Alexander M Rush

List of Publications by Year
in descending order

67 papers	7,559 citations

1478505

6 h-index

1372567

10 g-index

68 all docs	68 docs citations	68 times ranked	3245 citing authors

#	Article	IF	Citations
1	Transformers: State-of-the-Art Natural Language Processing. , 2020, , .		2,035
2	A Neural Attention Model for Abstractive Sentence Summarization. , 2015, , .		1,403
3	OpenNMT: Open-Source Toolkit for Neural Machine Translation. , 2017, , .		855
4	Abstractive Sentence Summarization with Attentive Recurrent Neural Networks. , 2016, , .		497
5	Bottom-Up Abstractive Summarization. , 2018, , .		364
6	Sequence-to-Sequence Learning as Beam-Search Optimization. , 2016, , .		254
7	Challenges in Data-to-Document Generation. , 2017, , .		254
8	Sequence-Level Knowledge Distillation. , 2016, , .		230
9	LSTMVis: A Tool for Visual Analysis of Hidden State Dynamics in Recurrent Neural Networks. IEEE Transactions on Visualization and Computer Graphics, 2018, 24, 667-676.	4.4	199
10	Seq2seq-Vis: A Visual Debugging Tool for Sequence-to-Sequence Models. IEEE Transactions on Visualization and Computer Graphics, 2019, 25, 353-363.	4.4	142
11	Learning Neural Templates for Text Generation. , 2018, , .		110
12	Commonsense Knowledge Mining from Pretrained Models. , 2019, , .		102
13	Learning Global Features for Coreference Resolution. , 2016, , .		93
14	Learning Anaphoricity and Antecedent Ranking Features for Coreference Resolution. , 2015, , .		79
15	GLTR: Statistical Detection and Visualization of Generated Text. , 2019, , .		71
16	Unsupervised Recurrent Neural Network Grammars. , 2019, , .		50
17	Compound Probabilistic Context-Free Grammars for Grammar Induction. , 2019, , .		46
18	Neural Linguistic Steganography. , 2019, , .		43

#	Article	IF	Citations
19	Datasets: A Community Library for Natural Language Processing. , 2021, , .		40
20	End-to-End Content and Plan Selection for Data-to-Text Generation. , 2018, , .		37
21	Parameter-Efficient Transfer Learning with Diff Pruning. , 2021, , .		33
22	How many data points is a prompt worth?. , 2021, , .		32
23	OpenNMT System Description for WNMT 2018: 800 words/sec on a single-core CPU. , 2018, , .		32
24	MASR: A Modular Accelerator for Sparse RNNs. , 2019, , .		30
25	Training for Diversity in Image Paragraph Captioning. , 2018, , .		29
26	Adapting Sequence Models for Sentence Correction. , 2017, , .		28
27	Visual Interaction with Deep Learning Models through Collaborative Semantic Inference. IEEE Transactions on Visualization and Computer Graphics, 2019, 26, 1-1.	4.4	26
28	EdgeBERT: Sentence-Level Energy Optimizations for Latency-Aware Multi-Task NLP Inference. , 2021, , .		26
29	On Adversarial Removal of Hypothesis-only Bias in Natural Language Inference. , 2019, , .		25
30	Donâ€™t Take the Premise for Granted: Mitigating Artifacts in Natural Language Inference. , 2019, , .		23
31	Torch-Struct: Deep Structured Prediction Library. , 2020, , .		22
32	Algorithm-Hardware Co-Design of Adaptive Floating-Point Encodings for Resilient Deep Learning Inference. , 2020, , .		22
33	Classical and quantum dynamics in the (non-Hermitian) Swanson oscillator. Journal of Physics A: Mathematical and Theoretical, 2015, 48, 055301.	2.1	21
34	Sequence-Level Mixed Sample Data Augmentation. , 2020, , .		21
35	9.8 A 25mm^2 SoC for IoT Devices with 18ms Noise-Robust Speech-to-Text Latency via Bayesian Speech Denoising and Attention-Based Sequence-to-Sequence DNN Speech Recognition in 16nm FinFET. , 2021, , .		19
36	Sentence-Level Grammatical Error Identification as Sequence-to-Sequence Correction. , 2016, , .		19

#	Article	IF	Citations
37	Block Pruning For Faster Transformers. , 2021, , .		19
38	GRIT: Generative Role-filler Transformers for Document-level Event Entity Extraction. , 2021, , .		18
39	Word Ordering Without Syntax. , 2016, , .		17
40	Coarse-to-Fine Attention Models for Document Summarization. , 2017, , .		17
41	An Embedding Model for Predicting Roll-Call Votes. , 2016, , .		15
42	Simple Unsupervised Summarization by Contextual Matching. , 2019, , .		14
43	Transforming Dependencies into Phrase Structures. , 2015, , .		13
44	Generating Abstractive Summaries with Finetuned Language Models. , 2019, , .		11
45	Adversarial Semantic Collisions. , 2020, , .		10
46	Scaling Hidden Markov Language Models. , 2020, , .		10
47	Posterior Control of Blackbox Generation. , 2020, , .		8
48	Entity Tracking Improves Cloze-style Reading Comprehension. , 2018, , .		8
49	Dual inductive link coil design for a neural recording system. , 2011, 2011, 6397-400.		7
50	LAN: A Materials Notation for Two-Dimensional Layered Assemblies. Journal of Chemical Information and Modeling, 2020, 60, 3457-3462.	5.4	7
51	Debugging Sequence-to-Sequence Models with Seq2Seq-Vis. , 2018, , .		7
52	A Constrained Viterbi Relaxation for Bidirectional Word Alignment. , 2014, , .		7
53	What is Learned in Visually Grounded Neural Syntax Acquisition. , 2020, , .		7
54	Semiclassical quantization for a bosonic atom-molecule conversion system. Physical Review A, 2015, 92, .	2.5	6

#	Article	IF	Citations
55	Classical-quantum correspondence in bosonic two-mode conversion systems: Polynomial algebras and Kummer shapes. Physical Review A, 2016, 93, .	2.5	6
56	Power and data for a wireless implanted neural recording system. , 2011, , .		5
57	Propagation of Gaussian Beams in the Presence of Gain and Loss. IEEE Journal of Selected Topics in Quantum Electronics, 2016, 22, 130-135.	2.9	5
58	GenNI: Human-AI Collaboration for Data-Backed Text Generation. IEEE Transactions on Visualization and Computer Graphics, 2022, 28, 1106-1116.	4.4	5
59	Template Filling with Generative Transformers. , 2021, , .		4
60	Low-Complexity Probing via Finding Subnetworks. , 2021, , .		4
61	Developmental Stage Classification of Embryos Using Two-Stream Neural Network with Linear-Chain Conditional Random Field. Lecture Notes in Computer Science, 2021, 12908, 363-372.	1.3	4
62	Electronic performance of a dual inductive link for a wireless neural recording implant. , 2011, 2011, 6348-51.		3
63	Improving Event Duration Prediction via Time-aware Pre-training. , 2020, , .		2
64	Antecedent Prediction Without a Pipeline. , 2016, , .		0
65	SM6: A 16nm System-on-Chip for Accurate and Noise-Robust Attention-Based NLP Applications : The 33[rd] Hot Chips Symposium â€" August 22-24, 2021. , 2021, , .		0
66	Sequence-to-Lattice Models for Fast Translation. , 2021, , .		0
67	Rationales for Sequential Predictions. , 2021, , .		0