



# NLP 2018 Highlights

By Elvis Saravia



# Table of Contents

Introduction .....	4
Research .....	5
Reinforcement Learning .....	5
Sentiment Analysis and Related Topics .....	7
AI Ethics and Security .....	9
Clinical NLP and ML .....	12
Computer Vision.....	15
Deep Learning and Optimization.....	17
Transfer Learning for NLP .....	19
AI Generalization .....	20
Explainability and Interpretability .....	21
AI Creativity.....	22
General NLP.....	26
Trends and Overviews .....	31
AI Annual Reports.....	31
Overviews .....	31
Resources.....	36
Dataset Releases .....	36
Learning Resources.....	39
Books .....	43
Newsletters and Podcasts .....	44
Open Source .....	45
Industry.....	53
Conversational AI .....	53
Health and Lifestyle .....	54
Machine Learning Tools.....	56
Financing, Advertising, and Law .....	58
General NLP and ML Stories.....	58
Miscellaneous .....	60
People and Society .....	60

Opinions..... 62  
Research Guides..... 65  
NLP and ML Guides..... 66

## Introduction

It has been a big year for the field of natural language processing (NLP) and for machine learning as a whole. There have been many trends and breaking stories, with state-of-the-art results and new interesting research directions emerging. We owe all of this progress to the brilliant researchers around the world and the millions of developers devoting their full time to improve tools that make it easier for everyone to learn and progress the field. We have witnessed a rise in transfer learning and other niche areas of research such as AI Ethics.

In this report, I will provide a summary of all the biggest NLP stories of the year (2018) coming from both academia and the industry. I hope this report serves as a guide for the researcher and developer that wishes to start learning about this field and also for the expert who wishes to review some of the latest advancements. The selected topics are based purely on personal observation, so it is highly possible that I missed other important stories. In fact, the stories shared here are an extended compilation of the NLP Newsletter published weekly on the [dair.ai](https://dair.ai) publication. I have made an effort to categorize the stories in the best possible way so that the report could benefit as many readers as possible. Please note that the report is meant to be non-technical for the purpose of reaching as many diverse readers as possible. The stories are mainly categorized by the following key topics: *AI ethics, research publications, trends, education, resources, industry*, and much more.

# Research

## Reinforcement Learning

*Many researchers are now considering reinforcement learning as the next frontier of machine learning. While this remains to be seen, it is important to familiarize with some of the advancements that have been occurring in this field in terms of generalization capabilities.*

- David Ha presents "[World Models](#)", a model which aims to answer the big question of whether an agent is able to learn inside its own dream environment. This is a very interesting work that leverages reinforcement learning to learn a policy that can solve a designated task by relying on features extracted from the world model. [Here](#) is a Python implementation of "World Models" using PyTorch.
- OpenAI [develops](#) a human-like robot hand to manipulate physical objects with dexterity using reinforcement learning algorithms.
- DeepMind's published a [paper](#) in Nature which discusses how grid-like representations in artificial agents enabled shortcuts in virtual environments through vector-based navigation.
- Google researchers have developed a state-of-the-art [method](#) called MnasNet, which is an automated neural network architecture search approach for designing mobile ML models using reinforcement learning.
- [OpenAI Five](#) is the agent trained on reinforcement learning that's defeating amateur human teams in the complex game of Dota 2.
- DeepMind develops a new technique called [PopArt](#) which enables a reinforcement learning agent to perform accurately in multitasking environments.
- [TextWorld](#) is an environment inspired by text-based games to create a learning environment for training reinforcement learning agents.

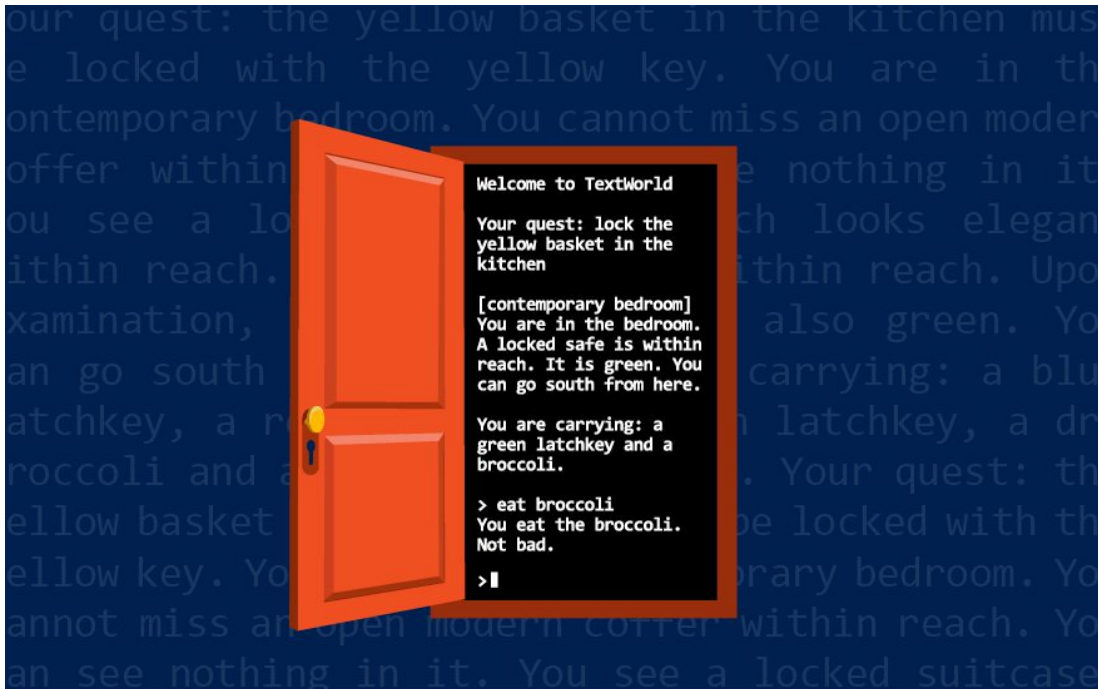


Figure: TextWorld environment ([source](#))

- Samsung's [CoZNet](#) is a deep reinforcement learning algorithm which achieved state-of-the-art results on two popular NLP challenges.
- AREL is an [approach](#) that uses adversarial reward learning for storytelling and addressing limitations of the metrics used for story evaluation (code and paper included).
- [Metacar](#) is a reinforcement learning environment for self-driving cars built on the browser using Tensorflow.js.
- Disney [research](#) reveals a method that automates a learning environment for developing control policies directly on the hardware of a modular legged robot.
- OpenAI releases [Gym Retro](#), a platform that provides 1000+ games for conducting reinforcement learning research.

## Sentiment Analysis and Related Topics

*It is clear that machine learning systems need more than logical understanding to be able to fully interact with humans at an emotional level. In this section, I summarize some of the recent advancements in the important field of sentiment analysis.*

- A recent [paper](#) proposed "*Hierarchical CVAE for Fine-Grained Hate Speech Classification*", a method for understanding hate speech across 40 hate groups and 13 different hate categories.
- A recent [paper](#) discusses how to obtain state-of-the-art results on a sentiment analysis task with simple Naive Bayes and Support Vector Machine variants. Although most of the work in the field of sentiment analysis today is mostly based on neural approaches, it's always important to pay attention to the model selection and features for whatever NLP tasks you are dealing with. This paper shows some of the different aspects of a model that you could pay attention to.
- On [combining domain knowledge and deep](#) learning for conducting sentiment classification.
- [Sarcasm detection](#) using deep convolutional neural networks. Another [paper](#) aims to understand sarcastic posts using attention models and deep neural networks.
- This [paper](#) discusses a simple method for performing style transfer on sentiment aspects.
- A new [study](#) proposed an approach for sarcasm detection using a multi-faceted approach: eye-tracking features, NLP, and deep learning algorithms. This work also emphasizes the importance of leveraging multimodal learning for natural language tasks. We expect to see more of these type of studies in the field of NLP and machine learning in the years that follow.
- A recent [paper](#) published in NAACL 2018, proposed a method that is able to model the flow of emotions over a book using simple RNNs.

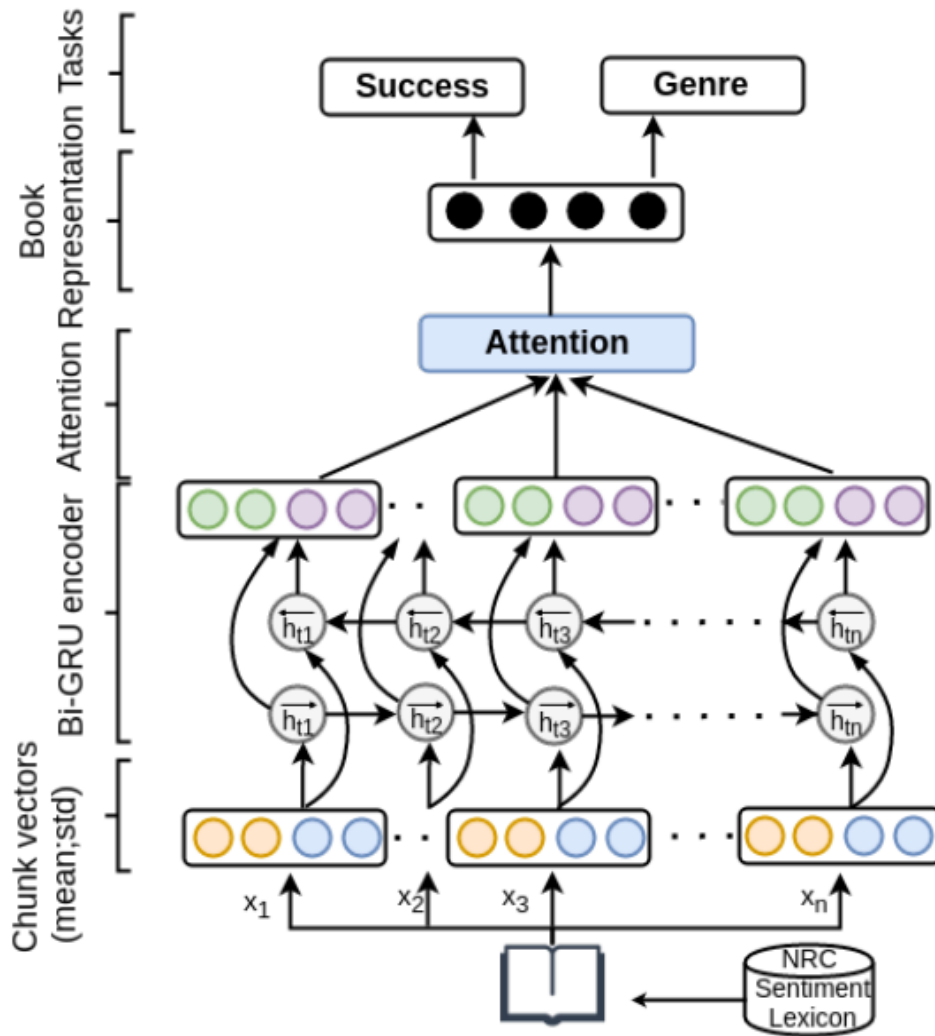


Figure: Multitask emotion flow model ([source](#))

- Researchers develop a method they called "[emotional chatting machine](#)" which consists of a chatbot that produces factually coherent answers while also being able to incorporate emotions such as sadness and disgust in a conversation.
- Lei Zhang et al. (2018) published a [survey paper](#) that provides a comprehensive overview of deep learning methods used for sentiment analysis.
- This interesting [paper](#) presents a quantitative analysis of the use of skin tone modifiers on emoji on Twitter.





Figure: Emoji variation ([source](#))

- A [bidirectional-asynchronous framework](#) is leveraged to generate meaningful emotional replies in conversations (something traditional language models fail to achieve due to generic replies).
- This [study](#) investigates emotion recognition in context using a computer vision method.

## AI Ethics and Security

*Once considered a niche field, this year saw AI ethics and security at the forefront of machine learning research. In this section, I summarize all the recent developments in the area of AI ethics and security.*

- A group of researchers from the University of California develops a novel [method](#) that preserves gender information in word vectors while compelling other dimensions to be free of gender influence. Their paper is called "*Learning Gender-Neutral Word Embeddings*" and aims to reduce gender bias present in natural language datasets.
- BBC recently hosted a [program](#) where AI ethics and morality are discussed with AI experts from all over the world. Some contributors to the program include Mustafa Suleyman, Nick Bostrom, and other students.

