



Intrinsic separation principles[☆]

Boris Houska

ShanghaiTech University, China



ARTICLE INFO

Article history:

Received 25 November 2022

Received in revised form 24 September 2024

Accepted 19 September 2025

Keywords:

Dual control

Information theory

Convex optimization

ABSTRACT

This paper is about output-feedback control problems for linear systems in the presence of given state-, control-, disturbance-, and measurement error constraints. Because the traditional separation theorem in stochastic control is inapplicable to such constrained systems, a novel information-theoretic framework is proposed. It leads to an intrinsic separation principle that can be used to break the dual control problem for constrained linear systems into a meta-learning problem that minimizes an intrinsic information measure and a robust control problem that minimizes an extrinsic risk measure. The theoretical results in this paper can be applied in combination with modern polytopic computing methods in order to approximate a large class of dual control problems by finite-dimensional convex optimization problems.

© 2025 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

1. Introduction

The separation principle in stochastic control is a fundamental result in control theory (Joseph & Tou, 1961; Lindquist, 1973; Wonham, 1968), closely related to the certainty-equivalence principle (van Warter & Willems, 1981). It states that – under certain assumptions – the problem of optimal control and state estimation can be decoupled.

For general control systems, however, the separation theorem fails to hold. Thus, if one is interested in finding optimal output-feedback control laws for such systems, one needs to solve a rather complicated dual control problem (Feldbaum, 1961). There are two cases where such dual- or output-feedback control problems are of interest:

- (i) The first case is that we have an uncertain nonlinear system – in the easiest case, without state- and control constraints – for which the information content of future measurements depends on the control actions. In practice, this dependency can often be neglected, because, at least for small measurement errors and process noise, and under certain regularity assumptions, the separation theorem holds in a first order approximation (Stengel, 1994). Nevertheless, there are some nonlinear systems that can only be stabilized if this dependency is taken into account (Filatov & Unbehauen, 2004; Sehr & Bitmead, 2019).
- (ii) And, the second case is that we have an uncertain linear system with state- and control constraints. Here, the

process noise and future measurement errors have to be taken into account if one wants to operate the system safely, for instance, by designing output-feedback laws that ensure constraint satisfaction for all possible uncertainty scenarios.

The current paper is about the second case. This focus is motivated by the recent trend towards the development of safe learning and control methods (Hewing, Wabersich, Menner, & Zeilinger, 2020; Zanon & Gros, 2021).

1.1. Literature review

Dual control problems have been introduced by Feldbaum in the early 1960s (Feldbaum, 1961). Mature game-theoretic and stochastic methods for analyzing such dual- and output feedback control problems have, however, only been developed much later. They go back to the seminal work of N.N. Krasovskii (Krasovskii, 1964; Krasovskii & Krasovskii, 1995) and A.B. Kurzhanskii, (Kurzhanskii, 1972, 2004). Note that these historical articles are complemented by modern set-theoretic control theory (Blanchini & Miani, 2003, 2015). Specifically, in the context of constrained linear systems, set-theoretic notions of invariance under output feedback can be found in the work of Dórea (Almeida & Dórea, 2021; Dórea, 2009), which focuses on the invariance of a single information set, and in the work of Artstein and Raković (Artstein & Raković, 2011), which focuses on the invariance of a collection of information sets. Moreover, a variety of set-theoretic output-feedback control methods for constrained linear systems have appeared in Bemporad and Garulli (2000), Brunner, Müller, and Allgöwer (2018), dos Reis de Souza, Efimov, Raïssi, and Ping (2022). These have in common that they propose to take bounds

[☆] The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Subhrakanti Dey under the direction of Editor Florian Dorfler.

E-mail address: borish@shanghaitech.edu.cn.

on measurement errors into account upon designing a robust predictive controller. In this context, the work of Goulart and Kerrigan must be highlighted (Goulart & Kerrigan, 2007), who found a remarkably elegant way to optimize uncertainty-affine output feedback control laws for constrained linear systems. A general overview of methods for output-feedback and dual model predictive control (MPC) can be found in Findeisen, Imsland, Allgöwer, and Foss (2003), Hovd and Bitmead (2005), Mayne, Rakovic, Findeisen, and Allgöwer (2009), Sehr and Bitmead (2019), and the reference therein.

1.2. Contribution

The three main contributions of this paper can be outlined as follows.

Meta Information Theory. While traditional information theories are based on the assumption that one can learn from accessible data, models for predicting the evolution of an uncertain control system require a higher level of abstraction. Here, one needs a prediction structure that is capable of representing the set of all possible future information states of a dynamic learning process without having access to future measurement data. Note that a comprehensive and thorough discussion of this aspect can be found in the above mentioned article by Artstein and Raković (Artstein & Raković, 2011), in which notions of invariance under output-feedback for collections of information sets are introduced. Similar to their construction, the current article proposes a meta information theoretic framework that is based on a class of information set collections, too. A novel idea of the current article in this regard, however, is the introduction of intrinsic equivalence relations that can be used to categorize information sets with respect to their geometric properties. This leads to an algebraic-geometric definition of meta information spaces in which one can distinguish between extrinsic and intrinsic information measures. Here, intrinsic information about a system is needed to predict what we will know about its states, while extrinsic information is needed to predict and assess the risk that is associated to control decisions.

Intrinsic Separation Principle. The central contribution of this paper is the introduction of the *intrinsic separation principle*. It formalizes the fact that the intrinsic information content of a constrained linear system does not depend on the choice of the control law. An important consequence of this result is that a large class of dual receding horizon control problems can be solved by separating them into a meta learning problem that predicts intrinsic information and a robust control problem that minimizes extrinsic risk measures. Moreover, the intrinsic separation principle can be used to analyze the existence of solutions to dual control problems under certain assumptions on the continuity and monotonicity of the objective function of the dual control problem.

Polytopic Dual Control. The theoretical results in this paper are used to develop practical methods to approximately solve dual control problems for linear systems with convex state- and control constraints as well as polytopic process noise and polytopic measurement error bounds. In order to appreciate the novelty of this approach, it needs to be recalled first that many existing robust output-feedback control methods, for instance the state-of-the-art output-feedback model predictive control methods in Findeisen et al. (2003), Mayne et al. (2009), are based on a set-theoretic or stochastic analysis of a coupled system-observer dynamics, where the control law depends on a state estimate. This is in contrast to the presented information theoretic approach to dual control, where control decisions are made based on the system's true information state rather than a state estimate. In fact, for the first time, this paper presents a polytopic

dual control method that neither computes vector-valued state estimates nor introduces an affine observer structure. Instead, the discretization of the control law is based on optimizing a finite number of control inputs that are associated to so-called extreme polytopes. The shapes, sizes, and orientations of these extreme polytopes encode the system's intrinsic information while their convex hull encodes the system's extrinsic information. The result of this discretization is a finite dimensional convex optimization problem that approximates the original dual control problem.

1.3. Overview

The paper is structured as follows.

- Section 2 reviews the main idea of set-theoretic learning and introduces related notation.
- Section 3 establishes the technical foundation of this article. This includes the introduction of meta information spaces and a discussion of the difference between intrinsic and extrinsic information measures.
- Section 4 introduces the intrinsic separation principle for constrained linear systems, see Theorem 1.
- Section 5 discusses how to resolve dual control problems by intrinsic separation, see Theorem 2.
- Section 6 presents methods for discretizing dual control problems using polytopic information set approximations. The main technical result is summarized in Theorem 3. A numerical case study is presented. And,
- Section 7 summarizes the highlights of this paper.

1.4. Notation

Throughout this paper, $\mathbb{K}^n$ denotes the set of closed subsets of $\mathbb{R}^n$, while $\mathbb{K}_c^n$ denotes the set of compact subsets of $\mathbb{R}^n$. It is equipped with the Hausdorff distance

$$d_H(X, Y) \stackrel{\text{def}}{=} \max \left\{ \max_{x \in X} \min_{y \in Y} \|x - y\|, \max_{y \in Y} \min_{x \in X} \|x - y\| \right\}$$

for all $X, Y \in \mathbb{K}_c^n$, where $\|\cdot\| : \mathbb{R}^n \rightarrow \mathbb{R}$ denotes a norm on $\mathbb{R}^n$, such that $(\mathbb{K}_c^n, d_H)$ is a metric space. This definition can be extended to $\mathbb{K}^n$ as follows: if the maxima in the above definition do not exist for $X, Y \in \mathbb{K}^n$, we set $d_H(X, Y) = \infty$. The pair $(\mathbb{K}^n, d_H)$ is called an extended metric space. Finally, the notation $\text{cl}(\cdot)$ is used to denote the closure, assuming that it is clear from the context what the underlying metric distance function is. For instance, if $\mathfrak{X} \subseteq \mathbb{K}^n$ denotes a set of closed sets, $\text{cl}(\mathfrak{X})$ denotes the closure of $\mathfrak{X}$ in $(\mathbb{K}^n, d_H)$.

2. Information spaces

An information space $(\mathcal{I}, d, \sqcap)$ is a space in which learning can take place. This means that $(\mathcal{I}, d)$ is an extended metric space that is equipped with a learning operator

$$\sqcap : \mathcal{I} \times \mathcal{I} \rightarrow \mathcal{I},$$

such that $(\mathcal{I}, \sqcap)$ is a semi-group. Among the most important examples for such spaces is the so-called set-theoretic information space, which is introduced below.

2.1. Set-theoretic learning

In the context of set-theoretic learning (Bertsekas & Rhodes, 1971; Blanchini & Miani, 2003; Witsenhausen, 1968), $\mathcal{I} = \mathbb{K}^n$ denotes the set of closed subsets of the vector space $\mathbb{R}^n$, while $d =$

d_H denotes the (extended) Hausdorff distance. Here, the standard intersection operator takes the role of a learning operator,

$$\sqcap = \cap,$$

recalling that the intersection of closed sets is closed. The motivation behind this definition can be outlined as follows: let us assume that we currently know that a vector $x \in \mathbb{R}^n$ is contained in a given set $X \in \mathbb{K}^n$. If we receive additional information, for instance, that the vector x is also contained in the set $Y \in \mathbb{K}^n$, our posterior information is that x is contained in the intersection of the sets X and Y , which is denoted by $X \cap Y$.

The above set-theoretic framework is compatible with affine functions of the form $f(x) = Ax + b$. Namely, if $X \in \mathbb{K}^n$ is a closed set, the image set,

$$f(X) \stackrel{\text{def}}{=} AX + b \stackrel{\text{def}}{=} \{Ax + b \mid x \in X\},$$

is closed, too. Here, the matrix A and the vector b are assumed to have compatible dimensions. Finally, the Minkowski sum

$$X + Y \stackrel{\text{def}}{=} \{x + y \mid x \in X, y \in Y\},$$

is defined for all sets $X, Y \in \mathbb{K}^n$.

2.2. Expectation and deviation

Expectation and deviation functions are among the most basic tools for analyzing learning processes (Doob, 1953). The expectation function is defined by

$$\forall X \in \mathbb{K}_c^n, \quad E(X) \stackrel{\text{def}}{=} \int_X x \, dx.$$

It is a continuous function on $\mathbb{K}_c^n$ that satisfies

$$E(AX + b) = AE(X) + b.$$

Similarly, a deviation function $D : \mathbb{K}_c^n \rightarrow \mathbb{R}$ is a continuous and radially unbounded function that satisfies

- (i) $D(X) \geq 0$,
- (ii) $D(X) = 0$ if and only if $X = \{E(X)\}$,
- (iii) $D(X) = D(X - E(X))$, and
- (iv) $D(X \cap Y) \leq D(X)$,

for all $X, Y \in \mathbb{K}_c^n$. While statistical learning models often use the variance of a random variable as a deviation measure, a more natural choice for D in the context of set theoretic learning is given by the diameter,

$$D(X) = \text{diam}(X) \stackrel{\text{def}}{=} \max_{x, y \in X} \|x - y\|.$$

A long and creative list of other possible choices for D can be found in Rockafellar and Uryasev (2013).

3. Meta learning

The definition of information spaces assumes data accessibility during learning operations. If one wishes to predict the future evolution of a learning process, however, one faces the problem that the future data is inaccessible. Therefore, this section proposes to introduce a meta information space in which one can represent the set of all possible posterior information states of a learning process without having access to its future data.

3.1. Information ensembles

The focus of this and the following sections is on the set-theoretic framework recalling that $\mathbb{K}_c^n$ denotes the set of compact subsets of $\mathbb{R}^n$. A set $\mathfrak{X} \subseteq \mathbb{K}_c^n$ is called an information ensemble of $\mathbb{K}_c^n$ if

$$\forall Y \in \mathbb{K}_c^n, \quad X \cap Y \in \mathfrak{X} \quad (1)$$

for all $X \in \mathfrak{X}$. Because $\emptyset = X \cap \emptyset \in \mathfrak{X}$, any information ensemble contains the empty set.

