CAUSAL GRAPH RECOVERY IN NEUROIMAGING THROUGH ANSWER SET PROGRAMMING

Anonymous authors
Paper under double-blind review

ABSTRACT

Learning directed causal graphs from time-series data poses significant challenges, especially in fMRI where slow sampling rate obscures fast neural interactions. This temporal mismatch leads to undersampling, which can make multiple graphs equally plausible. We address this problem by explicitly modeling undersampling effects when recovering causal graphs. Our approach employs answer set programming (ASP) to enforce domain-specific constraints and optimize soft observational constraints, thereby identifying a Markov equivalence class for the resulting graph solutions. By customizing an ASP solver to collect multiple near-optimal solutions, we obtain not only the single best-fitting graph but an equivalence class of high-scoring graphs for expert consideration. This method, called Real-world noisy RASL (RnR), can also act as a meta-solver: it refines the output of other causal discovery algorithms by accounting for undersampling biases. In simulations and empirical brain network data, RnR produces more accurate causal graphs than state-of-the-art methods, improving F1-scores by an average of 12% by reducing false connections. We demonstrate that RnR is robust to varying undersampling rates – maintaining high precision and recall even as sampling becomes sparser – whereas competing methods degrade significantly. Our results suggest that incorporating undersampling-aware constraints via ASP yields more reliable and interpretable brain connectivity estimates from fMRI time series, bridging the gap between neural dynamics and observational data.

1 INTRODUCTION

Causal inference from functional Magnetic Resonance Imaging (fMRI) data has emerged as a critical endeavor to understand the neural mechanisms underlying cognitive processes and behaviors. Researchers not only seek to identify active brain regions during tasks, but also unravel the causal relationships between these regions, often referred to as "effective connectivity" (Friston, 1994). Graphical causal models, such as causal Bayesian networks, have become a popular framework for this purpose, combining directed graphs with joint probability distributions to model the dependencies between different brain regions (Pearl, 2009). These models adhere to the Causal Markov Condition, which asserts that each node in a causal graph is conditionally independent of its non-descendants given its parents (Spirtes et al., 2001).

However, applying these models to fMRI data is fraught with challenges, particularly due to the mismatch between the temporal resolution of fMRI and the rapid timescale of neural processes. Typical sampling intervals in fMRI, ranging from one to three seconds, are much slower than the millisecond-level interactions between neurons, leading to significant undersampling (Valdes-Sosa et al., 2011). This undersampling often results in multiple causal graphs being statistically indistinguishable given observed data, forming what is known as a Markov Equivalence Class (Pearl, 2009; Spirtes et al., 2001). These problems are exacerbated by the indirect nature of the *Blood-Oxygen-Level Dependent* (BOLD)[1] signal, which reflects neural activity through complex and variable hemodynamic responses (Handwerker et al., 2004).

[1]BOLD signal is the measurement obtained from fMRI that reflects changes in blood oxygenation related to neural activity. When neurons fire, blood flow to that region changes; fMRI captures these changes at a slow timescale (on the order of 1–2 seconds per sample). Importantly, the BOLD signal is an indirect and delayed proxy of neural firing.

In addition, the inherent variability in the hemodynamic response across different brain regions and subjects, adds another layer of complexity. Variations in the time-to-peak of the BOLD response can lead to incorrect inferences about the direction of causality, particularly when using methods like Granger causality, which assumes a fixed temporal relationship between cause and effect (David et al., 2008; Seth et al., 2013). Although some studies suggest that Granger causality may be robust to certain variations in the hemodynamic response (Seth et al., 2013), the combination of measurement noise, undersampling, and hemodynamic variability often undermines the reliability of causal inferences from fMRI data.

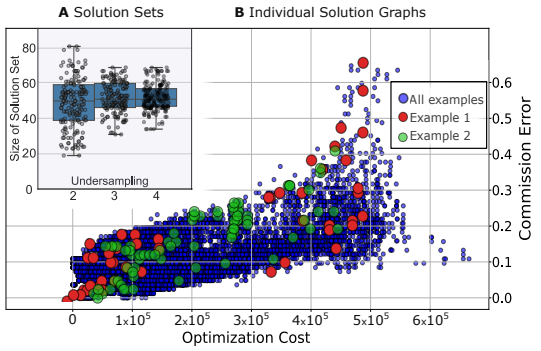

Figure 1: Top left(A): Size of optimization solution set across different undersamplings, repeated 100 times. Bottom right(B): Commission error of the solution vs. optimization cost for that solution. Solutions in one equivalent class are highlighted in red.

In response to these challenges, this paper proposes a novel approach called **RnR** that explicitly accounts for the effects of undersampling in the derivation of causal graphs. By employing constraint optimization through Answer Set Programming (ASP), we aim to identify the most probable causal graph from a set of potential candidates. ASP allows for the incorporation of domain-specific knowledge and constraints, facilitating the identification of not only a single graph but an equivalence class of possible graphs, thereby offering a more comprehensive understanding of the underlying causal structure (Gebser et al., 2012). We validate our approach using both simulated data and real fMRI data, demonstrating its superiority over existing methods . Our results suggest that ASP offers a powerful new tool for causal inference in neuroimaging, providing more accurate and intuitive insights into the brain's functional architecture.

