July 2019 · Philippe Lorenz and Kate Saslow

# Demystifying AI & AI Companies

## What foreign policy makers need to know about the global AI industry

Stiftung
Neue
Verantwortung

Think Tank für die Gesellschaft im technologischen Wandel

## Executive Summary

As a scientific discipline, Artificial Intelligence (AI) has been around since the late 1950s. During the past few years, however, a subdiscipline of AI known as Machine Learning (ML) has captured the imagination of investors, technologists and policy makers with break through achievements in the recognition of speech and images and its performance in playing complex games. Countries all over the world are adopting national strategies to position themselves to get a piece of the proverbial pie, setting up new public funds to support AI research, expert groups on AI ethics, or new initiatives to study the implications of AI on the future of work.

But AI is not only a hot topic for social and economic policy. It has also become an increasingly important topic in the foreign policy community. For example, as national governments try to strengthen their own AI innovation base, they are worried about domestic AI companies becoming targets of foreign takeovers. But while there has been a lot of talk about AI, the policy debates still suffer from a poor understanding of what AI actually is and means. This paper seeks to address this problem. It explains why most of the international conversation on AI is focused on Machine Learning (ML). In chapter one it breaks down ML into its core inputs – data, software, hardware, and human talent, explaining the importance of each of these inputs for building and applying AI. Chapter two introduces a model with three stages to assess the capabilities of AI companies: early stage ventures, larger companies using AI, and AI platforms.

The analytical framework introduced in this paper can help foreign policy analysts understand AI and develop tools to monitor trends, especially among companies. Foreign policy officers can leverage the distributed network of embassies for information gathering and analysis on the global distribution of ML's core ingredients and global AI companies. Furthermore, the framework allows foreign policy makers to inform government as a whole on the international AI developments and their economic and ultimately geostrategic implications.

## Table of Contents

| | |
|---|---|
| Introduction | 4 |
| Section 1: Understanding AI through its Core Components | 6 |
| What is AI? | 6 |
| Industrial Inputs for AI | 8 |
| Data | 9 |
| Software | 10 |
| Hardware | 13 |
| Talent | 18 |
| Section 2: Stages of AI Firm Maturity | 24 |
| Shared Traits of AI Companies | 24 |
| Stage I: Early Stage Ventures | 25 |
| Stage II: Domain Specific Application | 26 |
| Stage III: AI Platforms | 28 |
| Conclusion | 32 |

*SNV's Artificial Intelligence and Foreign Policy project was made possible by the generous support from the German Federal Foreign Office and the Mercator Foundation. The views expressed in this policy brief do not necessarily represent the official positions of the German Federal Foreign Office or the Mercator Foundation.*

*The authors would like to thank Matthias Plappert, Johannes Otterbach, Ben Scott, and Stefan Heumann for their inspiration and valuable feedback during the writing of this policy brief.*

## Introduction

The most important and at the same time least understood topic in modern geopolitics is Artificial Intelligence (AI). There are, of course, urgent and informed discussions among foreign policymakers about the dangers of autonomous weapons and the human rights implications of a wide range of AI technologies. There is a growing interest in the global competition for dominance in AI-driven technologies with most of the attention focused on the US and China. There is talk of creating a European model for AI – a third way that brings together both world-leading innovation and the values of European liberalism. But what exactly this European third-way might look like, what policies would lead to that outcome, and even what exactly we mean by "AI" are still unclear.

Despite the absence of a clear and unifying vision for where we are going, EU policymakers are forging ahead with economic policies to support national AI strategies. The pressure to act is intense. Within the last two years alone, different national governments have been scrambling to protect AI companies and startups from foreign takeovers. This has been most commonly attempted through tightening Foreign Direct Investment (FDI) screening mechanisms, or amending export control lists.

The term "AI" has appeared in FDI screening mechanisms within and outside of Europe, but not thoroughly defined. Instead, it is listed as an emerging technology that is deemed "critical." Other trends of economic protectionism have been seen with governments adding extremely specific technologies to increasingly growing export control lists, which fails to grasp the general-purpose nature of AI technologies. Currently, in order to put a technology on export control lists, it needs to be explicitly defined, including its uses. AI as a technology may be designed for one purpose originally, but because of its omni-use potential, could be applied in ways not fathomable by policy makers, or even the original designers. Both of these regulation strategies require significant capacity in policy makers tasked with defining what is critical and what should be protected.

In order to do these assessments properly and lay down the redlines effectively, policy makers need a deeper understanding of the complexity of AI as a technology. States will need a common nomenclature across government to collect knowledge and derive insight from monitoring global dynamics in the sector.

This paper argues that the foreign policy community needs a holistic understanding of AI as a starting point for more effective information gathering and monitoring, which can then inform policy-making – such as in the field of FDI screening or export controls. This work begins by correcting the near-total lack of clarity about what we mean in public policy debates when we say "AI." The term is thrown around as if its meaning were well defined and needed no further explication. This is an original sin that will corrupt all future policy decisions if we do not address it now. Consider the comparison with how we track the arms industry, specifying with great detail the precise characteristics of different munitions and tactical military equipment. Consider the scrutiny that we apply to pharmaceuticals and medical technology.

In contemporary policy debates, the term "AI" is often used to refer to any kind of advanced data processing technology that enables a novel product or service. But as a practical matter, we cannot build a policy agenda or a national strategy on such vague foundations. Instead, we should establish a definition of AI that is specific to how it works and the core technologies that make it different from anything that has come before. That is an evolving definition – and necessarily so. The definition of AI has changed periodically over the last few decades as technologies advanced and different research trends have come and gone. At the moment machine learning (ML) is the most dominant trend. And this is expected to remain the case for the foreseeable future. Thus, for the purpose of this paper, *AI will be defined as technologies that use variations of ML software and that are enabled through specialized hardware to power its most advanced forms.*

The analysis presented below lays out a basic framework for understanding and analyzing AI companies. Building on our foundational definition, there are two large categories of issues that we must identify and classify so that we may evaluate them for policy purposes. First, we need clear criteria for what kinds of technology products and services, which comprise the key inputs in the AI industry – unpacking this opaque term into clear industrial inputs – the focus of section 1 of this paper. Second, we need to identify stages of AI firm maturity to understand how these industrial inputs play out at the company and business level and affect global market dynamics. This is the focus of section 2.

**Section 1 – Industrial Inputs**
Data
Software
Hardware
Human Talent

**Section 2 – Stages of AI Firm Maturity**
- Early Stage Ventures
- Domain Specific Application
- AI Platforms

When we look at AI from the perspective of key inputs and maturity of development of AI companies, we gain several advantages. First and foremost we build capacity within policy making by breaking down a complex term into digestible pieces. By doing so, we present a framework that can leverage the information gathering activities and regional expertise of the foreign services and inform government as a whole on global developments of AI.

# Section 1: Understanding AI through its Core Components

### What is AI?

For the purpose of this policy brief, *AI will be defined as technologies that use variations of machine learning (ML) software and that are enabled through specialized hardware to power its most advanced forms.*

AI traditionally refers to the scientific pursuit of teaching machines to think like humans, or more simply, the automation of cognitive processes.[1] ML, not to be conflated with AI, is a subdiscipline of AI. ML adds to this traditional scientific approach, but we can demystify this aspiration by looking at how ML actually works.

In the most basic form, ML systems are made up of component parts: *data* and learning *algorithms* (code) are both fed into a software *model* through which inferences (probabilistic conclusions) are drawn, all of which is powered by special-purpose *hardware* (see Figure 1 below). The operational logic of ML is straightforward.[2] In image recognition, for instance, images of cyclists are specifically labeled according to the category “cyclists” by human annotators. In the training phase, such labeled data serve as answers (or targets) for the automation process. The goal of which is to get the software model to identify or “learn” patterns within the training data that will allow it to

1 François Chollet, *Deep learning with Python* (Shelter Island: Manning, 2018).

2 This comprehensive description is based on François Chollet, *Deep learning with Python* (Shelter Island: Manning, 2018).

map images of cyclists according to the specified category, with high levels of confidence. The data – hundreds or thousands of labeled images of cyclists – is first processed by algorithms that identify or "learn" patterns within that training data. These rules can then be applied to images of cyclists it has never seen before. In its real world application, the software model is then capable of assessing the probability that new data points fit the learned pattern (inference). For instance, this allows for automated image recognition of cyclists by software models that consume live camera-imagery in autonomous vehicles. This functional sequence in ML processes – learning through training data followed by inference about new data – is the basis for most AI technologies discussed in this paper.

