

Artificial Intelligence

Current challenges and Inria's engagement



Inria white papers look at major current challenges in informatics and mathematics and show actions conducted by our project-teams to address these challenges.

This document is the first produced by the Strategic Technology Monitoring & Prospective Studies Unit. Thanks to a reactive observation system, this unit plays a lead role in supporting Inria to develop its strategic and scientific orientations. It also enables the institute to anticipate the impact of digital sciences on all social and economic domains.

It has been coordinated by Bertrand Braunschweig with contributions from 45 researchers from Inria and from our partners. Special thanks to Peter Sturm for his precise and complete review.




Artificial Intelligence: a major research theme at Inria

1. Samuel and his butler	5
2. 2016, The year of AI?	8
3. The debates about AI	12
4. The challenges of AI and Inria contributions	16
4.1 Generic challenges in artificial intelligence	20
4.2 In machine learning	22
4.3 In signal analysis: vision, speech	32
4.4 In knowledge and semantic web	39
4.5 In robotics and self-driving vehicles	45
4.6 In neurosciences and cognition	51
4.7 In language processing	54
4.8 In constraint programming for decision support	56
4.9 In music and smart environments	59
5. Inria references: numbers	63
6. Other references for further reading	66
 Acknowledgment	 71





Samuel and his butler¹





7:15 a.m., Sam wakes up and prepares for a normal working day. After a quick shower, he goes and sits at the kitchen table for breakfast. Toi.Net², his robot companion, brings warm coffee and a plate of fresh fruits. *“Toi.Net, pass me the sugar please,”* Sam says. The robot brings the sugar shaker from the other end of the breakfast table – there is a sugar box in the kitchen cupboard but Toi.Net knows that it is much more convenient to use the shaker.

“Any interesting news?” Sam asks. The robot guesses she must find news that correspond to Sam’s topics of interest. She starts with football.

Toi.Net: *“Monaco beat Marseille 3-1 at home, it is the first time they score three goals against Marseille since the last twelve years. A hat trick by Diego Suarez.”*

Toi.Net: *“The Miss France election took place in Palais des Congrès, Miss Martinique was elected.”*

Sam: *“Please don’t bother me again with this kind of news, I don’t care about the misses ceremony.”*

Toi.Net: *“Alright. I won’t.”*

Toi.Net: *“The weather forecast for Paris is sunny in the morning, but there will be some heavy rain around 1:00p.m. and in the afternoon”*

1 The title of this section is a reference to Samuel Butler, a 19th-century English novelist, author of *Erewhon*, one of the first books to speculate about the possibility of an artificial intelligence grown by Darwinian selection and reproduction among machines.

2 Pronounce *‘toanet’*, after the name of the maid-servant in Molière’s *“The imaginary invalid”*

Toi.Net: *“Mr. Lamaison, a candidate for the presidency of the South-west region, declared that the unemployment level reached 3.2 million, its highest value since 2004.”*

Sam: *“Can you check this? I sort of remember that the level was higher in the mid 2010s.”*

Toi.Net (after two seconds): *“You’re right, it went up to 3.4 million in 2015. Got that from INSEE semantic statistics.”*

By the end of the breakfast, Sam does not feel very well. His connected bracelet indicates abnormal blood pressure and Toi.Net gets the notification. *“Where did you leave your pills?”* She asks Sam. *“I left them on the nightstand, or maybe in the bathroom.”* Toi.Net brings the box of pills, and Sam quickly recovers.

Toi.Net: *“It’s time for you to go to work. Since it will probably be raining when you go for a walk in the park after lunch, I brought your half boots.”*


An autonomous car is waiting in front of the house. Sam enters the car, which announces “I will take a detour through A-4 this morning, since there was an accident on your usual route and a waiting time of 45 minutes because of the traffic jam.”

Toi.Net is a well-educated robot. She knows a lot about Sam, understands his requests, remembers his preferences, can find objects and act on them, connects to the internet and extracts relevant information, learns from new situations. This has only been possible thanks to the huge progresses made in artificial intelligence (AI): speech processing and understanding (to understand Sam’s requests); vision and object recognition (to locate the sugar shaker on the table); automated planning (to define the correct sequences of action for reaching a certain situation such as delivering a box of pills located in another room); knowledge representation (to identify a hat trick as a series of three goals made by the same football player); reasoning (to decide to pick the sugar shaker rather than the sugar box in the cupboard, or to use weather forecast data to decide which pair of shoes Sam should wear); data mining (to extract relevant news from the internet, including fact checking in the case of the political declaration); her incremental machine learning algorithm will make her remember not to mention Miss France contests in the future; she continuously adapts her interactions with Sam by building her owner’s profile and by detecting his emotions.

Inria, with its 160+ project-teams in eight research centres, is active in all these scientific areas. This paper presents our views on the main trends and challenges in artificial intelligence and how our teams are actively conducting scientific research, software development and technology transfer around these challenges.



2016, the year of AI?



This is what a few Microsoft leading scientists said in recent interviews. **AI has become a popular topic** in the media and in scientific magazines, due to several achievements, many of them thanks to the developments in machine learning. Major companies including Google, Facebook, IBM, Microsoft but also car manufacturers such as Toyota, Volvo and Renault are active in AI research and plan to invest even more in the future, and many AI scientists are now leading the research laboratories of these and other companies.

AI research produced major achievements in the last decade, in several areas. The most publicised are those obtained in machine learning, thanks in particular to the development of deep learning architectures, multi-layered convolutional neural networks learning from massive volumes of data and trained on high performance computing systems. An example of such machine learning achievements is the solution of seven Atari games (bricks, space invaders etc.) by Google DeepMind, using the full pixel image as input to decide which action to take in

order to reach the highest expected return at the end of the game, and the more recent result obtained by the same team for the game of Go, beating the top player Lee Ledol in a series of five games, using a combination of Monte Carlo tree search, deep learning from real games, and reinforcement learning by playing with a clone of itself.

Other remarkable examples are:

- Automatic description of the content of an image (“a picture is worth a thousand words”), also by Google (<http://googleresearch.blogspot.fr/2014/11/a-picture-is-worth-thousand-coherent.html>)

- The results of Imagenet’s 2012 Large Scale Visualisation Challenge, won by a very large convolutional neural network developed by University of Toronto, and trained with Nvidia GPUs (<http://image-net.org/challenges/LSVRC/2012/results.html>)

- The quality of face recognition systems such as Facebook’s (<https://www.newscientist.com/article/dn27761-facebook-can-recognise-you-in-photos-even-if-youre-not-looking#.VYkVxFzjZ5g>)

- etc.

But this is only a subset of the results obtained with AI. The progress in robotics, in self-driving cars, speech processing, natural language understanding is also quite impressive.

.....
The progress in robotics, self driving cars, speech processing, natural language understanding is quite impressive.
.....

Examples are:

- The level of competence reached by robots in Darpa's Robotic Challenge, won by KAIST in 2015, where robots must drive a small car, open a door and pass through, close a valve, cut a hole in the wall with a drill, climb over debris or clear a path, and climb a flight of stairs (https://en.wikipedia.org/wiki/DARPA_Robotics_Challenge).
- Speech understanding is now considered a standard feature of smartphones and tablets with artificial companions including Apple's Siri, Microsoft's Cortana, Facebook's M, and others.
- Microsoft Skype Translator translates conversations in different languages in real time.
- Self-driving cars have driven thousands of kilometres without major incidents happening.

The image shows a Google search result for 'Mondol Kiri' in Paris. On the left, a white box indicates the restaurant is 'Fermé' (Closed) on Monday. Below this, there are links to TripAdvisor and Yelp, along with a note about 4.5/5 reviews. On the right, there is a map showing the location at Avenue Edouard Belin, a photo gallery, and a 'Voir les photos' button. The main section features the restaurant name 'Mondol Kiri', a 4.5/5 rating from Google, and a description: 'Sièges noirs et murs rouges ou en pierre ornés de tableaux créent un cadre cosy pour des plats cambodgiens.' It also lists the address (159 Avenue de Choisy, 75013 Paris), phone number (01 53 79 75 96), and hours (Closed today).

Figure 2 : Semantic Information added to Google search engine results

This section could not end without mentioning the results obtained in knowledge representation and reasoning, ontologies and other technologies for the semantic web and for linked data:

- **Google Knowledge Graph** improves the search results by displaying structured data on the requested search terms or sentences.
- **Schema.org** contains millions of RDF triplets describing known facts: search engines can use this data to provide structured information upon request.
- **The OpenGraph protocol** – which uses RDFa – is used by Facebook to enable any web page to become a rich object in a social graph.

Finally, another important trend is the recent opening of several technologies that were previously proprietary, in order for the AI research community to benefit from them but also to contribute with additional features.

Examples are:

- IBM's cognitive computing services for Watson, available through their Application Programming Interfaces, offers up to 20 different technologies such as speech-to-text and text-to-speech, concepts identification and linking, visual recognition and many others: <http://www.ibm.com/cloud-computing/bluemix/solutions/watson>





Figure 3: IBM Watson computer - © Clockready (CC BY-SA 3.0, via Wikimedia Commons)

- Google's TensorFlow software library for machine intelligence with neural networks – was put into open source: <https://www.tensorflow.org/>
- Facebook recently opensourced its Big Sur hardware design for running large deep learning neural networks on GPUs: https://code.facebook.com/posts/1687861518126048?__mref=message_bubble

All these positive achievements have been balanced by some concerns about the dangers of AI expressed by highly recognised scientists, which is the subject of the next section.



The debates about AI



Debates about AI really started in the 20th century – for example, think of **Isaac Asimov's Laws of Robotics** – but increased to a much higher level because of the recent progresses achieved by AI systems as shown above. The Technological Singularity Theory claims that a new era of machines dominating mankind will start when AI systems will become super-intelligent: *“The technological singularity is a hypothetical event related to the advent of genuine artificial general intelligence. Such a computer, computer network, or robot would theoretically be capable of recursive self-improvement (redesigning itself), or of designing and building computers or robots better than itself on its own. Repetitions of this cycle would likely result in a runaway effect – an intelligence explosion – where smart machines design successive generations of increasingly powerful machines, creating intelligence far exceeding human intellectual capacity and control. Because the capabilities of such a super intelligence may be impossible for a human to comprehend, the technological singularity is the point beyond which events may become unpredictable or even unfathomable to human intelligence”* (Wikipedia).

Advocates of the technological singularity are close to the **transhumanist movement**, which aims at improving physical and intellectual capacities of humans with new technologies. The singularity would be a time when the nature of human beings would fundamentally change, this being perceived either as a desirable event, or as a danger for mankind.

.....
*A new era
of machines
dominating
mankind will start
when AI systems
will become super-
intelligent*
.....

An important outcome of the debate about the dangers of AI has been the recent discussion on **autonomous weapons and killer robots**, supported by an open letter published at the opening of the IJCAI conference in 2015³. The letter, which asks for a ban of such weapons able to operate beyond human control, has been signed by thousands of individuals including Stephen Hawking, Elon Musk, Steve Wozniak and a number of leading AI researchers including some from Inria, contributors to this document.

Other dangers and threats that have been discussed in the community include: the financial consequences on the stock markets of high

³ See <http://futureoflife.org/open-letter-autonomous-weapons/>

frequency trading, which now represents the vast majority of orders placed, where supposedly intelligent software operate at a high rate leading to possible market crashes, as for the Flash Crash of 2010; the consequences of big data mining on privacy, with mining systems able to divulgate private properties of individuals by establishing links between their online operations or their recordings in data banks; and of course the potential unemployment caused by the progressive replacement of workforce by machines.




Figure 4: In the movie "Her" by Spike Jonze (Annapurna Pictures / Warner Bros, 2013), a man falls in love with his intelligent operating system

The more we develop artificial intelligence the greater the risk of developing only certain intelligent capabilities (e.g. optimisation and mining by learning) to the detriment of others for which the return on investment may not be immediate or may not even be a concern for the creator of the agent (e.g. moral, respect, ethics, etc.). There are many risks and challenges in the large scale coupling of artificial intelligence and people. In particular, if the artificial intelligences are not designed and regulated to respect and preserve humans, if, for instance, optimisation and performances are the only goal of their intelligence then this may be the recipe for large scale disasters where users are used, abused, manipulated, etc. by tireless and shameless artificial agents. We need to research AI at large including everything that makes behaviours intelligent and not only the most "reasonable aspects".

Dietterich and Horvitz lately published an interesting answer to some of these questions⁴. In their short paper, the authors recognise that the AI research community should pay moderate attention to the risk of loss of control by humans, because this is not critical in a foreseeable future, but should instead pay more attention to five near-term risks for AI-based systems, namely: bugs in software; cyberattacks; “The Sorcerer’s Apprentice”, that is, making AI systems understand what people intend rather than literally interpreting their commands; “shared autonomy”, that is, the fluid cooperation of AI systems with users, so that users can always take control when needed; and the socioeconomic impacts of AI, meaning that AI should be beneficial for the whole society and not just for a group of a happy few.

