

Efficient and Effective Query Auto-Completion

Giulio Ermanno Pibiri



ISTITUTO DI SCIENZA E TECNOLOGIE
DELL'INFORMAZIONE "A. FAEDO"

Simon Gog



Rossano Venturini



ACM Conference on Research and Development in
Information Retrieval (SIGIR), 2020

27/07/2020

Query Auto-Completion

Given a collection S of **scored** strings and a *partially completed* user query Q ,
find the **top- k** strings that “match” Q in S .

dream the|

dream theater

dream theater images and words

dream theater discografia

dream theater discography

dream theater scenes from a memory

dream theater awake

dream theater another day

dream theater logo

dream theater through her eyes

dream theater octavarium

Setting

We focus on **matching algorithms**, not ranking mechanisms:
we return the “most popular” results from a query log.

Many matching algorithms are possible, such as:
exact, prefix, pattern (substring), edit-distance...

Setting

We focus on **matching algorithms**, not ranking mechanisms:
we return the “most popular” results from a query log.

Many matching algorithms are possible, such as:
exact, prefix, pattern (substring), edit-distance...

greensboro north carol|

- greensboro north carolina trucking jobs
- greensboro north carolina postal route job
- greensboro north carolina liens
- greensboro north carolina hotels
- greensboro north carolina television
- greensboro north carolina office of unemployment
- greensboro north carolina chamber of commerce
- greensboro north carolina channel 2
- greensboro north carolina dmv
- greensboro north carolina motor vehicle

prefix

Setting

We focus on **matching algorithms**, not ranking mechanisms:
we return the “most popular” results from a query log.

Many matching algorithms are possible, such as:
exact, prefix, pattern (substring), edit-distance...

greensboro north carol|

greensboro north carolina trucking jobs

greensboro north carolina postal route job

greensboro north carolina liens

greensboro north carolina hotels

greensboro north carolina television

greensboro north carolina office of unemployment

greensboro north carolina chamber of commerce

greensboro north carolina channel 2

greensboro north carolina dmv

greensboro north carolina motor vehicle

prefix

greensboro north carol|

free online dating in greensboro north carolina

mayor of greensboro north carolina

greensboro north carolina trucking jobs

university of north carolina greensboro

university of north carolina at greensboro

clarion hotel greensboro north carolina

climate controlled storage greensboro north carolina

greensboro north carolina postal route job

homes for sale greensboro north carolina

fire extinguisher refill in greensboro north carolina

conjunctive

Conjunctive-Search

Return strings containing **all** the tokens in the prefix and **any** token prefixed by the suffix.

Build an inverted index where docids are assigned in decreasing score order: smaller docids are better.

Conjunctive-Search

Return strings containing **all** the tokens in the prefix and **any** token prefixed by the suffix.

Build an inverted index where docids are assigned in decreasing score order: smaller docids are better.

docids	completions
9	audi
6	audi a3 sport
3	audi q8 sedan
8	bmw
5	bmw x1
1	bmw i3 sedan
4	bmw i3 sport
2	bmw i3 sportback
7	bmw i8 sport

termids	terms	inverted lists
1	a3	⟨6⟩
2	audi	⟨3, 6, 9⟩
3	bmw	⟨1, 2, 4, 5, 7, 8⟩
4	i3	⟨1, 2, 4⟩
5	i8	⟨7⟩
6	q8	⟨3⟩
7	sedan	⟨1, 3⟩
8	sport	⟨4, 6, 7⟩
9	sportback	⟨2⟩
10	x1	⟨5⟩



Conjunctive-Search

Return strings containing **all** the tokens in the prefix and **any** token prefixed by the suffix.

Build an inverted index where docids are assigned in decreasing score order: smaller docids are better.

docids	completions
9	audi
6	audi a3 sport
3	audi q8 sedan
8	bmw
5	bmw x1
1	bmw i3 sedan
4	bmw i3 sport
2	bmw i3 sportback
7	bmw i8 sport

termids	terms	inverted lists
1	a3	<6>
2	audi	<3, 6, 9>
3	bmw	① 2, 4, 5, 7, 8>
4	i3	<1, 2, 4>
5	i8	<7>
6	q8	<3>
7	sedan	① 3>
8	sport	<4, 6, 7>
9	sportback	<2>
10	x1	<5>

bmw i3 sedan

Conjunctive-Search

Return strings containing **all** the tokens in the prefix and **any** token prefixed by the suffix.

