

Text Analysis with Enhanced Annotated Suffix Trees Algorithms and Implementation

Mikhail Dubov¹

National Research University Higher School of Economics
Computer Science faculty, Moscow, Russia

AIST'2015

Table of Contents

- 1 Annotated suffix trees
 - Suffix trees
 - Annotated suffix trees
 - AST relevance score
- 2 Algorithms
 - From suffix tries to suffix trees
 - From suffix trees to suffix arrays
- 3 Implementation
 - Package EAST
 - Synonym extraction
- 4 LM Monitor
 - Concept & architecture
 - Keyphrase reference graphs

Annotated suffix trees

- **Letter-based** method for text analysis
- **Annotated suffix trees:** full-text index
- **Basic computation:** relevance score of a keyphrase to the text collection indexed by AST
- **Range of applications:**
 - Text classification (e.g. spam filtering)
 - Feature extraction
 - Keyphrase analysis (stay tuned)

Suffix trees




Figure: Suffix tree for string “XABXAC”

- **Suffix tree** for a string S ($|S| = n$) is a rooted directed tree encoding all the suffixes of that string [3]
- The concatenation of edge labels on every path from the root node to one of the leaves makes up one of the suffixes of that string, i.e. $S[i \dots n]$.
- It is also required that each internal node has two or more children, and each edge is labeled with a non-empty substring of S .

Suffix trees

- Various $O(n)$ construction algorithms exist (Ukkonen, Weiner)
- Establishes a linear-time solution for the **exact pattern matching** problem
- Suffix tree is a **full-text index**

Annotated suffix trees




Figure: Annotated suffix tree for string "XABXAC"

- Extension: node labels
- Node label $f(v)$ indicates the number of entries of the substring on the path from root to v in the text collection

AST relevance score




Figure: Naive AST representation (as a trie) for a collection of 3 strings

- *Conditional probability of a node given its parent:*

$$\hat{p}(v) = \frac{f(v)}{f(\text{parent}(v))}$$

AST relevance score

Relevance score computation for keyword S in text collection T (described in terms of a **trie**, not a **tree**):

- For each suffix $S[i \dots n]$ of S , try to match it against the suffix tree $AST(T)$, starting at the root.
- If, for suffix s , we matched exactly k symbols in the tree, then


$$score_{suff}(s) = \frac{\sum_{i=1}^k \hat{p}(v_i)}{k},$$

where v_i is the i -th node on the matching path starting at the root (if $k = 0$, then $score_{suff}(s) = 0$).

- The final score for keyword S is obtained as


$$SCORE(S) = \frac{\sum_{i=1}^{|S|} score_{suff}(S[i :])}{|S|}.$$

AST relevance score: Example 1

 $T = [\text{"XAB"}, \text{"XAC"}, \text{"CAB"}]$

$$\begin{aligned} \text{SCORE}(\text{"ABC"}) &= \frac{\text{score}(\text{"ABC"}) + \text{score}(\text{"BC"}) + \text{score}(\text{"C"})}{3} = \\ &= \frac{(0.33 + 0.67)/2 + (0.22)/1 + (0.22)/1}{3} = 0.31 \end{aligned}$$

AST relevance score: Example 2

 $T = ["XAB", "XAC", "CAB"]$

$$\begin{aligned} \text{SCORE}("XYZ") &= \frac{\text{score}("XYZ") + \text{score}("YZ") + \text{score}("Z")}{3} = \\ &= \frac{(0.22)/1 + 0 + 0}{3} = 0.07 \end{aligned}$$

AST relevance score: Example 3



$$\text{SCORE}(\text{"Alice"}) = 0.32$$

$$\text{SCORE}(\text{"Bob"}) = 0.04$$

(Usually, $\text{SCORE} > 0.2$ is a strong evidence of relevance)

AST relevance score: alternatives & summary

- **Alternative solution:** count the number of occurrences of a keyword in the text collection
 - **Word-based** approach
 - **Requires at least normalization**, NLP involved
 - Can also use the **Levenstein distance** for more sensitivity
 - Relevance score definition & interpretation is not obvious
- **AST Relevance score:**
 - **Letter-based**, “fuzzy” approach
 - **Language-independent**, no NLP involved
 - **Interpretation:** *average conditional probability* of an occurrence of a single symbol of the input key phrase in the text collection

Table of Contents

- 1 Annotated suffix trees
 - Suffix trees
 - Annotated suffix trees
 - AST relevance score
- 2 Algorithms
 - From suffix tries to suffix trees
 - From suffix trees to suffix arrays
- 3 Implementation
 - Package EAST
 - Synonym extraction
- 4 LM Monitor
 - Concept & architecture
 - Keyphrase reference graphs

Enhanced annotated suffix trees

- In the original papers, AST was represented as a **trie** $\implies O(n^2)$ time & space complexity.
- Even when implemented properly with suffix trees, the AST construction time & space usage still has a large **hidden constant** behind $O(n)$.
- We propose an **enhanced** implementation that uses **suffix arrays**.