Inria is aware of these debates and acts as a research institute, engaged into science and technology transfer for societal welfare, and conscious of its responsibilities in front of the society. Informing the society and our governing bodies about the potentialities and risks of digital science and technologies is one of our missions.

Therefore, in the recent years, Inria:

- **Launched a reflexion** about ethics long before the threats of AI were subject of debates in the scientific society;

- **Contributed to the creation of Allistene’s CERNA**⁵, a think tank looking at ethics problems arising from research on digital science and technologies; the first recommendations report published by CERNA concerns the research on robotics;

- **Set up a body responsible for assessing the legal or ethical issues of research** on a case by case basis: the Operational Committee for the Evaluation of Legal and Ethical Risks (COERLE) with scientists from Inria and external contributors; COERLE’s mission is to help identify risks and determine whether the supervision of a given research project is required.

Moreover, Inria encourages its researchers to take part in the societal debates when solicited by press and media about ethical questions such as the ones raised on **robotics, deep learning, data mining and autonomous systems**.

This being said let us now look at the scientific and technological challenges for AI research, and at how Inria contributes to addressing these challenges: this will be the subject of the next section.

⁴ Dietterich, Thomas G. and Horvitz, Eric J., Rise of Concerns about AI: Reflections and Directions, Communications of the ACM, October 2015 Vol. 58 no. 10, pp. 38-40

⁵ Commission de réflexion sur l’Ethique de la Recherche en sciences et technologies du Numérique of Alliance des Sciences et Technologies du Numérique : <https://www.allistene.fr/cerna-2/>



The challenges of AI and Inria contributions



AI is a vast domain; any attempt to structure it in subdomains can be debated. We will use the keywords hierarchy recently proposed by the community of Inria team leaders in order to best identify their contributions to digital sciences in general. In this hierarchy, Artificial Intelligence is a top-level keyword with seven subdomains, some of them specific, some of them referring to other sections of the hierarchy: see the following table.

- **KNOWLEDGE**
 - _Knowledge bases
 - _Knowledge extraction & cleaning
 - _Inference
 - _Semantic web
 - _Ontologies
- **MACHINE LEARNING**
 - _Supervised Learning
 - _(Partially) unsupervised learning
 - _Sequential & reinforcement learning
 - _Optimisation for learning
 - _Bayesian methods
 - _Neural networks
 - _Kernel methods
 - _Deep learning
 - _Data mining
 - _Massive data analysis
- **NATURAL LANGUAGE PROCESSING**
- **SIGNAL PROCESSING**
 - _Speech
 - _Vision
 - Object recognition
 - Activity recognition
 - Search in image & video banks
 - 3D and spatiotemporal reconstruction
 - Objects tracking & movement analysis
 - Objects localisation
 - Visual servoing
- **ROBOTICS (including autonomous vehicles)**
 - _Design
 - _Perception
 - _Decision
 - _Action
 - _Robot interaction (environment/humans/robots)

- _ Robot fleets
- _ Robot learning
- _ Cognition for robotics & systems

■ **NEUROSCIENCES, COGNITIVE SCIENCES**

- _ Understanding & simulation of the brain & of the nervous system
- _ Cognitive sciences

■ **ALGORITHMICS OF AI**

- _ Logic programming & ASP
- _ Deduction, proof
- _ SAT theories
- _ Causal, temporal, uncertain reasoning
- _ Constraint programming
- _ Heuristic search
- _ Planning and scheduling

■ **DECISION SUPPORT**

Table 1: Inria keywords hierarchy for AI

In the following, Inria contributions will be identified by project-teams. The project-team is Inria’s original model for research: a project-team brings together, around a scientific personality, a group of researchers, research-lecturers, PhD students and engineers. They all have a common objective: to rise to a scientific and technological challenge in one of the institute’s priority research fields defined in the strategy plan. The maximum duration of a project-team is twelve years, eight years being the average. (The project-teams’ names will be written in **ORANGE CAPS**, so as to distinguish them from other nouns.)

FROM INRIA’S STRATEGIC PLAN “TOWARDS INRIA 2020”

Inria’s scientific strategy is built around two complementary dimensions, based on which the institute’s contribution will be articulated:

• **Digital sciences and technologies that serve individuals, society and knowledge.**

The major areas of work in which Inria is involved consider the individual as a central element in digital problems:

- The human as such: health and well-being
- The human and its environments: from the individual to society and

from habitats to the planet

- The human and knowledge: emergence, mediation and education

• **Priority scientific developments at the heart of our sciences:**

- Computing the future: models, software and digital systems

- Mastering complexity: data, networks and flows

- Interacting with the real and digital worlds: interaction, uses and learning

We are addressing key questions raised by other sciences and major fields of application, to which computer and mathematical sciences must contribute:

- **Health & well-being**

- **Energy & natural resources**

- **Environment & sustainable development**

- **Society & education**

After an initial subsection dealing with generic challenges, more specific challenges are presented in decreasing Inria investment order, that is, we start with domains where Inria has the highest number of active teams and researchers. This does not mean that the teams contributing to the last challenges are of lesser quality, but only that there are less Inria teams.

We do not provide definitions of AI and of subdomains: there is abundant literature about them. Good definitions can also be found on Wikipedia, *e.g.*:

https://en.wikipedia.org/wiki/Artificial_intelligence

https://en.wikipedia.org/wiki/Machine_learning

<https://en.wikipedia.org/wiki/Robotics>

https://en.wikipedia.org/wiki/Natural_language_processing

https://en.wikipedia.org/wiki/Semantic_Web

https://en.wikipedia.org/wiki/Knowledge_representation_and_reasoning

etc.

4.1 Generic challenges in artificial intelligence

The main generic challenges identified by Inria are as follows: (i) situated AI; (ii) user-in-the-loop; (iii) openness to other disciplines; (iv) multi-tasking; (v) validation and certification of AI systems.

Situated AI

AI systems must operate in the real world and interact with their environment: receive data from sensors, determine the context in which they operate, act on the real world; they have to behave autonomously and maintain their integrity in various conditions; these challenges imply that AI systems must deal with unstructured as well as with semantic data.

Human-in-the-loop

AI systems will deal with human users: they must be able to explain their behaviour, justify in some way the decisions that they make so that human users understand what is going on. If human users are not in such a position, there will be no or little trust in the systems, which won't be accepted. Moreover, AI systems need some flexibility and adaptation in order to deal with different users and different expectations. It is important to develop interaction mechanisms that support good communication and interoperation between humans and AI systems. Some Inria teams work on the collaboration between AI and man-machine interfaces, one of the desired features.

Openness to other disciplines

An AI will often be integrated in a larger system composed of many parts. Openness therefore means that AI scientists and developers will have to collaborate with specialists of other disciplines in computer science (*e.g.* modelling, verification & validation, networks, visualisation, man-machine interaction, etc.) to compose the wider system, and with non-computer scientists that contribute to AI *e.g.* psychologists, biologists (*e.g.* biomimetics), mathematicians, etc. A second aspect is the impact of AI systems on several facets of our life, our economy, and our society: collaboration with specialists from other domains (it would be too long to mention them, *e.g.* economists, environmentalists, biologists, lawyers, etc.) becomes mandatory.

Multitasking

Many AI systems are good at one thing but show little competence outside their focus domain; but real-life systems, such as robots must be able to undertake several actions in parallel, such as memorising facts, learning new concepts, acting on the real world and interacting with humans.

Validation and certification

A mandatory component in mission-critical systems, certification of AI systems, or their validation by appropriate means, is a real challenge especially if these systems fulfil the previous expectations (adaptation, multitasking, user-in-the-loop): verification, validation and certification of classical (*i.e.* non-AI) systems is already a difficult task – even if there are already exploitable technologies, some being developed by Inria project-teams – but applying these tools to complex AI systems is an overwhelming undertaking which must be approached if we want to put these systems in use in environments such as aircrafts, nuclear power plants, hospitals etc.

Other generic challenges

In addition to the previous challenges, the following desired properties for AI systems should trigger new research activities beyond the current ones: some are extremely demanding and cannot be addressed in the near term but nevertheless are worth considering.

Giving norms and values to AIs goes far beyond current science and technologies: for example, should a robot going to buy milk for his owner stop on his way to help a person whose life is in danger? Could a powerful AI technology be used for artificial terrorists? Current AI research is far from being able to address these demands.

The need for privacy is particularly relevant for AIs that are confronted with personal data, such as intelligent assistants/companions or data mining systems. This need is valid for non-AI systems too, but the specificity of AI is that new knowledge will be derived from private data and possibly made public if not restricted by technical means.

Scaling up: **AI systems must be able to handle vast quantities of data and of situations.** We have seen deep learning algorithms absorbing mil-

lions of data points (signal, images, video etc.) and large-scale reasoning systems such as IBM's Watson making use of encyclopaedic knowledge; however, the general question of scaling up for the many V's (variety, volume, velocity, vocabularies, ...) still remains.

Applications of AI

This is not really a challenge for AI, but it is important to emphasise that **AI systems contribute to address societal issues**: in environment and energy, in health and ambient assisted living, for transportation and smart cities, for business, etc. AI applications cover the whole spectrum of human activities. They can be beneficial to mankind and to economy, but they can also become threats if uncontrolled, see section 3.

4.2 Generic challenges in machine learning

Machine learning algorithms and systems made a steep progress in the recent years due to the availability of large volumes of data and of high performance computing, together with interesting progress in optimisation. A very powerful feature of deep learning is its capacity to learn the descriptors while clustering the data. However, there remain several limitations and challenges that we have grouped as follows: i) data sources; ii) symbolic vs. continuous representations; iii) continuous and never-ending learning; iv) learning under constraints; v) computing architectures; vi) unsupervised learning; vii) human-in-the-learning-loop, explanations.

Data sources

The challenges are to learn from heterogeneous data, available through multiple channels; to deal with uncertain information; to identify and process rare events beyond the purely statistic approaches; to work with knowledge sources as well as data sources, integrating models and ontologies in the learning process; and finally to obtain good learning performance with little data, in cases where big data sources are not common.

ORPAILLEUR is a project-team at Inria Nancy-Grand Est and Loria since the beginning of 2008. It is a rather large and special team as it includes computer scientists, but also a biologist, chemists, and a physician.

Life sciences, chemistry, and medicine, are application domains of first importance and the team develops working systems for these domains.

Knowledge discovery in databases – hereafter KDD – consists in processing a large volume of data in order to discover knowledge units that are significant and reusable. Assimilating knowledge units to gold nuggets, and databases to lands or rivers to be explored, the KDD process can be likened to the process of searching for gold. This explains the name of the research team: “Orpailleur” means gold digger in French. Moreover, the KDD process is iterative, interactive, and generally controlled by an expert of the data domain, called the analyst. The analyst selects and interprets a subset of the extracted units for obtaining knowledge units having a certain plausibility. Like a person searching for gold and having a certain knowledge of the task and of the location, the analyst may use his/her own knowledge but also knowledge on the domain of data for improving the KDD process.

A way for the KDD process to take advantage of domain knowledge is to be in connection with ontologies relative to the domain of data, for making a step towards the notion of knowledge discovery guided by domain knowledge or KDDK. In the KDDK process, the extracted knowledge units have still «a life» after the interpretation step: they are represented using a knowledge representation formalism to be integrated within an ontology and reused for problem-solving needs. In this way, knowledge discovery is used for extending and updating existing ontologies, showing that knowledge discovery and knowledge representation are complementary tasks and reifying the notion of KDDK.

Symbolic vs. continuous representations

Continuous representations allow machine-learning algorithms to approximate complex functions, while symbolic representations are used to learn rules and symbolic models. The main recent progresses have been for continuous representations; they leave reasoning away. One desired feature is to embed reasoning into continuous representations, that is, to find ways to make inferences on numeric data; on the other hand, in order to benefit from the power of deep learning, defining continuous representations of symbolic data can be quite useful, as has been done *e.g.* for text with word2vec and text2vec representations.

LINKMEDIA focuses on machine interpretation of professional and social multimedia content across all modalities. In this framework, artificial intelligence relies on the design of content models and associated learning algorithms to retrieve, describe and interpret messages edited for humans. Aiming at multimedia analytics, **LINKMEDIA** develops machine-learning algorithms primarily based on statistical and neural models to extract structure, knowledge, entities or facts from multimedia documents and collections. Multimodality and cross-modality to reconcile symbolic representations (e.g., words in a text or concepts) with continuous observations (e.g., continuous image or signal descriptors) is one of the key challenges for **LINKMEDIA**, where neural networks embedding appears as a promising research direction. Hoax detection in social networks combining image processing and natural language processing, hyperlinking in video collections simultaneously leveraging spoken and visual content, interactive news analytics based on content-based proximity graphs are among key subjects that the team addresses.

“User-in-the-loop” analytics, where artificial intelligence is at the service of a user, is also central to the team and raises challenges for humanly supervised machine-based multimedia content interpretation: humans need to understand machine-based decisions and to assess their reliability, two difficult issues with today’s data-driven approaches; knowledge and machine learning are strongly entangled in this scenario, requiring mechanisms for human experts to inject knowledge into data interpretation algorithms; malicious users will inevitably temper with data to bias machine-based interpretations in their favour, a situation that current adversarial machine learning can poorly handle; last but not least, evaluation shifts from objective measures on annotated data to user-centric design paradigms that are difficult to cast into objective functions to optimise.

