

On Injective Embeddings of Tree Patterns

Jakub Michaliszyn¹, Anca Muscholl², Sławek Staworko³, Piotr Wieczorek¹, and Zhilin Wu⁴

¹ University of Wrocław

² LaBRI, University of Bordeaux

³ Mostrare, INRIA Lille, University of Lille

⁴ State Key Laboratory of Computer Science, Institute of Software, Chinese Academy of Sciences

Abstract. We study three different kinds of embeddings of tree patterns: weakly-injective, ancestor-preserving, and lca-preserving. While each of them is often referred to as injective embedding, they form a proper hierarchy and their computational properties vary (from P to NP-complete). We present a thorough study of the complexity of the model checking problem i.e., is there an embedding of a given tree pattern in a given tree, and we investigate the impact of various restrictions imposed on the tree pattern: bound on the degree of a node, bound on the height, and type of allowed labels and edges.

1 Introduction

An embedding is a fundamental notion with numerous applications in computer science e.g., in graph pattern matching (cf. [4]). Usually, an embedding is defined as a structure-preserving mapping that is typically required to be injective. Tree patterns are a special class of graph patterns that found applications, for instance in XML databases [11, 1] where they form a functional equivalent of (acyclic) conjunctive queries for relational databases. Tree patterns are typically matched against trees and are allowed to use special descendant edges (double lines in Fig. 1) that can be mapped to paths rather than to single edges as it is the case with the standard child edges.

Traditionally, the semantics of tree patterns for XML is defined using non-injective embeddings [1, 15] (Fig. 1(a)), which is reminiscent of relational data. Since XML data has more structure, it makes sense to exploit the tree structure when defining tree pattern embeddings. In this context, it is interesting to consider injective embeddings [3, 8, 5, 11]. However, the use of descendant edges makes it cumbersome to define what exactly an injective embedding of a tree pattern should be, and consequently, different notions have been employed.

A *weakly-injective* embedding requires only the mapping to be injective and recent developments in graph matching suggest that such embeddings are crucial for expressing important patterns occurring in real life databases [5]. They are a natural choice when we do not wish to constrain in any way the vertical relationship of the images of two children of some node connected with descendant edges. However, descendant edges can be mapped to paths that interleave, which

means that even if there is a weakly-injective embedding between a tree pattern and a tree, there need not be a structural similarity between the tree and the tree pattern (Fig. 1(b)). This is contrary to the structure-preservation nature of embeddings and hence the prefix *weakly*. One could strengthen the restriction and prevent the embedding from introducing vertical relationships between the nodes, which gives us *ancestor-preserving* embeddings [3]. In this case two descendant edges are mapped into paths that might overlap at the beginning but eventually branch (Fig. 1(c)). Finally, we can go one step further and require the paths not to overlap at all, which translates to *lca-preserving* embeddings [8] i.e., embeddings that preserve lowest common ancestors of any pair of nodes (Fig. 1(d)).

Unfortunately, there is a lack of a systematic and thorough treatment of injective embeddings and there is a tendency to name each of the embeddings above as simply injective, which could be potentially confusing and error-prone. This paper fills this gap and shows that injective embeddings form a proper hierarchy and that their computational properties vary significantly (from P to NP-complete). This further strengthens our belief that the different injective embeddings should not be confused. More precisely, we study the complexity of the model checking problem i.e., given a tree pattern p and a tree t is there an embedding (of a given type) of p in t , and we investigate the impact of various restrictions imposed on the tree pattern: bound on the degree of a node, bound on the height, and type of allowed labels and edges.

Our results show that while lca-preserving embeddings are in P, both weakly-injective and ancestor-preserving embeddings are NP-complete. Bounding the height of the pattern practically does not change the picture but bounding the degree of a node in the pattern renders ancestor-preserving embeddings tractable while weakly-injective embeddings remain NP-complete. Our results show that the high complexity springs from the use of descendant edges: if we disallow them, the hierarchy collapses and all injective embeddings fall into P. On the other hand, the use of node label is not essential, the complexity remains unchanged even if we consider tree patterns using the wildcard symbol only, essentially patterns that query only structural properties of the tree.

Injective embeddings of tree patterns are closely related to a number of well-established and studied notions, including *tree inclusion* [12, 18], *minor containment* [16, 17], *subgraph homeomorphism* [2, 14], and *graph pattern matching* [5]. Not surprisingly, a number of our results is subsumed by or can be easily obtained from existing results, and conversely, there are some that are subsumed by ours (see Sec. 5 for a complete discussion of related work). The principal aim of this paper is, however, to catalog the different kinds of injective embeddings of tree patterns and identify what aspects of tree patterns lead to intractability. To that end, all our reductions and algorithms are new and the reductions clearly illustrate the source of complexity of injective tree patterns.

This paper is organized as follows. In Sec. 2 we define basic notions and in Sec. 3 we define formally the three types of injective embeddings of tree patterns. In Sec. 4 we study the model checking problem of the injective embeddings.

Discussion of related work is in Sec. 5 and in Sec. 6 we summarize our results and outline further directions of study. Some proofs have been moved to appendix.

2 Preliminaries

We assume a fixed and finite set of node labels Σ and use a wildcard symbol \star not present in Σ . A *tree pattern* [11, 1] is a tuple $p = (N_p, root_p, lab_p, child_p, desc_p)$, where N_p is a finite set of nodes, $root_p \in N_p$ is the root node, $lab_p : N_p \rightarrow \Sigma \cup \{\star\}$ is a labeling function, $child_p \subseteq N_p \times N_p$ is a set of child edges, and $desc_p \subseteq N_p \times N_p$ is a set of (proper) descendant edges. We assume that $child_p \cap desc_p = \emptyset$ and that the relation $child_p \cup desc_p$ is acyclic and require every non-root node to have exactly one predecessor in this relation. A *tree* is a tree pattern that has no descendant edges and uses no wildcard symbols \star .

