Artificial Intelligence Business

How you can profit from AI

by Przemek Chojecki, PhD

Table of contents

Introduction 5

Why Artificial Intelligence 7

Executive Summary. 7

Artificial Intelligence Paradigm 10

A short history of Artificial Intelligence 14

Basic terminology of Artificial Intelligence 16

Statistics related to AI. 19

Practical AI and how it is done 26

Research in Artificial Intelligence 28

Open-source community. 31

From research to applications 32

Cons of using AI 34

Powering Enterprises with AI 37

AI Maturity Levels 38

Solving maturity issues 44

Fostering a culture of innovation. 48

How to shift an existing culture 49

Hiring AI Talent 51

Building an innovative culture in enterprises. 52

Boosting Startups with Artificial Intelligence. 53

Make data-driven decisions fueled by AI 55

Automate your marketing efforts with AI 57

Improve your hiring process. 59

Startups powered by AI 59

One person enhanced with AI 61
One person startup 61
Using AI as an individual 62
Collect data 64

Trends in Artificial Intelligence 66
AI in retail 69
Manufacturing 73
Logistics 77
Robotics and Autonomous Vehicles 79
Robotic Process Automation 85
Image generation 89
Text generation and Chatbots 92
AI-powered education 95
AI in Healthcare 98
Cybersecurity powered by AI 106
Climate Change 108
Games and Reinforcement Learning 111
Hardware and beyond 116

Machine Learning Trends 124
Natural Language Processing 124
AutoML or automatic AI 125
One-shot learning and transfer learning 126
Reinforcement Learning 128
Computer Vision 129
Fundamental concepts 130
Pushing boundaries of machine learning 131

AI, Politics and Society 134
AI for social good 136
Public programs 138
Ethics and Regulations 141
Risks of using AI 144

Future of Artificial Intelligence 146
We are all cyborgs 146
Artificial General Intelligence 147

Introduction

We're living through a revolution. Artificial Intelligence is changing how we operate in the world and how smooth certain processes are. Just think about going on holidays. Multiple services allow you to find the most convenient flights and best hotels, you get personalized suggestions on what you might want to see, you go to the airport via one of the ride-sharing apps. At each of these steps, some AI algorithms are at work for your convenience.

I've started writing this book to fill in the gap. There are plenty of resources to learn machine learning and data science for technical people but no overviews, trends, discussion of applications, or anything at a more abstract level. On the other hand, those technical books are inaccessible to business people who don't want to go into coding themselves and want to stay at a more abstract idea level. The goal of this book is to address both issues and have a book that will be interesting both for data scientists and people with a business background.

That's why there's no code in this book, but I wasn't refraining from explaining certain more technical trends or discussing general toolset for artificial intelligence. The book is intended as a useful guide rather than a novel, that you have to read from the beginning until the end. Chapters are connected, but feel free to skip them whenever you want.

All in all, after reading this book, you'll know recent business applications of AI, understand how machine learning works and what to expect from it, what's trending, and how AI transforms every single business. Whether you're running a business, work at a large enterprise or in the public sector, this book will give you an overview of Artificial Intelligence as it is practiced today. I finish by discussing how we should integrate AI into our society, what are the risks of AI and how we can use AI to our benefit in the future.

Why Artificial Intelligence

Executive Summary

Artificial Intelligence is used in business through machine learning algorithms. Machine learning is a part of computer science focused on computer systems learning to perform a specific task without using explicit instructions, relying on patterns and inference instead.

Machine learning algorithms detect patterns and learn how to make predictions and recommendations by processing data, rather than by receiving explicit programming instructions ('if-then' loops). The algorithms improve over time with new data coming in, 'learning' through examples.

Machine learning is primarily used in:
- *predictions*: what will happen,
- *prescriptions*: what should be done to achieve goals,
- *descriptions*: what happened.

There are three main types of machine learning algorithms: supervised learning, unsupervised learning, and reinforcement learning.

Supervised learning uses training data and feedback from humans to learn the relationship of given inputs to a given output (for example, how the inputs “date” and “sales” predict customers’ preferences). Use it if you already know how to classify the input data and the type of behavior you want to predict, but you want to do it on new data.

Unsupervised learning explores input data without being given an explicit output variable (for example, explores customer sales data to identify patterns and classify them). Use it when you want to classify the data, but you’re unsure how to label the data yourself, or you want to discover hidden patterns.