Continuous and never-ending learning

Some AI systems are expected to be resilient, that is, to be able to operate on a 24/7 basis without interruptions. Interesting developments have been made for lifelong learning systems that will continuously learn new knowledge while they operate. The challenges are to operate online in real time and to be able to revise existing beliefs learned from previous cases, in a self-supervised way. Bootstrapping is an option for

such systems, where elementary knowledge learned in the first stages of operation is used to direct future learning tasks, such as in the NELL (never-ending language learning) system developed at Carnegie-Mellon University (<http://rtw.ml.cmu.edu/rtw/>).

Learning under constraints

Privacy is certainly the most important constraint that must be taken into account. The field of machine learning recently recognised the need to maintain privacy while learning from records about individuals; a theory of machine learning respectful of privacy is being developed by researchers such as Michael Jordan (<http://arxiv.org/abs/1210.2085>). In Inria, several teams work on privacy: especially **ORPAILLEUR** (see above) in machine learning, but also teams from other domains such as **PRIVATICS** (algorithmics of privacy) and **SMIS** (privacy in databases). More generally speaking, machine learning might have to cope with other external constraints such as decentralised data or energy limitations. Research on the wider problem of machine learning with external constraints is needed.

MAGNET

Machine learning, in one of its dimensions, is to identify regularities, patterns, in the data and represent them in a model. The concept of model is diverse: statistical, probabilistic or formal as a finite state machine etc. The purpose of learning is to then use this model to predict the values associated with new data. The approach differs in this from the attempt to reason, to represent and process knowledge, other conventional paradigms of AI.

In a traditional mathematical formalisation, however, the goal of machine learning is how to well approximate an unknown function that takes as input data representing the description of an object or an entity, and produces a result as a discrete or continuous value. This is for example to assign a positive, negative or neutral value to the representation of a short text in order to analyse sentiments in tweets or estimate the interest of a person for a new movie. Finding this function, these patterns require enough data. The more complex is the model, *e.g.* based on dependency relationships between data, the more difficult the task becomes. This is also the case when the function to learn does not calculate a scalar value but a whole structure such as a parsing tree of a sentence. **MAGNET** is particularly interested in the challenge of taking

into account structural relationships in the data, whether as input when the data set is shown in graph form, or whatsoever when it comes to predicting complex structure.

What was not possible twenty years ago becomes practical today. Machine learning observed spectacular improvements with the explosion of both computing power and massive data collection. However, what became possible with the abundance of data providing sufficient statistics to estimate complex functions raises issues of scalability and efficiency of algorithms. These fundamental questions are what the **MAGNET** team aims to address. A representation in graph form, whether given or adaptively calculated for a given task, allows explaining and operating an approximation of these dependencies. The issue of these dependencies is also very relevant in the case of language processing, for example when solving the co-reference problem, when trying to identify if two parts of a text refer to the same entity or the same person.

The challenges ahead are posed by new constraints: massive data is distributed, systems are decentralised. A challenge lies in the ability to exploit this massive data not anymore in terms of centralisation but instead as a set of autonomous and customised learning. The personal and private nature of the data adds an additional constraint limiting communication. Energy constraints obviously appear. Learning under such constraints in a data and processing network is a fundamental challenge and an alternative to the hyper centralisation observed today.

Computing Architectures

Modern machine learning (ML) systems need high performance computing and data storage in order to scale up with the size of data and with problem dimensions; algorithms will run on GPUs and other powerful architectures, data and processes must be distributed over many processors. New research must address how ML algorithms and problem formulations can be improved to make best usage of these computing architectures.

SIERRA addresses primarily machine learning problems, with the main goal of making the link between theory and algorithms, and between algorithms and high-impact applications in various engineering and scientific fields, in particular computer vision, bioinformatics, audio processing,

text processing and neuro-imaging. Recent achievements include theoretical and algorithmic work for large-scale convex optimisation, leading to algorithms that make few passes on the data while still achieving optimal predictive performance in a wide variety of supervised learning situations. Challenges for the future include the development of new methods for unsupervised learning and the design of learning algorithms for parallel and distributed computing architectures.

Unsupervised learning

Most remarkable results obtained with ML are based on supervised learning, that is, learning from examples where the expected output is given together with the input data. This implies prior labelling of data with the corresponding expected outputs and can be quite demanding for large-scale data. Amazon's Mechanical Turk (www.mturk.com) is an example of how corporations mobilise human resources for annotating data. But the vast majority of data exists without an expected output (i.e. without a desired annotation or a class name). Unsupervised learning algorithms should be developed in order to deal with this vast amount of unlabelled data; in some cases a small amount of supervision can be exploited so as to guide the unsupervised algorithm.

SEQUEL means "Sequential Learning". As such, **SEQUEL** focuses on the task of **learning in artificial systems** (either hardware, or software) that gather information along time. Such systems are named (learning) agents (or learning machines) in the following. These data may be used to estimate some parameters of a model, which in turn, may be used for selecting actions in order to perform some long-term optimisation task. The acquired data may result from an observation process of an agent in interaction with its environment (the data thus represent a perception). This is the case when the agent makes decisions (in order to attain a certain objective) that impact the environment, and thus the observation process itself.

Hence, in **SEQUEL**, the term sequential refers to two aspects:

- the sequential acquisition of data, from which a model is learned (supervised and non supervised learning),
- the sequential decision making task, based on the learned model (reinforcement learning).

Examples of sequential learning problems include supervised learning, unsupervised learning, reinforcement learning. In all these cases, they

mostly assume that the process can be considered stationary for at least a certain amount of time, and slowly evolving.

They wish to have any-time algorithms, that is, at any moment, a prediction may be required/an action may be selected making full use, and hopefully, the best use, of the experience already gathered by the learning agent.

The perception of the environment by the learning agent (using its sensors) is generally neither the best one to make a prediction, nor to take a decision (they deal with a Partially Observable Markov Decision Problem). So, the perception has to be mapped in some way to a better, and relevant, state (or input) space.

Finally, an important issue of prediction regards its evaluation: how wrong may they be when they perform a prediction? For real systems to be controlled, this issue can not be simply left unanswered.

To sum-up, in **SEQUEL**, the main issues regard:

- the learning of a model: we focus on models that map some input space RP to R ,
 - the observation to state mapping,
 - the choice of the action to perform (in the case of a sequential decision problem),
 - the performance guarantees,
 - the implementation of usable algorithms,
- all that being understood in a sequential framework.

Main applications are related to recommendation systems, and also prediction, supervised learning and clustering. Various methods are used, ranging from kernel methods to deep learning; non parametric methods are often favoured in **SEQUEL**'s works.

Human-in-the-learning-loop, explanations

The challenges are on the seamless cooperation of ML algorithms and users for improving the learning process; in order to do so, machine-learning systems must be able to show their progress in a form understandable to humans. Moreover, it should be possible for the human user to obtain explanations from the system on any result obtained. These explanations would be produced during the system's progression and could be linked to input data or to intermediate representations; they could also indicate confidence levels as appropriate.

LACODAM

Data science techniques have already shown a tremendous potential. However, they require combining a strong expertise both about the data at hand and the techniques to extract knowledge from data. Only few highly skilled experts have such expertise.

A promising way of research is thus to automate (part of) the data science process and make it accessible to a wider range of users.

Pioneering work in this direction has been undertaken with the “Data Science Machine” project of the MIT¹, which automates the delicate feature engineering process, in order to find the most relevant features in thousands of potential features, and exploits these to perform machine learning tasks. Another approach of interest is the “Automatic Statistician” from Cambridge University², which discovers complex regressions by exploring a search space of combinations of simple regression kernels, and produces natural language reports.

Both approaches are devices for supervised tasks: they have readily access to evaluation measures that allow them to automatically score their performance and improve it. The recent **LACODAM** team is interested in providing automation for more exploratory endeavours of data science, where the goal is to discover “interesting” knowledge from data, with prior information on the notion of interest. First, this requires using different methods than those mentioned before, such as pattern mining or clustering. Second, this requires keeping the user in the loop as only the user knows what he/she is interested in. An open problem is to find the best method for co-exploring the data by the system and the user. The solution will certainly require the collaboration of many different fields: not only data mining and machine learning, but also artificial intelligence, visualisation and human computer interfaces.

¹ James Max Kanter, Kalyan Veeramachaneni: *Deep feature synthesis: Towards automating data science endeavors*. DSAA 2015: 1-10

² David K. Duvenaud, James Robert Lloyd, Roger B. Grosse, Joshua B. Tenenbaum, Zoubin Ghahramani: *Structure Discovery in Nonparametric Regression through Compositional Kernel Search*. ICML (3) 2013: 1166-1174

Transfer learning

Transfer learning is useful when there is little data available for learning a task. It means using for a new task what has been learned from another task for which more data is available. It is a rather old idea (1993) but the results are modest because its implementation is difficult: it implies to conceptualise what the system has learned in the first place, but there is no general solution to this problem (what to conceptualise, how, how to re-use? ...).

Another approach for transfer learning is the procedure known as “shaping”: learning a simple task, then gradually complicate the task, up to the target task. There are examples of such processes in the literature, but no general theory.

Machine learning and decision support: feedback, causal models, representations

There exists an «irrational exuberance» about the big data era. This exuberance and the high level of (entrepreneurs, journalists, politicians) expectations might backfire and expose the scientific community to some big data winter, in the likeness of the AI winter.

TAO

I. Computer science tackling the underspecified: programming by feedback

Desirable functionalities will be increasingly complex and hard to specify as computational agents are expected to behave well in an open world: not only the targeted environments but also the performance indicators are partially unknown at design time. Additionally, the user doesn't want to read the manual. To face this challenge, the **TAO** team's approach will rely on two core enabling technologies, machine learning – modelling the user's intent based on his feedback, and Optimisation – exploiting this model to select an informative behaviour to display while minimising the expected time-to-solution.

Human-centred approaches raise the communication issue, displaying learned models in order to get an expert's feedback and making sense of it to fuel the learning and interaction process. Two main ways will be considered, based on visualisation on one hand, and on establishing uni-

fied representations (e.g. between manually defined ontologies and hierarchical concept distributions - related to direction III, ...), on the other.

II. Machine learning for decision making

A main usage of big data is to provide models in order to support decision-making. The point is that only causal models can support decisions: correlation-based models cannot. Example: children's marks are correlated with the presence of books in their home. But sending books in every household will **not** improve children's marks, everything else being equal.

Significant work has been done to formalise causality in terms of a machine learning problem: given a few hundreds of ML problems for which the causality status is known, the point is to learn the causality status. Extensions (in the multi-variate case and for temporal data) will be considered.

III. Changes of representation / deep learning

A lesson of the deep learning field is that the amount of data can make up for efficient representations of complex domains, e.g. in computer vision or natural language processing. New domains, e.g. in reinforcement learning, are under exploration; current research has barely scratched the surface.

Several research avenues will be considered:

- the optimisation algorithmic is at the crossroads of efficient optimisation (complying with information geometry and second order information about the optimisation landscape) and tractable algorithms (required to deal with embarrassingly large-dimensional models);
- the use of priors, e.g. about invariance operators in the problem domain, must be handled at various levels (initialisation; pre-processing modules in the line of spatial alignment in order to replace convolutional architectures) to tame the model complexity;
- the existence of related data distributions, in the line of transfer learning and domain adaptation, will be leveraged to build a unified representation (e.g. coming across the reality gap between in-situ and in-silico data in robotics).

4.3 Generic challenges in signal analysis: vision, speech

Signal analysis, in particular concerning vision and speech, benefited from recent progress in AI techniques. Deep learning systems won many challenges in vision and pattern recognition, the MobilEye vision system empowers Tesla cars' self-driving abilities, while voice-guided assistants such as Siri, Cortana, or Amazon Echo are put in use every day by millions of users.

The challenges in signal analysis for vision are: (i) scaling up; (ii) from still images to video; (iii) multi-modality; (iv) introduction of a priori knowledge.

THOTH

The quantity of digital images and videos available on-line continues to grow at a phenomenal speed:

home users put their movies on YouTube and their images on Flickr; journalists and scientists set up web pages to disseminate news and research results; and audio-visual archives from TV broadcasts are opening to the public. In 2018, it is expected that nearly 80% of the Internet traffic will be due to videos, and that it would take an individual over 5 million years to watch the amount of video that will cross global IP networks each month by then. Thus, there is a pressing and in fact increasing demand to annotate and index this visual content for home and professional users alike. The available text and audio metadata is typically not sufficient by itself for answering most queries, and visual data must come into play. On the other hand, it is not imaginable to learn the models of visual content required to answer these queries by manually and precisely annotating every relevant concept, object, scene, or action category in a representative sample of everyday conditions – if only because it may be difficult, or even impossible to decide a priori what are the relevant categories and the proper granularity level.

