"To imagine a language is to imagine a form of life"

THE GOOD A.I.

babbling alloewd®

** The power of intuition: envisioning AI to be useful, good and with a human (not perfect) touch.
ML, DL, CNN, Generative AI, new social systems, Blockchain, IoT, leadership/ownership.
Philosophy, linguistics, limits of thought.
Bookmarks of activity and thoughts around 2016-2018.

Babbled In		by
23.05.2017	**Generative adversarial nets**, by Bengio, Y., Courville, A.C., Goodfellow, I.J., Mirza, M., Ozair, S., Pouget-Abadie, J., Warde-Farley, D., & Xu, B. (2014) NIPS. (cited 463 times, HIC: 55 , CV: 0) Summary: We propose a new framework for estimating generative models via an adversarial process, in which we simultaneously train two models: a generative model G that captures the data distribution, and a discriminative model D that estimates the probability that a sample came from the training data rather than G. **High-Speed Tracking with Kernelized Correlation Filters**, by Batista, J., Caseiro, R., Henriques, J.F., & Martins, P. (2015). CoRR, abs/1404.7584. (cited 439 times, HIC: 43 , CV: 0) Summary: In most modern trackers, to cope with natural image changes, a classifier is typically trained with translated and scaled sample patches. We propose an analytic model for datasets of thousands of translated patches. By showing that the resulting data matrix is circulant, we can diagonalize it with the discrete Fourier transform, reducing both storage and computation by several orders of magnitude. **A Review on Multi-Label Learning Algorithms**, by Zhang, M., & Zhou, Z. (2014). IEEE TKDE, (cited 436 times, HIC: 7 , CV: 91) Summary: This paper aims to provide a timely review on multi-label learning studies the problem where each example is represented by a single instance while associated with a set of labels simultaneously. **How transferable are features in deep neural networks**, by Bengio, Y., Clune, J., Lipson, H., & Yosinski, J. (2014) CoRR, abs/1411.1792. (cited 402 times, HIC: 14 , CV: 0) Summary: We experimentally quantify the generality versus specificity of neurons in each layer of a deep convolutional neural network and report a few surprising results. Transferability is negatively affected by two distinct issues: (1) the specialization of higher layer neurons to their original task at the expense of performance on the target task, which was expected, and (2) optimization difficulties related to splitting networks between co-adapted neurons, which was not expected. **Do we need hundreds of classifiers to solve real world classification problems**, by Amorim, D.G., Barro, S., Cernadas, E., & Delgado, M.F. (2014). Journal of Machine Learning Research (cited 387 times, HIC: 3 , CV: 0) Summary: We evaluate 179 classifiers arising from 17 families (discriminant analysis, Bayesian, neural networks, support vector machines, decision trees, rule-based classifiers, boosting, bagging, stacking, random forests and other ensembles, generalized linear models, nearest-neighbors, partial least squares and principal component	CST

Babbled In		by
13.01.2017	**7 AI trends to watch in 2017**	CST

From tools, to research, to ethics, Ben Lorica looks at what's in store for artificial intelligence in 2017.

2016 saw tremendous innovation, lots of AI investment in both big companies and startups, and more than a little hype. What will 2017 bring?

1. Democratization of tools will enable more companies to try AI technologies.

A recent Forrester survey of business and technology professionals found that 58% of them are researching AI, but only 12% are using AI systems. This is partially because applied AI applications are only now starting to be realized, but it's also because right now AI is hard. It requires very specialized skills and a develop-it-yourself attitude.

But frameworks like Facebook's Wit.ai and Howdy's Slack bot are competing to become the Visual Basic of AI, promising point-and-click development of intelligent conversational interfaces to relatively unsophisticated developers. Tools like Bonsai, Keras, and TensorFlow (if you don't mind coding) simplify the implementation of deep learning models. And cloud platforms like Google's APIs and Microsoft Azure allow you to create intelligent apps without having to worry about setting up and maintaining accompanying infrastructure.

…

Babbled In		by
09.01.2017		CST

Deep Learning suits problems where the target function is complex and datasets are large but with examples of positive and negative cases. Deep Learning also suits problems that involve Hierarchy and Abstraction.

Abstraction is a conceptual process by which general rules and concepts are derived from the usage and classification of specific examples. We can think of an abstraction as the creation of a 'super-category' which comprises of the common features that describe the examples for a specific purpose but ignores the 'local changes' in each example. For example, the abstraction of a 'Cat' would comprise fur, whiskers etc. For Deep Learning, each layer is involved with detection of one characteristic and subsequent layers build upon previous ones. Hence, Deep Learning is used in situations where the problem domain comprises abstract and hierarchical concepts. Image recognition falls in this category. In contrast, a Spam detection problem that can be modelled neatly as a spreadsheet probably is not a complex problem to warrant Deep Learning

A more detailed explanation of this question can be found in THIS Quora thread.

