) with quantum-advantage algorithms within the PLANQK platform. Their agents coordinate classical and quantum resources, executing hybrid algorithms on commercial quantum hardware from IBM, D-Wave, QuEra, and others. These

systems aim to solve real-world industrial problems such as combinatorial optimization, material design, and data-driven classification tasks.

One of Kipu's most notable innovations is *ChatQPT*, an agent designed to assess whether a natural-language prompt requires quantum-advantage execution—specifically for higher-order unconstrained binary optimization (HUBO) problems. This leads to a novel metric: the *averaged k-local HUBO quantum-advantage threshold*, which aims to define the boundary between classical and quantum superintelligence in industrial reasoning contexts.

Kipu Quantum thus provides a commercial and application-driven angle to quantum agents, emphasizing practical deployment, hybrid integration, and user-centric agentic interfaces for quantum problem-solving.

While existing research demonstrates isolated advantages of quantum-enhanced learning, memory, optimization, and orchestration, these works remain largely fragmented across domains. In the following section, we move from literature synthesis to a system-level perspective and analyze how quantum computation can be systematically embedded into agentic architectures to enhance autonomous decision-making.

## 4. From Quantum to Agents: Enhancing Agency with Quantum Computing

Quantum computing offers unique advantages for solving problems that are computationally intensive or intractable for classical systems. By integrating quantum capabilities into agentic architectures, we can significantly expand the scope, speed, and intelligence of autonomous systems. In this section, we explore three key areas where quantum computing can enhance agency: search and optimization, reinforcement learning, and simulation-based decision-making.

### 4.1. Quantum Search and Optimization in Agent Systems

Many agentic tasks, such as complex scheduling, high-dimensional resource allocation, or constraint satisfaction under uncertainty, require efficient search and optimization. Classical agents often rely on heuristic or approximate methods due to the combinatorial explosion of possible solutions. Quantum computing introduces new paradigms that offer potential speedup and performance gains in these domains.

Grover's algorithm [8] and the family of Grover-based search approaches [9–11], for example, may offer a quadratic speedup for unstructured search problems, potentially allowing agents to locate target solutions with fewer evaluations. More generally, Quantum Approximate Optimization Algorithms (QAOA) [12] and their variants [13], along with quantum annealing techniques [14], have been proposed as candidates to address complex combinatorial optimization tasks, such as scheduling, vehicle routing, or energy distribution. Although the practical utility of these algorithms remains under active investigation, especially in the noisy intermediate-scale quantum (NISQ) era, hybrid quantum-classical search algorithms [15] are increasingly being seen as a scalable path forward for near-term applications.

By integrating quantum search modules into agentic reasoning cycles, agents can evaluate and refine action plans more effectively—especially in dynamic or multi-agent environments where decision spaces are large and rapidly evolving.

### 4.2. Quantum Reinforcement Learning

Reinforcement learning (RL) is central to agentic behavior, where agents learn optimal policies through trial and error. Quantum reinforcement learning (QRL) explores how quantum computation can accelerate or generalize the RL paradigm.

In model-free QRL, quantum-enhanced policy evaluation and value function estimation can improve learning efficiency, particularly in environments with high-dimensional or continuous state spaces. Hybrid architectures combine classical perception and actuation with quantum components for exploration and reward evaluation, enabling agents to converge faster to optimal strategies.

Model-based QRL further leverages quantum simulation (discussed below) to construct and evaluate internal models of the environment, supporting better foresight and planning. Recent research also explores quantum-inspired RL algorithms that use entanglement and superposition to encode exploration strategies not easily achievable classically.

These advancements make QRL a promising direction for agents operating in environments that demand fast adaptation and strategic decision-making.

#### 4.3. Quantum Simulation for Decision-Making Agents

Decision-making agents often require internal simulations of their environment to predict outcomes and assess potential actions. Quantum simulation provides a powerful framework for modeling complex physical systems, including those governed by quantum mechanics, non-linear dynamics, or stochastic processes.

For agents in domains such as chemistry, materials science, or quantum control, quantum simulators offer access to predictive models that are infeasible to compute classically. This allows agents to reason more effectively about molecular structures, reaction pathways, or sensor responses at the quantum level.

Moreover, quantum simulation enables probabilistic sampling of multiple futures in parallel. Agents can use these simulated outcomes to perform robust scenario analysis, optimize under uncertainty, or select actions with the highest expected utility.

In sum, quantum simulation might enhance the cognitive depth of agents, empowering them to understand and navigate environments with emerging complexity.

The integration of quantum computation into agentic reasoning expands the cognitive and operational capabilities of autonomous systems. However, the relationship is not unidirectional. As quantum systems themselves grow in complexity, they increasingly require intelligent coordination, control, and optimization. The following section therefore inverts the perspective and examines how agentic AI becomes an enabling infrastructure for scalable quantum computing.

#### 4.4. Conditions for Quantum Advantage in Agentic Architectures

While quantum algorithms such as Grover's search and variational quantum circuits offer theoretical speedups or representational advantages, practical performance benefits depend on several non-trivial conditions.

First, asymptotic speedups (e.g., quadratic improvements in unstructured search) become relevant only for sufficiently large problem sizes. In small decision spaces, such as the minimal prototypes presented in this work, classical enumeration remains competitive or superior due to quantum overhead and measurement noise.

Second, quantum reinforcement learning approaches may provide representational advantages in high-dimensional probability encoding. However, empirical superiority over classical deep reinforcement learning has not yet been consistently demonstrated under current NISQ hardware constraints.

Third, hybrid orchestration benefits arise not primarily from raw computational speed, but from structural integration: the ability to dynamically select, parameterize, and coordinate quantum subroutines within autonomous control loops.

Accordingly, the primary contribution of this work is not to claim immediate empirical superiority, but to formalize architectural conditions under which quantum-enhanced agency could become advantageous as hardware matures.

## 5. From Agents to Quantum: Enabling Quantum Systems Through Agency

As quantum technologies mature, their complexity and operational demands grow rapidly. Agentic AI offers a scalable and intelligent interface to manage, automate, and optimize quantum systems. In this section, we explore how agents can support and accelerate the deployment of quantum computing—from managing intricate workflows to designing circuits and orchestrating hybrid systems.

### 5.1. AI Agents for Quantum Workflow Management

Operating quantum systems involves multiple interdependent tasks: calibration, scheduling, execution, data handling, error tracking, and result interpretation. These tasks are typically performed manually or through rigid scripts, making the process inefficient and error-prone. AI agents can bring autonomy and intelligence to quantum workflow management.

A unique offer of agentic systems would be preserving the context of jobs along the stack (like properties of problems in question, algorithms used and hardware parameters) to make the most optimal and coherent choices. For example, such agents could be used to monitor system status and trigger recalibrations when performance metrics degrade, adaptively schedule quantum jobs between different QPUs based on hardware availability and real-time properties, manage queue priorities for cloud users, and so on.

By learning from historical data, agents can anticipate bottlenecks or failures and recommend preemptive actions—such as choice of algorithms, error-mitigation routines, mappings, etc. In multi-user or cloud-based quantum environments, agents act as coordinators that balance load, enforce policies, and ensure optimal utilization of scarce quantum resources.

Integrating agents into quantum software stacks (for example, Qiskit, Cirq and PennyLane) would enable dynamic, context-aware control over end-to-end quantum execution pipelines.

### 5.2. Automated Quantum Circuit Design and Optimization

Designing quantum circuits is a non-trivial task, often requiring deep expertise in quantum logic, gate synthesis, and noise-aware compilation. AI agents equipped with domain knowledge and optimization capabilities can significantly streamline this process.

Agents can generate circuit templates based on high-level problem specifications, apply transformations to reduce gate counts, and adapt circuits to specific hardware constraints (such as with qubit connectivity and coherence times). For example, selecting the optimal quantum circuit or ansatz for a given problem is a pivotal aspect of designing effective variational quantum algorithms (VQAs). The choice of ansatz influences the algorithm's expressibility, trainability, and compatibility with quantum hardware. However, the choice of quantum circuit ansatzes, detailing their intent, applicability, circuit diagrams, and implementation needs considerable effort [16,17]. This complexity exists for other algorithms like Hamiltonian Variational Ansatz (HVA), inspired by the Quantum Approximate Optimization Algorithm (QAOA) and adiabatic quantum computation [18], therefore, agents can pass the intended use case and design goals of a quantum circuit to Model Context Protocol (MCP) servers, which guide reinforcement learning or genetic algorithms. By incorporating this contextual intent, the system can iteratively refine circuits,

optimizing for performance metrics such as fidelity, circuit depth, and execution time which ensures a reduction in the effort needed in the current NISQ era.

Moreover, agents can explore novel circuit architectures that may be unintuitive to human designers by leveraging techniques such as quantum architecture search and symbolic reasoning. As quantum algorithms become increasingly complex and tailored to specific applications, agent-based automation becomes essential not only to accelerate development, but also to broaden accessibility. In this context, “non-experts” are individuals who have deep expertise in their own fields, such as chemists, materials scientists, or machine learning researchers, but who do not possess detailed knowledge of quantum circuit design, quantum error correction, or low-level quantum hardware constraints. These users may understand the applications of quantum computing in broad terms but lack the technical background to develop or optimize quantum circuits directly. By bridging this gap, intelligent agent systems can democratize quantum computing resources, enabling these domain specialists to deploy and benefit from quantum algorithms without the need to become quantum computing experts themselves.

One of the most promising areas, where Quantum Computers are expected to bring an advantage in the near term is quantum chemistry [19], following the advice of Richard Feynman to simulate nature by leveraging quantum mechanical principles [20]. As an example, imagine agents trained to detect cases with strong electron–electron correlation, where even standard classical methods which go beyond mean-field approximations, like Coupled Cluster with Singles, Doubles, and Perturbative Triples, CCSD(T), fail [21]. Here, agents could help [22] to automatically construct a Hilbert subspace around regions of concern in molecules, extract an effective Hamiltonian, qubit map it, decide which quantum algorithm is the most suitable, and finally design a quantum circuit that can be executed on most suitable quantum hardware platform.

### 5.3. Orchestration of Hybrid Quantum–Classical Systems

Most near-term quantum applications follow a hybrid model, combining classical computing with quantum subroutines. Examples include variational quantum algorithms, quantum machine learning workflows, and hybrid solvers. Orchestrating these workflows requires intelligent coordination of classical and quantum resources, real-time decision-making, and data-dependent branching logic.

AI agents are ideally suited to this role. They can manage hybrid task graphs, optimize data flow between classical and quantum modules, and make run-time adjustments based on intermediate results or hardware conditions. For instance, in a variational quantum eigensolver (VQE), an agent may adapt the classical optimizer strategy depending on the convergence rate or noise profile observed during quantum evaluations.

Agents also enable fault-aware orchestration, where alternative execution paths or fallback strategies are selected dynamically if a quantum task fails or underperforms. In distributed settings, agents can coordinate across multiple backends, selecting the best available quantum hardware for each task segment.

Through such orchestration, AI agents unlock practical performance gains and make hybrid quantum–classical workflows robust, adaptive, and efficient.

## 6. Defining Quantum Agents

As the fields of quantum computing and agentic AI begin to converge, a clear and rigorous definition of *quantum agents* becomes essential. Such a definition provides a foundation for system design, benchmarking, and theoretical exploration. In this section, we introduce the core properties of quantum agents, provide a formal definition, and

contrast them with classical agents. We also outline key design criteria for developing effective quantum–agentic systems.

### 6.1. Anatomy of a Quantum Agent

The classical agent paradigm is typically organized into three core components: *perception*, *processing*, and *action*. This structure provides a natural blueprint for the design of quantum agents. In the quantum context, each component is extended or reinterpreted to exploit the capabilities of quantum information processing.

