ISSN 2464-4617 (print)
ISSN 2464-4625 (online)
Quantum Informatics, Computing & Technology 2025
QC-Horizon 2025
https://www.qc-horizon.eu/

Bayesian Update Step Using Adiabatic Quantum Computing

Norman Bunk
Rheinische Friedrich-Wilhelms-Universität Bonn
Regina-Pacis-Weg 3
53113 Bonn,
Germany
s6nobunk@uni-bonn.de /
norman.bunk@gmx.de

Felix Govaers
Fraunhofer FKIE
Fraunhofer Str. 20
53343 Wchatberg,
Germany
felix.govaers@fkie.fraunhofer.de

ABSTRACT

We propose a quantum algorithm for the filtering step in Bayesian state estimation, based on adiabatic quantum computing (AQC). The approach embeds the likelihood function into the energy landscape of an Ising Hamiltonian, enabling the posterior distribution to emerge through quantum annealing. Unlike existing gate-based methods, our formulation avoids complex conditional operations and scales more naturally with the size of the state space. Numerical simulations confirm that the resulting distributions closely match the correct posterior, even with approximate parameter settings. This demonstrates the feasibility of AQC as a tool for probabilistic inference in high-dimensional filtering problems.

Keywords

adiabatic quantum computing, sensor data fusion, Bayes theorem

1 INTRODUCTION

Quantum computing has emerged as a disruptive computational paradigm with the potential to address classically intractable problems across various fields, including artificial intelligence (AI), optimization, and data processing. One promising area of application is sensor data fusion, a fundamental concept in robotics, surveillance, autonomous vehicles, and environmental monitoring, where information from multiple sources must be integrated to infer the state of a system. The standard mathematical framework for sensor fusion is Bayesian filtering, which recursively applies a prediction and an update step to estimate the posterior distribution over the state space.

While the prediction step often relies on models that are tractable in both classical and quantum domains, the Bayesian update step poses significant challenges [**?**]. It requires the conditioning of a prior distribution on new measurement data using the likelihood function. In high-dimensional state spaces, this operation quickly becomes computationally expensive. Quantum computers, by virtue of their ability to represent and manipulate high-dimensional probability distributions using amplitude encoding, have been proposed as a potential solution to this bottleneck.

To date, the implementation of a quantum algorithm for the Bayesian update step proposed in [3] is employing gate-based quantum circuits. This method encodes the posterior distribution by applying a sequence of state-conditioned unitary operations. While theoretically sound, the approach scales poorly, as it necessitates a distinct quantum gate for each possible state hypothesis. This results in exponential growth in circuit depth and complexity, rendering it impractical for real-world systems with continuous or large discrete state spaces.

In this paper, we introduce a novel quantum approach to the Bayesian update step using adiabatic quantum computing, specifically quantum annealing. Unlike gate-based quantum algorithms, our method encodes the Bayesian update directly into the energy landscape of an Ising-type Hamiltonian. The likelihood function is embedded into the cost function of the annealing process, enabling the system to evolve toward the most probable state configuration in a natural and scalable way. This formulation bypasses the need for complex conditional operations and opens a new direction for quantum inference methods.

Our approach fits into the growing body of literature at the intersection of quantum computing and artificial in-

telligence, particularly in the use of quantum hardware for probabilistic inference, optimization, and machine learning. While most existing work in quantum machine learning focuses on variational circuits and classification tasks, few have addressed the challenge of recursive inference under uncertainty, a critical requirement in real-time decision-making systems. Moreover, our contribution complements recent research that explores quantum-enhanced data fusion, offering a concrete mechanism for incorporating observational data in a probabilistically rigorous way.

We present a complete formulation of the Bayesian update using quantum annealing, validate our method on simulated sensor fusion tasks, and analyze its complexity and scalability compared to classical and gate-based quantum methods. Our results suggest that adiabatic quantum computing provides a promising and efficient alternative for Bayesian inference in high-dimensional sensor fusion scenarios.