Reinforcement learning learns to perform a task by trying to maximize rewards that you prescribe for its actions (for example, maximize returns of an investment portfolio). Use it when you have limited training data, and you cannot clearly define the end goal, or you want to explore possibilities without assuming what the solution might be.

The most common framework for doing machine learning is Python as a programming language. Experiments with machine learning models usually require access to powerful computers to ‘train’ algorithms. That’s why the additional cost of doing AI is the cost of the cloud when data scientists train their models. Those can range from a couple of hundred dollars per month to millions of dollars, depending on how heavy is the data and machine learning architecture. For most businesses the cost won’t exceed a couple of thousand dollars per month unless they want to invest heavily in AI capabilities and train their own models, rather than mostly use pre-trained, open-source solutions.

The most common architecture for machine learning algorithms is neural networks. You can think of them as Lego blocks of different sizes and colors that you can mix to build a specific construction. The basic parameter of a neural network is how many layers it has and how those layers interact with each other.

Deep learning is a subfield of machine learning which focuses on neural networks with at least 3 layers. Deep learning is the actual reason why AI is so popular today, as its applications in image or voice recognition are far better than classical methods. Neural networks combined with enough computing power give outstanding results on real-world data.

'Big data' is another buzzword used in the last decade often. Big data never had a proper definition, always meaning having more data than is possible to process using a single personal computer. That's why what we today understand as big data (petabytes of data) is far away from used to be big data just 10 years ago (terabytes) and how it will change in the next 10 years (exabytes).

As data is crucial for machine learning algorithms, 'big data' is coming back in organizations as a fundamental term to explore AI capabilities. Machine learning requires that the right set of data be applied to a learning process. You don't need big data to use machine learning algorithms, but big data can help you improve the accuracy of your algorithms.

That's why often it's not necessarily true that you need a lot of data to start experimenting with AI. Especially with the raise of reinforcement learning and techniques like 'one-shot learning,' AI is within reach for every single organization. The first step to benefit from AI is to prepare data by cleaning it and sorting it by human coworkers. Then machine learning engineers and data scientists will be able to take care of the rest.

INPUT	OUTPUT	APPLICATION
Voice Recording	Transcript	Speech Recognition
Image	Caption	Image Recognition
Recipe Ingredients	Customer Reviews	Food Recommendations
Historical Market Data	Future Market Data	Trading Bots
Drug chemical properties	Treatment efficacy	Pharma R&D
Transaction details	Is transaction fraudulent?	Fraud detection
Purchase history	Future purchase behavior	Customer retention
Faces	Names	Face Recognition
Car locations and speed	Traffic Flow	Traffic Lights

Artificial Intelligence Paradigm

Though it might seem like we've come a long way in the last ten years, which is true from a research perspective, the adoption of AI among corporations is still relatively low. According to some studies, only around 20% of companies are experimenting with AI. Most of the companies have only started to dabble with artificial intelligence.

This new era of information depends heavily on the knowledge, and we're currently missing a lot of experts.

There's a lot of fear and hype around AI, and it's crucial to have as many people as possible to know about AI in order to understand what's possible and what's not. Only an educated AI-wise society will be able to adopt the technology fully. The goal of this book is to make it happen sooner than later.

At business level decision-makers, project managers and executives need basic knowledge, understanding of the machine learning paradigm, which would allow them to apply practical algorithms to real use cases, driving business, and growing sales.

On the other hand, trust is a crucial ingredient in any system. Thus explainable AI, being able to explain why these automation systems, machine learning architecture, act in this way, will allow us to understand how it's going to affect all of us: co-workers, citizens, humans. For that purpose, we need to think about re-education when it comes to our organizations and implementing AI elements in schools from the very start of education.

AI will mostly enhance what we do on a daily basis. It's not going to be full automation most of the time, but an augmentation of how we do certain tasks and how we work. This way of collaboration will also require rewiring of how we think of machines.

That's why we also need to think about regulations. AI is the atomic energy of our times, and we can either use it to produce bombs or use it to produce energy. The standard paradigm of computing is based on 'if-then' loops. Coders give instructions to computers, supervising every single step of computing. Machine learning changed that completely. Coders don't have to code every single step to make the computer work. They can just build a general architecture, like a specific neural network, and supply data in order to 'train' this architecture - that is, let the computer system self-tune as it sees fit by analysing data. This approach is more similar to teaching a child a particular task or introducing a junior coworker to a particular business process for the first time. And as it is the case with children, it often takes time for them to grasp the concept fully. It's similar to machines. They need retraining and rebuilding parts of their architecture to excel at a given task.