The main idea of the **THOTH** project-team proposal is to **develop a new framework for learning the structure and parameters of visual models** by actively exploring large digital image and video sources (off-line archives as well as growing on-line content, with hundreds of thousands of images), and exploiting the weak supervisory signal provided by the accompanying metadata. This huge volume of visual training data will allow us to learn complex non-linear models with a

large number of parameters, such as deep convolutional networks and higher-order graphical models.

The main goal of **THOTH** is to automatically explore large collections of data, select the relevant information, and learn the structure and parameters of visual models. **There are three main challenges:** (i) designing and learning structured models capable of representing complex visual information; (ii) on-line joint learning of visual models from textual annotation, sound, image and video; and (iii) large-scale learning and optimisation. Another important focus is (iv) data collection and evaluation.

Scaling up

Modern vision systems must be able to deal with high volume and high frequency data: for example, surveillance systems in public places, robots moving in unknown environments, web search engines in images have to process huge quantities of data. Vision systems must not only process these data at high speed, but need to reach high levels of precision in order to free operators from checking the results and post-processing. Even precision rates of 99.9% for image classification on mission-critical operations are not enough when processing millions of images, as the remaining 0.1% will need hours of human processing.

From still images to video

Object recognition in images still is a challenge in the general sense, even though very promising results are available; activity recognition in videos and more generally scene understanding is a growing domain which triggers new research, as presented below for the **STARS** project-team:

STARS

Many advanced studies have been done in signal analysis and in particular in scene understanding during these last few years. Scene understanding is the process, often real time, of perceiving, analysing and elaborating an interpretation of a 3D dynamic scene observed through a network of sensors (*e.g.* video cameras). This process consists mainly in matching signal information coming from sensors observing the scene with models which humans are using to understand the scene. Based on that, scene understanding is both adding and extracting semantics from the sensor data characterising a scene. This scene can contain a

number of physical objects of various types (*e.g.* people, vehicle) interacting with each other or with their environment (*e.g.* equipment), more or less structured. The phenomenon to be understood can last few instants (*e.g.* the fall of a person) or even several months (*e.g.* the depression of a person), can be limited to a laboratory slide observed through a microscope or go beyond the size of a city. Sensors include usually cameras (*e.g.* omni-directional, infrared, depth), but also may include microphones and other sensors (*e.g.* optical cells, contact sensors, physiological sensors, accelerometers, radars, smoke detectors, smart phones).



Figure 5: activity recognition in a hospital - © Inria / Photo H. Raguez

Scene understanding is influenced by cognitive vision and it requires at least the melding of three areas: **computer vision, cognition and software engineering**. Scene understanding can achieve five levels of generic computer vision functionality of detection, localisation, tracking, recognition and understanding. But scene understanding systems go beyond the detection of visual features such as corners, edges and moving regions to extract information related to the physical world which is meaningful for human operators. Its requirement is also to achieve more robust, resilient, adaptable computer vision functionalities by endowing them with a cognitive faculty: the ability to learn, adapt, weigh alternative solutions, and develop new strategies for analysis and interpretation.

Concerning scene understanding, the **STARS** team has developed original automated systems to understand human behaviours in a large variety of environments for different applications:

- in metro stations, in streets and aboard trains: fighting, abandoned luggage, graffiti, fraud, crowd behaviour,
- on airport aprons: aircraft arrival, aircraft refuelling, luggage loading/unloading, marshalling,
- in bank agencies: bank robbery, access control in buildings, using ATM machines,
- homecare applications for monitoring older persons' activities or problems: cooking, sleeping, preparing coffee, watching TV, preparing pill box, falling,
- smart home, office behaviour monitoring for ambient intelligence: reading, drinking,
- supermarket monitoring for business intelligence: stopping, queuing, picking up objects,
- biological applications: wasp monitoring, etc.

To build these systems, the **STARS** team has designed novel technologies for the recognition of human activities using in particular 2D or 3D video cameras. More specifically, they have combined **3 categories of algorithms to recognise human activities**:

- Recognition engines using hand-crafted ontologies based on rules modelling expert knowledge. These activity recognition engines are easily extendable and allow later integration of additional sensor information when available [König 2015].
- Supervised learning methods based on positive/negative samples representative of the targeted activities which have to be specified by users. These methods are usually based on Bag-of-Words computing a large variety of spatio-temporal descriptors [Bilinski 2015].
- Unsupervised (fully automated) learning methods based on clustering of frequent activity patterns on large datasets which can generate/discover new activity models [Negin 2015].

- A. König, C. Crispim, A. Covella, F. Bremond, A. Derreumaux, G. Bensadoum, R. David, F. Verhey, P. Aalten and P.H. Robert. Ecological Assessment of Autonomy in Instrumental Activities of Daily Living in Dementia Patients by the means of an Automatic Video Monitoring System, *Frontiers in Aging Neuroscience* - open access publication - <http://dx.doi.org/10.3389/fnagi.2015.00098>, 02 June 2015.

- P. Bilinski and F. Bremond. Video Covariance Matrix Logarithm for Human Action Recognition in Videos. *The International Joint Conference on Artificial Intelligence, IJCAI 2015*, Buenos Aires, Argentina from July 25th to July 31st, 2015.

- F. Negin, S. Cosar, M. Koperski, and F. Bremond. Generating Unsupervised Models for Online Long-Term Daily Living Activity Recognition. *The 3rd IAPR Asian Conference on Pattern Recognition, ACPR2015*, Kuala Lumpur, Malaysia, 4-6 November, 2015.

Multi-modality

Understanding visual data can be improved by different means: on the web, metadata provided with images and videos can be used to filter out several hypotheses, and to guide the system towards the recognition of specific objects, events, situations. Another option is to use multimodality, that is, signals coming from various channels *e.g.* infrared, laser, magnetic data etc. It is also desirable to use a combination of auditory and visual signals (images or video) if available. For example, a challenge was proposed by ANR (Agence Nationale de la Recherche) and DGA (Direction Générale de l'Armement) in 2010 (<http://www.defi-repere.fr/>), devoted to multimodal identification of persons in news programs, using the image in which people were visible, pop-up texts in which the names of persons appeared, the soundtrack in which the voice of the speakers was recognisable, and the content of the speech signal in which the names of the people could be heard.

The research agenda of the **PERCEPTION** group is the investigation and implementation of computational models for mapping images and sounds onto meaning and actions. **PERCEPTION** team members address these challenging topics with an interdisciplinary approach that spans the following disciplines: computer vision, auditory signal processing and scene analysis, machine learning, and robotics. In particular, they develop methods for the representation and recognition of visual and auditory objects and events, audio-visual fusion, recognition of human actions, gestures and speech, spatial hearing, and human-robot interaction.

Research topics:

- **computer vision:** spatio-temporal representation of 2D and 3D visual information, action and gesture recognition, 3D sensors, binocular vision, multiple-camera systems,
- **auditory scene analysis:** binocular hearing, sound source localisation and separation, speech communication, sound-event classification,
- **machine learning:** mixture models, linear and non-linear dimension reduction, manifold learning, graphical models,
- **robotics:** robot vision, robot hearing, human robot interaction, data fusion.

WILLOW addresses fundamental computer vision problems such as three-dimensional perception, computational photography, and image and video understanding. It investigates new models of image content

(what makes a good visual vocabulary?) and of the interpretation process (what is a good recognition architecture?). Recent achievements include theoretical work on the geometric foundations of computer vision, new advances in image restoration tasks such as deblurring, denoising, or upsampling, and weakly supervised methods for the temporal localisation of actions in videos. **WILLOW** members collaborate closely with the **SIERRA** and **THOTH** teams at Inria, and researchers at places such as Carnegie-Mellon University, UC Berkeley, or Facebook AI Research, in efforts that reflect the strong synergy between machine learning and computer vision, with new opportunities in domains ranging from archaeology to robotics emerging as well. Challenges for the future include the development of minimally supervised models of visual recognition in large-scale image and video datasets using available metadata in the form of text or speech for example.

Introduction of a priori knowledge

Another option for improving vision applications is to introduce a priori knowledge in the recognition engine. The example below, taken from the **ASCLEPIOS** project-team, consists in adding information about the anatomy and pathology of a patient for better analysis of biomedical images; in other domains, contextual information, information about a situation, about a task, localisation data, etc. can be used for disambiguating candidate interpretations. However, the question of how to provide this a priori knowledge is not solved in the general case: specific methods and specific knowledge representations must be established for dealing with a target application in vision understanding.

The **ASCLEPIOS** project-team has three main objectives:

- **analysis of biomedical images** with advanced geometrical, statistical, physical and functional models,
- **simulation of physiological systems** with computational models built from biomedical images and other signals,
- **application of previous tools to medicine and biology** to assist prevention, diagnosis and therapy.

Focus on medical image analysis: the quality of biomedical images tends to improve constantly (better spatial and temporal resolution, better signal to noise ratio). Not only the images are multidimensional (3 spatial coordinates and possibly one temporal dimension), but also

medical protocols tend to include multi-sequence (or multi-parametric) and multi-modal images for each single patient. Despite remarkable efforts and advances during the past twenty years, the central problems of segmentation and registration have not been solved in the general case. It is our objective in the short term to work on specific versions of these problems, taking into account as much a priori information as possible on the underlying anatomy and pathology at hand. It is also our objective to include more knowledge on the physics of image acquisition and observed tissues, as well as on the biological processes involved.




Figure 6: Biophysical modelling based on medical image analysis - © Inria / Photo H. Raguet

The other research themes of the **ASCLEPIOS** project are biological image analysis, computational anatomy, and computational physiology.

The challenges in signal analysis for speech and sound have a lot in common with the previous list: scaling up, multimodality, introduction of prior knowledge are relevant for audio applications too. The target applications are speaker identification, speech understanding, dialogue – including for robots – and automatic translation in real time. In the case of audio signals, it is also mandatory to develop or to have access to high volume data for machine learning. Online incremental learning might be needed for real time speech processing.

MULTISPEECH is a joint research team between the University of Lorraine, Inria, and CNRS. It is part of department D4 “Natural language and knowledge processing” of Loria. Besides vision, audition is an essential modality that carries invaluable information for human-machine interaction (speech understanding, speaker and emotion recognition) and environment perception (sound event informing about activities, danger, etc.). Deep learning technology has become the state of the art and fuelled the dissemination of close-microphone voice interfaces in the last 5 years. Several challenges remain, including the use of distant microphones corrupted by reverberation and acoustic noise, the exploitation of huge amounts of unlabeled speech and audio data, the robustness to speaker and sound event variabilities, the exploitation of contextual information, and robust integration with other sensing modalities.

The **MULTISPEECH** team is currently investigating these challenges by developing **tailored deep learning architectures** for diverse subproblems, combining them into larger systems, and seeking to estimate and propagate the confidence from one system to another. Besides voice interfaces, the applications of this technology include monitoring devices for assisted living or smart cities, among others.

4.4 Generic challenges in knowledge representation and semantic web

From Tim Berners-Lee’s initial definition, “the semantic web is an extension of the current web in which information is given well-defined meaning, better enabling computers and people to work in cooperation”. The semantic tower builds upon URIs and XML, through RDF schemas representing data triplets, up to ontologies allowing reasoning and logical processing. Inria teams involved in knowledge representation, reasoning and processing address the following challenges in different manners: (i) dealing with large volumes of information from heterogeneous distributed sources; (ii) building bridges between massive data stored in data bases using semantic technologies; (iii) developing semantically based applications on top of these technologies.

Dealing with large volumes of information from heterogeneous distributed sources

With the ubiquity of the Internet we are now faced with the opportunity and challenge of moving from local artificial intelligent systems to

massively distributed artificial intelligences and societies. Designing and running reliable and efficient systems combining linked data from distant sources through workflows of distributed services remains an open problem. The data quality and their processes traceability, the precision of their extraction and capture, the correctness of their alignment and integration, the availability and quality of shared models (ontologies, vocabularies) to represent, exchange and reason on them, etc. all these aspects need to be addressed on a large scale and continuously.

The focal point of **EXMO** is the heterogeneity of knowledge and its representation. The opportunity brought by the web to share knowledge at a wide scale has made this concern critical. **EXMO** has developed extensive work on ontology matching in the context of the semantic web. Ontology matching consists of finding related elements in heterogeneous ontologies and expressing this as alignments.

The **EXMO** team has designed alignment languages together with their formal semantics and software support (the Alignment API). It has developed ontology matchers based on similarity measures or context. Finally, the team has been at the origin of the yearly Ontology Alignment Evaluation Initiative campaigns, which evaluates ontology matchers. More recent work has concerned data interlinking, i.e., matching data instead of ontologies, with the notion of link keys.

Another concern is the evolution of knowledge. **EXMO** has contributed to the revision of networks of ontologies. However, in order to achieve seamless knowledge evolution, future work will adapt experimental cultural evolution techniques to knowledge representation.