2 PROBLEM FORMULATION

Bayesian filtering is a recursive estimation method that updates the belief over a system state based on new observations. In its discrete form, the posterior distribution $p(x_k|Z_k)$ is computed by conditioning the prior $p(x_k)$ on the latest measurement z_k via the likelihood $p(z_k|x_k)$. This results in a pointwise multiplication followed by normalization:

$$p(x_k|Z_k) \propto p(z_k|x_k) \cdot p(x_k|Z_{k-1}) \tag{1}$$

Here, $x_k \in \mathrm{X} \subset \mathbb{R}^D$ denotes the discrete state at time k, and $Z_k = \{z_1, z_2, \ldots, z_k\}$ the measurement history up to time k.

The classical framework for Bayesian estimation and sensor data fusion is well established, with in-depth discussions available in foundational texts [5, 2]. In many practical applications, the state space is discretized into a finite grid, where each axis $d = 1, \ldots, D$ is divided into N_d bins of size Δ_d. The resulting joint distribution $p(x)$ over the full state space becomes a tensor of shape $N_1 \times N_2 \times \ldots \times N_D$, containing $\prod_{d=1}^{D} N_d$ entries. For high-dimensional problems, this representation quickly becomes computationally infeasible in terms of storage and processing requirements.

Quantum computing offers a compact representation for such distributions. Using amplitude encoding, a quantum state $|\psi\rangle$ of n qubits can represent a probability distribution over $N = 2^n$ states:

$$|\psi\rangle = \sum_{x=0}^{N-1} a_x |x\rangle, \quad \text{with} \quad |a_x|^2 = p(x) \tag{2}$$

This encoding allows parallel access to all probability amplitudes via quantum superposition. Measuring the quantum state yields samples according to the probability distribution $p(x)$, and repeated measurement allows estimation of statistical moments.

The objective of this work is to implement the Bayesian filtering update on such a quantum state by transforming an initial distribution (prior) into the posterior distribution through quantum evolution. In contrast to gate-based methods, we investigate an adiabatic quantum approach that leverages energy-based encoding to shape the final state distribution according to the desired posterior. The challenge lies in embedding the likelihood information into a cost function that can be realized through a time-dependent quantum Hamiltonian, while maintaining physical feasibility and computational scalability.

3 ADIABATIC QUANTUM COMPUTING

Adiabatic Quantum Computing (AQC) is a computational paradigm that solves optimization problems by exploiting the adiabatic theorem of quantum mechanics. The adiabatic theorem was first stated by Born and Fock in 1928 [1]; it asserts that a quantum system remains in its instantaneous eigenstate if the Hamiltonian changes sufficiently slowly and a finite spectral gap exists. Since then, the theorem has been analyzed and extended in a range of computational contexts [6]. Instead of executing a sequence of logic gates, AQC encodes the solution to a given problem into the ground state of a final, problem-specific Hamiltonian. The system is initialized in the ground state of a simple, easily preparable Hamiltonian and then evolved slowly toward the problem Hamiltonian. If the evolution is sufficiently slow and the system remains isolated, the adiabatic theorem guarantees that the system stays in its instantaneous ground state throughout the process.

Formally, the total Hamiltonian of the system at time $t \in [0, T]$ is defined as an interpolation between two Hamiltonians:

$$H(t) = (1 - s(t))H_0 + s(t)H_1, \quad s(0) = 0, \quad s(T) = 1 \tag{3}$$

Here, H_0 is the initial (driver) Hamiltonian, whose ground state is typically a uniform superposition over all computational basis states. A common choice is a transverse-field Hamiltonian, such as:

$$H_0 = -\sum_i X_i \tag{4}$$

where X_i denotes the Pauli-X operator acting on qubit i. The problem Hamiltonian H_1 encodes the cost function of the optimization problem in its ground state. For many problems of practical interest, H_1 can be expressed in the form of an Ising Hamiltonian:

$$H_1 = \sum_{i<j} J_{ij} Z_i Z_j + \sum_i h_i Z_i \tag{5}$$

where Z_i is the Pauli-Z operator on qubit i, and J_{ij}, h_i are problem-specific coefficients derived from the structure of the cost function. [4]

The system is evolved according to a time-dependent schedule $s(t)$, often chosen to be linear: $s(t) = t/T$. The total evolution time T must be large enough to ensure adiabaticity, i.e., to avoid excitations to higher energy levels. The minimum energy gap Δ between the ground state and the first excited state during the evolution determines the required runtime via the relation:

$$T \gg \frac{1}{\Delta^2} \tag{6}$$

In practice, small violations of adiabaticity lead to a final state that is not exactly the ground state of H_1, but rather a thermal distribution over low-energy states. This behavior can still be useful, particularly when the objective is to sample from a distribution shaped by the cost landscape rather than to find a single minimum.

