

1 A new disruptive technology is coming

Artificial intelligence is being considered as the new industrial revolution, heart of what some call industry 4.0. Well, Deep Learning is the engine of this process and in the following chapters we will talk about it extensively. But in this chapter we will situate the issue first, to see why artificial intelligence is already here and why it has come to stay.

1.1 Artificial Intelligence is changing our life

We are facing dizzying advances in the quality and device features of a wide range of everyday technologies: in the case of speech recognition, voice-to-text transcription has experienced incredible advances and is now available in different devices. We are interacting more and more with our computers (and all kinds of devices) by simply talking to them.

There have also been spectacular advances in natural language processing. For example, just by clicking on the *Google Translate*³⁷ micro icon (at the bottom left of the text box), the system transcribes into another language what is being dictated. *Google Translate* already allows you to convert sentences from one language to another in 32 language pairs, and offers text translation for more than 100.

In turn, the advances in computer vision are also enormous: now our computers, for example, can recognize images and generate textual descriptions of their content in seconds.

These three areas are crucial to unleash the improvements in robotics, drones or cars without driver, among many others areas. Artificial Intelligence is at the heart of all this technological innovation, which lately advances so quickly

³⁷ See https://support.google.com/translate/answer/6142468?hl=en&ref_topic=7011659 [Accessed: 18/06/2018]. Use Google Chrome browser.

thanks to Deep Learning.

And all this despite the fact that artificial intelligence has not yet been widely deployed and it is difficult to get an idea of the great impact it will have, just as in 1995 it was difficult to imagine the future impact of the internet. Back then, most people did not see how the internet would end up being as relevant to them and how it was going to change their lives.

It is said that the first industrial revolution used steam energy to mechanize production in the second half of the 18th century; the second revolution used electricity to boost mass production in the mid-nineteenth century while, in the third one, electronics and software were used in the seventies of the last century. Today we are facing a new source of value creation in the area of information processing where everything will change. In different forums, people already talk about a fourth industrial revolution, marked by technological advances such as Artificial Intelligence (Machine Learning, Deep Learning, etc.) and in which computers will be "even wiser".

1.2 Artificial Intelligence, Machine Learning and Deep Learning

Before continuing we will specify a bit what we understand by Artificial Intelligence, Machine Learning and Deep Learning, three terms that will appear very often throughout the book.

Artificial Intelligence

What do we mean when we talk about Artificial Intelligence? An extensive and precise definition (and description of its areas) is found in the book

*Artificial Intelligence, a modern approach*³⁸ written by Stuart Russell³⁹ and Peter Norvig⁴⁰, the most popular book on Artificial Intelligence in the university world and, without a doubt for me, the best starting point to have a global vision of the subject. But trying to make a more general approach (purpose of this book), we could accept a simple definition in which by Artificial Intelligence we refer to that intelligence shown by machines, in contrast to the natural intelligence of humans. In this sense, a possible concise and general definition of Artificial Intelligence could be the effort to automate intellectual tasks normally performed by humans.

As such, the area of artificial intelligence is a very broad scientific field that covers many areas of knowledge related to machine learning; even many more approaches not always cataloged as Machine Learning are included by my university colleagues who are experts in the subject. Besides, over time, as computers have been increasingly able to "do things", tasks or technologies considered as "smart" have been changing.

Furthermore, since the 1950s, Artificial Intelligence has experienced several waves of optimism, followed by disappointment and loss of funding and interest (periods known as AI winter⁴¹), followed by new approaches, success and financing. Moreover, during most of its history, Artificial Intelligence research has been dynamically divided into subfields based on technical considerations or concrete mathematical tools and with research communities that sometimes did not communicate sufficiently with each other.

