UCLouvain

École polytechnique de Louvain

Unveiling System Controllability from Noisy Data by Harnessing Network Structures

Authors: **Arnaud COLLETTE, Antoine SPRINGAEL**
Supervisors: **Gianluca BIANCHIN, Jean-Charles DELVENNE**
Reader: **Julien HENDRICKX**
Academic year 2024–2025
Master [120] in Mathematical Engineering

Abstract

In recent years, data-driven control has gained prominence as an alternative to traditional model-based strategies, particularly in contexts where constructing accurate system models is difficult. Controllability remains a cornerstone of control theory, as it determines the feasibility of designing effective control inputs. This work addresses the problem of assessing controllability from noisy input–state trajectory data, aiming to bridge the gap between theoretical guarantees and the challenges posed by real-world noise-affected systems.

The study begins with a review of the literature on data-driven control, system identification, and structural analysis, with an emphasis on the various definitions of controllability. Special attention is given to the data-informativity framework, which allows controllability to be tested directly from measured data. However, both theoretical analysis and numerical evidence highlight a key limitation of this approach: its lack of robustness in the presence of noise.

To address this issue, a hybrid method is introduced. Rather than applying controllability tests directly to noisy measurements, the structural sparsity of the system is first inferred using either hard thresholding or LASSO regression. The resulting estimate is then analyzed using structural controllability theory. Simulations conducted on both random and deliberately constructed systems show that this two-step approach improves robustness significantly, yielding high accuracy even under non-negligible noise.

Overall, this framework offers a practical and noise-resilient alternative for data-driven controllability assessment, and suggests new research directions at the intersection of structural inference, noisy data and controllability.

Acknowledgements

We would like to sincerely thank Prof. Jean-Charles Delvenne and Prof. Gianluca Bianchin for their time, valuable advice, feedback, and intuition throughout this year. Their guidance through the various steps of this thesis was instrumental in shaping the direction of this thesis.

We would also like to thank Prof. Julien Hendrickx for agreeing to be part of the jury for our thesis. His willingness to contribute his expertise and attention to this work is deeply appreciated.

We extend our gratitude to our respective families for their unconditional support during these five years, and especially throughout the completion of this thesis.

Finally, we would like to thank our mutual friends Lucas, Mary Jo, Juliette, Anaïs, Pierre and Cyril for their motivation and trust during this past year.

In addition, we would particularly like to thank Paul-Emile Lobet for his advice and careful review of this thesis.

Contents

List of Notations

x Vector notation. For $x \in \mathbb{R}^n$,

$$x = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}.$$

A Matrix notation. For $A \in \mathbb{R}^{n \times n}$,

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nn} \end{bmatrix}.$$

$\|\cdot\|_1$ Vector 1-norm. For $x \in \mathbb{R}^n$,

$$\|x\|_1 = \sum_{i=1}^{n} |x_i|.$$

$\|\cdot\|_2$ Vector 2-norm (Euclidean). For $x \in \mathbb{R}^n$,

$$\|x\|_2 = \sqrt{\sum_{i=1}^{n} |x_i|^2}.$$

$\|\cdot\|_F$ Frobenius norm. For $A \in \mathbb{R}^{m \times n}$,

$$\|A\|_F = \sqrt{\sum_{i=1}^{m} \sum_{j=1}^{n} |A_{ij}|^2}.$$

$\otimes$ Kronecker product. For $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{p \times q}$,

$$A \otimes B = \begin{bmatrix} a_{11}B & \cdots & a_{1n}B \\ \vdots & \ddots & \vdots \\ a_{m1}B & \cdots & a_{mn}B \end{bmatrix} \in \mathbb{R}^{(mp) \times (nq)}.$$

$\mathcal{N}(\mu, \Sigma)$ Multivariate normal distribution. For mean $\mu \in \mathbb{R}^n$ and covariance matrix $\Sigma \in \mathbb{R}^{n \times n}$,

$$x \sim \mathcal{N}(\mu, \Sigma)$$

Chapter 1

Introduction

"The laws of nature are written in the language of mathematics." —Galileo Galilei

From collective animal motion such as bird flocking [1] to climate forecasting [2], from aircraft design [3] to personalized recommendation systems [4], mathematics provides the language through which we understand and influence the world. In particular, mathematical modeling plays a pivotal role in science and engineering. A model offers a simplified quantitative abstraction of a system, allowing us to describe, analyze, and anticipate the evolution of the system over time. Their universality and abstraction enable application across disciplines, enabling progress in both theoretical and applied domains.

In the field of systems and control, models typically take the form of *dynamical systems*, i.e., mathematical objects that describe how a system evolves over time under the influence of external inputs. A widely used class of such systems is that of discrete-time Linear Time-Invariant (LTI) systems, described by the update equation:

$$x(t+1) = Ax(t) + Bu(t), \tag{1.1}$$

where $x(t) \in \mathbb{R}^n$ is the state vector at time t, $u(t) \in \mathbb{R}^m$ is the input vector at time t, and $A \in \mathbb{R}^{n\times n}$, $B \in \mathbb{R}^{n\times m}$ are matrices encoding the system dynamics and input interactions, respectively. While many systems in practice are nonlinear, time-varying, or described more naturally in continuous time, the LTI framework (whether in discrete or continuous form) remains a foundational tool in control theory due to its analytical tractability and broad applicability.

A central concept in control theory is the *controllability* of a system. Informally, a system is controllable if it is possible to drive its state from any initial configuration to any desired final configuration within finite time, using admissible inputs. This notion is formalized by the following definition.

Definition 1.1 (Controllability)**.** A discrete-time system of the form (1.1) is said to be *controllable* if, for any initial state $x(0)$ and any final state $x(f)$, there exists a finite sequence of inputs $(u(0), \ldots, u(T-1))$ such that $x(T) = x(f)$.

Controllability is not merely a theoretical curiosity; it is a cornerstone assumption behind many powerful tools and results in systems and control, such as state feedback design [5], Kalman decomposition [6], and optimal control [7]. For LTI systems, the classical test for controllability is given by the Hautus criterion or the rank condition involving the controllability matrix.

Theorem 1.2 (Kalman Rank Condition [8])**.** *The LTI system* (1.1) *is controllable if and only if the controllability matrix*

$$\mathcal{C} = \begin{bmatrix} B & AB & A^2B & \cdots & A^{n-1}B \end{bmatrix}$$

has full rank, i.e., $\operatorname{rank}(\mathcal{C}) = n$.

These classical methods rely on the availability of a fully specified model. They belong to the *model-based* paradigm of control theory, in which the analysis and design of control laws presuppose knowledge of the system matrices A and B. However, in many contemporary applications, this assumption is increasingly difficult to satisfy.

Indeed, modern systems are often large-scale, high-dimensional, and interdependent. Examples include power grids [9], communication networks [10], transportation systems [11], and biological networks [12]. In such systems, obtaining a precise model is prohibitively complex, or even fundamentally impossible. Sometimes the underlying physics is unknown or too intricate; other times the network structure is partially hidden, and the sheer scale of the network, with thousands of nodes and interconnections, renders comprehensive modeling intractable. In fact, it is widely acknowledged that the most time-consuming and error-prone step in deploying model-based control is system identification [13].

This modeling challenge has sparked a growing interest in alternative approaches, notably those based on *data* and *structure*. In data-driven methods, the focus shifts from building an explicit model to extracting information directly from observed input-output or input-state trajectories. In structural methods, only partial knowledge of the system is assumed, such as known interconnections between states and inputs, or the underlying network topology. By *data*, we mean measurements of a dynamical system, typically of its inputs and outputs. By *structure*, we refer to the pattern of zero and nonzero entries in the system matrices A and B.

This thesis explores how the complementary strengths of data and structure can be harnessed to overcome the limitations of traditional control methods. The main focus is on determining whether controllability can be reliably inferred when the available data are corrupted by noise. Classical data-driven techniques often fail in such settings. To address this limitation, structural information such as sparsity patterns or network topology is extracted from data and used as a substitute for exact numerical knowledge. More precisely, a unified framework is developed for *data-driven controllability* under noise, built upon *structural controllability* tests and *topology identification* algorithms, enabling robust inference of controllability and the design of control strategies in the absence of precise models.

1.1 From Data to Controllability

In many modern applications, obtaining an accurate first-principles model is impractical or impossible due to high dimensionality, unmodeled dynamics, or unknown interconnections. The traditional two-step approach, consisting of system identification followed by model-based controller design, can therefore be prohibitively expensive or lead to unreliable performance. As an alternative, a rapidly growing body of work seeks to bypass explicit model construction and operate directly on measured data, giving rise to the *data-driven control* paradigm. This shift reflects a broader movement in systems and control theory toward leveraging experimental data as the primary source of information, rather than relying on analytical models.

Data-driven control traces its roots to empirical tuning rules and adaptive architectures. Early milestones include classical PID tuning heuristics [14], adaptive control schemes [15], iterative feedback tuning [16], and the concept of unfalsified control [17]. More recently, these model-free concepts have evolved into a spectrum of approaches, including one-shot predictive control based on Hankel matrices [18] or not [19], model reference schemes [20], reinforcement learning for optimal control [21, 22], and optimization-based PID tuning [23]. Surveys such as [24] document the breadth of these techniques and classify them according to design objectives and data requirements.

At the core of many of these developments lies Willems' Fundamental Lemma [25], which asserts that the set of trajectories of a controllable, observable LTI system can be parametrized by a single persistently exciting input–output sequence [26]. This result has inspired a wide range of behavioral, optimization-based, and subspace-free control formulations [27, 28], enabling direct controller synthesis from data without identifying the system matrices.

In parallel to controller design, data-driven *analysis* of system-theoretic properties has gained attention. Stability, dissipativity, and performance can be verified using data-driven linear matrix inequalities and generalizations of the S-lemma [29, 30]. Controllability and observability are often assessed via rank conditions on Hankel matrices or trajectory-based tests [31]. Alternative approaches construct data-based estimates of the controllability and observability Gramians, which remain accurate even in the presence of measurement and process noise [32, 33].

Despite this progress, most data-driven methods impose strong assumptions: the unknown system must be controllable, observable, or stabilizable a priori, and the input used to collect data must be persistently exciting of sufficiently high order. These assumptions are frequently difficult to guarantee in realistic settings. To relax these requirements, van Waarde *et al.* introduced the *informativity framework*, which recasts control and analysis tasks as set-membership problems over all systems consistent with observed data [34, 35]. In this setting, data are said to be *informative for controllability* if every model compatible with the measurements is controllable; similarly, data are *informative for stabilization* if there exists a single feedback gain that stabilizes all consistent systems.

A particularly notable result in this line of work is a *data-driven Hautus test* for controllability that avoids the need for persistently exciting inputs, relying instead on simple rank conditions applied to measured data [34]. This contrasts with earlier Hankel-rank tests, which rely on strong excitation conditions.

In practice, however, even small amounts of noise can cause purely data-driven controllability tests to fail, resulting in false positives and undermining reliability. To address this limitation, data can be augmented with partial structural information, such as known sparsity patterns or network connectivity, to restore robustness. The remainder of this thesis investigates this hybrid *data-and-structure* paradigm: failure modes of classical and informativity-based controllability tests under noise are analyzed, and algorithms are developed to recover and exploit structural priors, enabling reliable controllability assessment when data alone are insufficient.

1.2 From Data to Network Structure

In many settings, the interconnection pattern of a dynamical network is either unknown or only partially available. Recovering this pattern, commonly referred to as structure or topology identification, is crucial for characterizing system dynamics, diagnosing faults, and designing distributed control strategies.

By *structure* is meant the zero/nonzero pattern of the entries in the system matrices A and B of (1.1), irrespective of their numerical values. These pattern matrices are denoted by $\mathcal{A} \in \{0, *\}^{n \times n}$ and $\mathcal{B} \in \{0, *\}^{n \times m}$, where each "$*$" indicates a strictly nonzero entry and "0" indicates a zero entry. The structure thus captures only which states or inputs influence which states (note that inputs do not influence other inputs). A given structure $(\mathcal{A}, \mathcal{B})$ admits infinitely many numerical realizations (A, B) sharing the same zero/nonzero pattern.

Example 1.3. The numerical pair

$$A = \begin{pmatrix} 1 & 0 \\ 0.5 & -0.3 \end{pmatrix}, \quad B = \begin{pmatrix} 0.7 \\ 0 \end{pmatrix}$$

is one realization of the pattern

$$\mathcal{A} = \begin{pmatrix} * & 0 \\ * & * \end{pmatrix}, \qquad \mathcal{B} = \begin{pmatrix} * \\ 0 \end{pmatrix}.$$

Any other pair (A', B') whose nonzero entries exactly match the "$*$" positions in $\mathcal{A}$ and $\mathcal{B}$ is also a valid realization of this structure.

In the context of data-driven controllability, the aim is to recover the zero/nonzero pattern of (A, B) directly from (noisy) input–state trajectories. Once this pattern is identified, structural controllability tests can be applied to determine whether the system is controllable, without estimating the exact numerical values. This contrasts with classical system identification, which fits a full-valued model to the data before checking controllability. Structure identification, by focusing solely on the binary pattern, proves significantly more resilient to measurement noise. This thesis demonstrates numerically that, even in the presence of noise, reliable controllability conclusions can be obtained by combining robust structure recovery with structural controllability theory. The goal is to extract the underlying structure from data alone, without requiring additional model assumptions. To the best of the author's knowledge, no comparable noise-tolerant, data-driven approach to controllability has yet been clearly established in the literature.

Necessary and sufficient conditions for structure identification have been established in [36], where it is shown that, given full knowledge of the system's transfer function, the binary pattern of (A, B) can be uniquely reconstructed. This framework was extended in [37], which introduces a graph simplification algorithm yielding equivalent necessary and sufficient conditions via left-invertibility tests on selected transfer submatrices. Further identifiability results, relying on controllability and observability assumptions, appear in [38].

While these contributions provide a solid theoretical foundation, they all require exact, full-frequency transfer-function data or perfect state-space information and assume noise-free measurements, assumptions that are deliberately avoided in this trajectory-based, noisy setting. Moreover, such requirements are often unrealistic in practical systems. In this thesis, the aim is to devise algorithms that recover a structure $(\mathcal{A}, \mathcal{B})$ consistent with the observed data, even when the identification is not unique. Rather than focusing on uniqueness, the approach emphasizes robustness, scalability, and practical viability. The literature on structure identification encompasses several mathematical approaches.

One prominent family of methods is the *node knockout* technique [39, 40]. In this approach, the system is represented as a graph, and each node (or pair of nodes) is "knocked out", for example, by forcing its state to zero. The resulting transfer function of the modified network is then analyzed via the graph Laplacian to infer the original connectivity. While this technique targets the binary pattern directly, it crucially assumes that the network is both controllable and observable, and is formulated for undirected graphs, whereas the focus here is on directed networks and on assessing controllability itself. Moreover, node knockout methods are inherently intrusive, as they require the ability to intervene on the system in specific ways, which may not be feasible or safe in many engineering contexts.

Alternative approaches employ *constrained Lyapunov equations* [41], which recover structure in diffusion-only systems (i.e. systems without inputs) by fitting Lyapunov matrices to steady-state covariances. Spectral methods have also been proposed: reconstruction of binary links using cross-spectral densities and causal Wiener filtering is developed in [42], while a power-spectral-density-based algorithm for binary reconstruction in directed networks is presented in [43]. These methods offer noise resilience in specific settings but either neglect inputs or rely on statistical averages, rather than directly exploiting time-domain trajectory data as considered here.

In contrast, the proposed framework operates directly on finite, noisy trajectory data, without relying on transfer-function information. This perspective naturally leads to convex optimization methods for structure recovery. A first idea, proposed in [44], is to test each candidate zero/nonzero pattern by enforcing its zeros and checking compatibility with the data. Although this exhaustive search guarantees identification of all structures consistent with the data, the number of possible patterns grows exponentially with system dimension, making it computationally infeasible even for moderately sized networks. To overcome this combinatorial explosion, modern approaches incorporate regularization techniques or leverage structural priors.

Another perspective casts structure identification as a *sparse* optimization problem: the sparsest pattern that minimizes the discrepancy between observed trajectories and the estimated system is sought. This formulation is studied in [45], which analyzes LASSO-based recovery and shows that the number of samples required for accurate identification grows only logarithmically with the system dimension. This is a highly encouraging result for scalability. The developments in this thesis build on these insights and further explore sparsity-promoting regularization techniques.

In addition, *hard-thresholding* methods are investigated, where small entries are iteratively zeroed out to enforce sparsity; such algorithms have been analyzed in [46]. By analyzing gradient-based estimators, LASSO-type penalties, and hard thresholding schemes, the goal is to recover the interaction pattern among system components using only trajectory data, even in the presence of measurement noise or unmodeled dynamics. In particular, the influence of key parameters, such as the number of observations, the noise level, and the system dimension, on the accuracy and reliability of the recovered structure is quantified.

1.3 From Network Structure to Controllability

Once the zero/nonzero pattern $(\mathcal{A}, \mathcal{B})$ of a system is identified, the next step is to determine whether this pattern alone guarantees controllability. Classical tests such as the Hautus criterion require numerical matrices and cannot be applied directly to binary patterns. Structural controllability theory addresses this limitation by providing graph-theoretic and algebraic tests that operate solely on the pattern matrices. These tests make it possible to assess controllability without requiring knowledge of the exact system dynamics.

In structural controllability, the pair $(\mathcal{A}, \mathcal{B})$ is represented by a directed graph whose vertices correspond to state and input nodes, and whose edges correspond to "$*$" entries in $\mathcal{A}$ or $\mathcal{B}$. A path in this graph indicates an influence from one node to another. The key problem is to characterize the conditions under which such a graph guarantees that all (or at least one) numerical realizations (A, B) consistent with the pattern are controllable.

Two notions emerge. A pattern is *strongly structurally controllable* if *every* numerical realization of $(\mathcal{A}, \mathcal{B})$ is controllable. In contrast, it is *weakly structurally controllable* if *at least one* realization is controllable. Although these definitions may seem opposed, weak structural controllability is a generic property: if a pattern is weakly controllable, then almost every realization is controllable.

Weak structural controllability was introduced by Lin [47] and further developed by Dion *and al.* [48]. It is shown that a pattern is weakly controllable if and only if its graph is spanned by a collection of disjoint cacti, a simple graph structure composed of cycles and paths, and if every state node is reachable from at least one input node. While conceptually insightful, these conditions do not yield efficient algorithms. A constructive criterion was provided in [49], which shows that weak controllability holds if and only if the bipartite graph between state and input nodes admits a perfect matching under a suitable input placement. This result was extended in [50] to the general LTI setting, leading to a polynomial-time test.

Strong structural controllability imposes a more stringent requirement. Matching-based conditions were first derived in [50] and independently in [51], while a complementary condition based on graph coloring appears in [51] and [52]. In these formulations, colors are assigned to nodes according to propagation rules that ensure every pattern edge can support controllability for all realizations. More recently, Jia *and al.* [53] introduced an extended notion of structure that allows for a third type of entry: one that may be either zero or nonzero depending on the parameter realization. This is formalized by pattern matrices $\mathcal{A} \in \{0, *, ?\}^{n\times n}$ and $\mathcal{B} \in \{0, *, ?\}^{n\times m}$, where "0" indicates a fixed zero, "$*$" a fixed nonzero value, and "?" an indeterminate entry that can be zero or nonzero. The presence of "?" entries introduces ambiguity in the interconnection structure, and can be interpreted as encoding two possible binary structures: one where the "?" is set to zero, and one where it is set to nonzero. Any test for strong structural controllability must therefore verify controllability under both configurations.

Example 1.4. Consider the pattern matrices

$$\mathcal{A} = \begin{pmatrix} * & 0 \\ ? & * \end{pmatrix}, \quad \mathcal{B} = \begin{pmatrix} * \\ 0 \end{pmatrix}.$$

The uncertain entry "?" at position $(2, 1)$ represents a possible connection that may or may not exist. This pattern encodes two standard binary structures:

$$\mathcal{A}' = \begin{pmatrix} * & 0 \\ 0 & * \end{pmatrix}, \qquad \mathcal{A}'' = \begin{pmatrix} * & 0 \\ * & * \end{pmatrix}.$$

To conclude strong structural controllability from this pattern, both $(\mathcal{A}', \mathcal{B})$ and $(\mathcal{A}'', \mathcal{B})$ must be controllable.

This extension captures uncertainty in the pattern itself, an important feature when recovering structure from noisy data. Note that classical binary structure theory is recovered as a special case when no "?" entries are present.

When only the binary pattern (or its uncertain extension with "?" entries) is available, structural controllability theory offers a powerful framework for assessing whether a system can be controlled, irrespective of the exact parameter values. This thesis develops efficient algorithms based on graph matchings and color propagation rules to verify both weak and strong structural controllability. The computational complexity of these algorithms is analyzed, along with their robustness to structural uncertainty. The controllability properties of given structures are also systematically investigated, including a comparison of the weak and strong notions. In particular, the analysis examines how these properties vary with the sparsity of the underlying system and identifies structural regimes where controllability is guaranteed under either criterion.

1.4 Outline and Relations Between Chapters

This thesis addresses the challenge of determining controllability from data in the presence of noise, a setting of growing importance in large-scale, complex, and partially observed systems. As noted in recent literature, many data-driven controllability tests, though theoretically elegant, are highly sensitive to even small levels of measurement noise and often produce false positives, incorrectly declaring a system controllable when it is not. Consequently, relying solely on numerical conditions extracted from noisy data can compromise the reliability of such assessments. This work investigates whether controllability can be reliably inferred from noisy trajectory data.

To approach this problem, the analysis adopts a complementary paradigm: *structural controllability*, which bases controllability conclusions on the structure of the system rather than its precise numerical values. Structural methods are naturally more robust to noise, as they require only information about zero/nonzero patterns. The central idea is to integrate both paradigms by recovering structural information from noisy data and applying structural controllability theory to enable robust conclusions.

This conceptual flow is illustrated in Figure 1.1, which summarizes three main paradigms of controllability. In the classical setting, controllability is inferred by first identifying a model and then applying algebraic tests such as the Kalman rank condition. In the data-driven approach, assessments are based directly on observed trajectories, using informativity-based or rank-based conditions without model identification. The approach proposed in this thesis follows a hybrid path: structure is recovered from data and subsequently used to assess controllability via structural methods (highlighted in blue in the figure).

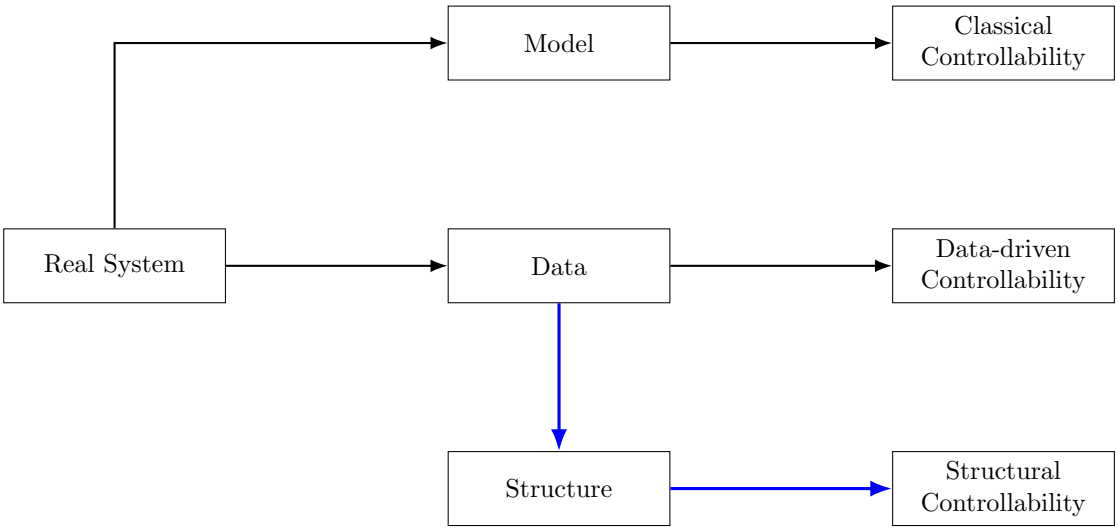

Figure 1.1: Paradigms of controllability: classical, data-driven, and structural.

Throughout this thesis, three notions of controllability are considered: classical, data-driven, and structural. Although closely related, these notions are not equivalent. Classical controllability depends on specific numerical realizations and must be reassessed whenever the system changes. Data-driven controllability, as defined by the informativity framework, considers all systems consistent with the observed data but can only confirm controllability; it cannot detect uncontrollability, which limits its reliability in the presence of noise. Structural controllability, by contrast, applies to all realizations matching a given structure: if the structure is strongly controllable, then every realization is controllable. However, this also implies that if even one realization is not controllable, the structure fails the strong test. Weak structural controllability relaxes this requirement but has its own limitations, particularly in ambiguous configurations.

The structure of this thesis reflects the conceptual progression from classical methods to hybrid data-and-structure-based approaches.

Chapter 1 motivates the need to move beyond traditional model-based control, highlighting the limitations of relying exclusively on accurate numerical models in large-scale, noisy, or partially known systems. It introduces the three paradigms of controllability (classical, data-driven, and structural) and reviews the relevant literature, including data-driven analysis, structural controllability theory, and network structure identification.

Chapter 2 presents the data-driven approach in detail through the informativity framework. It formalizes the notion of controllability from data, illustrates it with examples, and examines how measurement noise degrades the reliability of informativity-based tests.

Chapter 3 focuses on structural controllability. It introduces algorithmic tools to test both weak and strong structural controllability using graph-theoretic methods such as matchings and color propagation rules. The chapter also analyzes the computational complexity of these algorithms and explores how the sparsity of a system influences its controllability properties.

Chapter 4 addresses the problem of recovering system structure from data. It introduces algorithms that extract the zero/nonzero pattern of the system matrices from noisy input–state trajectories. Methods based on sparse optimization (such as LASSO and hard-thresholding) are discussed, along with theoretical results on sample complexity and noise robustness.

Finally, Chapter 5 integrates the techniques developed in the preceding chapters into a unified pipeline. Starting from noisy data, the system structure is recovered and controllability is assessed using structural theory. This chapter compares the proposed framework with classical and data-driven methods and applies it to various network examples to demonstrate its practical value and robustness.

Taken together, the thesis contributes a coherent methodology for inferring controllability from noisy data by integrating data-driven and structural perspectives into a robust and scalable framework.

Chapter 2

Data-Driven Control

Data-driven control refers to the class of methods that bypass explicit model identification and instead use measured data directly to analyze system properties or synthesize controllers. This approach is particularly appealing when accurate modeling is infeasible due to system complexity or partial observability. Among recent developments, one line of work stands out for its generality and rigor: the *data informativity framework*, introduced by van Waarde *and al.* [34], which recasts system-theoretic analysis as a property verification problem based entirely on trajectory data.

2.1 Data Informativity Framework

The following introduces the formalism of the data informativity framework developed in [34]. Rather than identifying a single system, this approach considers the entire set of systems that are consistent with the observed data. The central question is whether every such system satisfies a given property, such as controllability or stabilizability. If this is the case, the data are said to be informative for that property. This perspective enables robust verification based solely on measurements and forms the theoretical foundation for the results presented in this chapter.

2.1.1 System Class and Data-Consistent Set

Let Σ denote a general system class, such as linear time-invariant (LTI), linear time-varying (LTV), or nonlinear models. It is assumed that the true but unknown system $\mathcal{S} \in \Sigma$ belongs to this class. Let $\mathcal{D}$ denote a finite collection of observed data generated by $\mathcal{S}$. This data may include, for instance, full or partial measurements of inputs, states, or outputs, collected over one or several finite time intervals. The nature of the data depends on the system and the measurement setup.

The following defines the set of systems that are consistent with the available data.

Definition 2.1 (Data-Consistent Set)**.** Given the data $\mathcal{D}$, the set $\Sigma_{\mathcal{D}} \subseteq \Sigma$ is defined as the collection of all systems in Σ that are consistent with the data $\mathcal{D}$, that is, that could also have generated these data.

In other words, $\Sigma_{\mathcal{D}}$ contains every system that agrees with the measurements, according to the behavior specified by $\mathcal{D}$. Since only the data $\mathcal{D}$ are available, the true system $\mathcal{S}$ cannot be distinguished from any other element of $\Sigma_{\mathcal{D}}$ based on the data alone.

2.1.2 Data Informativity for Analysis

The problem of verifying whether a system-theoretic property holds, based solely on observed data, is now considered. Typical properties of interest include stability, controllability, observability, or dissipativity. Since the true system $\mathcal{S} \in \Sigma$ is unknown, only the data set $\mathcal{D}$ can be used to infer whether such a property is satisfied.

The following defines the set of systems within Σ that possess a given property.

Definition 2.2 (Property Set)**.** Let $\mathcal{P}$ be a system-theoretic property. The *property set* $\Sigma_{\mathcal{P}} \subseteq \Sigma$ is defined as the collection of all systems in Σ for which property $\mathcal{P}$ holds.

The objective is to determine whether the true system $\mathcal{S}$ belongs to $\Sigma_{\mathcal{P}}$. Since $\mathcal{S}$ cannot be distinguished from any element of $\Sigma_{\mathcal{D}}$, the only way to conclude that $\mathcal{S}$ satisfies the property is to verify that *every* system consistent with the data also satisfies it. This leads to the notion of informativity.

Definition 2.3 (Informativity for Analysis)**.** The data set $\mathcal{D}$ is said to be *informative* for the property $\mathcal{P}$ if

$$\Sigma_{\mathcal{D}} \subseteq \Sigma_{\mathcal{P}}.$$

That is, all systems consistent with the observed data possess the property $\mathcal{P}$.

This notion does not require exact identification. Even when multiple systems are compatible with the data, informativity holds as long as all such systems share the desired property. This situation is illustrated in Figure 2.1, where the data-consistent set lies entirely within the property set, thereby ensuring that the property holds for the true system. In the special case where $\Sigma_{\mathcal{D}} = \{\mathcal{S}\}$, informativity reduces to model identification followed by property verification.