An example of a tree pattern can be found in Fig. 1 (descendant edges are drawn with double lines). Sometimes, we use unranked terms to represent trees and the standard XPath syntax to represent tree patterns. XPath allows to navigate the tree with a syntax similar to directory paths used in the UNIX file system. For instance, in Fig. 1 p_0 can be written as $f/a[./b/c]/b$. In the sequel, we use p, p_0, p_1, \dots to range over tree patterns and t, t_0, t_1, \dots to range over trees.

Given a binary relation R , we denote by R^+ the transitive closure of R , and by R^* the transitive and reflexive closure of R . Now, fix a pattern p and take two of its nodes $n, n' \in N_p$. We say that n' is a *|*-child of n if $(n, n') \in child_p$, n' is a *||*-child of n if $(n, n') \in desc_p$, and n' is simply a *child* of n in p if $(n, n') \in child_p \cup desc_p$. Also, n' is a *descendant* of n , and n an *ancestor* of n' , if $(n, n') \in (child_p \cup desc_p)^*$. Note that descendantship and ancestorship are reflexive: a node is its own ancestor and its own descendant. The *depth* of a node n in p is the length of the path from the root node $root_p$ to n , and here, a path is a sequence of edges, and in particular, the depth of the root node is 0. The *lowest common ancestor* of n and n' in p , denoted by $lca_p(n, n')$, is the deepest node that is an ancestor of n and n' . The *size* of a tree pattern p , denoted $|p|$, is the number of its nodes. The *degree* of a node n , denoted $deg_p(n)$, is the number of its children. The *height* of a tree pattern p , denoted $height(p)$, is the depth of its deepest node.

The standard semantics of tree patterns is defined using non-injective embeddings which map the nodes of a tree pattern to the nodes of a tree in a manner that respects the wildcard and the semantics of the edges. Formally, an *embedding* of a tree pattern p in a tree t is a function $h : N_p \rightarrow N_t$ such that:

1. $h(root_p) = root_t$,
2. for every $(n, n') \in child_p$, $(h(n), h(n')) \in child_t$,
3. for every $(n, n') \in desc_p$, $(h(n), h(n')) \in (child_t)^+$,
4. for every $n \in N_p$, $lab_t(h(n)) = lab_p(n)$ unless $lab_p(n) = \star$.

We write $t \leq_{\text{std}} p$ if there exists a (standard) embedding of p in t . Note that the semantics of a descendant edge of the tree pattern is in fact that of a *proper descendant*: a descendant edge is mapped to a nonempty path in the tree.

3 Injective embeddings

We identify three subclasses of injective embeddings that restrict the standard embedding by adding one additional condition each. First, we have the *weakly-injective* embedding of p in t ($t \preceq_{\text{inj}} p$):

5'. h is an injective function i.e., $h(n_1) \neq h(n_2)$ for any two different nodes n_1 and n_2 of p .

Next, we have the *ancestor-preserving* embedding of p in t ($t \preceq_{\text{anc}} p$):

5''. $h(n_1)$ is an ancestor of $h(n_2)$ in t if and only if n_1 is an ancestor of n_2 in p , for any two nodes n_1 and n_2 of p . More formally, for any $n_1, n_2 \in N_p$

$$(h(n_1), h(n_2)) \in \text{child}_t^* \iff (n_1, n_2) \in (\text{child}_p \cup \text{desc}_p)^*.$$

Finally, we have the *lca-preserving* embedding of p in t ($t \preceq_{\text{lca}} p$):

5'''. h maps the lowest common ancestor of nodes n_1 and n_2 to the lowest common ancestor of $h(n_1)$ and $h(n_2)$ i.e., for any pair of nodes n_1 and n_2 of p we have $\text{lca}_t(h(n_1), h(n_2)) = h(\text{lca}_p(n_1, n_2))$.

In Fig. 1 we illustrate various embeddings of a tree pattern p_0 .

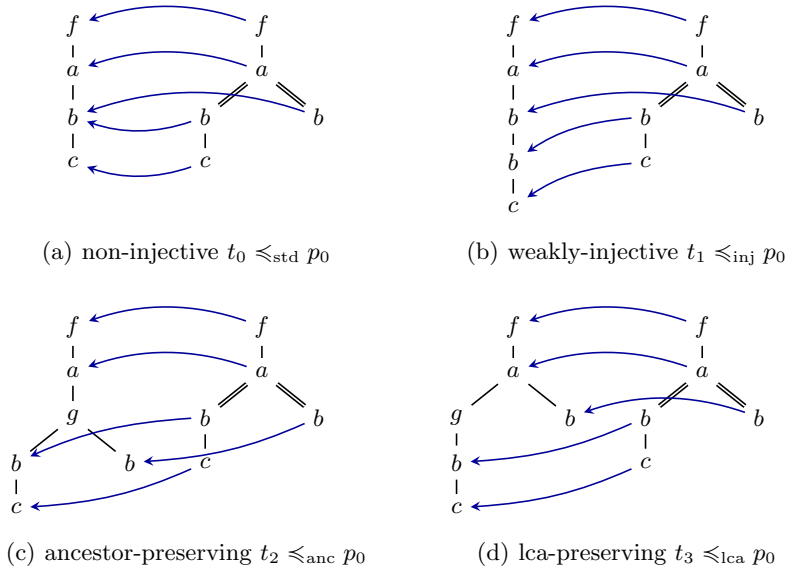


Fig. 1. Embeddings of a tree pattern p_0 .

We point out that injective embeddings form a hierarchy, and in particular, lca-preserving and ancestor-preserving embeddings are weakly-injective.

