Artificial Intelligence

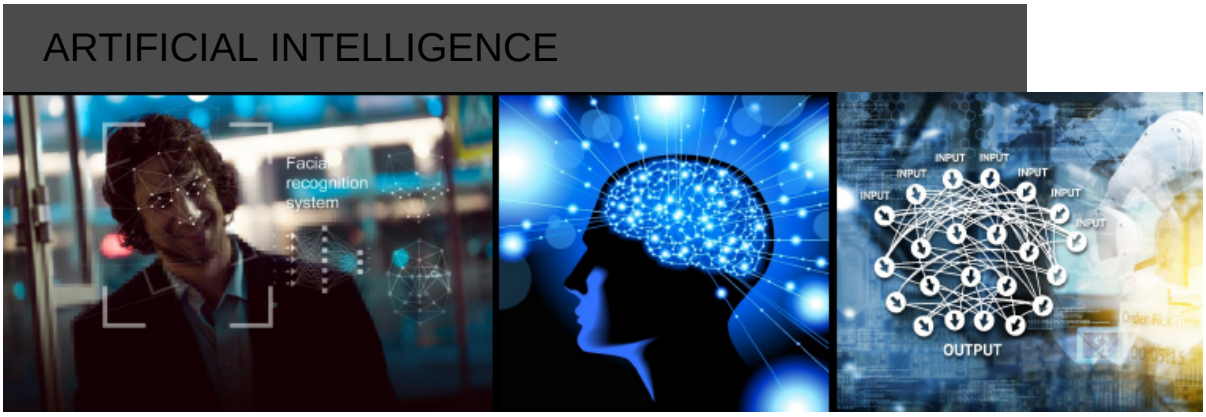

2020

2021-2022

2023-2024

2025+

AI ethics is a fast-emerging field that looks at how the design, development and implementation of AI could lead to unintended consequences, misuse or abuse of the technology. The ethical AI umbrella includes various aspects from bias and discrimination, data and information privacy, lack of transparency or explainability, responsibility and accountability, and the issue of workforce displacement due to automation. 2019 saw ethical AI come to the forefront, with technology vendors like Google and IBM developing ethical guidelines, regulators like EU driving GDPR regulations and ethical guidelines, and academics and non-profits working on AI research areas related to ethics.

Today, most of the AI being developed — including machine learning and deep learning — has a two-step process to it: training and inference. Training is where the model learns from the data, and inference is where the learned model is deployed and infers on the data it sees. Training a state-of-the-art AI model usually involves high-performance training hardware, which has largely been GPGPU (General Purpose Graphics Processing Unit) driven, with AI hardware compute performance growing roughly at 8x every year since 2012. But as trained AI models gets deployed across multiple use cases and industry verticals, there is a growing need for hardware that can perform the inference part of the compute. 2018-19 was a year where AI inference is a much bigger trend than AI training, as semiconductor companies including NVIDIA, Intel, Xilinx and a host of chip startups target AI inference — both in the cloud data center and the edge. The shift towards inference marks an important step in the maturity and deployment of AI across the enterprise.

Automated Machine Learning (AutoML) is the ability to automate AI model development tasks across data cleaning, feature engineering, model selection, and hyperparameter tuning. The ultimate goal for AutoML is to automate the grunt work, leaving the AI developers more time to think creatively around the application of AI in specific domains, or the ethical aspects of the AI deployment. AutoML could expand to cover the entire AI workflow at some point. AutoML capabilities have been rolling out across both AI cloud service providers like Google, Microsoft and Amazon, and also across traditional data science providers like DataRobot, Ayasdi and others. There are also a number of open-source AutoML frameworks including MLBox, Auto-Sklearn, TPOT, Auto-Keras, H2O and others. At Google I/O 2019, it was shown that AutoML already outperforms handcrafted models in a specific domain. AutoML has links to the field of meta learning, or learning to learn, where AI learns how to build AI.

Affective computing refers to the ability of computers to understand human emotions through analysis of voice, images, and other developing cues. Robots have largely been good at following instructions, but the integration of affective computing is bringing in an emotional component to robotics improving human computer interaction (HCI). Affective computing is being used in the Pepper robot to greet customers and respond to their facial expressions, while at the same time helping robots communicate with autistic children .

