Keeping Track of Latest Developments in AI with aipano.cse.ust.hk

The website aipano.cse.ust.hk provides a topic-based index to research papers published since 2000 at major AI and Machine Learning venues. It lets you keep track of the latest developments in AI and offers a fast way to find recent highly cited papers on various topics or published at a certain venue. Here are few screenshots of a session with the system to illustrate its functionalities.

Knowing the General Trends: At the front page, you can perform sort by "popularity: past year" to find the topics with the largest numbers of papers in the past year (2018):

If you sort the topics by "trend: 3 years", you will find the topics with the fastest increase in popularity the past three years:

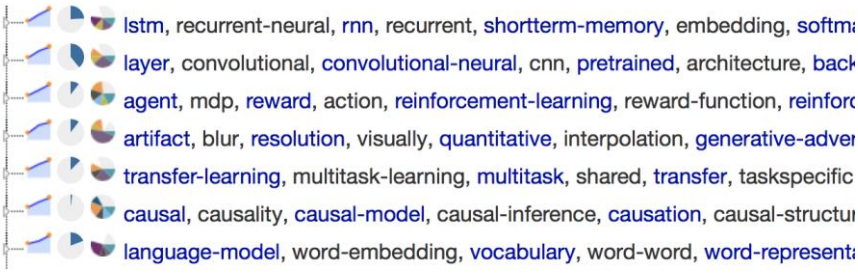

It is interesting to find the topic "causal, causality, causal-model, causal-inference, …" at the 6th position although the total number of papers on the topic is relatively small.

Keeping Track of the Latest Developments on a Topic: By clicking on a topic, you will see its trend curve and, more importantly, papers on the topic that you might want to add to your reading list. Here is the trend curve of the topic "layer, convolutional, convolutional-neural, cnn, pretrained, …":

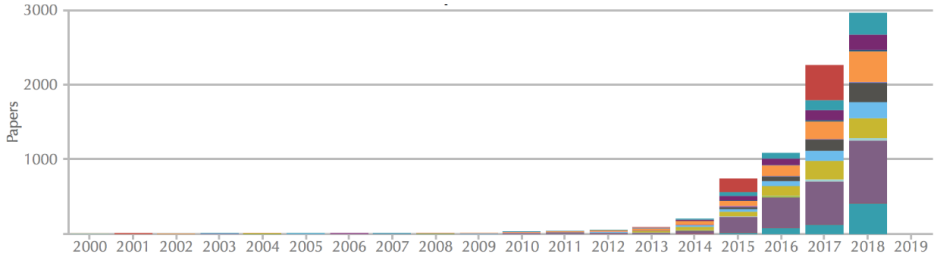

Nearly 3,000 papers were published on the topic in 2018! Here are the most cited ones, which are good candidates for a reading list:

CVPR	2018	Squeeze-and-Excitation Networks	425

CVPR	2018	Interpretable Convolutional Neural Networks	303
JMLR	2018	Quantized Neural Networks: Training Neural Networks with Low Precision Weights and Activations	266
CVPR	2018	Learning Transferable Architectures for Scalable Image Recognition	252
ICML	2018	Obfuscated Gradients Give a False Sense of Security: Circumventing Defenses to Adversarial Examples	176

It is possible to narrow down the candidates by requiring key words. Here are the most cited papers in 2018 that contain the key word “adversarial” in title.

ICML	2018	Obfuscated Gradients Give a False Sense of Security: Circumventing Defenses to Adversarial Examples	176
ICML	2018	CyCADA: Cycle-Consistent Adversarial Domain Adaptation	111
ICLR	2018	Towards Deep Learning Models Resistant to Adversarial Attacks	77
ICML	2018	Provable Defenses against Adversarial Examples via the Convex Outer Adversarial Polytope	68
CVPR	2018	AttnGAN: Fine-Grained Text to Image Generation With Attentional Generative Adversarial Networks	45

If you are interested in reinforcement learning, you might want to take a look at the topic “agent, mdp, reward, action, reinforcement-learning, …”:

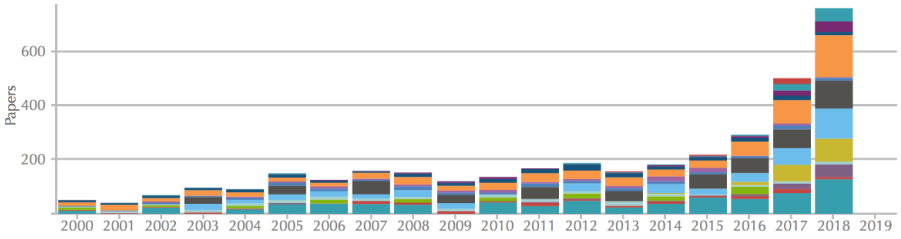

Here are the most cited papers on the topic published in 2018:

