The Creation of AI at FBK: IRST

- The Autonomous Province of Trento (PAT) had already established the University of Trento through the Trentino Institute of Culture in the 1960s. When it passed to the Government in the eighties, PAT decided, with President Senator Kessler, to invest in research, with the idea of offering new opportunities to Trentino and an avant-garde prospects to its young people. The IRST (Institute for Scientific and Technological Research), i.e. the scientific part of FBK, was thus founded more than thirty years ago. It became an important independent research institute dedicated to two topics: artificial intelligence and microsystems. It has had the great ability to stick to these topics, looking forward even at times when they were criticized.

- In the first phase, with its founder Luigi Stringa (formerly managing director of the Selenia - Elsag group), the goal was to build an integrated system of artificial intelligence at IRST. This also included an implementation part and therefore an intelligent service robot. The idea was that this long term goal of deep integration would bring with it an integration made in the MAIA system, which would favor the growth of a number of parallel research areas. A role was also planned for the Microsystems Division to manufacture smart sensors to be integrated into MAIA's robot.

- Thus frontier research groups were formed in artificial vision, speech recognition, machine reasoning, natural language processing, knowledge based systems, human factors. Each of these activities could count on a well-established researcher to coordinate each specific field. Many of the best young researchers from Italy and abroad were recruited, and in a short time the institute came to be ranked among the best research centers on the international scene. This allowed, among other things, to have a large number of projects accepted in the competition for European projects, and important industrial projects.

AI at the time of Stringa

Artificial intelligence (AI) was created with the aim of studying the theoretical foundations underlying human knowledge and reasoning, but at the same time also for the purposes of building computerized systems able to solve problems and perform reasoning operations typical of human beings. At the time of Stringa, the debate on the vision of a ***Strong AI*** *,* which tried to simulate the real processes of human reasoning on the computer - see, for example, deductive reasoning through tools such as formal logic or those of learning through neural networks - as opposed to the vision of a ***Weak AI***, according to which the important aspect was to build systems that could solve problems even with methods and techniques that did not simulate human reasoning, was still strong.

Be it weak or strong, AI research at that time was organized into **several rather sectorial and discrete strands** , such as the representation of knowledge and reasoning based on machine deduction, machine planning, cognitive models and processing of natural language, artificial vision and speech recognition (only to mention some very relevant ones). The encouraging early results in these different areas gave rise to a series of high expectations as to their impact on the market and society that were not fully met: it was the so-called "*AI winter*" partly due to lack of appreciation of the highly innovative results that AI was producing (only few know that current "window" systems on PCs are a result of AI) and partly due to the application of theories and techniques in a series of toy *examples*, i.e. examples not complex enough to have an impact in practical applications. A famous case is the one of the"blocks world" for machine planning in robotics, an artificial world where the task was to move blocks on a table with operations that could never fail.

Stringa created the IRST - like other important research centers such as the DFKI in Germany and the SRI in the US - as one of the forerunners who were the interpreters of a philosophy in which the idea was indeed to work on the theoretical foundations of artificial intelligence, but putting them into practice and testing them on problems in the real world, whose complexity was very different from the famous *toy examples*. An essential element of this approach was the integration of different techniques that at the time were viewed, studied and implemented in a completely separate way. In fact, the goal in IRST was to build an integrated system of artificial intelligence that found an experimental platform in MAIA (Advanced Model of Artificial Intelligence), an intelligent service robot that navigated autonomously in the building, interacted with people and was able to react to unexpected situations. The idea was that this long term goal of deep integration would bring with it an integration made in the MAIA system, which would favor the growth of a number of parallel research areas.

This global momentum slowly allowed artificial intelligence to come out of its winter and enter a **"new AI spring"** *stage, in which the* first systems were created that led to great results, both theoretical and practical. And since we have mentioned the case of research in machine planning and the *toy example* of the blocks world, let us recall the success in those years (1998) of NASA's "Deep Space 1" space mission, in which an AI system was able to control commands for space exploration *tasks* on a spaceship (see article "Remote Agent: to boldly go where no AI system has gone before ").

