

Thermodynamic Cost and Benefit of Memory

Susanne Still*

*Department of Information and Computer Sciences, and Department of Physics and Astronomy,
University of Hawaii at Mānoa, 1680 East-West Road, Honolulu Hawaii, USA*



(Received 20 June 2019; accepted 2 October 2019; published 6 February 2020)

This Letter exposes a tight connection between the thermodynamic efficiency of information processing and predictive inference. A generalized lower bound on dissipation is derived for *partially observable* information engines which are allowed to use temperature differences. It is shown that the retention of *irrelevant* information limits efficiency. A data representation method is derived from optimizing a fundamental physical limit to information processing: minimizing the lower bound on dissipation leads to a compression method that maximally retains relevant, predictive, information. In that sense, predictive inference emerges as the strategy that least precludes energy efficiency.

DOI: 10.1103/PhysRevLett.124.050601

Effective and meaningful processing of information is crucial for the operation of “Maxwell’s demon,” a thought experiment in which a sentient being operates a trap door in a split gas container to sort fast from slow molecules, attempting to defy the second law of thermodynamics [1]. The thermodynamic cost of the demon’s operation offsets any gains [2], ensuring that the second law is not violated. Details of this process have been discussed for over a hundred years [3], but recently, interest in understanding the thermodynamics of information processing has spiked with the increasing ability to control biomolecular machines, and the advent of nanotechnology [4–20].

Ninety years ago, a simple and elegant thought experiment was proposed by Léo Szilárd [21] to illuminate the concept behind Maxwell’s demon without having to discuss any detailed physiology of a sentient being. Szilárd considered a one-particle gas in a container with a movable wall in the middle. Knowing which side is empty allows for work extraction by isothermal expansion. Highly influential to this day (e.g., Refs. [22–28]), Szilárd’s simple information engine forms a basis for understanding the conversion between information and work, and fundamental thermodynamic limits to computation.

Importantly, only information about certain aspects of a system at hand can enable work extraction. What these aspects are is determined by the physics of any given setup. An observer has to make choices about what to measure, what part of the observable data to store in memory, and to which precision. The knowledge so acquired is then used to act. Interactive observers are often called “agents” in machine learning and robotics. Szilárd assumed all agent choices to be optimal, and all relevant degrees of freedom to be observable. These assumptions reverberate in the ensuing literature, limiting the discussion largely to the special case in which *all* of the captured information can, in principle, be turned into work [29].

But in many real-world scenarios encountered by agents, biological and artificial alike, not all degrees of freedom of an observed system are accessible, thus fundamentally limiting agents to act on partial knowledge. In the most general case, systems are only partially observable, and the agent has to make inferences to predict relevant quantities from available data. This predictive inference constitutes a core function of intelligent behavior. Relative costs and benefits of more or less complicated memories have to be weighed against each other. Intuitively, smarter choices should result in fewer losses, but how precisely should a general strategy be designed? Can it be derived from a physical principle?

This Letter reveals that performing predictive inference by discarding irrelevant information is a strategy that not only enables energy efficiency, but also emerges naturally from minimizing a lower bound on dissipation—demonstrating that an intelligent data representation strategy can be derived from the optimization of a fundamental physical limit to information processing.

Information engines.—A variety of models use information to extract work from a heat bath. This information either is acquired and memorized [21], or it is supplied from the outside, whereby the costs of running a memory can be ignored, and bits can be viewed as fuel [8]. We follow Szilárd’s lead, including the cost of information acquisition and decision-making in the engine’s energy bill. We generalize to partial observability and allow the use of temperature differences. A lower bound on dissipation encountered by generalized information engines is given in Eq. (13), together with a Szilárd-Carnot cycle that saturates the bound.

An information engine contains the following. (1) A (partially) observable system, with microstate denoted by the random variable Z with realizations $z \in \mathcal{Z}$. (2) An agent, implemented via another physical system (artificial

or biological) that turns measurements into a stable memory, denoted by the random variable M with realizations $m \in \mathcal{M}$. The memory is used to decide on a work extraction protocol. (3) A work extraction device that enables the agent to couple useful energy out of the system. This device is given by the physical setup, and determines which aspects of the system are relevant with respect to extracting work.