Because of this shift of paradigm, it became possible to automate more tasks and business processes than ever before. You don't have to program every single step of the process, predicting at each step what might happen and how to react to that. You can leave many of those details to algorithms, letting them see data and decide for themselves in each case how it should be solved.

Of course, that's theory. In practice, things can get messy. AI is not a magic wand, and it doesn't solve any problem you throw at it. You need good preparation to really benefit from artificial intelligence. This includes:

1. Clearly describing what business process you want to automate or optimize.

2. Defining what the output of the process is and how to distinguish between good or bad results. If that's not possible, or if this is a continuous process without an end, then define mid-steps and mid-results that are anticipated.
3. Defining what the input of the process is, that is what kind of data you take into account when looking for the output.
4. Acquiring large datasets related to the business process. Cleaning it by removing unnecessary parts and organizing it in one place and in one format (e.g. .doc files stored on a cloud).

Having done this work, you're ready to start hiring data scientists to build machine learning algorithms for you. Often this preliminary work will be revised and enhanced with new data and new insights, but you don't have to worry about it at the start of the process. The crucial part for you, as an executive, is being clear about what business process exactly you want to tackle with AI.

Explanation of this new machine learning paradigm and how to apply it in business is the main reason for writing this book. I believe that understanding how data science teams work, how machine learning models are constructed, and what they need to perform well, is crucial to be competitive amid today's technological revolution.

It is also crucial for legislators, politicians, and philosophers to understand how the machine learning community operates and what is possible with AI, to make legislations and laws working for everyone involved. There's definitely too much hype regarding AI when it comes to how fast it will disrupt human jobs as we know them. And even though I believe that will eventually happen, what we should prepare for is the next 5-10-20 years of growing dependence on AI. Every job will be enhanced by artificial intelligence, rather than replaced by it, and we shouldn't be scared of it. Examples in art and video games show that perfectly. We can learn a lot from the way machine learning systems look at our world and how they transpose it.

The most important thing is to stay open, embrace new technologies, and learn constantly. After all, what humans do best is adapt,

A short history of Artificial Intelligence

The term artificial intelligence was used the first time in 1955 by John McCarthy, a math professor at Dartmouth who organized the seminal conference on the topic the following year. In 1957 the economist Herbert Simon predicted that computers would beat humans at chess within 10 years (he was slightly wrong, it took 40). In 1967 the cognitive scientist Marvin Minsky said, "Within a generation, the problem of creating 'artificial intelligence' will be substantially solved." Simon and Minsky were both intellectual giants, but they were wrong about AI badly. These dramatic but wrong claims caused various repercussions in how people in the second half of 20th century thought about AI: more as a subject of a science fiction novel than actual science.

The idea of an Artificial Intelligence, automation of certain repetitive processes, dates back to the Cold War when US intelligence was trying to translate Russian documents and reports automatically. The initial optimism of the ’50s was then undermined by the underperformance of these early models and fundamental lack of progress past initial results. With a lack of optimism, funding was cut substantially, and the academic community turned away from AI, especially in the 1970s, when DARPA cut its funding. This period of lowered funding and loss of interest was later called the AI winter.

A renewal of AI came in the 1980s with LISP, a programming language, and LISP machines, computers optimized to run LISP code, which was the default language for doing AI research at that time. A couple of companies were producing those computers and selling them commercially with some initial successes, but eventually, they were overtaken by personal computers as we know them today. Again results were not good enough as to what was promised in research proposals. The second AI winter began.

The start of the new era of Artificial Intelligence is dated from 2009-2012. The community of AI researchers was steadily growing from the 1990s and the early 2000s, with larger grants and some interest from corporations, but the most significant catalyst for the current AI revolution came in 2009 with the creation of ImageNet, a large visual database designed to test image recognition algorithms on. Then on 30 September 2012, a convolutional neural network called AlexNet achieved a top-5 error of 15.3% in the ImageNet 2012 Challenge, more than 10.8 percentage points lower than that of the runner up, beating classical algorithms. This was made feasible due to the use of Graphics Processing Units (GPUs) during training, an essential ingredient of the deep learning revolution that was about to start. Suddenly people started to pay attention, not just within the AI community but across the technology industry as a whole and the current revolution began.