Proposition 3.1. *For any tree t and tree pattern p , 1) $t \preceq_{\text{lca}} p \Rightarrow t \preceq_{\text{anc}} p$, 2) $t \preceq_{\text{anc}} p \Rightarrow t \preceq_{\text{inj}} p$, and 3) $t \preceq_{\text{inj}} p \Rightarrow t \preceq_{\text{std}} p$.*

It is also easy to see that the hierarchy is proper. For that, take Fig. 1 and note that $t_0 \preceq_{\text{std}} p_0$ but $t_0 \not\preceq_{\text{inj}} p_0$, $t_1 \preceq_{\text{inj}} p_0$ but $t_1 \not\preceq_{\text{anc}} p_0$, and finally, $t_2 \preceq_{\text{anc}} p_0$ but $t_2 \not\preceq_{\text{lca}} p_0$. We point out, however, that the hierarchy of injective embeddings collapses if we disallow descendant edges in tree patterns.

Proposition 3.2. *For any tree t and any tree pattern p that does not use descendant edges, $t \preceq_{\text{inj}} p$ iff $t \preceq_{\text{anc}} p$ iff $t \preceq_{\text{lca}} p$.*

Furthermore, if we consider path patterns i.e., tree patterns whose nodes have at most one child, there is no difference between any of the injective embeddings and the standard embedding.

Proposition 3.3. *For any tree t and any path pattern p , $t \preceq_{\text{std}} p$ iff $t \preceq_{\text{inj}} p$ iff $t \preceq_{\text{anc}} p$ iff $t \preceq_{\text{lca}} p$.*

4 Complexity of injective embeddings

For a type of embedding $\theta \in \{\text{inj}, \text{anc}, \text{lca}\}$ we define the corresponding (unconstrained) decision problem:

$$\mathcal{M}_\theta = \{(t, p) \mid t \preceq_\theta p\}.$$

Additionally, we investigate several constrained variants of this problem. First, we restrict the degree of a node in the tree pattern by a constant $k \geq 0$,

$$\mathcal{M}_\theta^{\text{D} \leq k} = \{(t, p) \mid t \preceq_\theta p, \forall n \in N_p. \text{deg}_p(n) \leq k\}.$$

Next, we restrict the height of the tree pattern by a constant $k \geq 0$,

$$\mathcal{M}_\theta^{\text{H} \leq k} = \{(t, p) \mid t \preceq_\theta p, \text{height}(p) \leq k\}.$$

We also investigate the importance of labels in tree patterns as opposed to those that are label-oblivious and query only the structure of the tree i.e., tree patterns that use \star only.

$$\mathcal{M}_\theta^\star = \{(t, p) \mid t \preceq_\theta p, \forall n \in N_p. \text{lab}_p(n) = \star\}.$$

It is also interesting to see if disallowing \star may change the picture.

$$\mathcal{M}_\theta^\circ = \{(t, p) \mid t \preceq_\theta p, \forall n \in N_p. \text{lab}_p(n) \neq \star\}.$$

Finally, we restrict the use of child and descendant edges in the tree pattern.

$$\mathcal{M}_\theta^\perp = \{(t, p) \mid t \preceq_\theta p, \text{desc}_p = \emptyset\} \quad \text{and} \quad \mathcal{M}_\theta^\parallel = \{(t, p) \mid t \preceq_\theta p, \text{child}_p = \emptyset\}.$$

We make several general observations. First, we point out that the conditions on the various injective embeddings can be easily verified and every embedding is a mapping whose size is bounded by the size of the tree pattern. Therefore,

Proposition 4.1. $\mathcal{M}_\theta, \mathcal{M}_\theta^{D \leq k}, \mathcal{M}_\theta^{H \leq k}, \mathcal{M}_\theta^\star, \mathcal{M}_\theta^\circ, \mathcal{M}_\theta^\perp$, and $\mathcal{M}_\theta^\parallel$ are in NP for any $\theta \in \{\text{inj}, \text{anc}, \text{lca}\}$ and $k \geq 0$.

By Prop. 3.3, for path patterns we employ the existing polynomial algorithm [6].

Proposition 4.2. $\mathcal{M}_\theta^{D \leq 1}$ is in P for any $\theta \in \{\text{inj}, \text{anc}, \text{lca}\}$.

Finally, by Prop. 3.2 and Thm. 4.15, which shows the tractability of lca-preserving embeddings, we get the following.

Proposition 4.3. \mathcal{M}_θ^\perp is in P for any $\theta \in \{\text{inj}, \text{anc}, \text{lca}\}$.

4.1 Weakly-injective embeddings

Theorem 4.4. \mathcal{M}_{inj} is NP-complete.

Proof. We reduce SAT to \mathcal{M}_{inj} . We take a CNF formula $\varphi = c_1 \wedge \dots \wedge c_k$ over the variables x_1, \dots, x_n and for every variable x_i we construct two (linear) trees $X_i = x_i(\pi_1(\pi_2(\dots \pi_{k-1}(\pi_k)\dots)))$ and $\bar{X}_i = x_i(\bar{\pi}_1(\bar{\pi}_2(\dots \bar{\pi}_{k-1}(\bar{\pi}_k)\dots)))$, where $\pi_j = c_j$ if the clause c_j uses the literal x_i and $\pi_j = \perp$ otherwise, and analogously, $\bar{\pi}_j = c_j$ if the clause c_j uses the literal $\neg x_i$ and $\bar{\pi}_j = \perp$ otherwise. The constructed tree is

$$t_\varphi = r(X_1, \bar{X}_1, X_2, \bar{X}_2, \dots, X_n, \bar{X}_n)$$

and the constructed tree pattern is

$$p_\varphi = r[./Y_1][./Y_2] \dots [./Y_n][./c_1][./c_2] \dots [./c_k],$$

where $Y_i = x_i/\star/\star/\dots/\star$ with exactly k repetitions of \star . Figure 2 illustrates the




