



HAL
open science

Fuzzy Tree Mining: Go Soft on Your Nodes

Federico del Razo Lopez, Anne Laurent, Pascal Poncelet, Maguelonne Teisseire

► **To cite this version:**

Federico del Razo Lopez, Anne Laurent, Pascal Poncelet, Maguelonne Teisseire. Fuzzy Tree Mining: Go Soft on Your Nodes. IFSA: International Fuzzy Systems Association World Congress, Jun 2007, Cancun, Mexico. pp.145-154, 10.1007/978-3-540-72950-1_15 . lirmm-00272118

HAL Id: lirmm-00272118

<https://hal-lirmm.ccsd.cnrs.fr/lirmm-00272118>

Submitted on 20 Mar 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fuzzy Tree Mining: Go soft on your nodes

Federico Del Razo Lopez¹, Anne Laurent¹, Pascal Poncelet², and Maguelonne Teisseire¹

¹ LIRMM-CNRS UMR5506, Université Montpellier 2, 161 rue Ada, 34392 Montpellier, France {delrazo, laurent, teisseire}@lirmm.fr

² LIGI2P - EMA pascal.poncelet@ema.fr

Abstract. Tree mining consists in discovering the frequent subtrees from a forest of trees. This problem has many application areas. For instance, a huge volume of data available from the Internet is now described by trees (e.g. XML). Still, for several documents dealing with the same topic, this description is not always the same. It is thus necessary to mine a common structure in order to query these documents. Biology is another field where data may be described by means of trees. The problem of mining trees has now been addressed for several years, leading to well-known algorithms. However, these algorithms can hardly deal with real data in a soft manner. Indeed, they consider a subtree as *fully included* in the super-tree. This means that all the nodes must appear. In this paper, we extend this definition to fuzzy inclusion based on the idea that a tree is included to a certain degree within another one, this fuzzy degree being correlated to the number of *matching nodes*.

1 Introduction

Tree mining is a subfield of data mining aiming at discovering automatically all the subtrees that appear frequently in a database of trees. This research area has several applications, including the discovery of mediator schemas. The background in this research is mainly constituted by the work by Asai et al. and Zaki et al.[1,9,12,16,17]. This work addresses the problem of tree mining considering several ways to define when a tree S is included within another one T . Inclusion is then decided depending on the way ancestry and brotherhood are considered. In this respect, the authors distinguish between approaches where (i) either all the pair of connex nodes in the tree S must be found in T with no intermediate node, (ii) or some intermediate nodes are accepted. Figure 1 illustrates this difference.

The work from the litterature is then twofolded, considering both:

- the representation of the trees,
- the extraction of frequent subtrees.

It should be noted that designing efficient algorithms to tackle the problem of extracting frequent subtrees is highly correlated to the representation of the trees, as this representation may help scanning the trees. The process of extracting frequent subtrees is based on the Apriori process, which is a recursive




Fig. 1. Traditional tree inclusion

process. It can be divided into the two following steps: for each size of trees (i) generation of candidates and (ii) validation of candidates. A candidate is a tree that is considered as being potentially frequent. The candidates of size k are built based on the frequent subtrees of size $k - 1$. This family of methods is well-known and has been applied for tree mining. However, all the existing methods consider that a tree **is** or **is not** included within another one, which is too restrictive to be efficient and relevant. We first propose the concept of fuzzy tree mining, which has been introduced in [10]. This concept has been detailed in [11], where we have defined fuzzy ancestor-descendant relation (fuzzy vertical path). In this paper, we consider another way for softening the tree inclusion definition by considering that some nodes may be discarded (partial inclusion). In classical approaches, *all* the nodes of a subtree S must be included in a tree T if S is included in T . For instance, Figure 2 shows a tree S that will not be considered as being included in T . However, we argue that this is too restrictive when mining data from the real world when imperfections are often present. For instance, in Figure 2 S has 75% of its nodes included in T .




Fig. 2. Partial Inclusion

The challenging part of our work is that we want to remain efficient, in the framework of fuzzy data mining. The paper is organized as follows: Section refsec:back recalls the related work on tree mining and our previous work on dealing with fuzzy tree mining. Section 3 introduces the necessary definitions for dealing with partial inclusion. Section 4 introduces the algorithms we design for

extracting frequent subtrees from a tree database in a soft manner by considering partial inclusion. Finally, Section 5 concludes this work and presents our future working directions.

2 Background

In this section, we recall from the literature and from previous work the basic definitions of tree mining and the ways trees can be represented.

2.1 Tree Mining

A *tree* is a connected graph containing no cycle. A tree is composed by nodes, which are linked by edges such that there exists a particular node called *root* and such that all the nodes but the root are composed by sub-trees. A tree is said to be an *ordered tree* if the children from a node are ordered. A tree is said to be an *unordered tree* otherwise.

