Information-Theoretic Aspects in the Control of Dynamical Systems

by

Hugo Touchette

B.Sc., Université de Sherbrooke, Sherbrooke, Québec, Canada (1997)

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

Master of Science (without specification)

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2000

Signature of Author . . . .
Department of Mechanical Engineering
January 15, 2000

Certified by . . . .
Seth Lloyd
Associate Professor of Mechanical Engineering
Thesis Supervisor

Accepted by . . . .
Ain A. Sonin
Chairman, Departmental Committee on Graduate Students

Information-theoretic aspects in the control of dynamical systems

by

Hugo Touchette

Submitted to the Department of Mechanical Engineering
on January 15, 2000 in Partial Fulfillment of the Requirements for the
Degree of Master of Science (without specification)

Abstract

Information is an intuitive notion that has been quantified successfully both in physics and communication theory. In physics, information takes the form of entropy; information that one does not possess. From this connection follows a trade-off, most famously embodied in Maxwell's demon: a device able to gather information about the state of a thermodynamic system could use that information to decrease the entropy of the system. In Shannon's mathematical theory of communication, on the other hand, an entropy-like measure called the mutual information quantifies the maximum amount of information that can be transmitted through a communication channel. In this thesis, we bring together these two different aspects of information, among others, in a theoretical and practical study of control theory. Several observations indicate that such an information-theoretic study of control is possible and can be effective. One of them, is the fact that control units can be regarded intuitively as information gathering and using systems (IGUS): controllers gather information from a dynamical system by measuring its state (estimation), and then use that information to conduct a specific control action on that system (actuation). As the thesis demonstrates, in the case of stochastic systems, the information gathered by a controller from a controlled system can be quantified formally using mutual information. Moreover, it is shown that this mutual information is at the source of a limiting result relating, in the form of a trade-off, the availability of information in a control process and its performance. The consequences of this trade-off, which is very similar to the one in thermodynamics mentioned above, are investigated by looking at the practical control of various systems, notably, the control of chaotic systems. The thesis also defines and investigates the concept of controllability, central in the classical theory of control, from an information viewpoint. For this part, necessary and sufficient entropic conditions for controllability are proved.

Thesis Supervisor: Seth Lloyd
Title: Assistant Professor of Mechanical Engineering

Acknowledgments

I always had the intention to write this thesis using singular predicates instead of plural ones, and notably use 'I' instead of 'we', as a thesis is supposed to be a personal work. In fact, I even tried to convince some people to do the same, until I realized that writing such a piece of work is really impossible without the help of many people, beginning with my advisor, Professor Seth Lloyd. In writing *my* thesis, I actually, with all sincerity, present *our* work which extended over two years of collaboration, and, in my case, of fruitful learning. I want to thank him for the friendly environment of research, for proofreading this work and, above all, for the freedom I was granted in doing my work from the beginning of the program.

I also want to thank Richard Watson for reading carefully the entire manuscript, for suggesting many improvements, and helping to get my writing 'de-Frenched'. I owe the same to Lorenza Viola who read a summary article of the work presented in this thesis, and raised many important questions related to the physics of all this. Many thanks are also due to Valérie Poulin for interesting discussions on the mathematics of dynamical systems, and other subjects which have no relation whatsoever with science.

I received help in various other forms. Financially, first, from the National Science and Engineering Research Council of Canada (NSERC) in the form of an ÉS A scholarship, and from the d'Arbeloff Laboratory for Information Systems and Technology at MIT. The financial support of these organizations has been more than appreciated. Second, I want to thank the Centre de recherche en physique du solide (CRPS) of the Université de Sherbrooke, and especially Pr. André-Marie Tremblay, for the permission to access the supercomputing facilities of the CRPS.

Finally, I want to thank all the members of my family for their constant support, and for being interested about my work. I hope, one day, their insatiable thirst for understanding what I am doing will be satisfied. This thesis is dedicated to them. *Cette thèse (un mémoire en fait) leur est dédiée.*

Contents

1 Introduction **8**
- 1.1 Perspectives on information and control . . . 8
- 1.2 General framework and overview . . . 14
- 1.3 Related works . . . 17

2 Entropy in dynamical systems **18**
- 2.1 Basic results of information theory . . . 18
- 2.2 Interpretations of entropy . . . 22
- 2.3 Dynamics properties and entropy rates . . . 23

3 Information control **26**
- 3.1 Information-theoretic problems of control . . . 26
- 3.2 General reduced models . . . 28
- 3.3 Separation analysis . . . 31
- 3.4 Controllability . . . 34
- 3.5 Entropy reduction . . . 36
- 3.6 Optimality and side information . . . 40
- 3.7 Continuous extensions of the results . . . 43
- 3.8 Thermodynamic aspects of control . . . 46

4 Applications **48**
- 4.1 Controlling a particle in a box: a counterexample? . . . 48
- 4.2 Binary control automata . . . 51
- 4.3 Control of chaotic maps . . . 56

5 Conclusion **64**
- 5.1 Summary and informal extensions . . . 64
- 5.2 Future work . . . 65

List of Figures

1.1 The signalman's office. Trains are coming from the bottom of the railways. 10
1.2 (a)-(d) Exclusive N-to-1 junctions with $N = 2, 3, 4, 8$. (e) Non-exclusive 2-to-1 junction. 11

2.1 Venn diagrams representing the correspondence between entropy, conditional entropy and mutual information. 20
2.2 Discrete probability distribution resulting from a regular quantization. . 22

3.1 (a) Deterministic propagation of the state x_n versus (b) stochastic propagation. In (b) both discrete and continuous distributions are illustrated. The thick grey line at the base of the distributions gives an indication of the uncertainty associated with X_n. 27
3.2 Directed acyclic graphs (DAGs) corresponding to (a) open-loop and (b) closed-loop control. The states of the controlled system $\mathcal{X}$ are represented by X and X', whereas the state of the controller $\mathcal{C}$ and the environment $\mathcal{E}$ are C and E respectively. (c)-(d) Reduced DAGs obtained by tracing over the random variable of the environment. 29
3.3 Separation analysis. (a)-(b) open-loop control and (c) closed-loop control. The size of the sets, or figuratively of the entropy 'bubbles', is an indication of the value of the entropy. 32
3.4 Illustration of the separation analysis procedure for a binary closed-loop controller acting on a binary state system. 33
3.5 Entropy bubbles representing optimal entropy reduction in (a) open-loop control and (b) closed-loop control. 42

4.1 Different stages in the control of a particle in a box of N states. (a) Initial equiprobable state. (b) Exact localization. (c) Open-loop actuation $c = 0$. 49
4.2 Closed-loop control stages. (a) Initial equiprobable state. (b) Coarse-grained measurement of the position of the particle. (c) Exact location of the particle in the box. (d) Actuation according to the controller's state. 50
4.3 Apparent violation of the closed-loop optimality theorem. ($\bullet$) ΔH_{closed} as a function of N. ($\circ$) $\Delta H_{\text{open}} + I(X;C)$ versus N. Outside the dashed region $\Delta H_{\text{closed}} > \Delta H_{\text{open}} + I(X;C)$. The dashed line represents the corrected result when is ΔH_{open} is calculated according to the proper definition. 52

4.4 (a) $H(C)$ as a function of of the measurement error e and the initial parameter a. (b) $I(X;C)$ as a function of e and a. 53
4.5 (a) ΔH_{closed} as a function of e and a. (b) Comparison of $I(X;C)$ (top surface) and ΔH_{closed} (partly hidden surface). 55
4.6 (a) Realization of the logistic map for $r = cte = 3.78$, $x^* \simeq 0.735$. (b) Application of the feedback control algorithm starting at $n = 150$ with $\gamma = 7.0$. (c) Application of the control law from $n = 150$ with $\gamma = 7.0$ in the presence of a uniform partition of size $\Delta = 0.02$ in the estimation of the state. 58
4.7 (Left) Lyapunov spectrum $\lambda(r)$ of the logistic map. The numerical calculations used approximately 20 000 iterates of the map. (Right) Spectrum for the noisy logistic map. The definition of the 4 chaotic control regions (0, 1, 2, 3) is illustrated by the dashed boxes. The vertical dashed lines give the position of r_{mid}. 60
4.8 (a) Entropy $H_n = H(X_n)$ of the logistic map with $r = cte$ and $\lambda(r) > 0$. The entropy was calculated by propagating about 10 000 points initially located in a very narrow interval, and then by calculating a coarse-grained distribution. Inside the typical region where the distribution is away from the initial distribution (lower left part) and the uniform distribution (upper right part), $H(X_{n+1}) - H(X_n) \simeq \lambda(r)$. (b) Entropy for the logistic map controlled by a coarse-grained OGY controller starting at $n = 150$. For $n \gg 150$ (controlled regime), $H(X_n) \simeq \log \varepsilon = cte$. 61
4.9 Separation analysis for the coarse-grained controller in the optimal case where $\Delta H = 0$. The size of the measurement interval is given by $e^{H(X|C)}$ where $H(X|C) = H(X) - I(X;C)$. 62
4.10 (a) Ensemble average of the parameter r used during the simulations of the OGY control algorithm. (b) Variance of $\langle r \rangle$ (logarithmic scale). The horizontal axis for both graphs represents the number of cells used in the partition, i.e., $a = \varepsilon \Delta^{-1}$. 62
4.11 ($\circ$) Control interval ε as a function of the effective measured interval ε_m for the four control regions. The solid line shows the predicted relation between ε and ε_m based on the value of $\lambda(r_{\text{mid}})$. 63

5.1 (a)-(b) General circuit models of open-loop controllers $\mathcal{C}$ acting on a system $\mathcal{X}$ using transition functions denoted by f_c. (c)-(d) Circuit models for closed-loop control. The transition functions for the estimation process are denoted by f_e. (e) Example of an *environment-assisted* open-loop control: the open-loop controller with $H(C) = 0$ is able to recover the bit of error introduced by the environment using a Toffoli gate which uses the bit of mutual information between E and X. 67

The future belongs to those who can manipulate entropy; those who understand but energy will be only accountants... The early industrial revolution involved energy, but the automatic factory of the future is an entropy revolution. [1]

—Frederick Keffer

[1] Quoted by H.C. von Baeyer in [69]. Originally cited by Rogers in [54]. Apart from the date of publication of Roger's book (1960), the quote, unfortunately, cannot be dated exactly.

1 | Introduction

1.1. Perspectives on information and control

Information is a central component of most decision processes. Intuitively, information is needed to conduct a decision task in the same way that directions are needed to find our way to a precise location in a foreign city. In that case actually we simply *ask* for information, or go to an information desk, in order to spare us from the trouble of wandering desperately around. This information constitutes the 'decision environment', the geographic details of the city, which restricts our set of possible (situation-dependent) actions to the ones that serve effectively the purpose of the task.

In the context of control systems, the same observation can be applied to so-called *closed-loop* controllers, for they are decision agents of a specific kind whose actions on a dynamical system are influenced by the state of that system. In fact, from an information perspective, conventional closed-loop controllers, regardless of the physical nature of the process underlying the control, proceed in a universal way: *sensors* are first used to gather information about the state of the system to be controlled (estimation step); this information is processed according to a determined control strategy (decision step), and then fed back to *actuators* which try to 'update' or redirect the state of the system by augmenting its natural, non-controlled, dynamics (actuation step). Such control systems paired with 'information-processing' devices are found in many modern control systems, ranging from sophisticated automatic flight guidance systems to very simple servomechanisms such as thermostats. In the former example, the sensors are the various measurement devices (e.g., altimeter, speedometer) which provide the navigation unit with the necessary data to steer a plane, whereas in the latter the sensor is just a thermometer.

This thesis aims at clarifying this picture of control systems by proposing a novel framework which allows the information-theoretic study of general dynamical processes underlying control systems. The present work goes beyond the intuitive fact that information is required in control by demonstrating that it is possible to quantify exactly the amount of information gathered by a controller, and that there exists a direct relationship between the amount of information gathered and the performance of the controller. It also addresses the fact, not illustrated above, that controllers do not always need information to execute a specific task. A toaster, to take a trivial example, toasts and eject pieces of bread independently of its environment, and whether there is bread in the toaster or not. In fact, the class of *open-loop* controllers, which can be conceived as a subclass of closed-loop controls, includes all the control devices which do not need a continual input of information to work; like the toaster, they implement a control strategy based on general particularities of the system they are intended to control (e.g., geometric features, statistics), and not its actual state. Why open-loop control methods can work without information is a central issue of this work, along with the complementary question of why closed-loop methods require information.

From Wiener's signalman to Shannon's theory of information

The idea that control theory can be cast or formulated in an information-theoretic framework has been proposed by many scientists. Yet, none of them seem to have fully investigated the problem. Historically, one of the first who thought seriously about the role of information as a definable quantity in control systems is Norbert Wiener. In his book that founded the field of cybernetics[1] [71], Wiener notes that "the problem of control engineering and communication engineering were inseparables, and that they centered around [...] the fundamental notion of the message". In addition, in close relationship with what we noted in the introductory paragraphs, Wiener observes that

> for effective action on the outer world it is not only essential that we possess good effectors, but that the performance of these effectors be properly monitored back to the central nervous system, and that the readings of the monitors be properly combined with the other information coming in from the sense organs [...]. Something quite similar is the case in mechanical systems. [71]

To illustrate the possibility that indeed something similar happens in mechanical system, Wiener discusses an example involving a man sitting in a signal tower along a railroad, and whose job is to control a lever which activates a semaphore signal. For this situation, he notes, it is not sufficient for the signalman to assume blindly that the semaphore signal has followed his command. The signaling device may actually be defective for some reason, and in order to avoid a catastrophe, the signalman must be confirmed of his action either directly by looking at the semaphore or by referring to another device correlated with it. Wiener called this "chain of transmission and return of information" a chain of *feedback*. It appeared to him as being the central mechanism by which control processes function.

Apart from discussing this example and other similar feedback systems, Wiener thereafter did surprisingly little to build a quantitative description of control systems focused expressly on information. Although he was well aware of the developments of Shannon in communication theory [59] (to which Wiener contributed himself), and although cybernetics was originally intended to study questions of control and communication, Wiener never formalized completely his intuition about how information should be quantified in control. In fact, around the time his book on cybernetics was published, Wiener decided to pursue his studies in another direction, and embarked on

[1] The word *cybernetics* originates from the Greek $\chi\upsilon\beta\epsilon\rho\nu\acute{\eta}\tau\eta\varsigma$, which means *steersman*, in recognition to the first feedback control systems known as governors [71].

Figure 1.1: The signalman's office. Trains are coming from the bottom of the railways.

an ambitious programme of research concerned with the theory of continuous random processes and estimation. In retrospect, it is thus fair to say that Wiener was probably prevented from developing a complete theory of information because he devoted much of his time to these other subjects which, not surprisingly, proved to be fundamental to the fields of control engineering and information theory.

In this thesis, we shall continue the subject of information and control exactly where Wiener left it, as most of the results described in the next pages can be outlined by analyzing a modified version of his example of the signalman. Suppose now that the duty of the signalman is to actuate a junction of two railway tracks either in position a or b according to whether there is a train coming from track 1 or track 2 (see figure 1.1). In order to achieve this task properly, and so avoid accidents, the signalman must be positively informed of the track number (1 or 2) of the incoming train, and thus receive, for every train controlled, one bit of information. This is the case if he is in charge of two tracks. For only one track, no information is necessary, nor is control actually, and the signalman may just as well sleep! By way of generalization, we can easily convince ourselves that the control of an *exclusive* 3-to-1 junction, i.e., a junction for which only one path can be connected (figure 1.2), requires at least 2 bits of information per controlled train, the same number as a 4-to-1 junction, while an 8-to-1 junction, for example, requires 3 bits. In general, it can be guessed that at least $\lceil \log_2 N \rceil$ bits of information are needed to control an exclusive N-to-1 junction.

The occurrence of the logarithm measure in this analysis is by no means accidental. As was pointed by Hartly [27] and by Shannon [59, 60], the logarithm function is the most natural choice for a measure of information which (i) is a function of the number of *alternatives* of the system considered, and (ii) is an additive function of that number. It is additive in the sense that two systems whose number of states adds up to N^2 (Cartesian product of the states of both systems) should possess twice as much information as one system with N states. The fact that such a measure depends only on the number of alternatives is an important conceptual step in defining information: it demonstrates that semantic aspects of information referring to the concept of meaning are irrelevant in the engineering problem of communication. The only relevant aspects are syntactic,

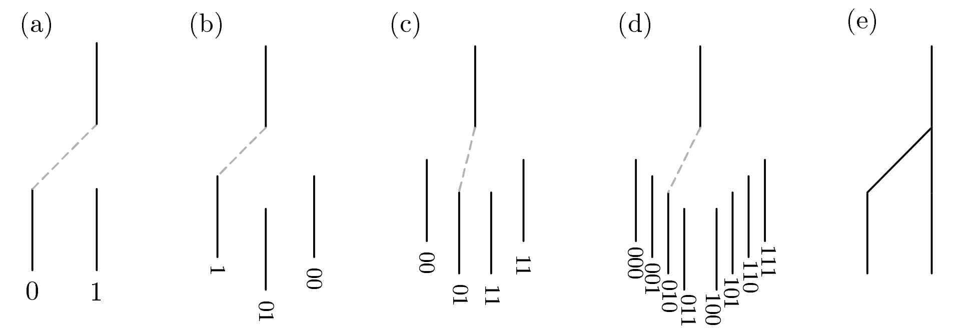

Figure 1.2: (a)-(d) Exclusive N-to-1 junctions with $N = 2, 3, 4, 8$. (e) Non-exclusive 2-to-1 junction.

which means that they refer to structural characteristics of information (e.g., number of messages and frequency of symbols in messages). Likewise, in the problem of control engineering, the meaning of the information processed by a control system plays no role in its performance. The only meaningful question is *how much* information there is?

Evidently, for an exclusive N-to-1 junction, 'how much' simply refers to the number N of incoming railways. Though this association is correct on its own, it is not the most general that one can consider. For instance, what happens if for the 2-to-1 junction only track 1 is used with probability one? Surely, in that case, no bit of information is needed to locate the position of the train. To face this eventuality, we have to take into account the statistics of the alternatives by replacing our original measure of the number of alternatives by the more general expression

$$H = -\sum_i p_i \log_2 p_i, \tag{1.1}$$

where p_i is the probability of the alternative i. The above quantity is known as the *binary entropy*, and was shown by Shannon [60] to correspond to the minimum *average* number of bits needed to encode a probabilistic source of N states distributed with p_i.[2] Intuitively, H can also be considered as a measure of uncertainty: it is minimum, and is equal to zero, when one of the alternatives appears with probability one, whereas it is maximum and equals to $\log_2 N$ when all the alternatives are equiprobable so that $p_i = N^{-1}$ for all i.

This interpretation of entropy is of foremost importance here. At the lowest level of description, controllers can arguably be thought of as devices aimed at reducing the uncertainty associated with a system. On the one hand, open-loop control methods attempt to reduce uncertainty about the state variables of a system, such as position

[2] Specifically, Shannon proved that the average size of any encoding of a data source is bounded below by its entropy. Moreover, this entropy bound can be achieved within 1 bit, thereby justifying the use of the ceiling predicate $\lceil \cdot \rceil$ in the previous page. See section 2.2 for more details.

or velocity, by acting on these variables, thereby increasing our information about their actual values. On the other hand, closed-loop control attempt the same process of reducing uncertainty with the difference that information about the system variables is used throughout the process. Returning to our example of the 2-to-1 exclusive junction, we see, for instance, that by gathering 1 bit of information, the signalman is able to identify with certainty the state of the outgoing train (on track or crashed), thus reducing the entropy by 1 bit.

Thermodynamics and the physics of information

Shannon's expression for the entropy of an information source is formally identical to Gibbs' formula for the entropy S of an equilibrium thermodynamic ensemble aside from the fact that, physically, S is measured in Joule per degree Kelvin and that i labels a physical (micro)state, not a message. For this reason, the quantity defined by Eq.(1.1) is sometimes called the *Gibbs-Shannon* entropy. This formal similarity has led many to suspect that physical entropy is just information in disguise, and that thermodynamics could be interpreted in information-theoretic terms. Conversely, the thermodynamic entropy change of a system, defined phenomenologically by Clausius as

$$\Delta S = \frac{\Delta Q}{T}, \tag{1.2}$$

should apply to information processes, therefore opening the way to a thermodynamics of information. In the above formula, ΔQ is the heat transferred in a reversible process from one physical system to another one, and T is the temperature at which this exchange process takes place [53].

In the recent years, this conjecture has found a firm ground in the observation that *information is physical*; information is not merely an abstract concept but inevitably is tied to a physical representation and must be necessarily be processed in accordance with the laws of physics [10, 32]. The earliest observations of the 'physical nature' of information are usually attributed to the works of James Clark Maxwell and Leo Szilard (cf. [35] for original references). Maxwell, in 1867 or possibly earlier, imagined a being whose ability to gather information about the positions and the velocities of the molecules of a gas would enable him to act on the molecules and create a temperature difference in the gas starting from an initial equilibrium state. Paradoxically, such a being would thus be able to decrease the entropy of a system without increasing the entropy elsewhere, in violation of the second law of thermodynamics.

Recall that, in one of its many formulations, the second law asserts exactly that it is impossible to construct a perfect engine (a perpetual machine) which would be able to extract, in a cyclic way, useful energy in the form of work from a heat reservoir without producing any other effect on the environment [53]. In another formulation, this laws posits equivalently that *the entropy of a closed system must increase or stay constant, but cannot, in any way, decrease.* In other words, if entropy decreases in some place, it must increase elsewhere.

After causing much confusion in physics, Maxwell's demon, as it has been called since, was shown by Szilard in 1929 to threaten the second law only in appearance. As noted in his seminal paper entitled "On the decrease of entropy in thermodynamic system by the intervention of intelligent beings", one bit of information gathered by the demon can only serve to decrease the thermodynamic entropy of the gas by at most one bit. In this case, the second law would still hold if there was a thermodynamic cost associated with information. In the years that followed this proposal, Brillouin, apparently unaware of Szilard's work, identified the measurement process, the act of gathering information, as being thermodynamically expensive. For many physical systems, he calculated that one bit of information can only be accessed if one bit of entropy is "dissipated" irreversibly in the form of heat [12, 13]. What Brillouin failed to show, however, is that his prognostic did not constitute a general result, for we know now that there exist non-dissipative measurement processes which do not involve the transfer of entropy in an environment (e.g., see [8]). So then, how can Maxwell's paradox be resolved?

The final word on this problem, provided by Bennett and Landauer [6, 7, 8, 9, 32], is that a thermodynamic cost has to be paid because the demon must erase the information gathered at each cycle of its operation in order to operate in a truly cyclic fashion. Hence, for each bit of information erased,

$$\Delta Q = k_B T \ln 2 \tag{1.3}$$

joules of heat should be dissipated, thereby restoring the validity of the second law (here $k_B \simeq 1.38 \times 10^{-23}$ joule/kelvin is the Boltzmann constant which provides the necessary conversion between energy units and information units, and the constant $\ln 2$ arises because S is customary defined in physics with natural base logarithms). This solution is based upon a general physical result known as Landauer's principle which states that many-to-one mappings of N bits of memory states, resulting physically from a compression of the phase space of the system serving as the representation of this memory, must lead to an entropy dissipation of N bits either in the form of heat, $\Delta Q = N k_B T \ln 2$, or as 'junk' entropy transferred to another system (cf. [32] and references therein). In other words, irreversible logical transformations on bits of information must be associated with dissipative physical processes. This must be so because the entropy of a closed system described by Hamiltonian mechanics is constant over time, as stated in Liouville's theorem [53]. As an important corollary, one-to-one mappings of logical states, i.e., reversible transformations preserving information, can be implemented physically as non-dissipative processes conserving entropy.