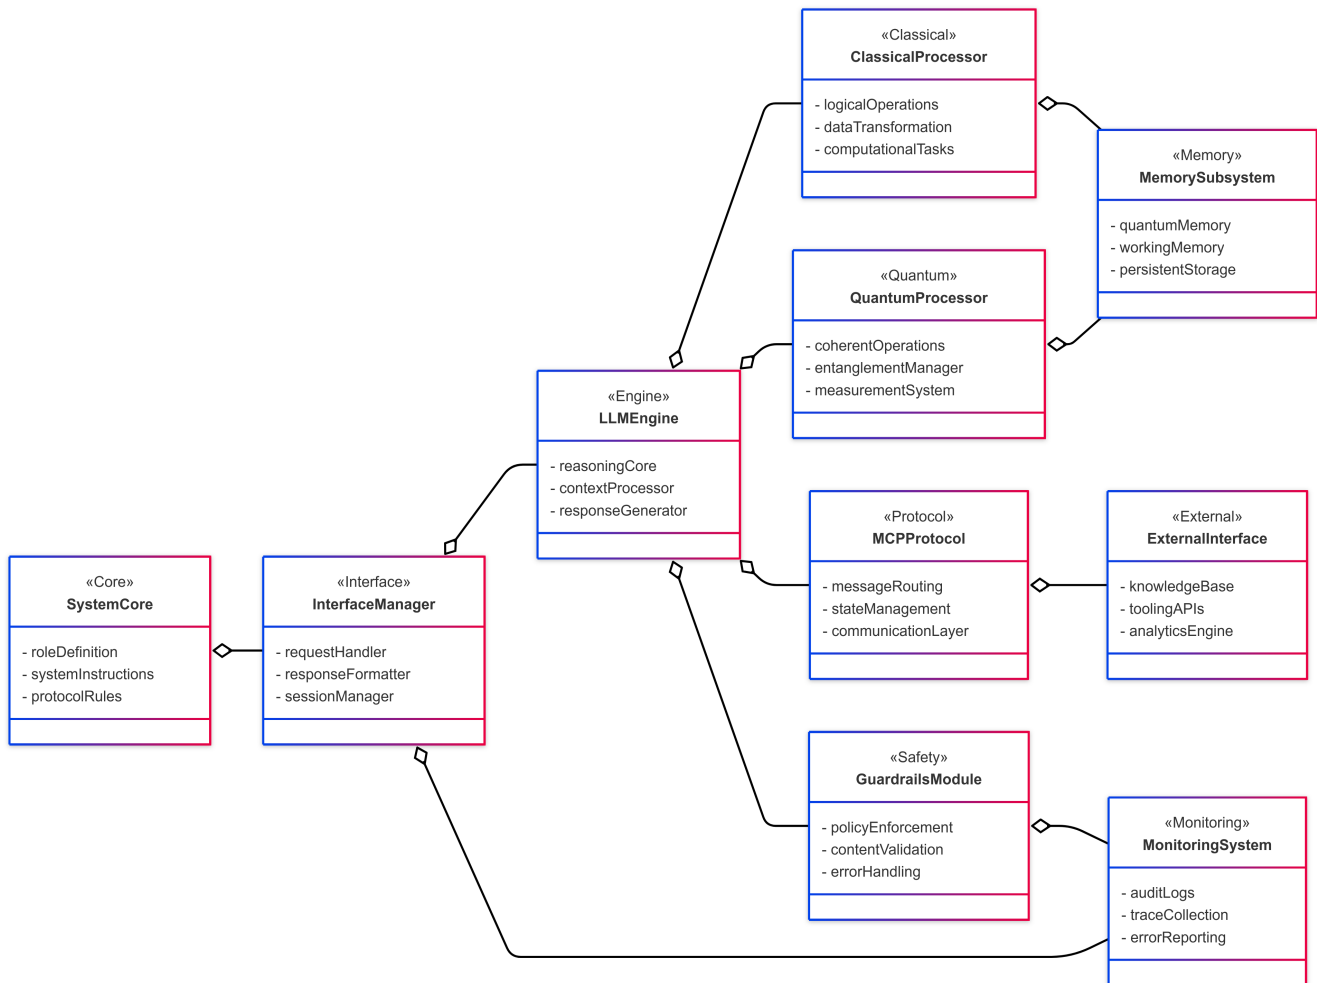
**Perception:** A quantum agent perceives its environment through both classical and quantum channels. Perception may include data streams from classical sensors, quantum measurements (e.g., projective measurements or POVMs), or direct access to quantum states via entangled inputs or quantum sensor arrays. For instance, a quantum agent operating in a lab environment may receive raw photonic qubit states, measurement statistics, or high-resolution data from quantum-enhanced magnetometers. Efficient encoding of this perceptual data into quantum memory or as parameterized input states is a crucial step toward enabling downstream quantum processing.

**Processing:** The reasoning and decision-making core of a quantum agent integrates quantum algorithms with classical control logic. This may involve hybrid quantum–classical cycles, where the agent uses variational algorithms (e.g., QAOA, VQE) to solve subproblems and reinforcement learning schemes to adapt policies. Quantum circuits may be dynamically constructed based on perceptual input, and optimized on-the-fly using meta-learning or neural architecture search. Processing tasks could include a high-dimensional search, probabilistic inference using quantum sampling, or policy updates based on quantum-advantage estimation. The agent must also manage decoherence times, noise models, and backend availability, orchestrated via an intelligent run-time scheduler. For example, an agent receiving high-dimensional sensor data may encode relevant features into a parameterized quantum state, invoke a variational quantum circuit to estimate an optimal action distribution, and then use the measured probabilities to update its decision policy within a classical reinforcement learning loop.

**Action:** The action component allows the quantum agent to interact with its environment. This includes initiating quantum operations (e.g., circuit execution on remote QPUs), controlling experimental setups (e.g., qubit calibration or laser tuning), or triggering communication protocols such as quantum teleportation or QKD. In multi-agent settings, actions may involve quantum communication with other agents or quantum resource negotiation. Actions may be represented as quantum gates, classical signals, or composite operations involving both digital and analog quantum interfaces.

Taken together, this architecture frames the quantum agent as an autonomous system that senses, reasons, and acts across both classical and quantum dimensions. Unlike traditional AI agents, quantum agents operate under constraints such as qubit decoherence, probabilistic measurement outcomes, and non-cloning, while gaining access to fundamentally new capabilities through superposition, entanglement, and quantum parallelism.

Figure 1 shows a Quantum Agent System Architecture designed as a modular framework combining classical and quantum components to enable agent behavior that can be checked and understood. The system is based on an interface manager that controls interactions between the main modules, external tools, and internal processes.



**Figure 1.** The anatomy of a Quantum Agent System Architecture: a modular framework combining classical logic, quantum operations, safety mechanisms, and external interfaces for intelligent, auditable agent behavior.

Classical tasks use the Model Context Protocol (MCP) for compatibility with existing AI systems. Quantum tasks run on the quantum processing unit (QPU), which handles complex calculations. A Knowledge Base provides context to support decision-making, while monitoring and auditing tools track agent behavior. Several protection layers are included to ensure reliability in critical applications.

This architecture is expected to evolve over time, following the quantum agent maturity model shown in Figure 2. The model has four levels that reflect improvements in quantum hardware and algorithms. Early levels include agents that combine near-term quantum technology with classical processes. Later levels focus on integrating quantum processing more deeply into decision-making. Each of these levels shows the growth of component's capabilities, example being Perception L1 vs. Perception L4. The subsequent sections will outline the specifics of the evolutionary process from the NISQ level to the Fully Quantum-Native Agent, and how the architectural components specified in Figure 1 will evolve. To make this evolution explicit, the architectural modules shown in Figure 2 correspond directly to the formal tuple components (Q, C, M, P, A) introduced below. Figure 3 describes how the overall agent capability progresses across maturity levels. In this way, the maturity model operationalizes the structural blueprint defined in Figure 1.

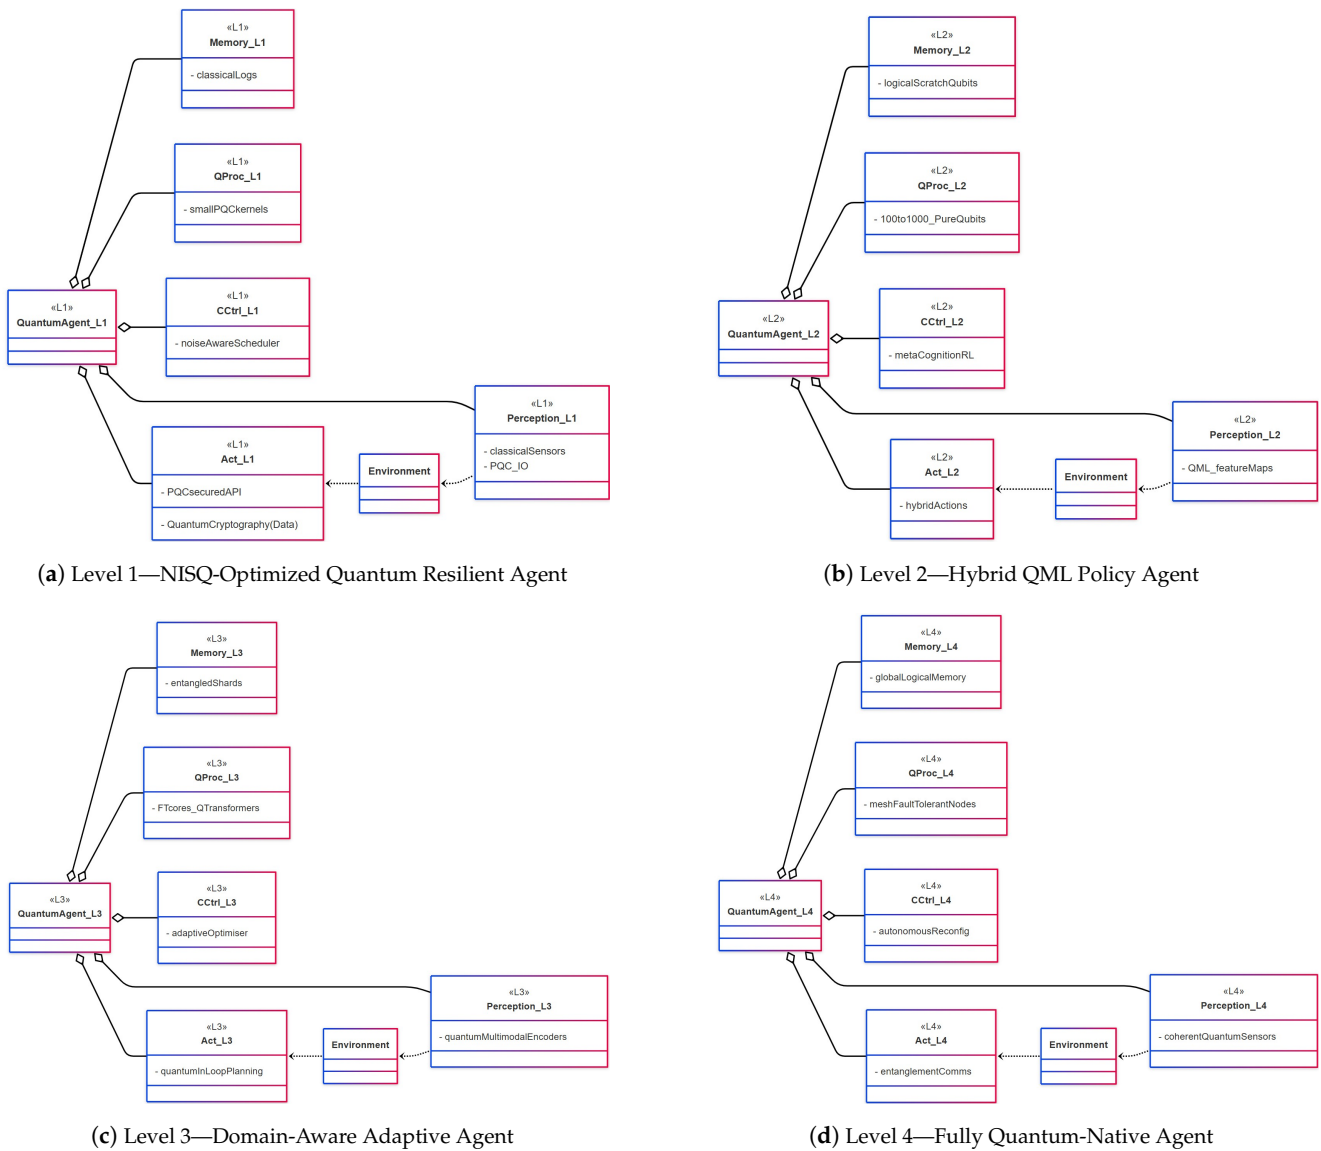


Figure 2. The anatomy of a quantum agent based on its maturity model.