To use AQC for a computational problem, one must translate the objective function into the energy landscape defined by H_1. In our case, this objective is the negative logarithm of the posterior distribution obtained from Bayesian filtering. For many models with Gaussian or quadratic structure, this results in a cost function that is quadratic in the discrete state variable x. This function must then be mapped to binary variables and reformulated as an Ising-type Hamiltonian, suitable for implementation on quantum annealers or other AQC platforms.

4 MAPPING THE BAYESIAN UPDATE TO AN ISING HAMILTONIAN

In order to implement the Bayesian update step within an adiabatic quantum computing (AQC) framework, the corresponding posterior distribution must be encoded into the ground state of a problem Hamiltonian. This requires the formulation of the Bayesian cost function in a quadratic form suitable for mapping to an Ising Hamiltonian.

4.1 Cost Function Formulation

The Bayesian filtering step corresponds to a pointwise multiplication of the prior and the likelihood function, followed by normalization. In log-domain and for Gaussian models, this operation becomes additive. The resulting cost function can be written as:

$$C(x) = P(x - x_\mu)^2 + Q(z - Ax)^2 \tag{7}$$

Here, x is the discrete latent variable, x_μ is the prior mean, z is the measurement, A is a linear mapping (e.g., a measurement matrix), and P, Q are weighting factors representing the inverse variances of the prior and likelihood distributions, respectively.

Expanding this cost function yields a quadratic polynomial:

$$\begin{aligned} C(x) &= P(x^2 - 2x x_\mu + x_\mu^2) + Q(z^2 - 2Axz + A^2x^2) \\ &= (P + QA^2)x^2 - 2(Px_\mu + QAz)x + (Px_\mu^2 + Qz^2) \end{aligned} \tag{8}$$

Defining the constants:

$$J = P + QA^2, \quad h = -2(Px_\mu + QAz), \quad C = Px_\mu^2 + Qz^2 \tag{9}$$

the cost function becomes:

$$C(x) = Jx^2 + hx + C \tag{10}$$

which is directly compatible with the energy form required for Ising-type Hamiltonians.

4.2 Ising Representation via Operator Mapping

To translate the above polynomial into a quantum operator acting on a register of qubits, we first define a binary operator that maps computational basis states to integers. For N qubits, let the operator $\tilde{Z}_i$ be:

$$\tilde{Z}_i = \frac{1}{2}(I - Z_i) \tag{11}$$

where Z_i is the Pauli-Z operator acting on qubit i. This maps:

$$\tilde{Z}_i |0\rangle = 0, \quad \tilde{Z}_i |1\rangle = |1\rangle \tag{12}$$

Each computational basis state $|b_0 b_1 \dots b_{N-1}\rangle$, with $b_k \in \{0,1\}$, represents a binary number

$x \in \{0, \ldots, 2^N - 1\}$. We construct the corresponding value operator X as:

$$X = \sum_{k=0}^{N-1} 2^{N-1-k} \tilde{Z}_k \tag{13}$$

This operator acts diagonally in the computational basis and returns the classical integer x encoded by the qubit register.

4.3 Problem Hamiltonian

Replacing x by the operator X in the cost function leads to the problem Hamiltonian:

$$H_p = JX^2 + hX + CI \tag{14}$$

This Hamiltonian is diagonal in the computational basis, and its ground state corresponds to the state $|x\rangle$ minimizing the cost function:

$$H_p |x\rangle = \left(Jx^2 + hx + C\right) |x\rangle \tag{15}$$

The offset term C is constant and does not affect the ground state. Thus, the optimization reduces to finding the value of x that minimizes the quadratic polynomial encoded in H_p, which corresponds to the Bayesian posterior mode under the assumed Gaussian prior and likelihood.