- The Google AI team is investing [efforts](#) to reduce bias in their language translation tools by incorporating gender-specific translations for gender-neutral words such as “doctor”, which is more critical in languages like Turkish.
- MIT [announces](#) \$1 billion initiative to help teach students and AI researchers about the social and ethical implications of AI.
- Want to keep track of the misuses of AI in society? This [repository](#) does a decent job of curating an updated list of potentially harmful AI systems appearing in the wild. Some notable examples include [Deep Fakes](#) and [Fake News Bots](#).
- OpenAI proposes an AI safety technique called [iterated amplification](#). The idea is that we can specify complicated behaviors and goals and be able to generate training signals that go beyond the traditional rewards and labels used in supervised and reinforcement learning, respectively. The interesting part is that humans are able to coordinate with an AI system in a framework that is able to scale up to higher level tasks based on a sample of smaller subtasks that are based on human judgment.
- Vered Shwartz releases a blog [post](#) where she discusses the importance of ethical machine learning and what questions we should be asking when building ML systems that make important decisions in the real world.
- AAIL releases [results](#) for the first ever research track on AI Ethics and Society.
- Google researchers conducted a [study](#) to identify several cases of bias brought about by the use of standard word embeddings when training machine learning models.

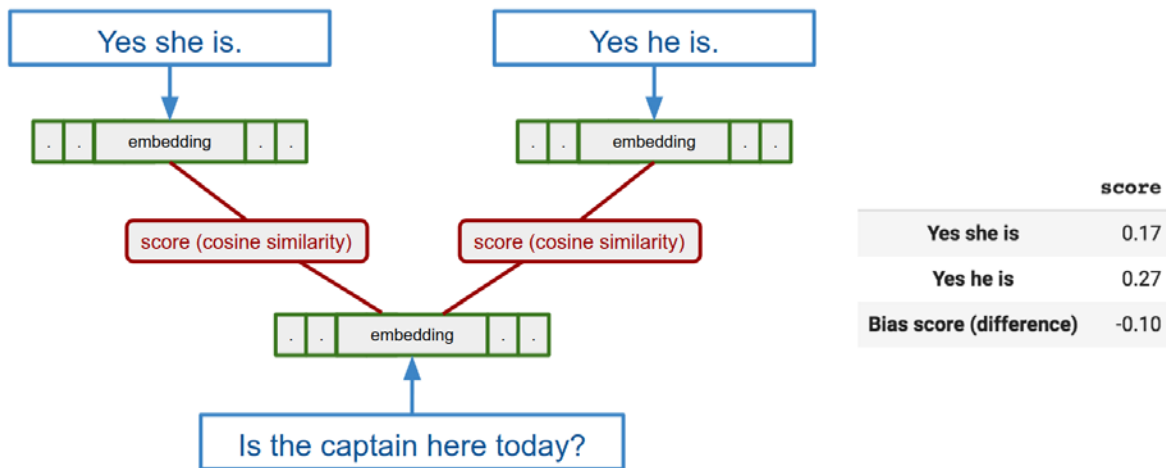


Figure: Gender bias ([source](#))

- Google [proposes](#) its new AI principles and practices.
- Irene Chen gives a [recap](#) of the important topics discussed at the Machine Learning for Health Care (MLHC) conference—from privacy to model robustness to clinical notes understanding.
- Rachel Thomas gives an [overview](#) of strategies to analyze and prevent bias in machine learning research.
- [Mirjalili et al. \(2018\)](#) propose semi-adversarial networks to enhance privacy by generating gender-neutral face images (code provided).
- [Here](#) is a 60-minute self-study training module on fairness released by Google as part of an initiative to help students get up-to-date on topics related to ethics in AI and fairness in machine learning.
- MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL) has launched a research-industry [collaboration](#) focused on building financial technologies to create new business models and improve security, among other goals.
- Sara Hooker [discusses](#) issues with the "data for good" movement.

- This [post](#) discusses how to achieve fairness in machine learning leveraging adversarial networks.
- NLP is already being used for clinical documentation, but what are the safety concerns for patients? Read more [here](#).

## Clinical NLP and ML

*While still a very young area of research, clinical NLP has become one of the most popular areas of research due to the social benefits of applications that leverage NLP techniques. Below I summarize the latest research and stories related to clinical NLP.*

- A [paper](#) presents an NLP-method to combine disparate resources and acquire accurate information about health providers.
- Radiology has recently seen significant gains in the use of deep learning methods for things like action recognition and automatic detection of brain injuries or skull fractures. AI researcher, Sasank Chilamkurthy, dives deep into the challenges faced when dealing with Head CT scans and medical imaging, and what creative ways can be used to address these challenges. As reported in this [blog post](#), the main challenge seems to be in the processing and preparation of the medical images as it requires different processing techniques compared to the more common computer vision or natural language datasets and tasks.
- Deep learning models were [used](#) to detect critical findings in head CT scans.
- Google has [developed](#) an algorithm to detect the spread of breast cancer.




Figure: Identifying tumor region ([source](#))

- We have seen the success of integrating different NLP techniques in designing conversational agents and recommendation systems, however, one of the promising areas where NLP will be heavily used in the future is in clinical informatics research. This [review paper](#) looks at the different NLP methods used in clinical research and the challenges involved in evaluating them.
- [James Zou et al. \(2018\)](#) released a tutorial and guide on how to apply deep learning for genomics via a Google Colab notebook.
- NLP researchers from National Tsing Hua University proposed a [novel method](#) for bipolar disorder prediction on social media based on an approach that leverages time-based features.
- A [syllabus](#) on clinical linguistics.
- There were many stories involving clinical applications that leverage artificial intelligence. In the context of this field, one of the questions that normally arise is “can AI beat doctors on a clinical exam?” A [piece](#) posted by Forbes covers more about his topic.

- Learn about how NLP can be used to [leverage](#) and unlock the unstructured healthcare datasets.
- A recent [study](#) conducted by Dr. Fei Fei Li and team proposes a machine learning model that tracks your face and voice features and is able to predict the severity of depression.
- Google AI [announces](#) improvements to their deep learning models used for diabetic retinopathy, one of the fastest growing causes of vision loss. The improvements include efforts to improve explainability and applicability in clinical settings.

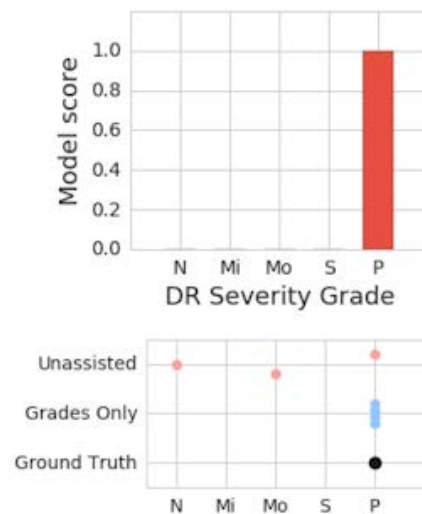


Figure: Model predictions ([source](#))

- Google AI opens new object detection competition which includes a massive training [dataset](#).
- Exciting [paper](#) on applying word embeddings to a massive source of multimodal medical data.
- This exciting [research](#) focuses on using deep learning methods to detect linguistic cues of Alzheimer's disease patients.



- An impressive [study](#) on how to collect high-quality data through search queries in developing nations, which may have some serious health benefits for society.
- [Detecting](#) social network mental disorders, such as internet overload and net compulsion, through a tensor decomposition approach.

## Computer Vision

*Even though computer vision is very different from NLP research, we rely on very similar methods to advance NLP systems. Computer vision research offers many inspirations for NLP research and in fact there have been efforts to bridge both domains to implement more robust systems. Learn more about these efforts from the highlights below.*

- Researchers at MIT [develop](#) an AI system that is able to achieve human-like reasoning to evaluate images. The system relies on an “attention mask” which enables a visual understanding of how the model reasons in several subtasks.
- MIT’s CSAIL team has developed a [system](#) named Dense Object Nets (DON) which is able to generate 3D visualizations and descriptions for objects it has never seen before.
- Researchers from Facebook Artificial Intelligence Research group developed [Detectron](#), a system that implements state-of-the-art object detection algorithms, including Mask R-CNN. The system was developed in Python and is powered by the Caffe2 deep learning framework.


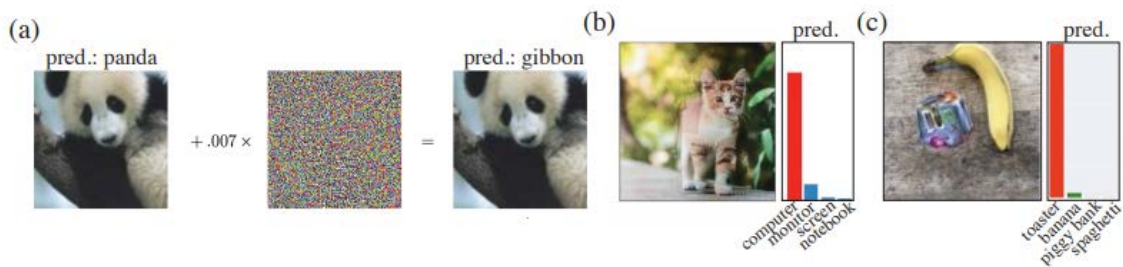


Figure: Example Mask R-CNN output ([source](#))

- One Pixel Attack is a Keras [implementation](#) of the paper called "One-pixel attack for fooling deep neural networks", which aims to investigate what would occur if one pixel in an image was modified by an attacker.
- An exciting [paper](#) shows how to conduct large-scale visual relationship understanding by mapping features from two modalities into a shared space.
- [On how to](#) learn joint multimodal embedding space with unpaired text and video data.
- Ian Goodfellow and others publish work on "[Adversarial Examples that Fool both Human and Computer Vision](#)". The study reports that adversarial examples which strongly transfer across computer vision models influence the classifications made by time-limited human observers.



- An interesting [paper](#) on generating images from scene graphs using graph convolution and GANs.
- Another interesting [research](#) by AI2 on generating "Scripts to Compositions to Videos".
- An interesting [paper](#) that proposes a technique which teaches machines the ability to learn textual-visual relationships.



## Deep Learning and Optimization

*Below I summarize some of the recent developments in deep learning that could have implications for NLP research.*

- DeepMind, in collaboration with Harvard professor Wouter Kool, releases [paper](#) investigating how human-decision makers deploy mental effort and how these insights can give way to opportunities and progress in recent artificial intelligence research.
- A paper introduces the concept of "[Group Normalization](#)", an effective alternative to batch normalization, considered a very important technique in deep learning.
- [Spherical CNNs](#) is a novel method for building convolutional neural networks.
- BAIR releases [article](#) discussing the trade-offs between recurrent neural networks and feed-forward neural networks when used to solve various tasks.
- A Facebook AI research team has developed a [new technique](#) that optimizes floating point math which makes AI models run 16 percent more efficient than when using standardized int8/32 math. All this means is that with such technique it will be possible to improve the speed of AI model training and simplify how models are quantized and deployed to production.
- Nature releases a [paper](#) describing a deep learning approach to predict earthquake aftershock locations. The model is also useful to understand the underlying physics behind the phenomena.
- DeepMind researcher develops an [approach](#) to enhance neural networks with Neural Arithmetic Logic Units (NALU) which are able to track time, perform arithmetic over images of numbers, count objects in images, among other capabilities.
- [DARTS](#) is a differentiable architecture search algorithm to efficiently design high-performance convolutional architectures for image classification.

- A paper called "[An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling](#)" discusses the differences between CNNs and RNNs for sequence modeling.
- Learn about [how](#) graph neural nets can help to infer latent relational structures and model multi-agent and physical dynamics.
- Google AI research team releases [paper](#) on how an improved version of RNN can be used to improve the accuracy of automatic interpretation of connectomics data.
- Distill releases new research called "[Feature-wise transformations](#)" to conduct context-based processing, which is a technique to process one source of information in the context of another. This method learns a task representation that captures the relationship between multiple sources of information, even in complex settings.