2 Background

We study recovery of the true causal graph $\mathcal{G}^1$ from an observed, undersampled measurement graph $\mathcal{H}$ when the sampling rate u is unknown. We first fix notation and core facts, then state the limitation that motivates our approach.

A directed dynamic causal model extends standard causal models (Pearl et al., 2000; Spirtes et al., 1993) by adding time: variables $V_{1:n}$ appear at times $t, t-1, \ldots$, and a first-order Markov structure is assumed so that $\mathbf{V}^t \perp\!\!\!\perp \mathbf{V}^{t-k} \mid \mathbf{V}^{t-1}$ for $k > 1$ (Spirtes et al., 2000). When sampling is slow relative to neural dynamics, fMRI produces an undersampled graph $\mathcal{G}^u$ that can differ from $\mathcal{G}^1$ (Danks & Plis, 2013; Gong et al., 2015). In the compressed representation, an edge $i \to j$ in $\mathcal{G}^u$ occurs iff there is a directed path of length u from i to j in $\mathcal{G}^1$; a bidirected edge $i \leftrightarrow j$ occurs if a common ancestor reaches both with equal path length $< u$ (Danks & Plis, 2013). This view clarifies which dependencies are artifacts of rate.

To address the structural implications of undersampling, several approaches have been developed. For instance, Danks & Plis (2013) explored the problem structurally, while Gong et al. (2015) provided a parametric approach for two-variable systems.

Dealing with latent confounders, such as those introduced by undersampling, has led to the development of various graph representations, including Partially-Observed Ancestral Graphs (PAGs) (Zhang, 2008) and Maximal Ancestral Graphs (MAGs) (Richardson & Spirtes, 2002). However, these frameworks often struggle with the complexities introduced by undersampling (Mooij & Claassen, 2020). Compressed graphs are effective for undersampling: they are 1–1 equivalent to dynamic ADMGs and allow computing rate effects by path length, and they outperform PAG/MAG-style formulations under unknown u (Plis et al., 2015; Abavisani et al., 2023).

The Rate-Agnostic Structure Learning (RASL) algorithm (Plis et al., 2015) addressed the challenge of causal inference from undersampled data by directly tackling undersampling while adopting a rate-agnostic approach that avoids assumptions about the undersampling rate. RASL was a big leap

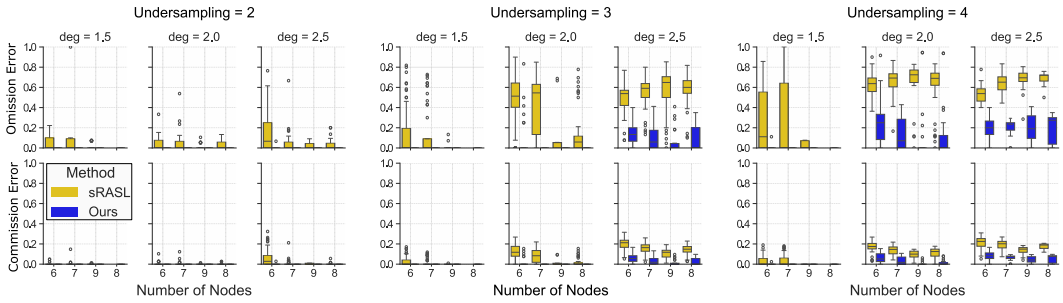

Figure 2: Normalized omission (top) and commission (bottom) errors for edge-breaking experiments with varying undersampling rates and graph densities, comparing the original approach with our improved sRASL-based method.

in computational efficiency. It cleverly checked all possible rates and used "stopping rules" to keep from re-exploring the same parts of a graph. Its main drawback, however, was that it didn't scale well to larger graphs, leaving room for improvement. Building on this foundation, newer approaches to causal learning have been developed that are far more scalable and efficient.

Building on these foundations, recent advancements in causal structure learning have led to the development of generalized rate-agnostic approaches that significantly enhance the scalability and efficiency of these methods. The reformulation of RASL into a constraint satisfaction framework, expressed using a declarative language, has shown considerable promise by enabling more efficient analysis of large graphs called Solver-based RASL (sRASL) (Abavisani et al., 2023). By incorporating constraints from strongly connected components (SCCs)[2] this approach offers a scalable and accurate method for causal inference from undersampled data. It can analyze graphs with over 100 nodes—a significant improvement on earlier methods that struggled with much smaller graphs. The sRASL method improves solving time 1000-fold while maintaining the same theoretical guarantees as its predecessors. However, sRASL shows limitations on real-world datasets, as the performance gains seen in simulations do not fully translate to noisy, practical conditions.

In Sanchez-Romero et al. (2019b), several methods were tested on synthetic BOLD data with feedback (e.g., Granger regression, MVAR, FASK, Two-Step). Many reached over 80% orientation precision and recall, including for 2-cycles. That study, however, did not consider that BOLD is undersampled. Here we revisit those algorithms under sparse sampling and compare them with an ASP-based approach (RnR).