³⁸ *Artificial Intelligence: A Modern Approach (AIMA)* · 3rd edition, Stuart J Russell and Peter Norvig, Prentice hall, 2009. ISBN 0-13-604259-7

³⁹ Stuart J. Russell. Wikipedia. [online]. Available at: https://en.wikipedia.org/wiki/Stuart_J._Russell [Accessed: 16/04/2018]

⁴⁰ Peter Norvig Wikipedia. [online]. Available at: https://en.wikipedia.org/wiki/Peter_Norvig [Accessed: 16/04/2018]

⁴¹ AI winter. Wikipedia. [online]. Available at: https://en.wikipedia.org/wiki/AI_winter [Accessed: 16/04/2018]

Machine Learning

As we said in the previous section, advances such as speech recognition, natural language processing or computer vision are crucial to trigger improvements in robotics, drones, driverless cars, among many other areas that are changing the near future. Many of these advances have been possible thanks to a family of techniques popularly known as Deep Learning, of which we will talk extensively. But first, for us to have a correct global image, I think it's interesting to specify that Deep Learning is a subpart of one of the areas of Artificial Intelligence known as Machine Learning.

Machine Learning, is in itself a large field of research and development. In particular, Machine Learning could be defined as the subfield of Artificial Intelligence that gives computers the ability to learn without being explicitly programmed, that is, without requiring the programmer to indicate the rules that must be followed to achieve their task; the computers do them automatically.

Generalizing, we can say that Machine Learning consists in developing for each problem a prediction "algorithm" for a particular use case. These algorithms learn from the data in order to find patterns or trends to understand what *the data tell* us and in this way build a model to predict and classify the elements.

Given the maturity of the research area in Machine Learning, there are many well-established approaches to Machine Learning. Each of them uses a different algorithmic structure to optimize the predictions based on the received data. Machine Learning is a broad field with a complex taxonomy of algorithms that are grouped, in general, into three main categories: Supervised Learning, Unsupervised Learning and Reinforcement Learning.

We mean that learning is supervised when the data we use for training includes the desired solution, called "label". Some of the most popular machine learning algorithms in this category are linear regression, logistic

regression, support vector machines, decision trees, random forest or neural networks.

On the other hand, when the training data do not include the labels, we refer to an Unsupervised Learning and it will be the algorithm which will try to classify the information by itself. Some of the best-known algorithms in this category are clustering (K-means) or principal component analysis (PCA).

We also talk about Reinforcement Learning when the model is implemented in the form of an agent that should explore an unknown space and determine the actions to be carried out through trial and error: the algorithm will learn by itself thanks to the rewards and penalties that it obtains from its “actions”. The agent must create the best possible strategy (policies) to obtain the greatest reward in time and form. This learning allows it to be combined with other types, and is now very fashionable since the real world presents many of these scenarios.

Basic Machine Learning terminology

In this section we will advance basic Machine Learning terminology that will allow us to keep a presentation script of the concepts of Deep Learning in a more comfortable and gradual way throughout the book.

In Machine Learning we refer to *label* to what we are trying to predict with a model. Instead, we call an input variable a *feature*.

A *model* defines the relationship between features and labels and has two clearly differentiated phases for the subject that concerns us:

- *Training* phase (or learning phase), which is when the model is created or learned, showing the examples of input that have been tagged; In this way, the model is able to iteratively learn the relationships between the features and labels of the examples.

- *Inference* phase (or prediction or phase), which refers to the process of making predictions by applying the model already trained to non-labeled examples.

Consider a simple example of a model that expresses a linear relationship between features and labels. The model could be expressed as follows:

$$y=wx+b$$

Where:

- y is the *label* of an input example.
- x is the *feature* of that input example.
- w is the slope of the line and that is what in general we call *weight*. It is one of the two parameters that the model has to learn during the training process to be able to use it later for inference.
- b is the intersection point of the line on the y-axis that we call *bias*. This is the other parameter that must be learned by the model.

