

Reasoning and decision: human or machine?



>>> Andreas HERZIG, CNRS senior scientist at the Institut de recherche en informatique de Toulouse (IRIT, joint UPS/CNRS/UT1/UT2/INPT lab)

How can a typically human activity such as reasoning be performed by machines? To answer this question we have to understand how our minds function. This is determined not only by physical reality, but also by our beliefs, goals and preferences.

We don't know everything, our beliefs may be erroneous and we live in a world that changes all the time. Despite all these difficulties we are able to reason and to make decisions quite successfully. Our abilities to learn and to revise our beliefs are crucial for this.

We not only reason about nature and physical reality, but also about the other human agents in our environment and with social reality as created by man: society and its institutions, such as norms and conventions, things that are allowed and things that are forbidden, contracts, etc. Indeed, language and communication may be considered to be the first of these institutions: in dialogue we exploit the fact that others know the conventions of communication, allowing them to interpret what we say. An example of a convention is to avoid contradictions and to stay focused on the subject.

Anthropomorphism

What we have said up to now applies to human agents, and one might wonder what all this has to do with computer science. However, it appears that the central concepts in our discourse apply to artificial agents too. Such an anthropomorphic point of view was adopted in Artificial Intelligence right from its beginnings in the 60s and in multi-agent systems in the 80s.

Of course, how agents reason and decide is not only investigated in computer science, but originally in the humanities and in social sciences, viz. in psychology, philosophy, linguistics, cognitive science, and economy. It is therefore a multidisciplinary object of study, and the researchers of IRIT very often collaborate with linguists and psychologists of the Toulouse labs "Cognition, Langues, Langage, Ergonomie" (CLLE, universit  Toulouse 2) and the economists of the "Toulouse School of Economics" (TSE, universit  Toulouse 1).

Theories of possibility and logics of modality

The overall aim is to implement artificial agents in computers. How can one relate concept analysis and computer programs? For several years now, the researchers working on the "Reasoning and Decision" theme at IRIT have adopted a classical research avenue, viz. formal approaches, allowing for rigorous and verifiable proofs. Their main tools are possibility theory and modal logics, and the procedure is to start by formal modeling (including concept analysis). This is then followed by an investigation of the mathematical properties and the development of automatic or semi-automatic reasoning procedures.

Contact: herzig@irit.fr

IRIT: Institut de recherche en informatique de Toulouse/
Toulouse Institute of Computer Science Research



>>> Rodin's The Thinker at Saint Di .
  Christian Amet/Creative Commons



>>> Nathalie Aussenac-Gilles, CNRS senior scientist, and Laure Vieu, CNRS scientist, researchers at the IRIT (UPS/CNRS/INP/UT1/UT2)

**“ You feel confused?
You don’t have a boyfriend?
No social life? No dreams?
You want to be more like
sitcom actors? Then go for
this TV show and watch it
over and over again! ”**

>>> Caption:

In opinion mining of short texts, such as forum comments, the main goal is to find out if the author has a favorable opinion of the discussed topic. The example illustrates the limitations of algorithms based only on word spotting. The advice given in the article indicates a positive opinion of the show. However, taking into account the whole context and with sufficient background knowledge, it can be seen that “ no social life ” characterizes an audience for which the author has no respect, so that recommending the show for that particular audience actually leads to a negative opinion about the show itself.

Language and representing knowledge

Reading a text, analyzing and summarizing it: these processes are not unique to humans anymore according to new work by researchers at the IRIT laboratory who are studying the automatic treatment of language

Among the first challenges faced by the Artificial Intelligence (AI) program was to reproduce the very human ability of producing and understanding linguistic messages. This is all the more difficult as such activity is thought of being a mark of human intelligence. Natural language processing makes linguistic analyses accessible to computer systems, from the handling of synonyms (lexical level), for instance, to the identification of verb complements (syntactic level), finding what a pronoun refers to (semantic level) or the relation between two sentences (discourse level). This, in turn, is used for information extraction, machine translation, or automated summary production. Collaboration between linguists and computer scientists has proved very useful in this respect, and researchers in the NLP at IRIT have a long standing partnership with linguistic groups at Université Toulouse II (CLLE-ERSS and Jacques Lordat Lab).

Information explosion

The path followed since the inception of the AI program has been tremendously transformed by the explosion of the amount of data available in natural language (whether written or oral) either published on the web or produced by companies or users. Managing this abundance opens up new avenues for testing hypotheses on languages, and the development of efficient tools for analysis. It also favours sharing of information, knowledge discovery and corroboration.