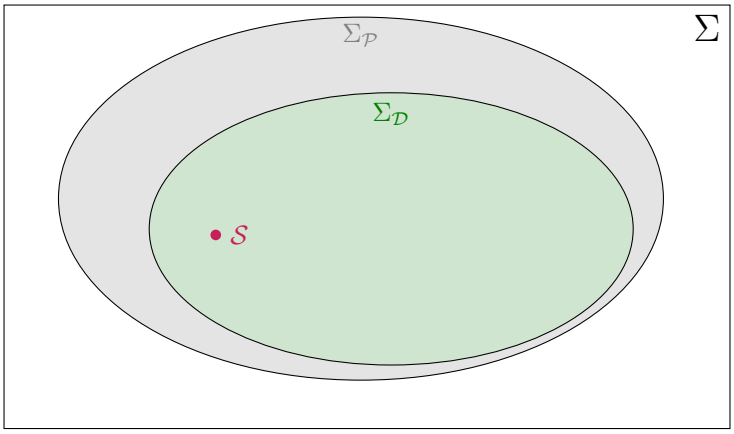

Σ : model class
$\mathcal{S}$: unknown system
$\mathcal{D}$: given data set
$\Sigma_{\mathcal{D}}$: data consistent systems
$\mathcal{P}$: system property
$\Sigma_{\mathcal{P}}$: systems with property $\mathcal{P}$

Figure 2.1: The data are informative for property $\mathcal{P}$ as $\Sigma_{\mathcal{D}} \subseteq \Sigma_{\mathcal{P}}$.

In general, however, it may occur that some systems in $\Sigma_{\mathcal{D}}$ satisfy $\mathcal{P}$ while others do not. In such cases, no definitive conclusion can be drawn from the data alone, since $\mathcal{D}$ is compatible with both systems that possess and systems that lack the property. This is illustrated in Figure 2.2, where $\Sigma_{\mathcal{D}}$ intersects but is not entirely contained in $\Sigma_{\mathcal{P}}$.

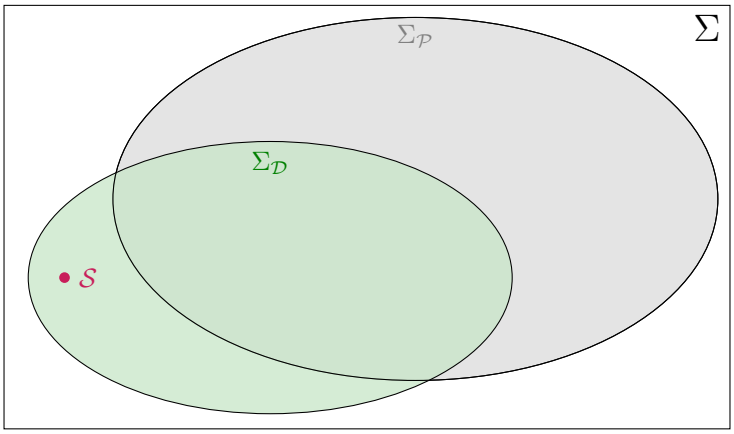

Σ : model class
$\mathcal{S}$: unknown system
$\mathcal{D}$: given data set
$\Sigma_{\mathcal{D}}$: data consistent systems
$\mathcal{P}$: system property
$\Sigma_{\mathcal{P}}$: systems with property $\mathcal{P}$

Figure 2.2: The data are not informative for property $\mathcal{P}$.

2.2 Data-Driven Analysis

To develop a concrete understanding of the data informativity framework, the analysis is restricted to the class of discrete-time linear time-invariant (LTI) systems with full input and state measurements. This setting serves as a foundational model for studying data-based controllability and enables explicit characterizations of the data-consistent set. Let the true system $\mathcal{S} = (A_s, B_s) \in \Sigma$ evolve according to

$$x(t+1) = A_s x(t) + B_s u(t), \tag{2.1}$$

where $x(t) \in \mathbb{R}^n$ is the state at time t and $u(t) \in \mathbb{R}^m$ the control input at time t. The dimensions n and m are assumed known, but the matrices $A_s \in \mathbb{R}^{n\times n}$ and $B_s \in \mathbb{R}^{n\times m}$ are unknown. The associated model class is given by

$$\Sigma := \left\{ (A,B) \in \mathbb{R}^{n\times n} \times \mathbb{R}^{n\times m} \;\middle|\; x(t+1) = Ax(t) + Bu(t) \right\}, \tag{2.2}$$

that is, the set of all discrete-time LTI input–state systems of compatible size.

2.2.1 Data Collection and Consistency

The state trajectory is generated by exciting the true system $\mathcal{S}$ over a finite horizon $t = 0, 1, \ldots, T$ with the input sequence

$$u(0),\ u(1),\ \ldots,\ u(T-1),$$

and recording the resulting states for a given initial condition:

$$x(0),\ x(1),\ \ldots,\ x(T).$$

The data are organized into block matrices as follows:

$$U_- := \begin{bmatrix} u(0) & u(1) & \cdots & u(T-1) \end{bmatrix} \in \mathbb{R}^{m\times T},$$

$$X := \begin{bmatrix} x(0) & x(1) & \cdots & x(T) \end{bmatrix} \in \mathbb{R}^{n\times (T+1)}.$$

The matrix X is further partitioned as

$$X_- \;=\; \begin{bmatrix} x(0) & x(1) & \cdots & x(T-1) \end{bmatrix} \;\in\; \mathbb{R}^{n\times T},$$

$$X_+ \;=\; \begin{bmatrix} x(1) & x(2) & \cdots & x(T) \end{bmatrix} \;\in\; \mathbb{R}^{n\times T},$$

so that the dynamics (2.1) take the compact form

$$X_+ = A_s X_- + B_s U_-. \tag{2.3}$$

Stacking input and state data yields the block matrix equation

$$X_+ = [A_s \;\; B_s] \begin{bmatrix} X_- \\ U_- \end{bmatrix}, \tag{2.4}$$

where $[A_s \;\; B_s] \in \mathbb{R}^{n\times(n+m)}$ is unknown. The full dataset is denoted $\mathcal{D} = (U_-, X)$, or equivalently (U_-, X_-, X_+) when needed. The set of systems consistent with the data is

$$\Sigma_{\mathcal{D}} := \{(A, B) \in \Sigma \mid X_+ = AX_- + BU_-\}.$$

This construction accommodates both long single trajectories and collections of shorter experiments by concatenating the associated matrices U_-, X_-, and X_+. The following example illustrates a data-consistent set in the LTI case.

Example 2.4. Let $n = 2$, $m = 1$, and consider the system $\mathcal{S} = (A_s, B_s)$ defined by

$$A_s = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, \qquad B_s = \begin{bmatrix} 1 \\ 0 \end{bmatrix}.$$

The system is excited over a time horizon $T = 2$, and the observed input–state trajectory is

$$U_- = \begin{bmatrix} 0 & 1 \end{bmatrix}, \quad X_- = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}, \quad X_+ = \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix}.$$

The corresponding data set is

$$\mathcal{D} = (U_-, X_-, X_+).$$

The data-consistent set is then given by

$$\Sigma_{\mathcal{D}} = \left\{ (A, B) \in \mathbb{R}^{2\times2} \times \mathbb{R}^{2\times1} \;\middle|\; X_+ = [A \;\; B] \begin{bmatrix} X_- \\ U_- \end{bmatrix} \right\},$$

which, in this case, becomes

$$\Sigma_{\mathcal{D}} = \left\{ (A, B) \in \mathbb{R}^{2\times2} \times \mathbb{R}^{2\times1} \;\middle|\; \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix} = \begin{bmatrix} A & B \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 0 & 1 \end{bmatrix} \right\}.$$

By construction, $\Sigma_{\mathcal{D}}$ contains precisely those systems that could have generated the observed data. Thus, it captures all information available for any subsequent informativity-based analysis or controller synthesis.

2.2.2 Informativity for System Identification

A natural question is whether the collected data $\mathcal{D}$ determine the underlying system uniquely. If this is the case, then the system matrices (A_s, B_s) can be recovered directly from the data, and standard model-based analysis or controller design techniques become applicable. This leads to the following definition.

Definition 2.5 (Informativity for Identification)**.** The data set $\mathcal{D} = (U_-, X)$ is said to be *informative for system identification* if

$$\Sigma_{\mathcal{D}} = \{(A_s, B_s)\},$$

i.e., if the data-consistent set is a singleton containing the true system.

Since the true model (A_s, B_s) always belongs to $\Sigma_{\mathcal{D}}$ by construction, informativity for identification amounts to asking whether it is the *only* model that fits the data. This depends entirely on the rank of the combined matrix of state and input trajectories.

Theorem 2.6 (Informativity for Identification [34])**.** *The data $\mathcal{D} = (U_-, X)$ are informative for identification if and only if*

$$\operatorname{rank}\begin{bmatrix} X_- \\ U_- \end{bmatrix} = n + m.$$

In that case, the unique model (A_s, B_s) can be recovered from the data via

$$[A_s \;\; B_s] = X_+ \begin{bmatrix} X_- \\ U_- \end{bmatrix}^{\dagger},$$

where † denotes the Moore–Penrose pseudoinverse.

This result shows that exact identification is possible whenever the data matrix has full row rank. However, the pseudoinverse is not unique when $T > n + m$, so the expression above yields the minimum-norm solution in that case.

2.2.3 Informativity for Controllability

The notion of informativity can be specialized to the property of controllability. In this case, the data are said to be informative for controllability if every system consistent with the data satisfies this property. This leads to the following definition.

Definition 2.7 (Informativity for Controllability)**.** Let $\Sigma_{\text{cont}} \subseteq \Sigma$ denote the set of all controllable systems in the model class Σ, i.e.,

$$\Sigma_{\text{cont}} = \{(A, B) \in \Sigma \mid (A, B) \text{ is controllable}\}.$$

The data set $\mathcal{D} = (U_-, X)$ is said to be *informative for controllability* if

$$\Sigma_{\mathcal{D}} \subseteq \Sigma_{\text{cont}}.$$

In other words, all systems compatible with the data must be controllable. A key result from [34] provides a simple, testable condition that guarantees this.

Theorem 2.8 (Data-Driven Hautus Test for Controllability [34])**.** *The data $\mathcal{D} = (U_-, X)$ are informative for controllability if and only if*

$$\operatorname{rank}\left(X_+ - \lambda X_-\right) = n \quad \forall \lambda \in \mathbb{C}.$$

This test mirrors the classical Hautus criterion [54], but applies directly to measured data and requires no model identification. Crucially, the condition may hold even when the true system is not uniquely identifiable. In such cases, the data may still reveal that all compatible models are controllable, as illustrated in the following example.

Example 2.9. Let $n = 2$, $m = 1$, $T = 2$, and consider the observed input–state data

$$X = \begin{bmatrix} 0 & 1 & 1 \\ 0 & 1 & -1 \end{bmatrix}, \qquad U_- = \begin{bmatrix} 1 & 1 \end{bmatrix}.$$

Then

$$X_- = \begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix}, \qquad X_+ = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}.$$

To apply the data-driven Hautus test, consider

$$X_+ - \lambda X_- = \begin{bmatrix} 1 & 1-\lambda \\ 1 & -1-\lambda \end{bmatrix} \quad \text{for all } \lambda \in \mathbb{C}.$$

This matrix has full rank for all $\lambda \in \mathbb{C}$ since its determinant equals $-2 \neq 0$, so the data are informative for controllability. However, they are not informative for system identification. Indeed, the combined matrix

$$\begin{bmatrix} X_- \\ U_- \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 0 & 1 \\ 1 & 1 \end{bmatrix}$$

has rank $2 < 3 = n + m$. Thus, multiple systems are consistent with the data.

One can show that the consistent set is

$$\Sigma_{\mathcal{D}} = \left\{ (A, B) \;\middle|\; B = \begin{bmatrix} 1 \\ 1 \end{bmatrix},\ A \begin{bmatrix} 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ -2 \end{bmatrix} \right\},$$

i.e., all (A, B) such that $a_{11} + a_{12} = 0$ and $a_{21} + a_{22} = -2$. Each such pair yields a controllability matrix of full rank:

$$\begin{bmatrix} B & AB \end{bmatrix} = \begin{bmatrix} 1 & \times \\ 1 & \times \end{bmatrix}$$

has rank 2 regardless of the precise entries of A. Hence, even without identifying the true system, the data allow us to conclude that all compatible models are controllable.

2.2.4 Limits of Certifying Uncontrollability

Although the data-driven Hautus test offers a direct criterion to certify controllability from data, it provides no mechanism to certify *uncontrollability*. If the test condition fails, it cannot be concluded that the underlying system is uncontrollable; only that the data are not informative for controllability.

This raises the question of whether the failure of informativity merely reflects insufficient data or indicates a more fundamental ambiguity: namely, whether both controllable and uncontrollable systems can explain the same trajectory. In such cases, even perfect, noise-free measurements do not resolve the system's controllability status. The following example illustrates this limitation.

Example 2.10. Consider two systems in Σ:

$$A_{\text{cont}} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad B_{\text{cont}} = \begin{bmatrix} 1 \\ 0 \end{bmatrix},$$

which is controllable, and

$$A_{\text{uncont}} = \begin{bmatrix} -2 & 2 \\ -1 & 1 \end{bmatrix}, \quad B_{\text{uncont}} = \begin{bmatrix} 2 \\ 1 \end{bmatrix},$$

which is not controllable. Suppose the collected data are:

$$U_- = \begin{bmatrix} 2 & 3 \end{bmatrix}, \quad X_- = \begin{bmatrix} 1 & 2 \\ 0 & 1 \end{bmatrix}, \quad X_+ = \begin{bmatrix} 2 & 4 \\ 1 & 2 \end{bmatrix}.$$

A direct computation shows that both systems satisfy

$$X_+ = AX_- + BU_-.$$

Hence, both $(A_{\text{cont}}, B_{\text{cont}})$ and $(A_{\text{uncont}}, B_{\text{uncont}})$ belong to $\Sigma_{\mathcal{D}}$. Since one is controllable and the other is not, it follows that $\Sigma_{\mathcal{D}} \nsubseteq \Sigma_{\text{cont}}$, and thus $\mathcal{D}$ is not informative for controllability. Importantly, this is not a trivial case: the input U_- is nonzero, and the observed data reflect valid and nondegenerate system behavior. Yet, despite the quality of the measurements, the underlying system could be either controllable or uncontrollable.

This example highlights an inherent asymmetry in the data-informativity framework: while it can certify that all systems consistent with the data are controllable, it cannot exclude the possibility of controllability. Consequently, informativity for controllability constitutes a *one-sided* test: failure to certify does not imply failure of the property.

2.3 Empirical Analysis

This section investigates the behavior of the data-driven Hautus test for controllability through numerical experiments. The analysis begins in the idealized setting of noise-free data, where expected outcomes are verified for both controllable and uncontrollable systems. The focus then shifts to the impact of perturbations, examining how measurement noise affects the test's reliability. This allows an evaluation of its robustness in practical scenarios, where exact data are rarely available. In addition, the effect of specific noise directions, such as those aligned with controllable or uncontrollable modes, is considered.

2.3.1 Computational Considerations

Although the data-driven Hautus condition provides a complete theoretical test, its implementation involves verifying a rank condition over the complex field. However, similar to the classical Hautus test where controllability is verified by checking a rank condition at the eigenvalues of the system matrix, this can be reduced to checking finitely many critical values of λ derived from the data [34]. In particular, the test is equivalent to the conditions:

$$\operatorname{rank}(X_+) = n, \tag{2.5}$$

and

$$\operatorname{rank}(X_+ - \lambda X_-) = n \quad \forall \lambda \neq 0 \text{ such that } \lambda^{-1} \in \sigma(X_- X_+^\dagger), \tag{2.6}$$

where $\sigma(M)$ denotes the spectrum of a matrix M, that is, the set of its eigenvalues, and $X_+^\dagger$ is any right inverse of X_+. This reformulation enables practical computation by evaluating a finite, data-dependent set of cases.

2.3.2 Singular Values as a Robust Alternative to Rank

Despite these simplifications, directly using matrix rank remains numerically fragile. The rank function is inherently discontinuous with respect to matrix entries, making it highly sensitive to small perturbations, such as measurement noise. This sensitivity can lead to unreliable conclusions when testing controllability from empirical data. The issue is illustrated by the following example.

Example 2.11. Let

$$A_\varepsilon = \begin{bmatrix} 1 & 1 \\ 1 & 1+\varepsilon \end{bmatrix}, \quad \varepsilon \in \mathbb{R}.$$

When $\varepsilon = 0$, the matrix has rank 1. For any $\varepsilon \neq 0$, its rank is 2. Hence, the rank changes discontinuously as $\varepsilon \to 0$, despite arbitrarily small perturbations.

This discontinuity limits the usefulness of rank as a diagnostic tool in the presence of noise or modeling uncertainty. A more robust alternative is provided by singular values, which vary continuously with matrix entries and quantify proximity to rank deficiency in a graded, numerically stable manner.

Definition 2.12 (Singular Values [55])**.** Let $A \in \mathbb{R}^{n\times m}$. The *singular values* of A are the nonnegative square roots of the eigenvalues of $A^\top A$, ordered as

$$\sigma_1(A) \geq \sigma_2(A) \geq \cdots \geq \sigma_{\min(n,m)}(A) \geq 0.$$

The largest and smallest singular values are denoted by $\sigma_{\max}(A) := \sigma_1(A)$ and $\sigma_{\min}(A) := \sigma_{\min(n,m)}(A)$, respectively.

Singular values are well-defined for all real matrices, regardless of shape. They provide a continuous measure of proximity to rank deficiency and are therefore more stable in data-driven settings. In the matrix A_ε, for example, the smallest singular value $\sigma_{\min}(A_\varepsilon)$ increases smoothly from zero, even though the rank jumps abruptly. This behavior offers a sensitive indicator of near-rank-deficiency and avoids the discontinuous nature of rank.

Lemma 2.13 (Singular Value Characterization of Rank [55])**.** *The number of nonzero singular values of a matrix A equals its rank.*

Since the data-driven Hautus test requires X_+ and $X_+ - \lambda X_-$ to be full row rank, this is equivalent to all singular values being strictly positive. In practice, one may verify that

$$\sigma_{\min}(X_+) > 0 \quad \text{and} \quad \sigma_{\min}(X_+ - \lambda X_-) > 0 \quad \forall \lambda \neq 0 \text{ such that } \lambda^{-1} \in \sigma(X_- X_+^\dagger),$$

thus obtaining a numerically stable certification of controllability.

If any of these minimal singular values approach zero, the associated matrix is nearly rank-deficient, indicating that the system lies close to the boundary between controllability and uncontrollability. In such cases, singular values enable a continuous and reliable assessment, unlike the abrupt and potentially misleading conclusions drawn from rank.

This interpretation aligns with the geometric formulation proposed by Eising [56], who defines the distance from a controllable pair (A, B) to the set of uncontrollable systems as

$$\inf_{\lambda \in \mathbb{C}} \sigma_{\min}\Big([\lambda I - A \;\; B]\Big).$$

This metric quantifies the robustness of the classical Hautus condition. In parallel, the singular values of X_+ and $X_+ - \lambda X_-$ can be viewed as data-driven analogues, reflecting the resilience of the system to uncertainty in the available measurements.

2.3.3 Noiseless Case

The empirical investigation begins in a noise-free setting to illustrate how the data-driven Hautus test behaves under ideal conditions. Although theoretical guarantees ensure the validity of the test, numerical examples help reveal its behavior near the boundary between controllability and uncontrollability.

Consider the system

$$A_s = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, \quad B_s = \begin{bmatrix} 1 \\ 0 \end{bmatrix}. \tag{2.7}$$

It can be easily checked that this system is not controllable, as the second state evolves independently of the input. Using trajectory data generated over a horizon $T = 100$, the smallest singular values involved in the data-driven Hautus test are computed:

$$\sigma_{\min}(X_+) = 8.82, \quad \sigma_{\min}(X_+ - \lambda X_-) = 0.$$

The vanishing singular value of $X_+ - \lambda X_-$ confirms that this matrix is rank-deficient, indicating that the data are not informative for controllability.

To examine the sensitivity of the test near the boundary between controllability and uncontrollability, a parameter $\alpha \in \mathbb{R}$ is introduced in the second entry of the input matrix. This defines a parameterized family of systems that includes both controllable and uncontrollable cases. The modified system is given by:

$$A_s = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, \quad B_s = \begin{bmatrix} 1 \\ \alpha \end{bmatrix}. \tag{2.8}$$

When $\alpha = 0$, the system coincides with the previously considered uncontrollable case, where the second state is unaffected by the input. For any $\alpha \neq 0$, controllability is restored, as both states become influenced by the control input. By continuously varying α, the smallest singular values of X_+ and $X_+ - \lambda X_-$ reveal how the system approaches the controllability boundary. This behavior is illustrated in Figure 2.3.

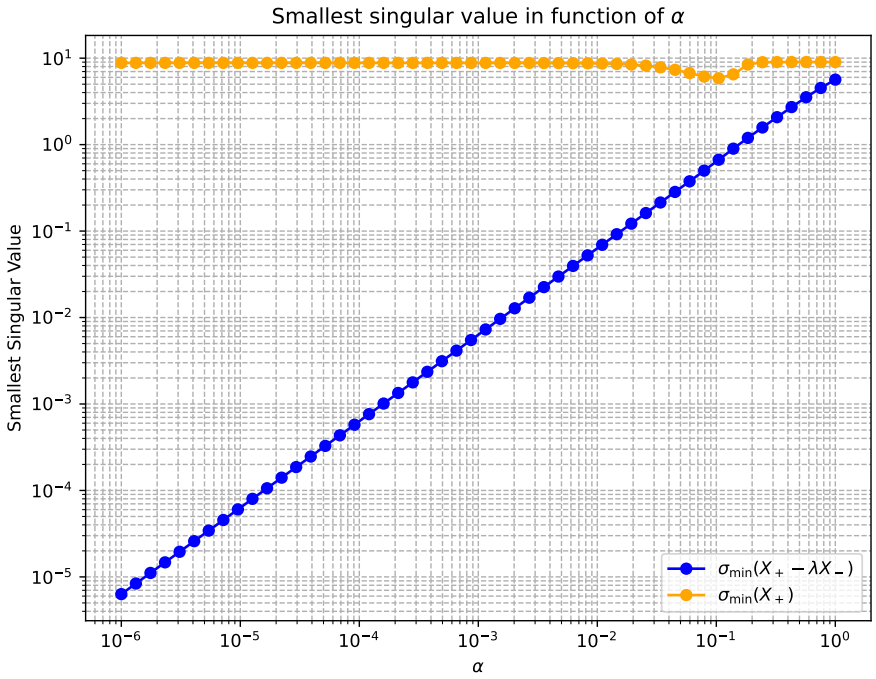

Figure 2.3: Evolution of the smallest singular values of X_+ and $X_+ - \lambda X_-$ as a function of α.

As shown in Figure 2.3, the smallest singular value $\sigma_{\min}(X_+)$ remains strictly positive and nearly constant for all values of α, including when the system is uncontrollable. This confirms that the rank of X_+ alone does not capture the loss of controllability.

In contrast, $\sigma_{\min}(X_+ - \lambda X_-)$ decreases smoothly as $\alpha \to 0$, vanishing precisely when the system becomes uncontrollable. This smooth variation confirms the usefulness of singular values in capturing how close a system lies to the boundary of controllability. It also demonstrates that the matrix $X_+ - \lambda X_-$ carries the structural information essential for assessing controllability, an aspect that becomes particularly valuable when dealing with noisy or incomplete measurements.

2.3.4 Noisy Case

The data-driven Hautus test is now examined in the presence of measurement noise. For this purpose, the system defined in (2.7) is reconsidered, with additive white Gaussian noise affecting the state trajectories. The true system dynamics become:

$$x(t+1) = A_s\, x(t) + B_s\, u(t) + w(t), \tag{2.9}$$

where

$$w(t) \sim \mathcal{N}(0, \Sigma_w^2 I_2),$$

with $\Sigma_w > 0$ denoting the standard deviation and $I_2 \in \mathbb{R}^{2\times 2}$ the identity matrix. Figure 2.4 illustrates the evolution of the smallest singular values of X_+ and $X_+ - \lambda X_-$ as the noise level increases, that is, for varying values of Σ_w.

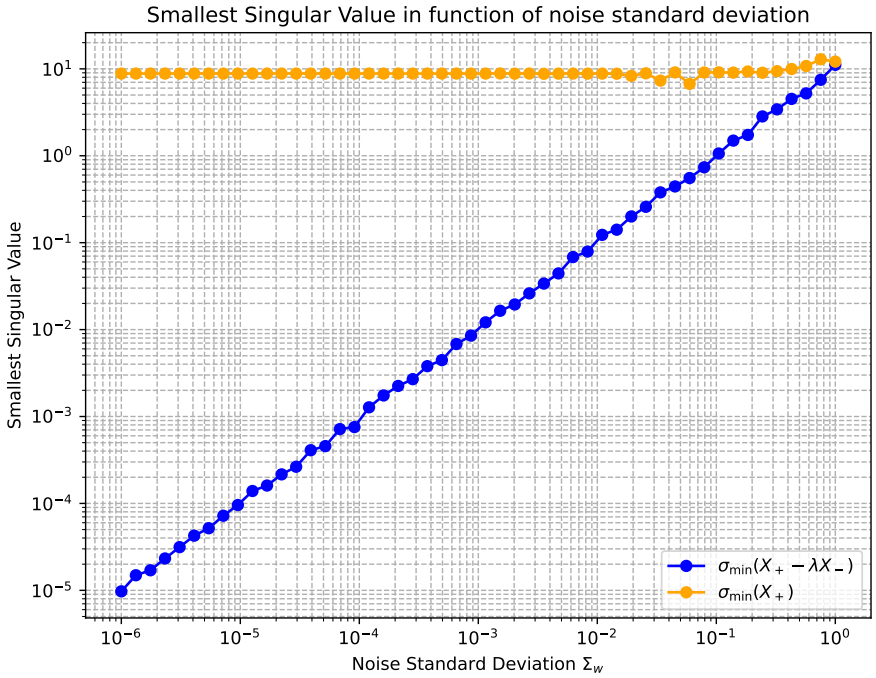

Figure 2.4: Evolution of the smallest singular values of X_+ and $X_+ - \lambda X_-$ as a function of the noise standard deviation Σ_w.

As shown in Figure 2.4, even when the true system is uncontrollable, adding noise perturbs the data matrices in a way that causes their smallest singular values to become strictly positive. This change artificially increases their apparent rank, as the noise randomizes the matrix columns and makes them appear linearly independent. As a result, the data-driven Hautus test incorrectly suggests controllability.

This behavior highlights a key limitation of the test: in the presence of noise, it may produce false positives, particularly when the system lies near the boundary of controllability. Nonetheless, the smallest singular value of $X_+ - \lambda X_-$ exhibits an almost linear dependence on the noise level, suggesting that it remains a meaningful and continuous indicator of proximity to uncontrollability. These observations raise the question of whether the location of the noise within the system also affects the test outcome.

To investigate this effect, the structure of the system is leveraged, as it contains both a controllable and an uncontrollable component:

$$A_s = \begin{bmatrix} A_{\text{contr}} \\ \hline A_{\text{uncontr}} \end{bmatrix} = \left[\begin{array}{cc} 0 & 0 \\ \hline 0 & 1 \end{array}\right], \quad B_s = \begin{bmatrix} B_{\text{contr}} \\ \hline B_{\text{uncontr}} \end{bmatrix} = \left[\begin{array}{c} 1 \\ \hline 0 \end{array}\right]. \tag{2.10}$$

Here, the first state is directly actuated and thus controllable, while the second evolves independently and is uncontrollable. To assess how noise location affects the test outcome, singular values are computed under two scenarios: when noise is selectively added either to the controllable or to the uncontrollable component. Figure 2.5 and Figure 2.6 illustrate the evolution of the smallest singular values of X_+ and $X_+ - \lambda X_-$ as the noise standard deviation Σ_w increases in each case.

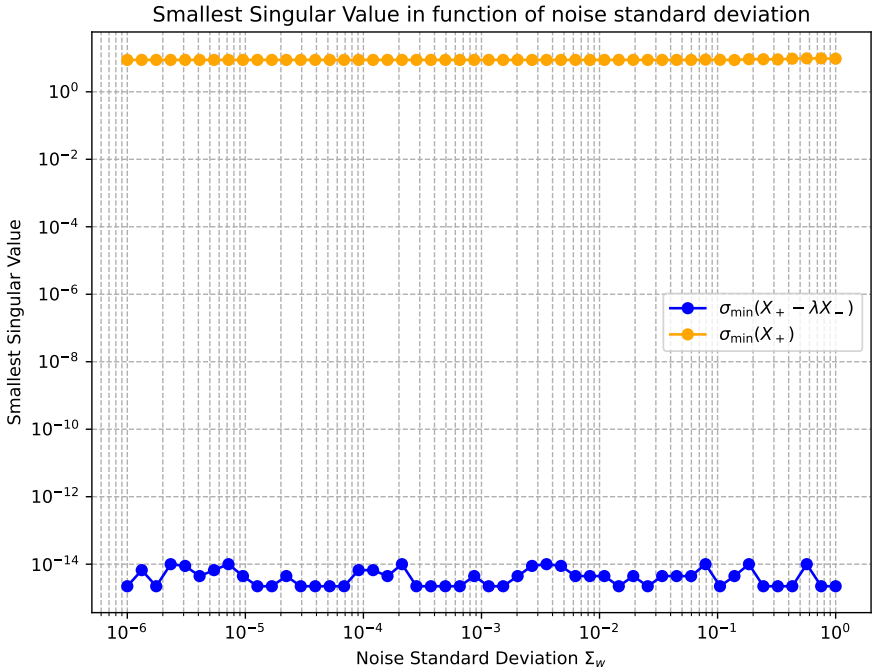

Figure 2.5: Evolution of the smallest singular values of X_+ and $X_+ - \lambda X_-$ in function of the noise standard deviation Σ_w for the controllable part.

