
GRACE-C: Generalized Rate Agnostic Causal Estimation via Constraints

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Graphical structures estimated by causal learning algorithms from time series data
2 can provide highly misleading causal information if the causal timescale of the
3 generating process fails to match the measurement timescale of the data. Existing
4 algorithms provide limited resources to respond to this challenge, and so researchers
5 must either use models that they know are likely misleading, or else forego causal
6 learning entirely. Existing methods face up-to-four distinct shortfalls, as they might
7 *a)* require that the difference between causal and measurement timescales is known;
8 *b)* only handle very small number of random variables when the timescale difference
9 is unknown; *c)* only apply to pairs of variables (albeit with fewer assumptions about
10 prior knowledge); or *d)* be unable to find a solution given statistical noise in the data.
11 This paper aims to address these challenges. We present an algorithm that combines
12 constraint programming with both theoretical insights into the problem structure
13 and prior information about admissible causal interactions to achieve speed up of
14 multiple orders of magnitude. The resulting system scales to significantly larger
15 sets of random variables (> 100) without knowledge of the timescale difference
16 while maintaining theoretical guarantees. This method is also robust to edge
17 misidentification and can use parametric connection strengths, while optionally
18 finding the optimal among many possible solutions.

19 1 Introduction

20 Dynamic causal models play a pivotal role in modeling real-world systems in diverse domains,
21 including economics, education, climatology, and neuroscience. Given a sufficiently accurate causal
22 graph over random variables, one can predict, explain, and potentially control some system; more
23 generally, one can understand it. In practice, however, specifying or learning an accurate causal
24 model of a dynamical system can be challenging for both statistical and theoretical reasons.

25 One particular challenge arises when data are not measured at the speed of the underlying causal
26 connections. For example, fMRI scanning of the brain measures bloodflow and oxygen level changes
27 in different brain regions, thereby indirectly measuring neural activity (which leads to increased
28 oxygen consumption). fMRI thus provides data about an important dynamical system, but these
29 measures take place (at most) every second while the brain's actual dynamics is known to proceed at a
30 faster rate [16], though we do not know how much faster. In general, when the measurement timescale
31 is significantly slower than the causal timescale (as with fMRI), learning can output importantly
32 incorrect causal information. For instance, if we only measure every other timestep in Figure 1,
33 then the true graph (top left) would differ from the data graph (top right). For example, we might
34 conclude that variable 2 directly influences variable 5, when variable 3 is the actual direct cause.
35 This type of error can lead to inefficient or costly methods of control. More generally, understanding
36 of a system depends on the causal-timescale (i.e., non-undersampled) causal relations, not the
37 measurement-timescale (apparent) relations.

38 In this paper, we consider the problem of learning the causal structure at the *causal* timescale from
 39 data collected at an unknown *measurement* timescale. This challenge has received significant attention
 40 in recent years [19, 7, 10, 18], but all current algorithms have significant limitations (see Section 2)
 41 that make them unusable for many real-world scientific challenges. Current algorithms show the
 42 theoretical possibility of causal learning from undersampled data, but their practical applicability
 43 is limited to small graph sizes, sometimes including only a pair of variables [7]. In contrast, we
 44 present a provably correct and complete algorithm that can operate on 100-node graphs and hence
 45 be potentially useful in biological and other domains for learning causal timescale structure from
 46 undersampled data.

47 2 Related Work And Notation

48 A directed dynamic causal model is a generalization of “regular” causal models [17, 23]: graph \mathbf{G} includes n distinct nodes
 49 for random variables $\mathbf{V} = \{V_1, V_2, \dots, V_n\}$ at both the current
 50 timestep t (\mathbf{V}^t), and also each previous timesteps (\mathbf{V}^{t-k}) in
 51 which there is a direct cause of some V_i^t . We assume that
 52 the “true” underlying causal structure is first-order Markov:
 53 the independence $\mathbf{V}^t \perp\!\!\!\perp \mathbf{V}^{t-k} \mid \mathbf{V}^{t-1}$ holds for all $k > 1$ ¹ (i.e.
 54 causal sufficiency assumption [24]). \mathbf{G} is thus over $2\mathbf{V}$, and the
 55 only permissible edges are $V_i^{t-1} \rightarrow V_j^t$, where possibly $i = j$.
 56 The quantitative component of the dynamic causal model is
 57 fully specified by parameters for $P(\mathbf{V}^t \mid \mathbf{V}^{t-1})$. We assume that
 58 these conditional probabilities are stationary over time, but the
 59 marginal $P(\mathbf{V}^t)$ need not be stationary.
 60

61 We denote the timepoints of the underlying causal structure
 62 as $\{t^0, t^1, t^2, \dots, t^k, \dots\}$. The data are said to be *undersampled*
 63 *at rate u* if measurements occur at $\{t^0, t^u, t^{2u}, \dots, t^{ku}, \dots\}$. We
 64 denote undersample rate with superscripts: the true causal
 65 graph (i.e., undersampled at rate 1) is \mathbf{G}^1 and that same graph undersampled at rate u is \mathbf{G}^u . To
 66 determine the implied \mathbf{G} at other timescales, the graph is first “unrolled” by adding instantiations of
 67 \mathbf{G}^1 at previous and future timesteps, where \mathbf{V}^{t-2} bear the same causal relationships to \mathbf{V}^{t-1} that \mathbf{V}^{t-1}
 68 bear to \mathbf{V}^t , and so forth. In this unrolled (time-indexed by t) graph, all \mathbf{V} at intermediate timesteps
 69 are not measured; this lack of measurement is equivalent to marginalizing out (the variables in)
 70 those timesteps to yield \mathbf{G}^u . This problem has been parametrically addressed by [7]. Yet, a very
 71 interesting approach proposed in the paper was demonstrated only on a 2-variable system. Although
 72 an interesting approach, it has not been developed further and made practical.

73 Various representations have been developed for graphs with latent confounders, including partially-
 74 observed ancestral graphs (PAGs) [20] and maximal ancestral graphs (MAGs) [26]. However, these
 75 graph-types cannot easily capture the types of latents produced by undersampling [14]. Instead, we
 76 use compressed graphs, along with properties that were previously proven for this representation [1].
 77 A condensed graph includes only \mathbf{V} , where temporal information is implicitly encoded in the edges.
 78 In particular, a condensed graph version \mathcal{G} of dynamic causal graph \mathbf{G} has $V_i \rightarrow V_j$ in \mathcal{G} iff $V_i^{t-1} \rightarrow V_j^t$
 79 is in \mathbf{G} . Undersampling (i.e., marginalizing intermediate timesteps) is a straightforward operation for
 80 compressed graphs: (1) $V_i \rightarrow V_j$ in \mathcal{G}^u iff there is a length- u directed path from V_i to V_j in \mathcal{G}^1 iff there
 81 is a directed path from V_i^{t-u} to V_j^t in \mathbf{G}^1 ; and (2) $V_i \leftrightarrow V_j$ in \mathcal{G}^u iff there exists length- $s < u$ directed
 82 paths from V_k to V_i , and to V_j , in \mathcal{G}^1 (i.e., V_k is an unobserved common cause in \mathbf{G}^1 fewer than u
 83 timesteps back). See Appendix for additional proofs. The bottom row of Figure 1 shows compressed
 84 graphs for the unrolled ones on the top row; the left shows the causal timescale and the right shows
 85 the graphs undersampled at rate 2.

