Scalable Quantum Research with Hierarchical Tensor Methods

Nahid Binandeh Dehaghani

Research Mentor: Prof. Rafal Wisniewski

Department of Electronic Systems
Aalborg University

DeiC Conference 2026
27–28 October 2026, Kolding, Denmark

My Research: Quantum Control and Computation

Quantum Control and Optimization

$\Psi_0 \longrightarrow \Psi_{\text{target}}$

Designing control strategies to steer quantum systems toward desired states and operations.

Quantum and Hybrid Algorithms

Quantum + Classical

Investigating quantum and hybrid quantum-classical methods for optimization and learning.

Scalable Quantum Modeling

Full $\longrightarrow$ Reduced

Developing compact representations of quantum systems with mathematical accuracy and control-performance guarantees.

Common Research Challenge

How can we make quantum computation and control **scalable, accurate, and reliable**?

Today's focus: Hierarchical Tucker methods for quantum optimal control.

The Scalability Challenge

Exponential Growth of Quantum Systems

An n-qubit pure quantum state is represented by

$$|\Psi\rangle = \sum_{j=0}^{2^n-1} c_j |j\rangle, \qquad c_j \in \mathbb{C}.$$

The number of complex amplitudes grows as 2^n.

System Size

Qubits	Amplitudes
10	1,024
20	1,048,576
30	1,073,741,824
40	1,099,511,627,776

Why Does This Matter?

- Large memory requirements.
- Expensive numerical simulation.
- Repeated state propagation during control optimization.

Research Question

Can we represent and control large quantum systems using fewer parameters while preserving accuracy?

Why This Matters for Digital Research Infrastructure

Computational Resources

High-dimensional quantum simulations create growing demands for memory and computational capacity.

Scalable Numerical Methods

Compact tensor representations offer a potential route to handling larger quantum models with fewer representation parameters.

Reliable Scientific Results

Mathematical guarantees help us understand how much we can trust approximate simulations and the resulting control decisions.

Developing scalable and trustworthy numerical methods can strengthen computational tools supporting emerging quantum technologies.

Tensor Networks for Scalable Quantum Modeling

What Are Tensor Network Methods?

Tensor network methods provide structured ways to represent high-dimensional data using smaller, interconnected tensors.
They are widely studied in quantum many-body physics and numerical simulation.

Why Are They Useful?

- Exploit underlying structure in high-dimensional quantum states.
- Can reduce representation complexity when suitable low-rank structure exists.
- Offer different representations for different computational problems.

Hierarchical Tucker (HT)

Hierarchical Tucker is one such representation. It organizes tensor components hierarchically and approximates them using low-rank factors.

Our Research Focus: Investigating HT approximations of quantum dynamics with mathematical guarantees on accuracy and control performance.

Hierarchical Tucker: A Compact Representation

From Quantum States to Tensors

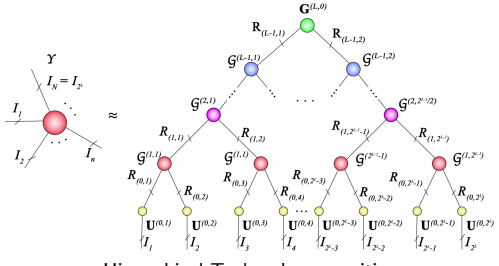
A quantum state can be reorganized into a multidimensional tensor:

$$\Psi \longleftrightarrow \mathcal{X} \in \mathbb{C}^{d_1 \times \cdots \times d_n}.$$

The Hierarchical Tucker Idea

- Represent a high-dimensional tensor using smaller components.
- Organize and combine these components hierarchically.
- Use low-rank approximations to retain the most important information.

Key parameter: HT rank controls the retained information and representation size.



Hierarchical Tucker decomposition

Reference: Fonał, K., Zdunek, R. (2019). Fast recursive nonnegative standard and hierarchical Tucker decomposition. IEEE Signal Processing Letters, 26(9).

For quantum states with suitable low-rank structure, HT can reduce the number of representation parameters while controlling approximation error.

From Compression to Certification

High-dimensional quantum system

Exponentially large state representation

$$\Downarrow$$

Hierarchical Tucker surrogate

Compact low-rank representation

But compression alone is not enough.

We must quantify how approximation affects the original quantum control problem.