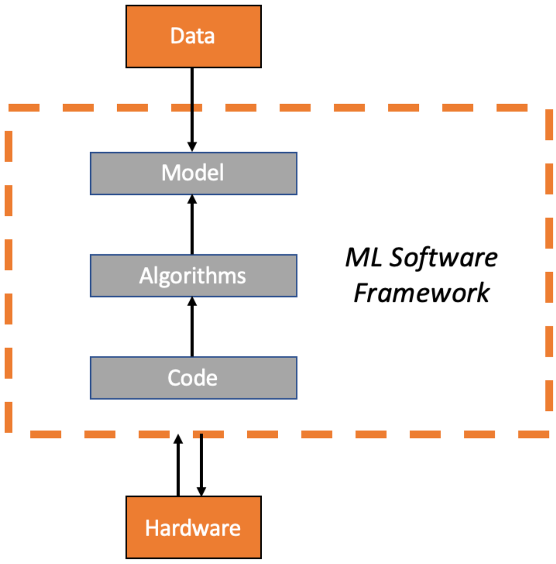


There are four basic categories[3] of ML algorithms or *learning styles*[4] that ML algorithms can adopt – supervised learning, unsupervised learning, rein-

3 Ibid.

4 Jason Brownlee, "A Tour of Machine Learning Algorithms," *Machine Learning Mastery*, November 25, 2013, https://machinelearningmastery.com/a-tour-of-machine-learning-algorithms/.

forcement learning and self-supervised learning.[5] *Supervised learning* is currently the most common application of ML. In supervised learning, the model learns to map the observed data to specific labels after processing a large set of example data annotated by human operators.[6] These systems are well suited to solve problems such as image classification, language translation and speech recognition – so called *perception*[7] tasks. Hence, ML is applied to industry use cases that aim to automate perception tasks, such as autonomous driving, speech and text translation, image recognition to name a few examples of industrial applications of supervised learning.

The scope of this paper will be primarily focused on supervised learning applications of ML using deep learning (DL) algorithms (more on this later), but other re-emerging categories of ML such as reinforcement learning must not be overlooked given the rapid pace at which ML as a field is advancing (particularly in Research and Development (R&D)[8]).

## Industrial Inputs for AI

AI technologies are comprised of different core inputs. The ensemble of industrial inputs – data, software, hardware, and human talent – together constitute what is generally referred to as "AI". Thus, monitoring AI means monitoring the interplay of these core inputs. Similarly policy making around AI means first and foremost understanding this interplay in order to make

5 *Unsupervised learning* is about finding correlations in data sets without relying on labels. This helps to make sense of data sets and as such is commonly applied to unstructured data. It is therefore often used as a precursor to applying supervised learning techniques to the data set. In reinforcement learning, an ML model receives input about its environment and makes decisions based on reward and punishment functions. For example, in a simulation space, autonomous driving software receives positive reinforcement for stopping at a red light and negative reinforcement for running it. After successful training, the learned behavior can be transferred to real-live applications which makes it a promising research avenue to pursue in data scarce environments such as in advanced robotics. In self-supervised learning, labels are not generated by human annotation but from the input data. For example, this technique is used to predict the next word in a sentence based on the previous words. For all, see François Chollet, Deep learning with Python (Shelter Island: Manning, 2018) and Jason Brownlee, "A Tour of Machine Learning Algorithms," *Machine Learning Mastery*, November 25, 2013, https://machinelearningmastery.com/a-tour-of-machine-learning-algorithms/.

6 Fabian Pedregosa et al., "Scikit-learn: Machine Learning in Python," Journal of Machine Learning Research 12, (October 2011): 2825-2830, http://www.jmlr.org/papers/volume12/pedregosa11a/pedregosa11a.pdf;
François Chollet, *Deep learning with Python* (Shelter Island: Manning, 2018).

7 "[...] extracting useful information from images, videos, sound, and more." Ibid. p. 315.

8 See for instance the accepted papers on reinforcement learning at the 2018 thirty-second Conference on Neural Information Processing Systems (NeurIPS), https://nips.cc/Conferences/2018/Schedule?bySubject=&selectedSubject=138.

informed decisions on restricting foreign investments or technology exports. For this reason, section 1 of the AI Monitoring Framework focuses on these industrial inputs.

### Data

The most advanced ML software is all but useless without data. Why the need for data? To paraphrase an analogy of ML developer and inventor of the Keras DL software framework – François Chollet: In today‘s information revolution, the steam engines are deep learning software frameworks and data is their new coal.[9] ML systems learn from representations in the data they process. This relationship is critical to the power of AI. As ML researcher Pedro Domingos wrote: „(...) a dumb algorithm with lots and lots of data beats a clever one with modest amounts of it.“[10]

Hence, there can be no ML – and consequently no ML ecosystem or industry – without vast quantities of data. However, it is not only about oceans of data, it is also about the quality of these data. Key criteria of high quality data consist first and foremost of the availability of so called *structured* data. That is data in clean, easy to follow, tabular or numeric format and typically labelled, for instance an image of a cat accompanied by the label „cat", or a particular wind speed numeric value tagged as „high winds". Additionally it is important to consider the volume of the data in relation to the size and complexity of the ML model. Sometimes even small data sets may be enough to adequately train a fairly small ML model[11] or they can be enlarged by adding synthetic[12] – artificially produced – data. Other important data traits to consider are variance, spurious relationships between variables, history and diversity of the data. Much like how a traditional statistician would need to understand the quality and characteristics of a data set, these are questions data scientists will ask when preparing the data to train ML models.[13]

9 François Chollet, *Deep learning with Python* (Shelter Island: Manning, 2018).

10 Pedro Domingos, "A Few Useful Things to Know about Machine Learning," Communications of the ACM 55, no. 10 (October 2012): 78-87, https://homes.cs.washington.edu/~pedrod/papers/cacm12.pdf.

11 François Chollet, *Deep learning with Python* (Shelter Island: Manning, 2018).

12 Ian Goodfellow et al., "Generative Adversarial Networks," arXiv:1406.26612014, (June 2014), https://arxiv.org/pdf/1406.2661.pdf.

13 Andrej Karpathy, "A Recipe for Training Neural Networks," Andrej Karpathy blog, April 25, 2019, http://karpathy.github.io/2019/04/25/recipe/.

## Software

Software for ML technologies is an interplay of specific algorithms, ML models, and data. Following the definition established in the introduction, all of the AI technologies discussed in this framework use some form of ML model. Contemporary ML models find patterns in input data. But, all the different forms of ML models are made up of code.

### Neural Networks

The specific code architecture responsible for present day ML applications such as self-driving cars and spam filters is an artificial neural network (NNs). NNs are layers of software that optimize data processing and inferential calculations. The importance of NNs for ML cannot be overstated. These are among the essential software building blocks of current cutting edge AI technologies. NNs are inspired by complex networks of neurons in biology.[14] Networks of neurons interact through the transmission of signals. In a network of different layers of neurons, there is a cascading effect to pass information along the electrical signal. Software NNs are merely a (loose) representation of this biological process.

Present day NN software architectures have their neurons arranged in stacks that can be several thousands of layers deep.[15] Consequently, such models are called deep neural networks. These models work by passing data through the network, much like the biological process. Certain neurons are activated when they detect specific features represented in the input data and consequently light up others and form a pattern associated with that piece of data. Through this process the information becomes more and more refined from very basic concepts to higher level representations of information. For example, in facial recognition software, the cascade moves from basic shapes to specific facial features such as an eye or chin. Depending on the pattern of neurons that were activated along the way through the NN, a certain neuron of the output layer will finally be activated and classify the information with a certain level of confidence (e.g. the output layer detects the face with a certain statistical probability). By way of this process, the NN acts as a sophisticated classifier of information.

14 Schmidhuber refers to cells discovered in the late 1950's and early 1960s in the visual cortex of cats as inspirational for deep neural networks. For a comprehensive overview of the history of Deep Learning see Jürgen Schmidhuber, "Deep Learning in Neural Networks: An Overview," arXiv:1404.7828, (April 2014), https://arxiv.org/pdf/1404.7828.pdf.