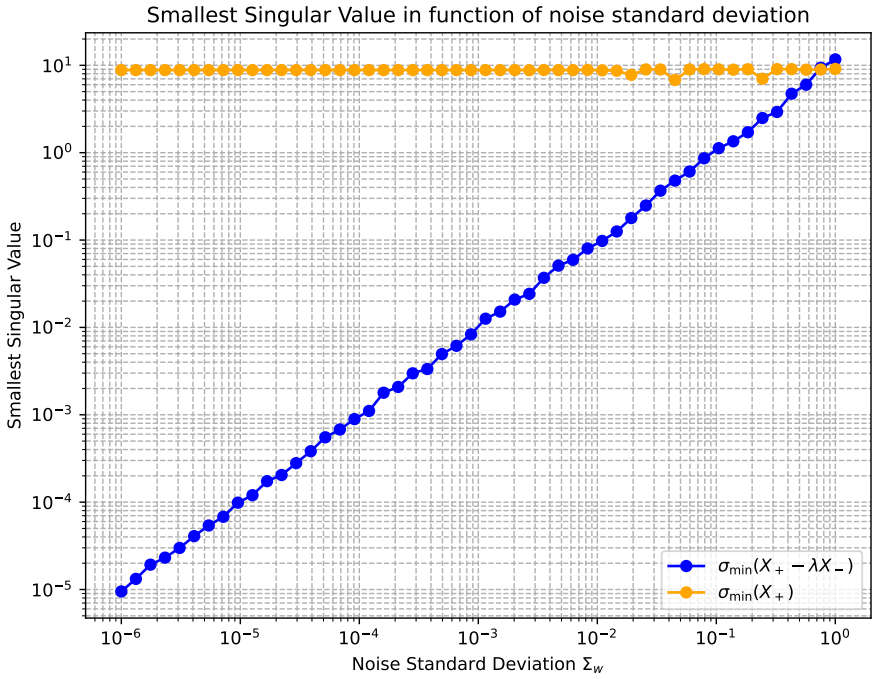

Figure 2.6: Evolution of the smallest singular values of X_+ and $X_+ - \lambda X_-$ in function of the noise standard deviation Σ_w for the uncontrollable part.

Previous observations indicate that $\sigma_{\min}(X_+ - \lambda X_-)$ is more informative than $\sigma_{\min}(X_+)$ for assessing controllability, especially near the boundary between controllable and uncontrollable regimes. Unlike $\sigma_{\min}(X_+)$, which remains nearly constant, the former responds sensitively to structural variations. It is therefore adopted as the primary diagnostic metric in the presence of noise.

Figure 2.5 shows that perturbing only the controllable state has little effect on the test outcome: $\sigma_{\min}(X_+ - \lambda X_-)$ stays close to numerical zero, and the test correctly identifies the lack of controllability. In contrast, Figure 2.6 reveals that injecting noise into the uncontrollable state significantly alters the result. Even modest disturbances are sufficient to inflate the singular value and trigger a false positive. This sensitivity to noise aligned with uncontrollable dynamics underscores a critical vulnerability of the test when applied to empirical data.

These findings underscore the susceptibility of informativity-based controllability tests to measurement noise, particularly when it affects uncontrollable modes. This lack of robustness motivates the integration of structural information into the analysis. The next sections develop a framework that combines data and structural insight to address this limitation.

Chapter 3

Structural Controllability

In many practical applications, the exact numerical values of system matrices are unavailable. However, the structural configuration, such as the presence or absence of interconnections, is often well known. Structural controllability offers a theoretical framework that abstracts from numerical values and focuses on the interconnection structure, typically modeled using patterns of zero, nonzero, or arbitrary entries. This chapter examines two central notions: weak structural controllability, introduced in [49, 50], which investigates whether at least one realization of the pattern yields a controllable system; and strong structural controllability, introduced in [50, 53], which requires that all realizations be controllable.

3.1 Pattern Matrices and Associated Graphs

To analyze systems for which only structural information is available, the framework of pattern matrices, as formalized in [53], is employed. This abstraction captures the sparsity structure of system matrices and enables controllability analysis without requiring precise parameter values. Importantly, it can also account for entries that may vanish under specific conditions, such as parameter cancellations.

Definition 3.1 (Pattern Matrix)**.** A pattern matrix is a symbolic matrix

$$\mathcal{M} \in \{0, *, ?\}^{p \times q},$$

where:

- 0 denotes a fixed zero entry,
- $*$ denotes a free but nonzero entry,
- ? denotes an arbitrary entry, possibly zero or nonzero.

This generalization is essential in real-world scenarios. A motivating example arises in electrical networks, where component values may cancel under specific configurations.

Example 3.2 (Circuit model with ambiguous interconnections [53]). Consider the circuit in Figure 3.1, which includes resistive, inductive, and capacitive components, along with a current-controlled voltage source with gain G. The physical parameters R, C_1, C_2, L are strictly positive but not precisely known.

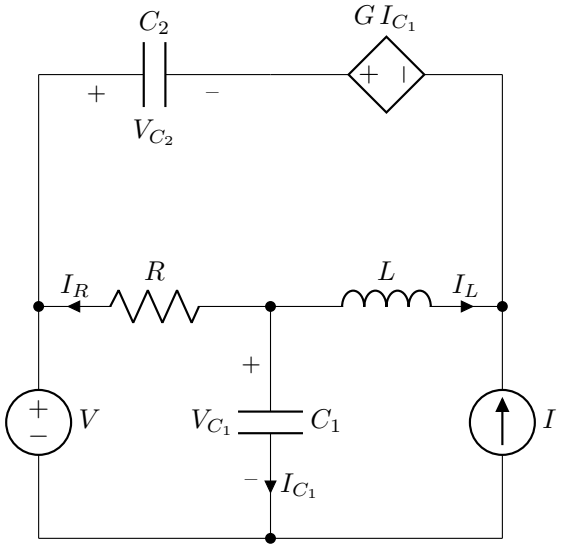

Figure 3.1: Example circuit.

According to Kirchhoff's laws, the system is described by the matrices

$$A = \begin{bmatrix} -\frac{1}{RC_1} & 0 & -\frac{1}{C_1} \\ 0 & 0 & -\frac{1}{C_2} \\ \frac{R-G}{RL} & \frac{1}{L} & -\frac{G}{L} \end{bmatrix}, \quad B = \begin{bmatrix} \frac{1}{RC_1} & 0 \\ 0 & -\frac{1}{C_2} \\ \frac{G-R}{RL} & 0 \end{bmatrix}.$$

Some entries, such as A_{11}, are structurally nonzero (i.e., nonzero for all admissible parameter values), some are fixed zeros, and others, such as A_{31}, may vanish depending on parameter relations, for instance when $R = G$. These cases motivate the use of '?' entries in the corresponding pattern matrices:

$$\mathcal{A} = \begin{bmatrix} * & 0 & * \\ 0 & 0 & * \\ ? & * & * \end{bmatrix}, \quad \mathcal{B} = \begin{bmatrix} * & 0 \\ 0 & * \\ ? & 0 \end{bmatrix}.$$

This illustrates how physical uncertainty naturally leads to pattern structures involving '?' entries, which fall outside the scope of classical structural controllability analysis based solely on '0' and '$*$'.

To formally describe all numerical systems consistent with a given pattern, the notion of a *pattern class* is introduced.

Definition 3.3 (Pattern Class)**.** Given a pattern matrix $\mathcal{M} \in \{0, *, ?\}^{p\times q}$, its *pattern class* is the set of all real matrices that satisfy the structural constraints encoded by $\mathcal{M}$:

$$\mathcal{P}(\mathcal{M}) = \left\{ M \in \mathbb{R}^{p\times q} \;\middle|\; \begin{array}{l} M_{ij} = 0 \text{ if } \mathcal{M}_{ij} = 0, \\ M_{ij} \neq 0 \text{ if } \mathcal{M}_{ij} = * \end{array} \right\}.$$

Entries marked with '?' impose no constraint and may take any real value.

The pattern class $\mathcal{P}(\mathcal{M})$ thus represents all numerical realizations consistent with a given symbolic structure. It provides a foundation for studying system-theoretic properties, such as controllability, independently of specific parameter values.

Definition 3.4 (Structured Discrete LTI System)**.** Let $\mathcal{A} \in \{0, *, ?\}^{n\times n}$ and $\mathcal{B} \in \{0, *, ?\}^{n\times m}$ be pattern matrices representing the structure of the state and input matrices. The pair $(\mathcal{A}, \mathcal{B})$ defines a *structured discrete-time LTI system* as the set of all systems of the form

$$x(t+1) = Ax(t) + Bu(t),$$

where $A \in \mathcal{P}(\mathcal{A})$ and $B \in \mathcal{P}(\mathcal{B})$. In this settings, the matrices A and B are not known exactly, but must conform to the symbolic patterns encoded in $\mathcal{A}$ and $\mathcal{B}$.

To illustrate how structured systems can be constructed from pattern matrices, consider the following example.

Example 3.5. Let $n = 2$ and $m = 1$. Define the pattern matrices:

$$\mathcal{A} = \begin{pmatrix} * & ? \\ 0 & * \end{pmatrix}, \qquad \mathcal{B} = \begin{pmatrix} ? \\ * \end{pmatrix}.$$

The corresponding pattern classes are given by:

$$\mathcal{P}(\mathcal{A}) = \left\{ A = \begin{pmatrix} a_{11} & a_{12} \\ 0 & a_{22} \end{pmatrix} \;\middle|\; a_{11}, a_{22} \neq 0,\ a_{12} \in \mathbb{R} \right\},$$

$$\mathcal{P}(\mathcal{B}) = \left\{ B = \begin{pmatrix} b_1 \\ b_2 \end{pmatrix} \;\middle|\; b_2 \neq 0,\ b_1 \in \mathbb{R} \right\}.$$

A particular realization within this structured family is:

$$A = \begin{pmatrix} 2 & -0.5 \\ 0 & 3 \end{pmatrix}, \quad B = \begin{pmatrix} 1 \\ 4 \end{pmatrix},$$

which yields the discrete-time system:

$$x(t+1) = \begin{pmatrix} 2 & -0.5 \\ 0 & 3 \end{pmatrix} x(t) + \begin{pmatrix} 1 \\ 4 \end{pmatrix} u(t).$$

Any other choice with $a_{11}, a_{22}, b_2 \neq 0$ and arbitrary a_{12}, b_1 also defines a valid realization of the structured system $(\mathcal{A}, \mathcal{B})$.

To analyze structural controllability using graph-theoretic tools, it is useful to associate each pattern matrix with a directed graph that encodes the structural constraints. This graph captures the influence structure implied by the entries of the matrix and serves as a central object in subsequent combinatorial tests.

Definition 3.6 (Associated Graph)**.** Let $\mathcal{M} \in \{0, *, ?\}^{p \times q}$ be a pattern matrix with $p \leq q$. Define two disjoint sets of directed edges:

$$E_* \;=\; \{\, (j,i) \in \{1,\dots,q\} \times \{1,\dots,p\} \,|\, \mathcal{M}_{ij} = * \},$$

$$E_? \;=\; \{\, (j,i) \in \{1,\dots,q\} \times \{1,\dots,p\} \,|\, \mathcal{M}_{ij} = ? \}.$$

The *associated graph* $G(\mathcal{M}) = (V, E)$ is defined by:

$$V = \{1, 2, \dots, q\}, \qquad E = E_* \cup E_?,$$

where edges in E_* are drawn as *solid* arrows, and those in $E_?$ as *dashed* arrows. A node j is said to be an *out-neighbor* of node i if $(i, j) \in E$, i.e., if there is a directed edge from i to j.

The following example illustrates how the associated graph can be constructed from a pattern matrix.

Example 3.7. Consider the pattern matrix $\mathcal{M} \in \{0, *, ?\}^{4 \times 5}$ given by:

$$\mathcal{M} = \begin{pmatrix} 0 & 0 & * & 0 & 0 \\ 0 & * & * & ? & * \\ * & 0 & ? & 0 & 0 \\ 0 & * & 0 & 0 & ? \end{pmatrix}.$$

The associated graph $G(\mathcal{M})$, shown in Figure 3.2, contains five nodes, each corresponding to a column of $\mathcal{M}$. Solid arrows represent entries marked '$*$', and dashed arrows represent entries marked '?'. For instance, since $\mathcal{M}_{4,2} = *$, there is a solid arrow from node 2 to node 4. Similarly, since $\mathcal{M}_{4,5} = ?$, there is a dashed arrow from node 5 to node 4. These edges indicate that node 4 is an out-neighbor of nodes 2 and 5. In the graph representation, node 5, which corresponds to an input, is shown in grey.

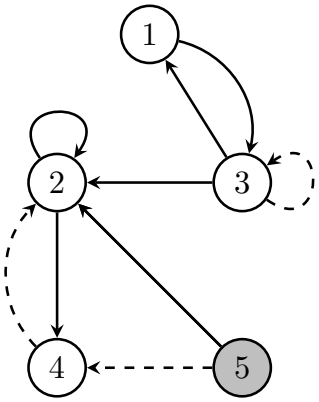

Figure 3.2: Associated graph $G(\mathcal{M})$ corresponding to the pattern matrix.

3.2 Weak Structural Controllability

Originally introduced by Lin [47], the notion of *structural controllability* focuses on whether a system can be controllable based solely on the pattern of zero and nonzero entries in its matrices, independently of their numerical values. In much of the literature, this concept is referred to simply as *structural controllability*. The qualifier "weak" has been adopted in more recent works to distinguish it from *strong structural controllability*, which imposes stricter conditions.

Definition 3.8 (Weak Structural Controllability)**.** The structured system $(\mathcal{A}, \mathcal{B})$ is *weakly structurally controllable* if there exists at least one realization (A, B) which is controllable, with $A \in \mathcal{P}(\mathcal{A})$ and $B \in \mathcal{P}(\mathcal{B})$.

Lin's framework assumes binary pattern matrices, whose entries are either fixed zeros or structurally free. In the present setting, more expressive structural information is allowed by including symbolic entries from $\{0, *, ?\}$. To recover the classical weak structural controllability tests, this general model can be translated into Lin's binary formalism by constructing equivalent pattern matrices as follows.

Let $\mathcal{A} \in \{0, *, ?\}^{n\times n}$ and $\mathcal{B} \in \{0, *, ?\}^{n\times m}$. Define binary matrices $\mathcal{A}' \in \{0, ?\}^{n\times n}$ and $\mathcal{B}' \in \{0, ?\}^{n\times m}$ by:

$$\mathcal{A}'_{ij} = \begin{cases} 0 & \text{if } \mathcal{A}_{ij} = 0 \\ ? & \text{otherwise} \end{cases}, \qquad \mathcal{B}'_{ij} = \begin{cases} 0 & \text{if } \mathcal{B}_{ij} = 0 \\ ? & \text{otherwise} \end{cases}.$$

The weak structural controllability of $(\mathcal{A}, \mathcal{B})$ in this general framework is then equivalent to that of $(\mathcal{A}', \mathcal{B}')$ in Lin's original sense. This equivalence holds because weak structural controllability is a *generic* property [57], meaning that if one admissible realization is controllable, then almost all are.

Definition 3.9 (Generic Property [57])**.** A system-theoretic property is said to be *generic* for a structured system if it holds for all realizations of the free parameters, except those lying in a proper algebraic variety, i.e., the zero set of one or more nontrivial real-coefficient polynomials in the parameter space.

This interpretation clarifies why structurally nonzero entries ($*$) can be treated as fully free (?) in this context: since the existence of a single controllable realization implies that almost all realizations are controllable, except for a negligible subset of pathological parameter values, no generality is lost by this simplification.

Several other properties in structural system theory, such as observability, decoupling, or pole assignability, are also known to be generic. Comprehensive surveys of these properties and their algebraic foundations can be found in [48] and [58]. The following sections introduce concrete algebraic and graph-theoretic criteria for verifying weak structural controllability.

3.2.1 Algebraic Characterization

Given a structured system $(\mathcal{A}, \mathcal{B})$, controllability is assessed by analyzing the controllability matrices of realizations (A, B) such that $A \in \mathcal{P}(\mathcal{A})$ and $B \in \mathcal{P}(\mathcal{B})$. This leads to a purely algebraic condition for weak structural controllability.

Definition 3.10 (Structured Controllability Matrix)**.** Let $\mathcal{A} \in \{0, *, ?\}^{n\times n}$ and $\mathcal{B} \in \{0, *, ?\}^{n\times m}$ be pattern matrices. The *structured controllability matrix* associated with the pair $(\mathcal{A}, \mathcal{B})$ is the symbolic matrix

$$\mathcal{C}(\mathcal{A}, \mathcal{B}) = \begin{bmatrix} \mathcal{B} & \mathcal{A}\mathcal{B} & \mathcal{A}^2\mathcal{B} & \cdots & \mathcal{A}^{n-1}\mathcal{B} \end{bmatrix},$$

where products are interpreted symbolically, preserving structural constraints.

This symbolic matrix reflects the structural pattern common to all controllability matrices of realizations with $A \in \mathcal{P}(\mathcal{A})$ and $B \in \mathcal{P}(\mathcal{B})$.

Definition 3.11 (Generic Rank [59])**.** The *generic rank* of a structured matrix $\mathcal{M}$, denoted g-rank$(\mathcal{M})$, is the highest rank that attained by any numerical realization $M \in \mathcal{P}(\mathcal{M})$. It corresponds to the rank obtained for almost all admissible choices of the free parameters.

Using this notion, the following criterion provides a necessary and sufficient condition for weak structural controllability.

Theorem 3.12 (Algebraic Test for Weak Structural Controllability [48])**.** *A structured system $(\mathcal{A}, \mathcal{B})$ is weakly structurally controllable if and only if the structured controllability matrix satisfies:*

$$\text{g-rank}\left(\mathcal{C}(\mathcal{A}, \mathcal{B})\right) = n.$$

In other words, the system is weakly structurally controllable if the controllability matrix has full row generic rank, meaning it attains rank n for almost all choices of the free parameters in A and B.

Example 3.13. Let $n = 3$, $m = 1$, and consider the pattern matrices:

$$\mathcal{A} = \begin{pmatrix} ? & ? & ? \\ ? & ? & ? \\ ? & ? & ? \end{pmatrix}, \qquad \mathcal{B} = \begin{pmatrix} ? \\ ? \\ ? \end{pmatrix}.$$

All entries are structurally free. By definition, any realization $A \in \mathcal{P}(\mathcal{A})$ and $B \in \mathcal{P}(\mathcal{B})$ can contain arbitrary (zero or nonzero) values. The corresponding structured controllability matrix is

$$\mathcal{C}(\mathcal{A}, \mathcal{B}) = \begin{bmatrix} \mathcal{B} & \mathcal{A}\mathcal{B} & \mathcal{A}^2\mathcal{B} \end{bmatrix} = \begin{pmatrix} ? & ? & ? \\ ? & ? & ? \\ ? & ? & ? \end{pmatrix}.$$

The structural pattern of $\mathcal{C}(\mathcal{A}, \mathcal{B})$ is fully free, and therefore its generic rank is 3. In addition, for almost all admissible choices of the free parameters, the numerical matrix $\mathcal{C}(A, B)$ has rank 3. Controllability fails only if the columns become algebraically dependent, which occurs on a proper algebraic variety in the parameter space. Hence, the system is weakly structurally controllable.

3.2.2 Graph–Theoretic Characterization

In addition to algebraic conditions, weak structural controllability can be characterized through purely combinatorial properties of the graph associated with the system [49, 50]. This graph-based approach enables controllability analysis by inspecting the structure of the pattern matrices directly. In particular, the concept of an *associated bipartite graph* plays a central role in this analysis.

Definition 3.14 (Associated Bipartite Graph)**.** Let $\mathcal{M} \in \{0, ?\}^{p \times q}$ be a pattern matrix with $p \leq q$. The *associated bipartite graph* is the undirected graph defined by

$$H(\mathcal{M}) = (V^+, V^-, E),$$

where the two disjoint node sets are:

$$V^+ = \{1, 2, \ldots, q\}, \qquad V^- = \{1', 2', \ldots, p'\},$$

and the edge set $E \subseteq V^+ \times V^-$ is given by:

$$E = \{\, (j, i') \in \{1, \ldots, q\} \times \{1', \ldots, p'\} \mid \mathcal{M}_{ij} = ? \,\}.$$

The primes on the nodes in V^-, which represent the rows of $\mathcal{M}$, serve to distinguish them from the column nodes in V^+. Since this characterization is developed in the classical weak controllability setting, where structurally nonzero entries are treated as free, all edges in the bipartite graph are drawn as solid lines.

To illustrate the construction of a bipartite graph from a structural pattern matrix, consider the following example.

Example 3.15. Let $\mathcal{M} \in \{0, ?\}^{4\times 5}$ be the pattern matrix

$$\mathcal{M} = \begin{pmatrix} ? & 0 & ? & 0 & 0 \\ 0 & ? & 0 & 0 & ? \\ ? & 0 & 0 & ? & 0 \\ ? & 0 & 0 & 0 & ? \end{pmatrix}.$$

The associated bipartite graph $H(\mathcal{M})$ is shown in Figure 3.3. Its two node sets are

$$V^+ = \{1, 2, 3, 4, 5\}, \qquad V^- = \{1', 2', 3', 4'\},$$

and each undirected edge connects column node $j \in V^+$ to row node $i' \in V^-$ whenever $\mathcal{M}_{ij} =?$.

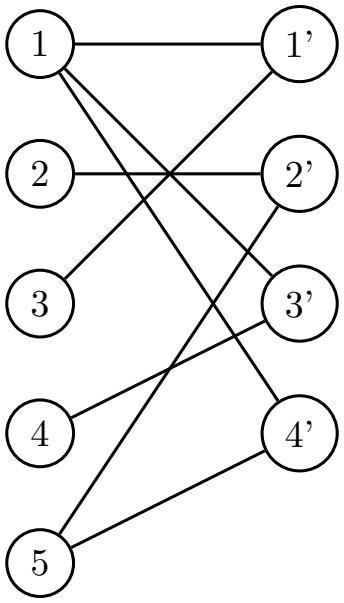

Figure 3.3: Bipartite graph $H(\mathcal{M})$ associated with the matrix in the example.

For example, since $\mathcal{M}_{3,1} =?$, node 1 is connected to node $3'$; and since $\mathcal{M}_{2,5} =?$, node 5 is connected to node $2'$. Each solid edge corresponds to a structurally free entry in the matrix $\mathcal{M}$.

To formulate a graph-theoretic test for weak structural controllability, two key notions from combinatorics and graph theory are used: matchings in bipartite graphs and accessibility in directed graphs. These concepts allow controllability conditions to be interpreted in purely structural terms. The definition below recalls the notion of a matching, which identifies disjoint connections between input and state nodes in the bipartite representation.

Definition 3.16 (Matching in a Bipartite Graph [60])**.** Let $H = (V^+, V^-, E)$ be a bipartite graph. A *matching* is a subset $M \subseteq E$ such that no two edges in M share a common endpoint. That is, for any distinct edges $(u_1, v_1), (u_2, v_2) \in M$, it holds that $\{u_1, v_1\} \cap \{u_2, v_2\} = \emptyset$. The *cardinality* of a matching is the number of edges it contains, denoted $|M|$.

This matching condition is complemented by a notion of accessibility, which ensures that each state can be influenced by at least one input node through a directed path in the associated graph.

Definition 3.17 (Accessibility in a Directed Graph)**.** Let $G = (V, E)$ be a directed graph, and let $S \subseteq V$ be a designated subset of nodes. A node $v \in V$ is said to be *accessible from* S if there exists a directed path in G from some node $s \in S$ to v. The set of accessible nodes from S is defined as

$$\{\, v \in V \mid \exists\, s \in S \text{ such that a directed path exists from } s \text{ to } v \,\}.$$

Remark 3.18. When S denotes the set of input nodes in a structured system, the corresponding notion is referred to as *input accessibility*. The system is said to be input-accessible if every state node is accessible from at least one input node.

The main graph-theoretic characterization of weak structural controllability can now be stated.

Theorem 3.19 (Graph-Theoretic Characterization of Weak Structural Controllability [50, 49])**.** *Let $(\mathcal{A}, \mathcal{B})$ be a structured system with $\mathcal{A} \in \{0, ?\}^{n\times n}$ and $\mathcal{B} \in \{0, ?\}^{n\times m}$. Let $G([\mathcal{A}\ \mathcal{B}])$ denote the directed graph associated with the combined pattern matrix $[\mathcal{A}\ \mathcal{B}]$, and let $H([\mathcal{A}\ \mathcal{B}])$ denote the corresponding bipartite graph. The structured system $(\mathcal{A}, \mathcal{B})$ is weakly structurally controllable if and only if the following two conditions are satisfied:*

1. ***Input Accessibility:*** *Every state node $i \in \{1, \ldots, n\}$ is accessible from the set of input nodes $\{n+1, \ldots, n+m\}$ in the directed graph $G([\mathcal{A}\ \mathcal{B}])$; that is, a directed path from an input node to node i exists.*

2. ***Full Matching:*** *The bipartite graph $H([\mathcal{A}\ \mathcal{B}])$ admits a matching of cardinality n, namely a set of n disjoint edges covering all state nodes $1', \ldots, n'$.*

This graph-theoretic characterization gives rise to efficient algorithmic procedures for verifying weak structural controllability. The first condition, input accessibility, can be checked by performing a single breadth-first search (BFS) or depth-first search (DFS) initiated from the set of input nodes in the digraph $G([\mathcal{A}\ \mathcal{B}])$. This procedure has linear time complexity, $\mathcal{O}(|V| + |E|)$.

The second condition, corresponding to the existence of a full matching in the bipartite graph $H([\mathcal{A}\ \mathcal{B}])$, constitutes a classical maximum matching problem. It can be solved efficiently using the Hopcroft–Karp algorithm [61], which runs in $O(|V|^{\frac{1}{2}} \cdot |E|)$ time.

A key implication is that weak structural controllability can be verified in polynomial time for general structured systems.

Remark 3.20. Although the verification of weak structural controllability is computationally tractable, the associated *input selection problem*, which consists in determining the *minimum number of input nodes* required to ensure structural controllability, is known to be *NP-hard* [62]. This task, often referred to as the *Minimum Driver Node Problem*, involves identifying the smallest input set such that all state nodes are accessible and a full matching exists.

The application of the graph-theoretic characterization is now illustrated through a complete example. The considered system satisfies both accessibility and full matching conditions, thereby confirming its weak structural controllability.

Example 3.21. Let $n = 3$ and $m = 1$. Consider the structured system $(\mathcal{A}, \mathcal{B})$ defined by:

$$\mathcal{A} = \begin{pmatrix} 0 & 0 & ? \\ ? & ? & 0 \\ 0 & ? & 0 \end{pmatrix}, \qquad \mathcal{B} = \begin{pmatrix} ? \\ 0 \\ ? \end{pmatrix}.$$

The associated directed graph $G([\mathcal{A}\ \mathcal{B}])$ is shown in Figure 3.4. The input node is labeled 4 and represented in grey, while state nodes are numbered 1, 2, and 3.

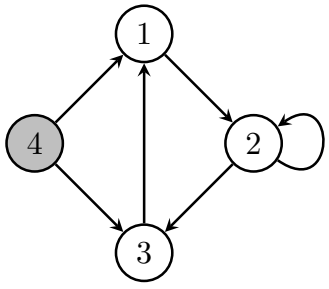

Figure 3.4: Directed graph $G([\mathcal{A}\ \mathcal{B}])$.

Input accessibility is verified by exhibiting directed paths from the input node to each state node. Figure 3.5 displays one such path per state node.

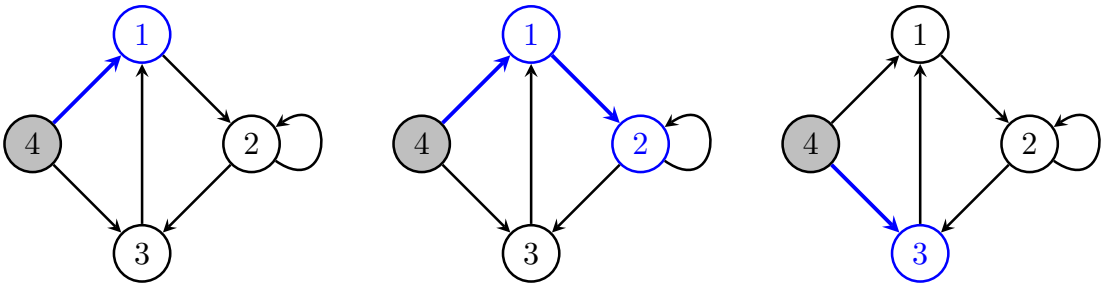

Figure 3.5: Accessibility of state nodes 1, 2, and 3 from input node 4.

Each state node is reachable from the input, confirming that the input accessibility condition is satisfied. The associated bipartite graph $H([\mathcal{A}\ \mathcal{B}])$ is shown in Figure 3.6. As per convention, the column indices define the set $V^+ = \{1, 2, 3, 4\}$, and the row indices define the set $V^- = \{1', 2', 3'\}$.

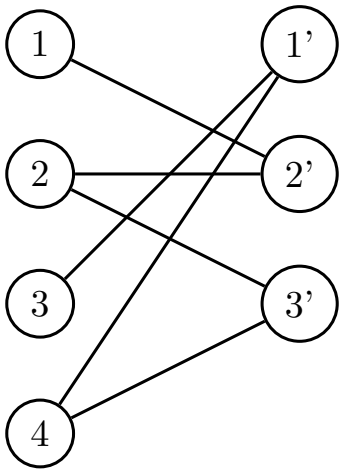

Figure 3.6: Associated bipartite graph $H([\mathcal{A}\ \mathcal{B}])$.

To verify the second condition in the graph-theoretic test, a matching that covers all nodes in V^- is required. One such matching is illustrated in Figure 3.7, where the selected edges are highlighted. It consists of the following three disjoint pairs:

$$M = \{(1, 2'),\ (3, 1'),\ (4, 3')\}.$$

Since the matching has cardinality $n = 3$ and covers all row nodes, the full matching condition is satisfied.