Fig. 2. Reduction to \mathcal{M}_{inj} for $\varphi = (x_1 \vee \neg x_3) \wedge (x_1 \vee \neg x_2 \vee x_3) \wedge (\neg x_1 \vee \neg x_2)$.

reduction for $\varphi = (x_1 \vee \neg x_3) \wedge (x_1 \vee \neg x_2 \vee x_3) \wedge (\neg x_1 \vee \neg x_2)$. We claim that

$$(t_\varphi, p_\varphi) \in \mathcal{M}_{\text{inj}} \iff \text{SAT}.$$

For the *if* part, we take a valuation V satisfying φ and construct a weakly-injective embedding h as follows. The fragment $[./Y_i]$ is mapped to \bar{X}_i if $V(x_i) =$

true and to X_i if $V(x_i) = \text{false}$. For the clause c_j we pick one literal satisfied by V and w.l.o.g. assume it is x_i i.e., c_j uses x_i and $V(x_i) = \text{true}$. Then, the embedding h maps the fragment $[./c_j]$ to the node c_j in the tree fragment X_i . Clearly, the constructed embedding is an injective function.

For the *only if* part, we take a weakly-injective embedding h and construct a satisfying valuation V as follows. If the fragment $[./Y_i]$ is mapped to X_i , then $V(x_i) = \text{false}$ and if $[./\bar{Y}_i]$ is mapped to \bar{X}_i , then $V(x_i) = \text{true}$. To show that φ is satisfied by V we take any clause c_j and check where h maps the fragment $[./c_j]$. W.l.o.g. assume that it is X_i and since h is weakly-injective, Y_i is mapped to \bar{X}_i , and consequently, $V(x_i) = \text{true}$. Hence, V satisfies c_j . \square

We observe that in the reduction above the use of the child edges in the tree pattern is not essential and they can be replaced by descendant edges.

Corollary 4.5. $\mathcal{M}_{\text{inj}}^{\parallel}$ is NP-complete.

Furthermore, the proof of Thm. 4.4 can be easily adapted to the bounded degree setting. Indeed, one can easily show that for any tree $t = r(t_1, \dots, t_k)$ and any tree pattern $p = r[./p_1] \dots [./p_m]$, $t \preceq_{\text{inj}} p$ if and only if $t' \preceq_{\text{inj}} p'$, where $t' = A_1(\dots A_m(t_1, \dots, t_k) \dots)$, $p' = A_1[./p_1] / \dots / A_m[./p_m]$, and A_1, \dots, A_m are new symbols not used in p . This observation, when applied to the tree pattern in the reduction above, allows to reduce the degree of the root node and to obtain a tree pattern of degree bounded by 2. Note, however, that this technique does not allow to reduce the degree of nodes in arbitrary tree patterns.

Corollary 4.6. $\mathcal{M}_{\text{inj}}^{\text{D} \leq k}$ is NP-complete for any $k \geq 2$.

A reduction similar to the one presented above can be used to construct patterns whose height is exactly 2.

Theorem 4.7. $\mathcal{M}_{\text{inj}}^{\text{H} \leq k}$ is NP-complete for any $k \geq 2$.

If we consider patterns of depth 1, where the children of the root node are leaves, then a diligent counting technique suffices to solve the problem.

Proposition 4.8. $\mathcal{M}_{\text{inj}}^{\text{H} \leq 1}$ is in P.

Proof. Fix a tree pattern p whose depth is 1 and a tree t . For $a \in \Sigma \cup \{\star\}$ we denote by p_a^{\perp} the number of a -labeled \perp -children of root_p , by p_a^{\parallel} the number of a -labeled \parallel -children of root_p , and by $t_a^{\leq i}$ and $t_a^{\geq i}$ the numbers of a -labeled nodes of t at depths $\leq i$ and $\geq i$ resp.

We attempt to construct a weakly-injective embedding of p to t using the following strategy: (1) we map the nodes of p_a^{\perp} to nodes of $t_a^{\leq 1}$, (2) we map the nodes of p_a^{\parallel} to nodes of $t_a^{\geq 2}$ and if $p_a^{\parallel} > t_a^{\geq 2}$, we map the remaining $p_a^{\parallel} - t_a^{\geq 2}$ nodes to the nodes of $t_a^{\leq 1}$, (3) we map the nodes of p_{\star}^{\perp} to the remaining nodes of t at depth 1, and (4) we map the nodes of p_{\star}^{\parallel} to the remaining nodes of t .

Clearly, this procedure succeeds and a weakly-injective embedding can be constructed if and only if the following inequalities are satisfied:

$$p_a^\downarrow \leq t_a^{\leq 1} \quad \text{for } a \in \Sigma, \quad (1)$$

$$p_a^\parallel \leq t_a^{\geq 1} - p_a^\downarrow \quad \text{for } a \in \Sigma, \quad (2)$$

$$p_\star^\downarrow \leq \sum_{a \in \Sigma} (t_a^{\leq 1} - p_a^\downarrow - \min(p_a^\parallel - t_a^{\geq 2}, 0)), \quad (3)$$

$$p_\star^\parallel \leq \left[\sum_{a \in \Sigma} (t_a^{\geq 1} - p_a^\downarrow - p_a^\parallel) \right] - p_\star^\downarrow. \quad (4)$$

Naturally, these inequalities can be verified in polynomial time. \square

Finally, we observe that while in the reductions above we use different labels to represent elements of a finite enumerable set, the same can be accomplished with patterns using \star labels only, where natural numbers are encoded with simple gadgets. The gadgets use the fact that a node of a tree pattern that has k \downarrow -children can be mapped by a weakly-injective embedding only to a node having at least k nodes. On the other hand, we can easily modify reduction from Thm. 4.4 yield tree patterns without \star nodes.

