
The Complexity of Bayesian Network Learning: Revisiting the Superstructure (Full Version)

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 We investigate the parameterized complexity of Bayesian Network Structure Learn-
2 ing (BNSL), a classical problem that has received significant attention in empirical
3 but also purely theoretical studies. We follow up on previous works that have
4 analyzed the complexity of BNSL w.r.t. the so-called *superstructure* of the input.
5 While known results imply that BNSL is unlikely to be fixed-parameter tractable
6 even when parameterized by the size of a vertex cover in the superstructure, here we
7 show that a different kind of parameterization—notably by the size of a feedback
8 edge set—yields fixed-parameter tractability. We proceed by showing that this
9 result can be strengthened to a localized version of the feedback edge set, and
10 provide corresponding lower bounds that complement previous results to provide a
11 complexity classification of BNSL w.r.t. virtually all well-studied graph parameters.
12 We then analyze how the complexity of BNSL depends on the representation of the
13 input. In particular, while the bulk of past theoretical work on the topic assumed
14 the use of the so-called *non-zero representation*, here we prove that if an *additive*
15 *representation* can be used instead then BNSL becomes fixed-parameter tractable
16 even under significantly milder restrictions to the superstructure, notably when
17 parameterized by the treewidth alone. Last but not least, we show how our results
18 can be extended to the closely related problem of Polytree Learning.

19 1 Introduction

20 Bayesian networks are among the most prominent graphical models for probability distributions. The
21 key feature of Bayesian networks is that they represent conditional dependencies between random
22 variables via a directed acyclic graph; the vertices of this graph are the variables, and an arc ab means
23 that the distribution of variable b depends on the value of a . One beneficial property of Bayesian
24 networks is that they can be used to infer the distribution of random variables in the network based
25 on the values of the remaining variables.

26 The problem of constructing a Bayesian network with an optimal network structure is NP-hard, and
27 remains NP-hard even on highly restricted instances [5]. This initial negative result has prompted
28 an extensive investigation of the problem’s complexity, with the aim of identifying new tractable
29 fragments as well as the boundaries of its intractability [29, 36, 30, 25, 14, 9, 22]. The problem—
30 which we simply call BAYESIAN NETWORK STRUCTURE LEARNING (BNSL)—can be stated as
31 follows: given a set of V of variables (represented as vertices), a family \mathcal{F} of *score functions* which
32 assign each variable $v \in V$ a score based on its *parents*, and a target value ℓ , determine if there exists
33 a directed acyclic graph over V that achieves a total score of at least ℓ^1 .

¹Formal definitions are provided in Section 2. We consider the decision version of BNSL for complexity-theoretic reasons only; all of the provided algorithms are constructive and can output a network as a witness.

34 To obtain a more refined understanding of the complexity of BNSL, past works have analyzed the
 35 problem not only in terms of classical complexity but also from the perspective of *parameterized*
 36 *complexity* [12, 8]. In parameterized complexity analysis, the tractability of problems is measured
 37 with respect to the input size n and additionally with respect to a specified numerical *parameter* k . In
 38 particular, a problem that is NP-hard in the classical sense may—depending on the parameterization
 39 used—be *fixed-parameter tractable* (FPT), which is the parameterized analogue of polynomial-time
 40 tractability and means that a solution can be found in time $f(k) \cdot n^{\mathcal{O}(1)}$ for some computable function
 41 f , or *W[1]-hard*, which rules out fixed-parameter tractability under standard complexity assumptions.
 42 The use of parameterized complexity as a refinement of classical complexity is becoming increasingly
 43 common and has been employed not only for BNSL [29, 36, 30], but also for numerous other
 44 problems arising in the context of neural networks and artificial intelligence [16, 44, 13, 19].

45 Unfortunately, past complexity-theoretic works have shown that BNSL is a surprisingly difficult
 46 problem. In particular, not only is the problem NP-hard, but it remains NP-hard even when asking
 47 for the existence of extremely simple networks such as directed paths [33] and is W[1]-hard when
 48 parameterized by the *vertex cover number* of the network [30]. In an effort to circumvent these lower
 49 bounds, several works have proposed to instead consider restrictions to the so-called *superstructure*,
 50 which is a graph that, informally speaking, captures all potential dependencies between variables [45,
 51 38]. Ordyniak and Szeider [36] studied the complexity of BNSL when parameterized by the
 52 structural properties of the superstructure, and showed that parameterizing by the *treewidth* [39]
 53 of the superstructure is sufficient to achieve a weaker notion of tractability called *XP-tractability*.
 54 However, they also proved that BNSL remains W[1]-hard when parameterized by the treewidth of
 55 the superstructure [36, Theorem 3].

56 **Contribution.** Up to now, no “implicit” restrictions of the superstructure were known to lead
 57 to a fixed-parameter algorithm for BNSL alone. More precisely, the only known fixed-parameter
 58 algorithms for the problem require that we place explicit restrictions on either the sought-after network
 59 or the parent sets on the input: BNSL is known to be fixed-parameter tractable when parameterized
 60 by the number of arcs in the target network [25], the treewidth of an “*extended superstructure graph*”
 61 which also bounds the maximum number of parents a variable can have [29], or the number of
 62 parent set candidates plus the treewidth of the superstructure [36]. Moreover, a closer analysis of the
 63 reduction given by Ordyniak and Szeider [36, Theorem 3] reveals that BNSL is also W[1]-hard when
 64 parameterized by the *treedepth*, *pathwidth*, and even the *vertex cover number* of the superstructure
 65 alone. The vertex cover number is equal to the vertex deletion distance to an edgeless graph, and
 66 hence their result essentially rules out the use of the vast majority of graph parameters; among others,
 67 any structural parameter based on vertex deletion distance.

68 As our first conceptual contribution, we show that a different kind of graph parameters—notably,
 69 parameters that are based on edge deletion distance—give rise to fixed-parameter algorithms for
 70 BNSL in its full generality, without requiring any further explicit restrictions on the target network
 71 or parent sets. Our first result in this direction concerns the *feedback edge number* (fen), which is the
 72 minimum number of edges that need to be deleted to achieve acyclicity. In Theorem 3 we show not
 73 only that BNSL is fixed-parameter tractable when parameterized by the fen of the superstructure, but
 74 also provide a polynomial-time preprocessing algorithm that reduces any instance of BNSL to an
 75 equivalent one whose number of variables is linear in the fen (i.e., a *kernelization* [12, 8]).

76 Since fen is a highly “restrictive” parameter—its value can be large even on simple superstructures
 77 such as collections of disjoint cycles—we proceed by asking whether it is possible to lift fixed-
 78 parameter tractability to a more relaxed way of measuring distance to acyclicity. For our second
 79 result, we introduce the *local feedback edge number* (lfen), which intuitively measures the maximum
 80 edge deletion distance to acyclicity for cycles intersecting any particular vertex in the superstructure.
 81 In Theorem 6, we show that BNSL is also fixed-parameter tractable when parameterized by lfen; we
 82 also show that this comes at the cost of BNSL not admitting any polynomial-time preprocessing
 83 procedure akin to Theorem 3 when parameterized by lfen. We conclude our investigation in the
 84 direction of parameters based on edge deletion distance by showing that BNSL parameterized by
 85 *treecut width* [32, 48, 17], a recently discovered edge-cut based counterpart to treewidth, remains
 86 W[1]-hard (Theorem 10). An overview of these complexity-theoretic results is provided in Figure 1.

87 As our second conceptual contribution, we show that BNSL becomes significantly easier when one
 88 can use an *additive representation* of the scores rather than the *non-zero representation* that was
 89 considered in the vast majority of complexity-theoretic works on BNSL to date [29, 36, 30, 25, 14, 22].




Figure 1: The complexity landscape of BNSL with respect to parameterizations of the superstructure. Arrows point from more restrictive parameters to less restrictive ones. Results for the three graph parameters on the left side follow from this paper, while all other $W[1]$ -hardness results follow from the reduction by Ordyniak and Szeider [36, Theorem 3].

90 The additive representation is inspired by known heuristics for BNSL [43, 42] and utilizes a succinct
 91 encoding of the score function which assumes that the scores for parent sets can be decomposed into
 92 a sum of the scores of individual variables in the parent set; a discussion and formal definitions are
 93 provided in Section 2. In Theorem 13, we show that if the additive representation can be used, BNSL
 94 becomes fixed-parameter tractable when parameterized by the treewidth of the superstructure (and
 95 hence under every parameterization depicted in Figure 1). Motivated by the empirical usage of the
 96 additive representation, we also consider the case where we additionally impose a bound q on the
 97 number of parents a vertex can accept; we show that the result of Theorem 13 also covers this case if
 98 q is taken as an additional parameter, and otherwise rule out fixed-parameter tractability using an
 99 intricate reduction (Theorem 15).

100 For our third and final conceptual contribution, we show how our results can be adapted for the
 101 emergent problem of POLYTREE LEARNING (PL), a variant of BNSL where we require that the
 102 network forms a polytree. The crucial advantage of such networks is that they allow for a more
 103 efficient solution of the inference task [37, 26], and the complexity of PL has been studied in several
 104 works [24, 22, 41]. We show that all our results for BNSL can be adapted to PL, albeit in some cases
 105 it is necessary to perform non-trivial modifications. Furthermore, we observe that unlike BNSL,
 106 PL becomes polynomial-time tractable when the additive representation is used (Observation 20);
 107 this matches the “naive” expectation that learning simple networks would be easier than BNSL
 108 in its full generality. As our concluding result, we show that this expectation is in fact not always
 109 validated: while PL was recently shown to be $W[1]$ -hard when parameterized by the number of
 110 so-called *dependent vertices* [24], in Theorem 21 we prove that BNSL is fixed-parameter tractable
 111 under that same parameterization.

112 2 Preliminaries

113 For an integer i , we let $[i] = \{1, 2, \dots, i\}$ and $[i]_0 = [i] \cup \{0\}$. We denote by \mathbb{N} the set of natural
 114 numbers, by \mathbb{N}_0 the set $\mathbb{N} \cup \{0\}$.

115 We refer to the handbook by Diestel [11] for standard graph terminology. In this paper, we will
 116 consider directed as well as undirected simple graphs. If $G = (V, E)$ is an undirected graph and
 117 $\{v, w\} \in E$, we will often use vw as shorthand for $\{v, w\}$; we will also sometimes use $V(G)$ to
 118 denote its vertex set. Moreover, we let $N_G(v)$ denote the set of *neighbors* of v , i.e., $\{u \in V \mid uv \in E\}$. We extend this notation to sets as follows: $N_G(X) = \{u \in V \setminus X \mid \exists x \in X : ux \in E(G)\}$. For
 119 a set X of vertices, let A_X denote the set of all possible arcs over X .
 120

121 If $D = (V, A)$ is a directed graph (i.e., a *digraph*) and $(v, w) \in A$, we will similarly use vw as
 122 shorthand for (v, w) . We also let $P_D(v)$ denote the set of *parents* of v , i.e., $\{u \in V \mid uv \in A\}$
 123 (there are sometimes called *in-neighbors* in the literature, while the notion of *out-neighbors* is defined
 124 analogously). In both cases, we may drop G or D from the subscript if the (di)graph is clear from the
 125 context. The *degree* of v is $|N(v)|$, and for digraphs we use the notions of *in-degree* (which is equal
 126 to $|P(v)|$) and *out-degree* (the number of arcs originating from the given vertex).

127 The *skeleton* (sometimes called the *underlying undirected graph*) of a digraph $G = (V, A)$ is the
 128 undirected graph $G' = (V, E)$ such that $vw \in E$ if $vw \in A$ or $wv \in A$. A digraph is a *polytree* if its
 129 skeleton is a forest.

130 When comparing two numerical parameters α, β of graphs, we say that α is more *restrictive* than β if
 131 there exists a function f such that $\beta(G) \leq f(\alpha(G))$ holds for every graph G . In other words, α is

132 more restrictive than β if and only if the following holds: whenever all graphs in some graph class
 133 \mathcal{H} have α upper-bounded by a constant, all graphs in \mathcal{H} also have β upper-bounded by a constant.
 134 Observe that in this case a fixed-parameter algorithm parameterized by β immediately implies a
 135 fixed-parameter algorithm parameterized by α , while $W[1]$ -hardness behaves in the opposite way.

136 **Problem Definitions.** Let V be a set of vertices and $\mathcal{F} = \{f_v : 2^{V \setminus \{v\}} \rightarrow \mathbb{N}_0 \mid v \in V\}$
 137 be a family of *local score functions*. For a digraph $D = (V, A)$, we define its score as follows:
 138 $\text{score}(D) = \sum_{v \in V} f_v(P_D(v))$, where $P_D(v)$ is the set of vertices of D with an outgoing arc to v
 139 (i.e., the *parent set* of v in D). We can now formalize our problem of interest [36, 25].

BAYESIAN NETWORK STRUCTURE LEARNING (BNSL)

140 **Input:** A set V of vertices, a family \mathcal{F} of local score functions, and an integer ℓ .
Question: Does there exist an acyclic digraph $D = (V, A)$ such that $\text{score}(D) \geq \ell$?

141 POLYTREE LEARNING (PL) is defined analogously, with the only difference that there D is addition-
 142 ally required to be a polytree [24]. We call D a *solution* for the given instance.

143 Since both V and \mathcal{F} are assumed to be given on the input of our problems, an issue that arises here
 144 is that an explicit representation of \mathcal{F} would be exponentially larger than $|V|$. A common way to
 145 potentially circumvent this is to use a *non-zero representation* of the family \mathcal{F} , i.e., where we only
 146 store values for $f_v(P)$ that are different than zero. This model has been used in a large number of
 147 works studying the complexity of BNSL and PL [29, 36, 30, 25, 22, 24] and is known to be strictly
 148 more general than, e.g., the bounded-arity representation where one only considers parent sets of
 149 arity bounded by a constant [36, Section 3]. Let $\Gamma_f(v)$ be the set of candidate parents of v which
 150 yield a non-zero score; formally, $\Gamma_f(v) = \{Z \mid f_v(Z) \neq 0\}$, and the input size $|\mathcal{I}|$ of an instance
 151 $\mathcal{I} = (V, \mathcal{F}, \ell)$ is simply defined as $|V| + \ell + \sum_{v \in V, P \in \Gamma_f(v)} |P|$.

152 Let $P_{\rightarrow}(v)$ be the set of all parents which appear in $\Gamma_f(v)$, i.e., $a \in P_{\rightarrow}(v)$ if and only if $\exists Z \in$
 153 $\Gamma_f(v) : a \in Z$. A natural way to think about and exploit the structure of inter-variable dependencies
 154 laid bare by the non-zero representation is to consider the *superstructure graph* $G_{\mathcal{I}} = (V, E)$ of a
 155 BNSL (or PL) instance $\mathcal{I} = (V, \mathcal{F}, \ell)$, where $ab \in E$ if and only if either $a \in P_{\rightarrow}(b)$, or $b \in P_{\rightarrow}(a)$,
 156 or both.

157 Naturally, families of local score functions may be exponentially larger than $|V|$ even when stored
 158 using the non-zero representation. In this paper, we also consider a second representation of \mathcal{F}
 159 which is guaranteed to be polynomial in $|V|$: in the *additive representation*, we require that for
 160 every vertex $v \in V$ and set $Q = \{q_1, \dots, q_m\} \subseteq V \setminus \{v\}$, $f_v(Q) = f_v(\{q_1\}) + \dots + f_v(\{q_m\})$.
 161 Hence, each cost function f_v can be fully characterized by storing at most $|V|$ -many entries of the
 162 form $f_v(x) := f_v(\{x\})$ for each $x \in V \setminus \{v\}$. To avoid overfitting, one may optionally impose an
 163 additional constraint: an upper bound q on the size of any parent set in the solution (or, equivalently, q
 164 is a maximum upper-bound on the in-degree of the sought-after acyclic digraph D).

165 While not every family of local score functions admits an additive representation, the additive model
 166 is similar in spirit to the models used by some practical algorithms for BNSL. For instance, the
 167 algorithms of Scanagatta, de Campos, Corani and Zaffalon [43, 42], which can process BNSL
 168 instances with up to thousands of variables, approximate the real score functions by adding up the
 169 known score functions for two parts of the parent set and applying a small, logarithmic correction.
 170 Both of these algorithms also use the aforementioned bound q for the parent set size. In spite of this
 171 connection to practice and the representation's streamlined nature, we are not aware of any prior
 172 works that considered the additive representation in complexity-theoretic studies of BNSL and PL.

173 As before, in the additive representation we will also only store scores for parents of v which yield a
 174 non-zero score, and can thus define $P_{\rightarrow}(v) = \{z \mid f_v(z) \neq 0\}$, as for the non-zero representation.
 175 This in turn allows us to define the superstructure graphs in an analogous way as before: $G_{\mathcal{I}} = (V, E)$
 176 where $ab \in E$ if and only if $a \in P_{\rightarrow}(b)$, $b \in P_{\rightarrow}(a)$, or both.

177 To distinguish between these models, we use $\text{BNSL}^{\neq 0}$, BNSL^+ , and BNSL_{\leq}^+ to denote BAYESIAN
 178 NETWORK STRUCTURE LEARNING with the non-zero representation, the additive representation,
 179 and the additive representation and the parent set size bound q , respectively. The same notation will
 180 also be used for POLYTREE LEARNING—for example, an instance of PL_{\leq}^+ will consist of V , a family

181 \mathcal{F} of local score functions in the additive representation, and integers ℓ, q , and the question is whether
 182 there exists a polytree $D = (V, A)$ with in-degree at most q and $\text{score}(D) \geq \ell$.

183 In our algorithmic results, we will often use $G = (V, E)$ to denote the superstructure graph of
 184 the input instance \mathcal{I} . Without any loss of generality, we will also assume that G is connected.
 185 Indeed, given an algorithm \mathbb{A} that solves BNSL on connected instances, we may solve disconnected
 186 instances of BNSL by using \mathbb{A} to find the maximum score ℓ_C for each connected component C of G
 187 independently, and we may then simply compare $\sum_{C \text{ is a connected component of } G} \ell_C$ with ℓ .

188 **Parameterized Complexity.** In parameterized algorithmics [8, 12, 35] the running-time of an
 189 algorithm is studied with respect to a parameter $k \in \mathbb{N}_0$ and input size n . The basic idea is to find
 190 a parameter that describes the structure of the instance such that the combinatorial explosion can
 191 be confined to this parameter. In this respect, the most favorable complexity class is FPT (*fixed-*
 192 *parameter tractable*) which contains all problems that can be decided by an algorithm running in
 193 time $f(k) \cdot n^{\mathcal{O}(1)}$, where f is a computable function. Algorithms with this running-time are called
 194 *fixed-parameter algorithms*. A less favorable outcome is an XP *algorithm*, which is an algorithm
 195 running in time $\mathcal{O}(n^{f(k)})$; problems admitting such algorithms belong to the class XP.

196 Showing that a problem is W[1]-hard rules out the existence of a fixed-parameter algorithm under
 197 the well-established assumption that $\text{W}[1] \neq \text{FPT}$. This is usually done via a *parameterized*
 198 *reduction* [8, 12] to some known W[1]-hard problem. A parameterized reduction from a parameterized
 199 problem \mathcal{P} to a parameterized problem \mathcal{Q} is a function:

- 200 • which maps **Yes**-instances to **Yes**-instances and **No**-instances to **No**-instances,
- 201 • which can be computed in time $f(k) \cdot n^{\mathcal{O}(1)}$, where f is a computable function, and
- 202 • where the parameter of the output instance can be upper-bounded by some function of the
 203 parameter of the input instance.

204 **Treewidth.** A *nice tree-decomposition* \mathcal{T} of a graph $G = (V, E)$ is a pair (T, χ) , where T is a tree
 205 (whose vertices we call *nodes*) rooted at a node r and χ is a function that assigns each node t a set
 206 $\chi(t) \subseteq V$ such that the following holds:

- 207 • For every $uv \in E$ there is a node t such that $u, v \in \chi(t)$.
- 208 • For every vertex $v \in V$, the set of nodes t satisfying $v \in \chi(t)$ forms a subtree of T .
- 209 • $|\chi(\ell)| = 1$ for every leaf ℓ of T and $|\chi(r)| = 0$.
- 210 • There are only three kinds of non-leaf nodes in T :
 - 211 – **Introduce node:** a node t with exactly one child t' such that $\chi(t) = \chi(t') \cup \{v\}$ for
 212 some vertex $v \notin \chi(t')$.
 - 213 – **Forget node:** a node t with exactly one child t' such that $\chi(t) = \chi(t') \setminus \{v\}$ for some
 214 vertex $v \in \chi(t')$.
 - 215 – **Join node:** a node t with two children t_1, t_2 such that $\chi(t) = \chi(t_1) = \chi(t_2)$.

216 The *width* of a nice tree-decomposition (T, χ) is the size of a largest set $\chi(t)$ minus 1, and the
 217 *treewidth* of the graph G , denoted $\text{tw}(G)$, is the minimum width of a nice tree-decomposition
 218 of G . Fixed-parameter algorithms are known for computing a nice tree-decomposition of optimal
 219 width [4, 27]. For $t \in V(T)$ we denote by T_t the subtree of T rooted at t .