Although in this simple model that we have represented there is only one input feature, in the case of Deep Learning we will see that we have many input variables, each with its w_i weight. For example, a model based on three features (x_1, x_2, x_3) can be expressed as follows:

$$y= w_1x_1+ w_2x_2+w_3x_3+b$$

Or, more generally, it can be expressed as:

$$y= \sum_i w_i x_i + b$$

which expresses the sum of the scalar product between the two vectors (\mathbf{X} and \mathbf{W}) and then adds the bias. The parameter bias b , in order to facilitate the formulation, it is sometimes expressed as the parameter w_0 (assuming a fixed additional entry of $x_0=1$).

In the training phase of a model, ideal values are learned for the model parameters (the w_i weights and the b bias). In supervised learning, the way to achieve this is to apply an automatic learning algorithm that obtains the value of these parameters by examining many labeled examples and try to determine values for these model parameters that minimize what we call *loss*.

As we will see throughout the book, the loss is a central concept in Deep Learning that represents the penalty of a bad prediction. That is, the loss is a number that indicates how bad a prediction has been in a particular example (if the prediction of the model is perfect, the loss is zero). To determine this value, as we will see later, in the training process the concept of the loss function will appear, and for the time being we can see how the mathematical function that aggregates the individual losses obtained from the input examples to the model.

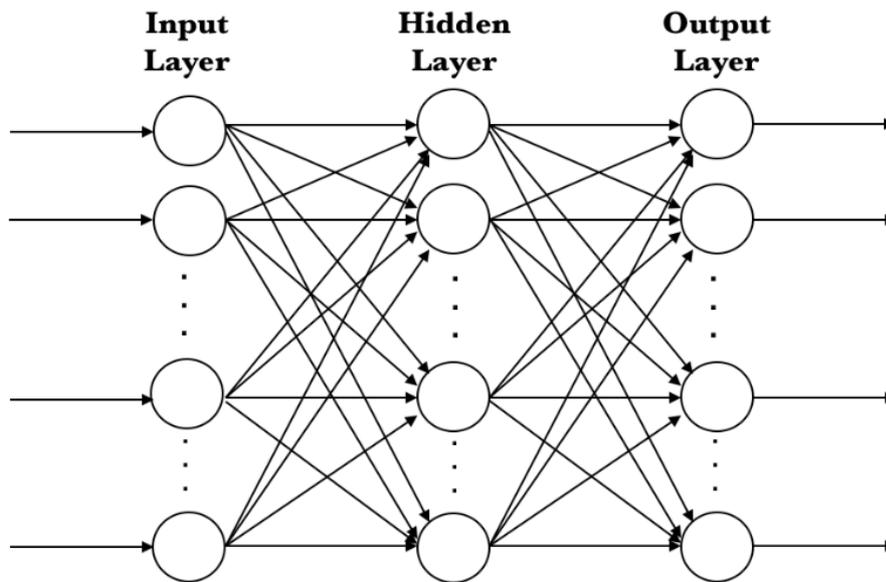
In this context, for now we can consider that the training phase of a model consists basically in adjusting the parameters (the weights w_i and the b bias) in such a way that the result of the loss function returns the minimum possible value.

Finally, we need to address the concept of overfitting in a model, which occurs when the model obtained is adjusted so much to the labeled examples of input that, later, the resulting model cannot make the correct predictions with new data examples that have never been seen before. Given the introductory nature of this book we will not go into this topic, but it is a central theme in Deep Learning.

Artificial neural networks and Deep Learning

A special case of Machine Learning algorithms are artificial neural networks, whose algorithmic structures allow models that are composed of multiple layers of processing to learn data representations with multiple levels of abstraction. This set of layers composed of "neurons" performs a series of linear and non-linear transformations to the input data to generate an output close to the expected (label). Supervised learning, in this case, consists in obtaining the parameters of these transformations (the w_i weights and the b bias) and attempting that these transformations produce an output that differ as little as possible to the expected output.