A major objective in the domain is to give some structure to the areas of knowledge that need to be formalized, by finding connections between relevant entities, their properties and the concepts they involve. This usually goes under the name of ontology. In the medical domain, many ontologies have been designed, often starting from thesaurii, such as MESH or UMLS, or from automated analyses of medical texts to jump-start the formalization process. These ontologies provide richer definitions of medical concepts, of relations between pathologies, treatments and anatomic aspects, etc. They are used to characterize the information available in patient files, good

practice guidelines or scientific papers, for verification and corroboration, and more globally to fuel decision processes with recent research results.

Opinion mining

The main activity of our team contributes to the formalization of different elements that take part in the semantics of a text. Our work is concerned with the globality of a text or a discourse, how lexical and syntactic elements combine but also how relations within and between sentences, paragraphs and other textual elements contribute to the construction of meaning. Research on ontologies includes these results to organize knowledge and build systems that can help to identify and extract information from texts. This also implies the development of tools to help model and maintain ontologies.

The end result combines knowledge and its linguistic expression. It is put to the test by a range of applications, mainly opinion mining (to understand such phenomena as corporate or private e-reputations, or product reviews), dialog analysis, text mining and automated ontology building. Our research also involves industrial partnerships (eg ACTIA Automotive, Synapse Development), institutional funding via ANR grants, for instance, for computer aided diagnostic of electronic failure, or geographical database integration.

Contact : aussenac@irit.fr, muller@irit.fr,
vieu@irit.fr

Probability Theory alone cannot handle all facets of uncertainty



>>> Didier DUBOIS and Henri PRADE, CNRS senior scientists at the Institut de Recherche en Informatique de Toulouse (IRIT, UPS/CNRS/UT1/ UT2/INPT)

People constantly deal with pieces of information that are incomplete, uncertain, inaccurate, and sometimes inconsistent. To face this issue, scientists naturally resort to probability theory. But this approach does not take into account where uncertainty comes from.

Probability theory as a tool for representing uncertainty has long existed, but it often does away with the fact that there may be several reasons for being uncertain. The main reason evoked is the variability of natural phenomena and of data coming from repeated measurements. Hence the frequentist's view of probabilities that is often taken for granted.

Thought lotteries

Another very common reason for uncertainty is the plain lack of information, which alone may prevent agents from knowing whether statements of interest are true or false. This kind of "poor" information often takes the form of an incomplete set of propositions in classical logic, or appears as an interval of possible numerical values of an ill-known quantity, or even as a set of attributes that does not allow for the precise description of an object. To understand such situations, we need another probabilistic concept, so-called subjectivism, where a probability does not reflect an ideal frequency, but represents, in a betting framework, the price of a thought lottery that yields an euro if the concerned event occurs. This concept claims that it makes sense to represent any incomplete information situation by means of a unique probability distribution. But this view can be challenged: such a representation is not scale-invariant, and does not depend on the origin of uncertainty. It is more natural to represent incomplete information by means of a mere set of possible values, often called the "epistemic state" in Artificial Intelligence.

Uncertainty theories

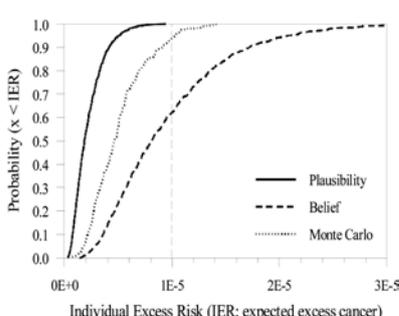
It makes sense to be able to tell variability and incompleteness apart in the scope of information processing. Several new theories try to address this issue, such as possibility theory, evidence theory and imprecise probability theory. In such settings, instead of representing any information state by means of a unique probability distribution, sets thereof, or even random sets, are used. The confidence attached to an event occurring is then quantified by a probability interval, measuring certainty by its lower bound and plausibility by its upper bound. Sometimes such quantification

is difficult to justify. All that can be said in some cases is that some events are more likely than others. One is then led to use qualitative representations of uncertainty, that turn out to be instrumental in dealing with exception-tolerant reasoning.

Possibilistic Logic

When information is incomplete, the human mind resorts to reasoning patterns that enable useful conclusions to be tentatively drawn despite incompleteness. Such conclusions can be questioned upon the arrival of new pieces of information. In this case, reasoning becomes non-monotonic and presupposes the truth of anything that is considered normal in the current informational context. This form of reasoning is not amenable to classical logic. It requires logic with embedded priorities, such as possibilistic logic. Such logic is also instrumental in the problem of merging partially inconsistent pieces of information.