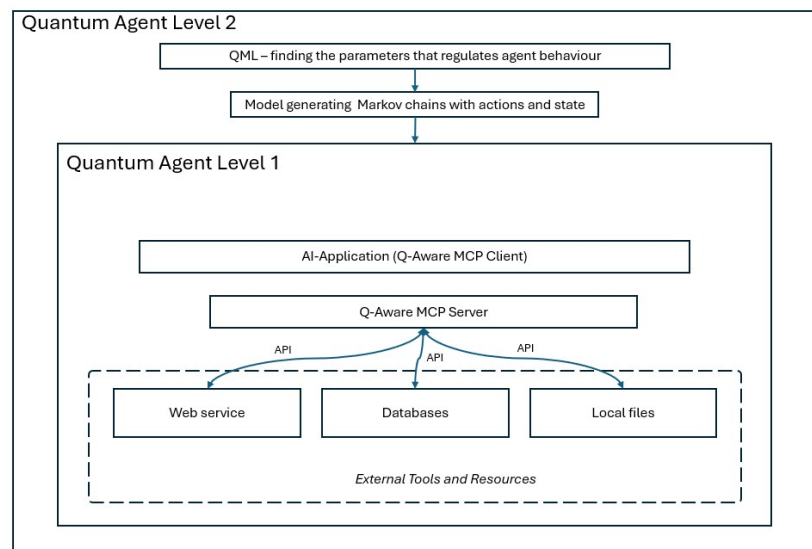


Figure 3. Quantum Agent—Level 2.

### 6.2. Core Properties and Formal Definition

A quantum agent is an autonomous system that integrates quantum computation or quantum information processing into its perception–decision–action loop. Unlike conventional agents, quantum agents possess access to quantum resources which, based on the quantum mechanical principles of superposition and entanglement, can make use of quantum parallelism to enhance or enable their cognitive and operational abilities.

We define a **quantum agent** as follows:

A *quantum agent* is an autonomous system characterized by the tuple  $(Q, C, M, P, A)$ , where:

- $Q$  is a set of quantum processing resources or devices (e.g., quantum processors, annealers, or simulators), along with the associated quantum assembly-based instructions or circuits used to operate them.
- $C$  is the classical control logic interfacing with  $Q$ .
- $M$  is a hybrid memory subsystem comprising classical memory and, where applicable, quantum memory for intermediate state storage, subject to quantum mechanical constraints such as the no-cloning theorem and typically requires careful control to preserve coherence and enable operations like teleportation, buffering, or delayed measurement.
- $P$  is a perception module receiving quantum or classical input from the environment.
- $A$  is an action module executing outputs, possibly involving quantum communication or control.

The defining characteristic of a quantum agent is that its reasoning, decision-making, learning, or sensing processes are enhanced—or made possible—by quantum operations that provide non-classical computational or informational advantages.

### 6.3. Comparison with Classical Agents

Quantum agents differ from classical agents not only in the hardware they use, but also in the nature of the computations they perform and the kinds of problems they can solve more efficiently.

Key distinctions include:

- **Computation:** Classical agents rely on Boolean logic and deterministic or probabilistic Turing machines; quantum agents utilize unitary transformations and quantum probability amplitudes.
- **Learning and Adaptation:** Quantum agents can exploit quantum-enhanced learning techniques, such as amplitude amplification in exploration or quantum kernel methods in classification.
- **Memory and Representation:** Quantum agents may store and manipulate information using quantum bits (qubits), allowing exponential state space representation compared to classical bitstrings.
- **Interaction with Environment:** Quantum agents may perform or interpret quantum measurements, interact with quantum environments (such as quantum sensors or systems), or communicate using quantum channels.

These differences imply that quantum agents may be suitable for tasks that are infeasible for classical agents, particularly in domains involving quantum physics, a high-dimensional search, or (time-sensitive) complex optimization.

### 6.4. Design Criteria for Quantum Agentic Systems

Building quantum agents involves careful architectural and algorithmic choices to ensure functional, scalable, and efficient systems. We propose the following design criteria:

1. **Quantum–Classical Integration:** Seamless communication between classical and quantum modules is essential. This includes latency-aware data exchange, co-scheduling, and error propagation handling.
2. **Modularity and Abstraction:** System components (such as quantum solvers, memory and sensors) should follow modular interfaces to allow flexible reuse and abstraction over underlying hardware.
3. **Resource Awareness:** Quantum agents must reason under constraints such as limited qubit count, coherence time, or access to remote QPUs, and adapt their strategies accordingly.
4. **Explainability and Debugging:** Given the non-intuitive nature of quantum processes, developers require tools to inspect, trace, and explain the behavior of quantum agents in interpretable terms.
5. **Scalability and Learning:** Quantum agents should generalize from experience and scale across increasingly complex environments, ideally through reinforcement learning or meta-learning mechanisms.
6. **Security and Trustworthiness:** Quantum agents interacting in open or adversarial environments must implement secure protocols and demonstrate verifiable behavior.

These criteria form the basis for the principled development of quantum–agentic platforms and guide future experimentation and deployment across real-world domains.

## 7. Control Paradigms in Quantum–Agentic Systems

In agentic quantum systems, the technical role of quantum computing can vary substantially depending on system architecture. In this section, we distinguish two fundamental operational modes: (1) *quantum-assisted agency*, in which the agent uses quantum resources to augment perception, reasoning, or action, and (2) *quantum-centric control*, where the quantum layer itself performs decision-making and directs agentic behavior. Understanding this distinction is crucial for designing intelligent architectures with meaningful autonomy and performance gains.

### 7.1. Quantum-Assisted Agency: Agents Using Quantum as a Subsystem

In this mode, the agent maintains control over the decision loop and uses quantum subsystems as specialized computational accelerators or perception modules. Technically, this involves:

- **Quantum Perception:** The agent invokes quantum sensing routines or performs quantum measurements to extract information from the environment. This may involve querying entangled states, performing quantum state discrimination, or executing variational sensors. The perception output is typically a classical summary (e.g., expectation values, sampled outcomes) passed to the agent’s decision core.
- **Quantum Reasoning:** The agent offloads computationally hard subroutines—such as combinatorial optimization, quantum sampling, or inference—to quantum processing units (QPUs). This is implemented via calls to cloud-accessible QPUs, variational quantum solvers (e.g., QAOA), or annealing-based optimization engines.
- **Quantum Action:** The agent may trigger quantum communication (e.g., QKD, entanglement distribution), perform quantum control actions (e.g., dynamic reconfiguration of a quantum network), or command quantum mechanical actuators in experimental setups.

Here, quantum functionality is modularized and abstracted behind APIs, with orchestration and interpretation remaining within the classical agent core. This mode offers

flexibility, transparency, and modular upgradability, but depends on the agent's ability to model and manage quantum behavior effectively.

### 7.2. Quantum-Centric Control: Quantum Mesh as Decision Substrate

In contrast, some architectures position quantum systems—such as entangled meshes or variational quantum processors—as autonomous or semi-autonomous decision substrates. In this case, the “agency” is partially embedded in the quantum system itself. This raises profound technical and philosophical implications.

- **Distributed Quantum Decision-Making:** In a quantum mesh or network of entangled nodes, decision-making can emerge from the structure and evolution of the entangled state itself. Quantum correlations are exploited to coordinate decentralized behavior, potentially without explicit classical communication. This could support quantum swarm intelligence, distributed control, or adaptive behavior in physical systems.
- **Quantum-Intrinsic Policies:** In variational quantum agents or quantum reservoir systems, policy functions may be implicitly encoded in a quantum state and updated through feedback and training. These systems blur the line between inference and evolution, as the decision is made not by a separate agent, but as a result of controlled quantum dynamics.
- **Measurement-Driven Action:** In some architectures, actions are determined directly by the measurement outcomes of quantum processes, e.g., collapse-based decision-making. Here, measurement results trigger state transitions in actuators or communication protocols without classical post-processing.

### 7.3. Blurring the Boundary: Hybrid Control and Shared Agency

In practice, many quantum–agentic systems will operate in a hybrid regime where agency is distributed across classical and quantum components. Rather than assigning control to a single substrate, decision-making becomes a collaborative process shaped by feedback loops between symbolic reasoning and quantum dynamics.

**Shared agency** emerges when classical agents define high-level goals or constraints, while quantum systems explore, evaluate, or realize possible solutions. The classical agent may query a quantum variational circuit to generate policy proposals, evaluate them based on non-quantum criteria (e.g., cost, ethics, or constraints), and then update the quantum system via parameter tuning or measurement selection. Conversely, the quantum layer may generate unexpected behaviors or correlations that lead the classical agent to revise its model of the environment.

Technically, hybrid agency often involves:

- **Co-dependent control loops**, where both classical and quantum layers exchange information asynchronously or in real time.
- **Meta-level orchestration**, where an outer agent coordinates the behavior of nested agents or modules—some of which may be quantum-enhanced or quantum-native.
- **Dual learning layers**, where reinforcement learning occurs simultaneously in both classical and quantum models, possibly with cross-adaptation.

Hybrid agency enables practical integration of quantum capabilities while preserving interpretability, robustness, and controllability in real-world quantum–agentic systems.

## 8. Maturity Model in Definition of Quantum Agents

The proposed maturity model is intended as a conceptual alignment framework rather than a predictive benchmark scale. In this work, the maturity discussion primarily serves to contextualize the architectural definition and the implemented NISQ-era prototypes, rather than to forecast near-term deployment of higher maturity levels. Its purpose is

to map architectural characteristics of quantum agents to broadly anticipated stages of quantum hardware development, as outlined in public technology roadmaps. The levels do not claim empirical validation or performance superiority, but instead provide a structured taxonomy for reasoning about increasing integration of quantum resources into agentic systems.

The growth of agentic AI must be aligned with the evolution of quantum hardware. Accordingly, we adopt the US Department of Energy’s quantum development timeline [23] as the primary reference for calibrating the maturity levels of our quantum agent model against projected hardware capabilities. The timeline and breakthroughs required to progress quantum agents through maturity levels are given in Table 1, whereas, Table 2 presents the threat categories across lifecycle phases for tool-using agents, showing vulnerabilities in creation, operation, updates, and general security.

**Table 1.** Timeline and breakthroughs required to progress quantum agents through maturity levels, aligned with eras in quantum computing.

Quantum Computing Era	Timeframe	Breakthroughs Needed	Quantum Agent Maturity Model Level
ine NISQ Devices and Quantum Error Correction (QEC) Demos	0–5 y	<ul style="list-style-type: none"> <li>• Demonstration of 10× suppression in logical error rates</li> <li>• Early implementations of QEC</li> <li>• Efficient near-term algorithms and error mitigation</li> <li>• Hardware-aware algorithm development</li> </ul>	Level 1: NISQ-Optimized Decision Agents
ine Small Error-Corrected Quantum Computers	5–10 y	<ul style="list-style-type: none"> <li>• Validation of logical DiVincenzo criteria</li> <li>• Development of quantum interconnects</li> <li>• Scale to 1000+ physical qubits below error threshold</li> <li>• Mid-circuit readout and low-latency VQC-gradient co-processor</li> </ul>	Level 2: Hybrid QML Policy Agents
ine Large Fault-Tolerant Quantum Computers	10–20 y	<ul style="list-style-type: none"> <li>• Scaling to 10,000+ physical qubits</li> <li>• Secure quantum co-processors</li> <li>• In-circuit quantum memory recall</li> <li>• Quantum–classical transformer compiler</li> </ul>	Level 3: Domain-Aware Adaptive Agents
ine Very Large Fault-Tolerant Systems	20+ y	<ul style="list-style-type: none"> <li>• Negligible logical error at any depth</li> <li>• Fully autonomous quantum systems with cross-node entanglement</li> <li>• Self-optimizing quantum clouds</li> </ul>	Level 4: Fully Quantum-Native Agent

We suggest interpreting the term “Quantum” here in its several faces: post-quantum cryptography, hybrid quantum–classical computing, quantum cryptography, and quantum resiliency.

The first level, ad hoc quantum agent, is a normal agent that is ready for Q-Day and, therefore, uses quantum resilient cryptography, either by applying post-quantum cryptography (PQC) for its duties when needed (e.g., PQC-TLS or quantum cryptography) or other approaches like what we proposed in our prototypes. Currently, the Model Context Protocol (MCP) is the primary communication protocol used by agents. However, it has the following security issues related to data protection and privacy [24]. A *Quantum Agent–Level 1* is therefore defined as an agent that uses a security-migrated MCP protocol to mitigate these issues and applies post-quantum cryptography (PQC) for compliance—or is capable of handling encrypted data, as demonstrated by Agent 3 in Section 11.

**Table 2.** Threat categories across lifecycle phases for tool-using agents, showing vulnerabilities in creation, operation, updates, and general security.