AI vs. Deep Learning vs. Machine Learning

Before we explore types of AI applications, we need to also discuss the differences between the three terms AI vs. Deep Learning vs. Machine Learning.

The term Artificial Intelligence (AI) implies a machine that can Reason. A more complete list or AI characteristics (source David Kelnar) is

Reasoning: the ability to solve problems through logical deduction

…

Babbled In		by
07.03.2017	"To integrate AI, you have to have an internal team of expert product people and engineers that know its application and are working very closely with the frontline teams that are actually delivering services," says Ian Crosby, cofounder and CEO of Bench, a digital bookkeeping provider. "When we are working AI into our frontline service, we don't go away to a dark room and come back after a year with our masterpiece. We work with our frontline bookkeepers day in, day out." In order to properly adapt to changing technologies, organizations are moving away from a top-down structure and toward multidisciplinary teams. In fact, 32% of survey respondents said they are redesigning their organizations to be more team-centric, optimizing them for adaptability and learning in preparation for technological disruption. Finding a balanced team structure, however, doesn't happen overnight, explains Crosby. "Very often, if there's a big organization, it's better to start with a small team first, and let them evolve and scale up, rather than try to introduce the whole company all at once."	CST

…

Babbled In		by
02.03.2017	**AI vs. Democracy (Feudalism 2.0 or Democracy 2.0)** Big data, artificial intelligence, cybernetics and behavioral economics are shaping our society—for better or worse. If such widespread technologies are not compatible with our society's core values, sooner or later they will cause extensive damage. They could lead to an automated society with totalitarian features. In the worst case, a centralized artificial intelligence would control what we know, what we think and how we act. We are at the historic moment, where we have to decide on the right path—a path that allows us all to benefit from the digital revolution. Therefore, we urge to adhere to the following fundamental principles: 1. to increasingly decentralize the function of information systems; 2. to support informational self-determination and participation; 3. to improve transparency in order to achieve greater trust; 4. to reduce the distortion and pollution of information; 5. to enable user-controlled information filters; 6. to support social and economic diversity; 7. to improve interoperability and collaborative opportunities;	CST

…

Babbled In		by
01.02.2017	**Artificial Intelligence** Report 2017	CST

IRC Reports Special: Künstliche Intelligenz

https://irc.siemens.de/rdds/rdds_content.php

The Challenges of Artificial Intelligence: A Literature Review
(more details and free table of content)

In the last two decades, researchers have found many advantages related to Artificial Intelligence (AI) in the performance of both service and manufacturing systems. It includes alternate methods to conventional techniques and there are mechanisms of integrated systems. Currently, many researches related to artificial intelligence are under process. At present, artificial intelligence has been used to solve various problems in several fields and has become more popular around the world. This field is extremely difficult to review either chronologically or thematically. AI has a long history of growing constantly and changing tremendously. Connected with the history of computers, artificial intelligence is the main concept behind the fifth generation computers. This report deals with the framework of artificial intelligence and focuses on fields related to artificial intelligence. Particularly this report describes about the recent growth in the field of artificial intelligence and its applications. Several issues of artificial intelligence are analyzed. It is concluded that still many researches are yet to be conducted for further growth in the field of Artificial Intelligence (AI).

…

Babbled In		by
16.02.2017	**Weak AI**	CST

Siri and Alexa could be considered AI, but generally, they are weak AI programs. Even advanced chess programs are considered weak AI. This categorization seems to be rooted in the difference between supervised and unsupervised programming. Voice-activated assistance and chess programs often have a programmed response. They are sensing for things similar to what they know, and classifying them accordingly. This presents a human-like experience, but that is all it is—a simulation. If you ask Alexa to turn on the TV, the programming understands key words like On and TV. The algorithm will respond by turning on the TV, but it is only responding to its programming. In other words, it does not comprehend any of the meaning of what you said.

Classification and recession can be similar to basic calculus; every Y is a function of X:

Y= f(x)

This is not what those cool sci-fi movies are using for their plots.

Strong AI

Featured in many movies, strong AI acts more like a brain. It does not classify, but uses clustering and association to process data. In short, it means there isn't a set answer to

…

Babbled In		by
20.01.2017	**What is Behavioral Psychology?** Let's define behavioral psychology. Behavioral psychology is the study of the connection between our minds and our behavior. Sometimes you will hear behavioral psychology referred to as behaviorism. The researchers and scientists who study behavioral psychology are trying to understand why we behave the way we do and they are concerned with discovering patterns in our actions and behaviors. The hope is that if we can use behavioral psychology to help us predict how humans will behave, we can build better habits as individuals, create better products as companies, and develop better living spaces as communities. … Traumatic events tend to trigger what Gilbert refers to as our "psychological immune systems." Our psychological immune systems promote our brain's ability to deliver a positive outlook and happiness from an inescapable situation. This is the opposite of what we would expect when we imagine such an event. As Gilbert says, "People are not aware of the fact that their defenses are more likely to be triggered by intense rather than mild suffering. Thus, they mis-predict their own emotional reactions to misfortunes of different sizes."	CST

