

Horizon 2020

 WORK PROGRAMME 2018 – 2020
 H2020-SC1-2019-Two-Stage-RTD

Project Acronym: **TO_AITION**

Grant Agreement N°: 848146

Project Full Title: **A high-dimensional approach for unwinding immune-metabolic causes of cardiovascular disease-depression multimorbidities**

Starting Date: 01/01/2020

Duration in months: 60

D3.2 Identification of driver nodes in disease network

Nature:	Report
Dissemination Level:	Public
Contractual Date of Delivery to the EC:	30 June 2023
Actual Date of Delivery to the EC:	28 July 2023
WP number and Title:	WP3 - Identification of common mechanisms, pathways and molecules driving co/multimorbidity through causal inference modelling
Lead Beneficiary:	UVA
Version 1.0:	Authored by UvA 20/07/2023 – Initial TOC
Version 2.0:	Reviewed by the partner UVA on 25/07/2023
Version 3.0:	Reviewed by the Consortium on 27 /07/2023– Final version
Final Version:	Submitted on 28/07/23

LIST OF BENEFICIARIES

Ben. No.	Beneficiary name	Short name	Country
1	IDRYMA IATROVIOLOGIKON EREUNON AKADEMIAS ATHINON	BRFAA	GREECE
2	UNIVERSITEIT VAN AMSTERDAM	UVA	NETHERLANDS
3	PIRKANMAA HOSPITAL DISTRICT	TAUH	FINLAND
4	UNIVERSITATSKLINIKUM BONN	UKB	GERMANY
5	PANEPISTIMIO IOANNINON	UOI	GREECE
6	MICRONIT HOLDING BV	MICRONIT	NETHERLANDS
7	GENOWAY	GENOWAY	FRANCE
8	RUPRECHT-KARLS-UNIVERSITAET HEIDELBERG	UHEI	GERMANY
9	STICHTING VUMC	VUMC	NETHERLANDS
10	UNIVERSYTET MEDYCZNY W LODZI	LODZ	POLAND

11	UNIVERSITAIR MEDISCH CENTRUM UTRECHT	UMCU	NETHERLANDS
12	UNIVERSITE DE GENEVE	UNIGE	SWITZERLAND
13	EXELIXIS RESEARCH MANAGEMENT & COMMUNICATION	EXELIXIS	GREECE
14	SOCIETE EUROPEENNE DE CARDIOLOGIE	ESC	FRANCE

Executive summary

One reason for inferring networks from data is to predict driver nodes: nodes that are opportunely placed in the network such that if their state changes, a large part of the network sees the effect. Typically, such driver nodes are ‘identified’ by using a network centrality algorithm applied to a network of correlations or of causal relationships. However, it remains uncertain whether this method offers a dependable approach to pinpointing driver nodes. Identifying driver nodes is not inherently a topological (centrality) issue, but rather, an interventional question.

Addressing this question proves challenging in practice, as we cannot conduct actual interventions based on a specific dataset within our consortium. Therefore, this report presents a series of simulated interventions on an extensive range of hypothetical symptom-disease models. From these computational models, we produce synthetic datasets and subsequently compare various driver node inference algorithms with the ground truth. The ground truth is established by emulating an intervention on each node and ranking the nodes on the scale of how the rest of the network was impacted by this intervention.

Our first set of results is negative. That is, using cross-sectional synthetic data generated from custom symptom-disease models, none of the conventional network centrality metrics for pairwise networks prove reliable in identifying driver nodes. This remains the case even when generalizing beyond pairwise networks by incorporating synergistic associations and considering less common centrality metrics for hypergraphs.

Our second set of results is more encouraging. If the datasets were to consist of a rich collection of time-series data for all involved variables, it would be feasible to reliably infer driver nodes. Nonetheless, even in this ideal scenario, we find that classic network centrality metrics are still not reliable predictors. As a result, we developed our own centrality metric, IMI, based on measuring short-term fluctuations. This metric proves reliable, but its application hinges on the dataset comprising densely sampled time series. Thus, the inference of driver nodes in the datasets currently available within the consortium, or more broadly within epidemiological cohorts, cannot be done reliably.

From these results, we delineate two main paths for future research. The first unresolved question is where the threshold lies between these two extremes (cross-sectional versus rich time-series) for it to become viable to identify driver nodes. It is clear that an extended number of follow-up waves is needed, but how many are adequate exactly? The second unresolved question is whether we can design a superior algorithm for identifying driver nodes based on hypergraphs (which incorporate synergistic associations). Given that the current set of centrality metrics for hypergraphs is not as diverse as for pairwise networks, the potential for developing new metrics may exist.

Table of Contents

1. General introduction	8
2. Exploring the effectiveness of Centrality Metrics in Identifying/ Predicting Influential Nodes in pairwise networks	9
2.1. Introduction:	9
2.2. Methods	9
2.2.1. Assessment	12
2.3. Results	13
2.4. Discussion	17
3. Identifying driver nodes based on higher-order associations (hypergraphs)	18
3.1. Preamble	18
3.2. Introduction	19
3.3. Methods	20
3.4. Conclusion	26
4. Dynamic importance of network nodes is poorly predicted by static structural features.	27
4.1. Preamble	27
4.2. Introduction	27
4.3. Theoretical background	33
4.3.1. Terminology	33
4.3.2. Causal interventions and dynamic importance	34
4.3.3. Intervention size	36
4.3.4. Intervention in kinetic Ising model	36
4.3.5. Measuring information flow	37
4.3.6. Mutual information and causality	38
4.3.7. Integrated mutual information	39
4.4. Methods and network data	40

4.4.1.	Network data	40
4.4.2.	Numerical methods	41
4.4.3.	Estimating $pSt0$ and $pSt0 + t St0$	42
4.4.4.	Time symmetry and mutual information	43
4.4.5.	Area under the curve estimation	43
4.4.6.	Sampling bias correction	43
4.4.7.	Driver node prediction and precision quantification	43
4.4.8.	Software	45
4.5.	Results	47
4.5.1.	Random network results	47
4.5.2.	Real-world network: psycho-symptoms results	48
4.5.3.	Limitations	52
4.6.	Conclusions	53
5.	General discussion, conclusions and next steps	53

List of Figures

Figure 1. Respectively Pearson correlation, Spearman correlation, and Kendall correlation coefficients matrices of Individual Centrality Metrics and Intervention impact node.	13
Figure 2. Coefficients of Centrality measures	15
Figure 3. Overview of centrality measures and coefficients	16
Figure 4. Feature importance: The values indicate the relative importance of each feature in the model.	17
Figure 5. Reconstructing a small synthetic network consisting of synergistic and pairwise interactions is impossible when only using pairwise edges.	19
Figure 6. A randomly generated hypergraph.	22
Figure 7. Histogram of node impact and centrality measures after log transformation	23
Figure 8. Node impact vs. centrality measures after log transformation	24
Figure 9. Box plot of correlation coefficients between node impact and centrality measures with log transformation	25
Figure 10. Centrality importance in predicting node impact from RFR (some after log transformation).	26
Figure 11.	29
Figure 12.	41
Figure 13.	45
Figure 14.	48
Figure 15.	49

List of Tables

Table 1	10
Table 2	14
Table 3	15
Table 4	16
Table 5. <i>Pearson's correlation between log(node impact) and centrality measures after log transformation.</i>	22
Table 6. <i>Metrics for prediction evaluation</i>	26

1. General introduction

This report fulfils Deliverable 3.2, titled “Identification of driver nodes in disease networks (M42)”, which is focused on the corresponding Task 3.2, titled “Inference of driver nodes of disease networks (M13-M42).”

A fundamental proposition in the To_Aition project is the assumption that the graphic representation of the comorbidity of Cardiovascular Disease (CVD) and Major Depression (MD) in the form of (hyper-)graphs not only illustrates the correlations between biomarkers and phenotypes, but also facilitates the prediction of driver nodes. These driver nodes, also known as leverage points, are nodes within a network that serve as strategic focal points for intervention, where the intended intervention would have a maximum impact on the network. This understanding is crucial when exploring preventative strategies. For example, a potential intervention method might involve 'blocking' pathways from one morbidity to another, thereby mitigating the overall burden.

The data sets used in this project are primarily cross-sectional and include biomarkers (like lipidomics and metabolomics), phenotypic variables (such as CVD outcomes or mental health scores), and potential confounders. In several tasks within WP2 and WP3, we devised novel methodologies to construct various types of networks among these variables. The links between variables represent either (partial) correlations or deduced causal relationships. These connections may be pairwise (dyadic: between two variables) or synergistic (polyadic: involving multiple variables and a single variable). The core question this report aims to answer is: "Can driver nodes be reliably inferred from these networks?"

In this deliverable, we present our findings from three different methodologies developed to address this question. Since the real data sets do not reveal a definitive truth regarding which variables are driver nodes, our initial step involves generating synthetic data sets using causal Bayesian Networks (BN; our 'toy models') with various characteristics. Each synthetic data set provides a ground truth, enabling us to determine whether a specific method for inferring driver nodes is effective. The primary idea is to gradually increase the complexity of these synthetic data sets to mimic a real data set in terms of size, value distribution, missingness, and correlations. Subsequently, the most effective driver node inference method is selected to be applied to the actual data set.

This progressive search begins with the simplest synthetic data: no selection bias, additive noise, a limited number of variables, few possible values per variable, and no latent confounders. Ideally, this should present the 'easiest' scenario for inferring driver nodes; if we cannot reliably infer driver nodes from such simple scenarios, it is unlikely that we can do so in more complex, realistic scenarios.

Regrettably, our studies, which we detail in the following sections, indicate that the network structures do not reliably predict driver nodes in our scenarios. More precisely, in the numerous different scenarios we explored,



we found virtually no correlation between a comprehensive list of node centrality metrics (which measure, in various ways, a node's central position in a given network) and the total impact of simulated interventions on a node. Our (negative) finding is robust against a range of methodological choices and network types.

The work presented in this deliverable has in the meanwhile been published in peer reviewed journals:

Elteren *et al.*, 2022; <https://doi.org/10.1016/j.physa.2022.126889>

Hourican *et al.*, 2023; doi.org/10.3389/fsysb.2023.1155599

Three more publications are currently in preparation or submitted.

2. Exploring the effectiveness of Centrality Metrics in Identifying/ Predicting Influential Nodes in pairwise networks

Centrality metrics are commonly used for identifying driver nodes in networks, but their effectiveness in capturing node importance in correlation networks is uncertain. This research aims to investigate the alignment between centrality metrics and intervention-induced significant nodes. We employ a systematic methodology involving a Bayesian model for simulation, hard interventions, computation of various centrality measures, correlation analysis, linear regression and nonlinear approaches to capture any relationship between centrality metrics and intervention output. We also explore the combination of centrality metrics and examine potential non-linear relationships using Machine Learning regression techniques.

2.1. Introduction:

Research Background. In a network, certain nodes play a critical role in shaping its dynamics and behavior. These influential nodes, often referred to as driver nodes, have a profound impact on the network's overall structure and functionality. Identifying these driver nodes is essential for understanding and manipulating network behavior effectively.

Research Objective. The primary objective of this research is to assess the alignment between centrality metrics and intervention-induced significant nodes. By comparing the nodes identified as significant through interventions with those identified based on centrality measures, we aim to investigate the effectiveness of centrality metrics in capturing the importance of nodes within the network.

2.2. Methods

Bayesian Model Definition. To establish a causal framework for the network, we define a Bayesian model that represents the relationships between variables. This model serves as the basis for subsequent analyses. By specifying the causal relationships, we can capture the influences and dependencies among the nodes in the

network. The Bayesian model allows us to estimate the parameters that govern these relationships and provides a solid foundation for further analysis.

Table 1

Bayesian model	number of nodes: 40	number of edges: 70
----------------	---------------------	---------------------

Simulation. After defining the Bayesian model, we proceed to simulate the behavior of the network based on this model. Through simulation, we generate data that closely resembles the dynamics of the network (number of samples: 100,000). This step involves generating multiple samples and assigning values to the variables according to their dependencies as defined in the Bayesian model. By simulating the network, we can study its behavior under various conditions and interventions, which enables us to gain a deeper understanding of its characteristics.

The dynamics of the process can be described as follows: For every sample and node, a series of steps is performed. Initially, the algorithm examines the predecessors of the current node within the graph, wherein predecessors refer to nodes that possess edges directed towards the current node. If any predecessors exist, their corresponding values are evaluated. Specifically, the values of the predecessors for the current sample are retrieved. Subsequently, the algorithm checks whether any of the predecessor values are non-zero. If at least one predecessor exhibits a non-zero value, this implies the existence of a connection leading to the current node. Based on the presence or absence of these connections, the algorithm assigns a value to the current node. If there are predecessors with non-zero values, the current node's value is set to 1. However, if there are no such predecessors, the value is randomly assigned as either 0 or 1. This process is iteratively repeated for all nodes within the graph and for each sample with the corresponding values for each node in every sample.

Intervention. To investigate the impact of interventions on the network, we perform interventions on each node individually. These interventions involve setting the value of a specific node to a desired value, representing a controlled change in the variable's value. We calculate the probability distribution of the target node based on the observed data when no specific evidence or intervention has been introduced and probability distribution/ intervention distribution of the target node following a specific intervention. The intervention involves setting the value of the target node to either 0 or 1. The intervention distribution represents the probability distribution of the target node under the specified intervention.

Next, the Hellinger distance between the original distribution of each node before intervention and the

intervention distribution after the intervention. The Hellinger distance provides a measure of dissimilarity between probability distributions and allows us to quantify the impact of interventions on the network.

Initial probability distribution and intervention probability distribution are conditional probability distributions (CPDs) specific to the target node/intervention node. These distributions are conditioned on the evidence or lack thereof, as well as the intervention that has been applied.

Centrality Measure Calculation. After generating the data, we calculate the phi coefficient values and an association network is constructed using these coefficients.