Lifecycle Phase	Threat Category	Description
Creation	Name Collision	Malicious servers can register with deceptive names to impersonate legitimate ones and intercept sensitive data.
Creation	Installer Spoofing	Unofficial installers might include malware or backdoors, compromising the user environment during setup.
Creation	Code Injection/Backdoors	Malicious code may be hidden in source code or dependencies, enabling persistent unauthorized access.
Operation	Tool Name Conflicts	Conflicting or malicious tool names can cause AI to invoke the wrong tool, leading to data leaks or unauthorized actions.
Operation	Slash Command Overlap	Duplicate command names (such as <code>/delete</code> ) across tools which may trigger unsafe operations or data loss.
Operation	Sandbox Escape	Poor sandboxing may allow tools to access host systems or data beyond their intended scope.
Update	Post-Update Privilege Persistence	Outdated privileges or tokens may remain active after updates, enabling unauthorized access.
Update	Re-deployment of Vulnerable Versions	Users may unintentionally install outdated versions with known vulnerabilities.
Update	Configuration Drift	Inconsistent configuration changes can introduce data exposure or excessive access rights.
General	Lack of Central Security Oversight	No unified platform for enforcing security policies leads to fragmented protection.
General	AuthN/AuthZ Gaps	Missing standard authentication and authorization mechanisms increase risk in multi-tenant setups.
General	Insufficient Monitoring	Lack of robust logging and alerting makes it hard to detect misuse or breaches.

Moreover, at this level, the operational context dictates the threat model, risk profile, and choice of counter-measures, most notably the way quantum and post-quantum cryptography are applied. For example, when safeguarding medical images, the agent could invoke techniques such as Chaotic Quantum Encryption (CQE) [25], thereby meeting the compliance requirements of frameworks including HIPAA (USA), the UK GDPR/Data Protection Act 2018, and the German BDSG/GDPR/SGB V.

A *Quantum Agent–Level 2* uses hybrid quantum machine learning to evaluate the explainability of the behavior of agentic AI and its adherence to responsible AI [26]. The architectural realization of this Level 2 agent is illustrated in Figure 3, which depicts the hybrid meta-cognition and quantum learning layers operating in conjunction with classical control modules. Using a Markov Chain model, the meta-cognition layer continuously provides a real-time set of states, actions, and the maximum Q-value as input to the meta-cognition layer. This layer uses QML to generate various future scenarios faster than the AI itself and evaluates their adherence to human-aligned values, offering to regulate parameters to correct potential non-compliance in *Quantum Agent–Level 1*. At this level, agents intelligently distribute tasks between classical and quantum modules to maximize efficiency, speed, and performance.

Levels 3 and 4 describe long-term theoretical horizons aligned with fault-tolerant quantum computing and should not be interpreted as near-term engineering predictions.

A *Quantum Agent–Level 3* integrates Quantum Transformers [27] and history-dependent (non-Markovian) simulation (i.e., simulation of dynamics with memory effects beyond standard Markovian assumptions) directly into its reasoning and decision-making processes by leveraging quantum devices for real-time sequence modeling and large lan-

guage model (LLM) tasks [28]. It replaces or accelerates core components of the Transformer architecture using quantum linear algebra techniques such as block encoding (embedding a matrix into a larger unitary operator so it can be processed by a quantum circuit) and quantum singular value transformation (QSVT) (a framework for applying polynomial transformations to a matrix spectrum via unitary circuits) [29], or utilizes parameterized quantum circuits (PQC) for QKV (queries (Q), keys (K), and values (V)) generation and attention calculation. This allows the agent to simulate and evaluate more complex simulation scenarios [30–32], supporting deeper and faster introspection, conversational modeling, and autonomous planning. Unlike Level 2, which uses QML to evaluate and regulate classical agentic behavior, Level 3 executes a hybrid or fully quantum self-attention mechanism on quantum hardware to directly generate, transform, and reason over-tokenized knowledge [33]. By executing inference through quantum circuits, it minimizes classical computation bottlenecks. For instance, Chain-of-Thought reasoning and large-scale search operations often incur significant classical compute overhead [34]. While quantum amplitude amplification provides a quadratic speedup for unstructured search problems [35], its direct applicability to complex reasoning tasks—such as Chain-of-Thought processing in language models—is more limited. Chain-of-Thought involves structured, multi-step inferential processes that extend beyond a pure search, often requiring hierarchical and sequential computation. However, quantum search algorithms could potentially accelerate specific subroutines within broader reasoning workflows, such as database lookups or combinatorial problem-solving steps, thereby reducing some classical computational bottlenecks. To bridge this gap in practice, hybrid quantum–classical agent architectures could incorporate dynamic orchestration modules that identify and offload suitable subroutines to quantum processors. This approach would allow agents to selectively integrate quantum resources, ensuring quantum acceleration only where it aligns with actual performance gains. Fully realizing these advantages in hybrid quantum–classical agents remains an active area of research.

A conceptual *Quantum Agent–Level 4* represents a hypothetical quantum-native autonomous system operating on fully fault-tolerant hardware with scalable quantum memory and negligible logical error rates. Unlike its predecessors, Level 4 agents possess quantum-enhanced sensory capabilities—processing high-dimensional data such as vision, sound, and natural language through quantum self-attention, quantum convolutional networks, and kernelized quantum similarity metrics. At its core, the agent maintains a persistent quantum memory that stores and retrieves entangled perceptual and conceptual states, enabling context-aware, lifelong learning with quantum coherence. Percepts from different modalities are encoded into a shared quantum latent space where cross-modal reasoning occurs via quantum singular value transformation (QSVT) [29], quantum kernel fusion [36], and quantum associative memory retrieval [37]. Using quantum variational policies and reinforcement circuits, the agent adaptively updates its internal goals and ethical boundaries through interaction with complex, uncertain environments.

Table 3, shows the maturity model mapped to our formalism. It specifies how each architectural component—quantum processing (Q), classical control (C), memory (M), perception (P), and action (A)—is incrementally enhanced from Level 1 to Level 4. The levels describe architectural evolution and are not tied to quantitative performance thresholds. It is important to emphasize that present-day quantum hardware remains in the NISQ regime, and current implementations are realistically confined to Level 1 and early Level 2 characteristics. Higher maturity levels serve as structured thought experiments that guide long-term architectural research rather than as claims of imminent deployment.

Within the next 5–10 years, realistic implementations of quantum–agentic systems are expected to remain within Level 1 and early Level 2 characteristics. These include:

- Hybrid quantum–classical orchestration of NISQ devices.
- Variational quantum circuits embedded in reinforcement learning loops for small-scale decision tasks.
- Agent-driven quantum workflow management and hardware-aware scheduling.
- Post-quantum cryptography integration and quantum-secure communication (e.g., QKD) within agent infrastructures.

**Table 3.** Maturity model of quantum agents using the formal tuple  $(Q, C, M, P, A)$ , where  $Q$  is the quantum processing unit,  $C$  the classical controller,  $M$  the memory subsystem,  $P$  the perception module, and  $A$  the action interface for execution and communication.

Tuple	Level 1	Level 2	Level 3	Level 4
	NISQ-Optimized Decision Agent	Hybrid QML Policy Agent	Domain-Aware Adaptive	Fully Quantum-Native
$Q$	NISQ-scale PQC kernels (2–433 qubits); error-mitigated gates	First logical qubits with mid-circuit measurement; variational kernels for policy gradients	Quantum-Transformer blocks on modular fault-tolerant cores; quantum-memory buses	Fault-tolerant universal QPU; cross-node entanglement
$C$	Classical control of PQC-secured MCP channels; noise-aware schedulers	Hybrid scheduler + meta-cognition using QML; live gradient feedback	Adaptive optimizer steering quantum self-attention via QSVT	Autonomous quantum control with dynamic reconfiguration
$M$	Encrypted classical logs; limited quantum scratch qubits	Classical memory + ephemeral logical qubit registers	Entangled memory shards for non-Markovian history	Logical quantum memory; global entangled model
$P$	Classical perception with PQC-protected I/O	Perception enriched by QML feature maps	Quantum multimodal encoders; history-dependent sensing	Coherent quantum sensors; continuous data streams
$A$	Classical actions over PQC channels	Hybrid actions tuned by QML alignment checks	Quantum-in-loop planning; amplitude-amplified search	Entanglement-based communication and quantum actuation

Fully quantum-native cognition, large-scale quantum memory integration, and transformer-level quantum attention mechanisms require fault-tolerant quantum computing and therefore remain as long-term research directions beyond the next decade. Accordingly, the core technical focus of this paper remains on the formal definition of quantum agents and the operationalization of Level 1 and early Level 2 characteristics through concrete prototypes. The higher maturity levels are included to provide structural completeness and long-term architectural orientation, but they are not required for the validation of the presented system models.

## 9. System Architectures and Platform Design

To realize the full potential of quantum agents, we need robust system architectures that tightly integrate quantum processing units (QPUs) with agentic reasoning components. These architectures must support bidirectional information flow, real-time decision-making, and adaptive control, while respecting the physical constraints of quantum hardware. In this section, we outline architectural principles and identify the key components of quantum–agentic platforms.

### 9.1. System Models and Communication Interfaces

A quantum–agentic system typically comprises three interacting layers: (1) the quantum computation layer, (2) the agentic decision layer, and (3) the environment interface layer. The quantum layer handles operations such as quantum state preparation, algorithm execution, and measurement. The agentic layer interprets perceptual inputs, updates internal representations, and makes goal-directed decisions potentially informed by quantum-

augmented reasoning. The interface layer enables sensing, actuation, and communication with external systems or other agents.

Communication between these layers is primarily classical, involving control signals, coordination logic, and measurement results. However, in distributed quantum–agentic systems—comprising multiple agents, each with its own quantum computation layer—quantum communication channels may be used between agents. These support functionalities such as entanglement distribution, quantum teleportation, and secure key exchange. In this context, quantum links are not internal to an individual agent’s layers, but rather operate between distinct agents. Middleware is required to abstract hardware-specific details and expose high-level APIs, allowing agentic modules to access quantum resources and coordinate both classical and quantum interactions seamlessly.

### 9.2. Quantum Memory, Control, and Sensing in Agents

A critical architectural challenge is how agents store and manage quantum information. Quantum memory components must maintain coherence over relevant timescales while remaining accessible for computation and control. This includes quantum registers, delay lines, and entangled memory networks. In mobile or distributed settings, memory management may rely on quantum repeaters and teleportation protocols.

Control systems in quantum agents must operate at multiple levels—from low-level gate operations to high-level behavioral planning. These controls can be implemented through adaptive feedback loops, where agentic reasoning modules analyze environmental inputs and measurement results to dynamically reprogram quantum circuits or reconfigure quantum communication strategies.

### 9.3. Security and Fault Tolerance in Quantum Agents

Robustness and trustworthiness are essential for agentic systems, particularly in high-stakes domains such as defense, critical infrastructure, or autonomous exploration. Quantum agents must account for both classical- and quantum-level threats and errors.

To ensure secure communication among agents in quantum-enhanced systems, security protocols should incorporate established quantum key distribution schemes such as BB84 [38], B92 [39], SARG04 [40], or entanglement-based protocols such as E91 [41] and BBM92 [42] which provide information-theoretic security based on quantum mechanics laws. Complementary quantum authentication methods including the MSW protocol [43] and emerging quantum digital signature schemes [44] can further safeguard the integrity and origin of the message. In distributed agent networks where coordination and trust are critical, quantum-secured consensus mechanisms such as Quantum Byzantine Agreement [45] or information theoretically secure adaptations of Byzantine Fault Tolerance offer resilience against adversarial interference. Together, these cryptographic tools form a foundation for building robust, tamper-resistant agentic systems in the quantum era.

Fault tolerance in quantum agents requires hybrid error correction strategies. While physical qubits are susceptible to decoherence and gate errors, logical qubits protected by quantum error correction codes (QECC) offer a path toward stable computation. Agent architectures must incorporate fault detection, diagnosis, and recovery routines—potentially using AI-driven meta-control loops to adapt to evolving system states and hardware constraints. Together, these architectural considerations form the technological backbone of quantum–agentic platforms and lay the groundwork for scalable, secure, and intelligent quantum-enhanced systems.

In the context of Figure 1, safety mechanisms refer to architectural safeguards such as run-time monitoring, policy validation layers, access control enforcement, and fault-detection routines that constrain agent behavior within predefined operational boundaries.

Auditable agent behavior denotes the ability to log decisions, quantum circuit invocations, parameter updates, and external interactions in a traceable manner, enabling post-hoc verification, compliance assessment, and debugging. These mechanisms are particularly important in hybrid quantum–classical settings, where probabilistic outcomes and hardware variability require transparent oversight.

## 10. Use Cases and Applications

Quantum agents offer transformative potential across diverse domains where intelligence, autonomy, and quantum processing converge. In this section, we illustrate key application areas where quantum–agentic platforms can provide substantial advantages—ranging from scientific discovery to mission-critical autonomous systems.