For more than 20 years, our team has been working on the construction of formal frameworks that are wide enough to enrich the expressive power of traditional approaches to uncertainty, both probabilistic and logic-based. For the last 10 years or so we have cooperated with public research laboratories in the area of environmental risk analysis, such as BRGM or IRSN. In risk analysis, the usual methodology is first to build a mathematical model of some potentially hazardous phenomenon. Then in a second step, we must check on the basis of data collected on-site in an area of concern, whether the probability of a risky event is above some acceptable, say, pollution, threshold. But the collected objective data are often incomplete and part of the information must come from expert opinion. We have developed risk analysis methods based on imprecise probabilities encoded as possibility distributions, so as to separately handle uncertainty due to known variability and uncertainty due to lack of data.



>>> Results for risk calculation application

Contacts: Didier.Dubois@irit.fr,
Henri.Prade@irit.fr

Game theory, emotion and trust: from social sciences to artificial intelligence



>>> Dominique LONGIN and Emiliano LORINI,
CNRS senior scientists
at the IRIT (UPS/CNRS/UT1/UT2/INP).

In recent times, game theory has begun to take into account both motivations and more and more complex emotions such as trust, guilt or shame. These are useful concepts for designing robots similar to humans.

During the last century, game theory became the dominant paradigm in social sciences for modeling and for explaining social interaction between human agents and economic agents (states, companies, banks, etc.). The goal of this theory is to explain and predict social actors' choices in strategic interaction contexts, that is, when the choice of a given agent depends on what other agents decide to do. Classical game theory is based on a very simple conceptual frame including the concepts of preference and action. But following the work of the economists John Charles Harsanyi and Robert Aumann, game theory began to include the concepts of knowledge and belief in order to model strategic interaction situations with incomplete information. More recently, the concepts of emotion and trust have become central to this theory. Empirical evidence and psychological theories show that emotions such as guilt, shame or regret affect human strategic decisions. Furthermore, it has been proved that trust plays a crucial role not only in individual decision making but also in social interaction by fostering cooperation.

Reasoning from emotions

In recent years, game theory has become the most used theoretical framework in the area of multi-agent systems (MAS). MAS are a part of artificial intelligence (AI), whose goal is to develop interaction models between artificial autonomous agents (for instance, models of cooperation and coordination, negotiation models, etc.). Similarly, emotion and trust have become central themes in the area of AI. Computational models of autonomous cognitive agents capable of reasoning by taking into account human users' emotions, and whose decisions are influenced by their own emotions, already exist. Several models of trust have been proposed in the area of MAS: statistical and socio-cognitive models of trust and reputation models. These models provide formal and



>>> "Les tricheurs" (Caravaggio, Kimbell Art Museum)

abstract specifications that can be used for developing several applications such as web services, reputation systems (EBay for instance), semantic web, embodied conversational agents and cognitive robotics.

Facial expressions

Work at the LILaC group at IRIT aims to develop logical models of social interaction based on game theory and on psychological theories of emotion and of trust. These models can be exploited as a base for implementing artificial agents capable of reasoning from concepts of emotion and trust during interaction with a human user or with other agents and whose strategic decisions are influenced by their emotions. For instance, a logical model of emotions based on counterfactual reasoning, such as regret or guilt, has been developed. This model has been recently used in the ANR project CECIL for expressing such emotions in a multimodal fashion (for example, facial and vocal expressions, and gestures).

Contacts : Emiliano.Lorini@irit.fr,
Dominique.Longin@irit.fr

My computer knows me so well

How much can computers learn to know us ? Machine learning is a research area where considerable progress has been made in the last 30 years.

Computers are now able to recognize their users' voice or handwriting. Computers have learnt to discriminate human features. Can a menu recommender system, for example, learn what its user likes in order to suggest dishes to her liking? As we can see, "knowing" the user in this case means knowing her culinary preferences.

Decision-making support

More generally, knowing the user's preferences must help improve the quality of the services offered by decision-making support systems or recommender systems. Such systems must often help a user choose some alternative among a huge number of possibilities, notably because of the combinatorial nature of the alternatives. The goal of the system is thus to guide its user in order to help her end up with her preferred alternative. The efficiency of the system for this task will be improved if it knows, at least partially, the user's preferences.