86 Given this framework, the overall causal learning challenge can now be restated as: given \mathcal{G}^u but not u
 87 (or given dataset \mathbf{D} at unknown undersample rate), what is the set of possible \mathcal{G}^1 ? There will often be
 88 many possible \mathcal{G}^1 for given \mathcal{G}^u , and so we use $\llbracket \mathcal{H} \rrbracket$ to denote the equivalence class of \mathcal{G}^1 that could
 89 yield \mathcal{H} (the given causal graph inferred from data \mathbf{D}) for some u . That is, $\llbracket \mathcal{H} \rrbracket = \{\mathcal{G}^1 : \exists u (\mathcal{G}^u = \mathcal{H})\}$

¹This assumption is relatively weak, as we do not assume that we measure at this “true” causal timescale. The system timescale can be arbitrarily fast to capture all connections.




Figure 1: Causal graph \mathbf{G}^1 and its undersampled version \mathbf{G}^2 : unrolled and compressed versions.

90 Various algorithms have been developed to infer $\llbracket \mathcal{H} \rrbracket$, each with distinctive shortcomings. There are
 91 2^{n^2} possible \mathcal{G}^1 , so perhaps unsurprisingly, this problem is NP-complete:

92 **Theorem 1** ([10][Theorem 1]). *Deciding whether a consistent \mathbf{G}^1 exists for a given \mathcal{H} is NP-complete,*
 93 *for all undersampling rates $u \geq 2$.*²

94 *Mesochronal Structure Learning (MSL)* [19] showed it
 95 is possible to learn $\llbracket \mathcal{H} \rrbracket$ in a non-brute force manner
 96 if we know u . Every edge in \mathbf{G}^u corresponds to one
 97 or more paths of length u in \mathbf{G}^1 , and so \mathbf{G}^1 can be
 98 constructed by identifying $u - 1$ intermediate nodes
 99 for each edge in \mathbf{G}^u . MSL searches the state space of
 100 possible identifications in a Depth-First Search (DFS)
 101 manner. Each identification implies a \mathbf{G}^1 , and if $\mathbf{G}^u =$
 102 \mathcal{H} , then $\mathbf{G}^1 \in \llbracket \mathcal{H} \rrbracket$. Otherwise, search continues. MSL
 103 backtracks in the DFS whenever some \mathbf{G}^u includes an edge
 104 that is absent from \mathcal{H} , as the candidate \mathbf{G}^1 and all
 105 its supergraphs cannot be in $\llbracket \mathcal{H} \rrbracket$.

106 Although [19] showed that the concept that causal inference
 107 from undersampled data is feasible, MSL is computationally
 108 intractable on even moderate-sized graphs. [10] used the implied constraints to develop an Answer
 109 Set Programming (ASP) [22, 15, 6, 11] method that formulated this causal inference challenge as a rule-based
 110 constraint satisfaction problem. ASP is a rule-based
 111 declarative constraint satisfaction paradigm that is well-
 112 suited for representing and solving various NP-hard
 113 problems (e.g. Theorem 1). In essence, the algorithm in
 114 [10] takes as input the measured causal graph \mathcal{H} ,
 115 determines the set of implied constraints on \mathbf{G}^1 , and then uses the general-purpose Answer Set Solver
 116 *Clingo* [5] to determine the set of possible \mathbf{G}^1 significantly faster than MSL. The same idea of using
 117 Boolean satisfiability solvers to integrate (in)dependent data constraints has been used for various
 118 other causal learning challenges [9, 25].

121 Although the method in [10] is significantly faster, one must specify the undersampling rate u (or else
 122 run the method sequentially for all possible u , thereby losing much of the computational advantage).
 123 In contrast, the *Rate-Agnostic (Causal) Structure Learning (RASL)* approach (with three different
 124 versions) [18] makes no such assumption. These algorithms are similar to MSL, but consider each
 125 possible u for some \mathbf{G}^1 . RASL reduces computational complexity with two additional stopping rules
 126 for given \mathbf{G}^1 : (1) if some \mathbf{G}^k has previously been seen, then further undersampling of \mathbf{G}^1 will not
 127 produce new graphs; and (2) if \mathbf{G}^k is not an edge-subset of \mathcal{H} for all k , then do not consider any
 128 edge-superset of \mathbf{G}^1 [18]. However, despite these improvements, RASL still faces memory and
 129 run-time constraints for even moderate numbers of nodes.

130 One key observation from all of these learning algorithms is the importance of *strongly connected*
 131 *components* (SCCs) [1]:

132 **Definition 2.1.** *An SCC in compressed graph \mathcal{H} is a maximal set of nodes $S \subseteq V$ such that, for every*
 133 *$X, Y \in S$ there is a directed path from X to Y .*

134 Note that the variables in a compressed graph \mathcal{H} can be fully partitioned based on SCC membership.
 135 SCCs can be highly stable, as the node-membership of an SCC will not change as we undersample,
 136 as long as the greatest common divisor (gcd) of the set of lengths of all simple loops (directed cycles
 137 without repeated nodes) in the SCC is 1:³

138 **Theorem 2** ([1][Theorem 3]). *S is an SCC in \mathbf{G}^u for all u iff $\text{gcd}(\mathcal{L}_S) = 1$ for $\text{SCC } S \in \mathbf{G}^1$*

139 In this paper, we develop *sRASL* (for solver-based RASL), a novel algorithm that leverages insights
 140 from multiple sources, such as the constraints implied by SCC stability (Theorem 2). We show that

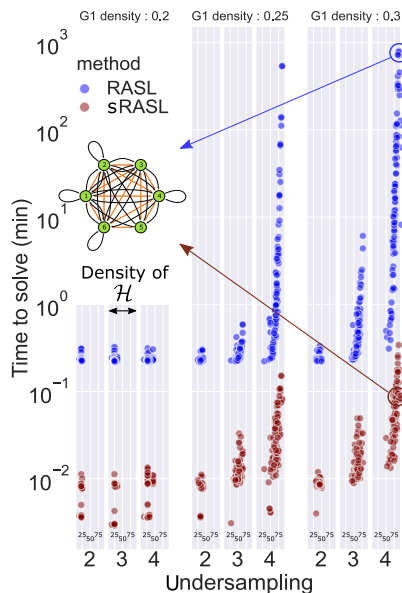


Figure 2: Comparison of sRASL (red) with previous state-of-the-art RASL (blue).

²Proof provided in [10]. In general, we omit previously published proofs.

³The condition easily holds, as it requires only (1) the graph is relatively dense with different loop lengths or (2) any node in the SCC has a self-loop (i.e., is autocorrelated).

141 sRASL significantly outperforms previous methods. The contributions of this paper are threefold:
 142 first, we reformulated the RASL algorithm from a search-based procedure to a constraint satisfaction
 143 problem encoded in a declarative language [4]. Second, this reformulation enables us to add additional
 144 constraints based on SCC structure, and thereby gain significant speed-up. Third, we ensure that
 145 sRASL provides a straightforward way to find approximate solutions when \mathcal{H} is an unreachable
 146 graph (i.e., when $\llbracket \mathcal{H} \rrbracket = \emptyset$). These advances collectively provide up to three orders of magnitude
 147 improvements in speed, thereby enabling causal inference given undersampling data involving over
 148 100 nodes. As a concrete example of the improvements, Figure 2 compares sRASL (red) with the
 149 previously-fastest RASL [18] method (blue) on the same graphs. The same input graph \mathcal{H} took
 150 RASL nearly 1000 minutes to compute $\llbracket \mathcal{H} \rrbracket$, but only 6 seconds for sRASL.

151 3 sRASL: Optimized ASP-based Causal Discovery

152 The sRASL algorithm takes as input a (potentially) undersampled graph \mathcal{H} , whether learned from data
 153 \mathbf{D} , expert domain knowledge, a combination of the two, or some other source. sRASL’s agnosticism
 154 about the source of the input graph enables wider applicability, as we can use whatever information is
 155 available [2]. In the asymptotic (data) limit, the sRASL output is the full $\llbracket \mathcal{H} \rrbracket$.