The phi coefficient (ϕ) is calculated using the following formula: $\phi = \sqrt{\chi^2 / n}$

where χ^2 represents the chi-squared statistic obtained from a contingency table of the variables, and n represents the total number of observations in the contingency table. The phi coefficient measures the strength of association between two categorical variables, with a value ranging from 0 (no association) to 1 (complete association).

To assess the importance and influence of individual nodes within the network, we calculate various centrality measures. These measures provide quantitative metrics that capture different aspects of node centrality and help us identify the most significant nodes in the network.

Betweenness Centrality: This measure identifies nodes that act as crucial bridges or intermediaries in the network. Nodes with higher betweenness centrality have a significant impact on the flow of information or influence between other nodes.

Closeness Centrality: This measure quantifies how close a node is to other nodes in the network. Nodes with higher closeness centrality have shorter average distances to other nodes and can efficiently spread information or influence within the network.

Eigenvector Centrality: This measure takes into account both the centrality of a node and the centrality of its neighboring nodes. Nodes with higher eigenvector centrality are well-connected to other highly influential nodes, making them more influential themselves.

PageRank is a centrality measure that originated from Google's web page ranking algorithm. It assesses the importance or significance of nodes within a network based on the number and quality of incoming links to those nodes. The underlying principle is that nodes with a higher number of high-quality incoming links are considered more important or influential within the network.

Enhancing Analysis Robustness. To ensure reliable results and enhance the robustness of our analysis, the entire process was conducted 500 times. This iterative approach allowed us to collect centrality values and node impact across multiple runs/networks, thereby providing a more comprehensive understanding of the



network dynamics.

2.2.1. Assessment

Correlation Analysis and Linear Regression. In the correlation analysis, we calculate the correlation matrix between each individual centrality metric and the "driverness" measure, which quantifies the size of the node's impact. By examining the correlation coefficients, we assess the strength of the relationship between centrality metrics and impact measures. A significant correlation would indicate that a single centrality metric alone can effectively capture the importance and influence of nodes within the network.

In addition to the correlation analysis, we also perform linear regression analysis. This allows us to investigate the extent to which centrality metrics can predict the impact measures. By fitting a linear regression model, we can estimate the coefficients of the centrality metrics and assess their statistical significance in predicting the "driverness" of nodes. This analysis provides further insights into the predictive power of individual centrality metrics.

Combination of Centrality Metrics and Linear Regression. In this step, we investigate the potential of combining multiple centrality metrics to enhance the prediction of driver nodes within the network. We employ linear regression analysis with all the centrality metrics as predictors to assess their collective predictive power in determining the "driverness" of nodes. By including all the centrality measures as predictors in the linear regression model, we aim to leverage the complementary information captured by each metric. The coefficients of the centrality metrics in the regression model provide insights into their individual contributions to predicting the impact measures. Furthermore, we evaluate the statistical significance of the coefficients to determine which centrality metrics are most influential in predicting driver nodes.

This analysis allows us to assess whether a combination of centrality metrics improves the prediction accuracy compared to using individual metrics alone. If the combined centrality metrics exhibit a stronger relationship with the impact measures than any individual metric, it suggests that a collective approach provides a more comprehensive understanding of node significance within the network.


Combination of Centrality Metrics and Nonlinear Regression. In order to explore the potential for a non-linear relationship between centrality metrics and the impact measures, we employ Random Forest regression. These methods allow us to capture complex interactions and non-linear patterns that may exist between the combined centrality metrics and the "driverness" of nodes. By utilizing Random Forest regression, we aim to uncover any hidden non-linear relationships that may exist beyond the linear regression analysis. These techniques have the capability to capture intricate interactions and non-linearities, which could provide deeper insights into the predictive power of combined centrality metrics.

By employing Random Forest regression with the combined centrality metrics as predictors, we aim to gain a comprehensive understanding of the relationship between centrality and impact in a non-linear context.

2.3. Results

Correlation between Individual Centrality Metrics and Intervention. The correlation analysis was conducted to examine the relationship between individual centrality metrics and the intervention outcome, which quantifies the impact of nodes. Pearson correlation, Spearman correlation, and Kendall correlation coefficients were calculated to assess the strength and direction of these relationships. However, the results showed no significant correlation between the centrality measures and the intervention outcome based on these measures.

Figure 1. Respectively Pearson correlation, Spearman correlation, and Kendall correlation coefficients matrices of Individual Centrality Metrics and Intervention impact node.



Individual Centrality Measures as Predictor for Intervention Output. A comprehensive analysis was conducted to examine the relationship between the centrality measures and an intervention output variable. The dataset was divided into subsets, with each centrality measure serving as a predictor variable and the intervention output as the target variable. Linear regression models were trained on the training set and applied to the testing set for prediction. Performance evaluation metrics, including the normalized mean squared error (NMSE) and R-squared score, were calculated to assess the predictive capability of the models. The obtained results, including the optimal alpha, L1 ratio, NMSE, R2 score, intercept, and coefficient for each centrality measure, are presented in Table 2. This analysis provides valuable insights into the predictive power of different centrality measures concerning the intervention output, facilitating the assessment of their effectiveness in capturing the significance of nodes within the network.

Table 2


Centrality Measure	Optimal Alpha	L1 Ratio	Normalized Mean Squared Error (NMSE)	R2 Score	Intercept	Coefficient
Degree Centrality	2.160819e-05	0.5	0.978863	0.021137	0.128328	0.110977
Betweenness	5.005945e-09	0.5	0.999871	0.000129	0.155269	38.199317
Closeness	1.934499e-05	0.5	0.979299	0.020701	0.129152	0.118996
Eigenvector	1.321950e-05	0.5	0.977198	0.022802	0.121761	0.321905
PageRank	2.167912e-06	0.5	0.976139	0.023861	0.109028	1.891030

Complete Set of Centrality Measures as Predictors for Intervention Output: A Linear Approach. In this analysis, the entire set of available centrality measures was collectively employed as predictors. A linear regression model was trained on the training set and subsequently used to make predictions on the testing set. Performance evaluation metrics, such as the normalized mean squared error (NMSE) and R-squared score, were computed to assess the accuracy of the models. The results of this analysis, including the L1 ratio, NMSE, R2 score, and intercept, are presented in Table 3. Furthermore, the coefficients corresponding to each centrality measure are visualized in Figure, providing additional insights into their individual contributions to the prediction model.

Table 3

L1 Ratio	Normalized Mean Squared Error (NMSE)	Root Squared Error (RMSE)	Mean Error	Mean Absolute Error (MAE)	R2 Score	Intercept
0.5	0.97971	0.257885		0.223292	0.02029	0.117958

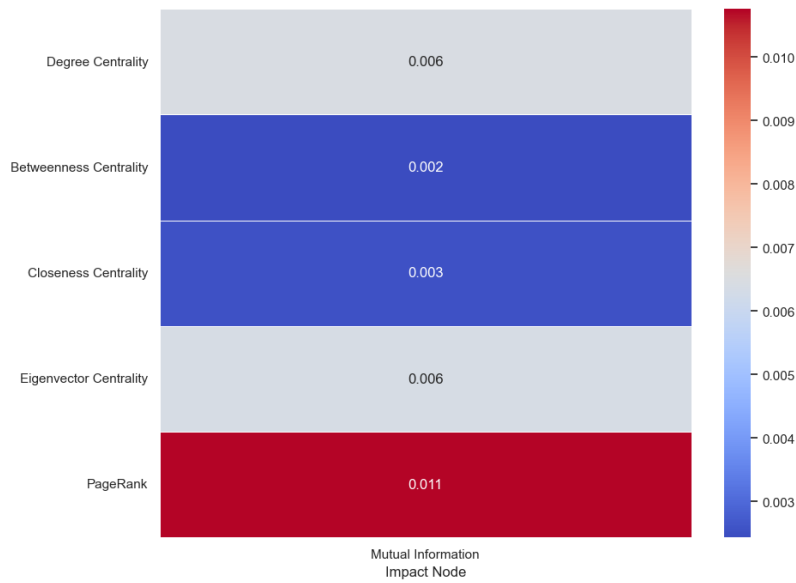
Figure 2. Coefficients of Centrality measures



Complete Set of Centrality Measures as Predictors for Intervention Output: A Non- Linear Approach

Using Mutual Information. To further explore the association between centrality metrics and the intervention outcome, mutual information was employed as an alternative measure of dependence. The mutual information captures the amount of information shared between two variables and can reveal non-linear associations. The mutual information matrix was visualized as a heatmap (data discretized using KBinsDiscretizer algorithm).

Figure 3. Overview of centrality measures and coefficients

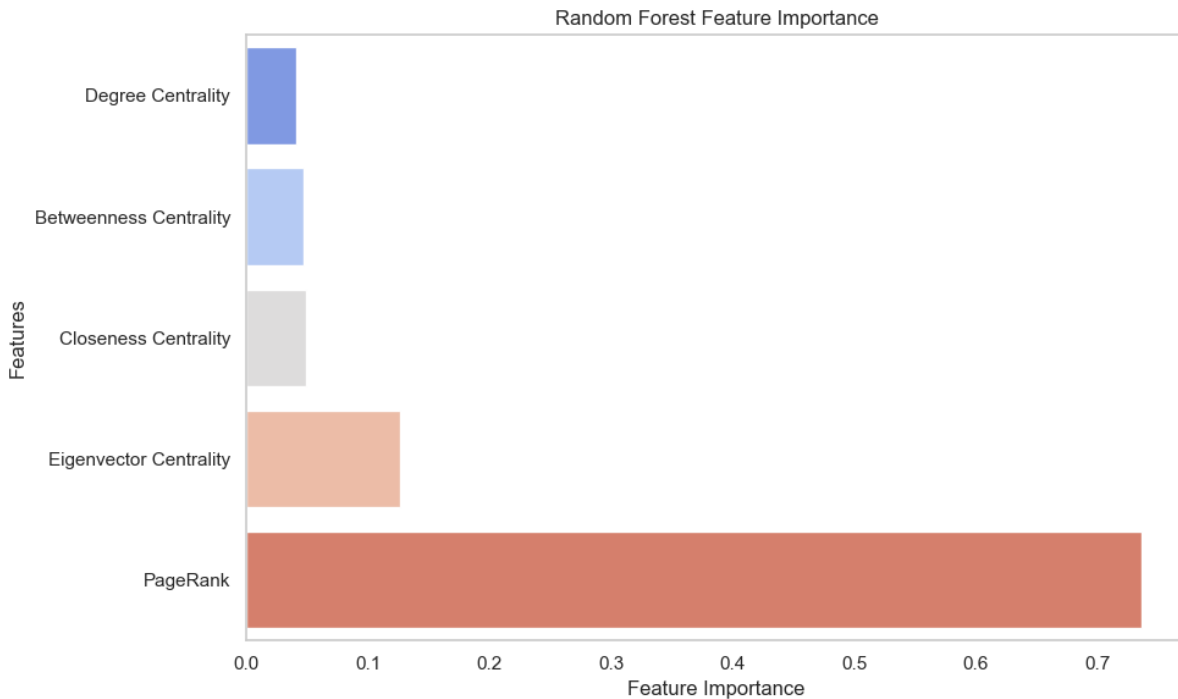


Random Forest Regression (RF). In this case, instead of using mutual information, we employed RF regression analysis to assess the association between centrality metrics and the intervention outcome. The evaluation metrics were calculated to evaluate the performance of the model. The feature importance values were then used to explore the relationship between the centrality metrics and the intervention outcome. In contrast to using mutual information, RF regression analysis provides a different perspective on the association between centrality metrics and the intervention outcome. It considers the predictive power of each metric based on the chosen evaluation metrics such as normalized mean squared error, root mean squared error, mean absolute error, and R2 score.

Table 4

Normalized Mean Squared Error (NMSE)	Root Mean Squared Error (RMSE)	Mean Absolute Error (MAE)	R2 Score
1.104	0.27	0.21	-0.104

Figure 4. Feature importance: The values indicate the relative importance of each feature in the model.



2.4. Discussion

The aim of this study was to evaluate the effectiveness of centrality metrics in identifying the significance of nodes within correlation networks. Our systematic methodology employed Bayesian modeling, interventions, centrality measure computation, correlation analysis, linear regression, and non-linear approaches utilizing machine learning regression techniques.

Our analysis found no significant correlation between individual centrality measures and intervention outcomes. This suggests that common centrality metrics, such as betweenness centrality, proximity centrality, eigenvector centrality, and PageRank, might not accurately represent the impact of nodes in correlation networks. Moreover, linear regression analyses using individual centrality measures as predictors of intervention outcomes failed to yield strong predictive results. The coefficients of each centrality measure were statistically insignificant. However, when using a linear regression model with the full set of centrality metrics as predictors, we observed a marginal improvement in prediction accuracy. The combined centrality metrics, according to the R-squared score, accounted for a small portion of the variability in intervention outcomes. While the coefficients for individual centrality metrics were statistically significant, their magnitudes were modest. Thus, although aggregated data from multiple centrality metrics can offer some insight into node relevance, the overall predictive power remains limited.

We also examined potential non-linear relationships using mutual information and random forest regression techniques. Neither method, however, revealed significant correlations or significantly improved prediction

accuracy. These results suggest that complex relationships and non-linear patterns may not hold significant influence. Our findings underscore the limitations of traditional centrality measures in identifying and predicting influential nodes in correlation networks. Despite the widespread use of centrality measures in various network studies, their efficacy in correlation networks remains unclear. Future research could explore alternative network metrics or develop new methodologies to quantify node relevance in correlation networks.

It's crucial to note that our conclusions are based on the methodologies and data used in this study. The effectiveness of centrality metrics in identifying key nodes can vary, depending on network characteristics and research context. Consequently, more research, coupled with validation using diverse datasets and network structures, is needed to deepen our understanding of the utility of centrality metrics in correlation networks.


In conclusion, this study illuminates the constraints of centrality metrics in signifying node importance in correlation networks. The findings suggest that centrality measures may not independently identify influential nodes, and their predictive capability in correlation networks is limited. Further research is required to explore additional measures and methodologies, enhancing our understanding of node relevance in correlation networks.