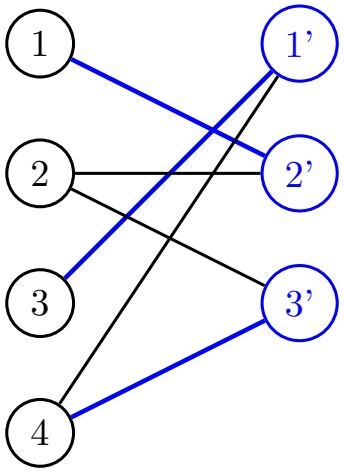

Figure 3.7: Full matching in $H([\mathcal{A}\ \mathcal{B}])$.

Together with input accessibility, this confirms that the system is weakly structurally controllable.

3.3 Strong Structural Controllability

In contrast to the weak setting, *strong structural controllability* requires that *every* realization consistent with the structural pattern yields a controllable system. This property is more conservative, as it ensures controllability under all admissible parameter combinations.

Strong structural controllability is studied within the full symbolic framework $\{0, *, ?\}$, as introduced in section 3.1. In this setting, entries marked with $*$ must be nonzero in all realizations, whereas entries marked ? may take arbitrary values. This distinction is critical when analyzing systems whose controllability must be preserved under all structurally consistent parameter variations, as examined in [53].

Definition 3.22 (Strong Structural Controllability)**.** The structured system $(\mathcal{A}, \mathcal{B})$ is said to be *strongly structurally controllable* if the pair (A, B) is controllable for all realizations with $A \in \mathcal{P}(\mathcal{A})$ and $B \in \mathcal{P}(\mathcal{B})$.

In other words, controllability must hold for all matrix pairs that respect the symbolic structure: no admissible choice of parameter values may lead to a loss of controllability.

The following example illustrates a structured system that is weakly, but not strongly, structurally controllable. This distinction highlights the more stringent requirements imposed by the strong notion.

Example 3.23. Let $n = 3$, $m = 1$, and consider the structured system $(\mathcal{A}, \mathcal{B})$ defined by:

$$\mathcal{A} = \begin{pmatrix} 0 & * & 0 \\ 0 & 0 & * \\ ? & 0 & ? \end{pmatrix}, \qquad \mathcal{B} = \begin{pmatrix} * \\ 0 \\ 0 \end{pmatrix}.$$

Entries marked $*$ are required to be nonzero in all realizations, while those marked ? may take arbitrary values. Consider the following admissible realization:

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

The corresponding controllability matrix is:

$$\mathcal{C}(A, B) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix},$$

which has full rank. Hence, the system is controllable for this specific realization, and the structured system is weakly structurally controllable. Now consider another realization in which the structurally free entry $\mathcal{A}_{3,1}$ is set to zero:

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

This choice remains consistent with the pattern matrix, since the entry $\mathcal{A}_{3,1} = 0$ is permitted under the ? symbol. The corresponding controllability matrix is:

$$\mathcal{C}(A, B) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix},$$

which has rank only 1. This realization therefore corresponds to an uncontrollable system. While the structured system admits a controllable realization (and is thus weakly structurally controllable), it is not controllable for all admissible parameter choices. Consequently, the system is not strongly structurally controllable.

The following sections introduce algebraic and graph-theoretic criteria for verifying strong structural controllability.

3.3.1 Algebraic Characterization

A key result established in [53] demonstrates that strong structural controllability can be verified by testing a symbolic rank condition directly on the pattern matrices. This approach does not require the enumeration of numerical realizations and instead relies entirely on the combinatorial structure encoded by $\mathcal{A}$ and $\mathcal{B}$. The central idea is to assess whether controllability is preserved under all admissible parameter assignments. This is accomplished by analyzing two symbolic matrices: the original pattern $[\mathcal{A}\ \mathcal{B}]$, and a modified version $[\bar{\mathcal{A}}\ \mathcal{B}]$, which adjusts for potential eigenvalue cancellations due to zero diagonal entries.

Definition 3.24 (Full Row Rank of a Pattern Matrix)**.** Let $\mathcal{M} \in \{0, *, ?\}^{p\times q}$ be a pattern matrix. The matrix $\mathcal{M}$ is said to have *full row rank* if every realization $M \in \mathcal{P}(\mathcal{M})$ satisfies $\operatorname{rank}(M) = p$.

To account for all possible eigenvalue configurations, including those affected by zero diagonal entries, a modified matrix $\bar{\mathcal{A}}$ is introduced.

Definition 3.25 (Modified Pattern Matrix $\bar{\mathcal{A}}$)**.** Let $\mathcal{A} \in \{0, *, ?\}^{n\times n}$ be a pattern matrix. The matrix $\bar{\mathcal{A}} \in \{0, *, ?\}^{n\times n}$ is obtained by modifying only the diagonal entries of $\mathcal{A}$ as:

$$\bar{\mathcal{A}}_{ii} := \begin{cases} * & \text{if } \mathcal{A}_{ii} = 0, \\ ? & \text{otherwise,} \end{cases} \quad \text{for all } i = 1, \ldots, n,$$

with all off-diagonal entries unchanged, i.e., $\bar{\mathcal{A}}_{ij} := \mathcal{A}_{ij}$ for all $i \neq j$.

This substitution enforces structural nonzeros on the diagonal, ensuring that all eigenvalue configurations are structurally admissible and avoiding singularities due to absent self-loops.

Theorem 3.26 (Algebraic Characterization of Strong Structural Controllability [53])**.** *Let $\mathcal{A} \in \{0, *, ?\}^{n\times n}$, $\mathcal{B} \in \{0, *, ?\}^{n\times m}$, and let $\bar{\mathcal{A}} \in \{0, *, ?\}^{n\times n}$ be defined as above. The structured system $(\mathcal{A}, \mathcal{B})$ is strongly structurally controllable if and only if both pattern matrices*

$$[\mathcal{A}\ \mathcal{B}] \quad \textit{and} \quad [\bar{\mathcal{A}}\ \mathcal{B}]$$

have full row rank.

The two rank conditions stated above are logically independent; neither implies the other. It is possible for $[\mathcal{A}\ \mathcal{B}]$ to have full row rank while $[\bar{\mathcal{A}}\ \mathcal{B}]$ does not, and vice versa. These cases reflect distinct structural obstructions to controllability: the former arises from general rank deficiency, while the latter is linked to pathological eigenvalue placement due to fixed zero entries on the diagonal. The following example illustrates a situation where the first rank condition is satisfied, but the second fails.

Example 3.27. Let

$$\mathcal{A} = \begin{pmatrix} * & * \\ 0 & 0 \end{pmatrix}, \quad \bar{\mathcal{A}} = \begin{pmatrix} ? & * \\ 0 & * \end{pmatrix}, \quad \mathcal{B} = \begin{pmatrix} * \\ * \end{pmatrix}.$$

In this case, the matrix $[\mathcal{A}\ \mathcal{B}]$ is generically of full row rank. Consider the realization

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

The resulting matrix $[\bar{A}\ B]$ fails to attain rank 2; its rank is only 1. This illustrates that the second rank condition can fail independently of the first.

A notable simplification arises when no diagonal entry of $\mathcal{A}$ is structurally zero. In this case, the second rank condition alone ensures strong structural controllability.

Corollary 3.28 (Zero-Diagonal Simplification [53])**.** *Suppose that none of the diagonal entries of $\mathcal{A}$ are fixed zeros. Then the structured system $(\mathcal{A}, \mathcal{B})$ is strongly structurally controllable if and only if the matrix $[\bar{\mathcal{A}}\ \mathcal{B}]$ has full row rank.*

Since $\bar{\mathcal{A}}$ is obtained by modifying only the zero diagonal entries of $\mathcal{A}$, the inclusion $\mathcal{P}(\mathcal{A}) \subseteq \mathcal{P}(\bar{\mathcal{A}})$ holds. Therefore, the first rank condition is implied by the second one.

Example 3.29. Let

$$\mathcal{A} = \begin{pmatrix} * & * \\ ? & * \end{pmatrix}, \quad \mathcal{B} = \begin{pmatrix} 0 \\ * \end{pmatrix}, \quad \bar{\mathcal{A}} = \begin{pmatrix} ? & * \\ ? & ? \end{pmatrix}.$$

Since no diagonal entry of $\mathcal{A}$ is fixed zero, each diagonal element of $\bar{\mathcal{A}}$ is structurally free. The combined pattern matrix becomes

$$[\bar{\mathcal{A}}\ \mathcal{B}] = \begin{pmatrix} ? & * & 0 \\ ? & ? & * \end{pmatrix},$$

which has full row rank for all realizations. Therefore, the system is strongly structurally controllable.

3.3.2 Graph–Theoretic Characterization

While the algebraic characterization provides exact and rigorous conditions for strong structural controllability, the associated symbolic rank tests are not always straightforward to verify in practice. A complementary perspective, based on graph theory, yields a purely combinatorial criterion that is both intuitive and computationally efficient. This approach relies on the notion of *colorability* of the directed graph associated with a pattern matrix. The procedure captures the propagation of influence through solid (nonzero) edges in the interconnection structure of the system.

Definition 3.30 (Colorability)**.** Let $\mathcal{M} \in \{0, *, ?\}^{p \times q}$ be a pattern matrix with $p \leq q$, and let $G(\mathcal{M}) = (V, E)$ denote its associated directed graph, where $V = \{1, \ldots, q\}$ is the vertex set and $E = E_* \cup E_?$ the edge set, with E_* and $E_?$ denoting the sets of solid and dashed edges, respectively. The graph $G(\mathcal{M})$ is said to be *colorable* if all nodes $1, \ldots, p$ (corresponding to the rows of $\mathcal{M}$) can be colored black via the following iterative process:

1. Initially, all nodes are colored white.
2. If a node $i \in V$ has exactly one white out-neighbor j, and the edge $(i, j) \in E_*$ is solid, then node j is colored black.
3. The process is repeated until no additional nodes can be colored.

The graph is deemed *colorable* if, upon termination, all nodes in $\{1, \ldots, p\}$ have been colored black.

The procedure described above is governed by a local mechanism, often referred to as the *color-change rule*, which reflects how structural influence can propagate deterministically along specific nonzero edges. Importantly, nodes $p + 1, \ldots, q$ (typically corresponding to input variables) are not required to be colored black. These nodes generally have no incoming edges and act as sources of control.

Theorem 3.31 (Full Rank via Colorability [53])**.** *A pattern matrix $\mathcal{M} \in \{0, *, ?\}^{p \times q}$ has full row rank for all realizations in its pattern class if and only if its associated graph $G(\mathcal{M})$ is colorable.*

This result enables a purely graph-theoretic reformulation of the algebraic test for strong structural controllability. In particular, verifying whether a structured system is strongly controllable amounts to checking the colorability of two associated pattern graphs.

Theorem 3.32 (Graph-Theoretic Strong Structural Controllability [53])**.** *The structured system $(\mathcal{A}, \mathcal{B})$ is strongly structurally controllable if and only if the associated graphs*

$$G([\mathcal{A}\ \mathcal{B}]) \quad \textit{and} \quad G([\bar{\mathcal{A}}\ \mathcal{B}])$$

are both colorable.

The colorability of a graph, as defined by the color-change rule, can be verified efficiently. Specifically, the procedure can be implemented in $\mathcal{O}(p \cdot q)$ operations by scanning each solid out-neighbor of every node at most once. This makes the graph-theoretic test both intuitive and computationally tractable for large-scale systems. The following example illustrates the application of this criterion in practice.

Example 3.33 (Colorability Verification). Let $n = 3$, $m = 2$, and consider the following pattern matrices:

$$\mathcal{A} = \begin{pmatrix} * & 0 & * \\ 0 & 0 & * \\ ? & * & * \end{pmatrix}, \quad \mathcal{B} = \begin{pmatrix} * & 0 \\ 0 & * \\ ? & 0 \end{pmatrix}.$$

The modified matrix $\bar{\mathcal{A}}$ is obtained by replacing diagonal entries of $\mathcal{A}$ according to the rule: set $\bar{\mathcal{A}}_{ii} = *$ if $\mathcal{A}_{ii} = 0$, and $\bar{\mathcal{A}}_{ii} =?$ otherwise. This yields:

$$\bar{\mathcal{A}} = \begin{pmatrix} ? & 0 & * \\ 0 & * & * \\ ? & * & ? \end{pmatrix}.$$

The associated graphs $G([\mathcal{A}\ \mathcal{B}])$ and $G([\bar{\mathcal{A}}\ \mathcal{B}])$ are displayed in Figure 3.8. Nodes 4 and 5, corresponding to inputs, are shown in grey.

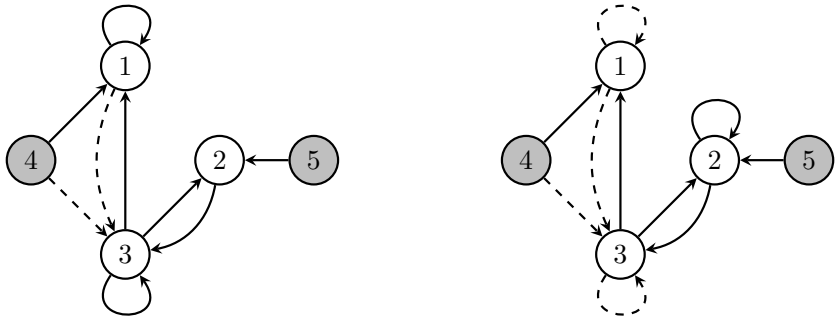

Figure 3.8: Graphs associated with $[\mathcal{A}\ \mathcal{B}]$ (left) and $[\bar{\mathcal{A}}\ \mathcal{B}]$ (right).

The color-change rule is now applied to $G([\mathcal{A}\ \mathcal{B}])$. The propagation proceeds in three steps as illustrated in Figure 3.9.

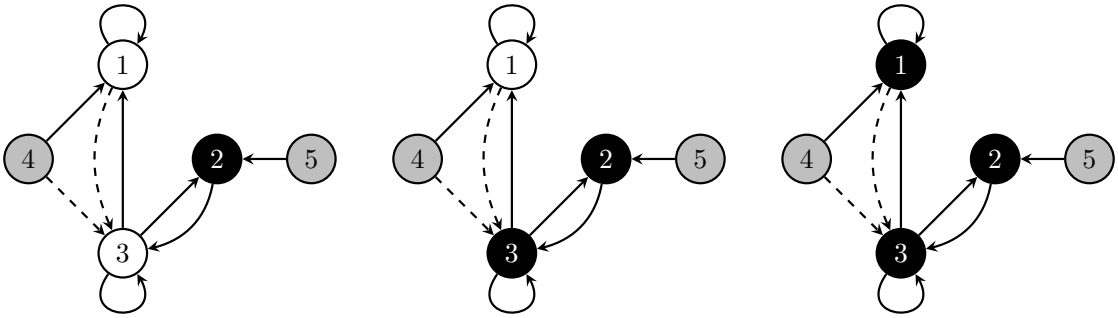

Figure 3.9: Color-change propagation on $G([\mathcal{A}\ \mathcal{B}])$.

In the first step (left), node 5 has a unique white out-neighbor, node 2, connected by a solid edge. Node 2 is therefore colored black. In the second step (center), node 2 has a unique white out-neighbor, node 3, which is then colored black. In the third and final step (right), node 3 has a unique white out-neighbor, node 1, which is also colored black.

At the end of the procedure, all nodes in $\{1, 2, 3\}$ are black, confirming that $G([\mathcal{A}\ \mathcal{B}])$ is colorable. The same color-change sequence applies to the graph $G([\bar{\mathcal{A}}\ \mathcal{B}])$, which is also colorable. According to the graph-theoretic characterization, this establishes that the structured system $(\mathcal{A}, \mathcal{B})$ is strongly structurally controllable.

3.4 Empirical Analysis

The preceding sections have presented rigorous algebraic and graph-theoretic conditions for verifying whether a structured system is weakly or strongly structurally controllable. Although these criteria are exact, their verification typically requires detailed structural information, which may be unavailable in uncertain environments. This raises the question of whether controllability can be inferred from coarse structural indicators alone, without explicitly checking rank or colorability.

This section explores that question from two complementary angles. First, it examines whether certain canonical interconnection patterns, such as directed cycles and directed stars, are structurally controllable by construction, regardless of specific parameter values. Second, it adopts a probabilistic perspective to assess how the likelihood of controllability varies with the overall edge density of the network.

Using both deterministic examples and Monte Carlo simulations, the empirical analysis provides insight into how global structural features, such as network topology and edge density, can influence controllability even in the absence of precise knowledge about the location or value of nonzero entries.

3.4.1 Controllability in Canonical Graph Topologies

Certain network structures possess inherent controllability properties that can be deduced without detailed analysis. This section illustrates that idea using two canonical interconnection patterns: directed cycles and directed stars. These examples show that structural controllability can, in some cases, be inferred directly from the topology of the graph, independently of the specific parameter values. Similar reasoning extends to other configurations, such as chains, trees, or complete graphs.

Directed Cycles with a Single Input

Consider the directed 3-cycle shown in Figure 3.10, where each state node influences the next in a closed loop. A single input (represented by node 4, shown in grey) is applied to state node 1.

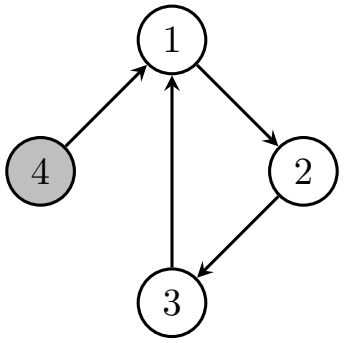

Figure 3.10: Directed 3-cycle with an input applied to node 1.

The associated structural matrices are:

$$\mathcal{A} = \begin{pmatrix} 0 & 0 & * \\ * & 0 & 0 \\ 0 & * & 0 \end{pmatrix}, \qquad \mathcal{B} = \begin{pmatrix} * \\ 0 \\ 0 \end{pmatrix}.$$

The corresponding controllability matrix is:

$$\mathcal{C}(A, B) = \begin{bmatrix} B & AB & A^2B \end{bmatrix} = \begin{pmatrix} * & 0 & 0 \\ 0 & * & 0 \\ 0 & 0 & * \end{pmatrix}.$$

Each structurally nonzero diagonal entry lies in a distinct row, guaranteeing full row rank for all realizations. Therefore, the system is strongly structurally controllable. This conclusion extends to any directed 3-cycle with a single input. Since the nodes of the cycle can be relabeled without altering the underlying interconnection pattern, the same argument applies regardless of the initial node ordering.

More generally, this property holds for any directed cycle of size n with a single input. By renumbering the nodes so that they follow the natural order of the cycle and placing the input at node 1, the system matrix $\mathcal{A}$ takes the form of a cyclic permutation matrix with a single nonzero entry per row and column, appropriately shifted. The corresponding controllability matrix $\mathcal{C}(A, B) = [B \; AB \; \ldots \; A^{n-1}B]$ becomes diagonal, with structurally nonzero entries along the diagonal, ensuring full rank. Therefore, every directed cycle of arbitrary size with a single input is strongly structurally controllable.

Directed Star

Now consider the directed star topology illustrated in Figure 3.11. It consists of node 1 at the center, with directed edges to leaf nodes 2, 3, and 4. The system has three inputs applied directly at nodes 1, 2, and 3 (input nodes are shown in grey). Node 4 is not directly actuated but is influenced through node 1.

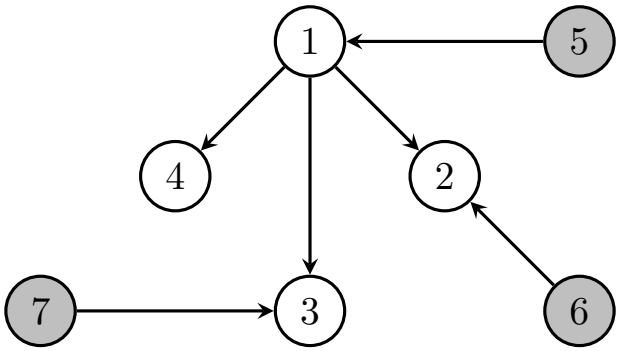

Figure 3.11: Directed star with 4 states and inputs at nodes 1, 2, and 3.

The structural matrices are:

$$\mathcal{A} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ * & 0 & 0 & 0 \\ * & 0 & 0 & 0 \\ * & 0 & 0 & 0 \end{pmatrix}, \qquad \mathcal{B} = \begin{pmatrix} * & 0 & 0 \\ 0 & * & 0 \\ 0 & 0 & * \\ 0 & 0 & 0 \end{pmatrix}.$$

The controllability matrix takes the form:

$$\mathcal{C} = \begin{bmatrix} B & AB & A^2B & A^3B \end{bmatrix} = \begin{pmatrix} * & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & * & 0 & * & 0 & \cdots & 0 \\ 0 & 0 & * & * & 0 & \cdots & 0 \\ 0 & 0 & 0 & * & 0 & \cdots & 0 \end{pmatrix}.$$

The first three rows are structurally covered by the input matrix B, while the last row is reached through the product AB, reflecting the influence path from node 1 to node 4. Each row contains a structural pivot in a distinct column, and all higher-order terms vanish since $A^k = 0$ structurally for $k \geq 2$. Therefore, the system is strongly structurally controllable.

This conclusion can also be understood in terms of structural matching. The edge from node 1 to node 4 ensures a path to the otherwise unactuated node, while the direct inputs to nodes 1, 2, and 3 provide control over the remaining states. More generally, a directed star with n nodes requires at least $n - 1$ input nodes, one at the center and the others at the leaves, to ensure strong structural controllability.

3.4.2 Controllability and Edge Density in Random Graphs

Beyond canonical motifs, it is natural to explore whether structural controllability can be inferred from global metrics such as *edge density*. Edge density is defined as the ratio of actual directed edges (either within the state-transition matrix $\mathcal{A}$ or from the input matrix $\mathcal{B}$) to the total number of possible edges, which equals $n(n+m)$ for a system with n states and m inputs. This notion quantifies how densely connected a system is, without regard to the specific placement of connections. Edge density is inversely related to sparsity: it corresponds to the proportion of nonzero entries (marked $*$ or ?) in the pattern matrices $(\mathcal{A}, \mathcal{B})$.

To empirically study the effect of edge density on controllability, a Monte Carlo simulation is conducted. For fixed system size n and various input dimensions $m \leq n$, $10,000$ random structured systems are generated at each level of edge density. For a given number of edges, the positions of nonzero entries are assigned uniformly at random across all $n(n+m)$ possible locations. Each realization is then tested for weak and strong structural controllability.

Figure 3.12 illustrates the relationship between edge density and the probability of structural controllability for systems of size $n = 3$ and input dimensions $m = 1, 2, 3$. The trends observed provide several key insights.

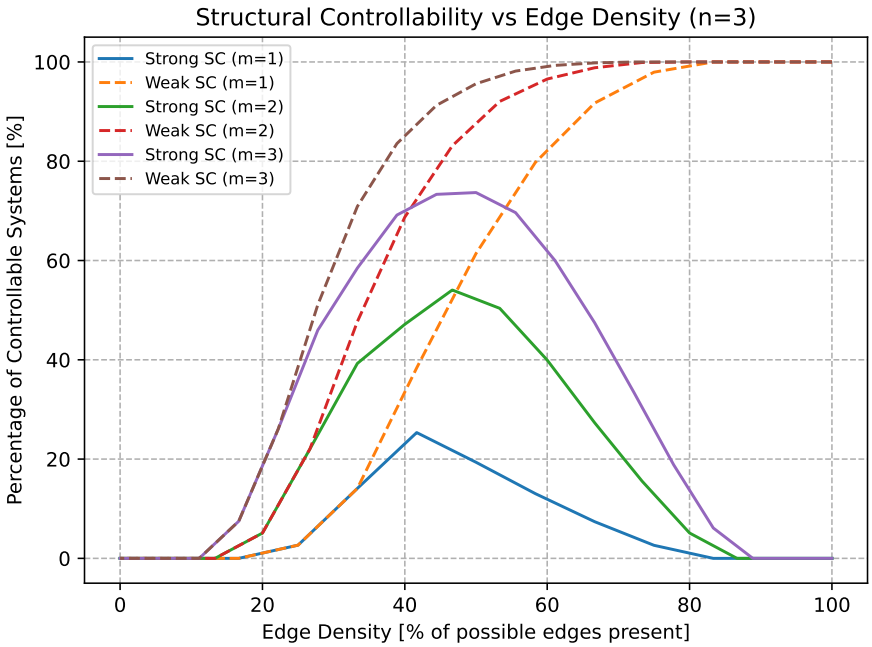

Figure 3.12: Controllability vs edge density for $n = 3$ and $m \in \{1, 2, 3\}$.

First, for very sparse systems (corresponding to edge densities below approximately 15%), the probability of controllability is essentially zero for both the weak and strong notions. This is expected: with too few edges, basic structural conditions such as the matching condition and input accessibility (for weak controllability), or the colorability condition (for strong controllability), cannot be satisfied. Algebraically, a system with many structural zeros is also likely to contain entire rows or columns of zeros in $(\mathcal{A}, \mathcal{B})$, making full row rank unattainable.

Focusing next on strong structural controllability, the curves reveal a distinct bell-shaped profile. As edge density increases beyond 15%, the probability of strong controllability rises steadily, reaching a peak between approximately 40% and 50%. Beyond this point, the probability decreases again and tends toward zero as the system approaches full connectivity (near 85% to 100% edge density). This behavior highlights a fundamental trade-off: both under-connected and over-connected systems are unlikely to be strongly controllable. In dense systems, the abundance of structural connections increases the likelihood of pathological realizations in which the controllability matrix fails to have full rank for all parameter values.

Interestingly, the location and height of the peak shift with the number of inputs m; larger m values consistently increase the maximum probability of strong controllability. This observation aligns with the intuition that additional inputs enhance the controllability of the system. Graphically, more input nodes yield more initial black nodes in the colorability test, facilitating control propagation. Algebraically, more inputs increase the structural richness of the matrix $\mathcal{B}$, increasing the likelihood that the controllability matrix spans all state directions.

In contrast, weak structural controllability exhibits a markedly different behavior. The corresponding curves increase monotonically with edge density, reaching 100% as the system becomes dense. This reflects the fact that weak controllability only requires the *existence* of a single favorable realization. Greater density offers more flexibility for satisfying the structural matching condition. As before, increasing m shifts the curves leftward, making weak controllability more likely at lower densities due to improved input accessibility.

These simulations highlight the critical role of edge density in determining the likelihood of structural controllability. For strong controllability, a "sweet spot" emerges, neither too sparse nor too dense, where the probability of controllability is maximized. For weak controllability, denser systems tend to be more favorable. In both cases, increasing the number of inputs systematically improves controllability.

The trends shown in Figure 3.13 for systems with $n = 6$ states broadly mirror those observed for $n = 3$ case, though they appear more pronounced and shift noticeably as system size increases.

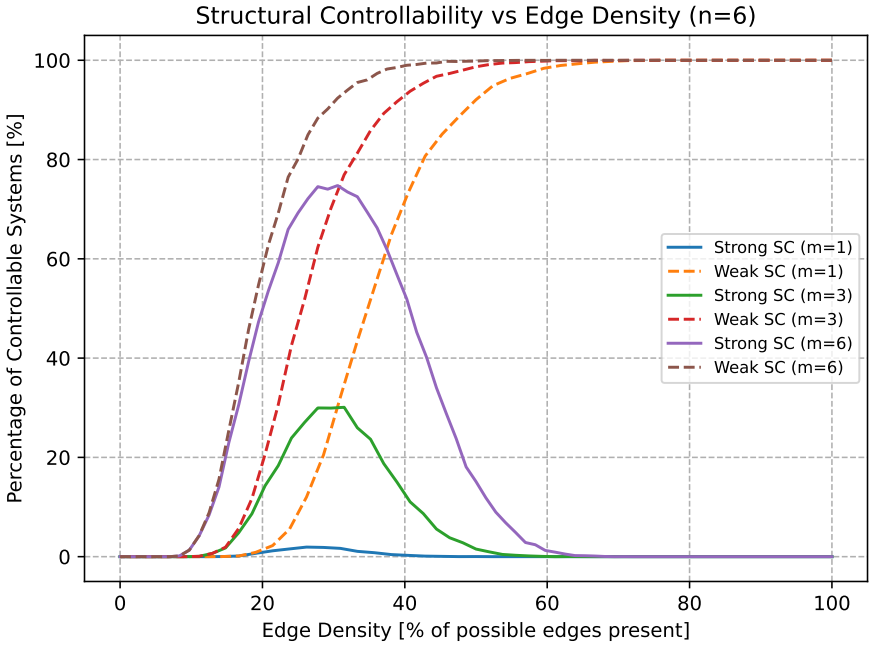

Figure 3.13: Controllability vs edge density for $n = 6$ and $m \in \{1, 3, 6\}$.

For strong structural controllability, the characteristic bell-shaped profile persists: the probability is near to zero at low edge densities, increases to a peak, and then declines again as connectivity increases. Compared to the $n = 3$ case, the peaks now occur at significantly lower edge densities (typically around 30% to 35% rather than 40% to 50%). This shift reflects the fact that larger systems admit more possible edge positions, so fewer connections are needed proportionally to meet the structural controllability conditions. As before, an excess of connections can introduce structural dependencies or redundancies that reduce the likelihood of achieving full rank across all realizations.

The number of inputs m remains a key factor. With only $m = 1$ input for $n = 6$ states, strong controllability is almost never achieved, regardless of edge density. This highlights the challenge of satisfying colorability or rank conditions when a single input must influence the entire system. In contrast, with $m = 6$ inputs (one assigned to each state), strong controllability is observed across a broad range of intermediate densities.