15 François Chollet, *Deep learning with Python* (Shelter Island: Manning, 2018).

**Deep Learning**

The capacity to draw representations from *deeper* stacks of layers transforms the basic ML model into a Deep Learning[16] (DL) model. The *deep* in DL is merely a reference to its architecture. DL leverages increasingly large and complex NNs, sometimes several hundreds or even thousands of layers *deep*. In turn, larger models can be trained on larger data sets which raises model accuracy. This is why deep neural networks are said to *scale* – their results get better with every additional data point. The head of AI at Tesla, Andrej Karpathy, believes that once a NN is well-configured, the only guaranteed way to improve its performance is by adding more data; and this will improve it indefinitely.[17] Well-configured means the NN is capable of learning and this requires exquisite engineering skills (more on the human factor below). But training larger models in turn requires more computational power. This is why specialized hardware must be combined with sophisticated software and large data sets in order to achieve the most advanced forms of AI – which currently are variations of deep learning.[18]

While this paper focuses on *deep* learning methods due to their ability to process sensory data and conduct perception tasks – tasks at which humans generally excel –, there is an equally important field of *shallow* ML known

16 The idea for using deep neural networks was around at least since the early 1980s. See Paul J. Werbos, "Applications of advances in nonlinear sensitivity analysis," Proceedings of the 10th IFIP Conference, New York City (September 1981): 762–770, http://werbos.com/Neural/SensitivityIFIPSeptember1981.pdf; Yann LeCun, "Une procédure d'apprentissage pour réseau a seuil asymmetrique (a Learning Scheme for Asymmetric Threshold Networks)," Proceedings of Cognitiva 85 (June 1985): 599-604, http://yann.lecun.com/exdb/publis/pdf/lecun-85.pdf; David E. Rumelhart et al. "Learning representations by back-propagating errors," Nature 323 (October 1986): 533–536, https://www.iro.umontreal.ca/~vincentp/ift3395/lectures/backprop_old.pdf.

17 Andrej Karpathy, "A Recipe for Training Neural Networks," *Andrej Karpathy blog*, April 25, 2019, http://karpathy.github.io/2019/04/25/recipe/.

18 To acquire a comprehensive understanding of the specialized hardware involved in ML, we encourage you to read the paper by Tim Hwang, "Computational Power and the Social Impact of Artificial Intelligence," SSRN (March 2018), http://dx.doi.org/10.2139/ssrn.3147971.

as gradient boosting.[19] This other realm of ML becomes important when tracking global talent distribution, and software and hardware use, although it is outside the scope of this paper.

### ML software frameworks

Each instance of a new AI technology does not produce an entire ML software framework from scratch. Once a classifier (learning algorithm) has proven proficiency in a specific application area, it is likely to find patterns in other data. Many ML or DL software frameworks (also known as *repositories*) are available for license or are available for free as open source code.[20] These offer pre-trained algorithms with a proven track record for accurate results in specific tasks. This is exactly what ML software frameworks such as TensorFlow, Keras, PyTorch, Caffe, and others provide to their users: the ability to either code specific ML algorithms from scratch, or to reuse pre-existing ones in a copy-and-paste fashion to apply them to new data in order to find new answers to problems. These repositories allow ML practitioners to acquire a specific application as a module that can plug into a larger model or product and allow them to choose from different programming languages such as Python or R.[21] This modularity of ML software further explains the general purpose applicability of ML.

19 Apart from DL approaches, data scientists and ML practitioners frequently apply an ML approach that relies on so called gradient boosting machines (shallow learning). It is a powerful approach for building predictive models. Gradient boosting machines rely on a combination of decision trees (*so called weak learners*) that in their ensemble are able to make better predictions. The technique is best applied to problem domains where lots of structured data is available, see François Chollet, *Deep learning with Python* (Shelter Island: Manning, 2018). Gradient boosting machines often get lost in the current focus on perception tasks – for which deep neural networks are better equipped. For a comprehensive overview on the gradient boosting approach, read the introduction by Jason Brownlee, “A Gentle Introduction to the Gradient Boosting Algorithm for Machine Learning,” *Machine Learning Mastery*, September 9, 2016, https://machinelearningmastery.com/gentle-introduction-gradient-boosting-algorithm-machine-learning/; another good resource is Ben Gorman’s explanation of gradient boosting “A Kaggle Master Explains Gradient Boosting,” *No Free Hunch*, January 23, 2017, http://blog.kaggle.com/2017/01/23/a-kaggle-master-explains-gradient-boosting/; For an explanation why decision trees fail at perception tasks, read the Quora answer provided by David Johnston “In what situations do neural networks outperform gradient boosting and random forest models on regular numeric and categorical data (non-image or text data) if any?,” *Quora*, June 11, 2018, https://www.quora.com/In-what-situations-do-neural-networks-outperform-gradient-boosting-and-random-forest-models-on-regular-numeric-and-categorical-data-non-image-or-text-data-if-any/answer/David-Johnston-147.

20 A popular open source ML framework is sci-kit learn, https://scikit-learn.org/stable/index.html.

21 Wikipedia, The Free Encyclopedia, s.v. “Comparison of deep-learning software,” https://en.wikipedia.org/wiki/Comparison_of_deep-learning_software.

This is why investors chase ML startups that provide industry solutions. It is their general applicability that makes NNs such a powerful tool. Depending on the objective, the overall framework may be a hybrid of many pre-existing DL frameworks integrated with new code into a novel application or service. In this way the simple building blocks of ML can be combined in ways that are extremely complex. But for the purposes of this framework it is necessary to identify and classify the core elements that serve as industry inputs. Software for ML technologies is an interplay of a specific algorithm, a model, and training data. The complexity of how these components interact forms the software used to power ML technologies.

## Hardware

Advancements in hardware in the very recent past have driven developments in ML. Steady increases in the availability of compute have enabled NNs to find patterns in ever larger troves of data. While CPUs (central processing units) enabled personal computers, machine learning requires special purpose chips to maximize performance. For the past decade, ML practitioners relied on chips for graphical processing provided by companies such as Nvidia and AMD, whose main audience was the gaming industry. Now, however, ML practitioners demand better, perfectly tailored chips for an array of ML use cases, which has spawned the creation of an ML hardware industry that adds to the complexity of the global semiconductor value chain. Thanks to their financial capabilities, big tech platforms have now become champions in this industry. Not only do the big platform driven tech companies have exclusive access to large datasets, they also invest heavily in developing their own hardware. These chips can be used for on-demand cloud computing, which is critical for both machine learning research and industry, and edge computing, which allows for on-device machine learning.[22]

### The importance of parallel computing for ML

Three distinct hardware categories are responsible for driving developments in ML: GPUs (graphics processing units), ASICs (application specific integrated circuits), and FPGAs (field programmable gate arrays).

22 Edge computing enables decentralized data processing – at the edge of the (local) network or „nearer the data“. Allan Alasdair, “Say “Hello” to Google Coral – Edge TPU, Google’s custom ASIC for machine learning has arrived,” *Medium*, March 4, 2019, https://blog.hackster.io/say-hello-to-google-coral-cdbb49183864; Tom Simonite, “Apple’s ‘Neural Engine’ Infuses the iPhone With AI Smarts,” *Wired*, September 13, 2017, https://www.wired.com/story/apples-neural-engine-infuses-the-iphone-with-ai-smarts/.

A GPU is a special purpose chip capable of processing computational tasks in parallel. With GPUs thousands of computer cores share the processing task among each other, compared to CPUs, where the infrastructure consists of much fewer, but faster cores. The latest desktop PCs have 12 cores.[23] The latest GPUs have 5,120 cores.[24] A CPU is thus sharply limited in handling computational tasks in a parallel manner.

ML requires parallel computing power, because it significantly reduces the amount of time needed to train a model. Parallel computing has been the key enabler for increasingly *deeper* learning processes. The breakthrough moment for GPUs and deep neural networks happened in 2012 when researchers trained a CNN[25] (a convolutional neural network, a form of DL especially suited for image recognition) to classify images using GPUs, and were able to complete the training process in about six days, compared to the weeks it would have taken with CPUs.[26] This allowed the CNN (consisting of 650,000 neurons) to be trained on an even larger dataset (around 1.2 million images).[27] This made the model's image classifier tremendously more accurate than other systems not using NNs.[28] This example highlights the interplay between the hardware, learning algorithms, and data. The researchers behind this advancement hinted toward this interplay saying, "All of our experiments suggest that our results can be improved simply by waiting for faster GPUs and bigger datasets to become available."[29] Both have hap-

23 For a list of the latest desktop CPUs, see "Dodeca-Core (12 cores)," WikiChip, accessed June 12, 2019, https://en.wikichip.org/wiki/dodeca-core.