Proposition 1. *If $\mathfrak{X} \subseteq \mathbb{K}_c^n$ is an information ensemble, then $\text{cl}(\mathfrak{X})$ is an information ensemble, too.*

Proof. Let $\mathfrak{X}$ be a given information ensemble and let $X_\infty \in \text{cl}(\mathfrak{X})$ be a given set in its closure. Then there exists a Cauchy sequence $X_1, X_2, \dots \in \mathfrak{X}$ such that

$$X_\infty \stackrel{\text{def}}{=} \lim_{k \rightarrow \infty} X_k \in \text{cl}(\mathfrak{X}).$$

Next, let $Y \in \mathbb{K}_c^n$ be an arbitrary compact set. The case $X_\infty \cap Y = \emptyset$ is trivial, since $\emptyset \in \mathfrak{X} \subseteq \text{cl}(\mathfrak{X})$. Next, if $X_\infty \cap Y \neq \emptyset$, then there exists for every $\xi \in X_\infty \cap Y$ an associated sequence $z_1(\xi) \in X_1, z_2(\xi) \in X_2, \dots$ with

$$\lim_{k \rightarrow \infty} z_k(\xi) = \xi.$$

This construction is such that the sets

$$Z_k \stackrel{\text{def}}{=} \text{cl}(\{z_k(\xi) \mid \xi \in X_\infty \cap Y\})$$

satisfy $Z_k = X_k \cap Z_k \in \mathfrak{X}$, since $\mathfrak{X}$ is an information ensemble. Consequently, it follows that

$$X_\infty \cap Y = \lim_{k \rightarrow \infty} Z_k \in \text{cl}(\mathfrak{X}).$$

Thus, $\text{cl}(\mathfrak{X})$ is an information ensemble, as claimed by the statement of the proposition. $\square$

Proposition 2. *Let $\mathfrak{X}$ be an information ensemble of $\mathbb{K}_c^n$. Then $(\mathfrak{X}, d_H, \cap)$ is an information space.*

Proof. Condition (1) implies that $X \cap Y \in \mathfrak{X}$ for all $X, Y \in \mathfrak{X}$. Thus, $(\mathfrak{X}, \cap)$ is a subsemigroup of $(\mathbb{K}_c^n, \cap)$. Moreover, d_H defines a metric on $\mathfrak{X}$. Consequently, $(\mathfrak{X}, d_H, \cap)$ is an information space. $\square$

Remark 1. Information ensembles differ from general set collections, as considered in Artstein and Raković (2011), in that Eq. (1) is imposed. This property is crucial for developing a coherent learning algebra: without it, $(\mathfrak{X}, \cap)$ would not generally be a subsemigroup of $(\mathbb{K}_c^n, \cap)$.

3.2. Extreme sets

A set $X \in \text{cl}(\mathfrak{X})$ of a given information ensemble $\mathfrak{X} \subseteq \mathbb{K}_c^n$ is called an extreme set of $\mathfrak{X}$ if

$$\forall Y \in \text{cl}(\mathfrak{X}) \setminus \{X\}, \quad X \cap Y \neq X.$$

The set of extreme sets of $\mathfrak{X}$ is denoted by $\partial\mathfrak{X}$. It is called the boundary of the information ensemble $\mathfrak{X}$. Clearly, we have $\partial\mathfrak{X} \subseteq \text{cl}(\mathfrak{X})$, but, in general, $\partial\mathfrak{X}$ is not an information ensemble. Instead, $\partial\mathfrak{X}$ can be interpreted as a minimal representation of the closure of $\mathfrak{X}$, because

$$\text{cl}(\mathfrak{X}) = \{Y \in \mathbb{K}_c^n \mid \exists X \in \partial\mathfrak{X}, Y \subseteq X\}.$$

Reversely, the closure of $\mathfrak{X}$ can be interpreted as the smallest information ensemble that contains $\partial\mathfrak{X}$.

3.3. Meta information spaces

Let $\mathbb{I}^n$ denote the set of closed information ensembles of $\mathbb{K}_c^n$. Similarly, let $\mathbb{I}_c^n$ denote the set of compact information ensembles of the information space $(\mathbb{K}_c^n, d_H, \cap)$. The meta learning operator $\cap : \mathbb{I}^n \times \mathbb{I}^n \rightarrow \mathbb{I}^n$ is defined by

$$\mathfrak{X} \cap \mathfrak{Y} \stackrel{\text{def}}{=} \{X \cap Y \mid X \in \mathfrak{X}, Y \in \mathfrak{Y}\}$$

for all $\mathfrak{X}, \mathfrak{Y} \in \mathbb{I}^n$. Let us define the distance function

$$\Delta_H(\mathfrak{X}, \mathfrak{Y}) \stackrel{\text{def}}{=} \max \left\{ \max_{X \in \mathfrak{X}} \min_{Y \in \mathfrak{Y}} d_H(X, Y), \max_{Y \in \mathfrak{Y}} \min_{X \in \mathfrak{X}} d_H(X, Y) \right\}$$

for all $\mathfrak{X}, \mathfrak{Y} \in \mathbb{I}_c^n$ such that $(\mathbb{I}_c^n, \Delta_H)$ is a metric space. The definition of Δ_H extends to $\mathbb{I}^n$ by understanding the above definition in the extended value sense. The proposition below shows that $(\mathbb{I}^n, \Delta_H, \cap)$ is an information space, called the meta information space of $(\mathbb{K}_c^n, d_H, \cap)$.

Proposition 3. *The triple $(\mathbb{I}^n, \Delta_H, \cap)$ is an information space. It can itself be interpreted as a set-theoretic information space in the sense that we have*

$$\forall \mathfrak{X}, \mathfrak{Y} \in \mathbb{I}^n, \quad \mathfrak{X} \cap \mathfrak{Y} = \mathfrak{X} \cap \mathfrak{Y}. \quad (2)$$

Proof. The proof of this proposition is divided into two parts: the first part shows that (2) holds and the second part concludes that $(\mathbb{I}^n, \Delta_H, \cap)$ is an information space.

Part I. Let $\mathfrak{X}, \mathfrak{Y} \in \mathbb{I}^n$ be given information ensembles. For any $X \in \mathfrak{X} \cap \mathfrak{Y}$ the relation $X \cap X = X \in \mathfrak{X} \cap \mathfrak{Y}$ holds. But this implies that

$$\mathfrak{X} \cap \mathfrak{Y} \subseteq \{X \cap Y \mid X \in \mathfrak{X}, Y \in \mathfrak{Y}\} = \mathfrak{X} \cap \mathfrak{Y}.$$

In order to also establish the reverse inclusion, assume that $Z \in \mathfrak{X} \cap \mathfrak{Y}$ is a given set. It can be written in the form $Z = X \cap Y$ with $X \in \mathfrak{X}$ and $Y \in \mathfrak{Y}$. We have $Z \subseteq X$ and $Z \subseteq Y$. Moreover, we have $Z \in \mathbb{K}_c^n$, since the intersection of compact sets is compact. Thus, since $\mathfrak{X}$ and $\mathfrak{Y}$ are information ensembles, (1) implies that $Z \in \mathfrak{X}$ and $Z \in \mathfrak{Y}$. But this is the same as saying that $Z \in \mathfrak{X} \cap \mathfrak{Y}$, which implies $\mathfrak{X} \cap \mathfrak{Y} \supseteq \mathfrak{X} \cap \mathfrak{Y}$. Together with the above reverse inclusion, this yields (2).

Part II. Note that $(\mathbb{I}^n, \cap)$ is a semigroup, which follows from the definition of intersection operations. Moreover, $(\mathbb{I}^n, \Delta_H)$ is, by construction, an extended metric space. Thus, $(\mathbb{I}^n, \Delta_H, \cap)$ is indeed an information space, as claimed by the statement of this proposition. $\square$

Corollary 1. *The triple $(\mathbb{I}_c^n, \Delta_H, \cap)$ is also an information space. It can be interpreted as a sub-meta information space of $(\mathbb{I}^n, \Delta_H, \cap)$.*

Proof. The statement follows from the previous proposition, since the intersection of compact sets is compact; that is, $(\mathbb{I}_c^n, \cap)$ is a subsemigroup of $(\mathbb{I}^n, \cap)$. $\square$

3.4. Interpretation of meta learning processes

Meta information spaces can be used to analyze the evolution of learning processes independently of data access. To illustrate this, we consider a set-theoretic sensor, which returns at each time instance a compact information set $X \in \mathbb{K}_c^1$ containing the scalar state x of a physical system, $x \in X$. If the absolute value of the measurement error of the sensor is bounded by 1, this means that $X \subseteq [a, a+2]$ for at least one lower bound $a \in \mathbb{R}$. The closed but unbounded information ensemble that is associated with such a sensor is given by

$$\mathfrak{X} = \{X \in \mathbb{K}_c^1 \mid \exists a \in \mathbb{R} : X \subseteq [a, a+2]\} \in \mathbb{I}^1. \quad (3)$$

It can be interpreted as the set of all information sets that the sensor could return when taking a measurement.

To model an associated meta learning process, we assume that $x \in [-3, 3]$; that is, our prior is given by

$$\mathfrak{X} = \{X \in \mathbb{K}_c^1 \mid X \subseteq [-3, 3]\}.$$

In such a situation, a meta learning process is – due to [Proposition 3](#) – described by an update of the form $\mathfrak{X}^+ = \mathfrak{X} \cap \mathfrak{Y} = \mathfrak{X} \cap \mathfrak{Y}$, where $\mathfrak{X}^+$ denotes the posterior,

$$\mathfrak{X}^+ = \left\{ X \in \mathbb{K}_c^1 \mid \begin{array}{l} \exists a \in \mathbb{R} : \\ X \subseteq [\max\{a, -3\}, 2 + \min\{1, a\}] \end{array} \right\}.$$

It is computed without having access to any sensor data.

3.5. Intrinsic equivalence

Equivalence relations can be used to categorize compact information sets with respect to their geometric properties. Here, we consider two sets $X, Y \in \mathbb{K}_c^n$ equivalent, writing $X \simeq Y$, if they have the same shape, size, and orientation. This means that

$$X \simeq Y \iff \exists a \in \mathbb{R}^n : X + a = Y.$$

The motivation for introducing this particular equivalence relation is that two information sets X and Y can be considered equally informative if they coincide after a translation.

Definition 1. Two information ensembles $\mathfrak{X}, \mathfrak{Y} \subseteq \mathbb{K}_c^n$ are called intrinsically equivalent, $\mathfrak{X} \sim \mathfrak{Y}$, if their quotient spaces coincide,

$$(\mathfrak{X} / \simeq) = (\mathfrak{Y} / \simeq). \quad (4)$$

The intrinsic equivalence relation $\sim$ from the above definition is – as the name suggests – an equivalence relation. This follows from the fact that $\mathfrak{X} \sim \mathfrak{Y}$ if and only if

$$\forall X \in \mathfrak{X}, \quad \exists a \in \mathbb{R}^n : X + a \in \mathfrak{Y}$$

$$\text{and } \forall Y \in \mathfrak{Y}, \quad \exists b \in \mathbb{R}^n : Y + b \in \mathfrak{X},$$

which, in turn, follows from (4).

Proposition 4. *If $\mathfrak{X}, \mathfrak{Y} \subseteq \mathbb{K}_c^n$ are intrinsically equivalent information ensembles, $\mathfrak{X} \sim \mathfrak{Y}$, their closures are intrinsically equivalent, too, $\text{cl}(\mathfrak{X}) \sim \text{cl}(\mathfrak{Y})$.*

Proof. [Proposition 1](#) ensures that the closures of $\mathfrak{X}$ and $\mathfrak{Y}$ are information ensembles, $\text{cl}(\mathfrak{X}) \in \mathbb{I}^n$ and $\text{cl}(\mathfrak{Y}) \in \mathbb{I}^n$. Next, there exists for every $X_\infty \in \text{cl}(\mathfrak{X})$ a convergent sequence of sets $X_1, X_2, \dots \in \mathfrak{X}$ such that $X_\infty = \lim_{k \rightarrow \infty} X_k$. Moreover, since $\mathfrak{X} \sim \mathfrak{Y}$, there also exists a sequence $a_1, a_2, \dots \in \mathbb{R}^n$ such that the sequence $Y_k \stackrel{\text{def}}{=} X_k + a_k \in \mathfrak{Y}$ remains in $\mathfrak{Y}$. Because $\mathfrak{X}$ and $\mathfrak{Y}$ are compact, the sequence of offsets a_k must be bounded. Thus, it has a convergent subsequence, $a_{j_1}, a_{j_2}, \dots \in \mathbb{R}^n$, with limit $a_\infty \stackrel{\text{def}}{=} \lim_{k \rightarrow \infty} a_{j_k} \in \mathbb{R}^n$. This construction is such that $X_\infty + a_\infty = \lim_{k \rightarrow \infty} \{X_{j_k} + a_{j_k}\} \in \text{cl}(\mathfrak{Y})$. A completely analogous statement holds after replacing the roles of $\mathfrak{X}$ and $\mathfrak{Y}$. Consequently, the closures of $\mathfrak{X}$ and $\mathfrak{Y}$ are intrinsically equivalent, which corresponds to the statement of the proposition. $\square$

3.6. Extrinsic versus intrinsic information

Throughout this paper, it will be important to distinguish between extrinsic and intrinsic information. Here, the extrinsic information of an information ensemble is encoded by the union of its elements, namely, the extrinsic information set. It describes present information. The extrinsic information content of an information ensemble can be quantified by extrinsic information measures:

Definition 2. An information measure $f : \mathbb{I}_c^n \rightarrow \mathbb{R}$ is called extrinsic, if there exist a function $g : \mathbb{K}_c^n \rightarrow \mathbb{R}$ with

$$\forall \mathfrak{X} \in \mathbb{I}_c^n, \quad f(\mathfrak{X}) = g\left(\bigcup_{X \in \mathfrak{X}} X\right).$$

In contrast to extrinsic information, the intrinsic information of an information ensemble $\mathfrak{X}$ is encoded by its quotient space, $\mathfrak{X}/\simeq$. It describes future information. We introduce the shorthand $\mathbb{Q}_c^n \stackrel{\text{def}}{=} \mathbb{I}_c^n/\simeq$.

