Pablo Sarró Sánchez

Detectability Transition in the Problem of Network Alignment

Degree Final Project

Supervised by Dr Marta Sales Pardo
and Dr Roger Guimerà Manrique

Bachelor's degree in Mathematical and Phyiscal Engineering

UNIVERSITAT ROVIRA i VIRGILI

Tarragona
2025

Abstract

The network alignment problem aims to find a correspondence between the nodes of two or more graphs, representing different observations of the same underlying system, that best preserves their structure and connectivity. This problem is central to various domains such as biology, social networks, and computer science, and becomes particularly challenging in noisy settings. In this work, we study the detectability of the ground truth alignment between two correlated random graphs as a function of noise. By adopting a Bayesian formulation and sampling alignments using Markov Chain Monte Carlo methods, specifically parallel tempering, we explore how the posterior distribution evolves and whether it concentrates near the true solution with respect to the noise level.

Our results reveal two key transitions: an algorithmic one, where recovery sharply degrades, and a detectability threshold beyond which recovery becomes impossible. Between these two thresholds, we identify a phase where the true alignment is still theoretically attainable but is algorithmically inaccessible. We interpret these findings through the perspective of statistical physics, analysing the energy landscape and the structural similarity of local minima with the optimal alignment. The results highlight a fundamental decoupling between low-energy configurations and structural proximity to the ground truth. This study contributes to the understanding of the fundamental limits of network alignment and the nature of its computational hardness.

Keywords: Network alignment, random graphs, detectability phase transition, Markov Chain Monte Carlo, parallel tempering, bayesian inference, statistical physics, energy landscape, hard phase.

Resumen

El problema del alineamiento de redes busca encontrar una correspondencia entre los nodos de dos o más grafos, que representan distintas observaciones de un mismo sistema subyacente, de manera que se preserve al máximo su estructura y conectividad. Este problema es central en diversos ámbitos como la biología, las redes sociales y la informática, y resulta especialmente desafiante en contextos con mucho ruido. En este trabajo, estudiamos la detectabilidad del alineamiento verdadero entre dos grafos aleatorios correlacionados en función del ruido. Adoptando una formulación bayesiana y muestreando alineamientos mediante métodos de Markov Chain Monte Carlo con parallel tempering, exploramos cómo evoluciona la distribución posterior y si esta se concentra cerca de la solución correcta en función del ruido.

Nuestros resultados revelan dos transiciones clave: una algorítmica, en la que la recuperación del alineamiento verdadero se pierde bruscamente; y un umbral de detectabilidad, a partir del cual la recuperación se vuelve imposible. Entre ambos umbrales, identificamos una fase en la que el alineamiento verdadero sigue siendo teóricamente alcanzable, pero es inaccesible algorítmicamente. Interpretamos estos hallazgos desde la perspectiva de la física estadística, analizando el conjunto de energías y la similitud estructural de los mínimos locales con el alineamiento óptimo. Los resultados muestran un desacoplamiento entre configuraciones de baja energía y su proximidad estructural al alineamiento real. Este estudio contribuye a comprender los límites fundamentales del alineamiento de redes y la naturaleza de su dificultad computacional.

Palabras clave: alineamiento de redes, grafos aleatorios, transición de fase de detectabilidad, Markov Chain Monte Carlo, parallel tempering, inferencia bayesiana, física estadística, conjunto de energías, fase difícil.

Resum

El problema de l'alineament de xarxes té com a objectiu trobar una correspondència entre els nodes de dos o més grafs, que representen diferents observacions d'un mateix sistema subjacent, de manera que es preservi al màxim la seva estructura i connectivitat. Aquest problema és molt important en diversos àmbits com la biologia, les xarxes socials i la informàtica, i esdevé especialment difícil quan aquestes observacions contenen molt de soroll. En aquest treball, estudiem la detectabilitat de l'alineament veritable entre dos grafs aleatoris correlacionats en funció del soroll. Adoptant una formulació bayesiana i mostrejant alineaments mitjançant mètodes de Markov Chain Monte Carlo, concretament el parallel tempering, explorem com evoluciona l'espai d'energies i si es concentra prop de la solució correct en funció del soroll.

Els nostres resultats mostren dues transicions clau: una algorísmica, en què la recuperació de l'alineament real es perd ràpidament; i un llindar de detectabilitat, a partir del qual la recuperació és impossible. Entre aquests dos llindars, identifiquem una fase en què l'alineament veritable encara és teòricament accessible però esdevé irrecuperable per l'algorisme. Interpretem aquests resultats des de la perspectiva de la física estadística, analitzant el conjunt d'energies i la similitud estructural dels mínims locals amb l'alineament òptim. Els resultats mostren un desacoblament entre les configuracions de baixa energia i la proximitat estructural amb l'alineament subjacent. Aquest estudi contribueix a la comprensió dels límits fonamentals de l'alineament de xarxes i de la naturalesa de la seva dificultat computacional.

Paraules clau: alineament de xarxes, grafs aleatoris, transició de fase de detectabilitat, Markov Chain Monte Carlo, parallel tempering, inferència bayesiana, física estadística, conjunt d'energies, fase difícil.

Acknowledgments

I would like to express my gratitude to Roger and Marta for their help, patience, and insightful guidance throughout this thesis.

I want to thank my colleagues from SEES:lab, especially to Teresa and Gemma for their kindness and the support they have offered me from day one.

I am also grateful to my academic tutor Sergio Gómez for his guidance and assistance throughout my bachelor's degree.

I would also like to thank Professor Àlex Arenas for our insightful conversations and for his encouragement throughout these four years.

I want to express my sincere appreciation to Toni Garijo, whose passion and commitment to teaching have been truly inspiring.

Finally, to my parents for their unconditional support.

All of them have, in one way or another, shaped me into the person I am today.

Contents

List of Figures

1 Introduction

1.1 Motivation and General Context of the Network Alignment Problem

Network alignment is a fundamental problem in network science [1], with applications spanning various domains, including neuroscience, biology and information technology. It involves finding a meaningful mapping between the nodes of two or more networks, which may represent different observations of the same underlying system. One of the most critical challenges in this field is detectability, the ability to correctly infer the alignment given noisy or incomplete information. Understanding the limits of detectability is essential, as it determines whether an alignment can be recovered with high confidence or if the problem is fundamentally unsolvable.

The network alignment problem arises in numerous real-world domains where we wish to find meaningful correspondences between the nodes of two or more networks. A well-known example is neuroscience, where researchers aim to align brain connectomes (complex graphs representing neural connectivity) across different individuals or experimental observations. In this context, successful alignment enables the identification of common functional or anatomical patterns across subjects, which is critical for understanding brain organisation and disorders [2]. However, when the noise or variability between observations exceeds a certain threshold, the true mapping between regions becomes undetectable. This means that the inferred mappings may reflect randomness rather than biologically meaningful structure.

A similar challenge arises in computational biology, particularly in the alignment of protein–protein interaction (PPI) networks, where aligning them across species can reveal evolutionary relationships. If alignment is successful, it can, for example, suggest which proteins in a newly sequenced organism might perform analogous functions to known proteins in a model species [3]. However, for large noise levels (due to high evolutionary divergence or incomplete data), the obtained correspondences lose biological significance, potentially leading to inaccurate conclusions.

Technological and transportation networks also benefit from alignment, for instance, when comparing different urban transit systems or analysing how infrastructure changes over time. A common example is road network inference, where researchers aim to align multiple noisy GPS traces to recover the underlying road structure [4]. In such contexts, excessive noise in the input data can cause meaningful patterns to become indistinguishable from random outcomes. As with other domains, falling into a non-alignable regime means that any real similarities between systems may go unnoticed.

In all of these scenarios, the existence of a detectability threshold establishes a fundamental limit: below this threshold, meaningful alignment is possible and can reveal deep structural insights; above it, the data lacks sufficient signal to distinguish the true alignment from misleading matches. Understanding this transition is crucial for correctly interpreting the results of alignment algorithms and for deciding whether alignment-based analysis is even feasible for a given dataset.

1.2 Objectives of the TFG

The objective of this project is to characterise the phase transition in the alignability of random graphs, both numerically and analytically, focusing on how the ability to recover a latent node mapping decreases as noise increases. This phenomenon, known as the detectability threshold, marks a fundamental transition between a regime where alignment is theoretically feasible and one where it

becomes statistically unattainable [5], [6].

Recent work in network inference has shown that fundamental limits arise not only from noise but also from structural constraints [7], [8]. For instance, in problems closely related to network alignment, such as model selection or link prediction, detectability transitions have been observed, marking sharp boundaries between informative and uninformative regions. These findings reinforce the idea that there exist intrinsic limitations to what can be inferred from network data, even in ideal algorithmic conditions. In this context, our investigation of the detectability threshold in network alignment contributes to a growing line of research that aims to rigorously characterise when and why inference fails.

While many studies explore alignment through the performance of specific algorithms [9], this work aims to propose a more general, approach-independent definition of detectability, grounded in the statistical properties of the underlying network ensemble rather than in any particular method. That is, we seek to identify conditions under which alignment is possible in principle, irrespective of the algorithm used.

To this end, we implement a Bayesian inference approach for network alignment, where the posterior distribution over alignments encodes our uncertainty about the correct matching given the observed graphs. Since the posterior is typically high-dimensional and multimodal, we sample from it using Markov Chain Monte Carlo (MCMC) enhanced by parallel tempering, which allows the exploration of a complex energy landscape more efficiently. Our approach enables us to map out the posterior structure and estimate the likelihood of successful recovery under varying noise levels.

Previous works have proposed other optimisation-based algorithms, but these methods often fall short because they tend to get trapped in local optima easily [10]. This suggests that the energy landscape may be too complex for purely optimisation approaches to be effective, reinforcing the value of probabilistic methods like the one used in this TFG.

1.3 Fundamental Concepts in Network Alignment Detectability

Detectability in network alignment is closely related to several key theoretical concepts that also appear in statistical physics and optimisation problems.

Alignment: It refers to the inferred mapping of node identities in different networks. The challenge lies in distinguishing correct alignments from random ones, particularly when dealing with high levels of noise.

Hard Phase: In many inference and optimisation problems, there exists a *hard phase*, a regime in which the correct solution remains statistically detectable in principle, but standard algorithms fail to recover it in practice. In the context of network alignment, the hard phase refers to a region where the ground truth is still energetically favoured by the posterior distribution but becomes computationally hard to access due to an ergodicity breaking, meaning the algorithm cannot efficiently explore all possible configurations over time.

Detectability Transition: A detectability transition occurs when a small change in the amount of available information (like an increase in the noise level of the networks) leads to a sudden shift from a regime where recovering the true alignment is theoretically possible to one where it is statistically infeasible. This behaviour is common in systems with many competing configurations and limited

information, such as disordered systems like spin glasses, magnetic materials where the interactions between neighbouring spins are randomly set. In such systems, small changes in temperature can lead to abrupt transitions: from a simple (ferromagnetic) phase where all spins align and the system shows global order, to a complex spin glass phase where the system gets trapped in metastable states due to the complex structure of the solution space.

Phase Transition: The detectability transition in network alignment can be described as a phase transition, similar to those observed in physical systems. However, rather than showing gradual changes, our problem exhibits behaviour characteristic of a first-order phase transition, which involves a discontinuous jump in the order parameter (alignability) as the control parameter (noise level) changes. This is in contrast to second-order transitions, like the classical ferromagnetic-to-paramagnetic transition at a critical temperature, where the change in behaviour is continuous. In first-order transitions, such as the liquid-to-gas transition under fixed pressure, the system abruptly switches between distinct phases. Similarly, in network alignment, small changes in noise can suddenly change the system from a phase where the ground truth is detectable to one where it becomes completely hidden among many energetically competitive configurations.

2 Theoretical Background — Network Alignment Problem

In this section, we present the theoretical foundations of the network alignment problem, which consists of inferring node correspondences across multiple observed networks that share an underlying structure. We begin by providing a mathematical formulation of the problem, followed by a probabilistic modelling framework that allows us to incorporate uncertainty and define optimal alignment strategies. These concepts will serve as the basis for the subsequent analysis of detectability thresholds and algorithmic approaches.

2.1 Mathematical Formulation of the Problem

We start by considering K networks of N nodes each and with adjacency matrices $\{A^{(k)}\}_{k=1}^{K}$. Each of them, corresponds to a graph $G^{(k)} = (V^{(k)}, E^{(k)})$ with a common node label set $V^{(k)}$, initially unknown or permuted across networks. We assume that these observed networks have a similar topological structure, which will allow us to map each node in an observation to another node in each of the other observations.

The goal of the network alignment problem is to find the most plausible mapping of node identities $\{\pi^{(k)}\}_{k=1}^{K}$ between the K networks in a way that best preserves the structural similarity between networks.

2.2 Probabilistic Formulation of the Network Alignment Problem

2.2.1 Problem Statement

We consider that all the K observed networks are directed and have binary edges ($A_{ij}^{(k)} \in \{0,1\}$, $\forall i, j \in [N]$, $\forall k \in [K]$[1]). We assume that each $A^{(k)}$ has been generated independently from some latent blueprint graph $L = (V, E)$ with $L_{ij} \in \{0,1\}$, and with some noise $f \in [0,1]$. Now let $\{\pi^{(k)}\}_{k=1}^{K}$ be the set of node mappings from each observation to the blueprint L, or formally, $\pi^{(k)} : V^{(k)} \to V$, $\forall k \in [K]$. That is, if i is a node in the k-th observation ($i \in V^{(k)}$), then $\pi^{(k)}(i)$ denotes its corresponding node in the blueprint L. The goal is to find the alignment $\{\pi^{(k)}\}_k$ and blueprint L that best explain the observations. In probabilistic terms, we aim to maximize the posterior probability

$$p(L, \{\pi^{(k)}\} | \{A^{(k)}\}),$$

thereby obtaining the most plausible blueprint L and alignment $\{\pi^{(k)}\}_k$ under our model.