220 **Graph Parameters Based on Edge Cuts.** Traditionally, the bulk of graph-theoretic research on
 221 structural parameters has focused on parameters that guarantee the existence of small vertex separators
 222 in the graph; these are inherently tied to the theory of *graph minors* [40, 39] and the vertex deletion
 223 distance. This approach gives rise not only to the classical notion of treewidth, but also to its
 224 well-known restrictions and refinements such as *pathwidth* [40], *treedepth* [34] and the *vertex cover*
 225 *number* [15, 28]. The vertex cover number is the most restrictive parameter in this hierarchy.

226 However, there are numerous problems of interest that remain intractable even when parameterized
 227 by the vertex cover number. A recent approach developed for attacking such problems has been to
 228 consider parameters that guarantee the existence of small edge cuts in the graph; these are typically
 229 based on the edge deletion distance or, more broadly, tied to the theory of *graph immersions* [48, 32].
 230 The parameter of choice for the latter is *treecut width* (tcw) [48, 32, 17, 18], a counterpart to
 231 treewidth which has been successfully used to tackle some problems that remained intractable when

232 parameterized by the vertex cover number [20]. For the purposes of this manuscript, it will be useful
 233 to note that graphs containing a vertex cover X such that every vertex outside of X has degree at
 234 most 2 have treecut width at most $|X|$ [20, Section 3].

235 On the other hand, the by far most prominent parameter based on edge deletion distance is the
 236 *feedback edge number* of a connected graph $G = (V, E)$, which is the minimum cardinality of a
 237 set $F \subseteq E$ of edges (called the *feedback edge set*) such that $G - F$ is acyclic. The feedback edge
 238 number can be computed in quadratic time and has primarily been used to obtain fixed-parameter
 239 algorithms and polynomial kernels for problems where other parameterizations failed [20, 3, 2, 47].

240 Up to now, these were the only two edge-cut based graph parameters that have been considered in
 241 the broader context of algorithm design. This situation could be seen as rather unsatisfactory in view
 242 of the large gap between the complexity of the richer class of graphs of bounded treecut width, and
 243 the significantly simpler class of graphs of bounded feedback edge number—for instance, the latter
 244 class is not even closed under disjoint union. Here, we propose a new parameter that lies “between”
 245 the feedback edge number and treecut width, and which can be seen as a localized relaxation of the
 246 feedback edge number: instead of measuring the total size of the feedback edge set, it only measures
 247 how many feedback edges can “locally interfere with” any particular part of the graph.

248 Formally, for a connected graph $G = (V, E)$ and a spanning tree T of G , let the *local feedback edge*
 249 *set* at $v \in V$ be

$$E_{\text{loc}}^T(v) = \{uv \in E \setminus E(T) \mid \text{the unique path between } u \text{ and } v \text{ in } T \text{ contains } v\}.$$

250 The *local feedback edge number* of (G, T) (denoted $\text{lfn}(G, T)$) is then equal to $\max_{v \in V} |E_{\text{loc}}^T(v)|$,
 251 and the *local feedback edge number* of G is simply the smallest local feedback edge number among
 252 all possible spanning trees of G , i.e., $\text{lfn}(G) = \min_{T \text{ is a spanning tree of } G} \text{lfn}(G, T)$.

253 It is not difficult to show that the local feedback edge number is “sandwiched” between the feedback
 254 edge number and treecut width. We also show that computing it is FPT.

255 **Proposition 1.** *For every graph G , $\text{tcw}(G) \leq \text{lfn}(G) + 1$ and $\text{lfn}(G) \leq \text{fen}(G)$.*

256 *Proof.* Let us begin with the second inequality. Consider an arbitrary spanning tree T of G . Then for
 257 every $v \in V(G)$, $E_{\text{loc}}^T(v)$ is a subset of a feedback edge set corresponding to the spanning tree T , so
 258 $|E_{\text{loc}}^T(v)| \leq \text{fen}(G)$ and the claim follows.

259 To establish the first inequality, we will use the notation and definition of treecut width from previous
 260 work [18, Subsection 2.4]. Let T be the spanning tree of G with $\text{lfn}(G, T) = \text{lfn}(G)$. We construct
 261 a treecut decomposition (T, \mathcal{X}) where each bag contains precisely one vertex, notably by setting
 262 $X_t = \{t\}$ for each $t \in V(T)$. Fix any node t in T other than root, let u be the parent of t in T . All
 263 the edges in $G \setminus ut$ with one endpoint in the rooted subtree T_t and another outside of T_t belong to
 264 $E_{\text{loc}}^T(t)$, so $\text{adh}_T(t) = |\text{cut}(t)| \leq |E_{\text{loc}}^T(t)| \leq \text{lfn}(G)$.

265
 266 Let H_t be the torso of (T, \mathcal{X}) in t , then $V(H_t) = \{t, z_1 \dots z_l\}$ where z_i correspond to connected
 267 components of $T \setminus t$, $i \in [l]$. In $\tilde{H}(t)$, only z_i with degree at least 3 are preserved. But all such z_i are
 268 the endpoints of at least 2 edges in $|E_{\text{loc}}^T(t)|$, so $\text{tor}(t) = |V(\tilde{H}_t)| \leq 1 + |E_{\text{loc}}^T(t)| \leq 1 + \text{lfn}(G)$.
 269 Thus $\text{tcw}(G) \leq \text{lfn}(G) + 1$. \square

270 **Theorem 2.** *The problem of determining whether $\text{lfn}(G) \leq k$ for an input graph G parameterized*
 271 *by an integer k is fixed-parameter tractable. Moreover, if the answer is positive, we may also output*
 272 *a spanning tree T such that $\text{lfn}(G, T) \leq k$ as a witness.*

273 *Proof.* Observe that since $\text{tcw}(G) \leq \text{lfn}(G) + 1$ by Proposition 1 and $\text{tw}(G) \leq 2 \text{tcw}(G)^2 +$
 274 $3 \text{tcw}(G)$ [17], we immediately see that no graph of treewidth greater than $k' = 2k^2 + 5k + 3$ can
 275 have a local feedback edge set of at most k . Hence, let us begin by checking that $\text{tw}(G) \leq k'$ using
 276 the classical fixed-parameter algorithm for computing treewidth [4]; if not, we can safely reject the
 277 instance.

278 Next, we use the fact that $\text{tw}(G) \leq k'$ to invoke Courcelle’s Theorem [6, 12], which provides a
 279 fixed-parameter algorithm for model-checking any *Monadic Second-Order Logic* formula on G when
 280 parameterized by the size of the formula and the treewidth of G . We refer interested readers to the
 281 appropriate books [7, 12] for a definition of Monadic Second Order Logic; intuitively, the logic

282 allows one to make statements about graphs using variables for vertices and edges as well as their
 283 sets, standard logical connectives, set inclusions, and atoms that check whether an edge is incident to
 284 a vertex. If the formula contains a free set variable X and admits a model on G , Courcelle’s Theorem
 285 allows us to also output an interpretation of X on G that satisfies the formula.

286 The formula ϕ we will use to check whether $\text{lfn}(G) \leq k$ will be constructed as follows. ϕ contains
 287 a single free edge set variable X (which will correspond to the sought-after feedback edge set). ϕ
 288 then consists of a conjunction of two parts, where the first part simply ensures that X is a minimal
 289 feedback edge set using a well-known folklore construction [31, 1]; this also ensures that $G - X$ is a
 290 spanning tree. In the second part, ϕ quantifies over all vertices in G , and for each such vertex v it
 291 says there exist edges e_1, \dots, e_k in X such that for every edge $ab \in X$ distinct from all of e_1, \dots, e_k ,
 292 there exists a path P between a and b in $G - X$ which is disjoint from v . (Note that since the path P
 293 is unique in $G - X$, one could also quantify P universally and achieve the same result.)

294 It is easy to verify that $\phi(X)$ is satisfied in G if and only if $\text{lfn}(G, G - X) \leq k$, and so the
 295 proof follows. Finally, we remark that—as with every algorithmic result arising from Courcelle’s
 296 Theorem—one could also use the formula as a template to build an explicit dynamic programming
 297 algorithm that proceeds along a tree-decomposition of G . \square

298 3 Solving $\text{BNSL}^{\neq 0}$ with Parameters Based on Edge Cuts.

299 In this section we provide tractability and lower-bound results for $\text{BNSL}^{\neq 0}$ from the viewpoint of
 300 superstructure parameters based on edge cuts. Together with the previous lower bound that rules
 301 out fixed-parameter algorithms based on all vertex-separator parameters [36, Theorem 3], the results
 302 presented here provide a comprehensive picture of the complexity of $\text{BNSL}^{\neq 0}$ with respect to
 303 superstructure parameterizations.

304 3.1 Using the Feedback Edge Number for $\text{BNSL}^{\neq 0}$

305 We say that two instances $\mathcal{I}, \mathcal{I}'$ of BNSL are *equivalent* if (1) they are either both **Yes**-instances or
 306 both **No**-instances, and furthermore (2) a solution to one instance can be transformed into a solution
 307 to the other instance in polynomial time. Our aim here is to prove the following theorem:

308 **Theorem 3.** *There is an algorithm which takes as input an instance \mathcal{I} of $\text{BNSL}^{\neq 0}$ whose super-*
 309 *structure has $\text{fen } k$, runs in time $\mathcal{O}(|\mathcal{I}|^2)$, and outputs an equivalent instance $\mathcal{I}' = (V', \mathcal{F}', \ell')$ of*
 310 *$\text{BNSL}^{\neq 0}$ such that $|V'| \leq 16k$.*

311 In parameterized complexity theory, such data reduction algorithms with performance guarantees are
 312 called *kernelization algorithms* [12, 8]. These may be applied as a polynomial-time preprocessing
 313 step before, e.g., more computationally expensive methods are used. The fixed-parameter tractability
 314 of $\text{BNSL}^{\neq 0}$ when parameterized by the fen of the superstructure follows as an immediate corollary
 315 of Theorem 3 (one may solve \mathcal{I} by, e.g., exhaustively looping over all possible DAGs on V' via a
 316 brute-force procedure). We also note that even though the number of variables of the output instance
 317 is polynomial in the parameter k , the instance \mathcal{I}' need not have size polynomial in k .

318 We begin our path towards a proof of Theorem 3 by computing a feedback edge set E_F of G of size k
 319 in time $\mathcal{O}(|\mathcal{I}|^2)$ by, e.g., Prim’s algorithm. Let T be the spanning tree of G , $E_F = E(G) \setminus E(T)$. The
 320 algorithm will proceed by the recursive application of certain reduction rules, which are polynomial-
 321 time operations that alter (“simplify”) the input instance in a certain way. A reduction rule is *safe* if it
 322 outputs an instance which is equivalent to the input instance. We start by describing a rule that will
 323 be used to prune T until all leaves are incident to at least one edge in E_F .

Reduction Rule 1. *Let $v \in V$ be a vertex and let Q be the set of neighbors of v with de-*
*gree 1 in G . We construct a new instance $\mathcal{I}' = (V', \mathcal{F}', \ell)$ by setting: **1.** $V' := V \setminus Q$; **2.***
 *$\Gamma_{\mathcal{F}'}(v) := \{\emptyset\} \cup \{(P \setminus Q) \mid P \in \Gamma_f(v)\}$; **3.** for all $w \in V' \setminus \{v\}$, $f'_w = f_w$; **4.** for every*
 $P' \in \Gamma_{\mathcal{F}'}(v)$:

$$f'_v(P') := \max_{P: P \setminus Q = P'} (f_v(P) + \sum_{v_{\text{in}} \in P \cap Q} f_{v_{\text{in}}}(\emptyset) + \sum_{v_{\text{out}} \in Q \setminus P} \max(f_{v_{\text{out}}}(\emptyset), f_{v_{\text{out}}}(v))).$$

324 **Lemma 4.** *Reduction Rule 1 is safe.*

325 *Proof.* For the forward direction, assume that \mathcal{I}' admits a solution D' , and let λ be the score
326 D' achieves on v . By the construction of \mathcal{I}' , there must be a parent set $Z \in \Gamma_f(v)$ such that
327 $Z \cap V' = P_{D'}(v)$ (i.e., Z agrees with v 's parents in D') and λ is the sum of the following scores:
328 (1) $f_v(Z)$, (2) the maximum achievable score for each vertex in $Q \setminus Z$, and (3) the score of $\{\emptyset\}$
329 for each vertex in $Z \cap Q$. Let D be obtained from D' by adding the following arcs: zv for each
330 $z \in Z$, and vq for each $q \in Q \setminus Z$ such that q achieves its maximum score with v as its parent. By
331 construction, $\lambda = \sum_{w' \in \{v\} \cup Q} f_w(P_D(w))$. Since the scores of D and D' coincide on all vertices
332 outside of $\{v\} \cup Q$ and D , we conclude that $\text{score}(D) = \text{score}(D')$, and hence \mathcal{I} is a **Yes**-instance.

333 For the converse direction, assume that \mathcal{I} admits a solution D . Let $D' = D - Q$. By the construction
334 of f'_v , it follows that $f'_v(P_{D'}(v))$ is greater or equal to the score D achieves on $\{v\} \cup Q$. Thus, D' is
335 a solution to \mathcal{I}' , and we conclude that Reduction Rule 1 is safe. \square

336 Observe that the superstructure graph G' obtained after applying one step of Reduction Rule 1 is
337 simply $G - Q$; after its exhaustive application we obtain an instance \mathcal{I} such that all the leaves of the
338 tree T are endpoints of E_F . Our next step is to get rid of long paths in G whose internal vertices
339 have degree 2. We note that this step is more complicated than in typical kernelization results using
340 feedback edge set as the parameter, since a directed path Q in G can serve multiple “roles” in a
341 hypothetical solution D and our reduction gadget needs to account for all of these. Intuitively, Q may
342 or may not appear as a directed path in D (which impacts what other arcs can be used in D due to
343 acyclicity), and in addition the total score achieved by D on the internal vertices of Q needs to be
344 preserved while taking into account whether the endpoints of Q have a neighbor in the path or not.
345 Because of this (and unlike in many other kernelization results of this kind [20, 46, 18]), we will not
346 be replacing Q merely by a shorter path, but by a more involved gadget.

347 **Reduction Rule 2.** Let a, b_1, \dots, b_m, c be a path in G such that for each $i \in [m]$, b_i has degree
348 precisely 2. For each $B \subseteq \{a, c\}$, let $\ell_{\max}(B)$ be the maximum sum of scores that can be achieved by
349 b_1, \dots, b_m under the condition that b_1 (and analogously b_m) takes a (c) into its parent set if and only
350 if $a \in B$ ($c \in B$). In other words, $\ell_{\max}(B) = \max_{D_B} \sum_{b_i | i \in [m]} f_{b_i}(P_{D_B}(b_i))$ where D_B is a DAG
351 on $\{b_1, \dots, b_m\} \cup B$ such that B does not contain any vertices of out-degree 0 in D_B . Moreover, let
352 $\ell_{\text{noPath}}(a)$ (and analogously $\ell_{\text{noPath}}(c)$) be the maximum score that can be achieved on the vertices
353 b_1, \dots, b_m by a DAG on a, b_1, \dots, b_m, c with the following properties: a (c) has out-degree 1, c (a)
354 has out-degree 0, and there is no directed path from a to b_m (from c to b_1).

355 We construct a new instance $\mathcal{I}' = (V', \mathcal{F}', \ell)$ as follows:

- 356 • $V' := V \cup \{b\} \setminus \{b_2 \dots b_{m-1}\}$;
- 357 • $\Gamma_{f'_v}(b) = \{B \cup \{b_1, b_m\} | B \subseteq \{a, c\}\}$ with scores $f'_b(B \cup \{b_1, b_m\}) := \ell_{\max}(B)$;
- 358 • The scores for a and c are obtained from \mathcal{F} by simply adding b to any parent set containing
359 either b_1 or b_m ; formally:
 - 360 – $\Gamma_{f'_v}(a)$ is a union of $\{P \in \Gamma_f(a) | b_1 \notin P\}$, where $f'_a(P) := f_a(P)$ and $\{P \cup \{b\} | b_1 \in$
361 $P, P \in \Gamma_f(a)\}$, where $f'_a(P \cup \{b\}) := f_a(P)$;
 - 362 – $\Gamma_{f'_v}(c)$ is a union of $\{P \in \Gamma_f(c) | b_m \notin P\}$, where $f'_c(P) := f_c(P)$, and $\{P \cup \{b\} | b_m \in$
363 $P, P \in \Gamma_f(c)\}$, where $f'_c(P \cup \{b\}) := f_c(P)$.
- 364 • $\Gamma_{f'_v}(b_1)$ contains only $\{a, b, b_m\}$ with score $\ell_{\text{noPath}}(a)$;
- 365 • $\Gamma_{f'_v}(b_m)$ contains only $\{c, b, b_1\}$ with score $\ell_{\text{noPath}}(c)$;
- 366 • for all $w \in V' \setminus \{a, b_1, b, b_m, c\}$, $f'_w = f_w$.

367 An Illustration of Reduction Rule 2 is provided in Figure 2. The rule can be applied in linear
368 time, since the 6 values of ℓ_{noPath} and ℓ_{\max} can be computed in linear time by a simple dynamic
369 programming subroutine that proceeds along the path a, b_1, \dots, b_m, c (alternatively, one may instead
370 invoke the fact that paths have treewidth 1 [36]).

371 **Lemma 5.** *Reduction Rule 2 is safe.*

372 *Proof.* Note that the superstructure graph of reduced instance is obtained from $G_{\mathcal{I}}$ by contracting
373 $b_2 \dots b_{m-1}$, adding b and connecting it by edges to a, c, b_1, b_m . We will show that a score of at least ℓ




Figure 2:
Top: The six possible scenarios that give rise to the values of ℓ_{max} (Cases 1-4) and ℓ_{noPath} (Cases 5-6).
Bottom: The corresponding arcs in the gadget after the application of Reduction Rule 2.

374 can be achieved in the original instance \mathcal{I} if and only if a score of at least ℓ can be achieved in the
375 reduced instance \mathcal{I}' .

376 Assume that D is a DAG that achieves a score of ℓ in \mathcal{I} . We will construct a DAG D' , called the
377 *reduct* of D , with $f'(D') \geq \ell$. To this end, we first modify D by removing the vertices $b_2 \dots b_{m-1}$
378 and adding b (let us denote the DAG obtained at this point D^*). Further modifications of D^* depend
379 only on $D[a, b_1 \dots b_m, c]$, and we distinguish the 6 cases listed below (see also Figure 2):

- 380 • case 1: D contains both arcs ab_1 and cb_m . We add to D^* arcs from a, c, b_1, b_m to b , denote
381 resulting graph by D' . As D' is obtained from DAG by making b a sink, it is a DAG as well.
382 Parent set of b in D' is $\{a, c, b_1, b_m\}$, so its score is $\ell_{max}(a, c) \geq \sum_{i=1}^m f_{b_i}(P_D(b_i))$, which
383 means that it achieves the highest scores all of b_i 's can achieve in D . The remaining vertices
384 in $V(D') \setminus \{b_1, b_m, b\}$ have the same scores as in D , so $f'(D') \geq f(D) = \ell$.
- 385 • case 2: D contains none of the arcs ab_1 and cb_m . To keep the scores of a and c the same as in
386 D , we add to D^* the arc ba iff D contains b_1a , add arc bc iff D contains $b_m c$. Furthermore,
387 we add arcs b_1b and b_mb and denote resulting graph D' . As D' is obtained from D by
388 making b a source and then adding sources b_1 and b_m , it is a DAG as well. The parent
389 set of b in D' is $\{b_1, b_m\}$, so its score is $\ell_{max}(\emptyset) \geq \sum_{i=1}^m f_{b_i}(P_D(b_i))$. Rest of vertices in
390 $V(D') \setminus \{b_1, b_m, b\}$ have the same scores as in D , so $f'(D') \geq f(D) = \ell$.
- 391 • case 3: D doesn't contain arc ab_1 , but contains cb_m and all the arcs $b_{i+1}b_i, i \in [m-1]$. We
392 add to D^* arcs cb, b_1b and b_mb . We also add ba iff D contains b_1a , to preserve the score
393 of a . Denote resulting graph by D' . D' can be considered as D where long directed path
394 $c \rightarrow b_m \rightarrow \dots \rightarrow b_1$ was replaced by $c \rightarrow b$ and then sources b_1 and b_m were added, so it is
395 a DAG. Arguments for scores are similar to cases 1 and 2.
- 396 • case 4: D doesn't contain arc cb_m , but contains ab_1 and all the arcs $b_i b_{i+1}, i \in [m-1]$.
397 This case is symmetric to case 3.
- 398 • case 5: D contains the arc ab_1 but does not contain the arc cb_m and at least one of the
399 arcs $b_i b_{i+1}, i \in [m-1]$ is also missing (i.e., there is no directed path from a to b_m). We
400 add to D' arcs bb_1 and $b_m b_1$. If $b_m c \in A(D)$, add also bc . Denote the resulting graph
401 D' . As D' is obtained from D^* by making b_1 a sink and b a source, it is a DAG. b_1 has
402 parent set $\{a, b, b_m\}$ in D' , so its score is $\ell_{noPath}(a) \geq \sum_{i=1}^m f_{b_i}(P_D(b_i))$. Rest of vertices
403 in $V(D') \setminus \{b_1, b_m, b\}$ have the same scores as in D , so $f'(D') \geq f(D) = \ell$.
- 404 • case 6: D contains the arc cb_m but does not contain the arc ab_1 and at least one of the arcs
405 $b_{i+1}b_i, i \in [m-1]$ is also missing. This case is symmetric to case 5.