Definition 3. An information measure $f : \mathbb{I}_c^n \rightarrow \mathbb{R}$ is called intrinsic, if there exist a function $g : \mathbb{Q}_c^n \rightarrow \mathbb{R}$ with

$$\forall \mathfrak{X} \in \mathbb{I}_c^n, \quad f(\mathfrak{X}) = g(\mathfrak{X}/\simeq).$$

Note that the definitions of E and D in Section 2.2 can be lifted to the meta information space setting by introducing their associated extrinsic expectation $\mathfrak{E}$ and extrinsic deviation $\mathfrak{D}$. They are given by

$$\mathfrak{E}(\mathfrak{X}) \stackrel{\text{def}}{=} E\left(\bigcup_{X \in \mathfrak{X}} X\right) \quad \text{and} \quad \mathfrak{D}(\mathfrak{X}) \stackrel{\text{def}}{=} D\left(\bigcup_{X \in \mathfrak{X}} X\right)$$

for all $\mathfrak{X} \in \mathbb{I}_c^n$. Note that $\mathfrak{E}$ and $\mathfrak{D}$ are continuous functions, which inherit the properties of E and D , namely, $\mathfrak{E}(A\mathfrak{X} + b) = A\mathfrak{E}(\mathfrak{X}) + b$ holds. Similarly, $\mathfrak{D}$ satisfies all axioms of a deviation measure in the sense that

- (i) $\mathfrak{D}(\mathfrak{X}) \geq 0$,
- (ii) $\mathfrak{D}(\mathfrak{X}) = 0$ if and only if $\mathfrak{X} = \{\{\mathfrak{E}(\mathfrak{X})\}\}$,
- (iii) $\mathfrak{D}(\mathfrak{X}) = \mathfrak{D}(\mathfrak{X} - \mathfrak{E}(\mathfrak{X}))$, and
- (iv) $\mathfrak{D}(\mathfrak{X} \sqcap \mathfrak{Y}) \leq \mathfrak{D}(\mathfrak{X})$,

for all $\mathfrak{X}, \mathfrak{Y} \in \mathbb{I}_c^n$. Extrinsic deviation measures need to be distinguished from intrinsic deviation measures. Here, a function $\mathfrak{D}^\circ : \mathbb{I}_c^n \rightarrow \mathbb{R}$, is called an intrinsic deviation measure if it is a continuous and intrinsic function with

- (i) $\mathfrak{D}^\circ(\mathfrak{X}) \geq 0$,
- (ii) $\mathfrak{D}^\circ(\mathfrak{X}) = 0$ if and only if $\mathfrak{X} \sim \{\{\mathfrak{E}(\mathfrak{X})\}\}$,
- (iii) $\mathfrak{D}^\circ(\mathfrak{X}) = \mathfrak{D}^\circ(\mathfrak{X} - \mathfrak{E}(\mathfrak{X}))$, and
- (iv) $\mathfrak{D}^\circ(\mathfrak{X} \sqcap \mathfrak{Y}) \leq \mathfrak{D}^\circ(\mathfrak{X})$,

for all $\mathfrak{X}, \mathfrak{Y} \in \mathbb{I}_c^n$. The second axiom is equivalent to requiring that $\mathfrak{D}^\circ$ is positive definite on the quotient space $\mathbb{Q}_c^n$. A practical example is given by

$$\forall \mathfrak{X} \in \mathbb{I}_c^n, \quad \mathfrak{D}_\infty^\circ(\mathfrak{X}) = \max_{X \in \mathfrak{X}} \max_{x, y \in X} \|x - y\|, \quad (5)$$

which turns out to be an intrinsic information measure.

Example 1. Let us revisit the tutorial example from Section 3.4, where we had considered the case that

$$\mathfrak{X} = \{X \in \mathbb{K}_c^1 \mid X \subseteq [-3, 3]\} \quad \text{and} \quad \mathfrak{X}^+ = \left\{X \in \mathbb{K}_c^1 \mid \begin{array}{l} \exists a \in \mathbb{R} : \\ X \subseteq [\max\{a, -3\}, 2 + \min\{1, a\}] \end{array} \right\}$$

denote the prior and posterior of a data-free meta learning process. If we set $D(X) = \text{diam}(X)$ and define $\mathfrak{D}$ and $\mathfrak{D}_\infty^\circ$ as above, then $\mathfrak{D}(\mathfrak{X}) = \mathfrak{D}(\mathfrak{X}^+) = 6$: since our meta learning process is not based on actual data, the extrinsic information content of the prior $\mathfrak{X}$ and the posterior $\mathfrak{X}^+$ must be the same. This is in contrast to the intrinsic deviation, $\mathfrak{D}_\infty^\circ(\mathfrak{X}) = 6 > 2 = \mathfrak{D}_\infty^\circ(\mathfrak{X}^+)$, which predicts that independent of the next measurement, the diameter of our posterior information set will be at most 2.

4. Intrinsic separation principle

The goal of this section is to formulate an intrinsic separation principle for constrained linear systems.

4.1. Constrained linear systems

We consider uncertain linear discrete-time systems,

$$\begin{aligned} x_{k+1} &= Ax_k + Bu_k + w_k \\ \eta_k &= Cx_k + v_k. \end{aligned} \quad (6)$$

Here, $x_k \in \mathbb{R}^n$ denotes the state, $u_k \in \mathbb{U}$ the control, $w_k \in \mathbb{W}$ the disturbance, $\eta_k \in \mathbb{R}^{n_v}$ the measurement, and $v_k \in \mathbb{V}$ the measurement error at time $k \in \mathbb{Z}$. The matrices A, B , and C and the state, control, disturbance, and measurement error constraints sets, $\mathbb{X} \in \mathbb{K}^n$, $\mathbb{U} \in \mathbb{K}_c^{n_u}$, $\mathbb{W} \in \mathbb{K}_c^{n_w}$, and $\mathbb{V} \in \mathbb{K}_c^{n_v}$, are given.

Example 2. Let us consider a system with matrices

$$A = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix}, \quad B = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}, \quad C = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}^T.$$

In this example, only the third state can be measured and there is no measurement error, $\mathbb{V} = \{0\}$. The seventh state can be controlled, while the first five states are uncontrollable. Because we have $CA = 0$, the controllable states are not observable in the sense of classical linear systems theory. Nevertheless, dual control is able to take into account information from measurements in order to make control decisions. For instance, if $\mathbb{W}$ is a polytope in $\mathbb{R}^7$, given by $\mathbb{W} = \text{conv}(w_1^+, w_1^-, w_2^+, w_2^-)$, $w_1^\pm = (\pm \frac{1}{2} \ 0 \ -\frac{1}{2} \ 0 \ 0 \ 0 \ 0)^T$, $w_2^\pm = (0 \ \pm \frac{1}{2} \ \frac{1}{2} \ 0 \ 0 \ 0 \ 0)^T$, measurements of the third state yield valuable information about the controllable state at future time instances. We will come back to this example further below, where this statement will be analyzed in more detail.

4.2. Information tubes

The sensor that measures the outputs Cx_k of (6) can be represented by the information ensemble

$$\mathfrak{V} \stackrel{\text{def}}{=} \{X \in \mathbb{K}_c^n \mid \exists \eta \in \mathbb{R}^{n_v} : \eta - CX \subseteq \mathbb{V}\}. \quad (7)$$

Since $\mathbb{V}$ is compact, $\mathfrak{V}$ is closed but potentially unbounded, $\mathfrak{V} \in \mathbb{I}^n$. If $\mathfrak{X} \in \mathbb{I}^n$ denotes a prior information ensemble of the state of (6) an associated posterior is given by $\mathfrak{X} \sqcap \mathfrak{V}$. Such posteriors are computed at every time step in order to allow for a progressive refinement of the system's information state whenever new measurements arrive. This motivates to introduce the function

$$F(\mathfrak{X}, \mu) \stackrel{\text{def}}{=} \left\{X^+ \in \mathbb{K}_c^n \mid \begin{array}{l} \exists X \in \mathfrak{X} \sqcap \mathfrak{V} : \\ X^+ \subseteq AX + B\mu(X) + \mathbb{W} \end{array} \right\},$$

which is defined for all $\mathfrak{X} \in \mathbb{I}^n$ and all $\mu : \mathbb{I}^n \rightarrow \mathbb{U}$. Let $\mathcal{U}$ denote the set of maps from $\mathbb{I}^n$ to $\mathbb{U}$. It is equipped with the supremum norm, $\|\mu\| \stackrel{\text{def}}{=} \sup_{X \in \mathbb{K}^n} \|\mu(X)\|$, such that $(\mathcal{U}, \|\cdot\|)$ is a complete metric space. As $\mu \in \mathcal{U}$ is potentially discontinuous, $F(\mathfrak{X}, \mu)$ is not necessarily closed. Instead, the following statement holds.

Proposition 5. If $\mathfrak{X}$, $\mathbb{U}$, $\mathbb{V}$, and $\mathbb{W}$ are closed, then the closure of the set $F(\mathfrak{X}, \mu)$ is for every $\mu \in \mathcal{U}$ a closed information ensemble, $\bar{F}(\mathfrak{X}, \mu) \stackrel{\text{def}}{=} \text{cl}(F(\mathfrak{X}, \mu)) \in \mathbb{I}^n$.

Proof. The statement of this proposition follows from Proposition 1 and the above definition of F . $\square$

Definition 4. An information ensemble $\mathfrak{X}_s \in \mathbb{I}^n$ is called control invariant for (6) if there exists a $\mu_s \in \mathcal{U}$ such that $\mathfrak{X}_s \supseteq F(\mathfrak{X}_s, \mu_s)$.

Definition 5. A sequence $\mathfrak{X}_0, \mathfrak{X}_1, \dots \in \mathbb{I}^n$ of information ensembles is called an information tube for (6) if there exists a sequence $\mu_0, \mu_1, \dots \in \mathcal{U}$ such that $\mathfrak{X}_{k+1} \supseteq F(\mathfrak{X}_k, \mu_k)$ for all $k \in \mathbb{N}$.

Definition 6. Let $\mu_0, \mu_1, \dots \in \mathcal{U}$ be any given control policy sequence. Its associated tight information tube $\mathfrak{X}_0, \mathfrak{X}_1, \dots \in \mathbb{I}^n$ is defined by the recursion

$$\forall k \in \mathbb{N}, \quad \mathfrak{X}_{k+1} = \bar{F}(\mathfrak{X}_k, \mu_k).$$

Example 3. Let us revisit Example 2 and choose the trivial control law sequence $\mu_k(X) = 0$. If the initial state is known to be 0, our prior is given by $\mathfrak{X}_0 = \{0\}$. In this example, the associated tight information tube can be worked out explicitly. It is given by

$$\begin{aligned} \mathfrak{X}_0 &= \{0\}, \quad \mathfrak{X}_1 = \{X \in \mathbb{K}_c^7 \mid X \subseteq \mathbb{W}\} \\ \mathfrak{X}_2 &= \left\{ X \in \mathbb{K}_c^7 \mid X \subseteq \left[\begin{pmatrix} 0 \\ 0 \\ 0 \\ -\frac{1+2\eta_1}{4} \\ -\frac{1-2\eta_1}{4} \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 0 \\ \frac{1+2\eta_1}{4} \\ \frac{1-2\eta_1}{4} \\ 0 \\ 0 \end{pmatrix} \right] + \mathbb{W} \right\} \\ \mathfrak{X}_3 &= \left\{ X \in \mathbb{K}_c^7 \mid X \subseteq \left[\begin{pmatrix} 0 \\ 0 \\ 0 \\ -\frac{1+2\eta_2}{4} \\ -\frac{1-2\eta_2}{4} \\ -\frac{1+2\eta_1}{4} \\ -\frac{1-2\eta_1}{4} \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 0 \\ \frac{1+2\eta_2}{4} \\ \frac{1-2\eta_2}{4} \\ \frac{1+2\eta_1}{4} \\ \frac{1-2\eta_1}{4} \end{pmatrix} \right] + \mathbb{W} \right\} \\ \mathfrak{X}_k &= \mathfrak{X}_3 \quad \text{for } k > 3. \end{aligned}$$

In particular, $\mathfrak{X}_3$ is a control invariant information ensemble. In order to understand the above expressions for $\mathfrak{X}_2$ and $\mathfrak{X}_3$, it is helpful to study the visualization in Fig. 1, which explains why the measurements η_1 and η_2 of the third state at times $k = 1$ and $k = 2$ affect our knowledge about the last four states at time $k = 3$. Since the third column of A comprises solely zeros, the measurements of the third state do not undergo direct propagation through the nominal system dynamics. Consequently, the parameters η_1 and η_2 solely influence the widths of the extreme sets of $\mathfrak{X}_2$ and $\mathfrak{X}_3$.