Figure: Feature-wise transformations ([source](#))

- Do you think it's possible to learn accurate cognitive models with no data and little to no human knowledge engineering? If you are curious you will find this [work](#) interesting.
- This [paper](#) goes into great details and investigates in depth why batch normalization works for training effective deep neural networks.
- This [paper](#) revisits how to make better choices with batch training in neural networks.

- This [paper](#) discusses how to properly conduct evaluations on deep semi-supervised learning algorithms.

## Transfer Learning for NLP

*Perhaps one of the biggest advancements in the field of NLP comes from the implementation of more efficient and effective transfer learning methods. Below I summarize important highlights in this research space.*

- So far this year we have witnessed many breakthroughs in the field of natural language processing, starting from models that learn deep contextualized features ([ELMo](#)) to models that leverage fine-tuned language modeling techniques to efficiently perform downstream tasks such as sentiment analysis ([ULMFit](#)). More recently, a new system has been proposed by Google AI Research team, called [BERT](#), which basically tweaks the way *contextualized features are learned in the Bi-LSTM layers* of a language representation model. BERT achieves state-of-the-art results on multiple tasks.
- Jeremy Howard and Sebastian Ruder proposed a transfer learning method called [Universal Language Model Fine-tuning for Text Classification](#) (ULMFit) which aims to introduce novel techniques for fine-tuning a language model. It is well known that training language models is a difficult and time-consuming task which the authors address, while at the same time achieving state-of-the-art results in tasks such as topic classification and question classification. Another interesting and important feature of ULMFit is that it also works for cases where there is a shortage of data, which is important when dealing with low-resource languages.
- William Wang and colleagues release [code and paper](#) for tackling a cross-lingual dialog state tracking problem using a simple but very efficient framework. The framework is based on a transfer learning technique that automatically generates semantic annotations in different languages and is able to track user's beliefs.
- FAIR releases a new method called [GLoMo](#) that leverages latent graph-based representations for enabling transfer learning to many downstream tasks.

- Language models are known to be inefficient in various aspects especially as it relates to efficiency. A recent [paper](#) introduces a novel theoretical framework that facilitates effective learning in language modeling while tying together word vectors and the output projection matrices.
- OpenAI proposes a transformer-based language mode called [OpenAI GPT](#) that is useful for a wide variety of NLP tasks.
- Another [paper](#) on transfer learning about how to uncover the cluster and relationships (disentangle) between computer vision tasks.

## AI Generalization

*While the goal of general AI seems far away, there were still many efforts this year to improve systems that show generalizable capabilities. Below I summarize the more interesting stories related to generalizable AI that we could borrow inspiration from.*

- Abilities tied to human intelligence such as abstract reasoning and planning require the ability to convert experience into concepts. An OpenAI research team has released a [technique](#), based on Boltzmann machines, that enables agents to learn and extract concepts such as "near" and "above" from specific tasks and then use these generated concepts to solve other tasks in various domains. For instance, they experimented with a 2D particle environment to solve tasks on a 3-dimensional physics-based robot. The cross-domain transfer is interesting and could allow for further analysis of concepts and language understanding.
- [Hyperbolic attention networks](#) that can match the complexity of data with hierarchical and power-law structure (useful for generalizing neural machine translation tasks).
- DeepMind's [paper](#) discusses how to measure abstract reasoning in neural networks. The problem of understanding the nature of generalization is an aspect of general-purpose learning systems that has long been a challenge for machine learning researchers.
- A recent [work](#) compares the robustness of humans and current convolutional deep neural networks (DNNs) on object recognition under twelve different types

of image degradations. The authors mainly test for generalization capabilities and how weaknesses observed in DNNs can be systematically addressed using a lifelong machine learning approach.




Figure: Distortion levels ([source](#))

→ In order to train machines to make better decisions and assertions that humans can understand, we must be able to teach them to reason similar to how humans reason on different experiences and interactions with their environment. In order to teach machines how to reason, there have been many advancements such as visual question answering (QA) and understanding visual relationships that provide clues on how to improve deep learning systems through embedding knowledge. In fact, more recent methods such as one-shot learning can be used to effectively reason based on past knowledge, but they are still limited in terms of generalizability and learning from rare events. One possible solution is to make use of augmented memory so that models can learn more efficiently and be able to reason faster. Read more on this type of systems [here](#).

## Explainability and Interpretability

*While it is not clear what the goals of explainable and interpretable systems are, there were many efforts to demystify some of the learning architectures used to build NLP systems.*

- This [paper](#) discusses when and why word embeddings are good for neural machine translation.
- Researchers propose a [method](#) that aims to better understand the predictions of an LSTM through an interpretation algorithm. They focus on the task of

sentiment classification and show that their method can identify words and phrases of contrasting sentiment and this contributes to the final prediction of the LSTM model. This is one of the first papers in NLP that I encountered addressing the problem of explainable AI.

- Learn about the latest [trends](#) in interpretable machine learning research.
- [Analyzing the behavior](#) of visual questions answering models to identify strengths and weakness.
- On [evaluating interpretability](#) methods that estimate feature importance in deep neural networks.
- Remove and Retrain ([ROAR](#)) is a benchmark that can be used to evaluate the accuracy of interpretability methods of deep neural networks.
- This [paper](#) presents a method that forces neural networks to behave as if they were trained on a completely different task. The method is called Adversarial Reprogramming of Neural Networks (ARNN).
- This [paper](#) shows a comprehensive comparison between recurrent vs. non-recurrent models for several NLP tasks. Aims to show that recurrent architectures are important for modeling hierarchical structure which is important in language.
- Published in ACL 2018, this [paper](#) aims to understand what exactly is being captured with sentence embeddings by proposing 10 probing tasks designed to capture simple linguistic features of sentences.
- This [paper](#) provides a deeper exploration of how deep learning architectures model the compositionality of text sequences.

## AI Creativity

*This year saw more creative machine learning applications than novel methods to learn from data. Below I list some of the more interesting research projects for this year.*

- Several research institutions and startups, in collaboration with OpenAI, are [building](#) an AI system that can learn curiosity without any guidance by only feeding it TV shows.
- An interesting [paper](#) that introduces a model for representing snippets of code as continuous distributed vectors (code2vec).

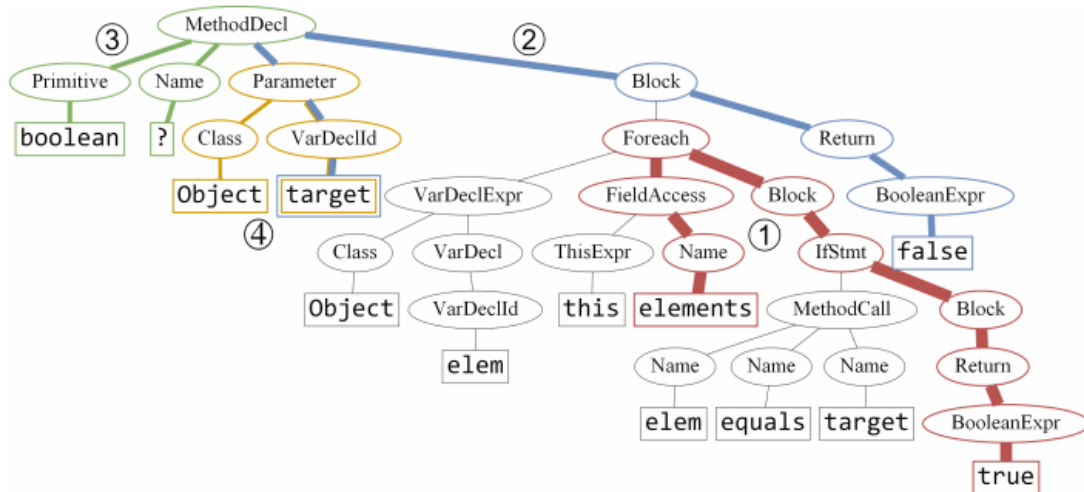


Figure: Attended paths ([source](#))

- Kakao Brain NLP releases a visualization tool called WordWeb which aims to visualize word vectors in many different languages. They use techniques such as fastText and t-SNE to develop the tool which is accessible [online](#).
- [Learn more](#) about bias and variance in machine learning with very neat visualizations.
- IBM releases a [tool](#) for visually debugging sequence-to-sequence models.
- A visual [guide](#) showing how decision trees work.
- Anna Huang builds the [Music Transformer](#), a music generation tool built which is able to generate harmonious, minute-long pieces based on a newly proposed self-attention mechanism

- [CariGANs](#) is a method proposed to perform photo-to-caricature translation by modeling geometric exaggeration and appearance stylization using GANs.
- Helena Sarin wrote an [article](#) for The Gradient explaining how she uses GANs to generate artistic pieces.
- [GAN Lab](#) is a visualization tool built on Tensorflow.js that teaches how GANs work and learn.
- [Learn](#) how NLP and machine learning are being used to train bots to create dialogue and perform improvisational comedy.
- An MIT research team is [developing](#) an AI algorithm that can automatically create pizza recipes and make graffiti. The idea with this research, besides being fun, is to help people develop deep intuitions about machine learning's capabilities and to help us in our creative endeavors.
- [Perception Engines](#) are a series of machine vision algorithms used to create different art prints that represent abstract shapes and the machines' understanding of the world.
- Microsoft has [trained](#) a conversational agent, known as Xiaolce, to convert pictures into Chinese poetry. This is part of an ambitious initiative to teach machines higher-level natural language generation abilities as compared to simple tasks such as generating caption.
- DeepMind proposes a [method](#) that attempts to make sense of the world around us by associating events to sounds, such as the wine glass shattering and the accompanying crash.
- Researchers propose a hierarchical convolutional neural network [approach](#) for automatically rating and evaluating the quality of academic papers.
- A recent [paper](#) presents a method for motion transfer using what's called a per-frame image-to-image translation with spatiotemporal smoothing. The authors also created a demonstration that received huge attention on YouTube showing



how professional dancing moves can be transferred to a target (amateur dancer) using pose detection.

→ Is it possible to generate faces from text? Using powerful expressive methods such as GANs, Animesh Karnewar [discusses](#) how he trained a model to generate faces from text descriptions alone.



Figure: Images generated by the text to face model ([source](#))

→ Distill releases new paper entitled “[Differentiable Image Parameterizations](#)” which describes a powerful, under-explored tool for neural network visualizations and art (highly recommended).

→ Facebook AI researchers [turn whistles](#) into orchestral music using style transfer.

→ On [capturing long-range structures](#) in music by stacking WaveNet autoencoders (listen to samples generated [here](#)).

→ This [method](#) is used to transcribe music through reinforcement learning.

→ [Libratus](#) is considered the world’s best AI-based poker player as covered by the Gradient.

## General NLP

*Below I provide a comprehensive list of many different types of papers that could help the reader to become more familiar with the latest advancements and emerging topics in the field.*

- Facebook researchers come up with a state-of-the-art [method](#) for machine translation that only relies on monolingual corpora which can be useful to deal with low-resource languages.
- Young and colleagues have released an update of their recent paper on "[Recent Trends in Deep Learning Based Natural Language Processing](#)." Their revised version includes the addition of SQuAD state-of-the-art results as reported in the current NLP literature.
- Bloomberg researcher, Yi Yang, releases [code and paper](#) for his new work on modeling convolutional filters with RNNs, which he claims naturally capture long-term dependencies and compositionality in language.
- Baidu presents [Deep Voice 3](#), a fully-convolutional attention-based neural text-to-speech system, which achieved state-of-the-art neural speech synthesis in orders of magnitude faster than current systems.
- [pair2vec](#) is a method for learning word embedding pairs that represent background knowledge about their implied relationships. The method is able to obtain gains on the SQuAD 2.0 dataset and can also be incorporated into a cross-sentence layer of existing inference models such as BiDAF.
- Baidu [releases](#) a machine translation algorithm called Simultaneous Translation with Anticipation and Controllable Latency (STACL), which performs simultaneous translation using a technique called controllable latency. Controllable latency works by using machine learning to anticipate speaker's next words instead of having to wait for a speaker to make a pause in their speech to carry out the translation.
- A [recent paper](#) released by the AI research group known as MILA from the Université de Montréal report several surprising observations on the limitations of GANs in the context of natural language generation (NLG) tasks. The main idea

of using GANs in place of maximum-likelihood (MLE) models trained with [teacher forcing](#) is that they don't suffer from the so-called "exposure bias". However, the authors found that the impact of exposure bias on sample quality is less severe than was previously claimed or thought.