406 The considered cases exhaustively partition all possible configurations of $D[a, b_1 \dots b_m, c]$, so we
 407 always can construct D' with a score at least ℓ . For the converse direction, note that the DAGs
 408 constructed in cases 1-6 cover all optimal configurations on $\{a, b_1, b, b_m, c\}$: if there is a DAG D''
 409 in \mathcal{T} with a score of ℓ' , we can always reverse the construction to obtain a DAG D' with score at
 410 least ℓ' such that $D'[a, b_1, b, b_m, c]$ has one of the forms depicted at the bottom line of the figure. The
 411 claim for the converse direction follows from the fact that every such D' is a reduct of some DAG D
 412 of the original instance with the same score. \square

413 We are now ready to prove the desired result.

414 *Proof of Theorem 3.* We begin by exhaustively applying Reduction Rule 1 on an instance whose
 415 superstructure graph has a feedback edge set of size k , which results in an instance with the same
 416 feedback edge set but whose spanning tree T has at most $2k$ leaves. It follows that there are at most
 417 $2k$ vertices with a degree greater than 2 in T .

418 Let us now “mark” all the vertices that either are endpoints of the edges in E_F or have a degree greater
 419 than 2 in T ; the total number of marked vertices is upper-bounded by $4k$. We now proceed to the
 420 exhaustive application of Reduction Rule 2, which will only be triggered for sufficiently long paths in
 421 T that connect two marked vertices but contain no marked vertices on its internal vertices; there are at
 422 most $4k$ such paths due to the tree structure of T . Reduction Rule 2 will replace each such path with
 423 a set of 3 vertices, and therefore after its exhaustive application we obtain an equivalent instance with
 424 at most $4k + 4k \cdot 3 = 16k$ vertices, as desired. Correctness follows from the safeness of Reduction
 425 Rules 1, 2, and the runtime bound follows by observing that the total number of applications of each
 426 rule as well as the runtime of each rule are upper-bounded by a linear function of the input size. \square

427 3.2 Fixed-Parameter Tractability of $\text{BNSL}^{\neq 0}$ using the Local Feedback Edge Number

428 Our aim here will be to lift the fixed-parameter tractability of $\text{BNSL}^{\neq 0}$ established by Theorem 3 by
 429 relaxing the parameterization to lfe . In particular, we will prove:

430 **Theorem 6.** *$\text{BNSL}^{\neq 0}$ is fixed-parameter tractable when parameterized by the local feedback edge
 431 number of the superstructure.*

432 Since lfe is a more restrictive parameter than lfn , this results in a strictly larger class of instances
 433 being identified as tractable. However, the means we will use to establish Theorem 6 will be funda-
 434 mentally different: we will not use a polynomial-time data reduction algorithm as the one provided
 435 in Theorem 3, but instead apply a dynamic programming approach. Since the kernels constructed
 436 by Theorem 3 contain only polynomially-many variables w.r.t. lfe , that result is incomparable to
 437 Theorem 6.

438 In fact one can use standard techniques to prove that, under well-established complexity assumptions,
 439 a data reduction result such as the one provided in Theorem 3 *cannot* exist for lfn . The intuitive
 440 reason for this is that lfe is a “local” parameter that does not increase by, e.g., performing a disjoint
 441 union of two distinct instances (the same property is shared by many other well-known parameters
 442 such as treewidth, pathwidth, treedepth, clique-width, and treecut width). We provide a formal proof
 443 of this claim at the end of Subsection 3.3.

444 As our first step towards proving Theorem 6, we provide general conditions for when the union of two
 445 DAGs is a DAG as well. Let $D = (V, A)$ be a directed graph and $V' \subseteq V$. Denote by $\text{Con}(V', D)$
 446 the binary relation on $V' \times V'$ which specifies whether vertices from V' are connected by a path in
 447 D : $\text{Con}(V', D) = \{(v_1, v_2) \subseteq V' \times V' \mid \exists \text{ directed path from } v_1 \text{ to } v_2 \text{ in } D\}$. Similarly to arcs, we
 448 will use $v_1 v_2 \in$ as shorthand for (v_1, v_2) ; we will also use trcl to denote the transitive closure.

449 **Lemma 7.** *Let D_1, D_2 be directed graphs with common vertices $V_{\text{com}} = V(D_1) \cap V(D_2)$, $V_{\text{com}} \subseteq$
 450 $V_1 \subseteq V(D_1)$, $V_{\text{com}} \subseteq V_2 \subseteq V(D_2)$. Then:*

- 451 • (i) $\text{Con}(V_1 \cup V_2, D_1 \cup D_2) = \text{trcl}(\text{Con}(V_1, D_1) \cup \text{Con}(V_2, D_2))$;
- 452 • (ii) If D_1, D_2 are DAGs and $\text{Con}(V_1 \cup V_2, D_1 \cup D_2)$ is irreflexive, then $D_1 \cup D_2$ is a DAG.

453 *Proof.* (i) Denote $R_i := \text{Con}(V_i, D_i)$, $i = 1, 2$. Obviously $\text{trcl}(R_1 \cup R_2)$ is a subset of $\text{Con}(V_1 \cup$
 454 $V_2, D_1 \cup D_2)$. Assume that for some $x, y \in V_1 \cup V_2$ there exists a directed path P from x to y in
 455 $D_1 \cup D_2$. We will show (by induction on the length l of shortest P) that $xy \in \text{trcl}(R_1 \cup R_2)$.

- 456 • $l = 1$: in this case there is an arc xy in some D_i , so $xy \in R_i \subseteq \text{trcl}(R_1 \cup R_2)$
- 457 • $l \rightarrow l + 1$. If P is completely contained in some D_i , then $xy \in R_i \subseteq \text{trcl}(R_1 \cup R_2)$.
- 458 Otherwise P must contain arcs $e \notin A(D_1)$, $f \notin A(D_2)$. Then there is $w \in V_{\text{com}} \subseteq V_1 \cup V_2$
- 459 between them. By the induction hypothesis $xw \in \text{trcl}(R_1 \cup R_2)$ and $wy \in \text{trcl}(R_1 \cup R_2)$,
- 460 so $xy \in \text{trcl}(R_1 \cup R_2)$

461 (ii) The precondition implies that the digraph $D_1 \cup D_2$ induced on $V_1 \cup V_2$ is a DAG. Assume that
 462 $D_1 \cup D_2$ is not a DAG and let C be a shortest directed cycle in $D_1 \cup D_2$. As D_1 and D_2 are DAGs, C
 463 must contain arcs $e \notin A(D_1)$, $f \notin A(D_2)$. So there are least 2 different vertices x, y from V_{com} in C .
 464 By (i) we have that $xy \in \text{trcl}(R_1 \cup R_2)$ and $yx \in \text{trcl}(R_1 \cup R_2)$, then also $xx \in \text{trcl}(R_1 \cup R_2)$,
 465 which contradicts irreflexivity. \square

466 Towards proving Theorem 6, assume that we are given an instance $\mathcal{I} = (V, \mathcal{F}, \ell)$ of $\text{BNSL}^{\neq 0}$ with
 467 connected superstructure graph $G = (V, E)$. Let T be a fixed rooted spanning tree of G such that
 468 $\text{lfn}(G, T) = \text{lfn}(G) = k$, denote the root by r . For $v \in V(T)$, let T_v be the subtree of T rooted at
 469 v , let $V_v = V(T_v)$, and let $\bar{V}_v = N_G(V_v) \cup V_v$. We define the *boundary* $\delta(v)$ of v to be the set of
 470 endpoints of all edges in G with precisely one endpoint in V_v (observe that the boundary can never
 471 have a size of 1). v is called *closed* if $|\delta(v)| \leq 2$ and *open* otherwise. We begin by establishing some
 472 basic properties of the local feedback edge set.

473 **Observation 8.** *Let v be a vertex of T . Then:*

- 474 1. For every closed child w of v in T , it holds that $\delta(w) = \{v, w\}$ and vw is the only edge
 475 between V_w and $V \setminus V_w$ in G .
- 476 2. $|\delta(v)| \leq 2k + 2$.
- 477 3. Let $\{v_i | i \in [t]\}$ be the set of all open children of v in T . Then $t \leq 2k$ and
 478 $\delta(v) \subseteq \cup_{i=1}^t \delta(v_i) \cup \{v\} \cup N_G(v)$

479 *Proof.* The first claim follows by the connectivity assumption on G and the definition of boundary.

480 For the second claim, clearly $\delta(r) = \emptyset$. Let $v \neq r$ have the parent u , and consider an arbitrary
 481 $w \in \delta(v) \setminus \{u, v\}$. Then there is an edge $ww' \in E(G)$ with precisely one endpoint in V_v and
 482 $ww' \neq uv$. Hence $ww' \notin E(T)$ and the path between w and w' in T contains v , and this implies
 483 $ww' \in E_{\text{loc}}^T(v)$ by definition. Consequently, $w \in V_{\text{loc}}^T(v)$. For the claimed bound we note that
 484 $|V_{\text{loc}}^T(v)| \leq 2|E_{\text{loc}}^T(v)| \leq 2k$.

485 For the third claim, let $w = v_i$ for some $i \in [t]$. As w is open, there exists an edge $e \neq vw$ between
 486 V_w and $V \setminus V_w$ in G . By definition of local feedback edge set, $e \in E_{\text{loc}}^T(v)$. Let x_w be the endpoint
 487 of e that belongs to V_w , then $x_w \in V_{\text{loc}}^T(v)$ and $x_w \notin V_{w'}$ for any open child $w' \neq w$ of v . But
 488 $|V_{\text{loc}}^T(v)| \leq 2k$, which yields the bound on number t of open children.

489 For the boundary inclusion, consider any edge c in G with precisely one endpoint x_v in V_v . Note that
 490 x_v can not belong to V_w for any closed child w of v . If $x_v \in V_{v_i}$ for some $i \in [t]$, then endpoints of c
 491 belong to $\delta(v_i)$. Otherwise $x_v = v$ and therefore the second endpoint of c is in $N_G(v)$. \square

492 With Observation 8 in hand, we can proceed to a definition of the records used in our dynamic
 493 program. Intuitively, these records will be computed in a leaf-to-root fashion and will store at each
 494 vertex v information about the best score that can be achieved by a partial solution that intersects the
 495 subtree rooted at v .

496 Let R be a binary relation on $\delta(v)$ and s an integer. For $s \in \mathbb{Z}$, we say that $(R : s)$ is a *record* for a
 497 vertex v if and only if there exists a DAG D on \bar{V}_v such that (1) $w \in V_v$ for each arc $uw \in A(D)$, (2)
 498 $R = \text{Con}(\delta(v), D)$ and (3) $\sum_{u \in V_v} f_u(P_D(u)) = s$. The records (R, s) where s is maximal for fixed
 499 R are called *valid*. Denote the set of all valid records for v by $\mathcal{R}(v)$, and note that $|\mathcal{R}(v)| \leq 2^{\mathcal{O}(k^2)}$.

500 Observe that if v_i is a closed child of v , then by Observation 8.1 $\mathcal{R}(v_i)$ consists of precisely two
 501 valid records: one for $R = \emptyset$ and one for $R = \{vv_i\}$. Moreover, the root r of T has only a single
 502 valid record $(\emptyset : s_{\mathcal{I}})$, where $s_{\mathcal{I}}$ is the maximum score that can be achieved by a solution in \mathcal{I} . The
 503 following lemma lies at the heart of our result and shows how we can compute our records in a
 504 leaf-to-root fashion along T .

505 **Lemma 9.** *Let $v \in V(G)$ have m children in T where $m > 0$, and assume we have computed $\mathcal{R}(v_i)$*
 506 *for each child v_i of v . Then $\mathcal{R}(v)$ can be computed in time at most $m \cdot |\Gamma_f(v)| \cdot 2^{\mathcal{O}(k^3)}$.*

507 *Proof.* Without loss of generality, let the open children of $v \in V(G)$ be v_1, \dots, v_t and let the
 508 remaining (i.e., closed) children of v be v_{t+1}, \dots, v_m ; recall that by Point 3. of Observation 8,
 509 $t \leq 2k$. For each closed child v_j , $j \in [m] \setminus [t]$, let s_j^\emptyset be the second component of the valid record for
 510 $\emptyset \in \mathcal{R}(v_j)$, and let s_j^\times be the second component of the valid record for the single non-empty relation
 511 in $\mathcal{R}(v_j)$. Consider the following procedure \mathbb{A} .

512 First, \mathbb{A} branches over all choices of $P \in \Gamma_f(v)$ and all choices of $(R_i, s_i) \in \mathcal{R}(v_i)$ for each
 513 individual open child v_i of v . Let $R_0 = \{pv \mid p \in P\}$ and let $R' = \bigcup_{j \in [t]_0} R_j$. If $\text{trcl}(R')$ is
 514 not irreflexive, we discard this branch; otherwise, we proceed as follows. Let R_{new} be the subset
 515 of R' containing all arcs uw such that $w \in V_v$. Moreover, let $s_{\text{new}} = f_v(P) + (\sum_{i \in [t]} s_i) +$
 516 $(\sum_{i \in [m] \setminus [t] \mid v_i \in P} s_i^\emptyset) + (\sum_{i \in [m] \setminus [t] \mid v_i \notin P} (\max(s_i^\emptyset, s_i^\times)))$.

517 The algorithm \mathbb{A} gradually constructs a set $\mathcal{R}^*(v)$ as follows. At the beginning, $\mathcal{R}^*(v) = \emptyset$. For
 518 each newly obtained tuple $(R_{\text{new}}, s_{\text{new}})$, \mathbb{A} checks whether $\mathcal{R}^*(v)$ already contains a tuple with R_{new}
 519 as its first element; if not, we add the new tuple to $\mathcal{R}^*(v)$. If there already exists such a tuple
 520 $(R_{\text{new}}, s_{\text{old}}) \in \mathcal{R}^*(v)$, we replace it with $(R_{\text{new}}, \max(s_{\text{old}}, s_{\text{new}}))$.

521 For the running time, recall that in order to construct $\mathcal{R}^*(v)$ the algorithm branched over $|\Gamma_f(v)|$ -
 522 many possible parent sets of v and over the choice of at most $2k$ -many binary relations R_i on the
 523 boundaries of open children. According to Observation 8.2, there are at most $3^{(2k+2)^2}$ options for
 524 every such relation, so we have at most $\mathcal{O}((3^{(2k+2)^2})^{2k} \cdot |\Gamma_f(v)|) \leq 2^{\mathcal{O}(k^3)} \cdot |\Gamma_f(v)|$ branches. In
 525 every branch we compute $\text{trcl}(R')$ in time $k^{\mathcal{O}(1)}$ and then compute the value of s_{new} using the
 526 equation provided above before updating $\mathcal{R}^*(v)$, which takes time at most $\mathcal{O}(m)$.

527 Finally, to establish correctness it suffices to prove following claim:

528 **Claim 1.** *$(R : s)$ is a record for v if and only if there exist $P \in \Gamma_f(v)$ and records $(R_i : s_i)$ for v_i ,*
 529 *$i \in [m]$, such that:*

- 530 • $\text{trcl}(\bigcup_{i=0}^t R_i)$ is irreflexive;
- 531 • $R_i = \emptyset$ for any closed child $v_i \in P$;
- 532 • $\sum_{i=1}^m s_i + f_v(P) = s$;
- 533 • $R = (\text{trcl}(\bigcup_{i=0}^t R_i))|_{\delta(v) \times \delta(v)}$.

534 *Moreover, if $(R : s) \in \mathcal{R}(v)$ then in addition:*

- 535 • $(R_i : s_i) \in \mathcal{R}(v_i)$, $i \in [t]$;
- 536 • for every closed child $v_i \notin P$, $s_i = \max(s_i^\emptyset, s_i^\times)$.

Proof of the Claim. (a) (\Leftarrow) Denote $V_i = V_{v_i}$ and $\bar{V}_i = \bar{V}_{v_i}$, $i \in [m]$. For every $i \in [m]$ there exists
 DAG D_i on \bar{V}_i such that all its arcs finish in V_i , $R_i = \text{Con}(\delta(v_i), D_i)$ and $\sum_{u \in V_i} f_u(P_{D_i}(u)) = s_i$.
 Denote by D_0 DAG on $V_0 = v \cup N_G(v)$ with arc set R_0 . We will construct the witness D of (R, s)
 by gluing together all D_i , $i \in [m]_0$.

We start from D_0 and DAGs of open children. Note that $\text{Con}(V_0, D_0) = R_0$ and
 $\text{Con}(\delta(v_i), D_i) = R_i$ for $i \in [t]$. Inductive application of Lemma 7 to DAGs D_i , $i \in [t]$,
 yields $\text{Con}(\bigcup_{i=1}^t \delta(v_i) \cup V_0, D^*) = \text{trcl}(\bigcup_{i=0}^t R_i)$. In particular, as $\delta(v) \subseteq \bigcup_{i=1}^t \delta(v_i) \cup V_0$ by
 Observation 8.3, we have that $\text{Con}(\delta(v), D^*) = (\text{trcl}(\bigcup_{i=0}^t R_i))|_{\delta(v) \times \delta(v)} = R$. As $\text{trcl}(\bigcup_{i=0}^t R_i)$
 is irreflexive, $D^* = \bigcup_{i=0}^t D_i$ is DAG by Lemma 7.

Now we add to D^* DAGs for closed children and finally obtain $D = \bigcup_{i=t+1}^m D_i \cup D^*$. For
 every closed child v_i , D_i is by Observation 8.1 the union of v and $D_i \setminus v$, plus at most one of arcs
 $vv_i, v_i v$ between them (recall $R_i = \emptyset$ for any closed child $v_i \in P$). Note that $D_i \setminus v$ can share only

v_i with D_0 and doesn't have common vertices with any other D_j . Therefore any directed path in D starting and finishing outside of V_i , $i > t$, doesn't intersect V_i . In particular, acyclicity of D^* and D_i , $i \in [m] \setminus [t]$, implies acyclicity of D ; $\text{Con}(\delta(v), D) = \text{Con}(\delta(v), D^*) = R$.

All the arcs in D_i finish in V_i , so parent set for every $x_i \in D_i$ in D is the same as in D_i , $i \in [m]$. Also parent set of v in D is the same as in D_0 . So

$$\sum_{u \in V_v} f_u(P_D(u)) = \sum_{i=1}^m \sum_{u \in V_i} f_u(P_{D_i}(u)) + f_v(P_{D_0}(v)) = \sum_{i=1}^m s_i + f_v(P) = s$$

(\Rightarrow) Let D be a witness for $(R : s)$, i.e. D is DAG on \bar{V}_v with all arcs finishing in V_v such that $\sum_{u \in V_v} f_u(P_D(u)) = s$ and $\text{Con}(\delta(v), D) = R$. For $i = 1 \in [m]$ define $D'_i = D[V_i]$ and let D_i be obtained from D'_i by deleting arcs that finish outside V_i . Note that $\cup_{i=1}^m D_i = D$. Let $R_i = \text{Con}(\delta(v_i), D_i)$, as in (\Leftarrow) we have that $R = \text{Con}(\delta(v), D) = \text{trcl}(\cup_{i=0}^t R_i)|_{\delta(v) \times \delta(v)}$. As D is DAG, $\text{trcl}(\cup_{i=0}^t R_i)$ is irreflexive and $R_i = \emptyset$ for any closed child $v_i \in P$. Local score for D_i is

$$s_i = \sum_{u \in V_i} f_u(P_{D_i}(u)) = \sum_{u \in V_i} f_u(P_{D'_i}(u)) = \sum_{u \in V_i} f_u(P_D(u))$$

So v_i has record $(R_i : s_i)$. Denote $P = P_D(v)$. Then:

$$s = \sum_{u \in V_v} f_u(P_D(u)) = \sum_{i=1}^m \sum_{u \in V_i} f_u(P_D(u)) + f_v(P_D(v)) = \sum_{i=1}^m s_i + f_v(P)$$

537 (b) Let $(R : s) \in \mathcal{R}(v)$ and all D, P, D_i, R_i, s_i are as in (a)(\Rightarrow). Assume that for some i (R_i, s_i) is
 538 not valid record of v_i . In this case v_i must have a record $(R_i : s_i + \Delta)$ with $\Delta > 0$. But then (a)(\Leftarrow)
 539 implies that v has record $(R : s + \Delta)$, which contradicts to validity of $(R : s)$

540

541 Assume that some closed $v_i \notin P$ has valid record $(R'_i, s_i + \Delta)$ with $\Delta > 0$. R' and R
 542 differ only by arc vv_i , so addition or deletion of the arc to D would increase the total score by $\Delta > 0$
 543 without creating cycles. This would result in record $(R : s + \Delta)$ and yield a contradiction with
 544 validity of $(R : s)$. \blacksquare \square

545 We are now ready to prove the main result of this subsection.

546 *Proof of Theorem 6.* We provide an algorithm that solves $\text{BNSL}^{\neq 0}$ in time $2^{\mathcal{O}(k^3)} \cdot n^3$, where $n = |\mathcal{I}|$,
 547 assuming that a spanning tree T of G such that $\text{lfn}(G, T) = k$ is provided as part of the input. Once
 548 that is done, the theorem will follow from Theorem 2.