A simple graphical approach to a Deep Learning neural network can be



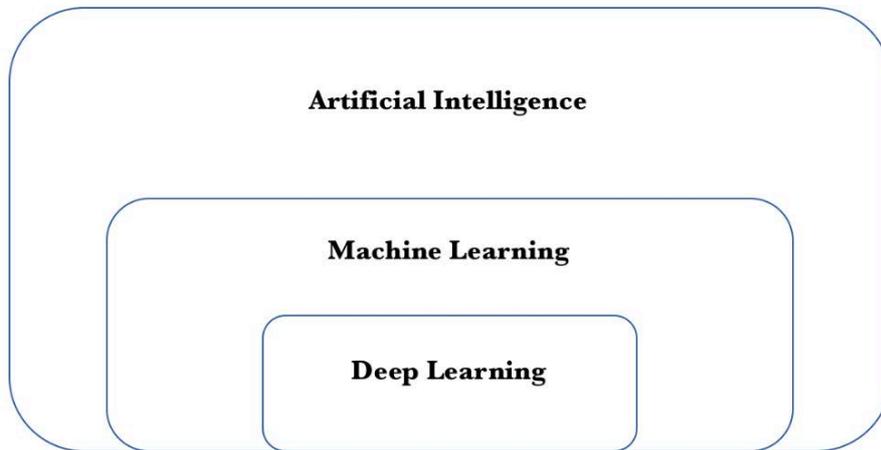
Specifically, in this image we represent an artificial neural network with 3 layers: an input layer that receives the input data and an output layer that returns the prediction made. The layers that we have in the middle are called hidden layers and we can have many, each one with a different number of neurons. We will see later that neurons, represented by the circles, will be interconnected with each other in different ways between the neurons of the different layers.

In general, today we are handling artificial neural networks with many layers, which are literally stacked one on top of the other; hence the concept of deep (depth of the network). Each of them is composed of many neurons, each with its parameters (the weights w_i and the bias b) which perform a simple transformation of the data that they receive from neurons of the previous layer to pass them to those of the posterior layer. The union of all these transformations carried out by the neurons of the network is what allows discovering complex patterns in the data.

Before finishing this section, I would like to give the reader a magnitude of the problem that involves programming the Deep Learning algorithms these days: different layers serve different purposes, and each parameter and hyperparameter matter a lot in the final result; this makes it extremely complicated when trying to refine the programming of a neural network model, looking more like an art than a science for those who enter the area for the first time. But this does not imply that it is something mysterious, although it is true that much remains to be investigated, but it simply takes many hours of learning and practice.

The next figure visually summarizes the intuitive idea that Deep Learning is just a part of Artificial Intelligence, although nowadays it is probably the most dynamic one and the one that is captivating the scientific community. And the same way I previously mentioned the work of Stuart Russell and Peter

Novig as the main book on Artificial Intelligence, when it comes to Deep Learning we can read an excellent book titled *Deep Learning*⁴², by Ian Goodfellow, Yoshua Bengio and Aaron Corville, which constitutes the grounds for learning about this topic more deeply.



1.3 Why now?

In just ten years⁴³, four of the five largest companies in the world by market capitalization have changed: Exxon Mobil, General Electric, Citigroup and Shell Oil are out and Apple, Alphabet (the parent company of Google), Amazon and Facebook have taken their place. Only Microsoft maintains its position. You have already realized that all of them dominate the new digital

⁴² Deep Learning. I Goodfellow, Y. Bengio and A Corville. MIT Press 2016. Also freely available online at <http://www.deeplearningbook.org> [consulta: 20/01/2018].

⁴³ The data that appear in this section are available at the time of writing this section (Spanish version of the book) at the beginning of the year 2018.

era in which we find ourselves immersed. We are talking about companies that base their power on Artificial Intelligence in general and Deep Learning in particular.

John McCarthy coined the term Artificial Intelligence in the 1950s, being one of the founding fathers of Artificial Intelligence along with Marvin Minsky. Also in 1958 Frank Rosenblatt built a prototype neuronal network, which he called the *Perceptron*. In addition, the key ideas of the Deep Learning neural networks for computer vision were already known in 1989; also the fundamental algorithms of Deep Learning for time series such as LSTM were already developed in 1997, to give some examples. So, why now this Artificial Intelligence boom?