$\|\Psi - \Psi^{(r)}\|$	$\|J - J_r\|$	$J(x_0, u_r^+) - V(x_0)$
Trajectory Accuracy	**Cost Accuracy**	**Control Performance**
How closely does the surrogate reproduce the quantum dynamics?	How accurately does the surrogate evaluate the control objective?	How close is the HT-designed control to optimal performance?

Our objective: Compact quantum representations with quantifiable approximation errors and certified control-performance guarantees.

Quantum Dynamics and Optimal Control

Controlled Quantum Dynamics

$$i\hbar \frac{d}{dt}\Psi(t) = \big(H_0 + u(t)H_1\big)\Psi(t)$$

H_0: natural dynamics $\qquad$ $u(t)H_1$: external control

Steering a Quantum System

Initial State Ψ_0

Control signal $u(t)$ $\longrightarrow$ Quantum evolution

Target State Ψ_{target}

Choose the control to drive the system toward a desired state or operation.

Optimal Control Objective

$$\min_{u(\cdot)} \left[\Phi\big(\Psi(T)\big) + \int_0^T L\big(\Psi(t), u(t)\big)\, dt \right]$$

Balance the final objective and the cost of control.

Quantum Dynamics with HT Compression

Full-Order Quantum Dynamics

For piecewise-constant controls, the state evolves as

$$\Psi_{k+1} = F_{\Delta t}(\Psi_k, u_k),$$

where $F_{\Delta t}$ denotes the exact time-step propagation under the applied control.

Hierarchical Tucker Surrogate

After each propagation step, apply HT truncation to retain a prescribed hierarchical rank:

$$\Psi_{k+1}^{(r)} = \Pi_r \left(F_{\Delta t}(\Psi_k^{(r)}, u_k) \right), \qquad \Psi_0^{(r)} = \Pi_r(x_0), \qquad k = 0, \ldots, N_T - 1.$$

Here, Π_r denotes the HT truncation operator.

The surrogate combines quantum-state propagation with rank-constrained HT approximation, enabling the study of accuracy and control performance as functions of the retained rank.

When Does HT Compression Work?

Key Assumption: Hierarchical Compressibility

HT compression works well when the most important information can be captured using relatively few components.

Mathematically, we assume that the hierarchical singular values decrease exponentially: $\boxed{\sigma_j^{(t)}(\Psi) \leq C_1 e^{-cj}}$

Here, j is the singular-value index and t denotes a non-root node of the dimension tree.

Rank-Dependent Approximation Accuracy

Under this assumption, the HT approximation admits an exponentially decreasing error bound:

$$\boxed{\|\Psi - \Pi_r(\Psi)\| \leq C_2 e^{-c' r}}$$

The HT rank r controls how much information is retained in the approximation.

For sufficiently compressible quantum states, **higher HT ranks lead to smaller theoretical approximation-error bounds**.

Trajectory and control guarantees additionally require uniform truncation-error bounds along the surrogate dynamics.

From Local Errors to Global Accuracy

How Do Errors Accumulate?

Each HT compression step introduces a small approximation error.

$$\boxed{\delta_{k+1} \leq \delta_k + \|e_k^{(r)}\|}$$

δ_k is the trajectory approximation error, and $e_k^{(r)}$ is the new truncation error introduced at step k.

For unitary quantum evolution, existing errors are not amplified by propagation itself.

Result: Local truncation errors accumulate additively.

Accuracy Over the Full Trajectory

Under suitable uniform HT approximation assumptions,

$$\boxed{\max_{0 \leq k \leq N_T} \|\Psi_k - \Psi_k^{(r)}\| \leq C_T e^{-\bar{c}r}}$$

r is the retained HT rank, and $C_T, \bar{c} > 0$ are constants independent of r for the fixed problem.

Result: Increasing the HT rank reduces the theoretical trajectory-error bound exponentially.

For a fixed simulation horizon and time discretization, we can bound the approximation error over the entire quantum trajectory.
The bound decreases exponentially with HT rank, provided the required approximation assumptions hold.

From Trajectory Accuracy to Cost Accuracy

Full-Order Control Cost

$$J(x_0, \mathrm{u}) = \sum_{k=0}^{N_T-1} L(\Psi_k, u_k) + \Phi(\Psi_{N_T})$$

HT-Surrogate Control Cost

$$J_r(x_0, \mathrm{u}) = \sum_{k=0}^{N_T-1} L(\Psi_k^{(r)}, u_k) + \Phi(\Psi_{N_T}^{(r)})$$

Both objectives evaluate the **same control sequence**, using different state trajectories.

