

Paolo Ferragina

List of Publications by Year in descending order

Source: <https://exaly.com/author-pdf/7828714/publications.pdf>

Version: 2024-02-01

117
papers

4,816
citations

236612

25
h-index

149479

56
g-index

130
all docs

130
docs citations

130
times ranked

2041
citing authors

#	ARTICLE	IF	CITATIONS
1	Indexing compressed text. Journal of the ACM, 2005, 52, 552-581.	1.8	538
2	TAGME. , 2010, , .		460
3	Compressed representations of sequences and full-text indexes. ACM Transactions on Algorithms, 2007, 3, 20.	0.9	271
4	The string B-tree. Journal of the ACM, 1999, 46, 236-280.	1.8	247
5	Optical recognition of motor vehicle license plates. IEEE Transactions on Vehicular Technology, 1995, 44, 790-799.	3.9	193
6	On the sorting-complexity of suffix tree construction. Journal of the ACM, 2000, 47, 987-1011.	1.8	175
7	Fast and Accurate Annotation of Short Texts with Wikipedia Pages. IEEE Software, 2012, 29, 70-75.	2.1	161
8	A personalized search engine based on web-snippet hierarchical clustering. , 2005, , .		142
9	Engineering a Lightweight Suffix Array Construction Algorithm. Algorithmica, 2004, 40, 33-50.	1.0	118
10	A framework for benchmarking entity-annotation systems. , 2013, , .		117
11	Compressing and indexing labeled trees, with applications. Journal of the ACM, 2009, 57, 1-33.	1.8	113
12	Boosting textual compression in optimal linear time. Journal of the ACM, 2005, 52, 688-713.	1.8	106
13	From TagME to WAT. , 2014, , .		105
14	GERBIL. , 2015, , .		104
15	The PGM-index. Proceedings of the VLDB Endowment, 2020, 13, 1162-1175.	2.1	101
16	Compression-based classification of biological sequences and structures via the Universal Similarity Metric: experimental assessment. BMC Bioinformatics, 2007, 8, 252.	1.2	89
17	A simple storage scheme for strings achieving entropy bounds. Theoretical Computer Science, 2007, 372, 115-121.	0.5	86
18	Lightweight Data Indexing and Compression in External Memory. Algorithmica, 2012, 63, 707-730.	1.0	85

#	ARTICLE	IF	CITATIONS
19	Compressed text indexes. <i>Journal of Experimental Algorithmics</i> , 2009, 13, .	0.7	84
20	q-gram based database searching using a suffix array (QUASAR). , 1999, , .		80
21	A public data set of spatio-temporal match events in soccer competitions. <i>Scientific Data</i> , 2019, 6, 236.	2.4	80
22	Topical clustering of search results. , 2012, , .		73
23	PlayeRank. <i>ACM Transactions on Intelligent Systems and Technology</i> , 2019, 10, 1-27.	2.9	72
24	Structuring labeled trees for optimal succinctness, and beyond. , 0, , .		64
25	The compressed permuterm index. <i>ACM Transactions on Algorithms</i> , 2010, 7, 1-21.	0.9	64
26	A Theoretical and Experimental Study on the Construction of Suffix Arrays in External Memory. <i>Algorithmica</i> , 2002, 32, 1-35.	1.0	55
27	An Alphabet-Friendly FM-Index. <i>Lecture Notes in Computer Science</i> , 2004, , 150-160.	1.0	55
28	Classification of Short Texts by Deploying Topical Annotations. <i>Lecture Notes in Computer Science</i> , 2012, , 376-387.	1.0	51
29	On sorting strings in external memory (extended abstract). , 1997, , .		46
30	The myriad virtues of Wavelet Trees. <i>Information and Computation</i> , 2009, 207, 849-866.	0.5	43
31	Wiser: A semantic approach for expert finding in academia based on entity linking. <i>Information Systems</i> , 2019, 82, 1-16.	2.4	43
32	A personalized search engine based on Web-snippet hierarchical clustering. <i>Software - Practice and Experience</i> , 2008, 38, 189-225.	2.5	41
33	Brotli. <i>ACM Transactions on Information Systems</i> , 2019, 37, 1-30.	3.8	35
34	An experimental study of a compressed index. <i>Information Sciences</i> , 2001, 135, 13-28.	4.0	33
35	On searching compressed string collections cache-obliviously. , 2008, , .		33
36	On compressing the textual web. , 2010, , .		33