An information engine runs cyclicly. A measurement is performed at time t_0^i , and written into memory (i denotes the i th cycle; depending on context, superscripts can be dropped). This process is implemented by a protocol that changes external control parameters on the memory between t_0^i and t_1^i . The protocol is a function of the observable data. During this process, the engine is connected to a heat bath at temperature T . An average amount of work W_M is done on the memory, and an average amount of heat, $-Q_M$, is dissipated (the convention is used that energy flows *into* a system are positive). The memory remains unchanged until t_3^i .

Work extraction starts at time t_2^i , and ends at t_3^i : a protocol is chosen as a function of the memory state and applied to the system. As a result, heat, in the average amount of Q_E , is absorbed from a heat bath at temperature T' , whereby an average amount of work, $-W_E$, is extracted. This is implemented using whichever work extraction device is given—in Szilárd's example it is the movable partition, combined with pulleys and weights [21]. The work extraction protocol is designed such that the memory has no exploitable correlations with the system after t_3^i .

The system is then prepared for the next cycle between times t_3^i and t_0^{i+1} . This preparation step is always the same procedure, and hence independent of the memory state—in Szilárd's example the partition gets reinserted into the middle of the box. Typically, it is assumed not to require work. At the start of the new cycle, the old memory state contains no correlations with the new state of the system, and is not used in the new cycle. The protocol by which the memory is set to a new state depends only on the observable data in the new cycle.

To calculate the engine's free energy changes during steps of this cycle, we consider the generalized free energy [30–35], $F_t = \langle E_t(m, z) \rangle_{p_t(m, z)} - kTH_t$, where $p_t(m, z)$ denotes the joint distribution, over system states, z , and memory states, m , at time t , $E_t(m, z)$ the energy, and $H_t = -\langle \ln[p_t(m, z)] \rangle_{p_t(m, z)}$ the Shannon entropy. (The shorthand $p(x) \equiv p(X = x)$ is used [36], and $\langle \cdot \rangle_{p(x)}$ is the average over $p(x)$.) Changes occur due to manipulation of the memory, $\Delta F_M \equiv F_{t_1} - F_{t_0}$, and due to work extraction, $\Delta F_E \equiv F_{t_3} - F_{t_2}$. The associated free energy differences,

$$\Delta F_A = W_A + Q_A - kT\Delta H_A; \quad A \in \{M, E\}, \quad (1)$$

contain an average change in energy, $\Delta E_A = W_A + Q_A$, which is the sum of the average work and heat, according

to the first law of thermodynamics, and they contain entropy changes, $\Delta H_M \equiv H_{t_1} - H_{t_0}$ and $\Delta H_E \equiv H_{t_3} - H_{t_2}$. To calculate those, it is useful to introduce the following notation:

(1) Decomposition: The system's microstate will be decomposed in two different ways, first, into observable components, X , versus nonobservable components, \bar{X} : $Z = (X, \bar{X})$. However, what can be observed is not necessarily the same as what can be manipulated. Denote by Y all components with the properties: (i) they can be manipulated between t_2 and t_3 , and (ii) they are predictable from t_0 to t_2 (the mutual information is nonzero). These are the components relevant for work extraction. Z can also be decomposed into those, versus all other components, \bar{Y} : $Z = (Y, \bar{Y})$.

(2) Manipulation: The system cannot be manipulated in a way that changes anything but Y . Thus,

$$p_{t_2}(\bar{y}|y, m) = p_{t_3}(\bar{y}|y, m) \equiv p(\bar{y}|y, m). \quad (2)$$

(3) Data representation: The stochastic map from observable data, x , to memory states, m , is independent of \bar{x} :

$$p_{t_1}(m|z) = p_{t_1}(m|x, \bar{x}) = p_{t_1}(m|x) \equiv p(m|x). \quad (3)$$

(4) Marginal distributions: All changes performed on the system do not change the marginal distribution at times t_k ($k = 1, \dots, 4$): $p_{t_k}(z) \equiv p(z)$, and therefore, $p_{t_k}(y) = \sum_{\bar{y}} p_{t_k}(y, \bar{y}) \equiv p(y)$ and $p_{t_k}(x) \equiv p(x)$. The preparation step simply introduces a hidden variable v , such that $p_{t_0}(y) = \sum_v p(y|v)p(v)$. If v is discovered, then the system appears in a nonequilibrium state best described by $p(y|v)$, which can be exploited during work extraction.