This formulation enables the implementation of the Bayesian update step as an energy minimization problem, suitable for execution on an adiabatic quantum computer.

5 THERMAL EFFECTS AND EFFECTIVE TEMPERATURE

In adiabatic quantum computing (AQC), the system ideally remains in the ground state of a slowly evolving Hamiltonian. However, in practice, non-adiabatic transitions may occur due to finite evolution time, especially near avoided level crossings. As a result, the final quantum state does not necessarily coincide with the exact ground state but instead resembles a thermal mixture over low-energy states. This effect can be modeled using a Gibbs distribution governed by an effective temperature. [7]

Formally, the probability of observing a state x with energy $E(x)$ is given by:

$$P(x) = \frac{1}{Z} \exp(-\beta_{\text{eff}} E(x)), \tag{16}$$

$$Z = \sum_x \exp(-\beta_{\text{eff}} E(x)) \tag{17}$$

Here, $\beta_{\text{eff}} = 1/T_{\text{eff}}$ denotes the inverse effective temperature that characterizes the deviation from ideal adiabaticity. The key challenge in AQC-based sampling is that this parameter is not externally controllable, but rather emerges from the physical dynamics of the system.

To derive the dependence of β_{eff} on system parameters, consider a two-level quantum system with a time-dependent Hamiltonian:

$$H(t) = \frac{\Delta}{2}\sigma_x + \frac{vt}{2}\sigma_z \tag{18}$$

where Δ is the minimum spectral gap and v is the rate at which the diabatic levels sweep past each other. According to Landau-Zener theory, the excitation probability to the first excited state is:

$$P_{\text{exc}} = \exp\left(-\frac{\pi\Delta^2}{2\hbar v}\right) \tag{19}$$

as shown in the derivation by Sun [8].
Assuming the ground and excited state energies are $E_0 = 0$ and $E_1 = \Delta$, and defining $p_0 = 1 - P_{\text{exc}}$, $p_1 = P_{\text{exc}}$, we compare the resulting state occupations to a thermal distribution:

$$\frac{p_1}{p_0} = \exp(-\beta_{\text{eff}}\Delta) \tag{20}$$

Solving for β_{eff} yields:

$$\beta_{\text{eff}} = \frac{1}{\Delta} \ln\left(\frac{1}{P_{\text{exc}}} - 1\right) = \frac{1}{\Delta} \ln\left(\exp\left(\frac{\pi\Delta^2}{2\hbar v}\right) - 1\right) \tag{21}$$

For small P_{exc} (i.e., in the adiabatic regime), we can approximate:

$$\ln(\exp(A) - 1) \approx A \quad \Rightarrow \quad \beta_{\text{eff}} \approx \frac{\pi\Delta}{2\hbar v} \tag{22}$$

This results in the proportionality:

$$\beta_{\text{eff}} \propto \frac{\Delta^2}{v} \tag{23}$$

In most practical implementations, the Hamiltonian is interpolated via a schedule $s(t)$, with $s(0) = 0$, $s(T) = 1$. For a linear ramp $s(t) = t/T$, we have $v \propto \dot{s}(t) = 1/T$, leading to:

$$\beta_{\text{eff}} \propto \Delta^2 T \tag{24}$$

This fundamental relationship illustrates that slower sweeps (larger T) yield lower effective temperatures and distributions that are more sharply peaked around the ground state. In the context of Bayesian filtering, this directly controls the concentration of the posterior approximation. Consequently, β_{eff} plays a critical role in tuning the quality and fidelity of the resulting distribution, and must be taken into account when designing quantum filtering procedures.