156 sRASL leverages the fact that connections between SCCs in \mathcal{H} must form a directed acyclic graph.
 157 More specifically: if $X \rightarrow Y$ with $X \in \mathbf{A}, Y \in \mathbf{B}$ for SCCs $\mathbf{A} \neq \mathbf{B}$, then $C \leftarrow D$ for all $C \in \mathbf{A}, D \in \mathbf{B}$.⁴
 158 Moreover, Theorem 2 provides the (weak) condition under which SCC membership is preserved
 159 under undersampling. These two observations imply that structural features potentially provide
 160 additional constraints beyond the obvious ones (See Section 4.3). In particular, if \mathcal{H} has a roughly
 161 modular structure—that is, the SCCs are not too large—then sRASL generates many more constraints
 162 than the algorithm of [10].

163 Listing 1 shows the Clingo (for a brief Introduction on Clingo and Answer Set Programming,
 164 refer to Appendix C) code of sRASL, which is based on exactly representing the conditioning and
 165 marginalization operations (defined in Section 2) in ASP. In the first line, we input the first-order
 166 graph-specific specification of \mathcal{H} (e.g., the edge $1 \rightarrow 10$ translates to `hdirected(1, 10)`). Line 2
 167 encodes the second-order structure of \mathcal{H} , including the partition of \mathbf{V} into SCCs. These predicates
 168 and basic descriptive information are added to the Clingo code (lines 3, 4, 5) in an automated way.⁵

169 `maxu` on line 3 specifies the maximum undersampling rate, as there is provably such a u where
 170 $\mathcal{G}^u = \mathcal{G}^k$ for all $k > u$, if we have the same condition that leads to stable SCC membership:

171 **Theorem 3** ([18][Theorem 3.1]). *If $\gcd(\mathcal{L}_S) = 1$ for all SCCs $S \subseteq \mathbf{V}$, then $\mathcal{G}^u = \mathcal{G}^{u+1}$ for all*
 172 *$u > f \leq n_F + \gamma + d + 1$.*

173 where γ is the transit number⁶, d is graph diameter⁷ and n_F is the Frobenius number.⁸ In practice, the
 174 plausible undersampling rate will often be much lower than the theoretical upper bound in Theorem 3.
 175 For example, consider fMRI data. The underlying rate of brain activity is generally thought to be
 176 ~ 100 milliseconds and fMRI devices measure approximately every two seconds. Hence, $u = 20$ is a
 177 plausible upper bound on undersampling in fMRI studies.⁹

178 Line 6 in Listing 1 stipulates that all edges in \mathcal{G}^1 are possible (by default), and so the output will
 179 contain any possible model that does not violate the integrity constraints of lines 11 – 16. Lines
 180 7 and 8 define paths of length L in the graph (i.e., an edge in \mathcal{G}^L). As described in Section 2:
 181 $X \rightarrow Y \in \mathcal{G}^u \iff X \rightsquigarrow^u Y \in \mathcal{G}^1$ where \rightsquigarrow^u is a path of length u . Line 10 similarly defines bidirected
 182 edges in \mathcal{G}^L : $X \leftrightarrow Y \in \mathcal{G}^u \iff \exists Z, l : (X \overset{l}{\leftarrow} Z \overset{l}{\rightarrow} Y \in \mathcal{G}^1)$.

⁴If $C \leftarrow D$, then by definition of SCC, there exists $\pi : X \leftarrow \dots \leftarrow C \leftarrow D \leftarrow \dots \leftarrow Y$. X, Y are thus mutually reachable so must be in the same SCC, contra $\mathbf{A} \neq \mathbf{B}$.

⁵The code is available at `removed for anonymity`

⁶Transient number is the length of the “longest shortest path” from a node that touches all simple loops of the SCC.

⁷Graph diameter the length of the “longest shortest path” between any two graph nodes.

⁸For set \mathbf{B} of positive integers with $\gcd(\mathbf{B}) = 1$, n_F is the max integer with $n_F \neq \sum_{i=1}^b \alpha_i B_i$ for $\alpha_i \geq 0$

⁹Of course, the actual undersample rate could be much lower than 20. Voxels typically contain 8 – 10 layers of neurons, so the “causal timescale of a voxel” could easily be as high as 1000 ms (i.e., $u = 2$).

```

1  %( * input graph edge specifications here * e.g.: hdirected(1,5) ... )
2  %( * input graph SCC specifications here * e.g.: sccsize(0, 5). scc(1, 0) ...)
3  #const n = 10, maxu = 20
4  node(1..n).
5  1 {u(1..maxu)} 1.
6  {edge1(X,Y)} :- node(X), node(Y).
7  directed(X, Y, 1) :- edge1(X, Y).
8  directed(X, Y, L) :- directed(X, Z, L-1),
9                        edge1(Z, Y), L <= U, u(U).
10 bidirected(X, Y, U) :- directed(Z, X, L), directed(Z, Y, L), node(X;Y;Z), X < Y, L
11                        < U, u(U).
12 :- directed(X, Y, L), not hdirected(X, Y), node(X;Y), u(L).
13 :- bidirected(X, Y, L), not hbidirected(X, Y), node(X;Y), u(L), X < Y.
14 :- not directed(X, Y, L), hdirected(X, Y), node(X;Y), u(L).
15 :- not bidirected(X, Y, L), hbidirected(X, Y), node(X;Y), u(L), X < Y.
16 % the following is only used when SCC accounting is enabled
17 :- edge1(X, Y), scc(X, K), scc(Y, L), K != L, sccsize(L, Z), Z > 1, not dag(K,L).

```

Listing 1: Clingo code for sRASL

```

1  :-~ directed(X, Y, L), no_hdirected(X, Y, W), node(X;Y), u(L). [W@1,X,Y]
2  :-~ bidirected(X, Y, L), no_hbidirected(X, Y, W), node(X;Y), u(L), X < Y.
3  [W@1,X,Y]
4  :-~ not directed(X, Y, L), hdirected(X, Y, W), node(X;Y), u(L). [W@1,X,Y]
5  :-~ not bidirected(X, Y, L), hbidirected(X, Y, W), node(X;Y), u(L), X < Y.
6  [W@1,X,Y]

```

Listing 2: Integrity constraints for turning sRASL algorithm into an optimization problem when they replace lines 11 through 14 in Listing 1

Lines 11 – 14 provide the core constraints, as they ensure that sRASL returns only \mathcal{G}^1 for which there exists u such that $\mathcal{G}^u = \mathcal{H}$. Line 16 adds the additional constraints based on impermissibility of cycles between SCCs. That is, if we consider each SCC as a *super-node*, Line 16 ensures that the edges of the directed acyclic graph (DAG) connecting SCCs in \mathcal{H} are not violated in the outputs.