In the context of control, dissipative processes are very important. They can be implemented physically to reduce the entropy of a system irrespective of its state, and thus underlie most open-loop controllers. From an informational point of view, the 2-to-1 exclusive junction considered previously in figure 1.2 is surely not a dissipative device. In fact, it is a reversible one if we consider the input state of the train to be an element of the set $\{1, 2\}$ and the output state as being in $\{1'$,crashed$\}$. Using these states, the action of the junction can be described by the diagram

$$\text{(incoming track)}\quad \begin{array}{l} 1 \;\rightarrow\; \begin{cases} 1' \text{ if junction is in state } a \\ \text{crashed otherwise} \end{cases} \\ 2 \;\rightarrow\; \begin{cases} 1' \text{ if junction is in state } b \\ \text{crashed otherwise} \end{cases} \end{array} \quad \text{(outcoming state)}$$

which is equivalent to the following logical table

x	c	x'	c'
1	a	$1'$	a
1	b	crashed	b
2	a	crashed	a
2	b	$1'$	b

where c is the state, a or b, of the junction (the controller), and where x is the state of the train. The condition of reversibility in this table amounts to the fact that the states x and c can be inferred uniquely with the knowledge of x' and c' alone, or, otherwise stated, the mapping $x, c \rightarrow x', c$ is one-to-one. A possible modification of this junction, which would make it dissipative, is to change it to a *non-exclusive* junction by merging both track 1 and 2 to the same outgoing track ($1'$) as shown in figure 1.2e. In this case, the action of the junction is just a 2-to-1 mapping of the form

$$\text{(incoming track)}\ \begin{array}{ccc} 1 & \rightarrow & 1' \\ 2 & \rightarrow & 1' \end{array}\ \text{(outgoing state)}$$

It maps the final state to $1'$ irrespectively of the initial state, which means for control purposes that no human intervention is needed for the junction to function correctly.

1.2. General framework and overview

The work presented in this thesis conveys the essence of the above results by seeking a quantitative characterization of the exchange of information between a controlled system and a controller. Such a characterization of information is achieved by representing the states of both of these systems as random variables, and by considering the control process itself as a stochastic process. Within this stochastic picture, it is then possible to define information-theoretic measures similar to Eq.(1.1).

Specifically, the processes that we shall be concerned with are autonomous *discrete-state, discrete-time* (DSDT) dynamical systems described by a one-dimensional stochastic equation of the form

$$X_{n+1} = f(X_n, C_n, E_n), \tag{1.4}$$

where $n = 0, 1, 2, \ldots$ is the *time*, and where the random variable X_n represents the *state vector* of the system. The possible values x_n of the state vector form a finite set $\mathcal{X}$ called the *state* or *phase space* of the system. The random variable C_n represents the *control variable* whose values $c_n \in \mathcal{C}$ are drawn from the *actuation set* $\mathcal{C}$, assumed to be finite as well. The way the values c_n are chosen depends on the information available to the controller which, at time-instant n, is a subset ζ_n of $\{X_0, X_1, \ldots, X_n, C_0, C_1, \ldots, C_{n-1}\}$ (no access to future information). Admissible control strategies for the controller now consist of all possible mappings $g : \zeta_n \to \mathcal{C}$ whose domain determines the information structure of the control process according to the following classes.

- *Open-loop* control. C_n has no explicit dependence on the values of X_n for all n and precedent values of the control, i.e., $\zeta_n = \emptyset$. The control may depend, however, on the statistics of X_n and other general characteristics of the system under control, such as geometric features.
- *Closed-loop* or *feedback* control. $\zeta_n = \{X_0, X_1, \ldots, X_n, C_0, C_1, \ldots, C_{n-1}\}$. Thus, the control is allowed to depend explicitly on the past values of the state vector, and past values of the control.

From these standard definitions, it is clear that open-loop control laws form a subset of the class of closed-loop control. Finally, the random variable E_n represents an external process (random or deterministic) influencing the control. Physically, this process can represent a noise component added to the system or may account for perturbations resulting from interactions with any other systems. The term *environment* is reserved for these systems which generally evolve according to their own, non-controllable, dynamics.

Our information-theoretic study of the system described by Eq.(1.4) shall proceed as follows. After a brief review, in chapter 2, of the basic results of Shannon's information theory, we present in chapter 3 two general probabilistic DSDT models of control. These are used in conjunction with several entropy-like quantities to derive fundamental limits on the controllability of physical systems. Two important results are obtained in this chapter. The first one generalizes the concept of controllability of control systems to allow the characterization of approximate controllability. It is shown that a necessary and sufficient entropic condition for a system to be perfectly controllable can be given in the form of an independency condition between the system to be controlled and any other interacting systems. The second result shows closed-loop control to be essentially a zero sum game in the sense of game theory: each bit of information gathered directly from a dynamical system by a controller can serve to decrease the entropy of the controlled system by at most one bit additional to the reduction of entropy attainable without such information (i.e., in open-loop control). This last result, as we shall see, generalizes the notion of Maxwell's demons to controllers and extends the scope of the second law of thermodynamics to control problems.

Despite the fact that our analysis is restricted to one-dimensional DSDT systems, it is very important to note that, given minor exceptions noted in section 3.7, all the results obtained apply equally to continuous-space (CS) systems, with state spaces of arbitrary finite dimension, and to continuous-time (CT) dynamics under specific conditions. The need to restrict our attention on DTDS is only for the purpose of simplifying the notations, and thus making the whole general modeling question more amenable to an effective formalization. In the same vein, the recourse to a stochastic formalism in treating control systems does not imply that our results are limited to this class of systems. Without loss of generality, the state of a dynamical system is represented here by a random variable in order to take into account the following situations.

- Most real-world systems are genuinely stochastic. Material properties, forces acting on mechanical systems, and the geometry of these systems exhibit spatial and temporal random fluctuations that are best characterized by probabilistic models. Also, a controller may be fed explicitly with noisy inputs, or can be bathed in a noisy environment.

- Deterministic systems often display sensitivity to initial data or control parameters which make them effectively stochastic at finite levels of observation and description. A finite precision sensor, for instance, may blur the divergence of sufficiently close trajectories of a dynamical system upon measurement, thereby inducing a randomization in the description of its motion.

- In certain cases, our study of dynamical systems may focus only on the long term behavior of trajectories. This is justified by the fact that a large class of systems display transient evolutions followed by asymptotic recurrent regimes. It is therefore appropriate to introduce probabilistic measures in the state space which reflect global and static behaviors in the limit $n \to \infty$.

- As emphasized mainly by Wiener, we are rarely interested in the performance of a control device for a single input. The need for versatility and robustness implies that controllers be tested against a set of situations or possibilities for which they are expected to perform satisfactorily. In this sense, the ensemble picture is just a way to take into account all these situations in the same analysis. For example, the ensemble describing the inputs of the junction depicted in figure 1.1 can be taken as the set $\{1,2\}$ with $p(1) = p(2) = 1/2$, even if the mechanism specifying the arrival state of the trains is known exactly. This is so because the only (overall) functionality requirement of the junction is to perform for both input states $\{1,2\}$, not for each state at a time.

To exemplify the results of chapter 3, and to further illustrate the generality of the stochastic picture, several applications are presented in chapter 4. The applications range from the control of discrete states of information by the use of logical gates subjected to noise to the control of continuous state chaotic maps. The working model for this latter application is a control algorithm for chaotic systems proposed by Ott, Grebogi and Yorke (see chapter 4 for references), and is applied to the well-known logistic map. Finally, in chapter 5, we point out some similarities of our approach with statistical game theory, and conclude with a list of suggestions for possible future research directions including possible applications to quantum systems.

1.3. Related works

Previous proposals and ideas about entropy, information and control can be classified *grosso modo* into two categories: those which studied problems of information, particularly the question of how entropy can be reduced in systems, but did not specifically address these problems in control theory; and those which considered, in a qualitative fashion, control as an information process without invoking the quantitative power of information theory. The physicists who solved the problem of Maxwell's demon, some problems of information in dynamical systems (e.g., see [56, 61, 23]), and a few researchers in information theory [18, 11] certainly belong to the former group. Wiener, on the other hand, was argued in the introduction to belong in the latter, perhaps along with several researchers in cybernetics [3, 55, 36], and Warren Weaver [60] who discussed informally the role of information in many contexts.

Of course, there are exceptions to this classification. One of the very few is a Russian physicist named R.P. Poplavskiĭ who, during the 70's, published two papers on the subject of information and entropy, and their possible application to control. The first paper, entitled "Thermodynamic models of information processes" [51], focuses on the role of information in measurement, but mentions in passing that "information is used for purposeful control, leading to a decrease of entropy [...]", an idea that he further developed in the second paper [52]. Although the essence of Poplavskiĭ's work is very similar to that which is presented here, his approach is somewhat incomplete with respect to control theory for two main reasons. First, Poplavskiĭ addresses mainly the problem of entropy in the measurement process, and leave aside processes of actuation which, as we show, can also be included in an information-theoretic study of control. Second, most of his results concerning the thermodynamics of measurement processes are based on Brillouin's work, and, for that reason, are misleading and even wrong.

Another proposal, similar to Poplavskiĭ's, was also put forward by A.M. Weinberg who suggested the relatedness of modern microprocessor control systems with Maxwell Demons in the article "On the relation between information and energy systems: a family of Maxwell's demons" [70]. In this paper, Weinberg discusses the duality between entropy and information, but does so in a very informal mode that does not fully address quantitatively the implications of the association controller-demon. To the author's knowledge, by far the most complete work proposing a formal picture of information in control systems is a paper published by Lloyd and Slotine [40]. In this paper, the authors review two measures of information, Shannon's entropy and algorithmic information, and study general problems of control and estimation to show how these measures can be applied usefully. This latter work, together with another article by Lloyd [37], provided the conceptual basis for a recent paper [68] developed here in greater details.

2 | Entropy in dynamical systems

The aim of the present chapter is to briefly review some of the probabilistic concepts and tools used in the subsequent chapters. We first define the concept of entropy and other related quantities, and give a list of their basic properties. Then, after exploring a few examples which enable us to interpret entropy as a correlative of information, we show how these concepts can be applied in the context of dynamical systems, either deterministic or stochastic. It is not our intention in this chapter to be complete; all the necessary definitions are given in as rigorous a way as allowed, but most of the results concerning entropy are presented without proofs. A modern and complete treatment of information theory, which provides all the proofs, can be found in [19] or in the seminal article of Shannon [59] (reprinted in [60]). We also refer to the excellent books by Beck and Schlögl [5], and by Nicolis [46] for a discussion of entropy in dynamical systems theory.

2.1. Basic results of information theory

Entropy of discrete random variables

Let X be a random variable whose value x is drawn according to a probability mass function $\Pr\{X = x\} = p_X(x)$ over the finite set of outcomes $\mathcal{X}$. The set $\mathcal{X}$ together with the probability distribution $p_X(x)$, denoted simply as $p(x)$ for convenience, form a *probabilistic ensemble*. Following these notations, the *entropy* of the random variable X (we also say of the ensemble or of the distribution) is defined as

$$H(X) \equiv -\sum_{x\in\mathcal{X}} p(x)\log p(x), \tag{2.1}$$

with the convention value $0\log 0 = 0$. The choice of the base for the logarithm function is arbitrary, and can always be changed using the identity $\log_a p = \log_a b \log_b p$. Here, unless otherwise noted, it is assumed to be 2 in which case the entropy is measured in *bits*. If the natural logarithm (base e) is used, the information is given in natural units or *nats* (1 nat $= \ln 2$ bits).

Properties:

1. *Positivity.* $H(X) \geq 0$ with equality if and only if (iff) X is a *deterministic* random variable; that is, if there exists an element $a \in \mathcal{X}$ such that $p(a) = 1$ while $p(x) = 0$ for all $x \neq a$.

2. *Concavity.* Entropy is a *concave* function of the probabilities, which means that $H(E[X]) \geq E[H(X)]$, where $E[\cdot]$ denotes the expectation value. In particular, for all $\lambda \in [0,1]$, and two arbitrary probability distributions p_1 and p_2, we have

$$H(\lambda p_1 + (1-\lambda)p_2) \geq \lambda H(p_1) + (1-\lambda)H(p_2). \tag{2.2}$$

3. *Maximum entropy.* $H(X) \leq \log|\mathcal{X}|$, where $|\mathcal{X}|$ denotes the cardinality (number of elements) of $\mathcal{X}$. The equality is achieved iff X is distributed *uniformly* (written as $X \sim \mathcal{U}$), i.e., iff $p(x) = |\mathcal{X}|^{-1}$ for all $x \in \mathcal{X}$.

4. *Function of a random variable.* $H(g(X)) \leq H(X)$ where g is a function of X. The equality holds iff g is a *one-to-one* function of X, in which case we write $X \asymp g(X)$. As a corollary, entropy is invariant under permutation of the elements of $\mathcal{X}$.

Joint entropy and conditional entropy

For a set $\mathcal{Z}$ containing the joint occurrences $\{z = (x,y) : x \in \mathcal{X}, y \in \mathcal{Y}\}$, the definition of entropy can be generalized to define the *joint entropy* of X and Y:

$$H(X,Y) \equiv -\sum_{x\in\mathcal{X},y\in\mathcal{Y}} p(x,y)\log p(x,y). \tag{2.3}$$

This expression, evidently, can be applied to any group of variables $X_1, ..., X_n$ to define the joint entropy $H(X_1, ..., X_n)$.

Now, given the *conditional distribution* $p(x|y) = p(x,y)/p(y)$ with $p(y) = \sum_x p(x,y)$ as the *marginal* of Y, we define the entropy of X conditionally on the value $Y = y$ as

$$H(X|Y=y) \equiv -\sum_x p(x|y)\log p(x|y). \tag{2.4}$$

On average, then, the *conditional entropy* $H(X|Y)$ of X given Y is

$$H(X|Y) \equiv \sum_y p(y)H(X|Y=y) = -\sum_{x,y} p(x,y)\log p(x|y). \tag{2.5}$$

Properties:

1. *Positivity.* $H(X,Y) \geq 0$ with equality iff X and Y are both deterministic. Moreover, $H(X|Y) \geq 0$ with equality iff X, conditioned on Y, is a deterministic variable, or equivalently iff $X = g(Y)$.

2. *Conditioning reduces entropy.* $H(X|Y) \leq H(X)$ with equality iff $p(x,y) = p(x)p(y)$ for all $x \in \mathcal{X}$, $y \in \mathcal{Y}$. Such random variables are said to be *statistically independent*, a condition denoted henceforth by $X \perp\!\!\!\perp Y$.

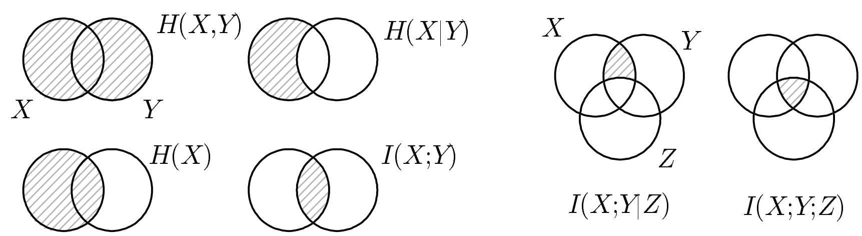

Figure 2.1: Venn diagrams representing the correspondence between entropy, conditional entropy and mutual information.

3. *Chain rule.* Using Bayes' rule $p(x|y)p(y) = p(y|x)p(x)$, we can derive $H(X,Y) = H(X) + H(Y|X) = H(Y) + H(X|Y)$. In general, we have $H(X_1, X_2, \ldots, X_n) = \sum_{i=1}^{n} H(X_i|X_{i-1}, \ldots, X_1)$.

4. *Subadditivity.* $H(X_1, X_2, \ldots, X_n) \leq \sum_{i=1}^{n} H(X_i)$ with equality iff all the X_i's are independent.

Mutual information

Given two random variables X and Y with joint distribution $p(x,y)$ and marginals $p(x)$ and $p(y)$ respectively, the *mutual information* between X and Y (or shared by X and Y) is defined to be

$$I(X;Y) \equiv \sum_{x,y} p(x,y) \log \frac{p(x,y)}{p(x)p(y)}. \tag{2.6}$$

Properties:

1. *Positivity.* $I(X;Y) \geq 0$ with equality iff $X \perp\!\!\!\perp Y$.

2. *Chain rules.* Using the chain rule for the joint entropy, we can write $I(X;Y) = H(X) - H(X|Y) = H(Y) - H(Y|X)$ and $I(X;Y) = H(X) + H(Y) - H(X,Y)$.

3. *Maximum information.* $I(X;Y) \leq \min[H(X), H(Y)]$.

Venn diagrams

The relationship between entropy, conditional entropy and mutual information can be summarized conveniently by means of Venn diagrams, as shown in the figure 2.1. This diagrammatic correspondence, which follows from the fact that entropy is a measure in the mathematical sense of the term, can be put in use to derive many chain rules and to define new quantities. For instance, figures 2.1 shows how one can define the *conditional mutual information*

$$I(X;Y|Z) = H(X|Z) - H(X|Y,Z), \tag{2.7}$$

and the *ternary mutual entropy*

$$I(X;Y;Z) = I(X;Y) - I(X;Y|Z). \tag{2.8}$$

Continuous entropy and coarse-graining

In the case where the set of outcomes $\mathcal{X}$ is a continuous set parameterized by the continuous variable x rather than a discrete set, all the above quantities can be replaced by expressions involving integrals instead of sums over the outcomes. As an example, the discrete entropy with this modification is replaced by the *differential* or *fine-grained* entropy

$$H(X) \equiv -\int_{\mathcal{X}} p(x) \log p(x)\ dx. \tag{2.9}$$

Differential entropies obey properties similar to that of discrete entropies, though a few exceptions must be noted [24, 28]:

1. The invariance under permutation property is replaced, in the continuous case, by $H(X+a) = H(X)$ for all constant a.
2. $H(X)$, $H(X,Y)$, $H(X|Y)$ and $I(X;Y)$ can be negative and/or infinite in certain cases (e.g., if $p(x)$ is a Dirac delta function).
3. $I(X;Y)$ is not necessarily upper bounded by $H(X)$ or $H(Y)$. Moreover, $H(X) \leq \log|\operatorname{supp} X|$, where $|\operatorname{supp} X| = \int_{\operatorname{supp} X} dx$ is the "volume" of the *support* $\operatorname{supp} X = \{x \in \mathcal{X} : p(x) > 0\}$. The equality is achieved in this last inequality iff X is uniform over its support, denoted by $X \sim \mathcal{U}(\operatorname{supp} X)$.
4. Some properties concerning continuous random variables must be understood as applying for *almost all points* of the density (p-almost everywhere). For example, the independence relation $X \perp\!\!\!\perp Y$ for the continuous variables X and Y is defined properly as $p(x,y) = p(x)p(y)$ for all points x and y except for a subset of points having measure zero. The same remark applies when two densities are said to be equal.

The continuous entropies can be related to their discrete analogs by considering what is called a *coarse-graining* of a continuous random variable (also known as a *discretization* or a *quantization*). Consider the random variable X with density $p(x)$ illustrated in figure 2.2. By dividing the range of X into non-overlapping bins of equal length Δ, it is possible to define a *coarse-grained* random variable $X^\Delta = x_i$, for $i\Delta \leq X < (i+1)\Delta$, and whose probability distribution is obtained by "averaging" the density in the bins

$$\Pr(X^\Delta = x_i) = p_i = \int_{i\Delta}^{(i+1)\Delta} p(x) dx. \tag{2.10}$$

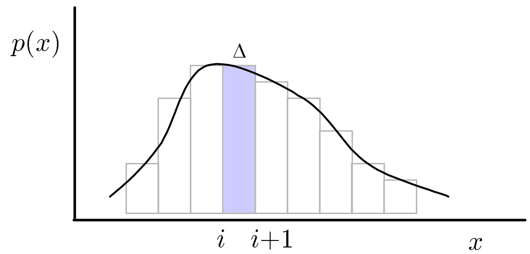

Figure 2.2: Discrete probability distribution resulting from a regular quantization.

The corresponding *coarse-grained* entropy thus equals

$$H(X^{\Delta}) = -\sum_i p_i \log p_i, \tag{2.11}$$

and can be shown to satisfy the limit $H(X^{\Delta}) + \log \Delta \to H(X)$ as $\Delta \to 0$ if $p(x)$ is Riemann integrable [19]. For Δ sufficiently small, this limit can be turned into the useful approximations

$$\begin{aligned} H(X^{\Delta}) + \log \Delta &\simeq H(X), && (2.12) \\ I(X^{\Delta}; Y^{\Delta}) &\simeq I(X;Y). && (2.13) \end{aligned}$$

See also [26] for further information on coarse-grained entropies.

2.2. Interpretations of entropy

As discussed in the introduction, the entropic quantities defined so far are important in that they capture quantitatively most of the intuitive connotations associated with the concept of information. As first pointed out by Shannon in the development of his theory of communication, entropy can be interpreted as a measure of uncertainty or missing information associated with a statistical ensemble, and enters naturally in the expression of many results relating the compression and the transmission of information. The following is a list of different interpretations of entropy and scenarios in which this quantity is used. The intuitive character of the scenarios should make them useful for understanding the results of the next chapter.

- *Uncertainty and variability.* Recall that the discrete entropy $H(X)$ is minimum and equal to zero when one of the probabilities is equal to unity, i.e., when there is no uncertainty left about the outcome of X. Conversely, $H(X)$ is maximum and equals $\log |\mathcal{X}|$ when all the alternatives of X are equiprobable; in this case, the ensemble is completely "disordered" since $p(x)$ gives no information about which outcome is more likely to happen. For a continuous random variable X, $H(X)$ can be thought as a measure of the "width" of the density $p(x)$ which, in turn, is a measure of the uncertainty or variability associated with X.

- *Noiseless coding.* As formalized in Shannon's source coding theorem (or noiseless channel theorem), $H(X)$ represents the minimum amount of resources (bits) required to encode unambiguously the ensemble describing X. More precisely, let C be a (uniquely decodable) *binary source code* for the ensemble $\{\mathcal{X}, p(x)\}$ consisting of codewords of length $l(x)$ associated with each outcome x. Then, the *average* codeword length, $E[l(X)] = \sum_x p(x)l(x)$, must satisfy $E[l(X)] \geq H(X)$. This agrees with the intuitive fact that 1 bit of information is needed to encode each toss of a fair coin and, equivalently, that 1 bit of information can be stored in a binary state system. In the case of a continuous random variable, one can essentially apply the same result by noting that n bits suffice to encode the uniform continuous variable $X \sim \mathcal{U}([0,1])$ quantized with $\Delta = 2^{-n}$ [19].

- *Correlation and observation.* Given two random variables X and Y, $I(X;Y)$ gives the average number of bits (shared information) that X and Y have in common. Also, from the chain rule $I(X;Y) = H(X) - H(X|Y)$, we can interpret $I(X;Y)$ as the reduction of entropy of X resulting from the observation of Y. In other words, $I(X;Y)$ represents the information on X gained by observing the random variable Y. For example, the observation of the outcome of a fair coin toss (1 bit of information) "reduces" the original uncertainty of 1 bit associated with the coin before observation.