3. Identifying driver nodes based on higher-order associations (hypergraphs)

3.1. Preamble

The previous Section we established that it not possible to reliably infer driver nodes based on networks of pairwise relationships (from one node to one other node). To see if generalizing the type of relationships from purely pairwise to higher order, we will now proceed to investigate driver node identification is achievable when applied to higher order associations, using a bespoke toy model which contains a mix of synergistic interactions (triplets) and pairwise interactions (Fig. 5). We show that, a argued in a recent published perspective article [doi.org/10.3389/fsysb.2023.1155599], it is paramount to include synergistic interactions when reconstructing networks. That is, the result of using 'classic' methods to reconstruct the ground truth network misses quite some interactions and results in three disconnected components, whereas the ground truth is actually a single connected component. This is exemplified in Figure 5, whereby the left-hand network of Fig. 5 is based on a synthetic data set is shown on the right-hand side of Fig. 5. Thus, including synergistic interactions is predicted to lead to better reconstructions, which in the context of this deliverable, should thereby lead to better predictions of driver nodes. For this reason, in this Section we will explore exactly that.

Figure 5. Reconstructing a small synthetic network consisting of synergistic and pairwise interactions is impossible when only using pairwise edges.



9

3.2. Introduction

In disease networks, driver nodes refer to specific genes, proteins or other biological entities that are highly influential and play a crucial role in the development or progression of diseases. Driver nodes are typically identified based on their topological properties and impacts on other nodes and the whole system, corresponding to structural and dynamical methods, respectively. Centrality measures are commonly used as structural methods to quantify the importance or influence of nodes, while central nodes are unnecessarily driver nodes in terms of their complex functionality. In this research, we use dynamical methods that analyze the influence of nodes on the networks' dynamics by simulating network dynamics and observing the impact of node perturbations. Nodes that significantly alter the dynamics or have a strong impact on other nodes are defined as driver nodes. By identifying and understanding the driver nodes, we can gain insights into the key biomarkers and regulatory mechanisms driving the disease.

However, directly quantifying nodes' influence in huge disease networks is challenging or even impossible due to their complex dynamics. Here we first develop smaller and simpler toy models whose dynamics can be easily simulated and modeled and calculate each node's influence on other nodes. Then the nodes' influences are compared with their centralities to see which centrality measure can well predict node influence. As centrality measures are well-documented and easy to calculate, we can use that centrality metric to infer driver nodes in real disease networks. Additionally, when adding synergistic associations to such a network, a hypergraph is retrieved, another problem comes that only few centrality metrics for hypergraphs exist and on top of that are poorly understood. The goal of this subproject is to integrate synergistic associations into a hypergraph in a meaningful way, as well as assess existing and newly proposed centrality metrics based on

their ability to infer either predictive power or interventional impact. This should provide an important step to interpret driver node identification in a multivariate setting.

Then the target is to check if any centrality measures can well predict node influence. For hypergraphs, five centrality measures, s -betweenness, s -closeness, s -eccentricity and s -harmonic (implemented in *Hypernetx* package) and clique eigenvector (implemented in *xgi* package) centrality measures are discussed in the analysis.

- S -betweenness centrality: a centrality measure for an s -edge(node) subgraph of H based on shortest paths. Equals the betweenness centrality of vertices in the edge(node) s -linegraph.
- S -closeness centrality: the closeness centrality of vertices in the edge(node) s -linegraph.
- eccentricity centrality: the length of the longest shortest path from a node to every other node in the s -linegraph.
- S -harmonic centrality: A centrality measure for an s -edge subgraph of H . A value equal to 1 means the s -edge intersects every other s -edge in H .
- Clique eigenvector centrality: the clique motif eigenvector centrality of a hypergraph.
- H -eigenvector centrality: the H -eigenvector centrality of a uniform hypergraph.

A linegraph is a graph in which each vertex in the s -line graph represents a hyperedge with at least s vertices in the hypergraph, and two vertices are linked in the s -line graph if their corresponding hyperedges intersect in at least s vertices in the hypergraph.

3.3. Methods

To accomplish the above objectives, we begin by designing an algorithm capable of randomly generating directed acyclic hypergraphs (DAGs). In these hypergraphs, each target node is limited to being an output of hyperedges only once. We utilize different types of logic gates to specify the dynamics of the hypergraph. For pairwise edges, 'IS' or 'NOT' gates are randomly chosen, while 'AND,' 'OR,' or 'XOR' gates are randomly assigned to hyperedges. Once the hypergraph dynamics are defined, the states of nodes can be determined using Bernoulli trials. To assess node influence, we identify influence paths of the node first based on the hypergraph structure, then intervene the node by changing the probability of observing 1 when doing the Bernoulli trials. Based on the influence paths, we can identify the nodes that are influenced by the intervention and update their states accordingly. To quantify the impact of the intervention on the affected nodes, we calculate the Hellinger distance. This distance measure captures the difference between the pre-intervention and post-intervention probability distributions of the node states. After computing the Hellinger distances for each affected node, we can determine the total impact of the node on the entire system by summing up these

distances. This cumulative impact provides an assessment of how significantly the node affects the overall dynamics of the system.

Figure 6 presents an illustration of a randomly generated hypergraph. The dynamics of the hypergraph are defined by assigning logic gates to each edge and hyperedge through a random process. Once the dynamics are established, we proceed to evaluate the node impact using the aforementioned method. To determine the centrality of each node within the hypergraph, we utilize published centrality measures. This analysis is conducted on 20 randomly generated hypergraphs, and the impact and centralities from all nodes in these hypergraphs are combined. Since metrics such as node impact, $s_{\text{betweenness}}$, clique eigenvector, weighted betweenness, weighted degree, and weighted eigenvector centralities exhibit histograms that roughly follow an exponential distribution, we apply a log transformation to these variables. Figure 7 displays the histograms of these metrics, some of which have undergone log transformation. To assess the predictive power of centrality measures in determining node impact, we calculate the correlation coefficient between node impact and various centrality measures after log transformation (Table 5). Additionally, scatter plots depicting the relationship between node impact and centrality measures are presented (Figure 8). The correlation coefficient serves as an indicator of the relationship between node impact and centrality measures. A high positive correlation suggests that centrality measures are effective predictors of node influence, while a low or negative correlation indicates a weak or potentially inverse relationship.

Figure 6. A randomly generated hypergraph.

(20 edges and hyperedges, 37 nodes in total, max. order 5)

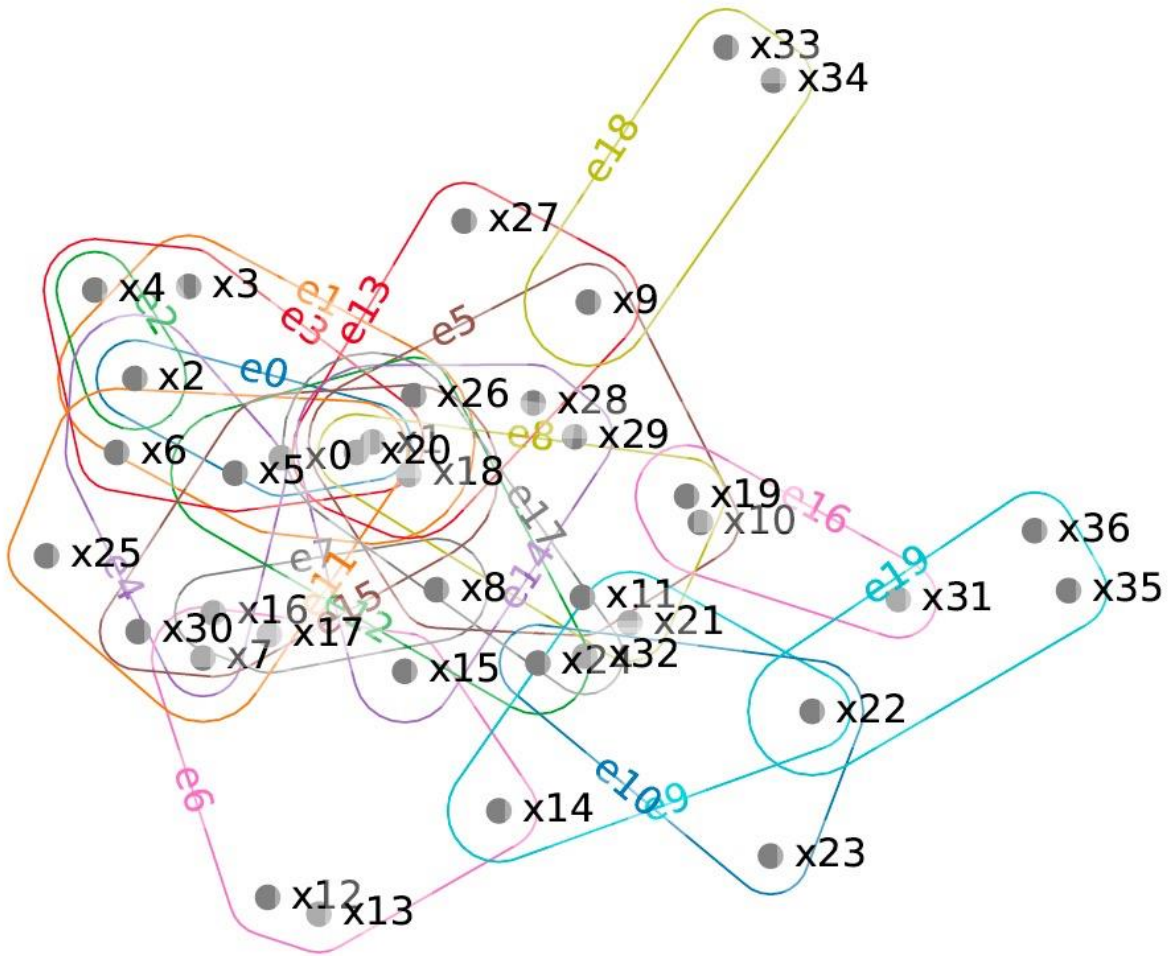


Table 5. Pearson's correlation between log(node impact) and centrality measures after log transformation.

Ln(s_bet-weenness)	S_close-ness	S_eccen-tricity	S-harmonic	Ln(clique eigenvector)	Ln(weighted betweenness)	Weighted closeness	Ln(weighted degree)	Ln(weighted eigenvector)
0.11	0.13	-0.07	0.18	0.15	0.09	0.17	0.16	0.13

Figure 7. Histogram of node impact and centrality measures after log transformation


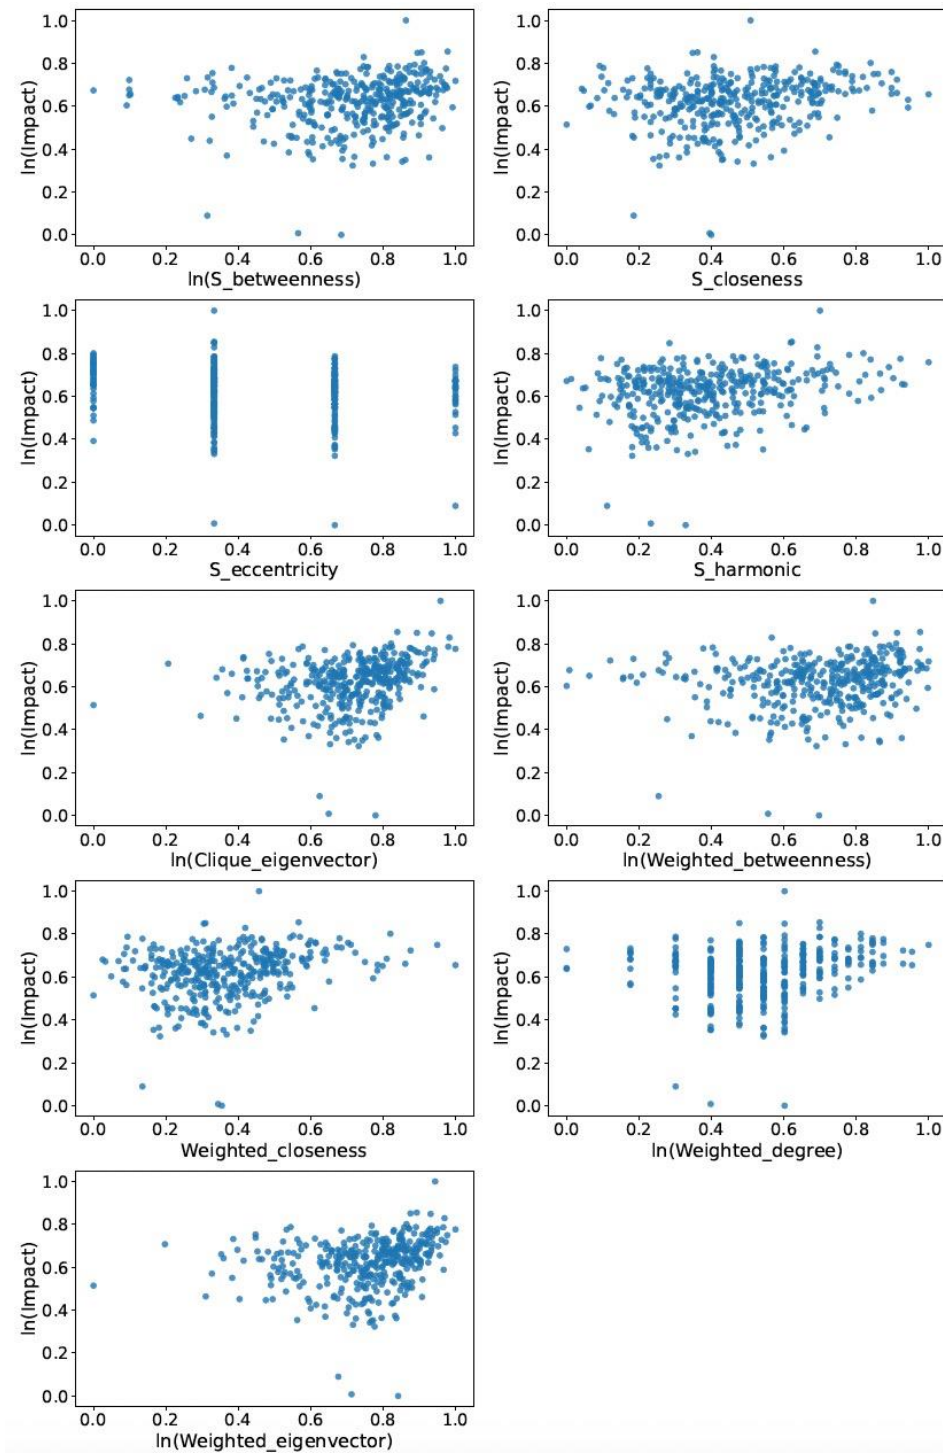



Figure 8. Node impact vs. centrality measures after log transformation



By repeating the analysis for 500 randomly generated hypergraphs, we obtain 500 correlation coefficient values between node impact and centrality measures after transformation. This approach allows us to assess the consistency and variability of the relationship between node impact and centrality measures across

different hypergraph instances. We plot the boxplot and violin plots for these 500 correlation coefficients. Fig. 9 shows that the correlation between node impact and s-eccentricity is explicitly lower than that with other centrality measures. The finding indicates that compared to s-eccentricity, other studied centrality measures can better predict node impact. Among them, the higher median value of clique eigenvector centrality suggests a bit higher predictive power of the centrality compared to other centralities. In addition, the taller boxplot for clique and weighted eigenvector centralities means larger range of correlation values between the centrality and node impact with log transformation, implying a less stable predictive power of the centrality in predicting node impact.