### 10.1. Scientific Discovery and Simulation

Scientific discovery is increasingly data-driven and computationally intensive. Quantum agents can act as intelligent co-explorers in scientific workflows, accelerating hypothesis testing, simulation, and model generation.

In quantum chemistry and materials science, quantum agents can orchestrate simulations of molecular structures or solid-state systems using quantum simulators while adaptively refining experimental parameters based on intermediate outcomes. This closed-loop approach shortens the discovery cycle by autonomously identifying promising compounds, reaction pathways, or material properties.

In high-energy physics, agents can be employed to optimize quantum circuits for simulating particle interactions or lattice gauge theories. Additionally, they may assist in the calibration and control of quantum sensors used in fundamental experiments, potentially improving their robustness and adaptability. This highlights the versatility of agent-based methods across emerging quantum technologies, including sensing.

### 10.2. Autonomous Systems in Quantum Environments

In environments where quantum systems must operate autonomously—such as in space-based platforms, quantum networking, or remote laboratories—quantum agents serve as critical enablers of resilience, adaptability, and autonomy.

For example, in satellite-based quantum communication systems, agents can manage entanglement distribution, schedule secure key exchanges, and adjust system parameters in response to environmental disturbances or orbital dynamics. Similarly, in distributed quantum sensing systems, agents can coordinate entangled or correlated quantum sensors to perform joint measurements, aggregate and interpret quantum-enhanced data, and dynamically recalibrate sensing units to maintain collective performance.

Autonomous quantum laboratories, or “self-driving labs,” benefit from quantum agents that design, execute, and refine experiments in real-time. These agents close the loop between data acquisition, hypothesis evaluation, and experiment control—accelerating progress and enabling new forms of automation in scientific research.

### 10.3. Secure and Adaptive Systems in Finance, Defense, and Logistics

Quantum agents can enhance security and adaptability in high-stakes domains such as finance, defense, and logistics, where speed, intelligence, and trustworthiness are paramount.

In finance, quantum agents can support risk modeling, portfolio optimization, and fraud detection using quantum-enhanced algorithms. Their ability to reason under uncertainty and process complex correlations gives them an edge in volatile or adversarial markets.

In defense, quantum agents can orchestrate secure communication channels using quantum key distribution (QKD), detect anomalies in sensor data, and coordinate cyber-physical systems under constrained and dynamic conditions. They can also manage quantum radar systems or act as intelligent edge nodes in secure battlefield networks.

In logistics and supply chain optimization, quantum agents can dynamically plan and reconfigure global transportation routes using quantum optimization methods. Their autonomy allows them to respond to disruptions (such as delays, shortages and cyberattacks) and coordinate across decentralized infrastructures.

These use cases demonstrate how quantum agents offer not just computational advantages but also intelligent autonomy—positioning them as a powerful technology across mission-critical and innovation-driven industries.

While these application domains illustrate the long-term potential of quantum–agentic systems, practical realization must begin with concrete implementations under current hardware constraints. To bridge this gap between theory and practice, the following section presents three prototype quantum agents that demonstrate early-stage feasibility in the NISQ era.

## 11. Prototype Quantum Agents: Demonstration of Early Capabilities

To assess the viability of quantum–agentic systems in the NISQ era, we developed three prototype agents that implement distinct quantum reasoning workflows. These prototypes align with Level 1 and Level 2 of our proposed maturity model. Potential applications include cognitive radio networks, online recommender systems, financial strategy selection, quantum edge AI for IoT devices, clinical trials, and adaptive cybersecurity. The developed prototypes are:

1. **Agent 1: Minimal Quantum Agent Using Grover’s Algorithm:** A proof-of-concept agent capable of action selection via Grover-based amplitude amplification, demonstrating basic search-based decision-making.
2. **Agent 2: Quantum Multi-Armed Bandit with Variational Policy:** A reinforcement learning prototype that uses a hybrid classical-quantum loop to learn optimal policies using variational quantum circuits.
3. **Agent 3: Adaptive Quantum Image Encryption Agent:** A perceptual agent that adapts its quantum encryption strategy (XOR, QFT, or scrambling) using entropy feedback via a policy circuit.

These agents illustrate targeted applications in quantum decision-making, perception-driven security, and learning, and are shown to operate within current hardware constraints using Qiskit and PennyLane frameworks.

### 11.1. A Minimal Quantum Agent Using Grover’s Algorithm

Let’s look at a minimal two-qubit four-actions quantum agent model that utilizes Grover’s search algorithm to identify the optimal action in a discrete action space. The environment consists of four possible actions, encoded as two-qubit basis states  $|00\rangle$  to  $|11\rangle$ . The agent is tasked with identifying the correct action, predefined as  $|10\rangle$ , through quantum amplitude amplification.

#### Agent Design and Action Encoding

The quantum agent, as shown in Figure 4, begins by applying Hadamard gates to each qubit to create a uniform superposition over all possible actions. An oracle circuit is constructed to mark the target state  $|10\rangle$  by applying a controlled-Z operation, conditioned on the qubit values corresponding to the target. This oracle serves as the problem-specific component of the Grover algorithm. Subsequently, the Grover diffuser is applied to amplify

the probability amplitude of the target state. Grover’s algorithm requires only a small number of iterations to achieve a high probability of success. For a search space of size  $N$ , the optimal number of iterations are calculated using the following expression [46]:

$$r \approx \left\lfloor \frac{\pi}{4} \sqrt{N} \right\rfloor$$

In our case, the agent operates on a two-qubit system representing  $N = 4$  possible actions. Substituting into the equation gives the following:

$$r \approx \left\lfloor \frac{\pi}{4} \cdot \sqrt{4} \right\rfloor = \left\lfloor \frac{\pi}{4} \cdot 2 \right\rfloor = \left\lfloor \frac{\pi}{2} \right\rfloor \approx 1.$$

Hence, a single Grover iteration is both mathematically sufficient and optimal to amplify the probability of the correct action in this minimal setup. Applying additional iterations would result in probability overshooting, reducing the likelihood of correct selection due to the sinusoidal nature of Grover’s amplitude amplification. The quantum circuit is then measured to collapse the superposition, and the action corresponding to the most frequent measurement outcome is selected as the agent’s decision. The pseudocode algorithm for this agent is given in Algorithm 1.


The Grover-based quantum agent correctly identified the optimal action (10) that matched the environment’s target, demonstrating reliable decision-making via quantum search as depicted in Figure 5. This approach demonstrates the utility of Grover’s algorithm in decision-making tasks, where the correct solution must be retrieved from a small unstructured search space with minimal queries.

---

**Algorithm 1** Quantum agent using Grover search

---

- 1: Define the correct action as bitstring  $a^* = 10$
  - 2: Initialize 2-qubit quantum register in state  $|00\rangle$
  - 3: Apply Hadamard gate on each qubit to prepare uniform superposition
  - 4: Construct oracle  $O$  such that  $O|a^*\rangle = -|a^*\rangle$
  - 5: Apply oracle gate:  $U_f \leftarrow \text{Oracle}$
  - 6: Apply Grover diffuser gate:  $D$
  - 7: Measure the quantum state in the computational basis
  - 8: Record measurement outcomes and select the most frequent bitstring
  - 9: **return** selected bitstring as agent’s chosen action
- 



**Figure 4.** Grover circuit for quantum agent

```

Quantum Agent chose action: 10
Environment's correct action was: 10
Was the agent correct? ✔
    
```

**Figure 5.** Output decision of Grover circuit quantum agent.

### 11.2. Quantum Multi-Armed Bandit Agent Using Variational Policy Circuits

Here we build a quantum learning agent that applies quantum variational circuits to solve the classical multi-armed bandit (MAB) problem. The objective is to identify and

exploit the arm with the highest expected reward among a set of four, using a gradient-trained quantum policy. Such MAB environments are common abstractions for sequential decision-making under uncertainty and this quantum multi-armed bandit agent can be used in cognitive radio networks, online recommender systems, financial strategy selection, quantum edge AI for IoT devices, clinical trials, and adaptive cybersecurity.

### 11.2.1. Problem Setup

Let's consider a stochastic environment comprising four actions (arms) encoded as two-bit strings  $\{00, 01, 10, 11\}$ . Each arm yields a binary reward (0 or 1) with a fixed but unknown probability. In our experimental setting, arm 10 is configured to be optimal with a reward probability of 0.8, while the others vary from 0.2 to 0.5.

The goal of the agent is to learn a probabilistic policy that increases its likelihood of selecting the optimal arm through experience. This setup offers a compelling testbed for quantum machine learning because it combines exploration (sampling from a quantum distribution) and exploitation (gradient updates that shape the quantum state).

### 11.2.2. Agent Design and Action Encoding

The agent's policy is encoded in a fixed, parameterized quantum circuit acting on the two qubits shown in Figure 6. Four trainable rotation angles are updated via gradient descent. The output distribution defines the sampling probability over four discrete actions and is obtained using `qml.probs`, which yields the probabilities of observing each basis state  $\{00, 01, 10, 11\}$ . The agent samples an action based on this distribution and updates the circuit parameters via an Adam optimizer, aiming to maximize the probability of reward-generating actions. The pseudocode algorithm for this agent can be seen in Algorithm 2.

---

#### Algorithm 2 Quantum multi-armed bandit agent using variational policy circuit

---

- 1: **Input:** Arm set  $\mathcal{A} = \{00, 01, 10, 11\}$ , reward probabilities  $p_a$
  - 2: Initialize quantum weights  $\theta \leftarrow [\pi/4, \pi/4, \pi/4, \pi/4]$
  - 3: Initialize cumulative reward  $R \leftarrow 0$ , episode count  $T \leftarrow 100$
  - 4: Initialize Adam optimizer with learning rate  $\eta$
  - 5: **for**  $t = 1$  to  $T$  **do**
  - 6:   Compute probabilities  $\pi_\theta(a)$  using quantum circuit with current  $\theta$
  - 7:   Sample action  $a_t \sim \pi_\theta(a)$
  - 8:   Observe binary reward  $r_t \sim \text{Bernoulli}(p_{a_t})$
  - 9:   Update cumulative reward  $R \leftarrow R + r_t$
  - 10:   Update estimated mean reward for  $a_t$  using incremental average
  - 11:   Define cost:  $J(\theta) \leftarrow -\pi_\theta(a_t)$  if  $r_t = 1$ , else  $+\pi_\theta(a_t)$
  - 12:   Update  $\theta \leftarrow \theta - \eta \nabla_\theta J(\theta)$  using Adam
  - 13: **end for**
  - 14: **Output:** Trained policy parameters  $\theta$ , action history, and reward trajectory
- 





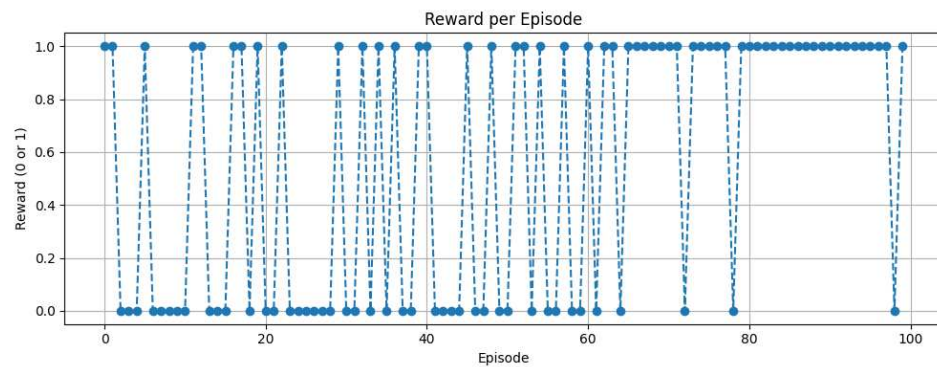
Figure 6. Quantum policy circuit used by the agent.

### 11.2.3. Results and Observations


The agent was trained over 100 episodes. At each step, it sampled an action, observed the binary reward from the environment, and performed a gradient update to reinforce favorable actions. Notably, action selection was *probabilistic* rather than greedy, which allowed sufficient exploration. Figure 7 shows the cumulative reward progression, highlighting steady improvement in expected returns. Figure 8 captures per-episode rewards, reflecting reward noise and policy evolution. Figure 9 confirms that the agent successfully identified the optimal arm 10 as the most frequently chosen.