These empirical trends are further supported by theoretical results presented in [63], which analyzes the number of strongly structurally controllable graphs as a function of the number of nodes n and the number of inputs m. They establish that when the number of control nodes grows sublinearly with the system size (i.e., $m \leq \beta n$ for some constant $\beta < 1$) the proportion of strongly controllable graphs tends to zero as n increases. This is consistent with the simulation results: for $m = 1$, the probability of strong controllability drops significantly between $n = 3$ and $n = 6$. In contrast, they show that if the number of inputs grows linearly with the number of nodes, the proportion of strongly controllable graphs remains bounded away from zero. This behavior is clearly observed when each node is actuated individually (m = n), where the probability of strong controllability remains stable even as the network size increases. These results confirm that the scaling of the number of inputs with respect to system size plays a critical role in maintaining strong structural controllability.

On the weak side, the behavior remains distinctly monotonic: the likelihood of weak structural controllability increases steadily with edge density and typically reaches 100% well before full connectivity is achieved. Crucially, there is no decline in probability at high densities, in contrast to the strong case. This is because weak controllability requires only the existence of a single favorable realization. As edge density increases, so does the number of structural degrees of freedom, making it easier to satisfy the graph-theoretic matching and input-accessibility conditions.

These results further support the conclusion that structural controllability is highly sensitive to both the connectivity pattern and the input configuration. From a practical perspective, this has important implications: for example, when $n = m = 6$, even in the absence of precise structural information, observing an edge density between 60% to 80% is sufficient to infer with high confidence that the system is weakly controllable. This highlights a key insight: structural controllability can, in many cases, be inferred probabilistically from global features alone, even when the structural model is imperfect or under-identified.

Chapter 4

Structure Identification

After analyzing controllability through both numerical data and structural patterns, attention is now directed toward the intermediate step that links these two perspectives: the identification of the structure of the system based on data. This chapter aims to develop methods for inferring the sparsity pattern, that is, the zero or nonzero structure of the system matrices, from input–state trajectories. This structural information is essential for applying the tools of structural controllability introduced in the previous chapter. Several optimization-based techniques are presented to recover the underlying structure of the system, both in ideal noise-free conditions and in more realistic scenarios involving measurement or process noise.

4.1 Problem Formulation

This section addresses the problem of identifying the structure of a linear dynamical system from trajectory data. The objective is to recover the structure of the matrices $A_s \in \mathbb{R}^{n\times n}$ and $B_s \in \mathbb{R}^{n\times m}$ governing the evolution of a discrete-time LTI system:

$$X_+ = A_s X_- + B_s U_-, \tag{4.1}$$

where $X_-, X_+ \in \mathbb{R}^{n\times T}$ and $U_- \in \mathbb{R}^{m\times T}$ are data matrices constructed from input–state trajectories, as described in Chapter 2. A natural approach consists in formulating a least-squares optimization problem, leading to the following objective function:

$$\min_{A\in\mathbb{R}^{n\times n},\, B\in\mathbb{R}^{n\times m}} \; f(A,B) = \|X_+ - AX_- - BU_-\|_F^2, \tag{4.2}$$

where $\|\cdot\|_F$ denotes the Frobenius norm, which measures the total squared error between predicted and observed next-state values. Once an optimal solution is obtained, the structured system $(\mathcal{A}, \mathcal{B})$ is inferred by examining which entries of (A, B) are approximately zero or nonzero.

A brute-force approach to structure identification consists in testing every possible sparsity pattern. For each combination of zero or nonzero entries, the constrained version of (4.2) is solved, and the compatibility of the resulting model with the data is assessed. This procedure guarantees the identification of all admissible structures, or the unique structure if identifiability holds. However, since each entry in A and B can independently be either zero or nonzero, the total number of possible patterns is $2^{n(n+m)}$, resulting in an exponential number of optimization problems. Such a method is therefore computationally intractable.

To address this combinatorial complexity, this chapter presents three alternative strategies designed to promote sparsity directly through the optimization process. These approaches enable efficient structure recovery without the need to exhaustively enumerate all possible zero or nonzero configurations.

4.1.1 Problem Transformation

The initial problem (4.2) is intuitive but can be difficult to manipulate directly, as the unknowns are stored separately across the matrices A and B. To address this, the entire optimization problem is reformulated as an equivalent standard vector least-squares problem:

$$\min_{y \in \mathbb{R}^{n(n+m)}} \quad \|z - \Phi y\|_2^2, \tag{4.3}$$

where y denotes a vector of unknowns, Φ is a design matrix constructed from the data, and z is the target vector. This formulation is well established in the context of least-squares regression. It allows for efficient computation of gradients and Hessians, and provides a natural setting for incorporating structural priors such as sparsity, which play a central role in the structure identification process. The derivation begins from the original objective function:

$$f(A, B) = \|X_+ - AX_- - BU_-\|_F^2.$$

By concatenating the matrices A and B horizontally, and stacking the data matrices X_- and U_- vertically, the objective function can be rewritten as:

$$f(A, B) = \left\| X_+ - \begin{bmatrix} A & B \end{bmatrix} \begin{bmatrix} X_- \\ U_- \end{bmatrix} \right\|_F^2.$$

In order to facilitate vectorization, the expression is transposed to operate in a row-wise format. This transformation is valid, as the Frobenius norm is invariant under transposition. By applying the distributivity of the transpose over matrix multiplication and subtraction, the expression becomes:

$$f(A, B) = \left\| X_+^T - \left(\begin{bmatrix} X_- \\ U_- \end{bmatrix} \right)^T \begin{bmatrix} A & B \end{bmatrix}^T \right\|_F^2.$$

The following quantities are introduced to simplify notation:

$$Z = X_+^T \in \mathbb{R}^{T\times n}, \quad \varphi = \begin{bmatrix} X_- \\ U_- \end{bmatrix}^T \in \mathbb{R}^{T\times(n+m)}, \quad Y = \begin{bmatrix} A & B \end{bmatrix}^T \in \mathbb{R}^{(n+m)\times n}.$$

With this notation, the objective function can be rewritten as:

$$f(A,B) = f(Y) = \|Z - \varphi Y\|_F^2.$$

Although this formulation is more compact, the variables Z and Y remain matrix-valued. To express the problem as a standard vector regression, both sides are now vectorized. Specifically, the following definitions are introduced:

$$z = \mathrm{vec}(Z^T) \in \mathbb{R}^{nT}, \qquad y = \mathrm{vec}(Y^T) \in \mathbb{R}^{n(n+m)}.$$

In expanded form, the vectors z and y are expressed as:

$$z = \begin{bmatrix} Z_1^T \\ Z_2^T \\ \vdots \\ Z_T^T \end{bmatrix} \in \mathbb{R}^{nT}, \qquad y = \begin{bmatrix} (Y^T)_1 \\ (Y^T)_2 \\ \vdots \\ (Y^T)_n \end{bmatrix} \in \mathbb{R}^{n(n+m)},$$

where Z_i denotes the i-th row of Z, and $(Y^T)_j$ denote the j-th row of the transposed parameter matrix Y^T. The regression matrix $\Phi \in \mathbb{R}^{nT\times n(n+m)}$ is then defined as:

$$\Phi = \begin{bmatrix} I_n \otimes \varphi_1 \\ I_n \otimes \varphi_2 \\ \vdots \\ I_n \otimes \varphi_T \end{bmatrix},$$

where each φ_t is the t-th row of φ, and $\otimes$ denotes the Kronecker product. This construction ensures that the vectorized product φY corresponds to Φy in the reformulated problem. The final expression of the objective function is:

$$f(y) = \|z - \Phi y\|_2^2. \tag{4.4}$$

Consequently, the original matrix least-squares problem:

$$\min_{A,B} \ \|X_+ - AX_- - BU_-\|_F^2$$

is fully equivalent to the following vector formulation:

$$\min_{y\in\mathbb{R}^{n(n+m)}} \ \|z - \Phi y\|_2^2.$$

This reformulation significantly simplifies both the analysis and the solution process. It provides access to a broad set of tools from convex optimization, with or without sparsity-inducing penalties, which are leveraged in the following sections to recover structured dynamics from data.

4.2 Least-Squares Regression Solving

The focus now shifts to solving the vector least-squares problem introduced in (4.4). Although this problem appears elementary, it plays a central role in the structure identification procedure, forming the basis for all sparsity-aware methods developed in the following sections. The problem can be addressed either through a closed-form solution or by means of iterative optimization techniques, each approach offering distinct trade-offs in terms of scalability, accuracy, and flexibility.

4.2.1 Theoretical Analysis

Prior to introducing specific solution methods, it is useful to examine the main theoretical properties of the objective function

$$f(y) = \|z - \Phi y\|_2^2.$$

This function is convex, as it results from the composition of the squared Euclidean norm, which is convex, with an affine transformation. Consequently, any local minimizer is also a global minimizer. This property guarantees convergence for a broad class of iterative algorithms, including those based on gradient descent.

The function is also differentiable and admits closed-form expressions for both its gradient and Hessian:

$$\nabla f(y) = 2\Phi^\mathsf{T}(\Phi y - z), \qquad \nabla^2 f(y) = 2\Phi^\mathsf{T}\Phi.$$

The Hessian is constant, symmetric, and positive semi-definite, confirming that the function is a convex quadratic function. When $\Phi^\mathsf{T}\Phi$ has full rank, the function becomes strictly convex, thereby guaranteeing the uniqueness of the global minimizer. This condition may be ensured under the so-called persistence of excitation property of the data [25]. To formalize the regularity of the gradient, the following definition is recalled.

Definition 4.1 (Lipschitz gradient [64])**.** Let $f : \mathbb{R}^n \to \mathbb{R}$ be a twice continuously differentiable function. If there exists a constant $L > 0$ such that

$$\|\nabla f(y_1) - \nabla f(y_2)\|_2 \leq L\|y_1 - y_2\|_2 \quad \text{for all } y_1, y_2 \in \mathbb{R}^n,$$

then the gradient ∇f is said to be *L-Lipschitz continuous*.

In the present setting, since the Hessian $\nabla^2 f(y)$ is constant, the gradient is globally Lipschitz continuous with constant:

$$L = \|\nabla^2 f(y)\|_2 = 2\|\Phi^\mathsf{T}\Phi\|_2.$$

This Lipschitz constant L plays a central role in selecting admissible step sizes and establishing convergence rates for iterative algorithms such as gradient methods.

4.2.2 Analytical Solution

The least-squares objective function

$$f(y) = \|z - \Phi y\|_2^2$$

admits a unique closed-form minimizer when the matrix $\Phi^\mathsf{T}\Phi$ is invertible. This solution is given by the normal equations, a classical result in linear estimation theory (see, e.g., [65]):

$$y^* = (\Phi^\mathsf{T}\Phi)^{-1}\Phi^\mathsf{T} z. \tag{4.5}$$

This expression yields the vector $y^* \in \mathbb{R}^{n(n+m)}$ that minimizes the squared error between the predicted and observed values. In practice, it is often preferable to avoid computing the inverse explicitly for reasons related to numerical stability. Instead, the normal equations

$$\Phi^\mathsf{T}\Phi\, y = \Phi^\mathsf{T} z, \tag{4.6}$$

are typically solved using factorization techniques, such as Cholesky decomposition when $\Phi^\mathsf{T}\Phi$ is positive definite, or QR decomposition when improved numerical robustness is required. If $\Phi^\mathsf{T}\Phi$ is not of full rank, the inverse does not exist. In such cases, the minimum-norm solution is obtained using the Moore–Penrose pseudoinverse, denoted by $(\Phi^\mathsf{T}\Phi)^\dagger$, leading to:

$$y^* = (\Phi^\mathsf{T}\Phi)^\dagger\Phi^\mathsf{T} z. \tag{4.7}$$

This analytical method is computationally efficient for systems of moderate dimension and serves as a useful baseline. However, its applicability decreases in high-dimensional settings or when structural constraints, such as sparsity, must be incorporated. In these cases, iterative solvers offer greater flexibility, both to accommodate larger systems and to integrate regularization terms or structural priors, as discussed in the next subsection.

4.2.3 Gradient Descent

As an alternative to the closed-form solution, the least-squares problem can be solved iteratively using gradient descent. Although this method is generally slower in idealized settings, it provides significant advantages in terms of flexibility and scalability. In particular, it becomes especially relevant when introducing regularization terms, such as sparsity-promoting penalties, which disrupt the separability exploited by analytical solutions. Based on the gradient expression established earlier, the standard gradient descent update is given by (see, e.g., [64]):

$$y_{k+1} = y_k - \frac{1}{L}\nabla f(y_k) = y_k - \frac{1}{2\|\Phi^\mathsf{T}\Phi\|_2}\,\Phi^\mathsf{T}(\Phi y_k - z),$$

where $L = 2\|\Phi^\top \Phi\|_2$ denotes the Lipschitz constant of the gradient. The iterations are typically terminated once the gradient norm falls below a prescribed tolerance:

$$\|\nabla f(y_k)\|_2 \leq \varepsilon,$$

where $\varepsilon > 0$ is a threshold specified by the user. In the ideal case where the data are noise-free and the model matches a linear system, the loss $f(y_k)$ can, in theory, converge to zero. More importantly, this iterative scheme provides a computational framework that can be naturally extended to structured recovery tasks.

4.2.4 Noiseless Case

In the absence of noise, a key question is whether the system matrices A_s and B_s, or at least their structure, can be uniquely identified from input–state trajectories. Uniqueness is essential, as it determines whether the recovered structure is the only one consistent with the data. As established in Theorem 2.6, a necessary and sufficient condition for identifiability in the noiseless setting is:

$$\operatorname{rank} \begin{bmatrix} X_- \\ U_- \end{bmatrix} = n + m. \tag{4.8}$$

This condition guarantees that the combined past states and inputs span the full parameter space, thereby enabling the unique recovery of A_s and B_s through the least-squares formulation introduced in (4.3).

To assess the practical implications of this identifiability condition, a Monte Carlo simulation is conducted to estimate the empirical probability of unique identification as a function of two key parameters: the number of observations T, and the sparsity level k. The sparsity k is defined as the total number of zero entries in the concatenated system matrix $[A_s\ B_s] \in \mathbb{R}^{n\times(n+m)}$. By construction, this quantity satisfies $0 \leq k \leq n(n+m)$, where $k = 0$ corresponds to a fully dense system, and $k = n(n+m)$ corresponds to a completely zero system.

In the simulations, each entry of A_s and B_s is independently sampled from the uniform distribution over the interval $[-1, 1]$, ensuring that the eigenvalues of A_s lie within the unit disk. This avoids instability from divergent dynamics and, although not required for identifiability, simplifies the numerical setting. To impose a prescribed sparsity level k, k positions are randomly selected in the matrix $[A_s\ B_s]$, and the corresponding entries are set to zero. The selection of zero positions is randomized independently in each trial, allowing a broad range of structural configurations to be sampled for each fixed value of k.

The influence of the number of observations is first examined. The data matrices $X_- \in \mathbb{R}^{n\times T}$ and $U_- \in \mathbb{R}^{m\times T}$ each contain T columns. Consequently, the rank of their vertical concatenation is upper-bounded by T:

$$\operatorname{rank}\begin{bmatrix} X_- \\ U_- \end{bmatrix} \leq T.$$

Therefore, if $T < n + m$, the rank condition for identifiability necessarily fails:

$$\operatorname{rank}\begin{bmatrix} X_- \\ U_- \end{bmatrix} < n + m,$$

and the identification problem becomes underdetermined. This reflects the fact that the number of independent data points is insufficient to recover the $n(n + m)$ parameters that define the system.

Next, the effect of sparsity is considered. When most entries in A_s and B_s are zero, the system exhibits fewer effective degrees of freedom, and certain states may display limited or no variation over time. This can lead to a loss of identifiability, even when $T \geq n + m$. To quantify this effect, small-scale systems with $n = 3$ and $m = 1$ are simulated. For each pair (k, T), 100,000 random systems are generated. The proportion of uniquely identifiable systems is reported in Figure 4.1.

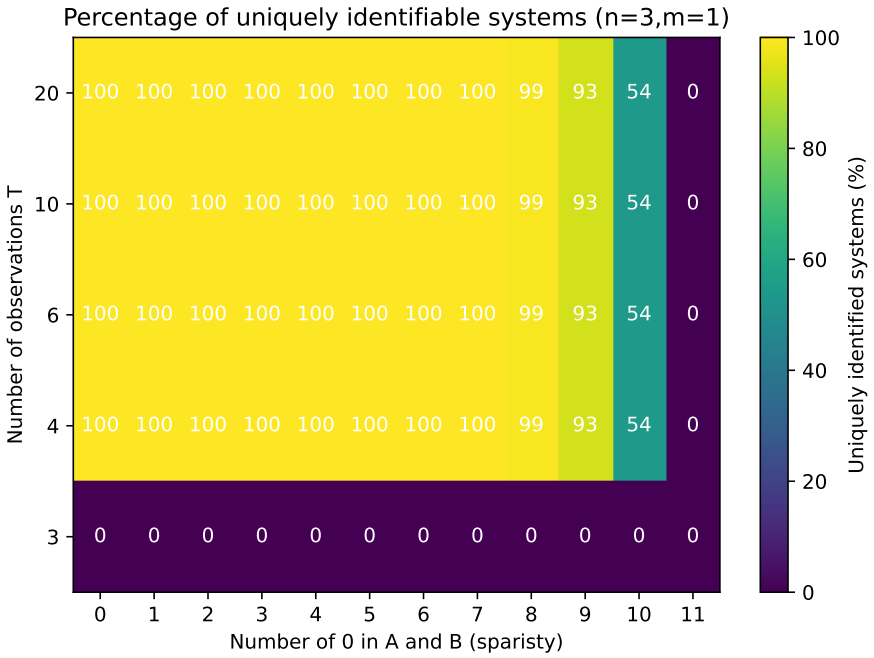

Figure 4.1: Percentage of uniquely identifiable systems as a function of the number of observations T and the sparsity level k.

The simulation results reveal several notable patterns. First, for any fixed sparsity level k, the proportion of uniquely identifiable systems saturates once the number of observations reaches $T = n + m$. Beyond this threshold, additional data have no further impact on identifiability. For example, when $k = 9$, approximately 54% of the systems are uniquely identifiable, and this proportion remains unchanged for all $T \geq n + m$. This confirms that, in the noiseless regime, increasing the number of observations beyond the minimal rank condition is unnecessary.

Second, as the system becomes sparser (i.e., as k increase), the probability of unique identification decreases. A sparsity threshold exists beyond which uniqueness is never guaranteed, regardless of the number of observations. In the present study, this threshold is observed around

$$k = (n - 2)(n + m),$$

which corresponds to the minimal number of zero entries required for certain pathological non-identifiable structures to arise.

Closer inspection of the non-unique cases reveals two recurrent structural patterns. In all observed failures, the system falls into one of two configurations. In the first case, at least two of the three state variables exhibit zero dynamics: their corresponding rows in both A_s and B_s are entirely zero. This configuration is referred to as *Form I*.

Example 4.2 (Form I). In this example, the dynamics of the first and third states are completely suppressed:

$$A_s = \begin{bmatrix} 0 & 0 & 0 \\ 0.1 & 0.85 & 0.58 \\ 0 & 0 & 0 \end{bmatrix}, \quad B_s = \begin{bmatrix} 0 \\ 0.2 \\ 0 \end{bmatrix}.$$

Because the trajectories of the first and third states remain constant over time, the corresponding rows in X_- are linearly dependent, and the data matrix fails to satisfy the rank condition. Consequently, the system is not uniquely identifiable.

The second failure mode arises when the entire state matrix A_s is zero. In this configuration, the system is purely input-driven, and state transitions are independent of prior state values. This case is referred to as *Form II*.

Example 4.3 (Form II). Here, the system evolves entirely due to inputs:

$$A_s = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad B_s = \begin{bmatrix} -0.98 \\ 0.2 \\ -0.54 \end{bmatrix}.$$

In this configuration, the matrix involved in the identifiability rank condition has linearly dependent columns, except possibly for the initial state. Since A_s is identically zero, all future state transitions are driven solely by the inputs. As a result, the trajectory data lack internal variability, and the rank condition is violated.

Both Forms I and II correspond to structurally degenerate systems that are trivially uncontrollable. Notably, Form I requires at least $(n-2)(n+m)$ zero entries in the system matrices, whereas Form II requires the entire matrix A_s to be zero, that is, n^2 zeros entries. Since

$$(n-2)(n+m) < n^2 \quad \text{for all } n \geq 3,$$

Form I occurs more frequently in random sampling. This structural asymmetry accounts for the predominance of Form I among non-identifiable systems observed in the simulations. These empirical observations support the following conclusions:

- If $k < (n-2)(n+m)$, the system is almost always uniquely identifiable, except for a few specially constructed counterexamples (see Example 2.10).
- If $k \geq (n-2)(n+m)$, all observed non-identifiable systems fall into Form I or II and are thus trivially uncontrollable.

These observations are encouraging: they indicate that the complete methodology, consisting of structure identification followed by structural controllability testing, is highly reliable in the noiseless regime, provided that the number of observations satisfies $T \geq n+m$. This hypothesis will be examined more extensively in later sections.

Finally, although it is theoretically possible for a system to admit both a controllable and an uncontrollable realization consistent with the same data (see Example 2.10), such ambiguous cases are exceedingly rare. In the simulations, which covered a wide range of sparsity levels and system realizations, no such cases were observed. These degenerate systems, which do not fall into Form I or II, remain non-identifiable but exhibit a subtler form of ambiguity. Additional experiments on higher-dimensional systems confirming this trend are provided in Appendix B.

To further validate the above findings and assess the behavior of the identification procedure under practical conditions, the performance of gradient descent is evaluated on the same class of systems with $n = 3$ and $m = 1$. Since empirical identifiability does not improve beyond $T = n + m$, the analysis is restricted to the critical range $T \leq n + m$, that is, when the number of observations is below the minimum required for identifiability.

For each pair (k, T), 100 independent random systems are generated, and the least-squares identification problem is solved using gradient descent. The accuracy of structure recovery is defined as the ratio of systems for which the recovered structure matches the ground truth, over the total number of systems tested. Results are presented in Figure 4.2, which illustrates how identifiability degrades with increasing sparsity and decreasing data volume.

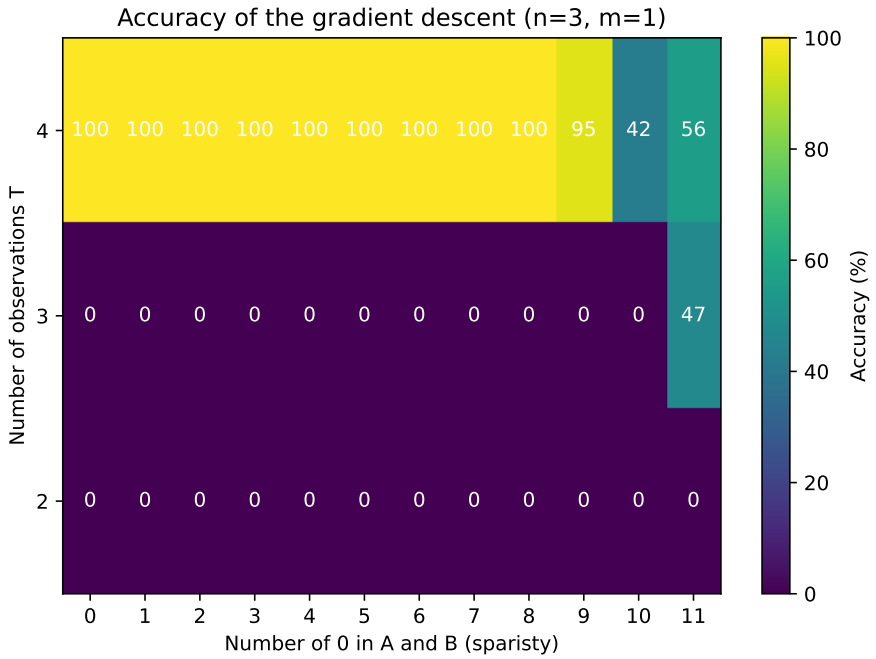

Figure 4.2: Structure identification accuracy of the gradient descent method across different sparsity levels k and observation lengths T.

The results confirm the observations made in the previous section. When the number of observations reaches the theoretical threshold $T = n + m$, and the system is not excessively sparse (specifically, when $k < (n-2)(n+m)$), the gradient descent algorithm consistently recovers the correct sparsity pattern, achieving perfect identification accuracy across all trials.

As the sparsity level increases beyond the threshold $k \geq (n-2)(n+m)$, recovery performance degrades. In these regimes, the systems contain few non-zero parameters, making it harder for the algorithm to detect active components reliably. In some cases, the optimization fails to converge to the correct solution. Nevertheless, the method often returns sparse estimates that remain structurally compatible with the ground truth, preserving the reliability of subsequent controllability tests.

In contrast, when $T < n + m$, the algorithm systematically fails to identify the correct structure. This aligns with the theoretical identifiability condition, which requires at least $n + m$ independent observations to uniquely determine the system matrices. This limitation is not a drawback of the algorithm itself, but rather a fundamental constraint of the problem. From a practical perspective, the requirement $T \geq n + m$ is a mild assumption that is typically satisfied in realistic data collection scenarios.

4.2.5 Noisy Case

The performance of the gradient descent algorithm is now examined in the presence of measurement or process noise. In this setting, the system evolves according to perturbed linear dynamics:

$$x(t+1) = A_s\, x(t) + B_s\, u(t) + w(t), \tag{4.9}$$

where $w(t) \in \mathbb{R}^n$ is an additive noise term modeled as white Gaussian noise, with

$$w(t) \sim \mathcal{N}(0, \Sigma_w^2 I_n).$$

Each component $w_i(t)$ is drawn independently from a zero-mean normal distribution with standard deviation $\Sigma_w > 0$, identical across all state dimensions.

To illustrate the impact of noise on structure recovery, a representative example is considered involving a system with a known sparsity pattern.

Example 4.4 (Effect of noise on structure recovery). The true system matrices A_s and B_s used to generate the data are:

$$A_{\text{s}} = \begin{bmatrix} 0.886 & 0 & 0.952 \\ 0 & 0.214 & -0.247 \\ 0.603 & 0 & 0 \end{bmatrix}, \quad B_{\text{s}} = \begin{bmatrix} 0.087 \\ 0.804 \\ 0 \end{bmatrix}.$$

Entries highlighted in blue denote exact zeros in the ground-truth system and define the underlying sparsity pattern. When data are generated without noise ($w(t) = 0$), gradient descent successfully recovers the correct structure without error.

Now consider the same system with additive white Gaussian noise of standard deviation $\Sigma_w = 0.001$. The recovered matrices are:

$$A_{\text{noise}} = \begin{bmatrix} 0.875 & 0.001 & 0.945 \\ 0.002 & 0.201 & -0.255 \\ 0.626 & 0.003 & 0.018 \end{bmatrix}, \quad B_{\text{noise}} = \begin{bmatrix} 0.086 \\ 0.813 \\ 0.007 \end{bmatrix}.$$

Entries shown in red indicate small non-zero values appearing at positions corresponding to exact zeros in the original system. These deviations, caused by noise, illustrate a key limitation of the unconstrained gradient descent approach: it is sensitive to perturbations and does not enforce sparsity in the estimated matrices. As a result, even minimal noise can introduce spurious entries that obscure the true structure.

Despite this sensitivity, the sparsity pattern remains visually apparent. Most originally zero entries are only slightly perturbed, and the dominant non-zero components are preserved in both location and magnitude.

This observation is encouraging. It suggests that although gradient descent alone does not guarantee exact structure recovery under noise, its output remains sufficiently close to the true system to serve as a starting point. The next section introduces a post-processing procedure that enables recovery of the correct sparsity pattern, even in the presence of noise.

4.3 Least-Squares Method with Thresholding

As shown in Example 4.4, in the presence of noise, the gradient descent solution may include small spurious entries at positions that should be structurally zero. A natural way to improve robustness is to eliminate these small-magnitude terms through a thresholding operation. The underlying idea is that true nonzero coefficients typically have larger magnitude than noise-induced artifacts. This motivates the use of *hard thresholding*, a simple yet effective technique widely employed in compressive sensing, particularly for sparse signal recovery [66].

4.3.1 Method Description

The thresholding procedure acts purely as a post-processing step: it is applied after solving the least-squares problem via gradient descent. No additional optimization is required. Given the solution vector y obtained from gradient descent, the $\mathcal{K}$ smallest entries in absolute value are set to zero, where $\mathcal{K} \in \{0, 1, \ldots, n(n+m)\}$ is a user-defined sparsity parameter. The remaining entries are left unchanged. For illustration, consider again the system from Example 4.4. When hard thresholding with $\mathcal{K} = 11$ is applied in the noisy case, the result is:

$$A_{\text{thr}} = \begin{bmatrix} 0 & 0 & 0.945 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \qquad B_{\text{thr}} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}.$$

Here, the 11 smallest absolute values in the vectorized form of $(A_{\text{noise}}, B_{\text{noise}})$ have been set to zero. This configuration is overly aggressive and removes relevant dynamics. In contrast, using $\mathcal{K} = 5$ yields:

$$A_{\text{thr}} = \begin{bmatrix} 0.875 & 0 & 0.945 \\ 0 & 0.201 & -0.255 \\ 0.626 & 0 & 0 \end{bmatrix}, \qquad B_{\text{thr}} = \begin{bmatrix} 0.086 \\ 0.813 \\ 0 \end{bmatrix}.$$

This threshold recovers the exact sparsity pattern of the original system, demonstrating the ability of hard thresholding to correct noise-induced errors when an appropriate value of $\mathcal{K}$ is selected.

A key practical challenge lies in selecting the appropriate threshold level $\mathcal{K}$. While the optimal value can often be identified visually in small-scale examples, defining a principled and automatic selection criterion remains non-trivial in general. The next section introduces strategies to determine an effective threshold $\mathcal{K}^*$ that ensures robust structure recovery under noisy conditions.

4.3.2 Hyperparameter Selection

Selecting an appropriate value of $\mathcal{K}$, denoted $\mathcal{K}^*$, is critical to the effectiveness of the thresholding procedure. This parameter directly governs the sparsity of the estimated model and therefore influences both the interpretability and the accuracy of the recovered structure. Under mild conditions, gradient descent yields a solution close to the true system, and a well-chosen value of $\mathcal{K}$ enables accurate recovery through thresholding.