Undoubtedly, the available computing has been the main trigger, as we have already presented previously. However, other factors have contributed to unleash the potential of Artificial Intelligence and related technologies. Next we are going to talk about the most important factors that have influenced it.

The data, the fuel for Artificial Intelligence

Artificial Intelligence requires large datasets for the training of its models but, fortunately, the creation and availability of data has grown exponentially thanks to the enormous decrease in cost and increased reliability of data generation: digital photos, cheaper and precise sensors, etc. Furthermore, the improvements in storage hardware of recent years, associated with the spectacular advances in technology for its management with NoSQL⁴⁴ databases, have allowed having enormous datasets to train Artificial Intelligence models.

⁴⁴ Wikipedia, NoSQL. [online]. Available at: <https://es.wikipedia.org/wiki/NoSQL> [Accessed: 15/04/2018]

Beyond the increases in the availability of data that the advent of the Internet has led to recently, specialized data resources have catalyzed the progress of the area. Many open databases have supported the rapid development of Artificial Intelligence algorithms. An example is the ImageNet⁴⁵ database, of which we have already spoken, freely available with more than 10 million images tagged by hand. But what makes ImageNet special is not precisely its size, but the competition that was carried out annually with it, being an excellent way to motivate researchers and engineers.

While in the early years the proposals were based on traditional computer vision algorithms, in 2012 Alex Krizhevsky used a Deep Learning neural network, now known as *AlexNet*, which reduced the error rate to less than half of what the winner of the previous edition of the competition got. Already in 2015, the winning algorithm rivaled human capabilities, and today Deep Learning algorithms far exceed the error rates in this competition of those who have humans.

But ImageNet is only one of the available databases that have been used to train Deep Learning networks lately; many others have been popular, such as: MNIST⁴⁶, STL⁴⁷, COCO⁴⁸, Open Images⁴⁹, Visual Question Answering⁵⁰, SVHN⁵¹, CIFAR-10/100⁵², Fashion-MNIST⁵³, IMDB

⁴⁵ The ImageNet Large Scale Visual Recognition Challenge (ILSVRC). [online]. Available at: www.image-net.org/challenges/LSVRC. [Accessed: 12/03/2018]

⁴⁶ MNIST [online]. Available at: <http://yann.lecun.com/exdb/mnist/> [Accessed: 12/03/2018]

⁴⁷ STL [online]. Available at: <http://ai.stanford.edu/~acoates/stl10/> [Accessed: 12/03/2018]

⁴⁸ See <http://cocodataset.org>

⁴⁹ See <http://github.com/openimages/dataset>

⁵⁰ See <http://www.visualqa.org>

⁵¹ See <http://ufldl.stanford.edu/housenumbers>

⁵² See <http://www.cs.toronto.edu/~kriz/cifar.html>

⁵³ See <https://github.com/zalandoresearch/fashion-mnist>

Reviews⁵⁴, Twenty Newsgroups⁵⁵, Reuters-21578⁵⁶, WordNet⁵⁷, Yelp Reviews⁵⁸, Wikipedia Corpus⁵⁹, Blog Authorship Corpus⁶⁰, Machine Translation of European Languages⁶¹, Free Spoken Digit Dataset⁶², Free Music Archive⁶³, Ballroom⁶⁴, The Million Song⁶⁵, LibriSpeech⁶⁶, VoxCeleb⁶⁷, The Boston Housing⁶⁸, Pascal⁶⁹, CVPPP Plant Leaf Segmentation⁷⁰, Cityscapes⁷¹.

It is also important to mention here Kaggle⁷², a platform that hosts competitions of data analysis where companies and researchers contribute and share their data while data engineers from around the world compete to create the best prediction or classification models.