The marginal probability of the memory derives only from the statistical average of measurement outcomes. Therefore, $p_{t_k}(m) = \sum_x p(m|x)p(x) \equiv p(m)$.

(5) Inference: Knowledge carried in memory about the relevant quantity Y derives from the statistical average over measurement outcomes:

$$p_{t_2}(y|m) = \sum_x p(y|x)p(m|x)p(x) \equiv p(y|m). \quad (4)$$

Here we used that if the measurement outcome itself is given, then the memory adds no information about the relevant quantity, i.e., $p(y|m, x) = p(y|x)$. The dynamics of the system determine $p(y|x)$.

Thermodynamic cost of memory.—At t_0 , system and memory are uncorrelated and the joint distribution factorizes: $p_{t_0}(m, z) = p_{t_0}(m)p_{t_0}(z) = p(m)p(z)$. After the memory is constructed,

$$p_{t_1}(m, z) = p_{t_1}(m|x, \bar{x})p_{t_1}(z) = p(m|x)p(z). \quad (5)$$

The entropy thus decreases by $\Delta H_M = H[M|X] - H[M] = -I[M, X]$, the amount of mutual information captured in memory about the observable data. This process happens

at temperature T and the free energy change associated with it, Eq. (1), is $\Delta F_M = W_M + Q_M + kTI[M, X]$. The second law implies that $W_M - \Delta F_M \geq 0$, and hence

$$-Q_M \geq kTI[M, X]. \quad (6)$$

Operating a memory requires at a minimum dissipation proportional to the amount of information retained. This is a known fact (e.g., Refs. [11,37,38] and references therein).

Thermodynamic gain derivable from a memory.—At the beginning of work extraction, the joint distribution, written with the decomposition $Z = (Y, \bar{Y})$, is

$$\begin{aligned} p_{t_2}(m, z) &= p_{t_2}(\bar{y}|y, m)p_{t_2}(y|m)p_{t_2}(m) \\ &= p(\bar{y}|y, m)p(y|m)p(m). \end{aligned} \quad (7)$$

At the end of work extraction, all correlations between Y and M are gone, and we have

$$\begin{aligned} p_{t_3}(m, z) &= p_{t_3}(\bar{y}|y, m)p_{t_3}(y)p_{t_3}(m) \\ &= p(\bar{y}|y, m)p(y)p(m). \end{aligned} \quad (8)$$

The entropy of the joint system thus increases by $\Delta H_E = H[Y] - H[Y|M] = I[M, Y]$. During this process, the engine is connected to a bath at temperature T' . The free energy change is $\Delta F_E = W_E + Q_E - kT'I[M, Y]$, and the second law implies that $W_E - \Delta F_E \geq 0$. Therefore,

$$Q_E \leq kT'I[M, Y]. \quad (9)$$

The amount of heat that can get absorbed and turned into work is limited by how much information is retained about the variable(s) relevant with respect to work extraction.

If the relevant quantities are fully observable, then all of the energetic cost of running the memory can in principle be recovered. But, in general, information captured in memory, $I_{\text{mem}} \equiv I[M, X]$, may contain only some bits which are predictive of the relevant quantity, that is, $I_{\text{rel}} \equiv I[M, Y]$, relevant, or predictive, information. The rest is *irrelevant information*, $I_{\text{irrel}} \equiv I_{\text{mem}} - I_{\text{rel}} \geq 0$, which is non-negative, as $p(y|m, x) = p(y|x)$ implies $I[M, X] - I[M, Y] = I[M, X|Y] \geq 0$. We thus see that $Q_E \leq kT'I[M, Y] \leq kT'I[M, X]$; which means that inequality (9) is more restrictive than the known ultimate bound, $Q_E \leq kT'I[M, X]$ (e.g., Refs. [11,39,40] and references therein), which can be reached only in the special case where relevant quantities are fully observable.