…

Babbled In		by
23.05.2017	AI R&D Papers to read: **Dropout: a simple way to prevent neural networks from overfitting**, by Hinton, G.E., Krizhevsky, A., Srivastava, N., Sutskever, I., & Salakhutdinov, R. (2014). Journal of Machine Learning Research, 15, 1929-1958. (cited 2084 times, HIC: 142 , CV: 536). Summary: The key idea is to randomly drop units (along with their connections) from the neural network during training. This prevents units from co-adapting too much. This significantly reduces overfitting and gives major improvements over other regularization methods **Deep Residual Learning for Image Recognition**, by He, K., Ren, S., Sun, J., & Zhang, X. (2016). CoRR, abs/1512.03385. (cited 1436 times, HIC: 137 , CV: 582). Summary: We present a residual learning framework to ease the training of deep neural networks that are substantially deeper than those used previously. We explicitly reformulate the layers as learning residual functions with reference to the layer inputs, instead of learning unreferenced functions. We provide comprehensive empirical evidence showing that these residual networks are easier to optimize, and can gain accuracy from considerably increased depth. **Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift**, by Sergey Ioffe, Christian Szegedy (2015) ICML. (cited 946 times, HIC: 56 , CV: 0). Summary: Training Deep Neural Networks is complicated by the fact that the distribution of each layer's inputs changes during training, as the parameters of the previous layers change. We refer to this phenomenon as internal covariate shift, and address the problem by normalizing layer inputs. Applied to a state-of-the-art image classification model, Batch Normalization achieves the same accuracy with 14 times fewer training steps, and beats the original model by a significant margin. **Large-Scale Video Classification with Convolutional Neural Networks** , by Fei-Fei, L., Karpathy, A., Leung, T., Shetty, S., Sukthankar, R., & Toderici, G. (2014). IEEE Conference on Computer Vision and Pattern Recognition (cited 865 times, HIC: 24 , CV: 239) Summary: Convolutional Neural Networks (CNNs) have been established as a powerful class of models for image recognition problems. Encouraged by these results, we provide an extensive empirical evaluation of CNNs on large-scale video	CST

23.05.2017 CST

knowledge repository, and features a probabilistic inference system that computes calibrated probabilities of fact correctness.

Scalable Nearest Neighbor Algorithms for High Dimensional Data, by Lowe, D.G., & Muja, M. (2014). IEEE Trans. Pattern Anal. Mach. Intell., (cited 324 times, HIC: 11 , CV: 69).

Summary: We propose new algorithms for approximate nearest neighbor matching and evaluate and compare them with previous algorithms. In order to scale to very large data sets that would otherwise not fit in the memory of a single machine, we propose a distributed nearest neighbor matching framework that can be used with any of the algorithms described in the paper.

Trends in extreme learning machines: a review, by Huang, G., Huang, G., Song, S., & You, K. (2015). Neural Networks, (cited 323 times, HIC: 0 , CV: 0)

Summary: We aim to report the current state of the theoretical research and practical advances on Extreme learning machine (ELM). Apart from classification and regression, ELM has recently been extended for clustering, feature selection, representational learning and many other learning tasks. Due to its remarkable efficiency, simplicity, and impressive generalization performance, ELM have been applied in a variety of domains, such as biomedical engineering, computer vision, system identification, and control and robotics.

A survey on concept drift adaptation, by Bifet, A., Bouchachia, A., Gama, J., Pechenizkiy, M., & Zliobaite, I. ACM Comput. Surv., 2014 , (cited 314 times, HIC: 4 , CV: 23)

Summary: This work aims at providing a comprehensive introduction to the concept drift adaptation that refers to an online supervised learning scenario when the relation between the input data and the target variable changes over time.

Multi-scale Orderless Pooling of Deep Convolutional Activation Features, by Gong, Y., Guo, R., Lazebnik, S., & Wang, L. (2014). ECCV(cited 293 times, HIC: 23 , CV: 95)

Summary: To improve the invariance of CNN activations without degrading their discriminative power, this paper presents a simple but effective scheme called multi-scale orderless pooling (MOP-CNN).

Simultaneous Detection and Segmentation, by Arbeláez, P.A., Girshick, R.B., Hariharan, B., & Malik, J. (2014) ECCV , (cited 286 times, HIC: 23 , CV: 94)

Summary: We aim to detect all instances of a category in an image and, for each instance, mark the pixels that belong to it. We call this task Simultaneous Detection and Segmentation (SDS).

...