Build an inverted index where docids are assigned in decreasing score order: smaller docids are better.

docids	completions
9	audi
6	audi a3 sport
3	audi q8 sedan
8	bmw
5	bmw x1
1	bmw i3 sedan
4	bmw i3 sport
2	bmw i3 sportback
7	bmw i8 sport

termids	terms	inverted lists
1	a3	<6>
2	audi	<3, 6, 9>
3	bmw	<1, 2, 4, 5, 7, 8>
4	i3	<1, 2, 4>
5	i8	<7>
6	q8	<3>
7	sedan	<1, 3>
8	sport	<4, 6, 7>
9	sportback	<2>
10	x1	<5>

<input type="text" value="bmw s"/>
bmw i3 sedan
bmw i3 sportback

Conjunctive-Search

Return strings containing **all** the tokens in the prefix and **any** token prefixed by the suffix.

Build an inverted index where docids are assigned in decreasing score order: smaller docids are better.

docids	completions
9	audi
6	audi a3 sport
3	audi q8 sedan
8	bmw
5	bmw x1
1	bmw i3 sedan
4	bmw i3 sport
2	bmw i3 sportback
7	bmw i8 sport

termids	terms	inverted lists
1	a3	<6>
2	audi	<3, 6, 9>
3	bmw	<1, 2, 4, 5, 7, 8>
4	i3	<1, 2, 4>
5	i8	<7>
6	q8	<3>
7	sedan	<1, 3>
8	sport	<4, 6, 7>
9	sportback	<2>
10	x1	<5>

Q bmw s
bmw i3 s edan
bmw i3 s portback
bmw i3 s port

Conjunctive-Search

bmw [7,9]

```
1 ConjunctiveSearch(prefix, [ℓ, r], k) :
2   intersection = index.IntersectionIterator(prefix)
3   results = [ ], heap = [ ]
4   for i = ℓ; i ≤ r; i = i + 1 :
5     heap.Append(index.Iterator(i))
6   heap.MakeHeap()
7   while intersection.HasNext() and !heap.Empty() :
8     x = intersection.Next()
9     while !heap.Empty() :
10      top = heap.Top()
11      if top.docid > x : break
12      if top.docid < x :
13        if top.NextGeq(x) < ∞ : heap.Heapify()
14        else : heap.Pop()
15      else :
16        results.Append(x)
17        if |results| == k : return results
18        break
19 return results
```

Conjunctive-Search

```
1 ConjunctiveSearch(prefix, [l, r], k) :
2   intersection = index.IntersectionIterator(prefix)
3   results = [], heap = []
4   for i = l; i ≤ r; i = i + 1 :
5     | heap.Append(index.Iterator(i))
6   heap.MakeHeap()
7   while intersection.HasNext() and !heap.Empty() :
8     | x = intersection.Next()
9     | while !heap.Empty() :
10    |   | top = heap.Top()
11    |   | if top.docid > x : break
12    |   | if top.docid < x :
13    |   |   | if top.NextGeq(x) < ∞ : heap.Heapify()
14    |   |   | else : heap.Pop()
15    |   | else :
16    |   |   | results.Append(x)
17    |   |   | if |results| == k : return results
18    |   |   | break
19 return results
```

bmw [7,9]



Heap-based approach:
(1) Much better than explicitly computing the union.
(2) Terms involved in union may be too many!

Conjunctive-Search

```
1 ConjunctiveSearch(prefix, [l, r], k) :  
2   results = [ ]  
3   intersection = index.IntersectionIterator(prefix)  
4   while intersection.HasNext() :  
5     x = intersection.Next()  
6     completion = Extract(x)  
7     if completion intersects [l, r] :  
8       results.Append(x)  
9       if |results| == k : break  
10  return results
```

**Forward
search**

Conjunctive-Search

```

1 ConjunctiveSearch(prefix, [l, r], k) :
2   results = [ ]
3   intersection = index.IntersectionIterator(prefix)
4   while intersection.HasNext() :
5     x = intersection.Next()
6     completion = Extract(x)
7     if completion intersects [l, r] :
8       results.Append(x)
9       if |results| == k : break
10  return results

```

Forward search

Q bmw 2|

terms	inverted lists	termids
1	4	0
2015	0, 5, 6	1
2016	3	2
2017	1, 2	3
3	0, 1, 3, 5	4
8	2, 6	5
a	5	6
audi	2, 5	7
bmw	0, 1, 3, 4, 6	8
i	0, 1, 3, 6	9
q	2	10
x	4	11

docids	completions	sets
5	audi a 3 2015	7, 6, 4, 1
2	audi q 8 2017	7, 10, 5, 3
4	bmw x 1	8, 11, 0
0	bmw i 3 2015	8, 9, 4, 1
3	bmw i 3 2016	8, 9, 4, 2
1	bmw i 3 2017	8, 9, 4, 3
6	bmw i 8 2015	8, 9, 5, 1