549 The algorithm computes $\mathcal{R}(v)$ for every node v in T , moving from leaves to the root:

550 • For a leaf v , compute $\mathcal{R}^*(v) := \{(R_P : f_v(P)) \mid P \in \Gamma_f(v), R_P = \{uv \mid u \in P\}\}$. This can
 551 be done by simply looping over $\Gamma_f(v)$ in time $\mathcal{O}(n)$. Note that $\mathcal{R}^*(v)$ is the set of all records
 552 of v , so we can correctly set $\mathcal{R}(v) := \{(R : s) \in \mathcal{R}^*(v) \mid \text{there is no } (R : s') \in \mathcal{R}^*(v) \text{ with}$
 553 $s' > s\}$.

554 • Let $v \in V(G)$ have at least one child in T , and assume we have computed $\mathcal{R}(v_i)$ for each
 555 child v_i of v . Then we invoke Lemma 9 to compute $\mathcal{R}(v)$ in time at most $m \cdot |\Gamma_f(v)| \cdot$
 556 $2^{\mathcal{O}(k^2)} \leq 2^{\mathcal{O}(k^2)} \cdot n^2$. \square

557 3.3 Lower Bounds for $\text{BNSL}^{\neq 0}$

558 Since lfn lies between fen and treecut width in the parameter hierarchy (see Proposition 1) and
 559 $\text{BNSL}^{\neq 0}$ is FPT when parameterized by lfn , the next step would be to ask whether this tractability
 560 result can be lifted to treecut width. Below, we answer this question negatively.

561 **Theorem 10.** $\text{BNSL}^{\neq 0}$ is $\text{W}[1]$ -hard when parameterized by the treecut width of the superstructure
 562 graph.

563 In fact, we show an even stronger result: $\text{BNSL}^{\neq 0}$ is $\text{W}[1]$ -hard when parameterized by the vertex
564 cover number of the superstructure even when all vertices outside of the vertex cover are required to
565 have degree at most 2. We remark that while $\text{BNSL}^{\neq 0}$ was already shown to be $\text{W}[1]$ -hard when
566 parameterized by the vertex cover number [36], in that reduction the degree of the vertices outside of
567 the vertex cover is not bounded by a constant and, in particular, the graphs obtained in that reduction
568 have unbounded treecut width.

569 *Proof of Theorem 10.* We reduce from the following well-known $\text{W}[1]$ -hard problem [12, 8]:

REGULAR MULTICOLORED CLIQUE (RMC)

570 Input: A k -partite graph $G = (V_1 \cup \dots \cup V_k, E)$ such that $|N_G(v)| = m$ for every $v \in V$
571 Parameter: The integer k
572 Question: Are there nodes v^i that form a k -colored clique in G , i.e. $v^i \in V_i$ and $v^i v^j \in E$
573 for all $i, j \in [k], i \neq j$?

571 We say that vertices in V_i have color i . Let $G = (V_1 \cup \dots \cup V_k, E)$ be an instance of RMC. We
572 will construct an instance (V, \mathcal{F}, ℓ) of $\text{BNSL}^{\neq 0}$ such that \mathcal{I} is a **Yes**-instance if and only if G is a
573 **Yes**-instance of RMC. V consists of one vertex v_i for each color $i \in [k]$ and one vertex v_e for every
574 edge $e \in E$. For each edge $e \in E$ that connects a vertex of color i with a vertex of color j , the
575 constructed vertex v_e will have precisely one element in its score function that achieves a non-zero
576 score, in particular: $f_{v_e}(\{v_i, v_j\}) = 1$.

577 Next, for each $i \in [k]$, we define the scores for v_i as follows. For every $v \in V_i$, let E_v be the set of all
578 edges incident to v in G , and let $P_i^v = \{v_e : e \in E_v\}$. We now set $f_{v_i}(P_i^v) = m + 1$ for each such
579 v ; all other parent sets will receive a score of 0. Note that $\{v_i \mid i \in [k]\}$ forms a vertex cover of the
580 superstructure graph and that all vertices outside of this vertex cover have degree at most 2, as desired.
581 We will show that G has a k -colored clique if and only if there is a Bayesian network D with score at
582 least $\ell = |E| + k + \binom{k}{2}$. (In fact, it will later become apparent that the score can never exceed ℓ .)

583 Assume first that G has a k -colored clique on $v^i, i \in [k]$, consisting of a set E_X of $\binom{k}{2}$ edges.
584 Consider the digraph D on V obtained as follows. For each vertex $v_i, i \in [k]$, and each vertex
585 v_e where $e \in E$, D contains the arc $v_e v_i$ if v_e is incident to v^i and otherwise D contains the arc
586 $v_i v_e$. This completes the construction of D . Now notice that the construction guarantees that each
587 v_i receives the parent set $P_i^{v^i}$ and hence contributes a score of $m + 1$. Moreover, for every edge e
588 not incident to a vertex in the clique, the vertex v_e contributes a score of 1; note that the number
589 of such edges is $|E| - km + \binom{k}{2}$; indeed, every v_i is incident to m edges but since $v^i, i \in [k]$,
590 was a clique we are guaranteed to double-count precisely $\binom{k}{2}$ many edges. Hence the total score is
591 $k(m + 1) + |E| - km + \binom{k}{2} = |E| + k + \binom{k}{2}$, as desired.

592 Assume that $\mathcal{I} = (V, \mathcal{F}, \ell)$ is a **Yes**-instance and let $s_{\text{opt}} \geq \ell = |E| + k + \binom{k}{2}$ be the maximum score
593 that can be achieved by a solution to \mathcal{I} ; let D be a dag witnessing such a score. Then all $v_i, i \in [k]$,
594 must receive a score of $m + 1$ in D . Indeed, assume that some v_i receives a score of 0 and let P_v be
595 any parent set of v_i with a score $m + 1$. Modify D by orienting edges $v_i v_e$ for every $v_e \in P_v$ inside
596 v_i . Now local score of v_i is $m + 1$, total score of the rest of vertices decreased by at most m (maximal
597 number of v_e that had local score 1 in D and lost it after the modification). So the modified DAG has
598 a score of at least $s_{\text{opt}} + 1$, which contradicts the optimality of s_{opt} . Therefore all $v_i, i \in [k]$, get
599 score $m + 1$ in D .

600 Let P_i be parent set of v_i in D , then $|P_i| = m$, $P_i = P_i^{v^i}$ for some $v^i \in V_i$. For every $v_e \in P_i$,
601 the local score of v_e in D is 0. Denote by E_{unsat} the set of all v_e that have a score of 0 in D . Every
602 v_e belongs to at most 2 different P_i and $P_i \cap P_j \leq 1$ for every $i \neq j$, so $|E_{\text{unsat}}| \geq km - \binom{k}{2}$. If
603 $|E_{\text{unsat}}| > km - \binom{k}{2}$, sum of local scores of v_e in D would be smaller than $|E| - km + \binom{k}{2}$, which
604 results in $s_{\text{opt}} < |E| + k + \binom{k}{2}$. Therefore $|E_{\text{unsat}}| = km - \binom{k}{2}$. But this means that $P_i \cap P_j \neq \emptyset$ for
605 any $i \neq j$, i.e. $v^i, i \in [k]$ form a k -colored clique in G . In particular $s_{\text{opt}} = \ell$. \square

606 For our second result, we note that the construction in the proof of Theorem 10 immediately implies
607 that $\text{BNSL}^{\neq 0}$ is NP-hard even under the following two conditions: (1) $\ell + \sum_{v \in V} |\Gamma_f(v)| \in \mathcal{O}(|V|^2)$

608 (i.e., the size of the parent set encoding is quadratic in the number of vertices), and (2) the instances
609 are constructed in a way which makes it impossible to achieve a score higher than ℓ . Using this,
610 as a fairly standard application of *AND-cross-compositions* [8] we can exclude the existence of an
611 efficient data reduction algorithm for $\text{BNSL}^{\neq 0}$ parameterized by lfe_n :

612 **Theorem 11.** *Unless $\text{NP} \subseteq \text{co-NP}/\text{poly}$, there is no polynomial-time algorithm which takes as*
613 *input an instance \mathcal{I} of $\text{BNSL}^{\neq 0}$ whose superstructure has lfe_n k and outputs an equivalent instance*
614 *$\mathcal{I}' = (V', \mathcal{F}', \ell')$ of $\text{BNSL}^{\neq 0}$ such that $|V'| \in k^{\mathcal{O}(1)}$. In particular, $\text{BNSL}^{\neq 0}$ does not admit a*
615 *polynomial kernel when parameterized by lfe_n .*

616 *Proof Sketch.* We describe an AND-cross-composition for the problem while closely following the
617 terminology and intuition introduced in Section 15 in the book [8]. Let the input consist of instances
618 $\mathcal{I}_1, \dots, \mathcal{I}_t$ of (unparameterized) instances of $\text{BNSL}^{\neq 0}$ which satisfy conditions (1) and (2) mentioned
619 above, and furthermore all have the same size and same target value of ℓ_1 (which is ensured through
620 the use of the polynomial equivalence relation \mathcal{R} [8, Definition 15.7]). The instance \mathcal{I} produced
621 on the output is merely the disjoint union of instances $\mathcal{I}_1, \dots, \mathcal{I}_t$ where we set $\ell := t \cdot \ell_1$, and we
622 parameterize \mathcal{I} by lfe_n .

623 Observe now that condition (a) in Definition 15.7 [8] is satisfied by the fact that the local feedback
624 edge number of \mathcal{I} is upper-bounded by the number of edges in a connected component of \mathcal{I} . Moreover,
625 the AND-variant of condition (b) in that same definition (see Subsection 15.1.3 [8]) is satisfied as
626 well: since none of the original instances can have a score greater than ℓ_1 , \mathcal{I} achieves a score of $\ell_1 \cdot t$
627 if and only if each of the original instances was a **Yes**-instance.

628 This completes the construction of an AND-cross-composition for $\text{BNSL}^{\neq 0}$ parameterized by lfe_n ,
629 and the claim follows by Theorem 15.12 [8]. \square

630 4 Additive Scores and Treewidth

631 While the previous section focused on the complexity of BNSL when the non-zero representation
632 was used (i.e., $\text{BNSL}^{\neq 0}$), here we turn our attention to the complexity of the problem with respect to
633 the additive representation. Recall from Subsection 2 that there are two variants of interest for this
634 representation: BNSL^+ and BNSL_{\leq}^+ . We begin by showing that, unsurprisingly, both of these are
635 NP-hard.

636 **Theorem 12.** *BNSL^+ is NP-hard. Moreover, BNSL_{\leq}^+ is NP-hard for every $q \geq 3$.*

637 *Proof.* We provide a direct reduction from the following NP-hard problem [23, 10]:

MINIMUM FEEDBACK ARC SET ON BOUNDED-DEGREE DIGRAPHS (MFAS)

638 **Input:** Digraph $D = (V, A)$ whose skeleton has degree at most 3, integer $m \leq |A|$.
Question: Is there a subset $A' \subseteq A$ where $|A'| \leq m$ such that $D - A'$ is a DAG?

639 Let (D, m) be an instance of MFAS. We construct an instance \mathcal{I} of BNSL_{\leq}^+ as follows:

- 640 • $V = V(D)$,
- 641 • $f_y(x) = 1$ for every $xy \in A(D)$,
- 642 • $f_y(x) = 0$ for every $xy \in A_V \setminus A(D)$,
- 643 • $\ell = |A| - m$, and
- 644 • $q = 3$.

645 Assume that (D, m) is a **Yes**-instance and A' is any feedback arc set of size m . Let D' be the DAG
646 obtained from D after deleting arcs in A' . Then $\text{score}(D')$ is equal to the number of arcs in D' ,
647 which is $|A| - m$, so \mathcal{I} is a **Yes**-instance. On the other hand, if \mathcal{I} is a **Yes**-instance of BNSL_{\leq}^+ ,
648 pick any DAG D' with $\text{score}(D') \geq \ell = |A| - m$. Without loss of generality we may assume that
649 $A(D') \subseteq A$, as the remaining arcs have a score of zero and may hence be removed. All the arcs in A
650 have a score 1 and hence the DAG D' contains at least $|A| - m$ arcs, i.e., it can be obtained from D

651 by deleting at most m arcs. Hence (D, m) is also a **Yes**-instance. To establish the NP-hardness of
 652 BNSL^+ , simply disregard the bound q on the input. \square

653 While the use of the additive representation did not affect the classical complexity of BNSL , it makes
 654 a significant difference in terms of parameterized complexity. Indeed, in contrast to $\text{BNSL}^{\neq 0}$:

655 **Theorem 13.** BNSL^+ is FPT when parameterized by the treewidth of the superstructure. Moreover,
 656 BNSL_{\leq}^+ is FPT when parameterized by q plus the treewidth of the superstructure.

657 *Proof.* We begin by proving the latter statement, and will then explain how that result can be
 658 straightforwardly adapted to obtain the former. As our initial step, we apply Bodlaender's algorithm [4,
 659 27] to compute a nice tree-decomposition (\mathcal{T}, χ) of $G_{\mathcal{I}}$ of width $k = \text{tw}(G_{\mathcal{I}})$. In this proof we use
 660 T to denote the set of nodes of \mathcal{T} and $r \in T$ be the root of \mathcal{T} . Given a node $t \in T$, let χ_t^\downarrow be the set of
 661 all vertices occurring in bags of the rooted subtree T_t , i.e., $\chi_t^\downarrow = \{u \mid \exists t' \in T_t \text{ such that } u \in \chi(t')\}$.
 662 Let G_t^\downarrow be the subgraph of $G_{\mathcal{I}}$ induced on χ_t^\downarrow .

663 To prove the theorem, we will design a leaf-to-root dynamic programming algorithm which will
 664 compute and store a set of records at each node of T , whereas once we ascertain the records for r
 665 we will have the information required to output a correct answer. Intuitively, the records will store
 666 all information about each possible set of arcs between vertices in each bag, along with relevant
 667 connectivity information provided by arcs between vertices in χ_t^\downarrow and information about the partial
 668 score. They will also keep track of parent set sizes in each bag.

669 Formally, the records will have the following structure. For a node t , let $S(t) =$
 670 $\{(\text{loc}, \text{con}, \text{inn}) \mid \text{loc}, \text{con} \subseteq A_{\chi(t)}, \text{inn} : \chi(t) \rightarrow [q]_0\}$ be the set of *snapshots* of t . The record \mathcal{R}_t
 671 of t is then a mapping from $S(t)$ to $\mathbb{N}_0 \cup \{\perp\}$. Observe that $|S(t)| \leq 4^{k^2} (q+1)^k$. To introduce the
 672 semantics of our records, let Υ_t be the set of all directed acyclic graphs over the vertex set χ_t^\downarrow with
 673 maximal in-degree at most q , and let $D_t = (\chi_t^\downarrow, A)$ be a directed acyclic graph in Υ_t . We say that the
 674 *snapshot of D_t in t* is the tuple (α, β, p) where $\alpha = A \cap A_{\chi(t)}$, $\beta = \text{Con}(\chi(t), D_t)$ and p specifies
 675 numbers of parents of vertices from $\chi(t)$ in D , i.e. $p(v) = |\{w \in \chi_t^\downarrow \mid wv \in A\}|$, $v \in \chi(t)$. We are
 676 now ready to define the record \mathcal{R}_t . For each snapshot $(\text{loc}, \text{con}, \text{inn}) \in S(t)$:

- 677 • $\mathcal{R}_t(\text{loc}, \text{con}, \text{inn}) = \perp$ if and only if there exists no directed acyclic graph in Υ_t whose
 678 snapshot is $(\text{loc}, \text{con}, \text{inn})$, and
- 679 • $\mathcal{R}_t(\text{loc}, \text{con}, \text{inn}) = \tau$ if $\exists D_t \in \Upsilon_t$ such that
 - 680 – the snapshot of D_t is $(\text{loc}, \text{con}, \text{inn})$,
 - 681 – $\text{score}(D_t) = \tau$, and
 - 682 – $\forall D'_t \in \Upsilon_t$ such that the snapshot of D'_t is $(\text{loc}, \text{con}, \text{inn})$: $\text{score}(D_t) \geq \text{score}(D'_t)$.

683 Recall that for the root $r \in T$, we assume $\chi(r) = \emptyset$. Hence \mathcal{R}_r is a mapping from the one-element
 684 set $\{(\emptyset, \emptyset, \emptyset)\}$ to an integer τ such that τ is the maximum score that can be achieved by any DAG
 685 $D = (V, A)$ with all in-degrees of vertices upper bounded by q . In other words, \mathcal{I} is a YES-instance
 686 if and only if $\mathcal{R}_r(\emptyset, \emptyset, \emptyset) \geq \ell$. To prove the theorem, it now suffices to show that the records can be
 687 computed in a leaf-to-root fashion by proceeding along the nodes of T . We distinguish four cases:

688 **t is a leaf node.** Let $\chi(t) = \{v\}$. By definition, $S(t) = \{(\emptyset, \emptyset, \emptyset)\}$ and $\mathcal{R}_t(\emptyset, \emptyset, \emptyset) = f_v(\emptyset)$.

689 **t is a forget node.** Let t' be the child of t in \mathcal{T} and let $\chi(t) = \chi(t') \setminus \{v\}$. We initiate by setting
 690 $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) = \perp$ for each $(\text{loc}, \text{con}, \text{inn}) \in S(t)$.

691 For each $(\text{loc}', \text{con}', \text{inn}') \in S(t')$, let $\text{loc}_v, \text{con}_v$ be the restrictions of loc', con' to tu-
 692 ples containing v . We now define $\text{loc} = \text{loc}' \setminus \text{loc}_v$, $\text{con} = \text{con}' \setminus \text{con}_v$, $\text{inn} =$
 693 $\text{inn}' \upharpoonright_{\chi(t)}$ and say that $(\text{loc}, \text{con}, \text{inn})$ is *induced* by $(\text{loc}', \text{con}', \text{inn}')$. Set $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) :=$
 694 $\max(\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}), \mathcal{R}_{t'}(\text{loc}', \text{con}', \text{inn}'))$, where \perp is assumed to be a minimal element.

695 For correctness, it will be useful to observe that $\Upsilon_t = \Upsilon_{t'}$. Consider our final computed value of
 696 $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn})$ for some $(\text{loc}, \text{con}, \text{inn}) \in S(t)$.

697 If $\mathcal{R}_t(\text{loc}, \text{con}, \text{inn}) = \tau$ for some $\tau \neq \perp$, then there exists a DAG D which wit-
 698 nesses this. But then D also admits a snapshot $(\text{loc}', \text{con}', \text{inn}')$ at t' and witnesses

699 $\mathcal{R}_{t'}(\text{loc}', \text{con}', \text{inn}') \geq \tau$. Note that $(\text{loc}, \text{con}, \text{inn})$ is induced by $(\text{loc}', \text{con}', \text{inn}')$. So in
700 our algorithm $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) \geq \mathcal{R}_{t'}(\text{loc}', \text{con}', \text{inn}') \geq \tau$.

701

702 If on the other hand $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) = \tau$ for some $\tau \neq \perp$, then there exists a
703 snapshot $(\text{loc}', \text{con}', \text{inn}')$ such that $(\text{loc}, \text{con}, \text{inn})$ is induced by $(\text{loc}', \text{con}', \text{inn}')$ and
704 $\mathcal{R}_{t'}(\text{loc}', \text{con}', \text{inn}') = \tau$. $\mathcal{R}_t(\text{loc}, \text{con}, \text{inn}) \geq \tau$ now follows from the existence of a DAG
705 witnessing the value of $\mathcal{R}_{t'}(\text{loc}', \text{con}', \text{inn}')$.

706 Hence, we can correctly set $\mathcal{R}_t = \mathcal{R}_t^0$.

707 **t is an introduce node.** Let t' be the child of t in \mathcal{T} and let $\chi(t) = \chi(t') \cup \{v\}$. We initiate by
708 setting $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) = \perp$ for each $(\text{loc}, \text{con}, \text{inn}) \in S(t)$.

709 For each $(\text{loc}', \text{con}', \text{inn}') \in S(t')$ and each $Q \subseteq \{ab \in A_{\chi(t)} \mid \{a, b\} \cap \{v\} \neq \emptyset\}$, we define:

- 710 • $\text{loc} := \text{loc}' \cup Q$
- 711 • $\text{con} := \text{trcl}(\text{con}' \cup Q)$
- 712 • $\text{inn}(x) := \text{inn}'(x) + |\{y \in \chi(t) \mid yx \in Q\}|$ for every $x \in \chi(t) \setminus \{v\}$
- 713 • $\text{inn}(v) := |\{y \in \chi(t) \mid yv \in Q\}|$

714 If con is not irreflexive or $\text{inn}(x) > q$ for some $x \in \chi(t)$, discard this branch. Other-
715 wise, let $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) := \max(\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}), \text{new})$ where $\text{new} = \mathcal{R}_{t'}(\text{loc}', \text{con}', \text{inn}') +$
716 $\sum_{ab \in Q} f_b(a)$. As before, \perp is assumed to be a minimal element here.