- *Transmission capacity.* For a communication channel of the form $X \to Y$ modeled by the transition probability $p(y|x)$, Shannon showed that $I(X;Y)$ quantifies the average maximum rate (capacity), measured in bits per use of the channel, at which one can transmit information (see [19] for a formal account of this result). If $X \perp\!\!\!\perp Y$, then observing Y yields no information about X; Y is completely random with respect to X, and no information can be transmitted between the two random variables. When $X \asymp Y$, on the other hand, $H(X) = H(Y)$ bits of information are gathered upon observing Y, which means that an average of $H(X)$ bits of information can be transmitted in a *noiseless* channel. Using this fact, $H(X|Y)$ may be interpreted as the *information loss* appearing, in the expression of the capacity, as a negative contribution to the maximum noiseless rate $H(X)$.

2.3. Dynamics properties and entropy rates

Once a specific model of the stochastic system under study is determined, it is possible to calculate the density $p(x_n)$ of the state of that system in time, and consequently its entropy $H(X_n)$ using Eq.(2.1) or (2.9). In principle, evolution equations for $p(x_n)$ and $H(X_n)$ can be derived from methods of dynamical systems theory (e.g., Frobenius equation, Fokker-Planck equation, and the Kolmogorov-Chapmann equation for Markovian processes [33]). For the present study, we will not touch upon the actual derivation of these equations, but will only be concerned with the effect of the dynamics on entropy.

To be more precise, consider a system such as the one described by Eq.(1.4). Let X be the state of the system at time n, and X' the state at a later time $n' > n$.

Using this notation, we say that the transition $X \to X'$ is *dissipative*, or simply *entropy decreasing*, if $H(X') < H(X)$, or equivalently if the change of entropy of the system over the transition, defined as

$$\Delta H = H(X) - H(X'), \tag{2.14}$$

is strictly positive. Conversely, the transition is called *entropy increasing* if $\Delta H < 0$. It is *entropy conserving* when $\Delta H = 0$. A *non-dissipative* transition is such that $\Delta H \leq 0$. (Note that ΔH is defined so that $\Delta H > 0$ corresponds to the desired situation where the system is controlled.)

Table 2.1 lists a few examples of systems whose dynamics fall into the three classes $\Delta H > 0$, $\Delta H = 0$, and $\Delta H < 0$. Let us note that most systems in this list, except for the ones with the asterisk *, are not *generic* in the sense that they do not always satisfy the requirements of the class they have been associated with. This arises because a system's entropy may decrease over a certain period of time and increase over another period. Also, ΔH may depend on the choice of X_n. In these cases, it is impossible to determine if a system is strictly dissipative or not. The listed systems do however satisfy the requirements in an average sense or for typical states. To illustrate this, consider the following one-dimensional CSDT map

$$x_{n+1} = f(x_n), \tag{2.15}$$

where $x \in \mathbb{R}$, $n = 0, 1, 2, \ldots$. Let us assume that we do not know anything about the system except that at a time n the iterate x_n lies in an interval ε_n centered around x_n, and at time $n+1$ it is in ε_{n+1}. According to our definition of continuous entropy, this implies that

$$\Delta H \simeq \log \varepsilon_n - \log \varepsilon_{n+1}. \tag{2.16}$$

For infinitesimal intervals, we have

$$\Delta H = \lim_{\varepsilon_n \to 0} \log \varepsilon_n - \log \varepsilon_{n+1} = -\log \left| f'(x_n) \right|. \tag{2.17}$$

where $f'(x_n) = \partial_x f|_{x_n}$. Hence, we see that ΔH depends locally on the value of the Jacobian of the map which, apart from linear systems, varies in general from point to point in the state space. Globally, we can also average the contribution of each Jacobian over the space and define, accordingly, the *Lyapunov exponent* of the map f

$$\lambda \equiv \int \rho(x) \log \left| f'(x) \right| dx, \tag{2.18}$$

which is a quantity that does not depend on time if the distribution $\rho(x)$ is the invariant density of the map obtained by solving the so-called Frobenius-Perron equation [33]. Using the expression for λ, we expect that $\Delta H = -\lambda$ should hold on average [61, 23, 58, 50, 67]. However, one can prove that this equality does not always hold (even on average) though, numerically, it can be observed that it holds for any typical choice of $p(x_n)$, and that $\Delta H \leq -\lambda$ for almost all $n \geq 0$ (see [34] and section 4.3). On average then $\Delta H \simeq -\lambda$.

Remarks:

1. The Lyapunov exponents can also be calculated as time averages instead of ensemble averages using the formula

$$\lambda = \lim_{N\to\infty} \frac{1}{N} \sum_{i=1}^{N} \ln \left| f'(x_i) \right| , \tag{2.19}$$

where the sum is taken over a discrete-time trajectory $x_1 x_2 \ldots x_N$. (Other formulae also exist for continuous-time dynamics. See [46]). The above formula and Eq.(2.18) are equal for *ergodic* systems [21].

2. For a D-dimensional system, there are D Lyapunov exponents λ_i, $i = 1, 2, \ldots, D$. In this situation, the relationship between ΔH and the λ_i's is more problematic since information about the evolution of a system can be lost in any direction of the space where $\lambda_i > 0$, i.e., in the directions for which $p(x_n)$ undergo an expansion over time. This suggests that ΔH should depend on the positive exponent, an observation embodied in the inequality $\Delta H \geq -h$, where $h = \sum_{i:\lambda_i>0} \lambda_i$ is the *Kolmogorov-Sinai entropy* [2]. This inequality has again been confirmed in numerical simulations [34].

Class	Type	System, evolution	Characteristics
$\Delta H < 0$	DSDT	N-to-M maps	$M > N$
	DSDT	Double stochastic MC	$\boldsymbol{p}' = P\boldsymbol{p}$, $\sum_i P_{ij} = \sum_j P_{ij} = 1$
	(CD)SDT	Noise-perturbed map	$X_{n+1} = f(X_n) + \xi_n$
	CSCT	Noise-perturbed flow	$\dot{X}(t) = f[X(t)] + \eta(t)$
	CS(CD)T	K system + coarse-grained	$\lambda_i > 0$ for some i
		hamiltonian systems	$\Delta H \lesssim h$
$\Delta H = 0$	DSDT	Permutation*	$\boldsymbol{p}' = P\boldsymbol{p}$,
			P is a permutation matrix
	CSCT	Hamiltonian systems*	Liouville theorem, $\sum_i \lambda_i = 0$
	(CD)S(CD)T	Quantum closed systems	
$\Delta H > 0$	DSDT	N-to-M maps	$M < N$
	CSCT	Damped system	$\dot{E} < 0$ (energy function)
			$\sum_i \lambda_i < 0$

Table 2.1: Examples of systems or dynamics for the three entropy classes. The asterisk * means that the example given satisfies systematically the conditions of its class (see text). In this table, $\boldsymbol{p} = (p_1, p_2, ..., p_n)^T$ denotes the probability (column) vector of a discrete state Markov chain (MC), and P is the transition probability matrix whose elements are $P_{ij} = \Pr(x_{n+1} = j | x_n = i)$. Also, ξ_n and $\eta(t)$ represent respectively discrete and continuous noise processes.

3 | Information control

3.1. Information-theoretic problems of control

In the context of DSDT systems, one general problem of control is to find a *history* $\{c_n\}$ of the control variables $c_n \in \mathcal{C}$, say for a time interval $[0, N]$, which forces the state x_n from an initial value x_0 to a final state x_N and, at the same time, minimizes or maximizes a *cost* function J along the trajectory connecting x_0 and x_N [65]. In a very broad sense, the measure J serves as a 'metric' for comparing the performance of different control histories. It is generally a scalar function of $\{x_n\}$ and $\{c_n\}$ whose purpose is to embody in a quantitative fashion the control goals, a criterion of stability, in addition to possible constraints on the control design (e.g., energy or power constraint, material specification, geometry, etc.) In principle, by incorporating all these aspects of control into one function, it is then possible to solve a single optimization problem

$$J(\{x_n^*\}, \{c_n^*\}) = \min_{\{c_n\}} J(\{x_n\}, \{c_n\}), \tag{3.1}$$

which yields an *optimal trajectory* $\{x_n^*\}$ for a given control history $\{c_n^*\}$ (figure 3.1a).

Now, for a system described by a stochastic evolution equation such as Eq.(1.4), even if we specify how the control variable c_n has to be chosen, it is impossible to predict the precise outcome of the cost function because of the presence of the random component E_n. This means that the deterministic cost function above cannot be minimized unambiguously by the choice of a control strategy. If, however, the statistics of the systems involved in the control are known to a certain extent, we may minimize the *expected value* of the cost function, e.g., using dynamic programming techniques [65, 4], thereby specifying a set of inputs $\{C_n\}$ that control an ensemble of trajectories which are optimal on average.

The fact that stochastic control strategies are intended to act upon an ensemble of trajectories rather than a single one is very important as it enables us to formulate problems of stochastic control which have no equivalent in deterministic control. One fundamental difference of stochastic control is that the state of a system cannot usually be "stochastically" controlled to take only one end value x_N. Instead, one might consider control of an *estimate* of the state, usually taken to be the expectation value $\hat{x}_n = E[X_n]$, and construct J so as to minimize the *average distance*

$$d(\hat{x}_N, x_N) = |E[X_N] - x_N|\,, \tag{3.2}$$

in addition to the *average error*

$$e(\hat{x}_N, x_N) = E[|X_N - x_N|]. \tag{3.3}$$

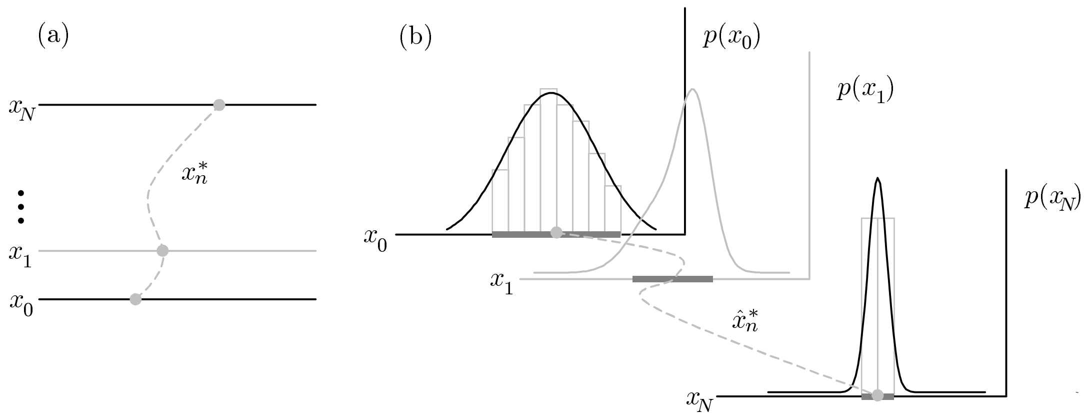

Figure 3.1: (a) Deterministic propagation of the state x_n versus (b) stochastic propagation. In (b) both discrete and continuous distributions are illustrated. The thick grey line at the base of the distributions gives an indication of the uncertainty associated with X_n.

The minimization of $e(\hat{x}_N, x_N)$ can be seen as a requirement of stability: in a deterministic sense, a state x is *stable* if any trajectory initiated by a point x_0, chosen in a small neighborhood of x, stays arbitrarily close to that state as $n \to \infty$ [64]. From a stochastic viewpoint, this requirement of stability translates readily into constraining the error within a certain interval. Accordingly, one problem that we shall be concerned with is to determine by how much the uncertainty of the final state X_N can be constrained for a specific control situation. By associating the measure of uncertainty with the entropy $H(X_n)$, this is equivalent to asking: given $H(X_0)$, the actuation set $\mathcal{C}$ and the type of controller (open-loop or closed-loop), what is the maximum entropy decrease $H(X_0) - H(X_N)$ achievable in the transition $X_0 \to X_N$?

For this problem, it must be stressed that the reduction of entropy we are referring to is the reduction associated with the entropy of the marginal distribution $p(x_n)$. For control purposes, it does not suffice to reduce the entropy of a system conditionally on the state Y of another system, as the inequality $H(X_n|Y) \leq H(X_n)$ would like to suggest. In particular, contrary to a statement commonly found in the physics literature, $H(X_n)$ cannot be reduced simply by 'measuring' X_n, for measuring apparatuses (classical ones at least) are not intended to affect or influence physically the dynamics of the observed system. As an example, one can imagine measuring the state X_n in such a way that, given the measurement outcome, X_n is determined with probability one. Yet, is the system really controlled after this measurement? Surely not. In fact, it is conceivable that the measurement has left the system's state unchanged and, as a result, any observer who does not know the outcome of the sensor will estimate the entropy of the controlled system from the data of $p(x_n)$ and obtain $H(X_n)$. Consequently, we see that what is required for a control action is to act directly on the dynamics of the controlled system through the use of actuators in order to reduce the entropy $H(X_n)$ for *all* observers unconditioned on the state of any external systems. (This is still true even if $p(x_n)$ is interpreted from a Bayesian perspective as a subjective belief measure on the state X_n. Subjective or not, $p(x_n)$ is the distribution that must be modified in a control process since X_n is the only state available to any observer of the system.)

Although the problem of controlling a dynamical system requires more than limiting its entropy, the ability to limit entropy is a necessary condition for control: the fact that a control process is able to localize x_n near a stable state x simply means that the controlled system can be constrained to evolve into states of low entropy. Another necessary requirement for control is that of reachability: given the initial (random) state, is it possible to reach a final state x_N at time instant N by manipulating the control inputs? In a stronger sense, we can also inquire about controllability; that is, given the initial state, is it possible to reach *any* final state? From these definitions, we see that another distinction has to be made for stochastic control, namely the need to allow for approximate reachability and approximate controllability. Probabilistically, a state x_N can be defined as *approximately reachable* from an initial state X_0 if, given the type of control, we can build a single trajectory connecting X_0 to x_N with non-zero probability. Clearly this definition can be extended to define approximate controllability.

In this chapter, we study these two problems, namely the problem of entropy reduction and the reachability-controllability problem (actually, in the reverse order). For both problems, we shall simplify our analysis further by restricting our attention to the study of transitions of the form $X \to X'$, where $X = X_n$ and $X' = X_{n+1}$. This restriction is totally justified for two reasons. First, the derivation of limiting bounds on ΔH for a single time-step automatically implies a bound for any multiple time-steps processes. Second, results concerning the controllability of a system can be formulated for a single time-step control history c_n, and can be generalized thereafter, when necessary, by associating c_n with a collective history $c_n c_{n+1} \cdots c_m$.

3.2. General reduced models

The stochastic models used for the description of a transition $X \to X'$ are shown in figure 3.2. In the jargon of statistics, the graphs shown in this figure are called *directed acyclic Bayesian networks*, or simply *directed acyclic graphs* (DAG) [49, 44, 17, 30, 38]. The vertices of these graphs correspond to random variables representing the state of a particular system, while the arrows give the probabilistic dependencies among the random variables according to the general decomposition

$$p(x_1, x_2, \ldots, x_N) = \prod_{j=1}^{N} p(x_j \,|\, \mathrm{parent}[x_j]), \tag{3.4}$$

where $\mathrm{parent}[x_j] \subset \{x_i\}$ is the set of direct parents of X_j. The acyclic condition means that no vertex is a descendant or an ancestor of itself, in which case we can order the vertices chronologically, i.e., from ancestors to descendants. This defines a causal ordering, and consequently a time line directed from left to right.

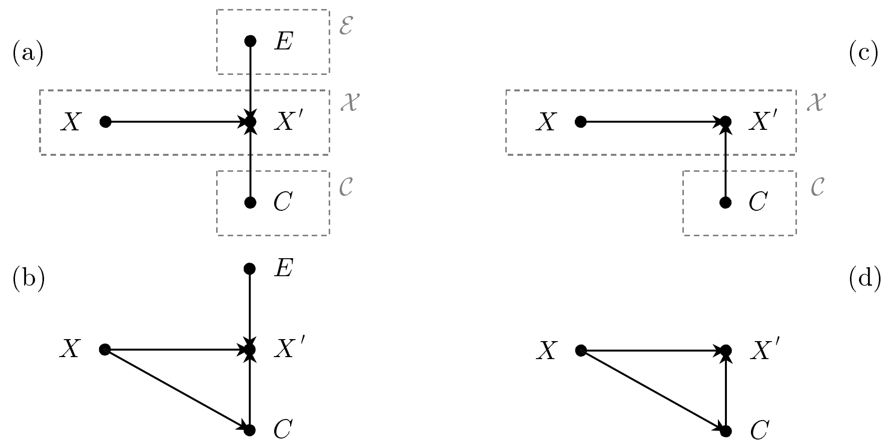

Figure 3.2: Directed acyclic graphs (DAGs) corresponding to (a) open-loop and (b) closed-loop control. The states of the controlled system $\mathcal{X}$ are represented by X and X', whereas the state of the controller $\mathcal{C}$ and the environment $\mathcal{E}$ are C and E respectively. (c)-(d) Reduced DAGs obtained by tracing over the random variable of the environment.

Let $\mathcal{X}$ denote the system to be controlled, and let $\mathcal{C}$ and $\mathcal{E}$ denote respectively the controller and the environment to which $\mathcal{X}$ is coupled.[1] The initial state X of $\mathcal{X}$ is distributed according to $p(x)$, $x \in \mathfrak{X}$, while the final state after the transition is denoted by X'. The state of the controller and the environment are respectively denoted by C and E, with $c \in \mathfrak{C}$ and $e \in \mathfrak{E}$. Using these notations and Eq.(3.4), the complete joint distribution $p(x, x', c, e)$ over the random variables of the graphs can be constructed. For instance, the complete joint distribution corresponding to the open-loop graph of figure 3.2a can be written as

$$p(x, x', c, e) = p(x)p(c)p(e)p(x'|x, c, e), \tag{3.5}$$

whereas the closed-loop graph of figure 3.2b is characterized by a joint distribution of the form

$$p(x, x', c, e) = p(x)p(c|x)p(e)p(x'|x, c, e). \tag{3.6}$$

In accordance with the definition of open-loop and closed-loop control given in the introduction, what distinguishes graphically both control strategies is the presence, for closed-loop control, of an explicit correlation link between X and C. This correlation can be thought of as a communication channel (*measurement* channel) that enables $\mathcal{C}$ to gather an amount of information identified formally as the mutual information $I(X;C)$. Operationally, we thus define closed-loop as having $I(X;C) \neq 0$; open-loop control, on the other hand, corresponds to the condition $I(X;C) = 0$, or equivalently $X \perp\!\!\!\perp C$.

[1] The calligraphic letters should not be confused with the script upright letters (Euler fonts), such as $\mathfrak{X}$ and $\mathfrak{C}$, which stand for the respective state spaces of the systems $\mathcal{X}$ and $\mathcal{C}$.

Furthermore, since in open-loop control the control variable C is entirely determined by the choice of $p(c)$, we define the following: an open-loop control strategy is called *pure* if C is a deterministic random variable, i.e., if $p(\hat{c})=1$ for one value c of C and $p(c)=0$ for all $c\neq\hat{c}$. An open-loop controller that is not pure is called *mixed* (we also say that a mixed controller implements a *mixture* of control actions).

The joint distributions (3.5) and (3.6) show that the effect of the environment and the controller is taken into account by a channel-like probability transition matrix $p(x'|x,c,e)$ parameterized by the control value c and by each *realization* e of the stochastic perturbation E appearing with probability $p(e)$. In effect, for each value c and e, the system $\mathcal{X}$ undergoes an evolution

$$X \xrightarrow{E,C} X', \tag{3.7}$$

referred to here as a *subdynamics*, which is assumed to be completely determined by the values of the states X, C and E. In other words, we assume that X' conditioned on X, C and E is a deterministic random variable whose value is specified with probability one by $p(x'|x,c,e)$. As a result of this assumption, we have the following.

Proposition 1. *If X' is a deterministic random variable conditioned on X, C and E, then $H(X'|X,C,E)=0$.*

Proof. For each triplet (x,c,e), $p(x'|x,c,e)=1$ for one value of x'. Thus, $H(X'|x,c,e)=0$. On average then

$$H(X'|X,C,E)=\sum_{x,c,e}H(X'|x,c,e)p(x,c,e)=0, \tag{3.8}$$

where $p(x,c,e)$ is obtained by summing Eq.(3.5) or (3.6) over x'. □

It is very important to note that the deterministic assumption on X', which only relies on the mathematical form of the evolution function f in Eq.(1.4), does not imply that $H(X')=H(X)$. That is to say, a system can be stochastic even if $p(x'|x,c,e)$ is a deterministic transition function. For example, consider an open-loop system with the binary states $x,x',c,e\in\{0,1\}$ described by

$$x'=\begin{cases} x\oplus e, & c=0\\ \bar{x}\oplus e, & c=1\end{cases}, \tag{3.9}$$

where $\oplus$ stands for modulo 2 addition, and $\bar{x}=x\oplus 1$ is the *complement* of x. If $E\sim\mathcal{U}$, it is easy to verify that $H(X')=\log 2$ for all $p(x)$ and $p(c)$, notwithstanding the fact that, obviously, $p(x|x,c,e)$ is deterministic. For a system such as Eq.(3.9), or any of the models of figure 3.2a-b, the deterministic/nondeterministic character of the dynamics can be revealed by the reduced matrix $p(x'|x,c)$ which is obtained by tracing out the variables of the environment

$$p(x'|x,c)=\sum_e p(x'|x,c,e)p(e). \tag{3.10}$$

(See figure 3.2c-d). The specific form of this *actuation* matrix depends on the subdynamics envisaged for the control process and the effect of the environment. Some of the control actions, for example, may correspond to control strategies forcing several initial states to a common stable final state, in which case the corresponding dynamics is entropy decreasing. Others can model uncontrolled transitions perturbed by the external noise component E leading to "fuzzy" actuation rules which increase the entropy. (This is valid only approximately, for ΔH may also depend on X.) Eq.(3.9) is an example of such a fuzzy control system which specifies the final state X', conditioned on X and C, only up to a certain probability different than one or zero. Explicitly, for this example, $p(x'|x,c) = 1/2$ for all values x', x and c.

Remarks:

1. As a shortcut, it is also possible to model a control process (stochastic or deterministic) by specifying directly the actuation matrix without including explicitly the random variable E. This approach in modeling a control process is sometimes adopted in this chapter.

2. A complete model of a control system should include two intermediate devices: a sensor $\mathcal{S}$ which is the measurement apparatus used to access the value of X, and an actuator $\mathcal{A}$ which is the device used to change the dynamics of $\mathcal{X}$. Here, $\mathcal{S}$ and $\mathcal{A}$ are merged in a single device $\mathcal{C}$ called the controller which fulfills the roles of estimation and actuation. From the viewpoint of information, this simplification amounts to a case where we may consider the sensor as being connected to the actuator by a perfect noiseless communication channel, so that $A \asymp S$ and $I(A;S) = H(A) = H(S)$. In this case, we are justified in replacing the random variables A and S by C with $H(C) = H(A) = H(S)$.