Conjunctive-Search

```

1 ConjunctiveSearch(prefix, [l, r], k) :
2   results = [ ]
3   intersection = index.IntersectionIterator(prefix)
4   while intersection.HasNext() :
5     x = intersection.Next()
6     completion = Extract(x)
7     if completion intersects [l, r] :
8       results.Append(x)
9       if |results| == k : break
10  return results

```

Forward search

Q **bmw 2**|

terms	inverted lists	termids
1	4	0
2015	0, 5, 6	1
2016	3	2
2017	1, 2	3
3	0, 1, 3, 5	4
8	2, 6	5
a	5	6
audi	2, 5	7
bmw	0, 1, 3, 4, 6	8
i	0, 1, 3, 6	9
q	2	10
x	4	11

docids	completions	sets
5	audi a 3 2015	7, 6, 4, 1
2	audi q 8 2017	7, 10, 5, 3
4	bmw x 1	8, 11, 0
0	bmw i 3 2015	8, 9, 4, 1
3	bmw i 3 2016	8, 9, 4, 2
1	bmw i 3 2017	8, 9, 4, 3
6	bmw i 8 2015	8, 9, 5, 1

Conjunctive-Search

```

1 ConjunctiveSearch(prefix, [l, r], k) :
2   results = [ ]
3   intersection = index.IntersectionIterator(prefix)
4   while intersection.HasNext() :
5     x = intersection.Next()
6     completion = Extract(x)
7     if completion intersects [l, r] :
8       results.Append(x)
9       if |results| == k : break
10  return results

```

Forward search



terms	inverted lists	termids
1	4	0
2015	0, 5, 6	1
2016	3	2
2017	1, 2	3
3	0, 1, 3, 5	4
8	2, 6	5
a	5	6
audi	2, 5	7
bmw	0, 1, 3, 4, 6	8
i	0, 1, 3, 6	9
q	2	10
x	4	11

docids	completions	sets
5	audi a 3 2015	7 , 6 , 4 , 1
2	audi q 8 2017	7 , 10 , 5 , 3
4	bmw x 1	8 , 11 , 0
0	bmw i 3 2015	8 , 9 , 4 , 1
3	bmw i 3 2016	8 , 9 , 4 , 2
1	bmw i 3 2017	8 , 9 , 4 , 3
6	bmw i 8 2015	8 , 9 , 5 , 1

Conjunctive-Search

```

1 ConjunctiveSearch(prefix, [l, r], k) :
2   results = [ ]
3   intersection = index.IntersectionIterator(prefix)
4   while intersection.HasNext() :
5     x = intersection.Next()
6     completion = Extract(x)
7     if completion intersects [l, r] :
8       results.Append(x)
9       if |results| == k : break
10  return results

```

Forward search



terms	inverted lists	termids
1	4	0
2015	0, 5, 6	1
2016	3	2
2017	1, 2	3
3	0, 1, 3, 5	4
8	2, 6	5
a	5	6
audi	2, 5	7
bmw	0, 1, 3, 4, 6	8
i	0, 1, 3, 6	9
q	2	10
x	4	11

docids	completions	sets
5	audi a 3 2015	7, 6, 4, 1
2	audi q 8 2017	7, 10, 5, 3
4	bmw x 1	8, 11, 0
0	bmw i 3 2015	8, 9, 4, 1
3	bmw i 3 2016	8, 9, 4, 2
1	bmw i 3 2017	8, 9, 4, 3
6	bmw i 8 2015	8, 9, 5, 1


Conjunctive-Search

```

1 ConjunctiveSearch(prefix, [l, r], k) :
2   results = [ ]
3   intersection = index.IntersectionIterator(prefix)
4   while intersection.HasNext() :
5     x = intersection.Next()
6     completion = Extract(x)
7     if completion intersects [l, r] :
8       results.Append(x)
9       if |results| == k : break
10  return results