717 Consider our final computed value of $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn})$ for some $(\text{loc}, \text{con}, \text{inn}) \in S(t)$.

718 For correctness, assume that $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) = \tau$ for some $\tau \neq \perp$ and is obtained from
719 $(\text{loc}', \text{con}', \text{inn}')$, Q defined as above. Then $\mathcal{R}_{t'}(\text{loc}', \text{con}', \text{inn}') = \tau - \sum_{ab \in Q} f_b(a)$. Construct
720 a directed graph D from the witness D' of $\mathcal{R}_{t'}(\text{loc}', \text{con}', \text{inn}')$ by adding the arcs specified in
721 Q . As $\text{con} = \text{trcl}(\text{con}' \cup Q)$ is irreflexive and D' is a DAG, D is a DAG as well by 7.
722 Moreover, $\text{inn}(x) \leq q$ for every $x \in \chi(t)$ and the rest of vertices have in D the same parents
723 as in D' , so $D \in \Upsilon_t$. In particular, $(\text{loc}, \text{con}, \text{inn})$ is a snapshot of D in t and D witnesses
724 $\mathcal{R}_t(\text{loc}, \text{con}, \text{inn}) \geq \mathcal{R}_{t'}(\text{loc}', \text{con}', \text{inn}') + \sum_{ab \in Q} f_b(a) = \tau$.

725 On the other hand, if $\mathcal{R}_t(\text{loc}, \text{con}, \text{inn}) = \tau$ for some $\tau \neq \perp$, then there must exist a directed
726 acyclic graph $D = (\chi_t^\downarrow, A)$ in Υ_t that achieves a score of τ . Let Q be the restriction of A to
727 arcs containing v , and let $D' = (\chi_{t'}^\downarrow \setminus v, A \setminus Q)$, clearly $D' \in \Upsilon_{t'}$. Let $(\text{loc}', \text{con}', \text{inn}')$ be the
728 snapshot of D' at t' . Observe that $\text{loc} = \text{loc}' \cup Q$, $\text{con} = \text{trcl}(\text{con}' \cup Q)$, inn differs from inn'
729 by the numbers of incoming arcs in Q and the score of D' is precisely equal to the score τ of D
730 minus $\sum_{(a,b) \in Q} f_b(a)$. Therefore $\mathcal{R}_{t'}(\text{loc}', \text{con}', \text{inn}') \geq \tau - \sum_{(a,b) \in Q} f_b(a)$ and in the algorithm
731 $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) \geq \mathcal{R}_{t'}(\text{loc}', \text{con}', \text{inn}') + \sum_{(a,b) \in Q} f_b(a) \geq \tau$. Equality then follows from the
732 previous direction of the correctness argument.

733 Hence, at the end of our procedure we can correctly set $\mathcal{R}_t = \mathcal{R}_t^0$.

734 **t is a join node.** Let t_1, t_2 be the two children of t in \mathcal{T} , recall that $\chi(t_1) = \chi(t_2) = \chi(t)$. By the
735 well-known separation property of tree-decompositions, $\chi_{t_1}^\downarrow \cap \chi_{t_2}^\downarrow = \chi(t)$ [12, 8]. We initiate by
736 setting $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) := \perp$ for each $(\text{loc}, \text{con}, \text{inn}) \in S(t)$.

737 Let us branch over each $\text{loc}, \text{con}_1, \text{con}_2 \subseteq A_{\chi(t)}$ and $\text{inn}_1, \text{inn}_2 : \chi(t) \rightarrow [q]_0$. For every $b \in \chi(t)$
738 set $\text{inn}(b) = \text{inn}_1(b) + \text{inn}_2(b) - |\{a \mid ab \in \text{loc}\}|$. If:

- 739 • $\text{trcl}(\text{con}_1 \cup \text{con}_2)$ is not irreflexive and/or
- 740 • $\mathcal{R}_{t_1}(\text{loc}, \text{con}_1, \text{inn}_1) = \perp$, and/or
- 741 • $\mathcal{R}_{t_2}(\text{loc}, \text{con}_2, \text{inn}_2) = \perp$, and/or
- 742 • $\text{inn}(b) > q$ for some $b \in \chi(t)$

743 then discard this branch. Otherwise, set $\text{con} = \text{trcl}(\text{con}_1 \cup \text{con}_2)$, $\text{doublecount} = \sum_{ab \in \text{loc}} f_b(a)$
 744 and $\text{new} = \mathcal{R}_{t_1}(\text{loc}, \text{con}_1) + \mathcal{R}_{t_2}(\text{loc}, \text{con}_2) - \text{doublecount}$. We then set $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) :=$
 745 $\max(\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}), \text{new})$ where \perp is once again assumed to be a minimal element.

746 At the end of this procedure, we set $\mathcal{R}_t = \mathcal{R}_t^0$.

747 For correctness, assume that $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) = \tau \neq \perp$ is obtained from $\text{loc}, \text{con}_1, \text{con}_2, \text{inn}_1, \text{inn}_2$
 748 as above. Let $D_1 = (\chi_{t_1}^\downarrow, A_1)$ and $D_2 = (\chi_{t_2}^\downarrow, A_2)$ be DAGs witnessing $\mathcal{R}_{t_1}(\text{loc}, \text{con}_1, \text{inn}_1)$
 749 and $\mathcal{R}_{t_2}(\text{loc}, \text{con}_2, \text{inn}_2)$ correspondingly. Note that common vertices of D_1 and D_2 are precisely
 750 $\chi(t)$. In particular, if D_1 and D_2 share an arc ab , then $a, b \in \chi(t)$ and therefore $ab \in \text{loc}$. On
 751 the other hand, $\text{loc} \subseteq A_1, \text{loc} \subseteq A_2$, so $\text{loc} = A_1 \cap A_2$. Hence inn specifies the number of
 752 parents of every $b \in \chi(T)$ in $D = D_1 \cup D_2$. Rest of vertices $v \in V(D) \setminus \chi(t)$ belong to
 753 precisely one of D_i and their parents in D are the same as in this D_i . As $\text{trcl}(\text{con}_1 \cup \text{con}_2)$ is
 754 irreflexive, D is a DAG by Lemma 7, so $D \in \Upsilon_t$. The snapshot of D in t is $(\text{loc}, \text{con}, \text{inn})$ and
 755 $\text{score}(D) = \sum_{ab \in A(D)} f_b(a) = \sum_{ab \in A_1} f_b(a) + \sum_{ab \in A_2} f_b(a) - \sum_{ab \in \text{loc}} f_b(a) = \text{score}(D_1) +$
 756 $\text{score}(D_2) - \text{doublecount} = \mathcal{R}_{t_1}(\text{loc}, \text{con}_1, \text{inn}_1) + \mathcal{R}_{t_2}(\text{loc}, \text{con}_2, \text{inn}_2) - \text{doublecount} = \tau$.
 757 So D witnesses that $\mathcal{R}_t(\text{loc}, \text{con}, \text{inn}) \geq \tau$.

758 For the converse, assume that $\mathcal{R}_t(\text{loc}, \text{con}, \text{inn}) = \tau \neq \perp$ and D is a DAG witnessing this. Let
 759 D_1 and D_2 be restrictions of D to $\chi_{t_1}^\downarrow$ and $\chi_{t_2}^\downarrow$ correspondingly, then by the same arguments as
 760 above $A(D_1) \cap A(D_2) = \text{loc}$, in particular $D = D_1 \cup D_2$. Let $(\text{loc}, \text{con}_i, \text{inn}_i)$ be the snapshot
 761 of D_i in $t_i, i = 1, 2$, then $\mathcal{R}_{t_i}(\text{loc}, \text{con}_i, \text{inn}_i) \geq \text{score}(D_i)$. By the procedure of our algorithm,
 762 $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) \geq \mathcal{R}_{t_1}(\text{loc}, \text{con}_1, \text{inn}_1) + \mathcal{R}_{t_2}(\text{loc}, \text{con}_2, \text{inn}_2) - \text{doublecount} \geq \text{score}(D_1) +$
 763 $\text{score}(D_2) - \sum_{ab \in \text{loc}} f_b(a) = \text{score}(D) = \tau$.

764 Hence the resulting record \mathcal{R}_t is correct, which concludes the correctness proof of the algorithm.

765 Since the nice tree-decomposition \mathcal{T} has $\mathcal{O}(n)$ nodes, the runtime of the algorithm is upper-bounded
 766 by $\mathcal{O}(n)$ times the maximum time required to process each node. This is dominated by the time
 767 required to process join nodes, for which there are at most $(2^{k^2})^3((q+1)^k)^2 = 8^{k^2} \cdot (q+1)^{2k}$ branches
 768 corresponding to different choices of $\text{loc}, \text{con}_1, \text{con}_2, \text{inn}_1, \text{inn}_2$. Constructing $\text{trcl}(\text{con}_1 \cup \text{con}_2)$
 769 and verifying that it is irreflexive can be done in time $\mathcal{O}(k^3)$. Computing doublecount and inn
 770 takes time at most $\mathcal{O}(k^2)$. So the record for a join node can be computed in time $2^{\mathcal{O}(k^2)} \cdot q^{\mathcal{O}(k)}$.
 771 Hence, after we have computed a width-optimal tree-decomposition for instance by Bodlaender's
 772 algorithm [4], the total runtime of the algorithm is upper-bounded by $2^{\mathcal{O}(k^2)} \cdot q^{\mathcal{O}(k)} \cdot n$.

773 Finally, to obtain the desired result for BNSL^+ , we can simply adapt the above algorithm by
 774 disregarding the entry inn and disregard all explicit bounds on the in-degrees (e.g., in the definition
 775 of Υ_t). The runtime for this dynamic programming procedure is then $2^{\mathcal{O}(k^2)} \cdot n$. \square

776 This completely resolves the parameterized complexity of BNSL^+ w.r.t. all parameters depicted
 777 on Figure 1. However, the same is not true for BNSL_{\leq}^+ : while a careful analysis of the algorithm
 778 provided in the proof of Theorem 13 reveals that BNSL_{\leq}^+ is XP -tractable when parameterized by the
 779 treewidth of the superstructure alone, it is not yet clear whether it is FPT —in other words, do we
 780 need to parameterize by both q and treewidth to achieve fixed-parameter tractability?

781 We conclude this section by answering this question affirmatively. To do so, we will aim to reduce
 782 from the following problem, which can be seen as a dual to the $\text{W}[1]$ -hard MULTIDIMENSIONAL
 783 SUBSET SUM problem considered in recent works [21, 18].

UNIFORM DUAL MULTIDIMENSIONAL SUBSET SUM (UDMSS)

784 **Input:** An integer k , a set $S = \{s_1, \dots, s_n\}$ of item-vectors with $s_i \in \mathbb{N}^k$ for every i
 with $1 \leq i \leq n$, a uniform target vector $t = (r, \dots, r) \in \mathbb{N}^k$, and an integer d .
Parameter: k .
Question: Is there a subset $S' \subseteq S$ with $|S'| \geq d$ such that $\sum_{s \in S'} s \leq t$?

785 We first begin by showing that this variant of the problem is $\text{W}[1]$ -hard by giving a fairly direct
 786 reduction from the originally considered problem, and then show how it can be used to obtain the
 787 desired lower-bound result.

788 **Lemma 14.** DMSS is $W[1]$ -hard.

789 *Proof.* The $W[1]$ -hard MULTIDIMENSIONAL SUBSET SUM problem is stated as follows:

MULTIDIMENSIONAL SUBSET SUM (MSS)	
Input:	An integer k , a set $S = \{s_1, \dots, s_n\}$ of item-vectors with $s_i \in \mathbb{N}^k$ for every i with $1 \leq i \leq n$, a target vector $t = (t^1, \dots, t^k) \in \mathbb{N}^k$, and an integer d .
Parameter:	k .
Question:	Is there a subset $S' \subseteq S$ with $ S' \leq d$ such that $\sum_{s \in S'} s \geq t$?

791 Consider its dual version, obtained by reversing both inequalities:

DUAL MULTIDIMENSIONAL SUBSET SUM (DMSS)	
Input:	An integer k , a set $S = \{s_1, \dots, s_n\}$ of item-vectors with $s_i \in \mathbb{N}^k$ for every i with $1 \leq i \leq n$, a target vector $t = (t^1, \dots, t^k) \in \mathbb{N}^k$, and an integer d .
Parameter:	k .
Question:	Is there a subset $S' \subseteq S$ with $ S' \geq d$ such that $\sum_{s \in S'} s \leq t$?

793 Given an instance $\mathcal{I} = (S, t, k, d)$ of MSS, we construct an instance $\mathcal{I}_d = (S, z - t, k, n - d)$ of
 794 DMSS, where $z = \sum_{s \in S} s$. Note that S' is a witness of \mathcal{I} if and only if $S \setminus S'$ is a witness of \mathcal{I}_d .
 795 The observation establishes $W[1]$ -hardness of DMSS.

796 Now it remains to show that DMSS is $W[1]$ -hard even if we require all the components of the target
 797 vector t to be equal. Let $\mathcal{I} = (S, t, k, d)$ be the instance of DMSS. We construct an equivalent
 798 instance $\mathcal{I}_{eq} = (S_{eq}, t_{eq}, k + 1, d + 1)$ of UDMSS with $t_{eq} = (d \cdot t_{max}, \dots, d \cdot t_{max})$, where $t_{max} =$
 799 $\max\{t^i : i \in [k]\}$. S_{eq} is obtained from S by setting the $(k + 1)$ -th entries equal to t_{max} , plus one
 800 auxiliary vector to make the target uniform: $S_{eq} = \{(a^1, \dots, a^k, t_{max}) \mid (a^1, \dots, a^k) \in S\} \cup \{b\}$,
 801 where $b = (dt_{max} - t^1, \dots, dt_{max} - t^k, 0)$.

802 For correctness, assume that \mathcal{I} is a **Yes**-instance, in particular, we can choose S' with $|S'| = d$
 803 and $\sum_{s \in S'} s \leq t$. Then $S'_{eq} = \{(a^1, \dots, a^k, t_{max}) \mid (a^1, \dots, a^k) \in S'\} \cup \{b\}$ witnesses that \mathcal{I}_{eq}
 804 is a **Yes**-instance. For the converse direction, let \mathcal{I}_{eq} be a **Yes**-instance, we choose S'_{eq} with $|S'_{eq}| =$
 805 $d + 1$ and $\sum_{s \in S'_{eq}} s \leq t_{eq}$. If $b \notin S'_{eq}$, sum of the $(k + 1)$ -th entries in S'_{eq} would be at least
 806 $(d + 1)t_{max}$, so b must belong to S'_{eq} . Then $S'_{eq} \setminus \{b\}$ consists of precisely d vectors with sum at most
 807 $t_{eq} - b = (t^1, \dots, t^k, dt_{max})$. Restrictions of these vectors to k first coordinates witness that \mathcal{I} is a
 808 **Yes**-instance. \square

809 **Theorem 15.** $BNSL_{\leq}^+$ is $W[1]$ -hard when parameterized by the treewidth of the superstructure.

810 *Proof.* Let $\mathcal{I} = (S, t, k, d)$ be an instance of UDMSS with $t = (r, \dots, r)$, and w.l.o.g. assume that
 811 r is greater than the parameter k . We construct an equivalent instance $(V, \mathcal{F}, \ell, r)$ of $BNSL_{\leq}^+$. Let
 812 us start from the vertex set V . For every $i \in [k]$, we add to V a vertex v^i corresponding to the i -th
 813 coordinate of the target vector t . Further, for every $s = (s^1, \dots, s^k) \in S$, we add vertices a_s, b_s and
 814 $s^1 + \dots + s^k$ many vertices $s_j^i, i \in [k], j \in [s^i]$. Intuitively, taking s into S' will correspond to adding
 815 arcs from s_j^i to v^i for every $i \in [k], j \in [s^i]$. The upper bound r for each coordinate of the sum in S'
 816 is captured by allowing v^i to have at most r many parents. Formally, for every $s \in S, i \in [k], j \in [s^i]$
 817 the scores are defined as follows (for convenience we list them as scores per arc): $f(s_j^i v^i) = 2$,
 818 $f(b_s a_s) = M_s = 2 \cdot \sum_{i \in [k]} s^i - 1$. We call the arcs mentioned so far *light*. Note that for every
 819 fixed $s \in S, \sum_{i \in [k]} \sum_{j \in [s^i]} f(s_j^i v^i) = 2 \cdot \sum_{i \in [k]} s^i = M_s + 1$ so the sum of scores of light arcs is
 820 $L = \sum_{s \in S} (2M_s + 1)$. We finally set $f(a_s s_j^i) = f(v^i b_s) = L$ for every $s \in S, i \in [k]$ and $j \in [s^i]$.
 821 Now the number of arcs yielding the score of L is $m = k|S| + \sum_{s \in S} \sum_{i \in [k]} s^i$; we call these arcs
 822 *heavy*. We set the scores of all arcs not mentioned above to zero and we set $\ell = mL + \sum_{s \in S} M_s + d$.
 823 This finishes our construction; see Figure 3 for an illustration. Note that the superstructure graph has
 824 treewidth of at most $k + 2$: the deletion of vertices $v^i, i \in [k]$, makes it acyclic.




Figure 3: An example of our main gadget encoding the vector $s = (1, 3, 2)$ with $k = 3$. Heavy arcs are marked in green, while purple and blue arcs are light.

Intuitively, the reduction forces a choice between using the blue edge or all the purple edges; the latter case provides a total score that is 1 greater than the former, but is constrained by the upper bound r on the in-degrees of v^1, v^2, v^3 .

825 For correctness, assume that $\mathcal{I} = (S, t, k, d)$ is a **Yes**-instance of UDMSS, let S' be a subset of
 826 S of size d witnessing it. We add all the heavy arcs, resulting in a total score of mL . Further, for
 827 every $s = (s^1, \dots, s^k) \in S'$, we add all the arcs $s_j^i v^i, i \in [k], j \in [s^i]$, which increases the total
 828 score by $M_s + 1$. For every $s \in S \setminus S'$, we add an arc $b_s a_s$, augmenting the total score by M_s .
 829 Denote the resulting digraph by D , then $\text{score}(D) = mL + \sum_{s \in S'} (M_s + 1) + \sum_{s \in S \setminus S'} M_s =$
 830 $mL + \sum_{s \in S} M_s + d = \ell$. We proceed by checking parent set sizes. Note that every s_j^i has precisely
 831 one incoming arc $a_s s_j^i$ in D , every a_s has at most one in-neighbour b_s and in-neighbours of every
 832 b_s are $v^i, i \in [k]$. Finally, for every $i \in [k], P_D(v^i) = \{s_j^i | s \in S', j \in [s^i]\}$ by construction, so
 833 $|P_D(v^i)| = \sum_{s \in S'} s^i \leq r$ as S' is a solution to UDMSS. Therefore all the vertices in D have at
 834 most r in-neighbours. It remains to show acyclicity of D . As any cycle in the superstructure contains
 835 v^i for some $i \in [k]$, the same holds for any potential directed cycle C in D . Two next vertices of
 836 C after v^i can be only b_s and a_s for some $s \in S$. In particular, by our construction, $s \in S \setminus S'$.
 837 Then, again by construction, D doesn't contain an arc $s_j^i v^i$ for any $i \in [k], j \in [s^i]$, so v^i is not
 838 reachable from a_s , which contradicts to C being a cycle. Therefore D witnesses that $(V, \mathcal{F}, \ell, r)$ is a
 839 **Yes**-instance.