2.2.2 Modelling of the Posterior Probability

We assume a probabilistic model [5] where each observed graph $A^{(k)}$ is generated from the blueprint L with some noise, such that for all positions where $L_{\pi^{(k)}(i)\pi^{(k)}(j)} = 1$, an edge isn't copied to $A_{ij}^{(k)}$ with

[1] In this context, $[N] := \{1, 2, \ldots, N\}$

probability q, and in positions where $L_{\pi^{(k)}(i)\pi^{(k)}(j)} = 0$, then an edge is created in $A^{(k)}_{ij}$ with probability p. In other words,

$$L_{\pi^{(k)}(i)\pi^{(k)}(j)} = 1 \longrightarrow \begin{cases} A^{(k)}_{ij} = 1, \text{ with probability } 1-q \\ A^{(k)}_{ij} = 0 \text{ with probability } q \end{cases}$$

$$L_{\pi^{(k)}(i)\pi^{(k)}(j)} = 0 \longrightarrow \begin{cases} A^{(k)}_{ij} = 1, \text{ with probability } p \\ A^{(k)}_{ij} = 0 \text{ with probability } 1-p \end{cases}$$

In this way, the probability of observing a non-edge or an edge in a certain position of an observation becomes:

$$p(A^{(k)}_{ij} = 0 | L_{\pi^{(k)}(i)\pi^{(k)}(j)}, p, q) =$$

$$= p(A^{(k)}_{ij} = 0 | L_{\pi^{(k)}(i)\pi^{(k)}(j)} = 0, p, q)\delta_{L_{\pi^{(k)}(i)\pi^{(k)}(j)},0} + p(A^{(k)}_{ij} = 0 | L_{\pi^{(k)}(i)\pi^{(k)}(j)} = 1, p, q)\delta_{L_{\pi^{(k)}(i)\pi^{(k)}(j)},1} =$$

$$= (1-p)\cdot\delta_{L_{\pi^{(k)}(i)\pi^{(k)}(j)},0} + q\cdot\delta_{L_{\pi^{(k)}(i)\pi^{(k)}(j)},1} = (1-p)^{1-L_{\pi^{(k)}(i)\pi^{(k)}(j)}} q^{L_{\pi^{(k)}(i)\pi^{(k)}(j)}},$$

where $\delta_{i,j}$ refers to the Kronecker delta, e.g.:

$$\delta_{L_{\pi^{(k)}(i)\pi^{(k)}(j)},1} = \begin{cases} 1, & \text{if } L_{\pi^{(k)}(i)\pi^{(k)}(j)} = 1 \\ 0, & \text{otherwise} \end{cases}$$

Similarly,

$$p(A^{(k)}_{ij} = 1 | L_{\pi^{(k)}(i)\pi^{(k)}(j)}, p, q) =$$

$$= p(A^{(k)}_{ij} = 1 | L_{\pi^{(k)}(i)\pi^{(k)}(j)} = 0, p, q)\delta_{L_{\pi^{(k)}(i)\pi^{(k)}(j)},0} + p(A^{(k)}_{ij} = 1 | L_{\pi^{(k)}(i)\pi^{(k)}(j)} = 1, p, q)\delta_{L_{\pi^{(k)}(i)\pi^{(k)}(j)},1} =$$

$$p\cdot\delta_{L_{\pi^{(k)}(i)\pi^{(k)}(j)},0} + (1-q)\cdot\delta_{L_{\pi^{(k)}(i)\pi^{(k)}(j)},1} = p^{1-L_{\pi^{(k)}(i)\pi^{(k)}(j)}}(1-q)^{L_{\pi^{(k)}(i)\pi^{(k)}(j)}}$$

So that,

$$p(A^{(k)}_{ij} | L_{\pi^{(k)}(i)\pi^{(k)}(j)}, p, q) = \begin{cases} (1-p)^{1-L_{\pi^{(k)}(i)\pi^{(k)}(j)}} q^{L_{\pi^{(k)}(i)\pi^{(k)}(j)}}, & \text{if } A^{(k)}_{ij} = 0 \\ p^{1-L_{\pi^{(k)}(i)\pi^{(k)}(j)}}(1-q)^{L_{\pi^{(k)}(i)\pi^{(k)}(j)}}, & \text{if } A^{(k)}_{ij} = 1 \end{cases}$$

As each edge of a given observation is generated independently in each position, the conditional likelihood on that observation $A^{(k)}$ is:

$$p(A^{(k)} | L, p, q, \pi^{(k)}) = \prod_{i,j} p(A^{(k)}_{ij} | L, p, q, \pi^{(k)}) = \prod_{i,j} p(A^{(k)}_{ij} | L_{\pi^{(k)}(i)\pi^{(k)}(j)}, p, q)$$
$$= q^{o^k_{10}} p^{o^k_{01}} (1-q)^{o^k_{11}} (1-p)^{o^k_{00}}$$

Note that the number of times that (for example) q appears in this product, is exactly the same as the number of positions i, j such that $L_{\pi^{(k)}(i)\pi^{(k)}(j)} = 1$ and $A^{(k)}_{ij} = 0$, that is, the number of positions where there's an edge in the blueprint but not in the observation $A^{(k)}$. Denote this number by o^k_{10}. The rest of the variables o^k_{XY} with $X, Y \in \{0, 1\}$ are defined in the same way for an observation $A^{(k)}$. Considering all observations, it can be written,

$$p(\{A^{(k)}\}|L, p, q, \{\pi^{(k)}\}) = \prod_{k=1}^{K} p(A^{(k)}|L, p, q, \pi^{(k)}) =$$

$$= \prod_{k=1}^{K} q^{o^k_{10}} p^{o^k_{01}} (1-q)^{o^k_{11}} (1-p)^{o^k_{00}} = q^{O_{10}} p^{O_{01}} (1-q)^{O_{11}} (1-p)^{O_{00}}$$

where now $O_{XY} = \sum_{k=1}^{K} o^k_{XY}$, or in other words, the total number of times that $L_{\pi^{(k)}(i),\pi^{(k)}(j)} = X$ and $A^{(k)}_{ij} = Y$ over all positions i, j and observations $\{A^{(k)}\}$, for $X, Y \in \{0, 1\}$. To obtain the desired posterior distribution, we use Bayes rule,

$$p(L, p, q, \{\pi^{(k)}\}|\{A^{(k)}\}) = \frac{p(\{A^{(k)}\}|L, p, q, \{\pi^{(k)}\})p(L, p, q, \{\pi^{(k)}\})}{p(\{A^{(k)}\})} =$$

$$= p(\{A^{(k)}\}|L, p, q, \{\pi^{(k)}\}) \cdot p(p, q) \cdot \frac{p(L, \{\pi^{(k)}\})}{p(\{A^{(k)}\})} =$$

$$= q^{O_{10}} p^{O_{01}} (1-q)^{O_{11}} (1-p)^{O_{00}} \cdot \frac{p^{\alpha_p - 1}(1-p)^{\beta_p - 1} q^{\alpha_q - 1}(1-q)^{\beta_q - 1}}{B(\alpha_p, \beta_p)B(\alpha_q, \beta_q)} \cdot \frac{p(L, \{\pi^{(k)}\})}{p(\{A^{(k)}\})} =$$

$$= \frac{q^{O_{10}+\alpha_q-1} p^{O_{01}+\alpha_p-1}(1-q)^{O_{11}+\beta_q-1}(1-p)^{O_{00}+\beta_p-1}}{B(\alpha_p, \beta_p)B(\alpha_q, \beta_q)} \cdot \frac{p(L, \{\pi^{(k)}\})}{p(\{A^{(k)}\})}$$

where a Beta prior has been considered for both p, q:

$$p \sim \text{Beta}(\alpha_p, \beta_p), \quad q \sim \text{Beta}(\alpha_q, \beta_q),$$

with $\alpha_p = 0.2, \beta_p = 10, \alpha_q = 0.5, \beta_q = 5$.

Finally, marginalising over p and q we obtain,

$$p(L, \{\pi^{(k)}\}|\{A^{(k)}\}) = \iint_{p,q} p(L, p, q, \{\pi^{(k)}\}|\{A^{(k)}\})dpdq =$$

$$= \frac{p(L, \{\pi^{(k)}\})}{p(\{A^{(k)}\})B(\alpha_p, \beta_p)B(\alpha_q, \beta_q)} \int_0^1 p^{O_{01}+\alpha_p-1}(1-p)^{O_{00}+\beta_p-1}dp \int_0^1 q^{O_{10}+\alpha_q-1}(1-q)^{O_{11}+\beta_q-1}dq =$$

$$= \frac{p(L, \{\pi^{(k)}\})}{p(\{A^{(k)}\})B(\alpha_p, \beta_p)B(\alpha_q, \beta_q)} \frac{\Gamma(O_{00}+\beta_p)\Gamma(O_{01}+\alpha_p)}{\Gamma(Kn_0+\alpha_p+\beta_p)} \frac{\Gamma(O_{10}+\alpha_q)\Gamma(O_{11}+\beta_q)}{\Gamma(Kn_1+\alpha_q+\beta_q)}$$

Where the Beta integral has been used:

$$\int_0^1 x^{a-1}(1-x)^{b-1}\,dx = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)},$$

and where $Kn_0 = O_{00} + O_{01},\ Kn_1 = O_{10} + O_{11}$, with n_0, n_1 the number of non-edges and edges of the blueprint L, respectively.

So finally:

$$p(L, \{\pi^{(k)}\}|\{A^{(k)}\}) \propto \frac{\Gamma(O_{00}+\beta_p)\Gamma(O_{01}+\alpha_p)}{\Gamma(Kn_0+\alpha_p+\beta_p)} \frac{\Gamma(O_{10}+\alpha_q)\Gamma(O_{11}+\beta_q)}{\Gamma(Kn_1+\alpha_q+\beta_q)} \tag{1}$$

2.2.3 Connection with Statistical Physics

The probabilistic formulation of the network alignment problem allows a statistical physics interpretation. Specifically, the posterior distribution over alignments, derived from the generative model of noisy network observations, can be compared to the Boltzmann distribution of a physical system, where the probability of a certain state is given as a function of that state's energy and temperature. This correspondence enables the application of powerful tools from statistical mechanics, especially about energy landscapes and sampling with temperature, to further analyse the alignment problem.

2.2.3.1 Energy Formulation

In analogy with statistical physics, with the posterior we can define an energy function $\mathcal{H}$ of each alignment. Specifically,

$$p(L, \{\pi^{(k)}\}|\{A^{(k)}\}) = \frac{e^{-\mathcal{H}}}{p(\{A^{(k)}\})} \longrightarrow \mathcal{H} = -\log\left(p(L, \{\pi^{(k)}\}|\{A^{(k)}\})\cdot(\{A^{(k)}\})\right),$$

where up to additive constants,

$$\mathcal{H} = -\log\left(\frac{\Gamma(O_{00}+\beta_p)\Gamma(O_{01}+\alpha_p)}{\Gamma(Kn_0+\alpha_p+\beta_p)} \frac{\Gamma(O_{10}+\alpha_q)\Gamma(O_{11}+\beta_q)}{\Gamma(Kn_1+\alpha_q+\beta_q)}\right) \tag{2}$$

In this formulation, high-probability alignments correspond to low-energy configurations, which is suitable for sampling methods such as Markov Chain Monte Carlo with Parallel Tempering.

With this, the complexity of the energy landscape induced by $\mathcal{H}$ is governed by the structural similarity between the networks, or formally, by the overlap values O_{XY}, which are a function of the noise level f.

2.2.3.2 Sampling from the Space of Alignments

Now we have obtained an expression for the posterior probability, from where one can efficiently sample from the space of all possible alignments $\{\pi^{(k)}\}$ and blueprints L. However, this is computationally infeasible due to the combinatorial explosion of possibilities. For this reason, the implementation focuses exclusively on sampling the alignments $\{\pi^{(k)}\}$, while treating the blueprint graph L as a function of the current sampled alignment. Specifically, L is estimated at each step using a majority voting rule, where L has an edge in a certain position only if the majority of the observations also have an edge in that position according to the mapping[2]. We denote this blueprint by $\hat{L}$. Therefore, the energy function defined in (2) can be thought of as a function of just the alignment $\mathcal{H} \mapsto \mathcal{H}(\{\pi^{(k)}\})$

To effectively explore the energy landscape of $\mathcal{H}$, we employ a Markov Chain Monte Carlo (MCMC) approach. In particular, parallel tempering (or replica exchange MCMC) [11] is well suited to sampling multimodal distributions like those encountered in network alignment.

Parallel tempering simulates multiple replicas of the system at different temperatures T, each associated with a modified posterior:

$$p_i(L, \{\pi^{(k)}\}|\{A^{(k)}\}) \propto e^{-\frac{\mathcal{H}}{T_i}}$$

Low-temperature chains concentrate around local optima (providing high-quality alignments), while high-temperature chains explore the configuration space more freely (allowing them to escape local minima). Replicas can exchange their configurations between MCMC steps, which enables a global exploration of the posterior, improving convergence and mixing. With this sampling method, we will explore the configuration space according to the posterior distribution $p(L, \{\pi^{(k)}\}|\{A^{(k)}\})$.

2.3 Detectability — Formal Definition

In the context of network alignment, detectability refers to the fundamental possibility of identifying the ground truth alignment $\{\pi^{(k)}_{true}\}$ between the observed networks $\{A^{(k)}\}$. This concept is strongly related to the energy landscape defined in the previous section, as it provides a quantitative criterion to determine when the true alignment remains statistically and computationally accessible.

Recall that in the energy-based formulation, each possible alignment $\{\pi^{(k)}\}$ is assigned an energy $\mathcal{H}(\{\pi^{(k)}\})$ based on its posterior probability. In this way, $\mathcal{H}(\{\pi^{(k)}_{true}\})$ is denoted as H_{true}, while the lowest energy sampled by an MCMC run will be denoted as H_{min}; in other words, the best possible explanation (according to the posterior) that the inference method can provide for the data, regardless of whether it corresponds to the ground truth.

