

Quantum thermodynamics and semidefinite optimization: Boltzmann, Fermi–Dirac, and Bose–Einstein frameworks

Michele Minervini,¹ Nana Liu,² Dhruvil Patel,³ and Mark M. Wilde¹

¹*School of Electrical and Computer Engineering,
Cornell University, Ithaca, New York 14850, USA*

²*Institute of Natural Sciences, School of Mathematical Sciences,
Ministry of Education Key Laboratory in Scientific and Engineering Computing,
and Global College, Shanghai Jiao Tong University, Shanghai 200240, China*

³*School of Computer Science, Cornell University, Ithaca, New York 14850, USA*

We show how quantum thermodynamics offers a unifying interpretation for a wide class of semidefinite programs (SDPs) that arise in quantum information. Three SDP variable constraints, namely trace-one density operators, operator-bounded measurements, and the unbounded positive semidefinite cone, admit thermodynamic regularizations associated with Boltzmann, Fermi–Dirac, and Bose–Einstein statistics, respectively. In each case the entropy-regularized dual is unconstrained and concave in a chemical-potential vector, and the primal optimum is a thermal operator of the matched statistics. The dual gradient and Hessian are thermal expectation values, enabling hybrid quantum–classical algorithms for solving a wide variety of SDPs. Building on the Boltzmann framework presented at QTML 2025 [1], this extended abstract presents a unified picture of quantum thermodynamics and SDPs that incorporates the more recent Fermi–Dirac [2] and Bose–Einstein [3] works, which natively address measurement-bounded SDPs (and in particular quantum hypothesis testing) and standard-form SDPs over the unbounded positive semidefinite cone, respectively.

Keywords: semidefinite programming, quantum thermodynamics, hybrid quantum–classical algorithms, quantum hypothesis testing, quantum Boltzmann machines

Introduction—Many fundamental problems in quantum information theory, such as state discrimination, channel capacities, entanglement quantification, and quantum hypothesis testing, can be cast as semidefinite programs (SDPs) [4]. Efficient SDP solvers are therefore a basic tool for quantum information research, with quantum-native solvers especially attractive when the SDP dimension grows exponentially with the underlying system size. In each of these contexts, the SDP variable belongs to a different set, dictated by which constituent of quantum mechanics is being optimized: density operators $\rho \in \mathcal{D}_d$, where $\mathcal{D}_d$ is the set of $d \times d$ density operators, for states; measurement operators $M \in \mathcal{M}_d$ satisfying $0 \preceq M \preceq I$, where $\mathcal{M}_d$ is the set of $d \times d$ measurement operators, for positive operator-valued measure (POVM) elements; and the unbounded positive semidefinite cone $\mathcal{P}_d := \{X \succeq 0\}$ for unnormalized observables. Here we observe that the eigenvalues of operators in these three sets obey precisely the three rules of the corresponding quantum statistics: trace-one normalization for Boltzmann, the Pauli exclusion principle $0 \leq m_j \leq 1$ for Fermi–Dirac, and unrestricted occupation numbers $x_j \geq 0$ for Bose–Einstein.

Drawing on this observation, our line of work develops thermodynamic frameworks for SDPs across these three distributions [1–3]. Adopting Jaynes’ statistical-mechanics mindset [5], in each framework the linear SDP objective is interpreted as a system energy, and the optimization is recast as a free-energy minimization at a strictly positive temperature $T > 0$, by regularizing the energy with a statistics-matched quantum entropy scaled by T .

The three frameworks therefore share a common thermodynamic skeleton. From this primal free-energy minimization, Lagrangian duality combined with a statistics-matched quantum relative entropy produces an unconstrained concave dual in a chemical-potential vector, whose unique primal optimum is a thermal operator of the matched statistics. The dual gradient and Hessian admit closed-form expressions as thermal expectation values, amenable to estimation by hybrid quantum–classical algorithms based on similar primitives. The three frameworks differ in the specific entropy involved, in the thermal-operator form, and consequently in the SDP families that they natively address. Taken together, they assemble a cogent thermodynamic interpretation of SDPs in quantum information, paired with a corresponding family of novel hybrid quantum–classical solvers.

Unified thermodynamic framework—For $c \in \mathbb{N}$, a general SDP is specified by Hermitian operators $\mathcal{Q} = (H, Q_1, \dots, Q_c)$ on a d -dimensional Hilbert space and constraint values $q = (q_1, \dots, q_c) \in \mathbb{R}^c$. The variable is constrained to a feasible set $\mathcal{C} \in \{\mathcal{D}_d, \mathcal{M}_d, \mathcal{P}_d\}$ matching one of the three quantum sets introduced above (see the first row of Table I). The SDP optimizes the linear objective $\text{Tr}[HX]$ over $X \in \mathcal{C}$, subject to the affine constraints $\text{Tr}[Q_i X] = q_i$ for all $i \in [c]$:

$$E(\mathcal{Q}, q) = \min_{X \in \mathcal{C}} \{ \text{Tr}[HX] : \text{Tr}[Q_i X] = q_i \forall i \in [c] \}. \quad (1)$$

Physically, we can consider H to be a Hamiltonian, each Q_i to be a non-commuting charge, and the eigenvalues of X play the role of expected occupation numbers under the matched statistics.