The method space has grown (e.g., LiNGAM (Shimizu et al., 2006), EC-GAN (Kim et al., 2021), amortized transformers (Paul et al., 2022), and RL-style models (Mnih et al., 2016; Salehi et al., 2021; Pamfil et al., 2020)). For clean comparison and interpretability, we benchmark the widely used set from Sanchez-Romero et al. (2019b): GIMME (Gates & Molenaar, 2010), MVGC (Barnett & Seth, 2009), MVAR (Bressler & Seth, 2003), and FASK (Sanchez-Romero et al., 2019a). These cover group and individual modeling, multivariate dependencies, and non-Gaussian orientation cues.

We then evaluate all methods on undersampled data and against ASP. This isolates robustness to undersampling. Among baselines, FASK performs well; RnR further improves accuracy by making undersampling explicit. Detailed results appear in Section 4.2.

3 METHODS

In this section, we present Real-world noisy RASL (RnR)[3], our enhanced structure learning framework, tailored for fMRI connectivity analysis. Our method, RnR, builds upon the sRASL framework

[2]Subsets of nodes in a directed graph where each node is reachable from every other node in the subset. In other words, an SCC is a "maximal loop" or cluster of nodes with mutual reachability. In our method, we allow nodes within an SCC to be interlinked (modeling instantaneous or cyclic relationships), but we enforce that the graph of SCCs (each SCC treated as a single super-node) has no directed cycles, making it a DAG.

[3]A note on naming: "RASL" stands for Rate-Agnostic Structure Learning, an approach introduced by Plis et al. (2015) for causal discovery without knowing the true sampling rate. sRASL for structural RASL used ASP to find a single optimal causal graph consistent with undersampled data Abavisani et al. (2023). Our

through several key innovations tailored for noisy, undersampled fMRI data. These contributions include: (1) utilizing custom made Answer Set Programming solver to retrieve all plausible, near-optimal graphs within a cost threshold instead of a single solution; (2) constraining the optimization with realistic edge density limits to ensure the resulting graphs are biologically plausible; (3) acting as a meta-solver to refine the outputs of other causal discovery algorithms by systematically handling undersampling effects; (4) employing a prioritized, multi-stage process that sequentially matches graph density, resolves links, and orients edges; and (5) applying an adaptive weighting scheme to the cost function to discourage spurious connections while retaining high-confidence ones.

3.1 Overview of the sRASL Framework

The original sRASL framework (Abavisani et al., 2023) is an ASP-based approach that takes an initial directed graph $\mathcal{H}$ and applies structural constraints to refine it. We follow Abavisani et al. (2023) and use Clingo[4] for our implementation. sRASL allows cycles within strongly connected components (SCCs), but requires the SCCs themselves to form a Directed Acyclic Graph (DAG). This DAG-over-SCCs structure prevents latent-level causal cycles. The original sRASL only returned a single optimal solution minimizing a cost function, but this posed limitations. Due to inherent uncertainty in the case of undersampling, the solution with the least cost is not necessarily the correct answer. Therefore, ignoring near-optimal alternatives in noisy data can be misleading.

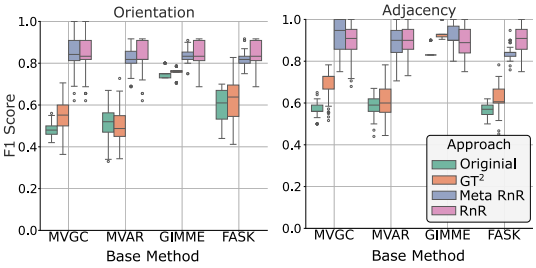

Figure 3: Performance comparison using Sanchez-Romero's data with and without the RASL meta-solver. Applying RASL on top of traditional methods improves accuracy by accounting for undersampling effects, with additional gains achieved by using PCMCI.

3.2 Optimization for Multiple Solutions

To address the limitation of single-solution output, we modified Clingo to return all graphs G satisfying:

$$C(G) \leq C(G^*) + \Delta, \tag{1}$$

where $C(G)$ is the cost of graph G, G^* is the optimal graph, and Δ is a small tolerance. We define the cost function as:

$$C(G) = \sum_{e \in H} w_e^+ \cdot \mathbb{I}[e \notin G] + \sum_{e \notin H} w_e^- \cdot \mathbb{I}[e \in G], \tag{2}$$

where $\mathcal{H}$ is the measurement graph (input), w_e^+ penalizes removal of edges in $\mathcal{H}$, w_e^- penalizes addition of edges not in $\mathcal{H}$, and $\mathbb{I}[\cdot]$ is the indicator function. This yields a manageable set of near-optimal solutions for downstream selection and interpretation.

3.3 Density Constraint for Realistic Graphs

Neuroscience knowledge can inform the expected density of functional brain networks. Empirical studies of brain connectivity (both functional and structural) show that brain graphs tend to be neither nearly empty nor fully connected; they often have a moderate density (for example, on the order of 10–30% of all possible edges, depending on node definition (Smith et al., 2011). After running sRASL, we observed that some solutions were unrealistically sparse or dense, likely because the

RnR extends this by considering real-world noise (hence "Real-world noisy RASL") and retrieving multiple solutions.