2.3.1 Noise and Detectability Criterion

For low noise levels f, the observations closely resemble the latent ground truth. Consequently, the posterior distribution favours the true alignment, which clearly stands out among the rest of them, making it attainable by the inference algorithm. However, as the noise level f in the observation model

[2] In case of a tie, a non-edge is introduced.

increases, the structure of the observed networks increasingly diverges from the latent ground truth. Therefore, the posterior distribution becomes broader, and alternative alignments may start to become more plausible than the true one.

This leads to the following formal detectability criterion:

$$\text{Ground truth alignment is } \begin{cases} \textbf{detectable}, & \text{if } H_{true} \leq H_{min} \\ \textbf{not detectable}, & \text{if } H_{true} > H_{min} \end{cases}$$

Furthermore, this condition establishes a phase boundary:

On the one hand, if $H_{true} \leq H_{min}$, then the true alignment is as good as (or better than) the best solution sampled by the inference algorithm. In this case, even if the inference is computationally hard, the space of alignments still energetically favours the ground truth among the rest of them. On the other hand, if $H_{true} > H_{min}$, then there exists at least one alternative alignment that explains the data better than the ground truth. This means that the posterior has begun to favour other alignments, making the true alignment unrecoverable.

From this criterion, the critical point of the phase transition will occur when $\boldsymbol{H_{true}=H_{min}}$.

2.3.2 Phase Transitions and Complexity

The previous allows for a phase diagram in terms of the noise level, which will be deeply studied in future sections.

For low noise levels (small f), the true alignment is easily sampled by the MCMC algorithm in a feasible computation time. In other words, $H_{min} = H_{true}$. At intermediate noise levels, the energy landscape becomes more complex, and it becomes computationally hard to sample from the ground truth, although theoretically possible. Thus, $H_{min} \geq H_{true}$. Finally, at high noise levels, other alignments become highly more sampled than the ground truth, regardless of the computational power provided, as the true explanation for the data becomes very unlikely in this regime. Hence, $H_{min} < H_{true}$.

The detectability transition in this framework corresponds to a shift from a phase in which the ground truth alignment is statistically recoverable to one in which it becomes indistinguishable from a multitude of alternative alignments that achieve lower energy. From the perspective of statistical physics, such transitions are well-known phenomena that take place as external control parameters are varied. In our setting, the analogous control parameter is the noise level f.

3 Results

3.1 Inference Implementation

Once we have derived an expression for the posterior distribution, the goal is to sample alignments from it, favouring those with higher posterior probability (i.e., lower energy). To achieve this, we implement a Markov Chain Monte Carlo (MCMC) algorithm, improved through the use of parallel tempering. This method, inspired by techniques in statistical physics, simulates several Markov chains (or *replicas*) at different temperatures. Each replica explores the alignment space according to its temperature, and they periodically exchange information via probabilistic swap operations.

3.1.1 Markov Chains and Metropolis Criterion

The algorithm begins by generating an initial alignment using a degree-based heuristic: nodes from each observation are ranked by degree (i.e., number of neighbours), and nodes with the same rank across observations are aligned to the same node in the blueprint. This provides a reasonable initialisation that aligns structurally similar nodes.

At each update, a new candidate alignment is proposed by selecting a random observation and swapping two of its nodes in the current alignment. Each replica then decides whether to accept this proposed move $\pi \to \pi'$ using the Metropolis-Hastings acceptance criterion. Let $\beta = 1/T$ be the inverse temperature of the replica, and $\Delta H = H(\pi') - H(\pi)$ the energy difference between the proposed and current states. The move is accepted with probability:

$$p_{acc} = \min(1, e^{-\beta \Delta H})$$

This criterion ensures that lower-energy alignments are always accepted, while higher-energy ones are accepted with a probability that decreases with the energy gap and increases with the temperature. Hence, high-temperature replicas (low β) tend to explore more freely, accepting even worse solutions with a high probability, while low-temperature replicas (high β) focus on refining alignments within low-energy regions.

Each MCMC step consists of performing $\lfloor \frac{N}{2} \rfloor \cdot K$ alignment proposals per replica, where N is the number of nodes in each observation and K is the number of observations. Once this is done for all replicas, an MCMC step has been completed.

3.1.2 Motivation for Replica Exchanges

However, we need a way of exchanging useful alignments between replicas, as we would like low-temperature ones to focus well on low-energy alignments, which are usually found by high-temperature replicas. In other words, we would benefit if replicas could communicate with each other and exchange their alignments after each MCMC step.

This is done with parallel tempering, where two adjacent replicas r and $r+1$ with current states π_r and

π_{r+1} exchange them with probability:

$$p_{swap} = \min(1, e^{-\Delta\beta\Delta H})$$

with $\Delta\beta = \beta_{r+1} - \beta_r$, $\Delta H = H(\pi_{r+1}) - H(\pi_r)$

This exchange mechanism allows low-temperature replicas (which explore the posterior more precisely) to benefit from the exploratory behaviour of high-temperature replicas (which more easily escape local minima). As a result, parallel tempering helps the overall sampler traverse different modes of the distribution more effectively than standard MCMC.

3.1.3 Adaptation to Network Alignment

To ensure efficient communication across all replicas, the exchange pattern alternates between MCMC steps. In even-numbered steps, swaps are attempted between replicas $1 \leftrightarrow 2,\ 3 \leftrightarrow 4,\ \ldots$, while in odd-numbered steps, swaps occur between $2 \leftrightarrow 3,\ 4 \leftrightarrow 5,\ \ldots$. This swap pattern allows information to diffuse across all replicas over time.

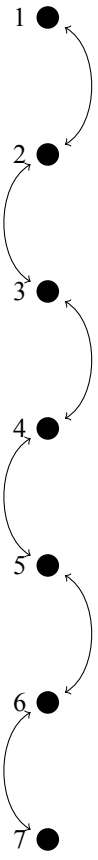

Figure 1. Illustration of the alternating swap pattern used in parallel tempering for the MCMC algorithm.

3.2 Numerical Experiments

To evaluate the performance of our inference algorithm in recovering alignments, we will use synthetic data with controlled noise.

3.2.1 Graph Generation Model

We are considering directed and binary graphs, i.e., where edge positions have value 0 or 1, indicating a non-edge or an edge in that position, respectively. The latent blueprint L is generated from an Erdős-Renyi random graph with $N = 100$ nodes and $p_e = 0.1$ edge probability. Then, we create $K = 2$ noisy observations $G^{(1)}, G^{(2)}$ with some noise f, and finally, we look at the problem of trying to recover the true alignment between both.

The noise process to create a copy $G^{(1)}$ with noise $f \in [0, 1]$ is divided into two phases.

1. Choose a proportion f of all the edges in L and remove them. In this way, $f|E|$ edges have been removed.

2. Select $f|E|$ non-edges in the remaining graph, and add edges in those positions. This new graph will be $G^{(1)}$, a copy of L with noise f.

This methodology ensures that edge density and average node degree are maintained on each observation. After creating these copies, we aim to recover the identity mapping. Each alignment will be represented as a matrix of the following form:

$$\begin{pmatrix} a_1^1 & \dots & a_1^N \\ a_2^1 & \dots & a_2^N \end{pmatrix}$$

Here, the first row is a permutation of the labels of all N nodes in the first observation (same for the second row), and each column groups together the 2 nodes (one from each observation) that are aligned with one another. That is, the nodes in the j-th column ($a_1^j \in G^{(1)}$ and $a_2^j \in G^{(2)}$) will be mapped to the same node in the blueprint, or also $\pi^{(1)}(a_1^j) = \pi^{(2)}(a_2^j)$.[3]

Taking into account that we are simulating the node labels in the observations to match those in the blueprint, the true alignment will correspond to the identity mapping:

$$\begin{pmatrix} 1 & 2 & \dots & N \\ 1 & 2 & \dots & N \end{pmatrix}$$

or, equivalently, any permutation of its columns.

With this information already defined, we run our MCMC algorithm for a total of 350.000 MCMC steps with 20 replicas at temperature $T_i = T_0^{s_i}$, with $T_0 = 1.05$ and where all $s_i \in [0, 30]$ are equally spaced[12]. In this way, the lowest temperature replica, which stores the best results of the run, will be at $T_1 = 1$.

[3] For this reason, note that any permutation of the columns in the matrix gives an equivalent alignment.

3.2.2 Run Examples

After running the algorithm for many different noise levels, three representative examples corresponding to noise levels $f = 0.20, 0.28$, and 0.46 have been chosen to initially show the behaviour of the algorithm and how does it sample the energy landscape. The following figures show the energy of the sampled alignments by each replica at every MCMC step.

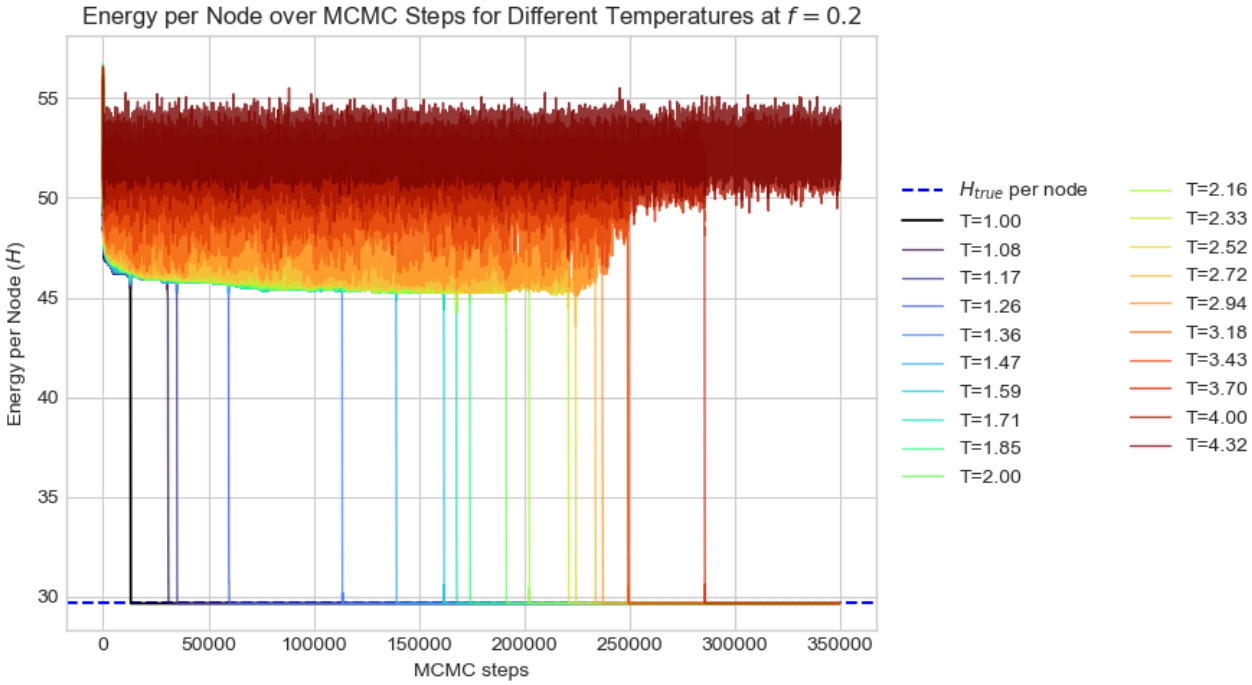

Figure 2. Energy exploration of a run with $f = 0.2$.

In this case, convergence takes place rapidly. As shown in Fig. 2, the replica at $T = 1$ samples from alignments at around $H \approx 46$ until finding the ground truth in about ~ 15.000 MCMC steps. Note that $H_{min} = H_{true}$ for this case.

In contrast, the run depicted in Fig. 3 fails to identify the true alignment and seems to be trapped in local minima around $H \approx 46$ during the whole run. In this case, $H_{min} > H_{true}$.

Finally, for Fig. 4, the behaviour is very similar to that in the previous phase, as it also remains stuck during the whole run, sampling from alignments at $H \approx 46$. However, the energy of the ground truth H_{true} has now become larger than that of many other alignments, so that $H_{min} < H_{true}$. As a result, the ground truth becomes energetically unfavourable, and instead, the algorithm consistently chooses alternative suboptimal alignments of lower energy.

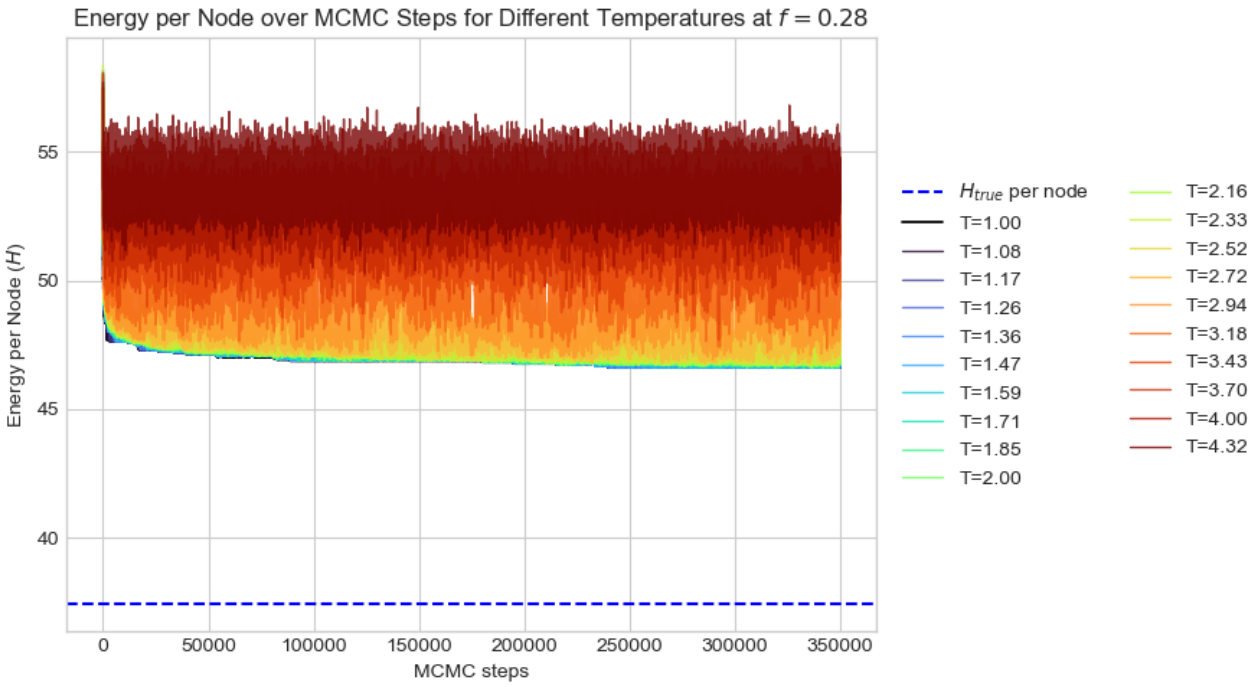

Figure 3. Energy exploration of a run with $f = 0.28$.