The ideas we use today in AI research like neural networks were to some extent known already 30 or 40 years ago. However, what was missing was enough data and enough computing power to process this data. Machine learning couldn't take off without access to large datasets, and thanks to the digital revolution we lived through in the early 2000s, suddenly many Internet companies emerged, older companies became digitized, and there was more data than any human could process. On the other hand, assuming Moore's law being true, the power of computers has been doubling every 18 months, reaching the necessary power to process big data as it was created, or at least enough of it to make neural networks work. And once that was established, everyone came to AI research again: governments, corporations, scientists.

Basic terminology of Artificial Intelligence

Let's now jump into basic terminology related to AI. When we say Artificial Intelligence we mostly mean machine learning - a domain of computer science that uses learning algorithms able to tune themselves on data provided by a user. The fundamental block of machine learning is neural networks. They are algorithmic systems based on simulating connected "neural units," loosely modeling the way that neurons interact in the brain.

As we have mentioned above, these computational models inspired by neural connections have been studied since the 1940s. They have returned to prominence with the rise of computer processing power able to cope with large training data sets and have been used to successfully analyze input data such as images, video, and speech. Deep learning is a subset of machine learning, where neural networks have many layers of neurons ("deep network"). The more layers you include in your machine learning model, the more computational power you need to train it. We talk about the architecture of a model, when we want to describe how many layers it has, how many neurons inside each layer, and how they are connected.

The most common neural networks appearing in applications are:

1. **Feedforward neural networks**: this is the simplest type of neural network. In this architecture, information moves in only one direction, forward, from the input layer, through the hidden layers (those between input and output), to the output layer. There are no loops in the network. The first single-neuron network was considered already in the 1950s. Advances in computing power and available data allowed this method to achieve great performance in the 21st century.
2. **Recurrent neural networks (RNNs)**: neural networks whose connections between neurons include loops. One of the most common examples of RNNs is LSTMs which are used in language processing tasks.
3. **Convolutional neural networks (CNNs)**: CNNs were originally invented as a superior neural network model

for handling cases in which the input data consisted of images. In 2012 CNNs were used for the winning entry in the ImageNet Large Scale Visual Recognition Competition, which sparked interest in machine learning again.

I also note generative adversarial networks (GANs) and reinforcement learning as two methods soon to be more common in commercial applications.

GANs use two neural networks competing against each other. They are often used for photo-realistic image generation: one network is trained to detect fakes, and the other is trying to fool the first one.

Reinforcement learning is an approach in machine learning based on giving rewards designed by developers to steer the machine into good behavior. Algorithms learn by trial and error. A notorious application of reinforcement learning is AlphaGo created by Google DeepMind, which was trained to play Go at a world-class level.

Any of these deep learning methods require thousands of data records for a model to train and achieve the desired accuracy. The authors of ‘Deep Learning’ book mention a general rule of thumb, that a supervised machine learning algorithm should achieve acceptable performance with around 5,000 labeled examples per category and match human-level performance when trained on at least 10 million labeled examples. Of course, it also depends on particular use cases and algorithms architecture. Sometimes more data isn’t that helpful if you don’t know how to feed it properly into machine learning models. On the other hand, sometimes machine learning techniques won't add more value than traditional statistical analytics. That's why it's essential to assess your level of technical development, look at your goals, and think about possible solutions without AI at first.

A lot of machine learning models used currently are trained through "supervised learning," which requires humans to label and categorize the underlying data. Nevertheless, the new methods like 'one-shot learning' show that in the future, we won't need that much data to train effective AI systems. One will need only a small set of labeled data and a good architecture in place. On top of that, autoML might improve AI even further without the need for human supervisors.

All that means that if an organisation wants to adopt AI successfully, it needs to start with assessing its technology stack and start by collecting data at scale. Linking data across various segments (customer, communication channel, platform) as well as controlling whether the right amount of data is given is crucial. A machine learning model can be 'overfitted' if it matches too well the test data but doesn't work in production, or 'underfitted' if it fails to capture essential features and thus fails to generalise.