- One of the key areas of NLP is the development of conversational agents (also referred to as chatbots) that are able to have smarter conversations. Building such type of AI technology requires understanding different aspects of a natural language dataset such as intent and entities, in addition to being able to resolve the difficult task of determining the intent of dialogue via back and forth conversation. The last phase is obviously the response generation phase, where the agent selects a response from a predefined set of answers or directly through automatically generated text. Read more [here](#) about how lang.ai, an NLP startup, is using unsupervised AI to induce intent, one of the challenging tasks of building natural language conversational agents.
- Find out how text generation can be done using an alternative method based on a hidden semi-markov model (HSMM) decoder, achieving similar performance to the standard encoder-decoder models. The proposed [model](#) provides a method which allows for more interpretability and control—something the authors claim is important in a text generation task.
- A [paper](#) proposes a system called Textual Analogy Parsing (TAP) to model higher-order meaning based on a frame-style meaning representation. The technique is able to specify what is shared and what is compared between component facts.
- Facebook AI research team releases a [paper](#) highlighting key findings on a method that dynamically selects the right types of embeddings for a specific task at hand. They call this concept dynamic meta-embeddings and it tends to outperform traditional word embedding algorithms on a variety of tasks.
- This year, we witnessed several approaches that aimed to build contextualized word embeddings which are frameworks that improves the traditional word embeddings by modeling context. ELMo was introduced as one of these successful methods and was presented in the paper "[Deep Contextualized Word Representations](#)".

- Researchers develop a machine learning [algorithm](#) which uses eye movements to predict personality traits.
- A paper called [“Character-Level Language Modeling with Deeper Self-Attention”](#) demonstrates how a character-level language model based on a transformer model outperforms RNN variants on two popular benchmarks.
- This survey [paper](#) discusses vector representations used in two main branches of sense representations (unsupervised and knowledge-based). It also highlights some of the major limitations behind vector space models.
- A recent [paper](#) proposes variational attention networks for retaining performance gains on machine translation and visual question answering tasks while boosting computational speed.
- [Deep INFOMAX](#) is a method to learn unsupervised representations with the training objective to maximize mutual information between part or all of the input and a high-level feature vector. The idea is that representations should be learned from the information content or a structural constraint rather than all the bits at once—decision that matters at the semantic level.

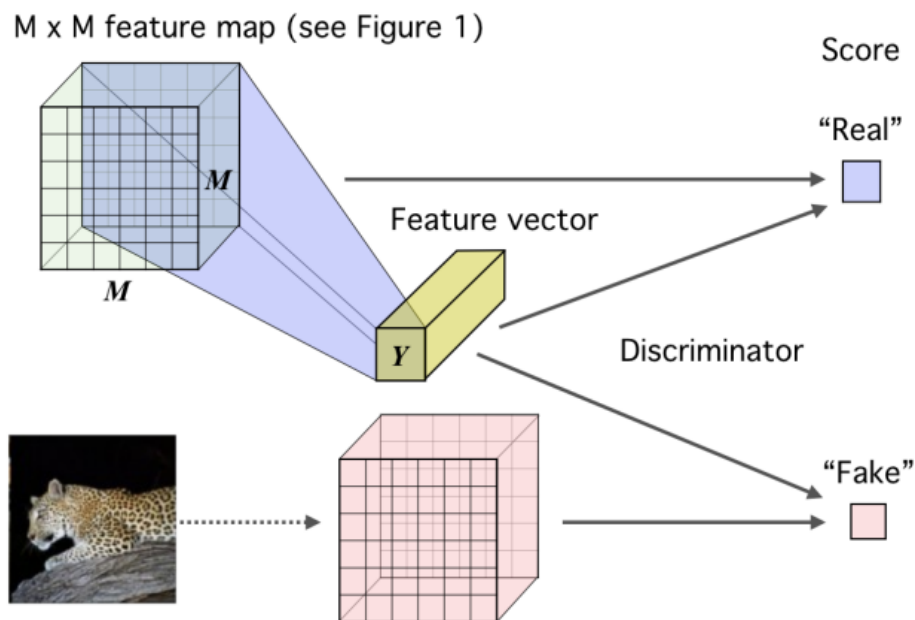


Figure: Deep InfoMax model ([source](#))

- NLP research is being used for all sorts of applications such as policymaking and clinical documentation research. This paper aims to introduce different methodologies and strategies that can be used to apply to textual data for conducting economic research.
- A novel spatio-textual clustering [algorithm](#) for clustering tweets using heterogeneous data.
- This [method](#) aims to evaluate the stability of embedding-based word similarities. The research, published in TACL, finds that nearest-neighbor distances are highly sensitive to small changes in the training corpus for a variety of algorithms.
- This [paper](#) introduces a la carte embedding, a simple and efficient alternative to the usual word2vec-based approaches for building word embeddings.
- Are BLEU and meaning representation in opposition? This important [paper](#) discusses this research question in great detail.
- This [paper](#) proposes distributed representations of hypertext documents (hyperdoc2vec) extracted from academic papers for two tasks (i.e., paper classification and citation recommendation).
- This [paper](#) describes the limitations of unsupervised bilingual dictionary induction that rely on x-lingual embedding models, which is an important topic in machine translation.
- This [paper](#) proposes a method to create accurate meta-embeddings from pre-trained source embeddings.
- This [work](#) investigates whether incorporating subword information in counting models is effective and how important it may be to learn representations of rare and out-of-vocabulary (OOV) words.
- This novel [method](#) aims to perform dependency parsing using a deep learning architecture and achieves state-of-the-art results on 21 treebank datasets with different dependency annotation schemas.

- This [research](#) discusses the possibility of predicting Twitter user socioeconomic attributes.
- This [article](#) discusses a method for "learning to represent words by how they're spelled", which helps to form representations of words that are unseen. This method essentially aims to deal with the out-of-vocabulary (OOV) problem.
- This [paper](#) provides detailed experiments and tips for training neural machine translation models using transformer model and Tensor2Tensor framework.
- This [paper](#) discusses a novel method for understanding online hate speech.
- A [paper](#) published by Google AI aims to explore the limits of language modeling.
- Amazon researchers have proposed a [language modeling](#) technique that adapts to different conversational contexts given an algorithm that analyzes a particular mathematical representation of a grammar's rules that calculates the probability that a grammar will produce any given string of words. This training strategy is important since acquiring a large body of training data to support a new capability may not be feasible in the real world. Thus, formal grammars are used to generate sample sentences for the language model instead.

# Trends and Overviews

## AI Annual Reports

*Below I list a few of the more interesting and comprehensive AI reports published this year.*

- The [AI Index 2018](#) report consists of a comprehensive list of trends for the year 2018. Readers are highly encouraged to read this report as it touches important themes such as published papers, public policy, media involvement, course enrollment, and AI ethics, among others.
- Nathan.ai published [State of AI Report](#) for 2018.
- If you are interested in diving deeper into AI ethics, the reader is strongly recommended to also refer to the [AI Now 2018](#) report. Topics in this report include race and gender bias, policy interventions, and other major problems brought about AI.
- The “[Malicious Use of AI](#)” report was released earlier this year and it covers topics such as security threats, digital security, physical security, interventions, prevention, policy solutions, among other security-related topics.
- A [free e-book](#) contains the current state of machine learning adoption in the enterprise. It discusses the different ML strategies and evaluation metrics used by the enterprise sector.
- Global AI talent [report](#) for 2018 describes how the AI talent pool is distributed across the world.

## Overviews

*In this section, I summarize a few interesting presentations that cover summaries of important areas of research in the NLP domain.*

- Yoshua Bengio [discusses](#) the implications of disentangled representations for higher-level cognition. He also discusses how natural language could be used as an additional hint about the abstract representations and disentangled factors which humans have discovered to explain their world.
- Marek Rei conducted [an analysis](#) earlier this the year reporting the fast growth of NLP and ML research. Even though the report was released early in the year (January 2018), it is possible to appreciate studies like these that aim to summarize the growth of the field. The figure below shows the number of publications for each of the top venues in the field. This coming year, we can expect an updated report.




Figure: NLP and ML 2017 Publication trend ([source](#))

- Earlier this year, Marek Rei [summarized](#) 57 of the latest and most important works on NLP. Topics included adversarial methods, neural machine translation, and transfer learning, among others. Even though the list was published at the beginning of the year, it can be used to compare how far we have come this year.
- Priya Dwivedi [covers](#) the building blocks of a question answering system build on top of [Stanford Question Answering Dataset \(SQuAD\)](#). She explains all the main



components that are needed to build an effective QA system — from the embedding layer to the output layer.

- André Martins summarizes his [notes](#) and discusses some of the most important topics presented in this year's [Annual Conference of the Association for Computational Linguistics](#) (ACL) and [International Conference in Machine Learning](#) (ICML).
- Christopher Manning and Ruslan Salakhutdinov released a series of [overview lectures](#) on deep learning methodologies and how to apply them to various datasets.
- Stephen Merity discusses attention and memory in deep learning networks in his "AI BY THE BAY" [talk](#).
- Here is Richard Socher's keynote [talk](#) on the recent progress in NLP and areas such as multitask learning.
- This [video](#) cover details on how to leverage knowledge graphs and deep learning for video classification and recommendation systems.
- Yisong Yue gives a broad overview of "[Imitation Learning](#)" techniques and applications (presented at ICML 2018).
- [Listen](#) to Bushra Anjum, Technical Lead at Amazon, discuss the importance of having an interdisciplinary mindset and focusing on core human values, such as creativity and empathy, when building technologies.
- Eugenio Culurciello, in a recent [post](#) published earlier this month, provides a nice history of neural network architectures and what motivated the design of each one (see chart below for an overview of architectures).




Figure: Model comparison ([source](#))

- DAIR.ai published a [post](#) which offers a light introduction to transfer learning for natural language processing and discusses key questions necessary to better understand this important research area.
- Lilian Weng [releases](#) a detailed article explaining the concept of Meta-Learning and how it can be used to build models that learn to adapt quickly to new environments given just a few examples.
- Sam Altman [discusses](#) progress in reinforcement learning.
- Elvis Saravia delivers a [talk](#) at PyCon 2018 Taiwan where he compares PyTorch and TensorFlow in the context of deep learning for NLP research.
- Dan Jurafsky, gives a [lecture](#) on how to process the language of policing and answers the question “Does This Vehicle Belong to You?”
- Jay Alammar provides a neatly illustrated [guide](#) to BERT, ELMo and other transfer learning techniques used to solve a variety of NLP tasks.
- Jacob Andreas [shares](#) his perspective on meaning representations in models without logical forms.
- [Slides](#) by Andrej Karpathy, on building the software 2.0 stack and what a machine learning IDE should contain.

- A comprehensive [overview](#) of deep learning for objection detection.
- Siva Reddy release [slides](#) on “The State of Natural Language Understanding”. I consider this a very important resource for those who want to learn more about NLU or need a brief introduction to the field.
- This [piece](#) discusses the differences between NLP and NLU.
- A collection of [notes](#) which includes the major highlights of ICML 2018.
- Ilya Sutskever, co-founder of OpenAI, gives a [talk](#) on Artificial General Intelligence at MIT.
- Christopher Manning [talks](#) about neural network that can reason.
- This comprehensive [report](#) details the current state of neural machine translation (NMT).