A second aspect is underlined by the Web which not only provides a universal application framework for the Internet but also a hybrid space where humans and software agents can interact on large scales and form mixed communities. **Millions of users and artificial agents now interact daily** in online applications resulting in very complex systems to be studied and designed. We need models and algorithms that generate justifications and explanations and accept feedbacks to support interactions with very different users. We need to consider complex systems including the users as an intelligent component that will interact with other components (e.g. artificial intelligence in interfaces, natural language interaction), participate to the process (e.g. human computing, crowdsourcing, social machines) and may be augmented by the system

(intelligence amplification, cognitive augmentation, augmented intelligence, extended mind and distributed cognition).

WIMMICS

Web applications (e.g. Wikipedia) provide virtual spaces where persons and software interact in mixed communities exchanging and using formal knowledge (e.g. ontologies, knowledge bases) and informal content (e.g. texts, posts, tags).

The **WIMMICS** team studies models and methods to bridge formal semantics and social semantics on the web. It follows a multidisciplinary approach to analyse and model these spaces, their communities of users and their interactions. It also provides algorithms to compute these models from traces on the web including, knowledge extraction from text, semantic social network analysis, argumentation theory.

In order to formalise and reason on these models, the **WIMMICS** team then proposes languages and algorithms relying on and extending graph-based knowledge approaches for the semantic web and linked data on the web - e.g. graph models of the Resource Description Framework (RDF). Together, these contributions provide analysis tools and indicators, and support new functionalities and management tasks in epistemic communities.

These research results are integrated, evaluated and transferred through generic software (e.g. semantic web factory CORESE) and dedicated applications (e.g. the search engines **DiscoveryHub** and **QAKIS**).

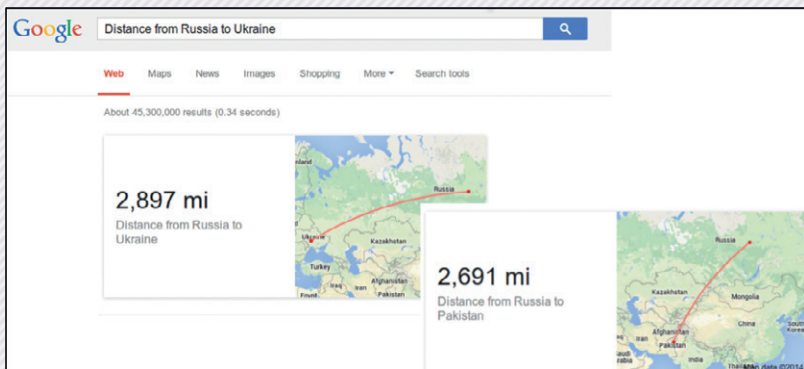


Figure 7: Without semantics, Russia appears closer to Pakistan than to Ukraine. From the paper «Why the Data Train Needs Semantic Rails» by Janowicz et al., AI Magazine, 2015.

Building bridges between massive data stored in data bases using semantic technologies

The semantic web addresses the massive integration of very different data sources (*e.g.* sensors of smart cities, biological knowledge extracted from scientific articles, event descriptions on social networks) and using very different vocabularies (*e.g.* relational schemas, lightweight thesauri, formal ontologies) in very different types of reasoning (*e.g.* decision making by logical derivation, enrichment by induction, analysis through mining, etc.). On the web, the initial graph of linked pages has been joined by a growing number of other graphs and is now mixed with sociograms capturing the social network structure, workflows specifying the decision paths to be followed, browsing logs capturing the trails of our navigation, service compositions specifying distributed processing, open data linking distant datasets, etc. Moreover, these graphs are not available in a single central repository but distributed over many different sources and some sub-graphs are public (*e.g.* [dbpedia http://dbpedia.org](http://dbpedia.org)) while others are private (*e.g.* corporate data). Some sub-graphs are small and local (*e.g.* a user's profile on a device), some are huge and hosted on clusters (*e.g.* Wikipedia), some are largely stable (*e.g.* thesaurus of Latin), some change several times per second (*e.g.* social network statuses), etc. **Each type of network of the web is not an isolated island**, they interact with each other: the social networks influence the message flows, their subjects and types, the semantic links between terms interact with the links between sites and vice-versa, etc. There is a huge challenge not only in finding means to represent and analyse each kind of graphs, but also means to combine them and combine their processing.

CEDAR

Making sense of "big data" requires interpreting it through the prism of knowledge about the data content, organisation, and meaning. Moreover, domain knowledge is often the language closest to the users, be they specialised domain experts or novice end users of a data-intensive application. Expressive and scalable tools for OBDA (Ontology-Based Data Access) are thus a key factor in the success of big data applications.

CEDAR works at the interface between knowledge representation formalisms (such as some description logics or classes of existential rules) and database engines. The team builds highly efficient OBDA tools with a particular focus on scaling up to very large databases; this can be seen as augmenting database engines with reasoning capabilities, and

deploying them in a cloud setting for scale. **CEDAR** also investigates novel ways of interacting with large, complex data and knowledge bases such as those referenced in the linked open data cloud (<http://lod-cloud.net>). Semantics is also investigated as a means to integrate and make sense of heterogeneous, complex content, in repositories of rich, heterogeneous Web data, in particular applied to journalistic fact checking.

GRAPHIK

Fully exploiting the various data available today requires taking into account their semantics, *i.e.*, knowledge. This widely recognised requirement gave rise to a novel paradigm, known as **Ontology-Based Data Access (OBDA)**, which takes advantage of domain ontologies when accessing data. In other words, databases become knowledge bases, in which an ontological layer is added on top of data. This added layer presents several interests: it allows to infer information not explicitly encoded in the data, to adapt the querying vocabulary to specific needs, and to access heterogeneous data sources in a uniform way. Over the last decade, a significant amount of research has led to a whole range of ontological languages offering different expressivity/complexity trade-offs. However, there is still much to do to obtain scalable systems that go beyond basic ontological knowledge. Moreover, work on OBDA has built on the assumption that data conform to relational structures, whereas increasing volumes of data-sources are not relational, like the NOSQL databases that sustain big data analytics. Whether the OBDA paradigm can be adapted to non-relational data management systems is still an open issue.

An important feature of knowledge-based techniques is their explanatory power, *i.e.* their potential ability to explain drawn conclusions. Being able to explain, justify or argue is a mandatory requirement in many AI applications in which the users need to understand the results of the system, in order to trust and control it. Moreover, it becomes a crucial concern with respect to ethical issues as soon as the automated decisions may impact human beings.

The main research domain of **GRAPHIK** is knowledge representation and reasoning. **The team is particularly interested in formalisms oriented towards data access**, like description logics and existential rules / Datalog+. The contributions of the team are on both theoretical and applied sides. The main issues currently tackled are ontology-based data access, with an emphasis on developing suitable formalisms with

practically efficient algorithms, as well as reasoning with imperfect information, with an emphasis on explaining and justifying conclusions drawn from inconsistent or conflicting data sources.

LINKS

The appearance of linked data on the web calls for novel database management technologies for linked data collections. The classical challenges from database research need to be now raised for linked data: how to define exact logical queries, how to manage dynamic updates, and how to automatise the search for appropriate queries. In contrast to mainstream linked open data, the **LINKS** project focuses on linked data collections in various formats, under the assumption that the data is correct in most dimensions. The challenges remain difficult due to incomplete data, uninformative or heterogeneous schemas, and the remaining data errors and ambiguities. They develop algorithms for evaluating and optimising logical queries on linked data collections, incremental algorithms that can monitor streams of linked data and manage dynamical updates of linked data collections, and symbolic learning algorithms that can infer appropriate queries for linked data collections from examples.

Research themes

The team develops algorithms for answering logical querying on heterogeneous linked data collections in hybrid formats, distributed programming languages for managing dynamic linked data collections and workflows based on queries and mappings, and symbolic machine learning algorithms that can link datasets by inferring appropriate queries and mappings.

Their main objectives are structured as follows:

- **Querying heterogeneous linked data.** They develop new kinds of schema mappings for semi-structured datasets in hybrid formats including graph databases, RDF collections, and relational databases. These induce recursive queries on linked data collections for which they investigate evaluation algorithms, static analysis problems, and concrete applications.
- **Managing dynamic linked data.** In order to manage dynamic linked data collections and workflows, they develop distributed data-centric programming languages with streams and parallelism, based on novel algorithms for incremental query answering; they study the propagation of updates of dynamic data through schema mappings, and investigate static analysis methods for linked data workflows.

- **Linking graphs.** Finally, they develop symbolic machine learning algorithms, for inferring queries and mappings between linked data collections in various graphs formats from annotated examples.

DYLISS

Experimental sciences undergo a data revolution due to the multiplication of sensors which allow for measuring the evolution of thousands of interdependent physical or biological components over time. When measurements are precise and various enough, they can be integrated in a machine learning framework to elucidate the role and function of components in the measured experimental system.

However, in many domains such as molecular biology, chemistry, and environment, the measurements are not sufficient to uniquely identify the role of each component. In this case, it becomes critical to confront the result of data analyses to knowledge corpora with are currently being structured through the linked open data initiative combined with semantic web technologies. For instance, there exist currently more than 1,500 knowledge-repositories about experimental science components stored in RDF format. A major challenge is to allow for exploration, sorting and mining of pools of candidates for the functions of components in experimental systems. To that end, the **DYLISS** team aims at developing logical programming techniques allowing to efficiently combine symbolic bi-clustering based of Formal Concept Analysis and semantic web technologies.

Other project-teams in this domain: **TYREX**, Grenoble, **ZENITH**, Montpellier.

4.5 Generic challenges in robotics and self-driving vehicles

Robotics combines many sciences and technologies, from the “lower level” mechanics, mechatronics, electronics, control, to the “upper level” of perception, cognition, collaboration and reasoning; in this section, even though artificial intelligence in robotics might imply to dig into the lower level functions for some processing features, we only deal with the upper levels, those which directly relate to the field of AI.

Recent progress made in robotics is impressive. Humanoid robots can walk, run, move in known and unknown environments, perform simple tasks like grasping objects or manipulating devices; bio-inspired robots

are able to mimic behaviours of a wealth of quite diverse living creatures (insects, birds, reptiles, rodents ...) and use these behaviours for efficiently solving complex problems; Boston Dynamics' Atlas (http://www.bostondynamics.com/robot_Atlas.html); biped robot can move efficiently in outdoor rough terrain and carry heavy objects, following the same company's four-legged robot BigDog.

.....
On the cognitive side, robots can play music, welcome visitors in shopping malls, converse with children.
.....

On the cognitive side, thanks to the progresses in speech processing, vision and scene understanding from many sensors, and thanks to the reasoning capacities implemented, robots can play music, welcome visitors in shopping malls, converse with children. With coordination features among a fleet of robots, they are able to play football together – but no robot team is yet able to beat a team of low-skilled humans. Autonomous vehicles are able to operate safely over long periods of time, and some countries and US states might allow them

to drive on public roads in the near future, even though a lot of open questions – including ethical ones – remain.

The challenges addressed by Inria teams developing research on robots and self-driving vehicles are: (i) situation understanding from multisensory input; (ii) reasoning under uncertainty, resilience; (iii) combining several approaches for decision making.

Situation understanding from multisensory input

For a robot to move in unknown areas, for a self-driving car in traffic, for a personal assistance robot such as Toi.Net (see section 1), it is essential to perceive the environment and to characterise the situation: this is done using input from multiple sensors (vision, laser, sound, internet, ..., road2car data in the case of vehicles). Situations can be simple symbols, ontologies, or more sophisticated representations of actors and objects present in an environment. A good characterisation of the situation can help the robot to make decisions - even in some case to infringe the law or a regulation for saving the car's passengers lives.

The project-team **RITS** is a multidisciplinary project at Inria, working on Robotics for Intelligent Transportation Systems. It seeks in particular to

combine the mathematical tools and techniques to design advanced intelligent robotics systems for autonomous and sustainable mobility.

Among the scientific topics covered:

- Multi-sensor signal processing (image processing, laser data, GPS data, IMU...) and data fusion,
- Advanced perception for environment modelling and understanding
- Vehicle control (acceleration, braking, steering),
- Wireless communications (vehicle-to-vehicle, vehicle-to-infrastructure),
- Large-scale traffic modelling and simulation,
- Control and optimisation of road transport systems,
- Development and deployment of automated vehicles (cyber cars, private vehicles,...).

The goal of these studies is to **improve road transportation** in terms of safety, efficiency, and comfort and also to minimise nuisances. The technical approach is based on driver's aids, going all the way to full driving automation. The project-team provides to the different teams cooperating with it, some important means such as a fleet of a dozen computer driven vehicles, various sensors and advanced computing facilities including a simulation tool.

An experimental system based on fully automated vehicles has been installed on the Inria grounds at Rocquencourt for demonstration purposes.



Figure 8: RITS self-driving cars - © Inria / Photo H. Raguez

Reasoning under uncertainty, resilience

Robots are active in the physical world and have to cope with defaults of many sorts: network shutdowns, defective sensors, electronic hazards etc. Some sensors provide incomplete information or have error margins generating uncertainty on the data. However, an autonomous mobile robot must perform its operation continuously without any human intervention, and for long periods of time. A challenge for robot architectures and software is to **deal with uncertain or missing information**, and to information only available at separate acquisition times. Anytime algorithms that provide an output on demand can be a solution in the case of fast decision-making needed even though the decision is not perfect.