Statistics related to AI

Following analysis done by AI State Index[1], let's review some statistics related to Artificial Intelligence, that will fully show

[1] Artificial Intelligence Index Report 2019, in this section we cite their numbers directly unless otherwise stated.

how important this market is becoming (or already is). It's crucial to understand that we are still early when it comes to applied AI, and most of those statistics will grow substantially in the upcoming years. The reason for that is most of the cutting-edge research is still far away from day-to-day business applications either because of the costs of the hardware or the required expertise to apply it. I expect the full AI boom to come within the next ten years when every company will need to implement AI elements to be competitive even at the local scale. This will come in pair with the democratisation of AI: the cost and difficulty of implementation of most algorithms will largely decrease. AI applications will be as available as general cloud storage is now.

Looking at Google Trends one can see that "cloud computing" appears in 2008 and then it is replaced by "big data" which starts taking off in 2011. "Machine learning" and "data science" begin to rise together in 2013, which matches the renewed interest in AI after the 2012 ImageNet competition.

Let's now look at different aspects of the Artificial Intelligence ecosystem.

Research

- Between 1998 and 2018, the share of AI papers among all papers published worldwide has grown three-fold, now accounting for 3% of peer-reviewed journal publications and 9% of published conference papers.
- The number of AI research papers surpassed 35,000 in 2019 worldwide as evaluated by looking at arXiv and AI

conferences. Most AI papers are published in North America and China.
– The number of patents related to AI is growing faster than the number of scientific papers. Most of the patents are within the computer vision subdomain and are registered in the US.
– Europe publishes the most AI papers. Papers published by American authors are cited 83% more than the global average.

Business and Funding

Global investments in AI and AI startups continue to rise. From a total of $1.3B raised in 2010 to over $40.4B in 2018 alone, funding has increased with an average annual growth rate of over 48% between 2010 and 2018.

'State of AI Report in 2019' claims that the number of AI companies that received funding is also increasing year by year, with over 3000 AI companies receiving funding in 2018. They calculated that between 2014 and 2019 (up to November 4th), a total of 15,798 investments of over $400K have been made in AI startups globally, with an average investment size of approximately $8.6M.

On the other hand, Crunchbase lists 13,650 AI companies as of May 2020, of which 97.8% are active. They have raised $19M on average, the median funding being $2.2M.

VC-driven private investments accounted for about half of total investments in AI in 2019, with M&A and Public Offerings taking the major share of the remaining half. However, private investment accounted for 92% of the number of deals, with M&A making up just over 4% of deals, and Minority stakes and Public offerings (IPOs) together accounting for 3%. These statistics show that most of the money goes to already successful startups.

"AI investment[2] is growing fast, dominated by digital giants such as Google and Baidu. Globally, we estimate tech giants spent $20 billion to $30 billion on AI in 2016, with 90 percent of this spent on R&D and deployment, and 10 percent on AI acquisitions. VC and PE financing, grants, and seed investments also grew rapidly, albeit from a small base, to a combined total of $6 billion to $9 billion. Machine learning, as an enabling technology, received the largest share of both internal and external investment." AI investments keep on growing in the last years.

Markets and Markets estimate that the AI market will be worth $190 billion by 2025.[3] We might hit this benchmark even sooner when you look at the above examples. To add even more examples:

- Open AI had a recent investment by Microsoft of $1 billion.
- SoftBank announces the second Vision Fund, which will be AI-focused and which will have $108 billion to invest.
- SAS is going to invest $1 billion in artificial intelligence over 3 years starting from 2019.

---

[2] McKinsey report on AI from June 2017

[3] https://www.marketsandmarkets.com/PressReleases/artificial-intelligence.asp%20.asp

The US federal government is projected to invest around $5 billion in AI R&D in fiscal 2020.

In the fiscal year 2018, the latest year in which complete contracting data is available, US federal agencies spent a combined $728 million on AI-related contracts, an almost 70% increase above the $429 million that agencies spent in fiscal 2017. Since the fiscal year 2000, the Pentagon has accounted for the largest share of AI spending of any federal agency ($1.85 billion), followed by NASA ($1.05 billion), and the departments of the Treasury ($267 million) and Health and Human Services ($245 million).

We can easily extrapolate that investments in AI and AI-related companies will only be growing in the next years as more research will be available for commercialisation. Also time of commercialisation might be shorter due to the growing talent pool and democratisation of AI.

Hiring

Hiring AI talent is hard, as the market is very competitive. Big tech companies have been actively buying AI startups, not just to acquire technology or clients but to secure qualified talent - this is usually called acquihire and is often practised in tech markets.