If sRASL initially returns the empty set (i.e., there are no suitable \mathcal{G}^1), then it is possible to run sRASL in an optimization mode instead to find optimal (though not perfect) outputs (see Section 4.5 for details). One potential reason for $[\mathcal{H}] = \emptyset$ is statistical noise or other errors in estimating or specifying \mathcal{H} .¹⁰ In such cases, sRASL finds the set of \mathbf{G}^1 that are, for some u , closest to \mathcal{H} by the objective function:

$$\mathcal{G}^{1*}, u^* \in \operatorname{argmin} \sum_{e \in \mathcal{H}} I[e \notin \mathcal{G}^u] \cdot w(e \in \mathcal{H}) + \sum_{e \notin \mathcal{H}} I[e \in \mathcal{G}^u] \cdot w(e \notin \mathcal{H}), \quad (1)$$

where the indicator function $I(c) = 1$ if the condition holds and zero otherwise. $w(e \in \mathcal{H})$ indicates the importance (i.e., reliability) of edge e ; $w(e \notin \mathcal{H})$ indicates the reliability of the absence of an edge. Since \mathcal{H} is an undersampled graph, it consists of directed and bidirected edges. We thus implement both $w(e \in \mathcal{H})$ and $w(e \notin \mathcal{H})$ as two pairs of $n \times n$ matrices, one pair for existence and absence of directed edges, and one pair for bidirected edges. To learn the optimal graph at the true causal timescale, for every \mathcal{G}^1 in the solutions set, the corresponding \mathcal{G}^u is compared to the input \mathcal{H} and penalized for the difference according to weights representing the reliability of the measurement timescale estimates.

In order to incorporate Equation 5 in Listing 1, we replace its exact integrity constraints (Lines 11-14) with the optimization formulation [5] in Listing 2. In Listing 2 we specify a weight for each edge (or lack thereof) in \mathcal{H} using \mathbb{W} and the importance of these weights can be specified for each integrity constraint using the $\mathbb{W}@i$ syntax with i being the importance.

¹⁰Note, among all possible graphs that have a combination of both directed (2^{n^2}) and bidirected ($2^{\binom{n}{2}}$) edges only a fraction may be obtained by undersampling a \mathcal{G}^1 .

204 3.1 sRASL Completeness and Correctness

205 sRASL exhibits significant improvements in computation time, so it is important to show that we do
206 not lose generality or theoretical guarantees. We demonstrate correctness and completeness using
207 the notion of a *direct encoding* of the problem (i.e., the space of solutions is fully characterized, and
208 any non-solution violates a constraint). We first prove (Appendix A) that we have provided a direct
209 encoding:

210 **Theorem 4.** *Listing 1 is a direct encoding of the undersampling problem.*

211 Clingo is a complete solver, based on CDNL (Conflict-Driven Nogood Learning) [3], itself based on
212 CDCL (Conflict-Driven Clause Learning) [12, 13]. [8][Theorem 2] and [9][Section 5.2] show that, if
213 the ASP encoding is the direct encoding of the problem, then ASP will produce the complete set of
214 solutions in the infinite sample space limit. In other words, Theorem 5 implies: since our algorithm
215 yields at least one sound solution, Clingo will produce all possible solutions. Therefore, soundness
216 results in completeness. That is, sRASL’s success is not due to heuristics or some incomplete or
217 not-everywhere-correct algorithmic step.¹¹

218 4 Results

219 A major virtue of sRASL is its empirical performance, so we now consider a range of simulations (to
220 ensure known ground truth) to understand this performance in more detail. For these experiments, we
221 used Clingo in parallel mode using 10 threads and computing on AMD EPYC 7551 CPUs. To cope
222 with the multiple repeated calculations and hundreds of graphs we have tested per parameter setting
223 all experiments were run on a slurm cluster which submits jobs to one of the 19 machines on the
224 same network. Each of the 19 nodes was equipped with 64 cores and 512 GB of RAM.

225 4.1 Comparing sRASL vs. RASL

226 We first compare sRASL with the existing RASL method (Figure 2). We generated 100 6-node SCCs
227 for each density in [0.2, 0.25, 0.3], and then undersampled each graph by 2, 3, and 4. We used 6-node
228 graphs as RASL struggles to handle larger graphs in reasonable time and space [18]. Each column of
229 Figure 2 consists of graphs of approximately same density (increasing density from left-to-right), and
230 subcolumns represent different undersample rates (for that density). As Figure 2 shows, sRASL is
231 typically three orders of magnitude faster than RASL, even on relatively small graphs.

232 4.2 Comparing Graph Size

233 It is perhaps unsurprising that sRASL runs much faster than RASL, as sRASL uses an ASP solver
234 (which were previously known to yield faster algorithms [10]). We next wanted to see just how
235 much larger the graphs could be. More generally, we aimed to better understand how sRASL’s
236 computational performance scales with the number of nodes for single-SCC graphs. The focus on
237 single SCCs is motivated by the theoretical need to understand the size-speed tradeoff, and also
238 scientific applicability since many real-world systems consist of tightly coupled factors with many
239 feedback loops (i.e., they are a single SCC). We consider multiple-SCC graphs in later subsections.

240 We generated 50 random single-SCC graphs each of 8, 16 and 32 nodes, all with average degree
241 of 1.4 outgoing edges per node. We then undersampled each graph by 2, 3 and 4, and used each
242 individual undersampled graph as input to sRASL. We used a 24-hour timeout (i.e., we stopped an
243 sRASL run if it did not finish in 24 hours). Figure 3 shows the increasing computational costs as both
244 number of nodes and undersample rate increase. Notably, sRASL was able to learn $[[\mathcal{H}]]$ for 32-node
245 single-SCC graphs, though it reached timeout for all \mathcal{H} at $u = 4$ 32-node graphs. That is, for low u ,
246 sRASL scales to much larger single-SCC graphs than RASL.

247 4.3 Comparing SCC Size

248 The other major innovation of sRASL is incorporation of constraints derived from the SCC structure.
249 We thus investigated the performance of sRASL on large, structured, multiple-SCC graphs. Many

¹¹Simulation testing provides further evidence. We found that sRASL and RASL produced identical outputs for 1000 different input graphs, and RASL is known to be correct and complete [18][Theorem 3.6].

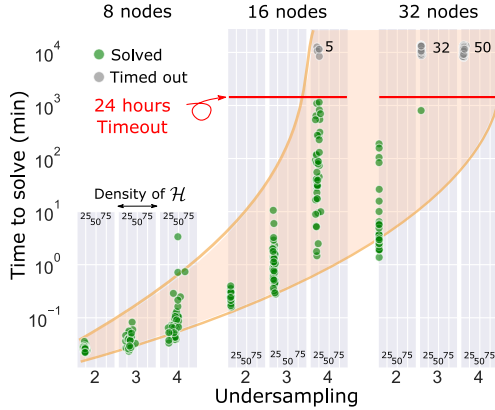


Figure 3: Time behavior of graphs of size 8, 16 and 32. The time out for this experiment indicated by the red line was 24 hours. Green dots represent graphs that has been computed within the 24-hours window. Gray represent graphs that could not be fully computed within 24-hours window.

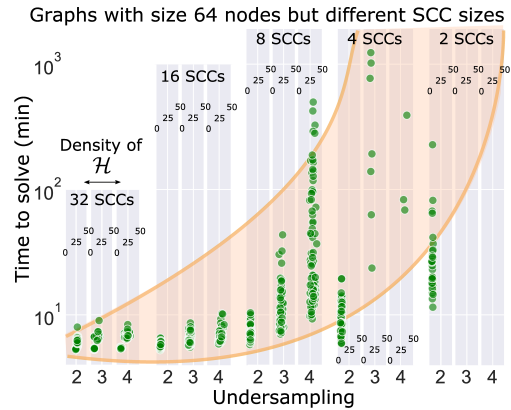


Figure 4: Time behavior of graphs of size 64 with various sub SCC sizes. The time out for this experiment was 24 hours (1440 Minutes).

250 real-world systems exhibit some degree of modularity, where there are dense or feedback connections
 251 within a module or subsystem, and relatively sparser connections between modules or subsystems.
 252 In theory, sRASL should perform well on these kinds of structures since it incorporates SCC-based
 253 constraints. Please refer to Appendix B for an ablation study on effect of using additional constraints
 254 for SCC structures.