3.3. Separation analysis

The reduced graph of figure 3.2c possesses a useful symmetry that enables us to separate the effect of the random variable X in the actuation matrix from the effect of the control variable C. From one perspective, the open-loop decomposition

$$p(x') = \sum_c p(c) \sum_x p(x'|x,c)p(x) \tag{3.11}$$

suggests that a mixed open-loop controller can be decomposed into pure actuations, one for each value of c, that takes the initial distribution $p(x)$ to a final distribution

$$p(x'|c) = \sum_x p(x'|x,c)p(x). \tag{3.12}$$

The final distribution $p(x')$ is then obtain by "averaging" over the control variable using $p(c)$. From another perspective, the same decomposition, re-ordered as follows

$$p(x') = \sum_x p(x) \sum_c p(x'|x,c)p(c), \tag{3.13}$$

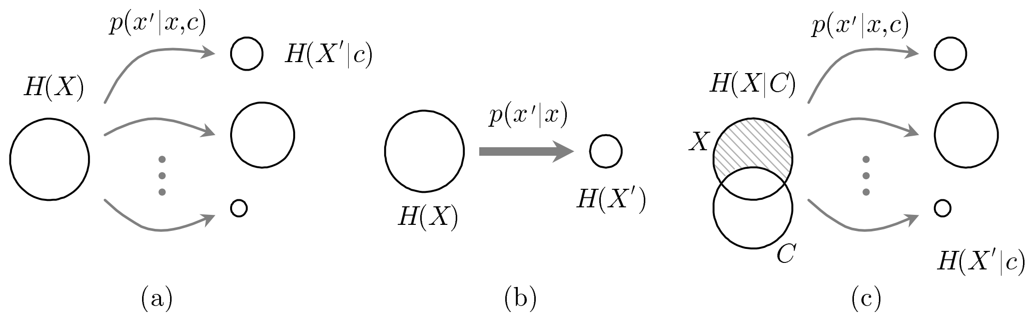

Figure 3.3: Separation analysis. (a)-(b) open-loop control and (c) closed-loop control. The size of the sets, or figuratively of the entropy 'bubbles', is an indication of the value of the entropy.

indicates that the overall action of a mixed controller can be seen as propagating $p(x)$ using

$$p(x') = \sum_x p(x'|x)p(x) \tag{3.14}$$

through an "average" actuation channel $p(x'|x) = \sum_c p(x'|x,c)p(c)$. In the former perspective, as shown in figure 3.3a, each actuation subdynamics can be characterized by a pure entropy reduction

$$\Delta H_c = H(X) - H(X'|c) \tag{3.15}$$

associated with the transitions

$$\begin{gathered} X \xrightarrow{c} X'|c \\ p(x) \overset{p(x'|x,c)}{\longrightarrow} p(x'|c) \end{gathered} \tag{3.16}$$

for all c. (The notation $X'|c$ stands for the random variable X' conditioned on the value $C = c$). In the latter perspective (figure 3.3b), the entropy reduction associated with the average actuation channel is exactly the total entropy change ΔH of the (possibly mixed) transition $X \to X'$.

For closed-loop control such a separation of $\mathcal{X}$ and $\mathcal{C}$ is a priori impossible, for C itself depends on the initial state X. Despite this fact, one can use Bayes' rule

$$p(x|c) = \frac{p(c|x)p(x)}{p(c)} \tag{3.17}$$

to invert the dependency between X and C in the closed-loop decomposition

$$p(x') = \sum_{x,c} p(x'|x,c)p(c|x)p(x), \tag{3.18}$$

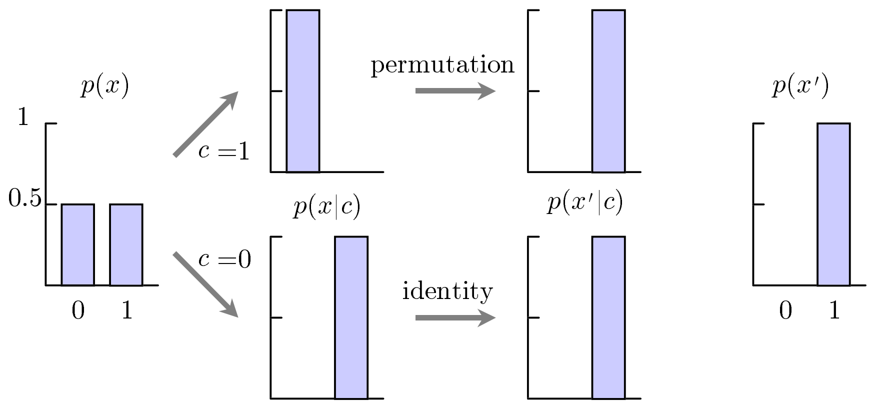

Figure 3.4: Illustration of the separation analysis procedure for a binary closed-loop controller acting on a binary state system.

so as to obtain

$$p(x') = \sum_c p(c) \sum_x p(x'|x,c)p(x|c), \tag{3.19}$$

where $p(c) = \sum_x p(c|x)p(x)$. This last equation shows that a closed-loop controller is essentially an open-loop controller acting on the basis of $p(x|c)$ instead of $p(x)$. Thus, given that c is fixed, the entropy $H(X'|c)$ can be calculated in the manner of Eq.(3.11) by replacing $p(x)$ with $p(x|c)$. In this case, the corresponding pure actuation that must be considered is

$$\begin{array}{c} X|c \xrightarrow{c} X'|c \\ p(x|c) \overset{p(x'|x,c)}{\longrightarrow} p(x'|c) \end{array}, \tag{3.20}$$

and is characterized by the entropy reduction $\Delta H'_c = H(X|c) - H(X'|c)$.

To illustrate this procedure more precisely, let us imagine that a binary state controller with $\mathcal{C} = \{0, 1\}$ "measures" a binary state system, $\mathcal{X} = \{0, 1\}$, in such a way that X conditioned on C is a deterministic random variable (figure 3.4). Moreover, let us assume that the controller is only allowed to use permutations of the state of $\mathcal{X}$ as actuation rules. For $\dim \mathcal{X} = 2$, this corresponds, as shown in figure 3.4, as using the identity transformation ($c = 0$) or single permutations $0 \leftrightarrow 1$ ($c = 1$). In that case, it is clear that after measuring the state of $\mathcal{X}$, the controller can control the state X' to a fixed value with probability one simply by permuting the measured value of X, i.e., the value for which $p(x|c) = 1$. Under this control action, the random variable X' conditioned on $C = c$ is forced to be deterministic for all c, implying that X' is deterministic as well, regardless of the statistics of C. In the situation depicted in the above figure where $p_X(0) = p_X(1) = 1/2$, this means that $\Delta H = 1$ bit. In open-loop control, only $\Delta H_C = 0$ can be achieved. Such a separation of closed-loop controllers will be very useful in the remainder.

3.4. Controllability

We are now in position to study the problem of controllability stated at the beginning of the chapter. In its simplest expression, we define a system to be *perfectly controllable* at a particular value x if for every $x' \in \mathcal{X}$ there exists a least one control value c such that $p(x'|x,c) = 1$. Let $\mathcal{C}_x$ denote the set of values c such that $p(x'|x,c) = 1$ over all x'. If we suppose that $\Pr\{\mathcal{C}_x\} \neq 0$, then as a necessary and sufficient condition for perfect controllability we have the following [40].

Proposition 2. *A system is perfectly controllable at x if and only if $p(x'|x) \neq 0$ for all x' and there exists a non-empty subset $\mathcal{C}_x \subseteq \mathcal{C}$ such that*

$$H(X'|x,C) = -\sum_{x' \in \mathcal{X}, c \in \mathcal{C}_x} p(x',c|x) \log p(x'|x,c) = 0. \tag{3.21}$$

Proof. If x is controllable, then for each control value $c \in \mathcal{C}_x$ there exists one value of x' such that $p(x'|x,c) = 1$. Thus $H(X'|x,c) = 0$ for $c \in \mathcal{C}_x$, and

$$H(X'|x,C) = \sum_{c \in \mathcal{C}_x} H(X'|x,c) p(c) = 0. \tag{3.22}$$

Also, $p(x'|x,c) = 1$ for all x', and for at least one c such that $p(c) \neq 0$, implies

$$p(x'|x) = \sum_c p(x'|x,c) p(c) \neq 0 \tag{3.23}$$

for all x'. Now, to prove the converse, note that if $p(x'|x) \neq 0$ for all x' then there exists at least one c for which $p(x'|x,c) \neq 0$. However, if in addition $H(X'|x,C) = 0$, then X' must be deterministic for some value of C. Hence, for all x' there exists a c such that $p(x'|x,c) = 1$. $\square$

From a perspective centered on information, Eq.(3.22) has the desirable feature of being interpretable as the residual uncertainty or uncontrolled variation left in the output X' when C is fixed for a given (pure) initial state x [40]. If one regards C as an input to a communication channel (the actuation channel), and X' as the channel output, then the degree to which X' can be controlled by fixing C can be identified with the information

$$I(X';C|x) = H(X'|x) - H(X'|x,C). \tag{3.24}$$

By analogy with what was presented in section 2.2, we see that $H(X'|x,C)$ can be interpreted as a *control loss* which appears as a negative contribution in $I(X';C|x)$, the number of bits of accuracy to which specifying the inputs specifies the outputs. In a global sense, one may also look at the controllability of a system over all the input $x \in \mathcal{X}$, and in that respect define

$$L \equiv H(X'|X,C) = \sum_x H(X'|x,C) p(x) \tag{3.25}$$

as the *average control loss* of the combined system $\mathcal{X}+\mathcal{C}$. A system is then perfectly controllable, for all the inputs x such that $p(x) \neq 0$, if $L = 0$ (recall that entropy is always positive); it is *approximately controllable* otherwise. The results in the remaining of this section relate $H(X'|X,C)$ with other quantities of interest.

Proposition 3. *Under the assumption that X' is deterministic conditioned on X, C and E, then $L \leq H(E)$ with equality iff $H(E|X',X,C) = 0$.*

Proof. Using the inequality $H(A) \leq H(A,B)$, and the chain rule for joint entropies, one can write

$$\begin{aligned} H(X'|X,C) &\leq H(X',E|X,C) && (3.26)\\ &= H(E|X,C) + H(X'|X,C,E). && (3.27)\end{aligned}$$

However, $H(X'|X,C,E) = 0$ by assumption, so that

$$\begin{aligned} H(X',E|X,C) &= H(E|X,C) && (3.28)\\ &= H(E) && (3.29)\end{aligned}$$

where the last equality follows from $E \perp\!\!\!\perp X, C$ (see the DAGs of figure 3.2). Now, from the chain rule

$$H(X',E|X,C) = H(X'|X,C) + H(E|X',X,C), \tag{3.30}$$

it is clear that the equality in Eq.(3.26) is achieved iff $H(E|X',X,C) = 0$. □

Intuitively, this result shows that the uncertainty associated with the control of $\mathcal{X}$ is upper bounded by the noise level introduced by the environment. In this sense, it seems clear that one goal of a control system is to protect $\mathcal{X}$ against $\mathcal{E}$ so as to ensure that $\mathcal{X}$ is minimally affected by the noise. (The other goal is to put the controlled system in a particular state.) Following the interpretation of the control loss, in the limit where $L = 0$ the state of the controlled system shows no variability given the initial state and the control action, even in the presence of a perturbing environment, and should thus be independent of the environment. This is the essence of the two next theorems which hold for the same conditions as proposition 3.4.2.

Theorem 4. $I(X';E|X,C) = L$

Proof. From the chain rule

$$I(A;B) = H(A) - H(A|B), \tag{3.31}$$

we can easily derive

$$I(X';E|X,C) = H(X'|X,C) - H(X'|X,C,E). \tag{3.32}$$

Thus, $I(X';E|X,C) = H(X'|X,C)$ where, again, we have used the deterministic condition on X'. □

Theorem 5. $I(X';X,C,E) = I(X';X,C) + L$

Proof. Using the chain rule of mutual information, Eq.(3.31), we can write

$$\begin{aligned} I(X';X,C,E) &= H(X') - H(X'|X,C,E) &(3.33)\\ &= H(X') - H(X'|X,C,E) \\ &\quad +H(X'|X,C) - H(X'|X,C) &(3.34)\\ &= I(X';X,C) + I(X';E|X,C). &(3.35)\end{aligned}$$

In the last equality we have used Eq.(3.32). Now, by substituting $I(X';E|X,C) = L$, we obtain the desired result. □

As a direct corollary of these two theorems, we have that a system is perfectly controllable on average, i.e., $L = 0$, if and only if $I(X';E|X,C) = 0$ or, equivalently, iff $I(X';X,C,E) = I(X';X,C)$. Hence, a necessary and sufficient entropic condition for perfect controllability is that the final state of the system, after control, is statistically independent of E given X, and C. In that case, the information $I(X';E|X,C)$ transferred from the environment to the system in the form of noise is zero. Alternatively, since $L = I(X';E|X,C)$ measures the information (in bits) about X' that have been irrecoverably lost in correlations with E, the condition $L = 0$ is equivalent to the statement that the environment does not “extract” information from the system. Thus, in view of $L = 0$, the addition of a controller $\mathcal{C}$ to a system $\mathcal{X}$ has to effect of constraining the state X' of entropy $H(X') = I(X';X,C)$ to a sort of *noiseless* subspace decoupled from the environment.

3.5. Entropy reduction

The emphasis in the previous section was on proving upper limits on $H(X'|X,C)$ and conditions for which the average control loss vanishes. A complete and practical study of control systems, however, cannot be restricted to the evaluation of the control loss alone, since this quantity is not a direct physical observable. One information-theoretic quantity that we know is more ‘physical’ is the entropy difference ΔH which enters in Eq.(1.2). Also, as noted in section 3.1, one of the goals of a control system is to minimize the uncertainty associated with X'. And insofar as this uncertainty must be the same for any observer of the controlled system, we argued that it should not be taken as a conditional entropy, but should physically correspond to $H(X')$. It seems, therefore, more desirable to find limiting bounds on the final marginal entropy $H(X')$ than to limit the control loss itself.

In addition to provide a more direct route to the physics of control, the study of $H(X')$ has the advantage that any upper bound on $H(X')$ limits *per se* the conditional entropy $H(X'|X,C)$. To understand this, let us note that, since conditioning reduces entropy, we must have the following nested inequalities:

$$H(X') \geq H(X'|C) \geq H(X'|X,C) \geq H(X'|X,C,E). \tag{3.36}$$

The quantity $H(X'|C)$ in this 'hierarchy' of entropies was introduced in the earlier section on separation analysis. In open-loop or closed-loop control, the conditional entropy $H(X'|C)$ corresponds to the average final entropy attained by all the subdynamics of index c, that is,

$$H(X'|C) = \sum_c p(c)H(X'|c), \tag{3.37}$$

where $H(X'|c)$ is the entropy of the distribution $p(x'|c)$ obtained by propagating $p(x)$ or $p(x|c)$, depending on the type of control (open-loop and closed-loop respectively), using the channel $p(x'|x,c)$. In this section we shall prove two important theorems related to the fact that $H(X') \geq H(X'|C)$, but before we do so, we note the following result.

Proposition 6. *$H(X'|C) \leq H(X) + L$ with equality iff $\mathcal{C}$ is an open-loop controller and $H(X|X',C) = 0$.*

Proof. Repeating essentially the same steps as in the proof of proposition 3.4.2, we have

$$\begin{aligned} H(X'|C) &\leq H(X',X|C) && (3.38)\\ &= H(X|C) + H(X'|X,C) && (3.39)\\ &\leq H(X) + L. && (3.40) \end{aligned}$$

Now, to find the condition which saturates the bound, note that the inequality (3.40) becomes an equality when $C \perp\!\!\!\perp X$ (open-loop control). Moreover, since

$$H(X',X|C) = H(X'|C) + H(X|X',C), \tag{3.41}$$

$H(X'|C) = H(X'|C)$ when $H(X|X',C) = 0$. The converse of the equality part follows in a similar fashion. □

We now show that ΔH in open-loop control is upper bounded by the maximum ΔH_c over all pure actuations of variable c. As a consequence, the maximum value of ΔH can always be attained by choosing a deterministic value for C; any mixture of the control variable either achieves the maximum or yields a smaller value. Explicitly, we have the two following results. (The second one was first stated without proof in [68]).

Proposition 7. *For open-loop control,*

$$\Delta H \leq \Delta H_C, \tag{3.42}$$

where $\Delta H_C = H(X) - H(X'|C)$. The equality is achieved iff $C \perp\!\!\!\perp X'$.

Proof. Using the property $H(X') \geq H(X'|C)$, we write directly

$$\begin{aligned} \Delta H &= H(X) - H(X') && (3.43)\\ &\leq H(X) - H(X'|C) && (3.44)\\ &= \sum_c p(c)\Delta H_c, && (3.45) \end{aligned}$$

where $\Delta H_c = H(X) - H(X'|c)$. Now, let us prove the equality part. If $C \perp\!\!\!\perp X'$, then $H(X') = H(X'|C)$ and $\Delta H = \Delta H_C$. Conversely, if $\Delta H = \Delta H_C$, we must have that $H(X'|C) = H(X')$, and thus $C \perp\!\!\!\perp X'$. (It must be noted that in closed-loop control ΔH cannot be maximized by choosing C since the control variable is determined by the condition $p(c) = \sum_x p(c|x)p(x)$, and $p(x'|c) \neq \sum_x p(x'|x,c)p(x)$.) □

Theorem 8. (Open-loop optimality). *In open-loop control,*

$$\Delta H \leq \max_{c \in \mathcal{C}} \Delta H_c, \tag{3.46}$$

where $\Delta H_c = H(X) - H(X'|c)$. *The equality can always be achieved for the pure controller* $C = \hat{c}$ *where* $\hat{c} = \arg\max_c \Delta H_c$.

Proof. Notice that the average of a set of numbers $\{a_i\}$ always sits in the interval $[\min a_i, \max a_i]$, so that

$$\min_c H(X'|c) \leq \sum_c p(c)H(X'|c) \leq \max_c H(X'|c), \tag{3.47}$$

and, therefore,

$$\begin{aligned} \Delta H &= H(X) - H(X') && (3.48)\\ &\leq H(H) - \min_c H(X'|c) && (3.49)\\ &= \max_c \Delta H_c. && (3.50) \end{aligned}$$

Also, if $C = \hat{c}$ with probability one, $\hat{c}$ being the value for which ΔH_c is maximum, then C must be such that $C \perp\!\!\!\perp X'$, and the average ΔH_C reduces to $\Delta H_{\hat{c}}$. (Note that $H(X') = \max_c H(X'|c)$ does not imply necessarily that $C \perp\!\!\!\perp X'$. Hence, $C \perp\!\!\!\perp X'$ cannot be a necessary and sufficient condition for having $\Delta H = \max_c \Delta H_c$.) □

From the standpoint of the controller, one major drawback of acting independently of the state of $\mathcal{X}$ is that often no information other than that available from the state X itself can provide a reasonable way to determine which subdynamics are optimal or even accessible given the initial state. For this reason, open-loop control strategies implemented independently of the state of the system, or solely on its statistics, usually fail to operate efficiently in the presence of noise because of their inability to react or be adjusted over time. In order to account for all the possible behaviors of a stochastic dynamical system, we have to use the information contained in its evolution by considering a closed-loop control scheme in which the state of the controller is allowed to be correlated with the state of $\mathcal{X}$.

For example, in the case of the signalman, the knowledge of the perfect localization of the incoming train (1 or 2) helped him to decide in which position, a or b, he should actuate the 2-to-1 exclusive junction. In fact, as noted before, $I(X;C) = 1$ was even a prerequisite for control in that case, since any 'open-loop' choice of the junction's state could result in generating with some probability the high risk state 'train crashed'. Yet, just as a non-exclusive 2-to-1 does not require information to function correctly, we expect that ΔH in closed-loop should not only depend on $I(X;C)$, but also on the reduction of entropy obtainable by open-loop control. The next theorem, which constitutes the main result of this thesis, exactly embodies this statement by showing that the maximum improvement that closed-loop control give over open-loop control is limited by the information obtained by the controller.

Theorem 9. (Closed-loop optimality). *The amount of entropy $\Delta H_{\mathrm{closed}}$ that can be extracted from $\mathcal{X}$ by a closed-loop controller satisfies*

$$\Delta H_{\mathrm{closed}} \leq \Delta H_{\mathrm{open}} + I(X;C), \tag{3.51}$$

where

$$\Delta H_{\mathrm{open}} = \max_{q(x),c} \Delta H \tag{3.52}$$

is the maximum entropy decrease that can be obtained by (pure) open-loop control over all input distributions $q(x)$, and $I(X;C)$ is the mutual information gathered by the controller upon observation of the system's state with distribution $p(x)$.

A proof of the result, based on the conservation of entropy of closed systems, was given in [68] following some interesting results found in [37]. Here, we present another proof based on separation analysis which has the advantage over ref.[37] to make explicit the conditions under which we obtain equality in the expression (3.51). These conditions are derived in the next section.

Proof. Given that ΔH_{open} is the optimal entropy reduction for open-loop control over any input distribution, we can write

$$H(X')_{\mathrm{open}} \geq H(X) - \Delta H_{\mathrm{open}}. \tag{3.53}$$

Now, using the fact that a closed-loop controller is formally equivalent to an open-loop controller (possibly mixed) acting on $p(x|c)$, we also have, for each value c,

$$H(X'|c) \geq H(X|c) - \Delta H_{\mathrm{open}}, \tag{3.54}$$

and, on average,

$$H(X'|C) \geq H(X|C) - \Delta H_{\mathrm{open}}. \tag{3.55}$$

(The distribution $p(x|c)$ is a legitimate open-loop input distribution, so the maximum ΔH_{open} must also apply in that case.) At this point, notice that $H(X') \geq H(X'|C)$ implies

$$H(X')_{\mathrm{closed}} \geq H(X|C) - \Delta H_{\mathrm{open}}. \tag{3.56}$$

Hence, we obtain

$$\begin{aligned}\Delta H_{\text{closed}} &\equiv H(X) - H(X')_{\text{closed}} &(3.57)\\ &\leq H(X) - H(X|C) + \Delta H_{\text{open}} &(3.58)\\ &= I(X;C) + \Delta H_{\text{open}}, &(3.59)\end{aligned}$$

which is the desired upper bound for a feedback controller. □

The closed-loop optimality theorem is a general statement about a given set of actuation rules used to control any initial condition X whose distribution is taken in the dense set $\mathcal{P}$ of all the possible input distributions. Specifically, the optimality theorem states that for an *arbitrary* initial state, the set of actuations can be used in a closed-loop fashion to achieve at best $\Delta H_{\text{closed}} = \Delta H_{\text{open}} + I(X;C)$, where ΔH_{open} is the maximum open-loop decrease of entropy obtainable over *any* input distribution. In a weaker sense, the optimality theorem can also be applied for a specific input distribution $p(x)$. In that case, the maximization problem of Eq.(3.52) can be restricted over the subset $\widetilde{\mathcal{P}} \subseteq \mathcal{P}$ containing the distributions $p_X(x)$ and $p(x|c)$ for all values c. This is so because, of all the possible input distributions, $p(x)$ and $p(x|c)$, for all c, are the only distributions entering in the reasoning of the above proof. (The fact that ΔH_{open} cannot be calculated using $p(x)$ alone arises because the maximum entropy decrease in open-loop control starting from the actual distribution $p(x)$ may differ from the maximum ΔH obtained by starting with $p(x|c)$. See section 4.1 for a specific example.)

The same remark also applies for the maximization over the pure actuations: in the statement of the optimality theorem, we demand that ΔH_{open} be calculated for a deterministic random variable C since we know that optimality, in the open-loop sense, is satisfied when $C = \hat{c}$ with probability one. However, if we want the optimality theorem to hold in the case of a specific $p(x)$, ΔH_{open} may not be calculated for a pure controller; in fact, it could be associated with ΔH_C or $\Delta H'_C$, depending on which one is the maximum.

