Your Computer is Stupid

The field of artificial intelligence is progressing. TU Vienna is a major center of this line of research. Several conferences on this topic will be held during the "Vienna Summer of Logic".

The ball does not fit into the suitcase, because it is too small. What is too small? The ball or the suitcase? For humans the answer is evident, but computers still cannot handle such simple questions. Research in artificial intelligence is bound to change that. Using logical methods, strategies are developed to approximate natural reasoning, based on common sense. This has nothing to do with building "artificial humans" – it is about creating clever, helpful computer tools. Internet search engines should become a lot more useful using artificial intelligence. Simple databases will be replaced by well-organized knowledge bases, which could for instance help doctors with medical diagnoses.

From Chess Computers to HAL 9000?

Playing chess ist often considered to be the ultimate benchmark for intelligence. When computers can beat chess grand masters, are we perhaps already close to the intelligent computers we know from science fiction movies, which may even surpass our intelligence? Already in the 1940s, the author Isaac Asimov thought about the "laws of robotics" that would be needed to organize the relationship between intelligent computers and humans. In Stanley Kubrick's "2001: A Space Odyssey", the hyper-intelligent computer HAL seizes control of a space mission, and even "Terminator" exhibits intelligent behaviour – even though he does not appear to be particularly empathic.

But even after several decades of computer science, we still cannot buy intelligent robotic aides. True intelligence is very different from solving chess problems. A chess computer just tries out many different possibilities, but that is not intelligent behaviour. In order to create real intelligence, new logical methods are needed.

Parrots and Penguins

In classical logic, new facts are derived from given statements. When additional information is provided afterwards, this should not change anything about the previously derived facts. From "birds can fly" and "this parrot is a bird", we can deduce: "This parrot can fly." Learning the additional piece of information "this parrot is called Kurt" does not change anything about him being able to fly. In logic, this is called "monotonicity".

However, things are not always this simple. From "All birds can fly" and "Alan is a bird", we may deduce that "Alan can fly", but when in addition we learn that "Alan is a penguin", then we are forced to withdraw our conclusion – a case of logical non-monotonicity. In this case, additional information did not just make us learn something new, but it forced us to discard a statement we had believed to be true. Humans can deal with these kinds of logical exceptions effortlessly, but in a classical, monotonic logic this is not possible.

The Right Logic for Every Application

"Today, we work with different kinds of logics, suited for different fields of application", says Professor Thomas Eiter (TU Vienna). Some of these logics are non-monotonic and can handle cases

such as the example with the penguin. Some logics give up the idea of a binary truth value. A statement does not necessarily need to be either true or false. “Very likely true” or “rather false” statements can also be used in logical calculations.

“An important part of intelligence is the ability to cope with incomplete or even contradictory information”, says Thomas Eiter. This is important for medical diagnostics, where a list of symptoms is used to determine the disease at hand. The course of a disease may be atypical, some symptoms may not show up, or additional symptoms may develop. Using classical logic, problems would arise quickly. **Today** powerful **logical** tools are **available which search** for possible diagnoses in a medical **knowledge-base** and estimate their respective probabilities.

Connecting Knowledge

Knowledge which is just stored on a hard drive has nothing to do with intelligence. The crucial part is adding structure to knowledge, adding meaning by connecting pieces of knowledge to other known facts. To a standard database it does not matter at all whether the words in column three are names of people or locations. An intelligent program knows that these two categories have completely different meaning and deals with them in different ways. The program establishes relations between pieces of content, it recognizes structure, and it may even be able to reconstruct missing information from known data.

An intelligent program may know that a suitcase belongs to the category of containers, and that objects can only fit if they are smaller than the container. From that it could deduce that it is the suitcase which is too small if the ball does not fit into it, and not the other way around. This kind of intuitive understanding is still missing in the computer programs we use every day. By human standards, they are rather stupid. But Thomas Eiter believes that this will change considerably in the years and decades to come. We will get used to a much more intuitive way of cooperating with computers than we find possible today.