Several strategies for selecting $\mathcal{K}$ have been proposed in the literature [67], each balancing sparsity and data fidelity in a distinct way. Here, a pragmatic approach is adopted: the sparsest thresholded model that still explains the observed data with sufficient accuracy is selected. More precisely, the following definitions are introduced:

- $\mathcal{H}_{\mathcal{K}}(y)$: the hard-thresholding operator that sets to zero the $\mathcal{K}$ smallest entries of y in absolute value;
- $\text{RMSE}_{\mathcal{K}} = \dfrac{1}{\sqrt{N}} \left\| z - \Phi \mathcal{H}_{\mathcal{K}}(y) \right\|_2$: the root mean square error associated with the thresholded model, where $N = nT$ is the lenght of the output vector z.

The RMSE provides a normalized measure of reconstruction error that accounts for both the number of time steps T and the system dimension n.

To ensure numerical stability and impose a smooth penalty on poor reconstructions, the selection criterion is defined using the negative exponential of the RMSE:

$$\mathcal{K}^* = \max\left\{\mathcal{K} \ : \ \mathrm{e}^{-\mathrm{RMSE}_{\mathcal{K}}} \geq \tau\right\}, \tag{4.10}$$

where $\tau \in (0, 1)$ is a tolerance parameter that controls the acceptable level of reconstruction error. The exponential function approaches 1 for models that closely fit the data, and decreases rapidly for models with high error. The resulting $\mathcal{K}^*$ corresponds to the sparsest model that preserves an adequate level of data fidelity.

Example 4.5 (Threshold selection for known sparsity). The system defined in Example 4.4 is revisited, with ground-truth matrices given by:

$$A_{\mathrm{s}} = \begin{bmatrix} 0.886 & 0 & 0.952 \\ 0 & 0.214 & -0.247 \\ 0.603 & 0 & 0 \end{bmatrix}, \qquad B_{\mathrm{s}} = \begin{bmatrix} 0.087 \\ 0.804 \\ 0 \end{bmatrix}.$$

This system contains exactly five structural zeros. For each candidate value of $\mathcal{K}$, the corresponding value of $\mathrm{e}^{-\mathrm{RMSE}_{\mathcal{K}}}$ ic computed and plotted in Figure 4.3. The pink horizontal line represents tolerance threshold $\tau = 0.9995$, and the vertical purple line indicates the optimal sparsity level $\mathcal{K}^* = 5$.

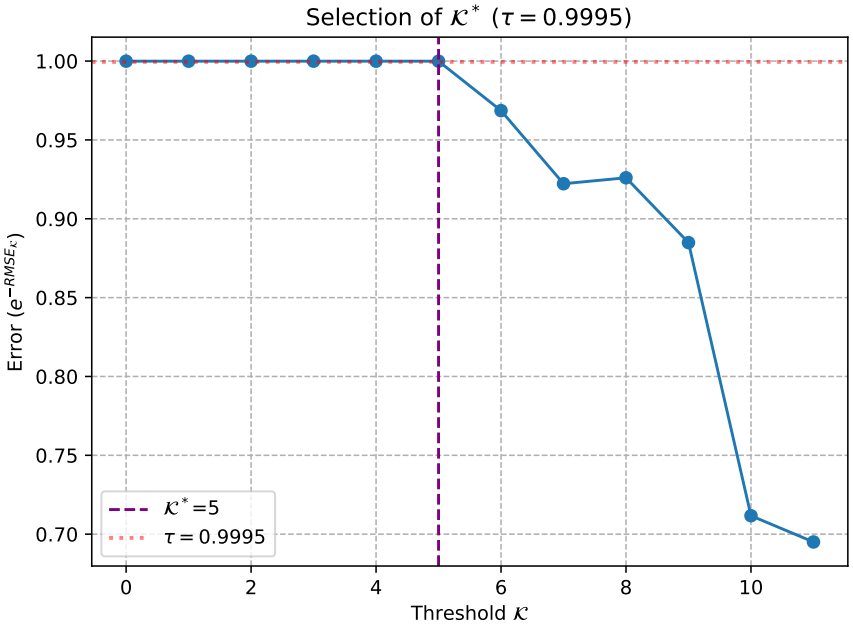

Figure 4.3: Exponential error metric $\mathrm{e}^{-\mathrm{RMSE}_{\mathcal{K}}}$ as a function of the threshold level $\mathcal{K}$ with $\tau = 0.9995$.

In this case, the selection rule correctly identifies the optimal sparsity level: $\mathcal{K}^* = 5$ is the largest value for which the reconstruction error remains above the threshold τ. Increasing the sparsity further, for instance by setting $\mathcal{K} = 6$, leads to a noticeable degradation in performance, indicating that relevant information has been removed. The detailed algorithm used to automatically select $\mathcal{K}^*$ is provided in Appendix A.

The effectiveness of this criterion depends on the choice of the threshold parameter τ, which controls the trade-off between data fidelity and model sparsity: higher values favor accurate reconstruction with denser models, while lower values encourage sparser solutions at the risk of underfitting.

To assess the influence of the threshold parameter τ, Monte Carlo simulations are performed across various sparsity levels and noise intensities. For each pair (k, τ), where k represents the number of zero entries in the true system matrices, 100 random systems are generated, and structure identification is carried out via gradient descent followed by hard thresholding. The resulting accuracy is reported in the heatmaps below. Figure Figure 4.4 shows results for low noise ($\Sigma_w = 0.0001$), and Figure 4.5 to higher noise ($\Sigma_w = 0.001$).

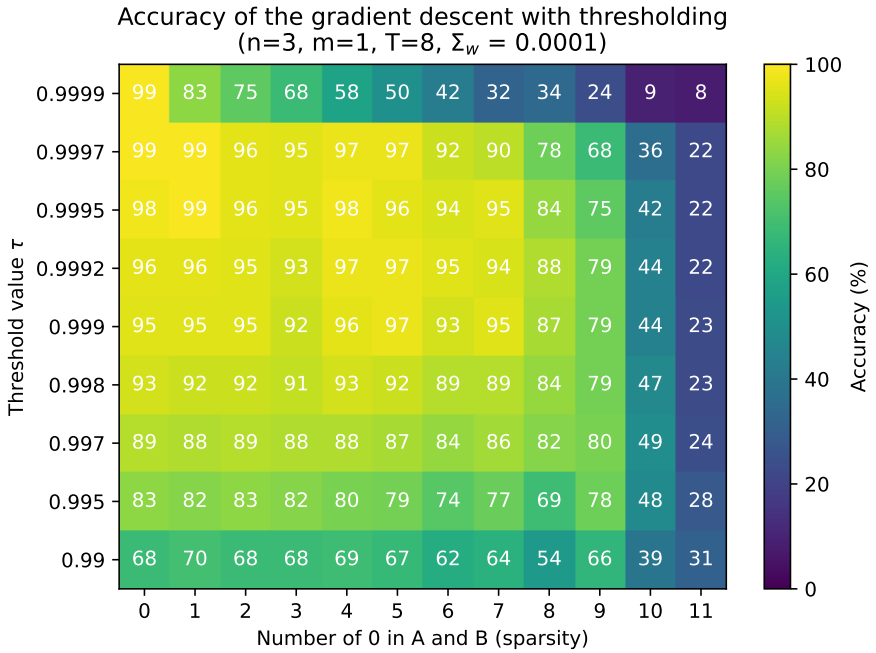

Figure 4.4: Structure identification accuracy as a function of τ for different sparsity levels k, with noise level $\Sigma_w = 0.0001$.

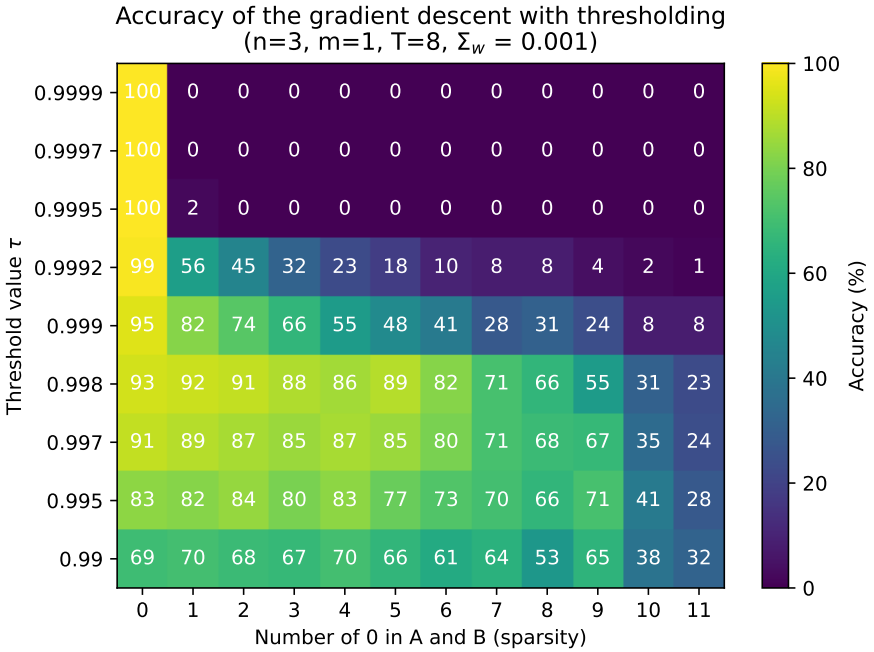

Figure 4.5: Structure identification accuracy as a function of τ for different sparsity levels k, with noise level $\Sigma_w = 0.001$.

The results in Figure 4.4 and Figure 4.5 highlight several trends. In low-noise regimes, higher thresholds such as $\tau = 0.9995$ yield the best results, especially for denser systems. As noise increases, however, this value becomes overly strict, causing relevant entries to be erroneously removed. Lower thresholds, like $\tau = 0.998$, offer more robust performance across a wide range of sparsity and noise levels.

Since the thresholding step is lightweight and purely post-processing, it is feasible to evaluate several values of τ in parallel. In the remainder of this work, the value $\tau = 0.9995$ is used by default in low-noise settings, while $\tau = 0.998$ serves as a robust alternative under noisier conditions.

Additional experiments reported in Appendix B confirm that this selection strategy generalizes well to higher-dimensional systems, with a natural decrease in accuracy due to increased estimation complexity. Although exact recovery of the true sparsity level $\mathcal{K}^*$ is not always achievable, particularly in the presence of noise with heavy-tailed perturbations, the thresholding procedure remains a principled and effective approach to structural recovery in practice.

4.3.3 Empirical Analysis

The effect of the number of observations T on structure identification accuracy is first examined. While $T = n + m$ is sufficient for unique recovery in the noiseless setting, the optimal number of observations under noise may differ. Understanding this dependence is essential for reliable practical application. The results for a fixed noise level are shown in Figure 4.6.

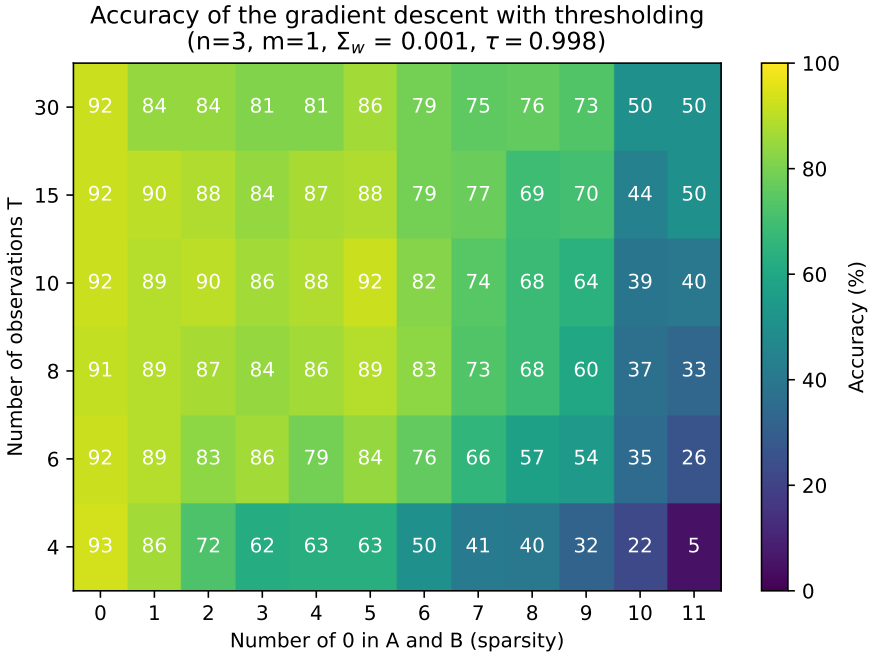

Figure 4.6: Structure identification accuracy of the gradient–descent with thresholding method as a function of T, for different sparsity levels k with noise level $\Sigma_w = 0.001$ and threshold $\tau = 0.998$.

Unlike the noiseless case, where $T = n + m$ suffices for identifiability, the presence of noise makes this threshold suboptimal. Accuracy improves substantially as T increases from 4 to 6 and continues to rise up to around $T = 15$. Beyond that, however, the performance plateaus or even slightly deteriorates.

This decline is due to the accumulation of noise over time. As T grows, each state becomes affected by a longer sequence of noise inputs, increasing the variance of $x(t)$. This makes the identification problem more ill-conditioned. In particular, for very large values such as $T = 30$, the trajectories become dominated by noise, making it harder to recover the true structure.

This effect is especially pronounced for sparse or moderately sparse systems. In contrast, hypersparse systems (although of limited practical interest) appear to benefit more from larger observation horizons. These observations confirm that extending the time horizon excessively tends to be detrimental.

To assess whether similar trends persist under lower noise levels, the analysis is repeated with a higher threshold $\tau = 0.9995$. The results are presented in Figure 4.7.

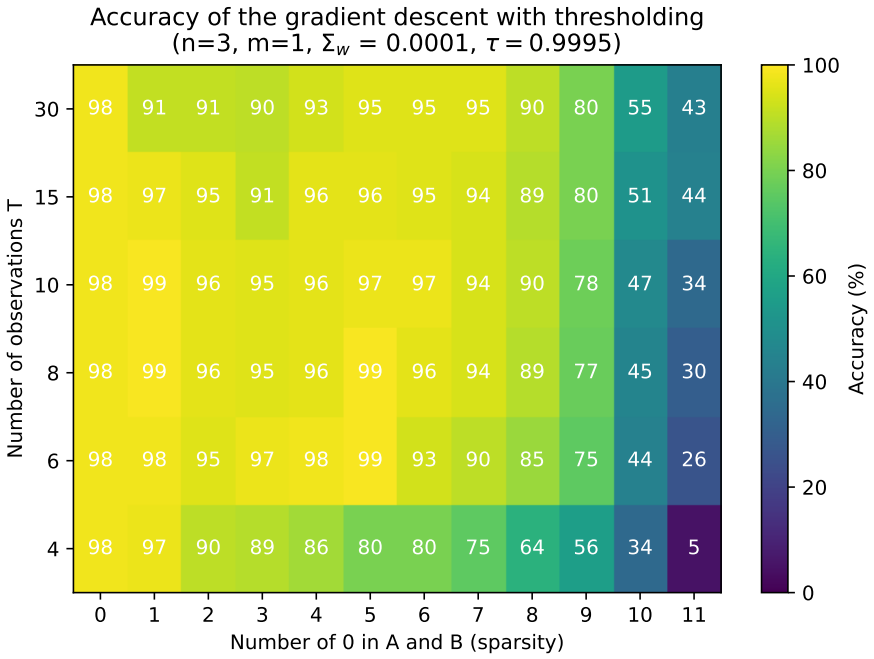

Figure 4.7: Structure identification accuracy of the gradient–descent with thresholding method as a function of T, for different sparsity levels k with noise level $\Sigma_w = 0.0001$ and threshold $\tau = 0.9995$.

The same qualitative behavior is observed: performance improves rapidly as T increases beyond $n + m$, but tends to plateau afterwards. When T becomes too large, a modest drop in accuracy is again observed due to the accumulation of noise.

Based on these observations, the number of observations is fixed to $T = 2(n + m)$ for the remaining experiments. This value offers a practical compromise: large enough to ensure identifiability, yet small enough to limit the effect of noise. In this setting, it corresponds to $T = 8$. The impact of noise magnitude Σ_w on identification accuracy is now evaluated.

To evaluate the impact of noise on structure identification, simulations are performed using the fixed settings from previous sections ($T = 2(n+m)$, $\tau \in \{0.998, 0.9995\}$). The parameter Σ_w denotes the standard deviation of the additive Gaussian noise, with higher values corresponding to increased noise levels.

For each pair (k, Σ_w), 100 random systems are generated, and structure identification is performed using gradient descent followed by hard thresholding. The resulting accuracy is reported in Figure 4.8.

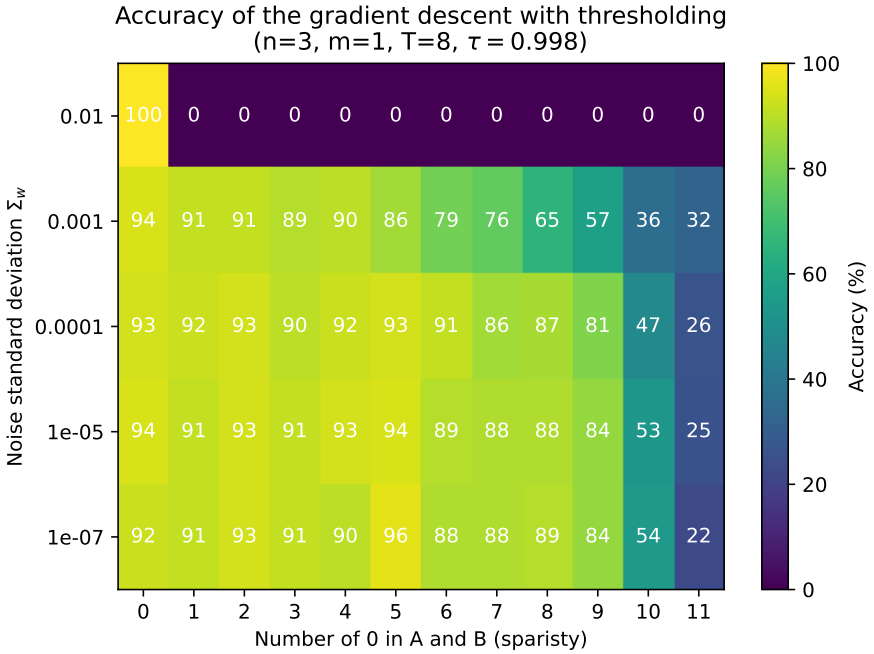

Figure 4.8: Structure identification accuracy of the gradient–descent with thresholding method as a function of noise level Σ_w, for different sparsity levels k with horizon $T = 8$ and threshold $\tau = 0.998$.

Figure 4.8 indicates that, with $\tau = 0.998$, the method maintains high accuracy across a broad range of noise levels. Accuracy decreases, however, for large noise magnitudes (e.g., $\Sigma_w = 0.01$), where the perturbations significantly affect the dynamics. As previously observed, hypersparse systems remain more difficult to identify accurately.

To assess whether these patterns hold under a more stringent threshold, results obtained with $\tau = 0.9995$ are reported in Figure 4.9.

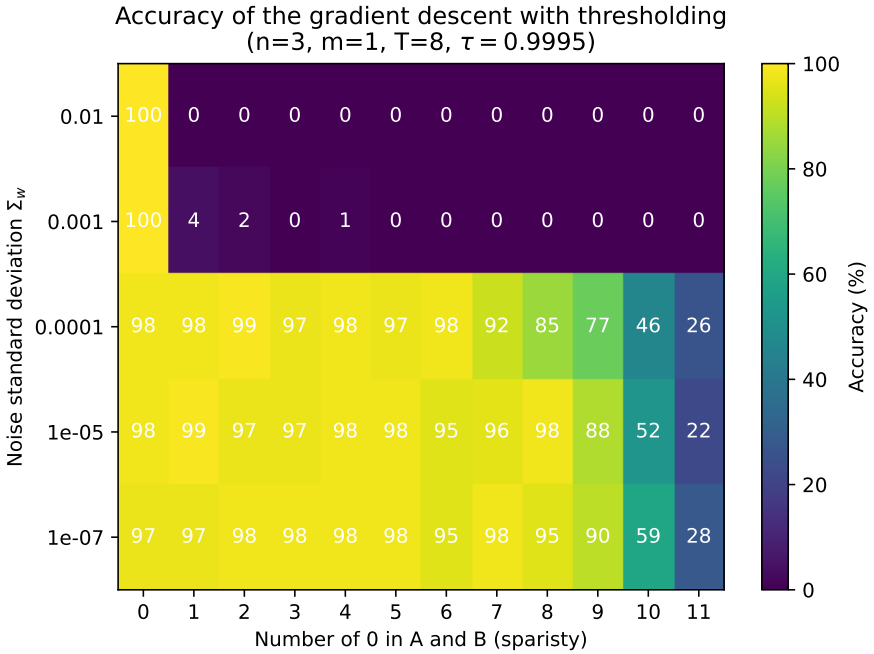

Figure 4.9: Structure identification accuracy of the gradient–descent with thresholding method as a function of noise level Σ_w, for different sparsity levels k with horizon $T = 8$ and threshold $\tau = 0.9995$.

Figure 4.9 shows that applying a stricter threshold ($\tau = 0.9995$) yields excellent results in low-noise conditions but leads to a marked decline in accuracy as the noise magnitude exceeds $\Sigma_w = 0.001$. This threshold is appropriate when prior knowledge indicates low levels of noise but may be too limiting in general. In settings where the noise level is unknown or expected to vary, a more moderate value of τ provides improved robustness.

In summary, the hard-thresholding approach demonstrates strong empirical performance under noisy conditions. With a moderate number of observations and for small-scale systems, it achieves over 90 % accuracy in structure recovery, provided the threshold τ is selected appropriately. These results support the use of post-processing via thresholding to mitigate the sensitivity of gradient descent to small perturbations, without incurring significant computational cost.

Additional experiments for higher-dimensional systems are presented in Appendix B. While the same qualitative trends are observed, accuracy declines with increasing system size, emphasizing the need for improved scalability.

4.3.4 Error Analysis

To better understand the behavior of the method in cases where it fails to recover the correct sparsity level $\mathcal{K}^*$, the nature of the errors is analyzed. Specifically, the objective is to assess whether the predicted sparsity $\mathcal{K}$ remains close to the true value or deviates significantly. Confusion matrices are computed for two representative scenarios. For each sparsity level k, 100 random systems are generated, and the corresponding predicted values of $\mathcal{K}$ are recorded following the thresholding step.

A first analysis is conducted under a robust configuration with moderate noise, where $\Sigma_w = 0.001$ and $\tau = 0.998$. The resulting confusion matrix is presented in Figure 4.10.

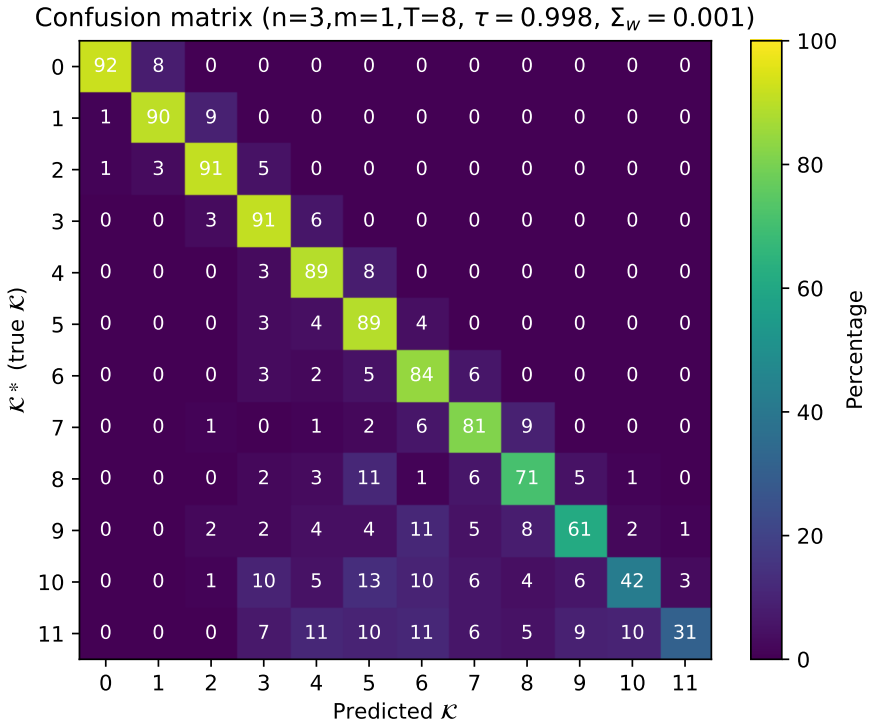

Figure 4.10: Confusion matrix for the gradient-descent with thresholding method, with $\Sigma_w = 0.001$ and $\tau = 0.998$.

For dense systems (i.e., small values of $\mathcal{K}^*$), the predicted sparsity $\mathcal{K}$ tends to remain tightly clustered around the true value. Even when misidentified, the deviations are minor, with most predictions located near the diagonal of the confusion matrix. As the true sparsity increases, the variability in predictions grows slightly, but remains relatively constrained for moderate levels of $\mathcal{K}^*$.

A more discriminative setting is then considered, characterized by very low noise with $\Sigma_w = 10^{-5}$ and a stricter threshold $\tau = 0.9995$. The corresponding results are shown in Figure 4.11.

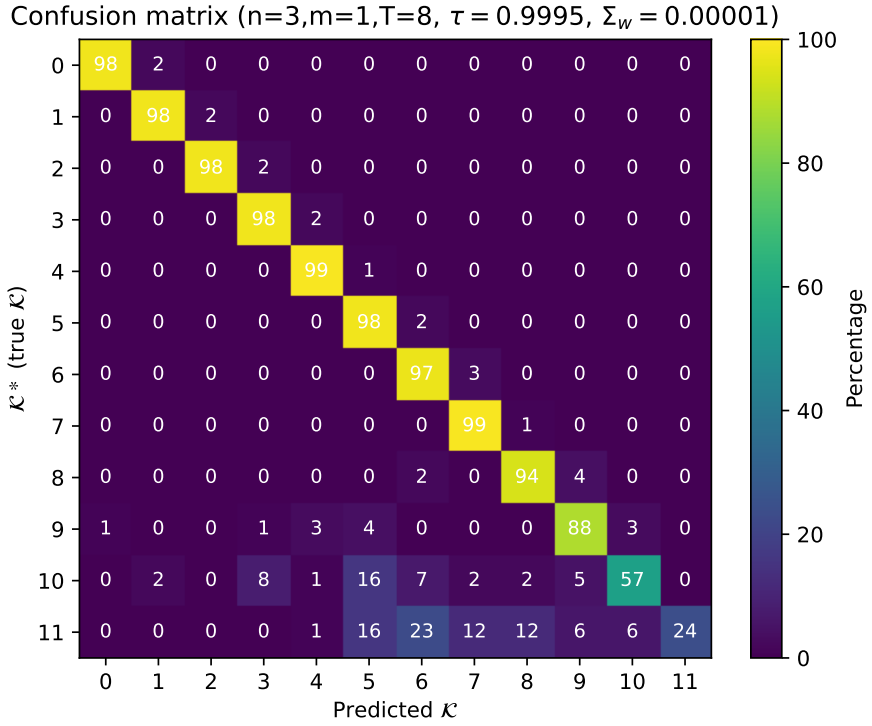

Figure 4.11: Confusion matrix for the gradient-descent with thresholding method, with $\Sigma_w = 10^{-5}$ and $\tau = 0.9995$.

In this low-noise regime, prediction accuracy improves further. For moderately sparse systems ($\mathcal{K}^* \in \{6, 7, 8, 9\}$), the predicted sparsity is highly concentrated along the diagonal, indicating consistent identification of the true structure. In contrast, for hypersparse systems ($\mathcal{K}^* \in \{10, 11\}$), the predictions become erratic, with values of $\mathcal{K}$ widely dispersed. This reflects the previously observed instability in these extreme cases.

These findings suggest that the method performs reliably when the underlying sparsity is moderate and noise remains limited. When errors occur, they tend to be of small magnitude, which is favorable for subsequent analysis tasks such as controllability assessment. As discussed in section 3.4, slight discrepancies in $\mathcal{K}$ rarely impact the controllability result. Performance degrades primarily in hypersparse regimes, which are of limited practical significance.

4.4 Gradient Descent with Regularization

To encourage sparsity directly within the optimization process, an alternative to hard-thresholding is to modify the objective function by introducing a regularization term. This approach leads to the class of penalized least-squares methods, among which the *LASSO* (Least Absolute Shrinkage and Selection Operator) is particularly well known [68].

The core idea is to favor solutions with fewer nonzero entries by penalizing the number of active coefficients in the parameter vector y, which includes both A and B. A direct formulation involves minimizing the following objective:

$$\min_{y \in \mathbb{R}^{n(n+m)}} \; f(y) = \|z - \Phi y\|_2^2 + \lambda \|y\|_0,$$

where $\|y\|_0$ counts the number of nonzero elements in y, and $\lambda > 0$ is a regularization parameter balancing reconstruction accuracy and sparsity.

The objective is to identify, among all solutions that explain the data well, the sparsest one. However, this formulation involves the ℓ_0-norm, making the problem combinatorial and NP-hard in general, except under strong assumptions.

To address this computational intractability, the ℓ_0 norm is commonly approximated by a convex surrogate. A popular and theoretically justified choice is the ℓ_1-norm, which leads to the LASSO formulation [69]:

$$\min_{y \in \mathbb{R}^{n(n+m)}} \; f(y) = \|z - \Phi y\|_2^2 + \lambda \|y\|_1. \tag{4.11}$$

This optimization problem is convex and can be solved efficiently using a variety of well-established algorithms. The ℓ_1 penalty encourages sparsity by shrinking small coefficients toward zero, effectively performing variable selection during the optimization. In the next sections, the algorithm used to solve this problem is described, followed by an evaluation of its performance on the structure identification task.