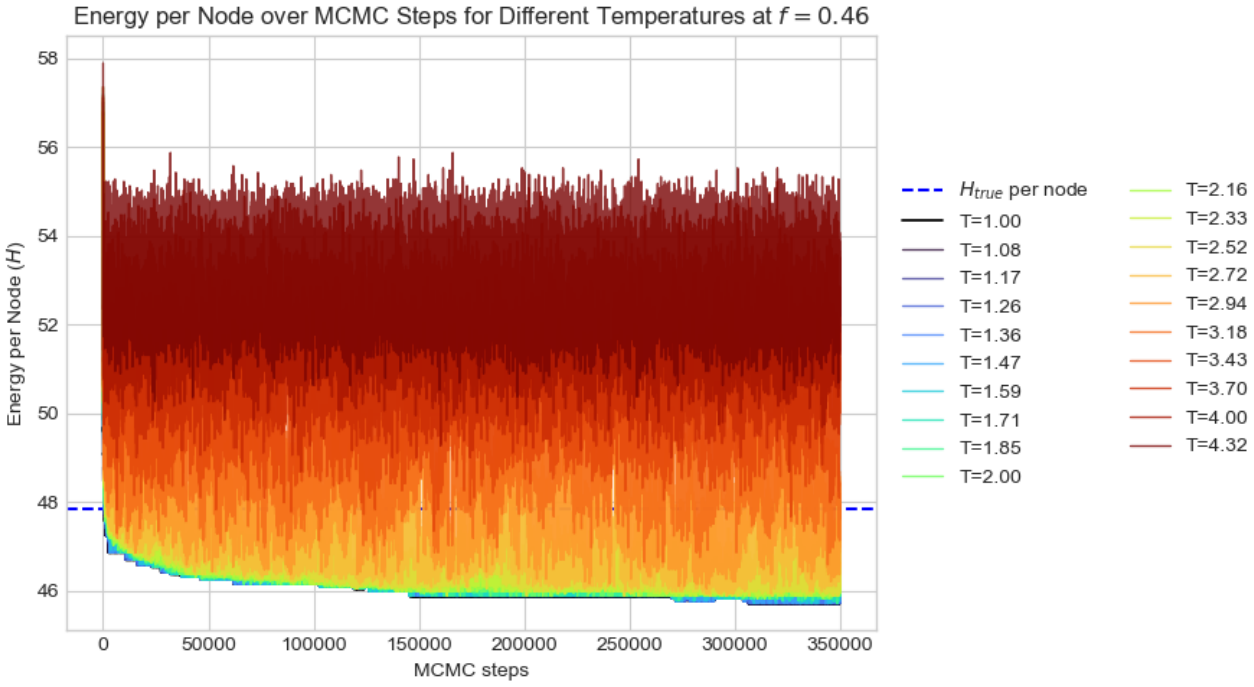

Figure 4. Energy exploration of a run with $f = 0.46$.

3.2.3 Empirical Phase Transition

With this, and taking into account the detectability criterion established in section 2.3.1, we can already find an empirical estimate for the detectability transition. Right below, we plot H_{min} and H_{true} averaged across all sigmas over all runs.

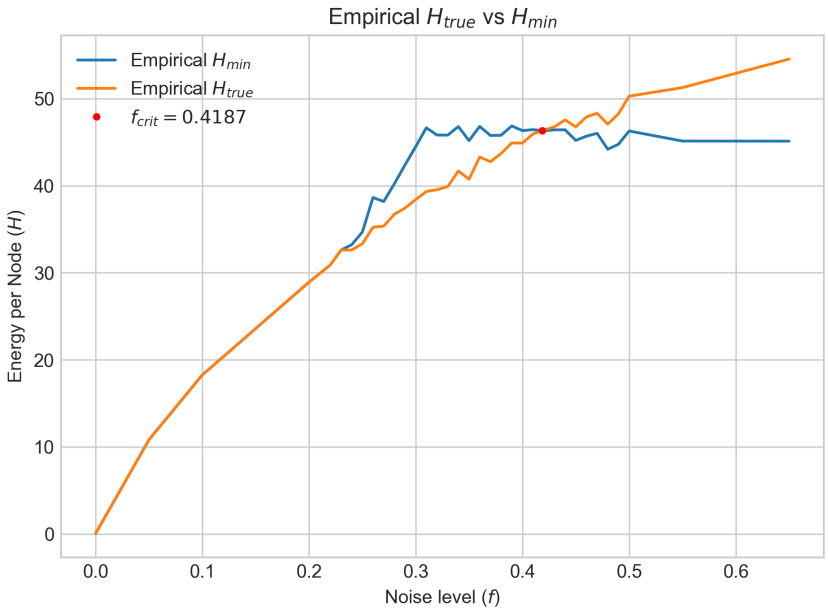

Figure 5. Averaged values of H_{min} and H_{true} across runs vs noise level f. Also, the empirical detectability transition is shown in red.

Fig. 5 shows that in the region $f \in [0, 0.23]$, we have $H_{min} = H_{true}$, implying that the algorithm consistently samples the ground truth alignment across all runs and noise levels in this regime (as expected when the noise is sufficiently low). Then, for noise levels $f \in [0.24, 0.3]$, we begin to observe that $H_{min} > H_{true}$, meaning that a portion of the runs now become unable to recover the true alignment, although the ground truth still remains the global minimum.

Beyond $f \approx 0.31$, the empirical value of H_{min} appears to stabilise around a constant value $H_{min} \approx 46$, suggesting that, at these higher noise levels, the algorithm becomes trapped in suboptimal regions of the configuration space, unable to locate the global minimum. Notably, even within the regime $f \in [0.31, 0.4187]$, we find that $H_{true} < H_{min}$, so that in this region, the true alignment is still theoretically detectable, but the algorithm is no longer able to reach it.

However, at approximately $f \approx 0.4187$, we observe a qualitative shift: $H_{min} < H_{true}$. This marks the point at which the ground-truth alignment stops being the global minimum of the energy function and consequently becomes statistically undetectable. This point can be interpreted as the empirical detectability transition, f_{crit}, where the system no longer energetically favours the true alignment over the rest of them.

It is important to distinguish between two relevant thresholds in this analysis.

First, we have what we will call the *Alignment Accuracy Transition*, which has been empirically

estimated to be around $f_{hard} \approx 0.31$. This represents the noise level beyond which the sampling process conducted by the MCMC algorithm consistently fails to recover the true alignment over all runs, despite it remaining the global optimum.

Then, we have the *Detectability Transition*, which is approximately $f_{crit} \approx 0.4187$. It marks the point after which the true alignment is no longer the global minimum of the energy function (2), so that it becomes theoretically unrecoverable.

This distinction emphasises the difference between algorithmic feasibility and statistical detectability. Let us now examine each of these two transition points in more detail in the following sections.

3.2.4 Alignment Accuracy Transition

In this section, we study the phase transition that occurs in the alignment accuracy of the best solution found by the algorithm as the noise level f increases. Specifically, we define alignment accuracy as the percentage of nodes correctly aligned with respect to the ground truth, and this will be our order parameter. For low noise levels, the algorithm is able to recover the true alignment, resulting in high accuracy. However, as f increases, the structural similarity between the networks becomes weaker and the algorithm struggles to identify the correct correspondence between nodes, causing a sharp drop in the alignment accuracy.

To quantify this transition, we run the alignment algorithm multiple times for increasing values of f, and for each run we compute the percentage of correctly aligned nodes in the best solution found. We then average these results to obtain the expected alignment accuracy as a function of noise.

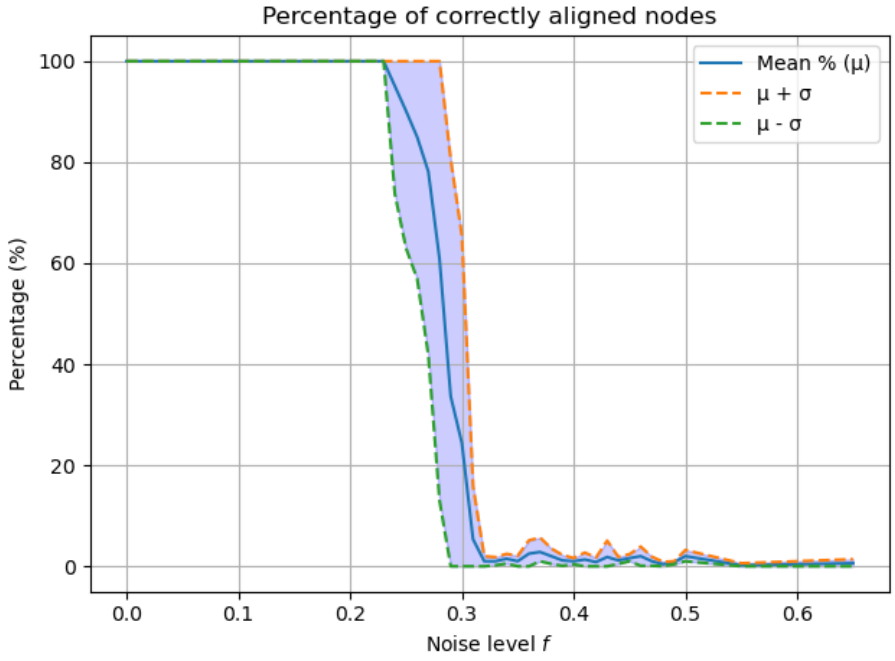

Figure 6. Mean alignment accuracy versus noise level f, with ± 1 standard deviation.

As Fig. 6 shows, the transition is characterised by a rapid decrease in this average accuracy, which we

observe to take place between $f \approx 0.24$ and $f \approx 0.3$. Consequently, below $f_{hard} \approx 0.31$, the alignment algorithm can recover the ground truth, while beyond it, the performance drops heavily, indicating a transition to a regime where alignment becomes significantly more difficult. Note that this value for the alignment accuracy transition aligns with the one estimated in the previous section.

A key observation from the results is the **absence of partially correct alignments**. That is, for each run, the alignment accuracy tends to fall into one of two extreme regimes: either the algorithm recovers the ground truth perfectly (achieving 100% accuracy), or it fails entirely and produces an energetically favourable but almost completely incorrect alignment (typically with only $1-3\%$ of correctly-aligned nodes). Intermediate cases, such as alignments with 30%, 50%, or 90% accuracy, are never the lowest-energy sampled alignments in practice.

This strongly suggests that the energy landscape doesn't contain near-optimal local minima, with the ground truth standing apart as a very distinct global minimum. That is, the true alignment differs widely from other stable regions of the configuration space, which contain local minima that, although energetically favourable, are structurally unrelated to it (this can be observed in more detail in section A). This kind of behaviour, where the algorithm either succeeds or fails completely, is characteristic of first-order phase transitions in physical systems.

3.2.5 Detectability Transition

Here, we focus on another phase transition of the network alignment problem: the fundamental limit of alignment detectability. This transition is not about the quality of the recovered alignment, but about whether it is even possible for the algorithm to identify the true alignment as a plausible solution.

To study this, we plot the alignability percentage across different noise levels f, defined as a binary variable that takes the value 100% if the energy of the true alignment satisfies $H_{\text{true}} \leq H_{\text{min}}$, and 0% otherwise. Intuitively, this measures whether the energy from the true alignment is energetically favoured among the rest or not.

By averaging this binary value over multiple runs for each noise level f, we obtain an empirical estimation of the probability that the true alignment is detectable. As illustrated in Fig. 7, alignability remains at 100% for low noise levels, but has a sharp drop around $f_{crit} \approx 0.42$, eventually reaching 0% after that noise level (which again matches the first estimate for the detectability transition, observed in Fig. 5).

This behaviour signals a fundamental shift: beyond this threshold, the ground truth alignment becomes inaccessible to the algorithm because other sampled alignments have lower energy values. In this regime, the true alignment becomes statistically undetectable, as the ground truth no longer corresponds to the lowest energy configuration found by the algorithm.

3.2.6 Comparing Both Transitions

The two transitions observed before highlight the same underlying phenomenon from different perspectives: the increasing difficulty of recovering the true alignment as noise increases.