Entering into an era of computation democratization

However, what happens if you do not have this computing capacity in your company? Artificial Intelligence has until now been mainly the toy of big technology companies like Amazon, Baidu, Google or Microsoft, as well as some new companies that had these capabilities. For many other businesses

⁵⁴ See <http://ai.stanford.edu/~amaas/data/sentiment>

⁵⁵ See <https://archive.ics.uci.edu/ml/datasets/Twenty+Newsgroups>

⁵⁶ See <https://archive.ics.uci.edu/ml/datasets/reuters-21578+text+categorization+collection>

⁵⁷ See <https://wordnet.princeton.edu>

⁵⁸ See <https://www.yelp.com/dataset>

⁵⁹ See <https://corpus.byu.edu/wiki>

⁶⁰ See <http://u.cs.biu.ac.il/~koppel/BlogCorpus.htm>

⁶¹ See <http://statmt.org/wmt11/translation-task.html>

⁶² See <https://github.com/Jakobovski/free-spoken-digit-dataset>

⁶³ See <https://github.com/mdeff/fma>

⁶⁴ See <http://mtg.upf.edu/ismir2004/contest/tempoContest/node5.html>

⁶⁵ See <https://labrosa.ee.columbia.edu/millionsong>

⁶⁶ See <http://www.openslr.org/12>

⁶⁷ See <http://www.robots.ox.ac.uk/~vgg/data/voxceleb>

⁶⁸ See <https://archive.ics.uci.edu/ml/machine-learning-databases/housing/housing.names>

⁶⁹ See <http://host.robots.ox.ac.uk/pascal/VOC/>

⁷⁰ See <https://www.plant-phenotyping.org/CVPPP2017>

⁷¹ See <https://www.cityscapes-dataset.com>

⁷² Kaggle [online]. Available at: <http://www.kaggle.com> [Accessed: 12/03/2018]

and parts of the economy, artificial intelligence systems have so far been too expensive and too difficult to fully implement the hardware and software required.

But now we are entering another era of democratization of computing, and companies can have access to large data processing centers of more than 28,000 square meters (four times the field of Barcelona football club (Barça)), with hundreds of thousands of servers inside. We are talking about Cloud Computing⁷³.

Cloud Computing has revolutionized the industry through the democratization of computing and has completely changed the way business operates. And now it is time to change the scenario of Artificial Intelligence and Deep Learning, offering a great opportunity for small and medium enterprises that cannot build this type of infrastructure, although Cloud Computing can offer it to them; in fact, it offers access to a computing capacity that previously was only available to large organizations or governments.

Besides, Cloud providers are now offering what is known as Artificial Intelligence algorithms as a Service (AI-as-a-Service), Artificial Intelligence services through Cloud that can be intertwined and work together with internal applications of companies through simple protocols based on API REST⁷⁴.

This implies that it is available to almost everyone, since it is a service that is only paid for the time used. This is disruptive, because right now it allows software developers to use and put virtually any artificial intelligence algorithm into production in a heartbeat.

⁷³ Empresas en la nube: ventajas y retos del Cloud Computing. Jordi Torres. Editorial Libros de Cabecera. 2011.

⁷⁴ Wikipedia. REST. [online]. Available at: https://en.wikipedia.org/wiki/Representational_state_transfer [Accessed: 12/03/2018]

Amazon, Microsoft, Google and IBM are leading this wave of AIaaS services, are put into production quickly from the initial stages (training). At the time of writing this book, Amazon AIaaS was available at two levels: predictive analytics with *Amazon Machine Learning*⁷⁵ and the *SageMaker*⁷⁶ tool for rapid model building and deployment. Microsoft offers its services through its *Azure Machine Learning* which can be divided into two main categories as well: *Azure Machine Learning Studio*⁷⁷ and *Azure Intelligent Gallery*⁷⁸. Google offers *Prediction API*⁷⁹ and the Google *ML Engine*⁸⁰. IBM offers AIaaS services through its *Watson Analytics*⁸¹. And let's not forget solutions that already come from *startups*, like *PredicSis*⁸² or *BigML*⁸³.