Lower bound on dissipation.—Over a cycle, an isothermal information engine (connected to a heat bath at temperature T throughout) dissipates an average amount of heat $-Q = -Q_M - Q_E$. From Eqs. (6) and (9), substituting $T' = T$, we see that the irrelevant information retained in memory sets a lower limit on dissipation:

$$-Q \geq kTI_{\text{irrel}} \geq 0. \quad (10)$$

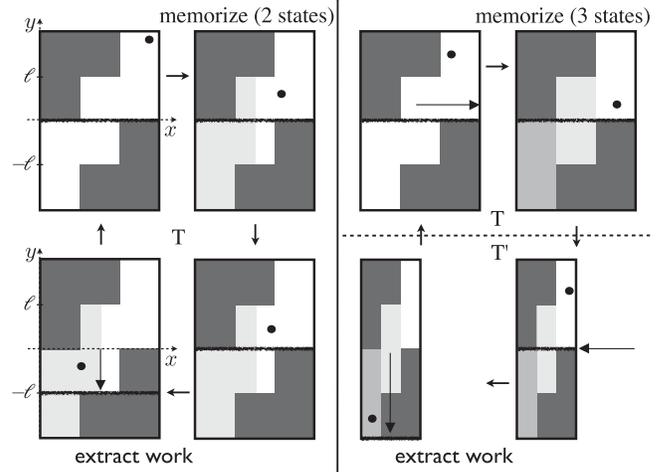


FIG. 1. Sketch of a partially observable Szilárd engine, run isothermally with a two-state memory (left panel), or with a three-state memory (right panel) using a temperature difference.

Example.—We recall Szilárd’s Gedankenexperiment [21] to build on it: one particle is trapped in a rectangular container with volume V . Let us look at the container from above, and insert a movable partition in the middle, at $y = 0$, parallel to the x axis. Knowing the sign of the particle’s y position allows us to extract, on average, work up to $kT \ln(2)$. But if we limit an agent to *only* observe the x position of the particle, then no work can be extracted, because we cannot infer from the x position alone which half of the container is empty.

Correlations can be introduced by making some areas inaccessible to the particle (Fig. 1, very dark grey). Here, the shape of the excluded areas is chosen to make it easy for the reader to calculate $I[X, Y]$ [41]. The lighter grey, and white areas in Fig. 1 signify that coarse graining along the x axis, which corresponds to possible memory states. We distinguish between two (left), or three (right) different states.

To extract work, we infer if the particle is in the upper or lower half, using the memory state m , via the inference $p(y|m)$. The wall is then moved to the half that is empty with higher probability. This will result in compression, rather than expansion, with probability $q(m)$. Therefore, we have to leave a fractional volume, $\rho(m)V$, unused. Given m , we extract, on average, $kT'([1 - q(m)] \ln[V - \rho(m)V/(V/2)] + q(m) \ln[\rho(m)V/(V/2)])$. This is maximized by choosing $\rho(m) = q(m)$, resulting in a total average extracted work of

$$\begin{aligned} -W_E &= kT' \sum_m p(m) (\ln(2) + [1 - q(m)] \ln[1 - q(m)] \\ &\quad + q(m) \ln[q(m)]) \end{aligned} \quad (11)$$

$$= kT'(H[Y] - H[Y|M]) = kT'I[M, Y]. \quad (12)$$

Because this isothermal transformation leaves the average energy of the one-particle gas unchanged, it saturates the bound in Eq. (9), as $-W_E = Q_E = kT'I[M, Y]$.

The particle's x position at the time of measurement contains information about the particle's y position at the time when work extraction begins, in the amount of $I[X, Y] = \frac{2}{3}\ln(2)$. This information can be captured by memorizing in which third of the box along the x axis the particle was found (Fig. 1, right panel, upper right drawing, lighter grey and white areas). On average, this costs at least $kT\ln(3)$, and allows us to extract at most $\frac{2}{3}kT\ln(2)$, because $\frac{2}{3}$ of the time we know for sure which side is empty (grey or white area in the drawing), otherwise (light grey area in the drawing) we can not extract any work.

If the temperature is fixed, then the dissipation encountered with a three-state memory is at least $kT(\ln(3) - 2\ln(2)/3) \simeq 0.64kT$, cf. Eq. (10). Less dissipation can be encountered with a two-state memory, costing at least $kT\ln(2)$, and capturing $I[M, Y] = \frac{5}{6}\ln(5) - \ln(3)$, all of which can be converted to work if $1/6$ of the total volume is left unused, e.g., by moving the wall between $\pm\ell$ in Fig. 1 (far left panels). This results in an average dissipative loss of at least $kT[\ln(6) - 5\ln(5)/6] \simeq 0.45kT$.