The notion of preferences has been studied in various domains, notably in psychology, social choice theory, micro-economics and decision theory. When studying preferences of individuals that are supposed to be rational, one can consider strict order relations: the relation "is preferred to" is then supposed to be irreflexive (M is never preferred to itself) and transitive (if M is preferred to N, itself preferred to O, then M is preferred to O).

Learning one's preferences

Of course, one can only learn a user's preferences if some data about her is available. Therefore, we assume that we are able to gather information about the user during her interaction with the system. In particular, we suppose that we can obtain pairwise comparisons between alternatives (between menus in our example). For instance, if at some point during the interaction

with the system, the user modifies a menu that has been suggested, we can record that she preferred the second menu. It is these examples of the user's preferences that will be the basis of the learning process. The problem is now to generalize these specific preferences, in order to obtain a total order relation over the possible menus. This induced relation can then be used to predict the user's preferences when the system suggests new menus.

As in any machine learning problem, the choice of the type of model that one tries to learn is crucial. We can try to learn ranking functions that associate a numerical value to every alternative. The ADRIA (Argumentation, Decision, Reasoning, Uncertainty and Machine Learning) research group at IRIT studies the induction of rules of the form: "If the main course is fish, then white wine is preferred to red wine, whatever the starter". We try to learn an order relation that is the transitive closure of pairwise comparisons implied by such rules. Depending on the type of rules that are authorized, one obtains classes of models of preferences of various richness, that can be learnt efficiently or not.

In particular, the ADRIA group has characterized the complexity of learning separable preferences (when preferences over, for instance, the main course, the wine, the starter and the dessert do not depend on one another). We have also proposed algorithms to learn lexicographic preferences - when some components of the menu are more important than others. These results have been obtained in collaboration with researchers from the LAMSADE - University Paris-Dauphine and from Mahasarakham University in Thailand.

Contact : Jerome.Mengin@irit.fr



>>> Jérôme MENGIN,
UPS assistant professor,
researcher at IRIT (UPS/CNRS/INP/UT1/UT2)

System security: look for weaknesses

To find security flaws, programs now think like crackers.



>>> Yannick CHEVALIER,
UPS assistant professor
at the Institut de recherche en informatique
de Toulouse (IRIT, UPS/CNRS/INPT/UT1/UT2)

Consider two University of Toulouse students, Alice and Bob, who share a bike. They don't meet regularly, but devised an ingenious protocol to use the bike. When Alice stops riding the bike, she puts her padlock on it. When Bob plans to use the bike, he also puts his padlock on it. This way, the next time Alice sees the bike with two padlocks on it, she can safely remove hers, leaving Bob's padlock on. He will then later be able to use the bike. To sum up, Alice and Bob have devised a set of rules such that, as long as they adhere to it, they will be able to use the bike, and the bike cannot be stolen. Computer scientists use cryptographic algorithms instead of padlocks to secure communications on the internet instead of a bike, but the idea is the same. The security analysis of cryptographic protocols consists in assessing where the devised set of rules is sufficient to provide a claimed guarantee. In plain text, is Alice and Bob's confidence that their protocol protects their bike well founded? While cryptographers are chiefly interested in a padlock's resistance (or rather on the security provided by a cryptographic algorithm), the members of the LILaC (**L**ogic, **I**nteraction, **L**anguage, and **C**alculus) work on logical flaws, i.e. Flaws that do not rely on lockpicking or on a shear, given the set of rules. More specifically, they search algorithms able to automatically find such flaws when they exist. For example, if Charlie knows Alice and Bob's protocol, he can wait until there's only Alice's padlock on the bike, put on another padlock looking like Bob's, and wait until Alice comes back and remove hers. Charlie will at this point be able to ride away with the bike.

Logical analysis

They work on a logical model of the system under scrutiny and of the goal properties. For example, a logical modeling of a function f (e.g. a decryption function) with one argument x would be a formula stating "for any message x , either x cannot be computed by the attacker or $f(x)$ can be computed by the attacker". Similarly, we state that a message $f(m)$ must remain confidential by the logical formula "the attacker cannot compute $f(m)$ ". These expressions are clauses of first-order logic. They express non-disjoint possible cases, and their meaning is that for every ground value of the variables at least one of the cases must be true. When there are several such clauses there is a risk that the cases are incompatible. For instance adding the clause "the attacker can compute m " yields an incompatible (logicians say *unsatisfiable*) set of clauses. What's interesting is that a set of clauses modelling

a system together with its purported properties is unsatisfiable when one of the properties does not hold. Consequently, as logicians, our goal is to determine whether a set of clauses modelling a system is unsatisfiable.

