Andrea Emilio Rizzoli – Director, Dalle Molle Institute for Artificial Intelligence (IDSIA USI-SUPSI), Switzerland



When we collaborate with any industry, we have to start by doing the painstaking work of understanding their needs and their language to create the bridge and to be able to talk to them

25.03.2021

Tags: Switzerland, IDSIA USI-SUPSI, Artificial Intelligence, Digital

Andrea Emilio Rizzoli of the Dalle Molle Institute for Artificial Intelligence, among the top 10 in the world in its field of research, outlines the focus of the Institute today, the synergies between AI and healthcare, and how pharma companies might structure their AI investments in the future.

Andrea, what first brought you to IDSIA and what made you decide to stay, being appointed director in December 2020?

I have been working at the Institute for a very long time, since 1996. Initially it was a part-time collaboration, which eventually became full-time in 2000. Having met the previous director, Professor Luca M. Gambardella, who is still with the Institute, he offered me the opportunity to work on a very interesting research project. So, I know IDSIA very well and I have seen it grow from a small institute with less than ten researchers to now over 80 and counting. As a matter of fact, we are expanding, as we are joining forces with professors from Università della Svizzera Italiana and their research groups.

AI has become a buzzword today. Can you tell us about the tipping point for IDSIA and what the focus of the Institute today is?

The tipping point for us dates back to 2010 when graphic processing units (GPUs), originally developed for video gaming, were made available to the wider community of researchers. People realized that they could use them to run parallel algorithms to train deep neural networks (DNNs). DNNs are everywhere now but in the 90s to the 2000s, they were not so common, simply because it used to be impossible to train them and to make them converge.

There was another important development. Professor Jürgen Schmidhuber, who had also been with IDSIA for a very long time as scientific director, developed a neural network architecture called 'long short-term memory' (LSTM) in the late-1990s that was able to manage data sequences and correlations. That became fundamental for applications such as voice and speech recognition, autonomous driving, etc. This theory developed by Professor Schmidhuber and his PhD student Sepp Hochreiter (now professor at Linz University in Austria), which enabled the convergence of these neural networks from a mathematical point of view, and the hardware made available at the beginning of the century, together opened the way to a number of applications that are now widely used. The year 2015 was when such DNNs really began to be used in a variety of fields and applications, but we have been working on them for a much longer time.

As a non-profit research institution, what challenges do you face when it comes to raising money for your research programs?

Our self-funding rate is about 72 percent. We tap into different resources, from basic research funds such as those provided by the Swiss National Science Foundation (SNSF), which grants us funds to study the principles and fundamentals of neural networks and data science. Then we have technology transfer and innovation funds such as the ones provided by Innosuisse, the Swiss Innovation Agency, which also offer us very interesting opportunities to partner with industry players to develop joint projects, the so-called "Innovation projects with an implementation partner".

There are also funding sources at the European level, both for basic research, such as the European Research Council grants and other schemes, and we are also quite active there.

Of course, we collaborate with the industry directly. For instance, we have an existing collaboration with Novartis. We are also working together with Hoffmann-Roche using natural language processing techniques to process disease related conversations on social media. They are interested to understand the knowledge and to process it to through machines while retaining the connection with the human counterpart.

Neural networks can often offer out-of-the-box solutions. For instance, in cancer, some cells become cancerous and others do not, so we are trying to develop a new breed of AI called *explainable AI*, where we try to explain why a particular decision had been reached and also to incorporate decades of human knowledge and experience from a physician into an AI procedure that can match it with analytical processes and expertise.

The other 28 percent we obtain from the academic institutions we belong to and for teaching activities. The universities also provided us with the beautiful building we are currently based in.

There are some that are quite opposed to the use of AI or this type of algorithm-driven approach to the healthcare sector. Before we talk about the advantages, what are the risks of this approach?

The ethics of AI have been discussed in many forums but there is a serious problem because when you train an algorithm with existing data, all these data contain all the decisions that had previously been made. For instance, the typical case is a credit evaluation. If people have previously decided to privilege certain groups of people, the algorithm will learn from that and will tend to reproduce the same biased decisions that have been made in the past. That is why it is dangerous to work on AI without thinking about how to avoid reproducing errors made in the past. There are also algorithms that evolve and learn, so if we do not give them guidance in terms of the ethical principles they should follow, then we might run the risk of creating rogue AI that could be detrimental or even dangerous for humans in the future. This is why there is so much attention about ethics and AI recently.

I think we have a lot to learn from the pharmaceutical sector in terms of the rigorous standards they have to comply with when they develop drugs, run clinical trials involving humans and report adverse effects. This model is definitely a good one for the AI sector to look at.

In that case, should we be looking at some sort of regulatory body for AI and other data science applications?

We already have some initiatives in this area. The GDPR at the European Union level is a first step in this direction. Professor Schmidhuber has expressed concern about what could happen in the future with the competition from the U.S. and China, who might not have the same ethics or concerns of a European institution. That is an assumption since there is no proof so far, but when it

comes to military research in the U.S. or China, it is never clear what they are working on.

A couple of years ago, a group of scientists and researchers in this area wrote an open letter to the United Nations urging AI not to be employed in the arms race. But military interests often come before such ethical considerations, as they have in other fields. As we know, inventions and ideas can be developed for good but used for evil purposes.

Looking at the life sciences, what are the kinds of requests that healthcare or pharma companies might come to the Institute for?

I mentioned the use of *explainable AI* before but in general, all the various applications of AI could be of interest to the pharma and healthcare sectors.