For example, starting in the 1990s, AI techniques were developed that allowed to safely design and validate highly critical systems (as in the case of applications for space, avionics or rail transport). Another element that favored the new AI spring was the fact that research increasingly took into account interaction and gave up the idea that everything had to be machine-run. The important thing was to build systems that would be able to interact with people and the surrounding environment, in the most natural way possible and to the end of working together to solve problems. Hence not only systems capable of performing functions typical of human intelligence and replacing man, but also capable of interacting with man, taking account of the characteristics of human cognition. In this way the interaction techniques in natural language, speech recognition and transcription and artificial vision techniques were incredibly improved. The advanced techniques of interaction and artificial vision gave a strong impulse to robotics, to the extent that some robotic platforms were put on the market that effectively carried out the tasks of the old MAIA.

It was definitely "spring", but not yet "summer". Applications were important, but still focused on specific and highly specialized areas, such as space, *safety critical* systems, analysis of medical records systems, transcripts from dictation in healthcare, or on robotics.

The AI summer

In recent years,**the *"*AI summer*"*** has finally exploded. Two are the key elements of this explosion: one relates to technology, the other one to the market.

The technological element lies in the advancement of machine learning techniques, especially ***deep learning***, i.e. *machine learning* based on different hierarchical layers of representation (hence "deep"). This has been possible thanks to the increased computational capacity of recent years and the availability of huge amounts of data necessary for training the learning system, data coming from the Internet of things, from the environment sensors, from wearable sensors as well as from social networks and from the web.

The other element is market. The new "big players" of information technology (Google, Amazon, Facebook, eBay, ...), have realized that, with this huge amount of data in their possession, they could train artificial intelligence algorithms that "deep-learn" and find relationships among the data that could hardly be discovered differently. Even more: we could understand how certain phenomena could evolve over time. Think of the possible applications for example in health treatment and prevention, if we could rely on artificial intelligence algorithms able to predict how a disease will evolve in a patient, or (even better) to detect early signs that a person risks getting sick from a disease and thus prevent the disease instead of treating it. There is evidence that a very strong commercial interest exists.

Moreover, in 2011 "Deep Mind" was created in London by some AI researchers; in 2014 Google bought it and it became Google Deep Mind. The agreement between Google Deep Mind and three major US hospitals (Chicago, Stanford and San Francisco) to analyze the data of their patients is recent. But consider all the applications that tap into artificial intelligence techniques, for example for smart cities and communities: based on the movements and behavior of people you can predict how certain areas of the city will evolve, how will prices rise or fall, how the areas will become more or less safe and livable and, more importantly, specific actions can be put in place to improve life in the cities and change the habits of citizens, making cities less polluted, more livable and safer. Or how the data that is produced in the production processes of Industry 4.0, the intelligent factory, can allow to optimize processes or do preventive maintenance. A summer that appears to be really... hot!

And what next?

What will follow the hot summer? We will have to be careful not to repeat the mistakes of the past. Following the trends and thinking that *deep learning* can solve all the problems could push us back if not into winter... at least into an AI autumn. Therefore, we believe we need to integrate pure learning techniques with modeling techniques and representation of knowledge, where neural networks work together with logical representations and statistical techniques together with formal models. Take the instance of image comprehension: we learn from examples, of course, but at some point learning techniques must be integrated with the fact that we are aware of certain rules and have learned certain constraints. It should not be necessary to train a neural network with millions of images every time to understand that donkeys do not fly (except perhaps in cartoons) ...

Take the case of the big challenge of the smart factory. Big data are very useful to understand when components of an industrial machine equipped with all the sensors are wearing off, and therefore to perform preventive maintenance. But the engineer who designed it or runs it has a model of the machine in ìtheir head that is equally important.

Once again AI must be careful not to slip into silo thinking.