4.4.1 Method Description

The algorithm used to solve the regularized problem introduced in (4.11) relies on a composite optimization framework. The objective is written as:

$$\min_{y} \; f(y) = F(y) + \varphi(y),$$

where $F(y)$ is a smooth function with Lipschitz-continuous gradient, and $\varphi(y)$ is a convex, possibly non-smooth term [64]. In this context, the components are:

$$F(y) = \|z - \Phi y\|_2^2, \qquad \varphi(y) = \lambda \|y\|_1.$$

This formulation belongs to the class of composite non-smooth optimization problems [70]. A well-established approach to solving such problems is the *Proximal Gradient Method* [71], which updates the current iterate y_k as follows:

$$y_{k+1} = \text{prox}_{\frac{1}{L}\varphi}\left(y_k - \frac{1}{L}\nabla F(y_k)\right), \quad \forall k \geq 0. \tag{4.12}$$

The proximal operator is defined by [72]:

$$\text{prox}_h(x) = \arg\min_{x'}\left\{\frac{1}{2}\|x' - x\|_2^2 + h(x')\right\}.$$

In the case of ℓ_1-regularization, this yields the soft-thresholding rule [73], which is applied element-wise. Specifically, for each component:

$$\tilde{y}^i = y_k^i - \frac{1}{L}\left(\nabla F(y_k)\right)^i,$$

where $\tilde{y}^i \in \mathbb{R}$ denotes the ith component of the vector $\tilde{y} \in \mathbb{R}^{n(n+m)}$. Then, the updated coordinate y_{k+1}^i is given by:

$$y_{k+1}^i = \begin{cases} \tilde{y}^i - \frac{\lambda}{L}, & \text{if } \tilde{y}^i > \frac{\lambda}{L}, \\ 0, & \text{if } |\tilde{y}^i| \leq \frac{\lambda}{L}, \\ \tilde{y}^i + \frac{\lambda}{L}, & \text{if } \tilde{y}^i < -\frac{\lambda}{L}. \end{cases}$$

The iterations are stopped when the change between two successive iterates becomes small, as measured by the gradient mapping criterion [64]:

$$L\|y_{k+1} - y_k\|_2 \leq \varepsilon,$$

where $\varepsilon > 0$ is a tolerance set by the user and L the Lipschitz constant.

The performance of the LASSO estimator critically depends on the choice of λ [74]. If λ is too small, the penalty is insufficient to remove noisy entries; if it is too large, true non-zero coefficients may be eliminated. Since the optimal value depends on unknown quantities such as the true sparsity and noise level, it is difficult to determine it analytically.

A practical and widely-used strategy is to select λ via K-fold cross-validation [75]. The idea is to divide the data into K consecutive blocks of equal length. For each candidate λ, the following steps are repeated K times:

- In the training phase, one block is left out and the LASSO problem is solved on the remaining $K-1$ blocks. This produces a candidate solution $\hat{y}(\lambda)$.
- In the validation phase, the prediction error of $\hat{y}(\lambda)$ is evaluated on the held-out block, using the root-mean-square error as the performance metric.

Repeating this process over all blocks yields K validation errors per value of λ. Averaging these gives the cross-validation curve $\mathrm{CV}(\lambda)$, and the optimal regularization strength is selected as the value minimizing this curve, i.e., $\lambda^* = \arg\min_\lambda \mathrm{CV}(\lambda)$. In practice, we use $K = 3$ or $K = 5$, depending on the total number of observations T, to ensure that each training set remains sufficiently large. This cross-validation procedure provides a fully data-driven value of λ, tuned to the actual noise and sparsity levels without requiring prior knowledge. The simulations presented in the next section rely on the optimal λ^* obtained through this method.

4.4.2 Empirical Analysis

The behavior of the LASSO method under noisy conditions is assessed, with a particular focus on its sensitivity to the number of observations. For each pair (k, T), 100 random systems are simulated, and the regularization parameter λ is selected using 3-fold cross-validation. The results are summarized in Figure 4.12.

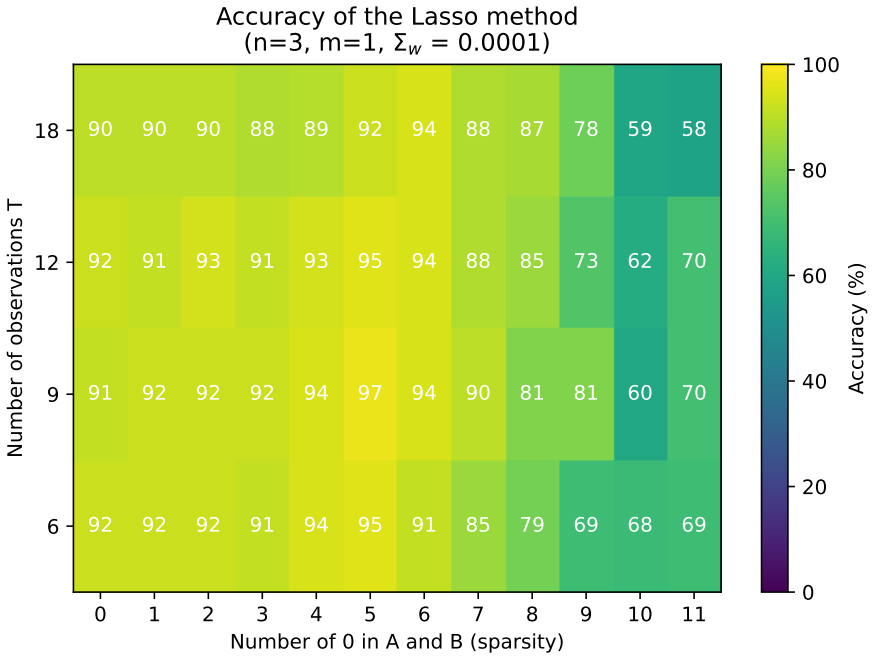

Figure 4.12: Structure identification accuracy of the LASSO method as a function of T, for different sparsity levels k and noise level $\Sigma_w = 0.0001$.

Figure 4.12 shows that the number of observations has limited impact on performance. While very small or large values of T may slightly reduce accuracy (e.g., $T = 18$), the method remains robust across the tested range.

To enable direct comparison with the hard-thresholding approach, the number of observations is fixed to $T = 3(n + m)$ for the remainder of the analysis. Cross-validation is performed using $K = 3$ contiguous blocks, ensuring that each training set contains $2(n + m)$ samples—consistent with previous experimental settings.

The performance of the LASSO method is evaluated under varying noise levels. For each pair (k, Σ_w), 100 random systems are generated, and structure identification is carried out using 3-fold cross-validation. The results are presented in Figure 4.13.

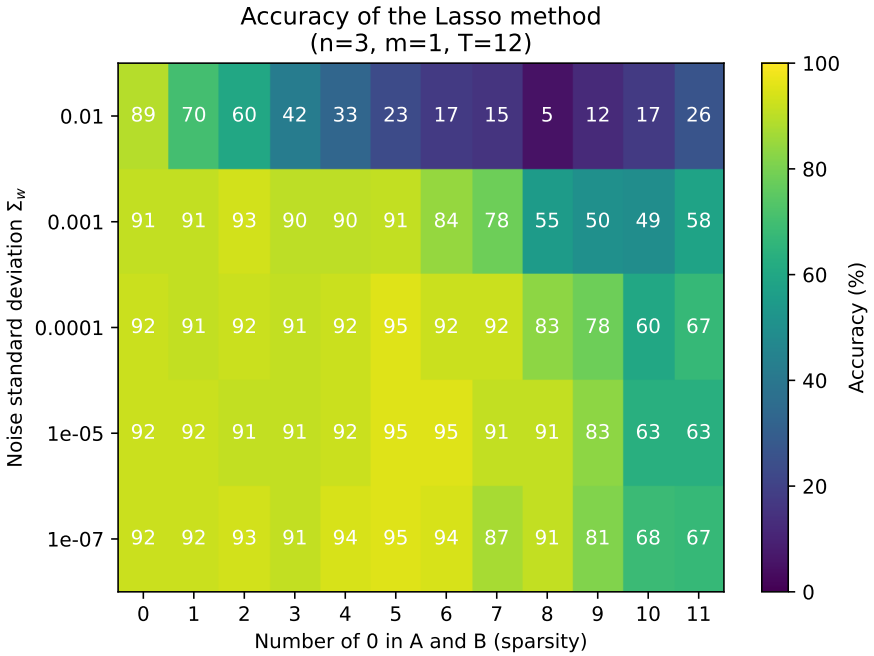

Figure 4.13: Structure identification accuracy of the LASSO method as a function of noise level Σ_w, for different sparsity levels k and $T = 12$.

Figure 4.13 shows that the method performs reliably across a wide range of noise levels. For most systems, it achieves over 90 % accuracy, and it handles hyper-sparse cases more effectively than the hard-thresholding approach. However, perfect accuracy is difficult to attain even under low noise. Achieving this would require using a very fine grid of λ values, effectively tuning the regularization for each individual instance. While possible, this significantly increases computational cost.

As already observed previously, at the highest noise level tested ($\Sigma_w = 0.01$), performance degrades slightly, but still remains superior to that of hard thresholding.

The LASSO framework has been introduced and evaluated in the presence of noise. Cross-validation allows for adaptive selection of the regularization parameter λ, resulting in accurate structure recovery in most scenarios. The next chapter extends the comparison between LASSO and hard thresholding to higher-dimensional systems and examines how both approaches affect controllability outcomes.

4.4.3 Error analysis

As with the hard-thresholding method, the deviation of LASSO from the true sparsity level $\mathcal{K}^*$ is analyzed when exact recovery fails. For each system, both the true number of zero entries and the number estimated by LASSO are recorded. These results are then aggregated into a confusion matrix.

For each sparsity level k, we simulate 100 random systems, using 3-fold cross-validation. The outcome is shown in Figure 4.14.

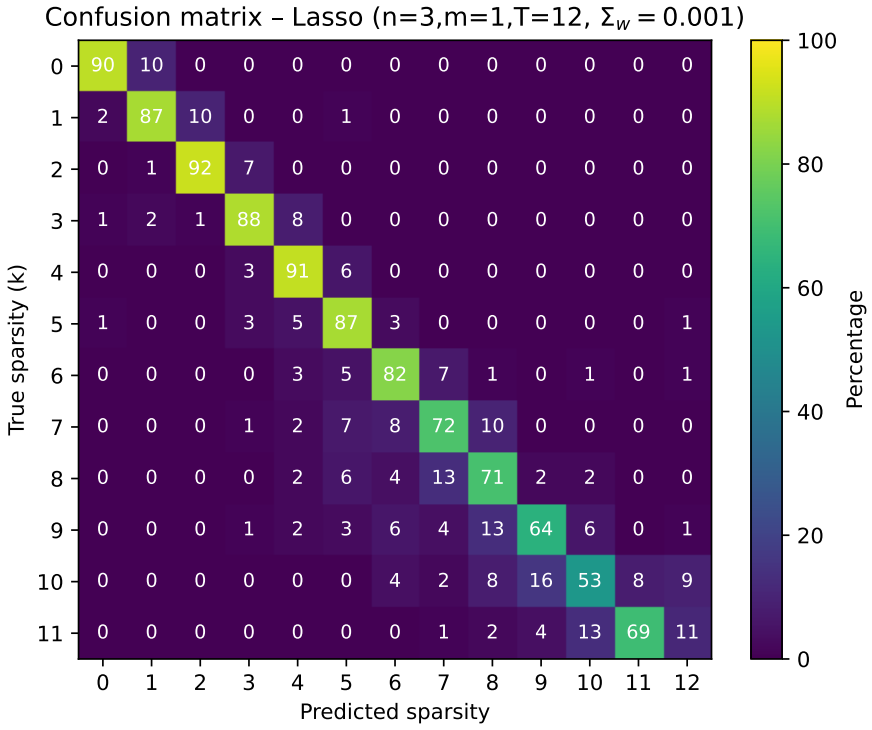

Figure 4.14: Confusion matrix for the LASSO method with $\Sigma_w = 0.001$ and $T = 12$.

The overall behavior is similar to that of the thresholding method. For dense systems, LASSO tends to predict a sparsity level close to the ground truth. In the moderately sparse regime, the estimates remain centered near the diagonal, albeit with slightly more dispersion.

The key distinction arises for ultra-sparse systems. When $k = 11$, LASSO retains 93 % of its predictions within the hyper-sparse band (10–12 zeros), whereas thresholding showed a broader spread in this regime. This confirms that LASSO handles extreme sparsity more consistently. In some cases, unlike hard-thresholding, the method may even yield completely null systems, i.e., with $k = 12$.

It is important to note that this analysis reflects only the number of zeros, not their exact positions. For LASSO, sparsity accuracy and exact structure accuracy can differ by up to 5 %. In other words, the algorithm may occasionally return a support with the correct number of non-zero entries, but located at incorrect indices. By contrast, with thresholding, the two notions typically coincided.

As before, results for higher-dimensional systems are available in Appendix B.

4.4.4 Comparison with the Hard Thresholding Method

To clarify the above observations, Figure 4.15 compares the performance of hard thresholding (for $\tau = 0.998$ and $\tau = 0.9995$) and the LASSO method under two noise levels: high noise ($\Sigma_w = 0.001$, shown in blue) and low noise ($\Sigma_w = 0.00001$, shown in orange). Accuracy is computed as in the previous graphs of this section.

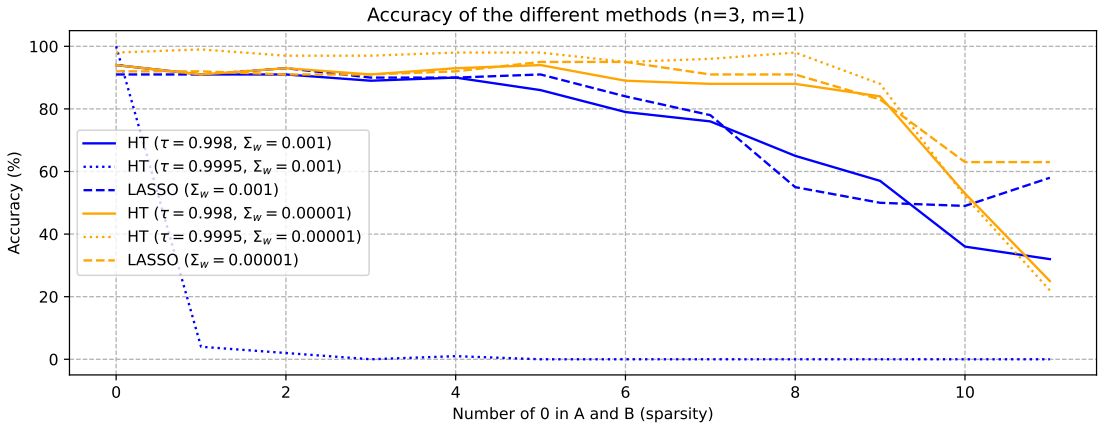

Figure 4.15: Accuracy of hard thresholding (HT) with two thresholds ($\tau = 0.998$ and $\tau = 0.9995$) and LASSO under two noise levels ($\Sigma_w = 0.001$ and $\Sigma_w = 0.00001$).

The figure confirms the trends discussed earlier. The hard thresholding method with $\tau = 0.998$ achieves comparable performance to the LASSO method across most sparsity levels, except in the hyper-sparse regime, where LASSO shows superior robustness.

Using a higher threshold ($\tau = 0.9995$) improves performance under low-noise conditions and even outperforms LASSO in most cases; again with the exception of hyper-sparse systems. However, this stricter thresholding becomes unreliable in moderate-to-high noise regimes, where it leads to significant accuracy loss.

Chapter 5

Numerical Experiments

This chapter evaluates the full pipeline: structure identification followed by structural controllability testing. While the previous chapter focused on the sparsity pattern recovery, the present analysis investigates whether the inferred structure preserves the controllability type of the original system. Even when recovery is imperfect, the predicted sparsity level typically remains close to the true value, except in hyper-sparse cases. As shown in Chapter 3, controllability is generally robust to small structural deviations, suggesting that minor identification errors may not alter the controllability outcome. This hypothesis is now tested empirically by assessing whether models obtained through hard thresholding or LASSO produce correct controllability classifications. Limitations of the approach, including sensitivity to parameter choices and intrinsic failure modes, are also discussed.

5.1 Simulations on Random Graphs

The complete identification pipeline (structure recovery followed by controllability testing) is evaluated on randomly generated systems. This full-scale assessment aims to measure the reliability of the approach when no prior structural assumptions are made and randomness affects all components of the problem.

For each sparsity level k, 100 independent systems are generated to ensure a broad exploration of the space of possible structures. Each system is uniquely defined by randomly assigning the positions of the k zero entries across the combined matrices (A_s, B_s), while the non-zero coefficients are sampled independently from a uniform distribution over $[-1, 1]$. The initial condition x_0 and the input sequence U_- are also drawn randomly from the same distribution. This ensures that every trial features distinct dynamics and excitation patterns, providing a representative sample of the system space at a given sparsity level.

This Monte Carlo strategy reduces the influence of stochastic noise and ensures that reported results reflect typical behavior rather than isolated cases. Averaging over many randomized trials mitigates sampling variability and provides a more accurate picture of how identification methods perform across diverse configurations. This setup allows for an evaluation of robustness without relying on idealized or specially constructed systems.

For each sampled system (A_s, B_s), the true structural controllability type (both weak and strong) is determined based on the criteria introduced in Chapter 3. Structure identification is then performed using the parameters specified in the previous chapter (such as the observation horizon T, the threshold τ, or the cross-validation scheme for LASSO) and examine whether the recovered structure yields the same controllability verdict. This end-to-end evaluation provides insight into the practical consequences of imperfect recovery and helps determine whether approximate models remain valid from a control-theoretic perspective.

5.1.1 Performance Under High Noise

The complete identification and controllability pipeline is first evaluated on small-scale systems under high process noise. The analysis considers systems with $n = 3$ states, $m = 1$ input, and a noise standard deviation of $\Sigma_w = 0.001$.

Both identification methods are applied using the configurations specified previously. For gradient descent followed by hard thresholding, the observation horizon is set to $T = 8$ with a threshold parameter $\tau = 0.998$. For LASSO regression, $T = 12$ is used, and the regularization parameter λ is selected via 3-fold cross-validation.

Figure 5.1 reports the accuracy of both methods across a range of sparsity levels. Each group of bars corresponds to a specific number of zero entries across A_s and B_s, from 0 (fully dense) to 11 (highly sparse).

The figure compares gradient descent with hard thresholding (labeled HT) and LASSO regression. Three accuracy metrics are shown for each method. The blue bars indicate the structure identification accuracy, defined as the percentage of cases where the recovered sparsity pattern (i.e., the zero positions) matches exactly the ground truth. The orange bars represent the accuracy in strong structural controllability: they reflect the proportion of systems for which the strong controllability status of the recovered model agrees with that of the true system. Finally, the green bars display the same metric for weak structural controllability.

This visual breakdown facilitates the evaluation of the performance of each method, not only in terms of raw identification accuracy, but also in preserving key system-theoretic properties under noisy conditions.

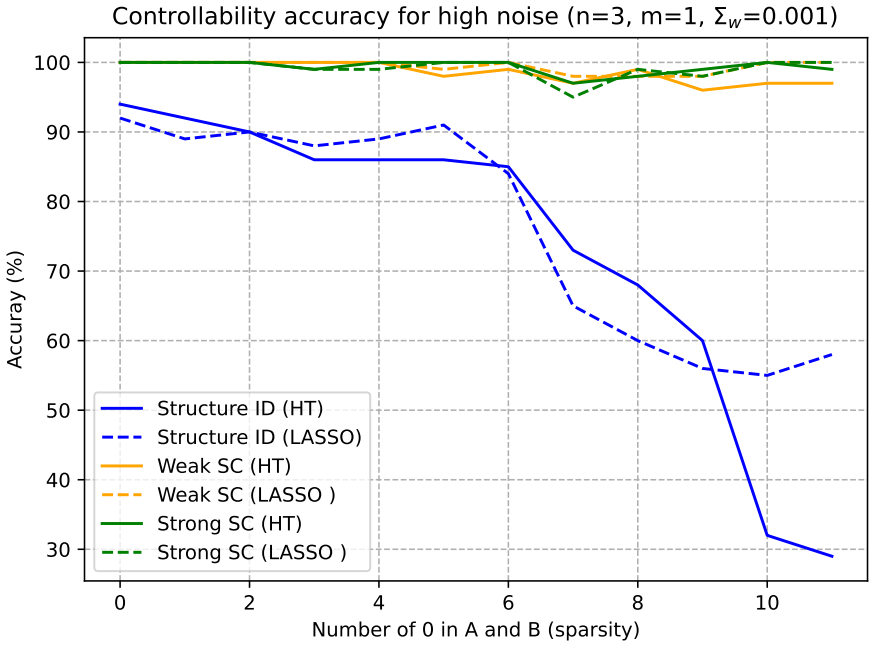

Figure 5.1: Controllability preservation rate for $n = 3$, $m = 1$, and $\Sigma_w = 0.001$.

As observed in previous chapters, both identification methods exhibit better performance on dense systems, that is, when the number of structural zeros k is small. In these regimes, the structure is easier to recover and both gradient descent with hard thresholding (HT) and LASSO regression return accurate sparsity patterns.

However, as the systems become increasingly sparse ($k \geq 10$), the structure identification accuracy drops significantly. In these hyper-sparse cases, LASSO consistently outperforms HT, achieving better recovery of the underlying sparsity even when noise is high. Despite this, both methods fall below 60 % structure accuracy in the most extreme sparsity settings.

The most striking observation comes from the orange and green curves, which represent the ability of the recovered models to preserve the original system's strong and weak structural controllability. Across the entire sparsity spectrum, both HT and LASSO maintain a prediction accuracy above 95 % for both controllability types.

This shows that, even when the structure is not recovered exactly, the essential control properties of the system are retained. This behavior can be traced back to the error analyses conducted earlier: when the methods fail, they tend to produce models with sparsity levels close to the true one (either slightly higher or slightly lower). As shown in Chapter 3, structural controllability typically evolves smoothly with sparsity. Therefore, small errors in the pattern often do not affect the controllability status, which explains the high level of agreement between ground truth and predictions.

In summary, even under high noise, the pipeline reliably captures the controllability character of the system, validating its use for practical purposes even when exact structure recovery is challenging.

To gain further insight into the reliability of the pipeline, a detailed error analysis is performed. Specifically, the predictions made by both methods are examined across all configurations; covering every sparsity level k and all 100 random trials per level. For each method, the frequency with which the recovered model correctly matches the controllability status (weak or strong) of the true system is evaluated.

The results are summarized in Table 5.1, which reports the classification performance of both identification methods with respect to weak and strong structural controllability. The four columns show the proportion of true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN). These quantities measure how often the predicted controllability status of the identified structure matches that of the ground-truth system.

Controllability type	TP	TN	FP	FN
Weak SC (HT)	56.75 %	41.83 %	1.33 %	0.08 %
Weak SC (LASSO)	56.67 %	42.75 %	0.42 %	0.17 %
Strong SC (HT)	6.25 %	93.08 %	0.50 %	0.17 %
Strong SC (LASSO)	6.33 %	92.83 %	0.75 %	0.08 %

Table 5.1: Classification results over all trials ($n = 3$, $m = 1$, $\Sigma_w = 0.001$).

From Table 5.1, it is observed that misclassifications are extremely rare across all categories, confirming the robustness of the method. In most cases, both HT and LASSO correctly identify the system's controllability status. When errors do occur, they tend to be false positives, that is, the recovered model is deemed structurally controllable (weakly or strongly) when the true system is not.

This asymmetry can be attributed to the nature of structure identification errors. As shown previously, incorrect recoveries usually involve small deviations in the sparsity pattern, and often result in models that are slightly denser than the original. Since controllability is typically more likely in denser systems (as they offer more interconnections), these structural overestimations bias the result toward declaring controllability, even when the true system is not.

5.1.2 Performance Under Low Noise

The analysis is now extended to the low-noise setting, while keeping the same system dimensions ($n = 3$, $m = 1$) and reducing the noise level to $\Sigma_w = 0.00001$. The identification procedures are applied using the previously established configurations: gradient descent with hard thresholding is run with $T = 8$ and $\tau = 0.9995$, while LASSO regression uses $T = 12$ and 3-fold cross-validation for selecting the regularization parameter λ.

Figure 5.2 presents the results, following the same structure as in the high-noise case. The plot shows, for each sparsity level, the structure identification accuracy (blue), as well as the preservation rate for strong (orange) and weak (green) structural controllability. This allows for direct comparison with Figure 5.1 and highlights how lower noise improves or shifts the performance profile of the two methods.

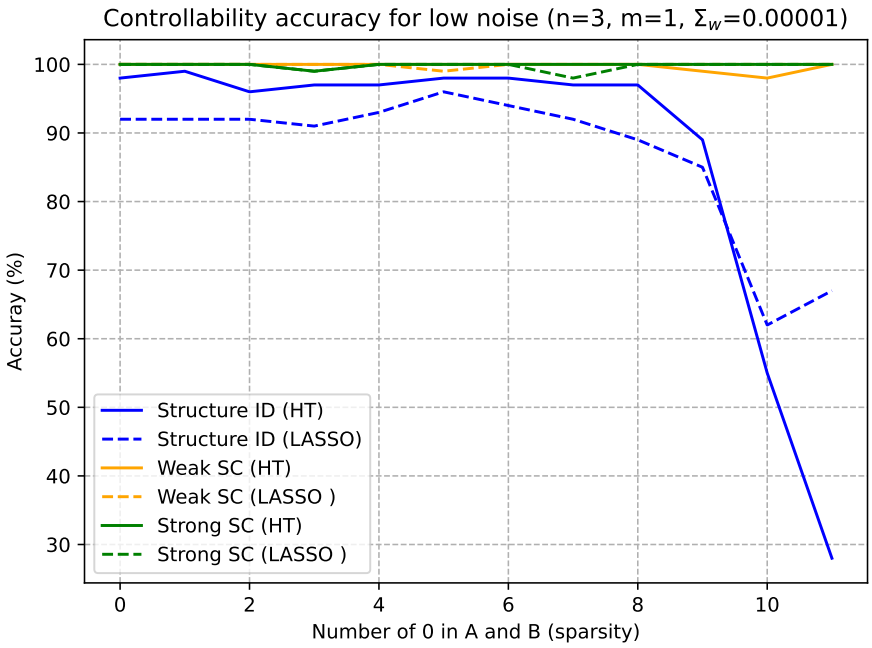

Figure 5.2: Controllability preservation rate for $n = 3$, $m = 1$, and $\Sigma_w = 0.00001$.

In this low-noise setting, both methods achieve high structure identification accuracy for dense systems, with exact recovery exceeding 90 % when $k \leq 7$. As sparsity increases, the identification becomes more difficult, and accuracy decreases, especially for HT in the hypersparse regime ($k \geq 10$). Nonetheless, LASSO continues to outperform HT in these cases, confirming its greater robustness to extreme sparsity.

Importantly, the structural controllability metrics (orange and green bars) remain consistently high across the entire range. Both HT and LASSO maintain weak and strong controllability agreement with the ground truth in more than 95 % of cases, even when the structure is not perfectly recovered. This suggests that the critical system-level properties are preserved despite minor structural inaccuracies.

To better understand the reliability of these findings, the full set of predictions is compiled into confusion matrices, which summarize the frequencies of true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN) across all trials. The results are presented in Table 5.2.

Controllability type	TP	TN	FP	FN
Weak SC (HT)	56.83 %	42.92 %	0.25 %	0.00 %
Weak SC (LASSO)	56.75 %	43.17 %	0.00 %	0.08 %
Strong SC (HT)	6.42 %	93.50 %	0.08 %	0.00 %
Strong SC (LASSO)	6.42 %	93.33 %	0.25 %	0.00 %

Table 5.2: Classification results over all trials ($n = 3$, $m = 1$, $\Sigma_w = 0.00001$).

These numbers reinforce the earlier conclusions. The rate of misclassification remains extremely low, particularly for false negatives. Most errors are again false positives, which occur when the inferred structure slightly overestimates the system's connectivity, causing it to be (incorrectly) labeled as controllable. This is consistent with the tendency of identification methods to slightly underestimate sparsity when noise is low.

In summary, the low-noise regime showcases the strengths of both identification techniques, with excellent structural and controllability recovery, particularly for LASSO in highly sparse systems. These results validate the use of data-driven structure inference in low-noise environments and highlight the reliability of the proposed pipeline in practice.

5.1.3 Comparison with Data-Driven Controllability Testing

In this experiment, the identification-based pipeline is compared to the data-informativity framework introduced by van Waarde *and al.* [34] and discussed in Chapter 2. Specifically, the data-driven Hautus test is applied directly to the measured input–state trajectories, without performing any system identification.

The experimental setup remains identical to the previous section: systems with $n = 3$ states, $m = 1$ input, and low process noise ($\Sigma_w = 0.00001$) are considered. Both identification strategies, gradient descent with hard thresholding (HT) and LASSO regression, are applied using the same parameters as previously defined.

Figure 5.3 reports the classification accuracy for weak structural controllability as a function of sparsity. Each curve indicates the proportion of systems for which the predicted controllability status (weakly controllable or not) matches the ground truth. The performance of the HT and LASSO identification pipelines is compared with the data-driven test based on the informativity framework. Additionally, the baseline proportion of truly weakly controllable systems at each sparsity level is included. This reference curve corresponds to a naive classifier that always assigns the weakly controllable label, regardless of the actual system properties.