#	ARTICLE	IF	CITATIONS
37	Give more data, awareness and control to individual citizens, and they will help COVID-19 containment. <i>Ethics and Information Technology</i> , 2021, 23, 1-6.	2.3	33
38	On Optimally Partitioning a Text to Improve Its Compression. <i>Algorithmica</i> , 2011, 61, 51-74.	1.0	26
39	A fully-dynamic data structure for external substring search. , 1995, , .		24
40	The SMAPH system for query entity recognition and disambiguation. , 2014, , .		23
41	The Engineering of a Compression Boosting Library: Theory vs Practice in BWT Compression. <i>Lecture Notes in Computer Science</i> , 2006, , 756-767.	1.0	20
42	Dynamic Text Indexing under String Updates. <i>Journal of Algorithms</i> , 1997, 22, 296-328.	0.9	19
43	A Piggyback System for Joint Entity Mention Detection and Linking in Web Queries. , 2016, , .		19
44	Two-dimensional substring indexing. <i>Journal of Computer and System Sciences</i> , 2003, 66, 763-774.	0.9	17
45	On the Bit-Complexity of Lempel-Ziv Compression. <i>SIAM Journal on Computing</i> , 2013, 42, 1521-1541.	0.8	17
46	On updating suffix tree labels. <i>Theoretical Computer Science</i> , 1998, 201, 249-262.	0.5	16
47	Optimal On-Line Search and Sublinear Time Update in String Matching. <i>SIAM Journal on Computing</i> , 1998, 27, 713-736.	0.8	16
48	RANDOMIZED EXTERNAL-MEMORY ALGORITHMS FOR LINE SEGMENT INTERSECTION AND OTHER GEOMETRIC PROBLEMS. <i>International Journal of Computational Geometry and Applications</i> , 2001, 11, 305-337.	0.3	16
49	The Anatomy of a Hierarchical Clustering Engine for Web-page, News and Book Snippets. , 0, , .		15
50	Compressed permuterm index. , 2007, , .		15
51	Engineering a Lightweight Suffix Array Construction Algorithm. <i>Lecture Notes in Computer Science</i> , 2002, , 698-710.	1.0	14
52	Suffix Arrays on Words. <i>Lecture Notes in Computer Science</i> , 2007, , 328-339.	1.0	14
53	On the bit-complexity of Lempel-Ziv compression. , 2009, , .		14
54	Query-log mining for detecting spam. , 2008, , .		13

#	ARTICLE	IF	CITATIONS
55	Dynamic Dictionary Matching in External Memory. Information and Computation, 1998, 146, 85-99.	0.5	12
56	A Two-Stage Framework for Computing Entity Relatedness in Wikipedia. , 2017, , .		12
57	The Myriad Virtues of Wavelet Trees. Lecture Notes in Computer Science, 2006, , 560-571.	1.0	12
58	Bicriteria Data Compression: Efficient and Usable. Lecture Notes in Computer Science, 2014, , 406-417.	1.0	12
59	NETME: on-the-fly knowledge network construction from biomedical literature. Applied Network Science, 2022, 7, 1.	0.8	12
60	The Anatomy of SnakeT: A Hierarchical Clustering Engine for Web-Page Snippets. Lecture Notes in Computer Science, 2004, , 506-508.	1.0	11
61	On the performance of learned data structures. Theoretical Computer Science, 2021, 871, 107-120.	0.5	10
62	Efficient dynamic method-lookup for object oriented languages. Lecture Notes in Computer Science, 1996, , 107-120.	1.0	10
63	Distribution-Aware Compressed Full-Text Indexes. Algorithmica, 2013, 67, 529-546.	1.0	9
64	The BioPrompt-box: an ontology-based clustering tool for searching in biological databases. BMC Bioinformatics, 2007, 8, S8.	1.2	8
65	<scp>Swat</scp>: A system for detecting salient Wikipedia entities in texts. Computational Intelligence, 2019, 35, 858-890.	2.1	8
66	On Computing Entity Relatedness in Wikipedia, with Applications. Knowledge-Based Systems, 2020, 188, 105051.	4.0	8
67	A Learned Approach to Design Compressed Rank/Select Data Structures. ACM Transactions on Algorithms, 2022, 18, 1-28.	0.9	8
68	An experimental study of priority queues in external memory. Journal of Experimental Algorithmics, 2000, 5, 17.	0.7	7
69	A data structure for a sequence of string accesses in external memory. ACM Transactions on Algorithms, 2007, 3, 1-23.	0.9	7
70	Data Structures: Time, I/Os, Entropy, Joules!. Lecture Notes in Computer Science, 2010, , 1-16.	1.0	7
71	Bicriteria data compression. , 2014, , .		6
72	Compressed Cache-Oblivious String B-Tree. ACM Transactions on Algorithms, 2016, 12, 1-17.	0.9	6