The first transition, identified in Fig. 6, occurs around $f_{hard} \approx 0.31$ and corresponds to the threshold

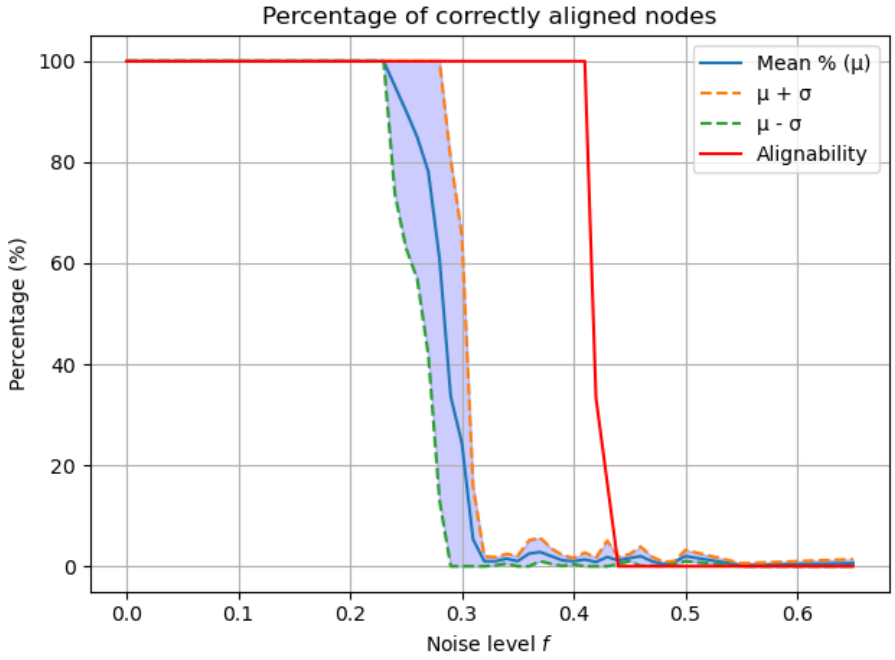

Figure 7. Alignability percentage in red as a function of noise level f. The alignment accuracy is also shown.

beyond which the algorithm quickly decays in performance, as evidenced by a sharp decline in the number of correctly aligned nodes of the best alignment sampled. This drop reflects a performance limitation of the algorithm: even though the ground truth may still be the global minimum of the energy landscape, the algorithm increasingly fails to find it due to the growing complexity of the energy landscape.

The second transition, shown in Fig. 7, takes place at a higher noise level, around $f_{crit} \approx 0.4187$. It marks the point beyond which the true alignment is no longer favourable energetically, since the algorithm can now sample from other alignments with lower energy. This indicates a fundamental loss of detectability.

Now, the empirical estimation of the detectability transition at $f \approx 0.4187$ serves as a motivation to try to determine it analytically, as a theoretical perspective on when the true alignment stops being energetically favourable would provide a deeper understanding of the detectability limit.

At the same time, the existence of a gap between the alignment accuracy transition and the detectability transition raises the question of what happens in the intermediate regime $f \in [f_{hard}, f_{crit}]$, where the true alignment remains the global minimum of the energy function but the algorithm consistently fails to find it. This suggests the possible existence of a hard phase, where the configuration space becomes fragmented and the system breaks ergodicity: the true alignment is theoretically reachable, but becomes inaccessible (even with limitless computational resources), unless the algorithm is initialised within a specific basin of attraction. Alternatively, this regime might not reflect a distinct phase but rather a region where the convergence time increases rapidly as $f \to f_{crit}$ and the observed failure is simply due to insufficient sampling time within the algorithm's maximum number of MCMC steps. These two

interpretations lead to different views of the system's behaviour, and will tell us whether the barrier we see is merely algorithmic or intrinsic to the system itself.

3.3 Explanation of the Transition

In the previous section, we observed that the transition in alignment accuracy occurs at a lower noise level than the detectability threshold. To understand the origin of the observed transitions, we will aim to theoretically estimate the detectability threshold via the detectability criterion in 2.3.1.

3.3.1 Theoretical Calculation of H_{true}

In this scenario, we continue to consider the case with $K = 2$ observations. Since the blueprint $\hat{L}$, constructed using the majority voting rule, includes only those edges that appear in both observations, it follows that $\boldsymbol{O_{10}=0}$; in other words, there are no positions where an edge is present in $\hat{L}$ but absent in either observation.

Besides, remember the relations:

$$O_{00} + O_{10} = Kn_0^*, \quad O_{01} + O_{11} = Kn_1^* \tag{3}$$

$$O_{00} + O_{01} = Kn_0, \quad O_{10} + O_{11} = Kn_1$$

with n_0^*, n_1^* being the number of non-edges and edges of each observation, respectively, and n_0, n_1 the non-edges and edges of the blueprint $\hat{L}$.

Given the possible number of node pairs in a directed graph is $N(N-1)$, and a probability of an edge being present p_e, we estimate:

$$n_0^* \approx N(N-1)p_{\text{non-edge}} = N(N-1)(1-p_e) \Rightarrow \boldsymbol{O_{00}} = Kn_0^* - O_{10} \boldsymbol{\approx 2N(N-1)(1-p_e)}$$

Similarly,

$$n_1^* \approx N(N-1)p_e \Rightarrow O_{01} + O_{11} = Kn_1^* \approx 2N(N-1)p_e$$

Here, O_{11} represents the number of nodes in an observation that are also edges in the blueprint. Since an edge is included in $\hat{L}$ if and only if it appears in both observations (each common edge will create an edge in $\hat{L}$), we can estimate O_{11} as twice the number of common edges between the two observations. That is, $O_{11} = 2n_1$. We will study and estimate this quantity by looking at how each observation G_1, G_2 is constructed from an initial underlying graph G. By the way this procedure is defined, the number of edges of G and each observation is the same. Thus, $|E| = n_1^* \approx N(N-1)p_e$ and $|\bar{E}| = n_0^* \approx N(N-1)(1-p_e)$.

We introduce some notation. Recall that n_1 is the number of edges in $\hat{L}$, or in other words, the number of common edges between the two observations. Define the random variables:

- $X =$ common edges that survive in both observations after Phase 1 (survival edges),
- $Y =$ common edges resulting from one edge surviving and the other being reintroduced in Phase 2,
- $Z =$ edges that are simultaneously created in both observations during Phase 2.

Then, $n_1 = X + Y + Z$.

The variable X follows a binomial distribution with success probability $p_X = (1-f)^2$, since each original edge must survive in both observations. Besides, X takes values between $\max(0, |E| - 2f|E|)$ and $|E| - f|E|$, depending on whether the removed edges in both observations are disjoint or identical.

As $X \sim \text{Bin}(|E|, p_X)$, it follows:

$$\mathbb{E}[X] = \sum_{x=\max(0,E-2f|E|)}^{E-f|E|} x\mathbb{P}(X=x) \approx n_X p_X = |E| \cdot (1-f)^2$$

Next, we consider Y. The probability that one edge survives in phase 1 is $1-f$, and the likelihood that one edge is removed and then reintroduced is $f\frac{f|E|}{|\bar{E}|+f|E|}$. since this can occur in either order, we obtain: $p_Y = 2f(1-f)\frac{f|E|}{|\bar{E}|+f|E|}$. Additionally, Y will take values from 0 to $2f|E|$.

Since $Y \sim \text{Bin}(|E|, p_Y)$:

$$\mathbb{E}[Y] = \sum_{y=0}^{2f|E|} y\mathbb{P}(Y=y) \approx n_Y p_Y = |E| \cdot 2f(1-f)\frac{f|E|}{|\bar{E}| + f|E|} = 2\frac{f^2|E|^2(1-f)}{|\bar{E}| + f|E|}$$

Lastly, we compute Z. The number of available places where two edges can be introduced in both observations is, on average, $|\bar{E}| + f^2|E|$, and the probability of both observations adding an edge at the same position is $p_Z = \left(\frac{f|E|}{|\bar{E}|+f^2|E|}\right)^2$. Note that Z will take values from 0 to $f|E|$.

Since $Z \sim \text{Bin}(|\bar{E}| + f^2|E|, p_Z)$:

$$\mathbb{E}[Z] = \sum_{z=0}^{f|E|} z\mathbb{P}(Z=z) \approx n_Z p_Z = (|\bar{E}| + f^2|E|) \cdot \left(\frac{f|E|}{|\bar{E}| + f^2|E|}\right)^2 = \frac{f^2|E|^2}{|\bar{E}| + f^2|E|}$$

Putting everything together, the expected number of common edges between both observations will be:

$$\mathbb{E}[n_1] = \mathbb{E}[X+Y+Z] = \mathbb{E}[X] + \mathbb{E}[Y] + \mathbb{E}[Z] \approx$$

$$\approx |E|(1-f)^2 + 2\frac{f^2|E|^2(1-f)}{|\bar{E}|+f|E|} + \frac{f^2|E|^2}{|\bar{E}|+f^2|E|}$$

So that,

$$\boldsymbol{O_{11} \approx 2\mathbb{E}[n_1] \approx 2N(N-1)p_e\left((1-f)^2 + 2\frac{f^2(1-f)}{\frac{1-p_e}{p_e}+f} + \frac{f^2}{\frac{1-p_e}{p_e}+f^2}\right)}$$

as $\frac{|\bar{E}|}{|E|} = \frac{N(N-1)(1-p_e)}{N(N-1)p_e} = \frac{1-p_e}{p_e}$. The result above, directly implies that,

$$\boldsymbol{O_{01} = Kn_1^* - O_{11} \approx 2N(N-1)p_e\left(1-(1-f)^2 - 2\frac{f^2(1-f)}{\frac{1-p_e}{p_e}+f} - \frac{f^2}{\frac{1-p_e}{p_e}+f^2}\right)}$$

Given the number of nodes of the graph N and edge density p_e, the previous gives an estimation for all values of O_{XY} with respect to the noise level f. Then, we can substitute these values in the expression for the energy of an alignment, and recalling that $Kn_0 = O_{00} + O_{01},\ Kn_1 = O_{10} + O_{11}$, and the values for the hyperparameters of the Beta prior distribution $\alpha_p = 0.2, \beta_p = 10, \alpha_q = 0.5, \beta_q = 5$, we can approximate $H_{true}(f)$ in terms of the noise level f for each N, p_e:

$$H_{true} = -\ln\left(\frac{\Gamma(O_{00}+\beta_p)\Gamma(O_{01}+\alpha_p)\Gamma(O_{10}+\alpha_q)\Gamma(O_{11}+\beta_q)}{\Gamma(Kn_0+\alpha_p+\beta_p)\Gamma(Kn_1+\alpha_q+\beta_q)}\right) \tag{4}$$

In this way, considering an underlying blueprint with $N = 100$ and $p_e = 0.1$, the values for O_{XY} are:

$$O_{00} \approx 17820$$
$$O_{01} \approx 1980\left(1-(1-f)^2 - 2\frac{f^2(1-f)}{9+f} - \frac{f^2}{9+f^2}\right)$$
$$O_{10} = 0$$
$$O_{11} \approx 1980\left((1-f)^2 + 2\frac{f^2(1-f)}{9+f} + \frac{f^2}{9+f^2}\right)$$

Substituting them into (4), we obtain a numerical expression for H_{true}. We would like to evaluate the approximation of this model, so we will plot $H_{true}(f)$ along with some experimental data points, where the average energy of the ground truth for 1000 blueprints every 0.025 noise level is displayed.

Fig. 8 indeed shows that the theoretical curve closely follows the empirical values of the transition point, confirming that the analytical estimation for $H_{true}(f)$ accurately captures the behaviour observed in simulations and thus validating our approximation of the overlap quantities O_{XY}.

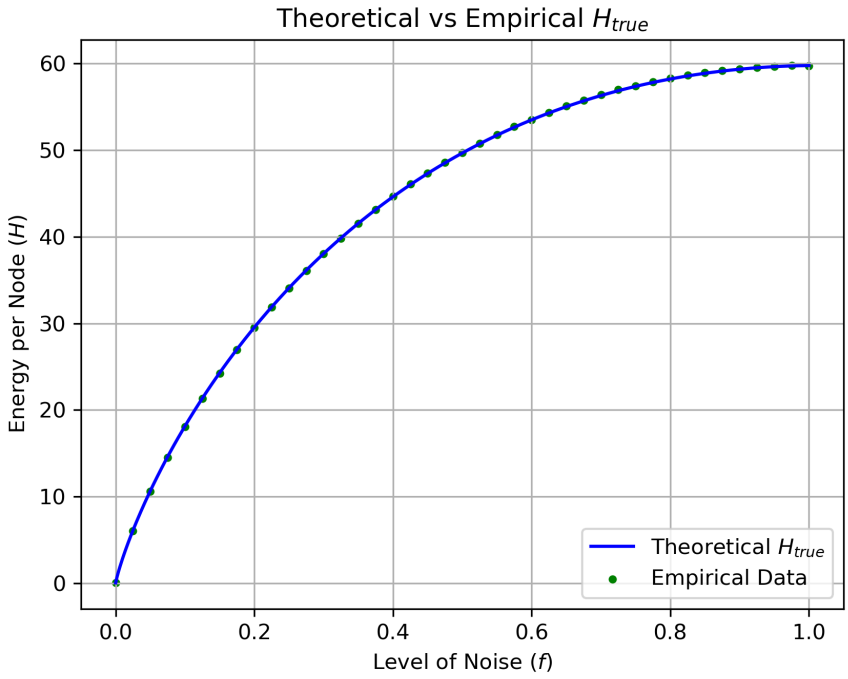

Figure 8. Theoretical H_{true} vs averaged empirical data of H

3.3.2 Energy Distribution and Theoretical Calculation of Transition Point

Having obtained a theoretical expression for H_{true}, we now analyse the distribution of energies across the alignment space to theoretically estimate H_{min} and consequently determine the critical noise level f_{crit} at which the true alignment stops being energetically favourable.

3.3.2.1 Gaussian Distribution Approximation

To start with, we would like to understand the shape of the underlying energy landscape. For this reason, we generate two independent observations from the same underlying graph at some noise level, e.g., $f = 0.4$. We then randomly sample 10^6 alignments from the entire space of possible configurations and compute the energy, recording the associated energy values for each.

As shown in Fig. 9, the kernel density estimate (KDE) of the sampled energies resembles a Gaussian distribution. To observe this, a normal distribution has been fitted using the empirical mean μ and standard deviation σ of the sampled energies.

To further support this observation, the experiment was repeated multiple times using different underlying graphs generated at the same noise level ($f = 0.4$). For each graph, 10^6 random alignments were sampled, and the corresponding energy values were recorded. After fitting a Gaussian to each resulting energy distribution, the empirical values of μ and σ were averaged across all runs, which provided initial estimates for this noise level.