# Resources

## Dataset Releases

*Below is a list of the most interesting and impressive dataset releases of the year.*

- A [paper](#) accepted at EMNLP 2018 presented a benchmark dataset for machine translation of noisy text (MTNT), consisting of Reddit comments and professionally sourced translations. It differs from previous datasets which are mostly synthetically generated.
- Quora has publicly released a dataset called [Question Sincerity](#), where the goal is to feed the data to a machine learning model and train it to be able to distinguish well-intentioned questions from provocative ones. The idea is to be able to capture those ill-intentioned questions, flag them, and remove them from the platform so as to reduce any harm that such information could cause to the community. The dataset was also released on [Kaggle](#) as a competition.
- Google releases an image captioning dataset called [Conceptual Captions](#). This dataset was released as part of a research paper presented at ACL 2018.
- [Gab.ai corpus](#) is a large-scale dataset for studying hate speech and toxicity on social media platforms.
- A new large-scale [dataset](#) and task for visual common sense reasoning has been made publicly available, with the goal to enable cognition-level understanding in AI systems.
- Earlier this year, [Gensim](#) (a popular machine learning tool) released a dataset that contains patent grants from the US Patent and Trademark Office.
- Facebook and New York University released a dataset called [XNLI](#) which was created for evaluating cross-lingual approaches to natural language understanding (NLU). XNLI also includes baselines that can assist researchers to create systems that understand multiple languages.

- [ShaRC](#) is a dataset that focuses on building end to end conversational question answering systems that support the addition of background knowledge, especially when the system tries to answer a more difficult question, where the answer is not directly in the text source.
- Question Answering in Context ([QuAC](#)), is a dataset consisting of data instances representing dialog between two crowd workers. It is useful for modeling, understanding, and participating in information seeking dialog.
- Amazon releases sales rank [dataset](#) for kindle and print books.
- [Dataset](#) released by Rada Mihalcea can be used to build fake news detection systems.
- Alexa scientist releases [FEVER](#), a public dataset containing 185,000 data instances useful for fact extraction and verification.
- IBM releases a [dataset](#) which includes information that can be leveraged to build more comprehensive QA systems based on knowledge and reasoning.
- In an effort to build more representative ML models and promote inclusiveness in AI, Google AI announces the [Inclusive Images Competition](#) on Kaggle. The challenge is to build robust image captioning tools that work even for images that contain underrepresented groups based on the [Open Images dataset](#).
- TwentyBN has released a massive video [dataset](#) (Something-Something V2) to help enable systems to have the ability of video understanding and visual common sense.
- A [new version](#) of FastText includes pre-trained word embeddings in 157 languages, a resource which is useful for those working on multilingual research.
- Microsoft releases a massive [collection](#) of free datasets for advancing NLP and machine learning research.

→ [SWAG](#) is a new natural language inference dataset which was presented at the EMNLP 2018 conference. The dataset consists of adversarial generated and human-verified question-answer pairs, which are useful to test language modeling, question-answering, or natural language inference systems.



Figure: SWAG dataset ([source](#))

→ DeepMind open sources their [dataset](#) used to train generative query networks (GQNs) for neural scene representation and rendering.

→ A song lyric toxicity [dataset](#) is made available accompanied by analysis and slides.

→ FAIR releases [dataset](#) which can be useful to train AI Agents to teach other visual navigation.

→ Here is a nice [dataset](#) which contains short jokes scraped from various websites.

→ A [dataset and model](#) for comprehending paragraphs describing processes.

→ Google releases a [dataset](#) and challenge for landmark recognition.

→ [DataTurks](#) offers several open datasets and offers a neat interface to explore them.

→ [CoDraw](#) is a dataset used for training a task that enables collaborative drawing between two agents using natural language understanding from dialog.

→ Microsoft [released](#) 125 million building footprints from the US as open data.

→ IBM [releases](#) a dataset of recorded debates containing 60 speeches (audio + ASR & human transcripts). Useful for conducting different NLP tasks such as argument detection and argument stance classification.

- Here is an alphabetical [list](#) of NLP related datasets.
- Google announces [Open Image V4](#) which contains 15.4 bounding-boxes for 600 categories on 1.9M images. Google claims this is the largest existing dataset with object location annotations.
- Amazon is working on making a large-scale fact extraction and verification dataset publicly available. Learn more about their efforts [here](#).

## Learning Resources

*Due to the high demand for NLP and ML talent in both industry and academia, this year we witnessed a large number of online learning resources made publicly available by different companies and research institutions. Below, I list some of the notable learning resources of 2018.*

- [Materials](#) for the deep learning and reinforcement learning summer school carried out in Toronto are now available (videos included). The materials include topics such as auto differentiation, interpretability, and language understanding.
- Elvis Saravia and Soujanya Poria release a project called [NLP-Overview](#) that is intended to help students and practitioners get a condensed overview of modern deep learning techniques applied to NLP, including theory, algorithms, applications, and state-of-the-art result.

between the input vectors without making the number of parameters exceptionally large like MV-RNN. RNTN is defined by:

$$p_i = \tanh \left( \begin{bmatrix} b \\ c \end{bmatrix}^T V^{[i,D]} \begin{bmatrix} b \\ c \end{bmatrix} + W \begin{bmatrix} b \\ c \end{bmatrix} \right)$$

where  $V \in \mathcal{R}^{2D \times 2D \times D}$  is a tensor that defines multiple bilinear forms.

## B. Applications

One natural application of recursive neural networks is parsing (Socher et al., 2011). A scoring function is defined on the phrase representation to calculate the plausibility of that phrase. Beam search is usually applied for searching the best tree. The model is trained with the max-margin objective (Taskar et al., 2004).

Based on recursive neural networks and the parsing tree, Socher et al. (2013) proposed a phrase-level sentiment analysis framework (Figure 19), where each node in the parsing tree can be assigned a sentiment label.

Figure: NLP Overview interface ([source](#))

- [Lecture material](#) is available for Sebastian Raschka's new ML course given at the University of Washington Madison.
- OpenAI releases a new educational package, called [Spinning Up in Deep RL](#), for those interested in learning about the topic of deep reinforcement learning. It includes an extensive list of algorithms and resources used to effectively train deep reinforcement learning algorithms.
- Yandex School of Data Analysis (YSDA) releases [material](#) for their new NLP course (GitHub repo).
- NYU announces a [course](#) called "Neural Aesthetics" for learning how to teach different artistic capabilities to neural networks.
- Stanford released a new course called "[TensorFlow for Deep Learning Research](#)". It contains full lecture notes and slides covering topics such as convnets, generative adversarial networks (GANs), transformer, Tensor2Tensor, and much more.
- Stanford also released [slides](#) for all the NLP related seminars for Fall 2018, which include emerging topics such as multi-task learning, semantic role labeling, and visual questions answering.



- Deeplizard releases a new [course](#) covering the fundamentals of neural networks and tensor math, taught purely with PyTorch.
- Shervine Amidi releases a neat website containing several deep learning and machine learning cheat-sheets. The [guides](#) are available in several languages.
- Berkley AI Research (BAIR) offers new "Intro to AI" [course](#).
- Here is a nice [list](#) of Machine Learning rules and best practices for deploying real-world ML-based apps provided by Google's ML team.
- Here is a [mini-course](#) on Deep Learning with PyTorch (lecture slides and code included).
- Bloomberg is offering a new [course](#) on "Foundations of Machine Learning".
- Facebook is investing [efforts](#) to teach machine learning in an intuitive way (in a six-part video series) to learners from all different backgrounds.
- The University of Texas is offering a new [course](#) on Linear Algebra hosted on edX.
- Sebastian Ruder releases [NLP Progress](#), a repository to keep track of state-of-the-art results in NLP research.
- Google releases new machine learning [course](#) for free.
- As part of the ML Education project, Google releases [ML Learning Guides](#) which are a set of tutorials to teach step-by-step machine learning essentials and best practices.
- Moustapha Cisse, the founder of Google AI Ghana, [introduces](#) new one-year intensive Master's Program for Machine Intelligence in Africa.
- [The Gradient](#) is a publication that aims to democratize AI through educational content.

- François Fleuret releases handouts and slides for his [deep learning course](#) which was taught using PyTorch 1.0.
- Washington University (in St. Louis) opens course "[T81-558: Applications of Deep Neural Networks](#)" (includes Jupyter notebooks).
- Yann LeCun, Mikael Henaff, and Alfredo Canziani released a [new course](#) on deep learning that aims to teach the latest techniques in deep learning and representation learning.
- Stanford released a new course for NLP called "[CS224n: NLP with Deep Learning](#)".
- Egor Polusmak releases a [machine learning course](#) which consists of topics such as visual data analysis with Python and unsupervised learning.

## Multitask Question Answering Network (MQAN)

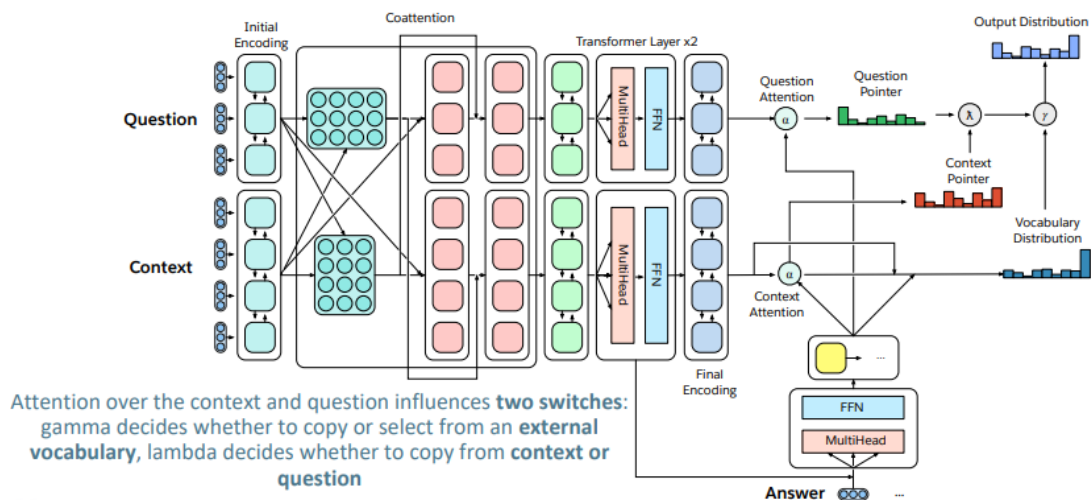


Figure: Multitask Question Answering Network ([source](#))

- [Deeplearning.ai](#) releases new [course](#) on sequence models, introducing topics such as GRUs, LSTMs, and Recurrent Neural Networks (RNNs). This specialization is taught by Andrew Ng.

## Books

*This year, we also saw a ton of book releases. Here I list some of the notable book releases.*

- Daniel Jurafsky and James H. Martin release the 3rd edition of their book called "[Speech and Language Processing: An introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition](#)". This edition covers a wide range of topics such as dialog systems and chatbots, question answering, and much more.
- Will Ratcliff discuss [methods](#) to capture the curiosity and imagination of your audience when delivering scientific presentations.
- Andrew Trask releases his new book entitled "[Grokking Deep Learning](#)" where he aims to teach deep learning and related mathematical concepts in a more intuitive way using Numpy (notebooks included).
- Terence Parr and Jeremy Howard released an online book called "[The Matrix Calculus You Need for Deep Learning](#)", which discusses the fundamental math you need for applying deep learning to various problems.
- "Fairness and Machine Learning: Limitations and Opportunities" is an online [textbook](#) discussing fairness in machine learning.
- Andrew Ng releases a [book](#) called "Machine Learning Yearning" where he teaches how to structure Machine Learning projects.
- Goku Mohandas [released](#) a set of notebooks (called Practical AI) that teaches how to program machine learning models using a practical approach via PyTorch.

