Understanding the Informational Structure of the Past-Future Bottleneck of a Controlled Linear System

Hippolyte Charvin, Nicola Catenacci Volpi and Daniel Polani

School of Physics, Engineering and Computer Science
University of Hertfordshire, UK

Who controls the past controls the future.
Who controls the present controls the past.

George Orwell

It is instructive to consider a variety of systems, especially complex ones, under the lens of information theory. This consideration permits one to observe structures and effects which otherwise escape consideration or modeling. More than that, however, it permits one to place "cost" metric on the dynamical behaviour that the systems exhibit. Furthermore, it permits one to possibly tune through a cascade of resolutions (defined by the informational capacity) to understand the structure of a system not in the "naive" metrics/topology of the dynamical system's original description (which may have many equivalent alternatives), but through the more abstract, yet more discerning characterization of information theory.

As the linear control channel is one of the central building blocks of any cybernetic system under consideration, one can hope to learn many important properties by considering the informational signature of such systems. In particular, the linear case can be understood as a first order approximation of non-linear systems near stable equilibria. Moreover, focusing on this simplified framework allows one to call on existing analytical tools for the information-theoretic study of stochastic dynamical systems.

Indeed, crucial work has been undertaken in this respect by the past-future bottleneck framework [Creutzig et al., 2009, Amir et al., 2015]. Considering a past-future information bottleneck means that one constrains the information that any controller/processor may infer about the future, given a particular (controlled) past. When one limits this information, this acts as a "squint" which puts a constraint on the complexity of the controlled dynamics that is actually being modelled. As that information is constrained, parts of the system dynamics are suppressed at defined levels of "squinting". These are the information-dominant components of the system.

Previous work made clear that this informational schematization of linear dynamical systems reduces to manipulating the singular value decomposition of the system's Hankel matrix. More precisely, the simplified system is built by manipulating the singular values, while, at the same time, leaving the singular spaces unchanged — the system's schematization thus becoming eventually a mere change in the relative importance of the singular spaces.

We propose the point of view that these singular spaces carry, each of them, a very specific dynamical structure: that each of them turns out to implement a linear dynamical system of its own, independent from the ones of other singular spaces. The global dynamical system is thus understood as modularly built up from independent sub-dynamical components, and the informational schematization as a manipulation of their relative prominence in the system's description, so as to make salient only the sub-dynamical components which carry the most predictive power.

As the informational constraints are increased, this "squinting" results, as mentioned above, in the total discarding of some of the low-predictive dynamical components. In particular, the control is *also* skeletonized into a "garbled" control, which increasingly only constrains the most relevant parts of the system.

These studies open up a whole domain of new possibilities for the analysis and interpretation of controlled linear systems and their decomposition into separate domains of dynamics. This has consequences for how such systems now can be recast.

However, the analysis of these dynamical components turns out to be subtle. For one, one needs to ensure that they indeed describe an actual controlled system (which is not per se obvious, as the past-future bottleneck does not impose a condition that the reduced system be actually a realizable system). We discuss under which conditions this is the case.

Furthermore, we describe several links between the information bottleneck reduction of the linear control system and the invariances and equivariances of the original system. Note that while some of these may be in the state space, some of these may be across times; in principle, some may be also mixing space and time transformations.

These links imply consequences on how information is and can be exchanged between the different subspaces of the system under consideration when certain structural properties hold. They also imply to which degree the control signal is able to affect particular degrees of freedom at a given resolution of the systems' permitted capacity; we will discuss these different types of phenomena, their links with the structural properties of the systems and give examples (some of them counterintuitive).

References

[Amir et al., 2015] Amir, N., Tiomkin, S., and Tishby, N. (2015). Past-future information bottleneck for linear feedback systems. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 5737–5742.

[Creutzig et al., 2009] Creutzig, F., Globerson, A., and Tishby, N. (2009). Past-future information bottleneck in dynamical systems. *Phys. Rev. E*, 79:041925.