```

Forward search



terms	inverted lists	termids
1	4	0
2015	0, 5, 6	1
2016	3	2
2017	1, 2	3
3	0, 1, 3, 5	4
8	2, 6	5
a	5	6
audi	2, 5	7
bmw	0, 1, 3, 4, 6	8
i	0, 1, 3, 6	9
q	2	10
x	4	11

docids	completions	sets
5	audi a 3 2015	7, 6, 4, 1
2	audi q 8 2017	7, 10, 5, 3
4	bmw x 1	8, 11, 0
0	bmw i 3 2015	8, 9, 4, 1
3	bmw i 3 2016	8, 9, 4, 2
1	bmw i 3 2017	8, 9, 4, 3
6	bmw i 8 2015	8, 9, 5, 1

Conjunctive-Search

```

1 ConjunctiveSearch(prefix, [l, r], k) :
2   results = [ ]
3   intersection = index.IntersectionIterator(prefix)
4   while intersection.HasNext() :
5     x = intersection.Next()
6     completion = Extract(x)
7     if completion intersects [l, r] :
8       results.Append(x)
9       if |results| == k : break
10  return results

```

Forward search

Q bmw 2
bmw i 3 2015
bmw i 3 2017

terms	inverted lists	termids
1	4	0
2015	0, 5, 6	1
2016	3	2
2017	1, 2	3
3	0, 1, 3, 5	4
8	2, 6	5
a	5	6
audi	2, 5	7
bmw	0, 1, 3, 4, 6	8
i	0, 1, 3, 6	9
q	2	10
x	4	11

docids	completions	sets
5	audi a 3 2015	7, 6, 4, 1
2	audi q 8 2017	7, 10, 5, 3
4	bmw x 1	8, 11, 0
0	bmw i 3 2015	8, 9, 4, 1
3	bmw i 3 2016	8, 9, 4, 2
1	bmw i 3 2017	8, 9, 4, 3
6	bmw i 8 2015	8, 9, 5, 1

Conjunctive-Search

```

1 ConjunctiveSearch(prefix, [l, r], k) :
2   results = [ ]
3   intersection = index.IntersectionIterator(prefix)
4   while intersection.HasNext() :
5     x = intersection.Next()
6     completion = Extract(x)
7     if completion intersects [l, r] :
8       results.Append(x)
9       if |results| == k : break
10  return results

```

Forward search

Q bmw 2
bmw i 3 2015
bmw i 3 2017

terms	inverted lists	termids
1	4	0
2015	0, 5, 6	1
2016	3	2
2017	1, 2	3
3	0, 1, 3, 5	4
8	2, 6	5
a	5	6
audi	2, 5	7
bmw	0, 1, 3, 4, 6	8
i	0, 1, 3, 6	9
q	2	10
x	4	11

docids	completions	sets
5	audi a 3 2015	7, 6, 4, 1
2	audi q 8 2017	7, 10, 5, 3
4	bmw x 1	8, 11, 0
0	bmw i 3 2015	8, 9, 4, 1
3	bmw i 3 2016	8, 9, 4, 2
1	bmw i 3 2017	8, 9, 4, 3
6	bmw i 8 2015	8, 9, 5, 1

Conjunctive-Search

```

1 ConjunctiveSearch(prefix, [l, r], k) :
2   results = [ ]
3   intersection = index.IntersectionIterator(prefix)
4   while intersection.HasNext() :
5     x = intersection.Next()
6     completion = Extract(x)
7     if completion intersects [l, r] :
8       results.Append(x)
9       if |results| == k : break
10  return results

```

Forward search

Q bmw 2
bmw i 3 2015
bmw i 3 2017
bmw i 3 2016

terms	inverted lists	termids
1	4	0
2015	0, 5, 6	1
2016	3	2
2017	1, 2	3
3	0, 1, 3, 5	4
8	2, 6	5
a	5	6
audi	2, 5	7
bmw	0, 1, 3, 4, 6	8
i	0, 1, 3, 6	9
q	2	10
x	4	11

docids	completions	sets
5	audi a 3 2015	7, 6, 4, 1
2	audi q 8 2017	7, 10, 5, 3
4	bmw x 1	8, 11, 0
0	bmw i 3 2015	8, 9, 4, 1
3	bmw i 3 2016	8, 9, 4, 2
1	bmw i 3 2017	8, 9, 4, 3
6	bmw i 8 2015	8, 9, 5, 1

Conjunctive-Search

```

1 ConjunctiveSearch(prefix, [l, r], k) :
2   results = [ ]
3   intersection = index.IntersectionIterator(prefix)
4   while intersection.HasNext() :
5     x = intersection.Next()
6     completion = Extract(x)
7     if completion intersects [l, r] :
8       results.Append(x)
9       if |results| == k : break
10  return results