While the more detailed three-state memory allows us to extract more work than the less detailed two-state memory, the relative benefit does not outweigh the relative increase in cost.

Access to two temperatures.—The situation changes if we allow the engine to form a memory at temperature T , but extract work at a higher temperature T' . Whether it is advantageous to use a two-state or a three-state memory now depends on the temperature ratio $\alpha \equiv T'/T > 1$. The added gain of the three-state memory outweighs the additional cost when $\alpha > \alpha^* = [\ln(3) - \ln(2)]/[\ln(3) + 2\ln(2)/3 - 5\ln(5)/6] \simeq 1.847$.

While an isothermally run information engine can, at best, recover the energy needed to run the memory, an engine that has access to two different temperatures can produce non-negative work output. The engine needs to be heated from T to T' between t_1^i and t_2^i , and cooled from T' to T between t_3^i and t_0^{i+1} . We consider “memory preserving” temperature changes that do not destroy the correlations between memory and system. No additional degrees of freedom get unlocked at the higher temperature, and heating and cooling steps taken together cannot result in a net influx of heat, so that the total heat dissipated on average over a cycle is $-Q \geq -Q_M - Q_E$. The bound in Eq. (10) is thus generalized (with Eqs. (6) and (9)):

$$-Q \geq k(TI_{\text{mem}} - T'I_{\text{rel}}). \quad (13)$$

Below, we give a protocol, akin to a Carnot process, that saturates the bound.

Finding a thermodynamically efficient information processing strategy.—The bound in Eq. (13) directly

informs an optimal way to represent available data. Recall that the observer's strategy is characterized by the stochastic map $p(m|x)$. This map determines not only the amount of memory retained, but also how much relevant information is captured. We can now find the best strategy as that which *allows for the smallest possible dissipation*, by minimizing the bound in Eq. (13). Mathematically, the optimization is equivalent to

$$\min_{p(m|x)} (I[M, X] - \alpha I[M, Y]) \quad (14)$$

$$\text{subject to: } \sum_m p(m|x), \quad \forall x. \quad (15)$$

The added constraints simply ensure normalization of the stochastic map $p(m|x)$. This optimization problem is a lossy data compression method called the Information Bottleneck method [42]. It is a fairly general, information theoretic, predictive inference method, containing within it other methods as special cases (for a review see, e.g., Ref. [43]). The temperature ratio α determines the relative importance of the added benefits of a more detailed data representation, as we have seen in the example. It can also be interpreted as a Lagrange multiplier that controls the trade-off between a concise summary and keeping relevant information [42].

A protocol to change between temperatures.—A generalized Szilárd information engine can be run in a process akin to the Carnot process. The data representation step is the same as before, an isothermal transformation at temperature T . This can be implemented with $\Delta E_M = 0$, implying that $W_M = -Q_M = kTI_{\text{mem}}$.

A memory dependent work extraction protocol is then chosen (see right side of Fig. 1), consisting of the following steps: (1) isolate the one-particle gas box from the heat bath and perform an isentropic (reversible and adiabatic) compression of the entire box from volume V to V' , raising the temperature to T' ; (2) connect to a heat bath at T' , and extract work isothermally by moving into the most probably empty direction, leaving an optimized fractional volume; (3) isolate the box and perform an isentropic expansion from V' to V , lowering the temperature back to T , and (4) remove the partition and re-insert in the middle. The work done in the isentropic steps one and three cancels, and overall this protocol turns the maximally possible amount of heat into work, $Q_E = -W_E = kT'I_{\text{rel}}$. In the isentropic steps, $V/V' = (T'/T)^{d/2}$, with d degrees of freedom. The right panel of Fig. 1 illustrates $T'/T > \alpha^* \simeq 1.847$, thus making the sketched three-state memory more efficient than the two-state memory [The container is treated as two-dimensional (thus $d = 2$) to improve visibility of the smaller boxes].

Engine efficiency.—This process saturates the lower bound on dissipation, Eq. (13). It produces net work in the amount of $-W_E - W_M = k(T'I_{\text{rel}} - TI_{\text{mem}})$, and absorbs heat at the higher temperature in the amount of $Q_E = kT'I_{\text{rel}}$. Thus, it has an efficiency of

$$\eta = 1 - \frac{T}{T'} \frac{I_{\text{mem}}}{I_{\text{rel}}} = \eta_C - \frac{T}{T'} \frac{I_{\text{irrel}}}{I_{\text{rel}}}. \quad (16)$$

The Carnot efficiency, $\eta_C = 1 - (T/T')$, is reduced in proportion to the ratio of irrelevant to relevant information.