What is important is that when we collaborate with any industry, we have to start by doing the painstaking work of understanding their needs and their language to create the bridge and to be able to talk to them. We have started working with other research institutes like the Oncology Institute of Southern Switzerland (IOSI), who have very in-depth knowledge in their respective fields. But before we can work with them, we have to understand their problems and personally, on our side, it is not imaginable that we will ever reach their knowledge depth on the subject. At the same time, they do not know our side well, so we need this exchange of experiences to create and drive new development.

For standard neural networks, for instance, you have to pretreat the data so that the machine can process and extract the relevant structures. It is not a black box where anyone can take any data, put it into the box, and generate beautiful results. There is much more work behind the scenes that has to be done before running the algorithms.

With health data, pseudonymizing the data is usually essential and in theory, the idea is simple: you create an ID for each patient and then link it with the dataset but it has been shown that in some situations, using the right combination of anonymized data, you can still trace it back to the original person. There is ongoing research in terms of making data totally anonymized. This has been done in other areas like network signals or smart meter measurements but I am not sure about the case of the health sector.

How should pharma companies be advancing their capabilities or investments in AI? Should this be done internally or through external collaborations?

I cannot really give advice to industry executives but I can take an example of our collaboration with UBS, a leading player in the global banking sector: they wanted to develop a tight collaboration with us, which led them to open a research and development center in Ticino with a group of scientists. At this center we also have people embedded within UBS to work collaboratively, and this is actually required since UBS wanted us to work with their data only on their premises so as to ensure the highest possible standards for data security and privacy, something that could also be relevant to the pharma sector.

Further advice also goes back to the topic strict collaboration. A pharma company should not outsource or delegate this type of knowledge or expertise entirely to external players, preferring instead a shared language and an exchange of domain expertise and information.

We are about to start a new EU-funded project on AI-driven drug discovery efforts in conjunction with Janssen, Bayer and AstraZeneca, that will be funding PhD students to work with us for 18 months, and 18 months with one of the Big Pharma companies, so their PhDs will involve AI research and its applications to drug R&D. This will breed a new type of scientists able to bridge the gap between AI and drug R&D.

On that note, how do you feel about the shortage of resources and talent because I presume you cannot produce armies of PhDs in your field, right?

That is a big issue, along with competition in terms of salaries for post-docs, because once you have a PhD in such a highly competitive and in-demand field like this one, you can really aim for very high salaries that we, as universities, are not able to afford. Big Pharma can, however, and this affects the sustainability in the field. If all the talents end up in the industry, who is going to teach the next generations?

What can be a solution to this? In Europe, translational science is not as strong, scientists are tight on resources and there is no revolving door where people move from research to industry and back to research again.

This is a problem that we have in Europe: we are maybe not as dynamic as in the U.S. and also in China, because China is all about being fast and bending the rules a little, maybe, to accelerate science. In Europe, it is sometimes positive that we have our own rules and processes, but I do not see an easy way out unless we have these revolving doors and more collaborations between industry and academia. The European industry does not fund academia the same way that the industry funds academia in the U.S. If we can obtain important grants from industries, we can pay competitive salaries or provide career paths that are also attractive for brilliant researchers. If the companies could obtain tax deductions on that, they could be motivated to move in this direction.

The researchers also like to do their own things, which may not be what the industry likes. You need to also maintain the right balance between fostering the interest of researchers, developing new ideas and new ways of solving problems while meeting the demands of an industry that seeks working and practical solutions that may not be so attractive from a theoretical point of view.

You mentioned China, who is trying to become the global leader in AI. Is there some national interpretation of AI or is there some human or boundary fluctuation? Do you have collaborations with China?

Not directly at the moment, but we have collaborations on scientific issues. As you said, China is seeking to be recognized as the leader in a number of areas of AI research, from academia to industrial applications. Their President aims at proving that new algorithms, new ideas, and new applications are emerging from China.

But science is not something that you can put a label on of whether it is, Chinese or American or European. Very often, some of the best ideas stem from collaborations between scientists of many countries, though they are all affiliated with certain institutions. It is a bit like Lionel Messi playing for Barcelona. Sure, Barcelona is a Spanish team but if you look at all the nationalities on the individual players, it does not look so Spanish after all.

What reputation does the Institute have globally?

We do have a reputation of excellence at the international level and we have to maintain it. I think it is important to build strategic alliances with other research institutions in Switzerland because we are a small country and we do not have the funds that other larger countries do. It is an important but tough job for us to maintain our visibility and keep the Institute on the map to represent

Switzerland in the global AI space.

Moving ahead, what strategic areas would you like the Institute to pursue?

For the moment, we want to focus on healthcare and the environment, because we are not that active in these areas so far.

The environment and climate change are important because the pharma industry might pay for research in the end but so far no one really seems willing to invest in preventing climate change. It is also an issue because AI currently has a negative impact on the environment because of the huge amount of energy needed to train AI algorithms. For instance, take the GPT-3, a powerful algorithm that is able to write newspaper articles. The electricity bill alone for the training of this algorithm tops 1 million U.S. dollars.

In that sense, I think it is important to map the value of AI algorithms. Of course, we can also develop AI algorithms to support the efficient management of resources. Many complex systems can be optimized using these algorithms. But we have to keep in mind the trade-offs between the energy invested in trying to solve the problem and the energy that will be saved from the solution. If it is a one-off application, maybe it is not worthwhile, but if it can be scaled up, there might be a lot of value.

A final message?

Professor Schmidhuber believes that at a certain stage, AI will become more intelligent than us. His dream is to be replaced by an algorithm that could become the AI scientist in his place. For him, we can conquer the universe through robots designed by humans. Humans obviously cannot even travel to our closest solar system at the moment, and we have to send robots instead.

Personally, this is not my conviction because I think that machines may become intelligent enough for us to talk to and understand, but they will never be humans. Humans are quite unique. I think AI is about the interaction between man and machine, which are two complementary elements that have to work together in order to be effective.

See more interviews