Reinforcement learning (RL) is a subset of machine learning algorithms that learn by exploring its environment. Unlike supervised learning that trains on labeled datasets, RL achieves its stated objective by getting positive or negative rewards for the actions that it takes. The environments in which RL works are usually simulated environments like games. AlphaGo uses RL (in combination with other techniques) and similar techniques have been used to have AI learn Atari games, or become champions at Poker. In the coming years, we are likely to see RL find uses in enterprise verticals such as manufacturing, aerospace, energy, transport, logistics applied to systems control and management problems.

Generative models are a branch of unsupervised learning which allows computers to generate data similar to the data you supply it — in essence re-creating data which could be images, video, text or speech. The most popular generative model is GANs (Generative Adversarial Networks) an d VAEs (Variational Autoencoders). Generative models for vision are being used to create realistic 3D worlds in games, create on-demand art, take Photoshop to the next level or simply generate artificial training data for AI. GANs are also giving rise to a market for synthetic content including synthetic avatars, synthetic celebrities and even synthetic actors. Generative models are also a cause for worry as they are likely to produce fake artworks, fake political propaganda, or fake celebrity videos.

Within generative models, language models like OpenAI's GPT-2 model, which is able to predict the next word given a series of previous words, have shown major performance gains. GPT-2 was trained on 8 million web pages with 1.5 billion parameters. GPT-2 has been shown to work in a 'zero-shot' setting, without any prior domain-specific training data, and outperform domain-specific models. GPT-2 has also shown promising results in language tasks like Q&A, reading comprehension, summarization and translation. OpenAI decided to only release a small subset of the model, citing AI safety and ethical issues related to malicious use of the technology like impersonating others and automating fake news production. Apart from GPT-2, there are other models like XLNet, BERT that are advancing language capabilities — all of which will have major implications for the field of writing both professionally and creatively. GPT-2 has also shown capabilities to generate code, which could change software development forever.

Improvement in ASICs like the use of novel graph-based AI architectures, combined with increasing ability to compress complex AI models in software will lead to a proliferation of Edge AI. It will become possible to process complex AI models at the source of the data, or the edge of the network rather than in the cloud. In areas like autonomous driving, robotics, IoT, industrial and medical vision, and mobile and personal computing, we are likely to see AI models be increasingly deployed on the edge, allowing for both inference as well as real-time, ongoing training. The combination of 5G and AI Edge Data Centers is a promising proposition, but Issues like data privacy, connectivity, and security will also dictate the distribution of compute in the edge vs cloud.

While AI is being used defensively to identify and tag new and emerging threats which are rapidly changing, AI can also be used offensively to generate new threats and bypass security systems. Trickbot and Doppelgangers were attacks in which the malware exhibited "hard coded" stealth tactics, obfuscation techniques, and locking mechanisms— but, in the future — could use AI and machine learning to perform a 'wait and watch' maneuver, staying dormant on a system and attacking at an appropriate time. Other areas where AI could be used in a cyberattack include AI-generated phishing emails or malware-inducing chatbots which can trick users into clickbaits. AI-enabled phishing and rogue chatbots are a genuine threat which could become exacerbated. Nation states are much more likely to possess and cultivate AI talent, or use proxies to create offensive cyber threats and tactics, which could be used against another nation state or even a large corporation.

Symbolic AI is an older brand of AI dating back to the 1980s that includes techniques like knowledge graphs, expert systems, and machine reasoning that use symbols, knowledge or specific instructions sets in order to help machines learn. Symbolic AI has largely been overtaken by non-symbolic AI techniques like machine learning and deep learning that use pattern recognition to learn from datasets, rather than using specific instructions or symbols. Symbolic AI, however, is making a comeback and the best way forward for deep learning is a combination with symbolic AI to create a global library of models that help us get to some form of Artificial General Intelligence (AGI).