Efficiency in regular heat engines is non-negative. Information engines encounter a loss in the memory making process that can outweigh any thermodynamic gain of having the memory, and efficiency is non-negative only when the fraction of relevant bits retained in memory is larger than the fraction of low to high temperature: $I_{\text{rel}}/I_{\text{mem}} \geq T/T'$.

The better an information engine is at solving the predictive inference task by memorizing relevant bits and eliminating irrelevant bits, the more efficient it can be.

Discussion.—We have seen that it is possible to derive a rather general data representation strategy from optimization of a fundamental physical limit to information processing. This approach can be pursued in a broader context.

Thermodynamic efficiency arguments apply, for instance, to the Boltzmann machine [44], another staple machine learning algorithm. Input patterns drive this neural network out of its parameter dependent equilibrium state, q_θ , to a nonequilibrium state, p . The associated additional free energy, $F_{\text{add}} = kT\mathcal{D}[p||q_\theta]$ [30–35] (\mathcal{D} denotes the Kullback-Leibler divergence), is dissipated during the relaxation process involved in predicting labels on new patterns. Those parameters θ are found that minimize $\mathcal{D}[p||q_\theta]$ [44], thereby minimizing a lower bound on the average dissipation encountered during prediction.

Besides energy efficiency, other important physical limits to information processing include speed, accuracy, and robustness, as well as fundamental trade-offs between them (e.g., Ref. [45]). One can ask which information processing strategies emerge when those limits are optimized. It is also interesting to see how this extends to quantum systems (such as treated in, e.g., Refs. [25,40,46–52]), and to investigate the consequences of considering single-shot quantities (see, e.g., Refs. [53,54] and references therein).

I am most grateful for support from the Foundational Questions Institute (Grants No. FQXi-RFP3-1345, and No. FQXi-RFP-1820 with the Fetzer Franklin Fund). I thank Rob Shaw for patiently providing feedback and inspiration, and Elan Stopnitzky for commenting on earlier versions of this manuscript. Helpful discussions with John Bechhoefer, Tobi Berger, Bill Bialek, Gavin Crooks, Josh Deutsch, Arne Grimsmo, Massimiliano Esposito, Chris Jarzynski, Matteo Marsili, Marcelo Magnasco, Seth Lloyd, Tom Ouldridge, and David Sivak were greatly appreciated.

*Corresponding author.

sstill@hawaii.edu

[1] J. C. Maxwell, Letter to P. G. Tait, 11 December 1867, in *Life and Scientific Work of Peter Guthrie Tait*, edited by