**Figure 7.** Cumulative reward across training episodes, indicating learning behavior.



**Figure 8.** Per-episode reward. Spikes correspond to successful rewards; fluctuations reflect stochastic environment.



**Figure 9.** Histogram of arm selection frequency. The agent converges to arm 10, the optimal choice.

This experiment demonstrates that a variational quantum circuit can successfully serve as a learnable policy in reinforcement learning. The agent efficiently learns the best action in a noisy reward landscape, reinforcing the viability of quantum-enhanced learning frameworks for adaptive decision-making tasks.

### 11.3. Quantum Agent for Adaptive Quantum Image Encryption

Conventional quantum image encryption (QIE) algorithms, though powerful, remain predominantly static in their design—fixed in their encryption logic, insensitive to image properties, and unaware of adversarial context. As quantum technologies mature, the need for intelligent encryption systems increases, i.e., the systems that are capable of dynamic adaptation, optimization, and real-time decision-making. In this vision, we put forth a class of *quantum agents* designed to autonomously orchestrate and optimize quantum image encryption workflows.

A quantum agent in this context refers to a quantum-aware system that actively learns to select encryption strategies, manage keys, and adapt security operations based on feedback from its environment. Unlike traditional encryption circuits that execute a predetermined sequence of operations, a quantum agent introduces a policy-driven abstraction, offering intelligence, flexibility, and situational awareness. The agent may be designed using hybrid quantum–classical reinforcement learning, variational quantum circuits, or Grover-based decision modules.

#### 11.3.1. Agent Functionalities in Intelligent Cryptography

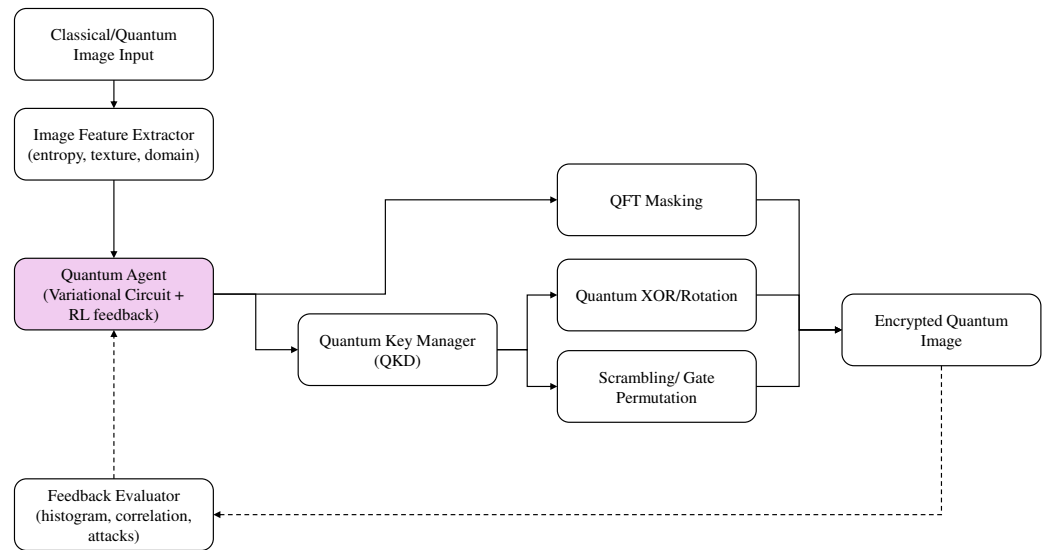
Five core functionalities could be performed by a quantum agent in intelligent cryptography, particularly, in the image encryption domain:

1. **Adaptive Basis and Key Selection:** The agent learns to choose initial quantum states and measurement bases based on image characteristics such as entropy, texture complexity, or domain sensitivity (e.g., medical vs. satellite imaging).
2. **Encryption Strategy Optimization:** The agent dynamically selects among multiple encryption primitives, such as pixel scrambling, quantum Fourier transforms (QFT), or gate-level permutations. The policy evolves over time based on prior encryption success, attack simulations, or circuit execution costs.
3. **Reward-Driven Reinforcement Learning:** Inspired by classical multi-armed bandits, the agent receives feedback from its environment in the form of decryption success, histogram uniformity, correlation metrics, or resistance to simulated attacks. These rewards shape future decisions using a gradient-trained or Grover-amplified policy circuit.
4. **Context-Aware Encryption Personalization:** A quantum agent may tailor encryption strength and style according to the image domain. High-sensitivity applications like biometric or military imaging would automatically invoke deeper, entangled encryption circuits, while generic content may use lightweight routines.
5. **Quantum Key Lifecycle Management:** The agent autonomously manages key generation, QKD channel selection, and entangled state recycling based on device-level fidelity, decoherence risk, or entropy thresholds.

#### 11.3.2. Architecture of the QIE Agent

To design a visionary, intelligent encryption scheme, we propose a *Quantum Agent for Adaptive Quantum Image Encryption* embedded within a hybrid framework. The architecture, illustrated in Figure 10, integrates classical or quantum image inputs, feature extraction, policy-driven quantum circuit selection, and feedback-guided learning. The agent operates on the premise that quantum operations such as QFT masking, quantum XOR, and


scrambling/gate permutation can provide complementary cryptographic transformations depending on image characteristics.



**Figure 10.** Architecture of the Quantum Agent for Adaptive Quantum Image Encryption.

### 11.3.3. Agent Design and Action Encoding

The Quantum Policy Agent is implemented as a two-qubit variational quantum circuit (VQC), serving as the decision-making core of the encryption pipeline. The architecture is shown in Figure 11, where two qubits are used to encode normalized image entropy features. These features are derived from local pixel intensity histograms and are mapped into the rotation angles of the input layer.



**Figure 11.** Quantum Policy Agent circuit. It encodes entropy into qubit rotations and produces a distribution over four encryption actions.

The circuit begins by applying parameterized rotations  $RY(x_0)$  and  $RY(x_1)$  on the two qubits, encoding the input feature vector. This is followed by a trainable variational layer that consists of two single-qubit gates  $RZ(\theta_0)$  and  $RZ(\theta_1)$  to encode the agent’s policy parameters, introducing adaptive phase shifts to each qubit. A CNOT gate is then applied to entangle the qubits, enabling the circuit to model correlated action probabilities. Finally, another pair of  $RY(\theta_2)$  and  $RY(\theta_3)$  gates is applied to modulate the amplitudes post-entanglement.

Upon measurement, the circuit outputs a two-qubit probability distribution over four possible outcomes,  $\{00, 01, 10, 11\}$ , corresponding to the encryption actions XOR, QFT, Scramble, and None, respectively. These action probabilities are treated as a quantum policy distribution, from which the agent samples one operation per image segment. The pseudocode algorithm of this QIE agent is given in Algorithm 3.

The circuit is trained using reinforcement learning to maximize a reward function based on the entropy of the encrypted image. Over multiple episodes, the weights  $\theta = \{\theta_0, \theta_1, \theta_2, \theta_3\}$  are updated via gradient-based optimization to increasingly favor actions that produce higher post-encryption entropy. This enables the agent to learn adaptive and content-sensitive encryption strategies based on the statistical structure of the input image.

---

**Algorithm 3** Quantum RL agent for adaptive image encryption
 

---

**Require:** Input image  $I$  of size  $N \times N$

- 1: Extract image feature:  $f \leftarrow \text{Entropy}(I)$
  - 2: Initialize quantum policy parameters  $\theta \in \mathbb{R}^4$
  - 3: Initialize optimizer  $\mathcal{O}$  and learning rate  $\eta$
  - 4: Convert image  $I$  to 4 bit segments  $\{s_1, s_2, \dots, s_T\}$
  - 5: Generate encryption key  $k \in \{0, 1\}^4$
  - 6: **for** each episode  $e = 1, \dots, E$  **do**
  - 7:   Preprocess feature vector  $x \leftarrow \text{ScaleToQuantumDomain}(f)$
  - 8:   Compute action probabilities  $p_a \leftarrow \text{PolicyQNode}(\theta, x)$
  - 9:   Sample action  $a \sim p_a$  where  $a \in \{\text{XOR}, \text{QFT}, \text{Scramble}, \text{None}\}$
  - 10:   **for** each segment  $s_i$  **do**
  - 11:     **if**  $a$  is XOR **then**
  - 12:        $e_i \leftarrow \text{QUANTUMXOR}(s_i, k)$
  - 13:     **else if**  $a$  is QFT **then**
  - 14:        $e_i \leftarrow \text{QFTENCRYPT}(s_i)$
  - 15:     **else if**  $a$  is Scramble **then**
  - 16:        $e_i \leftarrow \text{SCRAMBLEENCRYPT}(s_i)$
  - 17:     **else**
  - 18:        $e_i \leftarrow s_i$  ▷ No-op
  - 19:     **end if**
  - 20:   **end for**
  - 21:   Reconstruct encrypted image  $\hat{I} \leftarrow \text{SegmentsToImage}(\{e_1, \dots, e_T\})$
  - 22:   Evaluate reward  $r \leftarrow \text{Entropy}(\hat{I})$
  - 23:   Update policy:  $\theta \leftarrow \mathcal{O}.\text{Step}(\theta, -r)$
  - 24: **end for**
- 

### 11.3.4. Quantum Encryption Primitives

The proposed agent selects from a set of three quantum encryption operations, each implemented as a dedicated quantum circuit. These operations are designed to exploit different quantum properties—linearity, superposition, and permutation—to obfuscate pixel values and structure in gray-scale images. The choice of operation is learned dynamically based on image features and feedback from the encrypted output. Below we detail the construction and roles of these circuits in the agent’s workflow.

#### Quantum XOR:

The quantum XOR circuit, as shown in Figure 12, simulates modular addition using CNOT gates and ancillary qubits. This circuit loads the 4 bit segment of image data onto primary qubits, and a 4 bit quantum key is prepared on auxiliary qubits. A sequence of CNOT gates then applies a controlled bit-flip operation, implementing the XOR logic. This process yields a reversible, low-complexity encryption mechanism that is sensitive to key variations.

#### Quantum Fourier Transform (QFT):

The QFT encryption circuit transforms image segments into the frequency domain. As shown in Figure 13, a series of Hadamard gates and controlled phase rotations are applied hierarchically across qubits, encoding phase relations that decorrelate adjacent bit positions.

This transformation introduces significant diffusion, ideal for images with high texture or spatial redundancy. Within the agent’s training cycle, QFT often becomes the preferred choice due to its entropy-amplifying characteristics, particularly for high-entropy inputs where additional complexity yields diminishing returns in XOR-based masking.




Figure 12. Quantum XOR encryption circuit with ancilla qubits.

**Scrambling/Gate Permutation:**

Scrambling is achieved through a combination of SWAP and Pauli-X gates to reorder and invert specific qubit states as shown in Figure 14. This form of transformation does not require additional ancilla qubits and provides a light-weight but effective non-linear permutation of input data. The scrambling circuit is particularly useful in diversifying ciphertexts when redundancy is low but spatial patterns persist. For example, the agent tends to choose this action when the image exhibits moderate entropy with repetitive structural features, as it disrupts positional regularity without relying on complex quantum arithmetic.


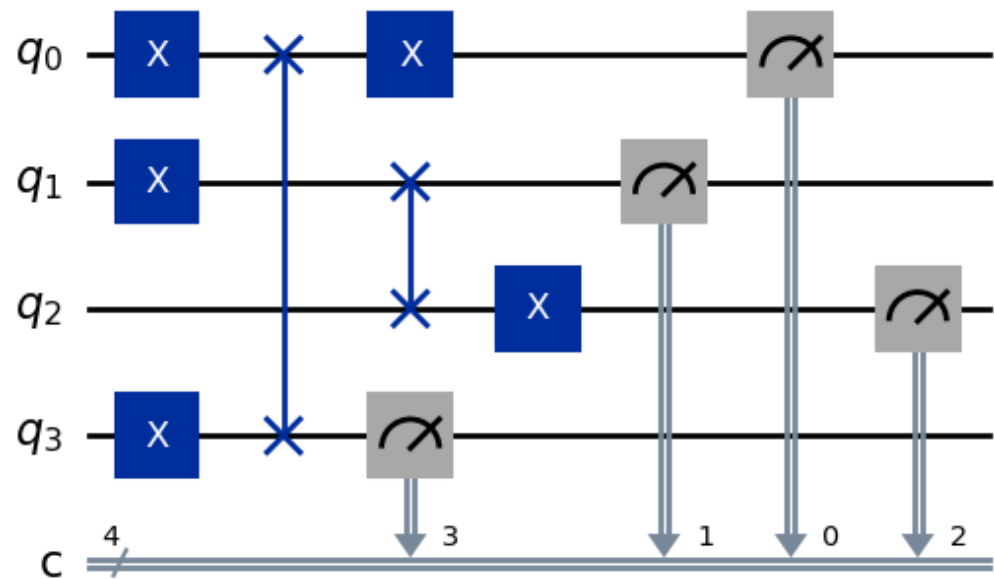


Figure 13. QFT encryption circuit for image segments.