Figure 9. Box plot of correlation coefficients between node impact and centrality measures with log transformation



We also employ random forest regression (RFR) to assess the relationship between node impact and centrality measures and analyze their feature importance in predicting node impact (see Fig. 10). It demonstrates that weighted degree, s-harmonic and clique eigenvector centralities can predict better node impact compared to other measures.

Table 6. Metrics for prediction evaluation


Mean squared error (MSE)	Mean absolute error (MAE)	R2 score
1.13	0.78	0.12

Mean Squared Error (MSE): MSE represents the average difference between the predicted values and real values. The small NMSE indicates small average error of prediction for target variable.

Mean Absolute Error (MAE): MAE represents the average absolute difference between the predicted and actual values. The lower MAE indicates that the model's predictions have a smaller average absolute error.

R2 Score: R2 score represents the proportion of the variance in the target variable that can be explained by the model. A negative R2 score suggests that the model is performing worse than a model that predicts the mean value of the target variable. A positive R2 score closer to 1 indicates a better fit.

Figure 10. Centrality importance in predicting node impact from RFR (some after log transformation).



3.4. Conclusion

Using synthetic data in which the 'ground truth' can be objectively established, we could determine that our hypothesis that data models that include synergistic interactions does not lead to better predictions of driver nodes.

4. Dynamic importance of network nodes is poorly predicted by static structural features.

4.1. Preamble

In the previous sections we have seen that, based on Bayesian Networks that generate cross-sectional synthetic data, the prediction of driver nodes is ineffective. Here we turn to a different data-generating mechanism: a specific type of Glauber dynamics, called the kinetic Ising spin model. This model is capable of generating time-series data. The question in this section is two-fold: (1) is it possible to identify the driver nodes based on static network features, now that the network describes a different type of model; and (2) is it possible to identify the driver node based on rich time-series data? The second question is important as a dot on the horizon: if one would have access to very rich longitudinal data, would that allow identification of driver nodes?

To answer these questions, we simulate the model on randomly generated networks as well as one real-world weighted network (psychometrics data set). The dataset, inferred from relatively rich time-series, contains self-reported variables that are relevant to major depression, such as feelings of sadness, bereavement, and sleeping quality. The dynamic impact of nodes is assessed by causally intervening on node state probabilities and measuring the effect on the systemic dynamics. The results show that structural features such as network centrality or connectedness are actually poor predictors of the dynamical impact of a node on the rest of the network.

A solution is however offered in the form of an information theoretical measure named integrated mutual information. The metric is able to accurately predict the driver node in networks based on observational data of non-intervened dynamics. We conclude that the driver node(s) in networks are not necessarily the most well-connected or central nodes. Indeed, the common assumption of network structural features being proportional to dynamical importance turns out to be false. Consequently, great care should be taken when deriving dynamical importance from network data alone. These results highlight the need for novel inference methods that take both structure and dynamics into account.

An extended report of this study has been published in *Physica A: Statistical Mechanics and its Applications* (<https://doi.org/10.1016/j.physa.2022.126889>). For this deliverable, references and referral to supplemental material have been limited to the essentials, and for a complete account the reviewer may consult the published manuscript.

4.2. Introduction

The dynamic importance of a node varies as a function of both the dynamics that exist on the network in

addition to its structural connectedness. This effect was rigorously shown by Harush and colleagues [doi.org/10.1038/s41467-017-01916-3]. In their study, the dynamic processes were varied while keeping the network structure the same. Both random generated networks as well as real-world networks were studied. Nodal importance was computed based on the “information flow” through a node. The flow through a node had a non-linear relation between its structural connectedness and the type of dynamics present in the system. These results highlight the non-intuitive and non-trivial interplay of the structure of a system and dynamics between the nodes of a system. However, the results required full knowledge of the system, i.e. both the network structure and the dynamics of the system are required to estimate the node with the highest dynamic importance. For a real-world system having both the dynamics available and the underlying network structure may be difficult. In addition, in their study the network structure was deemed constant. It remains unclear whether the proposed flow would generalize to systems with varying network structure. In addition, the results were obtained from steady-state. For dynamical processes an out-of-equilibrium approach highlights *how* trajectories generate systemic behavior. What is needed are methods that can detect the highest dynamic importance of a node from observations directly without knowing or assuming the underlying network or causal structure among variables.

A node’s dynamic (or causal) importance is traditionally inferred by means of interventions and counterfactuals [<https://doi.org/10.1111/phpr.12095>]. Through external interventions the behavior of the system may change. That is, an intervention on a part of the system causes a divergence of the system behavior proportional to its dynamic importance; high causal importance relates to a large change in the system behavior. It is common for studies to use causal interventions to apply so-called hard intervention strategies. In hard interventions a node is pinned to a state. This effectively removes all incoming connections to a node, removing all influences the system has on this node. The use of hard interventions is common in various disciplines such as gene-knockout experiments, epidemic spreading, network analysis and driver node identification.

Applying external interventions to identify nodes with high dynamical importance, referred to as driver nodes, is challenging for three reasons. First, hard interventions may cause a change in systemic behavior that does not occur in the non-intervened behavior (e.g. see fig 11) or may be impossible in practice (e.g. requiring infinite resources: doi.org/10.1016/j.physrep.2016.06.004). This complicates the theory forming on what mechanisms underlie observed systemic behavior. Second, most approaches for causal inference assume that the causal interactions follow a particular structure without (local) loops (doi.org/10.3758/s13428-017-0862-1). This assumption simplifies theoretical analysis and is in many cases justified under the assumption that causes precede their effects in time. However, this approach assumes that the underlying causal structure is known and or accessible. If the underlying causal structure is not known, it has to be inferred from data which prompts problems in terms of temporal resolution and scale. For example, in a causal process where $A \rightarrow B \rightarrow$


$C \rightarrow A$, the cycle may be unrolled over time. This effectively removes cycles. However, this requires the data to support time resolution such that it reflects unrolled (non-cyclical) causal processes (doi.org/10.3758/s13428-017-0862-1). Cyclical causal structures do occur in real-world systems such as ecosystems, biological systems, gene-regulatory systems and so on. Importantly, for many of these systems determining the correct temporal scale for causal interactions is non-trivial. Consequently, it remains an open question on how to perform causal inference applying these in general for cyclical causal events. Thirdly, the underlying causal structure may not be known or difficult to determine (doi.org/10.1038/s41598-019-43033-9). In addition, many dynamical systems are prohibited from analytical approaches to decompose each nodal dynamical importance directly due to the polyadic, often non-linear, interactions (doi.org/10.1038/s41598-019-43033-9).

For closed systems, it is possible to avoid these challenges to deduce the driver node by considering cross-sectional time series and computing the correlation of each node with the entire system out-of-equilibrium and without lossy compression (DOI: [10.1073/pnas.1704663114](https://doi.org/10.1073/pnas.1704663114)). Here, a driver node is the element of the system that is correlated the most with the entire system as a function of time. It has the maximum dynamic (causal) impact when intervened upon across all nodes. By taking the entire system into account, any internal confounding information is implicitly accounted for; there cannot exist any other node with more dynamic importance than the node with the highest correlation as the system is closed and external confounding is excluded.


Figure 11.

(A) Driver node inference. Dynamics are simulated using kinetic Ising spin dynamics. Causal interventions (soft and hard) are shown in the left column. Structural metrics (betweenness, closeness, degree, eigenvector centrality) each produce different driver node estimates. Integrated mutual information predicts the driver node for soft causal interventions. High causal interventions produce different system dynamics (see B for an example) and different driver nodes. This figure shows that dynamics interact with structure to produce non-trivial driver node estimates. In addition, causal intervention size (hard or soft) influences the observed system dynamics. (B) Effect of intervention size on system magnetization with kinetic Ising spin dynamics (see 2.1). Energy is added (gray arrow) to the nodal Hamiltonian of the blue node (\cdot). (top) Non-intervened system dynamics are shown with the distribution of the blue node (right). (middle) Soft intervention on node keeps similar dynamics as the non-intervened system dynamics. (bottom) Hard interventions yield profound different system dynamics compared to non-intervened dynamics. (C) Example of non-causal inflation of mutual information decay. The network structure is given by the graph inset; consisting of a system with 4 nodes in which each node is binary variable. Node 0 has a 50/50 distribution, and all other nodes copy the state of its predecessor. Node 0 has the largest causal influence as it can influence the state of all other nodes in the system (over time). The information content of node 3 is biased; it stores the information from 0 but has no downstream causal effects. Yet, its information decay is similar to 1 which has one downstream node. The true driver node (0) has the largest information decay over time and is not confounded.


A



B



C



Shannon information theory offers profound advantages over previous approaches for defining a measure of dynamical importance of driver node on the behavior of all other nodes. Firstly, and most prominently, Shannon mutual information can quantify statistical associations among variables without bias to specific forms of association. In particular, it equals zero if and only if the full probability distribution of the system state remains exactly unchanged regardless of the state of the node. Second, mutual information does not require *a priori* knowledge of the representational base of the system. That is, it allows for direct comparison among different systems that may have different units of measurement such as currency, density of animals, voltage per surface area and so on. Finally, it is defined for both discrete-valued and real-valued state variables.

The concept of measuring the dynamic importance of a node through information flow is not new. Colloquially, information flows from process X to process Y represents the existence of statistical coherence between the present information in Y and the past of X not accounted for by the past of Y . Various methods have been proposed in the past such as transfer entropy (doi.org/10.1016/j.compbio.2016.10.010). Although originally intended as a predictive measure, the notions of information and information flow can be extended to causal influence or dynamic importance. Previous research developed several measures and methods for determining how much information flow between two processes is truly causal; examples include (but not limited to) conditional mutual information under causal intervention, causation entropy, permutation conditional mutual information, time-delayed Shannon mutual information (e.g., doi.org/10.1038/s41598-023-32762-7; DOI: 10.1103/PhysRevE.97.052216).

These measures are commonly used to infer the information transfer between sets of nodes by possibly correcting for a third confounding variable (doi.org/10.3390/e22080854). That is, informational flows are used to determine how information is transferred or shared between pair or sets of variables. However, in polyadic settings most measures of information flow are prone to underestimate or overestimate nodal importance. Determining how much information flow is causal between source and sink variables in polyadic settings remains difficult due to the so-called synergetic and redundant information (doi.org/10.48550/arXiv.1602.01265)

Instead of focusing on the open problem of full information decomposition among variables, we focus here on the amount of information that a node shares with the entire system. A driver node is expected to have a corresponding high information flow from it to the system. We introduce a novel metric named integrated mutual information (IMI) based on time-delayed Shannon mutual information with a node and the entire system over time that captures driver nodes with the highest *causal* impact in ergodic systems. Additionally, we avoid the problem of confounding information synergy with redundancy by avoiding conditional mutual information. The consequence of this is that we quantify only *how much* causal impact *one single node* has on the *entire* network. In other words, we will not be able to infer which parts of the network are impacted more than others; nor how exactly the causal influence percolates through the network. This is nevertheless sufficient for the purpose of this study.

Using this approach it was previously shown analytically that the number of connections of a node does not necessarily scale monotonically with its dynamical importance in infinite-sized, locally-tree-like networks (i.e. networks without loops) (doi.org/10.1098/rsif.2013.0568). For those networks, nodes with high dynamic importance were not nodes with high degree (so-called hubs). Instead, the nodes with intermediate connectedness were identified with the highest dynamical importance. This study extends this prior research by numerically computing information flow in finite-sized, random networks, i.e. where small feedback loops

cannot be ignored, and studies the relation between connectivity and the dynamical importance of nodes in such networks.

The aim of the work in this chapter is to test the common (often implicit) hypothesis that the connectedness of a node is proportional to its dynamical importance. A node's dynamic importance is determined by simulating the (out-of-equilibrium) dynamics of a system under causal external interventions. The distance between the system state probability distributions with *and* without performing the intervention is used as a ground truth for a node's causal influence over time. The resulting causal impact score over time is compared with common network centrality metrics as well as our proposed information-based metric.