```

Forward
search

Q bmw 2
bmw i 3 2015
bmw i 3 2017
bmw i 3 2016

terms	inverted lists	termids
1	4	0
2015	0, 5, 6	1
2016	3	2
2017	1, 2	3
3	0, 1, 3, 5	4
8	2, 6	5
a	5	6
audi	2, 5	7
bmw	0, 1, 3, 4, 6	8
i	0, 1, 3, 6	9
q	2	10
x	4	11

docids	completions	sets
5	audi a 3 2015	7, 6, 4, 1
2	audi q 8 2017	7, 10, 5, 3
4	bmw x 1	8, 11, 0
0	bmw i 3 2015	8, 9, 4, 1
3	bmw i 3 2016	8, 9, 4, 2
1	bmw i 3 2017	8, 9, 4, 3
6	bmw i 8 2015	8, 9, 5, 1

Conjunctive-Search

```

1 ConjunctiveSearch(prefix, [l, r], k) :
2   results = [ ]
3   intersection = index.IntersectionIterator(prefix)
4   while intersection.HasNext() :
5     x = intersection.Next()
6     completion = Extract(x)
7     if completion intersects [l, r] :
8       results.Append(x)
9       if |results| == k : break
10  return results

```

Forward search

Q bmw 2
bmw i 3 2015
bmw i 3 2017
bmw i 3 2016

terms	inverted lists	termids
1	4	0
2015	0, 5, 6	1
2016	3	2
2017	1, 2	3
3	0, 1, 3, 5	4
8	2, 6	5
a	5	6
audi	2, 5	7
bmw	0, 1, 3, 4, 6	8
i	0, 1, 3, 6	9
q	2	10
x	4	11

docids	completions	sets
5	audi a 3 2015	7, 6, 4, 1
2	audi q 8 2017	7, 10, 5, 3
4	bmw x 1	8, 11, 0
0	bmw i 3 2015	8, 9, 4, 1
3	bmw i 3 2016	8, 9, 4, 2
1	bmw i 3 2017	8, 9, 4, 3
6	bmw i 8 2015	8, 9, 5, 1

Conjunctive-Search

```

1 ConjunctiveSearch(prefix, [l, r], k) :
2   results = [ ]
3   intersection = index.IntersectionIterator(prefix)
4   while intersection.HasNext() :
5     x = intersection.Next()
6     completion = Extract(x)
7     if completion intersects [l, r] :
8       results.Append(x)
9       if |results| == k : break
10  return results

```

Forward search

Q bmw 2
bmw i 3 2015
bmw i 3 2017
bmw i 3 2016
bmw i 8 2015

terms	inverted lists	termids
1	4	0
2015	0, 5, 6	1
2016	3	2
2017	1, 2	3
3	0, 1, 3, 5	4
8	2, 6	5
a	5	6
audi	2, 5	7
bmw	0, 1, 3, 4, 6	8
i	0, 1, 3, 6	9
q	2	10
x	4	11

docids	completions	sets
5	audi a 3 2015	7, 6, 4, 1
2	audi q 8 2017	7, 10, 5, 3
4	bmw x 1	8, 11, 0
0	bmw i 3 2015	8, 9, 4, 1
3	bmw i 3 2016	8, 9, 4, 2
1	bmw i 3 2017	8, 9, 4, 3
6	bmw i 8 2015	8, 9, 5, 1

Conjunctive-Search

```

1 ConjunctiveSearch(prefix, [l, r], k) :
2   results = [ ]
3   intersection = index.IntersectionIterator(prefix)
4   while intersection.HasNext() :
5     x = intersection.Next()
6     completion = Extract(x)
7     if completion intersects [l, r] :
8       results.Append(x)
9       if |results| == k : break
10  return results

```

Forward search

Forward-Search approach:
 (1) No heap management.
 (2) Need “direct” access to completions: Fwd or FC.

Q bmw 2
bmw i 3 2015
bmw i 3 2017
bmw i 3 2016
bmw i 8 2015

terms	inverted lists	termids
1	4	0
2015	0, 5, 6	1
2016	3	2
2017	1, 2	3
3	0, 1, 3, 5	4
8	2, 6	5
a	5	6
audi	2, 5	7
bmw	0, 1, 3, 4, 6	8
i	0, 1, 3, 6	9
q	2	10
x	4	11

docids	completions	sets
5	audi a 3 2015	7, 6, 4, 1
2	audi q 8 2017	7, 10, 5, 3
4	bmw x 1	8, 11, 0
0	bmw i 3 2015	8, 9, 4, 1
3	bmw i 3 2016	8, 9, 4, 2
1	bmw i 3 2017	8, 9, 4, 3
6	bmw i 8 2015	8, 9, 5, 1

Experiments

Machine equipped with Intel i9-9900K cores (@3.60 GHz),
64 GB of RAM, and running Linux 5 (64 bits).