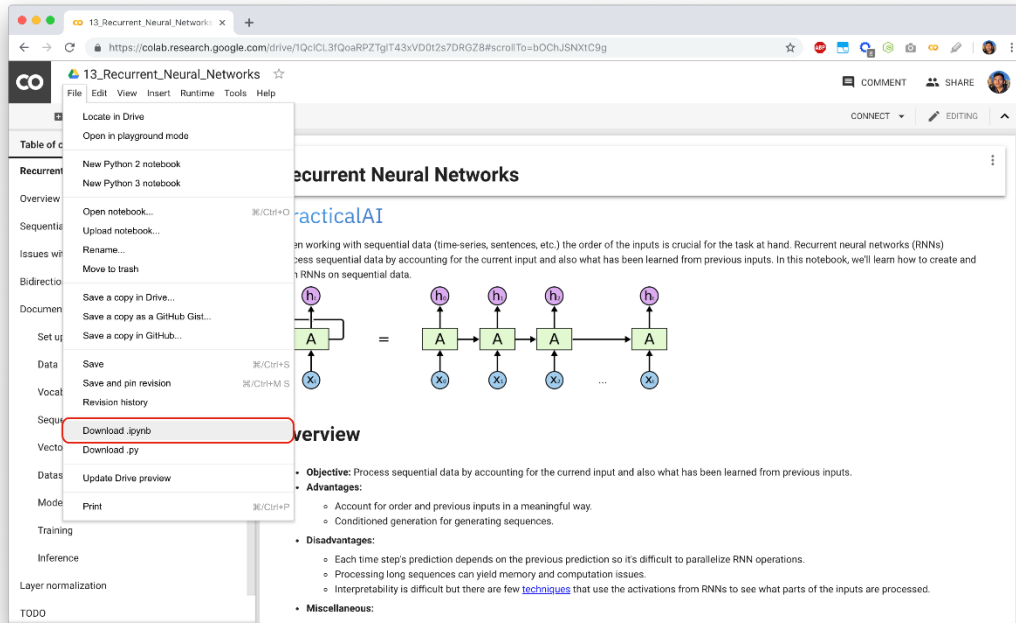


Figure: Practical AI notebooks ([source](#))

- Bharath Ramsundar and Reza Zadeh published a book on “[TensorFlow for Deep Learning](#)” which covers topics that range from linear regression to Reinforcement Learning.
- Ian Goodfellow publicly releases [LaTeX files](#) for the notations used in his “Deep Learning” book.

## Newsletters and Podcasts

*Over the past year or so, there has been a series of newsletters that aim to update readers on the latest trends and stories in AI and NLP. Here are a few resources which I came across this year, and which I strongly recommend the reader to follow.*

- [This week in AI](#) by Denny Britz
- [NLP News](#) by Sebastian Ruder
- [NLP Highlights](#) (podcast) by Matt Gardner and Waleed Ammar
- [Deep Learning Weekly](#)
- [nathan.ai newsletter](#) by Nathan Benaich

- [Alignment Newsletter](#)
- [AI Diary](#) by Elvis Saravia
- [The Creative AI Newsletter](#)
- [NLP Newsletter](#) by Elvis Saravia

## Open Source

*As described at the beginning of this report, the field wouldn't be where it is now if it were not for the great efforts made by developers as well. It's important to understand the role software engineering plays in research. In an effort to bring awareness, a few of the outstanding tools and software releases are summarized below.*

- Google AI releases [Dopamine](#), a TensorFlow-based framework that provides flexibility, stability, and reproducibility for new and experienced reinforcement learning researchers.
- TensorFlow announces [Tensorflow.js](#), a Web-GL accelerated, browser-based JavaScript library for training and deploying ML models on the web.
- TensorFlow releases a beginner's guide on how to apply probabilistic programming to real-world problems. It basically serves as a guide on how to leverage [TensorFlow Probability](#), a library built for scientists, statisticians, and ML researchers, to encode domain knowledge to understand data and make predictions.
- Google releases [Cirq](#), an open source tool which enables researchers to write quantum algorithms for quantum processors such as those used for NISQ computers.
- [GluonNLP](#) is a tool for reproducing state-of-the-art research in NLP. It provides several implementations and features that makes it easier to build NLP-based prototypes and products.
- The Facebook AI team has released [pytext](#), a set of NLP libraries and pre-built models built on top of PyTorch that provide functionalities to build NLP

applications that scale and are efficient at inference time. The library is mostly aimed at developers rather than researchers.

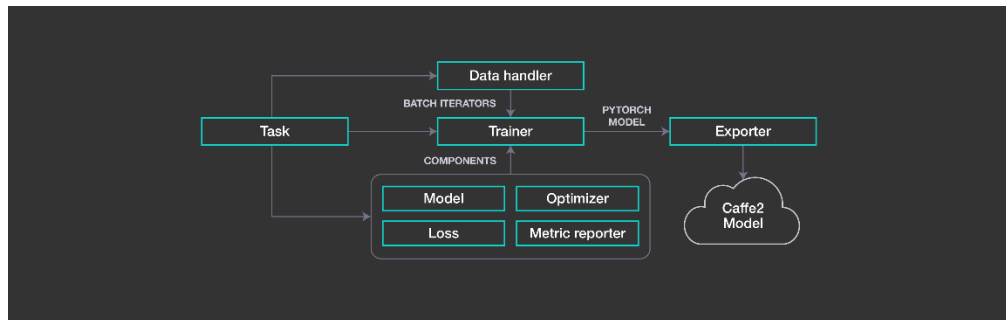


Figure: PyText framework ([source](#))

- Google AI releases [What-If Tool](#), a TensorBoard feature to help users better understand their machine learning models without writing code.
- The TensorFlow team releases [AutoAugment](#) which consists of a series of modules provided on TFHub that allows researchers to train better image models with fewer data using image augmentation tricks.
- Google releases [Dataset Search](#), a platform to quickly and efficiently search and find open datasets which have been uploaded to public sites such as personal websites and university profiles.
- Earlier last week, Facebook AI research team released [wav2letter](#) — “a simple and efficient end-to-end Automatic Speech Recognition (ASR) system.” The toolkit provides pre-trained models that can get you started right away with transcribing speech. The accompanying paper can also be found [here](#).
- The Facebook research team introduces [DensePose](#), a real-time approach for mapping human pixels from 2D images to 3D surface model of a human body.




Figure: DensePose-RCNN results ([source](#))

- The Natural Language Decathlon ([decaNLP](#)) benchmark offers a unique setting for studying general NLP models that can perform several natural language tasks.
- [DGL](#) is a library to build graph neural networks including Graph Convolutional Networks.
- With [TensorBoardX](#), now you can use TensorBoard with many other deep learning frameworks such as PyTorch, Chainer, MXNet, among others.
- Google's [MLPerf](#) is an effort that aims to build a common set of machine learning benchmarks to measure system performance for both training and inference from mobile devices to cloud services.
- Salesforce's Einstein AI team releases [TransmogrifAI](#), an AutoML library that focuses on accelerating machine learning developer productivity through automated machine learning for structured data.
- [Papers with Code](#) is a web tool for searching machine learning papers that contain open source code.
- The MXNet team releases [GluonCV](#), a deep learning toolkit that lets engineers and researchers quickly develop new algorithms and baselines (for image recognition, object detection, and semantic segmentation).

- [DeepSuperLearner](#) is an implementation of deep ensemble methods for tackling classification problems.
- [Pythia](#) is a modular framework that won the visual question answering challenge presented by FAIR.
- OpenAI announced an [efficient method](#), called gradient-checkpointing, to train memory-efficient and fast deep neural networks.
- You can find all of Facebook's open source projects [here](#). It includes PyTorch, Caffe2, ONNX, Tensor Comprehensions and much more.
- Google offers free access to GPUs on their online collaborative notebook service called [Google Colab](#).
- [ml5.js](#) is an open source web development tool, built on top of Tensorflow.js, that lets users easily access machine learning algorithms and models on the browser.
- FAIR releases [code](#) for the paper "[Colorless green recurrent networks dream hierarchically](#)". This is another work that aims to understand what recurrent architectures are learning and to what extent they are effective at modeling hierarchical structure.
- FAIR releases [PyTorch implementation](#) of "Poincaré Embeddings for Learning Hierarchical Representations". This [paper](#) was published in NeurIPS 2017 and proposed an approach for learning hierarchical representations of symbolic data by embedding them into hyperbolic space.





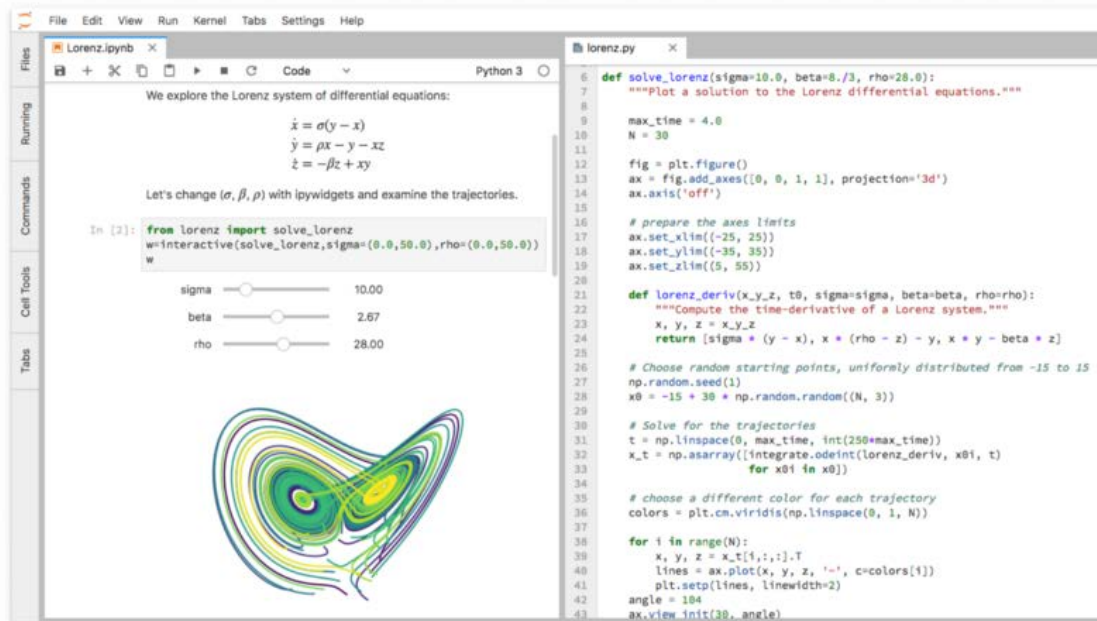


Figure: JupyterLab interface ([source](#))

- Peter Norvig releases [Lisp code](#) for the textbook "Paradigms of Artificial Intelligence Programming".
- TensorFlow [implementation](#) of the Capsule Network model implemented in the paper "[Dynamic Routing between Capsules](#)".
- [TextQL](#) is a library that allows you to execute against CSV or TSV text files.
- [PolygonRNN++](#) is a new interactive annotation tool for segmentation datasets.
- FAIR share open source [tools](#) for fastMRI with the aim to increase the efficiency and speed of processing MRI scans (up to 10x times faster).
- [CoinRun](#) is a training environment developed by OpenAI which provides a metric for an agent's ability to generalize across new environments, which is a challenging task in deep reinforcement learning.
- DeepMind releases [GraphNets](#), which is an easy-to-use library for training graph-based neural networks.

- [Stateoftheart.ai](#) is a platform for crowdsourcing state-of-the-art results in different areas of AI research such as reinforcement learning and natural language processing.
- [Muppet Show](#) is a web interface that allows you to explore how language is seen and analyzed by models such as BERT and ELMo. It consists of a visualization tool where you can query for specific tokens and find similar contexts to them.
- Jonathan Kummerfeld has developed an impressive [web page](#) to visualize and showcase a neural Part-of-Speech tagger using various deep learning toolkits such as Dynet, PyTorch, and TensorFlow.
- Zaid Alyafeai releases [code](#) for his image-to-image translation system which operates on the browser. The system implements the [pix2pix paper](#) and was developed to work on the browser using Tensorflow.js.
- Here is a beautiful website called [REGEXPER](#) which allows you to test your regular expressions on the browser.
- A new medical imaging framework ([Medical Torch](#)), based on PyTorch, has been open sourced by Christian Perone. The idea with this tool is *to* simplify the steps of pre-processing MRI data.
- Code release for [MojiTalk](#), which is a system used to generate emotional responses at scale.
- TensorFlow introduces [SPINN](#), a tool that enables natural language understanding in TensorFlow with eager execution.
- Sublime Text 3.1 is [released](#).
- [Test Tube](#) is a library to track and optimize deep learning experiments which allows researchers to easily log experiments and parallelize hyperparameter search.

- Facebook AI Research (FAIR) release "[Tensor Comprehensions](#)", which is a library that aims to bridge the gap between researchers and engineers who work together in ML research and implementations.
- [TorchFold](#) is a tool built on top of PyTorch that makes it easy to batch anything regardless of the complexity of your dynamic architectures.
- [NCRF++](#) is an open-source neural sequence labeling toolkit.
- HuggingFace introduces [NeuralCoref](#) which is a fast and efficient coreference resolution tool built with neural networks and SpaCy.