**Figure 14.** Scrambling circuit applying swaps and targeted X gates.

### 11.3.5. Reinforcement Learning Integration

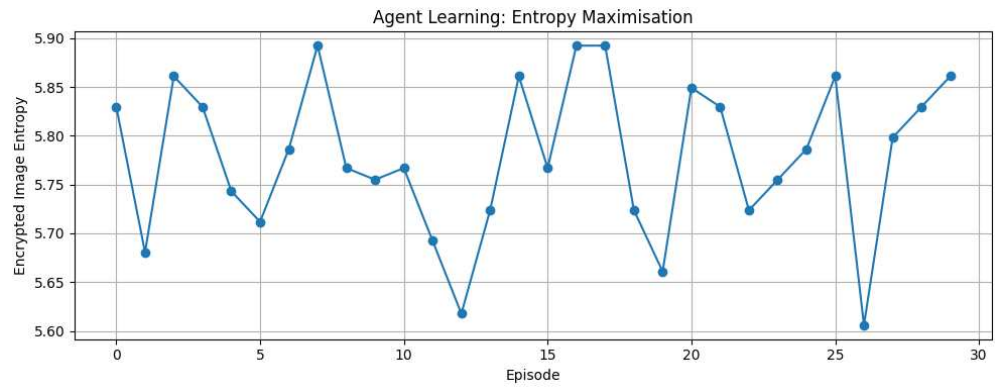
The agent is trained using a reinforcement learning loop with entropy maximization as the reward proxy. Over 30 episodes, the agent receives image entropy as input and selects one encryption action to apply to all image segments. The encrypted image's entropy is computed, and negative entropy is used as the loss function for gradient descent optimization of the VQC parameters. This continuous feedback loop encourages the agent to favor transformations that yield higher visual randomness in ciphertexts.

### 11.3.6. Empirical Results and Observations


As shown in Figure 15, the agent learns to prefer transformations/encryption methods that yield higher ciphertext entropy. The action selection histogram in Figure 16 reveals that QFT is frequently selected, which aligns with its strong entropy-enhancing properties. XOR and Scramble actions are chosen when the input entropy stabilizes, showcasing the agent's adaptivity.

The proposed agent integrates reinforcement learning with quantum encryption for dynamic, context-aware image protection. Unlike static schemes, it learns from ciphertext outcomes and adjusts operation preferences accordingly. This framework is especially promising for use cases in quantum-secured surveillance, medical imaging, and adaptive quantum watermarking. Future work may integrate multi-feature inputs (e.g., texture, frequency spectrum), and extend the operation space with more complex gate-level manipulations or noise-adaptive strategies.

It is important to note that the adaptive QIE agent introduces additional computational overhead compared to a static encryption scheme. A fixed encryption pipeline (e.g., always applying XOR or QFT) requires only a single deterministic transformation per image segment. In contrast, the adaptive agent executes a reinforcement learning loop that includes repeated policy circuit evaluations, probabilistic action sampling, entropy-based feedback computation, and parameter updates across multiple episodes. While this hybrid loop enables context-aware selection of encryption strategies and potential robustness improvements, it incurs higher execution cost due to iterative quantum circuit calls and classical optimization steps. The design therefore reflects a trade-off between computational efficiency and adaptive security capability.



**Figure 15.** Encrypted image entropy over episodes. The agent learns to favor transformations that yield high ciphertext entropy.



**Figure 16.** Frequencies of encryption operations selected during training. QFT dominates due to its entropy amplification.

#### 11.4. Scalability, Communication Cost, and Latency in Hybrid Classical–Quantum Agent Loops

Hybrid quantum agents incur quantum execution cost from circuit sampling and classical overhead from orchestration and learning updates. We summarize the dominant scaling behavior of the prototype workflows in Section 11.

##### Generic hybrid-loop cost model.

Let  $n$  be the number of qubits in the agent circuit,  $d$  its depth, and  $S$  the number of measurement shots used to estimate output probabilities. Over  $E$  agent episodes, the number of quantum circuit evaluations is

$$N_{\text{eval}} = E \cdot C_{\text{step}}, \tag{1}$$

where  $C_{\text{step}}$  is the number of circuit calls per decision step (typically one per action evaluation in the prototypes). Each evaluation returns  $n$  bit measurement samples, giving the classical readout cost

$$B_{\text{meas}} = \mathcal{O}(N_{\text{eval}} \cdot S \cdot n) \text{ bits.} \tag{2}$$

For variational agents with  $p$  trainable parameters, classical parameter updates scale as  $\mathcal{O}(N_{\text{eval}} \cdot p)$ .

**Latency decomposition.**

The per-decision latency may be written as

$$T_{\text{step}} = T_{\text{prep}} + T_{\text{queue}} + S T_{\text{shot}}(n, d) + T_{\text{post}} + T_{\text{upd}}, \quad (3)$$

where  $T_{\text{queue}}$  captures backend access delays and  $T_{\text{upd}}$  denotes classical learning updates. In hybrid workflows, repeated sampling and backend latency dominate overall response time.

**Prototype-specific scaling.**

- Grover decision agent: For an action space of size  $N = 2^n$ , the number of Grover iterations scales as  $\mathcal{O}(\sqrt{N})$ , with circuit depth increasing proportionally through repeated oracle and diffusion operations. Each amplified decision is then sampled to identify the dominant action.
- Variational multi-armed bandit agent: A two-qubit parameterized policy circuit outputs a four-action probability distribution. Each episode performs probability estimation via repeated shots followed by a classical parameter update, giving  $N_{\text{eval}} = \mathcal{O}(E)$  and communication dominated by measurement sampling.
- Adaptive quantum image encryption agent: The policy circuit similarly operates on two qubits to select among four encryption actions. An  $8 \times 8$  gray-scale image is decomposed into fixed-size bit segments, and encryption cost scales linearly with the number of processed segments per episode, while policy evaluation follows the same hybrid sampling model above.

Hence, the scalability is constrained primarily by the sampling budget  $S$ , the number of hybrid iterations  $E$ , and remote execution latency  $T_{\text{queue}}$ . These factors motivate practical execution strategies in hybrid loops: (i) batching multiple circuit evaluations per optimization step to reduce host–QPU call overhead, (ii) parallelising classical feature/reward computation with circuit preparation where possible, and (iii) reducing sampling cost via adaptive shot allocation (e.g., using fewer shots early in training and increasing shots only when policy updates stabilize).

**Reproducibility under Hardware Noise**

Results obtained from NISQ devices may vary across runs due to calibration drift, gate errors, and stochastic measurement noise. Consequently, empirical outcomes from hybrid quantum–classical workflows can differ depending on the backend state and execution conditions. The prototypes presented in this work therefore focus on architectural feasibility, while reproducible large-scale benchmarking will require controlled experiments on stable hardware configurations.

**11.5. Complexity and Run-Time Considerations**

The complexity expressions referenced in the prototype descriptions correspond to the theoretical algorithmic properties of the employed quantum methods (e.g., Grover search and variational quantum circuits). Since the implemented prototypes operate on very small NISQ-scale circuits and are primarily intended as architectural proof-of-concept demonstrations, empirical run-time measurements would largely reflect simulator or hardware latency rather than intrinsic algorithmic complexity.

**12. Challenges and Open Questions**

While the concept of quantum agents holds significant promise, its realization comes with technical, conceptual, and practical challenges. These span the scalability of quan-

tum hardware, the evaluation of agent performance in quantum settings, and the lack of standardized frameworks. In this section, we highlight key open questions that must be addressed to advance the field.

### 12.1. Scalability and Resource Constraints

Quantum computing remains limited by hardware constraints such as qubit count, coherence time, gate fidelity, and error rates. These limitations directly impact the feasibility of integrating quantum processing into agentic loops, especially when real-time responsiveness is required.

This includes dynamic scheduling of quantum operations, prioritizing critical computations, and offloading tasks to classical components where appropriate. Agents need to make resource-aware decisions that balance quantum advantage with practical constraints.

Open questions include: How can agents effectively orchestrate computation between quantum and classical subsystems? What frameworks enable modular and scalable integration as quantum hardware evolves? And how can we rigorously evaluate the performance of such hybrid agentic systems under the physical and algorithmic constraints of the NISQ era?

### 12.2. Agent Evaluation in Quantum Contexts

Traditional metrics for evaluating agents—such as performance, learning speed, robustness, and adaptability—may not fully capture the behavior of quantum agents.

Evaluating quantum agents thus requires a hybrid approach that considers both classical agent benchmarks and quantum-specific performance indicators. Moreover, it remains unclear how to fairly compare classical and quantum agents in shared environments, especially when problem instances are inherently quantum or when classical emulation is infeasible.

Key open questions include: What are meaningful benchmarks for quantum agents? How do we isolate the contribution of quantum components from overall system behavior? And how can we design environments and competitions that fairly evaluate hybrid intelligence?

### 12.3. Interoperability and Standards

As quantum agents evolve, interoperability becomes critical for integration across platforms, hardware backends, and AI frameworks. Currently, there is no unified abstraction layer for quantum–agentic systems, and development is often tightly coupled to specific quantum SDKs (such as, Qiskit, Cirq and Braket).

Quantum agents must interface with heterogeneous components: quantum processors, classical control logic, AI toolchains, and networking protocols. Without standardized APIs, formats, and design patterns, scalability and adoption are severely constrained.

Important open questions include: What abstraction layers can decouple agent logic from quantum hardware? How can we define standard interfaces for quantum reasoning modules? And what governance structures are needed to establish best practices and interoperability norms across the quantum–agentic ecosystem?

## 13. Conclusions and Outlook

This work introduced a structured framework for quantum–agentic systems, providing (i) a formal definition of quantum agents based on the tuple  $(Q, C, M, P, A)$ , (ii) a bidirectional control taxonomy distinguishing quantum-enhanced agency and agent-enabled quantum systems, and (iii) an architectural maturity model aligned with projected hardware development. These contributions are conceptual in nature and aim to establish a coherent systems-level perspective for integrating quantum computation into agentic AI.

In addition to the conceptual framework, we implemented three proof-of-concept prototypes operating under NISQ constraints. These prototypes demonstrate the feasibility of embedding quantum search, variational policy learning, and adaptive quantum encryption within agentic perception–decision–action loops. They serve as architectural demonstrations rather than empirical benchmarks of quantum advantage.

Importantly, the results presented here do not claim general performance superiority over classical approaches. Instead, they illustrate how quantum computational primitives can be systematically integrated into autonomous systems under current hardware limitations.

Future research will require rigorous comparative benchmarking, scalability analysis under increasing qubit counts, and empirical validation on emerging fault-tolerant hardware platforms. As quantum technology matures, the architectural principles outlined in this work may provide a structured foundation for evaluating when and where quantum-enhanced agency yields measurable advantages.

While the maturity model outlines a long-term evolutionary perspective, the empirical contribution of this work is confined to NISQ-feasible architectures and prototype implementations.

### 13.1. *The Future of Quantum–Agentic Intelligence*

Over the coming decade, continued advances in quantum hardware and hybrid control architectures may enable selected quantum–agentic components to transition from laboratory prototypes toward domain-specific applied systems.

We envision quantum–agentic intelligence evolving along three parallel trajectories:

- **Cognitive Expansion:** Agents will gain the ability to reason with uncertainty, simulate quantum phenomena, and make decisions that exploit quantum-enhanced optimization and learning.
- **Operational Autonomy:** In particular, near- to mid-term progress may focus on semi-autonomous orchestration layers, including adaptive scheduling, calibration assistance, and fault-aware workflow management under hybrid quantum–classical control.
- **Systemic Integration:** Quantum agents will operate within larger multi-agent ecosystems, integrating with cloud services, edge devices, and classical AI systems to form hybrid intelligent infrastructures.

### 13.2. *Research Roadmap*

In order to realize the vision of quantum–agentic systems, coordinated research is required across multiple domains, we outline the following road map:

1. **Foundational Theory:** Develop formal models and complexity analyses for quantum–agentic behaviors, including new frameworks for learning, planning, and interaction in quantum environments.
2. **Benchmarking and Evaluation:** Establish standardized metrics, testbeds, and datasets for evaluating quantum agents, both in simulation and on real hardware.
3. **Hardware–Software Co-Design:** Design agent-friendly quantum hardware and co-optimize control protocols to support tight integration of perception, computation, and actuation.
4. **Tooling and Abstractions:** Build high-level frameworks, middleware, and development tools that abstract quantum complexity while supporting agentic flexibility.
5. **Application Pilots:** Deploy quantum agents in domain-specific pilots (such as, quantum chemistry, logistics and cybersecurity) to test feasibility, performance, and user acceptance in operational settings.