C++ code available at

<https://github.com/jermp/autocomplete>

Statistic	AOL	MSN	EBAY
Queries	10,142,395	7,083,363	7,295,104
Uncompressed size in MiB	299	208	189
Unique query terms	3,825,848	2,590,937	323,180
Avg. num. of chars per term	14.58	14.18	7.32
Avg. num. of queries per term	7.87	8.15	73.02
Avg. num. of terms per query	2.99	2.99	3.24

Datasets

Experiments — Efficiency

		(a) AOL							(b) MSN						
		Query terms							Query terms						
%		1	2	3	4	5	6	7+	1	2	3	4	5	6	7+
Fwd	0	4	5	22	30	24	24	16	4	5	14	15	11	10	7
	25	2	97	70	41	30	25	16	1	39	34	18	13	10	7
	50	0	149	77	48	30	25	16	0	56	38	19	13	10	8
	75	0	150	76	48	30	25	16	0	57	37	19	12	10	7
FC	0	5	15	27	30	24	24	16	5	15	17	15	11	10	7
	25	3	251	110	45	31	25	16	2	101	51	19	13	10	8
	50	1	370	121	56	31	25	16	1	137	58	21	13	10	7
	75	0	375	121	57	32	25	16	0	137	57	21	13	10	7
Heap	0	55,537	29,189	30,498	22,431	17,713	16,474	13,312	7,626	12,459	11,964	8,921	6,164	5,749	5,686
	25	474	623	957	485	376	378	299	353	252	256	282	170	192	125
	50	1	251	178	251	229	123	178	10	73	70	109	84	66	54
	75	0	226	162	240	219	116	173	1	61	62	83	80	63	51
Hyb	0	286	2,718	1,673	965	634	503	413	53	1,626	915	477	307	270	237
	25	11	184	223	276	258	221	192	10	90	109	127	111	111	90
	50	10	126	185	270	250	217	186	7	53	97	122	107	108	87
	75	6	116	178	268	248	216	184	4	46	95	121	106	106	85

Top-**10** conjunctive-search query timings in μsec per query, by varying query length and percentage of the last query token.

Experiments — Effectiveness

	(a) AOL							(b) MSN						
	Query terms							Query terms						
%	1	2	3	4	5	6	7+	1	2	3	4	5	6	7+
0	17	107	207	327	295	270	270	27	139	252	283	325	206	248
25	19	178	246	373	298	155	356	23	231	310	297	333	190	200
50	23	227	302	440	364	213	524	27	243	313	320	359	251	208
75	41	282	362	504	424	257	882	44	284	364	357	407	319	236

Percentage of better scored results returned by conjunctive-search with respect to those returned by prefix-search for top-**10** queries.

Experiments — Space

	AOL		MSN		EBAY	
	MiB	bpc	MiB	bpc	MiB	bpc
Fwd	312	32.28	218	32.32	168	24.14
FC	266	27.51	185	27.42	140	20.13
Heap	254	26.25	177	26.25	139	19.99
Hyb	275	28.48	191	28.26	157	22.50

Space usage in total MiB and bytes per completion (bpc).

Experiments — Space

Statistic	AOL		MSN		EBAY	
Queries	10,142,395		7,083,363		7,295,104	
Uncompressed size in MiB	299		208		189	
Unique query terms	3,825,848		2,590,937		323,180	
Avg. nun	AOL		MSN		EBAY	
Avg. nun						
Avg. nun						
	MiB	bpc	MiB	bpc	MiB	bpc
Fwd	312	32.28	218	32.32	168	24.14
FC	266	27.51	185	27.42	140	20.13
Heap	254	26.25	177	26.25	139	19.99
Hyb	275	28.48	191	28.26	157	22.50

Space usage in total MiB and bytes per completion (bpc).

Experiments — Space

Statistic	AOL		MSN		EBAY	
Queries	10,142,395		7,083,363		7,295,104	
Uncompressed size in MiB	299		208		189	
Unique query terms	3,825,848		2,590,937		323,180	
Avg. nun	AOL		MSN		EBAY	
Avg. nun	MiB	bpc	MiB	bpc	MiB	bpc
Avg. nun						
Fwd	312	32.28	218	32.32	168	24.14
FC	266	27.51	185	27.42	140	20.13
Heap	254	26.25	177	26.25	139	19.99
Hyb	275	28.48	191	28.26	157	22.50

Space usage in total MiB and bytes per completion (bpc).