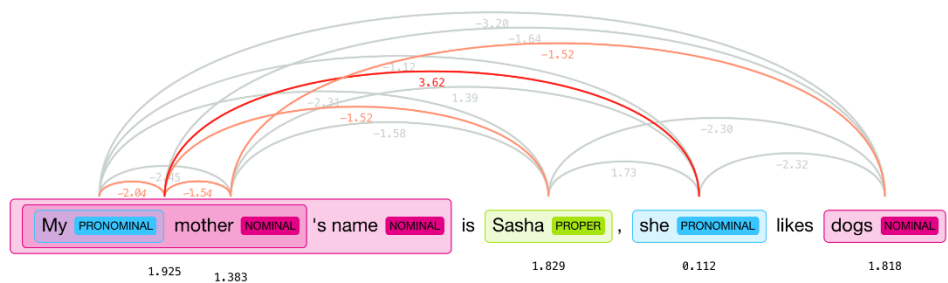


Figure: NeuralCoref snapshot ([source](#))

- [TextDistance](#) is a python library for comparing the distance between two or more sequences. Implements various algorithms such as Hamming, Jaccard index, etc.
- [DL-Text](#) is a repository containing code on how to pre-process textual data for deep learning models built with TensorFlow.

# Industry

## Conversational AI

*Conversational AI is a very important area of research in NLP because it encompasses the majority of NLP tasks that researchers investigate. Below I list some of the interesting and important stories related to conversational AI from the industry sector.*

- Wired.com released an [article](#) explaining how emotion recognition through voice will help machines to have a healthy relationship with humans.
- Mercedes-Benz released MBUX, a conversational agent which they consider achieves “revolutionary” and “best conversation you’ve ever had with a car” experience. They claim this is all made possible by advanced natural language understanding and AI technology. The idea is to allow the human to interact with a machine using natural phrases, a goal which a lot of companies in the conversational AI space seek to solve. Read more about this story [here](#).
- Hugging Face, NLP-based startup, [raises \\$4M](#) to create emotional intelligent chatbots.
- Wired also released an [article](#) earlier this year discussing the current state of Facebook’s virtual assistant (M) and why the company had plans to shut it down. This story was part of a bigger discussion in which it was believed that chatbots were just a fad and that the current state of NLP and AI technology couldn’t meet with the expectations of customers. A lot has changed since the release of this story but looking into what went wrong and how these technologies were improved over time helps to understand how challenging certain NLP tasks are.
- This [piece](#) covers how Passage AI is using NLP and deep learning, specifically bi-LSTMs, to train state-of-the-art conversational bots that can converse in all major languages, such as English, Spanish, and Chinese.
- Google researchers develop [techniques](#) to advance semantic textual similarity in conversation.



Figure: Semantic similarity in conversation ([source](#))

- Mastercard [claims](#) that conversational banking may already be mainstream. This is just another story which discussed the potential of conversational AI to disrupt the financial sector.
- Allen Institute for AI to [pursue](#) common sense for AI via Project Alexandria.
- [AskArvi](#) uses NLP and deep learning to understand customers' needs in order to provide insurance recommendations.
- Amazon's new [solutions](#) and APIs are set to offer machine learning for offline use, powered by a new compression algorithm, which can help to power in-vehicle infotainment systems and other mobile devices.
- Facebook files [patent](#) for e-Commerce messenger bot with NLP being the core technology behind the technology.
- Google releases [Google Duplex](#), a new technology for conducting natural conversations to carry out "real world" tasks over the phone. This was announced at Google IO 2018.

## Health and Lifestyle

*As mentioned earlier, there have been many types of applications in the health space. Below I list some of the notable related stories coming from the industry sector.*

- Jessica Kent [discusses](#) how NLP is being used to accurately analyze data from the electronic health records and improve care for heart failure patients.
- Anthem is [partnering](#) with doc.ai to leverage AI to help predict when people will experience allergies or allergy patterns.
- [Learn](#) how MIT researchers are building an infrastructure based on AI technologies to crowd-source data, which they use to build a system that can predict new drug-protein interactions.
- According to a recent [report](#), NLP and AI are being combined to help transform healthcare from physical interaction to high-quality clinical service free from physical co-location of the doctor. Essentially, NLP can be used to mine unstructured documents and make fast and efficient suggestions around clinical decision support.
- DeepMind claims to have developed an automated [solution](#) for detecting and diagnosing retinal disease using deep learning.
- Computer scientists developed a [text mining tool](#) based on NLP that determines cancer therapeutics given a patient's biomarkers.
- [Proven Beauty](#) is a startup that uses NLP to provide a personalized line of skin care products to its customers.
- On the [challenges](#) of using NLP to assist in improving clinic documentation.
- [BrightSign](#) is a smart glove that empowers speech-disabled people to communicate better.
- A health organization [leveraged](#) weakly supervised learning approaches to label large clinical datasets; then they were able to use deep learning, computer vision techniques to perform disease detection in CT scans, 150 times faster than human radiologists.
- Learn [how](#) Google is planning to encourage well-being with the use of technology.

- Automated psychotherapy bot, Woebot, [raises](#) \$8 million round of series A funding.
- Stitch's CEO [shares](#) how the company sells personal style using data science.
- [Linguamatics](#) is a library that offers tools for text-based drug discovery and exploration.

## Machine Learning Tools

*Below is a list of popular machine learning tools released this year.*

- Google launched [Cloud AutoML](#), which aims to bring easy access to AI services on the cloud to big businesses. According to Google, this is part of their mission to democratize AI.
- Google AI announces [BigQuery ML](#), a feature that allows data scientists to deploy machine learning models on large-scale structured and semi-structured datasets.
- Fast.ai [partners](#) with AWS Open Datasets to standardize, host, and release open datasets on an infrastructure with high reliability.
- Facebook brings [faster translation services](#) to its platform thanks to multilingual embeddings.
- [PyTorch Geometric](#) is a tool for implementing geometric deep learning with PyTorch.
- Simple [tests](#) show the shallowness of the state-of-the-art Google translation.
- Linguistics Agents Ltd. releases [NLP platform](#) for training of deep reinforcement learning agents.
- [Here](#) is an AI-based tool that helps make it easier to code video games. The new tool, named Commit Assistant, is offered by Ubisoft.



- Google Brain announces [Magenta.js](#), an API for generating music and art using TensorFlow.js
- Google presents [Google Text-to-Speech](#), which is powered by DeepMind's WaveNet.
- [Semantris](#) is a word association game powered by NLP and machine learning.



Figure 1: Semantris interface ([source](#))

- TensorFlow 1.9 was [released](#), which includes an excellent starting guide on how to use tf.keras and eager execution. Subsequently, TensorFlow 1.10.0 was released, which includes BigTable integration, deprecation of Distribution module, among other announcements. TensorFlow 2.0 is set to be released early next year (2019). You can look at some of the new features that we can expect from the new version in [this forum](#).

- PyTorch [hosted](#) its first developer conference where the community discussed research and production capabilities of the new release, PyTorch 1.0.
- Intel open sources a python library called [NLP Architect](#) which is useful for exploring state-of-the-art deep learning techniques for NLP.

## Financing, Advertising, and Law

*The highlights below relate to other areas such as financing, advertising and law.*

- Australian Securities and Investments Commission (ASIC) wants to [use NLP](#) for enforcing and regulating company and financial service laws.
- [Tumi](#) is using AI and NLP for ad targeting.
- Wall Street firms and other big start-ups are now [investing](#) heavily in applying NLP and machine learning to help consumers manage their money and plan for retirement (a sort of automated financial manager if you will).
- Artificial Intelligence is now [being used](#) as a tool to level the playing field between small and huge law firms.
- [Learn](#) more about how ML and NLP are being used in the financial sector to redefine services and create new markets.
- [Learn](#) how law firms are using NLP to change the way they do business.
- Through the "[Data Transfer Project](#)", Google wants to enable the portability of data, which involves giving users full control of their own data.

## General NLP and ML Stories

*Here is a list of more stories related to NLP coming from the industry sector.*

- DARPA is said to be [investing](#) up to \$2 billion on the AI Next initiative, which aims to enable machines with contextual reasoning and problem-solving capabilities.
- Renowned professor and researcher, Pedro Domingos, is [set](#) to become the managing director to lead the machine learning efforts at a New York-based investment corporate, D.E. Shaw Group.
- Microsoft's latest [patent](#) intends to take the phone caller ID feature to the next level using NLP. The system will be able to receive a call and identify caller intention and action.
- Observe.AI, a startup that uses AI to improve call center services, has [raised](#) \$8 million to increase efforts to automate the call center experience with machine learning and NLP technologies.
- Bahasa.ai, an Indonesian AI startup, gets [seed funding](#) to build more robust NLP/NLU technology to help businesses deploy their chatbot strategy.
- [On the Dot](#), a machine learning-based delivery startup, raises £8M to increase their efforts and advance product development.
- A Shanghai-based startup, Ins For Renaissance (IFR), [raised](#) \$18+ million to keep improving their insurance services through data mining and machine learning.

# Miscellaneous

## People and Society

*Here is a list of the most interesting stories related to how NLP impacts people and society.*

- The organ donation process is now being heavily automated using AI technologies. In a recent [article](#), it was reported that thousands of living kidney donors are already being identified by these algorithms. The process that also needs help is to decide the waiting list for kidney transplants.
- OpenAI [hosts](#) transfer learning contest using a “Sonic The Hedgehog” theme game.
- Learn more about how researchers are making machine learning algorithms fair in this [article](#) published in Nature.
- A recent [survey](#) produced interesting results that reflected the different views of people on AI technologies. For instance, 61% of people agreed that educating people about AI technologies is a responsibility that should be shared.
- [Listen](#) to Michael Nielsen talk about how important Open Access is for achieving impact and his experience in reaching 3.5 million readers with his free online book “Neural Networks and Deep Learning”.
- Here is a [piece](#) discussing how DeepMind is using AI to explore dopamine’s role in learning.
- OpenAI bots [fall](#) to Dota 2 human professional players at The International 2018.
- OpenAI [announces](#) their first batch of AI scholars.
- Motherboard released a [story](#) discussing FakeApp, which is an application that allows the ability to generate fake porn through a face-swapping algorithm powered by deep learning techniques. This story went viral because it helped to

raise awareness in the research community about how deep learning can be misused. Such an event also raised questions about the way we should democratize AI and protect privacy.

- [Learn](#) about how NLP can be used to detect cyber-attacks through understanding deep thoughts and linguistic cues found in resources already available on the internet (e.g., hacker chatrooms).
- Yann LeCun [steps down](#) as Facebook's head of AI research group, FAIR.
- Companies developing AI systems [sign agreement](#) at IJCAI pledging to avoid the development of lethal autonomous weapons.
- Lila Ibrahim [becomes](#) DeepMind's new chief operating officer.
- University of Washington professor, [Yejin Choinka](#), to lead \$125 million "common sense AI" initiative.
- Timely and interesting [TED talk](#) by Dirk Hovy on why computers are still struggling to understand us (video).
- Judea Pearl, an AI pioneer, gives his [advice](#) on how we should move AI forward. His main argument is that in order for machines to be truly intelligent we must teach them cause and effect.
- Google [expands](#) the Google Brain team in Tokyo.
- An [AI-generated painting](#) sells for \$432,000. This is the first ever AI-created artwork that sells in a prominent auction house. It also created controversy among community members who discussed the implications of open source algorithms being used for profit.
- Helping kids develop [language skills](#) using NLP research.
- Scientists [plan](#) huge European AI hub to compete with the US.

- A great [roundup](#) of the talk delivered by Rachel Thomas discussing the accessibility of AI.
- CVPR publishes a set of [educational materials](#) for learning how to properly conduct research and be involved in the community (videos and slides available). There has been a rising concern in the field of machine learning on the topic of properly disseminating research. This set of tutorials are a great resource for all types of researchers from all levels of academia.
- Andrew Ng continues his mission to build an AI-powered society with the formation of the [AI Fund](#).
- There has been an ongoing [debate](#) among two leading researchers, Gary Marcus and Yann LeCun, on whether AI needs more innate machinery.
- MIT Technology Review published an [interesting piece](#) which features Ian Goodfellow, regarded as the “The GANfather” for his invention of GANs.
- Novel antibiotic recipes could be hidden in medieval medical text (see more [here](#)).
- Learn how Ashton Kutcher and company [plan](#) to use NLP for fighting child sex trafficking.