ACL	2018	Deep Reinforcement Learning for NLP	295
AAAI	2018	Deep Reinforcement Learning That Matters	133
AAAI	2018	Rainbow: Combining Improvements in Deep Reinforcement Learning	104
AAAI	2018	Counterfactual Multi-Agent Policy Gradients	79
AAAI	2018	Emergence of Grounded Compositional Language in Multi-Agent Populations	76

An Investigation into Work on Causality: Here is the trend curve of the topic “causal, causality, causal-model, causal-inference, …”:

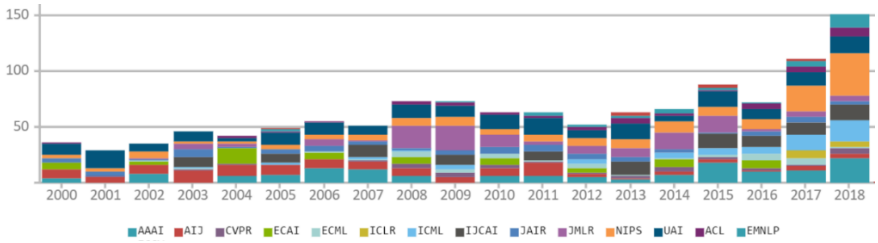

Here are the most cited papers on the topic in 2018, 2017 and since 2000:

Series	Year	Title	Cited
AAAI	2018	Measuring Conditional Independence by Independent Residuals: Theoretical Results and Application in Causal Discovery	201

AAAI	2018	Fair Inference on Outcomes	31
JMLR	2018	Learning Certifiably Optimal Rule Lists for Categorical Data	29
ICML	2018	Neural Relational Inference for Interacting Systems	20
JMLR	2018	Uncovering Causality from Multivariate Hawkes Integrated Cumulants	15

ICLR	2017	A Compositional Object-Based Approach to Learning Physical Dynamics	40
EMNLP	2017	A causal framework for explaining the predictions of black-box sequence-to-sequence models	29
ICLR	2017	Revisiting Classifier Two-Sample Tests	25
NIPS	2017	Avoiding Discrimination through Causal Reasoning	24
NIPS	2017	Counterfactual Fairness	22

UAI	2001	Direct and Indirect Effects	889
UAI	2001	Causes and Explanations: A Structural-Model Approach --- Part 1: Causes	648
JMLR	2006	A Linear Non-Gaussian Acyclic Model for Causal Discovery	585
UAI	2007	Causal Bounds and Instruments	567
JMLR	2007	Estimating High-Dimensional Directed Acyclic Graphs with the PC-Algorithm	500

Appearng Smart in Front of Students: A student is interested in zero-shot learning. I searched for "zero-shot" within several relevant topics and found the following papers:

From topic "transfer-learning, multitask-learning, multitask, …":

ICML	2015	An embarrassingly simple approach to zero-shot learning	287
CVPR	2016	Synthesized Classifiers for Zero-Shot Learning	90
AAAI	2016	Transductive Zero-Shot Recognition via Shared Model Space Learning	42
IJCAI	2016	Using Task Features for Zero-Shot Knowledge Transfer in Lifelong Learning	32
CVPR	2017	Learning a Deep Embedding Model for Zero-Shot Learning	28

From topic "lstm, recurrent-neural, rnn, recurrent, …":

CVPR	2016	Multi-cue Zero-Shot Learning with Strong Supervision	39
CVPR	2017	Learning a Deep Embedding Model for Zero-Shot Learning	28
ICML	2017	Zero-Shot Task Generalization with Multi-Task Deep Reinforcement Learning	16
EMNLP	2017	Zero-Shot Activity Recognition with Verb Attribute Induction	10
CVPR	2018	Multi-Label Zero-Shot Learning With Structured Knowledge Graphs	6

From topic "layer, convolutional, convolutional-neural, cnn, …":

NIPS	2013	Zero-Shot Learning Through Cross-Modal Transfer	532
ICLR	2014	Zero-Shot Learning by Convex Combination of Semantic Embeddings	222
ACL	2015	Hubness and Pollution: Delving into Cross-Space Mapping for Zero-Shot Learning	67
IJCAI	2015	Semantic Concept Discovery for Large-Scale Zero-Shot Event Detection	56
CVPR	2016	Multi-cue Zero-Shot Learning with Strong Supervision	39

Other Uses: Do your ever wonder how AIJ and JMLR differ in terms of topics? Here is what I

found on aipano for the past 5 years.