255 We tested the value of SCC-based constraints using graphs with 64 nodes that differed in their SCC
 256 structure. Specifically, we randomly generated 50 graphs each of: 32 size-2 SCCs; 16 size-4 SCCs; 8
 257 size-8 SCCs; 4 size-16 SCCs; or 2 size-32 SCCs. We then undersampled each graph by $u = 2, 3, 4$,
 258 and ran sRASL (again with a 24-hour timeout).

259 Figure 4 shows the computation time for these graphs, with increasing SCC size (and decreasing
 260 number of SCCs) from left to right. The first key observation is that sRASL successfully found
 261 $\llbracket \mathcal{H} \rrbracket$ for 64-node graphs, at least when there was some internal structure. Second, and relatedly, we
 262 observe a wide range of computation times for these graphs, even though all had the same number
 263 of nodes (64). We clearly see the impact of SCC structure, as sRASL was dramatically faster when
 264 there were many small SCCs, rather than a few large SCCs. The results in Figure 3 might seem to
 265 suggest an “upper bound” around 30 nodes for sRASL. But the results in Figure 4 make it clear that
 266 any potential “upper bound” is primarily on the number of nodes *in the SCCs*, rather than the total
 267 number of nodes in the graph.

268 4.4 Comparing Graph Size With Constant SCC Size

269 The previous results suggest that sRASL might be able to solve much larger graphs, as long as
 270 the SCCs are not overly large. More generally, the previous simulations showed that sRASL’s
 271 computational cost scales (at least) exponentially in the *size* of the SCC, but did not reveal how it
 272 scales in the *number* of SCCs.

273 We again generated 50 different graphs for each of several settings. We considered SCCs with 7, 8,
 274 and 10 nodes, and varied the number of SCCs within the graph (again for $u = 2, 3$, and 4). Figure 5
 275 shows the computational cost of sRASL, where each row includes graphs with SCCs of the same
 276 size, but the number of SCCs increasing from left-to-right. The critical observation here is that the
 277 time complexity grows approximately *linearly*, rather than exponentially (or worse). For example,
 278 the graph shown in Figure 5 has 98 nodes, but sRASL successfully computes $\llbracket \mathcal{H} \rrbracket$ in approximately
 279 20 minutes. (Recall that RASL took 17 hours to compute a graph with only 6 nodes.)

280 This simulation demonstrates that sRASL is usable on relatively large graphs, as long as there is
 281 appropriate internal structure. One might worry, though, whether real-world systems do not have the
 282 right structure. If we consider fMRI (brain) data, [21] recently aggregated a number of simulations of

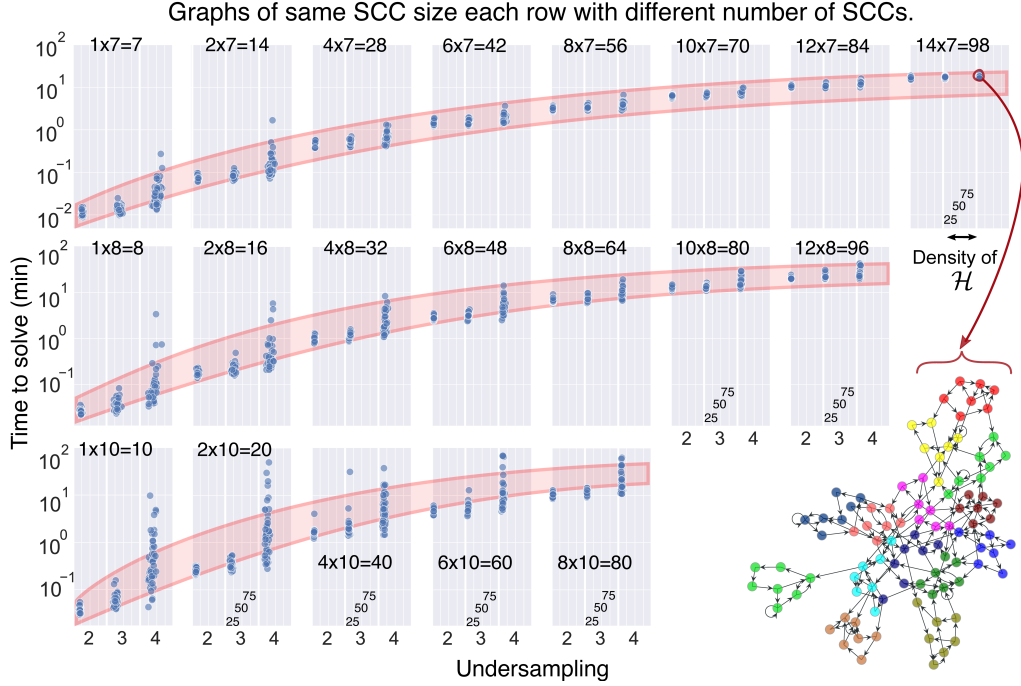


Figure 5: Time behaviour of graphs with the same SCCs sizes but with multiple number of SCCs. Top row graphs of SCC size 7 with 1, 2, ..., 14 number of SCCs. Middle row graphs of SCC size 8. Bottom row graphs of SCC size 10. Bottom right corner is an example of a structured graph with 98 nodes structured as 14 SCCs of size 7. Each color represents one Strongly Connected Component.

283 realistic causal graphs for brain processes studied with fMRI, and the largest SCC in these widely-
 284 accepted models has only seven nodes. Moreover, typical brain parcellations contain 50 – 100 regions
 285 (= nodes), and sRASL can easily handle graphs with 100 nodes if the SCC size is in the 8 – 10 range.

286 The results in this subsection suggest that we could potentially find $\llbracket \mathcal{H} \rrbracket$ for each larger graphs, as
 287 long as they were composed of reasonably-sized SCCs. However, we found that the Clingo language
 288 and solver seems to be limited in the number of atoms that it can handle. In our simulations, graphs
 289 of size 100 seem to be the limit for Clingo to handle all the predicates. An open question is whether
 290 sRASL can be optimized to produce fewer predicates (or Clingo improved to handle more atoms).

291 4.5 Optimization

292 Finally, we explored the optimization capability of Clingo. Recall that sometimes $\llbracket \mathcal{H} \rrbracket = \emptyset$ due
 293 to statistical errors or other noise in learning \mathcal{H} . Clingo can solve an optimization problem based
 294 on user-specified weights and priorities, and output a single solution with minimum cost function
 295 (along with u for this solution). In particular, we can use Clingo to find \mathcal{G}^1 whose \mathcal{G}^u (for some u)
 296 are closest (relative to the edge weights) to \mathcal{H} .¹²

297 In this simulation, we first randomly generate \mathcal{G}^1 and undersample it to a random u to get $\mathcal{G}^u = \mathcal{H}$
 298 such that $\llbracket \mathcal{H} \rrbracket \neq \emptyset$. We then assign weights to the edges of \mathcal{H} and randomly break one edge from it.
 299 We then run sRASL on this “broken” \mathcal{H} to learn a suitable \mathcal{G}^1 . Red bars in Figure 6 show the edge
 300 omission and commission errors for this approach. We see that, except for high undersamplings, the
 301 optimization capability of Clingo can be used to frequently retrieve the true \mathcal{G}^1 ; that is, this version
 302 of sRASL is robust to small errors in \mathcal{H} in many settings.