Resolution

Though the problem is conceptually simple---it suffices to try all possible instances of each variable and see whether each clause is satisfied---it cannot be solved by a machine in general because there is an infinite number of possibilities for each variable. Alan Robinson has however devised a principle he named resolution that permits to speed up the examination of all possible instances. It is based on the combination of the cases occurring in the clauses. For instance, if a clause states that " x cannot be computed by the attacker or $f(x)$ can be computed by the attacker" while another states that " m can be computed by the attacker", resolution on the first case of the first clause with the only case of the second clause yields a clause stating " $f(m)$ can be computed by the attacker". Using the resolution principle one adds new clauses to a set of clauses. It is guaranteed that if the initial set of clauses is unsatisfiable, one will end up adding an unsatisfiable clause. But if the set of clauses is satisfiable, the computation goes on indefinitely.

Unification

This principle is based on the detection of when two clauses state a fact and its negation. One thus needs to be able to compute when two facts have common instances. This computation is called *unification*. While in the basic case this problem is simple, the security analysis of protocols depends upon additional equalities constructions, expressing for instance that the state of the bike with two padlocks is the same regardless of whether Alice or Bob puts her/his padlock first. The existence of such equalities makes unification unsolvable by a machine. Given that both detecting when two facts have common instances and detecting when a set of clauses is unsatisfiable cannot be solved *in general* by a machine, the work on security in the LILaC team consists in finding classes of sets of clauses and equality theories that are general enough to specify the systems analyzed and specific enough to be solvable by algorithms. For further information on how this can be applied to find a flaw on Google's implementation of SAML, we refer the reader to our latest project website, <http://www.avantssar.eu> !



>>> A bike with two padlocks
awaits the reader in front of the IRIT

Contact: ychevali@irit.fr

Artificial Intelligence at work in negotiation



>>> Argumentation in AI Group, ADRIA and LILAC teams: Leïla Amgoud (research scientist, CNRS), Philippe BESNARD (CNRS), Claudette CAYROL (UPS professor), Sylvie DOUTRE (associate professor at UT1), Florence DUPIN de St-Cyr (UPS associate professor), Marie-Christine LAGASQUIE-SCHIEUX (associate professor at UPS) in the Toulouse Institute of Computer Science Research (Institut de recherche en informatique de Toulouse, IRIT, UPS/CNRS/UT1/INP)

Which side to take in a trial? Which is the best medical decision to take? When choices are difficult, artificial intelligence can provide valuable assistance in dissecting the foundations of an argument. It becomes a theoretical tool for analyzing and formalizing the interactions between rational agents, for example, in negotiation.

An opinion is justified by giving reasons that enforce or explain it. These reasons, called arguments, can take various forms, have different strengths, and are more or less relevant with regard to the thesis. Argumentation is a process that consists of evaluating and comparing arguments and counterarguments in order to select the more acceptable ones.

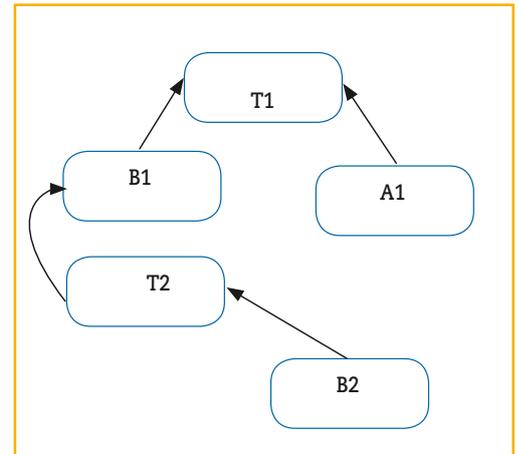
For an autonomous agent, it is a major component of reasoning, explanation of reasoning and decision support, especially in the presence of contradictory information.

The chances of reaching a consensus

Argumentation also plays an important role in multi-agent interactions in general and in particular for negotiating. Negotiation-based reasoning enables agents to explain their choices with arguments. In the light of a new argument, an agent may revise its beliefs and preferences, increasing the chances of reaching a consensus.

How to appreciate an argument? This is the central issue on which the group "Argumentation in AI" of IRIT is working.