C. G. Knott (Cambridge University Press, Cambridge, England, 1911), p. 213.

- [2] L. Brillouin, Maxwell's demon cannot operate: Information and entropy. I, *J. Appl. Phys.* **22**, 334 (1951).
- [3] A. Rex and H. S. Leff, in *Maxwells Demon: Entropy, Information, Computing*, edited by A. Hilger (Bristol, England, 1990).
- [4] T. Sagawa and M. Ueda, Minimal Energy Cost for Thermodynamic Information Processing: Measurement and Information Erasure, *Phys. Rev. Lett.* **102**, 250602 (2009); Erratum, *Phys. Rev. Lett.* **106**, 189901(E) (2011).
- [5] A. Bérut, A. Arakelyan, A. Petrosyan, S. Ciliberto, R. Dillenschneider, and E. Lutz, Experimental verification of Landauer's principle linking information and thermodynamics, *Nature (London)* **483**, 187 (2012).
- [6] D. Mandal and C. Jarzynski, Work and information processing in a solvable model of Maxwells demon, *Proc. Natl. Acad. Sci. U.S.A.* **109**, 11641 (2012).
- [7] T. Sagawa and M. Ueda, Nonequilibrium thermodynamics of feedback control, *Phys. Rev. E* **85**, 021104 (2012).
- [8] Z. Lu, D. Mandal, and C. Jarzynski, Engineering Maxwell's demon, *Phys. Today* **67**, No. 8, 60 (2014).
- [9] J. V. Koski, V. F. Maisi, T. Sagawa, and J. P. Pekola, Experimental study of mutual information in a Maxwell Demon, [arXiv:1405.1272](https://arxiv.org/abs/1405.1272).
- [10] Y. Jun, M. Gavrilov, and J. Bechhoefer, High-Precision Test of Landauer's Principle in a Feedback Trap, *Phys. Rev. Lett.* **113**, 190601 (2014).
- [11] J. M. R. Parrondo, J. M. Horowitz, and T. Sagawa, Thermodynamics of information, *Nat. Phys.* **11**, 131 (2015).
- [12] I. A. Martínez, E. Roldán, L. Dinis, D. Petrov, J. M. R. Parrondo, and R. A. Rica, Brownian carnot engine, *Nat. Phys.* **12**, 67 (2016).
- [13] J. Hong, B. Lambson, S. Dhuey, and J. Bokor, Experimental test of Landauers principle in single-bit operations on nanomagnetic memory bits, *Sci. Adv.* **2**, e1501492 (2016).
- [14] M. Gavrilov and J. Bechhoefer, Erasure Without Work in an Asymmetric Double-Well Potential, *Phys. Rev. Lett.* **117**, 200601 (2016).
- [15] M. Gavrilov, R. Chétrite, and J. Bechhoefer, Direct measurement of nonequilibrium system entropy is consistent with Gibbs-Shannon form, *Proc. Natl. Acad. Sci. U.S.A.* **114**, 11097 (2017).
- [16] E. Lathouwers and J. Bechhoefer, When memory pays: Discord in hidden markov models, *Phys. Rev. E* **95**, 062144 (2017).
- [17] A. Kumar and J. Bechhoefer, Nanoscale virtual potentials using optical tweezers, *Appl. Phys. Lett.* **113**, 183702 (2018).
- [18] G. Paneru, D. Y. Lee, T. Tlusty, and H. K. Pak, Lossless Brownian Information Engine, *Phys. Rev. Lett.* **120**, 020601 (2018).
- [19] T. Admon, S. Rahav, and Y. Roichman, Experimental Realization of an Information Machine with Tunable Temporal Correlations, *Phys. Rev. Lett.* **121**, 180601 (2018).
- [20] *The Energetics of Computing in Life and Machines*, edited by D. H. Wolpert, C. Kempes, P. F. Stadler, and J. A. Grochow (SFI Press, 2019).
- [21] L. Szilard, On the decrease of entropy in a thermodynamic system by the intervention of intelligent beings, *Z. Phys.* **53**, 840 (1929).