The pool of experts in machine learning is small, moreover, Microsoft, Amazon, Facebook, Google, and other tech giants have hired many of them. Companies have adopted M&A as a way to grab the top talent - typically those deals are valued at

$5 million to $10 million per person on an M&A deal (the lowest is usually $1 million per person). The shortage of talent and the cost of acquiring talent are underlined by a recent report that companies are seeking to fill 10,000 AI-related jobs and have budgeted more than $650 million for salaries. The US alone is opening over 7,000 AI-related jobs in 2019.[4] We can look at three aspects of the AI job market.

– **growth**: the rapid growth in AI hiring is also confirmed by job postings data from Burning Glass that shows the share of AI jobs (% of total jobs posted online) grew from 0.1% in 2012 to 1.7% in 2019 for Singapore. Similarly, in the US, the share of AI jobs increased from 0.3% in 2012 to 0.8% of total jobs posted in 2019.
– **demand**: machine learning jobs increased from 0.07% of total jobs posted in the US in 2010 to over 0.51% in October 2019.
– **salary**: compensation of senior engineers at large tech companies is approaching $1,000,000 of which about half is in company's stocks. At the other end of the spectrum, there's huge growth in $1.47/hour data labeling jobs.

All in all, hiring might be the biggest bottleneck for organisations to start deploying AI at scale. There are various ways to overcome this problem:

– outsource AI tasks to specialised software houses,
– use existing solutions and adapt them to your needs,

[4] https://enterprisersproject.com/article/2019/8/ai-artificial-intelligence-careers-salaries-7-statistics

- acquihire whole teams,
- offer a competitive environment for machine learning engineers.

We'll come back to these issues in the next chapters, while discussing how to use AI in organisations, be that large enterprises or startups.

Practical AI and how it is done

Artificial Intelligence in business is practical. When you think about neural networks, don't think about abstract mathematical structures, but rather computer systems that need data to learn business processes and how to operate within them.

Data Science is not a real science, it's an experimentation domain, where you need to constantly adjust, test, build prototypes from scratch, and rebuild what you have. It's a framework for approaching problems rather than a specific set of tools. This paradigm of using neural networks, statistics on steroids, is what makes AI both practically and theoretically complex, with such a broad range of applications, which we're going to cover in the next chapter.

So how Data Science or Artificial Intelligence is currently done? You could split the actual work into two parts, connected strongly with each other:

- implementation,
- research.

Implementation phase is focused on delivering practical solutions to a specific business problem. By using data from within your organisation, data scientists implement machine learning models to learn on this data. This phase is heavily focused on engineering aspects of data science:

- cleaning data
- feature extraction
- statistical analysis
- training neural networks
- setting up virtual machines and a general framework.

Research phase is about looking for possible tweaks, ameliorations, or totally new approaches to existing problems. It may consist of reading scientific papers, white papers from other organisations, browsing open-source code on Github, talking with fellow machine learning engineers, attending conferences. The goal is to broaden perspective and find new strategies to implement.

It's in general impossible to say what comes first, implementation or research, as the first steps of data scientists are often building the very first naive model, seeing how it works on given data, and then looking for other approaches and enhancements. For harder and more engaging projects, machine learning engineers might start with research, reading what's possible to find on a subject on the web, and only then choosing a couple of models to implement and try.

Nevertheless, data scientists spend most of the time in front of the computer, whether reading, writing code, or training machine learning models. What's often misunderstood in corporations is that usual sprints done in classical software development (lean startup method) are not always beneficial to finding solutions to more involved problems that require deep thinking. That's why the 20% rule of Google, allowing for 20% time off to work on software engineers' own projects, is so fruitful. Data scientists need to tinker and play around with ideas to boost their creativity.

Research in Artificial Intelligence

The research community in Artificial Intelligence can be split into three divisions:

– machine learning community
– ethics and social community
– business community

Machine Learning community is concerned primarily with research questions related to building machine learning models: from architecture through data to implementations. PhD in computer science or STEM field is necessary to participate actively in it.

Ethics and social community focuses on social ramifications of doing AI research and applying it in practice: from legislations to important questions or limits on what should be the goal of AI research. People in this community often work in social departments of universities, think tanks, or public institutions.