3.6. Optimality and side information

The fact that ΔH_{closed} may be greater than ΔH_{open} is in agreement with the intuitive observation that if a controller is accorded information regarding the state of $\mathcal{X}$, then a decision rule for selecting control strategies based on that information can improve upon a rule neglecting the information. In fact, feedback control can always be made to perform at least as well as open-loop control, simply by applying open-loop actuation rules as closed-loop ones. Improvement, though, is not a systematic requirement in control: it is always possible to design a pathological feedback controller which could randomize the state of a system, so that $H(X')_{\text{closed}} \geq H(X')_{\text{open}}$, even in the presence of a non-vanishing mutual information. Hence the need for an inequality in the expression (3.51). The precise statement of the optimality theorem is that, in the case where closed-loop control improves upon open-loop control, the improvement

$$\Delta\mathcal{H} \equiv \Delta H_{\text{closed}} - \Delta H_{\text{open}} \tag{3.60}$$

must be bounded by $I(X;C)$. Equivalently, the *efficiency* of a closed-loop control, defined as

$$\eta \equiv \frac{\Delta\mathcal{H}}{I(X;C)}, \tag{3.61}$$

must be such that $\eta \leq 1$. (Note that we can talk of improvement since $I(X;C) \geq 0$. This means that ΔH_{closed} cannot be systematically smaller that ΔH_{open}, in agreement with the basic observation, noted in section 1.2, that open-loop control strategies are a subset of closed-loop strategies.)

The closed-loop control strategies for which $\eta = 1$ are called *optimal*. Following the same line of reasoning as in the optimality theorem for open-loop control, it can be proved that a necessary, though not sufficient, condition for closed-loop optimality is the following.

Proposition 10. *If $C \perp\!\!\!\perp X'$ and $\Delta H_C = \Delta H'_C = \Delta H_{\text{open}}$, where*

$$\Delta H_C = H(X) - H(X'|C)_{\text{open}} \tag{3.62}$$

$$\Delta H'_C = H(X|C) - H(X'|C)_{\text{closed}}, \tag{3.63}$$

then $\Delta H_{\text{open}} = \Delta H_{\text{open}} + I(X;C)$. (Subscripts have been added to $H(X'|C)$ to stress the fact that $p(x'|c)$ in open-loop control may differ from $p(x'|c)$ obtained in a closed-loop fashion.)

Proof. Suppose that $C \perp\!\!\!\perp X'$ and $\Delta H_C = \Delta H'_C = \Delta H_{\text{open}}$. Then, using the open-loop optimality theorem, we can write

$$\begin{aligned} H(X')_{\text{open}} &= H(X'|C)_{\text{open}} = H(X) - \Delta H_C && (3.64)\\ &= H(X) - \Delta H_{\text{open}}, && (3.65)\\ H(X')_{\text{closed}} &= H(X'|C) = H(X|C) - \Delta H'_C && (3.66)\\ &= H(X|C) - \Delta H_{\text{open}}. && (3.67) \end{aligned}$$

Finally, substracting Eq.(3.65) from Eq.(3.67) yields

$$\begin{aligned} H(X')_{\text{open}} - H(X')_{\text{closed}} &= H(X) - H(X|C) && (3.68)\\ &= I(X;C). && (3.69) \end{aligned}$$

Hence the result. Note that $C \perp\!\!\!\perp X'$ and $\Delta H_C = \Delta H'_C = \Delta H_{\text{open}}$ are not necessary conditions for optimality simply because $H(X')_{\text{open}} = H(X) - \Delta H_{\text{open}}$ does not imply $C \perp\!\!\!\perp X'$ (similarly for the closed-loop case). □

Corollary 11. *Closed-loop optimality is also achieved if $p(c) = \sum_x p(c|x)p(x)$ equals one for one value c corresponding to $\hat{c}$, and if $\Delta H_C = \Delta H'_C = \Delta H_{\text{open}}$. In this case, actually, $I(X;C) = 0$.*

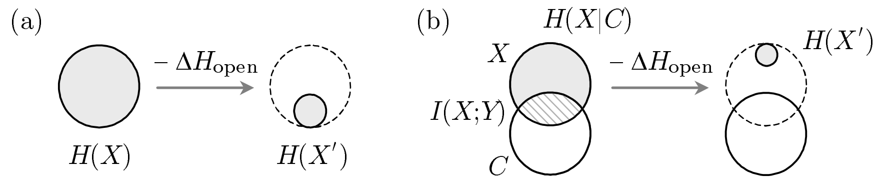

Figure 3.5: Entropy bubbles representing optimal entropy reduction in (a) open-loop control and (b) closed-loop control.

The result of the condition stated in proposition 3.6.1 is illustrated in figure 3.5 which shows how one can reduce the entropy $H(X|C)$ (shaded region in figure 3.5b) by ΔH_{open}, the maximum open-loop control entropy reduction (figure 3.4a). From the entropy bubbles, it is clear that the maximum improvement that closed-loop can give over open-loop control is limited by the information obtained by the controller (dashed area in the figure). Also, we can see that the maximum entropy decrease achievable, both in closed-loop and open-loop control, is $\Delta H = H(X)$. In that case, $H(X') = L = 0$.

In practice, several aspects of a control system will prevent a controller (open-loop or closed-loop) from being *maximally optimal*, i.e., from achieving $\Delta H = H(X)$. In the following, we list and discuss three main reasons that mainly pertain to closed-loop control.

- Measurement devices usually do not have access to the whole state space of the system they measure. Hence $I(X;C) \leq H(X)$. The loss $H(X|C)$ gives a rough measure of the "area" of the state space of $\mathcal{X}$ not covered by $\mathcal{C}$.

- Noiseless communication channels are a rare exception in the real world. This means that losses measured by $H(X|C)$ must necessarily occur in the measurement channel. Note that any post-processing of the information gathered by $\mathcal{C}$ cannot increase the mutual information $I(X;C)$, and thus improve the inferences that would result from the data. This comes as a result of the so-called *data processing inequality* [19] which states, basically, that for a double channel $X \to C \to C'$ with $p(x,c,c') = p(x)p(c|x)p(c'|c)$, we must have $I(X;C) \geq I(X;C')$. In particular, if $C' = g(C)$ then $I(X;C) \geq I(X;g(C))$. Thus, functions of the data C cannot increase the information about X.

- Control actuations that cover the whole state space of a system are very "costly" to implement. In practice, it is usually the case that actuators only modify the dynamics of a controlled system locally in space. To illustrate this, consider the equation of a *thermostat*

$$\dot{x}(t) = \alpha[x^* - x(t)]. \tag{3.70}$$

where $x(t) \in \mathbb{R}$. In this equation, the variable $x(t)$ on the right-hand side represents the feedback component $c(t) = x(t)$ of a control system which tries to constrain the state $x(t)$ to the target state x^*. By defining $y(t) = x^* - x(t)$, the solution of the above equation can be calculated easily:

$$y(t) = y(0)e^{-\alpha t}. \tag{3.71}$$

Now, by choosing $X \sim \mathcal{U}(d)$, the differential entropy of the system can be estimated to be

$$H(X(t)) \propto \log d - \alpha t. \tag{3.72}$$

Hence, even if $c(t) = x(t)$ (perfect measurement with $I(X;C) = \infty$), we see that ΔH is essentially limited by the dissipation factor α which, according to the difference equation

$$\Delta x \simeq \alpha[x^* - x(t)]\Delta t, \tag{3.73}$$

gives the approximate radius of action of the controller over a time interval Δt.

It is interesting to note for the data processing inequality that, even though the manipulation of C cannot increase $I(X;C)$, it is possible to access the state X using *side information*, i.e., any data Y that is correlated with X in such a way that $I(X;Y) \neq 0$, and possibly $I(X;Y) \geq I(X;C)$. What should enter in the inequality (3.51) is thus the mutual information between X and the variable used as the basis for the choice of the actuation rules (which is not necessarily the control variable). For example, a control system $\mathcal{C}_1$ could infer the state of a system $\mathcal{X}$ not directly by observation, but indirectly by solving the exact dynamics of $\mathcal{X}$ using the past knowledge of its state. In that case, the controller could act on the basis of the calculated state in the same way another control system, say $\mathcal{C}_2$, would operate by measuring $\mathcal{X}$. Thus, even though $\mathcal{C}_1$ does not interact directly with the controlled system, it should be considered as a closed-loop control system; from an outside perspective on the system $\mathcal{X} + \mathcal{C}_1$, $I(X;C_1) \neq 0$. (Controllers such as $\mathcal{C}_1$ are often referred to as *feed-forward* controllers.)

3.7. Continuous extensions of the results

The theorems presented in the two previous sections can be generalized beyond the DSDT models of figure 3.2 to continuous-states systems, multi-stages processes and continuous-time dynamics. For each of these three classes, the modifications to the results of sections 3.4 and 3.5 are as follows.

- *Continuous-state systems.* If $\mathfrak{X}$ and $\mathfrak{C}$ are continuous state spaces, then *a priori* the only modification that has to be made is to change the sums involved in the expressions of entropy and mutual information to integrals over the state variables. (Recall that $I(X;C)$ is still positive-definite for continuous random variables.) A more careful analysis of the results, however, reveals that this modification is not sufficient in the case of controllability, since the concept of a deterministic continuous random variable is somewhat ill-defined, and is not associated with the condition $H(X) = 0$. To circumvent this difficulty, we may extend the concept of controllability to continuous variables in the following manner: a system is defined to be perfectly controllable at $x \in \mathcal{X}$ if for every $x' \in \mathcal{X}$ there exists at least one control value c which forces the system to reach a small neighborhood of radius $\varepsilon > 0$ around x' with very high probability. From this definition, the connection with all of our previous results on controllability simply follows by coarse-graining the continuous variables X and X' using a "grid" of size $\Delta \gtrsim \varepsilon$, and by applying the necessary definitions to the discrete random variables X^Δ and X'^Δ. The recourse to this coarse-grained description of a continuous state system has also the advantage that $H(X^\Delta)$ and $I(X^\Delta; C^\Delta)$ are well-defined functions which cannot be infinite, contrary to their continuous analogs.

- *Higher-order processes.* For a process of the form $X_n \to X_{n+1} \to \cdots \to X_{n+m}$ controlled by the history $C^{m-1} = C_n, C_{n+1}, \ldots, C_{n+m-1}$, one applies the results on controllability with the only difference that the control variable C is replaced by the control history C^{m-1}. Now, for the results covering the section on entropy reduction, one can 'concatenate' the DAGs of figure 3.2 over several time steps in order to apply the optimality theorems on a one time-step basis, and in particular write
$$\Delta H^n_{\text{closed}} \leq \Delta H^n_{\text{open}} + I(X_n; C_n), \tag{3.74}$$
where $\Delta H^n = H(X_n) - H(X_{n+1})$. However, this sort of "memoryless" approximation is not the most general model of control we can think of. For instance, it does not say anything about the use of past information to reduce the entropy, nor it exploits non-Markovian features of coupled dynamics, such as memory effect or learning, to increase the information of the controller above what is permitted by $I(X_n; C_n)$. Work is ongoing along this direction.

- *Continuous-time dynamics.* The extension of our results on controllability to dynamics evolving continuously in time is direct: in this case, one essentially follows the directives of the previous paragraph on multi-steps processes, the history now being the continuous history $\{C(t) : t_0 \leq t \leq t_f\}$ over the interval $[t_0, t_f]$. For the results on entropy reduction, the generalization to continuous time is more problematic. On the one hand, the state $X(t)$ of a continuous-time system can be sampled at two discrete time instants, $X(t)$ and $X(t + \Delta t)$, separated by the interval Δt. In this case, the inequality (3.51) can be modified in the following manner:
$$\frac{\Delta H_{\text{closed}}}{\Delta t} \leq \frac{\Delta H_{\text{open}}}{\Delta t} + \frac{I(X(t); C(t))}{\Delta t}. \tag{3.75}$$
Now, in the limit $\Delta t \to 0$, we obtain well-defined rates for ΔH_{closed} and ΔH_{open}, assuming that $H(X(t))$ is a smooth and differentiable function of t. Unfortunately, the quantity
$$\lim_{\Delta t \to 0} \frac{I(X(t); C(t))}{\Delta t} \tag{3.76}$$
does not constitute a rate, for $I(X(t); C(t))$ is not a differential which goes to zero as $\Delta t \to 0$.

On the other hand, even our very definition of open-loop control, namely the requirement that $I(X;C)$ be equal to zero prior to control, fails to apply in continuous time. Indeed, the distinction between the estimation step and the actuation step for a continuous-time control process cannot be drawn in a meaningful way, other than by sampling the control process in time. Moreover, even open-loop controllers, which operate continuously over a given time interval, will be such that $I(X(t);C(t)) \neq 0$. One possible solution, in such a case, could be to replace the condition $I(X;C) = 0$ by the requirement $I(X(t);C(t)|\{C(t')\}) = 0$ for $t' < t$. More precisely, let $X(t-\Delta t)$, $X(t)$ and $X(t+\Delta t)$ be three consecutive sampled points of a trajectory $X(t)$. Also, let $C(t-\Delta t)$ and $C(t)$ be the states of a controller during a time interval in which the state of the controlled system is estimated. (The actuation step is assumed to take place between the time instants t and $t+\Delta t$.) Then, by redefining the entropy reductions as

$$\Delta H^t = H(X(t)|C^{t-\Delta t}) - H(X(t+\Delta t)|C^{t-\Delta t}), \tag{3.77}$$

where C^t represents the control history up to time t, we must have

$$\Delta H^t_{\text{closed}} \leq \Delta H^t_{\text{open}} + I(X(t);C(t)|C^{t-\Delta t}). \tag{3.78}$$

Now, since $I(X(t-\Delta t);C(t-\Delta t)|C^{t-\Delta t}) = 0$, we also have

$$\begin{aligned} \Delta H^t_{\text{closed}} &\leq \Delta H^t_{\text{open}} + I(X(t);C(t)|C^{t-\Delta t}) && (3.79)\\ &\quad -I(X(t-\Delta t);C(t-\Delta t)|C^{t-\Delta t}). && (3.80) \end{aligned}$$

Hence, by taking the limit $\Delta t \to 0$ on both sides of the above equation, we obtain the rate equation

$$\dot{H}_{\text{closed}} \leq \dot{H}_{\text{open}} + \dot{I}, \tag{3.81}$$

relating the rate at which entropy is dissipated

$$\dot{H} = \lim_{\Delta t \to 0} \frac{1}{\Delta t} \left[H(X(t)|C^{t-\Delta t}) - H(X(t+\Delta t)|C^{t-\Delta t}) \right], \tag{3.82}$$

with the rate at which the conditional information

$$\Delta I = I(X(t);C(t)|C^{t-\Delta t}) - I(X(t-\Delta t);C(t-\Delta t)|C^{t-\Delta t}) \tag{3.83}$$

is gathered upon estimation. The difference between the conditional information rate $\Delta I/\Delta t$ and the 'wrong' rate defined by Eq.(3.76) lies in the fact that ΔI represents 'new' or 'fresh' information gathered during the latest estimation stage of the control process: it does not include past correlations induced by the control history $C^{t-\Delta t}$. Finally, note that in this picture, there exists a meaningful definition of open-loop control, which is $\Delta I = 0$, simply because, in open-loop control, there is no direct estimation process.

3.8. Thermodynamic aspects of control

The closed-loop optimality theorem, re-written as

$$H(X')_{\text{closed}} \geq H(X')_{\text{open}} - I(X;C), \tag{3.84}$$

establishes a unavoidable limit to the power or performance of any control devices whose designs are based on the possibility to accede low entropy states. Specifically, since a uniform discrete random variable X distributed over $N(X)$ states has entropy $H(X) = \log N(X)$, it can be estimated that the approximate average number of states over which X' is distributed with non-zero probability is

$$N(X') \simeq 2^{H(X')}. \tag{3.85}$$

Thus, from Eq.(3.84), it can be inferred that

$$N(X')_{\text{closed}} \gtrsim 2^{-I(X;C)} N(X')_{\text{open}}. \tag{3.86}$$

In practice, the extend to which the entropy of a system can be decreased will depend on several factors and parameters determining the effect of a controller on a dynamical system. In addition to the physical factors mentioned in section 3.6, the performance of a controller may also be associated with the "cleverness" of an sensor-actuator system: by observing regularities in the behavior of a system, a controller can "learn" and adapt its strategy in order to enhance its prediction abilities, and therefore improve its effectiveness. Interestingly, this aspect of control can be put in correspondence with the modern analysis of Maxwell's demon [37, 73, 72, 74] (see also [14, 57, 47]). Here, controllers, just as Maxwell's demon, are complex and adaptive information gathering and using systems (IGUS) that interact with the controlled system, perceive regularities in the motion of that system, and compress that information in a utilizable form for a given control task.

As a direct consequence of this association, the "energetic" capabilities of a controller must be limited by thermodynamic laws similar to the ones we discussed in relationship with the Maxwell's demon problem. On the one hand, since the overall system $\mathcal{X} + \mathcal{C} + \mathcal{E}$ is a closed system whose dynamics is given in terms of a constant Hamiltonian, Landauer's principle immediately implies that if the controller is initially uncorrelated with the controlled system (open-loop control), then a decrease in entropy ΔH_{open} for the system $\mathcal{X}$ must be compensated for by an increase in entropy of at least ΔH_{open} for the controller and the environment, e.g., conveyed in the form of $\Delta Q_{\text{open}} = (k_B T \ln 2)\Delta H_{\text{open}}$ joules of heat.

A feedback controller, on the other hand, need not necessarily transfer entropy to an environment in decreasing the entropy of $\mathcal{X}$. For example, a controller can use entropy-conserving Hamiltonian dynamics to actuate the state of $\mathcal{X}$ using the mutual information $I(X;C)$, thereby reducing the entropy of $\mathcal{X}$ by an amount ΔH_{closed} without affecting the entropy of the environment ($\Delta H_{\text{open}} = 0$). In this scenario, in analog to what Szilard, Landauer and Bennett showed, the amount by which $\mathcal{C}$ can decrease the entropy of $\mathcal{X}$ without any effect on $\mathcal{E}$ must be bounded above by $I(X;C)$. Again, as argued before, the fact that $\mathcal{C}$ is able to decrease the entropy of $\mathcal{X}$ without affecting $\mathcal{E}$ does not entail the second law of thermodynamics. Ultimately, if the controller is to be a cyclic machine, then it must supply work

$$\Delta W_e = (k_B T \ln 2) I(X;C) \tag{3.87}$$

after each actuation step in order to erase the mutual information $I(X;C)$ into an environment at temperature T. This price of information corresponds exactly to Landauer's erasure cost mentioned earlier: it is the thermodynamic cost of information needed to rehabilitate the validity of the second law of thermodynamics when generalized to include information. Here, in accordance with Eq.(3.51), the second law can be written as

$$\Delta Q_{\text{tot}} \geq \Delta Q_{\text{open}} + \Delta Q_e, \tag{3.88}$$

where $\Delta Q_e = \Delta W_e$ is the heat dissipated upon erasure of $I(X;C)$ bits of information.

4 | Applications

This chapter offers a collection of applications of the formalism developed in this thesis and the results derived previously. We first begin by analyzing a simple system aimed at the control of a particle enclosed in a box. The setting of this example is very similar to the one considered originally by Szilard in its analysis of Maxwell's demon (see [35] for references), and is given here to illustrate in more details how the quantity ΔH_{open} is to be calculated. As a second application, we treat the control of binary states systems through the use of controllers apparented to logical gates. The model presented in this section concerns the well-known controlled-not (CNOT) gate modified here to allow its use in closed-loop control (both noiseless and noisy) of an ensemble of bits. This example is very instructive in that it illustrates in a simple fashion fundamental limitations on feedback control and the conditions for which controllers are optimal. Finally, we show how our analysis can be applied in the context of chaotic control to calculate the size of the interval at which a noisy closed-loop controller can constrain the state of a chaotic map. This application to nonlinear systems is very important for two main reasons: (i) nonlinear maps, and nonlinear systems in general, offer a natural framework for calculating information measures through the use of Lyapunov exponents; and (ii) chaotic actuation rules are usually entropy increasing, which means that $I(X;C)$ should be non-zero if we want $\Delta H_{\text{closed}} \geq 0$. The results of this last section constitute the first step towards a general information-theoretic study of noise in nonlinear systems. Other applications and open problems are suggested in the next and final chapter.

4.1. Controlling a particle in a box: a counterexample?

Consider a system $\mathcal{X}$ consisting of $N > 2$ states (N even) and a binary sensor-actuator $\mathcal{C}$ which consists of 2 states $c \in \{0, 1\}$. The states of the system are mapped onto the states of the controller, upon measurement, as follows:

$$p(c=0|x) = \begin{cases} 1, & x \leq N/2 \\ 0, & x > N/2 \end{cases} \tag{4.1}$$

$$p(c=1|x) = \begin{cases} 0, & x \leq N/2 \\ 1, & x > N/2 \end{cases} \tag{4.2}$$

That is, half of the states ($x \leq N/2$) are "observed" by the state $c = 0$, while the other half is "observed" by $c = 1$. Furthermore, we assume the following dynamics for the

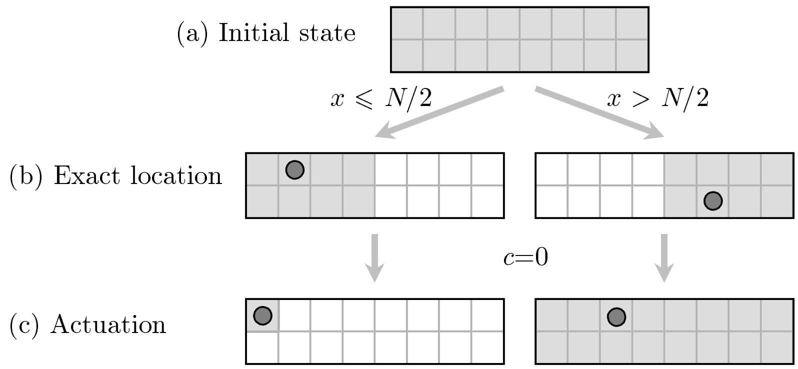

Figure 4.1: Different stages in the control of a particle in a box of N states. (a) Initial equiprobable state. (b) Exact localization. (c) Open-loop actuation $c = 0$.

whole system:

$$\text{Initial state} \qquad p(x) = N^{-1}, \text{ for all } x = 1, \ldots, N$$

$$\text{Actuation matrices} \qquad \begin{aligned} p(x'|x, c=0) &= \begin{cases} \delta_{x',1}, & x', x \leq N/2 \\ 1/N, & x', x > N/2 \end{cases} \\ p(x'|x, c=1) &= \begin{cases} 1/N, & x', x \leq N/2 \\ \delta_{x',N}, & x', x > N/2 \end{cases} \end{aligned} \tag{4.3}$$

Physically, the actuation matrix $p(x'|x, c)$ acts as a piston inserted in the middle of a box containing a single particle which is initially in one of the N states so that $H(X) = \log N$ (figure 4.1a). When $c = 0$, the piston is moved to the "left" part of the box and a particle x located in the right half of the box ($x > N/2$) is able to move throughout the entire volume of the box (figure 4.1c-d on the right), while if the particle is in the left part of the enclosure ($x \leq N/2$), it is compressed to one final position corresponding to $x' = 1$ (figure 4.1c). For $c = 1$, the complementary actuation is adopted: the piston is moved to the "right" part and a particle at $x > N/2$ is compressed to $x' = N$, and moves freely otherwise.