One of AI's biggest challenges is in the area of reasoning, which is the ability to infer facts from existing data, deriving a logical conclusion or thinking rationally like a human. Until recently, AI reasoning has used traditional methods like associative memory-based learning, expert systems or knowledge graphs to create knowledge representations and then perform different reasoning techniques. However, we are now witnessing the use of deep learning within reasoning with DeepMind using multiple neural network-based techniques that perform on relational reasoning. As AI reasoning develops we could use it in multiple ways including problem solving, mathematical proofs, agent-based scene understanding, and enabling improved explainability of AI models.

AI is largely controlled and run by large Internet companies who have access to massive datasets, the best AI talent, and large-scale, high-performance compute capabilities in their data centers to train AI models. One of the top language models, XLNet, developed by Google and CMU, is estimated to cost $245,000 to train based on the hardware requirements with 512 TPU V23 chips trained for 2.5 days. As a result, there is a growing gap in terms of access to high-performance hardware required to advance AI. The privacy issues around Internet companies harvesting user data to advance AI is another issue around centralized AI development. Decentralized AI that doesn't have control centralized amongst a few private institutions, or for that matter, nation states, is inevitable. Early grassroot efforts around Decentralized AI exist in Singularity NET, OpenMined, Ocean and Algorithmia. Decentralized AI uses multiple technologies from decentralized ledgers or blockchain, smart contracts, homomorphic encryption, and federated learning to achieve runtime, privacy and learning capabilities.

The path from narrow AI to strong AI or AGI (Artificial General Intelligence) requires AI models to learn from fewer examples, generalize their learnings across a variety of data, and be good at making sense of novel, unseen situations. Rather than training an AI solution on a specific task with specific data, meta learning is about using a variety of data on a variety of tasks, focusing on having a learner that learns new tasks and a meta-learner that trains the learner. Lifelong learning is also an objective of meta learning AI. Techniques like Model Agnostic Meta Learning (MAML) are promising techniques that will get us to having robust meta learning models and libraries for use in commercial AI.

The application of quantum computing in AI will revolutionize how we build and train AI models. Quantum computers liberate us from the limitation of performing computation in binary digits of '0' and '1', Qubits, quantum bits exist in multiple states of '0' and '1' at the same time. In the AI context, this means exponentially faster training of AI models— but in the meta learning context, it would allow a much more efficient way of processing through a much wider search space of models, accelerating the learning-to-learn process and getting us closer to AGI. IBM has already shown quantum computing being applied to a machine learning classification task, while Rigetti Computing has demonstrated quantum computing being used to run a clustering algorithm.

AI Safety is a fast-emerging field within AI, which is aimed at building, testing and deploying AI systems that do not harm humanity. AI ethics and safety go hand, although safety should be viewed from the lens of testing an AI system for robustness — just like you would test an airplane, car or bridge before it is deployed to identify points of failure and make sure the system is failsafe. As AI systems get deployed more widely, especially in mission-critical systems which involve 'life and death' scenarios, the need for AI Safety will become much more acute. AI safety has become more of a priority given the way in which fake news on social media is suspected to have influenced voters in the last election. Today, there are only a handful of AI safety experts across Stanford, DeepMind, OpenAI, Oxford FLI, CHAI.

Strategic decision-making for business is about combining analytical capabilities with real-world experience, intuition and reasoning. Company boards thrive on a varied set of skills and experience that human board members bring to the table helping companies to make the right decisions. As AI blends symbolic and non-symbolic AI, fills in the gaps around reasoning and intuition, improves performance on NLG tasks, and incorporates explainability, it will become an invaluable tool for company boards. Company boards will not only use AI as a tool for augmenting decisions, but as a board member in its own right. Companies like Zignal Labs that help in competitive intelligence dashboards and Domo who integrate operational intelligence are expected to have a role to play in this decision-making domain.

Rather than have a homogeneous AGI (Artificial General Intelligence) or strong AI emerge in the next 10 years, we are more likely to see strong AI in specific domains. Language is one domain where strong AI is likely to emerge first in relation to understanding, translating, speaking and writing most languages that exist today. Google has shown how its Neural Machine Translation (NMT) technique can allow AI to translate a language pair it has never seen before, but can generalize learnings from other known language pairs. This is a significant step, although there are limitations around 12 specific language pairs. Using meta learning techniques, improved hardware, and fusing NLG we should be able to realize a universal AI agent that can translate almost any language and do it in real-time.