303 A more complex approach to finding suitable solutions is to first run the optimization method to
 304 identify a solution \mathbf{G}_{opt}^1 and undersample rate u_{opt} . We can then undersample this solution \mathbf{G}_{opt}^1
 305 by u_{opt} to get \mathbf{G}_{opt}^u . We then use sRASL to obtain $\llbracket \mathbf{G}_{opt}^u \rrbracket$ (i.e., the full equivalence class of the undersampled
 306 graph that is “nearest” to \mathcal{H}). We then compute the error based on the minimum error among all

¹²If $\llbracket \mathcal{H} \rrbracket \neq \emptyset$, then this optimization will return a graph from $\llbracket \mathcal{H} \rrbracket$,

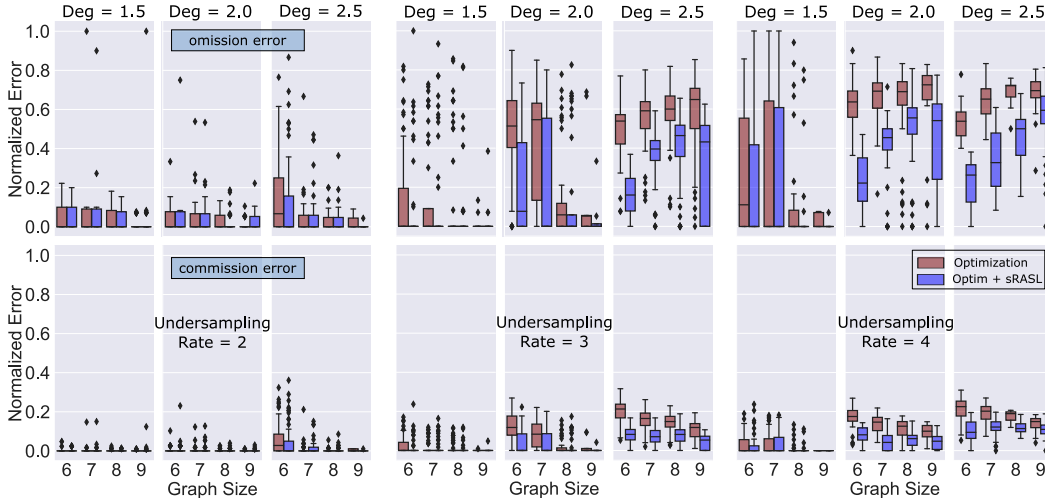


Figure 6: The omission (top) and commission (bottom) error of different graph sizes and undersampling of two, three and four from left to right.

307 $\mathbf{G}^1 \in \llbracket \mathbf{G}_{opt}^u \rrbracket$; that is, we ask whether the true graph was actually found. This approach is motivated
 308 by the intended use of sRASL by domain scientists, where the final decision on which graph in the
 309 equivalence class better suits the question is made by the scientist using the algorithm. Blue bars in
 310 Figure 6 show that this more complex method provides improved performance compared to regular
 311 optimization.

312 5 Conclusion and Discussion

313 Real-world scientific problems frequently involve measurement processes that operate at a different
 314 timescale than the causal structure of the system under study. As causal learning and analysis methods
 315 are increasingly used to address societal and policy challenges, it is increasingly critical that we
 316 use methods that reveal usable information (while also being clear when we *cannot* infer some
 317 information). Obviously, like any method, sRASL could yield information that is misused, but the
 318 aim here is to provide another useful tool in the scientists’ policy-makers’ toolboxes. If measurements
 319 occur at a slower rate than the causal influences, then causal discovery from those undersampled
 320 data can yield highly misleading outputs. Multiple methods have been developed to infer aspects
 321 of the underlying causal structure from the undersampled data/graph. However, the assumptions or
 322 computational complexities of those algorithms make them unusable for most real-world challenges.
 323 In this paper, we have developed and tested sRASL, a novel algorithm that is less subject to those
 324 same limitations. More specifically, sRASL provides all consistent solutions (without knowledge
 325 of exact undersampling rate) for large (100-node) graphs in a usable amount of time. sRASL also
 326 shows reasonable robustness to statistical error in the estimated graph by finding the closest consistent
 327 solution. Future research will focus on application of sRASL to actual neuroimaging data, and
 328 extensions to situations with multiple measurement modalities.

329 References

330 [1] DANKS, D., AND PLIS, S. Learning causal structure from undersampled time series. In *NIPS*
 331 *Workshop on Causality* (2013), vol. 1, pp. 1–10.

332 [2] DANKS, D., AND PLIS, S. Amalgamating evidence of dynamics. *Synthese* 196, 8 (2019), 3213–
 333 3230.

334 [3] DRESCHER, C., AND WALSH, T. Conflict-driven constraint answer set solving with lazy nogood
 335 generation. In *Twenty-Fifth AAAI Conference on Artificial Intelligence* (2011).

336 [4] FAHLAND, D., LÜBKE, D., MENDLING, J., REIJERS, H., WEBER, B., WEIDLICH, M., AND ZUGAL, S.
 337 Declarative versus imperative process modeling languages: The issue of understandability. In
 338 *Enterprise, Business-Process and Information Systems Modeling*. Springer, 2009, pp. 353–366.

- 339 [5] GEBSEER, M., KAUFMANN, B., KAMINSKI, R., OSTROWSKI, M., SCHAUB, T., AND SCHNEIDER, M.
340 Potassco: The Potsdam answer set solving collection. *Ai Communications* 24, 2 (2011),
341 107–124.
- 342 [6] GELFOND, M., AND LIFSCHITZ, V. The stable model semantics for logic programming. ICSLP,
343 1988.
- 344 [7] GONG, M., ZHANG, K., SCHOELKOPF, B., TAO, D., AND GEIGER, P. Discovering temporal causal
345 relations from subsampled data. In *International Conference on Machine Learning* (2015),
346 PMLR, pp. 1898–1906.
- 347 [8] HYTINEN, A., EBERHARDT, F., AND JÄRVISALO, M. Constraint-based Causal Discovery: Conflict
348 Resolution with Answer Set Programming. In *UAI* (2014), pp. 340–349.
- 349 [9] HYTINEN, A., HOYER, P. O., EBERHARDT, F., AND JÄRVISALO, M. Discovering cyclic causal models
350 with latent variables: A general SAT-based procedure. *arXiv preprint arXiv:1309.6836* (2013).
- 351 [10] HYTINEN, A., PLIS, S., JÄRVISALO, M., EBERHARDT, F., AND DANKS, D. A constraint optimization
352 approach to causal discovery from subsampled time series data. *International Journal of*
353 *Approximate Reasoning* 90 (2017), 208–225.
- 354 [11] LIFSCHITZ, V. The stable model semantics for logic programming, 1988.
- 355 [12] MARQUES SILVA, J., AND SAKALLAH, K. GRASP-A new search algorithm for satisfiability. In
356 *Proceedings of International Conference on Computer Aided Design* (1996), pp. 220–227.
- 357 [13] MARQUES-SILVA, J. P., AND SAKALLAH, K. A. GRASP: A search algorithm for propositional
358 satisfiability. *IEEE Transactions on Computers* 48, 5 (1999), 506–521.
- 359 [14] MOOIJ, J. M., AND CLAASSEN, T. Constraint-based causal discovery using partial ancestral graphs
360 in the presence of cycles. In *Conference on Uncertainty in Artificial Intelligence* (2020), PMLR,
361 pp. 1159–1168.
- 362 [15] NIEMELÄ, I. Logic programs with stable model semantics as a constraint programming paradigm.
363 *Annals of mathematics and Artificial Intelligence* 25, 3 (1999), 241–273.
- 364 [16] ORAM, M., AND PERRETT, D. Time course of neural responses discriminating different views of
365 the face and head. *Journal of neurophysiology* 68, 1 (1992), 70–84.
- 366 [17] PEARL, J., ET AL. Models, reasoning and inference. *Cambridge, UK: CambridgeUniversityPress*
367 *19* (2000), 2.
- 368 [18] PLIS, S., DANKS, D., FREEMAN, C., AND CALHOUN, V. Rate-agnostic (causal) structure learning. In
369 *Advances in neural information processing systems* (2015), pp. 3303–3311.
- 370 [19] PLIS, S., DANKS, D., AND YANG, J. Mesochronal structure learning. In *Uncertainty in artificial in-*
371 *telligence: proceedings of the... conference. Conference on Uncertainty in Artificial Intelligence*
372 (2015), vol. 31, NIH Public Access.
- 373 [20] RICHARDSON, T., AND SPIRITES, P. Ancestral graph Markov models. *The Annals of Statistics* 30, 4
374 (2002), 962–1030.
- 375 [21] SANCHEZ-ROMERO, R., RAMSEY, J. D., ZHANG, K., GLYMOUR, M. R., HUANG, B., AND GLYMOUR, C.
376 Estimating feedforward and feedback effective connections from fMRI time series: Assessments
377 of statistical methods. *Network Neuroscience* 3, 2 (2019), 274–306.
- 378 [22] SIMONS, P., NIEMELÄ, I., AND SOININEN, T. Extending and implementing the stable model semantics.
379 *Artificial Intelligence* 138, 1-2 (2002), 181–234.
- 380 [23] SPIRITES, P., GLYMOUR, C., AND SCHEINES, R. *Causation, Prediction, and Search*. Springer New
381 York, 1993.
- 382 [24] SPIRITES, P., GLYMOUR, C. N., SCHEINES, R., AND HECKERMAN, D. *Causation, Prediction, and*
383 *Search*. MIT press, 2000.
- 384 [25] TRIANTAFILLOU, S., TSAMARDINOS, I., AND TOLLIS, I. Learning causal structure from overlapping
385 variable sets. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence*
386 *and Statistics* (2010), JMLR Workshop and Conference Proceedings, pp. 860–867.
- 387 [26] ZHANG, J. Causal reasoning with ancestral graphs. *Journal of Machine Learning Research* 9
388 (2008), 1437–1474.