4.3. Intrinsic separation

The following theorem establishes the fact that the intrinsic equivalence class of tight information tubes does not depend on the control policy sequence.

Theorem 1. Let $\mu_0, \mu_1, \dots \in \mathcal{U}$ and $v_0, v_1, \dots \in \mathcal{U}$ be arbitrary feasible control policy sequences and let $\mathfrak{X}_0, \mathfrak{X}_1, \dots \in \mathbb{I}_c^n$ and $\mathfrak{Y}_0, \mathfrak{Y}_1, \dots \in \mathbb{I}_c^n$ be associated tight information tubes, such that

$$\mathfrak{X}_{k+1} = \bar{F}(\mathfrak{X}_k, \mu_k) \quad \text{and} \quad \mathfrak{Y}_{k+1} = \bar{F}(\mathfrak{Y}_k, v_k) \quad (8)$$

for all $k \in \mathbb{N}$. If the initial information ensembles are intrinsically equivalent, $\mathfrak{X}_0 \sim \mathfrak{Y}_0$, then all information ensembles are intrinsically equivalent; that is, $\mathfrak{X}_k \sim \mathfrak{Y}_k$ for all $k \in \mathbb{N}$.

Proof. The statement of the theorem can be proven by induction over k : the assumption $\mathfrak{X}_0 \sim \mathfrak{Y}_0$ is our induction start. Next, if $\mathfrak{X}_k \sim \mathfrak{Y}_k$, there exists for every $X_k \in \mathfrak{X}_k \cap \mathfrak{Y}$ an offset vector $a_k \in \mathbb{R}^n$ such that $Y_k = X_k + a_k \in \mathfrak{Y}_k$. Because $\mathfrak{Y}$ satisfies

$$\forall a \in \mathbb{R}^n, \quad \forall V \in \mathfrak{Y}, \quad V + a \in \mathfrak{Y},$$

it follows that $Y_k = X_k + a_k \in \mathfrak{Y}_k \cap \mathfrak{Y}$. Consequently, a relation of the form

$$\begin{aligned} AX_k + B\mu_k(X_k) + \mathbb{W} &= AY_k + (B\mu_k(X_k) - Aa_k) + \mathbb{W} \\ &= AY_k + Bv_k(Y_k) + \mathbb{W} - a_{k+1}, \end{aligned}$$

can be established, where a_{k+1} is given by

$$a_{k+1} \stackrel{\text{def}}{=} AX_k + Bv_k(Y_k) - B\mu_k(X_k) \in \mathbb{R}^n.$$

A completely symmetric relation holds after exchanging the roles of $\mathfrak{X}_k$ and $\mathfrak{Y}_k$. In summary, it follows that an implication of the form

$$\mathfrak{X}_k \sim \mathfrak{Y}_k \implies F(\mathfrak{X}_k, \mu_k) \sim F(\mathfrak{Y}_k, v_k)$$

holds. An application of Proposition 4 to the latter equivalence relation yields the induction step. $\square$

The above theorem allows us to formulate an intrinsic separation principle. Namely, Theorem 1 implies that the predicted intrinsic information content of a tight information tube does not depend on the choice of the control policy sequence with which it is generated. In particular, the tight information tubes from (8) satisfy

$$\forall k \in \mathbb{N}, \quad \mathfrak{D}^\circ(\mathfrak{X}_k) = \mathfrak{D}^\circ(\mathfrak{Y}_k)$$

for any intrinsic information measure $\mathfrak{D}^\circ$, independent of the choice of the control policy sequences.

Example 4. Examples 2 and 3 can be used to illustrate Theorem 1: let us consider the control law

$$\hat{\mu}(X) = -\frac{1}{2} - \left(\min_{x \in X} e_7^T x \right),$$

where $e_7 = [0, 0, 0, 0, 0, 0, 1]^T \in \mathbb{R}^7$ denotes the seventh unit vector. Its associated tight information tube

$$\begin{aligned} \hat{\mathfrak{X}}_0 &= \mathfrak{X}_0, \quad \hat{\mathfrak{X}}_1 = \mathfrak{X}_1, \quad \hat{\mathfrak{X}}_2 = \mathfrak{X}_2, \\ \hat{\mathfrak{X}}_3 &= \left\{ X \in \mathbb{K}_c^7 \mid X \subseteq \left[\begin{pmatrix} 0 \\ 0 \\ 0 \\ -\frac{1+2\eta_2}{4} \\ -\frac{1-2\eta_2}{4} \\ -\frac{1+2\eta_1}{4} \\ -\frac{1}{2} \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 0 \\ \frac{1+2\eta_2}{4} \\ \frac{1-2\eta_2}{4} \\ \frac{1+2\eta_1}{4} \\ -\eta_1 \end{pmatrix} \right] + \mathbb{W} \right\} \\ \hat{\mathfrak{X}}_k &= \hat{\mathfrak{X}}_3 \quad \text{for } k > 3 \end{aligned} \quad (9)$$

is intrinsically equivalent to the information tube from Example 3 for $\mu(X) = 0$. The difference between μ and $\hat{\mu}$ is visualized in the right part of Fig. 1. While μ maps the centers of the intersection sets to the same point, $\hat{\mu}$ re-locates these informations sets in dependence on their geometric properties, such that their lower edges all are aligned at the hyperplane $\{x \mid e_7^T x = -1/2\}$.




Fig. 1. LEFT: Visualization of the information set $\mathfrak{X}$ for our tutorial example (see Examples 2, 3, and 4) at any given time instance $k > 0$ in the $(x_1-x_2-x_3)$ -space. The four vertices (blue dots) are located at w_1^+ , w_1^- , w_2^+ , and w_2^- . MIDDLE-LEFT: Visualization of the posterior $\mathfrak{X} \cap \mathfrak{V}$: because we can measure the third state, $x_3 = \eta$, the sets of the collection $\mathfrak{X} \cap \mathfrak{V}$ correspond to all possible intersections of the gray shaded simplex with horizontal planes at height $\eta \in [-1/2, 1/2]$ and all compact subsets of these intersections. The green solid line, the blue shaded square, and the orange dashed line show, respectively, the intersection of the simplex with a horizontal plane at height $\eta = 1/2$, $\eta = 0$, and $\eta = -1/2$. MIDDLE-RIGHT: the information set $\mathfrak{X}^+ = F(\mathfrak{X}, \mu)$ in the (x_4-x_5) -plane at the next time instance. The centers of the dashed orange line, the light blue shaded square, and the green solid line are all mapped to the same point. RIGHT: The upper right plot visualizes the information set $\mathfrak{X}^{++} = F(\mathfrak{X}^+, \mu)$ in the (x_6-x_7) -plane, while the lower right part visualizes $\hat{\mathfrak{X}}^{++} = F(\mathfrak{X}^+, \hat{\mu})$, also in the (x_6-x_7) -plane, but by using the control law $\hat{\mu}$ from Example 4. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

5. Dual control

This section discusses assumptions under which dual control problems for linear systems can be separated into a meta learning and a robust control problem.

5.1. Receding horizon control

Dual control problems can be implemented in analogy to traditional model predictive control (MPC) methods. Here, one solves the online optimization problem

$$J(X_0) = \inf_{\mathfrak{X}, \mu} \sum_{k=0}^{N-1} L(\mathfrak{X}_k, \mu_k) + M(\mathfrak{X}_N) \quad (10)$$

$$\text{s.t.} \quad \begin{cases} \forall k \in \{0, 1, \dots, N-1\}, \\ F(\mathfrak{X}_k, \mu_k) \subseteq \mathfrak{X}_{k+1}, \quad X_0 \in \mathfrak{X}_0 \\ \mu_k \in \mathcal{U}, \quad \forall X_k \in \mathfrak{X}_k, \quad X_k \subseteq \mathbb{X} \end{cases}$$

on a finite time horizon $\{0, 1, \dots, N\}$, where 0 is the current time. The optimization variables are the feedback policies $\mu_0, \mu_1, \dots, \mu_{N-1} \in \mathcal{U}$ and their associated information tube, $\mathfrak{X}_0, \mathfrak{X}_1, \dots, \mathfrak{X}_N \in \mathbb{I}_c^n$. In the most general setting, the stage and terminal cost functions, $L : \mathbb{I}_c^n \times \mathcal{U} \rightarrow \mathbb{R}$ and $M : \mathbb{I}_c^n \rightarrow \mathbb{R}$, are assumed to be lower semi-continuous, although some of the analysis results below will be based on stronger assumptions. We recall that $\mathbb{X}$ denotes the closed state constraint set. The parameter $X_0 \in \mathbb{I}_c^n$ corresponds to the current information set. It is updated twice per sampling time by repeating the following steps online:

- (i) Wait for the next measurement η .
- (ii) Update the information set,
$$X_0 \leftarrow X_0 \cap \{x \in \mathbb{R}^n \mid \eta - Cx \in \mathbb{V}\}.$$
- (iii) Solve (10) and denote the first element of the optimal feedback sequence by $\mu_0^* \in \mathcal{U}$.
- (iv) Send $u^* = \mu_0^*(X_0)$ to the real process.
- (v) Propagate the information set,
$$X_0 \leftarrow AX_0 + Bu^* + \mathbb{W}.$$
- (vi) Set the current time to 0 and continue with Step (i).

Note that Step (iii) assumes that the “inf” operator in (10) can be replaced by a “min” and that an associated optimal feedback policy exists. Conditions under which this can be guaranteed are discussed in Section 5.4.

5.2. Objectives

Tube model predictive control formulations (Rawlings, Mayne, & Diehl, 2017; Villanueva, De Lazzari, Müller, & Houska, 2020; Villanueva, Müller, & Houska, 2024) use risk measures as stage cost functions. Here, a lower semi-continuous function of the form $R : \mathbb{K}_c^n \rightarrow \mathbb{R} \cup \{\infty\}$ is a risk measure if it satisfies the monotonicity condition

$$X \subseteq Y \implies R(X) \leq R(Y) \quad (11)$$

holds for all $X, Y \in \mathbb{K}_c^n$. Similarly, this paper proposes to call $\mathfrak{R} : \mathbb{I}_c^n \rightarrow \mathbb{R} \cup \{\infty\}$ an extrinsic risk measure if

$$\mathfrak{R}(\mathfrak{X}) = R\left(\bigcup_{X \in \mathfrak{X}} X\right) \quad (12)$$

for a lower semi-continuous function R that satisfies (11). By using the language of this paper, the traditional objective of dual control (Feldbaum, 1961) is to tradeoff between extrinsic risk and intrinsic deviation. This motivates to consider stage cost functions of the form

$$L(\mathfrak{X}, \mu) = \mathfrak{R}(\mathfrak{X}) + \tau \cdot \mathfrak{D}^\circ(\mathfrak{X}). \quad (13)$$

Here, $\mathfrak{R}$ denotes a lower semi-continuous extrinsic risk measure and $\mathfrak{D}^\circ$ a lower semi-continuous intrinsic information measure. For general nonlinear systems, the parameter $\tau > 0$ can be used to tradeoff between risk and deviation. In the context of constrained linear systems, however, such a tradeoff is superfluous, as formally proven in the sections below.

Remark 2. The stage cost function (13) can be augmented by a control penalty. For example, one could set

$$L(\mathfrak{X}, \mu) = \mathfrak{R}(\mathfrak{X}) + \tau \cdot \mathfrak{D}^\circ(\mathfrak{X}) + \mathfrak{C}(\mu), \quad (14)$$

where $\mathfrak{C} : \mathcal{U} \rightarrow \mathbb{R}$ models a lower semi-continuous control cost. This additional term does, however, not change the fact that the parameter τ does not affect the optimal solution of (10). Details about how to construct $\mathfrak{C}$ in practice will be discussed later on in this paper, see Section 6.

Remark 3. Dual control laws are, in general, less conservative than control laws that make decisions based on state estimates. In order to understand why this is so, we revisit Examples 2, 3, and 4 for the state constraint

$$\mathbb{X} = \left\{ x \in \mathbb{R}^7 \mid \begin{array}{l} x_7 \geq -1/2 \\ 2x_6 + x_7 \leq 1/2 \\ -2x_6 + x_7 \leq 1/2 \end{array} \right\}.$$


In this case, the dual control law $\hat{\mu}$ from Example 4 is the only feasible control law. It is implemented by setting

$$\forall k \geq 2, \quad u_k = \frac{1}{2} + \eta_{k-1} \quad (15)$$

implying that the control at time k depends on the state measurement η_{k-1} at time $k-1$. Because the state estimate at time k ,

$$\hat{x}_k = (0, 0, \eta_k, 0, 0, 0, 0)^\top$$

does not depend on η_{k-1} , it is, however, impossible to express the right-hand of (15) as a function of the current state estimate $\hat{x}_k$. Consequently, this example shows that state-estimation based control design structures can be sub-optimal, or, as in this case, even infeasible, as visualized below:



In both figures, the black triangle represents the feasible region $\mathbb{X}$ in the (x_6-x_7) -plane. As in the right part of Fig. 1, the green solid line, the blue shaded square, and the orange dashed line visualize selected extreme sets of the controlled information ensemble. The left set collection shows the smallest possible constraint violation that can be achieved with a state-estimation based controller. In Fig. 1's middle right image, the green horizontal information set and the orange vertical information set correspond to the same state estimate, making it impossible for a state estimation based controller to move these sets independently. However, the vertical orange and the horizontal green information sets do not fit into the feasible triangle if their centers align. This is in contrast to the optimal dual control law $\hat{\mu}$ visualized on the right side, which can fit all information sets into the feasible triangle by making independent control decisions based on the shapes, sizes, and orientations of these sets.