840 For the opposite direction, let $(V, \mathcal{F}, \ell, r)$ be a **Yes**-instance of BNSL_{\leq}^+ and let D be a DAG
 841 witnessing this. Then D contains all the heavy arcs. Indeed, sum of scores of all light arcs in \mathcal{F} is L ,
 842 so if at least one heavy arc is not in $A(D)$, then $\text{score}(D) \leq (m-1)L + L = mL < \ell$. For every
 843 $s \in S$, let $A^s = \{s_j^i v^i | i \in [k], j \in [s^i]\}$. If D doesn't contain an arc $b_s a_s$ and some of arcs from A^s ,
 844 the total score of $A(D) \cap A^s$ is at most $M_s - 1$. In this case we modify D by deletion of $A(D) \cap A^s$
 845 and addition of arc $b_s a_s$, which increases $\text{score}(D)$ and may only decrease the parent set sizes of
 846 $v^i, i \in [k]$. After these modifications, let $S'' = \{s \in S | D \text{ contains an arc } b_s a_s\}$. Note that whenever
 847 $s \in S''$, D cannot contain any of the arcs $s_j^i v^i, i \in [k], j \in [s^i]$, as this would result in directed cycle
 848 $v^i \rightarrow b_s \rightarrow a_s \rightarrow s_j^i \rightarrow v^i$. Therefore for every $s \in S$, D contains either an arc $b_s a_s$ (yielding the
 849 score of M_s) or all of arcs $s_j^i v^i, i \in [k], j \in [s^i]$ (yielding the score of $M_s + 1$ in total), so the sum of
 850 scores of light arcs in D is $\sum_{s \in S \setminus S''} (M_s + 1) + \sum_{s \in S''} M_s = \sum_{s \in S} M_s + |S \setminus S''|$, which should
 851 be at least $\ell - mL = \sum_{s \in S} M_s + d$. So $|S \setminus S''| \geq d$, we claim that $S' = S \setminus S''$ is a solution to
 852 $\mathcal{I} = (S, t, k, d)$. Indeed, for every $i \in [k], r \geq |P_D(v^i)| = |\{s_j^i | s \in S', j \in [s^i]\}| = \sum_{s \in S'} s^i$. \square

853 5 Implications for Polytree Learning

854 Here, we discuss how the results of Sections 3 and 4 can be adapted to POLYTREE LEARNING (PL).

855 **Theorem 3: Data Reduction.** Recall that the proof of Theorem 3 used two data reduction rules.
 856 While Reduction Rule 1 carries over to $\text{PL}^{\neq 0}$, Reduction Rule 2 has to be completely redesigned to
 857 preserve the (non-)existence of undirected paths between a and c . By doing so, we obtain:

858 **Theorem 16.** *There is an algorithm which takes as input an instance \mathcal{I} of $\text{PL}^{\neq 0}$ whose superstructure*
859 *has feedback edge number k , runs in time $\mathcal{O}(|\mathcal{I}|^2)$, and outputs an equivalent instance $\mathcal{I}' =$*
860 *$(V', \mathcal{F}', \ell')$ of $\text{PL}^{\neq 0}$ such that $|V'| \leq 24k$.*

861 *Proof.* Note that Reduction Rule 1 acts on the superstructure graph by deleting leaves and therefore
862 preserves not only optimal scores but also (non-)existence of polytrees achieving the scores. Hence we
863 can safely apply the rule to reduce the instance of $\text{PL}^{\neq 0}$. After the exhaustive application, all the
864 leaves of the superstructure graph G are the endpoints of edges in feedback edge set, so there can be
865 at most $2k$ of them. To get rid of long induced paths in G , we introduce the following rule:

866 **Reduction Rule 3.** *Let a, b_1, \dots, b_m, c be a path in G such that for each $i \in [m]$, b_i has degree*
867 *precisely 2. For every $B \subseteq \{a, c\}$ and $p \in \{0, 1\}$, let $\ell_p(B)$ be the maximum sum of scores that can*
868 *be achieved by b_1, \dots, b_m under the conditions that (1) there exists an undirected path between b_1*
869 *and b_m if and only if $p = 1$; (2) b_1 (and analogously b_m) takes a (c) into its parent set if and only if*
870 *$a \in B$ ($c \in B$).*

871 *We construct a new instance $\mathcal{I}' = (V', \mathcal{F}', \ell)$ as follows:*

- 872 • $V' := V \cup \{b, b'_1, b''_1, b'_m, b''_m\} \setminus \{b_1 \dots b_m\}$;
- 873 • $\Gamma_{f'}(b'_1) = \Gamma_{f'}(b''_1) = \Gamma_{f'}(b'_m) = \Gamma_{f'}(b''_m) = \emptyset$;
- 874 • *The scores for a (analogously c) are obtained from \mathcal{F} by simply replacing every occurrence*
875 *of b_1 by b'_1 and b''_1 (b_m by b'_m and b''_m), formally:*
 - 876 – $\Gamma_{f'}(a)$ is a union of $\{P \in \Gamma_f(a) \mid b_1 \notin P\}$, where $f'_a(P) := f_a(P)$ and
 - 877 $\{P \setminus b_1 \cup \{b'_1, b''_1\} \mid b_1 \in P, P \in \Gamma_f(a)\}$, where $f'_a(P \setminus b_1 \cup \{b'_1, b''_1\}) := f_a(P)$;
 - 878 – $\Gamma_{f'}(c)$ is a union of $\{P \in \Gamma_f(c) \mid b_m \notin P\}$, where $f'_c(P) := f_c(P)$, and
 - 879 $\{P \setminus b_m \cup \{b'_m, b''_m\} \mid b_m \in P, P \in \Gamma_f(c)\}$, where $f'_c(P \setminus b_m \cup \{b'_m, b''_m\}) := f_c(P)$.
- 880 • $\Gamma_{f'}(b)$ consists of eight sets, yielding corresponding scores f'_b : $\{a, c, b'_1, b''_1, b'_m, b''_m\} \rightarrow$
881 $l_1(\{a, c\})$, $\{b'_1, b''_1, b'_m, b''_m\} \rightarrow l_0(\{a, c\})$, $\{b'_1, b'_m\} \rightarrow l_1(\emptyset)$, $\emptyset \rightarrow l_0(\emptyset)$, $\{a, b'_1, b''_1, b'_m\} \rightarrow$
882 $l_1(\{a\})$, $\{b'_1, b''_1\} \rightarrow l_0(\{a\})$, $\{b'_m, b''_m\} \rightarrow l_1(\{c\})$, $\{b'_1, b'_m, b''_m, c\} \rightarrow l_0(\{c\})$.

883 Parent sets of b are defined in a way to cover all the possible configurations on solutions to \mathcal{I} restricted
884 to a, b_1, \dots, b_m, c ; the corresponding scores of b are intuitively the sums of scores that $b_i, i \in [m]$,
885 receive in the solutions. The eight cases that may arise are illustrated in Figure 4.

886 **Claim 2.** *Reduction Rule 3 is safe.*

887 *Proof.* We will show that a score of at least ℓ can be achieved in the original instance \mathcal{I} if and only if
888 a score of at least ℓ can be achieved in the reduced instance \mathcal{I}' .

889 Assume that D is a polytree that achieves a score of ℓ in \mathcal{I} . We will construct a polytree D' , called the
890 *reduct* of D , with $f'(D') \geq \ell$. To this end, we first modify D by removing the vertices b_1, \dots, b_m and
891 adding $b, b'_1, b''_1, b'_m, b''_m$. We also add arcs $b'_1 a$ and $b''_1 a$ ($b'_m c$ and $b''_m c$ correspondingly) if and only if
892 $b_1 a \in A(D)$ ($b_m c \in A(D)$). Let us denote the DAG obtained at this point D^* . Note that scores of a
893 and c in D^* are the same as in D . Further modifications of D^* depend only on $D[a, b_1 \dots b_m, c]$ and
894 change only the parent set of b . We distinguish the 8 cases listed below (see also Figure 4):

- 895 • case 1.1 (1.2): $ab_1, cb_m \in A(D)$, b_1 and b_m are (not) connected by path in D . We add
896 incoming arcs to b from $a, c, b'_1, b''_1, b'_m, b''_m$ (b'_1, b''_1, b'_m, b''_m only) resulting in $f'_b(P_{D'}(b)) =$
897 $l_1(\{a, c\})$ ($f'_b(P_{D'}(b)) = l_0(\{a, c\})$).
- 898 • case 2.1 (2.2): $ab_1, cb_m \notin A(D)$, b_1 and b_m are (not) connected by path in D . We add
899 incoming arcs to b from b'_1 and b'_m (leave D^* unchanged) yielding $f'_b(P_{D'}(b)) = l_1(\emptyset)$
900 ($f'_b(P_{D'}(b)) = l_0(\emptyset)$).
- 901 • case 3.1 (3.2): $ab_1 \in A(D)$, $cb_m \notin A(D)$, b_1 and b_m are (not) connected by path in D .
902 We add incoming arcs to b from a, b'_1, b''_1, b'_m (b'_1 and b''_1 only), then $f'_b(P_{D'}(b)) = l_1(\{a\})$
903 ($f'_b(P_{D'}(b)) = l_0(\{a\})$).
- 904 • case 4.1 (4.2): $ab_1 \notin A(D)$, $cb_m \in A(D)$, b_1 and b_m are (not) connected by path in D . The
905 cases are symmetric to 3.1 (3.2)

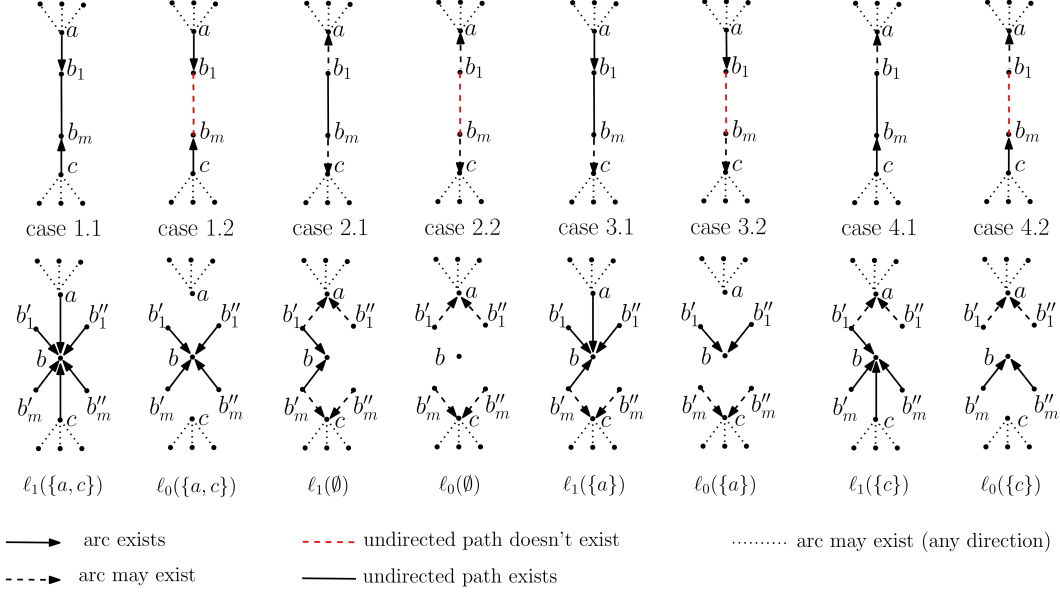


Figure 4: Top: The eight possible scenarios for solutions to \mathcal{I} . Bottom: The corresponding arcs in the gadget after the application of Reduction Rule 2' (the scores of b are specified below).

906 Note that D' contains a path between a and c if and only if D does. By definition of ℓ_0 and ℓ_1 , the
907 score of b in D' is at least as large as the sum of scores of b_i , $i \in [m]$, in D . Moreover, each vertex
908 in $V(D) \cap V(D')$ receives equal scores in D and D' . Hence D' is a polytree with $f'(D') \geq \ell$,
909 as desired. For the converse direction, note that the polytrees constructed in
910 cases 1.1-4.2 cover all optimal configurations which may arise in \mathcal{I}' : if there is a polytree D'' in
911 \mathcal{I}' with a score of ℓ' , we can always modify it to a polytree D' with a score of at least ℓ' such that
912 $D'[a, b'_1, b''_1, b, b'_m, b''_m, c]$ has one of the forms depicted at the bottom line of the figure. But every
913 such D' is a reduct of some polytree D of the original instance with the same score. \square

914 We apply Reduction Rule 3 exhaustively, until there is no more path to shorten. Bounds on the
915 running time of the procedure and size of the reduced instance can be obtained similarly to the case
916 of $\text{BNSL}^{\neq 0}$. In particular, every long path is replaced with a set of 5 vertices, resulting in at most
917 $4k + 4k \cdot 5 = 24k$ vertices. \square

918 **Theorem 6: Fixed-parameter tractability.** Analogously to $\text{BNSL}^{\neq 0}$ a data reduction procedure
919 as the one provided in Theorem 16 does not exist for $\text{PL}^{\neq 0}$ parametrized by l fen unless $\text{NP} \subseteq$
920 co-NP/poly , since the lower-bound result provided in Theorem 11 can be straightforwardly adapted
921 to $\text{PL}^{\neq 0}$. But similarly as for BNSL we can provide an FPT algorithm using the same ideas as in
922 the proof of Theorem 6. The algorithm proceeds by dynamic programming on the spanning tree T of
923 G with $\text{l fen}(G, T) = \text{l fen}(G) = k$. The records will, however, need to be modified: for each vertex
924 v , instead of the path-connectivity relation on $\delta(v)$, we store connected components of the *inner*
925 *boundary* $\delta(v) \cap V_v$ and incoming arcs to T_v . We provide a full description of the algorithm below.

926 **Theorem 17.** $\text{PL}^{\neq 0}$ is fixed-parameter tractable when parameterized by the local feedback edge
927 number of the superstructure.

928 *Proof.* As before, given an instance \mathcal{I} with a superstructure graph $G = G_{\mathcal{I}}$ such that $\text{l fen}(G) = k$,
929 we start from computing the spanning tree T of G with $\text{l fen}(G, T) = \text{l fen}(G) = k$; pick a root r in
930 T . We keep all the notations $T_v, V_v, \bar{V}_v, \delta(v)$ for $v \in V(T)$ from the subsection 3.2. In addition,
931 we define the *inner boundary* of $v \in V(T)$ to be $\delta_{in}(v) := \delta(v) \cap V_v$ i.e. part of boundary that
932 belongs to subtree of T rooted in v . The remaining part we call the *outer boundary* of v and denote by
933 $\delta_{out}(v) := \delta(v) \setminus \delta_{in}(v)$. For any set A of arcs, we define $\tilde{A} = \{uv | uv \in A \text{ or } vu \in A\}$. Obviously,

934 the claims of Observation 8 still hold. Moreover, for every closed v , $\delta_{in}(v)$ contains only v itself and
 935 $\delta_{out}(v)$ is either the parent of v in T or \emptyset (for $v = r$).

936 Let R_v be binary relation on $\delta_{in}(v)$, $A_v \subseteq \delta_{out}(v) \times \delta_{in}(v)$, s_v is integer. Then (R_v, A_v, s_v) is a
 937 *record* for v if and only if there exist a polytree D on \bar{V}_v with all arcs oriented inside V_v such that:

- 938 • $A_v = \{xy \in A(D) \mid x \in \delta_{out}(v), y \in \delta_{in}(v)\}$
- 939 • $R_v = \{xy \mid x, y \in \delta_{in}(v) \text{ are in the same connected component of } D[V_v]\}$
- 940 • $s_v = \sum_{u \in V_v} f_u(P_D(u))$

941 Note that R_v is an equivalence relation on $\delta_{in}(v)$, number of its equivalence classes is equal to
 942 number of connected components of $D[V_v]$ that intersect $\delta(v)$.

943 Record (R_v, A_v, s_v) is called *valid* if and only if s_v is maximal for fixed R_v, A_v among all the records
 944 for v . Denote by $\mathcal{R}(v)$ the set of all valid records for v , then $|\mathcal{R}(v)| \leq 2^{(2k+2)^2}$. Indeed, R_v and A_v
 945 can be uniquely determined by the choice of some relation on $\delta(v) \times \delta(v)$. As $|\delta(v)| \leq 2k + 2$, there
 946 are at most $2^{(2k+2)^2}$ possible relations.

947 The root r of T has a single valid record $(\emptyset, \emptyset, s_{\mathcal{T}})$, where $s_{\mathcal{T}}$ is the maximum score that can be
 948 achieved by a solution to \mathcal{L} . For any closed $v \neq r$, $\mathcal{R}(v)$ consists of precisely two valid records: one
 949 for $A_v = \emptyset, R_v = \{vv\}$ and another for $A_v = \{wv\}, R_v = \{vv\}$, where w is a parent of v in T .

950 We proceed by computing our records in a leaf-to-root fashion along T .

951 Let v be a leaf. Start by initiating $\mathcal{R}^*(v) := \emptyset$, then for each $P \in \Gamma_f(v)$ add to $\mathcal{R}^*(v)$ the triple
 952 $(\{vv\}, \{uv \mid u \in P\}, f_v(P))$. Note that $\mathcal{R}^*(v)$ is by definition precisely the set of all records for v , so
 953 we can correctly set $\mathcal{R}(v) = \{(R_v, A_v, s_v) \in \mathcal{R}^*(v) \mid s_v \text{ is maximal for fixed } R_v, A_v\}$.

954 Assume that v has m children $\{v_i : i \in [m]\}$ in T , where $v_i, i \in [t]$, are open and $v_i, i \in [m] \setminus [t]$,
 955 are closed. The following claim shows how (and under which conditions) the records of children of v
 956 can be composed into a record of v .

957 **Claim 3.** Let $P \in \Gamma_f(v)$, D_0 is a polytree on $V_0 = v \cup P$ with arc set $A_0 = \{uv \mid u \in P\}$, (R_i, A_i, s_i)
 958 are records for v_i witnessed by $D_i, i \in [m]$. Let A_{loc}^{in} be the set of arcs in $\bigcup_{i \in [t]_0} A_i$ which have both
 959 endpoints in V_v , $R = \text{trcl}(\tilde{A}_{loc}^{in} \cup \bigcup_{i \in [t]_0} R_i)$. Then $D = \bigcup_{i=0}^m D_i$ is a polytree if and only if the
 960 following two conditions hold:

- 961 1. $A_i = \emptyset$ for each closed child $v_i \in P$.
- 962 2. $\sum_{i=0}^t N_i - |A_{loc}^{in}| - \sum_{y \in Y} (n_y - 1) = N$, where
 - 963 • N is the number of equivalence classes in $\text{trcl}(\bigcup_{i \in [t]_0} (\tilde{A}_i \cup R_i))$
 - 964 • N_i is the number of equivalence classes in $R_i, i \in [t]$
 - 965 • Y is the set of endpoints of arcs in $\bigcup_{i \in [t]_0} A_i$ which don't belong to any $V_{v_i}, i \in [m]$.
 - 966 • For every $y \in Y$, n_y is the number of arcs in $A_0 \cup \dots \cup A_t$ having endpoint y .

967 In this case D witnesses the record (R_v, A_v, s_v) , where:

968 $R_v = R|_{\delta_{in}(v) \times \delta_{in}(v)}, A_v = (\bigcup_{i \in [t]_0} A_i)|_{\delta_{out}(v) \times \delta_{in}(v)}, s_v = \sum_{i=0}^m s_i + f_v(P)$.

969 If $(R_v, A_v, s_v) \in \mathcal{R}(v)$, then $(R_i, A_i, s_i) \in \mathcal{R}(v_i), i \in [m]$. Moreover, for any closed child $v_i \notin P$,
 970 there is no $(R'_i, A'_i, s'_i) \in \mathcal{R}(v_i)$ with $s'_i > s_i$.

971 We will prove the claim at the end, let us show how it can be exploited to compute valid records of
 972 v . We start from initial setting $\mathcal{R}^*(v) := \emptyset$, then branch over all parent sets $P \in \Gamma_f(v)$ and triples
 973 $(R_i, A_i, s_i) \in \mathcal{R}(v_i)$ for open children v_i . For each closed child $v_i \notin P$ take $(R_i, A_i, s_i) \in \mathcal{R}(v_i)$
 974 with maximal s_i , for each closed child $v_i \in P$ take $(R_i, A_i, s_i) \in \mathcal{R}(v_i)$ with $A_i = \emptyset$. Now the first
 975 condition of Claim3 holds, if the second one holds as well, we add to $\mathcal{R}^*(v)$ the triple (R_v, A_v, s_v) .

976 According to Claim 3, $\mathcal{R}^*(v)$ computed in such a way consists only of records for v and, in particular,
 977 contains all the valid records. Therefore we can correctly set $\mathcal{R}(v) = \{(R_v, A_v, s_v) \in \mathcal{R}^*(v) | s_v \text{ is}$
 978 maximal for fixed $R_v, A_v\}$.

979 To construct $\mathcal{R}^*(v)$ for node v with children $v_i, i \in [m]$, we branch over at most n possible parent
 980 sets of v and at most $2^{(2k+2)^2}$ valid records for every open child of v . Number of open children is
 981 bounded by $2k$, so we have at most $\mathcal{O}((2^{(2k+2)^2})^{2k} \cdot n) \leq 2^{\mathcal{O}(k^3)} \cdot n$ branches. In a fixed branch we
 982 compute scores for closed children in $\mathcal{O}(n)$, application of Claim 3 requires time polynomial in k .
 983 So $\mathcal{R}^*(v)$ is computed in time $2^{\mathcal{O}(k^3)} \cdot n^2$ that majorizes running time for leaves. As the number of
 984 vertices in T is at most n , total running time of the algorithm is $2^{\mathcal{O}(k^3)} \cdot n^3$ assuming that T is given
 985 as a part of the input.

986 *proof of Claim 3* (\Leftarrow). We start from checking whether $D = \cup_{i=0}^m D_i$ is a polytree. As the first
 987 condition implies that a polytree of every closed child v_i is connected to the rest of D by at most
 988 one arc $v_i v$ or $v v_i$, it is sufficient to check whether $D^t = \cup_{i=0}^t D_i$ is polytree. Number of connected
 989 components of D^t is $N' + N$, where N' is the total number of connected components of D_i that
 990 don't intersect $\delta(v_i), i \in [t]$. Note that D^t can be constructed as follows:

- 991 1. Take a disjoint union of polytrees $D'_i = D_i[V_i], i \in [t]_0$, then the resulting polytree has
 992 $N' + \sum_{i=0}^t N_i$ connected components.
- 993 2. Add arcs between D'_i and D'_j that occur in D for every $i, j \in [t]_0$, i.e. the arcs specified by
 994 A_{loc}^{in} . Resulting digraph is a polytree if and only if every added arc decreases the number
 995 of connected components by 1, i.e. the number of connected components after this step is
 996 $N' + \sum_{i=0}^t N_i - |A_{loc}^{in}|$.
- 997 3. Add all remaining vertices y of D together with their adjacent arcs in D . Note that such y
 998 precisely form the set Y , so D^t is a polytree if and only if we obtained a polytree after the
 999 previous step and every $y \in Y$ decreased it's number of connected components by $(n_y - 1)$,
 1000 i.e. the number $N' + N$ of connected components in D^t is equal to $N' + \sum_{i=0}^t N_i -$
 1001 $|A_{loc}^{in}| - \sum_{y \in Y} (n_y - 1)$. But this is precisely the condition 2 of the claim.