	Boltzmann [1]	Fermi–Dirac [2]	Bose–Einstein [3]
Set	$\rho \succeq 0, \text{Tr}[\rho] = 1$	$0 \preceq M \preceq I$	$X \succeq 0$
Eigenvalue rule	trace-one probabilities	Pauli exclusion: $0 \leq m_j \leq 1$	unrestricted occupations: $x_j \geq 0$
Entropy	$S(\rho) := -\text{Tr}[\rho \ln \rho]$	$S_{\text{FD}}(M) := -\text{Tr}[M \ln M + (I - M) \ln(I - M)]$	$S_{\text{BE}}(X) := -\text{Tr}[X \ln X - (X + I) \ln(X + I)]$
Relative entropy	$D(\rho \parallel \sigma) := \text{Tr}[\rho(\ln \rho - \ln \sigma)]$	$D_{\text{FD}}(M \parallel N) := D(M \parallel N) + D(I - M \parallel I - N)$	$D_{\text{BE}}(X \parallel Y) := D(X \parallel Y) - D(X + I \parallel Y + I)$
Primal optimum	$\rho_T(\mu) := \frac{e^{-K_\mu/T}}{Z_T(\mu)}$	$M_T(\mu) := (e^{K_\mu/T} + I)^{-1}$	$X_T(\mu) := (e^{K_\mu/T} - I)^{-1}$
Optimum’s name	Boltzmann (Gibbs) state	Fermi–Dirac thermal measurement	Bose–Einstein thermal operator
Target SDP	trace-normalized	operator-bounded	standard form; unbounded PSD cone

TABLE I. The three thermodynamic frameworks share a common skeleton, namely free-energy regularization, an unconstrained concave dual in the chemical-potential vector μ , and a primal optimum that is a thermal operator parametrized by μ through the grand canonical Hamiltonian $K_\mu = H - \mu \cdot Q$. They differ in the feasible set, the entropy that regularizes the SDP, and the resulting thermal-operator form; the framework-specific quantum primitives are discussed in the main text.

Inspired by the fact that physical systems operate at a strictly positive temperature, we minimize the free energy rather than the energy, obtaining a smooth approximation to the original SDP:

$$F_T(\mathcal{Q}, q) = \min_{X \in \mathcal{C}} \{ \text{Tr}[HX] - T S_s(X) : \text{Tr}[Q_i X] = q_i \} \quad (2)$$

at temperature $T > 0$. Here the index $s \in \{\text{B}, \text{FD}, \text{BE}\}$ runs over the three kinds of statistics (Boltzmann, Fermi–Dirac, Bose–Einstein), and S_s is the corresponding entropy: the von Neumann entropy for Boltzmann, the Fermi–Dirac entropy for the operator-bounded set, or the Bose–Einstein entropy for the unbounded cone, whose explicit expressions are reported in Table I.

Lagrangian duality, combined with a quantum relative entropy adapted to each set (Umegaki for Boltzmann, Fermi–Dirac for the operator-bounded case, Bose–Einstein for the unbounded one; see Table I for the expressions), yields an unconstrained concave dual

$$F_T(\mathcal{Q}, q) = \sup_{\mu \in \mathbb{R}^c} f_T^s(\mu), \quad (3)$$

whose Lagrange multipliers μ play the role of chemical potentials conjugate to the conserved charges Q_i . The unique primal optimum is a thermal operator $X_T^s(\mu^*)$ of the matched statistics, parametrized by the chemical-potential vector μ through the grand canonical Hamiltonian $K_\mu := H - \mu \cdot Q$; see the “Primal optimum” row of Table I for the explicit form in each framework. In all three frameworks, the dual derivatives admit closed forms as thermal expectation values:

$$\frac{\partial f_T^s}{\partial \mu_i} = q_i - \text{Tr}[X_T^s(\mu) Q_i], \quad (4)$$

and the Hessian has an analogous closed form involving thermal expectations of pairs of charges. Together with the strict concavity of f_T^s , these closed-form derivatives guarantee that gradient ascent and Newton’s method converge globally to the dual optimum. Each step of the optimization is thus, physically, the measurement of the charges $\mathcal{Q}$ on the thermal operator formed from K_μ/T . As $T \rightarrow 0$, $F_T \rightarrow E$, and the temperature T is our control knob, smoothly trading the precision of the approximation to the original SDP against the conditioning of the dual problem.