The work of this group covers both the formal aspects of argumentation and the use of this formalism for reasoning and decision problems. Regarding the formal aspects, the group studies the interaction between different types of arguments, the methods of comparison and criteria for judging whether an argument is acceptable, and the dynamics of a system of argument, that is to say the change induced by the inclusion of a new argument or the deletion of an existing argument.



>>> AN ARGUMENTATION DIALOGUE: Tom (T1): "for going downtown, my car is a good mode of transport".
 Bob (B1): "no, a car is a too dangerous mode of transport".
 Tom (T2): "no, my car is equipped with airbags".
 Bob (B2): "an airbag can explode!".
 Anne (A1): "anyway, there's too much traffic to use the car"

Trials and medical decision

Formalisms are used to explain decisions, classify objects, and to model negotiation between agents. Although our research is usually very "upstream", it leads to many practical uses and affects many different areas. Argumentation is used, for example, to model and analyze legal arguments (the transcript of a trial). There exist also numerous applications in medicine; for example an argument-based formalism has been used to produce a tool managing the exchange of tissue between hospitals in transplantation (European project ASPIC). Finally, it is often used in trade negotiations on the web.

Contact: lagasq@irit.fr

Thinking fast without too much effort



>>> Martin COOPER and Olivier GASQUET,
UPS professors and researchers at IRIT (UMR UPS/
CNRS/INP/UT1/UT2)

Formal models of artificial intelligence require the development of algorithms that automatically solve the problems they pose. An essential requirement is that they do not consume too much computing time or too much memory.

The reconstruction of the shape of a 3D object from a 2D line-drawing is an example of a problem that can be modeled in terms of numerical, ordinal or structural constraints together with preferences (for flat surfaces and right angles, for example). Solving these constraints allows a program to reconstruct the object or to detect its physical impossibility.

The CSP (“Constraint Satisfaction Problem”) has applications in many areas (such as in the aviation or automobile industries), but if constraint problems must be solved in real time, it is extremely difficult to guarantee a quick response time. Techniques known as “compilation” use off-line pre-processing to solve part of the problem (the model of the system to be diagnosed or the vehicle to be configured). Members of our research team work with industrial partners (Renault, Access Commerce) to integrate preprocessed structures into configuration systems for on-line sales.

Languages

In another line of research, the evolution of a discrete system can be modeled as a “transition system”, the transitions marking the change from one system state to another. Examples include executing a program, using a computer-security protocol, driving a vehicle, or evolving information (knowledge) within an artificial agent. “Can the system get blocked? Does it contain unnecessary states? Is it possible to reach a state satisfying a

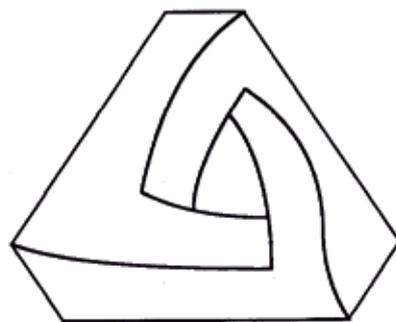
given condition?”. At IRIT we design formal “modal” languages in order to express such questions and to provide an automatic answer using calculations. Conversely, thanks to these modal languages, it is possible to describe the properties required of a system (this defines a modal logic). Appropriate algorithms can then determine the existence of a system, called a model, satisfying these requirements and even build such a model.

Tractable problems

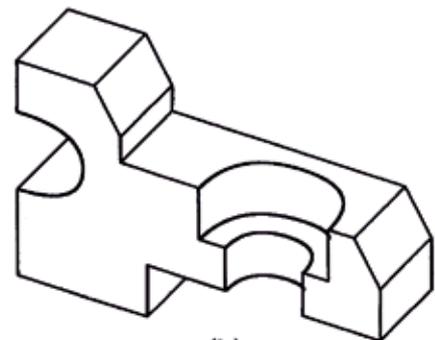
It is not only important to discover algorithms and optimize them, but also to search for classes of problems that can be solved by these algorithms. Particular emphasis is placed on the identification of so-called tractable classes of problems whose solution time does not increase exponentially. A joint research project with the University of Oxford has identified several new tractable classes.

Our work has also given rise to the production of free software: Toulbar2 (in collaboration with INRA <http://carlit.toulouse.inra.fr/cgi-bin/awki.cgi/ToolBarIntro>), which is a complete program for solving valued constraint problems and LoTREC (<http://www.irit.fr/Lotrec>), which is a development platform for modal logic and model building.

Contacts : cooper@irit.fr, gasquet@irit.fr



(a)



(b)

(a) An impossible object; (b) a possible object.