Business community focuses on applying cutting-edge research to business problems. Those may include manufacturing, drug design, cybersecurity, video games, and others. Researchers here work mostly at research labs of large organisations. PhD is not necessary, but often an additional advantage when it comes to looking for a job in those.

If one wants to become a researcher in AI, the standard road is via university, doing a PhD in computer science, and then becoming an assistant professor or a research fellow. Thanks to recent changes in how research works, for example how Amazon, Facebook, Google and similar large tech companies are participating in doing research, it often happens that freshly mint PhDs go directly to one of tech giants' research laboratories. It's also possible that they do simultaneously PhD and work at one of those companies which is beneficial to each party: a company sponsors a PhD, the university is relieved from costs, a PhD student has a job and does something relevant to the industry.

PhD thesis itself is a monograph discussing and solving an open problem or some case of it, using novel methods in an already established problem or inventing new problems related to existing knowledge. Some topics are more in fashion at a particular time than others and this relates to interests of particular professors or interest of the market (where the money is). Often during the time of doing a PhD a student publishes a couple of papers, which then consists of the main body of a PhD thesis.

For a PhD student, the most important is finding a good advisor with access to interesting problems, funds, and a research group. Interesting problems will allow him to do meaningful research, funds will allow him to travel to conferences and spend money on infrastructure, the research group will be invaluable for research discussions.

Going to conferences is a great way to connect with fellow scientists. The most popular and most prestigious machine learning conference is NIPS, Neural Information Processing Systems' annual meeting. The number of scientists applying with their papers to NIPS is growing by 30% each year, which also shows how lively is the machine learning community currently.

From the point of view of business that competes with academia for talent, the crucial aspect is creating a vibrant environment to do research in. Assigning free time to do any research is a good solution, but crucial is building a research group around a senior figure in the field. It was often the case that large organisations hired a professor from the machine learning department together with his PhD students as a way to start up a research community quickly. For example, this is what Uber did with poaching people from Carnegie Mellon's robotics department.[5]

A big problem for established institutions like banks or insurers is presenting themselves in an appealing way to potential

[5] https://www.theverge.com/transportation/2015/5/19/8622831/uber-self-driving-cars-carnegie-mellon-poached

machine learning employees. Crucial here is understanding that what's appealing to researches is being able to innovate, have freedom of thought, an atmosphere of openness, and hard problems at hand to solve. No one wants to be stuck with linear regressions all the time. It's often better to pose too hard problems than too easy problems to attract talent (think Tesla or SpaceX).

Excellent examples of good problems are on Kaggle (www.kaggle.com), where companies run data science challenges for their business problems setting a prize for top entries. Often these competitions are attended by thousands of teams. One of the most famous ones was a competition run by Netflix[6] to make better their recommendation algorithms. By putting the prize at $1M the competition attracted a lot of data scientists, put Netflix on a map of great tech companies to work at, and gave Netflix a lot of new research input relevant to their business operation.

Open-source community

Important from a business perspective and still largely underused by the business is the open-source community within machine learning. Much of research is available for free on GitHub, a repository of code, and can be picked up and used jointly with other pieces to build something unique you need. Never making a prototype was so fast and cheap as now. The open-source community is also an excellent source for potential hires as it accurately shows

[6] https://www.kaggle.com/netflix-inc/netflix-prize-data

what a given person is capable of by just looking at his or her code repository.

Business-wise supporting the open-source community has many advantages: access to the talent pool, staying informed about current research. Moreover, it can bring business leads. Recall the model of Red Hat which was responsible for maintaining Linux and then earning money via support and customisations. In the end, Red Hat was acquired by IBM in one of the largest tech acquisitions to date at massive $34 billion closed in 2019.

GitHub itself was acquired by Microsoft in 2018 for $7.5 billion,[7] and Kaggle was acquired by Google in 2017.[8] This not only shows how important open-source community is for business but actually that you can make a business out of open source efforts if you're able to deliver a great product and build a community of engaged users around it.

From research to applications

Having discussed how research is done in AI, it's now time to focus on applications. Assuming you already have a data science team in place and preliminary research on a problem you want to solve done, the next step is to gather and clean data. This process can be short if most of your business is digital with easy access to data, or long and painful if

[7] https://www.theverge.com/2018/10/26/17954714/microsoft-github-deal-acquisition-complete
[8] https://techcrunch.com/2017/03/08/google-confirms-its-acquisition-of-data-science-community-kaggle/