Remarkably, repeating this experiment for different noise levels f causes the resulting energy

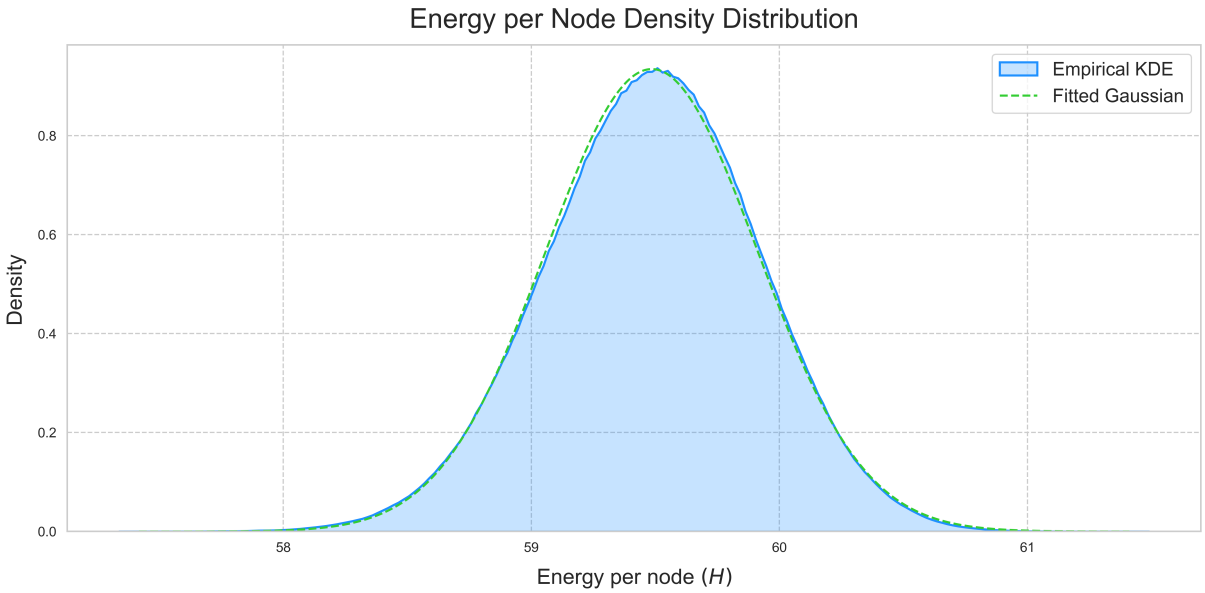

Figure 9. Empirical distribution of energies for 10^6 randomly sampled alignments at $f = 0.4$. A Gaussian curve fitted to the data using the empirical mean and standard deviation is also shown.

distributions to maintain their Gaussian shape, and the fitted parameters μ and σ remain practically unchanged. This suggests that the distribution of energies over the space of random alignments is essentially independent of the noise level f used to generate the graph observations, all of which can be approximated by a Gaussian distribution.

Consequently, to further estimate μ and σ that can serve as Gaussian parameters for any noise level, the experiment was extended to cover 50 evenly spaced noise levels between 0 and 1. For each of these values, the same procedure was followed: generating two observations, sampling random alignments, and computing the fitted Gaussian parameters. The final reference values, obtained after averaging across all 50 cases, were $\mu \approx 59.77$ and $\sigma \approx 0.429$. These averaged values will be used as the representative parameters of the energy distribution under random alignments throughout the remainder of the work.

The strong agreement between the KDE and the fitted Gaussian suggests that, in the limit of a large number of alignments, the energy distribution over the space of random alignments approximates a normal distribution. This provides a probabilistic interpretation of the energy landscape that will be useful for the next section.

3.3.2.2 Determining the Minimum Energy per node H_{min}

We have now observed that the density of all energy states aligns closely with the probability distribution function (PDF) $p_{\mu,\sigma}(x)$ of a Gaussian random variable. It is well known that the probability of sampling a value within $[a, b]$ can be found as:

$$\mathbb{P}(a \leq X \leq b) = \int_a^b p_{\mu,\sigma}(x)dx$$

Thus, the probability of sampling an energy in the range $[h_{min}, h_{max}]$ is:

$$\mathbb{P}\left(h_{min} \leq H \leq h_{max}\right) = \int_{h_{min}}^{h_{max}} p_{\mu,\sigma}(x)dx = \int_{h_{min}}^{h_{max}} \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} dx$$

Besides, for two directed, unweighted networks of N nodes, the total number of possible alignments is given by $N!$[4]. In this way, we can now find the minimum alignment energy that can be achieved for a given pair of networks. Since there are $N!$ total alignments, the lowest energy can be approximated as the first energy value where the cumulative probability of finding an energy between $-\infty$ and x is precisely $1/N!$. In this way, H_{min} can be found with respect to the number of nodes N in the graph from the following equation:

$$\int_{-\infty}^{H_{min}} \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} dx = \frac{1}{N!} \overset{\mu=59.77,\sigma=0.429,N=100}{\longrightarrow} H_{min} \approx 48.27 \tag{5}$$

3.3.2.3 Theoretical Calculation of Transition Point

With this, equating H_{min} to the theoretical H_{true} expression derived in (4) will give the transition point, as the phase transition takes place when $H_{true} = H_{min}$. Plotting both energies, we can visualise the transition.

As shown in Fig. 10, the theoretical phase transition is estimated to happen at $f \approx 0.4694$. However, this value overestimates the empirical transition point obtained in 3.2.3, which was around $f \approx 0.4187$. Since the approximations for the theoretical H_{true} agree well with empirical results (as shown in Fig. 8), the discrepancy must originate from an imprecise estimation of H_{min}.

3.3.3 Discrepancy Between Theoretical and Empirical Transition Point

As we concluded that the differences may be due to the theoretical calculation of H_{min}, we will evaluate the results obtained with the empirical data of the minimum energy sampled by the alignment in practice.

Theoretically, we obtained $H_{min} \approx 48.27$. However, empirical measurements suggest that H_{min} stabilises around $H_{min} \approx 46$ once the noise level surpasses $f_{hard} = 0.31$, as previously discussed in section 3.2.3. To refine this empirical estimate, we average all observed empirical H_{min} corresponding to noise levels after f_{hard}, obtaining an empirical mean of $H_{min} \approx 45.97$. This can be seen in Fig. 11,

[4]This arises because the first node from the first network can be paired with N nodes from the second one, node 2 with any of the remaining $N-1$ nodes, and so on until node N, which will be aligned with the remaining node to pair.

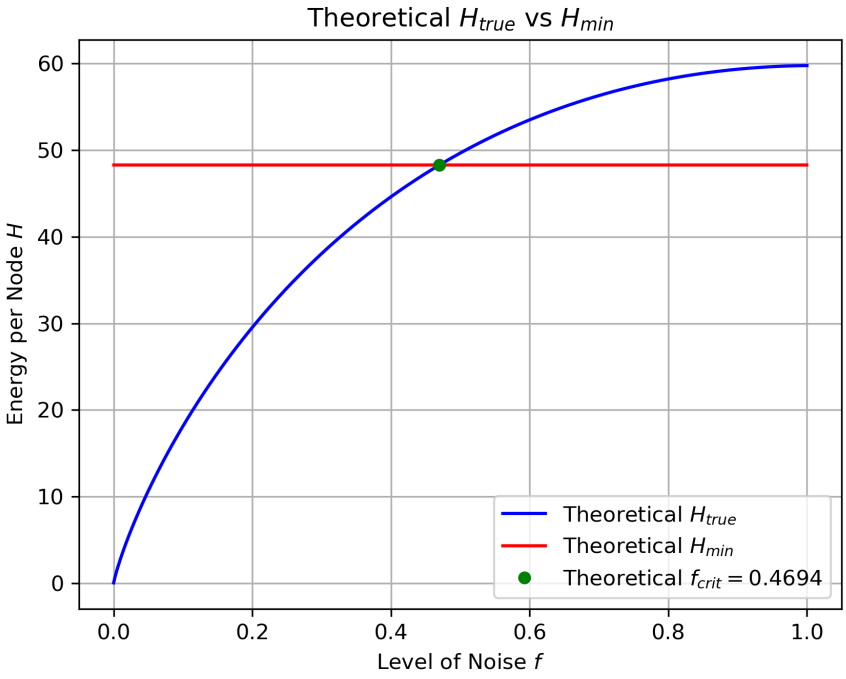

Figure 10. Plot of the theoretical values of H_{true} and H_{min} as a function of the noise level f for $N = 100$, $p_e = 0.1$

where the vertical gap between the theoretical and empirical lines highlights the inaccuracies between both approaches.

A possible explanation for this deviation lies in the assumption made earlier regarding the assumed shape of the energy distribution. In section 3.3.2.1, we approximated the distribution of alignment energies by a Gaussian distribution, which although visually similar, it deviates from a perfect Gaussian, as the latter is not symmetric.

This asymmetry comes from a fundamental constraint in the problem itself: the energy landscape has a fixed lower bound. As a result, extremely low energy values (which would exist under a Gaussian) cannot actually occur in the real system, so that the distribution is effectively cut off at this energy, leading to a shorter left tail (low-energy region) than that of the Gaussian.

This means that, while the Gaussian approximation may capture the bulk of the distribution well, it underestimates the probability mass near the minimum energy, thereby slightly overestimating the expected minimum energy in the hard phase. A more accurate estimation of H_{min} would require a refined model of the energy distribution that accounts for its asymmetry, something that lies beyond the scope of this TFG.

Correcting for this issue in the theoretical minimum energy, we substitute the mean empirical estimate $H_{min} \approx 45.97$ in our detectability criterion, and using this refined result, we find that the transition point now lies at approximately $f \approx 0.4240$. This is also plotted in Fig. 11 as the intersection point between the theoretical curve for H_{true} and the horizontal line at the mean empirical minimum energy. Now this corrected estimate aligns much more closely with the previous empirical transition point of

$f_{crit} \approx 0.4187$ found in section 3.2.3.

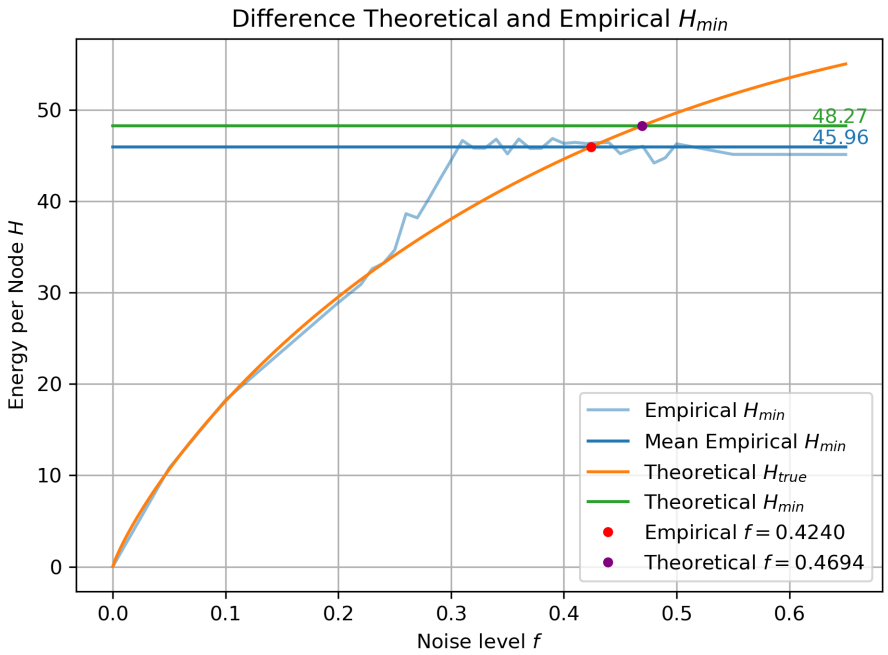

Figure 11. Plot of the theoretical and empirical H_{min}, together with the mean empirical H_{min}, the theoretical H_{true} and both detectability transition points.

3.3.4 Expected MCMC Discovery Time and Critical Scaling Behaviour

After theoretically analysing the detectability transition and the limitations of the assumptions for H_{min}, we now examine the dynamics of the sampling process itself. In particular, we focus on the expected number of MCMC steps required to discover the ground truth alignment (or a better one).

This is especially relevant in the intermediate region $f \in [f_{hard}, f_{crit}]$, where no run succeeds in finding the ground truth. Understanding whether this failure is due to ergodicity breaking (where large regions of the alignment space become effectively inaccessible to MCMC) or to a first-order phase transition (before which success is still possible but increasingly rare) can be studied by analysing how the discovery time scales with the noise level.

In this section, we estimate and interpret this quantity across different noise levels to gain insight into the underlying transition of the alignment problem.

We model the search process as a Poisson process, where each MCMC step is treated as a Bernoulli trial with success probability ρ (which represents the probability of finding either the ground truth or a lower-energy alignment at each step). Then, the number of MCMC steps until the first success follows a geometric distribution, which will be denoted by $\mathcal{X}$.

In this way, for a run $\mathcal{X}_i$, the probability of succeeding at MCMC step m_i is

$$p\left(\mathcal{X}_i = m_i \mid \rho\right) = (1-\rho)^{m_i - 1}\rho$$

and the probability of not succeeding after M steps is

$$p\left(\mathcal{X}_i > M \mid \rho\right) = \sum_{m=M}^{\infty} (1-\rho)^m \rho = (1-\rho)^M$$

Then, given a set of runs $\{\mathcal{X}_i\}$ with MCMC steps $\{m_i\}$, the likelihood $p(\{m_i\} \mid \rho)$ is just the product of the $p(\mathcal{X}_i \mid \rho)$ with the corresponding m_i (or the undefined $m_i \geq M$ when we do not succeed). Using Bayes theorem, we have that

$$p(\rho \mid \{m_i\}) \propto p(\{m_i\} \mid \rho) p(\rho)$$

and, assuming that $p(\rho) = \text{const.}$, we can calculate

$$p(\rho \mid \{m_i\}) \propto \prod_{i \in success} (1-\rho)^{m_i - 1}\rho \prod_{i \in fail} (1-\rho)^M = \rho^{N_{success}} (1-\rho)^{M_T - N_{success}}$$

$$\left.\frac{\partial \log p\left(\rho \mid \{m_i\}\right)}{\partial \rho}\right|_{\rho^*} = 0 \Longrightarrow \rho^* = \frac{N_{success}}{M_T}$$

where $N_{success}$ is the number of times we managed to converge and $M_T = \sum_i{}' m_i = \sum_{success} m_i + MN_{fail}$ (the $'$ in the sum indicates that m_i is set to the maximum M when the MCMC did not converge).