Open-loop control

In open-loop control, we choose $p_C(0) = 1$, $p_C(1) = 0$ as our control strategy. Using the control rules, the final state of the system can be calculated using the open-loop decomposition, Eq.(3.13), which yields

$$\begin{aligned} p(x') &= \frac{1}{N} \sum_x p(x'|x, c) = \frac{1}{N} \sum_{x \leq N/2} \delta_{x',1} + \frac{1}{N} \sum_{x > N/2} \frac{1}{N} \\ &= \frac{1}{2N} + \frac{1}{2} \delta_{x',1}. \end{aligned} \tag{4.4}$$

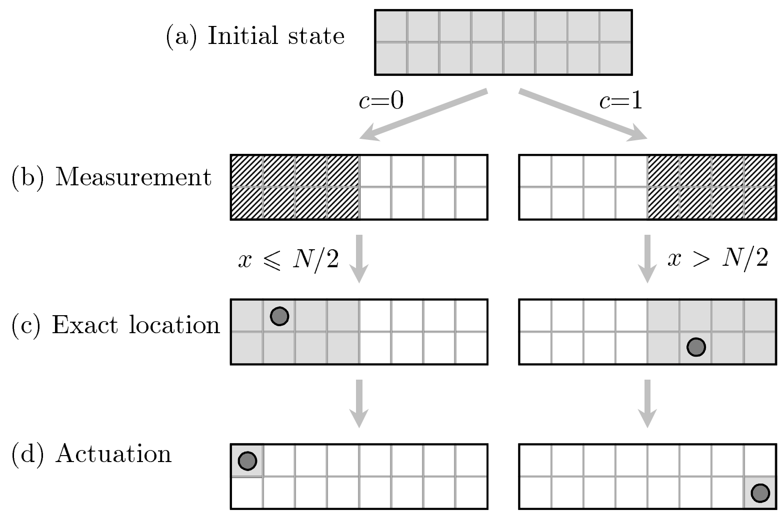

Figure 4.2: Closed-loop control stages. (a) Initial equiprobable state. (b) Coarse-grained measurement of the position of the particle. (c) Exact location of the particle in the box. (d) Actuation according to the controller's state.

We thus obtain

$$\begin{aligned} H(X')_{\mathrm{open}} &= -\left(\frac{1}{2N}+\frac{1}{2}\right)\log\left(\frac{1}{2N}+\frac{1}{2}\right) - \sum_{x'\neq 1}\frac{1}{2N}\log\frac{1}{2N} \\ &= -\left(\frac{1}{2N}+\frac{1}{2}\right)\log\left(\frac{1}{2N}+\frac{1}{2}\right) - \frac{N-1}{2N}\log\frac{1}{2N}. \end{aligned} \tag{4.5}$$

Closed-loop control

In closed-loop control, the actuation variable c is chosen according to Eqs.(4.1) and (4.2) (see figure 4.2). The mutual information in that case can be calculated easily since

$$p_C(0) = p_C(1) = \sum_x p(x)p(c=1|x) = \frac{1}{2}, \tag{4.6}$$

so that

$$\begin{aligned} I(X;C) &= \sum_{x,c} p(x)p(c|x)\log\frac{p(x)p(c|x)}{p(x)p(c)} = \sum_{x,c} p(x)p(c|x)\log\frac{p(c|x)}{p(c)} \\ &= \frac{1}{N}\sum_{x\leq N/2}[p(c=0|x)\log 2p(c=0|x) + p(c=1|x)\log 2p(c=1|x)] \\ &\quad +\frac{1}{N}\sum_{x>N/2}[p(c=0|x)\log 2p(c=0|x) + p(c=1|x)\log 2p(c=1|x)] \\ &= \log 2 \end{aligned} \tag{4.7}$$

Finally, the distribution $p(x')$ in closed-loop control becomes

$$\begin{aligned} p(x') &= \frac{1}{N}\sum_{x\leq N/2}\big[p(x'|x,c=0)p(c=0|x)+p(x'|x,c=1)p(c=1|x)\big] \\ &\quad +\frac{1}{N}\sum_{x>N/2}\big[p(x'|x,c=0)p(c=0|x)+p(x'|x,c=1)p(c=1|x)\big] \\ &= \frac{1}{N}\sum_{x\leq N/2}\delta_{x',1}+\frac{1}{N}\sum_{x>N/2}\delta_{x',N}=\frac{1}{2}\left(\delta_{x',1}+\delta_{x',N}\right). \end{aligned} \tag{4.8}$$

Thus, $H(X')_{\text{closed}}=\log 2$.

In light of these results, a contradiction seems to arise because

$$\Delta H_{\text{closed}} \nleq \Delta H_{\text{open}}+I(X;C) \tag{4.9}$$

for some values of N. In fact, figure 4.3 shows that from $N>6$ (outside the dashed region), ΔH_{closed} deviates significantly from what is allowed by the closed-loop optimality theorem. What has not been taken into account in the analysis, however, is that ΔH_{open} requires the calculation of $\max\Delta H$ over the set $\widetilde{\mathcal{P}}=\{p(x),p(x|c)\}$ since, as mentioned, the optimality theorem does not necessarily hold if both ΔH_{open} and ΔH_{closed} are calculated using $p(x)$. If one properly re-calculates the reduction of entropy in open-loop control, the value $\Delta H_{\text{open}}=\log N/2$ should be found, so that $\Delta H_{\text{closed}}<\Delta H_{\text{open}}+I(X;C)$ for all $N>2$. The value $\log N/2$ is achieved by an equiprobable distribution covering either the left half or the right half part of the box, which is exactly the distribution $p(x|c)$ after the measurement of the particle with the initial state $p(x)=N^{-1}$.

4.2. Binary control automata

In the example of the signalman, we noted that the gathering of one bit of information was necessary to control the ensemble $\{1,2\}$, $p(1)=p(2)=1/2$, to a single final state of zero entropy. Now, what happens if $p(1)\neq 1/2$, or if the controller has access to less than one bit of information? To answer these two questions, we consider the following toy model. Let the controlled system $\mathcal{X}$ have a binary state space $\mathcal{X}=\{0,1\}$ with $p_X(0)=a$ and $p_X(1)=1-a$, $0\leq a\leq 1$. Also, let $\mathcal{C}$ be a binary state controller $\mathcal{C}=\{0,1\}$ restricted to use the two actuation channels

$$p(x'|x,c=0):\begin{pmatrix}1 & 0\\ 0 & 1\end{pmatrix};\qquad p(x'|x,c=1):\begin{pmatrix}0 & 1\\ 1 & 0\end{pmatrix}. \tag{4.10}$$

(In matrix notation, columns are labelled by the values of X, and the rows are indexed by the values of X'. With this convention, the propagation $p(x')=\sum_x p(x'|x)p(x)$ is written as the matrix multiplication

$$\begin{pmatrix} p(x_1') \\ p(x_2') \\ \vdots \\ p(x_n') \end{pmatrix} = \begin{pmatrix} p(x_1'|x_1) & p(x_1'|x_2) & \cdots & p(x_1'|x_m) \\ p(x_2'|x_1) & & & \\ \vdots & & \ddots & \vdots \\ p(x_n'|x_1) & & \cdots & p(x_n'|x_m) \end{pmatrix} \begin{pmatrix} p(x_1) \\ p(x_2) \\ \vdots \\ p(x_m) \end{pmatrix} \tag{4.11}$$

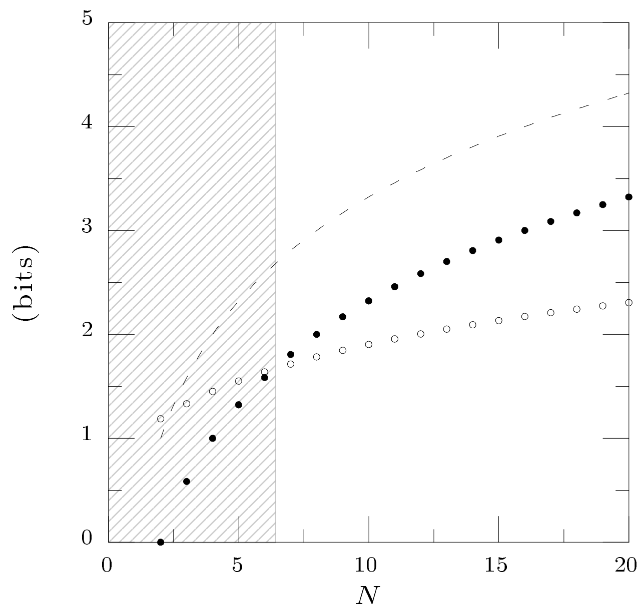

Figure 4.3: Apparent violation of the closed-loop optimality theorem. (•) ΔH_{closed} as a function of N. (○) $\Delta H_{\text{open}} + I(X;C)$ versus N. Outside the dashed region $\Delta H_{\text{closed}} > \Delta H_{\text{open}} + I(X;C)$. The dashed line represents the corrected result when is ΔH_{open} is calculated according to the proper definition.

known as the *Kolmogorov-Chapman* equation [25].) The first channel of Eq.(4.10) is called the *copy channel*, whereas the second one is the *anti-copy channel*. As an equivalent representation of the whole system $\mathcal{X}+\mathcal{C}$, we can think of the matrices of Eq.(4.10) as representing the transitions of a CNOT (reversible) gate

c	x	$c' = c$	x'
0	0	0	0
0	1	0	1
1	0	1	1
1	1	1	0

$=$

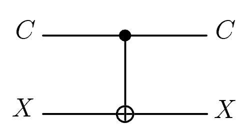

(4.12)

which takes the complement $\bar{x} = x \oplus 1 \mod 2$ of x whenever $c = 1$.

Now, since these actuation rules consist of permutations, $H(X')_{\text{open}} \geq H(X)$ with equality if a pure controller is used (either $c = 0$ or $c = 1$), or if $H(X) = H_{\max} = 1$ bit. These two conditions are in agreement with the optimality theorem for open-loop control, since for a deterministic random variable C or for $p_X(0) = p_X(1) = 1/2$, it can be verified that $X' \perp\!\!\!\perp C$. Thus, for open-loop control we have $\Delta H_{\text{open}} = 0$.

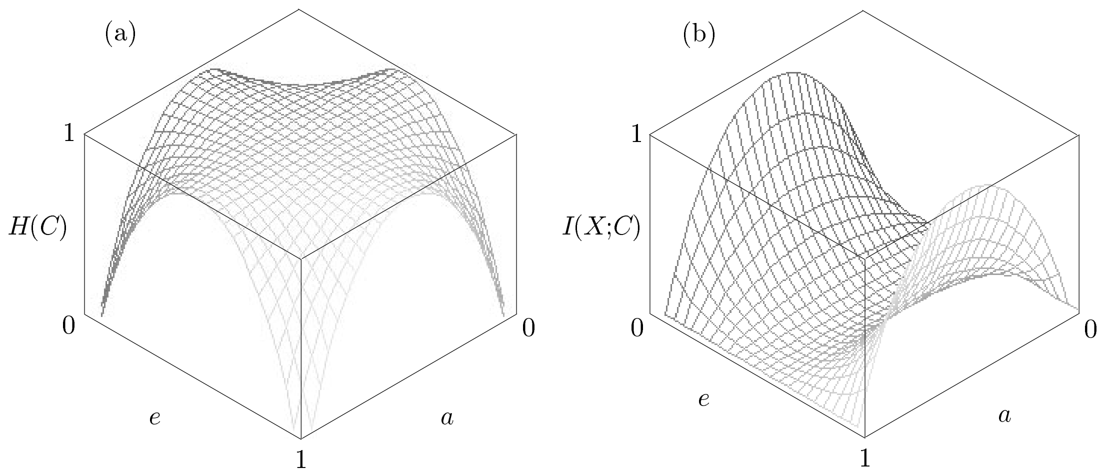

Figure 4.4: (a) $H(C)$ as a function of of the measurement error e and the initial parameter a. (b) $I(X;C)$ as a function of e and a.

Noisy measurement

In closed-loop control, we consider a measurement channel modeled by a binary symmetric channel (BSC) [19]:

$$p(c|x) : (1-e)\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} + e\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} 1-e & e \\ e & 1-e \end{pmatrix} \tag{4.13}$$

where $e \in [0,1]$ (see Eq.(4.11) to identify the elements of the matrix). This transition matrix represents a binary measurement channel in which the input is copied with probability $1-e$, and is changed otherwise, i.e., anti-copied, with probability e. Thus, it can be said that for $0 < e < 1/2$, the channel is a noisy copy channel, whereas it is a noisy anti-copy channel for $1/2 < e < 1$. The mutual information for this channel is

$$\begin{aligned} I(X;C) &= H(C) - \sum_{x=0,1} p(x)H(C|x) \\ &= H(C) - H(e), \end{aligned} \tag{4.14}$$

where

$$H(e) \equiv -e\log e - (1-e)\log(1-e). \tag{4.15}$$

Figure 4.4 shows the plots of $H(C)$ and $I(X;C)$. From these figures, it can be clearly seen that $\mathcal{C}$ gathers a maximum of information when e equals 0 or 1, in which case the measurement channel is noiseless, and thus $I(X;C) = H(X) = H(a)$. Also, $H(C) = 0$ (pure controller) for the four couples $(e,a) \in \{(0,0),(1,0),(0,1),(1,1)\}$. We thus expect the controller to be optimal at these points since $\Delta H'_C = \Delta H_C = \Delta H_{\text{open}} = 0$.

To verify this, we calculate the final state in closed-loop control using Eq.(3.18). The complete solution is $p_{X'}(0) = 1 - e$, $p_{X'}(1) = e$. Hence, $H(X') = H(e)$, and

$$\Delta H_{\text{closed}} = H(X) - H(e). \tag{4.16}$$

Figure 4.5a contains the plot of the entropy decrease ΔH_{closed} as a function of the error parameter e and the initial condition $p_X(0) = a$. This plot is compared, in figure 4.5b, with the mutual information $I(X;C)$. By analyzing these figures, a few results can be inferred. First, we see that, indeed, the controller is optimal for the four couples (e, a) mentioned above. Furthermore, it is also optimal along the lines $e = 0$, and $e = 1$ since $H(X') = 0$ when $X \asymp C$. More precisely, for $e = 0$,

$$p_{X'}(0) = 1; \qquad p_{X'}(1) = 0, \tag{4.17}$$

whereas for $e = 1$,

$$p_{X'}(0) = 0; \qquad p_{X'}(1) = 1. \tag{4.18}$$

(This is very similar to the example given in the section on separability.) In addition to these results, the controller is found to be optimal along the line $a = 0.5$. Unfortunately, this case cannot be understood from the optimality condition of closed-loop control because $I(X';C) \neq 0$.

Except for the mentioned cases, $\Delta H < 0$ even if $I(X;C) \neq 0$ (non-optimal controller). In fact, for $e \simeq 1/2$, it can be seen on figure 4.5b that $\mathcal{C}$ randomizes the state of $\mathcal{X}$ in such a way that $\Delta H < 0$. This does not come as a surprise since, for $I(X;C) \simeq 0$, the controller's action relies essentially on a random bit. Note, however, that the optimality theorem always holds true for this example. Indeed, figure 4.5b shows that

$$\begin{aligned} \Delta H_{\text{closed}} &\leq \Delta H_{\text{open}} + I(X;C) \\ &= I(X;C) \end{aligned} \tag{4.19}$$

for all values of e and a. Analytically, this can be verified by using Eqs.(4.14) and (4.16) in the optimality theorem to obtain the inequality $H(X) \leq H(C)$, which is always satisfied since the measurement channel is doubly stochastic. Note, finally, that $H(X')$ is independent of $H(X)$. Thus, by repeating the control action after a first transition $X \to X'$, we must end up with $H(X'') = H(X') = H(e)$ so that $\Delta H_{\text{closed}} = I(X;C) = 0$. Hence, applying the control action a second time does not change the state of the system.

Noisy actuations

Before concluding this section, let us study the case where the controller implements noisy actuation channels modeled by the following matrices

$$p(x'|x, c=0) \;:\; (1-d)\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} + d\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} 1-d & d \\ d & 1-d \end{pmatrix} \tag{4.20}$$

$$p(x'|x, c=1) \;:\; d\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} + (1-d)\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} d & 1-d \\ 1-d & d \end{pmatrix} \tag{4.21}$$

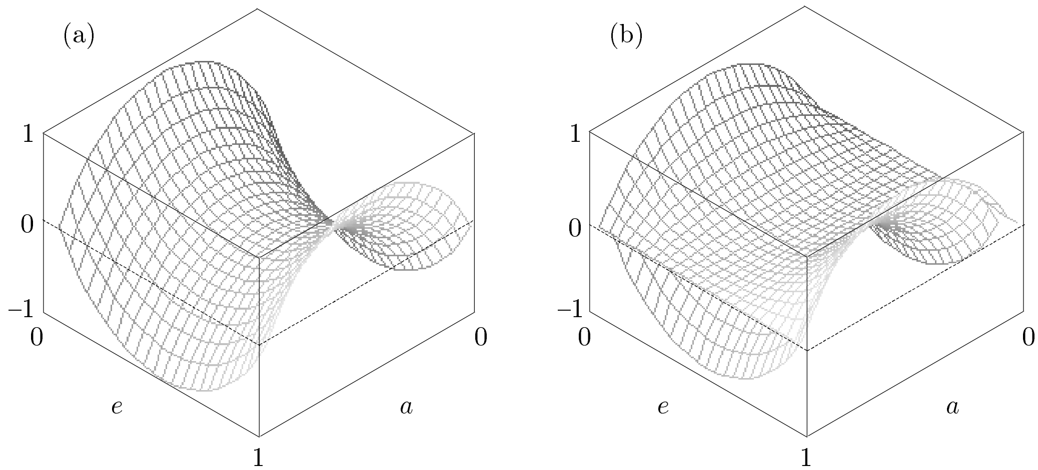

Figure 4.5: (a) ΔH_{closed} as a function of e and a. (b) Comparison of $I(X;C)$ (top surface) and ΔH_{closed} (partly hidden surface).

where $d \in [0,1]$ is the error parameter. These matrices being doubly stochastic, we expect to have $\Delta H < 0$ for almost all initial conditions $p(x)$ and values of d, so that $\Delta H_{\text{open}} < 0$. However, in general this guess would be wrong, since $X \sim \mathcal{U}$ gives rise to $\Delta H_{\text{open}} = 0$ for any value of d. This illustrates an important fact about entropy-increasing actuation rules: if we want to apply the optimality theorem for any initial random variable X, then ΔH_{open} must trivially be taken to equal zero in order to face the possibility of having $X \sim \mathcal{U}$. However, as remarked in chapter 3, if the distributions $p(x)$ and $p(x|c)$ are given, then ΔH_{open} can be evaluated on a case by case basis using the set $\widetilde{\mathcal{P}} = \{p(x), p(x|c)\}$. For the present example, by calculating ΔH_{open} either with $\mathcal{P}$ or $\widetilde{\mathcal{P}}$, it can be shown that, for a perfect binary symmetric measurement channel ($e = 0,1$), the controller is optimal only for an actuation channel with $d = 0,1$. This leads us to conjecture that controllers which use strictly entropy-increasing actuation rules can never be optimal, except for the trivial case where $H(C) = 0$.

To conclude, it should be noted that the controller with $e \neq 0,1$ and $d = 0,1$ (noisy measurement, perfect actuation) controls the final state X' to an entropy $H(e)$ which is independent of X. This is rather interesting considering that the controller with $e = 0,1$ and $d \neq 0,1$ (perfect measurement, noisy actuation) is such that $H(X') = H(d)$. This kind of complementarity between noise components affecting measurement and noise components affecting actuations is believed to be a general feature of controllers. Also, for arbitrary values of e and d, and for $H(X) = 1$ bit, it was observed that

$$\Delta H_{\text{closed}} = [1 - H(d)]I(X;C), \tag{4.22}$$

although no proof of this result was found.

4.3. Control of chaotic maps

We now consider the feedback control scheme proposed by Ott, Grebogi and Yorke (OGY) [48, 63], as applied to the *logistic map*

$$x_{n+1} = r x_n (1 - x_n), \tag{4.23}$$

with $0 \leq r \leq 4$, and $x_n \in [0, 1]$, $n = 0, 1, 2, \ldots$ (See [62] for a review of chaotic control). The OGY method, specifically, consists to apply to Eq.(4.23) small perturbations $r \to r + \delta r_n$ according to

$$\delta r_n = -\gamma (x_n - x^*), \tag{4.24}$$

whenever the state x_n falls into a small *control region* D in the vicinity of a *target* point x^*. As a result of ergodicity of the logistic map, the probability that the trajectory will enter *any* region D goes to one as $n \to \infty$. The point x^* is usually taken to be an unstable fixed point of the unperturbed map ($\delta r = 0$), satisfying the equation $f_r(x^*) = x^*$, that we want to stabilize. Moreover, the *gain* $\gamma > 0$ is fixed so as to ensure the stability of the control actions. The possible stable values for γ, explicitly, are determined by linear analysis as follows. In the vicinity of the point x^*, the map can be linearized to obtain

$$x_{n+1} = f_r(x^*) + \left.\frac{\partial f}{\partial x}\right|_{x=x^*} \delta x_n + O(\delta x_n^2) \tag{4.25}$$

for $r = cte$, where $\delta x_n = x_n - x^*$. Now, by allowing a variation $r \to r + \delta r_n$, the linearization becomes

$$\delta x_{n+1} = \left.\frac{\partial f}{\partial x}\right|_{x=x^*} \delta x_n + \left.\frac{\partial f}{\partial r}\right|_{x=x^*} \delta r_n = \partial_x f(x^*) \delta x_n + \partial_r f(x^*) \delta r_n \tag{4.26}$$

(neglecting all terms of order higher than 2). Thus, using Eqs.(4.24) and (4.26), we see that if we want to reach the state x^* from any state in one iteration of the map, then the gain should be set to the value

$$\gamma = \frac{\partial_x f(x^*)}{\partial_r f(x^*)}. \tag{4.27}$$

In that case, $\delta x_{n+1} = 0$, implying that $x_{n+1} = x^*$.

This procedure, clearly, is valid in the limit of a small control region $D \to 0$; in a more realistic fashion, D need not be infinitesimal if the gain γ is fixed to a constant value that differs slightly from the value prescribed by Eq.(4.27). Indeed, the only requirement for stability, and achievability of the control goals, is that $|\delta x_{n+1}| < |\delta x_n|$ for all n. Hence, γ can be chosen so as to satisfy the inequality

$$|\partial_x f(x^*) - \partial_r f(x^*) \gamma| < 1, \tag{4.28}$$

which is readily derived by substituting the expression $\delta r_n = -\gamma \delta x_n$ in the linearized equation of motion.