Theorem 4.9. $\mathcal{M}_{\text{inj}}^\star$ and $\mathcal{M}_{\text{inj}}^\circ$ are NP-complete.

4.2 Ancestor-preserving embeddings

Theorem 4.10. \mathcal{M}_{anc} is NP-complete.

Proof. To prove NP-hardness we reduce SAT to \mathcal{M}_{anc} . We take a formula $\varphi = c_1 \wedge c_2 \wedge \dots \wedge c_k$ over variables x_1, \dots, x_n and for every variable x_i we construct two trees: $X_i = x_i(c_{j_1}, \dots, c_{j_m})$ such that c_{j_1}, \dots, c_{j_m} are exactly the clauses satisfied by using the literal x_i , and $\bar{X}_i = x_i(c_{j_1}, \dots, c_{j_m})$ such that c_{j_1}, \dots, c_{j_m} are exactly the clauses using the literal $\neg x_i$. The constructed tree is

$$t_\varphi = r(X_1, \bar{X}_1, \dots, X_n, \bar{X}_n).$$

And the tree pattern (written in XPath syntax) is

$$p_\varphi = r[x_1] \dots [x_n][\text{//}c_1] \dots [\text{//}c_k].$$

An example of the reduction for $\varphi = (x_1 \vee \neg x_3) \wedge (x_1 \vee \neg x_2 \vee x_3) \wedge (\neg x_1 \vee \neg x_2)$ is presented in Fig. 3. The main claim is that $(t_\varphi, p_\varphi) \in \mathcal{M}_{\text{anc}}$ iff $\varphi \in \text{SAT}$. We




Fig. 3. Reduction to \mathcal{M}_{anc} for $\varphi = (x_1 \vee \neg x_3) \wedge (x_1 \vee \neg x_2 \vee x_3) \wedge (\neg x_1 \vee \neg x_2)$.

prove it analogously to the main claim in the proof of Theorem 4.4. The use of ancestor-preserving embeddings ensures that the fragments $[x_i]$ and $[./c_j]$ are not mapped to the same subtree of t_φ , and this reduction does not work for weakly-injective embeddings. \square

We point out that in the proof above, the constructed pattern has height 1.

Corollary 4.11. $\mathcal{M}_{\text{anc}}^{\text{H} \leq k}$ is NP-complete for every $k \geq 1$.

Also, the use of child edges is not essential and they can be replaced by descendant edges and the reduction does not use \star labels.

Corollary 4.12. $\mathcal{M}_{\text{anc}}^{\parallel}$ and $\mathcal{M}_{\text{anc}}^{\circ}$ are NP-complete.

Bounding the degree of a node in the tree pattern renders, however, checking the existence of an ancestor-preserving embedding tractable.

Theorem 4.13. For any $k \geq 0$, $\mathcal{M}_{\text{anc}}^{\text{D} \leq k}$ is in P.

Proof. We fix a tree t and a tree pattern p . For a node $m \in N_p$ we define $\Phi(m) = \{n \in N_t \mid t|_n \leq_{\text{anc}} p|_m\}$, where $t|_n$ is a subtree of t rooted at n (and similarly, we define $p|_m$). Naturally, $t \leq_{\text{anc}} p$ iff $\text{root}_t \in \Phi(\text{root}_p)$.

We fix a node $m \in N_p$ with children m_1, \dots, m_k , suppose that we have computed $\Phi(m_i)$ for every $i \in \{1, \dots, k\}$, and take a node $n \in N_t$. We claim that n belongs to $\Phi(m)$ if and only if the following two conditions are satisfied: 1) $\text{lab}_t(n) = \text{lab}_p(m)$ unless $\text{lab}_p(m) = \star$, 2) there is $(n_1, \dots, n_k) \in \Phi(m_1) \times \dots \times \Phi(m_k)$ such that a) n_i is not an ancestor of n_j for all $i \neq j$, b) $(n, n_i) \in \text{child}_t$ if $(m, m_i) \in \text{child}_p$, and c) $(n, n_i) \in \text{child}_t^+$ if $(m, m_i) \in \text{desc}_p$.

Since k is bounded by a constant, the product $\Phi(m_1) \times \dots \times \Phi(m_k)$ is of size polynomial in the size of t , and therefore, the whole procedure works in polynomial time too. \square

Finally, gadgets similar to those in Thm. 4.9 allow us dispose of labels altogether.

Theorem 4.14. $\mathcal{M}_{\text{anc}}^{\star}$ is NP-complete.

4.3 LCA-preserving embeddings

Theorem 4.15. \mathcal{M}_{lca} is in P.

Proof. We fix a tree t and a tree pattern p . For a node $m \in N_p$ we define $\Phi(m) = \{n \in N_t \mid t|_n \leq_{\text{lca}} p|_m\}$, where $t|_n$ is a subtree of t rooted at n (and similarly, we define $p|_m$). Naturally, $t \leq_{\text{lca}} p$ if and only if $\text{root}_t \in \Phi(\text{root}_p)$. We present a bottom-up procedure for computing Φ .