Combining several approaches for decision-making

A variety of data and information can be available for a robot to make a decision: data from different sensors, information about the environment in the form of a situation assessment, memories of past decisions made, rules and regulations implemented in the robot's memory: there is a need to combine these facts and data and to conduct hybrid reasoning from numeric data, continuous or discrete, and from semantic representations. Moreover, as seen above, this reasoning must also take uncertainty into account: the research on decision-making for robots has to address this challenge. One possible solution is unsupervised machine learning and reinforcement learning of situations and semantic interpretations.

The goal of the **LARSEN** team is to move robots outside of the research laboratories and manufacturing industries: current robots are far from being the fully autonomous, reliable, and interactive robots that could co-exist with us in our society and run for days, weeks, or months. While there is undoubtedly progress to be made on the hardware side, robotics platforms are quickly maturing and they believe **the main challenges to achieve their goal is now on the software side**. They want their software to be able to run on low-cost mobile robots that are not equipped with high-performance (high-cost) sensors or actuators, so that their techniques can realistically be deployed and evaluated in real settings, such as in service and assistive robotic applications. They envision that these robots will be able to cooperate with each other but also with intelligent spaces or apartments, which can be seen as robots spread in the environment.

The **LARSEN** team members will organise their research around two research lines: livelong autonomy and natural interaction with robotics systems. Concerning the livelong autonomy axis, the team members have identified two challenges to address. This first one is to endow robots with stable situation awareness in open and dynamic environments. The second challenge is to allow a robot to recover from physical damages. One approach proposed by the team is to let the robot discover how to continue to perform a task despite such damages, by new trial and error algorithms that will allow robots to learn with a much smaller number of trials (typically, a dozen). Another approach the team will address is to deploy several robots or a swarm of robots.

.....
An original challenge that will be addressed by the team will be to develop new ways for a robot to physically interact with people.
.....

Concerning natural interaction with robotics systems, the team will consider the challenge of making robots behave with a better degree of acceptability as they can today. An original challenge that will be addressed by the team will be to develop **new ways for a robot to physically interact with people**. This requires to make the robot react and adapt online to the human feedback, exploiting the whole set of measurable verbal and non-verbal signals that humans naturally produce during a physical or social interaction. An important objective will be to make a robot able to sense humans around it and to be able to understand their activity.

Personal Assistance: the PAL Inria Project Lab

During the last fifty years, the many progresses of medicine as well as the improvement of the quality of life have resulted in a longer life expectancy in the industrial societies. The increase of the number of elderly people is a matter of public health because although elderly people can age in good health, old age also causes embrittlement in particular on the physical side, which can result in a loss of autonomy. That will force to re-think the current model regarding the care of elderly people. Capacity limits in specialised institutes, along with the preference of elderly people to stay at home as long as possible, explain a growing need for specific services at home.

Ambient intelligence technologies and robotics could participate to this societal challenge. Inria launched an Inria Project Lab termed **PAL**

for Personally Assisted Living from 2010 to 2015. Using the skills and objectives of the involved teams, **PAL** has defined and addressed **four research themes**: i) Assessing the degree of frailty of the elderly, ii) Mobility of people, iii) Rehabilitation, transfer and assistance in walking and iv) Social interaction.

CHROMA

The overall objective of **CHROMA** is to address fundamental and open issues that lie at the intersection of the emerging research fields called “Human Centred Robotics” and “Multi-Robot Systems (MRS)”.

Their goal is to design algorithms and develop models allowing mobile robots to navigate and operate in dynamic and human-populated environments. **CHROMA** is involved in all decision aspects pertaining to (multi)robot navigation tasks, including perception and motion-planning. Their approach for addressing this challenge is to bring together probabilistic methods, planning techniques and multi-agent decision models. This is done in cooperation with other disciplines such as psycho-sociology for the purpose of taking into account human models.

Two main research themes of robotic navigation are addressed: i) Perception and situation awareness in human-populated environments, by focusing on Bayesian perception and sensor fusion ii) Scaling-up single and multi-robot motion-planning, by combining uncertainty modelling, decentralised (swarm) models and multi-agent sequential decision making.

The **CHROMA** team is also concerned with applications and transfer of the scientific results. They work with industrial and start-up partners. Main application domains concern autonomous vehicle driving (with Renault and Toyota), aerial robots for surveillance tasks and services robotics.

Other project-teams in this domain: **LAGADIC**, Rennes, **HEPHAISTOS**, Sophia-Antipolis.

4.6 Generic challenges in neurosciences and cognition

FLOWERS objective is to model open-ended development in humans and its applications in robotics, human-computer interaction and educational technologies.

A major scientific challenge in cognitive science and artificial intelligence is to understand **how organisms can acquire open and cumulative repertoires of skills over an extended time span**. The process of sensorimotor, cognitive and social development in human children is organised along ordered stages, and results from the complex interaction between the brain/body with its physical and social environment.

To advance the fundamental understanding of mechanisms of development, the **FLOWERS** team has developed computational and robotic models in strong collaboration with developmental psychology and neuroscience. The team has developed models of guided exploration processes that allow learners to collect data efficiently in high-dimensional multi-task spaces. This includes mechanisms of active learning and information seeking (also called curiosity-driven learning), imitation learning, and maturational release of degrees of freedom.




Figure 9: The Poppy open source 3-D printed humanoid robot - © Inria / Photo H. Raguet

Beyond leading to new theories and new experimental paradigms to understand human development in cognitive science, the team has also

explored how such models can find applications in robotics, human-computer interaction and educational technologies. In robotics, the team has shown how artificial curiosity combined with imitation learning can provide essential building blocks allowing robots to acquire multiple tasks through natural interaction with naïve human users, for example in the context of assistive robotics. The team also showed that **models of curiosity-driven learning can be transposed in algorithms for intelligent tutoring systems**, allowing educational software to incrementally and dynamically adapt to the particularities of each human learner, and proposing personalised sequences of teaching activities. In human-computer interaction, the team has shown how incremental learning algorithms can be used to remove the calibration phase in certain brain-computer Interfaces.

Information Seeking, Curiosity and Attention: Computational and Neural Mechanisms, Gottlieb, J., Oudeyer, P-Y, Lopes, M., Baranes, A. (2013) Trends in Cognitive Science, 17(11), pp. 585-596.

How Evolution may work through Curiosity-driven Developmental Process Oudeyer, P-Y and Smith. L. (in press) Topics in Cognitive Science.

Self-Organisation in the Evolution of Speech, 2nd edition (2016) Oudeyer, P-Y., Oxford University Press

Multi-Armed Bandits for Intelligent Tutoring Systems, Manuel Lopes, Benjamin Clement, Didier Roy, Pierre-Yves Oudeyer. Journal of Educational Data Mining (JEDM), Vol 7, No 2, 2015.

Exploiting task constraints for self-calibrated brain-machine interface control using error-related potentials, I. Iturrate, J. Grizou, J. Omedes, P-Y. Oudeyer, M. Lopes and L. Montesano. PLOS One, 2015.

Inverse Reinforcement Learning in Relational Domains, Thibaut Munzer, Bilal Piot, Mathieu Geist, Olivier Pietquin and Manuel Lopes. International Joint Conference on Artificial Intelligence (IJCAI'15), Buenos Aires, Argentina, 2015.

ARAMIS

Understanding brain function and its alterations requires the integration of multiple levels of organisation, operating at different spatial (from the molecular to the whole-brain level) and temporal (from the millisecond to the entire lifespan) scales and representing different types of biological processes (anatomical, functional, molecular and cellular processes). Many aspects of these processes can now be quantified in living human subjects and patients thanks to the development of various technologies including neuroimaging, electrophysiology, genomics, transcriptomics... A major aim of computer science is to develop approaches that can automatically learn relevant patterns from multimodal data generated by these techniques.

The aim of the **ARAMIS** project-team is to **develop computational and statistical approaches to learn from such multimodal brain data.** A

first line of research aims at modeling brain structure and its statistical variability within populations. This involves the development of segmentation methods to extract brain structures from images and statistical shape models (Gori et al., 2013), implemented in the freely-available software package Deformetrica⁶.

A second axis aims to **model the functional interactions between distant brain areas** that underlie cognitive processes. This is based on approaches that can model the organisation of complex brain networks (De Vico Fallani et al., 2014). They are applied to the design of new devices, brain-computer interfaces and neurofeedback, for the rehabilitation of neurological patients. Finally, the team develops methods to learn models of neurological diseases and their progression. They allow learning topographical (Cuingnet et al., 2013) and spatio-temporal patterns (Schiratti et al., 2015) that are characteristic of a given disease. They shall result in new tools for diagnosis, prognosis and treatment personalisation in disorders such as Alzheimer's disease, Parkinson's disease, epilepsy and cerebrovascular disorders.

Cuingnet R, Glaunès JA, Chupin M, Benali H, and Colliot O (2013). Spatial and Anatomical Regularization of SVM: A General Framework for Neuroimaging Data. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35, 682–696.

Gori P, Colliot O, WorbeY, Marrakchi-Kacem L, Lecomte S, Poupon C, et al. (2013). Bayesian Atlas Estimation for the Variability Analysis of Shape Complexes. In *Proc. Medical Image Computing and Computer Assisted Intervention*, Nagoya, Japan, pp. 267–274.

Schiratti J-B, Allasonniere S, Colliot O, and Durrleman S (2015). Learning spatiotemporal trajectories from manifold-valued longitudinal data. In *Neural Information Processing Systems*, Montréal, Canada.

De Vico Fallani F, Richiardi J, Chavez M, and Achard S (2014). Graph analysis of functional brain networks: practical issues in translational neuroscience. *Philosophical Transactions of the Royal Society of London Series B, Biological Sciences* 369.

MNEMOSYNE

At the frontier between integrative and computational neuroscience, the **MNEMOSYNE** team proposes to model the brain as a system of active memories in synergy and in interaction with the internal and external world and to simulate it as a whole and in situation.

Major cognitive and behavioural functions (eg. attention, recognition, planning, decision) emerge from adaptive sensorimotor loops involving the external world, the body and the brain. **MNEMOSYNE** studies, models and implements such loops and their interactions toward a fully

⁶ <http://www.deformetrica.org>

autonomous behavior. With such a “systemic” approach, they mean that such complex systems can only be truly apprehended as a whole and in natural behavioural situations. To design the functioning and learning characteristics of such models at the level of the neuronal circuitry and to implement them in systems interacting in loops with the world, they combine principles, methods and tools from different fields of science.

They model the main cerebral structures and flows of information in the brain (as in integrative and cognitive neuroscience), stressing the links between brain, body and environment (embodied cognition). They use distributed computing formalisms allowing us to implement such models at different levels of description (as in computational neuroscience).

MNEMOSYNE deploys its models at large scale (high performance computing), incarnate them in bodies interacting with the environment (autonomous robotics) and simulate them interactively with respect to events encountered by a (virtual/real) robot.

4.7 Generic challenges in language Processing

The field of Natural Language Processing (NLP) goes back to the 1950s. Yet it is still of crucial importance today for the new information society. Its goal is to process natural language texts, either for analysing existing texts or generating new ones, for the purpose of achieving human-like language processing for a range of tasks or applications. These applications, regrouped under the term **“language engineering”**, include machine translation, question answering, information retrieval, information extraction, data mining, reading and writing aid, and many others. From a more research-oriented point of view, empirical linguistics and digital humanities can also be viewed as application domains of NLP.

NLP is a transdisciplinary domain; it requires an expertise in formal and descriptive linguistics (to develop linguistic models of human languages), in computer science and algorithmics (to design and develop efficient programs that can deal with such models) and in applied mathematics (to automatically acquire linguistic or general knowledge). Processing natural language texts is a difficult task, in particular because of the large amount of ambiguity in natural language, the specificities of individual languages and dialects and because many users do not

necessarily conform to grammatical and spelling conventions, when such conventions exist.

The first decades of NLP mostly focused on symbolic approaches, also contributing major notions to Computer Science, especially in formal grammar theory and parsing techniques. Linguistic knowledge was mostly encoded in the form of manually developed grammars and lexical databases. Over the last two decades statistical and machine-learning based approaches have greatly renewed the field, bringing annotated corpora to centre stage, and significantly improving the state of the art. However symbolic approaches retain specific advantages, and best results are often obtained when leveraging all types of resources including linguistic knowledge within hybrid systems coupling symbolic and statistical techniques.

The **ALPAGE** project-team (Inria Paris & Université Paris-Diderot) is specialised in NLP and its applications. Its efforts are spread across all levels of linguistic analysis (morphology, syntax, semantics, discourse), thus allowing for a global multi-scale analysis of textual data. **ALPAGE** has developed **many language resources, especially for French, that are used internationally**. It has also contributed to improving the state of the art in areas such as symbolic, statistical and hybrid parsing (for French, English and other languages, including for non-standard variants such as those found on the Web 2.0), semantic analysis and discourse modelling. An emerging approach the team has started to explore is to couple semantic analysis with models of the «world» the text refers to (for example the pairing of a video game's artificial world with the semantic representation of its gamers' online chats).