The foregoing equations can now be combined in the following algorithm which describes in details the functioning of the OGY control scheme:

$$
\text{Initialization}\;\left[\begin{array}{ll}
\text{Choose} & \\
\quad x_0 \in [0,1] & ;\text{ initial condition} \\
\quad r \in [0,4] & ;\text{ unperturbed control parameter} \\
\quad \gamma \text{ according to Eq.(4.28)} & ;\text{ stable gain} \\
\quad D \ll 1 & ;\text{ control region} \\
\quad N & ;\text{ Number of control iterations} \\
\text{Solve} & \\
\quad x^* = f_r(x^*) \quad ;\text{ target point} &
\end{array}\right.
$$

$$
\text{Control loop}\;\left[\begin{array}{l}
\texttt{DO}\; n = 0, N \\
\qquad \texttt{IF}\; x_n \in D\; \texttt{THEN} \\
\qquad\qquad \text{Calculate } \delta r_n = -\gamma \delta x_n \quad ;\text{ controlled} \\
\qquad \texttt{ELSE} \\
\qquad\qquad \delta r_n = 0 \quad ;\text{ uncontrolled} \\
\qquad \texttt{END IF} \\
\qquad x_{n+1} = (r + \delta r_n) x_n (1 - x_n) \\
\texttt{END DO}
\end{array}\right.
$$

The plots in figure 4.6 illustrate the effect of this control algorithm on the logistic map for $r = 3.78$. The plots (a) and (b) represent, respectively, an uncontrolled and a controlled trajectory $\{x_n : 0 \leq n \leq 500\}$. For the controlled trajectory, the point $x^* = (r-1)/r \simeq 0.735$ was stabilized by applying the control law from the time instant $n = 150$ with the parameters $\gamma = 7.0$ and $D = [0.725, 0.745]$. Figure 4.6c also illustrates the performance of the feedback control when the variable x_n in the control law (4.24) is replace by the "coarse-grained" value

$$x_n^\Delta = \left\lceil \frac{x_n}{\Delta} \right\rceil \Delta, \tag{4.29}$$

obtained by using a uniform partition of the unit interval with cells of constant size Δ. In this case, and more generally for any values of Δ smaller than a certain threshold, the controller is able to localize the system within a constant interval ε enclosing x^*, provided that γ lies in the stable interval defined by Eq.(4.28). The latter part of this section will be devoted to the study of the properties of this interval using the closed-loop optimality theorem.

Chaotic control regions

For the purpose of chaotic control, all the accessible control actions determined by the values of δr_n, and correspondingly by the coordinates $x_n \in D$, can be constrained to be entropy-increasing by a proper choice of D. In other words, by choosing D conveniently, the Lyapunov exponents

$$\lambda(r) = \lim_{N\to\infty} \sum_{n=0}^{N-1} \ln |\partial_x f_r(x_n)|, \tag{4.30}$$

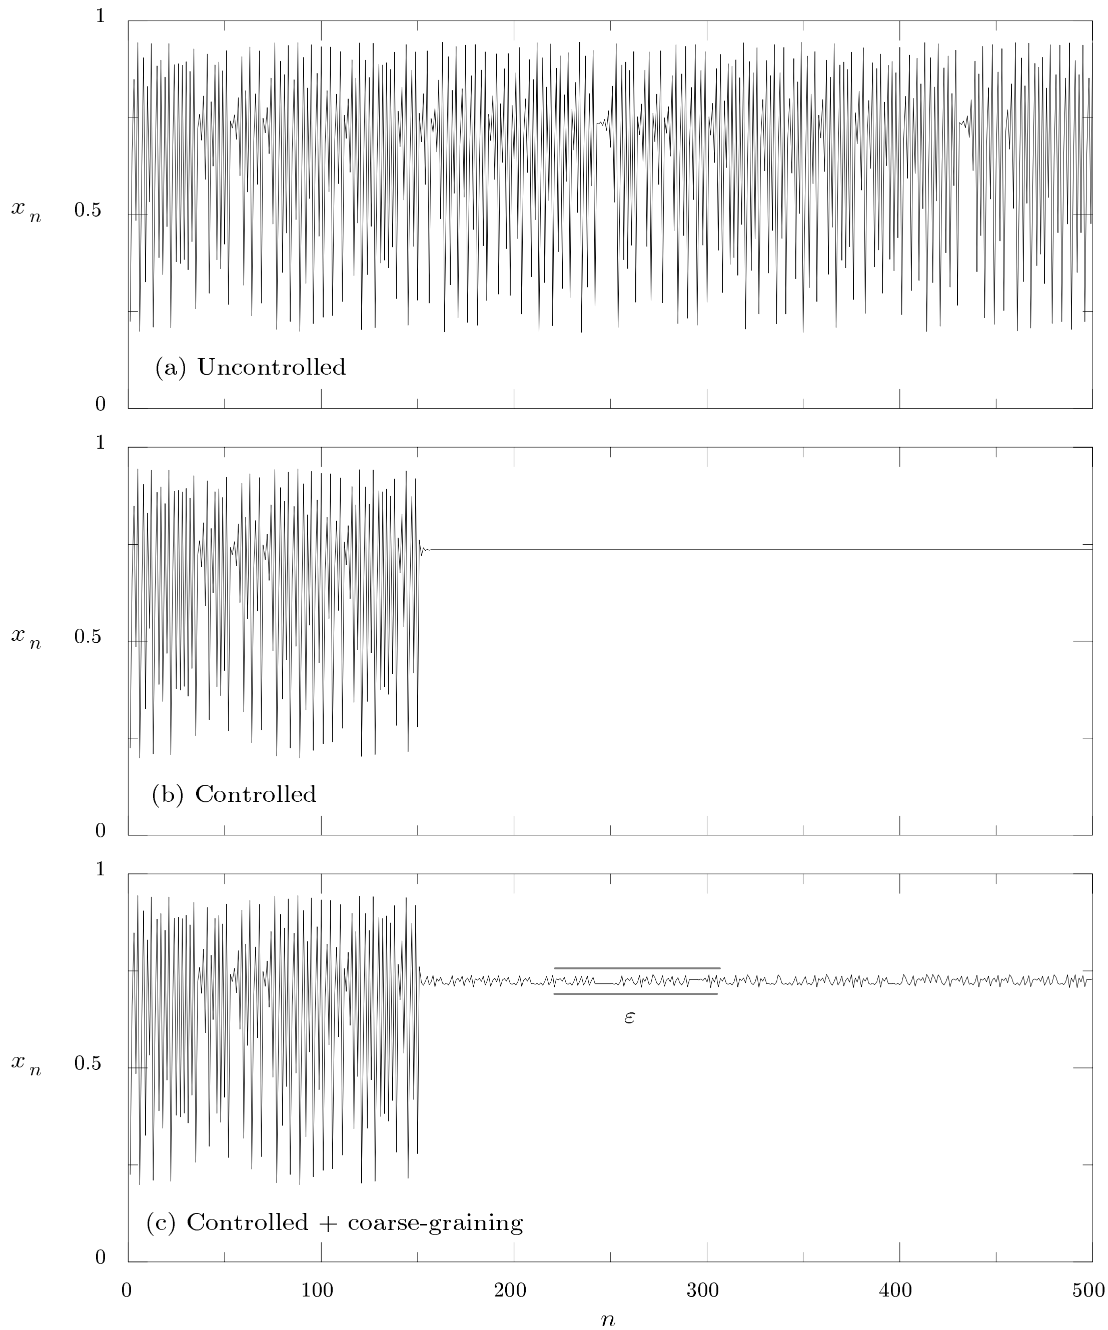

Figure 4.6: (a) Realization of the logistic map for $r = cte = 3.78$, $x^* \simeq 0.735$. (b) Application of the feedback control algorithm starting at $n = 150$ with $\gamma = 7.0$. (c) Application of the control law from $n = 150$ with $\gamma = 7.0$ in the presence of a uniform partition of size $\Delta = 0.02$ in the estimation of the state.

Regions	r_{min}	r_{max}	r_{mid}	Δr	λ_{mid} (base e)	x^*
R_0	3.565	3.625	3.595	0.03	0.1745	0.7218
R_1	3.63	3.735	3.6825	0.0525	0.3461	0.7284
R_2	3.74	3.825	3.7825	0.0425	0.4088	0.7356
R_3	3.858	4.0	3.929	0.071	0.5488	0.7455

Table 4.1: Parameters of the four control regions.

associated with any actuation indexed by r can be forced to be positive. This *a priori* is not obvious since, as figure 4.7 shows, the Lyapunov spectrum of the logistic map contains many negative Lyapunov exponents closely interspaced between positive values of $\lambda(r)$. In practice, however, the negative values appearing in the chaotic regions where $\lambda(r)$ seems mostly positive can be effectively suppressed either by adding a small noise component to the map or by coarse-graining the state space [20]. The effect of the addition of noise to the map is shown in figure 4.7. The spectrum of this figure, which has been obtained by computing numerically the sum (4.30) up to $N = 20000$ with an additive noise component of amplitude 10^{-3}, shows the disappearance of the negative exponents in four clearly defined regions denoted respectively by R_0, R_1, R_2 and R_3. The characteristics of these chaotic regions are listed in table 4.1. Among these, we have listed the boundaries $[r_{\text{min}}, r_{\text{max}}]$ of each region, in addition to the length Δr and the middle value, r_{mid}, of the intervals. The Lyapunov exponents (calculated in base e) and the fixed point associated with r_{mid} are also listed in the table.

Coarse-grained controller

Physically, the fact that $\lambda(r) > 0$ for all $r \in R_i$ $(i = 0, 1, 2, 3)$ implies that *almost* any initial uniform distribution for X covering an interval of size ε "stretches" on average by a factor $e^{\lambda(r)}$ after one iteration of the map with parameter r [58, 34]. This is true as long as the support of $p(x)$ is not too small or does not cover the entire unit interval, as it is usually the case during the control process. Now, for an open-loop controller, it can be seen that if $\lambda(r) > 0$, no control of the state is possible by using the OGY control algorithm. Indeed, without the knowledge of the position x_n at the time instant n, a controller merely acts as a perturbation during the transition $x_n \to x_{n+1}$, and the optimal control strategy then consists in using the smallest Lyapunov exponent available in order to achieve

$$\begin{aligned} \Delta H_{\text{open}} &= H(X_n) - H(X_{n+1})_{\text{open}} \\ &= \ln \varepsilon - \ln(\varepsilon e^{\lambda_{\text{min}}}) \\ &= -\lambda_{\text{min}} < 0, \end{aligned} \tag{4.31}$$

for a typical distribution (figure 4.8a). Following the closed-loop optimality theorem, it is thus necessary in a controlled situation where $\Delta H \geq 0$ (figure 4.8b) to have $I(X;C) \geq \lambda_{\text{min}}$ using a measurement channel with a capacity of at least λ_{min} nats. For example,

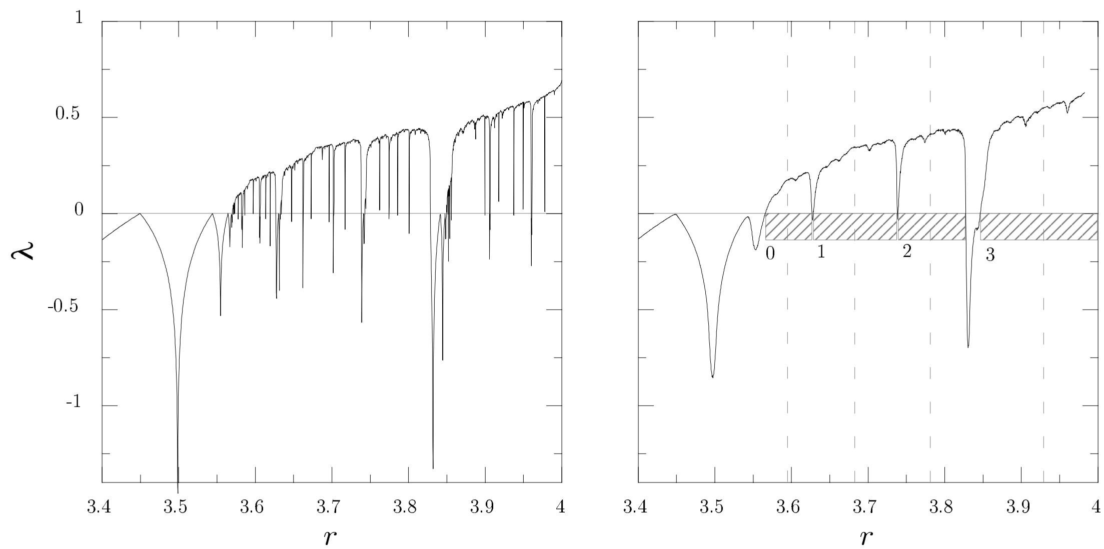

Figure 4.7: (Left) Lyapunov spectrum $\lambda(r)$ of the logistic map. The numerical calculations used approximately 20 000 iterates of the map. (Right) Spectrum for the noisy logistic map. The definition of the 4 chaotic control regions (0, 1, 2, 3) is illustrated by the dashed boxes. The vertical dashed lines give the position of r_{mid}.

if by using the OGY control scheme, we want to localize the trajectory generated by the logistic map uniformly within an interval of size ε, using a set of actuations from the chaotic control regions, then we need to measure the state x_n in Eq.(4.29) within an interval $\Delta \leq e^{-\lambda_{\min}}\varepsilon$. In a weaker sense, the fact that $\Delta H_{\text{open}} = -\lambda_{\min}$ also implies that an optimal controller in the regime where $\Delta H_{\text{closed}} = 0$ should be such that $I(X;C) = \lambda_{\min}$.

To understand this last observation, note that a perfect measurement of $I(X;C) = \ln a$ nats consists, for a uniform distribution with $|\text{supp}\, X| = \varepsilon$, in partitioning ε into a "measurement" subintervals of length $\varepsilon_m = \varepsilon/a$ (figure 4.9). By separation analysis, the controller under this partition then applies the same actuation $r^{(i)}$ for all the coordinates of the initial density lying in each of the subintervals i, thereby stretching each of them by a multiplicative factor $\exp[\lambda(r^{(i)})]$. In the optimal case, all the subintervals are controlled as if all the points of a given subintervals were directed toward x^* using the same optimal open-loop actuation associated with $\lambda_{\min}$ (figure 4.9), and the corresponding entropy change is thus given by

$$\begin{aligned}\Delta H_{\text{closed}} &= \ln\varepsilon - \ln\left(e^{\lambda_{\min}}\frac{\varepsilon}{a}\right)\\ &= -\lambda_{\min} + \log a. \end{aligned}\qquad (4.32)$$

This last equation is consistent with the closed-loop optimality theorem with $\Delta H_{\text{open}} = -\lambda_{\min}$, and gives the correct value for a, namely $a = e^{\lambda_{\min}}$, for $\Delta H_{\text{closed}} = 0$.

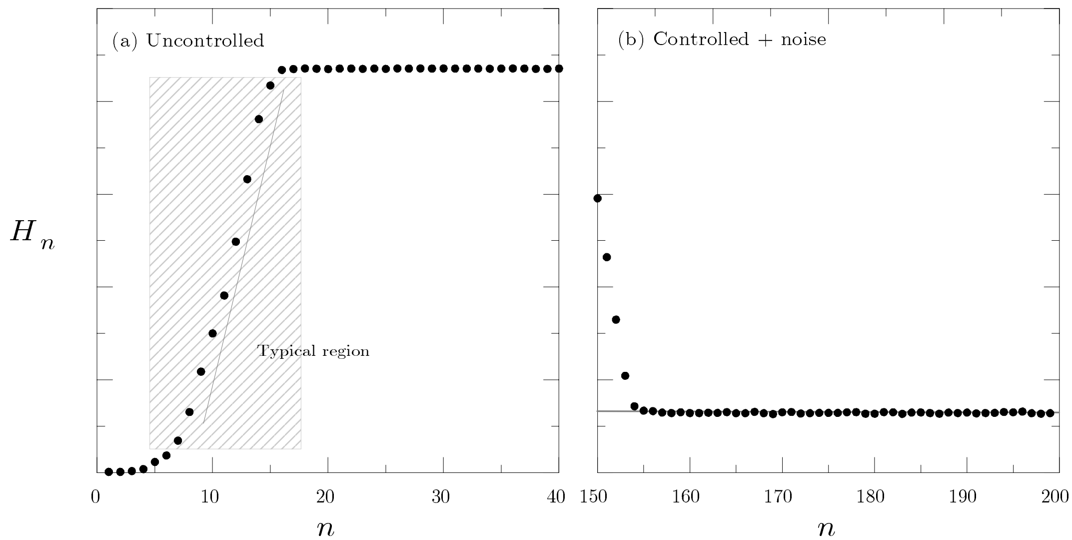

Figure 4.8: (a) Entropy $H_n = H(X_n)$ of the logistic map with $r = cte$ and $\lambda(r) > 0$. The entropy was calculated by propagating about 10 000 points initially located in a very narrow interval, and then by calculating a coarse-grained distribution. Inside the typical region where the distribution is away from the initial distribution (lower left part) and the uniform distribution (upper right part), $H(X_{n+1}) - H(X_n) \simeq \lambda(r)$. (b) Entropy for the logistic map controlled by a coarse-grained OGY controller starting at $n = 150$. For $n \gg 150$ (controlled regime), $H(X_n) \simeq \log \varepsilon = cte$.

Evidently, since the OGY algorithm does not control the points of different subintervals using the same Lyapunov exponents, the value $a = e^{\lambda_{\min}}$ only gives us a lower bound on the measurement interval ε_m at which one should measure X_n in order to achieve $\Delta H_{\text{closed}} = 0$. More realistically, $\lambda_{\min}$ could be replaced by $\lambda(\langle r \rangle)$, or $\lambda_{\text{mid}} = \lambda(r_{\text{mid}})$. (In practice, it was found during the simulations that $\langle r \rangle$ is well approximated by r_{mid}. See figure 4.10.) The fact that $e^{\lambda_{\text{mid}}}$ is a good approximation for a is illustrated in Figure 4.11. In this figure, the results of the numerical calculations of ε as a function of the effective measurement interval $\varepsilon_m = \varepsilon/a$ are plotted for the four control regions defined earlier. For each of the control region, the numerical simulations were carried out by constructing a coarse-grained density consisting of about 10^4 different points controlled by a gain γ fixed according to Eq.(4.28), and by choosing D in order to have $\delta r_n \leq \Delta r/2$ for all n. In doing so, we were able to restrain the possible values of $r + \delta r_n$ to lie in the control region R_i associated with the fixed points $x^*(r_{\text{mid}})$ of interest (see table 4.1). The solid line in each of the plots illustrates the average optimal "line"

$$\varepsilon = e^{\lambda_{\text{mid}}} \varepsilon_m \tag{4.33}$$

predicted by the optimality theorem with $\lambda_{\min}$ replaced by λ_{mid}. From the plots, it can be seen that, even by taking λ_{mid} as the actuation parameter, the data points for ε still

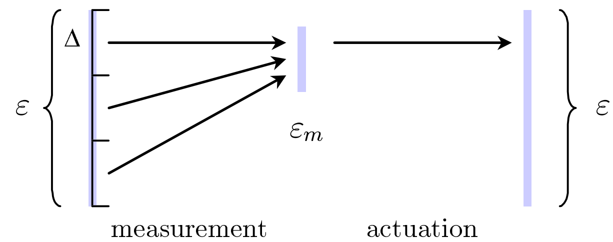

Figure 4.9: Separation analysis for the coarse-grained controller in the optimal case where $\Delta H = 0$. The size of the measurement interval is given by $e^{H(X|C)}$ where $H(X|C) = H(X) - I(X;C)$.

deviate from the relation predicted by Eq.(4.33). From what we know about optimality in closed-loop control, this can be partly understood by noticing that a controller based on the OGY scheme must become increasingly mixed (in a closed-loop sense) as $\Delta \to \varepsilon$, or equivalently as $a \to 1$. This is corroborated by the plot on the right hand side of figure 4.10 which shows the variance of $\langle r \rangle$. However, as figure 4.11 indicates, the deviation from the optimal line predicted by Eq.(4.33) does not increase monotonically with ε_m, a fact that cannot be explained by the optimality condition of closed-loop control. In the future, it will be interesting to better understand the conditions under which a closed-loop controller is optimal, and, in particular, to try to see if there exists a necessary and sufficient condition for closed-loop optimality.

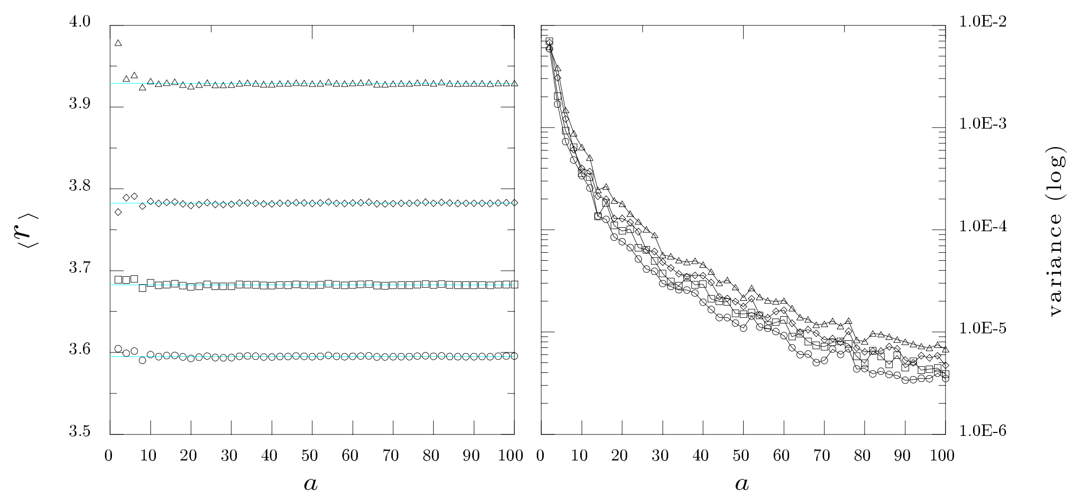

Figure 4.10: (a) Ensemble average of the parameter r used during the simulations of the OGY control algorithm. (b) Variance of $\langle r \rangle$ (logarithmic scale). The horizontal axis for both graphs represents the number of cells used in the partition, i.e., $a = \varepsilon\Delta^{-1}$.

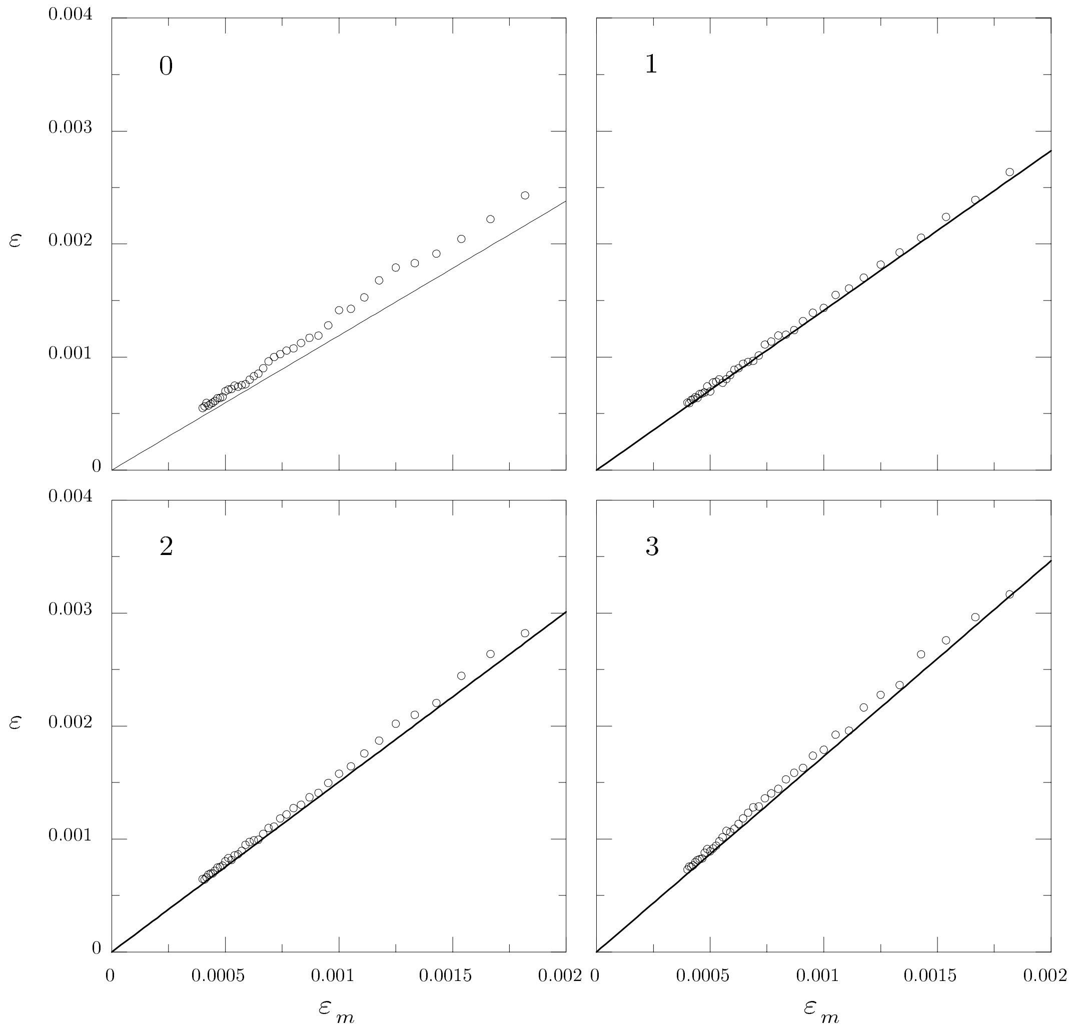

Figure 4.11: ($\circ$) Control interval ε as a function of the effective measured interval ε_m for the four control regions. The solid line shows the predicted relation between ε and ε_m based on the value of $\lambda(r_{\mathrm{mid}})$.