1002 Now, assuming that D is a polytree, we will show that it witnesses (R_v, A_v, s_v) . Parent sets of
 1003 vertices from each V_i in D are the same as in D_i , parent set of v in D is P . So $s_v = \sum_{i=0}^m s_i + f_v(P)$
 1004 is indeed the sum of scores over V_v in D .

1005 There are two kinds of arcs in D starting outside of V_v : incoming arcs to v and incoming arcs to the
 1006 subtrees of open children. Thus $A(D)|_{\delta_{out}(v) \times \delta_{in}(v)} = (\bigcup_{i \in [t]_0} A_i)|_{\delta_{out}(v) \times \delta_{in}(v)} = A_v$.

1007 Take any $u, w \in \delta_{in}(v), u \neq w$, note that u and w can not belong to subtrees of closed children.
 1008 So u and w are in the same connected component of $D[V_v]$ if and only if they are connected by
 1009 some undirected path π in the skeleton of D using only vertices from $D^t \cap V_v$. In this case R_i
 1010 captures the segments of π which are completely contained in $D_i[V_i], i \in [t]$. Rest of edges in π
 1011 either connect v to some $V_i, i \in [t]$, or have endpoints in different V_i and V_j for some $i, j \in [t]$. Edges
 1012 of this kind precisely form the set \tilde{A}_{loc}^{in} , so uw belongs to $R = \text{trcl}(\bigcup_{i \in [t]} R_i \cup \tilde{A}_{loc}^{in})$. Therefore
 1013 $R_v = R|_{\delta_{in}(v) \times \delta_{in}(v)}$ indeed represents connected components of $\delta_{in}(v)$ in $D[V_v]$.

1014 (\Rightarrow) Condition 1 obviously holds, otherwise D would contain a pair of arcs with the same endpoints
 1015 and different directions. In (\Leftarrow) we actually showed the necessity of condition 2 when 1 holds.

1016 For the last statement, assume that $(R_v, A_v, s_v) \in \mathcal{R}(v)$ but $(R_i, A_i, s_i) \notin \mathcal{R}(v_i)$ for some i . Then
 1017 there is $(R_i, A_i, s_i + \Delta) \in \mathcal{R}(v_i)$ for some $\Delta > 0$. Let D'_i be a witness of $(R_i, A_i, s_i + \Delta)$, then
 1018 $D' = \bigcup_{j \in [m] \setminus \{i\}} D_j \cup D'_i$ is a polytree witnessing $(R_v, A_v, s_v + \Delta)$. But this contradicts to validity
 1019 of (R_v, A_v, s_v) . By the same arguments records for closed children $v_i \notin P$ are the ones with maximal
 1020 s_i among two $(R_i, A_i, s_i) \in \mathcal{R}(v_i)$. \blacksquare \square

1021 As for treecut width, we remark that a recent reduction for $\text{PL}^{\neq 0}$ [24, Theorem 4.2] immediately
 1022 implies that the problem is $\text{W}[1]$ -hard when parameterized by the treecut width (the superstructure
 1023 graphs obtained in that reduction have a vertex cover of size bounded in the parameter, and the
 1024 vertices outside of the vertex cover have degree at most 2).

1025 **Theorem 13: Additive Representation.** We remark that, like BNSL^+ and BNSL_{\leq}^+ , a simple
 1026 reduction shows that PL_{\leq}^+ is NP-hard for a fixed value of q , in this case $q = 1$.

1027 **Theorem 18.** PL_{\leq}^+ is NP-hard when $q = 1$.

1028 *Proof.* We reduce from the classical HAMILTONIAN PATH problem. Given a graph G , we construct
 1029 an instance \mathcal{I} of PL_{\leq}^+ with $q = 1$ and the same vertex set. Whenever G contains an edge ab , we set
 1030 $f_a(b) = f_b(a) = 1$; all other cost functions are set to 0. ℓ is set to $|V| - 1$.

1031 Consider a solution D for \mathcal{I} . Since D is a DAG, it must contain a source; by construction, all other
 1032 vertices in D must have an in-degree of 1. This implies that the arcs of D form a Hamiltonian path
 1033 in G . Conversely, given a Hamiltonian path in G , one can construct a solution D by choosing one
 1034 endpoint of the path as the source and then adding all arcs along the path. \square

1035 Moreover, the dynamic programming algorithm for BNSL_{\leq}^+ parameterized by treewidth and q can
 1036 be adapted to also solve PL_{\leq}^+ . For completeness, we provide a full proof below; however one should
 1037 keep in mind that the ideas are very similar to the proof of Theorem 13.

1038 **Theorem 19.** PL^+ is FPT when parameterized by the treewidth of the superstructure. Moreover,
 1039 PL_{\leq}^+ is FPT when parameterized by q plus the treewidth of the superstructure.

1040 *Proof.* We begin by proving the latter statement, and will then explain how that result can be
 1041 straightforwardly adapted to obtain the former. As our initial step, we apply Bodlaender's algorithm [4,
 1042 27] to compute a nice tree-decomposition (\mathcal{T}, χ) of $G_{\mathcal{I}}$ of width $k = \text{tw}(G_{\mathcal{I}})$. We keep the notations
 1043 T, r , and $\chi_t^\downarrow G_t^\downarrow$ from the proof of Theorem 13. For any arc set A we denote $\tilde{A} = \{uw, wu \mid uw \in A\}$.

1044 We will design a leaf-to-root dynamic programming algorithm which will compute and store a
 1045 set of records at each node of T , whereas once we ascertain the records for r we will have the
 1046 information required to output a correct answer. The set of snapshots and structure of records will
 1047 be the same as in the proof of Theorem 13. However, semantics will slightly differ: in contrast to
 1048 information about directed paths via forgotten nodes, con will now specify whether vertices of the
 1049 bag belong to the same connected component of the partial polytree. Formally, let Ψ_t be the set
 1050 of all polytrees over the vertex set χ_t^\downarrow with maximal in-degree at most q , and let $D_t = (\chi_t^\downarrow, A)$ be
 1051 a polytree in Ψ_t . We say that the *snapshot* of D_t in t is the tuple (α, β, p) where $\alpha = A_{\chi(t)} \cap A$,
 1052 $\beta = A_{\chi(t)} \cap \{uw \mid u \text{ and } w \text{ belong to the same connected component of } D_t\}$ and p specifies numbers
 1053 of parents of vertices from $\chi(t)$ in D , i.e. $p(v) = |\{w \in \chi_t^\downarrow \mid wv \in A\}|$, $v \in \chi(t)$. We will call a
 1054 connected component of D_t *active* if it intersects $\chi(t)$. Note that the number of equivalence classes
 1055 of con is equal to the number of active connected components of D_t . We are now ready to define the
 1056 record \mathcal{R}_t . For each snapshot $(\text{loc}, \text{con}, \text{inn}) \in S(t)$:

- 1057 • $\mathcal{R}_t(\text{loc}, \text{con}, \text{inn}) = \perp$ if and only if there exists no polytree in Ψ_t whose snapshot is
 1058 $(\text{loc}, \text{con}, \text{inn})$, and
- 1059 • $\mathcal{R}_t(\text{loc}, \text{con}, \text{inn}) = \tau$ if $\exists D_t \in \Psi_t$ such that
 - 1060 – the snapshot of D_t is $(\text{loc}, \text{con}, \text{inn})$,
 - 1061 – $\text{score}(D_t) = \tau$, and
 - 1062 – $\forall D'_t \in \Psi_t$ such that the snapshot of D'_t is $(\text{loc}, \text{con}, \text{inn})$: $\text{score}(D_t) \geq \text{score}(D'_t)$.

1063 Recall that for the root $r \in T$, we assume $\chi(r) = \emptyset$. Hence \mathcal{R}_r is a mapping from the one-element
 1064 set $\{(\emptyset, \emptyset, \emptyset)\}$ to an integer τ such that τ is the maximum score that can be achieved by any polytree
 1065 $D = (V, A)$ with all in-degrees of vertices upper bounded by q . In other words, \mathcal{I} is a YES-instance
 1066 if and only if $\mathcal{R}_r(\emptyset, \emptyset, \emptyset) \geq \ell$. To prove the theorem, it now suffices to show that the records can be
 1067 computed in a leaf-to-root fashion by proceeding along the nodes of T . We distinguish four cases:

1068 **t is a leaf node.** Let $\chi(t) = \{v\}$. By definition, $S(t) = \{(\emptyset, \emptyset, \emptyset)\}$ and $\mathcal{R}_t(\emptyset, \emptyset, \emptyset) = f_v(\emptyset)$.

1069 **t is a forget node.** Let t' be the child of t in \mathcal{T} and let $\chi(t) = \chi(t') \setminus \{v\}$. We initiate by setting
 1070 $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) = \perp$ for each $(\text{loc}, \text{con}, \text{inn}) \in S(t)$.

1071 For each $(\text{loc}', \text{con}', \text{inn}') \in S(t')$, let $\text{loc}_v, \text{con}_v$ be the restrictions of loc', con' to tuples con-
 1072 taining v . We now define $\text{loc} = \text{loc}' \setminus \text{loc}_v, \text{con} = \text{con}' \setminus \text{con}_v, \text{inn} = \text{inn}'|_{\chi(t)}$ and set
 1073 $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) := \max(\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}), \mathcal{R}_{t'}(\text{loc}', \text{con}', \text{inn}'))$, where \perp is assumed to be a
 1074 minimal element. Finally we set $\mathcal{R}_t = \mathcal{R}_t^0$, correctness can be argued analogously to the case of
 1075 BNSL_{\leq}^+ .

1076 **t is an introduce node.** Let t' be the child of t in \mathcal{T} and let $\chi(t) = \chi(t') \cup \{v\}$. We initiate by
 1077 setting $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) = \perp$ for each $(\text{loc}, \text{con}, \text{inn}) \in S(t)$.

1078 For each $(\text{loc}', \text{con}', \text{inn}') \in S(t')$ and each $Q \subseteq \{ab \in A_{\chi(t)} \mid \{a, b\} \cap \{v\} \neq \emptyset\}$, we define:

- 1079 • $\text{loc} := \text{loc}' \cup Q$
- 1080 • $\text{con} := \text{trcl}(\text{con}' \cup \tilde{Q})$
- 1081 • $\text{inn}(x) := \text{inn}'(x) + |\{y \in \chi(t) \mid yx \in Q\}|$ for every $x \in \chi(t) \setminus \{v\}$
- 1082 • $\text{inn}(v) := |\{y \in \chi(t) \mid yv \in Q\}|$

1083 Let N and N' be the numbers of equivalence classes in con and con' correspondingly. If $N \neq N' +$
 1084 $1 - |Q|$ or $\text{inn}(x) > q$ for some $x \in \chi(t)$, discard this branch. Otherwise, let $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) :=$
 1085 $\max(\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}), \text{new})$ where $\text{new} = \mathcal{R}_{t'}(\text{loc}', \text{con}', \text{inn}') + \sum_{ab \in Q} f_b(a)$. As before, \perp is
 1086 assumed to be a minimal element here.

1087 Consider our final computed value of $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn})$ for some $(\text{loc}, \text{con}, \text{inn}) \in S(t)$.

1088 For correctness, assume that $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) = \tau$ for some $\tau \neq \perp$ and is obtained from
 1089 $(\text{loc}', \text{con}', \text{inn}'), Q$ defined as above. Then $\mathcal{R}_{t'}(\text{loc}', \text{con}', \text{inn}') = \tau - \sum_{ab \in Q} f_b(a)$. Construct a
 1090 directed graph D from the witness D' of $\mathcal{R}_{t'}(\text{loc}', \text{con}', \text{inn}')$ by adding v and the arcs specified in
 1091 Q . The equality $N = N' + 1 - |Q|$ guarantees that every such arc decreases the number of active
 1092 connected components by one, so D is a polytree. Moreover, $\text{inn}(x) \leq q$ for every $x \in \chi(t)$ and the
 1093 rest of vertices have in D the same parents as in D' , so $D \in \Psi_t$. In particular, $(\text{loc}, \text{con}, \text{inn})$ is a
 1094 snapshot of D in t and D witnesses $\mathcal{R}_t(\text{loc}, \text{con}, \text{inn}) \geq \mathcal{R}_{t'}(\text{loc}', \text{con}', \text{inn}') + \sum_{ab \in Q} f_b(a) = \tau$.

1095 On the other hand, if $\mathcal{R}_t(\text{loc}, \text{con}, \text{inn}) = \tau$ for some $\tau \neq \perp$, then there must exist a polytree
 1096 $D = (\chi_t^\downarrow, A)$ in Ψ_t that achieves a score of τ . Let Q be the restriction of A to arcs containing v ,
 1097 and let $D' = (\chi_t^\downarrow \setminus v, A \setminus Q)$, clearly $D' \in \Psi_{t'}$. Let $(\text{loc}', \text{con}', \text{inn}')$ be the snapshot of D' at
 1098 t' . Observe that $\text{loc} = \text{loc}' \cup Q, \text{con} = \text{trcl}(\text{con}' \cup \tilde{Q})$, inn differs from inn' by the numbers of
 1099 incoming arcs in Q and the score of D' is precisely equal to the score τ of D minus $\sum_{(a,b) \in Q} f_b(a)$.
 1100 Therefore $\mathcal{R}_{t'}(\text{loc}', \text{con}', \text{inn}') \geq \tau - \sum_{(a,b) \in Q} f_b(a)$ and in the algorithm $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) \geq$
 1101 $\mathcal{R}_{t'}(\text{loc}', \text{con}', \text{inn}') + \sum_{(a,b) \in Q} f_b(a) \geq \tau$. Equality then follows from the previous direction of the
 1102 correctness argument.

1103 Hence, at the end of our procedure we can correctly set $\mathcal{R}_t = \mathcal{R}_t^0$.

1104 **t is a join node.** Let t_1, t_2 be the two children of t in \mathcal{T} , recall that $\chi(t_1) = \chi(t_2) = \chi(t)$. We
 1105 initiate by setting $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) := \perp$ for each $(\text{loc}, \text{con}, \text{inn}) \in S(t)$.

1106 Let us branch over each $\text{loc}, \text{con}_1, \text{con}_2 \subseteq A_{\chi(t)}$ and $\text{inn}_1, \text{inn}_2 : \chi(t) \rightarrow [q]_0$. For every $b \in \chi(t)$
 1107 set $\text{inn}(b) = \text{inn}_1(b) + \text{inn}_2(b) - |\{a \mid ab \in \text{loc}\}|$. Let N_1 and N be the numbers of equivalence
 1108 classes in con_1 and $\text{trcl}(\text{con}_1 \cup \text{con}_2)$ correspondingly. If:

- 1109 • $\text{con}_1 \cap \text{con}_2 \neq \text{trcl}(\tilde{\text{loc}})$, and/or
- 1110 • $N - N_1 \neq \frac{1}{2} |\text{con}_2 \setminus \text{trcl}(\tilde{\text{loc}})|$, and/or
- 1111 • $\mathcal{R}_{t_1}(\text{loc}, \text{con}_1, \text{inn}_1) = \perp$, and/or
- 1112 • $\mathcal{R}_{t_2}(\text{loc}, \text{con}_2, \text{inn}_2) = \perp$, and/or
- 1113 • $\text{inn}(b) > q$ for some $b \in \chi(t)$

1114 then discard this branch. Otherwise, set $\text{con} = \text{trcl}(\text{con}_1 \cup \text{con}_2)$, $\text{doublecount} = \sum_{ab \in \text{loc}} f_b(a)$
 1115 and $\text{new} = \mathcal{R}_{t_1}(\text{loc}, \text{con}_1) + \mathcal{R}_{t_2}(\text{loc}, \text{con}_2) - \text{doublecount}$. We then set $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) :=$
 1116 $\max(\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}), \text{new})$ where \perp is once again assumed to be a minimal element.

1117 At the end of this procedure, we set $\mathcal{R}_t = \mathcal{R}_t^0$.

1118 For correctness, assume that $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) = \tau \neq \perp$ is obtained from $\text{loc}, \text{con}_1, \text{con}_2, \text{inn}_1, \text{inn}_2$
 1119 as above. Let $D_1 = (\chi_{t_1}^\downarrow, A_1)$ and $D_2 = (\chi_{t_2}^\downarrow, A_2)$ be polytrees witnessing $\mathcal{R}_{t_1}(\text{loc}, \text{con}_1, \text{inn}_1)$
 1120 and $\mathcal{R}_{t_2}(\text{loc}, \text{con}_2, \text{inn}_2)$ correspondingly. Recall from the proof of Theorem 13 that common
 1121 vertices of D_1 and D_2 are precisely $\chi(t)$, $\text{loc} = A_1 \cap A_2$ and inn specifies the number of parents
 1122 of every $b \in \chi(T)$ in $D = D_1 \cup D_2$. Numbers of active connected components of D and D_1 are
 1123 N and N_1 correspondingly. Observe that D can be constructed from D_1 by adding vertices and
 1124 arcs of D_2 . As $\text{con}_1 \cap \text{con}_2 = \text{trcl}(\text{loc})$, we can only add a path between vertices in $\chi(t)$ if it
 1125 didn't exist in D_1 . Hence $\frac{1}{2}|\text{con}_2 \setminus \text{trcl}(\text{loc})|$ specifies the number of paths between vertices in
 1126 $\chi(t)$ via forgotten vertices of $\chi_{t_2}^\downarrow$. The equality $N_1 - N = \frac{1}{2}|\text{con}_2 \setminus \text{trcl}(\text{loc})|$ means that adding
 1127 every such path decreases the number of active connected components of D_1 by one. As D_1 is
 1128 a polytree, D is a polytree as well, so $D \in \Psi_t$. The snapshot of D in t is $(\text{loc}, \text{con}, \text{inn})$ and
 1129 $\text{score}(D) = \sum_{ab \in A(D)} f_b(a) = \sum_{ab \in A_1} f_b(a) + \sum_{ab \in A_2} f_b(a) - \sum_{ab \in \text{loc}} f_b(a) = \text{score}(D_1) +$
 1130 $\text{score}(D_2) - \text{doublecount} = \mathcal{R}_{t_1}(\text{loc}, \text{con}_1, \text{inn}_1) + \mathcal{R}_{t_2}(\text{loc}, \text{con}_2, \text{inn}_2) - \text{doublecount} = \tau$.
 1131 So D witnesses that $\mathcal{R}_t(\text{loc}, \text{con}, \text{inn}) \geq \tau$.

1132 For the converse, assume that $\mathcal{R}_t(\text{loc}, \text{con}, \text{inn}) = \tau \neq \perp$ and D is a polytree witnessing this. Let
 1133 D_1 and D_2 be restrictions of D to $\chi_{t_1}^\downarrow$ and $\chi_{t_2}^\downarrow$ correspondingly, then $A(D_1) \cap A(D_2) = \text{loc}$,
 1134 in particular $D = D_1 \cup D_2$. Let $(\text{loc}, \text{con}_i, \text{inn}_i)$ be the snapshot of D_i in t_i , $i = 1, 2$.
 1135 $D = D_1 \cup D_2$ is a polytree, so any pair of vertices in $\chi(t)$ can not be connected by differ-
 1136 ent paths in D_1 and D_2 , i.e. $\text{con}_1 \cap \text{con}_2 = \text{trcl}(\text{loc})$. By the procedure of our algorithm,
 1137 $\mathcal{R}_t^0(\text{loc}, \text{con}, \text{inn}) \geq \mathcal{R}_{t_1}(\text{loc}, \text{con}_1, \text{inn}_1) + \mathcal{R}_{t_2}(\text{loc}, \text{con}_2, \text{inn}_2) - \text{doublecount} \geq \text{score}(D_1) +$
 1138 $\text{score}(D_2) - \sum_{ab \in \text{loc}} f_b(a) = \text{score}(D) = \tau$.

1139 Hence the resulting record \mathcal{R}_t is correct, which concludes the correctness proof of the algorithm.

1140 Since the nice tree-decomposition \mathcal{T} has $\mathcal{O}(n)$ nodes, the runtime of the algorithm is upper-bounded
 1141 by $\mathcal{O}(n)$ times the maximum time required to process each node. This is dominated by the time
 1142 required to process join nodes, for which there are at most $(2^{k^2})^3((q+1)^k)^2 = 8^{k^2} \cdot (q+1)^{2k}$ branches
 1143 corresponding to different choices of $\text{loc}, \text{con}_1, \text{con}_2, \text{inn}_1, \text{inn}_2$. Constructing $\text{trcl}(\text{con}_1 \cup \text{con}_2)$
 1144 and computing numbers of active connected components can be done in time $\mathcal{O}(k^3)$. Computing
 1145 doublecount and inn takes time at most $\mathcal{O}(k^2)$. So the record for a join node can be computed
 1146 in time $2^{\mathcal{O}(k^2)} \cdot q^{\mathcal{O}(k)}$. Hence, after we have computed a width-optimal tree-decomposition for
 1147 instance by Bodlaender's algorithm [4], the total runtime of the algorithm is upper-bounded by
 1148 $2^{\mathcal{O}(k^2)} \cdot q^{\mathcal{O}(k)} \cdot n$.