[4]Clingo is an open-source solver for Answer Set Programming (ASP) that finds solutions (called "answer sets") to a given set of logical rules and constraints. In our context, Clingo explores different possible graphs and identifies those that satisfy all our constraints while optimizing a cost function.

optimization only cared about matching $\mathcal{H}$ and not an absolute scale of connectivity. To correct this, we introduce an explicit density constraint in our ASP formulation:

```
1  countedge1(C):- C = #count { edge1(X, Y): edge1(X, Y), node(X),
     node(Y)}.
2  countfull(C):- C = n*n.
3  hypoth_density(D) :- D = 1000*X/Y,  countfull(Y), countedge1(X).
4  abs_diff(Diff) :- hypoth_density(D), Diff = |D - d|.
5  :~ abs_diff(Diff). [Diff@2]
```

Line $1-3$ calculates graph density in the present hypothetical solution, and lines $4, 5$ measure the difference from ground truth and penalize when value `Diff`> 0. This prevents trivial empty or all-to-all graphs and ensures structural realism based on known brain network densities (which are reasonably well-estimated and biologically bound).

3.4 Prioritized Optimization by Connection Type

To efficiently recover causal graphs under severe noise and undersampling, we adopt a prioritized multi-stage optimization strategy that decomposes the total cost into structurally meaningful components. Our aim is to minimize the overall cost as formulated in Equation 2 to guide acceptance of approximate solutions satisfying Equation 1. We decompose the total cost as:

$$C(G) = C_d(G) + C_b(G) + C_o(G), \tag{3}$$

corresponding respectively to penalties on density (edge count), bidirected structure, and edge orientation. More specifically, we use the following stepwise cost minimizations:

$$\min_{G \in \mathcal{G}} \quad C_d(G) = \lambda_d \cdot |\text{density}(G) - \text{density}(G^*)| \tag{4}$$
$$\text{Feasible set } \mathcal{F}_1: \quad \mathcal{F}_1 = \{G \in \mathcal{G} \mid C_d(G) \leq \epsilon_1\}$$

$$\min_{G \in \mathcal{F}_1} \quad C_b(G) = \lambda_b \cdot \#\{\text{bidirected mismatch to } G^*\} \tag{5}$$
$$\text{Feasible set } \mathcal{F}_2: \quad \mathcal{F}_2 = \{G \in \mathcal{F}_1 \mid C_b(G) \leq \epsilon_2\}$$

$$\min_{G \in \mathcal{F}_1, \mathcal{F}_2} \quad C_d(G) = \lambda_d \cdot \#\{\text{directed mismatch to } G^*\} \tag{6}$$

where G denote a candidate causal graph, $\mathcal{G}$ the space of all possible graphs, G^* the ground-truth. ϵ_1 and ϵ_2 are tolerance parameters. Priority weights $\lambda_d \gg \lambda_b \gg \lambda_o$ reflect the order of importance. This approach constrains the search space in a lexicographic manner: first matching global sparsity, then resolving stable structural backbones, and finally fine-tuning edge directions.

3.5 Adaptive Weighting Scheme for Edges

The weights w_e^+ and w_e^- in the cost function (Equation 2) critically influence which edges the solver prefers to keep or drop. We devise an adaptive weighting scheme based on the intuition that not all edges in the initial graph $\mathcal{H}$ are equally reliable, and similarly not all potential missing edges are equally implausible. Specifically, we leverage the strength of connection evidence (e.g., correlation or a statistical score) for each edge to set these weights: For each edge e that is present in the initial graph $\mathcal{H}$ (meaning the first-order method or data proposed this connection), we assign w_e^+ inversely proportional to the confidence or strength of e. For example, if e corresponds to a high correlation or a strong Granger causality score between two regions, we give it a high weight—the ASP solver will incur a large cost for removing this edge, so it will likely keep it unless necessary. Conversely, if e was only weakly supported (perhaps a borderline significant link), we set w_e^+ lower, indicating that dropping this edge is not heavily penalized. This adaptivity encodes our uncertainty: strong edges should be trusted more (harder to delete), weak edges can be pruned if they conflict with other constraints.

For each potential edge e that is absent in $\mathcal{H}$ (the initial method found no link), we set w_e^- to a maximum penalty by default, reflecting a strong bias against adding completely new edges unless absolutely necessary. In other words, we assume $\mathcal{H}$ is mostly correct in saying those edges are absent, so the solver will avoid introducing a new connection e because it would incur a high cost. However, there is one exception: if we have domain knowledge or secondary evidence that a certain missing edge might actually exist (e.g., known anatomical connection between those regions), we could lower w_e^- for that edge. In general, though, w_e^- is large to favor sparse solutions that do not introduce unsupported edges.

This weighting strategy was implemented by extracting pairwise correlation magnitudes from the time-series data for all region pairs. For instance, in our fMRI simulations, we computed the Pearson correlation for each pair of nodes; if an edge was present in $\mathcal{H}$, we mapped its correlation (or other score) to a weight in [0,1] range and then scaled it (the exact mapping can be provided in the Appendix). Edges not in $\mathcal{H}$ were given a uniform high weight (effectively a large constant). The ASP solver thus solves a weighted min-change problem, where changes to highly credible edges are very costly, and changes to dubious edges are cheaper.