5 | Conclusion

5.1. Summary and informal extensions

This thesis presented a novel formalism for studying control processes from a informational viewpoint. Following a general discussion of control systems, presented in the introductory chapter, it was realized that controllers can be described by models analogous to probabilistic communication channels *à la* Shannon. Specifically, given the starting state X of a system, the task of a controller can be described as applying a 'transmission' channel, called the *actuation* channel, to X in order to redirect the system towards another desired state X'.[1] In the process, the controller may operate in two different regimes: (i) the controller may select the actuation channel independently of X, in which case the control is said to be *open-loop*; or (ii) the controller can select the actuation channel based on some information regarding the starting state of the controlled system. This latter case, corresponding to *closed-loop* control, was modeled by considering an extra channel to the model, the *measurement channel*, relating the state X to a control variable C labeling the actuation channel to be used.

Using this picture of control, which depicts both the estimation and the actuation processes as communication channels transmitting bits of information, several results have been proved. Among these, it was shown that the concept of controllability can receive an information-based definition which can be used to derive necessary and sufficient entropic conditions for perfect and approximate controllability (section 3.4). Furthermore, it was shown that the amount of information gathered by a controller through the measurement channel can be quantified formally in terms of a quantity $I(X;C)$ known as the *mutual information*. Then, in the context of our general control models, it was proved that $I(X;C)$ must bound an improvement function defined as

$$\Delta\mathcal{H} = H(X')_{\text{open}} - H(X')_{\text{closed}}, \tag{5.1}$$

$H(X')_{\text{open}}$ and $H(X')_{\text{closed}}$ being, respectively, the final entropy of the controlled system after an open-loop controller and a closed-loop controller have been applied. In other words, the amount of entropy $\Delta H_{\text{closed}} = H(X) - H(X')_{\text{closed}}$ that can be "extracted" from a dynamical system using a closed-loop control system much be such that

$$\Delta H_{\text{closed}} \leq \Delta H_{\text{open}} + I(X;C), \tag{5.2}$$

[1] After completing the thesis, it was found that the control models proposed in chapter 3 are very reminiscent of a communication model known, in the information theory literature, as the *matching channel*. See [1] for more details.

where ΔH_{open} is the entropy reduction attained by restricting the control model to an open-loop system. This constituted our central result which was extended to continuous-state and continuous-time systems, and then illustrated extensively, in chapter 4, by looking at the control of various dynamical systems, including chaotic maps.

The basic argument behind the proof of the above equations is that a closed-loop controller acts on the basis of the information provided by the conditional distribution $p(x|c)$ which reflects the knowledge of X acquired during its estimation. An open-loop controller, on the other hand, must proceed to a control action solely on the basis of the distribution $p(x)$. Then, by comparing the effect of both $p(x|c)$ and $p(x)$ on the closed-loop and open-loop actuation channels, it can be seen that a closed-loop controller essentially acts as if the values of the state X were partitioned into (usually non-overlapping) sets distributed according to different distributions $p(x|c)$ indexed by the values of the controller's state. In this case, the action of the closed-loop controller can *virtually* be described as if the initial state of the system was distributed according to $p(x|c)$, and, consequently, as if the initial entropy of the system was $H(X|C)$, as opposed to the entropy $H(X)$ 'seen' by open-loop controllers. Hence the mentioned improvement in ΔH_{closed}.

As such, this argument is very similar to the one given by Kelly [31, 19] in its information-theoretic treatment of the problem of gambling on a horse race. The similarity between control and gambling is in fact not fortuitous: in both cases, the task is to apply strategies (control or gambling strategies) chosen from some *a priori* set of admissible strategies [43]. Then, as it is the case in control, we can envision to decide of a gambling strategy either blindly (open-loop strategy) or based on some knowledge or side information relevant to the gambling situation (closed-loop strategy). In this context, the result derived by Kelly is the following: for a gambling problem based on doubling rate optimization, there is a quantitative improvement of the 'gambling returns' proportional to $I(X;C)$ in the eventuality where side information is used. (See [19] for a precise statement of the problem.) In control, the only difference with gambling is that the cost function in the problem is chosen to be the entropy, and more precisely the entropy difference ΔH. As another parallel between control and game theory, it should be noted that the terms 'pure' and 'mixed' control strategies were defined exactly as in game theory [22]. An interesting result in that case, already found in game theory and re-derived independently in this thesis, is the optimality theorem of pure open-loop gambling/control strategies for concave 'cost' or 'reward' functions.[2] This result in game theory can be found in [29].

5.2. Future work

There surely exist other parallels that one can draw from our analysis of control systems, as well as other applications for which our formalism seems applicable and useful. A promising avenue of research, in that sense, is the study of more realistic, hence more

[2] Cost functions are also referred to as value functions, Lyapunov functions, or Lagrangians, depending on the field of study.

complicated, control systems (e.g., high-dimensional noisy linear systems). As another example, let us mention potential applications in the field of nanoscale control systems which are often described by stochastic models. By way of conclusion, we present in the following several topics of interest which have not been fully investigated in this thesis, but will be in the future.

- *Linear systems.* We mentioned in chapter 3 that our results could be extended to continuous-state and continuous-time dynamical systems. However, as pointed out, in the continuous-time limit a difficulty concerning the definition of the mutual information arises: $I(X(t);C(t))$ is not a quantity relating two end points in time. Thus, the limit $\Delta t \to 0$ cannot be taken in order to yield a meaningful rate equation, except for the case where one considers conditional entropies instead of marginal ones.

  For linear systems, an alternative solution to the continuous-time problem can be put forward: it involves the study of continuous-time linear differential equations which can always be expressed, in the frequency domain, as scalar equations of the form
  $$Y(w) = G(w)X(w), \tag{5.3}$$
  where $X(w)$ and $Y(w)$ represent respectively the input and the output signals, as functions of the frequency w, of a linear filter whose transfer function is $G(w)$. Using this representation, the signals $X(w)$ and $Y(w)$ can be sampled at different frequencies $w_1, w_2, \ldots, w_N$ to construct two random vectors
  $$\begin{aligned} Y(w) &\to \boldsymbol{Y} = [Y(w_1), Y(w_2), \ldots, Y(w_N)]^T & (5.4)\\ X(w) &\to \boldsymbol{X} = [X(w_1), X(w_2), \ldots, X(w_N)]^T, & (5.5)\end{aligned}$$
  related by the matrix equation $\boldsymbol{Y} = G\boldsymbol{X}$. In this last equation, G now represents the transfer matrix obtained by sampling the function $G(w)$. The next step then is to apply our formalism to these vectors as if they were representing a multivariate map. In particular, the entropy of the output signal can be calculated using [24, 28]
  $$\begin{aligned} H(Y(t)) \underset{N\to\infty}{\longrightarrow} H(\boldsymbol{Y}) &= -\int_{\boldsymbol{y}} p(\boldsymbol{y}) \log p(\boldsymbol{y})\ d\boldsymbol{y} \\ &= H(\boldsymbol{X}) + \log |G|, \end{aligned} \tag{5.6}$$
  where $|G|$ is the determinant of G, and
  $$I(X(t);Y(t)) \underset{N\to\infty}{\longrightarrow} I(\boldsymbol{X};\boldsymbol{Y}) = \int_{\boldsymbol{x}} \int_{\boldsymbol{y}} p(\boldsymbol{x},\boldsymbol{y}) \log \frac{p(\boldsymbol{x},\boldsymbol{y})}{p(\boldsymbol{x})p(\boldsymbol{y})}\ d\boldsymbol{y} d\boldsymbol{x}. \tag{5.7}$$
  Evidently, for definiteness, it should be noted that the signals $X(t)$ and $Y(t)$ must be bounded in time, say over an interval T, in which case $X(w)$ and $Y(w)$

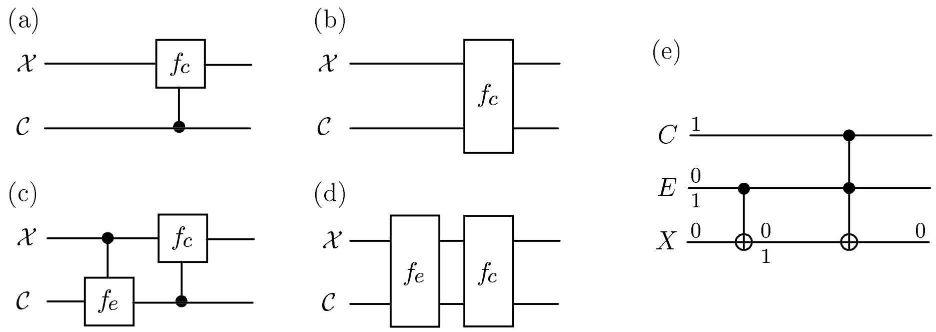

Figure 5.1: (a)-(b) General circuit models of open-loop controllers $\mathcal{C}$ acting on a system $\mathcal{X}$ using transition functions denoted by f_c. (c)-(d) Circuit models for closed-loop control. The transition functions for the estimation process are denoted by f_e. (e) Example of an *environment-assisted* open-loop control: the open-loop controller with $H(C) = 0$ is able to recover the bit of error introduced by the environment using a Toffoli gate which uses the bit of mutual information between E and X.

are also bounded in the frequency domain. This being said, we are justified, by Shannon's sampling theorem [24, 19], to restrict our study of the signals to a finite set of frequencies. This frequency representation of dynamical systems should be particularly useful for modeling linear systems disturbed by 'generic' noise processes such as Gaussian noise or white noise processes.

- *Generalizing the model: circuit models.* The models studied in this thesis implicitly assume that the controller's state is not changed during the actuation process. For future studies, it would be desirable to consider a more general setup which includes any possible dynamical effects that $\mathcal{X}$ or $\mathcal{E}$ can induce on $\mathcal{C}$. One convenient way of doing this is to model the dynamics of the systems involved in the control process as a logical circuit made of wires, representing the states of the systems, and gates which model the evolution of the states. Figure 5.1 displays a few examples of such circuit models. The circuit of figure 5.1a, for instance, can be thought of as the equivalent of the open-loop DAG of chapter 3, generalized in figure 5.1b to include possible effects of $\mathcal{X}$ on $\mathcal{C}$. Closed-loop control, on the other hand, is represented in figure 5.1c with its generalization in 5.1d.

  Such a generalization of the control models is very important if one wants to consider non-conventional types of controllers such as error-correcting codes or the controller shown in figure 5.1e. The circuit of this figure models, in a very abstract way, an open-loop control process similar to the spin-echo effect which, basically, undo the evolution of a combined system $\mathcal{X}+\mathcal{E}$ by using the correlations between $\mathcal{X}$ and $\mathcal{E}$. (See [41] for references on the spin-echo effect, and for an insightful information-theoretic discussion of the phenomenon.) In the circuit, specifically, the bit of noise introduced through the CNOT gate can be "erased" by reversing it exactly using a Toffoli gate such that $x = x \oplus 1 \bmod 2$ when $c = e = 1$. As such, $\mathcal{C}$ must be considered as an open-loop controller since $C = 1$ with probability one, in spite of the fact that the control process, as a whole, uses one bit of information!

- *Quantum controllers.* The circuits models presented in the previous paragraph is of particular interest for the generalization of control theory to quantum systems capable of existing in superpositions of quantum states [39, 45]. As a first guess, an information-theoretic formalism based on circuits seems to be directly applicable in the context of quantum mechanics, though a few questions remain open at this moment. One of them is related to the fact that the quantum mutual information of a bipartite quantum system $\mathcal{XY}$ can be larger than the quantum (von Neumann) entropy of $\mathcal{X}$ or $\mathcal{Y}$ [15]. Evidently, this is forbidden for classical systems, and the intriguing question then is: can this extra information be used to dissipate the entropy of a system above the limit imposed by the closed-loop optimality theorem? [66] Other related subjects might also be of interest: see, e.g., [16] for an interesting information-theoretic perspective on classical and quantum error correction, and [42] for an introduction to quantum gambling.

Bibliography

[1] R. Ahlswede and N. Cai. Information and control: Matching channels. *IEEE Trans. Inf. Th.*, 44:542–563, 1998.

[2] V. M. Alekseev and M. V. Yakobson. Symbolic dynamics and hyperbolic dynamic systems. *Phys. Rep.*, 75:287–325, 1.

[3] W. Ashby. Measuring the internal information exchange in a system. *Cybernetica*, VIII:5–22, 1965.

[4] A. Bagchi. *Optimal Control of Stochastic Systems*. Prentice Hall, New York, 1993.

[5] C. Beck and F. Schlögl. *Thermodynamics of Chaotic Systems*. Cambridge University Press, Cambridge, 1993.

[6] C. H. Bennett. Logical reversibility of computation. *IBM J. Res. Dev.*, pages 525–532, 1973.

[7] C. H. Bennett. The thermodynamics of computation: A review. *Int. J. Theor. Phys.*, 21:905–940, 1982.

[8] C. H. Bennett. Demons, engines and the second law. *Sci. Am.*, November 1987.

[9] C. H. Bennett. Notes on the history of reversible computation. *IBM J. Res. Dev.*, 32:16–23, 1988.

[10] C. H. Bennett and R. Landauer. The fundamental physical limits of computation. *Sci. Am.*, pages 48–56, July 1988.

[11] K. L. Boettcher. *An Information Theoretic Model of the Decision Maker*. M. Sc., Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 1981.

[12] L. Brillouin. *Science and Information Theory*. Academic Press, New York, 1956.

[13] L. Brillouin. *Scientific Uncertainty, and Information*. Academic Press, New York, 1964.

[14] C. M. Caves. Entropy and information: How much information is needed to assign a probability? In W. H. Zurek, editor, *Complexity, Entropy and the Physics of Information*, pages 91–115, Reading, 1991. Addison-Wesley.

[15] N. J. Cerf and C. Adami. Negative entropy and information in quantum mechanics. *Phys. Rev. Lett.*, 79:5194, 1997.

[16] N. J. Cerf and R. Cleve. Information-theoretic interpretation of quantum error-correcting codes. *Phys. Rev. A*, 56:1721–1732, 1997.

[17] E. Charniak. Bayesian networks without tears. *AI Magazine*, pages 50–63, Winter 1991.

[18] R. C. Conant. Laws of information which govern systems. *IEEE Trans. Syst. Man and Cyb.*, SMC-6:240–255, 1976.

[19] T. M. Cover and J. A. Thomas. *Elements of Information Theory.* John Wiley, New York, 1991.

[20] J. Crutchfield and J. Farmer. Fluctuations and simple chaotic dynamics. *Phys. Rep.*, 92:54–82, 1982.

[21] J.-P. Eckmann and D. Ruelle. Ergodic theory and strange attractors. *Rev. Mod. Phys.*, 57:617–656, 1986.

[22] R. A. Epstein. *The Theory of Gambling and Statistical Logic.* Academic Press, New York, 1977. QA271.E67 1977.

[23] J. D. Farmer. Information dimension and the probabilistic structure of chaos. *Z. Naturforsch*, 37a:1304–1325, 1982.

[24] R. G. Gallager. *Information Theory and Reliable Communication.* John Wiley, New York, 1968.

[25] R. G. Gallager. *Discrete Stochastic Processes.* Kluwer Academic, Boston, 1995.

[26] M. Gell-Mann and S. Lloyd. Information measures, effective complexity, and total information. *Complexity*, 2:44–52, 1996.

[27] R. Hartley. Transmission of information. *Bell Sys. Tech. J.*, page 535, July 1948.

[28] S. Ihara. *Information Theory for Continuous Systems.* World Scientific, Singapore, 1993.

[29] W. Jianhua. *The Theory of Games.* Clarendon Press, Oxford, 1988.

[30] M. I. Jordan, editor. *Learning in Graphical Models.* MIT Press, Cambridge, MA, 1999.

[31] J. Kelly. A new interpretation of information rate. *Bell Sys. Tech. Journal*, 35:917–926, 1956.

[32] R. Landauer. Information is physical. *Physics Today*, pages 23–29, May 1991.

[33] A. Lasota and M. Mackey. *Chaos, Fractals, and Noise.* Springer-Verlag, Berlin, 1994.

[34] V. Latora and M. Baranger. Kolmogorov-Sinai entropy rate versus physical entropy. *Phys. Rev. Lett.*, 82:520–523, 1999.

[35] H. Leff and A. Rex, editors. *Maxwell's Demon, Entropy, Information, Computing.* Princeton University Press, New Jersey, 1990.

[36] P. U. Lima and G. N. Saridis. Intelligent controller as hierarchical stochastic automata. *IEEE Trans. Syst. Man and Cyb.*, 29:151–163, 1999.

[37] S. Lloyd. Use of mutual information to decrease entropy: implications for the second law of thermodynamics. *Phys. Rev. A*, 39:5378–5386, 1989.

[38] S. Lloyd. Causality and information flow. In H. Atmanspacher and H. Scheingraber, editors, *Information Dynamics*, pages 131–142, New York, 1991. Plenum Press.

[39] S. Lloyd. Quantum controllers for quantum systems. LANL archive `quant-ph/9703042`.

[40] S. Lloyd and J.-J. Slotine. Information theoretic tools for stable adaptation and learning. *Int. J. Adapt. Cont. Sig. Proc.*, 10:499–530, 1996.

[41] S. Lloyd and W. Zurek. Algorithmic treatment of the spin-echo effect. *J. Stat. Phys.*, 62:819–839, 1991.

[42] D. A. Meyer. Quantum strategies. *Phys. Rev. Lett.*, 82:1052, 1999.

[43] D. Middleton. *An Introduction to Statistical Communication Theory.* IEEE Press, New York, 1996.

[44] R. E. Neapolitan. *Probabilistic Reasoning in Expert Systems: Theory and Algorithms.* Wiley, New York, 1990.

[45] R. J. Nelson. *Geometric Control of Quantum Mechanical and Nonlinear Classical Systems.* Ph.D., Mechanical Engineering, Massachusetts Institute of Technology, 1999.

[46] G. Nicolis. *Introduction to Nonlinear Science.* Cambridge University Press, Cambridge, 1995.

[47] K. Oexle. Data compressibility, physical entropy, and evolutionary a priori relation between observer and object. *Phys. Rev. E*, 51:2651–2653, 1995.

[48] E. Ott, C. Grebogi, and J. Yorke. Controlling chaos. *Phys. Rev. Lett.*, 64:1196, 1990.

[49] J. Pearl. *Probabilistic Reasoning in Intelligent Systems.* Morgan Kaufmann, San Mateo, 1988.

[50] H.-O. Peitgen, H. Jürgens, and D. Saupe. *Chaos and Fractals: New Frontiers of Science.* Springer-Verlag, Berlin, 1992.

[51] R. Poplavskiĭ. Thermodynamics models of information processes. *Sov. Phys. Usp.*, 18:222–241, 1975.

[52] R. Poplavskiĭ. Information and entropy. *Sov. Phys. Usp.*, 22:371–380, 1980.

[53] F. Reif. *Statistical and Thermal Physics.* McGraw-Hill, New York, 1965.

[54] E. M. Rogers. *Physics for The Inquiring Mind.* Princeton University Press, Princeton, 1960.

[55] E. Scano. Modèle de mécanisme informationnel. *Cybernetica*, VIII:188–223, 1965.

[56] R. Schack and C. M. Caves. Chaos for liouville probability densities. *Phys. Rev. E*, 53:3387–3401, 1996.

[57] B. W. Schumacher. Demonic heat engines. In J. J. Halliwell, J. Pérez-Mercader, and W. H. Zurek, editors, *Physical Origins of Time Asymmetry*, Cambridge, 1994. Cambridge University Press.

[58] H. G. Schuster. *Deterministic Chaos, 2nd ed.* VCH, Berlin, 1989.

[59] C. E. Shannon. A mathematical theory of communication. *Bell Sys. Tech. J.*, 27:379–423, 623–656, 1948.

[60] C. E. Shannon and W. Weaver. *The Mathematical Theory of Communication.* University of Illinois Press, Urbana, 1963.

[61] R. Shaw. Strange attractors, chaotic behavior, and information flow. *Z. Naturforsch*, 36a:80–112, 1980.

[62] T. Shinbrot. Progress in the control of chaos. *Adv. Physics*, 44:73–111, 1995.

[63] T. Shinbrot, C. Grebogi, E. Ott, and J. Yorke. Using small perturbations to control chaos. *Nature*, 363:441, 1993.

[64] J.-J. E. Slotine and W. Li. *Applied Nonlinear Control.* Prentice-Hall, New York, 1991.

[65] R. F. Stengel. *Optimal Control and Estimation.* Dover, New York, 1986.

[66] H. Touchette. Controlling information: cooling cbits and qbits. Term paper for the course 2.111, Quantum Computation, 1999. (Unpublished).

[67] H. Touchette. Lyapunov exponents, stability and information. Term paper for the course 2.152 Applied Nonlinear Control, Spring 1998. (Unpublished).

[68] H. Touchette and S. Lloyd. Information-theoretic limits of control. LANL archive `chao-dyn/9905039`. To be published in *Phys. Rev. Lett.* (2000).

[69] H. C. von Baeyer. *Maxwell's Demon: Why Warmth Disperses and Time Passes*. Random House, New York, 1998.

[70] A. M. Weinberg. On the relation between information and energy systems: A family of Maxwell's demons. *Interdisciplinary Science Reviews*, 7:47–52, 1982.

[71] N. Wiener. *Cybernetics: or Control and Communication in the Animal and the Machine*. MIT Press, Cambridge, MA, 1948.

[72] W. H. Zurek. Algorithmic randomness and physical entropy. *Phys. Rev. A*, 40:4731–4751, 1989.

[73] W. H. Zurek. Thermodynamic cost of computation, algorithmic complexity and the information metric. *Nature*, 341:119–124, 1989.

[74] W. H. Zurek. Algorithmic information content, Church-Turing thesis, physical entropy, and Maxwell's demon. In W. H. Zurek, editor, *Complexity, Entropy and the Physics of Information*, pages 73–89, Reading, 1991. Addison-Wesley.