From suffix tries to suffix trees

Ensure the linearity of our data structure:

- Suffix trie: one node per letter, $O(n^2)$ time & space
- Suffix tree: compacted edges, no chains, $O(n)$ time & space




Figure: Suffix trie




Figure: Suffix tree

From suffix tries to suffix trees

To construct *annotated suffix trees* in $O(n)$, simple preprocessing is needed:

Algorithm **LinearASTConstruction(C)**

Input. String collection $C = \{S_1, \dots, S_m\}$

Output. Generalized annotated suffix tree for C .

- ❶ Construct $C' = \{S_1\$1, \dots, S_m\$m\}$, where $\$i$ are unique characters that do not appear in $S_1 \dots S_m$.
- ❷ Construct a generalized suffix tree T for collection C' using a linear-time algorithm (e.g. the *Ukkonen algorithm*).
- ❸ **for** l **in** $leaves(T)$
- ❹ **do** set $f(l) \leftarrow 1$
- ❺ Run a postfix depth-first tree traversal on the suffix tree T .
For each inner node v , set $f(v) \leftarrow \sum_{u \in children(v)} f(u)$.

From suffix tries to suffix trees

One minor change in the suffix relevance score:




Figure: Suffix trie

$$score_{suff}(s) = \frac{\sum_{i=1}^k \hat{p}(v_i)}{k}$$




Figure: Suffix tree

If l is the number of symbols in the match, then

$$score_{suff}(s) = \frac{\sum_{i=1}^k \hat{p}(v_i) + l - k}{l}$$

From suffix trees to suffix arrays

- **Suffix array** for a string S ($|S| = n$) is an array of n integer numbers, enumerating the n suffixes of S in lexicographic order.

Table: Suffix array for string “XABXAC”
(the suffixes are not actually stored)

i	<i>suffix array</i>	$S[\text{suffix}[i]:]$
0	2	ABXAC
1	5	AC
2	3	BXAC
3	6	C
4	1	XABXAC
5	4	XAC

- Suffix arrays are more space efficient than suffix trees.

Enhanced suffix arrays

- Abouelhoda, Kurtz, & Ohlebusch [1] have shown that it is possible to systematically replace every algorithm that uses suffix trees with another one based on suffix arrays.
- Need to enhance the suffix array with two auxiliary arrays:
 - *lcp*-table for bottom-up traversal
 - *child*-table for top-down traversal
- Can be implemented to take no more than 10 bytes per input symbol (at least 20 for suffix trees)

Enhanced suffix arrays

Table: Enhanced suffix array for string "XABXAC"

<i>i</i>	<i>suffix array</i>	<i>lcp-table</i>	<i>child-table</i>			<i>S[suff[i]:]</i>
			1.	2.	3.	
0	1	0		1	2	ABXAC
1	4	1				AC
2	2	0	1		3	BXAC
3	5	0			4	C
4	0	0				XABXAC
5	3	2				XAC

Enhanced annotated suffix arrays

- We need to store annotations for suffix tree nodes
- The number of nodes in a suffix tree cannot exceed $(2n - 1)$
- After preprocessing, all the leaves will be annotated with 1, so there is no need to store these annotations explicitly
- We are left with at most $(n - 1)$ numbers to store \implies can introduce one more auxiliary array of length n (*annotation-table*)

Enhanced annotated suffix arrays

Table: Enhanced annotated suffix array for string "XABXAC"

<i>i</i>	<i>suffix array</i>	<i>lcp-table</i>	<i>child-table</i>			<i>annotation</i>	<i>S[suff[i]:]</i>
			1.	2.	3.		
0	1	0		1	2	6	ABXAC
1	4	1				2	AC
2	2	0	1		3		BXAC
3	5	0			4		C
4	0	0					XABXAC
5	3	2				2	XAC

Enhanced annotated suffix arrays




Figure: Annotated suffix tree for string "XABXAC"

Enhanced annotated suffix arrays

Node-to-array mapping – via **virtual lcp-trees**:

- Can be restored from the *lcp*-table
- Nodes correspond to the inner nodes of the suffix tree
- Nodes are represented as $\langle l, i, j \rangle$: the *lcp-value* l and the left and right boundaries of the *lcp-interval* (i, j)
- For each *lcp-interval* $v = \langle l, i, j \rangle$ there exists a unique index, $\text{index}(v) \in [0; n - 1]$, which is equal to the smallest k , such that $k > i$ and $\text{lcp}[k] = l$. It is this mapping that we use to store the inner node frequency annotations.