Undoubtedly, Artificial Intelligence will lead the next revolution. Its success will depend to a large extent on the creativity of the companies and not so much on the hardware technology, in part thanks to Cloud Computing.

An open-source world for the Deep Learning community

Some years ago, Deep Learning required experience in languages such as C++ and CUDA; Nowadays, basic Python skills are enough. This has been possible thanks to the large number of open source software frameworks that have been appearing, such as Keras, central to our book. These frameworks greatly facilitate the creation and training of the models and allow abstracting

⁷⁵ Amazon ML [online]. Available at: <https://aws.amazon.com/aml/> [Accessed: 12/03/2018]

⁷⁶ SageMaker [online]. Available at: <https://aws.amazon.com/sagemaker/> [Accessed: 12/03/2018]

⁷⁷ Azure ML Studio [online]. Available at: <https://azure.microsoft.com/en-us/services/machine-learning-studio/> [Accessed: 12/03/2018]

⁷⁸ Azure Intelligent Gallery [online]. Available at: <https://gallery.azure.ai> [Accessed: 12/03/2018]

⁷⁹ Google Prediction API [online]. Available at: <https://cloud.google.com/prediction/docs/> [Accessed: 12/03/2018]

⁸⁰ Google ML engine [online]. Available at: <https://cloud.google.com/ml-engine/docs/technical-overview> [Accessed: 12/03/2018]

⁸¹ Watson Analytics [online]. Available at: <https://www.ibm.com/watson/> [Accessed: 12/03/2018]

⁸² PredicSis [online]. Available at: <https://predicsis.ai> [Accessed: 12/03/2018]

⁸³ BigML [online]. Available at: <https://bigml.com> [Accessed: 12/03/2018]

the peculiarities of the hardware to the algorithm designer to accelerate the training processes.

The most popular at the moment are TensorFlow, Keras and PyTorch, because they are the most dynamic at this time if we rely on the contributors and commits or stars of these projects on GitHub⁸⁴.

In particular, TensorFlow has recently taken a lot of impulse and is undoubtedly the dominant one. It was originally developed by researchers and engineers from the Google Brain group at Google. The system was designed to facilitate Machine Learning research and make the transition from a research prototype to a production system faster. If we look at the Github page of the project⁸⁵ we will see that they have, at the time of writing this book, more than 35,000 commits, more than 1500 contributors and more than 100,000 stars. Not despicable at all.

TensorFlow is followed by Keras⁸⁶, a high level API for neural networks, which makes it the perfect environment to get started on the subject. The code is specified in Python, and at the moment it is able to run on top of three outstanding environments: TensorFlow, CNTK or Theano. Keras has more than 4500 commits, more than 700 contributors and more than 30,000 stars⁸⁷.

PyTorch and Torch⁸⁸ are two Machine Learning environments implemented in C, using OpenMP⁸⁹ and CUDA to take advantage of highly parallel infrastructures. PyTorch is the most focused version for Deep Learning and based on Python, developed by Facebook. It is a popular environment in this

⁸⁴ See <https://www.kdnuggets.com/2018/02/top-20-python-ai-machine-learning-open-source-projects.html>

⁸⁵ See <https://github.com/tensorflow/tensorflow>

⁸⁶ See <https://keras.io>

⁸⁷ See <https://github.com/keras-team/keras>

⁸⁸ See <http://pytorch.org>

⁸⁹ See <http://www.openmp.org>

field of research since it allows a lot of flexibility in the construction of neural networks and has dynamic tensors, among other things. At the time of writing this book, Pytorch has more than 12,000 commits, around 650 contributors and more than 17,000 stars⁹⁰.

Finally, and although it is not an exclusive environment of Deep Learning, it is important to mention Scikit-learn⁹¹, that is used very often in the Deep Learning community for the preprocessing of data⁹². Scikit-learn has more than 22500 commits, more than 1000 contributors and nearly 30,000 stars⁹³.