In contrast to other studies, the proposed metrics in this study do not make assumptions on network structure or type of dynamics. Small networks are used (up to 12 nodes) for which each node has an associated discrete state whose dynamics is governed by a stochastic update rule in discrete time. Smaller network sizes have the advantage of studying "network motifs" that are embedded in larger network structure; understanding causal flows of smaller structures may provide insights into larger systems consisting of a combination of smaller network structures. In addition, the use of small network sizes offers a numerical advantage for accurate estimation of causal and information measures. In total 16 Erdős-Rényi networks are generated (10 nodes), and one real-world weighted network (12 nodes) obtained from (doi.org/10.1037/abn0000028). Temporal dynamics of the nodes are simulated using kinetic Ising spin dynamics with Glauber updating for the purpose of demonstrating a case of non-trivial relation between network connectivity and dynamic importance, but our approach easily generalizes to any other dynamics or networks generation model.

The results show that nodes with generally high structural connectedness as measured by closeness, betweenness, eigenvector and degree centrality are not the *driver node(s)*. Our novel metric, IMI, achieved significantly better performance in predicting the driver node for non-intervened dynamical systems. In addition, and most importantly, hard causal interventions lead to causal flows that differ from the non-intervened dynamics. That is, as a function of external intervention, systemic behavior may result in "unnatural" system behavior that does not occur in the unperturbed system. The proposed metric IMI does not rely on the assumptions on dynamics, nor on assumptions on structural properties of the network. Therefore, the results of this study provide scientists of all fields with a novel, reliable and accurate metric for the identification of driver nodes.

4.3. Theoretical background

4.3.1. Terminology

In this paper, we consider a complex system as a set of discrete random variables $S = \{s_1, s_2, \dots, s_n\}$ with interaction structure $E = \{(s_i, s_k) | s_i, s_k \in S\}$, where each $s_i \in S$ has an alphabet A . This is also known as a (discrete) dynamical network. Please note that we use the term node and variable interchangeably referring to an element $s_i \in S$. The system chooses its next state S^t in discrete time with probability:

$$p(S^t | S^{t-1}, \dots, S^{t_0}) = p(S^t | S^{t-1}),$$

which is also known as a first-order Markov chain. More specifically, each discrete time step, a single variable $s_i \in S$ is chosen with uniform probability and updated. That is,

$$p(S^t | S^{t-1}, \dots, S^{t_0}) = p(S^t | S^{t-1}) = \prod_j p(s_j^t | S^{t-1}) = p(s_i^t | S^{t-1}).$$

For temporal dynamics, we adopt here the Metropolis-Hasting algorithm. By drawing a proposal state $S^{t+1} = X'$ from current state $S^t = X$ from a proposal distribution $g(X'|X)$ and accepting the new state X' with probability,

$$\begin{aligned} A(S^{t+1} = X', S^t = X) &= \min\left(1, \frac{p(S^{t+1} = X')g(S^{t+1} = X | S^t = X')}{p(S^t = X)g(S^{t+1} = X' | S^t = X)}\right) \\ &= \min\left(1, \frac{p(s_i^{t+1} = x')g(s_i^t = x | s_i^t = x')}{p(s_i^{t+1} = x)g(s_i^t = x' | s_i^t = x)}\right), \end{aligned}$$

where $x' \in A$. If the new state is not accepted, then the next state will be set to $S^{t+1} = X$. As the next state S^{t+1} is determined through considering updating a single variable $s_i \in S$, the next state X' will be generated through the proposal state x' drawn uniformly from the possible states A such that $g(x'|x) = g(x|x') = g(x') = \frac{1}{|A|}$. This means that $\frac{g(X'|X)}{g(X|X')} = \frac{g(x'|x)}{g(x|x')} = 1$.

For our experiments, we use the kinetic Ising model which is thought to fall in the same universality class as various other complex behaviors, such as (directed) percolation, diffusion, and many extensions of the model have been used as a base for opinion dynamics, modeling neural behavior and so on. For demonstrating our primary claim that high network connectivity does not necessarily lead to high dynamical importance, a single dynamics suffices as counterexample.

It is nevertheless important to emphasize that our proposed driver node inference method does not depend on the exact type of dynamics. That is, the choice for the kinetic Ising spin dynamics is arbitrary in that respect. Our methods only require that for a given dynamics the data-processing inequality is satisfied. More details

on the methods and its assumption will follow in below.

The kinetic Ising model consists of binary variables dictated by a Gibbs distribution that interact through nearest neighbor interactions. A prominent property of the Ising model in higher dimensions (two or more) is the phase transition from an ordered phase to a disordered phase by increasing the noise parameter β . For finite systems, the kinetic Ising model shows a continuous phase transition from ordered to unordered system regime (fig. 12 B). The Metropolis-Hastings update rule specifically for the kinetic Ising model equals:

$$p(\text{accept } X') = \frac{p(X')}{p(X)} = \begin{cases} 1 & \text{if } \mathcal{H}(X') - \mathcal{H}(X) < 0 \\ \exp(-\beta(\mathcal{H}(X') - \mathcal{H}(X))) & \text{otherwise,} \end{cases}$$

where $\mathcal{H}(S)$ is the system Hamiltonian defined as

$$\mathcal{H}(S) = - \sum_{i,j} J_{ij} s_i s_j - h_i s_i.$$

Here β is the inverse temperature $\frac{1}{k_b T}$ with Boltzmann constant k_b , J_{ij} is the interaction strength between variables s_i and s_j ; h_i represents external influence on node i . The matrix J effectively represents the network of the system: an edge between $s_i, s_j \in S$ exists if $|J_{ij}| > 0$. The networks considered in this study are undirected networks, which means that J is an $|S| \times |S|$ symmetric matrix. For the randomly generated networks the edges $J_{ij} \in \{0,1\}$; for the real network dataset the edge weights are positive and negative real numbers.

The β parameter can be seen as the (inverse) noise parameter in the system. Low values of β will induce each node in the system to detach from the influence of its neighbors, i.e. the probability of finding a node in a state $p(s_i = a), a \in A$ will tend to uniform distribution as $\beta \rightarrow 0$. In contrast, high values of β increases the influence a neighbor of a node may have on determining the node's next state.

4.3.2. Causal interventions and dynamic importance

We call a node a *driver node* for a dynamical system when it has the largest causal impact on the system dynamics. In brief, we will determine a node's causal impact by simulating a transient intervention and subsequently quantifying the difference between the system dynamics under intervention and without intervention. We expect that the impact of the driver nodes will penetrate deeper into the system and remain present longer than for nodes with lower causal impact. Here, we define causal impact by means of external intervention on a node. The external causal intervention \vec{e} on node $s_j \in S$ can be described as

$$p_{s_j}'(s_i^t | S^{t-1}) = p(s_i^t | S^{t-1}) + \vec{\epsilon} \delta_{ij},$$

where δ_{ij} is the Kronecker-delta, and $\dim(\vec{\epsilon}) = |A|$. In addition, $\sum_{i=0}^{|A|} \epsilon_i = 0$ and $\sum_{i=0}^{|A|} |\epsilon_i| = c$ for some $c \in (0,1]$. Note only those $\vec{\epsilon}$ are allowed that generate valid new probabilities, i.e. $0 \leq p' \leq 1$.

Relative to some equilibrium distribution $p(S^\tau)$, the effect of intervention $\vec{\epsilon}$ will result in a new system state equilibrium distribution $p'(S^\tau)$. Subsequently the intervention is removed at a random system state, after which the distribution of system states will gradually converge back to the original equilibrium. Nodes with higher dynamic importance will cause a larger difference in the system state probability distribution over time. Consequently, we quantify the causal impact of a node by integrating over time t the difference in system state distribution from the moment the intervention is released ($t = \tau$). Since our model is discrete in time, the integral becomes a summation. Thus, we define the causal impact of node $s_i \in S$ as

$$\begin{aligned} \Gamma(s_i) &= \sum_{t=\tau}^{\infty} \gamma(s_i^t) \Delta t \\ &= \sum_{t=\tau}^{\infty} \sum_{s_j \in S} D_{KL}(p'_{s_i}(s_j^t) | | p(s_j^t)) \Delta t \end{aligned}$$

where D_{KL} is the KL-divergence, and $\Delta t = 1$ throughout the paper. KL-divergence $D_{KL}(p | | q)$ quantifies the difference between probability distributions and is non-negative, invariant under affine parameter transformation and zero when $p = q$. For example, if $\Gamma(s_i) = 0$ the intervention on nodes i caused no difference in any state probabilities.

The driver node can then readily be defined as $\operatorname{argmax}_{s_i \in S} (\Gamma(s_i))$. A numerical implementation of driver node identification is given in section 4.4.7.

In this study we allow the intervention to evolve for some time period $\Delta t_{\text{nudge}} = \tau - t_0$ after which the nudge is removed (see 4.4.7 and fig. 12 C). The intervention will transiently bring the system out of equilibrium. For ergodic systems, the effect of the intervention will be lost from the system over time. Namely, $P_{s_i}'(S), \forall s_i$ will tend to equilibrium distribution $P(S)$ from τ on wards. The causal impact is computed relative to this τ (fig. 12 C). As a consequence, $\Gamma(s_i)$ will be finite for systems under study here, but may diverge for non-ergodic systems. The duration τ needs to be set appropriately that the intervention is allowed to percolate through the system. Here we used $\tau = 15$ for random networks consisting of 10 nodes and $\tau = 25$ for the real-world weighted network consisting of 12 nodes.

The time prior to τ is used for computing the causal impact on a node as the intervention could be disproportional affected by the act of changing the node distribution. It does not accurately reflect the causal impact of the node on the rest of the system. Only the decay after τ would be proportional to the causal impact a node has on all other nodes. The causal impact is therefore computed based on the decay of the causal impact relative to τ .

4.3.3. Intervention size

In experiments concerned with measuring causal flows in networks, often "hard" causal interventions are used to determine the causal impact of nodes. Hard causal interventions are those that (effectively) remove all inputs from a node. In contrast, soft interventions keep the existing causal inputs intact but add an additional causal effect to the dynamics of a variable. In general, the larger the added causal effect of the soft intervention, the more this intervention will 'overpower' the existing causal effects and hence the more the soft intervention converges to a hard intervention. In non-linear dynamical systems the intervention size is crucial, since even small interventions can have large effects, especially in the presence of bifurcations. In order to characterize the system without intervention as much as possible, one should thus aim to intervene with a minimal (but still measurable) effect. To illustrate this further, consider the system depicted in fig. 11 B, where each node is updated with Metropolis – Hastings rule. The figure shows the effect of intervention size on the observed system dynamics. Namely, when a hard intervention is used (bottom) the system magnetizes. In contrast, in the non-intervened system, the system magnetizes periodically, i.e. the system dynamics evolve with time periods of magnetization and metastable switches to the other side of the magnetization. Contrasting the non-intervened system (top plot) with the bottom plot (hard) intervention, shows how the metastable behavior disappears. Soft interventions (middle plot) maintain the meta-stable behavior of the system.

There will always be a minimal intervention size for which there is a measurable (i.e. non-zero) causal effect, given a finite amount of data. The causal impact will be maximal for hard interventions. As the intervention size is decreased so will causal impact. We hypothesize that there exists a lower bound for the soft interventions for which there will be a measurable causal impact. This minimal intervention size is dependent on the system's structure as well as the dynamics of the system and is difficult to determine *a priori*. Yet it is important to approach this minimum because if the intervention size is too large, then the intervened system dynamics will diverge from the non-intervened system dynamics, losing its representative capacity. Therefore, we estimate this minimum numerically, described further below.

4.3.4. Intervention in kinetic Ising model

In order to measure the causal intervention, transition probability intervention needs to be adjusted. For the kinetic Ising model, the transition probability can be modified by adding *energy* to the node which results in

the adjusted Hamiltonian

$$\mathcal{H}_{\text{nudge } s_k}(S) = - \sum_{i,j} J_{ij} s_i s_j - \sum_i h_i s_i - \eta s_i \delta_{ik}.$$

That is, an energy term is added to the intervened node s_k by letting it interact with an external 'spin state' η . Both soft and hard interventions were used. For hard interventions ($\eta \rightarrow \infty$) the nodal dynamics of the nudge node s_k are no longer influenced by nearest neighbor interaction. That is, the node state for the intervened node s_j will not change as a function of time. For soft interventions a grid-search was used with $\vec{\eta} := \{\eta: \eta = 0.55 + i, i \in \{1, \dots, 10\}\}$ to determine a minimal yet measurable intervention strength. Once determined the same η was applied to each node in turn and this process is repeated for all networks in the experiments.

4.3.5. Measuring information flow

Each node in a dynamical system can be considered as an information storage unit. For example in social networks gossip can be considered as information one person possesses. Similarly, disease can be present in one city while being absent in another. Over time through interaction, this information stored in a node will percolate throughout the system while at the same time decaying due to noise. The longer the information of a node stays in the system, the longer it could affect the system dynamics. Therefore, dynamic impact of a node is upper-bounded by the amount of information a node shares with the entire system.

How does one measure information stored in a node? A node s_i dictated by stochastic and ergodic dynamics can be considered a random variable. In Shannon information theory information is quantified in bits, i.e. yes/no questions concerning the outcome of a random variable. The average information that a random variable can encode is called entropy and is defined as:

$$H(s_i) = - \sum_{s_i=x} p(x) \log p(x).$$

Note all log are base 2 in this paper unless specified otherwise.

Entropy can also be interpreted as the amount of uncertainty of a random variable. In the extremes the random variable either conveys no uncertainty (i.e. a node always assumes the same state), or is randomly chosen between all possible states (uniform distribution). For example, consider a coin flip. One may ask how much information does a single coin flip encode? If the coin is fair, i.e. there is equal probability of the outcome being heads or tails, the amount of questions needed to determine the outcome is exactly 1. In other words, a fair coin encodes 1 bit of information. However, when the coin is unfair the information encoded is less than one. In the extreme case where the coin always turns up heads, the entropy is exactly 0.