24 Nvidia's Tesla V100. Its specifications are available here, "Nvidia's Tesla V100 GPU Accelerator," nvidia, accessed June 12, 2019, https://images.nvidia.com/content/technologies/volta/pdf/tesla-volta-v100-datasheet-letter-fnl-web.pdf.

25 Strongly influenced by foundational work on convolutional neural networks laid down by scientists around Yann LeCun in 1998. Yann LeCun et al., "Gradient-Based-Learning Applied to Document Recognition," Proceedings of the IEEE 86, no. 11 (November 1998): 2278-2324, http://yann.lecun.com/exdb/publis/pdf/lecun-01a.pdf.

26 Alex Krizhevsky et al., "ImageNet Classification with Deep Convolutional Neural Networks", Proceedings of the 25th International Conference on Neural Information Processing Systems, Volume 1 (December 2012), https://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf.

27 Ibid.

28 For more detail on the AlexNet we encourage you to read the explanation of Convolutional Neural Networks by Adit Deshpande, "The 9 Deep Learning Papers You Need To Know About (Understanding CNNs Part 3)", *Adit Deshpande Blog*, August 24, 2016, https://adeshpande3.github.io/adeshpande3.github.io/The-9-Deep-Learning-Papers-You-Need-To-Know-About.html.

29 Alex Krizhevsky et al., "ImageNet Classification with Deep Convolutional Neural Networks", Proceedings of the 25th International Conference on Neural Information Processing Systems, Volume 1 (December 2012), https://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf.

pened since 2012.

**Special purpose ML hardware**
ML can be subdivided into two phases: training and inference.[30] During training, the NN learns properties of the data and begins to classify the data according to patterns it finds. Tweaking deep NN to find patterns in very large datasets is computationally intensive. ML experts prefer GPUs over other special purpose chips for *training* a model because their architecture is easier to reconfigure than other special purpose ML hardware.[31] The next step is *inference*, where the model can be applied to real-world problems such as translating a document or prompting your digital assistant. While GPUs can be used for inference, ASICs and FPGAs are preferred as they are designed to process ML specific use-cases, which makes them run faster and consume less energy while processing live-data, therefore best designed for the inference phase.[32] While ASICs are truly application specific and cannot be reprogrammed after manufacturing (hence their name), FPGAs can be reconfigured after manufacturing according to the desired industry or application, however not as easily as GPUs can be.[33]

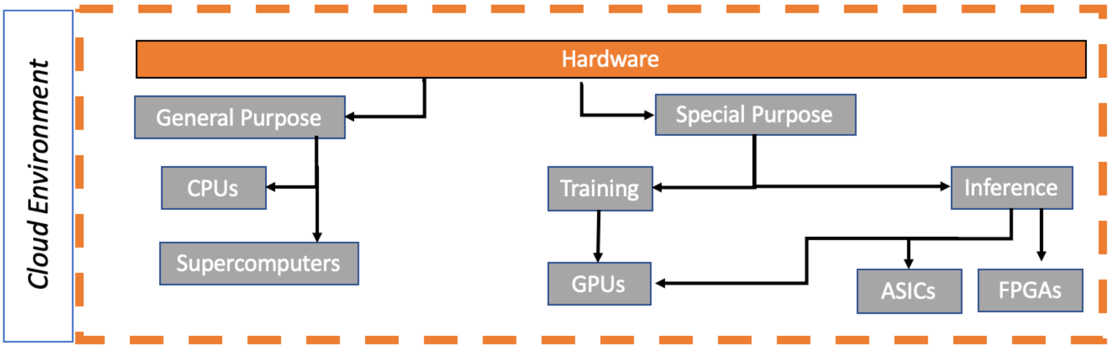


Tech companies such as Google and Amazon provide ML as-a-service on

30 Tim Hwang, "Computational Power and the Social Impact of Artificial Intelligence," SSRN (March 2018), http://dx.doi.org/10.2139/ssrn.3147971.

31 Ibid.

32 Although GPU manufacturers have reacted to this by devoting specific cores of their GPU's for inference (accelerators). See for instance Nvidia's Tesla V100 Tensor Core GPU, "Nvidia's Tesla V100 Tensor Core GPU," nvidia, accessed June 13, 2019, https://www.nvidia.com/en-us/data-center/tesla-v100/. It is providing 640 so called Tensor Cores. But, they share the trade-off with ASICs and FPGAs, as they are only limited programmable.

33 "Field Programmable Gate Array (FPGA): What is an FPGA," Xilinx, accessed June 13, 2019, https://www.xilinx.com/products/silicon-devices/fpga/what-is-an-fpga.html; see also What are field-programmable gate arrays (FPGA): https://docs.microsoft.com/bs-latn-ba/azure/machine-learning/service/concept-accelerate-with-fpgas.

their respective cloud platforms (translators, assistants, etcetera) and have begun to produce their own special purpose hardware that enable these services. Google's Tensor Processing Unit (TPU)[34] and Amazon's Inferentia[35] (both ASICs) and Intel's FPGAs to power specific Microsoft services on its cloud infrastructure are examples of tech firms developing tailor-made hardware to suit the services they provide.

**AI research at a juncture**

Access to large compute is not only a necessity for private sector companies, but it becomes more and more necessary for academic research as well. As cutting-edge AI research is increasingly dependent on access to large amounts of compute, AI research is at a juncture – separating researchers that can access large compute infrastructures from those who cannot. Access to high-performance computing infrastructures explains some of the most impressive recent developments in several sub-disciplines of machine learning. This is now achieved by actors outside classical academic research institutions, as seen with Google DeepMind's or OpenAI's reinforcement and unsupervised learning models.[36]

This could further tilt the advantage to tech-platforms that leverage their own compute infrastructure for applied research.[37] But at the same time, tech companies are now also advancing basic research. In the past, basic research of a general purpose technology would have predominantly been

34 Kaz Sato, "What makes TPUs fine-tuned for deep learning?," *AI & Machine Learning–Google Cloud Platform–TPUs*, August 30, 2018, https://cloud.google.com/blog/products/ai-machine-learning/what-makes-tpus-fine-tuned-for-deep-learning.

35 Gareth Halfacree, "Amazon announces high-performance Inferentia ML chip," *bit-tech*, November 29, 2018, https://bit-tech.net/news/tech/cpus/amazon-announces-high-performance-inferentia-ml-chip/1/; "AWS Inferentia", aws, accessed June 13, 2019, https://aws.amazon.com/machine-learning/inferentia/?nc1=h_ls.

36 David Silver et al., "A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play," arXiv:1712.01815, (December 2017), https://arxiv.org/pdf/1712.01815.pdf; "AlphaStar: Mastering the Real-Time Strategy Game StarCraft II," *DeepMind Blog*, January 24, 2019, https://deepmind.com/blog/alphastar-mastering-real-time-strategy-game-starcraft-ii/; OpenAI, Marcin Andrychowicz, Bowen Baker et al., "Learning Dexterous In-Hand Manipulation," arXiv:1808.00177, (January 2018), https://arxiv.org/pdf/1808.00177.pdf; Alec Radford, "Language Models are Unsupervised Multitask Learners," (February 2019), https://d4mucfpksywv.cloudfront.net/better-language-models/language_models_are_unsupervised_multitask_learners.pdf.

37 This is something that OpenAI's policy branch has often pointed out. See Jack Clark's weekly newsletter on new technological breakthroughs in AI and the resulting policy implications: *Import AI*, https://jack-clark.net/. For the respective newsletters that explain why access to „big compute" is crucial to being able to perform cutting-edge AI research: among others, see Import AI #121, #132, #135, #136.

pursued at academic research facilities.[38] Now, we witness corporate research teams at big tech firms significantly advancing AI research.[39] And this is achieved partly through exclusive access to large compute.