ALPAGE is also heavily involved in applied research, in two different ways. On the one hand, **ALPAGE** collaborates with the industrial world to develop NLP-enabled tools and technologies for tasks such as data mining (especially for noisy texts), information retrieval, information extraction and automatic text generation. On the other hand, **ALPAGE** is strongly involved in research on computational and empirical linguistic and on digital humanities. This includes work on **computational morphology** and on empirical syntax, including the study of the historical evolution of languages. Future directions could involve collaborations with researchers from Social Sciences (e.g. historians) as well as the development of rich language evolution models covering lexical, morphological and/or syntactic aspects.

Other project-team in this domain: **SEMAGRAMME**, Nancy.

4.8 Generic challenges in constraint programming for decision support

Constraint programming emerges in the eighties and develops at the intersection of artificial intelligence and operations research, of computer science and mathematics. Multidisciplinary by nature it keeps on using knowledge from various topics such as discrete mathematics, theoretical computer science (graph theory, combinatorics, algorithmics, complexity), functional analysis and optimisation, IT and software engineering.

.....
Computer aided decision and optimisation is becoming one of the cornerstones for providing assistance to all kinds of human activities.
.....

Constraint programming was identified in 1996 by the ACM as a strategic topic for Computer Science. The turn of the century has seen the development of optimisation technology in the industry (with notably Ilog, IBM, Dash and more recently Microsoft, Google and Dynadec) and the corresponding scientific field, at the border of constraint programming, mathematical programming, local search and numerical analysis. Optimisation technology is now assisting the public sector, companies and people to some extent for making decisions that use resources better and match specific requirements in an increasingly complex world.

Indeed, computer aided decision and optimisation is becoming one of the cornerstones for providing assistance to all kinds of human activities.

Today, with the pre-eminence of optimisation technology in most industrial sectors, we argue that quick and ad hoc solutions, often used today, cannot support the long-term development of optimisation technology and its broad diffusion. We also argue that **there should be a much more direct link between mathematical results and their systematic reuse in the main fields of optimisation technology.**

General Challenges

In spite of its importance, computer aided decision and optimisation suffer from a number of fundamental weaknesses that prevent from taking advantage of their full potential and hinder their progress and their capacity to deal with more and more complex situations. This can be mostly blamed on the diversity of involved disciplines, which are:

Design of effective solving techniques

On the one hand, computer science for providing languages, modeling tools and libraries. While focusing on providing flexible and powerful programming paradigms that can be easily deployed and maintained on modern architectures, it does not address the central question of how to come up in a systematic way with efficient methods for optimisation and decision problems. On the other hand, applied mathematics for the theory part. The focus is to come up with powerful abstractions that allow understanding the structure of a class of problems, independently of its practical and systematic uses in modern software components.

Spread out in distinct technological communities, each independently pushes its own solving paradigm like constraint programming, linear and integer programming, continuous optimisation, constraint-based local search (e.g., **COMET**). To some extent, most of these techniques exploit in different ways the same mathematical results, that are manually adapted to fit the main way to proceed of a given technology.

Thus, a first challenge encountered by constraint programming is the design of computer systems implementing in a transparent way effective solving techniques. Ideally, the user must be able to describe his problem in a high level modelling language without being concerned with the underlying solving mechanisms used. Such systems must also be independent both from any computer programming language and from any resolution engine.

In order to assist users, systems must also offer a **digital knowledge base in problem solving** that make available state of the art models and heuristics for a large set of well identified problems. Lastly, the user must have the ability to interpret the returned solutions, in particular within the context of over-constrained problems where it is necessary to partly relax some constraints, and that in the most realistic possible way.

Scaling up

A second challenge resides in **the speed of resolution especially in the context of large-scale data**. One has to adapt techniques such as generic consistency algorithms, graph algorithms, mathematical programming, meta-heuristics and to integrate them within the framework of constraint programming. This integration generates new questions such as the design of incremental algorithms, the automatic decomposition or the automatic reformulation of problems.

Complex industrial problems

Finally, a third challenge deals with **the use of constraint programming in the context of complex industrial problems**, especially when both discrete and continuous aspects are present. Complexity has multiple causes such as:

- the combination of temporal and spatial aspects, of continuous and discrete ones,
- the dynamic character of some phenomena inducing a modification of the constraints and data across time,
- the difficulty of expressing some physical constraints, *e.g.* load balancing and temporal stability,
- the necessary decomposition of large problems inducing significant performance losses of solutions.

TASC

Basic research of the **TASC** team is guided by the challenges raised before: to classify and enrich the models, to automate reformulation and resolution, to dissociate declarative and procedural knowledge, to develop modelling tools and to come up with solving tools that scale well.

On the one hand, classification aspects of this research are integrated within a knowledge base about combinatorial problem solving: the global constraint catalog. On the other hand, solving aspects are capitalised within the constraint solving system CHOCO. Lastly, within the framework of its activities of valorisation, teaching and of partnership research, the team uses constraint programming for solving various concrete problems.

The challenge is, on one side to increase the visibility of the constraints in the other disciplines of computer science, and on the other side to contribute to a broader diffusion of constraint programming in industry.

4.9 Generic challenges in music & smart environments

This section presents two project-teams which develop original work closely related to AI but cannot easily be allocated to the previous sections.

MUTANT

Music & artificial intelligence: paradigms for future challenges

Traditional signal processing and early probabilistic models for machine learning have led to significant improvement for speech and natural language processing especially in the 1990s but could never attain the same success in the field of music, whether for content information retrieval, music recognition systems, and computer assisted music practices. This is due to complex temporal patterns in music signals that cannot be easily abstracted, the heterogeneous nature of music information, and moreover the very adaptive nature of many musical concepts of interest (what really is music genre today?). Since the burst of the 21st century, scientific communities have challenged problems of music information retrieval at large and formed active communities that are comparable in size, industrial and societal impact to fields such as computer vision today. Most existing algorithms have taken opportunities in big data as granted and focused on off-line algorithms acting on large databases, despite the very fact that the majority of musical activity (production, listening, creating) is by nature online and through active learning and adaptation of human activities. **Today we have at our disposal algorithms in machine vision capable of understanding visual data as they arrive incrementally to the system.** Such a sea change is yet to occur for computer audition and computer musical understanding with applications in robotics and smart cities amidst music itself.

The **MUTANT** project-team has positioned itself on the lack in real-time music and audio understanding and interaction since its inception in 2013. **The artificial real-time machine listening software system Antescofo** developed by the team, capable of following musicians and decoding their musical parameters as they play, is composed of the simplest signal processing algorithms for data observation. Its robustness, proven in major world-class public performances, lies in the fact that parameter learning and adaptation occurs online and in real-time

(active learning) and its inherent information fusion capacities, mixing multiple fallible information sources with high uncertainty to achieve results. The robustness of the final results does not lie in employing the best or cleanest signal processing front-end but in intelligently handling multiple (fallible) sources of information as they arrive. **Dealing with unfiltered uncertainty in real-time is the major key in online intelligence** to achieve real-world applications. This is in analogy to human listening: we do not decode any musical information perfectly as we listen to, interact with or produce music. This is the approach adopted by the team to attain robust and industrial machine listening algorithms.



Figure 10: Using Antescofo for real time performance - © Inria / Photo H. Raguet

The question of music generation brings in a whole set of new challenges for artificial intelligence. How can we conceive machines capable of improvising music with or without humans (such as music style imitation) to the same level of excellence? To this end, one must explicitly consider the heterogeneous nature of musical time and information in artificial intelligence modeling. Decision making for such paradigms requires policy making both at the local level (e.g. considering semantic relations in music progression) and long-term levels (e.g. musical phrasing, and musical forms); and possibly by competitive and collaborative learning between various agents observing heterogeneous aspects of the same information source (pitch, timbre, rhythm, information flow, etc.). This problem alone requires important advances in machine learning and artificial intelligence in both online and offline forms.

The **MUTANT** team has elaborated **interactive learning agents with multiple-viewpoint learning** (employing reinforcement learning paradigms) capable of agnostically generating music in different styles without any prior knowledge or hand-engineering and also to improvise with live jazz musicians. But results are far from satisfactory and further research requires fluid architectures combining various approaches in knowledge discovery and decision making.

Furthermore, machine intelligence and real-time machine listening or musical understanding is of no use if it is not coupled to actions made through decisions. The machine listening in **Antescofo** allows our users to express concepts through potential perceptual experiences rather than mathematics. This capacity of expressing oneself with human concepts coupled to a high-level real-time language, allows them to build machines with complex temporal scenarios in an active environment. This is the conception of their team in realising **cyberphysical multi-media systems** that expert knowledge and systems become available to everyday users who can better express themselves with human concepts which under the hood is strongly coupled to critical real-time languages.

The objective of **PERVASIVE INTERACTION** is to develop the scientific and technological foundations for human environments that are capable of perceiving, acting, communicating, and interacting with people in order to provide services. The construction of such environments offers a rich set of problems related to interpretation of sensor information, learning, machine understanding, dynamic composition of components and man-machine interaction. **Their goal is to make progress on the theoretical foundations for perception and cognition**, as well as to develop new forms of man machine interaction, by using interactive environments as a source of example problems.

An environment is a connected volume of space. An environment is said to be “perceptive” when it is capable of recognising and describing things, people and activities within its volume. Simple forms of applications-specific perception may be constructed using a single sensor. However, to be general purpose and robust, perception must integrate information from multiple sensors and multiple modalities. **PERVASIVE INTERACTION** creates and develops machine perception techniques fusing computer vision, acoustic perception, range sensing and mechanical sensors to enable environments to perceive and understand humans and human activities.

An environment is said to be “active” when it is capable of changing its internal state. Common forms of state change include regulating ambient temperature, acoustic level and illumination. More innovative forms include context-aware presentation of information and communication, as well as services for cleaning, materials organisation and logistics. The use of multiple display surfaces coupled with location awareness offers the possibility of **automatically adapting information display** to fit the current activity of groups. The use of activity recognition and acoustic topic spotting offers the possibility to record a log of human to human interaction, as well as to provide relevant information without disruption. The use of steerable video projectors (with integrated visual sensing) offers the possibilities of using any surface for presentation, interaction and communication.

An environment may be considered as “interactive” when it is capable of interacting with humans using tightly coupled perception and action. Simple forms of interaction may be based on observing the manipulation of physical objects, or on visual sensing of fingers, hands or arms. Richer forms of interaction require perception and understanding of human activity and context. **PERVASIVE INTERACTION** has developed a novel theory for situation modeling for machine understanding of human activity, based on techniques used in cognitive psychology. **PERVASIVE INTERACTION** explores multiple forms of interaction, including projected interaction widgets, observation of manipulation of objects, fusion of acoustic and visual information, and systems that model interaction context in order to predict appropriate action and services by the environment.


For the design and integration of systems for perception of humans and their actions, **PERVASIVE INTERACTION** has developed:

- a framework for context aware services using situation models,
- Robust, view invariant techniques for computer vision using local appearance,
- a distributed autonomic software architecture for multimodal perceptual systems.

The experiments in **PERVASIVE INTERACTION** are oriented towards **developing interactive services for smart environments**. Application domains include activity monitoring services for healthy living, smart objects and services for the home, and new forms of man-machine interaction based on perception.



Inria references: numbers



Over the 2005-2015 period, Inria researchers published **more than 400 AI journal articles and more than 1000 AI conference papers**. The project teams' references, including other supports such as book chapters and Inria internal reports, can be found on <https://hal.inria.fr/> and on the research report website: <http://raweb.inria.fr/rapportsactivite/RA2015/index.html>. The following tables show lower bounds for numbers of publications and conference proceedings, since some journal articles and conference papers were not properly indexed in the Inria publications database before 2012.