6 CALIBRATION OF THE ENERGY SCALE

As shown in the previous section, the effective inverse temperature β_{eff} determines how sharply the final quantum state concentrates near the ground state. In the context of Bayesian filtering, this affects the shape of the posterior distribution obtained from a quantum annealing. If $\beta_{\text{eff}} \neq 1$, the resulting distribution deviates from the intended target density. To correct for this deviation, we introduce a scalar energy scaling factor α such that the product $\alpha \cdot \beta_{\text{eff}} = 1$, restoring the desired statistical form.

6.1 Motivation and Objective

To illustrate the effect of the effective inverse temperature on the resulting distribution, we assume that both the prior and the likelihood are Gaussian. In this case, the energy function used in the quantum annealing process corresponds to the negative logarithm of Gaussian density itself. The quantum system then evolves into a mixed state whose measurement statistics approximate a Boltzmann distribution with energy $E(x)$ and inverse temperature β_{eff}. As a result, the resulting distribution is:

$$P_{\text{sim}}(x) \approx \frac{1}{Z_{\text{eff}}} \exp(-\beta_{\text{eff}} E(x)) \tag{25}$$

For quadratic energies of the form $E(x) = \frac{(x-\mu)^2}{2\sigma^2}$, this becomes:

$$P_{\text{sim}}(x) \propto \exp\left(-\frac{\beta_{\text{eff}}}{2\sigma^2}(x-\mu)^2\right) \tag{26}$$

If $\beta_{\text{eff}} > 1$, the distribution is narrower than intended; if $\beta_{\text{eff}} < 1$, it is broader. To ensure that the quantum simulation reflects the desired distribution shape, we rescale the energy function prior to execution:

$$E'(x) = \alpha E(x) \tag{27}$$

Choosing $\alpha = 1/\beta_{\text{eff}}$ yields:

$$P_{\text{sim}}(x) \propto \exp\left(-\beta_{\text{eff}} \cdot \alpha E(x)\right) = \exp(-E(x)) \tag{28}$$

This correction aligns the quantum-measured distribution with the original posterior target.

6.2 Estimating β_{eff}

To determine the appropriate scaling factor α, we must first estimate β_{eff} empirically. This can be accomplished through histogram analysis of the measured quantum output. The procedure is as follows:

1. Run the quantum sweep using the unscaled Hamiltonian ($\alpha = 1$).
2. Measure the resulting distribution $P_{\text{sim}}(x)$.
3. Perform a logarithmic transformation:
$$\ln P_{\text{sim}}(x) = c - \beta_{\text{eff}} E(x) \tag{29}$$
where $c = -\ln Z_{\text{eff}}$ is an unknown constant offset.
4. Fit a straight line to the data $\{x_i, \ln P_{\text{sim}}(x_i)\}$ using least squares. The slope yields an estimate of $-\beta_{\text{eff}}$.

Only data points with $P_{\text{sim}}(x) > 0$ are considered in the fit to avoid singularities.

Once β_{eff} has been estimated, we define the energy scaling factor as:

$$\alpha = \frac{1}{\beta_{\text{eff}}} \tag{30}$$

This scaling is then applied to the Ising Hamiltonian before the next quantum run. In practice, this calibration step ensures that the simulated posterior matches the desired Bayesian update distribution in both shape and concentration. The correction is particularly important for accurately capturing uncertainties in probabilistic inference, where even slight distortions in variance can significantly affect the interpretation of the result.

7 NUMERICAL EVALUATION

To assess the performance of our adiabatic quantum filtering method, we compare the output distribution of the AQC-based approach with both the exact Bayesian posterior and an alternative quantum method referred to as *Quantum Flow*, which is based on a gate-based

realization of the posterior update in proposed this work [3].

Figure 1 shows the resulting probability distributions for a concrete one-dimensional filtering example. The prior distribution is Gaussian with mean $\mu = 9$ and standard deviation $\sigma = 3$. The likelihood is centered at $z = 17$ with $\sigma = 4$. The exact posterior is then analytically given by a Gaussian with mean $\mu = 11.880$ and $\sigma = 2.4$.

Our AQC-based approach was executed with an energy scaling factor $\alpha = 0.33$, yielding an empirical distribution with mean $\mu = 12.175$ and standard deviation $\sigma = 2.116$. The Quantum Flow approach produces a slightly lower mean $\mu = 10.53$ while matching the posterior width exactly with $\sigma = 2.4$.