**Compute per model**

The cost of compute and algorithmic design share an important relationship. Analyses show that in many current domains, more compute „seems to lead to predictably better performance, and is often complementary to algorithmic advances."[40] Thus, looking at the compute needed for an individual model can be an effective indicator to understand emerging ML technologies. Additionally, it serves as a way to recognize who is innovative in designing sophisticated algorithms that require less compute.

In order to gain an understanding of how much capital is needed to produce a specific ML application or reproduce a specific ML experiment, policy makers could calculate how much one company pays to train a model. Reference points in every calculation to determine the costs involved to produce ML models are the elapsed training time and the specific hardware infrastructure used for training and running the model. For example, take DeepMind's latest attempt of creating an autonomous agent capable of defeating humans at the ancient board game Go – AlphaGo Zero. The system was trained on an ensemble of GPUs and CPUs and engaged in self-play powered by TPUs. The experiment lasted for the duration of 40 days with an estimated cost in

38 The Bell Laboratories are the most famous exception to this rule – developing, among others, the programming languages C and C++, and producing an overall of nine Nobel Prizes. See, Wikipedia, The Free Encyclopedia, s.v., "Bell Labs", https://en.wikipedia.org/wiki/Bell_Labs. One of the pioneers of deep learning, Yann LeCun, did also advance ML as a discipline during his time at the AT&T Bell Laboratories from 1988-2002. See for instance the application of convolutional neural networks (CNNs) to shape recognition tasks in Yann LeCun et al., "Handwritten digit recognition with a back-propagation network," Advances in Neural Information Processing Systems 2 (1989): 396-404, https://papers.nips.cc/paper/293-handwritten-digit-recognition-with-a-back-propagation-network.pdf. See a video showing the CNN in action at recognizing handwritten digits (and get a feel for what cutting-edge ML research looked like back in 1993) "Convolutional Network Demo from 1993," *YouTube*, June 2, 2014, https://www.youtube.com/watch?v=FwFduRA_L6Q&feature=youtu.be.

39 Among others, Google AI, https://ai.google/, Google DeepMind, https://deepmind.com/research/, Facebook AI Research, https://research.fb.com/category/facebook-ai-research/, Baidu Research, http://research.baidu.com/, Alibaba Damo Academy, https://damo.alibaba.com/labs/.

40 Dario Amodei and Danny Hernandez, "AI and Compute," *OpenAI Blog*, May 16, 2018, https://openai.com/blog/ai-and-compute/.

the range of 25 to 35 million US dollars.[41] Thus, the sheer cost of compute needed to train and run an ML model is a huge barrier to entry for small AI firms, research labs and academia.

This example illustrates just how important talent is for the training process. The better the training algorithm, the less compute needed for training, simply due to successfully training models the first or second time around, rather than needing to revisit the drawing board constantly. Both hardware and algorithmic engineering can therefore contribute to the efficiency of a model.

### Talent

There is no "AI" without the people who develop and apply it. Thus this subsection is concerned with understanding the importance of human talent for ML research. Not only is their skill set relevant for continuously pushing the academic boundaries, but also for turning that knowledge into industry applications or software. This is where Chinese tech platforms seem to have an edge over their US counterparts[42], and where companies from the European Union seem to struggle the most. Similar to the previous sub-chapters, additional characteristics are described to help illustrate the importance of ML talent for academic and industry research.

The first thing that policy makers need to understand about people that work in ML is that, currently, such experts are scarce – and the demand for these experts is high. Different papers analyzing the global ML talent pool cite different numbers, but the impact is the same. A report issued by Tencent Research from 2017 puts the number of global ML talent at about 300,000: roughly one third working as researchers in academia, and two thirds pur-

41 Aidan Rocke, "The true cost of AlphaGo Zero," *Kepler Lounge*, March 24, 2019, https://keplerlounge.com/artificial/intelligence/2019/03/24/alpha-go-zero.html; Lisa Calhoun, "Google Artificial Intelligence ‚Alpha Go Zero' Just Pressed Reset On How To Learn," *Inc.* June 12, 2019, https://www.inc.com/lisa-calhoun/google-artificial-intelligence-alpha-go-zero-just-pressed-reset-on-how-we-learn.html; Dan Huang, "How much did AlphaGo Zero cost?," *Dansplaining*, accessed June 13, 2019, https://www.yuzeh.com/data/agz-cost.html.

42 "Speaker Interview – Kai-Fu Lee: Chairman and CEO of Sinovation Ventures," interview by Jordyn Dahl, *Speaker Interview*, Young China Watchers, September 2018, http://www.youngchinawatchers.com/speaker-interview-kai-fu-lee-chairman-and-ceo-of-sinovation-ventures/.

suing careers in industry.[43] In a recent report published by Tsinghua University, the number of global ML talent is pinned down at 204,000 people.[44]

**The specific skill set of ML experts**

Building ML algorithms requires a specific skill set: various programming languages, mathematics, data science, and data engineering, to name a few. In addition to pure computer science and statistics, a basic understanding of hardware components and architecture are of additional importance – especially given the relevance of compute for DL. This skill set is complex, and only rarely combined in one person. Therefore, ML specialists work in teams that consist of data scientists that design particular prediction models and data-engineers that help with the architecture necessary to reach the prediction levels. ML specialists themselves are concerned with programming ML architectures in DL software frameworks. Apart from an overall proficiency in applied mathematics such as linear algebra, this includes taking into account the specific data set that will be used to train the NN, subdividing it into a training and testing sets, choosing the classifier, and choosing the appropriate model for the data. In addition to data scientists and data engineers, ML specialists are supported by software developers that help with coding the final application or software. This pyramid structure – with ML specialists at the top, software developers in large numbers at the bottom, and data-scientists and -engineers in between – is the signature workforce structure of advanced tech clusters such as Silicon Valley. European policy makers must be aware of this. In order for European companies to be able to produce competitive ML products and services, appropriate talent to compose this pyramid must be made available to companies that seek to produce ML services and products in Europe.

**Number of PhDs in ML**

Top talent in ML is often associated with advanced degrees in computer sciences. Therefore mapping the number of PhDs in ML is a good way of keeping track of the global talent stock. Researchers around Jean-Francois Gagné from Canadian startup ElementAI are doing this in their yearly reports. In 2018 they found 22,000 PhDs "who are capable of working in AI research and applications."[45] Working in applications means that these experts pos-

---

43 James Vincent, "Tencent says there are only 300,000 AI engineers worldwide, but millions are needed," *the Verge*, December 5, 2017, https://www.theverge.com/2017/12/5/16737224/global-ai-talent-shortfall-tencent-report.

44 Tsinghua University, "China AI Development Report," July 2018, http://www.sppm.tsinghua.edu.cn/eWebEditor/UploadFile/China_AI_development_report_2018.pdf.

45 Fedor Karmanov and Simon Hudson, "Global AI Talent Pool Report 2018", *jfgagne*, 2018, https://jfgagne.ai/talent/.

sess the adequate software development capabilities to work either in industry or business.[46] In order to obtain these results, the researchers data-mined available information on LinkedIn. To discover the top-tier experts among those 22,000 PhDs, the researchers included search parameters that track their participation at well-known AI conferences. Given these advanced parameters only 5,400 PhDs belong to this elite group of researchers "well-versed enough in the technology to work with teams to take it from research to application"[47].

The 2019 edition is a good resource to understand global ML talent concentration and movement.[48] It goes deeper into the analysis by extending the LinkedIn search parameters and by also broadening the AI conferences participation analysis. 22,400 PhDs are found to have published at 21 AI conferences. Among those, 4,000 researchers have contributed work that had "major impact on the overall field measured by citation counts" in 2017-2018. Not surprisingly, US researchers are the most numerous among those 4,000 top-tier researchers (1,095), followed by researchers from China (255), the UK (140), Australia (80) and Canada (45). However, it also important to take biases in the data set into account as LinkedIn is much more widely used in the US than in Europe and even less so in countries like China. While these reports chose to look at PhD positions as a way to zoom in on ML talent, not all experts have PhDs. Conversely, many experts deliberately choose to leave academia for applied opportunities. What is universal to ML talent, however, is its current scarcity.