JOURNAL PAPERS

AI Communications	3
Artificial Intelligence	7
Computer Speech and Language	13
Computer Vision and Image Understanding	28
Constraints	17
Expert Systems with Applications	17
IEEE Robotics and Automation Magazine	8
IEEE Transactions on Evolutionary Computation	4
IEEE Transactions on Image Processing	45
IEEE Transactions on Intelligent Transportation Systems	14
IEEE Transactions on Knowledge and Data Engineering	15
IEEE Transactions on Pattern Analysis and Machine Intelligence	3
IEEE Transactions on Robotics	48
International Journal of computer vision	68
International Journal of Robotics Research	29
Journal of Artificial Intelligence Research	7
Journal of Data Semantics	6
Journal of Intelligent and Robotic Systems	3
Journal of Intelligent Systems	2
Journal of Language Technology and Computational Linguistics	1
Journal of machine learning Research	33
Journal of Web Semantics	7
Machine Vision and Applications	9
Pattern Recognition Letters	18
Revue d'Intelligence Artificielle	5
TAL - Traitement Automatique des Langues	20
Trends in Cognitive Sciences	1
Total	431

CONFERENCE PAPERS

Annual Conference of the International Speech Communication Association	43
Annual Conference on Neural Information Processing Systems	76
Conférence Francophone sur l'Apprentissage Automatique	11
Congrès Francophone de Reconnaissance des Formes et Intelligence Artificielle	43
European Conference on Artificial Intelligence	21
European Conference on Machine Learning	5
European Conference on Principles and Practice of Knowledge Discovery in Databases	2
European Semantic Web Conference	5
European Symposium on Artificial Neural Networks	8
Genetic and Evolutionary Computation Conference	71
IEEE Conference on Computer Vision and Pattern Recognition	60
IEEE Conference on Intelligent Transportation Systems	4
IEEE International Conference on Acoustics, Speech and Signal Processing	91
IEEE International Conference on Computer Vision	38
IEEE International Conference on Robotics and Automation	109
IEEE International Conference on Tools with Artificial Intelligence	17
IEEE RSJ International Conference on Intelligent Robots and Systems	110
International Conference on Artificial Evolution	7
International Conference on Autonomous Agents and Multiagent Systems	12
International Conference on Machine Learning	43
International Conference on Practical Applications of Agents and Multi-Agent Systems	5
International Conference on Principles and Practice of Constraint Programming	20
International Conference Parallel Problem Solving from Nature	22
International Joint Conference on Artificial Intelligence	24
International Semantic Web Conference	14
International Workshop on Principles of Diagnosis	5
Journées d'Extraction et Gestion des Connaissances	58
Journées Francophones sur les Systèmes Multi-Agents	7
National Conference on Artificial Intelligence	9
Traitement Automatique du Langage Naturel	18
Workshop on machine learning for System Identification	1
Total	959



Other references for further reading



This section contains other references identified to be among the most relevant for further reading, grouped in categories.

Generic AI

Cap Digital. **Intelligence Artificielle : technologies, marchés et défis**. Document # 15-167, www.capdigital.com/publications, 2015

Ernest Davis and Gary Marcus. **Commonsense Reasoning and Commonsense Knowledge in Artificial Intelligence**. Communications OfThe ACM Vol. 58 No. 9, 2015

Jonathan Grudin. **AI and HCI: Two Fields Divided by a Common Focus**. AI magazine, 30(4), 48-57, 2008

Kevin Kelly. **The Three Breakthroughs That Have Finally Unleashed AI On The World**. <http://www.wired.com/2014/10/future-of-artificial-intelligence/>, 2014

Pierre Marquis, Odile Papine, Henri Prade (eds). **Panorama de l'Intelligence Artificielle. ses bases méthodologiques, ses développements**. 3 vols. Cepaduwès, 2014

Stuart Russell and Peter Norvig. **Artificial Intelligence: A Modern Approach**. <http://aima.cs.berkeley.edu>

Alan Turing. **Intelligent Machinery, a Heretical Theory**. Philosophia Mathematica 4 (3): 256-260, 1996 - Original article from 1951

Terry Winograd. **Shifting viewpoints: Artificial intelligence and human-computer interaction**. Artificial Intelligence 170(18):1256-1258, 2006

Debates about AI

Ronald C. Arkin. **The Case for Ethical Autonomy in Unmanned Systems**. Journal of Military Ethics, 9(4), 12/2010

Samuel Butler. **Erewhon**. Free eBooks at Planet eBook.com, 1872

Dominique Cardon. **A quoi rêvent les algorithmes**. Seuil, 2015

Thomas G. Dietterich and Eric J. Horvitz. **Rise of Concerns about AI: Reflections and Directions**. Communications of the ACM | Vol. 58 No. 1, October 2015

Moshe Vardi. **On Lethal Autonomous Weapons**. Communications of the ACM, vol. 58 no. 12, December 2015

Machine learning

Martin Abadi et al. **Large-Scale machine learning on Heterogeneous Distributed Systems**. Software available from tensorflow.org, 2015

Christopher Bishop. ***Pattern Recognition and Machine Learning***. Springer, 2006

Léon Bottou. ***From machine learning to machine reasoning: an essay***, Machine Learning, 94:133-149, January 2014

Yann Le Cun. ***The Unreasonable Effectiveness of Deep Learning***. Facebook AI Research & Center for Data Science, NYU. <http://yann.lecun.com>, 2015

Sumit Gulwani, William R. Harris, and Rishabh Singh. ***Spreadsheet Data Manipulation Using Examples***. Communications of the ACM, Vol. 55 no. 8, 2012

Michael I. Jordan and Tom M. Mitchell. ***Machine learning: trends, perspectives, and prospects***. Science, Vol. 349 Issue 6245, 2015

James Max Kanter and Kalyan Veeramachaneni. ***Deep Feature Synthesis: Towards Automating Data Science Endeavors***. International IEEE/ACM Data Science and Advance Analytics Conference, 2015

Tom M. Mitchell. ***The Discipline of Machine Learning***. CMU-ML-06-108 School of Computer Science Carnegie Mellon University, 2006

Volodymyr Mnih et al. ***Human-level control through deep reinforcement learning***. Nature 518, 529–533, 2015

Michèle Sebag. ***A tour of Machine Learning: an AI perspective***. AI Communications, IOS Press, 27 (1), pp.11-23, 2014

Vision

Peter Auer et al. ***A Research Roadmap of Cognitive Vision***.

ECVision : the European Research network for Cognitive Computer Vision Systems. <https://www.researchgate.net>, 2005

Nicholas Ayache. ***Des images médicales au patient numérique***, Leçons inaugurales du Collège de France. Collège de France / Fayard, March 2015

Slawomir Bak. ***Human Reidentification Through a Video Camera Network***. Computer Vision and Pattern Recognition. PhD, Université Nice Sophia Antipolis, 2012

Svetlana Lazebnik, Cordelia Schmid, Jean Ponce. ***Spatial pyramid matching***. Sven J. Dickinson and Alès Leonardis and Bernt Schiele and Michael J. Tarr. Object Categorization: Computer and Human Vision Perspectives, Cambridge University Press, pp. 401-415, 2009

Sancho McCann, David G. Lowe. ***Efficient Detection for Spatially Local Coding***. Lecture Notes in Computer Science Volume 9008 pp. 615-629. 2015

Farhood Negin, Serhan Cosar, Michal Koperski, François Bremond. ***Generating Unsupervised Models for Online Long-Term Daily Living Activity Recognition***. Asian conference on pattern recognition (ACPR 2015), 2015

Oriol Vinyals, Alexander Toshev, Samy Bengio & Dumitru Erhan. **Show and Tell: A Neural Image Caption Generator**. <https://arxiv.org/pdf/1502.03044>, 2015

Knowledge representation, semantic web, data

Nathalie Aussenac-Gilles, Fabien Gandon. **From the knowledge acquisition bottleneck to the knowledge acquisition overflow: A brief French history of knowledge acquisition**. International Journal of Human-Computer Studies, Elsevier, vol. 71 (n 2), pp. 157-165, 2013

Tim Berners-Lee, James Hendler and Ora Lassila. **The Semantic Web**. Scientific American, May 2001

Meghyn Bienvenu, Balder Ten Cate, Carsten Lutz, Frank Wolter. **Ontology-Based Data Access: A Study through Disjunctive Datalog, CSP, and MMSNP**. ACM Transactions on Database Systems, Vol. 39, No. 4, Article 33, 2014

Fabien Gandon. **The three 'W' of the World Wide Web call for the three 'M' of a Massively Multidisciplinary Methodology**. Valérie Monfort; Karl-Heinz Krempels. 10th International Conference, WEBIST 2014, Barcelona, Spain. Springer International Publishing, 226, Web Information Systems and Technologies, 2014

Janowicz, K.; Hitzler, P.; Hendler, J.; and van Harmelen, F. **Why the Data Train Needs Semantic Rails**. AI Magazine, 36(5-14), 2015

Antonella Poggi et al. **Linking Data to Ontologies**. Journal on data semantics X pp. 133-173. Springer-Verlag Berlin, Heidelberg, 2008

Robotics and self-driving cars

J. Christian Gerdes, Sarah M. Thornton, **Implementable Ethics for Autonomous Vehicles**. Autonomes Fahren: Technische, rechtliche und gesellschaftliche Aspekte. Springer, Berlin, 2015

Philippe Morignot, Joshue Perez Rastelli and Fawzi Nashashibi, **Arbitration for Balancing Control between the Driver and ADAS Systems in an Automated Vehicle: Survey and Approach**. 2014 IEEE Intelligent Vehicles Symposium (IV), 2014

Developmental Robotics

Antoine Cully, Jeff Clune, Danesh Tarapore & Jean-Baptiste Mouret. **Robots that can adapt like animals**. Nature Vol 521 503-507, 2015

Pierre-Yves Oudeyer. **Developmental Robotics**. Encyclopaedia of the Sciences of Learning, N.M. Seel ed., Springer References Series, Springer, 2012

AI and cognition

Jacqueline Gottlieb, Pierre-Yves Oudeyer, Manuel Lopes and Adrien Baranes. ***Information-seeking, curiosity, and attention: computational and neural Mechanisms.*** Trends in Cognitive Science (2013) 1-9, 2013

Douglas Hofstadter & Emmanuel Sander. ***L'analogie, cœur de la pensée.*** Ed. Odile Jacob, 2013

Steels, Luc. ***Self-organisation and selection in cultural language evolution.*** In Luc Steels (Ed.), Experiments in Cultural Language Evolution, 1 – 37. Amsterdam: John Benjamins, 2012

Natural language, speech, audio, music

Kenneth Church. ***A Pendulum Swung Too Far.*** Linguistic Issues in Language Technology – LiLT. Vol. 2, Issue 4, 2007

Arnaud Dessein, Arshia Cont. ***An information-geometric approach to real-time audio segmentation.*** IEEE Signal Processing Letters, Institute of Electrical and Electronics Engineers, 20 (4), pp.331-334, 2013

G. Hinton, L. Deng, D. Yu, G.E. Dahl, A. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T.N. Sainath, B. Kingsbury. ***Deep neural networks for acoustic modeling in speech recognition: the shared views of four research groups.*** IEEE Signal Processing Magazine, 29(6):82-97, 2012

Acknowledgment

Researchers in Inria project-teams and centres who contributed to this document (were interviewed, provided text, or both):

Abiteboul Serge, DAHU project-team, Cachan

Ayache Nicholas, head of ASCLEPIOS project-team, Sophia-Antipolis

Bach Francis, head of SIERRA project-team, Paris

Beldiceanu Nicolas, head of TASC project-team, Nantes

Boujemaa Nozha, advisor on big data for the Inria President

Braunschweig Bertrand, director of Saclay-Île-de-France research centre

Brémond François, head of STARS project-team, Sophia-Antipolis

Charpillet François, head of LARSEN project-team, Nancy

Colliot Olivier, head of ARAMIS project-team, Paris

Cont Arshia, head of MUTANT project-team, Paris

Cordier Marie-Odile, LACODAM project-team, Rennes

Crowley James, head of PERVASIVE INTERACTION project-team, Grenoble

De La Clergerie Eric, ALPAGE project-team, Paris

De Vico Fallani Fabrizio, ARAMIS project-team, Paris

Euzenat Jérôme, head of EXMO project-team, Grenoble

Gandon Fabien, head of WIMMICS project-team, Sophia-Antipolis

Giaavito Jean-Louis, MUTANT project-team, Paris

Gilleron Rémi, MAGNET project-team, Lille

Giraudon Gérard, director of Sophia-Antipolis-Méditerranée research centre

Gravier Guillaume, head of LINKMEDIA project-team, Rennes

Gros Patrick, director of Grenoble-Rhône-Alpes research centre

Guittou Pascal, member of the technology watch and prospective group

Horaud Radu, head of PERCEPTION project-team, Grenoble

Manolescu Ioana, head of CEDAR project-team, Saclay

Moisan Sabine, STARS project-team, Sophia-Antipolis

Mugnier Marie-Laure, head of GRAPHIK project-team, Montpellier

Nashashibi Fawzi, head of RITS project-team, Paris

Niehren Joachim, head of LINKS project-team, Lille

Oudeyer Pierre-Yves, head of FLOWERS project-team, Bordeaux

Pietquin Olivier, SEQUEL project-team, Lille

Ponce Jean, head of WILLOW project-team, Paris

Preux Philippe, head of SEQUEL project-team, Lille

Roussel Nicolas, head of MJOLNIR project-team, Lille

Sagot Benoit, head of ALPAGE project-team, Paris

Schmid Cordelia, head of THOTH project-team, Grenoble

Schoenauer Marc, co-head of TAO project-team, Saclay

Sebag Michèle, co-head of TAO project-team, Saclay

Seddah Djamé, ALPAGE project-team, Paris

Siegel Anne, head of DYLISS project-team, Rennes

Sturm Peter, deputy scientific director


Termier Alexandre, head of LACODAM project-team, Rennes

Thonnat Monique, director of Bordeaux-Sud-Ouest research centre

Tommasi Marc, head of MAGNET project-team, Lille

Toussaint Yannick, ORPAILLEUR project-team, Nancy

Vincent Emmanuel, MULTISPEECH project-team, Nancy



Domaine de Voluceau, Rocquencourt BP 105
78153 Le Chesnay Cedex, France
Tél. : +33 (0)1 39 63 55 11
www.inria.fr