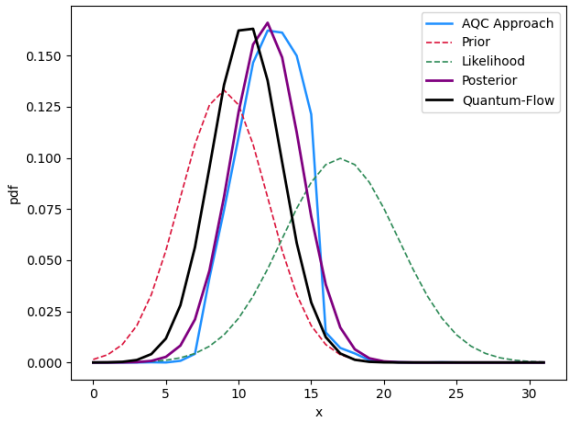

Figure 1: Comparison of the output distribution from the AQC-based approach and the gate-based Quantum Flow method against the analytical posterior.

Method	Mean	Std
Posterior (analytical)	11.880	2.400
Quantum Flow	10.530	2.400
AQC Approach	12.175	2.116

Table 1: Comparison of the mean and standard deviation of the posterior approximation obtained by different quantum methods.

Both quantum methods closely approximate the correct posterior distribution. The AQC-based solution slightly overestimates the posterior mean but provides a sharper distribution. The Quantum Flow approach matches the target standard deviation but slightly underestimates the mean. These results highlight that the proposed adiabatic method can yield accurate posterior approximations from raw quantum measurement statistics without explicit gate synthesis of the update rule.

It is important to emphasize that the energy scaling factor $\alpha = 0.33$ used in this example was not fine-tuned, but rather estimated heuristically based on a single calibration run. Despite this lack of optimization, the resulting AQC distribution matches the shape and location of the posterior with reasonably high accuracy.

This indicates that the proposed adiabatic approach is robust under approximate calibration and does not require precise knowledge of the effective temperature a priori. Further improvement could be achieved by systematically analyzing the energy gap Δ along the annealing path and its influence on the effective temperature. This would allow for targeted adjustment of the ramp duration or energy scaling to better align the resulting distribution with the theoretical posterior.

8 CONCLUSION

We have presented a novel approach to the Bayesian filtering update using adiabatic quantum computing. By encoding the likelihood into an Ising-type Hamiltonian, the posterior distribution emerges from the final quantum state after annealing. Numerical results demonstrate that even without fine-tuning, the resulting distributions approximate the target posterior with a close match. This suggests that AQC offers a scalable and physically grounded alternative to gate-based implementations of probabilistic inference. Future work may explore optimization of annealing schedules, analysis of spectral gaps, and hardware implementation on existing quantum annealers.

9 REFERENCES

[1] M. Born and V. Fock. Beweis des adiabatensatzes. *Zeitschrift für Physik*, 51(3–4):165–180, 1928.

[2] F. Govaers. *Theory and Methods for Distributed Data Fusion Applications*. Radar, Sonar and Navigation. The Institution of Engineering and Technology (IET), 2023.

[3] F. Govaers. On a quantum realization of the bayesian filtering using the log homotopy flow. In *Proceedings of the 2024 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, pages 1–6, 2024.

[4] T. Kadowaki and H. Nishimori. Quantum annealing in the transverse ising model. *Physical Review E*, 58(5):5355–5363, 1998.

[5] W. Koch. *Tracking and Sensor Data Fusion: Methodological Framework and Selected Applications*. Mathematical Engineering. Springer, 2013.

[6] C. C. McGeoch. *Adiabatic Quantum Computation and Quantum Annealing: Theory and Practice*. Morgan & Claypool Publishers, 2014.

[7] A. Perdomo-Ortiz, M.-H. Yung, et al. High-quality thermal gibbs sampling with quantum annealing hardware. *Physical Review Applied*, 17(4):044046, 2022.

[8] C. Sun. Derivation of the landau-zener formula via functional equations, 2025.