Promedas is a PRObabilistic MEdical Diagnostic Advisory System – a computer program that figures out how likely certain medical diagnoses are from a given set of symptoms and test results. It addresses the hugely complex task faced by doctors every day: how to ensure that even unusual or obscure causes of symptoms are not overlooked. Whether the doctor is a specialist in one area or a general medical practitioner, keeping up-to-date and remembering the myriad relationships between illness and symptoms is often tremendously challenging. Promedas does not make a diagnosis, but it suggests alternatives to the doctor and based on the information it has been given, shows exactly how likely it is for the patient to be suffering from each condition.

Since the beginnings of Artificial Intelligence, there has been a dream of creating a computerised expert, filled with our knowledge, and able to infer new conclusions of its own. These computer programs became known as expert systems, often using elaborate trees constructed of rules: “if symptom A and not symptom B then ask about C, if symptom C then diagnosis is D”. But unlocking the knowledge held within scientific journals and coded in the neurons of specialists was not so easy. In medicine, for example, medical conditions and symptoms are not either true or false. Not all people experience the same symptoms, and the appearance of some symptoms is hugely more significant than others. There can also be a large number of symptoms, leading to rules with excessive numbers of variables. So the old “decision tree” methods were often cumbersome and made bad decisions, sometimes ruling out possibilities for no other reason than data being presented in an unexpected order.

The failure of these rationalist “good old fashioned AI” (GOFAI) methods led to many scientists rethinking the ideas. Clearly an expert system needed to represent knowledge in some form, and clearly that knowledge needed to be used with data to infer some form of decision. But how best to achieve these goals?

One common solution was to use fuzzy logic, where the binary true or false rules were turned into linguistic variables such as “partially true” or “mostly false”. But even fuzzy logic still suffered from problems: it might enable the expert system to define degrees of truth, but in fields such as medicine, a partially true diagnosis is not ideal. Instead, it would be much more useful if the probability of a diagnosis being true was provided. For example, while a runny nose might result in a “partially true” diagnosis that you have a virus in a fuzzy logic system, it would be more helpful if the system could infer that you have a certain probability of suffering from several different illnesses, some much more probable than others. The solution, as exemplified by the medical decision support system Promedas, was to use Bayesian inference rules.

Thomas Bayes was a mathematician born in London in 1702. Amongst his works, he wrote about probability. Instead of being concerned with, say, the probability of drawing a black ball from a bag of a certain number of black and white balls, Bayes was interested in the inverse probability of the event. In other words, if you had drawn more black balls compared to white balls from the bag, what was the probability that the bag contained more black balls and so you would draw another black one next? Given a hypothesis (for example that the next ball will be black) and some information about which balls have been picked previously, Bayes figured out the maths to infer the probability that the hypothesis was true or not. Several decades later, Frenchman Laplace developed these ideas further, creating a more general version of the Bayes theorem for use in astronomy and physics.

Amazingly, two centuries later, systems such as Promedas now use Bayesian inference to achieve decision support. It’s ideal for an application such as medicine where we have plenty of evidence that certain symptoms tend to be observed for specific illnesses.

Promedas is the result of many years of development by researchers such as Wim Wiegerinck, led by Martijn Leisink and Bert Kappen of Radboud University Nijmegen in the Netherlands. It now contains Bayesian inference rules that cover a huge area of the medical domain. The Bayesian approach offers a major advantage. In the words of Martijn Leisink: "The clear advantage of the probabilistic representation is the natural way that different diseases or findings influence each other. If one finding leads to two different diagnoses these are difficult to merge in a rule based system. At least one additional rule is necessary. In the probabilistic setting, making use of the basic probabilistic rules, it is immediately clear how to combine evidence and variables."

The success of Promedas relies on its careful structuring of the dependencies (which findings imply which diseases). It organises its information as a tree in a three-layered noisy "OR" model. The layers of the tree correspond to risk factors such as occupation or drug use, possible diseases and the tests and symptoms. Each node in the network (either a risk factor, disease or test result) is linked according to specific probabilities of cause and effects, with some risk factors likely to cause some diseases, and some diseases likely to cause some symptoms and results of tests. By representing each node as a "noisy OR" Promedas simplifies and speeds up the inference process, making the assumption that each cause can behave independently but multiple causes combine to make outcomes more probable. With a given model and data about the risk factors and test results, the probabilities of different diseases can then be inferred.

> This is the future for medicine with all the sub specialists. They need a program that looks further away than their sub specialisation.

As with all such systems, often the main bottleneck is simply the input of data. Many patient records are incomplete or are not in the right kind of format to enable easy input, and typing in the results of tens or hundreds of different observations and tests can be laborious. As more patient records become stored electronically, this will become less of an issue but in the near future the use of Promedas is likely to be restricted to those mysterious cases where the doctor needs some new ideas. It's an important role, for specialists who have chosen to focus on one specific area of medicine may become less knowledgeable about other areas. The vast amount of expert knowledge in Promedas means that all physicians of all specialisations will have access to the same up-to-date specialist information. There can be no doubt that a list of possible diagnoses of varying probabilities for a patient makes an excellent decision support system, for it may suggest rarer alternatives that could be confirmed by additional tests.

The first large-scale trial will begin in early 2008 at the University of Utrecht, Netherlands. While results are still not perfect, initial experiences by doctors are very positive. In the words of Dr Jan Neijt of the University of Utrecht (the physician who has so far provided all of the medical knowledge for Promedas), "…used with reason it is always helpful… This is the future for medicine with all the sub specialists. They need a program that looks further away than their sub specialisation."

Perhaps one day systems such as Promedas will become as ubiquitous as the stethoscope, with their databases updated as new medical findings are published. We can never remove uncertainties from medicine, but with decision support systems we can ensure that all decisions made by our doctors are as well-informed as possible.

Resources:
Promedas (including online demonstraton): http://www.promedas.nl/
Promedas publications:
http://www.snn.ru.nl/nijmegen/publicatie.php3?projekt=Promedas
PASCAL: http://www.pascal-network.org/