But as we have already advanced, there are many other frameworks oriented to Deep Learning. Those that we would highlight are Theano⁹⁴ (Montreal Institute of Learning Algorithms), Caffe⁹⁵ (University de Berkeley), Caffe2⁹⁶ (Facebook Research) , CNTK⁹⁷ (Microsoft), MXNET⁹⁸ (supported by Amazon among others), Deeplearning4j⁹⁹, Chainer¹⁰⁰ , DIGITS¹⁰¹ (Nvidia), Kaldi¹⁰², Lasagne¹⁰³, Leaf¹⁰⁴, MatConvNet¹⁰⁵, OpenDeep¹⁰⁶, Minerva¹⁰⁷ and SoooA¹⁰⁸ , among many others.

⁹⁰ See <https://github.com/pytorch/pytorch>

⁹¹ See <http://scikit-learn.org>

⁹² See <http://scikit-learn.org/stable/modules/preprocessing.html>

⁹³ See <https://github.com/scikit-learn/scikit-learn>

⁹⁴ See <http://deeplearning.net/software/theano>

⁹⁵ See <http://caffe.berkeleyvision.org>

⁹⁶ See <https://caffe2.ai>

⁹⁷ See <https://github.com/Microsoft/CNTK>

⁹⁸ See <https://mxnet.apache.org>

⁹⁹ See <https://deeplearning4j.org>

¹⁰⁰ See <https://chainer.org>

¹⁰¹ See <https://developer.nvidia.com/digits>

¹⁰² See <http://kaldi-asr.org/doc/dnn.html>

¹⁰³ See <https://lasagne.readthedocs.io/en/latest/>

¹⁰⁴ See <https://github.com/autumnai/leaf>

¹⁰⁵ See <http://www.vlfeat.org/matconvnet/>

¹⁰⁶ See <http://www.opendeep.org>

¹⁰⁷ See <https://github.com/dmlc/minerva>

¹⁰⁸ See <https://github.com/laonbud/SoooA/>

An open-publication ethic

In the last few years, in this area of research, in contrast to other scientific fields, a culture of open publication has been generated, in which many researchers publish their results immediately (without waiting for the approval of the peer review usual in conferences) in databases such as arxiv.org of Cornell University (arXiv)¹⁰⁹. This implies that there are numerous softwares available in open source associated with these articles, which allow this field of research to move tremendously quickly, since any new discovery is immediately available for the whole community to see it and, if it is the case, build on top a new proposal.

This is a great opportunity for users of these techniques. The reasons for research groups to openly publishing their latest advances can be diverse. For example, articles rejected in main conferences in the area can propagate solely as a preprint on arxiv. This is the case of one key paper for the advancement of Deep Learning written by G. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever and R. Salakhutdinov that introduced Dropout mechanism¹¹⁰. This paper was rejected from NIPS in 2012¹¹¹.

Or Google, when publishing the results, consolidates its reputation as a leader in the sector, attracting the next wave of talent, which is one of the main obstacles to the advancement of the topic.

Advances in algorithms

Thanks to the improvement of the hardware that we have already presented and to having more computing capacity by the scientists who were

¹⁰⁹ See <https://arxiv.org>

¹¹⁰ G. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever and R. Salakhutdinov "Improving neural networks by preventing co-adaptation of feature detectors" <https://arxiv.org/pdf/1207.0580.pdf>

¹¹¹ See <https://twitter.com/ChrisFiloG/status/1009594246414790657>

researching in the area, it has been possible to advance dramatically in the design of new algorithms that have allowed to overcome important limitations detected in the previous algorithms. For example, not until many years ago it was very difficult to train multilayer networks from an algorithm point of view. But in this last decade there have been impressive advances with improvements in activation functions, use of pre-trained networks, improvements in training optimization algorithms, etc. Today, algorithmically speaking, we can train models of hundreds of layers without any problem.