Enhanced annotated suffix arrays

Algorithm **LinearEASACONSTRUCTION(C)**

Input. String collection $C = \{S_1, \dots, S_m\}$

Output. Enhanced suffix array for C with substring frequency annotations.

- ① Construct a string $S = S_1\$_1 + \dots + S_m\$_m$, where $\$_i$ are unique termination symbols.
- ② Construct a suffix array A for string S using a linear-time algorithm (e.g. the *Kärkkäinen-Sanders algorithm*) and two auxiliary arrays: *lcp-array* and *child-array*.
- ③ Simulate a postfix depth-first tree traversal on the suffix array A . At each of the *virtual inner nodes*, corresponding to an *lcp-interval* $v = \langle l, i, j \rangle$, where $i < j$, set
$$\text{annotation}[\text{index}(v)] = \sum_{u \in \text{children}(v)} \text{annotation}[\text{index}(u)] + \#(\langle l, i, j \rangle : i = j).$$

Enhanced annotated suffix arrays: Experimental results




Figure: Experimental results

- Implementation: Python 2.7
- 10x less memory – due to suffix arrays + the *Numpy* library

Table of Contents

- 1 Annotated suffix trees
 - Suffix trees
 - Annotated suffix trees
 - AST relevance score
- 2 Algorithms
 - From suffix tries to suffix trees
 - From suffix trees to suffix arrays
- 3 **Implementation**
 - Package EAST
 - Synonym extraction
- 4 LM Monitor
 - Concept & architecture
 - Keyphrase reference graphs

Package EAST

- EAST = “*Enhanced Annotated Suffix Trees*”
- Open-source:
`https://github.com/msdubov/AST-text-analysis`
- Registered in Python Package Index and is easy to install:
`$ pip install EAST`

Package EAST

- Provides command-line user interface:

```
$ east keyphrases table <keyphrases_list.txt>  
  <path/to/the/text/collection/>
```

- Can be used as a Python library:

```
>>> from east.ast.base import AST  
>>> ast = AST.get_ast(['XAB', 'XAC', 'CAB'])  
>>> ast.score('ABC')  
0.3148148148148149
```

Synonym extraction

- **EAST** implements one language-dependent feature: **synonym extraction**
- **Motivation:** Relevance scores should be similar, say, for "*plant taxonomy*" and "*plant classification*", even if the latter can be rarely found in the text collection.
- **Algorithm:** distributional synonym extraction algorithm based on that by Lin [4], which employs the so-called dependency triples (w_1, r, w_2) (idea: *similar texts appear in similar contexts*)
- **Domain-specific synonyms** are likely to be found with this context-based approach

Synonym extraction

- Dependency triples extraction is done by **Yandex Tomita parser** (based on grammatical templates like *"adjective + substantive"* or *"verb + arverb"*)
- Grammar:

```
S -> adj_mod_of interp (Relation.adj_mod_of::...) |  
    adv_of interp (Relation.adv_of::norm="inf") |  
    adv interp (Relation.adv::norm="inf") |  
    ...
```

```
adj_mod_of -> Adj<gnc-agr[1]> Noun<gnc-agr[1]>;  
adv_of -> Adv Verb;  
adv -> Verb Adv;  
...
```

Synonym extraction

- Synonyms extracted from a text collection from the "Izvestia" newspaper:
 - "head" ("глава") \Leftrightarrow "CEO" ("гендиректор")
 - "high" ("высокий") \Leftrightarrow "low" ("низкий")
 - ...
- Low precision is not very critical: among synonymous key phrases we chose the one that has $\max_{w \in \text{syn}(S)} \text{SCORE}(w)$
- To extract synonyms before computing relevance scores:

```
$ east -s keyphrases table <keyphrases_list.txt>  
  <path/to/the/text/collection/>
```


Table of Contents

- 1 Annotated suffix trees
 - Suffix trees
 - Annotated suffix trees
 - AST relevance score
- 2 Algorithms
 - From suffix tries to suffix trees
 - From suffix trees to suffix arrays
- 3 Implementation
 - Package EAST
 - Synonym extraction
- 4 LM Monitor
 - Concept & architecture
 - Keyphrase reference graphs

LM Monitor: Concept

LM Monitor = “*Latent Meaning Monitor*” [2]

- ① Web crawling
 - **RuNeWC**: Russian Newspaper Web Corpus
 - 5 sources available now
- ② Keyphrase analysis
 - Keyphrases are provided by the user
 - Using the AST relevance scores for these keyphrases, a **keyphrase reference graph** is built
 - Text visualization tool