5.3. Separation of meta-learning and robust control

This section shows how the dual control problem (10) can be decomposed into an intrinsic meta-learning and an extrinsic robust control problem. We assume that

- (i) the stage cost function L has the form (13),
- (ii) the function $\mathfrak{R}$ is an extrinsic risk measure,
- (iii) the function $\mathfrak{D}^\circ$ is intrinsic and $\tau \geq 0$, and
- (iv) the function M is extrinsic and monotonous,

$$\mathfrak{X} \subseteq \mathfrak{Y} \implies M(\mathfrak{X}) \leq M(\mathfrak{Y}).$$

In this context, the solution of the meta learning problem is given by the recursion

$$\begin{aligned} \forall k \in \mathbb{N}, \quad \mathfrak{Y}_{k+1} &\stackrel{\text{def}}{=} \bar{F}(\mathfrak{Y}_k, v_k) \\ \text{with} \quad \mathfrak{Y}_0 &\stackrel{\text{def}}{=} \{X \in \mathbb{K}_c^n \mid X \subseteq X_0\}, \end{aligned} \quad (16)$$

for a constant sequence $v_0, v_1, \dots \in \mathcal{U}$. For simplicity of presentation, we assume $0 \in \mathbb{U}$ such that we can set $v_k(X) = 0$ without loss of generality. Due to Theorem 1, L satisfies

$$L(\mathfrak{X}_k) = \mathfrak{R}(\mathfrak{X}_k) + \tau \cdot \mathfrak{D}^\circ(\mathfrak{Y}_k)$$

along any optimal tube of (10). Hence, (10) reduces to a robust control problem whose objective and constraint functions are extrinsic, while the shapes, sizes and orientations of the optimal

information tube's sets are constants, given by (16). In summary, the contribution of intrinsic information to the objective value of (10),

$$J_I(X_0) \stackrel{\text{def}}{=} \tau \cdot \sum_{k=0}^{N-1} \mathfrak{D}^\circ(\mathfrak{Y}_k),$$

depends on X_0 but it does not depend on the choice of the control law. It can be separated from the contribution of extrinsic objective terms, as elaborated below.

5.4. Existence of solutions

In order to discuss how one can — after evaluating the meta-learning recursion (16) — rewrite (10) in the form of an extrinsic robust control problem, a change of variables is introduced. Let $\mathcal{B}_k$ denote the set of bounded functions of the form $c_k : \mathfrak{Y}_k \rightarrow \mathbb{R}^n$. It is a Banach space with respect to its norm $\|c_k\| \stackrel{\text{def}}{=} \sup_{X \in \mathfrak{Y}_k} \|c_k(X)\|$. Due to Theorem 1, any tight information tube $\mathfrak{X}_0, \mathfrak{X}_1, \dots \in \mathbb{I}_c^n$, started at $\mathfrak{X}_0 = \mathfrak{Y}_0$, is intrinsically equivalent to the precomputed tube $\mathfrak{Y}_0, \mathfrak{Y}_1, \dots \in \mathbb{I}_c^n$ and

$$\mathfrak{X}_k = \{Y + c_k(Y) \mid Y \in \mathfrak{Y}_k\} \quad (17)$$

for suitable translation functions $c_k \in \mathcal{B}_k$. In the following, we introduce the auxiliary set

$$\mathcal{C}_k \stackrel{\text{def}}{=} \left\{ (c, c^+) \mid \begin{array}{l} \forall X \in \partial[\mathfrak{Y}_k \cap \mathfrak{Y}], \\ A c(X) - c^+(A X + \mathbb{W}) \in (-B \mathbb{U}) \end{array} \right\}$$

recalling that ∂ denotes the boundary operator that returns the set of extreme sets of a given information ensemble. Because $\mathbb{U}$ is compact, $\mathcal{C}_k \subseteq \mathcal{B}_k \times \mathcal{B}_{k+1}$ is a closed set. Additionally, we introduce the shorthands

$$\mathcal{R}_k(c_k) \stackrel{\text{def}}{=} \mathfrak{R}(\{Y + c_k(Y) \mid Y \in \mathfrak{Y}_k\})$$

$$\text{and } \mathcal{R}_N(c_N) \stackrel{\text{def}}{=} M(\{Y + c_N(Y) \mid Y \in \mathfrak{Y}_N\}).$$

Since we assume that $\mathfrak{R}$ and M are lower-semicontinuous on $\mathbb{I}_c^n$, the functions $\mathcal{R}_k : \mathcal{B}_k \rightarrow \mathbb{R}$ are lower semi-continuous on the Banach spaces $\mathcal{B}_k$. They can be used to formulate the extrinsic robust control problem¹

$$\begin{aligned} J_E(X_0) = \min_{c_0, c_1, \dots, c_N} & \sum_{k=0}^{N-1} \mathcal{R}_k(c_k) + \mathcal{R}_N(c_N) \\ \text{s.t.} & \begin{cases} \forall k \in \{0, 1, \dots, N-1\}, \\ (c_k, c_{k+1}) \in \mathcal{C}_k, \quad c_0 \equiv 0, \\ \forall Y \in \mathfrak{Y}_k, \quad Y + c_k(Y) \subseteq \mathfrak{X}, \end{cases} \end{aligned} \quad (18)$$

which can be used to find the optimal solution of (10).

Theorem 2. Let $\mathbb{X} \in \mathbb{K}^n$ be a closed set, let $\mathbb{U} \in \mathbb{K}_c^{n_u}$, $\mathbb{V} \in \mathbb{K}_c^{n_v}$, and $\mathbb{W} \in \mathbb{K}_c^{n_w}$ be compact sets, let L be given by (13) with $\mathfrak{R}$ and M being set-monotonous and lower semi-continuous extrinsic risk measures, and let $\mathfrak{D}^\circ$ be an intrinsic lower semi-continuous information measure. Then the following statements hold.

- (i) Problem (18) admits a minimizer or is infeasible.
- (ii) Problem (10) is intrinsically separable; that is,

$$J(X_0) = J_E(X_0) + J_I(X_0).$$

- (iii) If $c_0, c_1, \dots, c_N$ is a minimizer of (18), its associated sequence of information ensembles, given by (17), is an optimal information tube of (10).

¹ If the sets $\mathbb{U}$ and $\mathbb{X}$ and the functions $\mathcal{R}_k$ are convex, (18) is a convex optimization problem.

Proof. Because the objective functions $\mathcal{R}_k$ of (18) are lower semicontinuous and since the feasible set of (18) is closed under the listed assumptions, it follows directly that this optimization problem admits a minimizer or is infeasible. Next, a relation between (10) and (18) needs to be established. We divide the proof into three parts.

Part I. Let us assume that $\mathfrak{X}_0, \mathfrak{X}_1, \dots, \mathfrak{X}_N \in \mathbb{I}_c^n$ is a tight information tube for given $\mu_0, \mu_1, \dots, \mu_{N-1} \in \mathcal{U}$,

$$\forall k \in \{0, 1, \dots, N-1\}, \quad \mathfrak{X}_{k+1} = \bar{F}(\mathfrak{X}_k, \mu_k).$$

Due to Theorem 1, there exist functions $c_k : \mathcal{B}_k \rightarrow \mathbb{R}^n$ such that $\mathfrak{X}_k = \{Y + c_k(Y) \mid Y \in \mathfrak{Y}_k\}$. The goal of the first part of this proof is to show that $(c_k, c_{k+1}) \in C_k$. Because the information tube is tight, we have $AX + B\mu_k(X) + \mathbb{W} \in \partial\mathfrak{X}_{k+1}$ for all $X \in \partial[\mathfrak{X}_k \cap \mathfrak{V}]$. Since any set $Y \in \partial[\mathfrak{Y}_k \cap \mathfrak{V}]$ is mapped to an extreme set $X = Y + c_k(Y) \in \partial[\mathfrak{X}_k \cap \mathfrak{V}]$, it follows that

$$\begin{aligned} & A(Y + c_k(Y)) + B\mu_k(X) + \mathbb{W} \in \partial\mathfrak{X}_{k+1} \\ \implies & \underbrace{(AY + \mathbb{W})}_{\in \partial\mathfrak{Y}_{k+1}} + \underbrace{(c_k(Y) + B\mu_k(X))}_{\in \mathbb{R}^n} \in \partial\mathfrak{X}_{k+1} \end{aligned}$$

for any such pair (X, Y) . But this is only possible if

$$c_k(Y) + B\mu_k(X) = c_{k+1}(AY + \mathbb{W}). \quad (19)$$

Since $\mu_k(X) \in \mathbb{U}$ and since the choice of $Y \in \partial[\mathfrak{Y}_k \cap \mathfrak{V}]$ is arbitrary, it follows from (19) that $(c_k, c_{k+1}) \in C_k$.

Part II. The goal of the second part of this proof is to reverse the construction from the first part. For this aim, we assume that we have functions $c_k : \mathcal{B}_k \rightarrow \mathbb{R}^n$ that satisfy the recursivity condition $(c_k, c_{k+1}) \in C_k$ for all $k \in \{0, 1, \dots, N-1\}$ while the sets $\mathfrak{X}_k$ are given by (17). Since every set $X \in \mathfrak{X}_k \cap \mathfrak{V}$ is contained in at least one extreme set $\bar{X} \in \partial[\mathfrak{X}_k \cap \mathfrak{V}]$, there exists for every such X a set $\bar{Y} \in \partial[\mathfrak{Y}_k \cap \mathfrak{V}]$ with $X \subseteq \bar{X} = \bar{Y} + c_k(\bar{Y})$. Note that this is equivalent to stating that there exists a function $\Sigma_k : \mathfrak{X}_k \cap \mathfrak{V} \rightarrow \partial[\mathfrak{Y}_k \cap \mathfrak{V}]$ that satisfies

$$\forall X \in \mathfrak{X}_k \cap \mathfrak{V}, \quad X \subseteq \Sigma_k(X) + c_k(\Sigma_k(X)). \quad (20)$$

It can be used to define the control laws

$$\mu_k(X) \stackrel{\text{def}}{=} B^\dagger [c_{k+1}(A\Sigma_k(X) + \mathbb{W}) - Ac_k(\Sigma_k(X))], \quad (21)$$

where $B^\dagger$ denotes the pseudo-inverse of B . Because we assume $(c_k, c_{k+1}) \in C_k$, we have $\mu_k(X) \in \mathbb{U}$ and

$$\begin{aligned} & AX + B\mu_k(X) + \mathbb{W} \\ & \stackrel{(20),(21)}{\subseteq} A\Sigma_k(X) + \mathbb{W} + c_{k+1}(A\Sigma_k(X) + \mathbb{W}) \in \mathfrak{X}_{k+1} \end{aligned}$$

for all $X \in \mathfrak{X}_k \cap \mathfrak{V}$, where the latter inclusion follows from (17) and the fact that $A\Sigma_k(X) + \mathbb{W} \in \mathfrak{Y}_{k+1}$. Consequently, we have $\mathfrak{X}_{k+1} \supseteq F(\mathfrak{X}_k, \mu_k)$.

Part III. The construction from Part I can be used to start with any feasible information tube of (10) to construct a feasible sequence $c_0, c_1, \dots, c_N$ such that

$$\begin{aligned} J_E(X_0) & \leq \sum_{k=0}^{N-1} \mathcal{R}_k(c_k) + \mathcal{R}_N(c_N) \\ & = \sum_{k=0}^{N-1} L(\mathfrak{X}_k, \mu_k) + M(\mathfrak{X}_N) - J_I(X_0). \end{aligned} \quad (22)$$

Thus, we have $J_E(X_0) + J_I(X_0) \leq J(X_0)$. Similarly, the construction from Part II can be used to start with an optimal solution of (18) to construct a feasible point of (10), which implies $J_E(X_0) + J_I(X_0) \geq J(X_0)$. Thus, the second and the third statement of the theorem hold. $\square$

5.5. Recursive feasibility and stability

Feasible invariant information ensembles $\mathfrak{X}_s \in \mathbb{I}_c^n$ exist if and only if the optimization problem

$$\min_{\mathfrak{X}_s, \mu_s} L(\mathfrak{X}_s, \mu_s) \quad \text{s.t.} \quad \begin{cases} F(\mathfrak{X}_s, \mu_s) \subseteq \mathfrak{X}_s, \\ \mu_s \in \mathcal{U}, \\ \forall X \in \mathfrak{X}_s, \quad X \subseteq \mathbb{X} \end{cases} \quad (23)$$

is feasible. By solving this optimization problem, one can find optimal invariant information ensembles. In analogy to terminal regions in traditional MPC formulations (Rawlings et al., 2017) invariant information ensembles can be used as a terminal constraint, $\mathfrak{X}_N \subseteq \mathfrak{X}_s$. If (10) is augmented by such a terminal constraint, recursive feasibility can be guaranteed. If one chooses the terminal cost M such that

$$\min_{\mu \in \mathcal{U}} L(\mathfrak{X}, \mu) + M(\bar{F}(\mathfrak{X}, \mu)) \leq M(\mathfrak{X})$$

for all $\mathfrak{X} \in \mathbb{I}_c^n$, the objective value of (10) descends along the trajectories of its associated closed-loop system. Assuming L is continuous and positive definite, this condition serves as a basis for constructing Lyapunov functions. The specifics of these constructions are omitted here, as they resemble the construction of terminal costs in conventional Tube MPC methods (Rawlings et al., 2017; Villanueva et al., 2020).