## Opinions

*Opinions are usually not a favorite pass time for research scientists, but it's important to understand that in order to progress any field there has to be open discussion and then consensus among community members. Opinions somewhat encourages health and interesting debates. I also make space for some of the opinion pieces that called my attention this year.*

- Professor Michael Jordan published a thoughtful [piece](#) on the current state of AI where he argues that the revolution of AI hasn't happened yet. Based on the popularity and reach of this article, it is perhaps one of the most important AI opinionated pieces of the year.

- [Learn](#) a bit about how AI companies want to embed common sense into machine learning algorithms to be able to improve performance on certain critical tasks such as driving autonomous cars and performing medical diagnoses
- Will conversational AI become mainstream? According to a [report](#), these types of technologies will save businesses more than \$8 billion annually by 2020.
- Insightful and detailed [article](#) published in the “Sorta Insightful” blog, which provides discussion on whether deep reinforcement learning works or not.
- Towards the end of 2018, the board of the Neural Information Processing Systems (NIPS) conference released the results from a survey they had issued earlier this year. Based on the feedback they received, they made the final decision not to change the name of the conference which sparked controversy and heated debate within the community. Leaders in the field continued to raise concerns about the conference name and even started their own petitions for a name change as an act of protest against the conference. The name change eventually happened. Even though we must always keep a high interest in safeguarding the progress of science, we also need to take care of the people doing the actual science, and if takes changing the name of a prominent conference, then I believe that was the right choice to make. We should do our best to make other underrepresented groups feel more welcome to the field. Read more of this story [here](#).
- Professor Toby Walsh has recently [published](#) a new fiction book entitled “2062: The World That AI Made” in which he writes about what AI technology will become in the year 2062.
- Jack Clark, communication director of OpenAI, [argues](#) how Facebook’s DensePose algorithm can be used for “real-time surveillance” which brings forth troubling implications for the future use of this type of ML systems.
- Check out this interesting Twitter NLP [debate](#) from leaders in academia discussing the problems of meaning and whether it can be extracted from language models or other deep learning methods.

- In a candid [talk](#) delivered at the EmTech conference, AI researcher, Zachary Lipton discusses why he believes that AI has been over-hyped and how to promote a much healthier conversation around AI.
- Find out the math insights behind 10 types of GAN algorithms in this [video](#).
- Hal Daume shares his [thoughts](#) on why diversity of speakers in conferences matter.
- Sebastian Ruder [discusses](#) why he thinks NLP's ImageNet moment has arrived.
- Nando de Freitas [discusses](#) the need for taking AI studies that deal with emotion and sentiment more seriously. The main argument is that there is a need to move away from the notion that emotion can only be labeled using positive and negative labels.
- This [article](#) explains what the relationship between meaning and sound is, and why people feel annoyed by words like "moist".
- A new [study](#) reports that AI will contribute around \$15.7 trillion to the global economy by 2030.
- Rachel Thomas provides a [nice introduction](#) to neural architecture search and why this sub-field has been considered important in making machine learning more accessible.
- Here is another important [work](#) from the Facebook AI team related to debiasing machine learning.
- Hannah Wallach explains her [viewpoint](#) about the differences between computer science and social science. The argument is that computational social science needs to be taken more seriously and that we still need social scientists to be able to properly apply machine learning to study social phenomena in an ethical manner.
- A philosophical [view](#) of word embeddings and how they are deeply and linguistically motivated.



- Interesting [article](#) discussing useful lessons learned in reproducing a deep reinforcement learning paper.
- This [article](#) debates about whether the heavy use of emoji are ruining the way young people use the English language to communicate?
- Great podcast [episode](#) by Microsoft Research on how to make simple models more accurate and accurate models more intelligible or interpretable.

## Research Guides

*There were many discussions around the topic of conducting more responsible research. Here I list a few resources and guides for the reader that is interested to know more about effective ways to communicate and report scientific findings.*

- Yoshua Bengio shares his [experience](#) on creating a research lab and discusses the key lessons and best practices ML researchers should be paying attention to.
- Elvis Saravia [discusses](#) the challenges of AI communication by reviewing the current technologies and strategies used to communicate science-related topics.
- [Learn](#) more about tools such as iris.ai and Dimensions.ai, which aim to make it possible to search for scholarly text using modern NLP and ML techniques.
- Zachary Lipton [discusses](#) four major problems trending in machine learning research (a must-read for researchers in all levels of academia).
- This guide shows [how](#) to conduct socially responsible NLP research.
- COLING 2018 release blog [article](#) which talks about how to improve your presentation skills for academic conferences.
- On [challenges](#) related to the reproducibility of machine learning algorithms and research.

- Prodigy releases [best practices](#) for NLP projects, which include discussions and guides on how to better carry out your NLP research projects.
- Delip Rao [discusses](#) six major issues in neural machine translation. He also addresses the issues with irresponsible AI reporting.
- [Learn](#) more about abstraction in a neural network in this fun talk delivered at Eyeo Festival by David Ha.
- Distill releases a [blog post](#) where they reflect on all the challenges they have experienced in their first year operating as a new kind of scientific journal.
- [Learn](#) how to cite datasets used to conduct linguistics research. Includes a standard that can be used for both NLP and AI research since there are no available guidelines.
- On [reproducing](#) a deep learning paper and everything you need to know about this process.
- Manubot is an open source [tool](#) that enables open online collaborative writing.
- In machine learning research, we rely a lot on human evaluation, but how do we make it cheaper? Here is an informative [article](#) explaining plausible ways that this can be done.

## **NLP and ML Guides**

*And last but not least, I provide a list of some of the most interesting and helpful NLP and ML guides that taught me to implement interesting ML techniques. There are many resources and websites you can visit to find interesting tutorials, here I only provide the ones I came across.*

- Luis Serrano releases [video lecture](#) on matrix factorization and how this technique can be used for Netflix movie recommendations.

- DeepMind researcher, Shakir Mohamed, [releases](#) an impressive set of slides where he introduces foundations, tricks, and algorithms needed for probabilistic thinking.
- Deep reinforcement learning is an important area of research today. Take a look at Joyce Xu [survey](#) on the latest trends in advanced reinforcement learning. She covers topics that go beyond DQN/A3C and introduces a class of reinforcement learning known as hierarchical reinforcement learning. In addition, the survey discusses other important areas of research, such as memory and attention, and how they can be incorporated into the reinforcement paradigm of machine learning.
- This year's EMNLP conference, one of the top venues on all things related to NLP, featured groundbreaking research papers that ranged from capsule networks applied to NLP tasks and a remarkable number of new datasets such as [HotpotQA](#) and [SWAG](#). Some wonderful people like [Sebastian Ruder](#), [Claudia Hauff](#), and [Patrick Lewis](#) put together their thoughts and reviews/highlights on some of the important works that were presented at the conference.
- Do you often get frustrated with getting your latest batch of research papers and SOTA results that are relevant to the topics you are interested in studying? If you are, Chip Huyen built a handy little crawler called [sotawhat](#) to help you efficiently query for the latest abstracts and summaries from arXiv.
- The Machine Learning Tokyo group has open sourced a [series](#) of GAN models implemented in both Keras and PyTorch.
- In NLP research, we typically employ word embeddings as the input representation of a model. This [article](#) discusses in great detail how word embeddings contribute to neural networks and NLP systems as a whole.
- The Gradient releases [article](#) where Ana Marasović discusses how different research strategies can be used to tackle several of the limitations encountered in current deep learning methods for NLP.
- This [guide](#) shows how to transform words via semantic similarity through a word embedding approach. Mimics similar approach used in image morphing.

- Emmanuel Ameisen published a very popular [step-by-step guide](#) showing how 90% of NLP problems can be solved.
- [How](#) question answering systems are built at LinkedIn (slides presented at KDD 2018).
- Vered Shwartz [discusses](#) deep learning in NLP in this introductory post, which is aimed for non-experts or NLP enthusiasts.
- [Here](#) is a light introduction to natural language processing (NLP) by Niklas Donges.
- Check out this impressive [list](#) of tutorials on some of the most outstanding ML research of the past few years, such as DeepStack, InfoGAN, and AlphaGo Zero (highly recommended).
- Access all ICML 2018 tutorials [here](#) (videos).
- All ICLR talks for 2018 are available [here](#).
- Check out all the [talks](#) given at NAACL 2018 (videos).
- Complete [slides](#) for the EMNLP 2018 keynotes have been made publicly available.
- NeurIPS 2018 highlights and other relevant information can be found [here](#).
- PyTorch [implementation](#) of how to learn distributed sentence representations.
- Learn how to [speed up](#) your Python NLP modules by 50-100 times faster.
- An updated [list](#) of some of the most influential papers in deep learning (summaries included).
- Zaid Alyafeai teaches how to build Keras models and migrate them to the browser (using TensorFlow.js) in this very easy to read [article](#).

- This illustrated and animated [guide](#) helps to explain backpropagation through an intuitive and storytelling approach.
- Step-By-Step [implementation](#) of the "Attention Is All You Need" paper using PyTorch.
- Great [tutorials](#) on how to get started with the AllenNLP library to build and train your NLP models.
- [How](#) to obtain a good corpus for grammar correction via translation.
- Jay Alammar published a great [post](#) explaining the Transformer in an illustrative way. The Transformer is a model based on attention and helps to speed up model training.
- Beginner's [tutorial](#) on how to train and visualize word vectors (embeddings).
- This [tutorial](#) teaches how to efficiently train sequence to sequence models for neural machine translation with TensorFlow (Colab tutorial).
- Researcher develops a potential [way](#) to detect spelling error using word embeddings.
- A tutorial on "[Introduction to learning to trade with reinforcement learning](#)".
- Have you ever wondered what other types of applications you can apply word2vec to that doesn't involve NLP tasks? Find out from this [piece](#).
- If you haven't started to use TensorFlow high-level APIs, here is an [informative talk](#) discussing how to easily get started.
- On using CNNs to [classify resumes](#).
- Hiromi Suenaga provides detailed [summaries](#) of the deep learning lessons given by the fast.ai lecturers.

- Will Gannon provides his [list](#) of the 12 best free NLP and ML educational resources.
- HuggingFace published an impressive [piece](#) teaching how to implement Meta-Learning (learning to learn) for NLP using PyTorch.
- An excellent [tutorial](#) by Jaan Altosaar on everything there is to know about variational autoencoders.
- Here is a nice [piece](#) providing details on the most important and conventional text processing techniques important in text mining.
- Researchers at spaCy [discuss](#) how they perform multi-lingual natural language understanding with their in-house tools.
- Anima Anandkumar releases [notes](#) and slides on the problem of data scarcity in deep learning and discusses ways to mitigate it.
- Learn [how](#) to apply machine learning for music exploration and analysis.
- Mariya Yao wrote a [detailed article](#) explaining four approaches to NLP, which include topics such as distributional approaches, frame-based approaches, model-theoretical approaches, and interactive learning.
- Google releases machine learning [glossary](#) in Spanish, French, Korean, and Mandarin.
- This illustrated [guide](#) covers everything you need to know about the algorithms and methods used to build AlphaGo Zero.
- TechCrunch published a [story](#) discussing how deep learning combined with synthetic data can contribute to the democratization of AI. Synthetic data has also been heavily discussed in the community as an important area of research since deep learning techniques demand large-scale data, which for specific applications is still lacking.

- A great [guide](#) showing how to apply transfer learning to NLP and computer vision tasks.
- Facebook AI team obtain record-breaking [results](#) on image recognition using pre-training with hashtags on 3.5b images.
- The TensorFlow team shows how to implement a [real-time human pose estimation model](#) in the browser using TensorFlow.js.
- Learn how to ship neural network based application on iOS using PyTorch in this simple and easy to follow [guide](#) by Stefano Attardi.
- Here is a [comprehensive list](#) and comparison of web (text) annotation editors, which are all very useful for conducting and organizing research.