This adaptive weighting greatly improved the accuracy of the recovered graphs. It naturally handles the uncertainty in fMRI: for example, if two regions had a very low measured correlation yet some method erroneously connected them, RnR will likely remove that edge (low penalty to drop) in favor of satisfying structural constraints. On the other hand, if two regions are strongly correlated but the initial method missed the connection (say due to a too-conservative threshold), RnR will still be unlikely to add it because of the high default cost for new edges. However, if adding that edge is crucial to satisfy the DAG or density constraint, then the solver might do so, incurring the cost but yielding a more consistent graph. In summary, adaptive weights allow RnR to "decide" which edges to trust and which to doubt, rather than treating all edges uniformly.

3.6 RnR as a Meta-Solver for Undersampling Effects

A significant challenge in causal learning from fMRI is handling unobserved common causes (latent confounders) that arise from undersampling. For example, if two brain regions A and B are both driven by an unmeasured region or stimulus (common cause) that was not captured due to slow sampling, standard algorithms might either link A–B with directed edges in a loop or miss the connection altogether. Bidirected edges (often denoted $A \leftrightarrow B$) can represent unresolved causality due to a hidden common cause in causal graphs. Some time-series causal discovery methods (e.g., Granger causality (Granger, 1969), VAR models (Lütkepohl, 2005), and the PCMCI algorithm (Runge et al., 2019)) are capable of producing bidirected or contemporaneous links. However, many algorithms that ignore undersampling (e.g., standard PC or GES variants, Granger causality in its basic form, etc.) will not output bidirected edges at all, potentially misrepresenting the true connectivity.

We address this issue by using RnR as a meta-solver: RnR takes as input any initial graph $\mathcal{H}$ produced by a first-order method, but then refines it. If the initial method did not account for undersampling, $\mathcal{H}$ might contain spurious patterns (e.g., a two-node cycle $A \to B$ and $B \to A$) or be missing bidirectional links. We therefore augment $\mathcal{H}$ before running ASP: for any pair of nodes that form a mutual directed cycle ($A \to B$ and $A \leftarrow B$) in $\mathcal{H}$, we interpret this as evidence of a possible latent confounder. We then modify $\mathcal{H}$ by adding a bidirected edge $A \leftrightarrow B$. Additionally, to encode uncertainty about direction of influence, we include both directed edges $A \to B$ and $A \leftarrow B$ in $\mathcal{H}$ (if not already present), but assign them a low weight w^- (penalty) for inclusion.

The ASP solver thus starts with an enriched graph containing a bidirectional link and two low-confidence directed links between A and B. The solver can decide, through optimization, whether the final solution should keep the bidirected connection, one of the directed connections, or both, depending on what best satisfies all constraints and minimizes cost. Essentially, we are giving the solver the flexibility to explain an observed correlation between A and B as a latent confounding (bidirected edge), as a direct causal relationship (one of the directed edges), or as a mix of both, whichever is most consistent with the data and constraints. By using RnR as a post-processor in this way, we "inject" undersampling awareness into any algorithm's output. In Section 4.2, we demonstrate this meta-solver capability by applying RnR to outputs from the method of Sanchez-Romero et al. (2019a) and from PCMCI, yielding improved results.

3.7 Implementation and Practical Considerations

ASP Implementation: We developed the RnR ASP model using Clingo's input language. To summarize, the program encodes: (1) Facts: all possible directed edges between nodes (except self-connections) are initially considered. The input graph $\mathcal{H}$ is given as facts (e.g., `edge(A,B)` if $A \to B$ in $\mathcal{H}$). (2) Cost rules: rules that contribute to cost for including or excluding edges (as per Equation 2, using Clingo's `#minimize` or `:~` syntax with weights w^+ and w^-). (3) Constraints: (3.a) No directed cycle among SCCs (implemented via a reachability check: we introduce a predicate `reach(X,Y)` that is true if there is a path, and add a constraint `:- reach(X,X).` to forbid cycles); (3.b) Graph density bounds as described in Section 3.4; (3.c) any known prior knowledge (e.g., we add `:- not edge(A,B)` if we know $A \to B$ must exist).

We provide the full ASP code in Supplementary Material, and a simplified excerpt is shown in Figure 2 of the Appendix for illustration.

Computational Efficiency: Running Clingo with optN on our problem sizes (graphs of up to 10 nodes in our simulations) was fast (under a few seconds) for each instance. For larger graphs (e.g., 20-50 nodes), ASP solving could become slower. However, our modular constraints (SCC-based DAG) significantly reduce complexity by disallowing many cyclic possibilities upfront. Further optimization or heuristic solving is a direction for future work, but for the scope of our experiments, the approach was tractable.

Validation of Multiple Solutions: Importantly, retrieving multiple solutions allowed us to examine the stability of the recovered graph. We found that often many of the top solutions shared most edges in common, differing only in a few uncertain connections. This gave us confidence that our method identifies a core reliable structure plus a small set of alternative edges. We illustrate this in Figure 1A, which plots the number of solutions vs. undersampling rate: even at high undersampling, the solution set remains of modest size (tens of graphs), and in Figure 1B we show the relationship between solution cost and edge errors. Notably, we observed a positive correlation between optimization cost and edge commission error (false positive rate) across solutions. However, some higher-cost solutions had low commission error, indicating that near-optimal graphs can sometimes be almost as accurate as the optimal one. This justifies examining a small pool of top solutions. In practice, we found that the best solution (by ground-truth error) was often among the top 10 by cost. Therefore, our strategy is to take the top-N solutions from Clingo (we used $N = 10$) and then select the final graph by an additional criterion (e.g., best fit on a validation dataset or expert judgment). This approach yielded robust identification /of the true causal graph in our tests.