The three statistics, side by side—Each framework is specified by its feasible set, its entropy, and its primal-optimum form. The three triples are summarized in Table I, and we expand on each one in turn below.

Boltzmann framework [1]. Already covered in our QTML 2025 presentation, we interpret the eigenvalues of ρ as a probability distribution (trace-one normalization), the standard quantum-mechanical assignment, and identify the primal optimum $\rho_T(\mu)$ as a non-Abelian grand canonical thermal state, that is, a *quantum Boltzmann machine* (see Table I). Its constrained free-energy minimization is precisely the training step of a quantum Boltzmann machine on a target observable, connecting the framework directly to quantum Boltzmann machine learning [6, 7].

Fermi–Dirac framework [2]. We interpret the eigenvalues $0 \leq m_j \leq 1$ of M as fermionic occupation numbers (Pauli exclusion), which reveals the natural Fermi–Dirac structure of measurement-bounded SDPs and identifies the primal optimum $M_T(\mu)$ as a Fermi–Dirac thermal measurement (see Table I for the explicit expression). The framework handles quantum hypothesis testing problems, that is, symmetric and asymmetric binary hypothesis testing, binary classification, and composite hypothesis testing, as special cases, with the Fermi–Dirac thermal measurement approaching the Helstrom–Holevo optimal measurement as $T \rightarrow 0$. The Fermi–Dirac framework also introduces

Fermi–Dirac machines, a quantum machine learning paradigm complementary to quantum Boltzmann machines, in which the trained object is a parametrized Fermi–Dirac thermal measurement rather than a parametrized quantum state.

Bose–Einstein framework [3]. We interpret the eigenvalues $x_j \geq 0$ of X as bosonic occupation numbers (no exclusion principle, allowing arbitrarily many bosons per eigenmode), which reveals the natural Bose–Einstein structure of SDPs in standard form, in which the variable is an unbounded positive semidefinite operator and no a priori bound on its trace is required, and identifies the primal optimum $X_T(\mu)$ as a Bose–Einstein thermal operator (see Table I for the explicit expression). We show that the approximation error of the regularized SDP to the original one can be made independent of the global Hilbert-space dimension, depending instead on the spectral structure of the grand canonical Hamiltonian K_μ^* at the optimum, specifically on its ground-space degeneracy and spectral gap. We also introduce the Bose–Einstein quantum relative entropy $D_{\text{BE}}(X\|Y)$, a Bregman divergence well-defined on the unbounded positive semidefinite cone, faithful and additive under direct sums, and satisfying a restricted monotonicity property under affine maps that model bosonic Gaussian channels (attenuator, amplifier, additive noise).

Hybrid quantum–classical algorithms—The closed-form gradient (4) and the analogous Hessian formula reduce the optimization, in all three frameworks, to the estimation of thermal expectation values of the input observables on the thermal operator of K_μ/T . In all three settings the hybrid quantum–classical estimator combines the same three ingredients: (i) *Hamiltonian simulation* of $e^{-iK_\mu t}$ for an evolution time t ; (ii) the *Hadamard test*, which extracts the relevant trace from the simulated dynamics; and (iii) *classical random sampling* of certain parameters of the estimator. The framework-specific elements complete the recipe: the Boltzmann framework additionally calls a Gibbs-state preparation subroutine to prepare $\rho_T(\mu)$, on which Hamiltonian simulation, Hadamard tests, and a high-peak-tent random sampling of evolution times then realize the gradient and Hessian estimators [1]; the Fermi–Dirac framework realizes the Fermi–Dirac thermal measurement directly on a hybrid qubit–qumode register via Schrödingerization [8, 9] and the power of one qumode [10], then combines this primitive with Hamiltonian simulation and Hadamard tests [2]; and the Bose–Einstein framework expands the thermal operator as a geometric series and converts each term by sampling the evolution time t from a Cauchy distribution, with Hadamard tests extracting the relevant traces [3]. In all three cases, the overall structure of the hybrid algorithm is the same. Concavity of the dual then guarantees ε -suboptimality of stochastic gradient ascent and second-order Newton variants, with runtimes polynomial in $1/\varepsilon$ and c (framework-specific spectral parameters aside).

Highlights and applications across the three frameworks—The three frameworks share their thermodynamic interpretation but offer complementary strengths, dictated by the structure of the SDP at hand.

(B) *Boltzmann*. The Boltzmann framework is the natural choice when the SDP variable is a quantum state; it underlies quantum Boltzmann machine learning [6, 7] (numerically benchmarked in [11]), and yields a hybrid quantum–classical SDP solver that polynomially improves on existing quantum SDP solvers [12] when the dual variable has large norm.