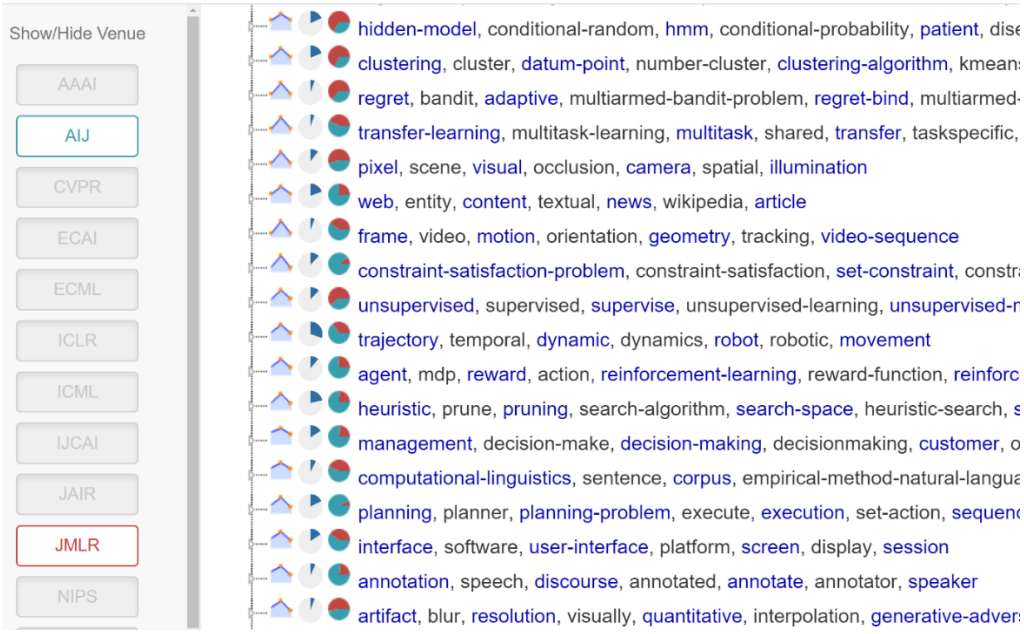

All papers: A special "All papers" topic is placed at top of the web page. It let's you see the trend of the entire AI and Machine Learning field since Year 2000 and the most influence papers.

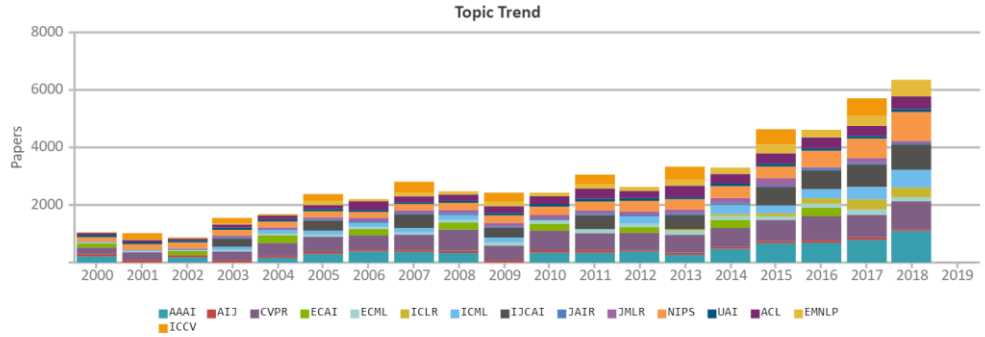

Show 10 entries Search: Search paper by word, year, and/or venue

Series	Year	Title	Cited
NIPS	2012	ImageNet Classification with Deep Convolutional Neural Networks	34948
NIPS	2001	Latent Dirichlet Allocation	25783
CVPR	2005	Histograms of Oriented Gradients for Human Detection	22410
CVPR	2001	Rapid Object Detection using a Boosted Cascade of Simple Features	16962
ICCV	2001	Robust Real-Time Face Detection	16562
JMLR	2011	Scikit-learn: Machine Learning in Python	14173
JMLR	2003	An Introduction to Variable and Feature Selection	12325
NIPS	2013	Distributed Representations of Words and Phrases and their Compositionality	11178
ICLR	2015	Very Deep Convolutional Networks for Large-Scale Image Recognition	10831
JMLR	2014	Dropout: A Simple Way to Prevent Neural Networks from Overfitting	10045