4 Results

4.1 Improved Edge-Breaking Experiment

We replicated and extended the edge-breaking experiment from Abavisani et al. (2023) to demonstrate the robustness of our approach on graphs with intentionally broken edges. In this experiment, we generated causal ground truth graphs $\mathcal{G}^1$ and undersampled them at various rates, then simulated noise by randomly deleting an edge. The goal was to test the ability of our method to recover the true graph structure, compared to the original sRASL. The results, depicted in Figure 2, show that our method consistently achieved lower omission and commission errors[5] compared to the original approach, even under high undersampling conditions. This illustrates that our improved algorithm can more effectively recover the true graph structure, including robustness against edge-breaking perturbations.

4.2 BOLD simulation data: Enhancing Causal Inference with RASL as a Meta-Solver

We further evaluated RnR on simulated data from Sanchez-Romero et al. (2019b). These data have been widely used within the neuroscience community, but were generated with small, simple

[5] A commission error is a false positive (i.e., incorrectly added edge). An omission error is a false negative (i.e., incorrectly omitted edge). We report these separately to characterize different types of errors (adding spurious links vs. missing real links).

graphs lacking real loops and using an idiosyncratic data generation method. We demonstrate that the methods that they study, including ones developed for these data (e.g., (Sanchez-Romero et al., 2019a)), do not account for the effects of undersampling, which can bias causal inference.

Our approach applied RnR as a meta-solver: after running Sanchez-Romero's original methods, we applied RnR to the resulting graph to explicitly incorporate the effects of undersampling. This adjustment led to improved accuracy, as RnR optimized the causal graph structure to better reflect the underlying dynamics. In essence, RnR served as an enhancement layer, correcting for undersampling effects ignored by previous methods (Figure 3).

Additionally, we explored the use of PCMCI as an alternative to standard methods like SVAR and Granger Causality (GC). Cook et al. (2017) had previously demonstrated that SVAR and GC perform well for scenarios with isochronal bidirected edges, which arise due to undersampling, as well as directed edges. However, we observed that PCMCI, initially introduced by Moneta et al. and later significantly improved by Runge, performs better in these scenarios. Therefore, we incorporated PCMCI with RASL, achieving more accurate results by effectively handling undersampling and improving causal inference. Figure 3 illustrates these improvements.

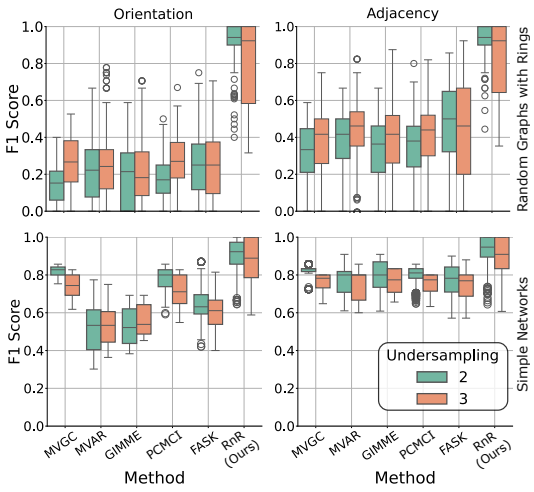

Figure 4: Comparison of Sanchez-Romero's simple data generation approach with larger, more complex VAR-generated graphs. Sanchez-Romero's data lacks real loops and is limited in scope, which may affect the generalizability of causal inference results.

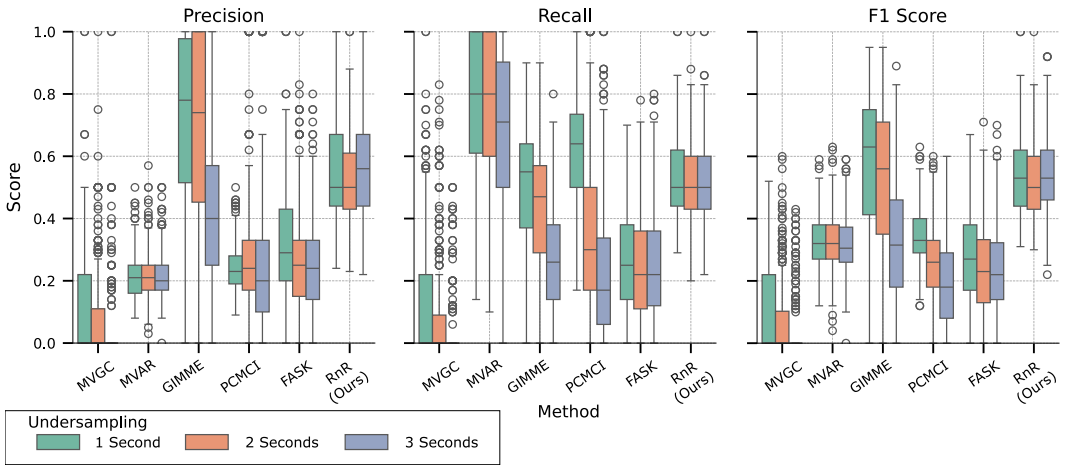

Figure 5: Impact of different undersampling rates (1s, 2s, 3s) on BOLD signal preservation in fMRI data simulated with the balloon model. Minimal error is observed with 1-second undersampling, whereas larger intervals degrade accuracy in all methods that don't account for undersampling effect. Our method RnR accounts for this effect and does not suffer loss from undersampling.