1149 Finally, to obtain the desired result for PL^+ , we can simply adapt the above algorithm by disregarding
 1150 the entry inn and disregard all explicit bounds on the in-degrees (e.g., in the definition of Ψ_t). The
 1151 runtime for this dynamic programming procedure is then $2^{\mathcal{O}(k^2)} \cdot n$. \square

1152 The situation is, however, completely different for PL^+ : unlike BNSL^+ , this problem is in fact
 1153 polynomial-time tractable. Indeed, it admits a simple reduction to the classical minimum edge-
 1154 weighted spanning tree problem.

1155 **Observation 20.** PL^+ is polynomial-time tractable.

1156 *Proof.* Consider an the superstructure graph G of an instance $\mathcal{I} = (V, \mathcal{F}, \ell)$ of PL^+ where we assign
 1157 to each edge $ab \in E(G)$ a weight $w(ab) = \max f_a(b), f_b(a)$, and recall that we can assume w.l.o.g.
 1158 that G is connected. Each spanning tree T of G with weight p can be transformed to a DAG D
 1159 over V with a score of p and whose skeleton is a tree by simply replacing each edge ab with the
 1160 arc ab or ba , depending on which achieves a higher score. On the other hand, each solution to \mathcal{I}
 1161 can be transformed into a spanning tree T of the same score by reversing this process. The claim
 1162 then follows from the fact that a minimum-weight spanning tree of a graph can be computed in time
 1163 $\mathcal{O}(|V| \cdot \log |V|)$. \square

1164 This coincides with the intuitive expectation that learning simple, more restricted networks could be
 1165 easier than learning general networks. We conclude our exposition with an example showcasing that
 1166 this is not true in general when comparing PL to BNSL. Grüttemeier et al. [24] recently showed that
 1167 $\text{PL}^{\neq 0}$ is $\text{W}[1]$ -hard when parameterized by the number of *dependent vertices*, which are vertices with
 1168 non-empty sets of candidate parents in the non-zero representation. For $\text{BNSL}^{\neq 0}$ we can show:

1169 **Theorem 21.** $\text{BNSL}^{\neq 0}$ is fixed-parameter tractable when parameterized by the number of dependent
 1170 vertices.

1171 *Proof.* Consider an algorithm \mathbb{B} for $\text{BNSL}^{\neq 0}$ which proceeds as follows. First, it identifies the set
 1172 X of dependent vertices in the input instance $\mathcal{I} = (V, \mathcal{F}, \ell)$, and then it branches over all possible
 1173 choices of arcs with both endpoints in X , i.e., it branches over each arc set $A \subseteq A_X$. This results
 1174 in at most 3^{k^2} branches, where $k = |X|$. In each branch and for each vertex $x \in X$, it now finds
 1175 the highest-scoring parent set among those which precisely match A on X , i.e., it first computes
 1176 $\Gamma_f^A(x) = \{P \in \text{parentsets}(x) \mid \forall w \in X \setminus \{x\} : w \in P \iff wp \in A\}$ and then computes
 1177 $\text{score}^A(x) = \max_{P \in \Gamma_f^A(x)} (f_x(P))$. It then compares $\sum_{x \in X} \text{score}^A(x)$ to ℓ ; if the former is at
 1178 least as large as the latter in at least one branch then \mathbb{B} outputs “Yes”, and otherwise it outputs no.

1179 The runtime of this algorithm is upper-bounded by $\mathcal{O}(3^{k^2} \cdot k \cdot |\mathcal{I}|)$. As for correctness, if \mathcal{I} admits
 1180 a solution D then we can construct a branch such that \mathbb{B} will output “Yes”: in particular, this must
 1181 occur when A is equal to the arcs of the subgraph of D induced on X . On the other hand, if \mathbb{B} outputs
 1182 “Yes” for some choice of A , we can construct a DAG D with a score of at least ℓ by extending A as
 1183 follows: for each $x \in X$ we choose a parent set $P \in \Gamma_f^A(x)$ which maximizes $f_x(P)$ and we add
 1184 arcs from each vertex in $P \setminus X$ to x . The score of this DAG will be precisely $\sum_{x \in X} \text{score}^A(x)$,
 1185 which concludes the proof. \square

1186 6 Concluding Remarks

1187 Our results provide a new set of tractability results that counterbalance the previously established
 1188 algorithmic lower bounds for BAYESIAN NETWORK STRUCTURE LEARNING and POLYTREE
 1189 LEARNING on “simple” superstructures. In particular, even though the problems remain $\text{W}[1]$ -hard
 1190 when parameterized by the vertex cover number of the superstructure [36, 24], we obtained fixed-
 1191 parameter tractability and a data reduction procedure using the feedback edge number and its localized
 1192 version. Together with our lower-bound result for treecut width, this completes the complexity map
 1193 for BNSL w.r.t. virtually all commonly considered graph parameters of the superstructure. Moreover,
 1194 we showed that if the input is provided with an additive representation instead of the non-zero
 1195 representation considered in previous theoretical works, the problems admit a dynamic programming
 1196 algorithm which guarantees fixed-parameter tractability w.r.t. the treewidth of the superstructure.

1197 This theoretical work follows up on previous complexity studies of the considered problems, and as
 1198 such we do not claim any immediate practical applications of the results. That being said, it would be
 1199 interesting to see if the polynomial-time data reduction procedure introduced in Theorem 3 could be
 1200 adapted and streamlined (and perhaps combined with other reduction rules which do not provide a
 1201 theoretical benefit, but perform well heuristically) to allow for a speedup of previously introduced
 1202 heuristics for the problem [43, 42], at least for some sets of instances.

1203 Last but not least, we’d like to draw attention to the *local feedback edge number* parameter introduced
 1204 in this manuscript specifically to tackle BNSL. This generalization of the feedback edge set has not
 1205 yet been considered in graph-theoretic works; while it is similar in spirit to the recent push towards
 1206 measuring the so-called *elimination distance* of a graph to a target class, it is not captured by that
 1207 notion. Crucially, we believe that the applications of this parameter go beyond BNSL; all indications
 1208 suggest that it may be used to achieve tractability also for purely graph-theoretic problems where
 1209 previously only tractability w.r.t. fen was known.

1210 References

1211 [1] Samuel Frederic Barr. Courcelle’s theorem: Overview and applications. Master’s thesis, 2020.

- 1212 [2] Matthias Bentert, Roman Haag, Christian Hofer, Tomohiro Koana, and André Nichterlein.
1213 Parameterized complexity of min-power asymmetric connectivity. *Theory of Computing Systems*,
1214 64:1158–1182, 2020.
- 1215 [3] Nadja Betzler, Robert Brederbeck, Rolf Niedermeier, and Johannes Uhlmann. On bounded-
1216 degree vertex deletion parameterized by treewidth. *Discret. Appl. Math.*, 160(1-2):53–60,
1217 2012.
- 1218 [4] Hans L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth.
1219 *SIAM J. Comput.*, 25(6):1305–1317, 1996.
- 1220 [5] David Maxwell Chickering. Learning Bayesian networks is NP-complete. In *Learning from*
1221 *data (Fort Lauderdale, FL, 1995)*, volume 112 of *Lecture Notes in Statist.*, pages 121–130.
1222 Springer Verlag, 1996.
- 1223 [6] Bruno Courcelle. The monadic second-order logic of graphs. i. recognizable sets of finite graphs.
1224 *Inf. Comput.*, 85(1):12–75, 1990.
- 1225 [7] Bruno Courcelle and Joost Engelfriet. *Graph Structure and Monadic Second-Order Logic - A*
1226 *Language-Theoretic Approach*, volume 138 of *Encyclopedia of mathematics and its applications*.
1227 Cambridge University Press, Cambridge, 2012.
- 1228 [8] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin
1229 Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.
- 1230 [9] Sanjoy Dasgupta. Learning polytrees. In Kathryn B. Laskey and Henri Prade, editors, *UAI '99:*
1231 *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence, Stockholm,*
1232 *Sweden, July 30 - August 1, 1999*, pages 134–141. Morgan Kaufmann, 1999.
- 1233 [10] Camil Demetrescu and Irene Finocchi. Combinatorial algorithms for feedback problems in
1234 directed graphs. *Inf. Process. Lett.*, 86(3):129–136, 2003.
- 1235 [11] Reinhard Diestel. *Graph Theory, 4th Edition*, volume 173 of *Graduate texts in mathematics*.
1236 Springer, 2012.
- 1237 [12] Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*. Texts
1238 in Computer Science. Springer, 2013.
- 1239 [13] Eduard Eiben, Robert Ganian, Iyad Kanj, and Stefan Szeider. The parameterized complexity of
1240 cascading portfolio scheduling. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer,
1241 Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Infor-*
1242 *mation Processing Systems 32: Annual Conference on Neural Information Processing Systems*
1243 *2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 7666–7676, 2019.
- 1244 [14] Gal Elidan and Stephen Gould. Learning bounded treewidth bayesian networks. In Daphne
1245 Koller, Dale Schuurmans, Yoshua Bengio, and Léon Bottou, editors, *Advances in Neural*
1246 *Information Processing Systems 21, Proceedings of the Twenty-Second Annual Conference on*
1247 *Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 8-11,*
1248 *2008*, pages 417–424. Curran Associates, Inc., 2008.
- 1249 [15] Jirí Fiala, Petr A. Golovach, and Jan Kratochvíl. Parameterized complexity of coloring problems:
1250 Treewidth versus vertex cover. *Theor. Comput. Sci.*, 412(23):2513–2523, 2011.
- 1251 [16] Robert Ganian, Iyad A. Kanj, Sebastian Ordyniak, and Stefan Szeider. Parameterized algorithms
1252 for the matrix completion problem. In Jennifer G. Dy and Andreas Krause, editors, *Proceedings*
1253 *of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan,*
1254 *Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*,
1255 pages 1642–1651. PMLR, 2018.
- 1256 [17] Robert Ganian, Eun Jung Kim, and Stefan Szeider. Algorithmic applications of tree-cut width.
1257 In Giuseppe F. Italiano, Giovanni Pighizzini, and Donald Sannella, editors, *Mathematical*
1258 *Foundations of Computer Science 2015 - 40th International Symposium, MFCS 2015, Milan,*
1259 *Italy, August 24-28, 2015, Proceedings, Part II*, volume 9235 of *Lecture Notes in Computer*
1260 *Science*, pages 348–360. Springer, 2015.

- 1261 [18] Robert Ganian, Fabian Klute, and Sebastian Ordyniak. On structural parameterizations of the
1262 bounded-degree vertex deletion problem. *Algorithmica*, 83(1):297–336, 2021.
- 1263 [19] Robert Ganian and Sebastian Ordyniak. The complexity landscape of decompositional param-
1264 eters for ILP. *Artif. Intell.*, 257:61–71, 2018.
- 1265 [20] Robert Ganian and Sebastian Ordyniak. The power of cut-based parameters for computing
1266 edge-disjoint paths. *Algorithmica*, 83(2):726–752, 2021.
- 1267 [21] Robert Ganian, Sebastian Ordyniak, and Ramanujan Sridharan. On structural parameterizations
1268 of the edge disjoint paths problem. *Algorithmica*, 83:1605–1637, 2021.
- 1269 [22] Serge Gaspers, Mikko Koivisto, Mathieu Liedloff, Sebastian Ordyniak, and Stefan Szeider. On
1270 finding optimal polytrees. *Theor. Comput. Sci.*, 592:49–58, 2015.
- 1271 [23] Fanica Gavril. Some np-complete problems on graphs. In *Proceedings of the 11th Conference*
1272 *on Information Sciences and Systems*, pages 91–95, 1977.
- 1273 [24] Niels Grüttemeier, Christian Komusiewicz, and Nils Morawietz. On the parameterized complex-
1274 ity of polytree learning. In *Proceedings of the 30th International Joint Conference on Artificial*
1275 *Intelligence (IJCAI-21)*, 2021. To appear; preprint available at <https://arxiv.org/abs/2105.09675>.
- 1276 [25] Niels Grüttemeier and Christian Komusiewicz. Learning bayesian networks under sparsity
1277 constraints: A parameterized complexity analysis. In Christian Bessiere, editor, *Proceedings of*
1278 *the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages
1279 4245–4251. ijcai.org, 2020.
- 1280 [26] Haipeng Guo and William Hsu. A survey of algorithms for real-time bayesian network inference.
1281 In *Working Notes of the Joint AAAI/UAI/KDD Workshop on Real-Time Decision Support and*
1282 *Diagnosis Systems*, 2002.
- 1283 [27] T. Kloks. *Treewidth: Computations and Approximations*. Springer Verlag, Berlin, 1994.
- 1284 [28] Yasuaki Kobayashi and Hisao Tamaki. Treedepth parameterized by vertex cover number. In
1285 Jiong Guo and Danny Hermelin, editors, *11th International Symposium on Parameterized and*
1286 *Exact Computation, IPEC 2016, August 24-26, 2016, Aarhus, Denmark*, volume 63 of *LIPICs*,
1287 pages 18:1–18:11. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.
- 1288 [29] Janne H. Korhonen and Pekka Parviainen. Exact learning of bounded tree-width bayesian
1289 networks. In *Proceedings of the Sixteenth International Conference on Artificial Intelligence*
1290 *and Statistics, AISTATS 2013, Scottsdale, AZ, USA, April 29 - May 1, 2013*, volume 31 of *JMLR*
1291 *Workshop and Conference Proceedings*, pages 370–378. [JMLR.org](http://jmlr.org), 2013.
- 1292 [30] Janne H. Korhonen and Pekka Parviainen. Tractable bayesian network structure learning with
1293 bounded vertex cover number. In Corinna Cortes, Neil D. Lawrence, Daniel D. Lee, Masashi
1294 Sugiyama, and Roman Garnett, editors, *Advances in Neural Information Processing Systems*
1295 *28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015,*
1296 *Montreal, Quebec, Canada*, pages 622–630, 2015.
- 1297 [31] Alexander Langer, Felix Reidl, Peter Rossmanith, and Somnath Sikdar. Practical algorithms for
1298 MSO model-checking on tree-decomposable graphs. *Comput. Sci. Rev.*, 13-14:39–74, 2014.
- 1299 [32] Dániel Marx and Paul Wollan. Immersions in highly edge connected graphs. *SIAM J. Discret.*
1300 *Math.*, 28(1):503–520, 2014.
- 1301 [33] Christopher Meek. Finding a path is harder than finding a tree. *J. Artif. Intell. Res.*, 15:383–389,
1302 2001.
- 1303 [34] Jaroslav Nesetril and Patrice Ossona de Mendez. *Sparsity - Graphs, Structures, and Algorithms*,
1304 volume 28 of *Algorithms and Combinatorics*. Springer Verlag, 2012.
- 1305 [35] Rolf Niedermeier. *Invitation to Fixed-Parameter Algorithms*. Oxford Lecture Series in Mathe-
1306 matics and its Applications. Oxford University Press, Oxford, 2006.

- 1307 [36] Sebastian Ordyniak and Stefan Szeider. Parameterized complexity results for exact bayesian
1308 network structure learning. *J. Artif. Intell. Res.*, 46:263–302, 2013.
- 1309 [37] Judea Pearl. *Probabilistic reasoning in intelligent systems - networks of plausible inference*.
1310 Morgan Kaufmann series in representation and reasoning, Morgan Kaufmann, 1989.
- 1311 [38] Eric Perrier, Seiya Imoto, and Satoru Miyano. Finding optimal bayesian network given a
1312 super-structure. *J. Mach. Learn. Res.*, 9(9):2251–2286, 2008.
- 1313 [39] Neil Robertson and P. D. Seymour. Graph minors. II. Algorithmic aspects of tree-width. *J.*
1314 *Algorithms*, 7(3):309–322, 1986.
- 1315 [40] Neil Robertson and Paul D. Seymour. Graph minors. i. excluding a forest. *J. Comb. Theory, Ser.*
1316 *B*, 35(1):39–61, 1983.
- 1317 [41] Javad Safaei, Ján Manuch, and Ladislav Stacho. Learning polytrees with constant number of
1318 roots from data. In Stephen Cranefield and Abhaya C. Nayak, editors, *AI 2013: Advances in*
1319 *Artificial Intelligence - 26th Australasian Joint Conference, Dunedin, New Zealand, December*
1320 *1-6, 2013. Proceedings*, volume 8272 of *Lecture Notes in Computer Science*, pages 447–452.
1321 Springer, 2013.
- 1322 [42] Mauro Scanagatta, Giorgio Corani, Cassio P. de Campos, and Marco Zaffalon. Learning
1323 treewidth-bounded bayesian networks with thousands of variables. In Daniel D. Lee, Masashi
1324 Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett, editors, *Advances in Neural*
1325 *Information Processing Systems 29: Annual Conference on Neural Information Processing*
1326 *Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 1462–1470, 2016.
- 1327 [43] Mauro Scanagatta, Cassio P. de Campos, Giorgio Corani, and Marco Zaffalon. Learning
1328 bayesian networks with thousands of variables. In Corinna Cortes, Neil D. Lawrence, Daniel D.
1329 Lee, Masashi Sugiyama, and Roman Garnett, editors, *Advances in Neural Information Process-*
1330 *ing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December*
1331 *7-12, 2015, Montreal, Quebec, Canada*, pages 1864–1872, 2015.
- 1332 [44] Kirill Simonov, Fedor V. Fomin, Petr A. Golovach, and Fahad Panolan. Refined complexity of
1333 PCA with outliers. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of*
1334 *the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long*
1335 *Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages
1336 5818–5826. PMLR, 2019.
- 1337 [45] Ioannis Tsamardinos, Laura E. Brown, and Constantin F. Aliferis. The max-min hill-climbing
1338 bayesian network structure learning algorithm. *Mach. Learn.*, 65(1):31–78, 2006.
- 1339 [46] Johannes Uhlmann and Mathias Weller. Two-layer planarization parameterized by feedback
1340 edge set. *Theor. Comput. Sci.*, 494:99–111, 2013.
- 1341 [47] René van Bevern, Till Fluschnik, and Oxana Yu. Tsidulko. Parameterized algorithms and data
1342 reduction for the short secluded s-t-path problem. *Networks*, 75(1):34–63, 2020.
- 1343 [48] Paul Wollan. The structure of graphs not admitting a fixed immersion. *J. Comb. Theory, Ser. B*,
1344 110:47–66, 2015.

1345 Checklist

- 1346 1. For all authors...
- 1347 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s
1348 contributions and scope? [Yes]
- 1349 (b) Did you describe the limitations of your work? [Yes] The limitations are given by
1350 assumptions in theorem formulations.
- 1351 (c) Did you discuss any potential negative societal impacts of your work? [N/A] This is
1352 a purely theoretical contribution that provides new insights into the complexity of a
1353 prominent problem in AI, and as such we do not see any conceivable negative societal
1354 impacts of this work.

- 1355 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
1356 them? [Yes]
- 1357 2. If you are including theoretical results...
- 1358 (a) Did you state the full set of assumptions of all theoretical results? [Yes]
1359 (b) Did you include complete proofs of all theoretical results? [Yes] Due to space con-
1360 straints, full proofs are provided in the supplementary material.
- 1361 3. If you ran experiments...
- 1362 (a) Did you include the code, data, and instructions needed to reproduce the main experi-
1363 mental results (either in the supplemental material or as a URL)? [N/A] The paper is
1364 purely theoretical.
- 1365 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
1366 were chosen)? [N/A]
- 1367 (c) Did you report error bars (e.g., with respect to the random seed after running experi-
1368 ments multiple times)? [N/A]
- 1369 (d) Did you include the total amount of compute and the type of resources used (e.g., type
1370 of GPUs, internal cluster, or cloud provider)? [N/A]
- 1371 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 1372 (a) If your work uses existing assets, did you cite the creators? [N/A]
1373 (b) Did you mention the license of the assets? [N/A]
1374 (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
1375
- 1376 (d) Did you discuss whether and how consent was obtained from people whose data you're
1377 using/curating? [N/A]
- 1378 (e) Did you discuss whether the data you are using/curating contains personally identifiable
1379 information or offensive content? [N/A]
- 1380 5. If you used crowdsourcing or conducted research with human subjects...
- 1381 (a) Did you include the full text of instructions given to participants and screenshots, if
1382 applicable? [N/A] We didn't use neither crowdsourcing nor conducted research with
1383 human subjects.
- 1384 (b) Did you describe any potential participant risks, with links to Institutional Review
1385 Board (IRB) approvals, if applicable? [N/A]
1386 (c) Did you include the estimated hourly wage paid to participants and the total amount
1387 spent on participant compensation? [N/A]