Theoretical Guarantee

Under Lipschitz cost regularity and uniform HT truncation accuracy,

$$\boxed{|J(x_0, \mathrm{u}) - J_r(x_0, \mathrm{u})| \leq C_J e^{-\bar{c}r}}$$

$$\forall\, \mathrm{u} \in \mathcal{U}^{N_T}.$$

The approximation bound holds **uniformly over all admissible control sequences**.

Consequently, the surrogate objective approaches the full-order objective as the retained rank increases, in the sense of the stated error bound.

Exponential trajectory-error bounds lead to exponential cost-error bounds, providing the foundation for certified control performance.

Certified Near-Optimality of HT-Based Control

The Main Question

If we design a control using the simplified HT model, how well will it perform on the original quantum system?

1. Optimal-Value Accuracy

The difference between the optimal costs of the original and HT models is bounded by

$$\boxed{|V(x_0) - V_r(x_0)| \leq C_J e^{-\bar{c}r}}$$

The theoretical optimal-value error bound decreases exponentially with HT rank.

2. Control-Performance Guarantee

For a globally optimal HT-surrogate control $u_r^\star$,

$$\boxed{J(x_0, u_r^\star) - V(x_0) \leq 2C_J e^{-\bar{c}r}}$$

When evaluated on the original quantum system, the HT-designed control has a bounded optimality gap.

Under suitable approximation assumptions and global surrogate optimality, we establish **rank-dependent guarantees** on control performance for the original quantum system.

From Accuracy to Representation Complexity

Full-State vs. HT Representation

For bounded local dimensions and hierarchical ranks at most r, the HT representation has a coefficient count of

$$\underbrace{2^n}_{\text{Full state}} \longrightarrow \underbrace{O(nr^3)}_{\text{HT representation}}$$

Connecting Accuracy to Representation Size

The HT rank controls the trade-off between representation size and approximation accuracy.
Under our theoretical assumptions, a sufficient rank for a desired control-performance tolerance ε can be chosen to scale as $r = O\left(\log \frac{1}{\varepsilon}\right)$.

Combining these estimates gives a sufficient HT representation coefficient-count bound: $\boxed{N_{\mathrm{HT}} = O\left(n \log^3 \frac{1}{\varepsilon}\right)}$

For sufficiently compressible quantum states, HT can provide a substantially smaller representation with controlled approximation accuracy.

Conditional on uniform rank-dependent approximation assumptions and globally optimal surrogate controls.

Numerical Validation: XXZ Spin Chain

We evaluate our approach using a controlled XXZ spin chain, a model of interacting quantum spins.

- 8 spins, corresponding to 256 complex state amplitudes.
- Quantum evolution over 15 time steps.
- Bounded external control signals.

Optimal Control Objective

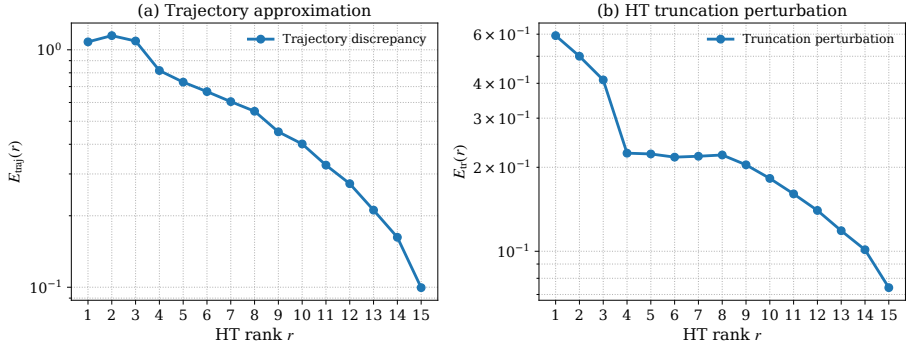
Our goal is to drive the quantum system toward a desired target state while limiting the control effort.

$$J(x_0, \mathrm{u}) = 1 - |\langle \Psi_{\mathrm{tar}}, \Psi_{N_T} \rangle|^2 + \rho \Delta t \sum_{k=0}^{N_T - 1} u_k^2$$

Numerical Evaluation

- Optimize controls using the HT surrogate.
- Evaluate HT-designed controls on the original full-order dynamics.
- Compare with an empirical best-found reference.