(FD) *Fermi–Dirac*. The Fermi–Dirac framework is the natural choice for measurement-bounded SDPs. Quantum hypothesis testing problems map onto it directly, and Fermi–Dirac thermal measurements provide a smooth, gradient-trainable substitute for the sharp Helstrom–Holevo threshold measurement, with provable closeness in the low-temperature limit. The same primitive has already proven useful beyond hypothesis testing: it underlies a recently proposed soft principal-component analysis with a one-shot quantum calibration [13]. Furthermore, the *Fermi–Dirac machines* introduced by the framework supply an alternative quantum machine learning paradigm in which the trained object is a parametrized measurement rather than a parametrized state, complementing quantum Boltzmann machines; they have recently been reinterpreted as quantizations of classical neurons [14].

(BE) *Bose–Einstein*. The Bose–Einstein framework targets SDPs in standard form. Because the cone $X \succeq 0$ is unbounded, no a priori upper bound on the trace of the primal variable is required. We proved that when the grand canonical Hamiltonian at the optimum is gapped, the approximation error and runtime depend only on its ground-space degeneracy and spectral gap, not on the full Hilbert-space dimension. The framework additionally yields the Bose–Einstein quantum relative entropy as a stand-alone information-theoretic primitive on the unbounded positive semidefinite cone.

Unified picture. Taken together, the three frameworks offer a novel thermodynamic interpretation, and a corresponding family of hybrid quantum–classical solvers, for a wide class of SDPs in quantum information, with natural connections to quantum machine learning paradigms based on parametrized thermal operators such as quantum Boltzmann machines and the newly introduced Fermi–Dirac machines and Bose–Einstein machines.

[1] N. Liu, M. Minervini, D. Patel, and M. M. Wilde, [Quantum thermodynamics and semi-definite optimization](#), arXiv:2505.04514 [quant-ph] (2025), [arXiv:2505.04514 \[quant-ph\]](#).

- [2] N. Liu and M. M. Wilde, [Fermi–Dirac thermal measurements: A framework for quantum hypothesis testing and semidefinite optimization](#), arXiv:2603.04061 [quant-ph] (2026), arXiv:2603.04061 [quant-ph].
- [3] M. Minervini, N. Liu, and M. M. Wilde, [Bose–Einstein thermal operators for semidefinite optimization](#), arXiv:2605.27228 [quant-ph] (2026), arXiv:2605.27228 [quant-ph].
- [4] P. Mironowicz, Semi-definite programming and quantum information, [Journal of Physics A: Mathematical and Theoretical](#) **57**, 163002 (2024).
- [5] E. T. Jaynes, Information theory and statistical mechanics, [Physical Review](#) **106**, 620 (1957).
- [6] D. Patel and M. M. Wilde, Natural gradient and parameter estimation for quantum boltzmann machines, [Physical Review A](#) **112**, 052421 (2025).
- [7] D. Patel, D. Koch, S. Patel, and M. M. Wilde, [Quantum Boltzmann machine learning of ground-state energies](#), arXiv:2410.12935 [quant-ph] (2025), arXiv:2410.12935 [quant-ph].
- [8] S. Jin, N. Liu, and Y. Yu, Quantum simulation of partial differential equations: Applications and detailed analysis, [Physical Review A](#) **108**, 032603 (2023).
- [9] S. Jin, N. Liu, and Y. Yu, Quantum simulation of partial differential equations via Schrödingerization, [Physical Review Letters](#) **133**, 230602 (2024).
- [10] N. Liu, J. Thompson, C. Weedbrook, S. Lloyd, V. Vedral, M. Gu, and K. Modi, Power of one qumode for quantum computation, [Physical Review A](#) **93**, 052304 (2016).
- [11] M. Minervini, M. Chin, J. Kupperman, N. Liu, I. Luo, M. Ly, S. Rethinasamy, K. Wang, and M. M. Wilde, [Constrained free energy minimization for the design of thermal states and stabilizer thermodynamic systems](#) (2026).
- [12] J. van Apeldoorn, A. Gilyén, S. Gribling, and R. de Wolf, Quantum SDP-solvers: Better upper and lower bounds, [Quantum](#) **4**, 230 (2020).
- [13] Y. Yuan, M. Minervini, M. M. Wilde, and N. Liu, [Quantum principal component analysis without eigenvector recovery](#) (2026), arXiv:2605.27942 [quant-ph].
- [14] A. He, N. Liu, and M. M. Wilde, [Fermi–Dirac machines as quantizations of neurons](#), arXiv:2605.24386 [quant-ph] (2026), arXiv:2605.24386 [quant-ph].