#	ARTICLE	IF	CITATIONS
73	Static and dynamic parallel computation of connected components. Information Processing Letters, 1994, 50, 63-68.	0.4	5
74	String Search in Coarse-Grained Parallel Computers. Algorithmica, 1999, 24, 177-194.	1.0	5
75	Improved Dynamic Text Indexing. Journal of Algorithms, 1999, 31, 291-319.	0.9	5
76	Compressed Indexes for String Searching in Labeled Graphs. , 2015, , .		5
77	Motif-Raptor: a cell type-specific and transcription factor centric approach for post-GWAS prioritization of causal regulators. Bioinformatics, 2021, 37, 2103-2111.	1.8	5
78	Early outcome detection for COVID-19 patients. Scientific Reports, 2021, 11, 18464.	1.6	5
79	A note on updating suffix tree labels. Lecture Notes in Computer Science, 1997, , 181-192.	1.0	5
80	Recognition by constructive neural algorithms. Pattern Recognition Letters, 1993, 14, 997-1007.	2.6	4
81	A Technique to Speed Up Parallel Fully Dynamic Algorithms for MST. Journal of Parallel and Distributed Computing, 1995, 31, 181-189.	2.7	4
82	On compact representations of All-Pairs-Shortest-Path-Distance matrices. Theoretical Computer Science, 2010, 411, 3293-3300.	0.5	4
83	On the weak prefix-search problem. Theoretical Computer Science, 2013, 483, 75-84.	0.5	4
84	Cache-Oblivious Comparison-Based Algorithms on Multisets. Lecture Notes in Computer Science, 2005, , 305-316.	1.0	4
85	On Optimally Partitioning a Text to Improve Its Compression. Lecture Notes in Computer Science, 2009, , 420-431.	1.0	4
86	SMAPH. ACM Transactions on Information Systems, 2019, 37, 1-42.	3.8	4
87	On the Parallel Dynamic Dictionary Matching problem: New results with applications. Lecture Notes in Computer Science, 1996, , 261-275.	1.0	4
88	THREE TECHNIQUES FOR PARALLEL MAINTENANCE OF A MINIMUM SPANNING TREE UNDER BATCH OF UPDATES. Parallel Processing Letters, 1996, 06, 213-222.	0.4	3
89	AN EREW PARM ALGORITHM FOR UPDATING MINIMUM SPANNING TREES. Parallel Processing Letters, 1999, 09, 111-122.	0.4	3
90	Bicriteria Data Compression. SIAM Journal on Computing, 2019, 48, 1603-1642.	0.8	2

#	ARTICLE	IF	CITATIONS
91	Tree Compression and Indexing. , 2008, , 964-968.		2
92	Suffix Tree Construction in Hierarchical Memory. , 2008, , 922-925.		2
93	PaTre: A Method for Paralogy Trees Construction. Journal of Computational Biology, 2003, 10, 791-802.	0.8	1
94	Information processing at work: On energy-aware algorithm design. , 2010, , .		1
95	Beyond the bag-of-words paradigm to enhance information retrieval applications. , 2011, , .		1
96	Indexing Compressed Text. , 2009, , 1442-1448.		1
97	Document Aboutness via Sophisticated Syntactic and Semantic Features. Lecture Notes in Computer Science, 2017, , 441-453.	1.0	1
98	On the Weak Prefix-Search Problem. Lecture Notes in Computer Science, 2011, , 261-272.	1.0	1
99	Web Search. , 2013, , 107-137.		1
100	Boosting Textual Compression. , 2016, , 228-232.		1
101	Recognition of hand-written rotated digits by neural networks. Machine Vision and Applications, 1995, 8, 351-357.	1.7	0
102	Burrows-Wheeler Transform. , 2008, , 112-116.		0
103	Guest Editorial: Selected Papers of European Symposium of Algorithms. Algorithmica, 2014, 70, 365-367.	1.0	0
104	Burrows-Wheeler Transform. , 2016, , 250-255.		0
105	Locality Filtering for Efficient Ride Sharing Platforms. IEEE Transactions on Intelligent Transportation Systems, 2022, 23, 7785-7804.	4.7	0
106	Experimenting SnakeT: A Hierarchical Clustering Engine for Web-Page Snippets. Lecture Notes in Computer Science, 2004, , 543-545.	1.0	0
107	Boosting Textual Compression. , 2008, , 97-101.		0
108	Distribution-Aware Compressed Full-Text Indexes. Lecture Notes in Computer Science, 2011, , 760-771.	1.0	0

#	ARTICLE	IF	CITATIONS
109	Burrows-Wheeler Transform. , 2014, , 1-7.		0
110	Compressing and Indexing Structured Text. , 2014, , 1-8.		0
111	Boosting Textual Compression. , 2015, , 1-7.		0
112	Suffix Tree Construction in Hierarchical Memory. , 2016, , 2149-2154.		0
113	Indexing Compressed Text. , 2017, , 1-8.		0
114	Text Compression. , 2018, , 4070-4072.		0
115	Indexing Compressed Text. , 2018, , 1861-1868.		0
116	Recognition of hand-written rotated digits by neural networks. Machine Vision and Applications, 1995, 8, 351-357.	1.7	0
117	BioTAGME: A Comprehensive Platform for Biological Knowledge Network Analysis. Frontiers in Genetics, 2022, 13, 855739.	1.1	0