We fix a node $m \in N_p$ with children m_1, \dots, m_k , suppose that we have computed $\Phi(m_i)$ for every $i \in \{1, \dots, k\}$, take a node $n \in N_t$, and let n_1, \dots, n_ℓ be its children. We claim that n belongs to $\Phi(m)$ if and only if the following

two conditions are satisfied: 1) $lab_t(n) = lab_p(m)$ unless $lab_p(m) = \star$ and 2) the bipartite graph $G = (X \cup Y, E)$ with $X = \{m_1, \dots, m_k\}$, $Y = \{n_1, \dots, n_\ell\}$, and

$$E = \{(m_i, n_j) \mid (m, m_i) \in child_p \wedge n_j \in \Phi(m_i) \vee (m, m_i) \in desc_p \wedge \exists n' \in \Phi(m_i). (n_j, n') \in child_t^*\},$$

has a matching of size k . In the construction of E we use the expression $(n_j, n') \in child_t^*$ because a \parallel -child m_i of m needs to be connected with proper descendants of n and these are descendants of n_j 's. We finish by pointing out that a maximum matching of G can be constructed in polynomial time [10]. \square

5 Related work

Model checking for tree patterns has been studied in the literature in a variety of variants depending on the requirements on the corresponding embeddings. They may, or may not, have to be injective, preserve various properties like the order among siblings, ancestor or child relationships, label equalities, etc. In this paper, we consider unordered, injective embeddings that additionally may be ancestor- or lca-preserving.

Kilpeläinen and Mannila [12] studied the *unordered tree inclusion* problem defined as follows. Given labeled trees P and T , can P be obtained from T by deleting nodes? Here, deleting a node u entails removing all edges incident to u and, if u has a parent v , replacing the edge from v to u by edges from v to the children of u . The unordered tree inclusion problem is equivalent to the model checking for ancestor-preserving embeddings where the tree pattern contains descendants edges only. [12] shows NP-completeness for tree patterns of height 1. Moreover, [14] shows that the problem remains NP-complete when all labels in both trees are \star or when degrees of all vertices except root are at most 3. These two results subsume our Thm. 4.10 and 4.14. [12, 14] show also the tractability of the problem when the degrees of all nodes in the tree pattern are bounded. Thm. 4.13 generalizes this to allow also for child edges in the tree patterns.

The tree inclusion problem is a special case of the *minor containment* problem for graphs [16, 17]: given two graphs G and H , decide whether G contains H as a minor, or equivalently, whether H can be obtained from a subgraph of G by edge contractions, where contracting an edge means replacing the edge and two incident vertices by a single new vertex. For trees, edge contraction is equivalent to node deletion. Since minor containment is known to be NP-complete, even for trees, this gives another proof of NP-completeness for ancestor-preserving embeddings.

Valiente [18] introduced the *constrained unordered tree inclusion* problem where the question is, given labeled trees P and T , whether P can be obtained from T by deleting nodes of degrees one or two. The polynomial time algorithm given there is based on the earlier results on subtree homeomorphisms [2] where unlabeled trees are considered. The constrained unordered tree inclusion

is equivalent to model checking of lca-preserving embeddings where all edges in the tree pattern are descendants. Our Thm. 4.15 slightly generalizes the above result allowing also for child edges in the pattern.

David [3] studied the complexity of ancestor-preserving embeddings of tree patterns with data comparison (equality and inequality) and showed their NP-completeness. Although we show that ancestor-preserving embeddings are NP-complete even without data comparisons, the reductions used in [3] construct tree patterns of a bounded degree, which shows that adding data comparisons indeed increases the computational complexity of the model checking problem.

Recently, Fan et al. [5] studied 1-1 p -homomorphisms which extend injective graph homomorphisms by relaxing the edge preservation condition. Namely, the edges have to be mapped to nonempty paths. However, neither the internal vertices nor edges within the paths have to be disjoint. In case of trees, 1-1 p -homomorphisms correspond to the weakly-injective embeddings that we consider in this paper. By reduction from *exact cover by 3-sets problem* they have shown NP-completeness of model checking in the case where the first graph is a tree and the second is a DAG. We improve over this result in Thm. 4.4 and 4.9.

When embeddings have to preserve the order among siblings, model checking becomes much easier. The *ordered tree inclusion problem* was initially introduced by Knuth [13, exercise 2.3.2-22] who gave a sufficient condition for testing inclusion. The polynomial time algorithms from [12] is based on dynamic programming and at each level may compute the inclusion greedily from left-to-right thanks to the order preservation requirement. The tree inclusion is also related to the *ordered tree pattern matching* [9], where embeddings have to preserve the order and child-relationship, but they do not have necessarily to preserve root.

6 Conclusions and future work

We have considered three different notions of injective embeddings of tree patterns and for each of them we have studied the problem of model checking. Table 1 summarizes the complexity results. All our results extend to embed-

	\leq_{std}	\leq_{inj}	\leq_{anc}	\leq_{lca}
unconstrained	P [6]	NP-c. (Thm. 4.4)	NP-c. (Thm. 4.10) [12, 14]	P (Thm. 4.15)
k -bounded degree $\frac{k \geq 2}{k = 1}$		NP-c. (Cor. 4.6)	P (Thm. 4.13)	
		P (Prop. 4.2)		
k -bounded height $\frac{k \geq 2}{k = 1}$		NP-c. (Thm. 4.9)	NP-c. (Thm. 4.10) [12]	
		P (Prop. 4.8)		
★ labels only		NP-c. (Thm. 4.9)	NP-c. (Thm. 4.14) [14]	
no ★ labels		NP-c. (Thm. 4.9)	NP-c. (Cor. 4.12) [12]	
no -edges		NP-c. (Cor. 4.5)	NP-c. (Cor. 4.12) [12, 14]	
no -edges		P (Prop. 4.3)		

Table 1. Summary of complexity results

dings between pairs of tree patterns, used for instance in static query analysis [15]. Although some of our results are subsumed by or can be easily obtained from existing results, our reductions and algorithms are simple and clean. In particular, we show intractability with direct reductions from SAT.

In the future, we would like to find out whether there is an algorithm for checking lca-preserving embeddings that does not rely on constructing perfect matchings in bipartite graphs. The exact bound on complexity of non-injective embeddings of tree patterns is a difficult open problem [7] and it would be interesting if establishing exact bounds on tractable cases of injective embeddings is any easier.