6. Polytopic approximation methods

This section discusses how to solve the dual control problem (10) by using a polytopic approximation method. For this aim, we assume that $\mathbb{V}$ and $\mathbb{W}$ are given convex polytopes, while $\mathbb{X}$ and $\mathbb{U}$ are convex sets.

6.1. Configuration-constrained polytopes

Polytopic computing (Fukuda, 2020) is the basis for many set-theoretic methods in control (Blanchini & Miani, 2015). Specifically, tube model predictive control methods routinely feature parametric polytopes with frozen facet directions (Raković, Kouvaritakis, Findeisen, & Cannon, 2012; Raković, Levine, & Açikmeşe, 2016). In this context, configuration-constrained polytopes are of special interest, as they admit a joint parameterization of their facets and vertices (Houska, Müller, & Villanueva, 2025; Villanueva et al., 2024). They are defined as follows: let $Y \in \mathbb{R}^{m \times n}$ and $G \in \mathbb{R}^{n \times m}$ be matrices that define the parametric polytope

$$\begin{aligned} \forall y \in \mathcal{G}, \quad P(y) & \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n \mid Yx \leq y\} \\ \text{on} \quad \mathcal{G} & \stackrel{\text{def}}{=} \{y \in \mathbb{R}^m \mid Gy \leq 0\}; \end{aligned}$$

and let $\Lambda_1, \Lambda_2, \dots, \Lambda_v \in \mathbb{R}^{n \times m}$ be vertex maps, such that

$$P(y) = \text{conv}(\Lambda_1 y, \Lambda_2 y, \dots, \Lambda_v y) \iff y \in \mathcal{G},$$

where $\text{conv}(\cdot)$ denotes the convex hull. The condition $y \in \mathcal{G}$ is called a configuration-constraint. It restricts the parameter domain of P to a region on which both a halfspace and a vertex representation is possible. Details on how to construct the template matrix Y together with the cone $\mathcal{G}$ and the matrices Λ_i can be found in Villanueva et al. (2024).

6.2. Polytopic information ensembles

As pointed out in Section 3.2, the minimal representation of a closed information ensemble $\mathfrak{X} \in \mathbb{I}_c^n$ is given by its set of extreme sets, denoted by $\partial\mathfrak{X}$. This motivates to discretize (10),

by introducing a suitable class of information ensembles, whose extreme sets are configuration-constrained polytopes. In detail,

$$\mathfrak{P}(z) \stackrel{\text{def}}{=} \{ X \in \mathbb{K}_c^n \mid \exists y \in \mathbb{P}(z) : X \subseteq P(y) \}$$

defines such a class of information ensembles with

$$\mathbb{P}(z) \stackrel{\text{def}}{=} \{ y \in \mathbb{R}^m \mid Gy \leq 0, Zy \leq z \} \subseteq \mathcal{G}$$

being used to parameterize convex subsets of $\mathcal{G}$. The choice of $Z \in \mathbb{R}^{l \times m}$ influences the polytopic discretization accuracy and $z \in \mathbb{R}^l$ denotes its parameter. Note that $\mathfrak{P}(z) \in \mathbb{I}_c^n$ is for any such z a compact but potentially empty information ensemble.

6.3. Polytopic meta learning

Traditional set-theoretic methods face a variety of computational difficulties upon dealing with output feedback problems, as summarized concisely in [Blanchini and Miani \(2015, Chapter 11\)](#). The proposed meta learning framework has the potential to overcome these difficulties, because [Proposition 3](#) alleviates the need to intersect infinitely many information sets for the sake of predicting the evolution of a learning process. Instead, it is sufficient to compute one intersection at meta level. If our prior information about the system's state is represented by a polytopic ensemble, $\mathfrak{P}(z)$, the posterior $\mathfrak{P}(z) \cap \mathfrak{V} = \mathfrak{P}(z) \cap \mathfrak{V}$ needs to be computed, where $\mathfrak{V}$ is given by [\(7\)](#). Since $\mathbb{V}$ is assumed to be a polytope, $\mathfrak{V}$ can be written in the form

$$\mathfrak{V} = \{ X \in \mathbb{K}_c^n \mid \exists y \in \mathcal{G} : X \subseteq P(y), Z_1 y \leq \bar{v} \},$$

as long as the template matrices Y , G , and $Z_1 \in \mathbb{R}^{l_1 \times m}$ as well as the vector $\bar{v} \in \mathbb{R}^{l_1}$ are appropriately chosen. Here, we construct the matrix $Z = (Z_1^\top, Z_2^\top)^\top$ such that its first $l_1 \leq l$ rows coincide with Z_1 . This particular construction of Z ensures that the intersection

$$\mathfrak{P}(z) \cap \mathfrak{V} = \mathfrak{P}(\zeta) \quad \text{with} \quad \begin{cases} \zeta_1 = \min(z_1, \bar{v}) \\ \zeta_2 = z_2 \end{cases}$$

can be computed explicitly, where $\min(z_1, \bar{v})$ denotes the componentwise minimizer of the vectors z_1 and $\bar{v}$. The latter condition is not jointly convex in z and ζ , but our constructions are based on the convex relaxation

$$\begin{pmatrix} \bar{v} \\ z_2 \end{pmatrix} \leq \begin{pmatrix} \zeta_1 \\ \zeta_2 \end{pmatrix} \implies \mathfrak{P}(z) \cap \mathfrak{V} \subseteq \mathfrak{P}(\zeta). \quad (24)$$

Note that the conservatism that is introduced by this convex relaxation is negligible if the measurement error set $\mathbb{V}$ is small. In fact, for the exact output feedback case, $\mathbb{V} = \{0\}$, we have $\min(z_1, \bar{v}) = \bar{v}$, since the measurements are exact and, as such, always informative.

6.4. Extreme vertex polytopes

In analogy to the construction of the domain $\mathcal{G}$, a configuration domain $\mathcal{H} = \{ z \in \mathbb{R}^l \mid Hz \leq 0 \}$ can be chosen. In detail, by using the methods from [Villanueva et al. \(2024\)](#), a matrix $H \in \mathbb{R}^{l \times n_H}$ and matrices $\Omega_1, \dots, \Omega_{\bar{v}} \in \mathbb{R}^{m \times l}$ can be pre-computed, such that

$$\mathbb{P}(\zeta) = \text{conv}(\Omega_1 \zeta, \Omega_2 \zeta, \dots, \Omega_{\bar{v}} \zeta) \iff \zeta \in \mathcal{H}.$$

This has the advantage that the vertices $\Omega_j \zeta$ of the polytope $\mathbb{P}(\zeta)$ are known. In order to filter the vertices that are associated to extreme polytopes, the index set

$$\mathbb{J} \stackrel{\text{def}}{=} \{ j \in \{1, 2, \dots, \bar{v}\} \mid P(\Omega_j \zeta) \in \partial[\mathfrak{P}(\zeta)] \}$$

is introduced. Its definition does not depend on the choice of the parameter $\zeta \in \mathcal{H}$, because the normal cones of the vertices of $\mathbb{P}(\zeta)$ do not depend on $\zeta \in \mathcal{H}$.

Definition 7. The polytopes $P(\Omega_j \zeta)$, with $j \in \mathbb{J}$, are called the extreme vertex polytopes of $\mathfrak{P}(\zeta)$.

Note that the convex hull of the extreme vertex polytopes encodes the extrinsic information of $\mathfrak{P}(\zeta)$,

$$\text{conv}(\{ P(\Omega_j \zeta) \mid j \in \mathbb{J} \}) = \bigcup_{X \in \mathfrak{P}(\zeta)} X. \quad (25)$$

The latter equation follows from the fact that the vertices of the extreme polytopes of $\mathfrak{P}(\zeta)$ are contained in the convex hull of the vertices $\Lambda_i \Omega_j \zeta$ of the extreme vertex polytopes, with $i \in \{1, \dots, v\}$ and $j \in \mathbb{J}$.

6.5. Polytopic information tubes

This section shows that assigning one extreme control input per extreme vertex polytope discretizes the control law without conservatism. Let us assume that $\mathbb{W} = P(\bar{w})$ is a polytope with given parameter $\bar{w} \in \mathcal{G}$. Moreover, we assume that the vertices of $\mathbb{P}(\cdot)$ are enumerated in such a way that $\mathbb{J} = \{1, 2, \dots, |\mathbb{J}|\}$, where $|\mathbb{J}| \leq \bar{v}$ denotes the cardinality of $\mathbb{J}$. We introduce the convex auxiliary set

$$\mathcal{F} \stackrel{\text{def}}{=} \left\{ (z, z^+) \mid \begin{array}{l} \exists (\zeta, \xi, u) \in \mathbb{R}^l \times (\mathbb{R}^m)^{|\mathbb{J}|} \times \mathbb{U}^{|\mathbb{J}|} \\ \forall i \in \{1, \dots, v\}, \forall j \in \mathbb{J}, \\ \bar{v} \leq \zeta_1, \quad z_2 \leq \zeta_2, \\ YA\Lambda_i \Omega_j \zeta + YBu_j + \bar{w} \leq \xi_j, \\ G\xi_j \leq 0, \quad H\zeta \leq 0, \quad Z\xi_j \leq z^+ \end{array} \right\}.$$

The rationale behind the convex constraints in this definition can be summarized as follows.

- We start with the current information ensemble $\mathfrak{P}(z)$.
- The constraints $\bar{v} \leq \zeta_1$ and $z_2 \leq \zeta_2$ subsume [\(24\)](#).
- The constraint $H\zeta \leq 0$ ensures that the vertices of $P(\Omega_j \zeta)$ are given by $\Lambda_i \Omega_j \zeta$, with $i \in \{1, \dots, v\}$.
- The extreme controls u_j are used to steer all vertices of $P(\Omega_j \zeta)$ into the auxiliary polytope $P(\xi_j)$.
- And, finally, the constraints $G\xi_j \leq 0$ and $Z\xi_j \leq z^+$ ensure that $P(\xi_j)$ is contained in $\mathfrak{P}(z^+)$.

Theorem 3. Let $\mathcal{F}$ and $\mathfrak{V}$ be defined as above, recalling that $\mathbb{W} = P(\bar{w})$ denotes the uncertainty set and that $\mathbb{U}$ is assumed to be convex. Then, the implication

$$(z, z^+) \in \mathcal{F} \implies \exists \mu \in \mathcal{U}, \quad \mathfrak{P}(z^+) \supseteq F(\mathfrak{P}(z), \mu)$$

holds for all $z, z^+ \in \mathbb{R}^l$.