Let $\mathcal{L} = \{a, b, c, \dots\}$ be a set of labels. A *labeled ordered tree* is a tree $T = (r, N, B, L, \preceq)$ where: r is the root, N is the set of nodes, B is the set of edges such that $B \subseteq V^2$, $(L : N \rightarrow \mathcal{L})$ is a mapping from the set of labels \mathcal{L} to the set of nodes N , and \preceq is an ordered relation between brother nodes.

Tree Mining refers to the process of extracting all the subtrees that appear frequently in a database D of trees. The frequency is computed using the notion of support: Given a database D , the *support* of a tree S is the proportion of trees from the database where S is embedded:

$$Support(S) = \frac{\# \text{ of trees where } S \text{ is embedded}}{\# \text{ of trees in } D}$$

S is said to be frequent if $Support(S) \geq \sigma$ where σ is a user-defined minimal support threshold.

Several kinds of tree inclusion can be defined [2], depending on the way ancestors and siblings are considered. For instance, [17] defines the inclusion as follows:

Definition 1 *A tree S is embedded into a tree T if there exists an injective and total function $\phi : N_S \rightarrow N_T$ such as for all $n, m \in N_S$:*

- ϕ keeps the labels: $L_S(n) = L_T(\phi(n))$;
- ϕ keeps the relations ancestor-descendant: $(n, m) \iff (\phi(n), \phi(m))$;
- ϕ keeps order relations: $(n \preceq_S m) \iff (\phi(n) \preceq_T \phi(m))$.

As highlighted in [10], fuzzy data mining can help when mining frequent subtrees from a tree database. Four ways to soften classical approaches has been proposed:

- *ancestor-descendant degree*: in classical approaches, a node *is* or *is not* an ancestor of another one. In our approach, we propose to indicate by a degree between 0 and 1 to which extent a node is an ancestor of another one, meaning that if there are too many nodes between them, then this degree will decrease
- *sibling ordering degree*: in classical approaches, nodes *are* or *are not* searched in the initial order. In our approach, we propose to indicate by a degree the sibling disorder.
- *partial inclusion*: in classical approaches, all the nodes from the candidate must be in the tree. In our approach, we propose to soften this rule by considering the degree to which the nodes are embedded in the tree.
- *Node similarity*: in classical methods, a node label *is* or *is not* the same as another one. In our approach, we propose to soften this by indicating by a degree to which extent two nodes are similar (*e.g.* based on a taxonomy).

The *ancestor-descendant degree* has been studied in [11]. In the rest of this paper, we focus on the partial inclusion.

Algorithms Several algorithms have been designed to address the problem of tree mining: *TreeMiner* in [17], *FreqT* in [1], *Chopper* [13], *FreeTreeMiner* [4] and *CMTreeMiner* [3]. All of them is based on a levelwise process consisting of the following two recursive steps: generation of candidates and validation of candidates. This process is recursive as it starts from the candidates that contain only one node, to discover the frequent 1-node subtrees, which are used to build the 2-node candidates, and so on.

These two steps have been studied. The generation of candidates is either based on methods that build trees containing n nodes by considering one tree containing $n-1$ nodes and adding another node, or is based on methods that mix two subtrees containing n nodes and sharing $n-1$ nodes to build a new candidate subtree containing $n+1$ nodes.

The validation aims at checking whether a tree is embedded within another one. Several approaches have been proposed. In our previous work, we have defined some algorithms that are based on the idea of *anchoring*: we try to anchor the root of the subtree until we find a node that matches. Then the following nodes are tested until (i) it is no more possible to find some remaining nodes for matching, or (ii) an incompatibility has been detected or (iii) the subtree fully matches.

2.2 Tree Representation

Several ways of representing trees have been proposed to support the algorithms cited above. The representation impacts the two steps discussed above (generation and validation of candidates). However, it may be the case that the representation is too rich and requires too much memory. We have thus proposed in previous work a low-memory representation of trees: RSF. This representation is defined below.

other node addresses are computed by concatenating the father address with: 1 if it is a child (left-hand path) and 0 otherwise (right-hand path), as shown on Fig. 2.2 c).




Fig. 4. Example of a Binary Tree Transformation and Node Addressing

3 FTMnodes: Definitions

In this paper, we formally extend the definition of tree inclusion to partial inclusion based on the number of nodes that are matched. Partial inclusion is defined as follows:

Definition 2 A tree S is partially embedded into a tree T with a degree $\delta(S, T)$ if there exists an injective and total function $\phi : N_S \rightarrow N_T \cup \perp$ such as for all $n, m \in N$:

- ϕ keeps the labels: $L_S(n) = L_T(\phi(n))$ or $\phi(n) = \perp$;
- ϕ keeps the relations ancestor-descendant: $(n, m) \iff (\phi(n), \phi(m))$ or $(\phi(n), \phi(m)) = \perp$;
- ϕ keeps the order relations:
 $(n \preceq_S m) \iff (\phi(n) \preceq_T \phi(m))$ or $(\phi(n) \preceq_T \phi(m)) = \perp$;
- $\delta(S, T) = \frac{|\{s \in S : \phi(s) \neq \perp\}|}{\# \text{ of nodes in } T}$.