References

1. S. Amer-Yahia, S. Cho, L. V. S. Lakshmanan, and D. Srivastava. Tree pattern query minimization. *VLDB J.*, 11(4):315–331, 2002.
2. M.-J. Chung. $O(n^{2.55})$ time algorithms for the subgraph homeomorphism problem on trees. *J. Algorithms*, 8(1):106–112, 1987.
3. C. David. Complexity of data tree patterns over XML documents. In *MFCS*, pages 278–289, 2008.
4. W. Fan. Graph pattern matching revised for social network analysis. In *ICDT*, March 2012.
5. W. Fan, J. Li, S. Ma, H. Wang, and Y. Wu. Graph homomorphism revisited for graph matching. *Proc. VLDB Endow.*, 3(1-2):1161–1172, September 2010.
6. G. Gottlob, C. Koch, and R. Pichler. XPath processing in a nutshell. *SIGMOD Record*, 32(2):21–27, 2003.
7. M. Götz, C. Koch, and W. Martens. Efficient algorithms for descendant-only tree pattern queries. *Inf. Syst.*, 34(7):602–623, November 2009.
8. B. Habegger and D. Debarbieux. Integrating data from the web by machine-learning tree-pattern queries. In *OTM Conferences (1)*, pages 941–948, 2006.
9. C. M. Hoffmann and M. J. O’Donnell. Pattern matching in trees. *J. ACM*, 29(1):68–95, 1982.
10. J. E. Hopcroft and R. M. Karp. An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM J. Comput.*, 2(4):225–231, 1973.
11. P. Kilpeläinen and H. Mannila. Query primitives for tree-structured data. In *Combinatorial Pattern Matching (CPM)*, pages 213–225, 1994.
12. P. Kilpeläinen and H. Mannila. Ordered and unordered tree inclusion. *SIAM J. Comput.*, 24(2):340–356, 1995.
13. D. E. Knuth. *The Art of Computer Programming, Volume I: Fundamental Algorithms*. Addison-Wesley, 1968.
14. J. Matousek and R. Thomas. On the complexity of finding iso- and other morphisms for partial k -trees. *Discrete Mathematics*, 108(1-3):343–364, 1992.
15. G. Miklau and D. Suciu. Containment and equivalence for a fragment of XPath. *Journal of the ACM*, 51(1):2–45, 2004.
16. N. Robertson and P. D. Seymour. Graph minors. II. Algorithmic aspects of tree-width. *J. Algorithms*, 7(3):309–322, 1986.
17. N. Robertson and P. D. Seymour. Graph minors. XIII. The disjoint paths problem. *J. Comb. Theory, Ser. B*, 63(1):65–110, 1995.
18. G. Valiente. Constrained tree inclusion. *J. Discrete Algorithms*, 3(2-4):431–447, 2005.

Appendix: Omitted proofs

Proposition 3.1. *For any tree t and tree pattern p , 1) $t \preceq_{\text{lca}} p \Rightarrow t \preceq_{\text{anc}} p$, 2) $t \preceq_{\text{anc}} p \Rightarrow t \preceq_{\text{inj}} p$, and 3) $t \preceq_{\text{inj}} p \Rightarrow t \preceq_{\text{std}} p$.*

Proof. Assume that $t \preceq_{\text{lca}} p$ and let h be a lca-preserving embedding. Consider any n_1, n_2 such that $(h(n_1), h(n_2)) \in \text{child}_t^*$. Since h is lca-preserving and $\text{lca}(h(n_1), h(n_2)) = h(n_1)$, $\text{lca}(n_1, n_2) = n_1$ and therefore n_1 is an ancestor of n_2 . So h is ancestor-preserving.

For the proof of 2, consider p, t such that $t \preceq_{\text{anc}} p$ and let h be an ancestor-preserving embedding. We show that h is injective. Assume that there are n_1, n_2 such that $h(n_1) = h(n_2)$. Since h is ancestor-preserving, $(h(n_1), h(n_2)) \in \text{child}_t^*$, and $(h(n_2), h(n_1)) \in \text{child}_t^*$, $(n_1, n_2) \in (\text{child}_p \cup \text{desc}_p)^*$ and $(n_2, n_1) \in (\text{child}_p \cup \text{desc}_p)^*$, so $n_1 = n_2$. $t \preceq_{\text{inj}} p$.

Finally, Implication 3 follows from the fact that any injective embedding is an embedding. \square

Proposition 3.2. *For any tree t and any tree pattern p that does not use descendant edges, $t \preceq_{\text{inj}} p$ iff $t \preceq_{\text{anc}} p$ iff $t \preceq_{\text{lca}} p$.*

Proof. Assume that p does not use descendant edges. By Proposition 3.1, it is enough to prove that $t \preceq_{\text{inj}} p$ implies $t \preceq_{\text{lca}} p$.

Let $t \preceq_{\text{inj}} p$ and h be an embedding from p to t . It is easy to see that h is an isomorphism on a substructure of t , and therefore h is lca-preserving and $t \preceq_{\text{lca}} p$. Together with Proposition 3.1 it implies all the equivalences. \square

Proposition 3.3. *For any tree t and any path pattern p , $t \preceq_{\text{std}} p$ iff $t \preceq_{\text{inj}} p$ iff $t \preceq_{\text{anc}} p$ iff $t \preceq_{\text{lca}} p$.*

Proof. Assume that p is a path pattern and $t \preceq_{\text{std}} p$ and let h be an embedding from p to t . Consider any nodes n, m of p such that there is a path from n to m . Clearly, $\text{lca}(n, m) = n$. By Properties 2 and 3 of the definition of embeddings, there is also a path from $h(n)$ to $h(m)$, hence $\text{lca}(h(n), h(m)) = h(n)$. Therefore h is lca-preserving and $t \preceq_{\text{lca}} p$. By Proposition 3.1 we obtain the required equivalence. \square

Theorem 4.7. $\mathcal{M}_{\text{inj}}^{\text{H} \leq k}$ is NP-complete for any $k \geq 2$.