Proof. Let us assume that $(z, z^+) \in \mathcal{F}$. As discussed in [Section 6.3](#), the inequalities $\bar{v} \leq \zeta_1$ and $z_2 \leq \zeta_2$ in the definition of $\mathcal{F}$ ensure that $\mathfrak{P}(z) \cap \mathfrak{V} \subseteq \mathfrak{P}(\zeta)$. Moreover, there exists for every $X \in \mathfrak{P}(\zeta)$ a $y \in \mathbb{P}(\zeta)$ with $X \subseteq P(y) \in \partial[\mathfrak{P}(\zeta)]$. Next, since we enforce $H\zeta \leq 0$, y is in the convex hull of the extreme vertices. That is, there exist scalar weights $\theta_1, \theta_2, \dots, \theta_{|\mathbb{J}|} \in [0, 1]$ with $\sum_{j \in \mathbb{J}} \theta_j = 1$ and $y = \sum_{j \in \mathbb{J}} \theta_j \Omega_j \zeta$, keeping in mind that these weights depend on X . They can be used to define the control law

$$\mu(X) \stackrel{\text{def}}{=} \sum_{j \in \mathbb{J}} \theta_j u_j \in \mathbb{U} \quad \text{and} \quad \xi \stackrel{\text{def}}{=} \sum_{j \in \mathbb{J}} \theta_j \xi_j \in \mathcal{G}$$

where $u_1, u_2, \dots, u_{|\mathbb{J}|} \in \mathbb{U}$ are the extreme control inputs and $\xi_1, \xi_2, \dots, \xi_{|\mathbb{J}|} \in \mathcal{G}$ are the auxiliary variables that satisfy the constraints from the definition of $\mathcal{F}$. Note that this construction is such that the vertices of the polytope $P(y)$, which are given by $\Lambda_i y$, satisfy

$$A\Lambda_i y + B\mu(X) + w = \sum_{j \in \mathbb{J}} \theta_j [A\Lambda_i \Omega_j \zeta + Bu_j + w] \in P(\xi),$$

where the latter inclusion holds for all $w \in W$. Consequently, since this holds for all vertices of $P(y)$, we have

$$AX + B\mu(X) + \mathbb{W} \subseteq AP(y) + B\mu(X) + \mathbb{W} \subseteq P(\xi).$$

Moreover, the above definition of ξ and the constraints $G\xi_j \leq 0$ and $Z\xi_j \leq z^+$ from the definition of $\mathcal{F}$ imply that $Z\xi \leq z^+$ and $P(\xi) \in \mathfrak{P}(z^+)$. But this yields $F(\mathfrak{P}(z), \mu) \subseteq \mathfrak{P}(z^+)$, which completes the proof. $\square$

6.6. Polytopic dual control

In order to approximate the original dual control problem (10) with a convex optimization problem, we assume that the stage and terminal cost functions have the form

$$L(\mathfrak{P}(z), \mu) = l(z, u) \quad \text{and} \quad M(\mathfrak{P}(z)) = m(z)$$

for given convex functions l and m , where the stacked vector $u = (u_1^T, u_2^T, \dots, u_{|\mathbb{J}|}^T)^T$ collects the extreme control inputs. Due to Theorem 3 a conservative approximation of (10) is given by

$$\begin{aligned} \min_{z, \zeta, \xi, u} \quad & \sum_{k=0}^{N-1} l(z_k, u_k) + m(z_N) \\ \text{s.t.} \quad & \begin{cases} \forall k \in \{0, \dots, N-1\}, \\ \forall i \in \{1, \dots, v\}, \forall j \in \mathbb{J}, \\ \bar{v} \leq \zeta_{k,1}, \quad z_{k,2} \leq \zeta_{k,2}, \\ YA\Lambda_i\Omega_j\zeta_k + YBu_{k,j} + \bar{w} \leq \xi_{k,j}, \\ G\xi_{k,j} \leq 0, \quad H\zeta_k \leq 0, \quad Z\xi_{k,j} \leq z_{k+1}, \\ u_{k,j} \in \mathbb{U}, \quad \Lambda_i\Omega_j\zeta_k \in \mathbb{X}, \quad Z\hat{y} \leq z_0. \end{cases} \end{aligned} \quad (26)$$

Since $\mathbb{U}$ and $\mathbb{X}$ are convex sets, this is a convex optimization problem. Its optimization variables are the parameters $z_k \in \mathbb{R}^l$ of the polytopic information tube, the associated extreme control inputs $u_{k,j} \in \mathbb{U}$ and the auxiliary variables $\zeta_k \in \mathcal{H}$ and $\xi_k \in \mathcal{G}$, needed to ensure that

$$\forall k \in \{0, 1, \dots, N-1\}, \quad F(\mathfrak{P}(z_k), \mu_k) \subseteq \mathfrak{P}(z_{k+1}).$$

Here, $X_0 = P(\hat{y})$ denotes the information set at the current time, modeled by the parameter $\hat{y} \in \mathcal{G}$. The constraint $Z\hat{y} \leq z_0$ corresponds to the initial value condition $X_0 \in \mathfrak{P}(z_0)$. Additionally, it is pointed out that the extrinsic information content of the auxiliary ensemble $\mathfrak{P}(\zeta) \supseteq \mathfrak{P}(z) \cap \mathfrak{V}$ overestimates the extrinsic information content of $\mathfrak{P}(z)$. Thus, the extrinsic state constraints can be enforced by using the implication chain

$$\begin{aligned} \forall i \in \{1, \dots, v\}, \forall j \in \mathbb{J}, \quad \Lambda_i\Omega_j\zeta_k \in \mathbb{X} \\ \implies \bigcup_{X \in \mathfrak{P}(\zeta)} X \subseteq \mathbb{X} \implies \bigcup_{X \in \mathfrak{P}(z)} X \subseteq \mathbb{X}. \end{aligned} \quad (27)$$

Finally, (26) is solved online whenever a new information set $X_0 = P(\hat{y})$ becomes available, denoting the optimal extreme controls by $u_{k,j}^*$. A corresponding feasible control law can then be recovered by setting $\mu_0^*(X_0) \stackrel{\text{def}}{=} \sum_{j \in \mathbb{J}} \theta_j^*(X_0) u_{0,j}$, where the scalar weights $\theta_j^*(X_0)$ can, for instance, be found by solving the convex quadratic programming problem

$$\begin{aligned} \theta^*(X_0) \stackrel{\text{def}}{=} \underset{\theta \geq 0}{\operatorname{argmin}} \quad & \left(\sum_{j \in \mathbb{J}} \theta_j^* u_{0,j} \right)^2 \\ \text{s.t.} \quad & \begin{cases} \sum_{j \in \mathbb{J}} \theta_j \Omega_j \zeta_0^* = \hat{y} \\ \sum_{j \in \mathbb{J}} \theta_j = 1, \end{cases} \end{aligned}$$

although, clearly, other choices for the weights θ_j^* are possible, too. Finally, the receding horizon control loop from Section 5.1 can be implemented by using the above expression for μ_0^* , while the information set update and propagation step can be implemented as in Blanchini and Miani (2015).

Remark 4. By solving the convex optimization problem

$$\begin{aligned} \min_{z^s, \zeta^s, \xi^s, u^s} \quad & l(z^s, u^s) \\ \text{s.t.} \quad & \begin{cases} \forall i \in \{1, \dots, v\}, \forall j \in \mathbb{J}, \\ \bar{v} \leq \zeta_1^s, \quad z_2^s \leq \zeta_2^s, \\ YA\Lambda_i\Omega_j\zeta^s + YBu_j^s + \bar{w} \leq \xi_j^s, \\ G\xi_j^s \leq 0, \quad H\zeta^s \leq 0, \quad Z\xi_j^s \leq z^s, \\ u_j^s \in \mathbb{U}, \quad \Lambda_i\Omega_j\zeta^s \in \mathbb{X} \end{cases} \end{aligned} \quad (28)$$

an optimal control invariant polytopic information ensemble can be computed.

6.7. Structure and complexity

Problem (26) admits a tradeoff between the computational complexity and the conservatism of polytopic dual control. In detail, this tradeoff can be adjusted by the choice of the following variables.

- (i) The number of facets, m , the number of vertices, v , and the number of configuration constraints, n_G , depend on our choice of Y and G . The larger m is, the more accurately we can represent the system's intrinsic information content.
- (ii) The number of information ensemble parameters, l , the number of extreme vertex polytopes, $|\mathbb{J}|$, and the number of meta configuration constraints, n_H , depends on how we choose Z and H . The larger $|\mathbb{J}|$ is, the more degrees of freedom we have to parameterize the optimal dual control law.

In contrast to these numbers, the number of controls, n_u , is given. If we assume that $\mathbb{U}$ and $\mathbb{X}$ are polyhedra with n_U and n_X facets, these number are given by the problem formulation, too. Additionally, we recall that N denotes the prediction horizon of the dual controller. The number of optimization variables n_{opt} and the number of constraints, n_{con} , of Problem (26) are given by

$$n_{\text{opt}} = (2N + 1)l + N|\mathbb{J}|(n_u + m)$$

$$n_{\text{con}} = N(l + |\mathbb{J}|(n_G + n_H) + l + n_U + v(m + n_X)) + l.$$

In this context, however, it should also be taken into account that the constraints of (26) are not only sparse but also possess further structure that can be exploited via intrinsic separation. For instance, the algebraic-geometric consistency conditions

$$\begin{aligned} GY &= 0, \quad \Lambda_i Y = \mathbb{1}, \\ HZ &= 0, \quad \Omega_j Z = \mathbb{1}, \quad \text{and} \quad Z_1 Y = 0 \end{aligned} \quad (29)$$

hold for all $i \in \{1, \dots, v\}$ and all $j \in \mathbb{J}$, which can be used to re-parameterize (26), if a separation of the centers and shapes of the information sets is desired.

Last but not least, more conservative but computationally less demanding variants of (26) can be derived by freezing some of its degrees of freedom. For instance, in analogy to Rigid Tube MPC (Langson, Chrysoschoos, Raković, & Mayne, 2004; Raković et al., 2012), one can implement a *Rigid Dual MPC* controller by pre-computing a feasible point $(z^s, \zeta^s, \xi^s, u^s)$ of (28). Next, we set

$$\begin{aligned} z_k &= ZY\bar{x}_k + z^s, \quad \zeta_k = ZY\bar{x}_k + \zeta^s, \\ \text{and } u_{j,k} &= \bar{u}_k + u_j^s, \quad \xi_{k,j} = Y\bar{x}_{k+1} + \xi_j^s, \end{aligned}$$

where $\bar{x}$ and $\bar{u}$ denote a central state- and a central control trajectory that are optimized online, subject to $\bar{x}_{k+1} = A\bar{x}_k + B\bar{u}_k$. By substituting these restrictions in (26) and by using (29), the resulting online optimization problem can be simplified and written in the form

$$\min_{\bar{x}, \bar{u}} \quad \sum_{k=0}^{N-1} \bar{\ell}(\bar{x}_k, \bar{u}_k) + \bar{m}(\bar{x}_N)$$

$$\text{s.t. } \begin{cases} \forall k \in \{0, \dots, N-1\}, \\ A\bar{x}_k + B\bar{u}_k = \bar{x}_{k+1}, \bar{u}_k \in \bar{\mathcal{U}}, \\ Z\hat{y} \leq ZY\bar{x}_0 + z^s, \bar{x}_k \in \bar{\mathcal{X}}. \end{cases} \quad (30)$$

Problem (30) can be regarded as a conservative approximation of (26). The sets

$$\bar{\mathcal{X}} \stackrel{\text{def}}{=} \left\{ \bar{x} \in \mathbb{R}^n \mid \begin{array}{l} \forall i \in \{1, \dots, \nu\}, \forall j \in \mathbb{J}, \\ \bar{x} + \Lambda_i \Omega_j \zeta^s \in \mathbb{X} \end{array} \right\}$$

$$\text{and } \bar{\mathcal{U}} \stackrel{\text{def}}{=} \left\{ x \in \mathbb{R}^n \mid \forall j \in \mathbb{J}, \bar{u} + u_j^s \in \mathcal{U} \right\}$$

take the robustness constraint margins into account, while $\bar{\ell}$ and $\bar{m}$ are found by re-parameterizing the objective function of (26). Problem (30) is a conservative dual MPC controller that has—apart from the initial value constraint—the same complexity as certainty-equivalent MPC. Its feedback law depends on the parameter $\hat{y}$ of the initial information set $X_0 = P(\hat{y})$.

6.8. Numerical illustration

We consider the constrained linear control system

$$\begin{aligned} A &= \frac{1}{4} \begin{pmatrix} 6 & 4 \\ 1 & 3 \end{pmatrix}, \quad B = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad C = (1 \ 0), \\ \mathbb{X} &= \{x \in \mathbb{R}^2 \mid x_2 \geq -45\}, \quad \mathcal{U} = [-55, 55], \\ \mathbb{W} &= [-1/2, 1/2]^2 \subseteq \mathbb{R}^2, \quad \mathbb{V} = [-1, 1]. \end{aligned} \quad (31)$$

In order to setup an associated polytopic dual controller for this system, we use the template matrices

$$Y = \begin{pmatrix} 1 & & & & & \\ & 1 & & & & \\ & & 1 & & & \\ -1 & & & & & \\ -1 & -1 & & & & \\ & & & -1 & & \end{pmatrix} \quad \text{and} \quad G = \begin{pmatrix} -1 & 1 & -1 & & & \\ & -1 & 1 & -1 & & \\ & & -1 & 1 & -1 & \\ & & & -1 & 1 & -1 \\ -1 & & & & -1 & 1 \\ 1 & -1 & & & & -1 \end{pmatrix}$$

setting $m = \nu = n_G = 6$. Here, Y and G are stored as sparse matrices: the empty spaces are filled with zeros. By using analogous notation, we set

$$Z = \begin{pmatrix} 1 & & & & & \\ & 1 & & & & \\ & & 1 & & & \\ & & & 1 & & \\ & & & & 1 & \\ & & & & & 1 \end{pmatrix}, \quad H = \begin{pmatrix} & -1 & -1 & 1 & & 1 \\ -1 & 1 & 1 & -1 & & -1 \\ & 1 & -1 & 1 & -1 & \\ & & 1 & -1 & & -1 \\ 1 & 1 & -1 & & -1 & \\ & -1 & & 1 & -1 & 1 \\ -1 & 1 & & -1 & 1 & -1 \\ & 1 & & -1 & 1 & -1 \\ 1 & & & -1 & 1 & -1 \end{pmatrix},$$