From this definition, it is possible to define the support of a subtree, as follows:

Definition 3 Given a database D and a tree S , the support of S in D is given by:

$$\text{Support}(S) = \text{Agg}_{T \in D}(\delta(S, T))$$

where Agg is a function of aggregation.

For instance, we may use Ordered Weighted Aggregators (also known as OWA) [15]. An OWA operator of dimension n is a mapping

$$F : R^n \rightarrow R$$

that has an associated n vector $W = (w_1, w_2, \dots, w_n)^T$ such that $w_i \in [0, 1]$ and $\sum_{i=1}^n w_i = 1$. We have $F(a_1, a_2, \dots, a_n) = \sum_{j=1}^n w_j \cdot b_j$ where b_j is the j^{th} largest value of the a_i .

For instance, the average may be applied:

$$Support(S) = \frac{\sum_{T \in D} d(S, T)}{\# \text{ of trees in } D}$$

In fact, we consider a thresholded Σ -count so that

- a tree cannot be considered as being embedded within another one if the number of embedded nodes is too low,
- the degree to which a tree is embedded within another one is taken into account.

We thus have:

Definition 4 *Given a database D , a threshold τ and a tree S , the support of S in D is given by:*

$$Support(S) = \sum_{T \in D} (\alpha_\tau(\delta(S, T)))$$

where

$$\alpha_\tau(x) = \begin{cases} 0 & \text{if } x > \tau \\ x & \text{otherwise} \end{cases}$$

4 FTMnodes: Algorithms

Note that in the classical case, mining totally included trees allows to cut in the database scan since whenever a node cannot be matched, there is no need looking for the other ones. In our approach, outliers are accepted, which may be considered as a drawback considering scalability. However, it is still possible to cut off the search when the proportion has been overpassed.

As defined previously, we consider that a tree cannot be considered as being embedded within another one if the number of matching nodes is not greater than a user-defined threshold τ . This definition not only guarantees the quality of the research from a semantic point of view, but it also guarantees the scalability of our approach. Indeed, it is then possible to draw the property of anti-monotonicity which is the basis of levelwise algorithms. We have the following properties:

Considering that the first n nodes of tree S matched to nodes from T , and that $\pi\%$ of the nodes of S have been matched, then the first $n + 1$ nodes of S cannot be embedded in T to a proportion greater than π .

This property comes from the fact that if it has not been possible to match ν nodes among the first n nodes of S , then the number of nodes being not matched when going ahead in the process to the first $n + 1$ nodes will either be equal or will be greater (equal to $\nu + 1$).

As a consequence, whenever the threshold τ is overpassed, the process can be stopped for this path as it will never be considered in the thresholded \sum -count.

Note that it may be the case that a subtree is included within another one in different manners, as illustrated by Figure 5. In this case, the best degree of inclusion will be considered and it is found by maintaining all the possible ways of inclusion until all the solutions have been considered as studied in [6] when considering fuzzy sequential patterns.




Fig. 5. Several ways of including S in $\{T_1, T_2, T_3\}$ with $\tau = 0.75$ (at least 3 nodes out of 4)

The following process is thus considered in our approach (algorithm 1 and 2):

- anchoring
- for each possible anchor, for each node n in S to be matched
 - scan the nodes of T until n is matched, start another way to find the other possible matches,
 - if no match is possible then go to the next node in n and increment the number of mismatched nodes
 - if the number of mismatched nodes is greater than the threshold τ or if T has been fully scanned, then discard this anchor
- compute the best inclusion from non discarded anchoring paths

Note that our approach is consistent, meaning that if $\tau = 1$ (i.e. all the nodes must be mapped), then our algorithms are exactly the same as the ones defined in the crisp case [5].