The information shared between a node state s_i and a system state S can be quantified by mutual information. Mutual information can be informally thought of as a non-linear correlation function which inherits its properties from the Kullback-Leibler divergence. Formally, mutual information quantifies the reduction in uncertainty of random variable X by knowing the outcome of random variable Y :

$$\begin{aligned} I(X:Y) &= \sum_{x \in X, y \in Y} p(x, y) \log \frac{p(x, y)}{p(x)p(y)} \\ &= H(X) - H(X|Y), \end{aligned}$$

where $p(x)$ and $p(y)$ are the marginals of $p(x, y)$ over X and Y respectively, and $H(X|Y)$ is the conditional entropy of X . The conditional entropy $H(X|Y)$ is similar to the entropy; it quantifies the reduction in uncertainty of the outcome X by knowing the outcome of Y . Please note that the yes/no question interpretation even applies to continuous variables; although it may take an infinite amount of questions to determine the outcome of a continuous random variable.

4.3.6. Mutual information and causality

Under some contexts, mutual information can be interpreted as being causal. For example, for a Markovian, isolated system consisting of two nodes $S = \{s_i, s_j\}$ where s_j causally depends on the previous state of s_i , encoded by a conditional probability distribution. It can be readily shown that the causal impact (eqref:causal_impact) of s_i on s_j then reduces to mutual information. In this study, the entire system is considered as a *downstream* node. Importantly, the system is assumed to be measured in its entirety: no confounding variables are allowed. Within this context, the node with the highest integrated mutual information contains more causal information than correlation information than any other node in the system. Since there cannot exist any other node with more mutual information over time by construction of the system, this node must be the driver node. For any other node the measured mutual information can be (significantly) conflated (fig 11 C). Note that for directed graphs, these methods will not work as by construction the highest mutual information may lead to wrongful identification of the driver node.

Additionally, for directed graphs the arrow of time causes a different interpretation on what correlation with the driver node means in cases without unobserved variables. For $t_0 + t$ where $t > 0$, integrated mutual information may correspond to correlation information of a driver node sender, whereas for $t < 0$, the mutual information relates to the correlation with a driver node sending information. The systems in this study, however, are systems without confounding possessing and undirected network structure. Consequently, information flows symmetrically across the interaction of two nodes and the driver node corresponds to the node with the highest integrated information.

4.3.7. Integrated mutual information

In a network of nodes each causal relation (edge) is obviously not isolated, so confounding variables exist. Therefore, the mutual information between a node state and a future state $I(S^{t_0-t}; s_i^{t_0})$ cannot be interpreted purely causal in general. Namely, this mutual information could in principle be created purely by another variable $s_j^{t_0-\omega}$ influencing both $s_i^{t_0}$ and S^{t_0-t} , even if there exists no causal influence from s^{t_0} to any other variable.

Assuming the network itself is isolated, the only node for which the mutual information with a future system state could not have been fully created by a confounding variable is the driver node. That is, the driver node has the largest mutual information with the future system state. If this were fully induced by a confounding variable, then a different node would have to have even larger mutual information with the same future system state. In addition, this must hold for all future system states. By definition of the driver node, this cannot be true under the condition that the system is isolated.

This point is illustrated in Fig. 11 C. Note that for all other nodes, however, it is possible for its mutual information value to be inflated due to non-causal correlations. This may result to a non-zero mutual information $I(s_i^{t_0-t}; s_j^t)$ among the two variables even if they do not depend on each other in causal manner (Fig 11). Consequently, we define the integrated mutual information (IMI) for node s_i as the driver node, i.e. the node with the largest causal impact over time. The information flow over time $\sum_{t=t_0}^{\infty} I(s_i^{t_0+t} | S^{t_0}) \Delta t$ is causal in closed systems lacking confounding variables (fig.11C), where S^{t_0} is the system state at some time t_0 and $s_i^{t_0-t}$ is the state of a node t away from that system state. At time $t = 0$ the value equals $I(s_i^{t_0}; S^{t_0}) = H(s_i^t)$ for any node.

Here, undirected networks are considered. Due to detailed balance for undirected networks, there exists a time symmetry in terms of variable dynamics. This means that for the systems considered in this study $I(s_i^{t_0-t}; S_0^t) = I(s_i^{t_0+t}; S_0^t)$ (see appendix 8.2.2 in <https://doi.org/10.1016/j.physa.2022.126889>). It is computationally easier to compute $I(s_i^{t_0+t}; S_0^t)$ rather than the reverse (see section. ref:sec:dist_est). For directed graphs, however, the meaning and interpretation of integrated mutual information changes depending on the direction in time it is computed (see **appendix 8.2.2** in <https://doi.org/10.1016/j.physa.2022.126889>). This effect is outside the scope of the present study and will be the subject of future studies. For undirected graphs, the causal impact is time invariant and is equal forward and backward in time.

$$\mu(s_i) = \sum_{t=t_0}^{\infty} I(s_i^{t_0-t}; S^{t_0}) \Delta t,$$

For all ergodic Markovian systems, the delayed mutual information $I(s_i^{t_0-t}; S^{t_0})$ will *always* decay to zero as $t \rightarrow \infty$ (doi.org/10.48550/arXiv.1602.01265)

. This decay is monotonic, which follows from the data-processing inequality and states that information can never increase in Markov chains without external information injection (see appendix 8.1 in <https://doi.org/10.1016/j.physa.2022.126889>). The question is *how fast this decay takes place for each node* (fig. 11 C), and consequently how much *informational impact* the node will have on the system.

Taking the entire system as a downstream 'node', $I(S^{t+1}; s_i^t)$ represents pure causal information flow for the driver node as long as there are no confounding variables. For the rest of the study, we will therefore use the notation $I(s_i^{t_0+t}; S^{t_0})$.

4.4. Methods and network data

4.4.1. Network data

Random networks. In total 16 Erdős-Rényi random networks were generated consisting of 10 nodes each. Each random network is generated by first drawing a random connection probability uniformly from [0,1], and then creating each possible undirected edge with probability r . Out of these 16 networks, ~82 percent had a single connected component (fig. 12). Each edge had unitary weight.

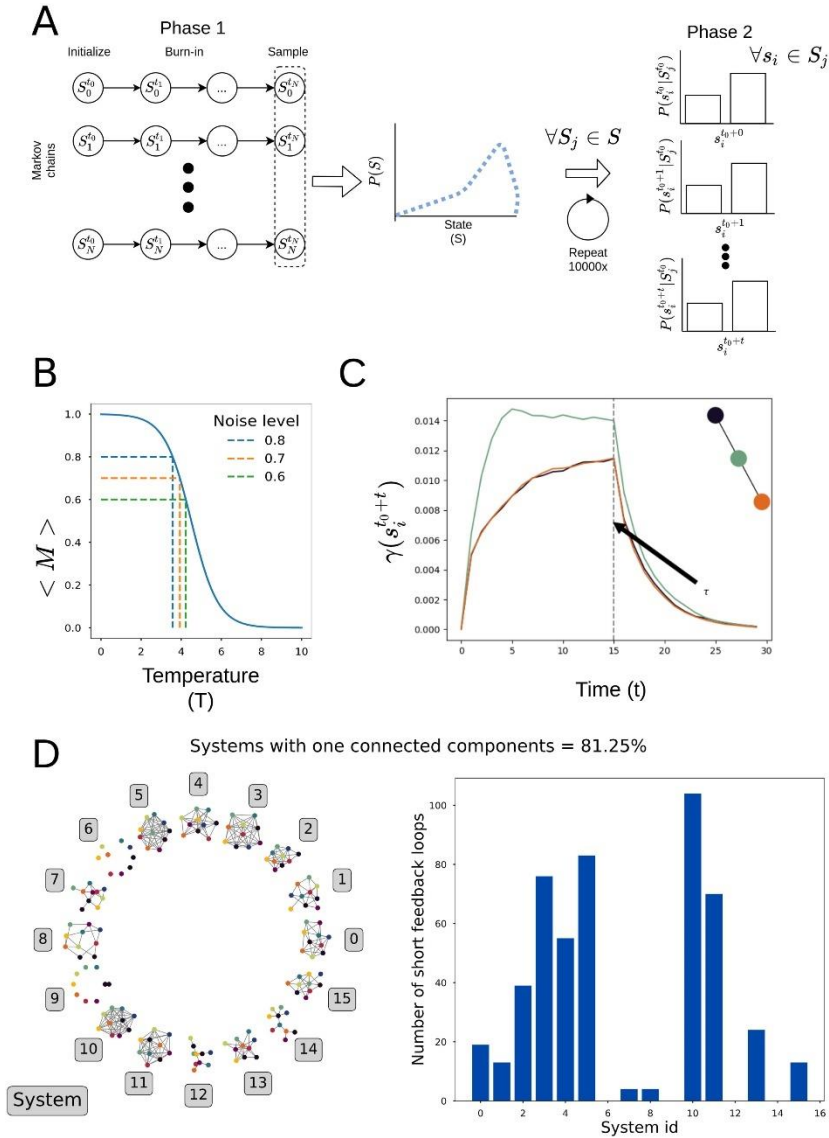
Real-world network: psychosymptoms. In addition to the generated random networks, we also test a small weighted network inferred from real data. This network differs from the random networks in that it is weighted and reflecting interactions among variables as a consequence of a real-world process, as well as reflecting inferred interactions from real data. The network data originates from the Changing Lives of Older Couples (CLOC) and compared depressive symptoms assessed via the 11-item Center for Epidemiologic Studies Depression Scale (CES-D) among those who lost their partner (N=241) with still-married control group (N=274) (<https://doi.org/10.1037/abn0000028>). Each of the CES-D items were binarized with the aid of a causal search algorithm using Ising model developed by and represented as a node with weighted connections. For more info on the procedure see doi.org/10.1038/srep05918 and <https://doi.org/10.1037/abn0000028>. The 11 CES-D items are (abbreviated names used in the remainder of this text in brackets): 'I felt depressed' (depr), 'I felt that everything I did was an effort' (effort), 'My sleep was restless' (sleep), 'I was happy' (happy), 'I felt lonely' (lonely), 'People were unfriendly' (unfr), 'I enjoyed life' (enjoy), 'My appetite was poor' (appet), 'I felt sad' (sad), 'I felt that people disliked me' (dislike), and 'I could not get going' (getgo).

4.4.2. Numerical methods

Magnetization matching. A prominent feature of the (kinetic) Ising model is the phase change that occurs as a function of noise (fig.12B). In this paper we tested whether the amount of noise would (i) affect which nodes becomes the driver node in the system, and (ii) whether the correct driver node could be predicted using either IMI or network centrality metrics. We tested three levels of noise (fig 12 B); a low noise level (80% of the maximum magnetization), a medium noise level (70%), and a high noise level (60%) were used. This magnetization matching was achieved by estimating the magnetization curve as a function of β^{-1} numerically.

Figure 12.

(A) Phase 1 initializes 10000 independent random Markov chains and equilibrated the chains for time steps. The state distribution is estimated at over the 10000 chains. In phase 2 for system states conditional distributions are estimated for time-steps from which mutual information is estimated. (B) Illustration of temperature matching. The graph depicts the conceptual transition between an ordered state (aligned spins) to a disordered state for the kinetic Ising on complex networks. The noise level was matched to the magnetization ratio of the max magnetization. For increase in temperature the noise level increases. (C) Illustration of applied interventions in an undirected 3 state system (top right corner). Each node is nudged according to the equation cited above. The time prior to “pushes” the system dynamics out of equilibrium. From onwards the causal impact decays proportional to the causal impact a node has on the rest of the system. (D) Generated Erdős–Rényi networks. (left) Structure of each network, the number indicates the system id. (right) Number of triangles (feedback loops) for each system.



4.4.3. Estimating $p(S^{t_0})$ and $p(S^{t_0+t}|S^{t_0})$

For each noise level, N independent Markov chains are run for simulation with 1000 steps (fig. 13 A). Each chain was initialized with random state distribution over nodes. We are interested in out-of-equilibrium dynamics. For the Ising model the node dynamics are symmetric in equilibrium, i.e. each node has maximum entropy. For any given trajectory, the time-scale for the system to move from a fully magnetized state, to the opposite magnetized state will take longer than the dynamics that are exerted on local time scales. For short time-scales nodes with higher degree be more affected by large scale changes in their environment than lower degree nodes. For larger time-scales high-degree nodes will be frozen, and do not reflect short-time scale dynamics. In light of this, the simulations are conducted by sampling one side of the magnetization. This has the added benefit that it reduces the sign of the intervention, i.e. interventions are made to introduce higher entropy in each node which depends on the average magnetization in the system.

Each simulation step executes the following:

- Pick a node at random from the system with equal probability;
- Compute energy using eq. [eqref:energy](#);
- Flip the node state with probability eq. [eqref:hastings](#).
- From this set, the equilibrium distribution over states $p(S^{t_0})$ was constructed in the form of sample of N system states.

For this sample of system states, Monte-Carlo methods are similarly used to construct the conditional $p(s_i^{t_0+t} | S^{t_0})$. For each of the sample states $S_i \in S^{t_0}$, the procedure above was repeated 100000 times for 100-time steps in the psychosymptoms and 30 time steps for the random generated networks. All numerical experiments were repeated $n_{\text{trials}} = 20$ times to provide confidence intervals for the results in all nudge conditions and temperature settings (noise conditions).

4.4.4. Time symmetry and mutual information

Thusfar, the definition of causal impact and IMI is ambiguous to the whether t is positive or negative. Namely, if the node state $s_i^{t \pm 1}$ is captured in forward in time or backward in time with respect to some state S^t . For undirected networks with Ising spin dynamics there exists time symmetry with respect to how causal influence flows through the network due to detailed balance (see appendix 8.2.2 in <https://doi.org/10.1016/j.physa.2022.126889>). However, for directed networks this is not the case. The results obtained here are obtained using forward simulation in time only. The detailed balance condition ensures that the results would be symmetric when simulating the system backwards in time.

4.4.5. Area under the curve estimation

The mutual information over time and KL-divergence over time were scaled for visual purposes in the range $[0,1]$ per trial set. A double exponential, $y = a \exp(-b(t - c)) + d \exp(-e(t - f))$, was fitted to estimate these curves and subsequently IMI (eq. [eqref:information_impact](#) and [eqref:causal_impact](#)) using least squares regression. The kernel showed to be a good fit as indicated by the low fit error.