**Global ML talent distribution**

Other reports seek to map the global distribution of existing ML experts by the impact they have on the field – in this case publications and citations. In an analysis published by Tsinghua University in 2018, ML talent is subdivided into two separate categories: "international AI-talent"[49] and "international top-AI talent". The latter takes into account experts' h-factors – an indicator of academic impact that relates a researcher's publications to his or her citations. Researchers counted as top-AI need to have an "H-index score

46 Ibid.

47 Ibid.

48 Jean-François Gagné, Grace Kiser, Yoan Mantha, "Global AI Talent Pool Report 2019", *jfgagne*, 2019 https://jfgagne.ai/talent-2019/.

49 People that belong to the first category are described as "researchers possessed of creative research ability and technical expertise in their research area and active in AI research with innovative outcomes." Tsinghua University, "China AI Development Report," July 2018, http://www.sppm.tsinghua.edu.cn/eWebEditor/UploadFile/China_AI_development_report_2018.pdf.

[that] ranks [them] among the top 10% of international AI researchers."[50] After having established these categories, the availability of both talent groups are mapped according to their distribution by (among others) global regions, universities, and enterprises.

In total numbers, the US is leading the report's statistics with about 28,536 total AI-talent, followed by China (18,232) and India (17,384). European countries are ranked behind India, and lead by Germany (9,441), followed by the UK (7,998), and France (6,395).[51] European countries perform especially well when top-AI talent is taken into account. In this ranking, it is again the US (5,158) that is placed first. European countries score right behind, led by the UK (1,177), followed by Germany (1,119), France (1,056) and Italy (987). China places sixth with 977 "top-AI" experts. Although there may not be a talent deficiency in Europe, there are pull factors that draw this talent to other countries.

Global distribution of talent is affected by the availability of the other two components of ML: software and hardware. All of the previous statistics show that talent flocks to certain countries more than others, and this can be explained by the sophistication and availability of software and data environments and hardware infrastructures with which AI talent can work. Sophisticated hardware infrastructures and data environments are more than anything an indicator of financial strength. This financial strength in industry generally means incomparable salaries for AI talent, which also is a huge draw from academia. According to an analysis in the New York Times, AI experts – either with a PhD or in the early years of their careers – have starting salaries within a range of 300,000-500,000 US dollars or more and sometimes own company shares.[52]

**Brain drain**

As long as DL is the driving force behind developments in ML technology, the availability of the specific ingredients that we have already introduced – software, talent, and hardware – will be determining factors not only for creating industry applications but also for performing academic research.

50 ibid.

51 ibid.

52 Cade Metz, "A.I. Researchers Are Making More Than $1 Million, Even at a Nonprofit," *The New York Times*, April 19, 2018, https://www.nytimes.com/2017/10/22/technology/artificial-intelligence-experts-salaries.html.

ML talent is scarce, well educated, highly mobile[53] and therefore capable of moving across borders. The attractiveness of the AI ecosystem has a major impact on where talent chooses to relocate. In addition to these core components, access to a community of peers that pursue interesting research avenues in ML is equally important.

Brain drain manifests in two ways. First, it can pull researchers out of academia and drive them into the private sector, since private sector research labs currently provide the adequate data environment and hardware infrastructure necessary to pursue cutting-edge ML research. Brain drain can also be characterized by researchers and industry specialists alike that leave their countries in order to seek better research or applied industry opportunities abroad. The first – sectoral – brain drain may at least prove beneficial to business and industry adoption of ML technologies. For instance, the 2018 Turing Award recipients all have affiliations with the largest US tech platforms and are responsible for their adoption of ML technologies.[54] This raises concerns that a lot of top research is moving from publicly funded institutions into corporate entities where progress and research processes are usually less open and accessible to the broader public. The second – cross border – brain drain can put domestic and regional AI bases at risk.

Section 1 gave an overview of the key inputs in the AI industry consisting of data, software, hardware and human talent. Through these different inputs, it is also clear why the scarce human talent follows flows of hardware, data, and of course, capital. Given the complexity of machine learning technologies, access to these ingredients is not only an advantage but a must.

With these inputs, the general purpose utility of ML can supercharge business processes and industries. What is more: industry incumbents can embrace ML while new market entrants may use ML to tap into pre-existing business cases – see for instance tech companies' pivoting towards the provision of mobility services. Because powerful actors in an AI ecosystem may come in different shapes and sizes, it is important for the foreign policy community to understand the different stages of development of an AI company,

53 "(...) with about one-third of of researchers working for an employer based in a country that was different from the country where they received their PhD", Jean-François Gagné, Grace Kiser, Yoan Mantha, "Global AI Talent Pool Report 2019", *jfgagne*, 2019 https://jfgagne.ai/talent-2019/.

54 Cade Metz, "Turing Award Won by 3 Pioneers in Artificial Intelligence," *The New York Times*, March 27, 2019, https://www.nytimes.com/2019/03/27/technology/turing-award-ai.html.

and how companies from different stages can interact within an economy in a way that illustrates the strength of that AI ecosystem.

# Section 2: Stages of AI Firm Maturity

After giving foreign policy makers an introduction to the basic components of ML systems, it is also important to assess the strengths and capabilities of AI companies. In addition to answering the question *what is AI?*, analysts should be able to identify and evaluate particularly interesting AI firms. The categorization introduced in this section, along with a basic understanding of AI according to the core components, will allow foreign policy makers to inform the whole of government on important developments regarding AI companies.

Section 2 lays out these stages of AI firm maturity by first explaining the necessary traits shared by AI firms (in order to qualify as an AI firm), and then by describing additional classifiers that set each stage apart from the next. The presence or absence of these traits is clearly shown at the beginning of each stage with a checklist. These checklists are based on different characteristics of companies that reflect the depth of adoption and innovation of ML in their products and services, as well as the degree to which they may pursue development of in-house ML technologies. These stages of AI firm maturity are not to say that a company moves through them chronologically, but rather that depending on access to the necessary inputs, an AI company can be classified in one of three ways.

We differentiate between three types of AI companies: I) early stage ventures, II) domain specific application, and III) AI platforms. The difference between these stages comes not only from the access to the four components presented above, but also the ability to *apply* these inputs and *scale* them. This is where access to capital and domain-specific expertise become crucial.

## Shared Traits of AI Companies

Even if AI companies differ in their maturity and varying access to specific ML ingredients, they must share two commonalities in order to qualify as *AI companies*: they must focus their research and development efforts on creating specific core ML driven products, and their core revenue stream must be based on these ML driven applications. To fulfil the first trait, a company could, for example, focus on natural language processing (NLP) algorithms. This may even present novel categories of products and services as seen with how Big Tech companies now offer an array of AI assistants largely enabled by NLP technologies. To assess whether an AI company fulfils the second trait, policy makers must consider the company's revenue stream. This revenue stream would need to feed from the provision of core ML technologies embodied in the products or services it provides. Only by fulfilling these two traits can a company identify as an AI company.

Furthermore, there are other classifiers that companies at distinct stages of maturity have access to or lack. These are addressed in depth in the following subsections as well as illustrated with check lists.

### Stage I: Early Stage Ventures

| Characteristics of an AI Firm? | Stage I |
|---|---|
| Pursuit of ML idea or research question | ✓ |
| Seed capital | ✓ |
| Revenue stream based on ML in products/services | ✓ |
| Access to talent | ✓ |
| Basic data environment | ✓ |

In the early ventures, key actors conceptualize ideas and develop research questions pursuing core AI research activities. A group of people that form a spin-off from academia into a startup with an approach on how to implement specific ML-algorithms, or young companies that pursue the same goal could serve as examples. Such actors pursue basic or use-case specific research into core ML technologies and are located at the very beginning of a spectrum that involves different actors trying to solve a specific problem with ML tools. Once they reach out to public or private investors and secure venture funding, then ideally they start to produce a revenue stream based on these ML technologies. At this point, these incorporated ventures qualify as AI companies by possessing the two shared traits.

Startups in stage I may have the software engineering skills from STEM backgrounds or formal computer- and data-science education needed to engage in building ML systems to solve specific problems. They may have established strong links to research networks with other peers that work in academia or industry. They can likely apply their skills to different problem sets in target industries, but usually lack domain-specific knowledge that comes from industry-specific data. Gaining access to industry data will become crucial for these ventures in order to qualify as stage II. Once the venture's ML models are trained on domain specific data – proprietary or open source – specialized hardware may become a crucial input in order to process the data, but on-demand compute will largely be available for service on cloud platforms.