- [22] W. H. Zurek, Maxwells demon, Szilards engine and quantum measurements, in *Frontiers of Nonequilibrium Statistical Physics* (Springer, 1986), pp. 151–161.
- [23] M. O. Magnasco, Szilard’s heat engine, *Europhys. Lett.* **33**, 583 (1996).
- [24] S. Lloyd, Quantum-mechanical Maxwells demon, *Phys. Rev. A* **56**, 3374 (1997).
- [25] S. W. Kim, T. Sagawa, S. De Liberato, and M. Ueda, Quantum Szilard Engine, *Phys. Rev. Lett.* **106**, 070401 (2011).
- [26] S. Vaikuntanathan and C. Jarzynski, Modeling Maxwells demon with a microcanonical Szilard engine, *Phys. Rev. E* **83**, 061120 (2011).
- [27] K. Proesmans, C. Driesen, B. Cleuren, and C. Van den Broeck, Efficiency of single-particle engines, *Phys. Rev. E* **92**, 032105 (2015).
- [28] R. A. Brittain, N. S. Jones, and T. E. Ouldridge, Biochemical Szilard engines for memory-limited inference, *New J. Phys.* **21**, 063022 (2019).
- [29] It has been discussed that measurement errors limit how much information can be obtained, and thus used, e.g. in Ref. [7], but here we focus on structural constraints that dictate limited exploitable correlations between the observable data and the variables that need to be known to extract work.
- [30] R. Shaw, *The Dripping Faucet as a Model Chaotic System* (Aerial Press, 1984).
- [31] B. Gaveau and L. S. Schulman, A general framework for non-equilibrium phenomena: The master equation and its formal consequences, *Phys. Lett. A* **229**, 347 (1997).
- [32] G. E. Crooks, Beyond Boltzmann-Gibbs statistics: Maximum entropy hyperensembles out-of-equilibrium, *Phys. Rev. E* **75**, 041119 (2007).
- [33] K. Takara, H.-H. Hasegawa, and D. J. Driebe, Generalization of the second law for a transition between non-equilibrium states, *Phys. Lett. A* **375**, 88 (2010).
- [34] S. Deffner, J. P. Paz, and W. H. Zurek, Quantum work and the thermodynamic cost of quantum measurements, *Phys. Rev. E* **94**, 010103 (2016).
- [35] M. Lostaglio, D. Jennings, and T. Rudolph, Description of quantum coherence in thermodynamic processes requires constraints beyond free energy, *Nat. Commun.* **6**, 6383 (2015).
- [36] Random variables are denoted by capital letters, realizations by small letters.
- [37] R. Landauer, Irreversibility and heat generation in the computing process, *IBM J. Res. Dev.* **5**, 183 (1961).
- [38] S. Still, Lossy is lazy, in *Proceedings of the Seventh Workshop on Information Theoretic Methods in Science and Engineering* (University of Helsinki CS Series of Publications B-2014-4, 2014), pp. 17–21.
- [39] T. Sagawa and M. Ueda, Second Law of Thermodynamics with Discrete Quantum Feedback Control, *Phys. Rev. Lett.* **100**, 080403 (2008).
- [40] P. Strasberg, G. Schaller, T. Brandes, and M. Esposito, Quantum and Information Thermodynamics: A Unifying Framework Based on Repeated Interactions, *Phys. Rev. X* **7**, 021003 (2017).
- [41] This is perhaps a bit crude, a more natural shape might be a triangular area which would be equivalent to tilting the partition in the original Szilárd box by an angle, where the angle determines how much relevant information is accessible from knowing the y position.
- [42] N. Tishby, F. Pereira, and W. Bialek, The Information Bottleneck method, in *Proceedings of the 37th Annual Allerton Conference*, edited by B. Hajek and R. S. Sreenivas (University of Illinois, 1999), pp. 368–377. Available at <http://xxx.arXiv.cornell.edu/abs/physics/0004057>.
- [43] S. Still, Information Bottleneck approach to predictive inference, *Entropy* **16**, 968 (2014).
- [44] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, A learning algorithm for Boltzmann machines, *Cogn. Sci.* **9**, 147 (1985).
- [45] K. Proesmans and C. Van den Broeck, The underdamped Brownian duet and stochastic linear irreversible thermodynamics, *Chaos* **27**, 104601 (2017).
- [46] S. Deffner and E. Lutz, Quantum Speed Limit for Non-Markovian Dynamics, *Phys. Rev. Lett.* **111**, 010402 (2013).
- [47] J. J. Park, K.-H. Kim, T. Sagawa, and S. W. Kim, Heat Engine Driven by Purely Quantum Information, *Phys. Rev. Lett.* **111**, 230402 (2013).
- [48] S. Lorenzo, R. McCloskey, F. Ciccarello, M. Paternostro, and G. M. Palma, Landauers Principle in Multipartite Open Quantum System Dynamics, *Phys. Rev. Lett.* **115**, 120403 (2015).
- [49] M. Weilenmann, L. Kraemer, P. Faist, and R. Renner, Axiomatic Relation Between Thermodynamic and Information-Theoretic Entropies, *Phys. Rev. Lett.* **117**, 260601 (2016).
- [50] J. Goold, M. Huber, A. Riera, L. del Rio, and P. Skrzypczyk, The role of quantum information in thermodynamics—A topical review, *J. Phys. A* **49**, 143001 (2016).
- [51] S. Deffner and S. Campbell, Quantum speed limits: From Heisenbergs uncertainty principle to optimal quantum control, *J. Phys. A* **50**, 453001 (2017).
- [52] O. Abah and M. Paternostro, Implications of non-Markovian dynamics on information-driven engine, [arXiv: 1902.06153](https://arxiv.org/abs/1902.06153).
- [53] N. Datta, J. M. Renes, R. Renner, and M. M. Wilde, One-shot lossy quantum data compression, *IEEE Trans. Inf. Theory* **59**, 8057 (2013).
- [54] P. Faist and R. Renner, Fundamental Work Cost of Quantum Processes, *Phys. Rev. X* **8**, 021011 (2018).