Proof. We show how to build, for a given instance φ of SAT problem, a pattern p_φ and a tree t_φ such that $t_\varphi \preceq_{\text{inj}} p_\varphi$ if and only if φ is satisfiable. Let $\varphi = c_1 \wedge c_2 \wedge \dots \wedge c_k$ be an instance of SAT over variables x_1, \dots, x_n . We set $\Sigma = \{a, c_1, \dots, c_k, s_1, \dots, s_n\}$.

For each i , we define the tree X_i as follows. Its root is a node x_i^p and it is connected to $k+1$ nodes, namely $x_i^n, p_i^1, p_i^2, \dots, p_i^k$. Node x_i^n has $k+2$ successors, namely $s_i, n_i^1, \dots, n_i^{k+1}$. All other nodes have no successors.

The tree t_φ consists of a root r and its n disjoint successors — X_1, \dots, X_n (see Fig. 4).

Now we define the labeling of t_φ . Let c_{i_1}, \dots, c_{i_l} be the clauses with the positive occurrence of x_i , and $c_{j_1}, \dots, c_{j_{l'}}$ be the clauses with the negative occurrence

of x_i . For all $s \leq l$, we label p_s with c_{i_s} . Similarly, for all $s \leq l'$ we label n_s by c_{j_s} . We label s_i by s_i and all other nodes by a .

The pattern p_φ is as presented at Fig. 4. Clearly, its depth is bounded by 2.

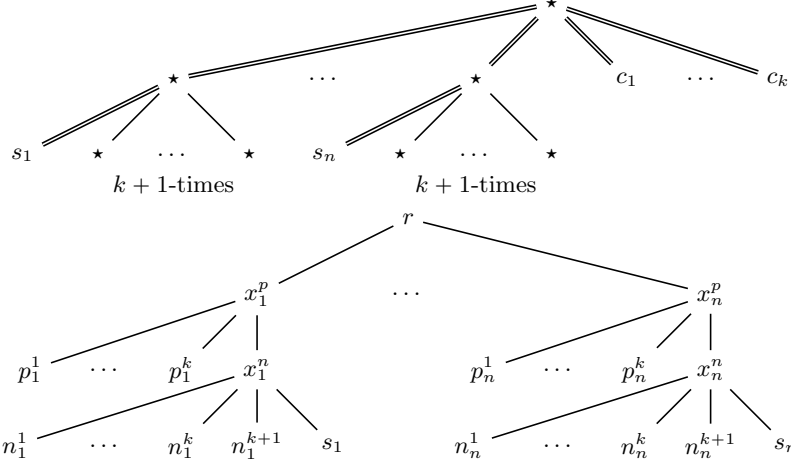


Fig. 4. The pattern p_φ (at the top) and the tree t_φ (without the labeling).

Assume that $t_\varphi \leq_{\text{inj}} p_\varphi$ and let h be the corresponding embedding. Let Y be the set of all the successors of root_{p_φ} labeled by \star in p_φ and $h(Y)$ be the image of Y . A quick check shows that for each i there is exactly one node from $\{x_i^p, x_i^n\}$ in $h(Y)$. We define the valuation for φ such that x_i is positive if $x_i^n \in h(Y)$ and negative otherwise.

Consider any clause c_s and let m be the node in p_φ labeled by c_s . Assume that $h(m) = p_i^j$ for some i, j . It means that x_i occurs positively in c_s and that x_i^p does not belong to $h(Y)$ — otherwise, if $x_i^p = h(m')$ for some m' , then all successors of x_i^p would be results of h applied to successors of m' , contradicting the facts that $h(m) = p_i^j$ and h is injective. Therefore, c_s is satisfied.

The proof that if φ is satisfiable then $t_\varphi \leq_{\text{inj}} p_\varphi$ should now be straightforward. \square

Theorem 4.9. $\mathcal{M}_{\text{inj}}^\star$ and $\mathcal{M}_{\text{inj}}^\circ$ are NP-complete.

Proof. For the $\mathcal{M}_{\text{inj}}^\star$ case, we simply adjust the proof of Theorem 4.7, taking the advantage of the fact that all the nodes with labels different than \star are in leaves.

For each $s \in \mathbb{N}$, we define tree T_s^k that consists of two nodes with $k+3$ successors and a path connecting them of length s (see Fig. 5). Note that in the original t_φ does not contain any node of degree $\geq k+3$.

We replace all nodes (in t_φ and p_φ) labeled by c_s by T_s^k and all nodes labeled by n_s by T_{k+s}^k . Then, in t_φ , we replace all labels by a . It is readily checked that for any $i \neq j$ there is no embedding from T_i^k to T_j^k , and since t_φ contains no




Fig. 5. The tree T_s^k .

nodes with degree at least $k+3$, there is an embedding from the modified pattern to the modified tree if and only if φ is satisfiable.

For the \mathcal{M}_{inj}° case, we simply replace all \star in the pattern defined above by a , the only label present in the tree. \square

Theorem 4.14. \mathcal{M}_{anc}^\star is NP-complete.

Proof. We modify the proof of Theorem 4.10. First, we adjust the tree and the pattern by adding one node below each x_i , label it by x_i and label old x_i by a . We also replace r by a and all a in the pattern by \star (see Fig. 6).




Fig. 6. Adjusted reduction to \mathcal{M}_{anc} for $\varphi = (x_1 \vee \neg x_3) \wedge (x_1 \vee \neg x_2 \vee x_3) \wedge (\neg x_1 \vee \neg x_2)$.

The obtained tree and pattern have the following property: the only nodes that are not labeled by \star or a are leaves. By virtually the same way as in Theorem 4.9 we can replace them by trees T_s . \square