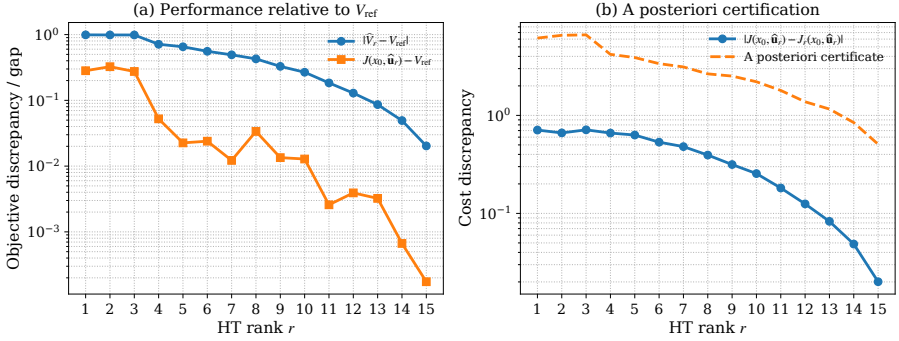
How Does HT Rank Affect Accuracy?



(a) Difference between the original and HT-approximated quantum trajectories.
(b) Error introduced by HT truncation at each rank.

Both approximation errors generally decrease as the retained HT rank increases.
At $r = 16$ (not shown), the errors reach numerical precision.

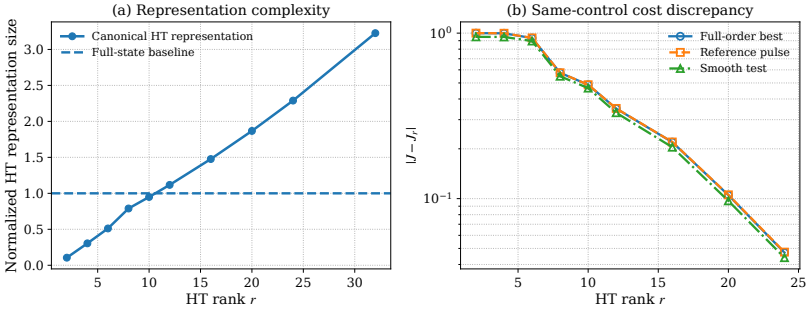
Does HT-Based Control Perform Well?



(a) Control-performance gaps relative to an empirical reference.
(b) Observed cost approximation errors and their mathematical upper bounds.

Full-order control-performance gaps generally decrease as the HT rank increases.
Trajectory-specific error bounds also control the observed cost discrepancies.

Balancing Representation Size and Accuracy



10-spin quantum system; three fixed admissible control signals.
(a) HT coefficient count compared with the full-state size.
(b) Cost approximation error for different HT ranks.

HT uses fewer representation coefficients than the full state through $r = 10$ in this experiment.
Increasing rank generally improves accuracy but also increases representation size.

Conclusions and Future Directions

Main Contributions

- Developed HT-based quantum control with **mathematical accuracy guarantees**.
- Established **rank-dependent near-optimality bounds** under suitable assumptions.
- Demonstrated the **representation–accuracy trade-off** through numerical experiments.

Why HT Matters for Quantum Computing

- **Compact quantum modeling:** Represent suitable multi-qubit states using fewer coefficients.
- **Reliable control design:** Quantify how approximation errors affect control performance.
- **Quantum gate synthesis:** Explore compressed simulation and optimization of quantum operations.

Takeaway: HT offers a promising framework for scalable and accuracy-aware quantum control, with potential applications in quantum computing.

References

1. N. Binandeh Dehaghani, R. Wisniewski, and A. P. Aguiar,
   "Certified Quantum Schrödinger Control via Hierarchical Tucker Models,"
   IEEE Control Systems Letters, vol. 10, pp. 1453–1458, 2026.

2. N. Binandeh Dehaghani, R. Wisniewski, J. Berberich, and A. P. Aguiar,
   "Rank-Dependent Error Bounds and Near-Optimality in Quantum Control via Hierarchical Tucker Surrogates,"
   arXiv preprint arXiv:2609.14686, 2026.

Thank You!

Questions and Discussion

Nahid Binandeh Dehaghani

Aalborg University, Denmark

nahidbd@es.aau.dk

Acknowledgment

I gratefully acknowledge the financial support of DeiC for my postdoctoral research.