5 Conclusion

In this paper, we have detailed our previous work on fuzzy tree mining by giving the necessary definitions and algorithms in order to address the partial inclusion. Partial inclusion is a big deal in tree mining as it is not possible to consider full matches in real applications. However, it is necessary to remain scalable as the volumes of data being considered in real databases is huge. We thus design solutions based on levelwise algorithms, which consider anti-monotonic properties that guarantee the scalability. The algorithms presented here are currently

```

Data:  $S$  //subtree to validate,
          $T$  //tree from database
Result:  $true$  //if  $S$  is embedded within  $T$ 

 $\mathcal{M}$  // mapping set of  $S$  within  $T$ ;
foreach  $node\ m \in N_T$  do
   $n \leftarrow root(S)$ ;
  if  $L(n) = L(m)$  then
     $\mathcal{M} \leftarrow \mathcal{M} \cup M$ ;
return the best inclusion  $\{M \in \mathcal{M} | MIN\{M.mismatchedNodes\}\}$ ;

```

Algorithm 1: ANCHORING

```

Data:  $S$  //subtree to validate,  $T$  //tree from database,
          $n, m$  //anchoring point,  $n \in S$  to  $m \in T$ ,
          $M$  //occurrence of  $S$ 

 $M[n] \leftarrow m$ ;
 $n \leftarrow n + 1$ ;
if  $n \leq |S|$  then
   $P \leftarrow \{w: w \in T \text{ such as } ancestor(w) = M[ancestor(n)] \text{ and } L(w) = L(n)\}$ ;
  if  $P = \emptyset$  then
    foreach  $node\ w \in P$  do
       $\mathcal{M} \leftarrow \mathcal{M} \cup M$ ;
  else
     $M.mismatchedNodes \leftarrow M.mismatchedNodes + 1$ ;
    if  $M.mismatchedNodes \geq \tau$  then
      exit;
else
  exit;

```

Algorithm 2: PARTIALINCLUSIONDEGREE

implemented, and it is possible to conclude that this approach allows the extraction of more frequent subtrees (as fuzziness is introduced) while remaining scalable. Future work includes the comparison of the results depending on the choices of the aggregation function. This comparison will be lead both on the quality of frequent subtrees and on runtime, as some aggregation functions are easier to compute than other ones.

References

1. T. Asai, K. Abe, S. Kawasoe, H. Arimura, and H. Sakamoto. Efficient substructure discovery from large semi-structure data. In *2nd Annual SIAM Symposium on Data Mining, SDM2002*, Arlington, VA, USA, 2002. Springer-Verlag.
2. Y. Chi, R. R. Muntz, S. Nijssen, and J. N. Kok. Frequent subtree mining - an overview. *Fundamenta Informaticae XXI*, pages 1001–1038, 2005.
3. Y. Chi, Y. Yang, and R. Muntz. Cmtreefinder: Mining both closed and maximal frequent subtrees. In *The Eighth Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD'04)*, 2004.
4. Y. Chi, Y. Yang, and R. R. Muntz. Indexing and mining free trees. In *International Conference on Data Mining 2003 (ICDM2003)*, 2003.
5. F. Del Razo, A. Laurent, P. Poncelet, and M. Teisseire. Rsf - a new tree mining approach with an efficient data structure. In *Proceedings of the joint Conference: 4th Conference of the European Society for Fuzzy Logic and Technology (EUSFLAT 2005)*, pages 1088–1093, 2005.
6. C. Fiot, A. Laurent, and M. Teisseire. From crispness to fuzziness: Three algorithms for soft sequential pattern mining. *IEEE Transactions on Fuzzy Systems*. (to appear), 2007.
7. D. Huffman. A method for the construction of minimum-redundancy codes. In *Proceedings of the Institute of Radio Engineers*, 1952.
8. D. Knuth. *The Art of Computer Programming, Volume 1: Fundamental Algorithms*. Addison-Wesley, 1973.
9. M. Kuramochi and G. Karypis. Frequent subgraph discovery. In *IEEE International Conference on Data Mining (ICDM)*, 2001.
10. A. Laurent, M. Teisseire, and P. Poncelet. *Fuzzy Data Mining for the Semantic Web: Building XML Mediator Schemas*, chapter 12. Elsevier, E. Sanchez(ed.), To appear, 2006.
11. S. Sanchez, A. Laurent, P. Poncelet, and M. Teisseire. Fuzbt: a binary approach for fuzzy tree mining. In *Proceedings of the 11th IPMU International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems (IPMU 2006)*, 2006.
12. A. Termier, M.-C. Rousset, and M. Sebag. Treefinder, a first step towards XML data mining. In *IEEE Conference on Data Mining (ICDM)*, pages 450–457, 2002.
13. C. Wang, Q. Yuan, H. Zhou, W. Wang, and B. Shi. Chopper: An efficient algorithm for tree mining. *Journal of Computer Science and Technology*, 19:309–319, May 2004.
14. M. A. Weiss. *Data Structures And Algorithm Analysis In C*. Addison Wesley, 1998.
15. R. Yager. Families of owa operators. *Fuzzy Sets and Systems*, 57(3):125 – 148, 1993.
16. X. Yan and J. Han. gspan: Graph-based substructure pattern mining. In *IEEE Conference on Data Mining (ICDM)*, 2002.

17. M. J. Zaki. *Efficiently Mining Frequent Trees in a Forest*. In *In KDD'02*, Edmonton, Alberta, Canada, 2002. ACM.