389 Checklist

- 390 1. For all authors...
- 391 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s
392 contributions and scope? [Yes] See Sections 3.1 and 4.
- 393 (b) Did you describe the limitations of your work? [Yes] See Sections 4.3 and 4.4.
- 394 (c) Did you discuss any potential negative societal impacts of your work? [Yes] See
395 Section 5.
- 396 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
397 them? [Yes]
- 398 2. If you are including theoretical results...
- 399 (a) Did you state the full set of assumptions of all theoretical results? [Yes] See footnote
400 number 3.
- 401 (b) Did you include complete proofs of all theoretical results? [Yes] For theorems directly
402 borrowed from other papers, we have cited the papers where the proofs are present. For
403 other theorems, we have provided proofs.
- 404 3. If you ran experiments...
- 405 (a) Did you include the code, data, and instructions needed to reproduce the main ex-
406 perimental results (either in the supplemental material or as a URL)? [Yes] We have
407 provided a URL to complete reproducible code and data. However, we obscured the
408 link until after the review for anonymity.
- 409 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
410 were chosen)? [Yes] In Section 4 we specified the hyperparameters and why we chose
411 them.
- 412 (c) Did you report error bars (e.g., with respect to the random seed after running experi-
413 ments multiple times)? [Yes] See Figure 6. We have also shown the variations of our
414 method by running on 50 or 100 random graphs(See Section 4).
- 415 (d) Did you include the total amount of compute and the type of resources used (e.g., type
416 of GPUs, internal cluster, or cloud provider)? [Yes] See first paragraph of Section 4.
- 417 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 418 (a) If your work uses existing assets, did you cite the creators? [Yes] We use Clingo and
419 we have cited the reference. See Section 2 paragraph 6.
- 420 (b) Did you mention the license of the assets? [N/A]
- 421 (c) Did you include any new assets either in the supplemental material or as a URL? [No]
- 422 (d) Did you discuss whether and how consent was obtained from people whose data you’re
423 using/curating? [N/A]
- 424 (e) Did you discuss whether the data you are using/curating contains personally identifiable
425 information or offensive content? [N/A]
- 426 5. If you used crowdsourcing or conducted research with human subjects...
- 427 (a) Did you include the full text of instructions given to participants and screenshots, if
428 applicable? [N/A]
- 429 (b) Did you describe any potential participant risks, with links to Institutional Review
430 Board (IRB) approvals, if applicable? [N/A]
- 431 (c) Did you include the estimated hourly wage paid to participants and the total amount
432 spent on participant compensation? [N/A]

433 A Appendix

434 We start with proving some results used in conversion of the DBN structures to their compressed
435 graph representations.

436 **Lemma 1.** *For all u , G_u contains no directed edges between variables at the same time step.*

437 *Proof.* $u = 1$ holds by assumption for G_1 . For $u > 1$, every directed edge corresponds to a directed
438 path of length u in G_1 . Since all directed edges in G_1 are from $t - 1$ to t (or more generally, from $t - k$
439 to $t - (k + 1)$), every directed path in G_1 is from an earlier time step to the current one. Hence, no
440 directed edge in G_u can be from V_i^t to V_j^t . \square

441 **Lemma 2.** *If the Markov order of G_1 is 1, then the Markov order of all G_u is also 1 (relative to*
442 *measurement at rate u).*

443 *Proof.* The Markov order of a dynamic causal graph is the smallest m such that \mathbf{V}^t is independent of
444 \mathbf{V}^{t-r} given $\mathbf{V}^{t-1}, \dots, \mathbf{V}^{t-m}$ for all $r > m$. If the Markov order of G_1 is 1, then all paths from \mathbf{V}^{t-r} to \mathbf{V}^t
445 must be blocked by \mathbf{V}^{t-1} for $r > 1$. Since graphical structure is replicated across timesteps, it follows
446 that all paths from \mathbf{V}^{t-r} to \mathbf{V}^t must be blocked by \mathbf{V}^{t-u} for $r > u$. Therefore, the Markov order of G_u
447 is u , which corresponds to Markov order 1 for measurements at rate u . \square

448 The following theorem demonstrates correctness of our ASP algorithm.

449 **Theorem 5.** *Listing 1 is a direct encoding of the undersampling problem.*

450 *Proof.* We will prove this by contradiction. Let us call the undersampled input graph to the algorithm
451 \mathcal{H} , considering that is the undersampled version of a graph \mathbf{G}_{true}^1 at rate u_{true} . By definition, every
452 directed edge in \mathcal{H} corresponds to a path of length u_{true} in \mathbf{G}_{true}^1 . Similarly, every bidirected edge in
453 \mathcal{H} corresponds to an unobserved common cause fewer than u_{true} timesteps back (refer to Section 2 for
454 exact definition). Line 7 – 11 in Listing 1 considers all such \mathbf{G}^1 s without exclusion. Let us call the
455 set all the pairs of graphs and corresponding undersampling rates u described by Listing 1 \mathbf{S} .