4.4.6. Sampling bias correction

Empirical estimates for mutual information are inherently contaminated due to sampling bias. In order to correct this, Panzeri-Treves correction was applied. This method offers a good performance in terms of signal-to-noise and computational complexity.

4.4.7. Driver node prediction and precision quantification

It is possible for two or more nodes to have exactly equal network structure features as well as node dynamics.

For example, consider a ring structure where each node has the exact same connectivity and all nodes have the same dynamics. In this case each node must have the same causal effect, and it would be impossible to disentangle these nodes causally from one another. Similarly graphs that are similar, e.g. show high degree of structural similarity but are not isomorphic, this causal separation may prove difficult for finite samples in stochastic settings as was discussed in ref:sec:dist_est. Consequently, we applied a parametric bootstrap procedure to estimate driver node sets (algorithm eqref:alg:driver_set).

Driver set estimation. The area under curve values, e.g. integrated mutual information and causal impact, were resampled to generate bootstrap distributions (see appendix 8.7 in <https://doi.org/10.1016/j.physa.2022.126889>). This creates a confidence interval for the integrated mutual information and causal impact. From these bootstrap distributions a driver node set is estimated (see algorithm of appendix 8.7 in <https://doi.org/10.1016/j.physa.2022.126889>). The bootstrap procedure constructs driver node set Λ iteratively by comparing the overlap ϕ with of the bootstrap for each variable with the distribution of the estimated driver node. For all experiments $\phi = 0.5$. The driver node distribution was taken as the bootstrap distribution with the highest mean. Variables will be included to the driver node set Λ if the overlap between its bootstrap distribution i and driver node bootstrap distribution j was $\phi_{ij} > 0.5$. In total $N = 1e4$ bootstrap trial were constructed of size $M = n_{\text{trials}} = 20$; for each of the trials the average was computed. For each variables a Gaussian distribution was estimated over the N bootstrap. This distribution was used for computing the overlap with the driver node distribution.

The bootstrap procedure cannot be applied to the network centrality metrics as there exists only one centrality rank assignment per network structure. Therefore, the driver nodes as inferred by maximum centrality metrics is the set of nodes (Λ) whose centrality metric (f) equals this maximum value, i.e.

$$\Lambda_{cent} = \operatorname{argmax} f_{cent}(x) := \{x \in S: f_{cent}(s) \leq f(x) \forall s \in S\}$$

where f is the centrality function which assigns a real value to each node in the system. For degree centrality,

$$f_{deg}(a_i) = \sum_j a_{ij}$$

where a_{ij} is the weighted connectivity between node i and j in the adjacency matrix A of the network. If $a_{ij} > 0$ node i and j are connected. In this study degree, betweenness, closeness or eigenvector centrality were used. (see appendix 8.6 in <https://doi.org/10.1016/j.physa.2022.126889> for the formal definitions for the centrality measures).

Ground truth comparison. The ground truth values are the driver node estimations for the causal impact bootstrap distribution. Each estimator also generated a driver node set estimation. That is for integrated mutual information, degree centrality, closeness centrality, betweenness centrality, eigenvector centrality,

the bootstrap distribution generates an estimated driver set. To evaluate the performance of these estimators, an overlap score was computed with the ground-truth (causal impact estimation) using the Jaccard similarity metric:

$$J = \frac{A \cap B}{A \cup B}.$$

A Jaccard score of 1 means perfect overlap, i.e. the driver node set identified by causal impact and predictor set by IMI or one of the centrality metrics are identical. Conversely, a Jaccard score of 0 means completely disjoint driver-sets.

For every network, intervention size and temperature the similarity metric we computed the Jaccard score per predictor. Additionally, the relative performance ratio of the IMI predictor was evaluated by

$$R_{cent} = J_{IMI} - J_{cent}$$

where J_{cent} is the Jaccard score of the structural metrics (betweenness, closeness, betweenness, eigenvector centrality). This performance indicator falls within $[-1, 1]$ range: A value of -1 would indicate that the centrality metric correctly identified the driver node set; a score of 0 would indicate equal performance for the driver node identification between a centrality metric and IMI; a score of 1 would indicate a correct identification of the driver node for IMI, but a false identification of the centrality metric. The ratios were bootstrapped ($N = 100000$) and tested for significance at $\alpha = 0.01$. The driver node inferred performance is bound between $(0, \infty)$.

4.4.8. Software


A general toolbox was developed for analyzing any discrete systems using IMI, e.g., Susceptible-Infected-Recovered, Random Boolean networks. The core engine is written in 3.07a with python 3.9.4 and offers C/C++ level performance¹, for more information see https://github.com/cvanelteren/information_impact.

Figure 13.

(A) Jaccard score per system (see Fig.12D for the various network structures) as a function of intervention size. Each row depicts an increase in noise indicated by the label on the right side of the plot. On average integrated mutual information is a better predictor for the driver node than the centrality metrics used in this study (see also Fig. 14). Additionally, centrality metric tends to become predictor for hard interventions (δ), whereas this is not the (generally) so for integrated mutual information. Hard interventions can cause different causal dynamics not present in unperturbed system dynamics (see B/C). Integrated mutual information is based on observations of the system. The Jaccard score indicates that integrated mutual information can infer driver nodes for unperturbed dynamics. (B) Example of typical experimental results. The figures highlight a selection from the generated random networks form A for the low noise condition. (top) The graph structure of the

¹ `codebase`

system and the mutual information decay curves (second from top). (Second bottom and bottom) Show the causal impact decay for soft intervention and hard intervention. The gray dotted line indicates where the nudge is removed from the system, the produces causal decay proportional to a node’s causal importance. For soft causal interventions the inferred driver node based on integrated mutual information (second from top) is predictive for the true causal driver node (second from bottom) with soft causal interventions. Importantly, the causal driver node may change as a function of the intervention size (second bottom vs bottom plot). The centrality metrics tend to not correctly identify the driver node (see C for an example). (C) Mutual information decay ($I(S_i^{t+\Delta t}, S_i^t)$) (top) and causal impact (for minimal soft intervention ($\eta^* = 1.1$) (middle) and hard intervention ($\eta = \infty$) (bottom)).



4.5. Results

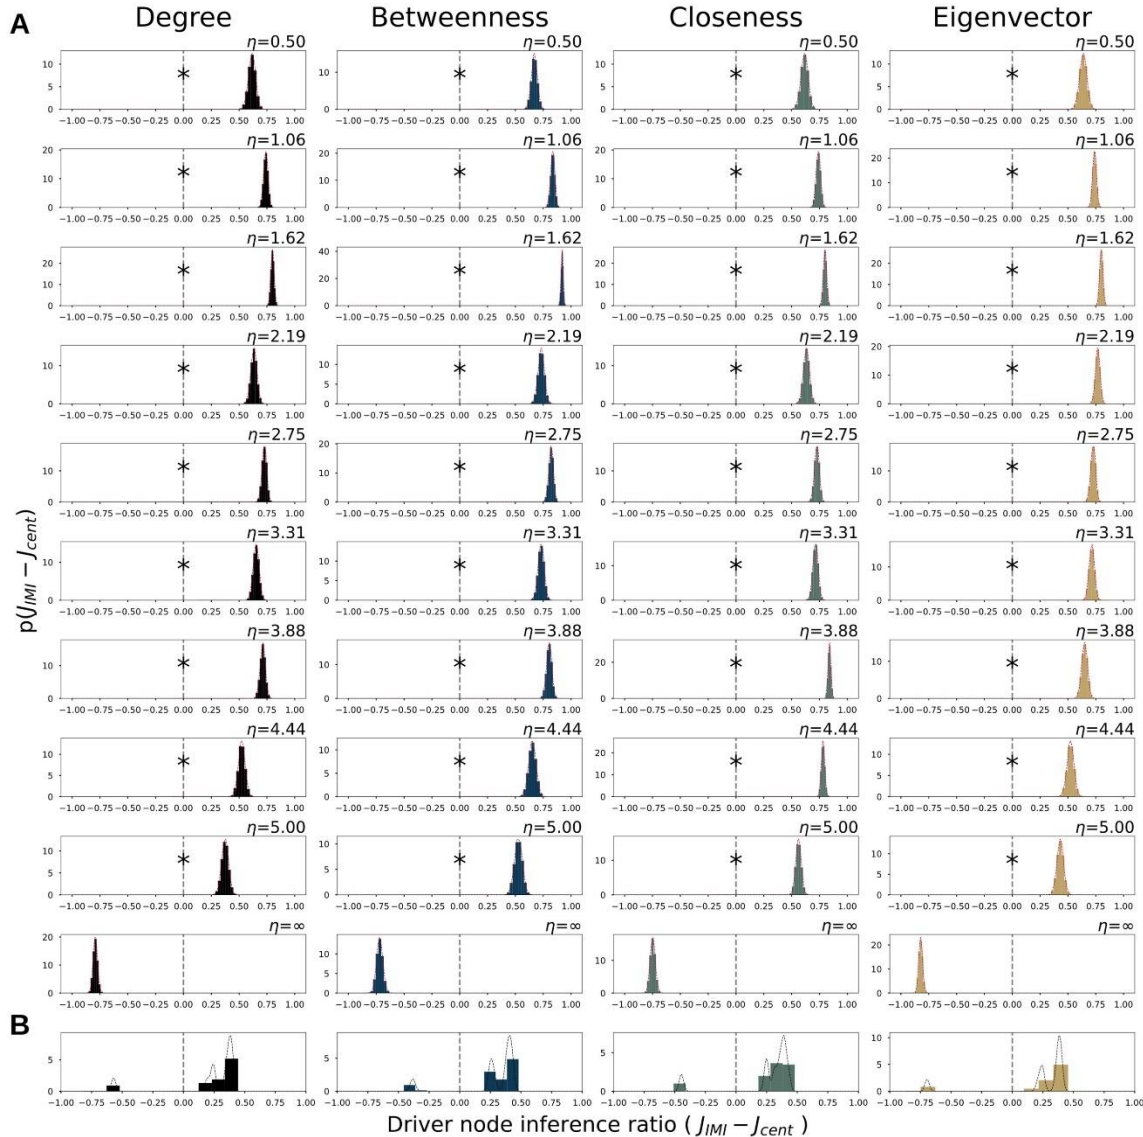
4.5.1. Random network results

Driver node inference accuracy is depicted in fig. 14 A for both IMI and the network centrality measures. Three crucial observations can be made from the Jaccard scores. First, for nearly all systems there exists an intervention size for which IMI obtains a Jaccard score of 1 (perfect true driver node inference), whereas this is not always the case for the centrality metrics, see e.g. system 3, 7, 15. Secondly, IMI is predictive nearly only for soft intervention sizes, i.e. intervention sizes smaller or of similar order of magnitude as the existing forces acting on nodes (i.e. their degree in the J interaction matrix). In contrast, centrality metrics are mainly predictive for hard interventions. The statistical results reflect these observations (fig. 14 A and B).

In figure 14 B average decay curves are shown for IMI (top) and different intervention sizes (middle) and hard interventions (bottom) for the systems 2, 5, 12. Their network structure is depicted in fig. 15 B (top). Comparing hard interventions (bottom) with a more moderate intervention strength (middle) in system 2, shows that the driver node can significantly differ depending on the intervention size applied to the system. The order of the causal importance is noticeably different for the minimal intervention η^* versus the strong intervention $\eta = \infty$ (middle and bottom plot in figure 15 B).

Figure 14.

Bootstrap results for driver node inference score per intervention size and for all interventions and noise levels(B). A kernel density is fitted for each distribution and integrated over the interval $[-1, 0]$. An inference ratio of 1 indicates that IMI is better at predicting the driver node than the structural metric and vice versa for a score of 0. Significant intervals are indicated by an asterisk * (). Except for the hard interventions () intervention size IMI is a better predictor than any of the network structural features.



4.5.2. Real-world network: psycho-symptoms results

The psychosymptoms system reveals similar results to the random networks (fig. 16). Namely, for medium to high noise level IMI yielded significantly higher Jaccard scores than centrality metrics for low causal intervention (fig. 16 C/D, $p \ll 0.01$), while not for hard interventions. In contrast, hard interventions yield different causal structures altogether which do not reflect non-intervened dynamics (fig. 16 A/B). For example, the true driver node under hard interventions identifies 'dislike' (medium and high noise) to be the driver node (fig 16 A). Whereas for low causal intervention 'sad' is identified as driver node (fig. 16 A). This implies that intervention itself has impact on what causal structure is observed and that the intervention can show

systemic behavior not present in the non-intervened system. The soft intervention of $\eta = 0.1$ was too low to provide proper resolution for identification of driver nodes in the low noise setting (fig. 16A). Hard intervention in low noise condition did provide a different driver node than the soft intervention. The grid-search for optimal η^* was insufficient for the psychosymptom network and should be investigated in future studies.

In addition, the change in driver node observed in figure ref:fig:psycho A highlights one major flaw in centrality metrics: they cannot account for a change in driver nodes due to a change in dynamics. The implicit assumption on dynamics that each centrality metric holds provides only one estimate per network structure. In contrast, IMI does not depend on what mechanisms generate system dynamics. Instead, it uses the distribution dictated by these system dynamics which match the true driver nodes.


Fried and colleagues postulated that 'lonely' was the gateway from which information spreads through the network, i.e. bereavement was embodied mainly by 'loneliness' which then percolated its effect to the other symptoms (doi.org/10.1037/abn0000028). Since the data was cross-sectional, the comparison with the results from this study relies on the assumption that binary dynamics are representative of the absence and presence of psychological symptoms. If correct, the results from this study give a causal perspective on the associative results from (doi.org/10.1037/abn0000028). e results from this study postulate that 'depr', 'lonely' and 'sad' have similar causal effect for moderate to high thermal noise.

It is important to emphasize that a quantification is given in terms of absolute effect size and not directed effects. This means that nudging for instance 'sleep' has some effect X on the psycho-symptom network, in what direction that effect is, or whether it has a positive or negative effect on the bereavement score / cognitive load of the patient is not clear, and should be the subject of future studies (see appendix 8.4 in <https://doi.org/10.1016/j.physa.2022.126889>).