Hence, actors in stage I may have access to specific talent and a basic data environment, but they lack domain-specific data sets and the talent pyramid characteristic of AI companies. Early stage ventures must be monitored by policy makers as they strengthen an innovation base and potentially attract further capital.

### Stage II: Domain Specific Application

| Characteristics of an AI Firm? | Stage I | Stage II |
|---|---|---|
| Pursuit of ML idea or research question | ✓ | ✓ |
| Seed capital | ✓ | ✓ |
| Revenue stream based on ML in products/services | ✓ | ✓ |
| Access to talent | ✓ | ✓ |
| Basic data environment | ✓ | ✓ |
| Venture Capital | | ✓ |
| Domain-specific knowledge or data | | ✓ |
| In-house AI accelerators | | ✓ |
| Targeted acquisitions of stage I companies | | ✓ |
| Talent pyramid | | ? |

Companies in the domain specific application stage (stage II) pursue their R&D efforts on creating specific core ML driven products and services and their revenue stream is based on these ML driven applications, thus qualifying as AI companies. In contrast to stage I, stage II companies apply their ML technologies to industry specific use cases or problem sets. In order to do so, the most critical input they need is domain specific knowledge and ultimately domain specific data needed to train and run their models.

But, domain specific knowledge is hard to come by, often held exclusively by established actors in the specific industry in question.The data in question must be high in volume and high in quality – meaning, structured or labeled, low in variance, diverse, and data that provides a certain history. Industry data is highly dynamic, so domain specific knowledge is important to constantly monitor, clean and organize the data in question. Lastly, it is also important to run quality assurance on the performance of ML models, which is best done by industry insiders with this domain knowledge.

Apart from data, there are other crucial inputs for such AI companies at this particular stage. They need access to hardware in order to train their NNs. Adequate funding provided, this can be achieved by renting GPU or other specialized ML hardware contingents via cloud services.[55] Apart from hardware, this stage of development also requires industry insiders who understand current problems within an industry and where innovation of ML can unleash new potential. If this insider knowledge is leveraged through inter-personal networks to research facilities, then stage II companies can introduce new revenue streams powered by ML technologies. An example of this is PrognoNetz, a research partnership between Karlsruhe Institute of Technology and industry partners – in this case meteorologists and power grid providers – whose algorithms were able to use historical weather data (gathered by the German Weather Service) and sensory data from the grid to first better understand the conditions around the power cables and then to optimize current carrying capacity and power grid infrastructure.[56]

In order to get access to domain specific knowledge, stage II companies with tech expertise will try to attract (or poach) talent from the industry domain, to which they tailor their models. Companies that have lots of domain specific industry data, but that lack ML expertise might think about creating AI accelerators in order to strengthen research efforts and the adoption of ML to products and services within the organization. These AI accelerators are also beneficial for the creation of an AI talent pipeline. An example is the

55 See a list of cloud computing providers, compiled by Jeff Hale, “deep-learning-cloud-providers,” *GitHub*, November 15, 2018, https://github.com/discdiver/deep-learning-cloud-providers/blob/master/list.md. For more information on the costs of various cloud computing providers, see Jeff Hale ”Best Deals in Deep Learning Cloud Providers,” *Medium*, October 29, 2018, https://towardsdatascience.com/maximize-your-gpu-dollars-a9133f4e546a.

56 “PrognoNetz – An Intelligent Ampacity Forecast for Overhead Lines”, Karlsruhe Institute of Technology, accessed June 13, 2019, https://www.itiv.kit.edu/english/6518.php.

German auto supplier Bosch. The company has recently announced its AI strategy which includes hiring 3,000 AI experts in the next two years.[57] These experts will be housed by the Bosch Center for Artificial Intelligence, which was established in early 2017.[58]

Stage II companies can also take a more aggressive path to gaining ML expertise through acquisitions. This type of strategy was one driving force of tighter FDI screening mechanisms all over the world. Activities within stage II companies around an AI talent pipeline – both built in-house or poached through acquisition – merit attention from policy makers in order to understand where innovation bases are strengthened and how a growing AI talent pool may allow companies to transcend traditional market boundaries and tap into new business cases.

In addition to specialists with domain specific industry knowledge, stage II companies need a talent structure that is capable of engineering ML systems. This can only be achieved by a mix of different professions that support each other. The setup resembles the pyramid structure presented in the talent subsection of section 1. It needs ML engineers that handle ML frameworks and software libraries, data scientists that are devoted to statistical models and algorithms that improve the application's accuracy, and data engineers responsible for engineering these models, and finally software developers involved in coding the application. This talent distribution is characteristic of tech hubs and challenging to achieve for non tech firms.

### Stage III: AI Platforms

| Characteristics of an AI Firm | Stage I | Stage II | Stage III |
|---|---|---|---|
| Pursuit of ML idea or research question | ✓ | ✓ | ✓ |
| Seed capital | ✓ | ✓ | ✓ |
| Revenue stream based on ML in products/services | ✓ | ✓ | ✓ |
| Access to talent | ✓ | ✓ | ✓ |

57 Handelsblatt Today, Twitter, January 30, 2019, https://twitter.com/HandelsblattGE/status/1090553570943033345.

58 "Bosch Center for Artificial Intelligence", Bosch, accessed June 13, 2019, https://www.bosch-ai.com/.

| | | | |
|---|---|---|---|
| Basic data environment | ✓ | ✓ | ✓ |
| Venture Capital | | ✓ | ✓ |
| Domain-specific knowledge or data | | ✓ | ✓ |
| Talent pyramid | | ? | ✓ |
| Development and provision of own hardware for training and inference | | | ✓ |
| AI platform business cases | | | ✓ |
| Targeted acquisitions of stage I and II companies | | | ✓ |
| In-house R&D centers | | | ✓ |

Due to its general-purpose nature, ML technologies can be applied to a wide range of industry products and services. With its application, market dynamics are changing and further shifting competitive advantages to tech companies that pursue platform models, giving them access to proprietary data and the financial means to invest in computing power. They apply ML to their vast troves of data, create products and services from it and offer these on their respective cloud platforms. Amazon's AWS is a good example of this. Amazon went from an online marketplace to a one-stop shop for ML practitioners – allowing them to choose their preferred ML software framework, apply it to their data already stored on AWS cloud services, and even use cloud compute to train and run their models.[59]

Furthermore, the general applicability of ML allows tech firms to tap into legacy firms' business cases – for example, in the auto industry where different tech companies are involved in the creation of mobility services that threaten car manufacturers' traditional business cases. These changing dynamics will become critical to EU legacy firms over time. But they also have tremendous implications for smaller companies that pursue business cases centered on ML technologies, which are prime targets for acquisition. This in turn can weaken an AI ecosystem's innovation base if that technology or

59 "TensorFlow on AWS", AWS, accessed June 13, 2019, https://aws.amazon.com/tensorflow/?nc1=h_ls. For a deep dive on the integration of ML software frameworks and cloud services, see Rebecca Wettemann, *TensorFlow on AWS* (Boston: Nucleus Research, 2018), https://d1.awsstatic.com/whitepapers/nucleus-tensorflow-2018.pdf.

talent is transferred abroad. The European Union is no stranger to this phenomenon. Brain drain is a big concern for the European innovation ecosystem. So while the interactions between stage I and stage II firms with tech or non-tech companies are important to watch, it is equally important to understand that it is problematic when such interaction turns into acquisition. This has strong implications for Europe's innovation base.

The acquisition of smaller (stage I or II) actors with the motivation to incorporate – as seen with the acquisition of London based startup DeepMind by Google in 2014 –, incubate, or exclusively benefit from these sources by larger firms is when oligopolies or monopolies, too-big-to-fail, kill zones[60], and market distortions might arise. DeepMind started with reinforcement learning but now provides solutions to Google ranging from text to speech synthesis, spam filters, translation, photo feature search, and simulation learning in advanced robotics – in brief, across the entire range of products and services Google offers.[61] Hence, acquisitions can be problematic and should be monitored carefully – especially when they target companies that are critical to Europe's innovation base. This further underlines the need for foreign policymakers to collect information on stage I developments and especially the transition of ML companies from stage I to stage II, as such companies represent an AI ecosystem's innovation base and are likely suspects to be targeted by AI platforms.