As the expected value for a geometric distribution is given by $\mathbb{E}[\mathcal{X}] = 1/\rho$, we can see that the expected number of MCMC steps until success that maximises the posterior distribution is:

$$\mathbb{E}[\mathcal{X}] = \frac{1}{\rho^*} = \frac{M_T}{N_{success}}$$

Given all the experiments, the results are shown in Fig. 12. It illustrates the expected number of MCMC steps $\mathbb{E}[\mathcal{X}]$ as a function of the noise level f, on a logarithmic scale. As expected, the discovery time grows with increasing noise, reflecting the growing difficulty of finding the ground truth or a better alignment.

Importantly, for noise levels above $f_{hard} \approx 0.31$, all observed runs failed to find the ground truth or any better alignment, and thus no successful events were recorded in the interval $f \in [0.31, 0.41]$. As a result, there is no empirical data in this intermediate region. However, by looking at the increase in $\mathbb{E}[\mathcal{X}]$ before this gap, it seems to diverge at the detectability transition value, and not anywhere before it, which would suggest the presence of a hard phase where the ground truth is theoretically recoverable but inaccessible due to ergodicity breaking of the system.

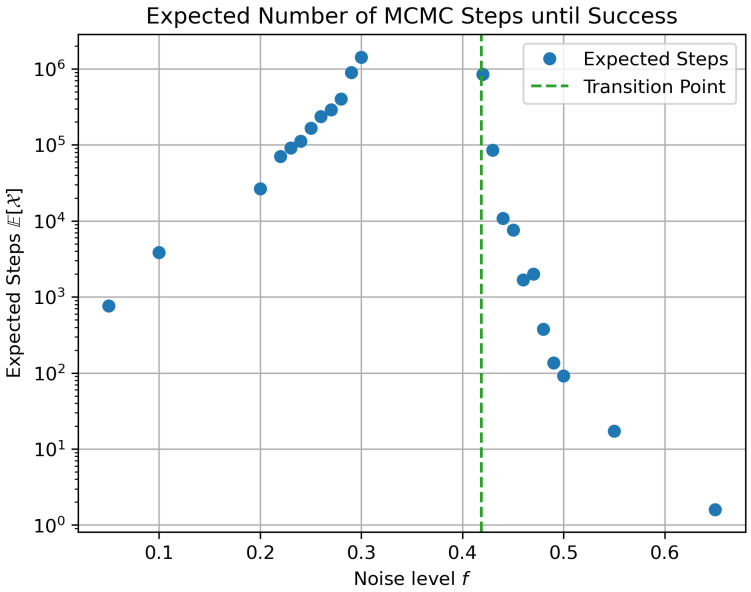

Figure 12. Expected Number of MCMC Steps $\mathbb{E}[\mathcal{X}]$ in log scale with respect to noise level f, and the detectability transition point.

In contrast, a divergence in $f_{crit} = 0.4187$ is consistent with a first-order phase transition, where the system remains ergodic (though increasingly hard) until a sharp threshold is reached, beyond which recovery becomes essentially impossible. Thus, the lack of data in the intermediate region is likely caused by the maximum number of MCMC steps being insufficiently large, which has been set to $M = 3.5 \cdot 10^5$. Thus, increasing the number of steps would eventually lead to the success of the algorithm in finding the ground truth.

Moreover, the observed divergence pattern near the transition suggests the possibility of fitting a power-law of the form:

$$\mathbb{E}[\mathcal{X}] \propto (f_{crit} - f)^{-\gamma}, \quad \text{as } f \to f_{crit}^{-}$$

This form of divergence is a well-established feature in statistical physics, where quantities such as correlation lengths, susceptibilities, or relaxation times often diverge near phase transitions following power laws. In our context, it provides a way to estimate the critical exponent γ and better characterise the scaling behaviour of the MCMC sampling process near the detectability transition value.

The results of this fit are shown in Fig. 13. The data is well approximated by a power law with parameters $A \approx 0.9$, $\gamma \approx 6.7$. This scaling behaviour suggests that the expected number of MCMC steps required to recover the ground truth alignment increases abruptly as the noise level approaches f_{crit} from below. The high value of the critical exponent γ implies a very steep divergence, which is indicative of a rapidly growing sampling difficulty near the transition point. This type of divergence is characteristic of systems undergoing a first-order (discontinuous) phase transition, further supporting this interpretation of the detectability threshold.

Besides, this observation supports the idea that the gap may be due to a limited maximum number of

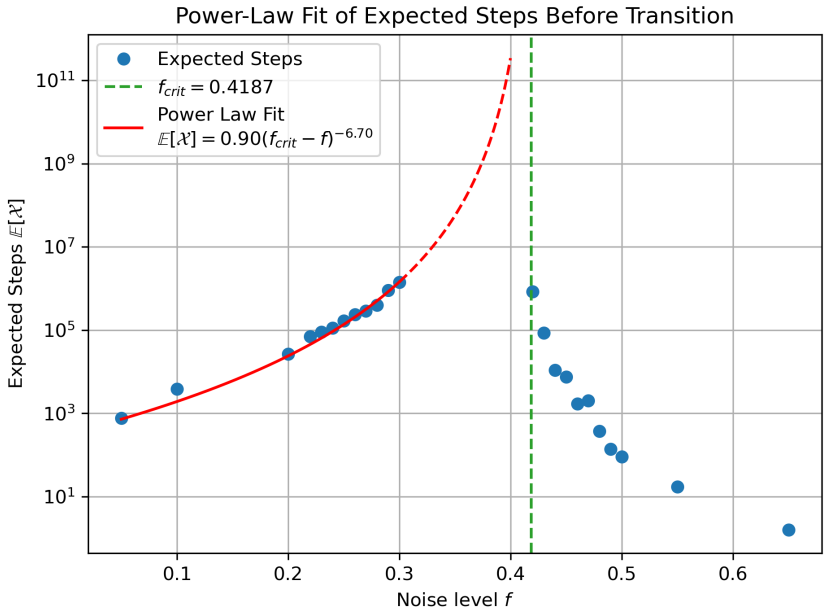

Figure 13. Power-law fit of the expected number of MCMC steps $\mathbb{E}[\mathcal{X}]$ as a function of the noise level f for values below the critical threshold $f_{crit} = 0.4187$.

MCMC steps M allowed in the simulations. For instance, having allowed $M = 10^7$ MCMC steps, successful runs could potentially have been observed at higher noise levels (up to approximately $f \approx 0.33$, according to the power-law extrapolation).

Additionally, the extrapolated section of the fit (dashed line in the plot) extends beyond the fitted region and visually suggests how the divergence would continue to increase towards the critical point. This analysis indicates that the divergence in $\mathbb{E}[\mathcal{X}]$ occurs precisely at f_{crit}, with no evidence of an earlier divergence that could correspond to an intermediate ”hard phase”. Of course, this does not rule out the possible existence of a narrow hard phase, but according to our data from the experiments, the divergence seems to take place at the noise level f_{crit}.

Finally, it is important to note that this power-law behaviour is an empirical fit, and although it aligns with similar systems in statistical physics, a more careful derivation would be required to fully characterise the nature of this transition.

3.3.5 Energy–Structure Decoupling Near the Phase Transition

We have seen that the algorithm often fails to recover the ground truth, especially for noise levels close to the phase transition. In this section, we examine why this difficulty arises in the first place and what is the fundamental problem that we face when sampling from the space of alignments.

To address this, we did additional simulations on noise levels close to the phase transition, where the algorithm was no longer initialised from a degree-based heuristic, but rather from an alignment structurally similar to the true alignment. That is, the initial alignment corresponded to the ground truth

but with some proportion of mismatched nodes, keeping the rest of them correctly aligned. With this, we aim to study the behaviour of the algorithm when initialised with a structurally similar alignment, rather than with an energetically favourable one. This was tested with several proportions, but a representative example with 60 misaligned nodes has been plotted below.

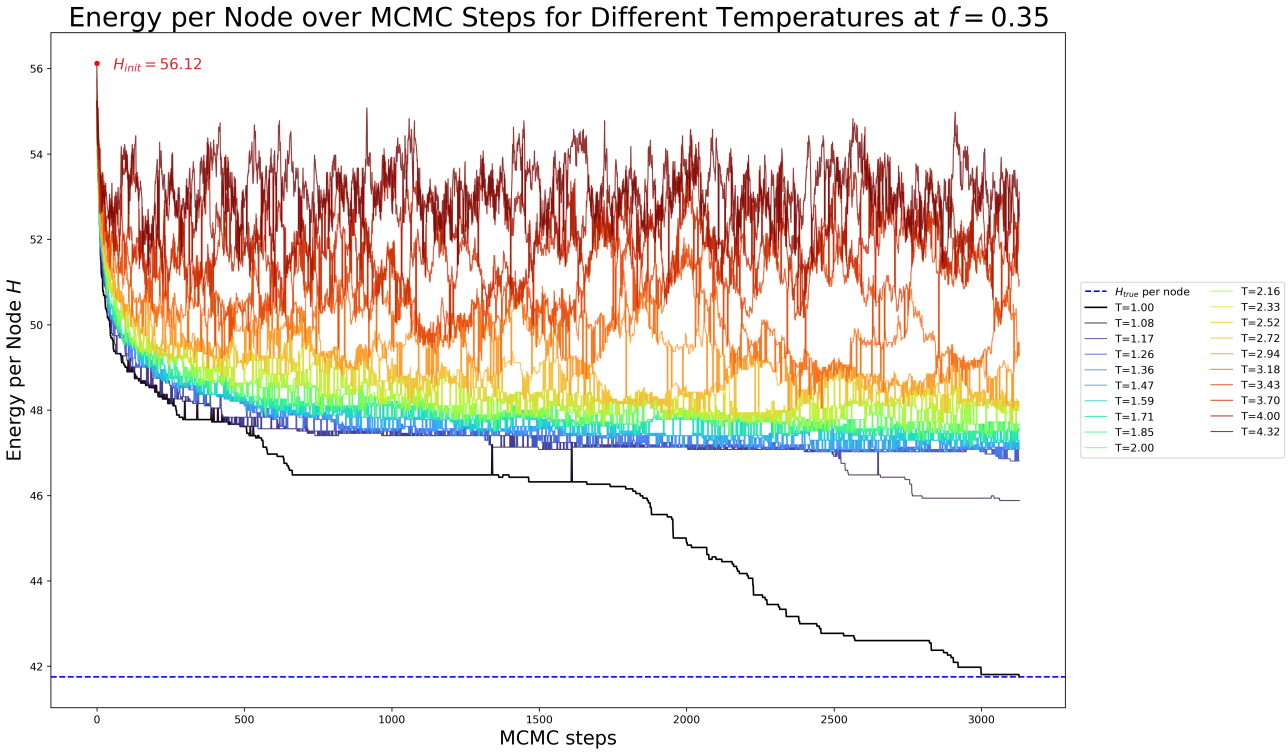

Figure 14. Energy sampling of a run with $f = 0.4$ and 60 mismatched nodes.

Looking at Fig. 14, we see that despite the high energy associated with the initial configuration ($H_{min} \approx 56.12$), the algorithm is able to recover the ground truth alignment in only around 3000 MCMC steps for a noise level of $f = 0.35$, where successful recovery is extremely rare, as shown in previous sections.

This highlights a crucial decoupling between structural similarity to the ground truth and low-energy configurations, especially for high noise levels. One might expect that alignments close to the ground truth (with many well-aligned nodes) would correspond to low-energy configurations. However, as the noise level increases, this intuition fails: structurally similar alignments (like H_{init} in Fig. 14) often exhibit higher energy than structurally unrelated alignments, which thus become energetically favoured.

Consequently, the MCMC algorithm, which samples from the alignment space by favouring low-energy configurations, becomes trapped in regions that, despite their low energy, are structurally distant from the ground truth. As a result, it fails to explore structurally promising alignments that could lead to the true alignment (as Fig. 14 shows), simply because their high energy makes them unlikely to be sampled.

In practice, for high noise levels, this decoupling causes the energy landscape to be increasingly uninformative about structural proximity to the ground truth (i.e., close alignments will have an increasingly higher energy than the true alignment). As noise increases, structurally similar configurations are assigned high energies, making them statistically unlikely for the MCMC algorithm. As the latter is guided by energy, it ignores these crucial regions, leading to a form of ”blindness” where

recovery is theoretically possible but algorithmically hard due to the misalignment between energy favourability and structural similarity.

In summary, this phenomenon explains the ”confusion” or difficulty faced by the MCMC algorithm near the transition point. Although the solution space still contains configurations that are close to the ground truth, the increasing noise causes the energy landscape to misrepresent their utility for recovering the true alignment.

4 Conclusion

In this work, we have studied the network alignment problem from a probabilistic perspective, focusing on the detectability of the ground-truth alignment under varying levels of noise representing network similarity. Our approach is grounded in a Bayesian formulation, which defines the posterior distribution over possible alignments based on a given set of observations sharing an underlying blueprint. This probabilistic perspective allows us to formally distinguish between regimes in which the ground truth is theoretically recoverable and those in which the true alignment is statistically unattainable.

To explore the structure of the energy landscape given by the posterior distribution, we have implemented a sampling-based inference method using Markov Chain Monte Carlo (MCMC), enhanced with parallel tempering to improve performance. This allowed us to examine the behaviour of the system across different noise levels (quantified by the parameter f) and temperatures. Our results reveal the existence of two distinct transitions: an *algorithmic transition*, corresponding to a threshold where alignment remains theoretically possible but becomes computationally challenging due to insufficient computation time, and a *detectability transition*, marking the point beyond which the true alignment becomes theoretically and effectively impossible to recover.