Figure 15.

(A) Causal impact versus integrated mutual information for different noise levels and intervention sizes. The inferred driver node is indicated above each figure with a dot. For low noise level the numerical estimates were too noisy and driver estimation was inaccurate. For intermediate to high noise level, causal impact scale linearly with integrated mutual information for soft interventions. Causal driver node was wrongly predicted for hard interventions by integrated mutual information. Notice how the causal flows change as a function of intervention size (top versus bottom plot). (B) Driver node identification in real-world network of psychosymptoms for medium noise as a function network features (top 4 subplots) and integrated mutual information (middle), and true causal driver nodes per intervention size (bottom). Integrated mutual information correctly identified the driver node with soft interventions, but not for high interventions. Only betweenness centrality correctly identified the true driver node for high causal interventions. (C) Bootstrap distributions for driver node inference score as function of intervention size across noise levels. A score of 1 indicates that integrated mutual information identified the driver node correctly, but the structural metric didn't, and vice versa for a score of -1. In red kernel density estimates are indicated and integrated over $[-1, 0]$. The asterisk (*) indicates a significant better inference score for integrated mutual information ($\alpha < .01$). Integrated mutual information was a better predictor than any of the network features for soft interventions.

(D) Bootstrap distribution for driver node inference score as function of temperature and nudge size. The integrated mutual information was a significant better predictor across noise and intervention size ($\alpha < 0.05$ indicated by *).



Structural metrics can be considered as implicitly assuming a particular dynamic. For example, the eigenvector centrality has a clear analytical connection with linear dynamics, such as a simple diffusion process. Betweenness centrality on the other hand can be considered to assume that causation between nodes is transmitted mainly along the shortest paths between pairs of nodes (see appendix 8.6 in <https://doi.org/10.1016/j.physa.2022.126889>), such as in message or package routing. Finally, many more network centrality metrics have been developed based on specific dynamics, such as current and flow centrality measures. Thus, such network structural metrics could indeed be used, but a careful analysis of the dynamic equations is required to assess which centrality measure (if any) turns out appropriate.

The results from this study show that high degree nodes with Ising spin dynamics tend to actually *not* be a driver node for small interventions. In this model it can be understood since high degree nodes exhibit "frozen" behavior; their large(r) number of neighbors effectively sum up to a constant and strong force towards the majority state. Consequently, a constant (soft) intervention will be relatively ineffective with hubs. For non-intervened Ising dynamics, it makes intuitive sense that the dynamics are driven by the nodes which are neither "frozen" (hubs) nor are poorly connected (low-degree nodes), which is reflected by the soft

intervention setting. This insight is specific to the Ising system but illustrates that the information flow and causal flow of a system is not merely determined by the connectedness of a node in the network. Rather, the connectedness of the entire system in addition to inter-node dynamics combined are crucially important for causal flows and driver node identification.

The results further imply that for unperturbed dynamics structural metrics may not be predictive for determining the causal importance of nodes. Soft interventions revealed a different causal structure than hard interventions (fig. 15). For hard interventions, structural metrics do become predictive for causal importance for the networks studied here. However, the dynamics of these systems are shown to deviate from the non-intervened dynamics, causing system dynamics that are not representative of the non-intervened system. This can be seen in the causal influence in figure 14 and 15 where the causal influence for hard interventions is opposite to soft causal interventions. Consequently, if the aim is to provide understanding to the information flows for non-intervened dynamics, then hard interventions are not preferred.

In addition, the results imply that in order to achieve maximum impact for a fixed 'intervention' budget (injected energy), choosing the high degree nodes is not necessarily optimal. The adjusted Hamiltonian (introduces a bias on low degree nodes. Here, a causal intervention was performed by adding fixed energy to the Hamiltonian. For fixed intervention size $\eta = c$, the causal impact on the nodal distribution will be relatively higher for nodes with low degree than nodes with higher degree. Higher degree nodes may have higher causal effect in principle if the same probability mass is changed. However, moving the same probability mass scales non-linearly in kinetic Ising. For a limited 'intervention budget' it is preferred to locate those elements of the system that reach maximal causal effect.

For some systems, however, the inferred driver node never matched the true driver node(s) (e.g. system 8 in fig. 14). This could be due to two main reasons. First a grid-search was applied for the intervention size. It was argued that there would be a minimal intervention η^* which would lead to a measurable effect. It is possible that the parameter space used here missed the intervention size that provided enough resolution to accurately determine the driver node(s). Second, the numerical procedure for driver node inference could be optimized. The overlap of distribution was set to $\phi = 0.5$, to infer the driver node set \mathcal{A} , to prevent false positives for driver node identification due to noise in observed system states. However, the choice of parameter was not optimized and could lead to ambiguous driver node inference. System 8 in particular was most nodes in the network had a similar connectivity pattern, and as such causal isomorphy occurs, i.e. the causal importance of a node is indistinguishable from any other node in the system. The bootstrap estimates led to high level but not perfect of overlap (see appendix 8.7 in <https://doi.org/10.1016/j.physa.2022.126889>). Consequently, if ϕ was set differently, the inferred driver nodes could be improved. In future studies, we aim to further look into what how this numerical procedure could be improved for inferring driver nodes in

dynamical networks particularly for causal isomorphic nodes.

4.5.3. Limitations

The systems considered are discrete and ergodic. IMI assumes that the data-processing inequality holds for the system (section 4.3.7). The data-processing inequality in ergodic systems ensures that $I(s_i^{t_0+t}; S^{t_0})$ monotonically approaches zeros as $t \rightarrow \infty$ (see appendix 8.1 in <https://doi.org/10.1016/j.physa.2022.126889>). As a consequence, IMI will always be finite for ergodic systems. For non-ergodic systems, the data-processing inequality can only guarantee that $I(s_i^{t_0+t}; S^{t_0})$ never increases as a function of t . Namely, the data-processing inequality ensures that no local manipulation of information may increase the information content of a signal. This implies that as $t \rightarrow \infty$, IMI may not converge for nodes with non-zero baselines. For these non-ergodic systems, however, it may be possible determine the driver nodes by considering finite time-scales or subtracting the asymptotic γ value and reporting it separately.

Furthermore, this study only considered information flows in undirected systems. For directed systems, the largest integrated mutual information may in fact not be the driver node due to correlation. In appendix 8.2 (see <https://doi.org/10.1016/j.physa.2022.126889>) an integrator node is considered, where the integrator stores the system information but lacks any causal outgoing connections. The information storage capacity in this node will be higher than any other node due to its increased alphabet. If this integrator node encodes a copy operator, its information decay will result in the highest integrated mutual information while lacking any causal potential. In the systems considered here, this situation cannot occur. That is, for undirected systems information flow has equal potential along an edge. Any causal effect a node has in undirected dynamical systems will percolate to all neighbors of a node. For directed graphs the information flow is not time symmetric as detailed balance cannot be ensured. Consequently, the use of forward mutual information $I(s_i^{t_0+t}; S^{t_0})$ may lead to false estimates of driver nodes in these cases. For undirected graphs, if there would exist a node with higher integrated mutual information, its causal impact will also be higher and as such it would be the driver node. The asymmetry of information flow, and the compound nature of both correlation and causation restricts the use of integrated information undirected graphs.

In addition, the focus on short-time scales biases information processing to less constrained nodes. By measuring the system out-of-equilibrium, the system dynamics for the kinetic Ising model has the tendency to "freeze" the behavior of nodes as a function of increasing degree. For undirected systems, this assumption seems sensible as the higher causal potential for hubs forms at the same time a higher constraint for switching its state. The flip probability of node s_i in the kinetic Ising model as a function of degree will tend to zero given that the neighborhood is congruent with the node's current state. Measuring the system in equilibrium will therefore result in nodes with higher degree to have lower information decay. In contrast, out-of-equilibrium

a node state will have lower variability and consequently less causal effect on the instantaneous system states. Therefore, the driver node identification through integrated mutual information is only valid for non-equilibrium kinetic Ising spin dynamics with undirect network structures. Since the Ising spin dynamics falls within a larger universality class, such as (directed) percolation. This implies that the results from this study may apply to different dynamical systems sharing the same core assumptions.

Systems of size at most $n = 12$ were used. The size of the system was chosen due in order to provide high reliability of the probability distributions. In addition, larger graphs can be decomposed in various different network motifs (doi.org/10.1038/nrg2102). It is believed that these motifs form the "computational" building blocks for larger complex systems under the assumption of nearest neighbor interactions. That is, understanding the motifs would gain insights in how a macroscopic property emerges from local interactions. The aim of this study was to relate structural connectedness to dynamic importance; the exact nature or occurrence of motifs were not the focal point. In real-world systems, however, it is exactly the composition and interaction of these motifs that are vital to complex systems. Here, the motifs were implicit on the real-world network and the generated structures. We leave it up to future work to map out the driver nodes of common network motifs in different dynamics and relate the structural importance to the dynamical importance.

4.6. Conclusions

Our results indicate that dynamic importance cannot necessarily be reliably inferred from network structural features alone, demonstrated here using kinetic Ising spin dynamics. The goal of this paper was to show that structural methods can provide unreliable estimates of the driver node in dynamical systems. The results from this study show that the common assumption of structurally central or well-connected nodes being simultaneously dynamically most important is not necessarily true. This implies that we cannot abstract away the dynamics of a dynamic system before inferring driver nodes. The proposed information theoretic metric, integrated mutual information (IMI), was better able to identify the driver node for non-intervened dynamics in systems. Importantly, IMI does not require knowing the dynamics equations and/or the network structure of the system. It is instead calculated directly from a cross-section of time-series of the system without interventions. The proposed metric could potentially be useful in applications with rich data sets and where performing interventions are infeasible or impractical.

5. General discussion, conclusions and next steps

As demonstrated in this report, we must conclude that our initial working hypothesis could not be confirmed. That is, for now, it was not possible to establish a reliable method of predicting driver nodes as an inferential

tool to assess causal probability of biomarkers. However, we emphasize that even if no reliable method for inferring driver nodes would be found during the To_Aition project, this does not suggest that our efforts to construct networks were futile (e.g., as evidenced by the results of T2.1 to T2.3). The networks can still provide valuable insights for prediction tasks, provide insights into potential mediating or confounding pathways between morbidities, and provide network-level insights into plausible mechanisms (causal discovery, WP3). Our findings only imply that top driver nodes cannot be reliably inferred -- using the datasets that we have currently available.

We also found that, if the datasets would have been much richer in the temporal domain (rich longitudinal data with higher numbers of repeated measures), that it would then indeed be possible to infer driver nodes (Section 4). It remains to be seen what would be the *minimal* requirements of a dataset in order to be able to infer driver nodes, but from that study it became at least clear to us that the data requirements would well exceed the properties of our current datasets (especially in the temporal domain). That is, a handful of 'waves' was not sufficient (not shown), and it seems that the number of measurements over time in a time-series dataset would be at least on the order of 100 (which matches the self-reported data that Fried et al. used to infer the network of the psychometrics dataset), but preferably more.

Despite this negative result, we expect that causal discovery and causal inference techniques (WP3) will still be able to give mechanistic insights into comorbidities. For instance, while it may not be possible to infer 'the' driver node in a network, it might still be possible to identify a group of biomarkers that play a pivotal role in the pathways between morbidities. Within such an identified group of biomarkers we would perhaps not be able to tell which is more important than the others, but this is not needed if the goal is to give direction for further research and data gathering specifically focused on the identified biomarkers and the pathways that they may form. Finally, our results may give direction at a meta-level: namely, future data gathering efforts should be informed by algorithms such as ours to decide, for instance, if more effort is to be placed on including more subjects versus gathering rich time-series per subject, if the goal is to identify driver nodes.

Our studies indicate several potential paths for further exploration, calling for the development of new methodologies or additional data collection and harmonization:

Cross-sectional Data and Directionality. A key issue is the inherent limitation of cross-sectional data in revealing the directionality of causal relationships. For instance, even if we could discern a causal structure from cross-sectional data, it provides minimal information about the direction of these causal relationships. For example, with three variables, the same correlation network could result from multiple causal structures, such as confounding ($A \leftarrow B \rightarrow C$), colliding ($A \rightarrow B \leftarrow C$), or sequential ($A \rightarrow B \rightarrow C$ or the reverse). As the network size increases, the number of corresponding causal network structures also rises exponentially. This situation is problematic because each causal network structure causes different impacts of simulated



interventions on the nodes, thereby leading to different driver node(s).

Thus, one possible approach would be to enhance the datasets used for network construction and causal inference. Adding a temporal dimension could provide the most valuable enrichment, a strategy we're currently exploring. However, such strategies won't apply to the cross-sectional and longitudinal datasets used in TO_AITION, even if they effectively address the aforementioned inference limitations.

Centrality Metrics for Hypergraphs. The discovery of synergistic associations leads to hypergraphs (detailed in this deliverable in Section 3), where edges can have a higher order than the usual pairwise edges. Given that hypergraphs are theoretically richer than pairwise networks, we initially hypothesized they would more accurately predict driver nodes. However, research into centrality metrics for hypergraphs is less extensive than for pairwise networks, and techniques for the latter can't directly apply to the former. The existing centrality metrics we have explored often make assumptions based on mathematical convenience or elegance. For example, they may analyze uniform sub-hypergraphs independently instead of a single heterogeneous hypergraph (containing multiple edge orders, such as pairs, triplets, quadruplets, etc.) and then aggregate the vector of node centralities. This approach limits the model because causal pathways in reality can likely cross edges of multiple orders. Another drawback is that all existing centrality metrics regard hyperedges as undirected, ignoring the directional nature of causal relationships.

However, it's not straightforward to determine which edge orders should interact and under what conditions. For instance, the influence might depend on the overlap of two higher-order edges, or the type of multivariate association encoded by the edges. This challenge doesn't even consider the equally important problem of devising an efficient algorithm to compute the centrality metric. Discovering higher-order edges from data is already a demanding computational task, let alone computing paths of higher-order edges based on intricate conditions.