6. **Ethical and Societal Implications:** Investigate the broader implications of autonomous quantum systems, including transparency, accountability, and the governance of intelligent quantum infrastructure.

By aligning efforts across academia, industry, and policy, we can accelerate the emergence of quantum–agentic intelligence as a robust, ethical, and scalable foundation for next-generation autonomous systems.

**Author Contributions:** Conceptualization, M.T.; Methodology, M.T. and M.S.K.; Software, M.S.K.; Validation, W.J.B. and M.S.K.; Formal Analysis, M.T. and W.J.B.; Investigation, M.S.K.; Data Curation, S.D.; Writing—Original Draft Preparation, M.T., E.S., S.D. and M.S.K.; Writing—Review & Editing, E.S. and W.J.B.; Visualization, S.D.; Supervision, W.J.B.; Project Administration, E.S. and M.T.; Funding Acquisition, W.J.B. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research received no external funding.

**Data Availability Statement:** The raw data supporting the conclusions of this article will be made available by the authors on request.

**Conflicts of Interest:** Eldar Sultanow is an employee of Capgemini Germany. The paper reflects the views of the scientists and not the company.

## References

1. Klusch, M. Toward Quantum Computational Agents. In *International Workshop on Computational Autonomy*; Springer: Berlin/Heidelberg, Germany, 2003.
2. Saggio, V.; Asenbeck, B.E.; Hamann, A.; Strömberg, T.; Schiansky, P.; Dunjko, V.; Friis, N.; Harris, N.C.; Hochberg, M.; Englund, D.; et al. Experimental quantum speed-up in reinforcement learning agents. *Nature* **2021**, *591*, 229–233. [[CrossRef](#)] [[PubMed](#)]
3. Elliott, T.J.; Gu, M.; Garner, A.J.; Thompson, J. Quantum Adaptive Agents with Efficient Long-Term Memories. *Phys. Rev. X* **2022**, *12*, 011007. [[CrossRef](#)]
4. Thompson, J.; Riechers, P.M.; Garner, A.J.; Elliott, T.J.; Gu, M. Energetic advantages for quantum agents in online execution of complex strategies. *arXiv* **2025**, arXiv:2503.19896. [[CrossRef](#)]
5. Yun, W.J.; Kwak, Y.; Kim, J.P.; Cho, H.; Jung, S.; Park, J.; Kim, J. Quantum Multi-Agent Reinforcement Learning via Variational Quantum Circuit Design. *arXiv* **2022**, arXiv:2203.10443. [[CrossRef](#)]
6. Klusch, M.; Lässig, J.; Müssig, D.; Macaluso, A.; Wilhelm, F.K. Quantum Artificial Intelligence: A Brief Survey. *Künstliche Intell.* **2024**, *38*, 257–276. [[CrossRef](#)]
7. Kipu Quantum. Kipu Quantum—Application- and Hardware-Specific Quantum Computing. 2026. Available online: <https://kipu-quantum.com/> (accessed on 23 February 2026).
8. Grover, L.K. Quantum mechanics helps in searching for a needle in a haystack. *Phys. Rev. Lett.* **1997**, *79*, 325. [[CrossRef](#)]
9. Gilliam, A.; Woerner, S.; Gonciulea, C. Grover adaptive search for constrained polynomial binary optimization. *Quantum* **2021**, *5*, 428. [[CrossRef](#)]
10. Benchasattabuse, N.; Satoh, T.; Hajdušek, M.; Van Meter, R. Amplitude amplification for optimization via subdivided phase oracle. In *Proceedings of the 2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*; IEEE: New York, NY, USA, 2022; pp. 22–30.
11. Biron, D.; Biham, O.; Biham, E.; Grassl, M.; Lidar, D.A. Generalized Grover search algorithm for arbitrary initial amplitude distribution. In *Proceedings of the First NASA International Conference, QCQC'98, Palm Springs, CA, USA, 17–20 February 1998*; Lecture Notes in Computer Science; Springer: Berlin/Heidelberg, Germany, 1999; pp. 140–147.
12. Farhi, E.; Goldstone, J.; Gutmann, S. A Quantum Approximate Optimization Algorithm. *arXiv* **2014**, arXiv:1411.4028. [[CrossRef](#)]
13. Blekos, K.; Brand, D.; Ceschini, A.; Chou, C.H.; Li, R.H.; Pandya, K.; Summer, A. A review on quantum approximate optimization algorithm and its variants. *Phys. Rep.* **2024**, *1068*, 1–66. [[CrossRef](#)]
14. Das, A.; Chakrabarti, B.K. *Quantum Annealing and Related Optimization Methods*; Springer Science & Business Media: Berlin/Heidelberg, Germany, 2005; Volume 679.
15. Rosmanis, A. Hybrid quantum-classical search algorithms. *ACM Trans. Quantum Comput.* **2024**, *5*, 1–18. [[CrossRef](#)]
16. Guo, X.; Muta, T.; Zhao, J. Quantum Circuit Ansatz: Patterns of Abstraction and Reuse of Quantum Algorithm Design. *arXiv* **2024**, arXiv:2405.05021.
17. Qin, J. Review of Ansatz Designing Techniques for Variational Quantum Algorithms. *arXiv* **2022**, arXiv:2212.04913. [[CrossRef](#)]

18. Wiersema, R.; Zhou, C.; de Sereville, Y.; Carrasquilla, J.F.; Kim, Y.B.; Yuen, H. Exploring Entanglement and Optimization within the Hamiltonian Variational Ansatz. *PRX Quantum* **2020**, *1*, 020319. [[CrossRef](#)]
19. Clinton, L.; Cubitt, T.; Flynn, B.; Gambetta, F.M.; Klassen, J.; Montanaro, A.; Piddock, S.; Santos, R.A.; Sheridan, E. Towards near-term quantum simulation of materials. *Nat. Commun.* **2024**, *15*, 211. [[CrossRef](#)]
20. Feynman, R.P. Simulating physics with computers. *Int. J. Theor. Phys.* **1982**, *21*, 553–555. [[CrossRef](#)]
21. Erhart, L.; Yoshida, Y.; Khinevich, V.; Mizukami, W. Coupled cluster method tailored with quantum computing. *Phys. Rev. Res.* **2024**, *6*, 023230. [[CrossRef](#)]
22. Zou, Y.; Cheng, A.H.; Aldossary, A.; Bai, J.; Leong, S.X.; Campos-Gonzalez-Angulo, J.A.; Choi, C.; Ser, C.T.; Tom, G.; Wang, A.; et al. El Agente: An Autonomous Agent for Quantum Chemistry. *arXiv* **2025**, arXiv:2505.02484. [[CrossRef](#)]
23. U.S. Department of Energy, Office of Science. *Quantum Information Science Applications Roadmap for the U.S. Department of Energy*; Technical Report; U.S. Department of Energy: Washington, DC, USA, 2024.
24. Hou, X.; Zhao, Y.; Wang, S.; Wang, H. Model Context Protocol (MCP): Landscape, Security Threats, and Future Research Directions. *arXiv* **2025**, arXiv:2503.23278. [[CrossRef](#)]
25. Khan, M.S.; Ahmad, J.; Al-Dubai, A.; Pitropakis, N.; Ghaleb, B.; Ullah, A.; Khan, M.A.; Buchanan, W.J. Chaotic quantum encryption to secure image data in post quantum consumer technology. *IEEE Trans. Consum. Electron.* **2024**, *70*, 7087–7101. [[CrossRef](#)]
26. Kenthapadi, K.; Lakkaraju, H.; Rajani, N. Generative ai meets responsible ai: Practical challenges and opportunities. In Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Long Beach, CA, USA, 6–10 August 2023; pp. 5805–5806.
27. Khatri, N.; Matos, G.; Coopmans, L.; Clark, S. Quixer: A quantum transformer model. *arXiv* **2024**, arXiv:2406.04305. [[CrossRef](#)]
28. Zhang, H.; Zhao, Q. A Survey of Quantum Transformers: Approaches, Advantages, Challenges, and Future Directions. *arXiv* **2025**, arXiv:2504.03192. [[CrossRef](#)]
29. Gilyén, A.; Su, Y.; Low, G.H.; Wiebe, N. Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics. In Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, Phoenix, AZ, USA, 23–26 June 2019; pp. 193–204.
30. Seneviratne, A.; Walters, P.L.; Wang, F. Polynomial time and space quantum algorithm for the simulation of non-Markovian quantum dynamics. *arXiv* **2024**, arXiv:2411.18168. [[CrossRef](#)]
31. Recio-Armengol, E.; Eisert, J.; Meyer, J.J. Single-shot quantum machine learning. *Phys. Rev. A* **2025**, *111*, 042420. [[CrossRef](#)]
32. Wilms, A.; Ohff, L.; Skolik, A.; Eisert, J.; Khatri, S.; Reiss, D.A. Quantum reinforcement learning of classical rare dynamics: Enhancement by intrinsic Fourier features. *arXiv* **2025**, arXiv:2504.16258. [[CrossRef](#)]
33. Liu, J.; Liu, M.; Liu, J.P.; Ye, Z.; Wang, Y.; Alexeev, Y.; Eisert, J.; Jiang, L. Towards provably efficient quantum algorithms for large-scale machine-learning models. *Nat. Commun.* **2024**, *15*, 434. [[CrossRef](#)] [[PubMed](#)]
34. Chen, Z.; Wang, S.; Tan, Z.; Fu, X.; Lei, Z.; Wang, P.; Liu, H.; Shen, C.; Li, J. A Survey of Scaling in Large Language Model Reasoning. *arXiv* **2025**, arXiv:2504.02181. [[CrossRef](#)]
35. Brassard, G.; Høyer, P.; Mosca, M.; Tapp, A. Quantum amplitude amplification and estimation. *arXiv* **2002**, arXiv:quant-ph/0005055v1. [[CrossRef](#)]
36. Vedaie, S.S.; Noori, M.; Oberoi, J.S.; Sanders, B.C.; Zahedinejad, E. Quantum multiple kernel learning. *arXiv* **2020**, arXiv:2011.09694. [[CrossRef](#)]
37. Trugenberger, C.A. Probabilistic quantum memories. *Phys. Rev. Lett.* **2001**, *87*, 067901. [[CrossRef](#)] [[PubMed](#)]
38. Bennett, C.H.; Brassard, G. Quantum cryptography: Public key distribution and coin tossing. *Theor. Comput. Sci.* **2014**, *560*, 7–11. [[CrossRef](#)]
39. Bennett, C.H. Quantum cryptography using any two nonorthogonal states. *Phys. Rev. Lett.* **1992**, *68*, 3121. [[CrossRef](#)] [[PubMed](#)]
40. Scarani, V.; Acin, A.; Ribordy, G.; Gisin, N. Quantum Cryptography Protocols Robust against Photon Number Splitting Attacks for Weak Laser Pulse Implementations. *Phys. Rev. Lett.* **2004**, *92*, 057901. [[CrossRef](#)] [[PubMed](#)]
41. Ekert, A.K. Quantum cryptography based on Bell’s theorem. *Phys. Rev. Lett.* **1991**, *67*, 661. [[CrossRef](#)]
42. Bennett, C.H.; Brassard, G.; Mermin, N.D. Quantum cryptography without Bell’s theorem. *Phys. Rev. Lett.* **1992**, *68*, 557. [[CrossRef](#)]
43. Barnum, H.; Crépeau, C.; Gottesman, D.; Smith, A.; Tapp, A. Authentication of quantum messages. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science*; IEEE : New York, NY, USA, 2002; pp. 449–458.
44. Wallden, P.; Dunjko, V.; Kent, A.; Andersson, E. Quantum digital signatures with quantum-key-distribution components. *Phys. Rev. A* **2015**, *91*, 042304. [[CrossRef](#)]

45. Weng, C.X.; Gao, R.Q.; Bao, Y.; Li, B.H.; Liu, W.B.; Xie, Y.M.; Lu, Y.S.; Yin, H.L.; Chen, Z.B. Beating the fault-tolerance bound and security loopholes for byzantine agreement with a quantum solution. *Research* **2023**, *6*, 0272. [[CrossRef](#)] [[PubMed](#)]
46. Github. *Qiskit Textbook*; Github: San Francisco, CA, USA, 2023.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.