What is happening at a conference? The "All papers" topic also lets you track the latest

developments at a particular venue:

Show 10 entries Search: NIPS 2018

Series	Year	Title	Cited
NIPS	2018	Are GANs Created Equal? A Large-Scale Study	77
NIPS	2018	Natasha 2: Faster Non-Convex Optimization Than SGD	56
NIPS	2018	Visualizing the Loss Landscape of Neural Nets	55
NIPS	2018	Glow: Generative Flow with Invertible 1x1 Convolutions	42
NIPS	2018	Improving Exploration in Evolution Strategies for Deep Reinforcement Learning via a Population of Novelty-Seeking Agents	35
NIPS	2018	Isolating Sources of Disentanglement in Variational Autoencoders	33
NIPS	2018	Adversarially Robust Generalization Requires More Data	30

Comparing Conferences: A comparison of the most cited papers from ICLR and NIPS indicates that ICLR is catching up with NIPS in terms of impact:

ICLR	2018	Ensemble Adversarial Training: Attacks and Defenses	77
ICLR	2018	Towards Deep Learning Models Resistant to Adversarial Attacks	77
ICLR	2018	Progressive Growing of GANs for Improved Quality, Stability, and Variation	74
ICLR	2018	A Deep Reinforced Model for Abstractive Summarization	39
ICLR	2018	Many Paths to Equilibrium: GANs Do Not Need to Decrease a Divergence At Every Step	34
ICLR	2018	On the State of the Art of Evaluation in Neural Language Models	30
ICLR	2018	On the Information Bottleneck Theory of Deep Learning	29
ICLR	2018	Parameter Space Noise for Exploration	28
ICLR	2018	Countering Adversarial Images using Input Transformations	26
ICLR	2017	Understanding deep learning requires rethinking generalization	299
ICLR	2017	Energy-based Generative Adversarial Networks	226
ICLR	2017	Bidirectional Attention Flow for Machine Comprehension	225
ICLR	2017	Adversarial examples in the physical world	211
ICLR	2017	Towards Principled Methods for Training Generative Adversarial Networks	211
ICLR	2017	Adversarial examples in the physical world	211
ICLR	2017	Adversarial Feature Learning	206
ICLR	2017	Adversarially Learned Inference	196
ICLR	2017	Dynamic Coattention Networks For Question Answering	179
ICLR	2016	Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks	1440
ICLR	2016	Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning	774
ICLR	2016	Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding	757
ICLR	2016	Multi-Scale Context Aggregation by Dilated Convolutions	689
ICLR	2016	Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs)	629

NIPS	2018	Are GANs Created Equal? A Large-Scale Study	77
NIPS	2018	Natasha 2: Faster Non-Convex Optimization Than SGD	56
NIPS	2018	Visualizing the Loss Landscape of Neural Nets	55
NIPS	2018	Glow: Generative Flow with Invertible 1x1 Convolutions	42
NIPS	2018	Improving Exploration in Evolution Strategies for Deep Reinforcement Learning via a Population of Novelty-Seeking Agents	35
NIPS	2018	Isolating Sources of Disentanglement in Variational Autoencoders	33
NIPS	2018	Adversarially Robust Generalization Requires More Data	30
NIPS	2018	How Does Batch Normalization Help Optimization?	30
NIPS	2018	Realistic Evaluation of Deep Semi-Supervised Learning Algorithms	25
NIPS	2017	Improved Training of Wasserstein GANs	308
NIPS	2017	Dual Path Networks	220
NIPS	2017	Attention is All you Need	217
NIPS	2017	Deep Sets	111
NIPS	2017	Self-Normalizing Neural Networks	103
NIPS	2017	Dynamic Routing Between Capsules	83
NIPS	2017	What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?	80
NIPS	2017	Unsupervised Image-to-Image Translation Networks	80
NIPS	2017	Learning Disentangled Representations with Semi-Supervised Deep Generative Models	68
NIPS	2017	Prototypical Networks for Few-shot Learning	65
NIPS	2016	Improved Techniques for Training GANs	1481
NIPS	2016	R-FCN: Object Detection via Region-based Fully Convolutional Networks	988
NIPS	2016	InfoGAN: Interpretable Representation Learning by Information Maximizing Generative Adversarial Nets	821
NIPS	2016	Binarized Neural Networks	617
NIPS	2016	Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering	551
NIPS	2016	Matching Networks for One Shot Learning	478

In summary, aipano provides a panoramic view of the AI literature. It lets you keep track of the latest developments in AI and offers a fast way to find recent highly cited papers on various topics or at a particular venue.