Although acquisitions by AI platforms may secure innovation and successfully shut out competition, it is nonetheless an aggressive approach. AI platforms also pursue a more elegant strategy to securing their innovation base by building their own research centers. Nearly every AI platform has created its own (AI) R&D center. Among the most prominent ones are Google Brain, Facebook AI Research, Baidu Research, and Alibaba Damo. Their efforts are aimed at creating innovation from within. But in parallel, these centers serve to attract top tier ML researchers. And this strategy is paying off. Talent is flocking to corporate environments that offer the above mentioned ML in-

60 The Economist, "American tech giants are making life tough for startups," *The Economist*, June 2, 2018, https://www.economist.com/business/2018/06/02/american-tech-giants-are-making-life-tough-for-startups;
The Economist, "How to fix what has gone wrong with the Internet," *The Economist*, June 28, 2018, https://www.economist.com/special-report/2018/06/28/how-to-fix-what-has-gone-wrong-with-the-internet.

61 Amir Efrati, "What's Really Behind Google's DeepMind Bet," *The Information*, January 27, 2014, https://www.theinformation.com/articles/What-s-Really-Behind-Google-s-DeepMind-Bet; and see the portrait of DeepMind by Thomas Macaulay, "Google DeepMind: the story behind the world's leading AI startup," *techworld*, March 7, 2019, https://www.techworld.com/startups/google-deepmind-what-is-it-how-it-works-should-you-be-scared-3615354/.

dustrial inputs. AI platforms have them in abundance. Products that result from such private sector R&D centers have now become essential building blocks to the companies' own product pipelines and to the global ML industry itself. Consider Google's deep learning software library TensorFlow – it provides a crucial foundation for creating DL technologies. It started out at Google Brain and creates a revenue stream through licensing fees. Another example are special purpose ML chips such as Amazon's Inferentia and Google's Tensor Processing Units (TPUs). These chips are available on the companies respective cloud computing environments and are used in the overall ML development and product cycle for training and inference.

Although stage II companies, such as Bosch, are trying to match these efforts to build in-house ML expertise through AI accelerators, AI platforms dramatically outperform them; not only do AI platforms have the first-mover advantage, they also leverage AI's general-purpose nature in order to horizontally integrate new business cases rather than focusing efforts on domain-specific use cases as stage II companies do. With Waymo, for example, Alphabet explores autonomous driving and mobility services, previously uncharted territories for tech platforms.

AI platform companies dominate not only with regards to their financial strength but also in terms of how dependent other companies seem to be on their products and services. For ML, these dependencies become obvious with on-demand cloud computing, ML software frameworks, and ML as-a-service products.

Equipped with an understanding of both which industry inputs are vital for a healthy AI ecosystem, as well as how these inputs interact in firms at different stages of maturity allows foreign policy makers to track the strength, competitive advantage, and market dynamics within different economies. Having foreign services monitor different innovation bases around the world by analyzing the two categories presented can give governments a deeper understanding of both domestic and foreign ML companies.

## Conclusion

The rhetoric around Artificial Intelligence (AI) in policy today remains inconsistent and vague. A deeper understanding of AI is urgently needed. It will improve the debate among the foreign policy community and thus provide the basis for better monitoring of key trends that can then inform policy making. The first section gives a clear definition of contemporary approaches to AI, namely machine learning (ML). Additionally it explains the industry inputs necessary to produce ML technologies: data, software, hardware, and talent. Understanding ML as a composition of these four inputs allows policy makers to grasp its economic relevance as well as recognize the significance of actors that produce and apply this general purpose technology.

The second section builds on this basic understanding of AI based on its core inputs, applying it to an analysis of the question: *what is an AI company?* Looking at the company level is important as innovation around ML disproportionately stems from private sector entities. This framework defines an AI company as one that focuses R&D efforts on creating specific core ML driven products, and one that has a revenue stream based on these ML driven applications. Next, this framework gives policy makers an overview of the market dynamics surrounding AI firms and the dependencies of the industrial inputs, by explaining three different stages of AI firm maturity: early stage ventures, domain specific application, and AI platforms.

Foreign service officers are instrumental in gathering and analyzing political and economic matters with regional and on-the-ground expertise and then reporting this information back to the government headquarters. The pace of AI development and deployment is fast; its economic importance substantial. But foreign services' resources to actively monitor global developments are currently limited. Responsibilities tend to be distributed among a handful of embassies. The same was true for cyber diplomacy a few years ago. As the significance of cyber related issues grew, so did the amount of resources and diplomatic activity.

Policy makers around the world have begun to understand the importance of AI. This is reflected by numerous international initiatives and fora devoted to AI Governance. But these fora and initiatives lack teeth and governments still lack a common nomenclature along which information on global developments in AI can be monitored. Leveraging the international network of embassies and the strength of foreign services – information gathering and analysis – foreign services can significantly help government headquarters shape both foreign and domestic policy around AI. Based on the two sections of the framework, foreign servants can collect data in a comprehensive way, rather than just tightening FDI screening mechanisms or adding buzzwords to export control lists – two strategies that continue to miss the general-purpose and complex nature of AI. The framework developed here allows foreign policy makers to gain a basic understanding of ML, to assess AI capabilities of companies around the world, and communicate these findings with headquarters so that governments can have an informed picture of global economic dynamics surrounding AI.

## About Stiftung Neue Verantwortung

Think Tank at the intersection of technology and society

The Stiftung Neue Verantwortung (SNV) is an independent think tank that develops concrete ideas as to how German politics can shape technological change in society, the economy and the state. In order to guarantee theindependence of its work, the organisation adopted a concept of mixed funding sources that include foundations, public funds and businesses. Issues of digital infrastructure, the changing pattern of employment, IT security or internet surveillance now affect key areas of economic and social policy, domestic security or the protection of the fundamental rights of individuals. The experts of the SNV formulate analyses, develop policy proposals and organise conferences that address these issues and further subject areas.

## About the authors

Philippe Lorenz leads the SNV's Artificial Intelligence and Foreign Policy project, assessing the implications of artificial intelligence for international relations. In his work, Philippe Lorenz focuses on policy areas that are undergoing major changes due to advanced technologies. In the project Labor Market 4.0, for example, he analyzed policy developments on the German labor market under the influence of increasing digitization. Here, he paid particular attention to the changes in occupational profiles due to the influence of automation technology. Another focus of his work was the influence of digitization on the design of the German energy transition. Philippe studied law at the University of Passau and International Relations at the Rhine-Waal University of Applied Sciences in Kleve.

### Contact the author

Philippe Lorenz
Project Director "AI and Foreign Policy"
plorenz@stiftung-nv.de
+49 (0)30 81 45 03 78 94
Twitter: @Ph_Lorenz

Kate Saslow works as project manager for three themes: Artificial Intelligence and Foreign Policy; EU Cyber Resilience; and Machine Learning and Cyber Security. In „Artificial Intelligence and Foreign Policy“ Kate looks at AI and specifically machine learning's (ML) effects on economic and trade dynamics, advancing an international human rights framework to apply to ML development, and analyzing global AI governance initiatives to understand Europe's position in international affairs. As part of the EU Cyber-Direct Project, Kate analyzes European cyber resilience policies to inform a debate on where Europe can foster strategic partnerships to push forward a global cyber resilience regime. With the Transatlantic Cyber Forum project, Kate manages the project on machine learning and information security as a way to understand potential for adversarial threats to ML systems and how this could be understood by policy makers as a danger to national security.

Kate holds a Master of International Affairs degree from the Hertie School of Governance, a public policy university in Berlin. Here she specialized in topics of digitalization, economic development, and governance, and concluded her studies with a capstone thesis on artificial intelligence and gender inequality in the United States labor economy. She also holds a Bachelor's degree in International Relations from the University of Michigan, Ann Arbor.

## Contact the author

Kate Saslow
Project Manager “AI and Foreign Policy” and
“International Cybersecurity Policy” and “EU Cyber-Direct”
ksaslow@stiftung-nv.de
+49 (030) 814 503 7884
Twitter: @kate_saslow

# Imprint

Stiftung Neue Verantwortung e. V.

Beisheim Center
Berliner Freiheit 2
10785 Berlin

T: +49 (0) 30 81 45 03 78 80
F: +49 (0) 30 81 45 03 78 97

www.stiftung-nv.de
info@stiftung-nv.de

Design:
Make Studio
www.make-studio.net

Layout:
Johanna Famulok

Free Download:

www.stiftung-nv.de