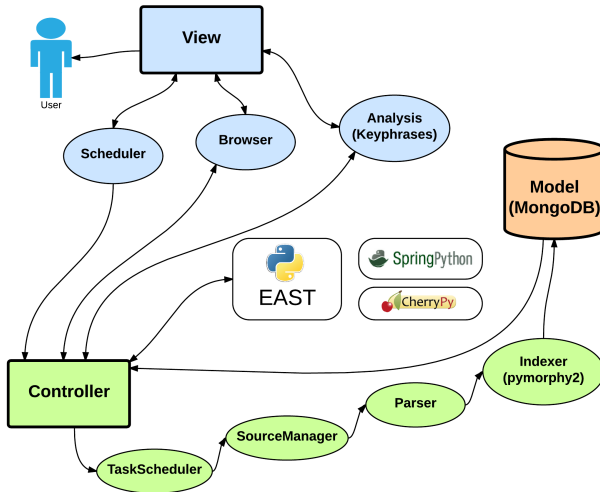
LM Monitor: Development



НАЦИОНАЛЬНЫЙ ИССЛЕДОВАТЕЛЬСКИЙ
УНИВЕРСИТЕТ

- ① Research group “Text analysis and visualization methods”
- ② Head: Boris Mirkin (Sc.D, prof.)
- ③ Staff: Bachelor/Master/PhD students

LM Monitor: Architecture



LM Monitor: System

LM Monitor

Home

Schedule

Corpus

Keyphrases


Data browser


Raw articles & Corpus

Database stats

Sources	5
Articles	4335
Lemmata	64435
Word usages	2669362

Browse articles





All sources ▾

Keypphrase reference graphs

Keypphrase reference graphs:

- Model **directed relations** between keyphrases
- Nodes are keyphrases
- For a keyphrase A ,
 - $r \in [0; 1]$ is a *relevance threshold*: if $SCORE_{AST(T)}(A) > r$, then A is considered to be relevant to text T (usually $r = 0.2$)
 - $F(A) = \{ T : SCORE_{AST(T)}(A) > r \}$
- $c \in [0; 1]$ is a *confidence threshold*
- For keyphrases A and B , if $\frac{|F(B) \cap F(A)|}{|F(A)|} \geq c$, then there is an edge in the graph from keyphrase A to keyphrase B (usually $c = 0.6$)
- $A \rightarrow B$ is like an *associative rule*

LM Monitor: Keyphrase reference graphs

LM Monitor Home Schedule Corpus Keyphrases

Keyphrase analysis Publication-Keyphrases table & Reference graph

Keyphrases

Ввод автоматизированного про

Ввод новых технологий

Ввод опционной программы дл

Выплата купона

Выпуск облигаций

Выпуск пресс релизов с полож

LM Monitor: Keyphrase reference graphs



List

No file chosen

Articles



Referral confidence:

Relevance threshold:

Support threshold:

Articles retrieved: 565



November 2014



Su	Mo	Tu	We	Th	Fr	Sa
26	27	28	29	30	31	1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29
30	1	2	3	4	5	6

Publication-Keyphrase table

Reference graph

LM Monitor: Keypphrase reference graphs




Figure: Keypphrase reference graph built for Oct/Nov 2014

LM Monitor: Keyphrase reference graphs

Id	Name
0	Congress of the United States
1	President of the United States
3	Legislative Power
4	Executive Power

Id	Name
0	Le Président de la République
1	Le Gouvernement
2	Le Parlement
3	Le Conseil constitutionnel





Figure: Keyphrase reference graph built USA/France constitutions

Future work

- Automated graph analysis (central nodes visualization etc.)
- Temporal graph analysis (how do graphs change over time?)
- Better support for synonyms

-  *Abouelhoda, M. I.* Replacing Suffix Trees with Enhanced Suffix Arrays / M. I. Abouelhoda, S. Kurtz, E. Ohlebusch // Journal of Discrete Algorithms, Amsterdam: Elsevier. – 2004. – № 2. – pp. 53-86.
-  *Dubov, M.* Automatic Russian Text Processing System / M. Dubov, B. Mirkin, A. Shal // Open Systems. DBMS – 2014. – v. 22 № 10 – pp. 15-17
-  *Gusfield, D.* Algorithms on Strings, Trees, and Sequences: Computer Science and Computational Biology / D. Gusfield. – Cambridge University Press, 1997.
-  *Lin, D.* Automatic Retrieval and Clustering of Similar Words. / D. Lin // Proceedings of the 17th International Conference on Computational Linguistics. – 1998. – pp. 768-774.



Mirkin, B. Method of annotated suffix tree for scoring the extent of presence of a string in text / B. Mirkin, E. Chernyak, O. Chugunova // Business-Informatics – 2012. – № 3(21). – pp. 31-41.