Take-away Messages

- Conjunctive-search overcomes the limited effectiveness of prefix-search by returning more and better scored results.
- While prefix-search is very fast (less than 3 μ sec per query on average), conjunctive-search is more expensive and costs between 4 and 500 μ sec per query depending on the size of the query.
- Our optimized implementation of conjunctive-search substantially outperforms the use of a classical as well as blocked inverted index with small extra, or even less, space.

Thanks for your attention!

Prefix-Search

```
1 Complete(query, k) :  
2   prefix, suffix = Parse(dictionary, query)  
3   if prefix was not found : return [ ]  
4   [l, r] = dictionary.LocatePrefix(suffix)  
5   if [l, r] is invalid : return [ ]  
6   [p, q] = completions.LocatePrefix(prefix, [l, r])  
7   if [p, q] is invalid : return [ ]  
8   topk_ids = RMQ([p, q], k)  
9   strings = ExtractStrings(topk_ids)  
10  return strings
```

Prefix-Search

```
1 Complete(query, k) :  
2   prefix, suffix = Parse(dictionary, query)  
3   if prefix was not found : return [ ]  
4   [l, r] = dictionary.LocatePrefix(suffix)  
5   if [l, r] is invalid : return [ ]  
6   [p, q] = completions.LocatePrefix(prefix, [l, r])  
7   if [p, q] is invalid : return [ ]  
8   topk_ids = RMQ([p, q], k)  
9   strings = ExtractStrings(topk_ids)  
10  return strings
```

docids	completions
9	audi
6	audi a3 sport
3	audi q8 sedan
8	bmw
5	bmw x1
1	bmw i3 sedan
4	bmw i3 sport
2	bmw i3 sportback
7	bmw i8 sport

termids	terms
1	a3
2	audi
3	bmw
4	i3
5	i8
6	q8
7	sedan
8	sport
9	sportback
10	x1

🔍 **bmw i3 s**

Prefix-Search

```
1 Complete(query, k) :  
2   prefix, suffix = Parse(dictionary, query)  
3   if prefix was not found : return []  
4   [l, r] = dictionary.LocatePrefix(suffix) 1  
5   if [l, r] is invalid : return []  
6   [p, q] = completions.LocatePrefix(prefix, [l, r])  
7   if [p, q] is invalid : return []  
8   topk_ids = RMQ([p, q], k)  
9   strings = ExtractStrings(topk_ids)  
10  return strings
```

docids	completions
9	audi
6	audi a3 sport
3	audi q8 sedan
8	bmw
5	bmw x1
1	bmw i3 sedan
4	bmw i3 sport
2	bmw i3 sportback
7	bmw i8 sport

termids	terms
1	a3
2	audi
3	bmw
4	i3
5	i8
6	q8
7	sedan
8	sport
9	sportback
10	x1

Q **bmw i3 s**

Prefix-Search

```
1 Complete(query, k) :  
2   prefix, suffix = Parse(dictionary, query)  
3   if prefix was not found : return []  
4   [l, r] = dictionary.LocatePrefix(suffix) 1  
5   if [l, r] is invalid : return []  
6   [p, q] = completions.LocatePrefix(prefix, [l, r]) 2  
7   if [p, q] is invalid : return []  
8   topk_ids = RMQ([p, q], k)  
9   strings = ExtractStrings(topk_ids)  
10  return strings
```

docids	completions
9	audi
6	audi a3 sport
3	audi q8 sedan
8	bmw
5	bmw x1
1	bmw i3 sedan
4	bmw i3 sport
2	bmw i3 sportback
7	bmw i8 sport

termids	terms
1	a3
2	audi
3	bmw
4	i3
5	i8
6	q8
7	sedan
8	sport
9	sportback
10	x1

Q **bmw i3 s**

Prefix-Search

```

1 Complete(query, k):
2   prefix, suffix = Parse(dictionary, query)
3   if prefix was not found: return []
4   [l, r] = dictionary.LocatePrefix(suffix) 1
5   if [l, r] is invalid: return []
6   [p, q] = completions.LocatePrefix(prefix, [l, r]) 2
7   if [p, q] is invalid: return []
8   topk_ids = RMQ([p, q], k) 3
9   strings = ExtractStrings(topk_ids)
10  return strings

```

Docids are assigned in decreasing score order: top-k algorithm reduces to RMQ.