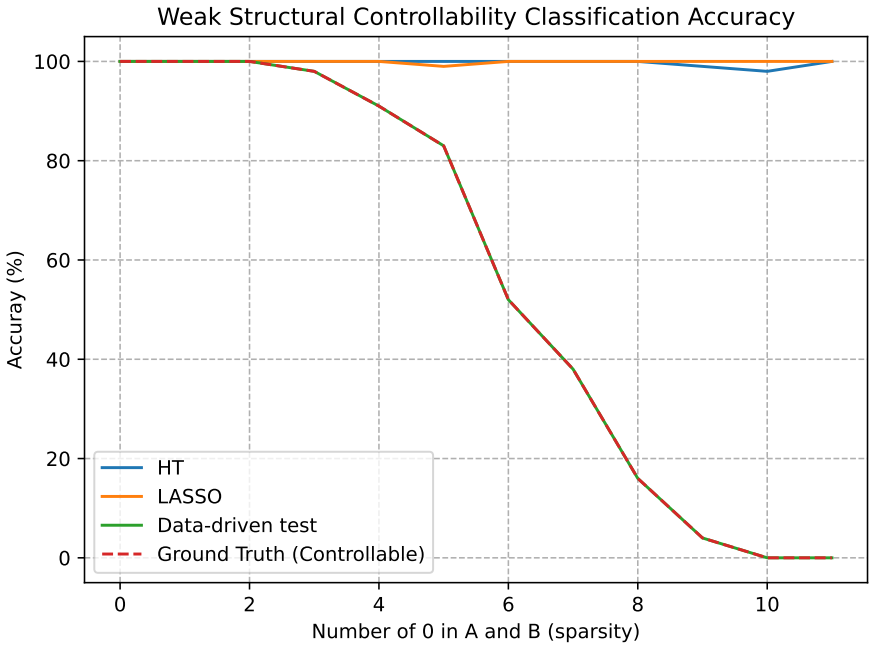

Figure 5.3: Accuracy of weak structural controllability classification for systems with $n = 3$, $m = 1$, and $\Sigma_w = 0.00001$.

Consistent with the findings of the previous section, both HT and LASSO achieve near-perfect accuracy across all sparsity levels, indicating that identification errors do not significantly affect the outcome of the controllability test. In contrast, the data-driven test aligns with the baseline curve, revealing that under noisy conditions, even minimal, the data-driven Hautus test systematically reports controllability. Introducing a threshold on the smallest singular value could potentially improve robustness, but this would require prior knowledge of the noise level, as the singular value is observed to increase approximately linearly with the noise magnitude (see Figure 2.4).

This phenomenon, analyzed in Chapter 2, shows a critical limitation: the data-informativity framework can yield overly optimistic conclusions when applied to noisy trajectories, effectively behaving as a constant classifier for weak controllability.

5.2 Limits of Structural Controllability

A central limitation of the random graph simulations presented so far lies in the nature of weak structural controllability itself. When system matrices are generated with random nonzero entries, any structure that satisfies the conditions for weak structural controllability will, with probability one, yield a numerically controllable system. This follows from the fact that uncontrollable parameter choices form a set of measure zero within the admissible space of free parameters.

As a result, the only uncontrollable systems observed in the random simulations correspond to structures that are not weakly structurally controllable from the outset. This indicates that the experiments do not naturally cover systems that are structurally controllable in the weak sense but become uncontrollable due to particular numerical configurations. Nevertheless, such pathological cases remain theoretically relevant and merit further investigation.

To investigate this specific class of systems (those that are weakly structurally controllable but numerically uncontrollable) system matrices A and B are constructed to embed a deliberately uncontrollable mode. A fixed direction $v = [1, -1, 1]^\top$ and scalar λ_v are chosen such that the conditions $v^\top A = \lambda_v v^\top$ and $v^\top B = 0$ hold. Under this construction, the scalar projection $z(t) = v^\top x(t)$ evolves independently of the input, rendering it unresponsive to the control signal $u(t)$. Consequently, the dynamics of $z(t)$ are autonomous, following $z(t+1) = \lambda_v z(t)$, and the system becomes uncontrollable along the associated mode.

What makes this construction subtle is that the zero-nonzero pattern of A and B still satisfies the graph-based conditions for weak structural controllability. The system's uncontrollability arises solely from the exact alignment of numerical values, not from its structural support. In algebraic terms, the controllability matrix $\mathcal{C} = [B\ AB\ A^2B\ \dots A^{n-1}B]$ has full structural rank but lacks full numerical rank.

This construction is used to simulate two categories of uncontrollable systems: those that are both numerically and structurally uncontrollable, and those that are structurally controllable in the weak sense but rendered uncontrollable through deliberate numerical tuning. In both scenarios, the identification pipeline, whether based on hard thresholding or LASSO, systematically classifies the systems as weakly structurally controllable. This outcome confirms that the approach inherits the same limitations as the underlying graph-theoretic framework: it does not account for numerical degeneracies that compromise controllability while remaining consistent with structural criteria.

Chapter 6

Conclusion

Grounded in the framework of data-driven control, and more specifically the concept of data informativity introduced by van Waarde *and al.* [34], this Master's thesis has addressed the challenge of assessing the controllability of linear time-invariant systems using only noisy observations of their input–state trajectories. To overcome the limitations introduced by noise, we proposed an intermediate two-step strategy: first identifying the system's structural sparsity pattern, and then testing controllability through structural condition.

Chapter 1 provided the context and motivation for this thesis, emphasizing the critical role of controllability in practical applications. It also presented a comprehensive literature review across the three core mathematical domains central to this work: data-driven control, structural controllability, and structure identification. By situating the thesis within the broader history of control theory and discussing the main contributions and limitations of existing approaches, the chapter laid the conceptual groundwork for the developments that follow.

In Chapter 2, we presented the data informativity framework in detail. We supported the theoretical exposition with examples and numerical investigations, particularly focusing on the robustness of the data-driven Hautus test. These experiments highlighted a major limitation: the framework lacks resilience to noise, which may result in false positive conclusions regarding controllability.

To address this limitation, we turned to structural methods. Chapter 3 reviewed key notions from the literature, such as weak and strong structural controllability, and introduced new numerical results showing how these properties evolve with edge sparsity in random LTI systems. These insights supported the idea that structural features can buffer the effects of noise in a meaningful way.

Chapter 4 developed two techniques for recovering the underlying sparsity pattern of a system from data: a hard-thresholding method based on least-squares estimation, and a regularized regression approach using LASSO. We analyzed the performance of both methods empirically, showing that even when exact recovery fails, the predicted structures often retain key system properties. We also analyzed the nature of the errors to better understand the limitations of these techniques.

In Chapter 5, we evaluated the full identification pipeline (structure recovery followed by controllability testing) on both random and carefully designed pathological systems. Our experiments showed that both hard-thresholding and LASSO can reliably identify the correct controllability status of the original system, even when the structural recovery is imperfect. On small systems, both methods achieved over 97% accuracy in controllability classification. Furthermore, comparisons with the data-informativity framework confirmed its theoretical limitations under noisy data, highlighting the advantage of the proposed hybrid method.

Overall, this thesis achieves its initial objectives: it proposes a robust and effective methodology for assessing system controllability using only trajectory data, without requiring full model identification. By bridging structural and data-driven paradigms, it provides a new approach to combining theoretical insights with the challenges of real-world noise.

This work also opens several avenues for future research. Since the method combines ideas from multiple mathematical domains, further refinement of each component could enhance performance or generality. For example, integrating data-driven filtering techniques to suppress noise, extending the method to higher-dimensional systems, or addressing partial observability through techniques such as compressive sensing would make the approach more broadly applicable. Additionally, a deeper analysis of pathological cases, where systems appear controllable under noise but are not, could improve robustness guarantees and help delineate the limits of inference from data alone.

Declaration of Generative AI and AI-Assisted Technologies in the Writing Process

During the preparation of this work, the authors used *ChatGPT 4.0* from *OpenAI* to verify grammar and spelling, enhance sentence structure, improve clarity, fluency and ensure coherence and readability throughout the manuscript. It is important to note that *ChatGPT 4.0* did not contribute any information, data, or original content to this work: its role was solely to assist in refining the existing text. After each use of this tool, the authors thoroughly reviewed and edited the content as needed and take full responsibility for the final content of the publication.

Bibliography

[1] C. W. Reynolds, "Flocks, Herds and Schools: A Distributed Behavioral Model," *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*, pp. 25–34, 1987.

[2] P. Bauer, A. Thorpe, and G. Brunet, "The quiet revolution of numerical weather prediction," *Nature*, vol. 525, no. 7567, pp. 47–55, 2015.

[3] M. V. Cook, *Flight Dynamics Principles: A Linear Systems Approach to Aircraft Stability and Control*, Butterworth-Heinemann, 2012.

[4] F. Ricci, L. Rokach, and B. Shapira, "Recommender systems: Techniques, applications, and challenges," *Recommender Systems Handbook*, pp. 1–35, 2021.

[5] K. J. Åström and R. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*, Princeton University Press, 2021.

[6] D. Boley, "Computing the Kalman decomposition: An optimal method," *IEEE Transactions on Automatic Control*, vol. 29, no. 1, pp. 51–53, 2003.

[7] J. Doyle, "Robust and Optimal Control," *Proceedings of the 35th IEEE Conference on Decision and Control*, pp. 1595–1598, 1996.

[8] R. E. Kalman, "Mathematical Description of Linear Dynamical Systems," *Journal of the Society for Industrial and Applied Mathematics, Series A: Control*, vol. 1, no. 2, pp. 152–192, 1963.

[9] M. N. I. Sarkar, L. G. Meegahapola, and M. Datta, "Reactive power management in renewable rich power grids: A review of grid-codes, renewable generators, support devices, control strategies and optimization algorithms," *IEEE Access*, vol. 6, pp. 41458–41489, 2018.

[10] S. H. Low and D. E. Lapsley, "Optimization flow control. I. Basic algorithm and convergence," *IEEE/ACM Transactions on Networking*, vol. 7, no. 6, pp. 861–874, 1999.

[11] C. Antoniou, H. N. Koutsopoulos, and G. Yannis, "Dynamic data-driven local traffic state estimation and prediction," *Transportation Research Part C: Emerging Technologies*, vol. 34, pp. 89–107, 2013.

[12] U. Alon, *An Introduction to Systems Biology: Design Principles of Biological Circuits*, Chapman and Hall/CRC, 2019.

[13] L. Ljung, "System Identification," *Signal Analysis and Prediction*, Boston, MA: Birkhäuser Boston, pp. 163–173, 1998.

[14] J. G. Ziegler and N. B. Nichols, "Optimum settings for automatic controllers," *Transactions of the American Society of Mechanical Engineers*, vol. 64, no. 8, pp. 759–765, 1942.

[15] K. J. Åström, "Adaptive control," *Mathematical System Theory: The Influence of R. E. Kalman*, Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 437–450, 1995.

[16] H. Hjalmarsson, M. Gevers, S. Gunnarsson, and O. Lequin, "Iterative feedback tuning: theory and applications," *IEEE Control Systems Magazine*, vol. 18, no. 4, pp. 26–41, 1998.

[17] M. G. Safonov and T. Tsao, "The unfalsified control concept and learning," *IEEE Transactions on Automatic Control*, vol. 42, no. 6, pp. 843–847, 1997.

[18] J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer, "Data-driven model predictive control with stability and robustness guarantees," *IEEE Transactions on Automatic Control*, vol. 66, no. 4, pp. 1702–1717, 2020.

[19] J. Coulson, J. Lygeros, and F. Dörfler, "Data-enabled predictive control: In the shallows of the DeePC," *2019 18th European Control Conference (ECC)*, pp. 307–312, 2019.

[20] L. Campestrini, D. Eckhard, A. S. Bazanella, and M. Gevers, "Data-driven model reference control design by prediction error identification," *Journal of the Franklin Institute*, vol. 354, no. 6, pp. 2628–2647, 2017.

[21] S. Bradtke, "Reinforcement learning applied to linear quadratic regulation," *Advances in Neural Information Processing Systems*, vol. 5, 1992.

[22] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu, "On the sample complexity of the linear quadratic regulator," *Foundations of Computational Mathematics*, vol. 20, no. 4, pp. 633–679, 2020.

[23] M. Fliess and C. Join, "Model-free control," *International Journal of Control*, vol. 86, no. 12, pp. 2228–2252, 2013.

[24] Z.-S. Hou and Z. Wang, "From model-based control to data-driven control: Survey, classification and perspective," *Information Sciences*, vol. 235, pp. 3–35, 2013.

[25] J. C. Willems, P. Rapisarda, I. Markovsky, and B. L. M. De Moor, "A note on persistency of excitation," *Systems & Control Letters*, vol. 54, no. 4, pp. 325–329, 2005.

[26] H. J. van Waarde, C. De Persis, M. K. Camlibel, and P. Tesi, "Willems' fundamental lemma for state-space systems and its extension to multiple datasets," *IEEE Control Systems Letters*, vol. 4, no. 3, pp. 602–607, 2020.

[27] I. Markovsky and P. Rapisarda, "Data-driven simulation and control," *International Journal of Control*, vol. 81, no. 12, pp. 1946–1959, 2008.

[28] J. Berberich, A. Koch, C. W. Scherer, and F. Allgöwer, "Robust data-driven state-feedback design," *2020 American Control Conference (ACC)*, pp. 1532–1538, 2020.

[29] H. J. van Waarde, M. K. Camlibel, J. Eising, and H. L. Trentelman, "Quadratic matrix inequalities with applications to data-based control," *SIAM Journal on Control and Optimization*, vol. 61, no. 4, pp. 2251–2281, 2023.

[30] A. Koch, J. Berberich, and F. Allgöwer, "Verifying dissipativity properties from noise-corrupted input-state data," *2020 59th IEEE Conference on Decision and Control (CDC)*, pp. 616–621, 2020.

[31] V. K. Mishra, I. Markovsky, and B. D. O. Anderson, "Data-driven tests for controllability," *IEEE Control Systems Letters*, vol. 5, no. 2, pp. 517–522, 2021.

[32] Z. Wang and D. Liu, "Data-based controllability and observability analysis of linear discrete-time systems," *IEEE Transactions on Neural Networks*, vol. 22, no. 12, pp. 2388–2392, 2011.

[33] D. Liu, P. Yan, and Q. Wei, "Data-based analysis of discrete-time linear systems in noisy environment: Controllability and observability," *Information Sciences*, vol. 288, pp. 314–329, 2014.

[34] H. J. van Waarde, J. Eising, H. L. Trentelman, and M. K. Camlibel, "Data informativity: A new perspective on data-driven analysis and control," *IEEE Transactions on Automatic Control*, vol. 65, no. 11, pp. 4753–4768, 2020.

[35] H. L. Trentelman, H. J. van Waarde, and M. K. Camlibel, "An informativity approach to the data-driven algebraic regulator problem," *IEEE Transactions on Automatic Control*, vol. 67, no. 11, pp. 6227–6233, 2022.

[36] J. Gonçalves and S. Warnick, "Necessary and sufficient conditions for dynamical structure reconstruction of LTI networks," *IEEE Transactions on Automatic Control*, vol. 53, no. 7, pp. 1670–1674, 2008.

[37] H. J. van Waarde, P. Tesi, and M. K. Camlibel, "Necessary and sufficient topological conditions for identifiability of dynamical networks," *IEEE Transactions on Automatic Control*, vol. 65, no. 11, pp. 4525–4537, 2019.

[38] H. J. van Waarde, P. Tesi, and M. K. Camlibel, "Topology identification of heterogeneous networks of linear systems," *2019 IEEE 58th Conference on Decision and Control (CDC)*, pp. 5513–5518, 2019.

[39] M. Nabi-Abdolyousefi and M. Mesbahi, "Network identification via node knockout," *IEEE Transactions on Automatic Control*, vol. 57, no. 12, pp. 3214–3219, 2012.

[40] M. Suzuki, N. Takatsuki, J.-I. Imura, and K. Aihara, "Node knock-out based structure identification in networks of identical multi-dimensional subsystems," *2013 European Control Conference (ECC)*, pp. 2280–2285, 2013.

[41] H. J. van Waarde, P. Tesi, and M. K. Camlibel, "Topology reconstruction of dynamical networks via constrained Lyapunov equations," *IEEE Transactions on Automatic Control*, vol. 64, no. 10, pp. 4300–4306, 2019.

[42] D. Materassi and M. V. Salapaka, "On the problem of reconstructing an unknown topology via locality properties of the Wiener filter," *IEEE Transactions on Automatic Control*, vol. 57, no. 7, pp. 1765–1777, 2012.

[43] S. Shahrampour and V. M. Preciado, "Topology identification of directed dynamical networks via power spectral analysis," *IEEE Transactions on Automatic Control*, vol. 60, no. 8, pp. 2260–2265, 2014.

[44] Y. Yuan, G. B. Stan, S. Warnick, and J. Gonçalves, "Robust dynamical network structure reconstruction," *Automatica*, vol. 47, no. 6, pp. 1230–1235, 2011.

[45] S. Fattahi and S. Sojoudi, "Data-driven sparse system identification," *2018 56th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 462–469, 2018.

[46] H.-S. Lee, J.-W. Lee, W.-J. Song, and S. E. Kim, "Adaptive algorithm for sparse system identification based on hard-thresholding techniques," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 67, no. 12, pp. 3597–3601, 2020.

[47] C.-T. Lin, "Structural controllability," *IEEE Transactions on Automatic Control*, vol. 19, no. 3, pp. 201–208, 1974.

[48] J.-M. Dion, C. Commault, and J. van der Woude, "Generic properties and control of linear structured systems: A survey," *Automatica*, vol. 39, no. 7, pp. 1125–1144, 2003.

[49] Y.-Y. Liu, J.-J. Slotine, and A.-L. Barabási, "Controllability of complex networks," *Nature*, vol. 473, no. 7346, pp. 167–173, 2011.

[50] A. Chapman and M. Mesbahi, "On strong structural controllability of networked systems: A constrained matching approach," *2013 American Control Conference*, pp. 6126–6131, 2013.

[51] M. Trefois and J.-C. Delvenne, "Zero forcing number, constrained matchings and strong structural controllability," *Linear Algebra and its Applications*, vol. 484, pp. 199–218, 2015.

[52] N. Monshizadeh, S. Zhang, and M. K. Camlibel, "Zero forcing sets and controllability of dynamical systems defined on graphs," *IEEE Transactions on Automatic Control*, vol. 59, no. 9, pp. 2562–2567, 2014.

[53] J. Jia, H. J. van Waarde, and H. L. Trentelman, "A unifying framework for strong structural controllability," *IEEE Transactions on Automatic Control*, vol. 66, no. 1, pp. 391–398, 2020.

[54] H. L. Trentelman, A. A. Stoorvogel, and M. Hautus, "Control theory for linear systems," *Applied Mechanics Reviews*, vol. 55, no. 5, pp. B87–B87, 2002.

[55] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 2nd ed., Cambridge University Press, 2012.

[56] R. Eising, "Between controllable and uncontrollable," *Systems and Control Letters*, vol. 4, no. 5, pp. 263–264, 1984.

[57] R. Shields and J. Pearson, "Structural controllability of multiinput linear systems," *IEEE Transactions on Automatic Control*, vol. 21, no. 2, pp. 203–212, 1976.

[58] G. Ramos, A. P. Aguiar, and S. Pequito, "An overview of structural systems theory," *Automatica*, vol. 140, p. 110229, 2022.

[59] X. Liu, H. Lin, and B. M. Chen, "Structural controllability of switched linear systems," *Automatica*, vol. 49, no. 12, pp. 3531–3537, 2013.

[60] J. A. Bondy and U. S. R. Murty, *Graph Theory with Applications*. London: Macmillan, 1976.

[61] J. E. Hopcroft and R. M. Karp, "An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs," *SIAM Journal on Computing*, vol. 2, no. 4, pp. 225–231, 1973.

[62] A. Olshevsky, "Minimal controllability problems," *IEEE Transactions on Control of Network Systems*, vol. 1, no. 3, pp. 249–258, 2014.

[63] T. Menara, G. Bianchin, M. Innocenti, and F. Pasqualetti, "On the number of strongly structurally controllable networks," *Proceedings of the 2017 American Control Conference (ACC)*, pp. 340–345, 2017.

[64] Y. Nesterov, *Lectures on Convex Optimization*, vol. 137. Berlin: Springer, 2018.

[65] G. Strang, *Introduction to Linear Algebra*. Wellesley-Cambridge Press, 2022.

[66] T. Blumensath and M. E. Davies, "Iterative hard thresholding for compressed sensing," *Applied and Computational Harmonic Analysis*, vol. 27, no. 3, pp. 265–274, 2009.

[67] Y. Ida, S. Kanai, A. Kumagai, T. Iwata, and Y. Fujiwara, "Fast Iterative Hard Thresholding Methods with Pruning Gradient Computations," *Proceedings of the Thirty-Eighth Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2024.

[68] J. Ranstam and J. A. Cook, "LASSO regression," *Journal of British Surgery*, vol. 105, no. 10, pp. 1348–1348, 2018.

[69] D. L. Donoho, "For most large underdetermined systems of linear equations the minimal ℓ_1-norm solution is also the sparsest solution," *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, vol. 59, no. 6, pp. 797–829, 2006.

[70] A. Bagirov, N. Karmitsa, and M. M. Mäkelä, *Introduction to Nonsmooth Optimization: Theory, Practice and Software*. Cham, Heidelberg: Springer International Publishing, 2014.

[71] A. I. Chen and A. Ozdaglar, “A fast distributed proximal-gradient method,” *Proceedings of the 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 601–608, 2012.

[72] N. Parikh, and S. Boyd, “Proximal algorithms,” *Foundations and Trends in Optimization*, vol. 1, no. 3, pp. 127–239, 2014.

[73] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM Journal on Imaging Sciences*, vol. 2, no. 1, pp. 183–202, 2009.

[74] R. Tibshirani, “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 58, no. 1, pp. 267–288, 1996.

[75] T.-T. Wong and P.-Y. Yeh, “Reliable accuracy estimates from k-fold cross validation,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 32, no. 8, pp. 1586–1594, 2019.

Appendix A

Hard Thresholding Algorithm

Concretely, the vector $y \in \mathbb{R}^{n(n+m)}$ subjected to thresholding is obtained by solving the least-squares problem (4.4) using gradient descent. Once the iterations converge to a solution satisfying the prescribed stopping criterion, this estimated vector y is post-processed to recover a sparse structure.

To that end, a hard thresholding procedure is employed. This method eliminates the $\mathcal{K}$ smallest entries in magnitude from the vector y, where $\mathcal{K}$ is automatically selected to balance reconstruction accuracy and sparsity. The selection criterion is based on the RMSE of the reconstructed outputs, as described earlier.

Algorithm 1 Hard Thresholding with Automatic Selection

Set $\mathcal{K}^* \leftarrow 0$, $y_{\mathcal{K}^*} \leftarrow y$
Compute the absolute values of y: $|y|$
for $\mathcal{K}_i = 0$ to $n(n+m)$ **do**
    Identify indices $\mathcal{I}_{\mathcal{K}_i}$ of the $\mathcal{K}_i$ smallest entries in $|y|$
    Set $y_{\text{thr}} \leftarrow y$
    for each $i \in \mathcal{I}_{\mathcal{K}_i}$ **do**
        Set $y_{\text{thr}}[i] \leftarrow 0$
    end for
    Compute $\text{RMSE}_{\mathcal{K}_i} = \frac{1}{\sqrt{N}} \|z - \Phi y_{\text{thr}}\|_2$
    if $\exp(-\text{RMSE}_{\mathcal{K}_i}) \geq \tau$ **then**
        $\mathcal{K}^* \leftarrow \mathcal{K}_i$
        $y_{\mathcal{K}^*} \leftarrow y_{\text{thr}}$
    end if
end for
Reshape $y_{\mathcal{K}^*}$ into $A_{\mathcal{K}^*}, B_{\mathcal{K}^*}$
return $A_{\mathcal{K}^*}, B_{\mathcal{K}^*}$

Appendix B

Higher Dimension Analysis

The results presented below are based on Monte Carlo simulations conducted with 50 random samples per sparsity level k.

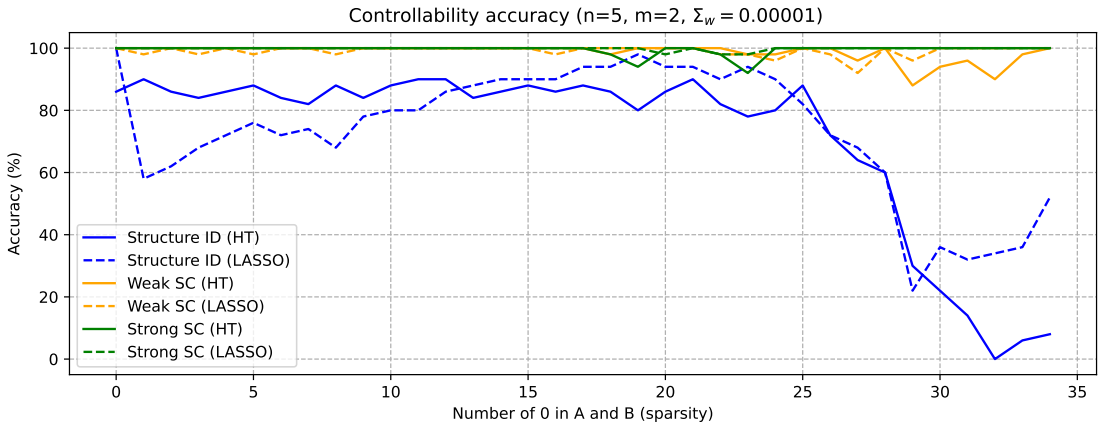

Figure B.1: Controllability preservation rate for $n = 5$, $m = 2$, and $\Sigma_w = 0.00001$.

The figure shows the performance of both identification methods on higher-dimensional systems with $n = 5$ states and $m = 2$ inputs. A slight divergence in behavior is observed: the hard-thresholding method achieves higher accuracy on dense systems, whereas LASSO performs better in sparse regimes, consistent with observations in smaller systems. With respect to controllability classification, both methods maintain high accuracy overall. Nevertheless, a slight drop is noted for hard thresholding in the very sparse setting.

A more detailed analysis of the structure identification performance follows.

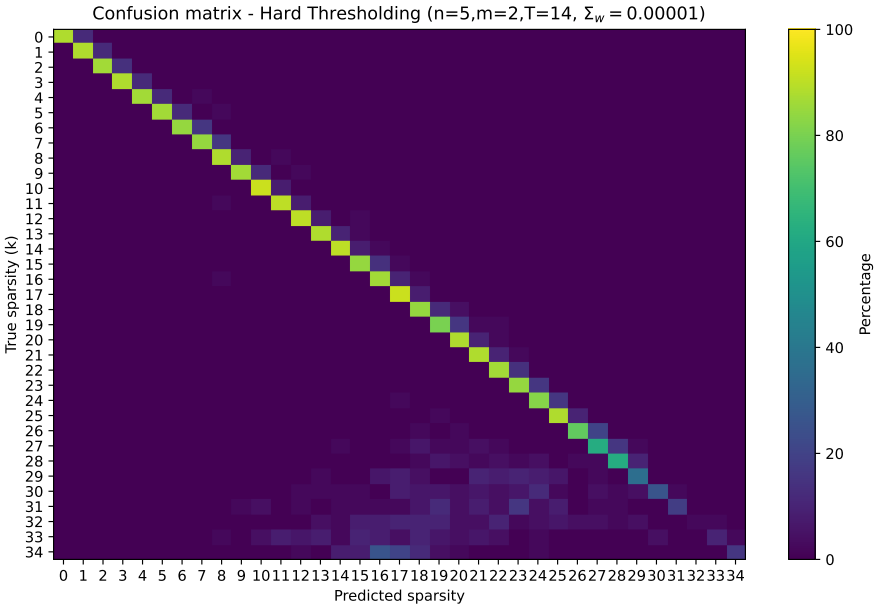

Figure B.2: Confusion matrix for the hard-thresholding method.

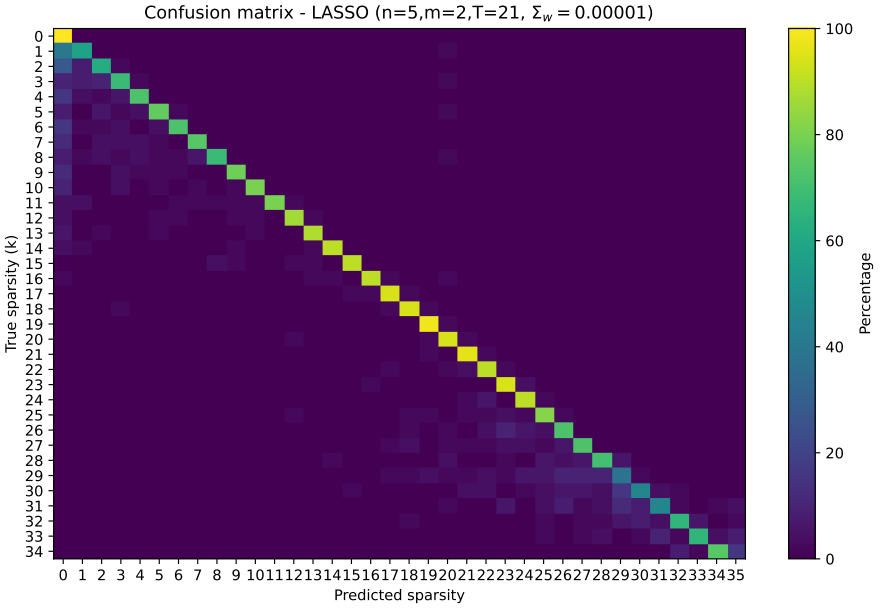

Figure B.3: Confusion matrix for the LASSO method.

The results confirm that for dense systems, the thresholding method achieves higher accuracy than LASSO regression, whereas the opposite holds in sparse regimes. In both cases, the predicted structures remain concentrated near the correct sparsity level k. However, LASSO exhibits a clear limitation in terms of computational efficiency: the cross-validation procedure becomes significantly slower in higher-dimensional settings. In contrast, the thresholding method retains low computational cost.

Overall, the simulations indicate that both approaches behave similarly to the small-scale case, although with a moderate decline in accuracy. Experiments conducted under increased noise levels display the same qualitative patterns, with a slightly more pronounced accuracy reduction.

UNIVERSITÉ CATHOLIQUE DE LOUVAIN
École polytechnique de Louvain
Rue Archimède, 1 bte L6.11.01, 1348 Louvain-la-Neuve, Belgique | **www.uclouvain.be/epl**