456 Let us assume there is a pair \mathbf{G}_a^1 and u_a that is in \mathbf{S} but if we undersample \mathbf{G}_a^1 by u_a , let us call it \mathbf{G}_a^u ,
457 will not be the same as \mathcal{H} . If \mathbf{G}_a^u has an extra directed (bidirected) edge, this will contradict with
458 line 12(13) of Listing 1. Similarly, if \mathcal{H} has a directed (bidirected) edge that is not present in \mathbf{G}_a^u ,
459 it will contradict with line 14(16). Therefore, Listing 1 is a direct encoding of the undersampling
460 problem. \square

461 B The Effects of Accounting for SCCs In sRASL

462 In this section, we show the results of additional experiments on the effects of accounting for strongly
463 connected components (SCCs) when the graph has a modular structure (i.e., consists of several
464 interconnected strongly connected components). For this experiment, we generated 50 random
465 graphs sized 8 to 15 with multiple SCCs as described in Table 1. Then on the same set of graphs,
466 we ran sRASL once with using our additional constraints for SCC structures and once without
467 accounting for the modular structure. We limited the computational resources available to each run to
468 24 hours time cutoff with a RAM limit of 50 GB. The results presented in Figure 8 show that using
469 additional constraints to account for SCC structure dramatically reduces the time and memory needed
470 to compute equivalent classes for undersampled graphs. Furthermore, the difference between time
471 and memory requirements to solve for these graphs with and without constraints for SCCs increases
472 for larger graphs as the computational requirements for the latter grow at a much faster pace. This
473 result allows us to handle much larger graphs as shown in Figure 5 of the main paper.

Table 1: Number of SCCs and nodes per SCC of the graphs in the benchmark dataset

Num Nodes	8	9	10	11	12	13	14	15
Num SCCs	2	3	3	3	3	3	3	3
SCC Sizes	4,4	3,3,3	3,3,4	3,4,4	4,4,4	4,4,5	4,5,5	5,5,5

474 C Brief Introduction on clingo and Answer Set Programming (ASP)

475 `clingo` [5] combines a grounder `gringo` and a solver `clasp`. `clingo` is a declarative programming
476 system based on logic programs and their answer sets, used to accelerate solutions of computationally
477 involved combinatorial problems. The grounder converts all parts of a `clingo` program to “atoms,”
478 (grounds the statements) and the solver finds “stable models.” In ASP, the answer set is a model in

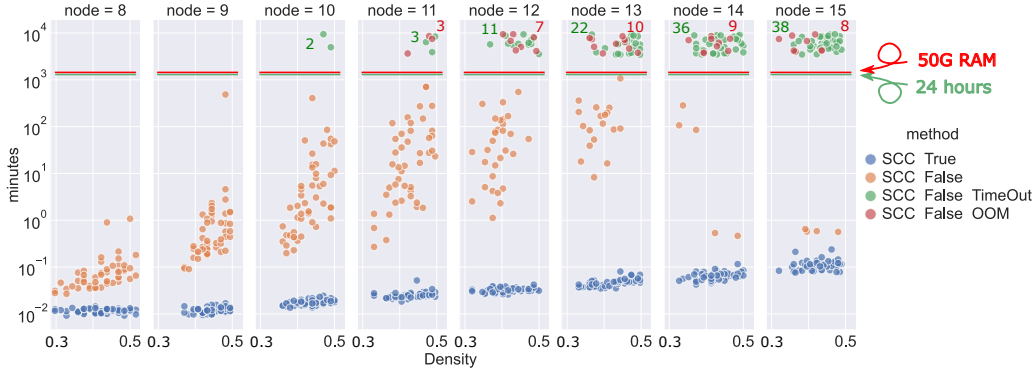


Figure 7: Time behavior of the same set of graphs when solved with and without accounting for additional constraints accounting for the SCC structure. While sRSL most of the 15-node graphs in a 24 hours period without the SCC constraints due to either timeout or Out Of Memory error(OOM), the longest it takes to solve a 15-node graph with SCC constraints is 14 seconds. None of the graphs failed to compute the complete equivalence class within the time and memory allocated when solved accounting for the SCC structure.




Figure 8: A knowledge of a definite presence of an edge in \mathcal{G}^1 between, for example, nodes 3 and 4, i.e. $V_3^1 \rightarrow V_4^{+1}$, can be easily encoded by adding ‘`edge1(3,4).`’ to Listing 1. In this experiment, we have added knowledge about a pair of arbitrary selected edges of \mathcal{G}^1 to the problem specification (orange dots) and compared the run time with the ASP specification that does not include this additional information about the solution (blue dots). The time out for the new computation was set to 1 hours and the examples were all the same as the ones already shown in Figure 1. The speed up with the additional constraints is clearly visible on the plots.

479 which all the atoms are derived from the program and each “answer” is a stable model where all the
 480 atoms are simultaneously true.

481 A general `clingo` program includes three main sections, which we show below using our algorithm
 482 as an example:

483 1. **Facts:** these are the known elements of the problem. For example, the input to Listing 1 is a
 484 graph for which we know the edges. A directed edge from node 1 to node 5 is in \mathcal{H} translates to
 485 `hdirected(1,5)` (line 1) or if node 1 is part of the SCC number 2, we state this fact in `clingo` by
 486 `scc(1,2)` (line 2).

487 2. **Rules:** much like an `if-else` statement, a rule in `clingo` consists of a body and a head, formatted
 488 as `head :- body`. If all the literals in the body are true, then the head must also be true. Rules can
 489 include variables (starting with capital letters), and they are used to derive new facts after grounding.
 490 For example:

$$\text{directed}(X, Y, 1) \text{ :- } \text{edge1}(X, Y). \quad (2)$$

491 means that for any instantiations of the variables X and Y , if we have an edge from X to Y , there is a
 492 directed path from X to Y of length 1. Before this line, if the model contained the fact `edge1(2,3)`,
 493 this line would generate a new fact: `directed(2,3,1)`.

494 Another type of rule is the “choice rule” that describes all the possible ways to choose which atoms
 495 are included in the model. For example, in line 5 of Listing 1 we used a choice rule to state that the

496 undersampling rate u can be anything from 1 to $maxu$. The cardinality constraint:

$$\{u(1..20)\}. \quad (3)$$

497 will generate 2^{20} different models (they will not all actually be generated if they conflict with other
 498 predicate in each model, or else it would not be possible). In each of these 2^{20} models, one subset
 499 of all possible atoms generated with this choice rule exists ($\phi, \{u(1)\}, \{u(1), u(2)\}, \dots$). An
 500 example of an unconstrained choice rule is line 6 in Listing 1, where we want to generate one model
 501 for each possible way edges can be present in a graph between two nodes X and Y . We can also limit
 502 the choice rule. In our problem, only one undersampling rate is present at each solution. We limit the
 503 cardinality constraint to have only one member in each model:

$$1 \{u(1..20)\} 1. \quad (4)$$

504 the 1 on the left is the minimum instantiations of this atom in the model and the 1 on the right is
 505 the maximum. Therefore, we only generate $\binom{20}{1} = 20$ models with this rule, namely one for each
 506 undersampling rate. Having several choice rules will multiply the number of generated models by
 507 each choice rule.

508 **3. Integrity Constraints:** if choice rules are to generate new models, integrity constraints are there
 509 to remove the wrong models from the answers set. More specifically, an integrity constraint is of the
 510 form:

$$:- L_0, L_1, \dots. \quad (5)$$

511 where literals L_0, L_1, \dots cannot be simultaneously positive. For example, in line 16 of Listing 1, we
 512 have:

$$:- \text{edge1}(X, Y), \text{scc}(X, K), \text{scc}(Y, L), K \neq L, \quad (6)$$

$$\text{sccsize}(L, Z), Z > 1, \text{not dag}(K, L).$$

513 for cases where the graph consists of several SCCs that are connected using a DAG. If the SCCs are
 514 connected by a cyclic directed graph, then the whole graph will become one big Strongly Connected
 515 Component. Integrity constraint 6 states that if there is not a directed edge from a node in SCC K to a
 516 node in SCC L as part of the initial DAG, there cannot be such $\text{edge1}(X, Y)$ from node X to node
 517 Y, if node X is in SCC K and node Y is in SCC L.