docids	completions
9	audi
6	audi a3 sport
3	audi q8 sedan
8	bmw
5	bmw x1
1	bmw i3 sedan
4	bmw i3 sport
2	bmw i3 sportback
7	bmw i8 sport

Annotations: A red box highlights docids 1, 4, and 2. A red '3' is to the left of this box. A red '2' is to the right of the row containing 'bmw x1'. A red box highlights the completions 'bmw i3 sedan', 'bmw i3 sport', and 'bmw i3 sportback'.

termids	terms
1	a3
2	audi
3	bmw
4	i3
5	i8
6	q8
7	sedan
8	sport
9	sportback
10	x1

Annotation: A red box highlights terms 'sedan', 'sport', and 'sportback'. A red '1' is to the right of this box.

Q **bmw i3 s**

Prefix-Search

```

1 Complete(query, k) :
2   prefix, suffix = Parse(dictionary, query)
3   if prefix was not found : return [ ]
4   [l, r] = dictionary.LocatePrefix(suffix) 1
5   if [l, r] is invalid : return [ ]
6   [p, q] = completions.LocatePrefix(prefix, [l, r]) 2
7   if [p, q] is invalid : return [ ]
8   topk_ids = RMQ([p, q], k) 3
9   strings = ExtractStrings(topk_ids)
10  return strings

```

Docids are assigned in decreasing score order: top-k algorithm reduces to RMQ.

docids	completions
9	audi
6	audi a3 sport
3	audi q8 sedan
8	bmw
5	bmw x1
①	bmw i3 sedan
4	bmw i3 sport
2	bmw i3 sportback
7	bmw i8 sport

Annotations: A red box highlights docids 1, 4, and 2. A green box highlights the completions for docid 1. A red box highlights the completions for docids 4 and 2. A red '2' is next to docid 5, and a red '3' is next to the highlighted docids.

termids	terms
1	a3
2	audi
3	bmw
4	i3
5	i8
6	q8
7	sedan
8	sport
9	sportback
10	x1

Annotation: A red box highlights terms 7, 8, and 9, with a red '1' next to it.

bmw i3 sedan

Prefix-Search

```

1 Complete(query, k) :
2   prefix, suffix = Parse(dictionary, query)
3   if prefix was not found : return [ ]
4   [l, r] = dictionary.LocatePrefix(suffix) 1
5   if [l, r] is invalid : return [ ]
6   [p, q] = completions.LocatePrefix(prefix, [l, r]) 2
7   if [p, q] is invalid : return [ ]
8   topk_ids = RMQ([p, q], k) 3
9   strings = ExtractStrings(topk_ids)
10  return strings

```

Docids are assigned in decreasing score order: top-k algorithm reduces to RMQ.

docids	completions
9	audi
6	audi a3 sport
3	audi q8 sedan
8	bmw
5	bmw x1
①	bmw i3 sedan
4	bmw i3 sport
②	bmw i3 sportback
7	bmw i8 sport

Annotations: A red box highlights docids 1 and 2. A green box highlights the completions for docids 1 and 2. A red box highlights the completions for docids 7, 8, and 9.

termids	terms
1	a3
2	audi
3	bmw
4	i3
5	i8
6	q8
7	sedan
8	sport
9	sportback
10	x1

Annotation: A red box highlights termids 7, 8, and 9.

🔍 **bmw i3 s**

bmw i3 sedan

bmw i3 sportback

Prefix-Search

```

1 Complete(query, k) :
2   prefix, suffix = Parse(dictionary, query)
3   if prefix was not found : return [ ]
4   [l, r] = dictionary.LocatePrefix(suffix) 1
5   if [l, r] is invalid : return [ ]
6   [p, q] = completions.LocatePrefix(prefix, [l, r]) 2
7   if [p, q] is invalid : return [ ]
8   topk_ids = RMQ([p, q], k) 3
9   strings = ExtractStrings(topk_ids)
10  return strings

```

Docids are assigned in decreasing score order: top-k algorithm reduces to RMQ.

docids	completions
9	audi
6	audi a3 sport
3	audi q8 sedan
8	bmw
5	bmw x1
3	bmw i3 sedan
4	bmw i3 sport
2	bmw i3 sportback
7	bmw i8 sport

termids	terms
1	a3
2	audi
3	bmw
4	i3
5	i8
6	q8
7	sedan
8	sport
9	sportback
10	x1

Q bmw i3 s
bmw i3 sedan
bmw i3 sportback
bmw i3 sport