4.3 Analyzing the Limitations of Sanchez-Romero's Data Generation

Despite its widespread acceptance, Sanchez-Romero's data generation approach is limited in scope and complexity. It employs simple, small graphs without real loops and is generated using a specific, contrived process. To demonstrate the limitations of this dataset for broader causal inference applications, we conducted experiments with larger VAR models on graphs with many variables and more complex structures, showcasing the limitations of Sanchez-Romero's setup (Figure 4).

This underscores the importance of considering more realistic and complex data when testing causal inference methods.

4.4 Impact of Undersampling on BOLD Signal in fmri Data

We also examined the effect of undersampling on time series data generated from a VAR model and then processed through a balloon model to simulate the BOLD response in fMRI data (Buxton et al., 1998). Given the smooth nature of the BOLD signal, undersampling by one-second intervals often leads to minimal information loss, while undersampling by larger intervals (e.g., two or three seconds) introduces significant inaccuracies. Figure 5 illustrates how different undersampling rates impact the preservation of temporal information in BOLD data, underscoring that aggressive undersampling can obscure meaningful connectivity patterns.

5 Conclusions and Future Directions

Undersampling is a critical yet often overlooked issue in the analysis of time series data, particularly in the context of fMRI. In this paper, we addressed the undersampling problem directly by incorporating its effects into the derivation of causal graphs using ASP. Our approach goes beyond traditional methods by explicitly accounting for the temporal disconnect between neural processes and fMRI sampling rates. By doing so, we not only identified the most probable causal graph but also provided an equivalence class of potential graphs, enabling a more nuanced understanding of the underlying causal structures.

Our results, validated on simulated fMRI data underscore the importance of addressing undersampling in neuroimaging studies. We demonstrated that our ASP-based method outperforms existing techniques, particularly in scenarios where undersampling distorts the true causal relationships between neural groups. Additionally, by comparing our approach to other algorithms, including those studied by Sanchez-Romero et al. (2019b), we highlighted that while methods like FASK show potential, they still fall short in fully capturing the complexities introduced by undersampling.

The undersampling issue is not merely a technical challenge but a fundamental barrier to accurate causal inference in neuroscience. Ignoring it risks drawing erroneous conclusions about brain function and connectivity. Our work represents a significant advancement in this area, providing a robust and scalable solution that can improve the accuracy of causal inference in the presence of undersampling.

As the scientific community continues to explore the neural mechanisms underlying cognition and behavior, it is imperative that the undersampling problem is given the attention it deserves. Addressing this issue will be crucial in ensuring that the inferences drawn from fMRI data truly reflect the underlying neural dynamics, leading to more reliable and meaningful insights into brain function.

One challenge for solving the optimization problem with ASP is ensuring a reasonable initial estimated graph $\mathcal{H}$. The estimation errors at the measurement time-scale may inflate the estimation errors at the causal timescale. However, simply selecting the estimator that minimizes the errors in $\mathcal{H}$, as we have done in this paper, may not be the optimal strategy. Not all errors in $\mathcal{H}$ have the same effect on the quality of estimation and developing methods that consider that interplay is a promising future direction.

Further optimization of these approaches to enable work with larger graphs may open potential new domains where our methods may be applied. Although RnR is highly capable for working with reasonably sized graphs, extending the number of nodes by an order of magnitude could broaden the range of potential applications. This includes further practical application of the methods in the study of brain function via fMRI as well as other dynamic modalities.

References

Mohammadsajad Abavisani, David Danks, and Sergey Plis. Grace-c: Generalized rate agnostic causal estimation via constraints. In *The Eleventh International Conference on Learning Representations*, 2023.

Lionel Barnett and Anil K Seth. Granger causality and transfer entropy are equivalent for gaussian variables. *Physical Review Letters*, 103(23):238701, 2009.

Steven L Bressler and Anil K Seth. Multivariate autoregressive modeling of fmri time series. *NeuroImage*, 19(2):361–375, 2003.

Richard B Buxton, Eric C Wong, and Lawrence R Frank. Dynamics of blood flow and oxygenation changes during brain activation: the balloon model. *Magnetic resonance in medicine*, 39(6): 855–864, 1998.

John W Cook, David Danks, and Sergey M Plis. Learning dynamic structure from undersampled data. In *Proceedings of the UAI Causality Workshop*, 2017.

David Danks and Sergey Plis. Learning causal structure from undersampled time series. In *NIPS Workshop on Causality*, volume 1, pp. 1–10, 2013.

Olivier David, Isabelle Guillemain, Sandrine Saillet, Sebastien Reyt, Colin Deransart, Christoph Segebarth, and Antoine Depaulis. Identifying neural drivers with functional mri: an electrophysiological validation. *PLoS biology*, 6(12):e315, 2008.