After a theoretical study on the latter transition, we were able to identify and understand the clear gap between both transition points. While the true alignment remains energetically favourable near the detectability transition, the algorithm struggles to recover it, not due to ergodicity breaking, but because of practical limitations: a finite number of MCMC steps may be insufficient to escape local minima in a complex energy landscape. This interpretation is supported by a detailed analysis of the phase transition, where a power-law fit suggested a first-order divergence in the expected number of MCMC steps until success at the critical transition point f_{crit}.

Importantly, our results also reveal a partial decoupling between energy and structure similarity: low-energy alignments aren’t usually close to the ground truth structurally, which becomes increasingly noticeable as the noise level increases. This explains the fundamental reason why recovering the ground truth becomes hard for the algorithm, as the energy landscape does not guide it reliably toward increasingly more similar alignments, but toward lower energies, which we characterised to be structurally very different from the ground truth. Overall, these findings clarify the behaviour of the network alignment problem and reinforce the need for more efficient inference methods that can better explore the posterior distribution near criticality.

Ethical and Social Responsibility Considerations (CT7 Competency)

This work focuses on the theoretical and computational study of a fundamental problem in complex networks and statistical physics, without involving personal, social, or environmental data. As such, the indicators associated with the CT7 competency, gender equality, environmental sustainability, social responsibility, and professional ethics, do not have a direct application in the development of this project. Nevertheless, the work has been carried out with a strong commitment to scientific integrity, methodological rigor, and general ethical principles.

5 Bibliography

[1] R. Tang, Z. Yong, S. Jiang, *et al.*, "Network alignment," *Physics Reports*, 2025. DOI: 10.48550/arXiv.2504.11367.

[2] M. Venkatesh and colleagues, "Network alignment and similarity reveal atlas-based topological differences in structural connectomes," *Network Neuroscience*, vol. 5, no. 3, pp. 711–733, Sep. 2021, ISSN: 2472-1751. DOI: 10.1162/netn_a_00199.

[3] C.-Y. Ma and C.-S. Liao, "A review of protein–protein interaction network alignment: From pathway comparison to global alignment," *Computational and Structural Biotechnology Journal*, vol. 18, no. 4, pp. 2647–2656, 2020, ISSN: 2001-0370. DOI: https://doi.org/10.1016/j.csbj.2020.09.011.

[4] X. Xie, K. B.-Y. Wong, H. Aghajan, P. Veelaert, and W. Philips, "Road network inference through multiple track alignment," in *Proceedings of the 20th SIGSPATIAL International Conference on Advances in Geographic Information Systems*, vol. 72, 2016, pp. 93–108. DOI: 10.1016/j.trc.2016.09.010.

[5] T. Lázaro, R. Guimerà, and M. Sales-Pardo, "Probabilistic alignment of multiple networks," *Nature Communications*, vol. 16, no. 1, p. 12, Apr. 2025, ISSN: 2041-1723. DOI: 10.1038/s41467-025-59077-7.

[6] H. Du, S. Gong, and R. Huang, "The algorithmic phase transition of random graph alignment problem," *arXiv preprint arXiv:2206.14464*, 2023. DOI: 10.48550/arXiv.2307.06590.

[7] O. Fajardo-Fontiveros, I. Reichardt, H. R. De Los Ríos, J. Duch, M. Sales-Pardo, and R. Guimerà, "Fundamental limits to learning closed-form mathematical models from data," *Nature Communications*, vol. 14, no. 1, p. 3692, 2022. DOI: 10.48550/arXiv.2204.02704.

[8] O. Fajardo-Fontiveros, R. Guimerà, and M. Sales-Pardo, "Node metadata can produce predictability crossovers in network inference problems," *Physical Review X*, vol. 13, no. 2, p. 021 030, 2023. DOI: 10.48550/arXiv.2103.14424.

[9] M. El-Kebir, J. Heringa, and G. W. Klau, "Sparse global network alignment as a special case of quadratic assignment," *Bioinformatics*, vol. 8, no. 4, pp. 1035–1051, 2015. DOI: 10.3390/a8041035.

[10] N. Mamano and W. Hayes, "Sana: Simulated annealing network alignment applied to biological networks," in *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, vol. 15, IEEE, 2016, pp. 1–17. DOI: 10.48550/arXiv.1607.02642.

[11] Y. Iba, N. Saito, and A. Kitajima, "Multicanonical mcmc for sampling rare events: An illustrative review," *Annals of the Institute of Statistical Mathematics*, vol. 190, no. 3, pp. 611–645, Apr. 2014. DOI: 10.1007/s10463-014-0460-2.

[12] W. D. Vousden, W. M. Farr, and I. Mandel, "Dynamic temperature selection for parallel tempering in markov chain monte carlo simulations," *Monthly Notices of the Royal Astronomical Society*, vol. 455, no. 2, pp. 1919–1937, Nov. 2015, ISSN: 1365-2966. DOI: 10.1093/mnras/stv2422.

Appendices

A Structural Similarities of Local Minima to the Ground Truth

As the noise level f increases, the MCMC algorithm tends to become trapped in local minima of the energy, failing to recover the ground truth. In this section, we study the characteristics of these local minima, investigating their structural relationship with the ground truth. Even if these alignments are energetically favourable, do they correspond to accurate node pairings, or are they simply misleading configurations?

To address this, we analyse the percentage of correctly aligned nodes throughout the MCMC run for the replica at $T = 1$ for the previous examples. The resulting plots will provide an informative measure of how close the algorithm's solution is to the true alignment.

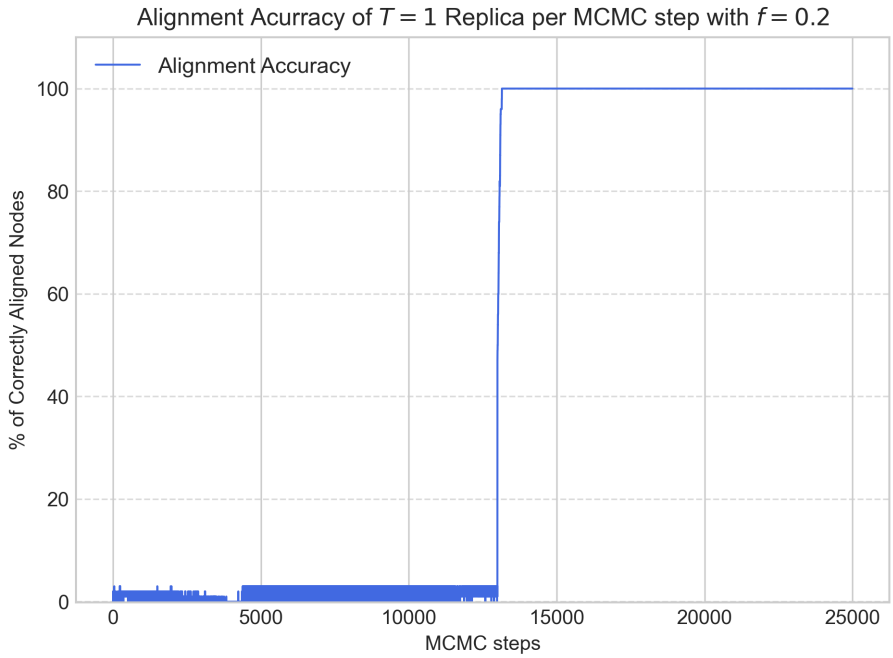

Figure 15. Alignment Accuracy of the lowest-temperature replica for low noise level.

Fig. 15 shows the alignment dynamics for $f = 0.2$. Initially, the algorithm fluctuates around configurations with a relatively low percentage of correctly aligned nodes. These configurations correspond to local minima that, while energetically favourable, are seen to have little resemblance to the ground truth. However, once the algorithm collapses to the global minimum, the percentage of correctly aligned nodes jumps abruptly to 100% and remains there for the rest of the run. This again reaffirms the observation made in section 3.2.4, where the majority of sampled alignments either correspond to local minima (unrelated with the ground truth) or the true alignment.

We now repeat this analysis for higher noise levels to assess the similarity of the local minima sampled there with the ground truth. Moreover, since the percentage of well-aligned nodes varies a lot for these phases, a moving average is also plotted as well to show the overall trend.

As it can be seen in Fig. 16, the percentage of correctly aligned nodes fluctuates around 1% throughout

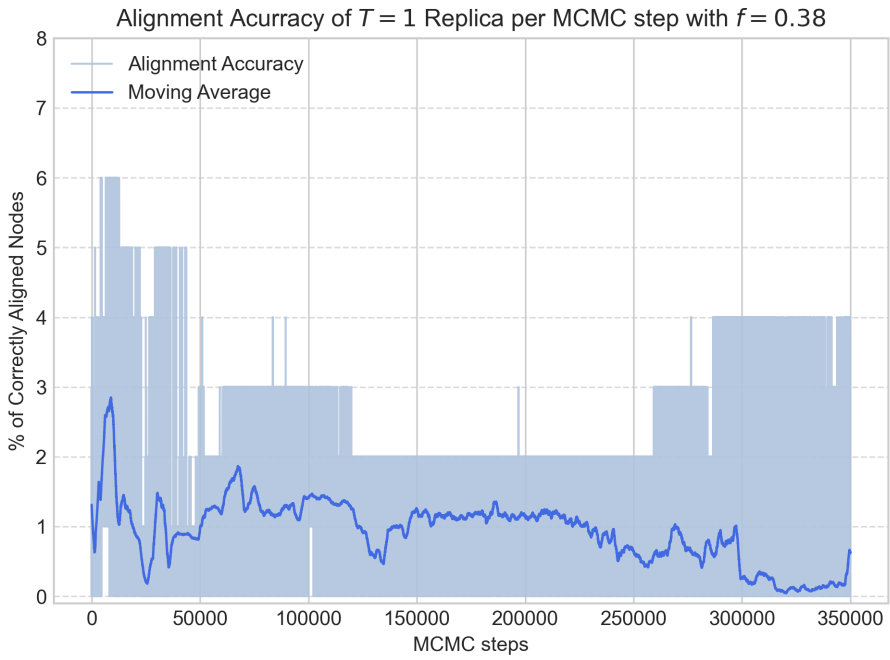

Figure 16. Alignment Accuracy of the lowest-temperature replica for a noise level near the transition.

the entire MCMC run, decreasing even more as the number of steps increase. For this reason, although these alignments minimise the energy locally, it becomes evident that they correspond to configurations that don't resemble the true alignment at all.

This result is particularly significant because it confirms that low-energy states are not necessarily informative. To contextualise these values, consider that in a random alignment of N nodes, the expected number of correctly aligned nodes is exactly 1,[5] which corresponds to a 1% of well-aligned nodes. This matches our observations for the previous, suggesting that they are essentially as uninformative as random alignments and lie very far from the ground truth..

In Fig. 17, the situation repeats itself. The percentage of correctly aligned nodes remains consistently low, often oscillating around $1 - 2\%$, which are values similar to those observed below the transition phase. As discussed in previous sections, in this phase, the ground truth becomes energetically indistinguishable from other low-energy configurations. The low percentage of correctly aligned nodes further confirms that the algorithm is no longer capturing any meaningful structure. The alignments found are essentially random, with accuracy figures barely above the theoretical random baseline of 1%.

All of this reinforces the previous point that, although local minima are energetically favourable, they are very different from the ground truth, which is supported by their low percentage of well-aligned nodes.

[5] Set $\mathcal{X}_i$ to be 1 if node i is correctly aligned and 0 otherwise. The expected number of correctly aligned nodes will be $\mathbb{E}\left[\sum_{i=1}^{N} \mathcal{X}_i\right] = \sum_{i=1}^{N} \mathbb{E}[\mathcal{X}_i] = \sum_{i=1}^{N} 1/N = 1$.

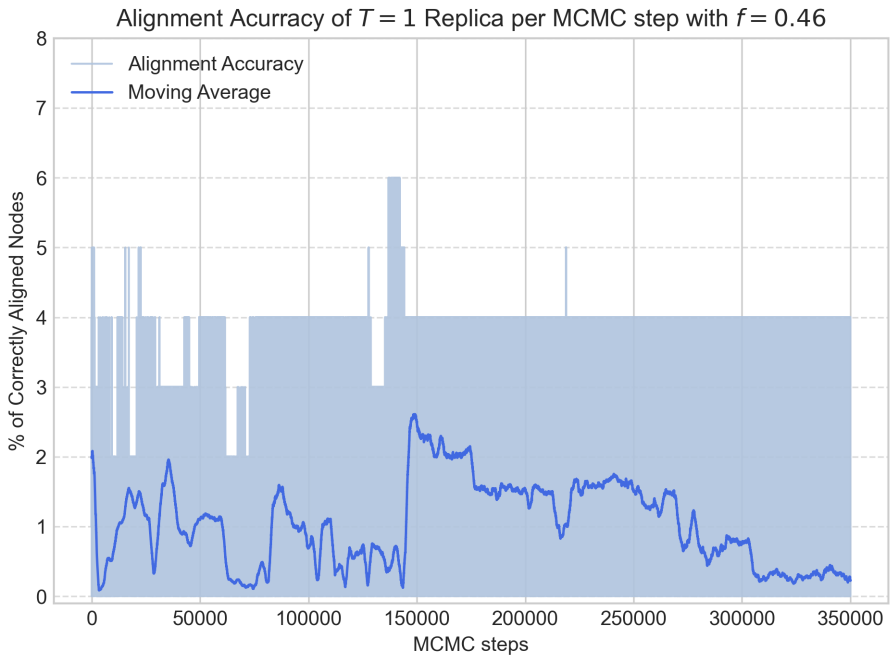

Figure 17. Alignment Accuracy of the lowest-temperature replica for a noise level beyond the transition.

B Code Availability

The implementation of the methods used in this thesis is available in a public GitHub repository. The code is structured to separate reusable functions from execution scripts and relies on precomputed data stored in CSV files to produce the figures and results included in the report. T he full repository is available at: `https://github.com/PabloSarro/Detectability-Transition-NAP`.