$l = 8$, and $n_H = 9$, which can be used to represent six dimensional meta polytopes with $6 + 8 = 14$ facets and $\bar{\nu} = 68$ vertices. They have up to $|\mathbb{J}| = 60$ extreme vertex polytopes. The first row of Z corresponds to the block matrix Z_1 . It can be used to represent the set $\mathfrak{V}$ by setting $\bar{\nu} = 2$, since the diameter of $\mathbb{V}$ is equal to 2. Moreover, due to our choice of Y and $\mathbb{W}$, we have

$$\bar{w} = [1/2, 1, 1/2, 1/2, 1, 1/2]^T \in \mathcal{G}.$$

Next, we construct the stage cost as in (14). We set

$$\mathfrak{R}(\mathfrak{X}) \stackrel{\text{def}}{=} \sum_{i=1}^6 \left(\max_{\{x\} \in \mathfrak{X}} Y_i x \right)^2 + 50 \cdot \sum_{i=1}^2 \max_{\{x, x'\} \in \mathfrak{X}} (x_i - x'_i)^2$$

$$\text{and } \mathfrak{D}^\circ(\mathfrak{X}) \stackrel{\text{def}}{=} \sum_{i=1}^2 \left(\max_{\{x\} \in \mathfrak{X}} \max_{x, x' \in \mathfrak{X}} |x_i - x'_i| \right)^2.$$

This particular choice of $\mathfrak{R}$ and $\mathfrak{D}^\circ$ is such that

$$\tau(z) \stackrel{\text{def}}{=} \mathfrak{R}(\mathfrak{P}(z)) \quad \text{and} \quad \mathfrak{d}^\circ(z) \stackrel{\text{def}}{=} \mathfrak{D}^\circ(\mathfrak{P}(z))$$

are convex quadratic forms in z that can be worked out explicitly. Namely, we find that

$$\tau(z) = \left(\sum_{i=1}^6 z_{i+2}^2 \right) + 50 \cdot [(z_3 + z_6)^2 + (z_5 + z_8)^2]$$

and $\mathfrak{d}^\circ(z) = z_1^2 + z_2^2$. We use the control penalty function

$$c(u) = \sum_{i=1}^{|\mathbb{J}|} \left[u_i^2 + 50 \cdot \left(u_i - \frac{1}{|\mathbb{J}|} \sum_{j=1}^{|\mathbb{J}|} u_j \right)^2 \right]$$

in order to penalize both the extreme inputs as well as the distances of these extreme inputs to their average value. The stage cost is given by $\ell(z, u) = \tau(z) + \tau \cdot \mathfrak{d}^\circ(z) + c(u)$, where we set the intrinsic regularization to $\tau = 0.01$.

The optimal invariant information ensemble $\mathfrak{P}(z_s)$ is found by solving (28), see the left plot of Fig. 2. The light blue shaded hexagon corresponds to the union of all sets in $\mathfrak{P}(z_s)$. The 60 extreme vertex polytopes of $\mathfrak{P}(\zeta_s)$, given by $P(\Omega_j \zeta_s)$ for $j \in \{1, 2, \dots, 60\}$, are clustered at the vertices of the extrinsic hexagon (they partly obscure each other; not all 60 are clearly visible). We visualize them in different shades of gray. As the optimal solution happens to satisfy $\mathfrak{P}(z_s) \cap \mathfrak{V} = \mathfrak{P}(\zeta_s)$, the convex relaxation (24) does not introduce conservatism. A closed-loop simulation of the polytopic dual controller (26) is started with the initial information set $X_0 = [17, 23] \times [17, 23]$ using the prediction horizon $N = 10$ while the terminal cost is set to

$$m(z_N) = \begin{cases} 0 & \text{if } z_N \leq z_s \\ \infty & \text{otherwise} \end{cases}$$

in order to enforce recursive feasibility. The right plot of Fig. 2 shows an extrinsic image of the first predicted tube; that is, the union of the sets $\mathfrak{P}(z_k)$ along the optimal solution of (26) for the above choice of X_0 .

Remark 5. The proposed polytopic dual control method optimizes feedback laws that depend on the system's information state. Such dual control laws have the potential to be less conservative than robust output feedback laws that are based on state estimates or observer equations with affine structure, as considered in Goulart and Kerrigan (2007) or (Mayne et al., 2009). The overall conservatism of set-theoretic control methods depends, however, not only on the observer structure but also on the control law discretization and the set arithmetic, which complicates a general comparison (Blanchini & Miani, 2015; Villanueva et al., 2024). Nevertheless, for Example 2–4, the proposed polytopic set propagation scheme can find the exact optimal dual control law if all extreme vertex polytopes are representable by the chosen template. As discussed in Remark 3, the dual optimal control law associated to this example cannot be represented by a state-estimation based control law.

7. Conclusions

This paper has presented a set-theoretic approach to dual control. It is based on meta information spaces that enable a data-free algebraic characterization of both the present and the future information content of learning processes. In detail, an intrinsic equivalence relation has been introduced in order to separate the




Fig. 2. LEFT: Visualization of the 60 extreme vertex polytopes of $\mathfrak{P}(z_s)$, colored in gray. The convex hull of these extreme polytopes is colored in light blue. It corresponds to the union of all sets in $\mathfrak{P}(z_s)$, as predicted by (25). RIGHT: Visualization of the extrinsic tube (light blue shaded polytopes) that corresponds to the first prediction of the dual controller (26) with horizon $N = 10$. These sets correspond to the union of the sets of the information ensembles $\mathfrak{P}(z_k)$. The union of the sets in the first information ensemble, $\mathfrak{P}(z_0)$, happens to coincide with the initial information set $X_0 = [17, 23] \times [17, 23]$. Active state constraints are colored in red. The union of the sets of the optimal invariant information ensemble is shown in dark blue. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

computation of the future information content of a constrained linear system from the computation of its robust optimal control laws. An associated intrinsic separation principle is summarized in Theorem 1. It is the basis for analyzing the existence of solutions of a large class of dual control problems under certain structural and continuity assumptions that are summarized in Theorem 2.

For the first time, this paper has presented a polytopic dual control method for constrained linear systems that is based on convex optimization. In contrast to existing robust output-feedback control schemes, this method optimizes control laws that depend on the system's information state. This alleviates the need to make control decisions based on state estimates or observer equations that induce a potentially sub-optimal feedback structure. Instead, (26) optimizes a finite number of control inputs that are associated to the extreme vertex polytopes of the predicted information ensembles.

A numerical case study for a system with two states has indicated that (26) can be solved without numerical problems for moderately sized problems. For larger systems, however, the computational complexity of accurate dual control can become exorbitant. In anticipation of this problem, this paper has outlined strategies towards reducing the computational complexity at the cost of more conservatism. For example, the Rigid Dual MPC problem (30) has essentially the same complexity as a comparable certainty-equivalent MPC problem.

References

- Almeida, T. A., & Dórea, C. E. T. (2021). Output feedback constrained regulation of linear systems via controlled-invariant sets. *IEEE Transactions on Automatic Control*, 66(7).
- Artstein, Z., & Raković, S. V. (2011). Set invariance under output feedback: A set-dynamics approach. *International Journal of Systems Science*, 42(4), 539–555.
- Bemporad, A., & Garulli, A. (2000). Output-feedback predictive control of constrained linear systems via set-membership state estimation. *International Journal of Control*, 73(8), 655–665.
- Bertsekas, D. P., & Rhodes, I. B. (1971). Recursive state estimation for a set-membership description of uncertainty. *IEEE Transactions on Automatic Control*, 16, 117–128.
- Blanchini, F., & Miani, S. (2003). Stabilization of LPV systems: State feedback, state estimation, and duality. *SIAM Journal on Control and Optimization*, 42(1), 76–97.
- Blanchini, F., & Miani, S. (2015). *Set-theoretic methods in control. Systems & control: Foundations & applications*. Birkhäuser Boston, Inc., Boston, MA.
- Brunner, F. D., Müller, M. A., & Allgöwer, F. (2018). Enhancing output-feedback MPC with set-valued moving horizon estimation. *IEEE Transactions on Automatic Control*, 63(9), 2976–2986.
- Doob, J. L. (1953). *Stochastic processes*. Wiley.
- Dórea, C. E. T. (2009). Output-feedback controlled-invariant polyhedra for constrained linear systems. *Proceedings of the 48th IEEE Conference on Decision and Control, Shanghai*, 5317–5322.
- Feldbaum, A. A. (1961). Dual-control theory (i-IV). *Automation and Remote Control*, 21, pages 1240–1249 and 1453–1464, 1960; and 22, pages 3–16 and 129–143.
- Filato, N. M., & Unbehauen, H. (2004). *Adaptive dual control*. Springer.
- Findeisen, R., Imsland, L., Allgöwer, F., & Foss, B. A. (2003). State and output feedback nonlinear model predictive control: An overview. *European Journal of Control*, 9, 190–207.
- Fukuda, K. (2020). *Polyhedral computation*. ETH Zürich Research Collection.
- Goulart, P., & Kerrigan, E. (2007). Output feedback receding horizon control of constrained systems. *International Journal of Control*, 80(1), 8–20.
- Hewing, L., Wabersich, K. P., Menner, M., & Zeilinger, M. N. (2020). Learning-based model predictive control: Toward safe learning in control. *Annual Review of Control, Robotics, and Autonomous Systems*, 3, 269–296.
- Houska, B., Müller, M. A., & Villanueva, M. E. (2025). Polyhedral control design: Theory and methods. *Annual Reviews in Control*, 60, Article 100992.
- Hovd, M., & Bitmead, R. R. (2005). Interaction between control and state estimation in nonlinear MPC. *Modeling, Identification, and Control*, 26(3), 165–174.
- Joseph, P. D., & Tou, J. T. (1961). On linear control theory. *AIEE Transactions on Applications and Industry*, 80, 193–196.
- Krasovskii, N. N. (1964). On the theory of controllability and observability of linear dynamic systems. *Journal of Applied Mathematics and Mechanics*, 28(1), 1–14.
- Krasovskii, A. N., & Krasovskii, N. N. (1995). *Control under lack of information*. Birkhäuser, Boston.
- Kurzanskii, A. B. (1972). Differential games of observation. *Doklady Akademii Nauk SSSR*, 207(3), 527–530.
- Kurzanskii, A. B. (2004). The problem of measurement feedback control. *Journal of Applied Mathematics and Mechanics*, 68, 487–501.
- Langson, W., Chrysoschoos, I., Raković, S. V., & Mayne, D. Q. (2004). Robust model predictive control using tubes. *Automatica*, 40(1), 125–133.
- Lindquist, A. (1973). On feedback control of linear stochastic systems. *SIAM Journal on Control*, 11, 323–343.
- Mayne, D. Q., Rakovic, S. V., Findeisen, R., & Allgöwer, F. (2009). Robust output feedback model predictive control of constrained linear systems: time varying case. *Automatica*, 45, 2082–2087.
- Raković, S., Kouvaritakis, B., Findeisen, R., & Cannon, M. (2012). Homothetic tube model predictive control. *Automatica*, 48(8), 1631–1638.
- Raković, S., Levine, W. S., & Açikmeşe, B. (2016). Elastic tube model predictive control. In *American control conference (ACC)*, 2016 (pp. 3594–3599). IEEE.
- Rawlings, J. B., Mayne, D. Q., & Diehl, M. M. (2017). *Model predictive control: Theory, computation, and design*. Madison, WI: Nob Hill Publishing.
- dos Reis de Souza, A., Efimov, D., Raïssi, T., & Ping, X. (2022). Robust output feedback model predictive control for constrained linear systems via interval observers. *Automatica*, 135(109951).
- Rockafellar, R. T., & Uryasev, S. (2013). The fundamental risk quadrangle in risk management, optimization and statistical estimation. *Surveys in Operations Research and Management Science*, 18, 33–53.
- Sehr, M. A., & Bitmead, R. R. (2019). *Probing and duality in stochastic model predictive control. Handbook of Model Predictive Control*. Control Engineering, pages 125–144, Birkhäuser.
- Stengel, R. (1994). *Optimal control and estimation*. Dover Publications, New York.
- Villanueva, M. E., De Lazzari, E., Müller, M. A., & Houska, B. (2020). A set-theoretic generalization of dissipativity with applications in tube MPC. *Automatica*, 122(109179).

- Villanueva, M. E., Müller, M. A., & Houska, B. (2024). Configuration-constrained tube MPC. *Automatica*, 163, Article 111543.
- van Warter, H., & Willems, J. C. (1981). The certainty equivalence property in stochastic control theory. *IEEE Transactions on Automatic Control*, AC-26(5), 1080–1087.
- Witsenhausen, H. S. (1968). Sets of possible states of linear systems given perturbed observations. *IEEE Transactions on Automatic Control*, 13, 556–558.
- Wonham, W. M. (1968). On the separation theorem of stochastic control. *SIAM Journal on Control*, 6(2), 312–326.
- Zanon, M., & Gros, S. (2021). Safe reinforcement learning using robust MPC. *IEEE Transactions on Automatic Control*, 66(8), 3638–3652.



Boris Houska is an associate professor at ShanghaiTech University. He received a diploma in mathematics from the University of Heidelberg in 2007 and a Ph.D. in electrical engineering from KU Leuven in 2011. From 2012 to 2013, he was a postdoctoral researcher at the Centre for Process Systems Engineering at Imperial College London. Boris Houska's research interests include numerical optimization, optimal control, robust and global optimization, and fast model predictive control algorithms.