Karl J Friston. Functional and effective connectivity in neuroimaging: a synthesis. *Human brain mapping*, 2(1-2):56–78, 1994.

Kathleen M Gates and Peter C Molenaar. Group iterative multiple model estimation (gimme) for group-level structural equation modeling of fmri data. *NeuroImage*, 42(2):416–426, 2010.

Martin Gebser, Benjamin Kaufmann, and Torsten Schaub. Conflict-driven answer set solving: From theory to practice. *Artificial Intelligence*, 187:52–89, 2012.

Mingming Gong, Kun Zhang, Bernhard Schoelkopf, Dacheng Tao, and Philipp Geiger. Discovering temporal causal relations from subsampled data. In *International Conference on Machine Learning*, pp. 1898–1906. PMLR, 2015.

Clive WJ Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society*, pp. 424–438, 1969.

Daniel A Handwerker, John M Ollinger, and Mark D'Esposito. Variation of bold hemodynamic responses across subjects and brain regions and their effects on statistical analyses. *Neuroimage*, 21(4):1639–1651, 2004.

Youngjun Kim, Juneyoung Lee, Woo-Seok Kim, Kangwook Lee, and Edward J Lee. Ec-gan: Estimating effective connectivity via generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14131–14141, 2021.

Helmut Lütkepohl. *New introduction to multiple time series analysis*. Springer Science & Business Media, 2005.

Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. *International Conference on Machine Learning*, pp. 1928–1937, 2016.

Joris M Mooij and Tom Claassen. Constraint-based causal discovery using partial ancestral graphs in the presence of cycles. In *Conference on Uncertainty in Artificial Intelligence*, pp. 1159–1168. PMLR, 2020.

Roland Pamfil, Stefan Bauer, Natalia Díaz-Rodríguez, Bernhard Schölkopf, and Jonas Peters. Dynotears: Structure learning from time-series data. *International Conference on Learning Representations*, 2020.

Subhrajit Paul, Prathosh AP Roy, and Deepali Jha. Amortization transformer for brain effective connectivity estimation from fmri data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 17815–17824, 2022.

Judea Pearl. *Causality*. Cambridge university press, 2009.

Judea Pearl et al. Models, reasoning and inference. *Cambridge, UK: CambridgeUniversityPress*, 19:2, 2000.

Sergey Plis, David Danks, Cynthia Freeman, and Vince Calhoun. Rate-agnostic (causal) structure learning. *Advances in neural information processing systems*, 28, 2015.

Thomas Richardson and Peter Spirtes. Ancestral graph Markov models. *The Annals of Statistics*, 30 (4):962–1030, 2002.

Jakob Runge, Peer Nowack, Marlene Kretschmer, Seth Flaxman, and Dino Sejdinovic. Detecting and quantifying causal associations in large nonlinear time series datasets. *Science Advances*, 5 (11):eaau4996, 2019. doi: 10.1126/sciadv.aau4996. URL https://www.science.org/doi/abs/10.1126/sciadv.aau4996.

Mahdi Salehi, Hamid R Karimi, Soroush Honari, and Axel P Hill. Calltif: Whole-brain causal discovery using fmri data. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 2078–2084, 2021.

Ruben Sanchez-Romero, Michael W Cole, and David M Schnyer. Estimating effective connectivity by jointly modeling neural dynamics and structural constraints. *NeuroImage*, 200:243–256, 2019a.

Ruben Sanchez-Romero, Joseph D Ramsey, Kun Zhang, Madelyn RK Glymour, Biwei Huang, and Clark Glymour. Estimating feedforward and feedback effective connections from fMRI time series: Assessments of statistical methods. *Network Neuroscience*, 3(2):274–306, 2019b.

Anil K Seth, Paul Chorley, and Lionel C Barnett. Granger causality analysis of fmri bold signals is invariant to hemodynamic convolution but not downsampling. *Neuroimage*, 65:540–555, 2013.

Shohei Shimizu, Patrik O Hoyer, Aapo Hyvärinen, and Marko Kerminen. A linear non-gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(Oct):2003–2030, 2006.

Stephen M Smith, Karla L Miller, Gholamreza Salimi-Khorshidi, Matthew Webster, Christian F Beckmann, Thomas E Nichols, Joseph D Ramsey, and Mark W Woolrich. Network modelling methods for fmri. *Neuroimage*, 54(2):875–891, 2011.

Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search*. Springer New York, 1993. doi: 10.1007/978-1-4612-2748-9. URL https://doi.org/10.1007/978-1-4612-2748-9.

Peter Spirtes, Clark N Glymour, Richard Scheines, and David Heckerman. *Causation, Prediction, and Search*. MIT press, 2000.

Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, prediction, and search*. MIT press, 2001.

Pedro A Valdes-Sosa, Alard Roebroeck, Jean Daunizeau, and Karl Friston. Effective connectivity: influence, causality and biophysical modeling. *Neuroimage*, 58(2):339–361, 2011.

Jiji Zhang. Causal reasoning with ancestral graphs. *Journal of Machine Learning Research*, 9: 1437–1474, 2008.